

Second Edition

**Best
Practice
in
Inventory
Management**

Tony Wild



Best Practice in Inventory Management

This Page Intentionally Left Blank

Best Practice in Inventory Management

Second edition

Tony Wild



BUTTERWORTH
HEINEMANN

OXFORD AMSTERDAM BOSTON LONDON NEW YORK PARIS
SAN DIEGO SAN FRANCISCO SINGAPORE SYDNEY TOKYO

Butterworth-Heinemann
An imprint of Elsevier Science
Linacre House, Jordan Hill, Oxford OX2 8DP
225 Wildwood Avenue, Woburn, MA 01801-2041

First published by John Wiley & Sons 1997
Second edition published by Elsevier Science Ltd 2002

Copyright © 2002, Elsevier Science Ltd. All rights reserved

No part of this publication may be reproduced in any material form (including photocopying or storing in any medium by electronic means and whether or not transiently or incidentally to some other use of this publication) without the written permission of the copyright holder except in accordance with the provisions of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London, England W1T 4LP. Applications for the copyright holder's written permission to reproduce any part of this publication should be addressed to the publisher

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication Data

A catalogue record for this book is available from the Library of Congress

ISBN 0 7506 5458 9

For information on all Butterworth-Heinemann publications visit our website at www.bh.com

Composition by Genesis Typesetting, Rochester, Kent
Printed in Great Britain

Contents

Preface	ix
Introduction	1
1 The basis of inventory control	4
The role of inventory management	4
Objectives for inventory control	7
Profit through inventory management	9
Reasons for the current stock	11
2 Customer service	17
Meeting customer requirements	17
Measuring availability	18
Demand management	25
3 Managing the inventory	33
Using Pareto analysis for control	33
Stock cover	43
Practical methods of reducing stockholding	51
The approach – how to reduce stocks successfully	54
4 Just-in-time management	59
The zero inventory philosophy	59
JIT environment	64
Advantages of JIT	71
Stock control using JIT	74

5	Organization and management	78
	Where stock control fits into the organization	78
	Responsibilities and targets	80
	Inventory valuation	87
	Skills and systems	92
6	Safety stocks	96
	Learning from history	96
	Normal demand patterns	100
	Evaluation of safety stocks	104
7	Setting the right stock levels	112
	Simple assessment of review levels	112
	Managing lead times	117
	Target stock levels	124
8	The changing role of purchasing	129
	Modern supply practice	129
	Supply partnerships	137
	The ordering process	139
	Order quantities	141
	Purchasing processes	148
9	Forecasting demand	152
	Options for assessing demand	152
	Causes of forecasting inaccuracy	158
	Methods of improving forecasts	162
10	Historical forecasting techniques	164
	Basic forecasting techniques	164
	Weighted averages	170
	Choosing the best forecast	178
	Monitoring forecasts	181
11	Advanced forecasting methods	187
	More forecasting tools	187
	Forecasting for seasonal sales	192
	Other methods	195

12	Material requirements planning – an alternative to forecasting	197
	Avoiding uncertainty	197
	Material requirements planning	199
	Master planning	205
	Batch sizes	210
13	The future – inventory and logistics	212
	The basis of the lean supply chain	212
	Logistics	217
	Review	228
	Questions and answers	231
	Index	263

This Page Intentionally Left Blank

Preface

In the distant past when I ascended from being a scientist to the heights of inventory analyst, I was surprised at the lack of helpful written work on the subject. Although there are now many books, manuals and papers, there does not seem to be much written on the very important subject of 'how to do inventory control faced with the real inventory, customers, and an assortment of inventory records and styles of suppliers'.

Yet better control of stock (inventory) can give major benefits to the profitability of all companies. This book is aimed at showing how good inventory control can be used in practice. It is a result of working continuously on inventory control with a large number of companies over many years. It contains the distilled techniques which have been tried out and proven to work.

This book contains two main ingredients, namely:

- the basics of inventory management as covered in the Institute of Operations Management (IOM) Diploma syllabus
- the application of these techniques to real inventory management.

The applications are what really count, since the knowledge is not at all interesting without application. The experience of the author has been to use the basics to provide powerful changes in inventory and profitability. In several cases millions of pounds of inventory value have been saved and in others customer service has been greatly improved. The topics here have been tried, tested and approved. They are all a matter of common sense, but which piece of common sense to use in a particular situation is a matter that needs deeper understanding.

The book is a driver's manual for Inventory Controllers. It covers the working of the engine (how inventory control techniques work) how to use the controls (what the techniques do and how to manage them), and how to get the best out of the vehicle (how to optimize inventory). Understand the text and it will show the way to guaranteed improved inventory control, reduced stock levels and higher availability.

In preparation for this book there have been contributions from the many thousands of people to whom I have presented courses on these ideas, and who have helped to hone the concepts. My thanks also to those companies who have let me loose on their inventory, carrying out consultancy projects which have extended my knowledge and proved the effectiveness and scope for different techniques.

Central to the development of inventory techniques has been the Institute of Operations Management (formerly BPICS) which has been a channel for professional debate, friendship, and technical information. This book has been created with the support of Richard Turner as Secretary of the Institute.

The evolution of this book has been a two-handed operation. The concepts have been blended into a logical text through the work of myself and Elaine Duckworth, who has shared the task of developing a useful and readable handbook.

Dr Tony Wild

Introduction

- Why we need professional inventory management
- Structure of the book
- Focus on developing theory that works in practice

Everyone is an inventory controller, at home and at work. We all keep food, clothes, domestic items, paper, pens and many other goods. We also have shortages and emergency purchases. Some people regularly have to throw out the contents of the refrigerator because they have been there a while and changed in character. So inventory control is a natural occupation that everyone undertakes, some more successfully than others.

There are many other activities that people do, such as sport, carpentry, music and medicine, and for these it is generally accepted that a level of technical expertise is needed. In carpentry and music there are people who become highly skilled through their natural talents. Success appears to come from a mixture of talent, determination, practice and knowledge. In medicine, particularly in surgery, we are more likely to rely on those who have studied the subject and have a deep knowledge of both theory and practice. Lawyers cannot function without knowledge; often a common-sense approach does not work for them.

As with any profession, inventory control (or the subject of Professional Inventory Management) has a body of techniques and knowledge which differentiates the professional from the do-it-yourself (DIY) enthusiast. This expertise is the result of 100 years of development and refining. Now there are professional inventory managers with recognized qualifications, including Diploma in Production and Inventory Management (DPIM) and Certificate in Operations Management (COM). There are still colleagues who believe that inventory can be managed using a clear brain and common sense. It is unlikely that they would be happy with the same competency in their doctor or financial adviser.

If you want to do better you need to train.

This book is for the inventory control practitioner. With the techniques described many people have been able to manage their stock of inventory so that their customers are happier and so are the accountants. The reduction of inventory value, the avoidance of unnecessary work and the improvement of customer service can be accomplished at the same time through simple application of the techniques discussed. Improved inventory management has been shown to halve stocks and improve service at the same time.

The techniques described are the basic concepts which inventory controllers should have at their fingertips. The more complex theoretical approaches can be left to those developing sophisticated systems, or those with excellent inventory control (and over the years I have not met any of these!). Inventory practitioners should be able to use this book to understand the best approaches and then to apply them to their own circumstances. Simple application of the methods is most successful, while modifications usually result in less effective outcomes.

The text covers the syllabus for the technical qualifications including DPIM and COM. It discusses good practice, inventory theory and the practical application of the techniques. The structure of the book is to start on inventory structure with target setting in Chapters 1 and 2, and how to structure inventory in Chapter 3. This chapter shows how a big impact on inventory levels can be made easily. The approach continues to give major benefits in tightening control. Continuing the theme of restructuring stock leads to the ultimate in stock management – the zero inventory approach – just in time (JIT), discussed in Chapter 4. Gradual absorption of the philosophy throughout inventory management is proving a great benefit, and it is worth examining how to utilize JIT in each stock environment. The rate of progress is often governed by the attitude and/organization in a company, and the development of lean inventory management practice is reviewed in Chapter 5.

The discussion then gets down to the detail of item-level inventory management, the essence of real inventory management. Chapters 6 to 8 are the elixir of item-level stock control, the secret recipe for success. Chapter 6 concentrates on how to keep out of trouble – a rational approach to the risks of stockouts, Chapter 7 shows how to use this information to set stock levels and Chapter 8 discusses the options for ensuring suppliers are agreeable and flexible.

The third aspect of inventory management, one which is usually carried out poorly, is forecasting. Forecasting methods are relatively well advanced,

but there is less practical application of these in most companies. This is an opportunity, discussed in Chapter 9, while the basic techniques are discussed in Chapter 10. Chapter 10 also includes the techniques used to ensure that forecasts produce sensible results. In practice there are items where it is worth managing the stock in more detail, so in Chapter 11 there are some more techniques which are commonly used and provide good forecasts.

The one item at a time approach to inventory management is really not appropriate in many situations. Users often require a range of items to do a job. For example, servicing the car requires filters, plugs, oil, etc., all available at one time. Inventory controllers should therefore balance the stocks to meet the structure demand. Unfortunately, the usual techniques do not help, and so the approach described in Chapter 12 is applied. The technique, developed over many years in manufacturing, is an important tool for general inventory management, although the application does not need the complexity used in manufacturing.

The inventory manager who has a good knowledge of these topics can then look round for new challenges. Industry has not been slow to provide them. In addition to the increasing pressure on inventory value and availability levels, there are new dimensions being added. The new challenges are in logistics and supply chain management. The balancing of inventory in one store is not sufficient; now the inventory, distribution costs and manufacturing efficiency have all to be optimized at the same time. Chapter 13 reviews the opportunities in this area.

Chapters 3, 6 and 7 are the bedrock of inventory practice. From experience, inventory management as described in this book works well, and the purpose of these pages is to enable more individuals to use the methods with confidence.

Readers outside the UK may find that a little translation is required. For instance in America the term 'inventory management' is applied to the whole subject, whereas in the UK it is normally known as 'stock control'. The terms 'inventory' and 'stock' are both used in the text to mean the same (in the UK inventory has a financial flavour and stock an operational context). It is hoped that this and other nuances of translation do not cloud the understanding of this technical and fascinating subject, whatever it is called.

The basis of inventory control

- The conflicts which beset inventory controllers.
- What inventory is there to achieve?
- How good inventory management will improve profitability.
- Analysing current stockholding.

The role of inventory management

The success of a venture depends on its ability to provide services to customers or users and remain financially viable. For an organization which is supplying goods to its customers, the major activity is to have suitable products available at an acceptable price within a reasonable timescale. Many parts of a business are involved in setting up this situation. Initially these are the marketing and design departments. Then purchasing, and in some cases, manufacture is involved. For an item which is already in the marketplace, the main activity is providing a continuity of supply for the customers.

Inventory control is the activity which organizes the availability of items to the customers. It co-ordinates the purchasing, manufacturing and distribution functions to meet the marketing needs. This role includes the supply of current sales items, new products, consumables, spare parts, obsolescent items and all other supplies.

Inventory enables a company to support the customer service, logistic or manufacturing activities in situations where purchase or manufacture of the items is not able to satisfy the demand. Lack of satisfaction could arise either because of the speed of purchasing or manufacturing is too protracted, or because quantities cannot be provided without stocks.

Stock control exists at a crossroads in the activities of a company (Figure 1.1). Many of the activities depend on the correct level of stock being held, but the definition of the term 'correct level' varies dependent upon which activity is defining the stock. Stock control is definitely a balancing act between the conflicting requirements of the company, and the prime reason




Figure 1.1 Which direction?

for the development of inventory management is to resolve this conflict in the best interest of the business. A conventional supply organization will have many departments including sales, purchasing, finance, quality assurance, contracts and general administration. In some cases there will also be manufacturing, distribution or support services or a variety of industry-specific activities. Each of these has a particular view of the role of stock control. Note the terms 'inventory' and 'stock' are used to mean the same thing in this book.

Sales consider that good stock control enables the company to have available any item which will meet immediate sales for as large a quantity as demanded: this requires large stock. For service companies where parts service is involved, the control of stock at the customer interface is traditionally left with the person carrying out the service, and this has led to overstocking and poor control. Similarly in distribution, the effect of bulking up shipments will lead to high stock levels, and a compromise has to be reached.

Inventory control keeps balancing conflicting requirements.

Purchasing consider that stock control provides the opportunity for goods to be purchased so that optimum prices can be obtained. Buying items in bulk often reduces the purchase price and it improves the efficiency within the purchasing department. The store is a means of keeping the bulk purchase items after buying advantageously.

Finance departments have a problem with stock because it consumes vast amounts of working capital and upsets the cash flow. One benefit of stock from a financial standpoint is that provisions can be made in case the stock turns out to be unsaleable, and this value can be adjusted to modify the profit figure in times of good or bad financial results. However, the existence of these provisions in the first place is detrimental to the finances of the company.

Quality management normally has the effect of slowing down the progress of stock while the necessary checks are made. This means that quality and inventory personnel effectively work in opposition. The gradual introduction of formal quality standards for supply and manufacture has reduced this conflict in most organizations. Supplier conformance has been improved enabling intrusive checking to be minimized.

General management see stock control, rightly, as a source of information. Some managers consider that they should be able to use stock control to give an immediate supply of information, statistics and forecasts. This can result in large amounts of unstructured work collecting, analysing and providing information.

The traditional view of manufacturing companies has been that large batches reduce the direct production costs. Manufacturing management tends to aim more for plant and labour efficiency and allow high stocks in order to avoid the disruptions caused by shortages, breakdowns and changing customer demand.

Good stock controllers (or materials management) will keep the stock down as long as they are responsible for stocks of everything, including raw materials, finished and semi-finished stocks, consumables, tools, work on the shop floor and, in fact, all inventory items.

Do not confuse high stock with good availability.

These views are now subject to question. Those who agree with these ideas should be considering a rethink and change to a consensus on stock management. The development of stock control practice has been given impetus by the move toward total quality management (TQM) and JIT concepts. These have focused on the possibility that good management, particularly, is both desirable and profitable, and has brought into question all the attitudes which are presented above as alternative views of stock control.

The responsibility for maintaining the correct balance is normally left to inventory management. (The actual name for the activity varies but the task remains the same: the controller could be materials or inventory manager,

stock planner or controller, or even considered as a buyer or logistics controller.)

Demand for the product results from changes in market and financial forces and the amount and type of stock control varies in tune with this. Stock control is a dynamic activity, and the successful inventory manager has to ensure that the balance is kept right. This requires both communications skills and professional inventory techniques. Operating methods should be continuously revised to reflect the changes and systems should be altered to suit new situations and operating policies.

Objectives for inventory control

The first consideration is the overall objective of the work of stock control. Like all other activities in the company, inventory management has to contribute to the welfare of the whole organization. The logistic operation must aim to 'contribute to profit by servicing the marketing and financial needs of the company'. The aim is not to make all items available at all times, as this may well be detrimental to the finances of the company. The normal role for stock control is to 'meet the required demand at a minimum cost'.

Our aim towards long-term profitability has then to be translated into operational and financial targets which can be applied to our daily operations. The purpose of the inventory control function in supporting the business activities is to optimize the three targets:

- customer service
- inventory costs
- operating costs.

The most profitable policy is not to optimize one of these at the expense of the others. The inventory controller has to make value judgements. If profit is lacking, the company goes out of business in the short term. If customer service is poor, then the customers disappear and the company goes out of business in the longer term. Balancing the financial and marketing aspects is the answer: the stock controller has a fine judgement to make.

The first target, customer service, can be considered in several ways, depending on the type of demand (discussed in Chapter 2). In a general stores environment the service will normally be taken as 'availability ex stock', whereas in a supply to customer specification, the service expected would be delivery on time against customer requested date.

The second target, inventory costs, requires a minimum of cash tied up in stock. This target has to be considered carefully, since there is often the

feeling that having any stock in stores for a few months is bad practice. In reality, minimizing the stock usually means attending to the major costs: very low-value items are not considered a significant problem. Low inventory can also be considered in terms of space, or other critical resource. Where the item is voluminous, or the stores space restricted, the size of the items will also be a major consideration.

The third target, avoiding operating costs, has become more of an issue as focus has been placed on inventory management. The prime operating costs are the stores operations, inventory control, purchasing and the associated services. The development of logistics, linking distribution costs with inventory, has added this new set of transportation costs to the analysis.

Optimizing the balance of these three objectives is the focus of stock control. The better the balance the greater the profits provided for the company (as will be shown in the next section, profit through inventory management). The improvement in stock control has been slow and gradual, created by new technology, financial need and competitive pressure. Those companies who can tighten their control faster than the average will flourish, but those which do not keep up with the average, even if they are improving, will gradually dwindle. The way to achieve much is relatively simple, and is outlined in the remainder of this book. The trick of the good inventory controller is to meet the objectives simultaneously, not one at a time, and of course 'the better the control the smaller the cost, the lower the stock levels, and the better the customer service'.

In reality: more inventory = worse delivery on time.

One of the dichotomies of inventory control is that at item level, the more stock the better the availability. (This is discussed quantitatively in Chapter 6, in the section on evaluation of safety stocks.) However for the whole inventory, experience has shown that the businesses with the highest stock are often those which have the worst availability. These observations are not in conflict if the causes of this are considered. Stockouts result from holding too little stock for the offending lines, because the forecasts, monitoring or controls are inadequate. High stock levels arise because too much stock has been purchased, through bad forecasting, monitoring or controls. High stock and poor availability are caused simultaneously as a result of poor control. The problem rests with the inventory controller, and the solution is in improved techniques. By using the techniques discussed in the remainder of this book the reader can reduce stockholding by 25 per cent even where the control is already reasonable.

Profit through inventory management

There is obviously a link between financing stock and profits. This can be seen most clearly using an example.

An example of the effect of stock on profitability

There are two companies in the same industry, M. Tight Ltd. and The Slack Company. Each has annual sales of £5 million and employs eighty-five people. They both have fixed assets of £2.5 million in plant and buildings, and operate in an industry where the value of customer debtors is balanced by the credit owed to suppliers.

In their organization the two companies differ. Tight is a small, independent company and Slack is a small part of a multinational group. Although they manage to achieve similar customer service, M. Tight Ltd. has concentrated on inventory control and has reduced stocks to £500 000, while the Slack Company have poorly defined responsibilities in this area and hold £2.5 million worth of stock.

The Slack Company have to finance an extra £2 million inventory, and they can borrow this from their head office. However, the company expects a 15 per cent return on investment and so charges Slack £300 000 per annum in interest. They also have the costs of controlling, holding and organizing the extra stock which is stored in an adjacent warehouse. These warehousing costs are shown in Table 1.1.

Looking at company finance, the size of expected profit will depend on the size of the company. Therefore, to compare how well different

Table 1.1 Warehousing costs

Cost element	Cost (£ p.a.)
Stores: extra area rent, heating etc.	500
Stores, equipment; trucks and racking depreciation	1 000
Maintenance	1 000
Stock obsolescence	5 000
Storekeeper's wages and employment costs	14 000
Annual stocktaking	500
Control systems and computer time	1 000
Total	23 000
Plus financing charge £2 million @ 15%	300 000
Total extra cost	323 000

businesses are doing, the profit has to be scaled to company size. The two most common ways of doing this are:

$$\text{Return on sales} = \frac{\text{Profit}}{\text{Annual sales}}$$

$$\text{Return on assets} = \frac{\text{Profit}}{\text{Assets employed}}$$

The assets include both:

- fixed assets (plant, buildings, land, etc.)
- variable assets (stock and debtors).

Now examine the Profit and Loss account for the two companies. The difference between sales revenue and costs is £250 000 for each company.

M. Tight Ltd has a return on turnover of 5 per cent, making £250 000 profit on the £5 million sales.

Slack has extra costs of £323 000 and thus makes a loss of £73 000. Good stock control has made the difference between profit for M. Tight Ltd and loss for The Slack Company.

The next year M. Tight Ltd made £500 000 profit on the same £5 million turnover, a 'return on sales' of 10 per cent. The Slack Company had a profit of £500 000 – £323 000 = £177 000 because of their annual stockholding costs.

Although M. Tight Ltd has a return on sales of 10 per cent, The Slack Company only get

$$\frac{177\,000}{5\,000\,000} = 3.5 \text{ per cent}$$

This shows how important it is to the profit of a company, that stockholding costs are minimized.

That is a surprisingly significant difference and, when we look at return on capital, the situation is even more dramatic.

Both companies have plant and buildings worth £2.5 million, with total assets consisting of the fixed element of plant and machinery, working capital tied up in inventory and creditors who are offset by outstanding debtors.

M. Tight Ltd also has £500 000 worth of stock, making a total capital investment of £3 million.

The return on assets for M. Tight Ltd is therefore

$$\frac{500\,000}{3\,000\,000} = 16.67 \text{ per cent}$$

The Slack Company has £2.5 million of stock making a total investment of £5 million. Their return on assets is therefore

$$\frac{177\,000}{5\,000\,000} = 3.5 \text{ per cent.}$$

Hold a minimum of the right items to be profitable.

By running the same business with two different inventory plans, we can earn anything from 17 per cent to 3 per cent for our trouble.

M. Tight Ltd has cash available for investment, wages, better working conditions, and to share out to the pension schemes, insurance companies and other investors in their company.

The Slack Company has little cash to spare other than to pay off the bank. Imagine the situation: poor profitability leading to disillusionment, poor working conditions and fears for the future.

Reasons for the current stock

It is easy to assess the effectiveness of inventory control in a company. All that is required is a visit to the stores. The stock level is a result of the effectiveness of the stock control, but what is there is probably not ideal! Instead of considering the detailed ordering history for each item (and there are some interesting causes for stockholdings), the inventory can be classified according to why it is now in stock. An examination of the factors which give rise to the stock in a typical store reveals causes that are detailed in the following sections.

Your current stock shows what is going wrong.

Purchase order quantity

There are several aspects which affect the ideal batch size. For example, most stores have items that are bought in bulk but sold in smaller quantities.

At any time there is a balance of receipts in stock – a half-empty box or pallet. This is usually an acceptable situation for keeping stocks at the right level, but that stock nevertheless still contributes to inventory value.

The size of logistic inventory is often determined by the transport methods chosen and the batch size of delivery which results. Changing the delivery method can change the stock levels in the stores.

Where items are being manufactured, the rate of production in one stage of the process may be greater than at other stages. The ideal solution to this would be Kanban (see the section on pull systems in Chapter 4), but often there is a stock of items manufactured slowly awaiting to fill a process load. Alternatively a bulk process, or one where items are produced at a fast rate, will create stocks which are subsequently used up at a much slower rate.

Frequent, small deliveries keep stock low.

Safety stock

In many situations customers do not provide information of their demands far enough in advance. To compensate for this problem there are two tactics. The first is to organize the customer to give more forewarning. Is this impossible? Manufacturers of specialist, make to order or highly sought after products may make their customers wait until the product is ready. In more competitive situations, however, the customer has more power and may not be prepared to wait. This puts the pressure on suppliers to reduce the lead times and better to forecast demand. With better knowledge of the customer's demand, the unexpected peaks would no longer be a surprise and stocks could be varied to cover these occasions. The second option is to hold sufficient stock to cope with unexpected or excess demand. Safety stock (or buffer stock as it has been called) is there to cover our inability to predict demand. The inventory manager has an investment choice. Either invest in inventory or invest in information. Those who can accurately predict the customer requirements do not need safety stock.

Let the source of the risk hold the inventory.

A secondary role for safety stock is to compensate for failure on the supply side, non-delivery of purchased goods, information failures, technical breakdown or, even, industrial action. You cannot be too careful! Generally the approach to these situations is to monitor the situation continually, and there are several standard ways of doing this (see lead times and schedules in Chapter 7).

Market change

Customers change their mind, contracts are lost, markets vary. These all cause excess or even obsolete stocks. But this should not be of a significant value. The changes that cause a build-up of inventory are either step changes or gradual changes. A gradual change is the most common, and leads to lots of minor excesses of stock. The cause is normally a forecasting method that is not reactive enough, and the simple cure is often to use a more advanced forecasting method. Methods discussed in Chapter 11 and used in the right way can avoid this situation.

A sudden reduction in demand can only be anticipated through understanding or information. Input from sales personnel is important in avoiding excesses as well as stockouts. Sales can warn of impending major changes, even if they are not good at identifying the size or the timing of changes. This is the best way to avoid having items remaining after demand is satisfied.

Detect trends and react automatically.

Obsolescence

In addition to excess stock caused by customers, a significant amount of obsolete stock can be caused within the company itself. Consider several situations:

- 1 The marketing department put an item on promotion without informing the inventory controller. The demand increases as the promotion starts; inventory control see demand escalate and stocks run out. They therefore order more. The promotion ends just as the new stock arrives. The company now has three years' worth of stock cover at the normal usage rate, and the major customers have already bought extra during the promotion.
- 2 The design department has replaced the old version of a product with an improved one; the sales team were extremely pleased with the new product and started selling it immediately. However, there was a large amount of the old product left in stock which was not sold.
- 3 The marketing department has just launched the new season's catalogue. It contains some new items, and some old ones have been deleted. There are still good stocks of some of the deleted items.
- 4 The sales director has bought a job lot of items which are going to be sold as a new line.

All these are common occurrences in some businesses. They all cause self-inflicted obsolete stocks and can be avoided by improving communications between departments. The responsibility for organizing this communication revolution lies with stock control, since the other departments rarely understand the challenges of inventory management or the concept of lead time. (Refer to the excess and obsolete items section in Chapter 6.)

Poorly defined responsibility

When buyers procure to optimize purchasing costs and operations, and sales sell to suit customers, in the middle are the stores, accumulating stock and controlling the conflict between the two activities. Stock control in this case becomes simply recording and storing. Inventory management should have the responsibility for assessing the demand pattern, determining the appropriate stock levels to support this and then instructing purchasing what and when to reorder.

Planned inventories

Cyclic demand, on a monthly, quarterly or annual basis, may be supported most economically by constantly holding high stock levels. In the stores there may be stock which is building up to satisfy a demand event or which is stored to cover a definite requirement. The feature of these is that they cover situations where the demand events are certain and in excess of the supply capabilities, as, for example, in industries where there is an annual holiday shutdown. Stock has to be provided to cover customer demand over this period, and it is not possible to acquire the extra quantity at the last minute to cover the shutdown. The stock is therefore gradually increased over several weeks to maintain continuity. This effect can be complicated by different stages in a supply chain taking holidays at different times.

Layout and location of stores

The effect of a stock location must be analysed in the context of the whole business. In a company where the store is situated inconveniently or a long way from the users, the users tend to keep their own stock. In fact, they might keep some just in case the stores have run out when they want some. This lack of co-operation and control can result in satellite stores being set up and excess stock being held.

A worse situation results if different departments independently order the same items. In this case, there are totally independent stock systems, usually controlled by amateur stock controllers whose major interest is another aspect of the business and, as a consequence, stocks are excessive.

Fortunately where this happens the items are normally low-value consumables.

Where a company has more than one store (such as in distributed warehousing), a similar range of stock items are held in each, resulting in duplication of stock. If there are slow-moving items in one store (try to find a store without slow-moving items!) then the same item can be found in the other store – still moving slowly, and total stock across all the stores can exceed the annual demand. This is a natural result of distributed stock. The safety stock required will increase as the square root of the number of distributed stores increases; so managers had better think carefully about setting up local stock points. The increase of stock (by the square root) is based on the assumption that the local stores are part of an integrated, well-controlled inventory system. If this is not the case in practice, then the total stock can go through the roof (literally).

Company strategy

Some inventory controllers do not have complete authority to maximize availability and minimize costs. This may be because top management has imposed some constraints such as an extra stockholding imposed through lack of confidence in the ability of the stock system to provide adequate customer service. It may result from a commitment to customers to keep an agreed quantity locally or consignment stocks at customers' premises.

In the normal course of events the stock control system will provide the desired availability by adjusting the control parameters and should not need overriding by management edict. It is up to inventory management staff to instil confidence in their ability to control stock excesses.

System and control

The time taken to process information can substantially affect the level of stockholding. Consider a stores where an issue is being made. This causes the stock to fall below the ordering trigger level. How long does it take before the supplier gets to know that a new delivery is required? On a conventional order processing system the procedure is:

- Enter issue on system. (Delay one hour?)
- Print out list of items which need ordering. (Next morning?)
- Confirm that items do need ordering. (Rest of the day?)
- Print purchase orders. (Next day?)
- Get authorization for purchase. (When director is next available!)
- Send order to supplier. (Two days by second class mail?)

These processes could easily take a week, or much longer if the individual lines are batched up into a large order. If the supplier is a local stockist, then the item could arrive the same day, but the total supply lead time would be one week (internal) plus one day (external), and the stock level would have to be correspondingly large to compensate for this lead time.

For efficiency use real time procedures.

Slow systems cause extra stock

The solution is to use the computer systems properly, without paperwork, and to use communications technology, (fax or e-mail) for ordering.

If the stock is analysed according to the headings above, the causes of stock excess become obvious. Analysis will show that actual stock is not held in the right balance to meet marketing and financial targets in the most effective way. By establishing a reason for this, a plan can be devised and action taken to ensure the lowest stockholding levels necessary to meet company targets.

Initially, the reason for stockholding has to be determined. It could be to achieve good customer service, cheap purchase, bulk discount, security or a large number of other reasons. The stockholding strategy depends directly on the company objectives and policy. Company policy defines the service which it intends to provide to customers, what investment to make in stock, how orders are placed on manufacturers or other principles on which to base supply plans. This is the first ground rule of stock control – to have a policy worked out by top management to guide the operation. The more professional the stock control operation, the closer the policy needs to be defined so that targets are set against which the activity can be run. This is a generic, broad-brush approach, but useful for finding out where to put most effort in stock management.

Key points

- Consider the objectives and need for inventory.
- Target availability, inventory value and operating costs simultaneously.
- Create profits by avoiding non-essential inventory.
- Determine the causes of stock and use techniques to reduce it.

Customer service

- The key to successful business.
- Focus on customer delight.

Objectives

- Identify the purpose of customer service.
- Define the alternative measurements of availability.
- Analyse the components of administering customer service.
- Discuss the role of forecasting.

Meeting customer requirements

The focus on customer service has gradually changed over the years and suppliers are now becoming really interested in customer service instead of just talking about it. Customer service is a complex subject of its own, but there are two main aspects, namely customer relations and availability of service or items.

Please the customer and measure performance.

Customer relations

Customer relations is about keeping the customer happy. It requires interpersonal skills to ensure that the customers have the correct level of expectation of supply and that they are happy with their purchase so that there is potential for repeat business and wider sales.

This aspect of customer service has gradually emerged because it is the differentiating factor between many companies. The product can be similar from a wide variety of vendors, the vendor which is most successful is the one where the quality of customer relations is best.

Fix the customer. Don't just supply the item.

Customers may have a perception of suppliers which varies greatly from their actual performance. It does not matter whether a supplier is good or bad, what does matter is what the customer thinks about them.

Customers can look for a match between their own style of company and the style of the supplier. If there is a match, then relations are likely to be good; if not, then the co-ordination requires more work. The major factors affecting customer relations are outside the scope of inventory control, in sales and sales order processing. However, there has to be the same customer-oriented approach in stock control, which, like the company, needs to maintain credibility.

The second aspect of customer service – availability – is a key target for stock control. It requires technical management.

Measuring availability

Availability

The entire reason for stockholding is to have items available but, despite the key nature of customer service to the success of the business, some companies do not quantify proper stockholding. They simply take notice of complaints or other oblique assessments of service. To manage professionally there is a need for solid facts.

For each item in stock the risk of stockout can be reduced by increasing the stockholding. The larger the investment in inventory for an item the better the service. This is shown in Figure 2.1 for an individual stock item. The curve shows that it is relatively inexpensive to give a reasonable level of availability, but the value of stock increases very rapidly when we try to achieve very good service. A 100 per cent service is not possible because it needs an infinite amount of stock! However, 99.99 per cent availability is possible, but very expensive and no one can tell the difference from 100 per cent.

Availability: get the facts, then act.

There is obviously a trade off between the investment and the availability of items from the shape of the curve in Figure 2.1. For each item, the stock policy can be adjusted to give the appropriate service.

The overall service to the customers, or other stores users, can be measured from success in meeting demands (on-time availability) on a daily, weekly or monthly basis. It can be calculated by adding the service




Figure 2.1 Availability of a stock item

provided on all the individual items, taking into account the issue frequencies for the total range of items. We can define availability as:

$$\text{Availability} = \frac{\text{Demand satisfied}}{\text{Total demand}}$$

This is the criterion which should be applied to all areas of inventory to monitor how well the investment in inventory is doing.

When the whole range of stock items is considered, the minimum stock levels result from all the items being under control at their optimum service level. Basically, each item should be on the same curve in Figure 2.1. The only way to reduce the stock beyond this balance, while maintaining service, is to improve the total control method. After all, the aim is toward the top left corner of the graph – high service with low stock across the whole of the inventory.

*High stock goes with poor availability.
They are caused by poor control.*

Measurement of availability in practice

Within this definition of availability there are a number of ways by which the service level can be measured, depending upon the type of demand.

For companies supplying single items to individual users as a result of individual telephone or mail orders, then the definition could be:

$$\text{Availability} = \frac{\text{Total number of items supplied}}{\text{Total number of items ordered}}$$

For a department which is sending large export orders, these are often covered by a 'letter of credit' or the whole shipment must be sent in one crate. In this case the service level would be:

$$\text{Availability} = \frac{\text{Total number of complete orders supplied}}{\text{Total number of orders}}$$

For businesses where the demand is not required to be satisfied immediately, care must be taken to measure availability at the time required for dispatch, not the time of receiving the order. This gives the inventory controller the opportunity to fill the demand from elsewhere rather than from stock.

Care must also be taken as to the definition of 'on-time' delivery. Delivery should really be measured as the time the customer receives the stock rather than when it leaves the stores. 'On time' is when the customer needs the stock. Where there is a history of poor delivery the customers are quite likely to request a delivery date well in advance of their real need, in order to have a chance of having the items when they need them.

Measure availability in line with typical customer needs.

The type of customer also affects the measurement of availability. For example, if a customer requires ten and the stock is eight what customer service can be provided? Is the answer:

- 0 per cent because the demand was not completed (non-stock/manufacturing answer)
- 80 per cent because eight out of ten can be supplied (normal answer)
- 100 per cent because the balance can be provided before the customer needs them (for stockholding customers).

Different situations have different priorities. To have an appropriate measure of availability, the service must be defined in one of the above ways

Table 2.1 Pitfalls of measuring availability

Part no.	Demand	Shortage	Service level (%)	No. of orders	Shortage	Service level (%)
A12	1000	10	99.0	5	1	80.0
B25	10	5	50.0	1	1	0
Average			74.5			40.0
Sum	1010	15	98.5	6	2	66.7

– the one which suits most customers. In Table 2.1 only two items are being considered. The demand for item A12 was ten orders totalling 1000 items of which 990 were sent on time and ten items were not available for one order. This means that 80 per cent of the orders (and 99 per cent of the items) were sent on time. The second item is a slow mover and there was only one order (for ten items) and only five were dispatched on time giving a 50 per cent service of parts (and complete failure in terms of complete orders).

How good is the service? Out of 1010 items requested only fifteen were short, so the number of items dispatched was 98.5 per cent of the total. However, in terms of item order fill, deliveries of item A12 were 99 per cent full but B25 was only 50 per cent full, giving an average per line of 74.5 per cent. If the customer is looking for complete orders, then out of the total of six orders two were short, corresponding to a 66.7 per cent success rate. As an order fill per line, this was 80 per cent for A12 and 0 for B25, so the average was 40 per cent. This illustrates that different assessment of levels of service which can often occur between customers and suppliers. In this example the supplier is claiming 98.5 per cent delivery on time, while the customer is measuring it at 40 per cent! It is important to agree a common goal between suppliers and customers, and to establish what the customers need in terms of availability. The appropriate measure can then be used for monitoring.

Availability policy

Customer service can be structured or focused on a particular group of customers or a market sector using alternative availability policies. Options which could be used are:

- the same availability across all the products
- minimizing the total cost of service

- concentrating on the most valuable customers
- enhancing the service on the most sensitive products
- greatest availability on the most profitable products, or
- better service on the major turnover items, reduced service on slow movers.

These alternatives will provide different availabilities for different products, with a management decision setting availability levels.

Whichever alternative is chosen, proper control of stock requires the inventory manager to regularly monitor availability.

A better measure takes account of 'lateness'.

Back-order measurement

The measurement of availability discussed in the measurement of availability in practice section should always be used, but a further measurement is necessary for sophisticated stock control.

Once an item has missed its delivery date, it is out of the calculation. This introduces a further factor into the equation: the priority given to satisfying back orders. Consider a case where there are two demands for an item, one which is past the due date and one which is to be delivered on time. In such circumstances, the availability target calculation requires us to fulfil the current order, not the old one, since the current one improves the availability factor, but the other has already failed. In practice there is normally a need to give precedence to customers whose delivery is late. Therefore a more rigorous measure of availability would be to measure how many orders are late and how late they are.

An example of how this can be approached is developed from Table 2.2 (and Figure 2.2) which shows the dispatch position for two companies, A

Table 2.2 Comparison of dispatches for companies A and B (items dispatched in week 37: analysis of lateness on orders dispatched)

	Company A	Company B
Delivered on time	10	30
1 week late	60	15
1 month late	20	40
3 months late	10	15
Total dispatches	100	100




Figure 2.2 Comparison of dispatches for companies A and B

and B. Which company is providing the better service? From our service-level (level of fill – LOF) approach company B is the better because it gives 30 per cent LOF against 10 per cent LOF for A. However, there is a better chance of getting the items quickly from company A. Perhaps A is better after all.

A measure of customer service which takes into account the overall performance, not just the ‘ex stock’ part, can be developed by using a weighted average. In Table 2.3 (and Figure 2.3) the arrears (or ‘back orders’) have been analysed by categories of lateness. By multiplying the number of orders outstanding by the weeks of lateness, a back-order week figure has been calculated. This gives an overall figure of merit of 45 arrear weeks, corresponding to an average for arrears of 2.1 weeks per late item. By monitoring weeks of lateness against a target level, inventory controllers can assess the level of service achieved. Obviously, the best practice is to ensure that the bulk of demands are satisfied on time because those few items which are long overdue give a negative contribution to the figure of

Table 2.3 Distribution of arrears

Time overdue (T weeks)	No. of items (N items)	Weighted average (T × N)
1	8	
2	6	
3	4	
4	2	
5	1	
Total	21	45




Figure 2.3 Distribution of arrears

merit. If the back orders are satisfied rapidly, the level of service can be measured in units of days or hours rather than weeks.

In practice the order backlog viewed on the system may not be all true demand. It often consists of:

- real arrears which are customer orders slightly overdue and a real priority
- virtual arrears which are often long-standing customer orders without any urgency to them; these virtual arrears are orders which the customer does not really want, or where they have bought an alternative and not cancelled the order for the original item.

Outstanding orders need regular review to ensure that virtual orders are either rescheduled or removed. If virtual arrears build up they have a distorting effect on the back-order measurement.

Verify dates required before calculating arrears.

The extent of arrears can be measured in several ways. The number of items in arrears is the simplest measure, and the aim is to focus the control on the important items. Measurement in terms of value is not usually recommended since the aim is to provide customer service. Routine assessment of arrears should be carried out in a similar manner to the monitoring of availability.

Demand management

Order processing

Entering the demand and making sure the customer is happy.

Stock should, of course, fill customer demand. In practice this includes both having stock available and dispatching it. The physical handling of the stock is a stores operation's responsibility, while in many companies the actual order processing forms part of the stock control function. Demand management consists of a series of steps:

- 1 Order receipt.
- 2 Order processing.
- 3 Estimating delivery times to customers.
- 4 Production of dispatch information.
- 5 Customer feedback.

Order receipt

This is the initial compilation of orders into a form which can be used to dispatch the goods at the right time to the customer. This information must be:

- complete
- specific
- accurate.

The information can arrive by post, telephone, electronically or by word of mouth. A first requirement is that the company has a good item identification system which enables sufficiently unique identification to satisfy a discerning customer. This has to be formalized into data for processing.

Stock control requires an identification code for each item, ideally provided by the customer. To assist this, some companies provide customers with catalogues and order forms to enable customers to specify precisely. It used to be customary for counter-sales customers of all types not to know code numbers, and to rely on stores personnel to remember them. This situation is no longer acceptable – if customers want the correct item, then they should expect to specify their requirement properly. An order form not only enables the correct data to be collected, it also presents

it in a standard format, which can be arranged to match computer input or stores picking procedures.

Order processing

This is the entry of the order into the system. A second requirement is that the processing of orders translates, reliably, speedily and without creating errors, the customer order into company information. This process has to be fast, simple and, above all, accurate. For computer systems this should be done as a one-stage process with the source (phone, fax, letter, etc.) being input directly into the system. This avoids errors and delays. Where printed customer orders are required, these should be printed out from the data input to the system.

Ensure that the data input is accurate.

For non-computer applications the source document (or a copy of it) should be used for stock picking. With technology, electronic data interchange (EDI) can create a demand in the system without any work. This is good as long as the demand is easy to fulfil. Stock controllers need the opportunity to vet orders, which means that all EDI orders should be formally accepted for dispatch. Application of the ‘manage by exception’ principle suggests that most orders are easy to satisfy and can be accepted without question – automatically. It is only those orders for fast delivery, for large quantities or for non-standard items which need control. Criteria can be set up to identify these so that the majority, the other standard orders, can be accepted without question. The stock control parameters for establishing these criteria are discussed later.

Estimating delivery times to customers

Estimating delivery times is the activity which causes a high level of customer confidence to be gained or lost. Promises have to reflect the real stock and replenishment situation or else customer service will be poor. (See the following section ‘Estimating delivery times.’)

Production of despatch information

The despatch information will consist of two stages:

- picking and delivery note production
- customer invoicing.

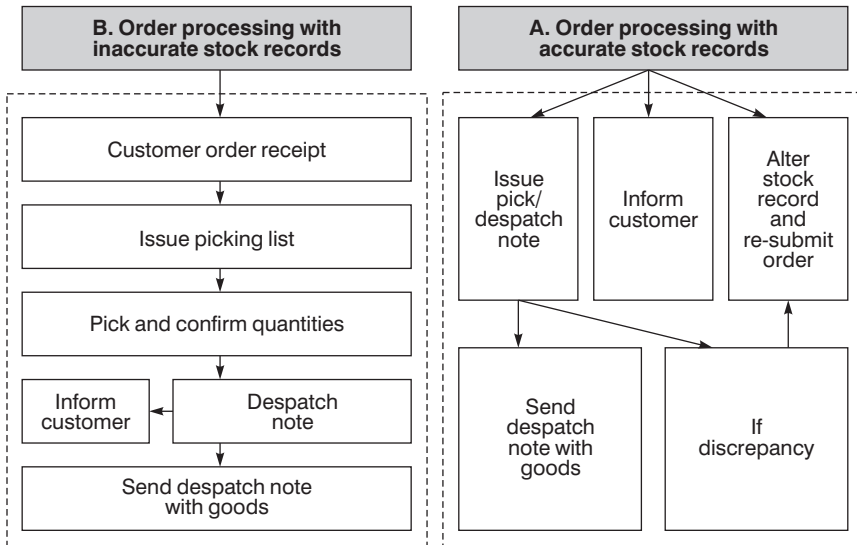


Figure 2.4 Stock item order processing

These stages should be carried out sequentially. The financial (invoicing) systems should not be allowed to interfere with the stock movement (dispatch) system, as invoicing is not part of the stock control system. Two alternative paperwork processes can be used. A delivery note can be prepared as soon as the order is processed, as long as the stock records are accurate as shown in Figure 2.4(B). Where records are somewhat suspect, a safer procedure is the one shown in Figure 2.4(A). This usually causes extra paperwork.

Customer feedback

Customer feedback to report on progress and to keep the customer feeling involved in the supply is very important from a customer relations standpoint (see previous section, ‘Meeting customer requirements’). Customers often require acknowledgement of order and confirmation of price as well as a dispatch note and invoice. This confirmation should form part of the normal information processing procedures.

Consuming forecast demand

In almost all situations business works to a forecast until the real orders arrive. However, real orders rarely match the forecast and the question is raised as to whether an actual order received is really part of the forecast.

Consuming the forecast = take out a forecast demand and replace with the customer order.

The ideal situation is when a company accurately forecasts demand and makes that quantity available. As customer orders arrive, available quantity is allocated and dispatched. When the quantity for the current period has all been allocated, then customers are provided with items in the next period. This situation is shown in Figure 2.5. Here the forecast in periods 1 and 2 (shown as the curve) is fully allocated and there are orders for periods 3 and 4 which have already been received. The demand in periods 5 and 6 is all forecasts as no orders have been received which need satisfying so far ahead. Orders are constantly being received and product is allocated to these customers. This consumes the forecast quantity and builds up the allocations. This approach ensures that a company does not allocate more than has been provisioned through the forecast. In some fast-turnaround businesses with period length of a day or less, this situation is acceptable and almost inevitable. In other businesses more flexibility is required and there has to be:

- safety stock to meet high demand, or
- extra supply sources available to meet peaks, or
- resources reserved for emergency supply (e.g. allocation of only 90 per cent of capacity in manufacture leaving 10 per cent for last minute orders).




Figure 2.5 Consuming the forecast

Where a company agrees to fill every order from customers at a time to suit the customer, then it is important to forecast carefully, otherwise delivery achievement will be poor, much management time will be spent chasing problems, panic purchases will be made and extra deliveries will be required, causing costs to be high.

The differences between forecast and actual demand are measures of the quality of the forecast (see Chapter 6). Actions arising from order processing, besides accepting the customer's request, could be:

- arrange increase or decrease in stock levels
- bring forward or delay supply
- offer earlier or later delivery date to customer
- reject order or offer alternative item
- provide part shipment.

Limit order book to achievable capacity.

An example

A supplier of fasteners has trouble with widely fluctuating demand levels on some products. This is because the customer thinks 'I need a box of fixings for this job' and orders enough to do perhaps two months' work. Consequently the supplier gets hit by two-months' demand and then nothing. This is acceptable if there are many customers, but for the slower moving fasteners demand can be very erratic. By understanding the supply problem, the customer can be encouraged to order enough to start the job off with sufficient supplies, while the supplier, who understands the user's situation, can provide the range of fasteners but can get away with much less stock. Of course, the situation must be carefully monitored and the balance must arrive before the initial stock has been used up.

Large quantity orders can be filtered out using simple criteria (see Chapter 6) so that real demand priorities can be fulfilled. Demand can thus be managed through the order processing activity.

Estimating delivery times

For supply to order items, and for stockouts, an estimate of the delivery date is usually required. Sales are often not aware of the status of the supply

and the order book, and so are badly placed to give more than a blanket estimate of delivery. The stock control or purchasing departments do have the information and can provide dates which are based on the actual current situation.

It is the responsibility of inventory management to provide delivery and availability information in support of customers' orders. The time taken to provide an item to a customer depends on the balance between stock currently held and on order and the quantity already allocated to customer requirements (see Chapter 12). If the order being accepted is allowed to use items already allocated to other customers, then the order can be fulfilled sooner. In general the priority sequence is 'first come first served', but there are often cases where certain customers gain priority. Sales may consider that promising the items regardless of ability to deliver will secure orders and keep inventory control people on their toes. This can lead to a mistrust of the optimistic delivery promises made and poorer customer service. The solution to this situation is for inventory management to provide customer delivery dates using priority rules specified in advance by the marketing department. Where there are exceptional circumstances, then priorities can be altered within the total availability profile. The request to 'ring the supplier and get a faster delivery this once' should be dismissed on all but the odd occasion, since it usually results from sales people not being able to negotiate successfully.

Monitoring delivery dates can become complex. Dates which can be considered as appropriate are:

- 1 When the customer placed the order.
- 2 When the customer requested delivery.
- 3 When the customer needed the item (usage date).
- 4 When delivery was originally promised.
- 5 When delivery is now promised.
- 6 When the item is dispatched.
- 7 When the item is received by the customer.

It is also useful to know:

- 8 Who changed the delivery date.
- 9 How many times has it been changed.
- 10 When the delivery date was changed.

Decide carefully what is 'on-time' delivery.

The importance of these dates varies from one market to another and one company to another. There are so many dates and potential monitors, that companies have to choose which ones are important in their own circumstances. Providing too many management controls is confusing and dilutes the important ones. The prime controls would be delivery on time measured as '7-2' and '7-4'.

The difference between the customers' request for delivery and expected delivery dates (2-1) is their perception of lead time. This may be a general indication of market expectations. The difference between the customer order date and the date they received the goods (7-1) is the actual customer lead time. Unfortunately, most companies measure their dispatch dates (6) rather than the customer receipt dates, so there could be a difference between the two views of delivery performance. Many carriers now have sophisticated tracking systems that have the potential to provide suppliers with customer receipt information to use in their control systems.

Customers may leave an interval between their requested supply date (2) and the time they actually need the items (3), to compensate for delivery problems. If information is available to monitor this, it can be used to improve collaboration between supplier and customer. The monitoring of delivery performance, as with all other inventory monitoring systems, exists to provide the best factual information for management and control. Monitoring systems should therefore differentiate between when the customer wanted the goods (2), when the delivery was originally promised (4) and when it is now promised (5). It is useful to measure credibility (7-4), to ensure that the supply systems are working correctly. Measuring performance against current delivery promise (6-5) can fool a company into thinking falsely that it is performing well. The current delivery promise may bear no relation to when the customer requires the item, and may be the latest in a succession of promises as the delivery date slips back. The estimated delivery date can often be changed so that 100 per cent performance is achieved. Therefore, monitoring success against current delivery promise is not likely to be a reliable measure of effectiveness. Monitoring of slippage – the number of times the delivery promise has changed (8) – can also be revealing in showing whether the company is providing customers with good information and service. However, the reason for late delivery against the original plan could be simply that the customer does not require the goods until later. It is therefore useful to examine (9) what proportion of the changes are due to poor supply and how many are due to customers changing their mind.

Key points

- Find out how to delight the customer – go for added value.
- Measuring availability is essential for success (more important than financial reasons).
- Choose an availability measurement that suits the typical customer.
- Ensure that demand-processing procedures are effective.
- Keep the customer informed and involved.

Managing the inventory

- The use of Pareto analysis to save time and give control
- ABC for managing inventory
- Stock cover provides a monitoring tool
- Minimizing effort and inventory value
- A guide to reducing inventory levels

Using Pareto analysis for control

Applying effective control

In the stores there are a wide variety of items, with a stock record for each. Some have high value and others are very cheap. The high-value items are normally controlled tightly, whereas the low-value items are not treated as carefully and are issued in bulk in approximate quantities. Most effort should be put into managing the items which are most important for achieving the inventory targets. In inventory control the best results are gained by organizing effort correctly. There is not sufficient management time to maintain detailed control of all the individual items. If the immediate aim is to reduce stockholding costs then studying the stock of low-value items is unlikely to be the best place to start unless the sales volume is very large. If service is the aim, then attention to a few fast-moving lines often provides the bulk of the improvement required.

This simple principle is embodied in Pareto's Law, which is illustrated by the curve in Figure 3.1. It is also called the 80/20 rule because 80 per cent of the effect is provided by 20 per cent of the cause. In Figure 3.1 80 per cent of stock value is caused by 20 per cent of stock lines. The principle can be applied to many different areas of activity: 80 per cent of the purchased items come from 20 per cent of the suppliers, and 20 per cent of sales lines give 80 per cent of turnover. For a warehouse, 80 per cent of the space is occupied by 20 per cent of the lines. This simple fact makes it obvious which lines should be received little and often.




Figure 3.1 Pareto curve

Pareto analysis is the technique which forms the basis of inventory control thinking and is an important management principle which can be applied to minimize effort and to obtain best results. It can also be applied for time management, credit control and many other areas of control. To gain best control, effort has to be directed to the most important areas. The Pareto curve (see Figure 3.1) is often called the '80/20 rule', but the values can be read off at any convenient point. For example, the graph shows that 50 per cent of the product lines account for 97 per cent of the sales (or that the other 50 per cent only provide 3 per cent of the sales – a worrying thought). The shape of the Pareto curve arises from the range of volumes and values combined in a statistical distribution. The shape of the curve does not always give exactly an 80/20 relationship, but this does not affect the principles of applying Pareto analysis to inventory management.

Stores carry a wide variety of items, with a stock record for each. Some have high value and others are very cheap. The high-value items are normally controlled tightly, whereas the low-value items are not treated as carefully and are issued in bulk in approximate quantities. Most effort should be put into managing items which are most important for achieving

inventory targets. In inventory control the best results are gained by organizing effort correctly. There is not sufficient management time to maintain detailed control of all individual items. If the immediate aim is to reduce stockholding costs then studying the stock of low-value items is unlikely to be the best place to start unless the sales volume is very large. If service is the aim, then attention to a few fast-moving lines often provides the bulk of the improvement required.

Use Pareto analysis to save time and get results.

Stores contain items ranging from main products to washers and labels, with a stock record for each. High stock value items need to be closely controlled, whereas minor items need not be treated as carefully. To control the resources of the company most effectively our effort and controls should be biased towards high cost areas. Pareto analysis formalizes our efforts to do this. It states that the majority of the effect is produced by a small proportion of the cause. (80 per cent of the effect is due to 20 per cent of the cause.)

The application of this to a stores stock control means that 80 per cent of the total stock value is made up from 20 per cent of the total stock items as stated. The other 80 per cent of stock items contribute only 20 per cent to the total inventory value (shown in Figure 3.1.) In a stock reduction exercise the majority of our cost saving will be gained by decreasing stocks of those few major items.

Eighty per cent of stock is in 20 per cent lines – so reduce the high stock values.

Example

Consider a stock of 12 000 types of items in store. Pareto's Law shows that for a stock value of £800 000 we find that 2400 items account for £640 000 of inventory. The remaining 9600 items are worth only £160 000. Therefore, by concentrating on the 2400, control over the total value will be tight. If 2400 is rather too many to review individually, then the Pareto curve (Figure 3.1) shows that 5 per cent of items account for 55 per cent of cost so 600 items contribute £442 000 to the total stock costs. Again by working on these 600 items carefully the overall stock value can be controlled or decreased.

ABC analysis

Pareto analysis by the current stock level is good for reducing stock levels, but a more consistent classification is required when focusing on the management of inventory. The current stock does not necessarily show which items are important for the business. In fact there may be some important items where the current level of stock is low because Stores are awaiting an impending delivery. On the other hand some items may have a high stock value simply because no one is buying these. It is therefore usual to rank the items according to the annual turnover. The annual turnover is given by

$$\text{Annual usage} \times \text{Unit cost}$$

It is not too important whether the unit cost is the standard cost, latest cost or an average as long as it is consistent across all the items. Annual usage has to be adjusted in the case of new or obsolescent items to reflect the future expected demand rate rather than the historical one.

Pareto analysis of this data shows that 80 per cent of the value of demand is for 20 per cent of the moving items. (There is often a number of items in stock for which the demand is zero, and therefore these items are not included in this turnover analysis.) For some businesses the 80/20 rule is not obeyed exactly, but the use of Pareto analysis is important for all inventory.

To use Pareto analysis properly requires the classification of stock by issue value, and the simplest way is to use ABC classes. These can be defined as:

- A = 10 per cent of stock numbers, giving 65 per cent of turnover.
- B = 20 per cent of stock numbers, giving 25 per cent of turnover.
- C = 70 per cent of stock numbers, giving 10 per cent of turnover.

*ABC analysis is an excellent technique
for achieving objectives.*

This is illustrated in Figure 3.2.

It is important to ensure that ABC analysis is based on turnover, but less important that the exact percentages are adhered to. In some instances a further classification (D) is useful to include a large number of very low turnover items. This enables the number of stock lines included in the A, B and C classes to be reduced to manageable numbers.




Figure 3.2 ABC analysis

The purpose of ABC analysis is not to provide different types of service but to provide service with the least amount of cost and effort. Different systems of control are used for the three categories of stock. As the A items carry the most value, accurate systems are required to control them. On the other hand, the C items are low turnover value but form the bulk of the inventory. For these the main requirement is to ensure that stock is available to meet demand.

A class – there aren't many, so control tightly.

The control requirements for each category are shown in Table 3.1. Category A items have a disproportionate amount of time and effort used on them and have to be controlled tightly using systems in conjunction with market expertise and product knowledge to maintain stock at the lowest appropriate level.

B class – let the system manage these.

Table 3.1 ABC inventory control

Characteristics	Policy	Method
A items Few items Most of turnover	Tight control Personal supervision Communication JIT approach – balanced safety stock	Frequent monitoring Accurate records Sophisticated forecasting Service-level policy
B items Important items Significant turnover	Lean stock policy Use classic stock control Fast appraisal method Manage by exception	Rely on sophisticated method Calculated safety stocks Limit order value Computerized management and exception reporting
C items Many items Low turnover value (Few movements or low value items)	Minimum supervision Supply to order where possible Large orders Zero or high safety stock policy	Simple system Avoid stockouts and excess Infrequent ordering Automatic system

For category B items computerized techniques are most appropriate. The number of items involved and the lower values make it a waste of time to use specialist skill which could be working on category A. The computer can maintain control through statistics and deal with the complex calculations using the computer forecasting models which are most important for B items. The use of management by exception is also important for B class.

The minor sales items, category C, should be controlled by a simple system which enables supply to be obtained with a minimum of administration. However, the control system for C items must be reliable and not result in stockouts or large excesses. An investment in extra stock of C class items is inexpensive but can greatly simplify the problems of controlling large numbers of stock lines. This is an appropriate policy for the faster moving C class items. For the very slow moving, higher value C class items a purchase to order policy should be adopted where possible, or if there is only one customer, they can hold the stock themselves and be responsible for reordering as required.


*C class – don't take risks, be lazy.
High but not excessive inventory.*

The most effective stock control systems are based upon ABC analysis combined in a common-sense manner with the other techniques we shall be discussing later.

ABC analysis is the basis for the total control of stock. It is also used as the basis for perpetual inventory stores control where annual stocktaking is avoided by routine counting of a few stock parts each week.

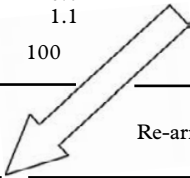
An example of Pareto analysis in action is given in Tables 3.2–3.4 and Figure 3.3. Table 3.2 shows a number of different stock items, their unit cost and annual usage in terms of quantity of value.

Table 3.2 Example of Pareto analysis




Item	Annual usage (units)	Unit cost (£)	Annual turnover (£)	Annual turnover (%)	Rank
A12	21	7	147	2.1	5
B23	105	11	1155	16.3	2
C34	2	15	30	0.4	10
D45	50	5	250	3.5	4
E56	9	14	126	1.8	6
F67	394	12	4728	66.6	1
G78	5	8	40	0.6	9
H89	500	1	500	7.0	3
I90	11	4	44	0.6	8
J01	3	25	75	1.1	7
Total			7095	100	

Table 3.3 Classification by usage value



Item	Annual usage (units)	Unit cost (£)	Annual turnover (£)	Annual turnover (%)	Rank	Class	Cumulative percentage
F67	394	12	4728	66.6	1	A	66.6
B23	105	11	1155	16.3	2	B	82.9
H89	500	1	500	7.0	3	B	89.9
D45	50	5	250	3.5	4	C	93.4
A12	21	7	147	2.1	5	C	95.5
E56	9	14	126	1.8	6	C	97.3
J01	3	25	75	1.1	7	C	98.4
I90	11	4	44	0.6	8	C	99.0
G78	5	8	40	0.6	9	C	99.6
C34	2	15	30	0.4	10	C	100
Total			7095	100			



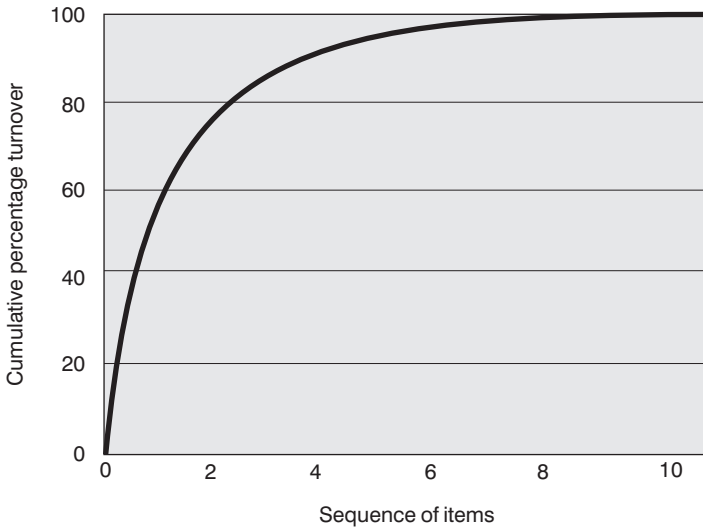


Figure 3.3 The cumulative Pareto curve

Items are then ranked in order of size of annual turnover value. This is displayed in Table 3.3 with the items in descending order. The cumulative annual usage value and percentage of turnover are also calculated. Items are then classified into A, B and C by assessing the number of items in each category and their percentage of total turnover value.

In the Table 3.4, classes A, B and C are compared as percentages of number of items and total value. The table shows that the seven items in class C have a turnover of £712, whereas the only item in class A has a turnover of £4764.

Table 3.4 Summary of ABC analysis

Classification	Percentage of items ¹	Percentage of value	Value per class
A	10.0	66.6	4728
B	20.0	23.3	1655
C	70.0	10.1	712
Total	100.0	100.0	7095

Note: ¹ABC analysis is carried out only for items with usage.

Saving time

Applying the Pareto principle is a way of balancing inventory, stock availability and critical resource spent on each item. How the law can be applied depends on what the critical resource is considered to be. The critical resource for all inventory controllers is time, because of the large amount of information required for tight control and the wide variety and quantity of stock held. The Pareto principle shows that 80 per cent of the time is spent doing 20 per cent of the jobs and a significant time saving can be made if a small reduction can be achieved in these jobs. They may be very frequent short jobs (such as keying stores issues into a computer) or more infrequent, long-winded jobs (such as writing a major management report).

Practical considerations in using ABC analysis

The ABC classification is a simple tool to enable the stock manager to control a large number of items in a limited amount of time. This simple approach is one of the most powerful tools employed to reduce stock value and to decrease the workload of busy inventory managers and purchasers. Experience with this technique over a large number of companies has suggested that there are some pitfalls and some practical ways to circumvent them. The situations which arise are typically:

- 1 Too many A class items.
- 2 Large numbers of lines (D class).
- 3 Non-moving lines (O class).
- 4 Fixed stock level items (F class).
- 5 Non-stock items.

Use ABC to save time.

Number of A class items

A class items are supposed to be reviewed on a daily basis, or a weekly basis at least. To be practical, there has to be an upper limit of, say, 300 A class items per inventory planner. Even where the planner manages many thousands of items, the number of A class has to be kept small.

D class items

For normal inventories of 3000 lines, the ABC classification works well. Where there are over 10 000 part numbers there should be a modifica-

tion to the classification system so that the vast bulk of low-value turnover items are dumped into a further class, D. The D class contain the lowest turnover lines, say 50 per cent of the active item numbers which contribute only 2–3 per cent of the total turnover. The ABCD classification now is as follows:

- A class: 5 per cent of moving lines (300 per planner maximum).
- B class: 10 per cent of moving lines.
- C class: 35 per cent of moving lines.
- D class: 50 per cent of moving lines.

The classification can alternatively be carried out by turnover value, for example:

- A class: 45 per cent of turnover.
- B class: 30 per cent of turnover.
- C class: 22 per cent of turnover.
- D class: 3 per cent of turnover.

These classifications can be varied to suit the exact shape of the Pareto distribution for the inventory to be managed. The principles of use of the technique do not change even if the distribution is not 80/20.

Non-movers

So far in the discussion the moving inventory has been classified, and the non-movers have been ignored. Companies do sometimes have a need to keep non-moving items, or items where the movement is so slow that they appear as non-movers in the recording systems. These items will probably be subject to the stock cleansing discussed in the section on turnover of stock. In the interim they should be provided with a separate classification. Normally they are given class 'O' or 'X' which signifies that they should not be ordered again.

Use ABCOFZ as a practical control method.

Fixed classification

As Pareto analysis is to be used for ordering, there are a few items where the stock level should not respond to usage rate. (Unfortunately in many older stock management systems, the stock control parameters are all like that!)

For example, the employment of two new maintenance fitters in a factory requires them to be kitted out with protective clothing, tools, tool boxes and a variety of items. The use of these items is likely to fall rather than rise as a result of that action since there is less likelihood of further recruitment. By putting these items in a separate classification (say F), they can be segregated and identified. The system can then identify that the stock level parameters from classification F are not updated.

Non-stock items

The decision as to which items are in the stock range is an arbitrary one and depends on the particular inventory policy and market conditions. Customers do not usually consider their requirements as 'stock' or 'non-stock' items and the difference to them is only in the lead time provided by the supplier. In fact, non-stock items are continuously being taken into the stock range and stock items being deleted. For this reason it is useful to include non-stock items in the ABC analysis. This can either be done by including them within the ABC classes or, more commonly, to have a separate classification (say, Z) for non-stock items. This separate classification then has the feature that no stock is ordered from suppliers unless the stock cover is negative (i.e. there is a customer requirement but no stock). It is very useful to include all goods in the inventory management system as it provides unified records identification and control over all goods, and makes management and analysis easy.

Stock cover

Turnover of stock

The current stock levels in the various stores throughout the company may not all be at the ideal levels, as we have seen. The purpose of controlling the inventory is to drive the stocks toward their proper level which is determined by the characteristics of supply and demand patterns. The major factors are:

- supply lead time
- average demand rate
- variability of demand
- supply frequency
- customer delivery time allowed.

There are also practical considerations such as:

- reliability of the supplier
- criticality of the item
- availability of item from other sources.

The concept of 'balance' is most important in ensuring that the maximum service is produced from a minimum of stockholding cost.

The best level of service will be provided if there is an equal chance of all items being available for the customer. High stocks of one item and no stocks of another will reduce the overall availability and increase the inventory cost. The percentage availability target (see Chapter 2) should therefore be the same for all items in the stock range unless there are specific marketing reasons for favouring some customers or products.

If the cost of managing the inventory is also considered, then the value of items being ordered and controlled becomes important. Valuable inventory management resources should be confined to the most cost-effective jobs. This means that time should be spent on high turnover value items and not wasted on items whose value is insignificant. The best balance of inventory leads to an optimization of costs and service over the full range of stock lines within the time available.

The inventory performance of each item can be monitored using a figure of merit for stock balance. This is the 'stock cover', which is defined as:

$$\text{Stock cover} = \frac{\text{Current stock} \times 52}{\text{Forecast annual usage}}$$

This gives the result in terms of weeks on hand. The same answer could be calculated for a month's predicted usage and multiplying the stock by 4.2 instead of 52. This is more convenient sometimes, but should only be used when demand is consistent. It is also convenient to use historical average usage in this equation. (See Chapters 10 and 11.)

A sample of stock items is shown in Table 3.5. Which of these items requires attention first?

Stock cover gives an insight into the priority for action. It is not an infallible guide, but it does indicate where review is required. The first instinct would be to look at 1P4 and 1P8 in Table 3.5, where the stock cover is small. However, there may not be a problem with these since a delivery

Table 3.5 Stock cover

Item code	Stock (units)	Annual usage (units)	Stock cover (weeks)
1P1	250	2000	6.50
1P2	700	1625	22.40
1P3	500	400	65.00
1P4	15	1000	0.78
1P5	20	25	41.60
1P6	40	250	8.32
1P7	500	200	130.00
1P8	8	400	1.04
1P9	6	40	7.80
1P10	65	20	169.00

may be arriving. At the other end of the scale 1P3, 1P7 and 1P10 have over a year’s worth of stock, so means of reducing this level will have to be found. Stock cover shows whether the stock is ‘in the right ballpark’: if it is not correct, attention to the items which are obviously well outside acceptable stock levels will keep the shape of the inventory reasonable. Stock cover is used by inventory controllers because it is easily understood in terms of usage rates and lead times. Stock cover is the time in which the stock will run out at average usage rate.

Use stock cover for analysis but not control.

As well as being a crude analysis tool for each stock item, stock cover is also an important tool for measuring the total inventory. Overall stock cover is calculated from the total value of stock divided by the annual issue value and multiplied by 52 to change the answer to weeks. In many distribution stores the answer is between one and eight weeks. In fast-moving consumer goods, one to eight days is more appropriate. Financial managers are often more interested in the use of funds, and therefore measure the effectiveness of inventory management using ‘stock turnover’ or ‘stockturn’. This is just the reciprocal of the stock cover, taken on a value basis for the complete stockholding.

$$\text{Stock turnover} = \frac{\text{Value of annual usage}}{\text{Value of stock}}$$

This calculation gives the number of times the stock would be used up per year.

Example of stockturn calculation

Value of stock in the stores is £150 000

Issues for the last twelve months amount to £900 000

Stockturn is therefore $900\,000 \div 150\,000 = 6$.

This means that the stock value would be used up completely six times per year so that the stock cover for the total stock will be two months, or by the stock cover calculation as

$$\text{Stock cover} = 150\,000 \times 52/900\,000 = 8.67 \text{ weeks.}$$

Stockturn is based on historical data and is used for financial reporting. Stock cover is an inventory management tool for planning stockholding and can be based on known data and the forecast usage rate, so that the stock will meet the expected demand for the item. When the stock level is being assessed for accounting purposes, the ratio uses the historical usage rate, which enables a conservative view to be taken of the stock level. Although this sometimes leads to a divergence of views on the necessary stockholding, the assessed future demand should always be used when controlling stock or placing orders.

Setting stock targets

For good stock balance the stock cover of all the items should have similar value. In practice differences in the variability of usage and the order cycles leads to a range of acceptable values for the items. Stock cover should not be used for working out reorder levels – there are proper accurate ways of doing this (see Chapter 7). Stock cover ratios can be used to calculate the broad ranges of weeks' cover which would be needed for inventory items. For instance it is unlikely that more than twenty-six weeks' worth of stock is planned for any inventory item, therefore a figure of more than twenty-six could be the boundary between 'OK.' and 'needs attention'.

As inventory control should be tightest for the A class items, these are the ones which can be controlled down to lower stock cover figures (leading to

the paradoxical situation of holding lowest stocks of the best sellers!) whereas extra stock of the minor C class items adds little to stock value and significantly reduces the work of controlling.

In practice the stocks could have control limits to avoid extremes of inventory, and an allowable stock cover range can therefore be set by the ABC inventory classes in a ratio which is theoretically 1:3:7. An illustration of the acceptable ranges for category A, B and C items is shown as:

- A class: between one and four weeks.
- B class: between two and eight weeks.
- C class: between three and twenty weeks.

A stockturn ratio (weeks of stock) for all the items in the category should lie in the ranges shown (see Figure 3.4.)




Figure 3.4 Ranges of stock cover

In theory the distribution of the population of items within a class is shown in Figure 3.4. Also shown in this figure is what the curve looks like in practice. (Note that the horizontal scale is logarithmic.) The theoretical curve allows for the variability of demand, the usage of the safety stock for some items and some slow movers, causing some of the items in a class to fall naturally outside the expected limits. The actual population curve when plotted against cover has an entirely different shape. Some items have high stock cover ratios, and since the stock value has to be limited, some of the other items have very low stocks to compensate. The shape of this curve is caused by the response systems initiated by low stock levels and the less effective action which normally is seen when the stock levels for some of the




Figure 3.5 Expected and actual stock profiles

items rise over the maximum. The profiles of stock turn ratio in Figure 3.5 are typical, even where inventory management practices are good.

The weeks of stock cover shown in Figure 3.4 indicate how long it would take at best to reduce stock levels. If the stock reduction is to be made through C class items, the time taken to reduce the levels will be in excess of twenty weeks. For A class items the whole process should be completed within four weeks.

Pareto stock balance

Pareto analysis enables us to take the discussion of stock balance a step further. Each Pareto class, calculated according to annual turnover, can have its own target months of stock range, as shown in Figure 3.4. The use of these two techniques (ABC analysis and stock cover) together is fundamental in forming the basis of the stockholding policy. Another illustration of the power of the Pareto rule is shown when the overall value of stock has to be reduced to meet a new target. Say the inventory investment currently amounts to six weeks' of stock and this needs reducing by a minimum of 10 per cent, then ABC analysis of the inventory items would follow the pattern given in Figure 3.5.

In Table 3.6 the A, B and C classes have average stockturns of 2.0, 5.0 and 11.5 weeks, and therefore constitute a well controlled inventory. If these amounts of stock cover appear to be over generous, then the substitution of the word 'day' for the word 'week' in the heading will suit all but those involved in JIT supply (see Chapter 4). The classes contribute 65 per cent, 25 per cent and 10 per cent to the inventory value, respectively. This means that the total stock values arising from A, B and C classes are equal.

There are several options for decreasing the stock value. One way is to cut A items down to an average of 1.5 weeks of stock. This new situation is

Table 3.6 ABC stock cover

Class	Moving lines (%)	Stock cover (weeks)	Turnover (%)	Weeks of value
A	0.1	2.0	0.65	1.30
B	0.2	5.0	0.25	1.25
C	0.7	11.5	0.10	1.15
Total				3.70

Table 3.7 Options for reducing stock value

Class	Moving lines (%)	Stock cover (weeks)	Turnover (%)	Weeks of value	Reduction in weeks of value	Number of lines involved
A	0.1	1.5	0.65	0.975	0.325	100
B	0.2	5.0	0.25	1.250		
C	0.7	8.0	0.10	0.800	0.350	700
Total				3.025		

shown in the A class calculation in Table 3.7. The net result is a stock reduction of 0.33 weeks.

An alternative way to achieve the same result is to reduce the stock of C class parts down to 8 weeks (from 11.5 weeks). The new contribution to stock cover is $8 \times 0.10 (= 0.8)$ weeks of value, a reduction of 0.35 weeks.

Tighter control of A items is the better alternative because:

- the change required is less (0.5 week as opposed to 3.5 weeks' cover)
- there are fewer lines involved, so less disruption of planning
- less work will be involved – 100 stock items instead of 700 (last column of Table 3.7)
- the effect will be more rapid.

This last point can be proved by looking at the stock cover for each of the ranges in Figure 3.5. The average taken for stock cover for A, B, and C, are 2.0, 5.0 and 11.5 weeks respectively. If the classes were perfectly balanced, and the supply was stopped, then the A class would run out after 2 weeks, the B Class after 5 weeks and the C Class after 11.5 weeks. Stock reduction through natural usage will take at least these respective times. As stocks are not often perfectly balanced, the true time to reduce the stock is normally over twice this long.

*Using ABC and stock cover together
saves time and inventory.*

Purchase order patterns based on Pareto analysis have been devised to save stock value and management time. This is well illustrated by the ordering process shown in Table 3.8. Consider an inventory of 1000 different lines. If each line is purchased each month from the suppliers, then there are 1000 lines ordered each month, or 12 000 per year, and a stock cover of just over two weeks plus safety stock. (Stock goes up by 4.2 weeks' worth on receipt and then reduces to nil, so the average is 2.1 weeks.) Safety stock can be omitted from the discussion since it is taken to be independent of the order pattern. The order pattern is A weekly, B fortnightly.

Table 3.8 Number of purchase orders with ABC inventories

Class	Moving lines (%)	Stock cover (weeks)	Turnover (%)	Weeks of value	Orders per year ¹
A	10	1	65	0.33	5 200
B	20	4	25	0.50	2 600
C	70	10	10	0.50	3 640
Total	100		100	1.33	11 440

Note: ¹Per 1000 lines.

C then weekly gives a reduction in both inventory and orders.

Further, if the A Class items were ordered twice as frequently and the C class half as frequently, what would be the effect?

The result is that:

- the total stock value is reduced from 2.1 weeks' cover down to 1.55 weeks' cover (excluding safety stock in each case)
- the number of line orders placed is also reduced from 12 000 per year down to 9750.

This shows that a reduction in effort and a decrease in stock are brought about at the same time. These results are typical and are borne out in practice. The effect on customer service from these changes is negligible because the safety stock remains the same in each instance. (The availability is based on the calculations shown in Chapter 6.) Stock discussed in Table 3.8 could be the same as overall inventory shown in Figure 3.4, where the total inventory includes the safety stock in the calculation.

Practical methods of reducing stockholding

The approach

Going back to the basic principle of

$$\text{Decrease in stock} = \text{Output} - \text{input}$$

shows that the way to reduce stock is to decrease input and increase output. A normal range of stock in stores comprises fast moving parts, a wealth of obsolescent and special items and fluctuating quantities of other parts, some in short supply, some overprovisioned. In this situation there is potential for decreasing stock by a significant amount, probably 30 per cent over eighteen months.

For many items there is little manoeuvrability to change the unit cost significantly. However this should be considered, since a 10 per cent reduction in price will cause a corresponding drop in inventory value in the long term. The average value of a stock line is given by multiplying

$$\text{Unit cost} \times \text{Average stock.}$$

The average stock quantity is typically halfway between maximum and minimum¹ stock. The minimum should be the safety stock and the maximum occurs immediately after a delivery. If the stock falls to safety stock level before delivery then:

$$\begin{aligned} \text{Average stock} &= \frac{(\text{max.} + \text{min.})}{2} \\ &= \frac{(\text{Safety stock} + \text{Order quantity}) + \text{Safety stock}}{2} \\ &= \text{Safety stock} + \frac{(\text{Order quantity})}{2} \end{aligned}$$

See Figure 3.6.

¹ The 'minimum' on some systems is set to the review level for triggering resupply even though the stock continues to fall below this level until the delivery arrives. Here, the 'minimum' is the lowest average stock level immediately before delivery.

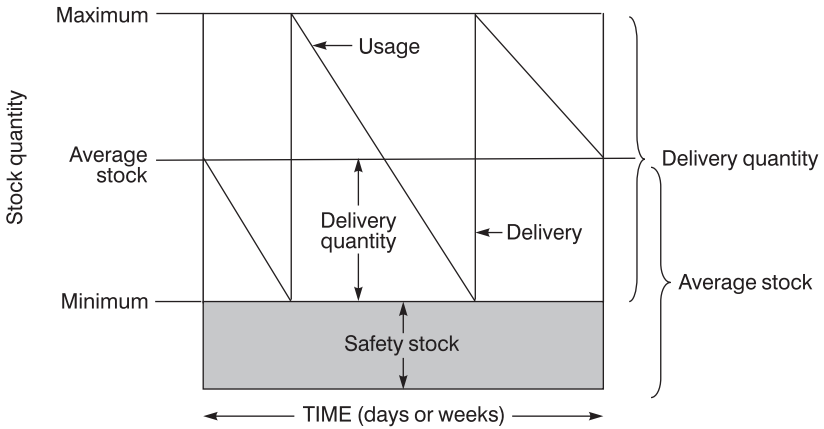


Figure 3.6

In order to reduce the average stock, therefore, there needs to be a reduction in either

- order quantity, or
- safety stock.

Batch quantity

The lowest inventory will result from receiving the same order size into stock as the issues to customers and, of course, this will be at a minimum if the order size is small. For those items which are individual, then the minimum order size is one. The parameter for order quantity on this basis is therefore the simple message ‘buy one’. This is also the obvious answer for the majority of slow moving stock items. It is not so obvious when stated in the form: ‘hold the least stock cover for the most important stock items’, a consequence of good stock control and also of JIT supply. Minimizing the delivery quantities of the A class items (or those contributing to the highest stock value) is normally the most fruitful way of reducing stock value.

Supplier delivery quantities are more important than lead times.

Safety stock

Minimizing the safety stock could result in lower customer service. Because the balance of stock is never perfect and the levels constantly need to be

changed, availability provided by the stock management needs to be monitored continuously. Therefore, the first step in controlling the high-value stocks is to ensure that the safety stock is sufficient to provide the service required and not more. Achieving a higher availability level than necessary is very expensive. The availability for each of the high-value items has to be matched to the requirements of the market. Balancing these stocks is most important. As the safety stock depends on the availability required, the lead time and the variability of demand, the next topic to consider is lead time.

Little and often for A class.

Lead times

By negotiating shorter lead times, the safety stock can be reduced. Modern purchasing practice should eliminate random demand ‘over the counter’ but it does require collaboration, forecasting and supply chain management so that the supplier can arrange to make items available by the demand date with a much shorter lead time. Lead time can be reduced by negotiation, and in the context of reducing stock investment, it is the few major stock items which will offer the most opportunity to negotiate.

The third factor is demand pattern. If the demand were ten every week consistently, then there would be no need for safety stock. (There are people from some areas of business who would question this simple statement, since the holding of extra stock would enable extra orders to be serviced should they arise. This extra stock is not safety stock but an investment in opportunity – or folly – and should be considered separately.) The question to be answered is ‘what opportunity is there to smooth out the demand pattern?’

Example: a case of a pan

The stock records for a range of materials sold by a major supplier of aluminium were typically as follows:

Week	1	2	3	4	5	6	7	8	9	10	11	12
Demand	225	2625	75	100	200	50	125	100	0	2050	300	125

The items were aluminium circles (flat round sheets) in various sizes used for making pans. The conventional stock level calculations suggested very high stock levels in order to cope with the demand peaks in weeks 2 and 10. This was introducing a high level of stock for these items, and inflating the inventory value.

Investigation of this order pattern showed that across the range of products, these peaks were caused by one major customer who had an order cycle of two or three months and this was causing the large demands. There are several solutions to this problem:

- **Understand demand** – *greater understanding of the demand pattern improves the forecast. Since safety stock is required to compensate for forecast inaccuracy, the improvement of forecasting reduces safety stock. There is, therefore, the opportunity to try to smooth customer demand by getting the customer to buy little and often.*
- **Flatten demand** – *remove sporadic demand from the stock by buying to order and giving the customer a longer lead time.*
- **Agree scheduled supply** – *It would be better to organize a continuous supply by scheduling delivery.*

The reduction project

Order quantities and safety stocks should be examined to ensure that they are synchronized, starting with the item with the most stock value and continuing down the stock value Pareto curve. A target should be set so that the stock reduction for each item can be judged against it. For example, an overall stock value reduction of 20 per cent can be achieved through a 25 per cent reduction in stock of the top 20 per cent of lines. Starting from the highest value stock line, a percentage reduction should be attempted for all lines. This is unlikely to be achieved for some items, but it is essential to concentrate on achieving the target for the highest value lines or the overall target will not be met. The value under-achievement on some items has to be compensated by extra value savings on other lines.

The approach – how to reduce stock successfully

Stock reduction is usually undertaken as a project, and concentration should be given to achieving the objective in a short time. By examining

the first few lines in the Pareto curve, major savings can be achieved in a short timescale. The next step within each area of activity is to find the major cost items. By Pareto Analysis it can be assumed that 80 per cent of the value of stock exists in 20 per cent of the item types. Items with the most stock value are likely to be a mixture of fast movers with reasonably high unit value and high-value items where there are relatively few demands per year. Pareto analysis shows that the way to reduce the stock value is to concentrate on the high stock value items, whether they are slow moving or fast moving. The value of inventory is the same for each and the financial investment is the same. In fact it is easier to vary the inventory levels for the faster moving items. As far as basic Pareto classification is concerned, there is no difference between the stock for fast-moving and slow-moving items with the same stock value. The differences will be managed through the individual stock calculations for the items.

High stock value fast movers reduce the fastest.

The stock of a fast-moving item is likely to be consumed quickly and consistently though servicing a wide variety of demands from a large customer base. The inventory profile for a fast mover is therefore large numbers of demands taking a small proportion of the inventory. Demand is relatively stable. For high-value items, there are relatively few issues to a smaller number of customers. This means that the demand is more unstable for two reasons:

- As the stock lasts longer, the risk of demand changing within the stock cover period is greater and therefore there is an enhanced risk of the stock remaining as obsolete or excess.
- The demand comes from only a few customers, if any of them reduces their offtake, this will have a much more significant effect than if there were more customers.

It is therefore especially important to manage risks on the slow-moving, A class items by forecasting, ordering little and often, and collaborating with customers and suppliers. The A items are selected for special control, applying a scheduling approach where possible and negotiating closely with the supplier. The supplier may be keeping a stock of these items specially or can be persuaded to do so if desirable. The delivery quantity for A class must be cut to a minimum and reviewed each time an order is placed.

The stock reduction project

Specific stock reduction can be achieved through a dedicated project. Continuous stock reduction is a background job for all inventory managers, but since insufficient focus is given to it, arranging a project tends to be more effective. The following steps may be taken:

- 1 Identify the potential for reducing stock.
 - (a) Find the total stock value.
 - (b) Find the annual material usage.
 - (c) Calculate the total stock cover for the various stores or categories of stock.

The results can then be tabulated, to identify which areas of stock have the best potential for significant value reduction, under the following categories: stock type, stock value, calculated stockturn, target stockturn.

The analysis of stock type can be generalized for distributors. Typical classifications are: market sector, product type, supplier, or other obvious classes. For manufacturers classifications can be into raw material, work in progress, bought-out parts, finished goods, consumables, etc.

All that is required is to reduce average stock value of significant stock lines. With this analysis it is easier to see which areas to concentrate on in a stock balancing programme. By assessing targets based on experience an initial plan can be launched, to be followed by a detailed analysis of appropriate targets at a later time.

- 2 The next step is to classify stock items into ABC.
- 3 For the few A class items, it is important that the minimum stock is held for the availability required. This is where the major effort in stock reduction should be focused. A items have to be reviewed individually and the supply arrangements reconsidered. Significant reductions are to be sought, through supply practice and better demand forecasting.
- 4 The reorder level and safety stock depend largely on the lead time. It is, therefore, important that lead times reflect current trading conditions and are kept low through contact with suppliers (especially for A class items).
- 5 For the medium price items, i.e. B class items, computer monitoring of safety stocks and order sizes will enable smaller batches to be ordered and safety stock to be reduced.

- 6 Avoid ordering more than (say) three months' supply of anything. Setting a limit ensures that the amount of slow-moving stock is reduced. Very long lead time items will have several purchase orders outstanding simultaneously.
- 7 Low value items purchased in category C are often standard products which can be obtained off the shelf. By proper planning, these need not be kept in stock. They can be purchased to meet demand or provided by a supplier on consignment. However, an efficient reorder procedure is necessary or else the works van will be touring the neighbourhood continuously picking up odds and ends.
- 8 From the stock turnover figures it becomes obvious that some items have too much stock cover, and are simply contributing to stock value. They are either obsolete or have very low sales. From our introduction to this section the formula shows that this stock has to be disposed of in order to have a low inventory presence.

The first step is to decide which stock is obsolete and to remove any chance of it being reordered by marking the records accordingly. The most profitable way to dispose of stock is through the servicing department. They can often sell obsolete lines to customers with old machines at full sales value or at an offer price near to it.

Stores are also likely to hold proprietary products such as components, fittings, motors, consumables, bearings and packaging which can often be sold back to the supplier at a low price. With high inflation over the past few years this return value can equal the original purchase price and so there may be no loss of assets involved.

By concentrating on major cost items the value of obsolete stock can be reduced. It generally takes time to negotiate the sale of items and therefore can be viewed as an ongoing project. Where these approaches still leave a significant number of useless items in stock, then disposal for scrap is the simplest course, and there should be an approved budget for this. The loss of assets through scrapping stock reflects directly on profit margins, and overzealous disposal projects can mean a period of poor financial results.

Armed with a budget for scrapping, items can be thrown out of stores and written off until the budget is used up. As stores are often short of space and control of slow movers is tedious, it is convenient to scrap a large number of the lowest value items in stock, thus reducing the number of stock lines and consequently freeing most space.

- 9 Design and planning departments can help to reduce stock requirements by very large amounts. All they need to do is to use the same items widely rather than marginally different ones for each application.

A stock rationalization exercise is normally fruitful for standard items such as electronic components, motors, fasteners, raw materials, packaging items, gearwheels and tools. This can be carried out initially by inventory control and then passed on to technical people to continue.

- 10 The order is the cause of increased stock. Therefore it is necessary when undertaking a stock reduction exercise to look at every order placed to ensure that:
 - (a) The item is required at the time purchased.
 - (b) A minimum is ordered.
 - (c) There is no stock existing which can be used instead.
- 11 To maintain overall control of the stock reduction project reliable management information must be available on a monthly or weekly basis for such things as the value of items on order, stock cover and availability. By plotting them on a graph the effect of these actions can be measured continuously.

Key points

- ABC analysis provides the best tool for saving time and structuring inventory.
- Stock cover enables inventory managers to look out for likely problems.
- Stock cover is not used for setting safety stock levels.
- Inventory is reduced chiefly by using Pareto analysis and controlling supply.

Just-in-time management

- The opportunity for low inventory
- Improving the operating conditions
- How to use just-in-time supply
- Identifying the benefits

The zero inventory philosophy

Conventional and JIT approaches

Stock exists because items have been bought before they are required. It is normally uncertainty or overcaution that causes inventory. The principle of JIT is simply that we have items when they are needed and none when they are not needed.

JIT supply is a result of high-quality supply.

The idea may be simple but the application of JIT has given the opportunity to decimate stockholding without affecting customers. Instead of trading availability and stockholding as discussed in Chapter 2, the trade-off is between organization and stockholding. The better organized and controlled the supply chain, the less inventory is required.

Companies which are considering how JIT can work in their business, or avoiding it in a Canute-like manner, should realize that JIT is an outcome of other techniques, not a technique of its own. It is the logical aim of tight inventory control, effective process planning and plant design, workforce motivation, cost reduction, logistics and even material requirements planning (MRP). The optimization of these together inevitably leads to the JIT approach. The elements of JIT are the techniques to be developed, for example:

- supply what is required
- supply the quality required

- reduce lead times
- organize effectiveness
- use all the expertise available (i.e. people who do the jobs plus technical backup).

These are the fundamental changes which lead to JIT – all good inventory management techniques. The most important of these improvements has been in quality, particularly in management processes – quality of records, supply, delivery, forecasting and target setting.

Holding inventory shows we don't have control.

Stock controllers have been sceptical about the efficiency of frequent deliveries of smaller batches without investigating all the options and the potential. There are many instances where the accepted delivery quantity is now much smaller than it was a few years ago, and is destined to continue to be reduced. Some of the contrasting features of managing conventional stock control and JIT control are shown in Table 4.1.

Table 4.1 Contrasts between conventional and JIT inventory control

Conventional	Just in time
Satisfied with the status quo	Continuous improvement
Lead time is fixed	Reducing lead time is a continuing challenge
Product range is a sales issue	Product range reduction is an inventory issue
Management provide methods	Operators are responsible for practices
Stock in case of customer demand	Purchased to meet demand rate
Convenient purchase batch size	Buy singly or small quantities

Just in time works as a pull system whereas conventional stock control and material requirements planning are essentially push systems. In the push system the stock is provided for the next stage of supply, e.g. buying items for sales to sell or starting manufacture without having the total production path clear. The philosophy with push systems could often be thought of as packing so much into the warehouse or production plant that they will have to send some items out because there is only limited space!

With a pull system the first action in the chain is that the item is demanded. To satisfy this demand there is an item in stock. As soon as this stock is used up, then another item is supplied, either from outside or from

a production process. The supply is organized so that the effective supply lead time is very short (under one day) and the quantity supplied is small – a few minutes up to a week’s worth of demand. In production the demand filters backwards through the processes creating a demand down the bill of materials, and ending up with issue of one product’s worth of raw material.

Continuous improvement processes are associated with JIT including product quality, process efficiency, information systems and operating value-added activities more effectively while eliminating non-value-added activities.

Who can use JIT?

Quality process improvement is usually thought of as a continuous journey of improvement, with no definite ending, since there is always more potential. From the point of view of material flow the principle of JIT is obviously ideal but it is often difficult to implement in practical situations unless the conditions are right. Of course, the right conditions do not happen by accident and anyone looking for the benefits has to work hard to create the appropriate situation. Just-in-time supply should be considered as a quality process, although most of the objections to JIT are based on lack of quality in some aspect of supply or demand. If stock levels are incorrect, this is often the result of complacency or lack of understanding (usually both). There is no perfect solution to stockholding but, like any other quality improvement process, JIT operations gradually develop an existing unsatisfactory situation into an improved one. A decrease in lead times and simplification of processes should be the aim of all inventory managers. From the viewpoint of JIT, time is a value-added commodity and wasting it is unprofitable. The more time saved the better, and continuous improvement means reducing the timescales.

JIT results from best practice in inventory control.

The definition of just in time presented so far can apply to any material management process which actively minimizes timescales. The purist would think that there is more to JIT than this simple concept and there are some specific concepts for achieving this reduction in timescales, particularly:

- desire to improve
- simplification
- demand-led supply (pull)

- quality conformance
- devolution of responsibility.

If these concepts are put into practice, then an operation has a JIT philosophy (most of these concepts are fundamental to good inventory management anyway). Supported by improvements in communications and driven by the need for better service and lower costs, the influence of JIT has been felt in all types of business and has fuelled change.

JIT in manufacture and warehousing

Just in time was taken up by some sectors of manufacturing originally as a space efficient method of production. For some supply chains, the application of JIT is natural because of the nature of the products or the processes, e.g. in fast-cycle manufacture, where a stage in the production process takes very little time. In a high-volume manufacture, delays between processes would cause large stockholding (or, more specifically, work in progress). In other businesses JIT techniques are less easy to apply. Retail shops find it important to hold stock so that the customer can have a choice of several options, so JIT supply would not be appropriate for this stage of supply without changing the nature of the relationship. If the customer were buying the same item by mail order, then there is a very good opportunity to use JIT techniques.

Although the original impetus for JIT was for space saving, the greatest benefits are financial. In Chapter 1 (section on profit through inventory management) the benefits of lower stock on profits were illustrated. This magnitude of improvement is likely to occur for a second time as a result of applying JIT.

Consider the example used in Chapter 1 with each of the companies having a £5 million turnover. Often companies find that half their costs are in purchased materials; let us assume that the same is true for these companies.

If these companies are distributors, then the stock cover for The Slack Company is 52 weeks and for M. Tight Ltd is 10.4 weeks. If the application of JIT on selected lines (product types) could reduce the overall stock cover to 1 week, then this would result in an inventory holding of only £48 000, a saving of £202 000 even for M. Tight Ltd. The stock cover and relative inventory investment are shown in Table 4.2 for a turnover of £5 million. The decrease in inventory values is large when the initial stock cover is high. However, as stock cover becomes low the potential savings are lower and

Table 4.2 Warehouse stockholding

Stock cover	Inventory investment (£)
52 weeks	2500
10.4 weeks	500
6 weeks	288
4.2 weeks	202
1 week	48
1 day	10
1 hour	1.3

the risks greater. It is then not the value of inventory that is the driving factor toward zero inventory, but the operational cost savings.

Applying the same logic to a manufacturing situation shows additional benefits through the savings in planning resource. The operational situation is more complex in manufacturing. The material cost starts out as, say, 50 per cent of turnover but then collects added value until it is worth typically 70 per cent of sales value. For simplicity let us assume that there is an equal balance of inventory throughout manufacture, so that the average inventory value is 60 per cent of sales value.

Lead time should be reduced – it results from a planning problem.

A further aspect of manufacturing is the number of jobs in progress. It is typical for a company with a turnover of £5 million to issue about fifty new jobs to production each week. (The logic of the following discussion works just as well for companies with different numbers of jobs and turnover values.)

Using this figure a manufacturer’s performance table can be created as shown in Table 4.3. Stock cover in manufacturing is required for the duration of average throughput time. For The Slack Company the £2.5 million inventory equates to 43.3 weeks’ cover and a total of 2167 jobs in work (assuming that all the inventory value is in work in progress – WIP). For Tight Ltd the WIP equates to 8.7 weeks’ cover and 433 jobs.

The advantage of reduced lead times in manufacturing is twofold, as can be seen from Table 4.3. Inventory investment is reduced as well as the number of jobs and the complexity of the materials management. It takes a great deal of effort to control 433 jobs, but much less to track fifty.

Table 4.3 Effect of throughput time on manufacturing inventory

Throughput time (weeks)	Inventory investment (£)	Number of jobs in progress
43.34 weeks	2500.0	2167.00
8.66 weeks	416.0	433.00
6 weeks	288.0	300.00
4.2 weeks	202.0	210.00
1 week	48.0	50.00
1 day	10.0	10.00
1 hour	1.3	1.35

Table 4.3 shows that major cost savings go hand in hand with reductions in manufacturing complexity. Inventory cost savings in particular can be achieved by cutting the lead time down to a few weeks, days or hours – and then benefits flow from other operating costs, such as less shop-floor space required, improved flexibility, less expediting, more control and less material handling. There is a similar simplification in warehousing where lower stocks means less to count and less to control.

JIT environment

The ideal situation

Over the years the perception of lead time has changed in most businesses. Lead times have shortened due to the combined influence of customers and accountants, the former wanting faster service and the latter lower inventory. This has been a gradual culture change, but the move to true JIT requires a further leap in thinking. An attitude of ‘we’re doing alright, don’t risk it!’, has to be replaced by ‘this is what we need to do, now how do we make it happen?’ – i.e. continuous development. Just-in-time individuals have to create an environment where things do not go wrong and their positive attitude is the basis of JIT philosophy, whose essence is reliability and consistency. Ideally this reliability must pervade the whole of the material control operation, not just one stage. Performance is dependent upon a number of factors, for example suppliers, the market, the company itself and the individuals within it, and these factors will always throw up a problem of some sort. (These should be considered as the challenges of everyday stock control.)

Business needs to use JIT to remain competitive.

In high-volume manufacture it is easy to see how items can be transferred from operation to operation, and lead times and cycle times can be kept very low. In distribution operations where the demand is variable and uncertain, then the application of JIT is rather more difficult, but not impossible.

Typical features of the ideal company where JIT concepts could be applied would be:

- narrow product range
- manufacturer
- high volume
- stable market
- influential
- good quality management
- local, reliable suppliers
- dependent suppliers
- fast-cycle processes
- affluent.

The one major element which was left out of this ideal list is personal commitment, and although comparison with one's own current business profile might prompt many to discard JIT as inappropriate, a positive attitude will prompt others to consider how close their business is to these ideals and how JIT can work in reality. This is a major factor because successful JIT depends on commitment by a core of people at operational level and equal dedication of managers and directors. As in all good management, it is useful to have monitors for success. These are local targets decided upon by the local team to set and manage their own performance and criteria for success. Quality and throughput will be part of these measures.

Repetitive processes

Just in time can be seen to work best in repetitive processes, whether in warehousing or in manufacturing. Of course, almost all of business is based on continuity of supply, so this is not really a major constraint. For a usage of twenty per month, a traditional approach could be to buy twenty at a time, whereas JIT would require one per day. Such a small order rate may not appear economic for one item, but it would be when dealing with a large number of items.

Consider a packing or manufacturing line that each week completes

- 2000 of item a
- 3000 of item b
- 4000 of item c
- 2500 of item d
- 1000 of item e.

Instead of working on each type of item in turn, it would be better to provide each item at the rate at which they are required. This can be achieved in two ways:

- produce at a balanced hourly rate across the whole range
- produce to order and adjust the rate to suit demand.

For the first option the rate of production is consistent. The supply rate is shown in Table 4.4. For this option to work, there has to be a very consistent demand or a buffer stock somewhere in the system. This is generally the way that JIT operates in warehousing and manufacture where internal planning is on a fast-cycle basis, but customer demand is on a slower-cycle basis so that demand fluctuations can be smoothed out over the customer supply cycle.

Examples of this situation are the automotive industry and high technology equipment manufacture. In these cases the supply can be geared to a constant rate of output and the resources organized to efficient and cost effective supply (and production where appropriate). In Table 4.4 the options for the stores or production unit are to have a single throughput channel which completes an item every 11 seconds, or to have a multiple of channels which deal with a single or a limited range of items.

Table 4.4 Continuous supply cycle times

Item	Throughput per week	Throughput per day	Cycle time (secs)
A	2 000	400	68
B	3 000	600	45
C	4 000	800	34
D	2 500	500	54
E	1 000	200	135
Total	12 500	2500	11

By making manufacturing repetitive planning capacity and materials is easy.

Example

In the automotive industry there is generally one production line making the variety of processes at the same time but on one car at a time (with a parallel line for flexibility and offsetting the effects of changeover or plant failure). In a mail-order warehouse the capital investment per line is not so great, and stock items are more diverse, so that generally there is a number of specialized packing lines.

In warehousing it is important to apply JIT where it is easiest to use, and the conditions discussed earlier in this chapter point toward the A class items with high usage, more consistent demand patterns and significant benefits from inventory reduction. The items sold can be split into two types:

- those controlled conventionally
- those using JIT.

The major items will be JIT controlled and the wide variety of low turnover value items controlled conventionally.

When we have achieved full control through conventional means (see Chapter 7 for methods) the control for the second option (JIT) becomes an extension of this. Here the customer requires the items on a fast-cycle basis. This means that fluctuations in demand have to be taken up by varying the capacity available. This is less efficient for the use of resources, but may avoid inventory build-up elsewhere in the supply chain.

JIT supply

Processes can be considered to be suitable for JIT if the timescale is in hours rather than days. For most suppliers this requires special arrangements for feedback of demand and issue of items within a short timescale. Effective JIT supply requires the same approach as that used for single source suppliers (see Chapter 8): there has to be a longer-term agreement on anticipated throughput levels, quality management and commitment by the supplier to provide an infallible supply of specified items. The supplier should be treated like an integral part of the organization.

One of the changes in attitude that this brings is the approach to transport. Often the suppliers organize delivery so that they incur the cost. With JIT supply there is an advantage if the customer collects the items at a time convenient to them (the pull system being extended in this way through the transport system), since the transport cost would be incurred in the final product cost whichever company pays for it.

For a fast pull-type relationship to work, then the supplier has to be local to the customer, otherwise the reactivity is not good enough. Location is of significant importance, and some suppliers have moved to the vicinity of a major customer in order to ensure close cooperation.

JIT and distribution

Just in time requires more frequent deliveries from suppliers but this does not automatically lead to higher transport costs, although it could do so if no changes are made to mode and arrangement of supply. Simple examples of where there is definitely no extra transport cost are:

- when the freight is charged simply by weight
- when a wider variety of items is put into each shipment and these compensate for the reduced amount of bulk orders
- where freight is charged at a flat contract rate.

Costs might increase where:

- the same delivery method is used more frequently
- batches of items are bought separately from different suppliers
- a flat rate per consignment is charged
- transport is arranged one journey at a time.

Change transport method and get frequent supply.

So these situations are to be avoided.

Just-in-time philosophy requires close co-operation with suppliers of distribution services as well as of goods. By moving toward longer-term, higher-volume commitments with suppliers, prices can be reduced based on the total traffic volume over an extended period. The distributor requires a continuity of demand and the opportunity to work out routes, loads and schedules well ahead and this will reduce the operating costs as long as the business is significant for the distribution partner.

Many people are sceptical about increasing the number of deliveries without increasing the overall costs. Others who have actually tried

increasing the delivery frequency have found distribution partners who can provide the service without increasing financial disadvantage, and there is a growing number of companies working this way.

The general cost savings produced by JIT are discussed below. The economics of frequent deliveries can be added to these savings for delivery processes. Some of the typical benefits are:

- 1 Smaller delivery quantity means smaller loads, which enable different transport methods to be used (parcel post instead of lorry load). Small bulk transport often costs considerably less, since the operating costs are reduced.
- 2 Frequent fixed route deliveries enable a carrier to have a base load upon which to build other business.
- 3 Frequent deliveries do not require planning and management effort once set up.
- 4 Discounts on contracts can be very significant.
- 5 Standard pack quantities and containers are smaller, less expensive and recycle faster.

JIT saves costs if it is organized properly.

The secret of success in saving costs on frequent JIT deliveries is to:

- plan carefully so that delivery schedules are not altered at the last minute (this is easier with JIT since the planning horizon is relatively short)
- be bold – radical changes in delivery methods are required if costs are to be minimized (this often means a complete change in transport method)
- negotiate carefully – transport costs are adjustable because they contain a large fixed element of cost. The apportionment of this cost (e.g. running a lorry from A to B) can be carried out in any way. Transport suppliers may be persuaded to apportion the costs away from routine major deliveries
- be open – the size frequency and costs of delivery required should be shared with the transporter. Their relevant costs should also be known in order to ensure that their service is viable.

Quality management

Concurrent with the development of JIT has been a great move to improve quality. This has taken two forms.

First there was the move toward consistency through the development of quality standards (including ISO 9000) which were instrumental in providing reliable performance not only for the products provided, but also for procedures and paperwork. These standards have changed attitudes within industry and there is now an expectation that items provided shall be correct to specification.

The second was the development of the TQM philosophy, which does not accept current quality and strives always to provide gradually improving quality. When this philosophy is applied to the operating processes of the organization, the key concerns are the elimination of waste and operating effectiveness. Total quality management then becomes a practical approach to productivity improvement, standardization of procedures and packaging, improvement of the quality of delivery on time, record accuracy information and operating effectiveness, and the complete operating processes of the company.

Suppliers should pay a penalty for poor quality.

Quality management principles integrate well into stock control methods and operations because it is a basic assumption of stock control systems that the items supplied are fit for their purpose. If a delivery batch is not acceptable, then it has to be replaced immediately for the stock control calculations to be correct. Delivery in smaller batch quantities often reduces the overall effect of rejected supply since it impacts on fewer customers' orders and leaves less time before the next delivery. Non-conformances will cause extra work, especially:

- creation and holding of extra records as a cross-check
- quality inspection on goods inward (also causing booking delays)
- sales personnel requiring physical stock checks rather than relying on the recorded quantities
- time spent investigating and checking
- arranging replacements and priority deliveries.

Formal specification of inventory management procedures enables them to be reviewed and identified for quality management processes. As time is costly, the allocation of time to administration has to be done sparingly. The cost and benefit accrued by maintaining full records and checking procedures has to be balanced against their added value benefits. Non-essential activities have to be avoided (such as paperwork and manual systems where there are computer records). Procedures have to be

improved and rationalized continuously, to maintain the impetus of the quality improvement.

Customer support

Supplier quality is not only measured by the acceptability of the product, but also by the customer support provided. Rating of suppliers should include:

- delivery performance
- packaging and labelling
- information systems and communication
- support
- style
- flexibility.

Supplier quality issues are discussed fully in Chapter 8.

Advantages of JIT

Operational benefits

How can it be more efficient to deliver in small quantities, manufacture in small batches and increase the amount of administration? These are often the queries from those brought up in the traditional school of thinking. Equally, someone who only knows JIT would ask:

- Why do you buy things when you don't need them?
- How do you know what the demand is so far ahead?
- What use is a warehouse?

The accounts of a conventional company show a large investment in inventory. In the case of the two companies in Chapter 1, the profits were increased by a low stock philosophy. If M. Tight Ltd were to adopt a JIT philosophy this could bring the stockholding down even further to one or two days' worth, say £20 000 total stock value. This would release a further £480 000 cash for the business and increase the return on assets from 16.7 per cent to 20 per cent as a minimum. (As a move to JIT will eliminate the stores and some overheads, the real improvement

will be greater than calculated here.) In this case each company had a large investment in fixed assets. For a distributor with rented warehousing the increases in profitability are very large, even if the average stock value turns out to be weeks' worth rather than days' worth.

The operational benefits arising from JIT are:

- inventory investment
- 'supply to order' instead of 'provision for stock' (see Chapter 7, section on managing lead times)
- easier forecasting so less slow moving stock
- better flexibility
- simplified administration
- waste elimination
- less scrap should there be a problem.

For each of these operational benefits there is a corresponding cost benefit which can be offset against any additional costs which arise. These additional costs usually occur because methods have not been changed to suit JIT. For instance, if an item is delivered in a batch once per month, it can be invoiced, delivery documentation produced and payment made. If through a change to JIT the item is delivered every day, then it would not be sensible to place a purchase order for each load, to raise delivery paperwork and to arrange separate payments for each load. Information is still required for control, but the information system has to be re-created to meet the new conditions. In the short term this may not be possible and, so, extra costs can arise. As JIT embodies the process of continuous improvement, the inefficiency will eventually be eliminated.

Efficiency benefits

Just in time supports the continuous output of a variety of different items, which in turn creates a variation of items going through the processes (picking, manufacture, packing, etc.) Flexibility is essential, but conventional thinking is that changeovers are to be avoided since they take up time and use resources. The JIT approach is to find the output rates as illustrated in Table 4.4 and then determine how to achieve the throughput and mix. Changeover time is therefore considered as a variable and this is a basic difference in the JIT approach. Normally the changeover time is associated with a particular product.

Example

Conventional logic is that if it takes ten minutes to change over for a particular item type, there is no point in making just one (which takes fifteen seconds); better to make at least eighty to be efficient. Accountants usually back up this type of simplistic logic with financial break-even analysis (see for example economic order quantity –EOQ – in Chapter 8). This results in large batches being produced and high stocks being held. The real situation is that changeover time is a variable depending on what the change is from and to.

Fast changeovers give the opportunity for lower stockholding. On a packing line, the change from one type of item can be fast if they are packed in the same type of box, but slow if the box size has to be changed. The secret of success is therefore to always run the right job sequence so that the number of short changeovers will be large and long changeovers small. This requires planning and load balancing which is difficult when forecasting for weeks ahead, but easier when looking ahead a few hours as required by JIT.

Cyclic sequencing reduces overheads.

Assuming that the items in Table 4.4 are all processed on the same process line, then one item needs to be completed every 11 seconds (5.6 per minute). If the process is carried out on a cyclic basis as shown in Table 4.5, then the cycle is repeated every 4.5 minutes. The number of each item in the cycle is shown, together with the output rate required to cope with the average demand. If the 4.5-minute cycle is to be carried out most efficiently, then the changeover times between the items need to be reduced. There are two ways in which the changeover time can be reduced, namely:

- Improve the planning of production sequences.
- Reduce the time per changeover.

Considering these two options, the improvement of product planning of sequences can make most changeovers minor. Within a short repetitive cycle it is much easier to sequence the production to optimize the setup time than it would be with conventional priority planning. The quantity of each item can be varied to suit the demand required if this varies, and items

Table 4.5 Cyclic production

Line	Throughput per minute	Number per cycle
Item A	0.9	4
Item B	1.3	6
Item C	1.8	8
Item D	1.1	5
Item E	0.4	2
Total	5.6	25
Times per complete cycle = 4.5 mins		

with the same setup can be substituted should this be required by the programme.

Changing the production philosophy from weekly batch quantities (see Table 4.4) to cyclic JIT production (Table 4.5) will reduce stock levels by a factor of twenty, and also save overall cost. Just in time organizations often have small management teams with few management levels. This is a result of the JIT process itself. The planning operation is simple and the management is largely a repetitive activity which does not require a great deal of managing, especially as the direct process control is carried out by the direct operators.

From an accountancy viewpoint, if the cost of supply is split under the headings of materials, labour and overhead, the major impact of JIT can be in the lower overheads resulting from the simpler operations. The proportion of overheads applied to the direct costs are much lower in JIT operations.

Stock control using JIT

Lead time reduction

Inventory control has to work under the conditions imposed by the market. To get the stock levels right calculations are based on the real information which is available and not on idealized values. This situation can be improved because JIT gives potential to improve, particularly the potential to shorten supply lead times. In working with a dedicated supplier effective lead times can be changed. The supplier and customer can work together and share the risk.

Lead times are a result of queuing, distribution and sometimes process time (see Chapter 7.) For JIT suppliers, the demand is assured by the

agreement with the customers, but the demand timing and quantity varies. As the JIT demand is frequent, the effects of these variations cause only minor stockholding and delays – days' worth not weeks' worth. Small stocks enable the supplier to meet the small daily fluctuations of the customer very rapidly. Lead times can therefore be reduced.

To set up an effective JIT supply:

- the customer must *smooth the forecast*
- the supplier needs a regular demand for the item in order to *estimate the average demand rate*
- the supplier also needs a good estimate of the long-term demand for their planning in order to *understand the supplier's constraints*.

JIT results from making other things reliable.

Suppliers can extend some aspects of their service and struggle to meet other aspects:

- Greater understanding improves the match between the needs of the two companies and *builds confidence*.
- Trust is the major ingredient for effective collaboration in order to *maximize supplier resources*.
- If suppliers know the volume of demand then they can react much more quickly to individual requirements. For instance a supplier of painted items could often paint and deliver within a few hours, as long as the demand for the item is expected and the choice of paint colours has been agreed previously. This enables them to offer *flexibility in demand changes*.
- As short-term demand is always variable, an agreement is needed to ensure the extent to which demand can be raised or reduced. This avoids failures in the JIT supply. It is therefore necessary to *communicate and monitor*.

Close liaison between the two parties involved is essential. It involves regular meetings for assessment and review between people at operational level.

Reducing the lead time is a continuous process starting with the A class items. These have the greatest cost benefit and are also amongst the simplest to change. They have most importance to the businesses and are among the easiest to forecast.

Pull systems

The essential difference between the JIT approach to material control and the stock planning approach is in the mechanism for triggering replenishment. Conventional inventory control predicts demand, works out how much stock is required to supply it and then procures that quantity. Just-in-time operations only acquire stock as a result of demand. This can be thought of as a 'pull' process, with the issue of an item triggering the demand for another – the number of triggers depends upon the rate at which demands are received. Stock-level systems purchase items to cover forecast demand (in larger quantities) and then try to supply the items to the customers (a 'push' system). The advantage of the pull system is that only sufficient stock is held to meet the immediate demand for a few hours or days. The effective supply lead time is arranged to be short (JIT) and the deliveries frequent, therefore the stock can be very low.

In inventory management simplicity of operation is the key to success. Kanban is an operational method, consisting simply of a ticket, which achieves a pull system from within the company or from a supplier. When an item is demanded a request must be passed back to the source to provide another. This information can be a kanban card or a simple signalling system. The card will identify the item, the quantity required (ideally one, but possibly more) and where it is required. It will also state times so that performance can be monitored. This ticket or traveller has to be provided to the source very rapidly if the supply is to be prompt.

The ticket is held with the physical goods and is an easily maintained and simple method of informing the previous stage of supply. Using the principle of the ticket, there are other trigger mechanisms which can be used instead, such as coloured lights or sound. Fax or network communications can be used for external supply. The way information is transferred is immaterial so long as the process is instantaneous, completely reliable and accurate.

An obvious application of JIT is for manufacturers producing large quantities of similar products. They can set up a flow line and have in-line stocks of a single item or a small batch between each process stage. The trigger for manufacture is simply 'use one so make one' starting from the last process and feeding back gradually to the first. The restructuring of production into a flow line is a most effective way to manufacture. Where the product range is wider and the process routes differ, cellular production is more applicable, and the pull system is used between the cells and by the cell operators for processes within the cells.

For purchased items the supply is triggered in the same way but with larger batches. The delivery batch would be expected to meet the demands

for a few hours or a day. As the lead time is normally over an hour, a supplier kanban does not run out before signalling. It has a residual stock to cover the demand during the supply lead time.

Just in time has made a major impact on manufacturing and the reduction of lead times is changing the whole of inventory control. The changes will provide competitive advantage for those companies with the best ongoing developments of JIT.

Key points

- JIT supply is the aim of best practice in inventory control.
- JIT is a result of improving the quality of everything and organizing 'flow'.
- JIT business is fundamentally different from the traditional. The key concepts are:
 - pull – only buy when needed (kanban)
 - one – only make or buy small quantities
 - flow – balance movement throughout supply chain
 - simplicity – operating controls should be simple and visible
 - quality – high reliability for supply, products and control processes.
- Apply JIT to where it is most beneficial – A class.

Organization and management

- Set the direction
- Identify objectives and targets
- Manage stock value
- How to value and depreciate inventory

Where stock control fits into the organization

Nowadays the inventory controller has extra responsibilities. The job has grown out from storekeeping to become a management discipline of maintaining the inventory to meet company policy. Gradually this responsibility has broadened to include all stock whether inside or outside the stores. The place of inventory management in the organization has altered as responsibilities have changed. This is illustrated in Figure 5.1.

The original role of stock control was part of the physical stores control shown in the structure in Figure 5.1A and its position in an organization was lowly, working for either purchasing, finance or, in manufacturing

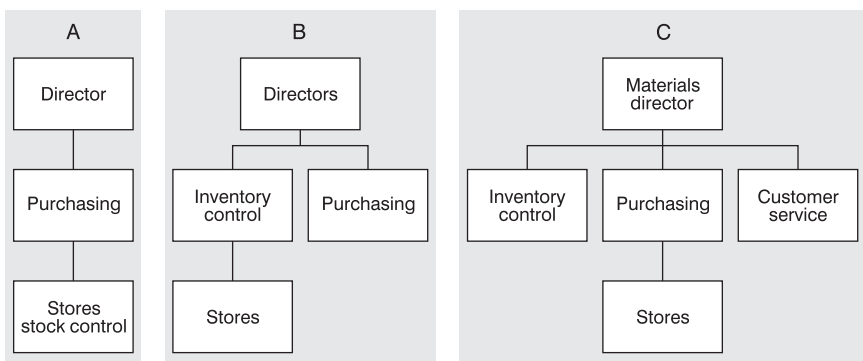


Figure 5.1 Changing inventory management structures

companies, for production. In this structure stock acquisition was initiated by buying and sales pressure, and there was no clear responsibility for, or focus on, inventory control. The tightening of company objectives changed the approach and separated stores and stock control activities. Inventory management is now accepted as an operational activity and works alongside purchasing, as illustrated in Figure 5.1B.

Inventory management is a senior management activity.

This arrangement has gradually evolved into structure C as pressure increases for better performance. The central role of inventory management has been recognized in many businesses, both in distribution and manufacturing, and material management or logistics directors have been appointed to co-ordinate various aspects of inventory management into one supply activity. Stock control encompasses:

- purchasing – especially reordering (but see Chapter 8)
- forecasting demand
- planning inventory against company targets
- stock allocation and delivery promising
- monitoring and controlling service and inventory
- production planning and control (for manufacturing).

These are organizational activities. The physical management of inventory, including stores, dispatch and movement then becomes treated as a separate problem to be managed by a parallel organization. A similar relationship can arise in manufacturing between the production control department, who carry out the planning, and production supervision itself, which seeks to work to the plans provided by the production control department.

Where there is a separation of the stock and stock control activities, the warehouse become responsible not only for managing the movement and holding of stock, but also for providing information to the rest of the company. It is particularly important for inventory management that the information provided is accurate. Checking stock physically is very time-consuming, and the recording systems should be sufficiently accurate to make this unnecessary apart from stock auditing.

*Stores control takes care of inventory.
Inventory control sets the amount to keep.*

Stock control in a professional organization can be carried out at a different location from the stores. Thus stock controllers can manage site stocks or multi-location storage as long as there are real-time communications with the stores. If these activities are split, warehousing can be located at the most economic location for transport and customers and inventory control can be located independently at the most convenient place for management.

Inventory management is in the process of evolution. Targets are being expanded to include extra aspects of operations, cost and service and there is much potential for development (see Chapter 13).

Responsibilities and targets

Rationalizing the problem

Inventory controllers have to take a simplified view of the stock items and treat the majority of them in the same way. In driving a car there is a distinction between the driver (who may not be able to repair the car) and the mechanic (who may not be able to drive). In stock control, most people are the drivers but the stock control vehicle is not well designed, mechanics are few and, so, the driver is often forced to make improvements. The driver who understands the mechanics can always get better performance out of a vehicle.

Inventory management has the same principles and techniques – only the situation is different.

Every item has its own characteristics, often with a unique story – ‘do you remember when the supplier’s warehouse burned down?’ ‘The customer was frozen up for five weeks?’ ‘It was put on a boat to Antarctica by mistake?’ In the past individuals managed stock through their memory and understanding but this is not enough to meet the tighter targets of modern inventory management even though such understanding is an important additional tool to use for A class items (and a waste of time for C class).

The stock controller should largely ignore the details of the item and consider the characteristics of the demand pattern and the supply, despite the widely different types, sizes, shapes, usage, origins and characteristics of the lines, but whatever the item is, inventory control remains the same. An arguable point? What about liquids, explosives, perishable foodstuffs, pine trees? They are all very different. True but, despite the differences, their supply and demand characteristics are not that different: they depend on

market conditions and not on the product. Just as financial controls are applicable across different industries, so is inventory management. The characteristics of the stock determine the mix of inventory techniques to use, but the components are the same:

- demand forecasting
- determining safety stocks
- negotiating supply patterns
- avoiding slow-moving stock.

Inventory techniques are common for all businesses, and their skilful mixing for aggregate and individual item inventory control allows many different types of stock to be managed. Simplicity is the key, but this is not the simplicity of systems which give superficial answers. It is the simplicity of operating processes resulting from rationalization of procedure, systems and approach. It arises from, say, using the same system for placing all orders and controlling all the stock; in having the same item identification throughout the business; in treating priorities through an organized system. The elements of the simple procedures have to be supported by complex tools such as software packages for forecasting, modelling and communicating, but these need not concern the daily operation as long as they work correctly. The actual logic of such tools has to be intricate to cope with the complicated nature of inventory planning.

Simplification should be considered as an ongoing challenge for development of good stock-keeping. The following areas can usually be improved:

- single stores philosophy
- inventory system applied to everything
- single operating procedure
- paperless information systems
- purchasing as part of inventory management
- real-time systems.

By rationalizing these activities, the professional stock controller can free time to work on the most beneficial topics, and improve the level of control.

External objectives

Discussion earlier in this book was about setting inventory targets for availability, inventory value and operating costs. The balance between these

targets is a matter of inventory policy. The management of inventory relies upon routine and regular monitoring of these values. The objectives set for inventory management should be chosen to meet the immediate demands of the business, and specifically:

- customer service
 - measured as stock availability or delivery on time (percentage against own and customers' targets)
 - analysis of age distribution of lateness
 - assessment of customer satisfaction – survey results
- inventory investment
 - value
 - stock cover ratio
- operating costs
 - warehousing and inventory operations costs
 - cost per transaction, movement and purchase.

Key objectives – 100 per cent availability, zero inventory, no operating costs.

The process of setting targets usually begins by measuring current success. Current levels of achievement are assessed to see whether they are good or bad. (If inventory control staff consider that they have the correct stock level, this is often because of their complacency and not because of adequate achievement assessment.) Targets can then be set through negotiation with those who are affected by them and with those who have to achieve them. The three target dimensions are:

- quantity – how much the monitored values will change
- date – the stages by which the quantities will be changed
- responsibility – who will be accountable for achieving the results?

The prime aim is to meet all the inventory targets simultaneously, and not to sacrifice one to meet another. Measurement methods should meet the prime business needs (in Chapter 2 alternative ways of measuring customer service were discussed) and should be chosen to give the most satisfactory results.

Identify all your real objectives.

Routine monitoring of these items is required on a weekly or monthly basis. The results, both operational and financial, should be available within three

days of the period end, otherwise the information cannot be used for control (although it is commonly used for apportioning blame).

Managers should also be aware that targets are set as part of company policy and the art of correct management is to set targets which are both achievable and challenging. They should reflect an expectation of performance, not an achievement under ideal conditions. This can be difficult when trying to co-ordinate planning with other departments. A good method of improving the accuracy of planning is to give two targets, one optimistic and one pessimistic, and to ensure that performance lies between the two. By adjusting the two limits, a target range can be identified which is both attainable consistently, and which is acceptable for planning and co-ordination with other activities.

Internal targets

External targets have to be supported by more detailed objectives within an inventory management department and these should include personal objectives for individuals which all contribute to departmental targets.

Customer service can be structured by setting availability targets for different types of items. Some service is likely to be ex-stock with different availability levels, while other service will be supplied on a range of lead times. Information required to give fast direct performance measures include:

- number of lines out of stock
- number of customers and orders affected by shortages or late deliveries
- analysis of shortage according to cause
- monitoring of customer complaints
- details of customer-care initiatives.

Set quantitative targets – or you won't know when you've achieved them.

Stock levels will differ for these product ranges and have different expectations in terms of stockturns. Items with a JIT-type of operation have very good stockturns while slow-moving items such as service parts can inevitably achieve only relatively poor results. In addition to the inventory value, the investment can be controlled by managing purchase commitment and write-off value as discussed below.

We need to aim for several targets at the same time.

Operating costs are increasingly becoming the responsibility of individual inventory controllers. The elements of cost are controlled through cost analyses and budgets, and monitoring operational activity levels often provides the basis for cost reduction. These controls include:

- value of transactions or number of transactions per day (e.g. customer orders processed, stock reconciliations done, adjustments to stock levels, etc.)
- processing delays or queue lengths (e.g. allocated customer orders)
- cost per transaction by type
- performance against budget for each activity.

Other targets are required to complete the management of inventory activities. Prime external targets do not necessarily cover all the requirements for good inventory management, and there are additional aspects which are important as well: requirements for individuals to manage change and control projects; other business objectives to be met, including quality management, credibility, legal, motivational, health and safety, communication and process development. The real requirements for inventory control are numerous and, where possible, targets are needed to evaluate and monitor these objectives. Inventory management must also meet general business objectives. Targets need to be set in the areas of:

- personnel development planning
- improvements in systems and procedures
- quality development for products, service and administration
- management of change processes
- supplier audit
- environmental issues
- presentation, style and image of the department.

For well-managed inventory control, all the requirements of the market, the operating environment and the supply chain have to be considered. Greater understanding of the real objectives will improve the effectiveness of the business as a whole and the worth of inventory management within the business.

To reduce stock, keep value of receipts less than issues each week.

Purchase commitment

One way of maintaining control over inventory value is to balance the values of input and output. If the value of input is kept to below the value of output, then stock reduces (as discussed in Chapter 3, in the section on practical methods of reducing stockholding). The value of supplier deliveries should be less than the forecast value of customer dispatches in each time period. Supplier deliveries are a result of purchase order and schedule commitment, so it is the ordering process which has to be kept under budgetary control. The control process relies on collecting information on:

- the forecast value of demand for each time period (week or month)
- the amount of schedule or order commitment already made for delivery in each time period
- the typical value of last minute (emergency) orders which are committed each time period.

A policy on stock reduction, or increase, has also to be defined in terms of what value change in inventory level is considered practical or desirable each time period. From this information it is easy to calculate:

- 1 The maximum delivery value for each time period (= Forecast customer demand – Value of stock reduction target per period).
- 2 The maximum value of order commitment for delivery per period (= Maximum delivery book value – Typical value of emergency orders).
- 3 The maximum value available for new orders placed for delivery during that period (= Maximum value of order book commitment – Amount of order book commitment already made).

This maximum value available for new orders is the ceiling for total order value for delivery in the time period. If this value is exceeded by placing a further supply order, then that order should be rescheduled into a later time

Table 5.1 Order book values

Period	Supplier deliveries already ordered	Value of order book available to place
Next month	185 000	5 000
1 month ahead	130 000	60 000
2 months ahead	60 000	130 000

period (or another order delayed) or else the resulting stock value will increase.

Example

The current order book value with suppliers for delivery over the next three months is as shown in Table 5.1, and the maximum order book value for delivery in each time period is calculated in the final column. The task for the inventory controller in M. Tight Ltd is to ensure that further purchases are placed so that the planned deliveries next month are not more than £5000 and then £60 000 for one month ahead, and £130 000 for the month after that. By monitoring delivery book value in real time, an inventory manager can maintain the stock within the required budget. The control of stock reduction using purchase commitment is illustrated in Table 5.2. (See also Tables, 5.3, 5.4 and 5.5.)

Table 5.2 Inventory management targets – stock value

Location	This week/month			Last week/month			Target	
	Inventory value	Slow movers value	Weeks on hand	Inventory value	Slow movers value	Weeks on hand	Inventory value	Weeks on hand
Central								
NW store								
SW store								
NE store								
SE store								

Table 5.3 Inventory management targets – availability

Location	This week/month	Last week/month	Target	
	Availability (%)	Availability (%)	Availability (%)	Date
Central				
NW store				
SW store				
NE store				
SE store				

Table 5.4 Inventory management targets – supply

Purchasing activity	This week/month		Last week/month		Target	
	Value		Value		Value	% success
Orders placed						
Goods received						
Late deliveries						
Emergency orders						
No. of orders	Qty		Qty			

Table 5.5 Inventory management targets – outstanding demand

Location	This week/month		Last week/month		Target	
	Overdue	More than x days	Overdue	More than x days	Overdue	More than x days
Central						
NW store						
SW store						
NE store						
SE store						

Inventory valuation

Unit cost of stock lines

Accounting procedures for stock are usually organized by the accounting functions of a company. The costing options which are commonly used are explained below. It is important to have a consistent costing system which values all items. There are several ways of updating inventory cost when new items are purchased by changing the unit cost. The preferred costing method is average value, although standard costing provides better budgetary control. The costing method should not influence the actual items being picked, in almost all circumstances picking should be organized as a first in first out (FIFO) process.

Don't associate costing methods with individual stock items.

First in first out

Stock is valued at its purchase value. The oldest stock is assumed to be used first (as is required by good inventory practice) and the stock value is

therefore the total of the most recent purchases. First in first out is best used as an accounting procedure, but not for identifying which items in stores to pick. Stores stock rotation should be arranged through the warehousing control system.

Last in first out (LIFO)

The issues are valued at the most recent purchase price, leaving the remaining inventory valued at a previous (generally lower) value. The effect is to minimize the profit on the stock being sold and to minimize the value of the remaining inventory. This method of valuation can be used to reduce the profit reported by a company and decrease the valuation of the stock. It is, however, only a valuation technique and not appropriate for organizing stock movements. For example, there is a stock of ten of an item valued at £80 each, and a new batch of five has been bought at £120. The stock value is therefore £1400. If a customer buys two, then the stock will be reduced to thirteen. By LIFO these two issues will be valued at £120 each, and the new stock value for the thirteen items will be £1160. If the items sell at £150 each, the recorded gross profit is £60. (Using FIFO, the remaining stock of thirteen would be valued at £1240 and the gross profit would be £140.)

Replacement value

Stock is valued at the current price for buying replenishments. This enables a simple rule to be applied to sales price, (such as 'sell at purchase price plus 50 per cent'). It increases the value of stockholding but requires market information to manage.

Standard cost

Fixed item value is calculated financially and held the same for a long time, usually a year, giving a good stable valuation. Differences from the standard cost are considered as good or bad variances. Standard costs are used widely in manufacturing, as they can account for typical costs and overheads.

Average value

A running average is the safest stock valuation method. The cost of new deliveries is added to the total stock value and the total value spread over

the new total stock. For example, there is a stock of twenty of an item with a unit cost of £10. Ten more are purchased at a cost of £13 each. The total value is now $20 \times £10 + 10 \times £13 = £330$. The stock is now thirty, so the average cost is $£330/30 = £11$. This is the new average cost and any demand will reduce the stock value by £11 per item.

Stock-keeping unit (SKU)

Stock-keeping units are especially useful in distributed warehousing. These SKUs acknowledge the fact that the stock value can depend on where the item is held. If a customer requires an item which can be supplied either from the main stores or from a satellite store and the cost of distribution is £15 from either source, where should the supply be made from? Since the stock in the satellite store has already been transported once (probably at a cost of £15), then it would be cheaper to supply from the main store since the satellite stock has already cost more. Stock-keeping units are stock at a specified location taking into account the transport and purchase costs. It is normal to talk about a 3000 SKU warehouse, which is equivalent to saying that there are 3000 lines in that remote location. The actual stock value element of cost in an SKU could be calculated by any of the preceding methods.

SKUs take into account logistic costs.

These alternative valuation methods are illustrated by an example in Table 5.6. The ‘cost of the next item to be issued’ is the unit value of the item remaining in stock which will be issued next to fulfil customer demand. Each of the valuations give slightly different unit costs. For LIFO and FIFO methods, allocation of the costs depends on working out from which supplier delivery the next sales item should be taken.

In Table 5.6 the standard cost set by the accountants is £24, but purchase costs are higher. There will be a purchase ‘variance’ when the outstanding orders are received:

- Order 1 $10 \times £2 = £20$
- Order 2 $10 \times £4 = £40$
- Order 3 $10 \times £0 = £0$

At a demand of ten the total (cumulative) variance = £60.

Table 5.6 Alternative methods of determining unit stock value

Date (week)	Transaction			Stock balance (quantity)	Cost of next item to be issued					
	Type	Quantity	Unit cost		FIFO	LIFO	Standard cost	Average cost	Replacement value	
1	Stock check ¹		10	20	10					
	Customer demand	3		7	20	20	24	20.0	20	
3	Supply order 1		10	26						26
	Customer demand	2		5	20	20	24	20.0	26	
4	Supplier delivery		10	26	15	20	26	24	24.0	26
	Customer demand	4		11	20	26	24	24.0	26	
5	Customer demand	1		10	20	26	24	24.0		
	Supply order 2		10	28						28
6	Customer demand	6		4	26	26	24	24.0	28	
	Customer demand	3		1	26	26	24	24.0	28	
7	Supplier delivery		10	28	11	26	28	24	27.6	28
	Customer demand	4		7	26	28	24	27.6	28	
9	Customer demand	4		7	26	28	24	27.6	28	
	Supply order 3		10	24						
10	Customer demand	2		5	28	28	24	27.6	24	

Note: ¹Existing stock of ten is valued at £20 each.

Aggregate stock valuation

Stores stock can be valued either at full value given by costing at one of the above costing methods, or by a written down value to acknowledge the fact that some stock is no longer usable or saleable. There must be an amount of money deducted from the stock value to account for this. There are two ways of achieving this: depreciation and provision.

Depreciation

The first is to depreciate the value of stock of each item. The longer an item is in stock the less it must be saleable. The rate at which depreciation is allocated will depend on the type of stock. Fashion goods can lose value very fast. Aircraft spare parts depreciate very slowly. The simplest and most common way of calculating depreciation is to deduct a portion of the item value each year it remains in stock. Taking 25 per cent off the original value each year would mean that items have no value after four years, and would call forth acknowledgement by the accountants that inventory control have made the expected mistakes and have squandered company profit on useless stock. It is not a healthy situation when inventory managers are forced to minimize the value of stockholding because of poor control.

Depreciation steals profit to cover for mistakes.

Writing down the value of non-moving stock is unfortunate but essential. The accounting code of practice requires stock to be valued at 'the lesser of its purchase cost and its realizable value'. This means that items bought at a discount have to be valued at a discount, and items which have no likely use have to be valued at zero value.

Provision

A second way of reducing the value of stock is to allow a provision for the total value likely to be lost in a year. This provision can be based on a detailed depreciation calculation or a simple assessment. Whereas depreciation is applied to individual items, provisions are applied at a global level. This enables the provisions 'reserve' to be used for whatever accidents happen to cause stock obsolescence (sudden changes in legislation, contract loss, etc.).

Provisions are financial judgements. They can therefore be used in situations where there is an impending change which will affect inventory, but which is not predictable from historical information.

Professional stock controllers do not use write down as a method of reducing stock value because this only transfers funds from company profits to pay for the write down. (The proper methods for reducing inventory value were discussed in Chapter 3 in the section on how to reduce stocks successfully.) For stock control purposes it is therefore good inventory management practice to value stock at its full value and to apply the depreciation reserve or provisions later for the financial accounts. This ensures that the slow-moving, high-value stock receive some priority in the control activities.

Skills and systems

Systems should not only be designed to perform a task, but also to enable the user to achieve the objectives most conveniently. The term ‘user-friendly’ is often applied but seldom true, and the usability of a system should be geared to suit the type of person expected to use it, and the knowledge, skill and time they are likely to devote to it. Where a complex system is in use, all operatives must be fully trained to understand it, otherwise they will not use it efficiently.

The more professional the user, the more complex the computer system can be and the better the results achieved by its operatives. This is true from the user options point of view. A routine operator may only require an uncluttered transaction screen with the system set up for fast input or enquiry, showing the relevant information and nothing else. However, if information is to be accessed by less frequent users, they will still need to have a thorough understanding of what the system can be made to do. If they do not understand it, they will not be able to use it effectively. Further, although systems should be straightforward and user-friendly, their underlying logic should be complex enough to cope with the many variables which can occur. For instance, a user with a seasonal demand pattern will have to disregard a system only providing a simple minimum stock level to help in reordering. What they require is a seasonal forecasting system using percentage weighting, base series or Fourier analysis. Just as it is not necessary to understand the circuit for electronic ignition if you want to drive a car, so it is not necessary for the user to understand underlying computer logic as long as the method can be applied, the answer is good and the forecast type is easy to switch on and off.

Two distinct features for choosing a successful system are:

- apparent simplicity geared to user skill
- system sophistication geared to the importance of the application.

Successful implementation develops the right balance between systems, procedures and the user’s knowledge of a particular application. As an example, take a computerized parts list. A customer requires an item but does not know the part number. The old hands in stock control with many years experience will just quote the item number and look at the stock record, because they know it. A new starter requires support in finding the number. How can this be achieved? There are basically three methods, namely:

- training
- support procedures
- system features.

These different approaches are illustrated in Figure 5.2.

Training

If the user has an understanding of the application and memory of the codes, the system will work effectively. Training is essential for providing the understanding and knowledge. Inventory personnel need specific job

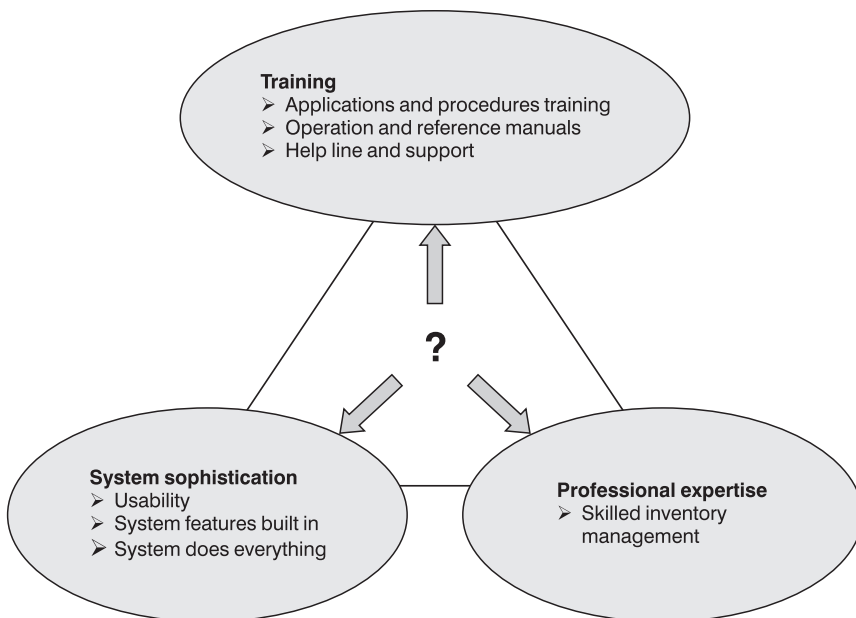


Figure 5.2 Inventory control options

skills which enable them to gain the full benefits of their systems and contribute to operating effectiveness. Investment in training is minuscule compared with the direct inventory savings.

Support procedures

Where the user of the system is unlikely to have knowledge of the application (e.g. technical knowledge by sales clerks), the company can provide the necessary information through supporting procedures. These can be paperwork or other aids which enable the user to follow a logical sequence with which to provide or gain the necessary information.

An example of this could be the person who often receives enquiries from customers needing advice on which stock item to buy. This person has either to gain sufficient technical knowledge of the stock to give advice or pass the item on to an expert, or to have the information provided for them in an appropriate form. A simple method of supporting this person is to provide a 'decision tree' for routine queries consisting of a set of questions. The answer to one question directs the user to the next question until the query is resolved. A person without detailed knowledge of a subject can still be effective using such support procedures.

Development of these procedures has to be overseen by a highly skilled person. Thus there is an initial cost, but the procedure can then run and advice can be given without the same level of expertise. This is a cost saving where the situation occurs regularly.

System features

Stock management systems can be considered as either stock-recording systems or stock control systems. Stock-recording systems concentrate on physical movements, recording issues, receipts and orders. Many of these systems also value the stock and keep a transaction history. 'Maximum' and 'minimum' levels can be input by the inventory controller and these levels and the order quantities are left to the discretion of the user.

*Make the system produce relevant information,
not large lists.*

Stock control systems take this data and provide decision support for inventory managers. Stock can be analysed into ABC classes and the historical information used to calculate the (weighted) average demand and the variability. Once policy on customer service and supply lead times has been decided, the system calculates the safety stocks and order quantities.

The outputs of a stock control system are:

- purchase schedules and orders
- expediting notes for suppliers
- management control information
- exceptions – unfulfilled customer orders
 - forecast predicting the wrong demand level
 - wrong supplier lead times
 - excess and obsolete stock analyses
 - options for special investigation identified by the user.

A successful system will produce information which can be used without amendment by the inventory controller. Some data may need modifying because the full information was not provided to the system, but this should be rare if the system is operated properly. In practice, most systems fail to provide this full specification resulting in time being wasted on routine analysis and interpretation. This time could be much better employed on customer and supplier liaison, and on ensuring that the system parameters are accurate. It is therefore important to use a system with stock control features, and to use these features to manage the majority of inventory.

Key points

- Identify the responsibility for stock management.
- Set quantitative targets (cover all aspects).
- Value all inventory.
- Design systems and information for the user.

Safety stocks

- Analysis of historical data
- Demand distribution
- Safety stock calculations – MAD and SD
- Customer service factors to set availability levels
- Determining obsolete and excess stocks

Learning from history

Collecting historical information

Stock is held either because it is convenient to buy in bulk, or because the item is required faster than the supply can provide it. In the latter case there is some uncertainty as to the quantity required, and so some safety stock is needed. In Chapter 2 the balance between the conflicting requirements of good service, low inventory costs and small operating costs were discussed. Now the means of achieving the controls can be examined.

The amount of safety stock held in an organization depends upon three main factors:

- the variability of demand
- the reliability of supply
- the dependability of transport.

The general approach to this situation is to set stock levels to cover the normal variability of demand and to adjust the other two variables so that they are relatively insignificant. Quality initiatives have improved the supply situation but it is usually found that the major uncertainty is caused by customers and their unpredictable requirements.

Store records show the movement of stock in and out of the warehouse and such historical information is essential for evaluating what level of stock to hold (unless the customer provides firm orders). However, the best guide

to the appropriate stockholding is the amount of demand, rather than the number of issues.

Note: When using demand history in forecasting, it is conventional to neglect the demand information from the current, incomplete period unless the demand is extremely large.

Individual stock movements logged on the stock records have to be analysed into movements per time period, normally weekly, so that weekly usage statistics are produced for each stock item. In cases where the demand is small or very variable, then monthly periods (time buckets) can be used. For fast-moving and rapidly varying demand (such as chart music), daily time buckets are required.

Collect data into time buckets.

The first step in analysing demand history is to collect the total demand for each period, as shown in Table 6.1. Sometimes such a table is used to estimate the stock levels and order requirements, whereas a little more understanding would enable accurate stock levels to be established.

Table 6.1 Demand history

Period	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Demand	23	35	28	19	34	25	41	15	39	33	28	48	31	38	28

The demand pattern can be displayed on a frequency distribution graph or histogram which shows how many periods the demand was at a certain rate. Figure 6.1 shows that the demand was between twenty-five and thirty at three times in Table 6.1 and it was between forty and forty-five once. This histogram gives a good idea of the usual demand rate and the spread of rates. In Figure 6.1 rate intervals of five per period were chosen (11–15, 16–20, 21–25, etc.) and the figures on the horizontal axis represent the maximum figures in the range). As can be seen the most common demand rate is 31–35.

If the demand comes from a similar type of customer, then the variations in orders will generally follow the same characteristics, although they are unpredictable. These demands form part of a population. If this assumption is correct, then the properties of this whole population can be used to predict the demand pattern. The most common demand pattern is the




Figure 6.1 Demand distribution

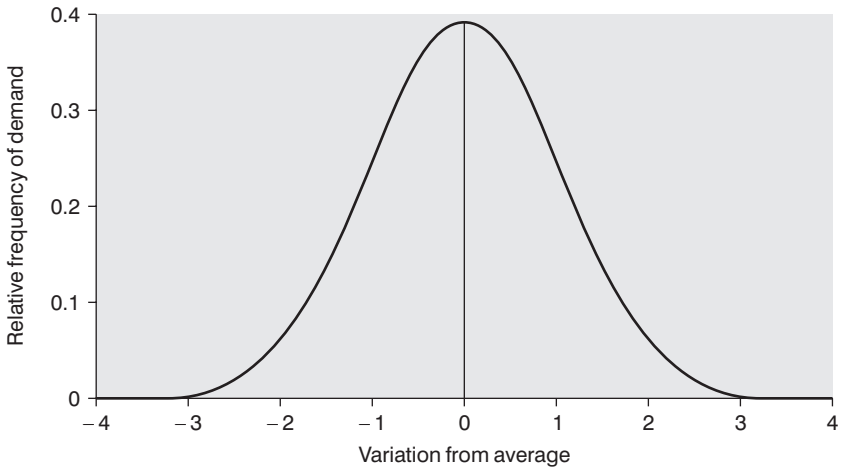


Figure 6.2 Normal distribution

Gaussian or 'normal' distribution shown in Figure 6.2. This is an idealized form of Figure 6.1 for a very large number of periods of demand. Using the small amount of information given in a table such as Table 6.1, the shape of the curve in Figure 6.2 can be calculated quite accurately.

Demand patterns can be assumed to be 'normal'. (Even if they are not, taking them to be normal is much better than simpler methods of control.) As can be seen from the distribution graphs, most of the time the demand is near to the average. Sometimes, but less often, it is much more or much

less. The vertical scale of the graph is frequency – how often that value occurs. The horizontal scale is really measured from the middle (the mean) and shows the values at which the populations occur. In practice, careful batching of the data is needed to produce a distribution like the one shown. In general, there is just a small amount of data to work with and this is used to fit a normal distribution.

A key property of the normal distribution is that it can be described using only two parameters, the average and the average width. Calculation of the average value is relatively simple. Measuring the width is more difficult.

Analyse the demand into true average and variability.

It is no use taking the highest and lowest figures, as these are usually the result of atypical situations (stockouts, one-off demands, holidays, etc.) It is better to take values which are, say, a quarter of the way in from the extremes of the sample (quartiles) or apply a more rigorous approach which uses all the available data. In general the measure of width used is standard deviation (SD), although in stock control a simpler approach is used, namely mean absolute deviation, MAD (see section on standard deviation below).

Of course more complicated demand patterns do arise. Occasionally there is more than one Normal distribution pattern mixed in the demand for an item. For example, a product could be sold in a local market in single units, whereas for export it is boxed in fifties. The demand pattern will then consist of two normal distributions, one with small fluctuations in the demand pattern, and the other with a more sporadic pattern.

Reasons for safety stock

The normal amount of stockholding can be determined by statistical methods which rely on history to predict the future, and assume no change in circumstances during periods ahead. Safety stock is primarily to cover random variations in demand, but it can also cover many other situations such as:

- supply failure
- production shortfall
- transport failure
- slow, unreliable or inaccurate information
- and any other source of disruption of service.

Safety stock is the buffer between supply and demand. It decouples customer service from manufacturing, and enables each to operate independently and more effectively.

Normal demand patterns

Stock availability

There is the need first to identify the major demand and supply variables discussed in this chapter and to quantify their variability. The need to cover for changeability of demand is a central task of inventory control. Unexpected demand requires a level of safety stock to cover for it. In fact using the normal distribution, the unexpected demand can be anticipated, or at least taken into account.

The more safety stock the better the customer service. Going back to the sample data in Table 6.1, a stock of 100 would probably be far too many and a stock of thirty-five too few. With a stock of forty-five it looks as though there will only be a shortage once in fifteen periods. But a surer way of finding this out would be to use all fifteen sets of data to give an accurate normal distribution, and then to find the risk of running out of stock.

Consider a supplier who arrives at the start of each period and tops up the stock to the required amount. One of the items supplied has the usage history shown in Table 6.1 and a unit cost of £1000. If the average usage amount, thirty-one, is held in stock at the start of each time period, then the expectation would be that the stock would run out during half the time periods, since the demand exceeds the average half the time. (This happens seven out of the fifteen times in the sample data.) If there were no stock at all then the stock would run out every time, and of course, if there were an infinite number in stock, then there would be 100 per cent availability. As the stock increases, the availability increases, and this relationship is shown in Figure 6.3. As safety stock is added above the average demand, the availability increases rapidly, but once the availability becomes high, then it requires a lot of stock to improve the situation.

Standard deviation

There are several ways of finding the typical variability, but there are two measurements which are generally used, namely standard deviation and mean absolute deviation. These are different ways of measuring the same thing, just as wealth can be measured in dollars or pounds, but the total value is the same. There are relative merits for both measurements.

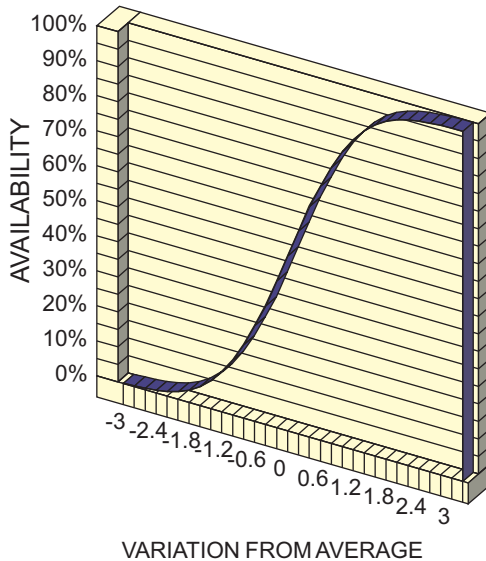


Figure 6.3 Stock and availability

Standard deviation is theoretically the correct measurement, but people find it more difficult to understand and work with. Standard deviations should be built into computer systems. Mean absolute deviations are much more user-friendly. They are easy to understand and, arguably, easier to bias. (The reason for wanting to bias them will be discussed in Chapter 10.) Mean absolute deviation is a good enough measure in practice. The standard deviation will be described first.

*Use SDs if they are in the software
or if you like mathematics.*

The basic premise for safety stock is that the differences between the forecast and the actual demand are random. These random variations in demand usually follow a ‘normal’ distribution (see Figure 6.2). This is symmetrical about the average value. The size of the variations differ from one stock item to another and, consequently, the width of the normal distribution differs also. This width, measured by the SD, is calculated from the sum of the squares of x , where x is the error, the difference between the actual value and the average (or the variations of the individual demand figures from the average). Mathematically this is:

$$SD = \sqrt{[(1/N) \times \sum xi^2]}$$

The SD is calculated by adding up the squares of the errors from a sample of N values; in our case periods of history. Using the history given in Table 6.1, this gives an SD value of 8.4. If there is only a little history, then N is low and the value of the SD will be relatively inaccurate. This inaccuracy (the standard error of the SD) is given by

$$\text{A standard error of SD} = \sqrt{\text{SD}/N}.$$

Mean absolute deviation

Use MADs for ease or to weight for more accuracy.

The MAD is a simple assessment of the variability of the demand pattern. Like SD, MAD is the basis from which safety stock is calculated. It is easiest to explain using an example.

Example

For a demand during successive periods of 5, 14, 6, 15 the average demand is 10.

The swings about that average are -5, +4, -4, +5.

Adding up these values gives 0, which shows that 10 is the true average. For measuring the variability, it is the size of the errors not the signs which matter, so forget the signs and add up the numbers. This gives 18 in total, or an average of 4.5 for the four periods, and is called adding the 'absolute deviations' or 'moduli'.

A second demand pattern has the form 9, 12, 8, 11 for the four periods of history, and an average of 10 as in the first example.

By the same analysis. the absolute deviations are now 1, 2, 2, 1 and the MAD 1.5 (6/4). Obviously the safety stock required in this case will be considerably smaller than in the previous example, in the ratio 4.5 to 1.5.

The MAD is calculated by this method, or by taking the last value of MAD and updating it with the latest absolute deviation. The formula for MAD is:

$$\text{MAD} = \frac{\text{Sum of absolute deviations from mean}}{\text{Number of periods included in the sum}}$$

For a normal distribution. MAD = 0.8 SD.

The MAD is easily calculated but to avoid the nuisance of recalculating each time, and to give a better value, an exponentially smoothed MAD, is normally used (see Chapter 10, section on weighted averages).

Choice of deviation measuring technique

With any statistical distribution, such as the difference between forecast and actual demand, there is a variety of standard ways for assessing the dispersion, such as quartiles, MAD and SD. Approaches using the greatest and least demands are unreliable since the extreme demands usually result from an abnormal occurrence. Basing the variability (and safety stock) on a stock cover calculation is a common, but incorrect assessment. This assumption gives completely the wrong stock balance.

There is a more serious debate about the relative merits of MAD and standard deviation. Mean absolute deviation is easy to understand, to calculate and to set up on a computer. Standard deviation is theoretically more correct, is used throughout statistics and is available as a pre-programmed function on spreadsheets. It is more difficult to calculate and understand.

For improved safety stock evaluation, the assessment of variability should be biased towards using the most recent data and the value should be exponentially weighted (see Chapter 10, section on improved values for MAD). This can be done very effectively with MAD, whereas the SD has to be converted to the variance ($= SD^2$) in order to use the weighting technique. Exponentially weighted MADs are therefore preferable to SD. Their use results in improved safety stocks, and offers a better selection of forecasting techniques to meet future demands (see 'Choosing the best forecast' section in Chapter 10).

The smoothing factor (see Chapter 10, 'Selection of the smoothing constant, α ') most commonly used for MADs is 0.1, since there is normally no rapid change in the demand variability. In practice this may prove to be a little insensitive and 0.15 is often found to work better.

For seasonal demands, the variations in size of the demands are likely to be greater at the maximum and smaller at the minimum. The MAD can be considered as a percentage of the demand level. Using this 'percentage MAD' approach enables an appropriate and up-to-date variability to be maintained and provides the appropriate safety stock levels. The use of exponentially weighted variances should give better values in theory, but is seldom used in practice. Table 6.2 shows a simple summary of the options. There is an arbitrary quality scale included to illustrate the relative merits of the techniques.

Table 6.2 Summary of deviation measurement techniques

Historical evaluation technique	Comment	Relative quality rating
Stock cover	Only for the amateur A major improvement on continuous inspection	1
MAD	Simple to understand and calculate Good for most situations	10
SD	Theoretically more accurate Available on spreadsheets	13
Smoothed MAD	Gives better results Also suitable for seasonal demand Easy to put into spreadsheets and software	25
Smoothed variance	Should give best values	30

For appropriate safety stocks, interpret the historical data, measure demand not sales, and respond to changing variability in demand.

Evaluation of safety stocks

Risk measurement

The reason for having to use SDs and MADs is that they are the only sensible ways to work out stock levels. The shape of the normal distribution curve measured in terms of SDs (or MADs) is always the same. This is the shape of the curves shown in Figures 6.2 and 6.3. These graphs indicate that when stocks are held at average usage, 50 per cent service is given. The other values behind Figure 6.3 are given by the customer service factors shown in Table 6.3. The customer service factors are shown in the table for both SDs and MADs. Either measure is acceptable and will give similar answers.

Adding one MAD's worth of safety stock to the average requirement means that customer service increases from 50 per cent to 79 per cent; or when stocks are increased by one SD, the service increases to 84 per cent, so there is a 16 per cent chance of stockout. By increasing the buffer stock to twice the SD, service is increased to 97 per cent. Additional stockholding increases the service toward 100 per cent but the cost in terms of inventory level is very high. Three SDs' worth gives 99.86 per cent and four SDs 99.99 per cent, But these are increases of only 2.1 per cent and 0.13 per cent respectively for the same investment in stock that

Table 6.3 Customer service factors (assuming a normal distribution)

Desired service level (% periods without stockout)	Multiply SD by	Multiply MAD by
50.00	0.00	0.00
75.00	0.67	0.84
79.00	0.80	1.00
80.00	0.84	1.05
84.13	1.00	1.25
85.00	1.04	1.30
89.44	1.25	1.56
90.00	1.28	1.60
93.32	1.50	1.88
94.00	1.56	1.95
94.52	1.60	2.00
95.00	1.65	2.06
96.00	1.75	2.19
97.00	1.88	2.35
97.72	2.00	2.50
98.00	2.05	2.56
99.00	2.33	2.91
99.18	2.40	3.00
99.50	2.57	3.20
99.70	2.75	3.44
99.86	3.00	3.75
99.90	3.09	3.85
99.93	3.20	4.00
99.99	4.00	5.00

gave 34 per cent for the first SD and 13.6 per cent for the second. This raises the question of which service will be too expensive to give, and whether the customer can tell the difference between 99.86 per cent and 99.99 per cent availability.

The availability level is measured by how often the designed stock level will meet all the demands during the period. (Other measures would be 'how many items were short during a year?' or 'how many customer orders were not completed?')

In Table 6.4 the stock is unable to meet demand in three periods. With review level at 5.0, average stock is 3.7 units (ignoring negatives).

In Table 6.5 all demand is fulfilled with safety stock of two MADs (four items). With the review level at 9.5, average stock is 6.5 units.

An example using stock records

Taking an example, Table 6.4 shows a conventional stores movement record for one item, with receipts, dispatches and supply orders for a stores. The average usage per period is 2.5. If the policy is to purchase sufficient stock to cover three periods, then the order quantity will be 7.5, rounded to 8. If it takes two periods' lead time for the stock to be replenished, then replenishment orders need to be placed when the stock falls below 5 items. This is called the 'review level'¹ because when the stock falls to this level the decision should be taken on whether to order. All too frequently an order is raised at some arbitrary 'minimum stock level' which bears no relationship to the current usage rate. As a result of assessing the demand, a purchaser will change the review level, place an order, or both.

The stock levels resulting from this are shown in the 'stock' column in Table 6.4. The orders are raised each time the stock falls below 5. (This review level is taken as a constant over the short time shown.) Notice that there is a shortage in periods 6, 7 and 10, resulting from the large number of sales in period 6 which coincided with the stock being below the stock review level. It is prudent to hold safety stocks to cover for these occurrences.

The amount of stock can be assessed from a statistical analysis of the demand pattern as discussed in the previous sections. The total sum of the variations from the mean is 20 for the 10 periods (demand variability in Table 6.4.) This gives a MAD of 2. If the safety stock is now increased by 4 from 5 to 9, this should give added protection of 2 MADs. ($2 \times 2 + 4$). This corresponds to a 94.5 per cent safety stock. The stock review level is also increased from 5 to 9 and the orders have been placed immediately at this level. The new pattern is shown in Table 6.5. The extra two stock has reduced the number of depleted weeks from 2 down to nil. It also meant that it was possible to supply all of the 25 items needed. With no safety stock it was possible to issue only 21. This calculation will be refined below.

Of course, from the statistics, this level of safety stock will cause a shortage in 5.5 per cent of the order periods and this may not be good enough. Additional safety stock can be added to give the service level required. The higher the safety stock, the more costly is the stockholding, not only because money is invested in the stock, but also because an investment in stock may prove fruitless if the item is superseded or becomes obsolete.

Note: ¹ Up to this point it has been referred to as 'reorder level'.

Table 6.4 Stock record

Period	Demand	Stock	Purchase orders	Demand variability
0		8		
1	1	7		1.5
2	5	2	Place order for 8	2.5
3	0	2	↓	2.5
4	3	7	Delivery of this order	0.5
5	2	5		0.5
6	6	-1	Place order for 8	3.5
7	1	-2	↓	1.5
8	0	6	Delivery of this order	2.5
9	6	0	Place order for 8	3.5
10	1	-1	↓	1.5
Average	2.5	37	MAD	2

Table 6.5 Stock record with safety stock

Period	Demand	Stock	Purchase orders	Demand variability
0		8		
1	1	7	Place order for 8	1.5
2	5	2	↓	2.5
3	0	10	Delivery of this order	2.5
4	3	7		0.5
5	2	5	Place order for 8	0.5
6	6	7	↓	3.5
7	1	6	Delivery of this order	
8	0	6	Place order for 8	1.5
9	6	8	↓	2.5
10	1	7	Delivery of this order	3.5
Average	2.5	6.5		2

Additional safety stock

Safety stock is normally calculated on the basis of unknown demand. The quality of supply from the manufacturers may also prompt an increase in safety stock, either because the lead time for an item may vary between placing one order and the next, necessitating extra stockholding, or because the supplier's delivery performance against the lead time is unreliable and delivery is more often late than early. An allowance in stock has to be made for this factor. If stock has to be held against poor supply, then theoretically

a supply performance MAD could be calculated and a safety stock worked out. This supply safety stock would not be additional to the demand safety stock, since the same safety stock would cater for either high demand or late delivery, but not both. It is found that the addition of MAD from different causes can be added using Pythagoras' theorem (in a similar way to adding variances).

Make the supply effective. Don't hold stock to compensate.

In dealing with the supplier, the supply pattern should be a result of compromise for both organizations. Whoever is the cause of the extra safety stock should incur the cost of it. Minimum stock is held when each party can trust and accept the information provided by the order.

Slow-moving items

For a large number of stock items the usage rate is low. This includes obsolescent items, non-standard lines and spares. For these the 'normal' distribution is squashed against zero usage and is replaced by a skewed distribution. The chance of an item being required is governed by Poisson's distribution law. Compared with the normal distribution, the Poisson distribution has a longer tail where there is an enhanced (but small) chance of high demand during a period. To cover for this, a higher stock would be required where high levels of availability are needed.

Demand for slow movers is a Poisson distribution.

An important feature of the Poisson distribution is that

$$\text{Standard deviation} = \sqrt{\text{Average demand}}$$

This relationship simplifies the review level formula for slow-moving items. When using the Poisson approximation, the usage events have to be independent (i.e. issues in ones to unconnected demand). If this is not the case, then

$$\text{Standard deviation of sales orders} = \sqrt{\text{Average number of sales orders.}}$$

Using the calculations developed in this chapter and applying the lead times discussed in the next provides the basic method for inventory management for the majority of stock.

Excess and obsolete items

In practice, stocks are never balanced ideally and some items become excessively slow moving. These items can be described as 'excess' or 'obsolete'. One of the areas which generates little interest and lots of cost is obsolescent items. The definition of obsolete should be that there is no use for the item. This information is normally only available when the customer changes to a direct and improved substitute. In many instances, goods are classed as obsolete as a result of a long period with no demand: so the practical definition is usually that an item is obsolete if it has not moved for 'x' months.

Have the system identify obsolete items, then remove them.

The selection of 'x' varies with the type of item and market, and the position in the supply chain. It is relatively easy in fast-moving consumer goods to get rid of items which have not moved for three months and this may well be the definition in this environment. Suppliers of heavy machinery and spare parts, where world demand has been zero for two years, may not regard demand to have finished and they will still desire to keep stock (only one!). In this case failure to supply quickly may have important strategic or economic effects, and supply lead times are long. A mechanical part on a machine which fails every seven years or so is not obsolete, but has a very low usage rate. The decision to purchase such slow-moving stocks is one problem, the decision to throw away existing stock is entirely another. Each company has to decide on a proper obsolescence policy for disposing of non-moving items.

Excess items

The 'not moved since' analysis can cope with non-movers, but there are other useless items, namely the excess items. If the stock of an item is 1000 and the usage is one per month, the item is certainly not obsolete, but may be before the stock is used up. Therefore the definition of excess could be 'Excess items have stock of more than y months' usage'. Again the value of 'y' depends on product, market, position in the supply chain, and also on the Pareto class of the item, since the stockturn definition is a limit on the stock investment.

Figure 6.4, with cut off values 'x' and 'y', shows simple control of excess and obsolete stock. It allows accountants to identify how much money to put aside to cover likely stock disposal. Writing down inventory value is, of course, an admission by the stock controllers that they failed!

In theory the values of x and y could be different for each item. For example, a stock of helicopter rotor assemblies could be held in stock for a long time before use, but there is a mandatory requirement to change and repair these units, so the stock is not obsolete, just slow moving. Whether stock is excess depends on whether the items are now units (bought too early) or repaired units (which would be in inventory anyway). In practice, therefore, the cut-offs shown in Figure 6.4 are used, and some allowance is made for items with stock of one.

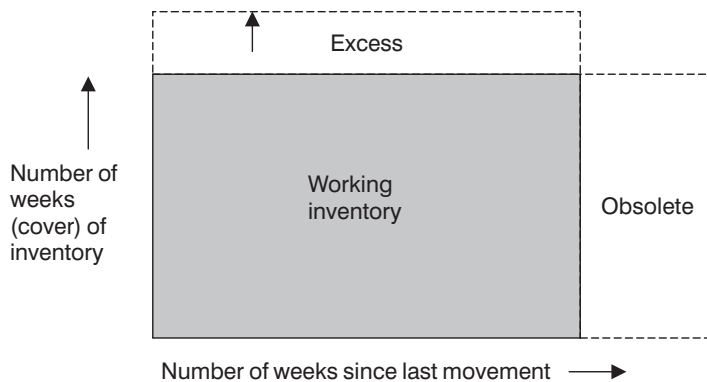


Figure 6.4 Definition of excess and obsolete stock

The first step in dealing with the phase-out of stock lines is to identify early the lines in this terminal condition. These lines can then be included in the ABC classification as 'O' or other suitable flag to inhibit normal ordering. The disposal of excess and obsolete stock has to be carried out carefully and professionally to avoid a large loss in profits. Before selling a company, the value of obsolete stock items has to be reduced since it is illegal to overvalue the company and the stock can be a major part of this value.

Key points

- Safety inventory compensates for unknown demand.
- Supply problems should be removed, not buffered.
- Demand can be split into average and variability.
- Variability (safety stock) should be calculated using MAD (or SD) but not estimated.
- Customer service factors should be used to decide the level of availability necessary.
- Identification of obsolete and excess items will enable old inventory to be used up more effectively.

Setting the right stock levels

- Order patterns with safety stock
- Successful calculation of review level
- When to expedite
- The effect of supply lead times
- Co-ordinating supply with target stock levels

Simple assessment of review levels

The basic calculation

Our aim in stock control is to maintain a balanced inventory so that customer service for each stock item is maintained within its proper limits. The time at which we can influence the stock levels most effectively is when we order, and it is at this point that the major opportunity occurs for ensuring a balanced inventory. This is the major control mechanism for ensuring a balanced inventory. It is also the point at which customer satisfaction or excess inventories are created. The technique for calculating the right time to order inventory is crucial to balancing these conflicting pressures.

Originally stock-recording systems used to include stock control levels as ‘minimum’ and ‘maximum’ stock levels. These basic concepts are still applied today as a simplistic approach where tight control over stock is not necessary. The use of a minimum stock level for order control is not sensible, since the minimum occurs immediately before delivery. Items have to be ordered well in advance of this, so control is through a ‘reorder level’ not a minimum stock. However, a minimum stock level is vital for ensuring that there is warning of low stocks. In a stock control system, therefore, there is a need for both reorder level and minimum stock level. These are called the ‘review level’ and ‘safety stock’ (see Figure 7.1).

The review level is triggered by the information from the system that stock is low (i.e. it has fallen below the ‘reorder level’) and makes a decision




Figure 7.1 Stock controls

as to whether to order. Where the demand has increased or lead times increased, orders are necessary; where the lead time is reduced or demand has diminished, then there is no need to purchase. Since demand can go up or down, the trigger level is treated in practice as a review level.

Minimum stock has also been replaced by the safety stock which is calculated as discussed in Chapter 6, and allowing for lead time (the analysis is shown below). Maximum stock is a result of negotiation with suppliers on delivery quantities. The maximum stock level should be a result, rather than a direct control. It is therefore necessary to control the delivery quantities directly by using three parameters to carry out simple and effective stock control, as shown in Figure 7.1, namely:

- review level
- safety stock
- delivery quantity.

The starting point is the calculation of the safety stocks. A quick look at the historical demand for the two items illustrated in Figure 7.2 shows that the second will require a great deal more safety stock than the first if the customers are going to be equally satisfied. The average demand rate is the same in each case. The more the variability of demand, the more safety stock is required. Therefore the safety stock does not depend on the demand rate. (This means that safety stock calculated as number of weeks' demand is simply wrong.) The safety stock for periodic delivery items is

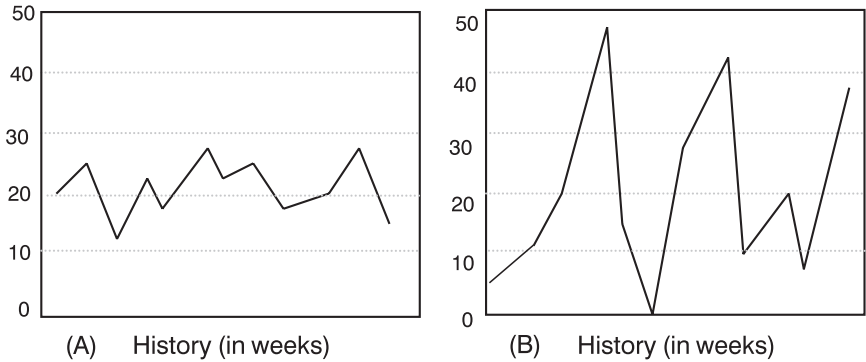


Figure 7.2 Historical demand patterns

given by multiplying the customer service factor by the standard deviation (or MAD) of the demand during the period.

Safety stock depends on demand variability and not demand rate.

If the demand is thirty per month and the supply lead time is two weeks, then there had better be fifteen in stock or on order or else there will be a shortage. This is because the time to order is given by:

$$\text{Review level} = \text{Usage rate} \times \text{Supply lead time}$$

and to be safe:

$$\text{Review level} = \text{Usage rate} \times \text{Supply lead time} + \text{Safety stock}$$

The safety stock calculation discussed in Chapter 6 showed that the stock depends on the variability of demand (MAD or SD) and the availability level required (customer service factor). Although the effect of supply lead

Table 7.1 Variation of safety stock with lead time

	Per day	Per week	Per month	Per year
SD scaled proportional to lead time	0.4	2	8	100
SD scaled as the square root of lead time	0.9	2	4	14

time was not discussed, the safety stock also depends on the supply lead time.

Example

For an item with a demand of five per week plus or minus two (i.e. MAD = 2) what is likely to be the variation in demand per day, per month or per year? Two options are suggested for the answer in Table 7.1. The average demand per period increases linearly with time period. The uncertainty of the demand, however, does not increase at the same rate. The average is more likely to be achieved over a long period of time than over a short period. The annual sales of an item are often similar to the year before, but from day to day the demand will vary widely. The weekly sales will also be quite widely dispersed but the monthly results are quite consistent. If the SD (or MAD) were to increase in proportion to the time period (see Table 7.1), then the inaccuracy in the long-term forecast would be great. However, because of the 'compensating errors' over longer time periods, the deviations are relatively small. In fact it is found (and can probably be proved) that the SD only increases as the square root of the time period. This is shown as the last line of Table 7.1.

This same effect can be illustrated in Figure 7.2B. A short sample could give a wide variety of answers to the average use, whereas a longer sample will give a more reliable answer. In other words, the longer the sample period (in the case under discussion – the lead time) the less the relative safety stock need be. From Chapter 6 this would be a square root relationship. As the variation in demand is largely a statistical effect, the square root relationship occurs very frequently in stock control.

The full formula for safety stock is therefore:

$$\text{Safety stock} = \text{Customer service factor} \times \text{SD} \times \sqrt{\text{lead time}}$$

For values of customer service factor, see Table 6.3. This equation can then be used in the review level calculation above.

The full calculation of review level can be carried out for the item with the demand history shown at the start of Chapter 6 (Table 6.1). The average was thirty-one and the MAD was 6.8. Some practical information is now

required, namely, the supplier's lead time and the availability level to be offered to customers.

Assuming that the lead time is three periods and the required customer service is 90 per cent, the customer service factor for 90 per cent customer service is 1.6 MAD from Table 6.3. First the safety stock has to be calculated:

$$\begin{aligned}\text{Safety stock} &= \text{Customer service factor} \times \text{MAD} \times \sqrt{\text{lead time}} \\ &= 1.6 \times 6.8 \times \sqrt{3} \\ &= 18.84\end{aligned}$$

From this the review level can be calculated:

$$\begin{aligned}\text{Review level} &= (\text{Usage rate} \times \text{Lead time}) + \text{Safety stock} \\ &= (31 \times 3) + 18.84 \\ &= 111.84\end{aligned}$$

This means that if the stock falls to 111, then more need ordering (assuming there are not enough on order already to cover the 111). If it is decided that 90 per cent service is not good enough and 98 per cent is more appropriate, then the customer service factor becomes 2.56, the safety stock 30.15 and the review level 123.15.

In reality, all the items do not need to be in stock to provide availability for customers as long as they are on order and due for delivery before the current stock runs out. In the example above, if the stock is currently sixty then there is a problem. However if there are already 100 on order, there is no need to raise another order. Therefore the condition for ordering is when:

$$\text{Current stock} + \text{Supply orders outstanding} < \text{Review level}$$

(To be exact, the supply orders which count here are only those which are due for delivery within the supply lead time.) The purpose of the review level is to trigger replenishment orders at the right time. Of course, triggering the order at the right time does not necessarily ensure that there is stock in the stores. The review level calculation just described is the most important piece of stock control theory.

Reordering triggers orders but does not eliminate stockouts.

When to expedite supplies

From the purchasing point of view the supplier has to deliver on time and should be chased if the delivery is late. From a stock control point of view, we rely on the supplier delivering at a fixed lead time. However, if the supply is late, it does not necessarily cause a problem if the rate of issue has not been as fast as expected (and the stock is still relatively high). In other circumstances the demand may be higher than average and it would be convenient to receive delivery early. In the example discussed before, a 90 per cent availability means that once in every ten time periods a stockout is to be expected. When this happens it would be useful to try to enhance availability by rushing in supplies. Therefore, a warning is needed for when the physical stock is getting low, not when the supply is late. As the supply should, on average, arrive when the stock gets down to the safety stock level (see Figure 7.1), then the safety stock level is a suitable trigger level to remind the supplier that goods should be arriving. Sometimes this reminder will be after the proper lead time (if delivery is late), or it could be before the item is due when the demand is above average. In either case there is still the safety stock to use before there is crisis. The condition for chasing supplier is therefore when current stock is less than safety stock. This should be a standard routine printout from the system used for automatic expediting.

Chase suppliers when you start using safety stock.

Managing lead times*Lead time*

Lead time is the time between a shortage occurring and items being available to maintain supply. The lead time for supply is a balance between what the supplier wants and what the customer wants. If the suppliers are powerful, they will specify a long lead time (such as a one year firm programme). If the customers have the upper hand, they will expect to be supplied immediately, without any commitment. Both of these extreme cases increase the total operating costs. The compromise is a lead time which gives the suppliers sufficient supply time but do not force the customers to commit themselves further than they can predict. The accuracy of their forecasts of requirements reduces rapidly the further ahead they look. Long lead times and large batch supply produce an order pattern which is irregular and subject to change. Stocks are needed where the sales lead time is shorter than the supply lead time. If this can be altered

so that supply lead time becomes less than sales lead time, the stockholding can be avoided.

Lead times for purchasing are governed by production and transport. Lead times for manufacturing result from capacity, material availability and production planning. Lead times required by customers depend upon availability and supply policy (make to order, assemble to order, stockhold).

Components of lead time

Lead time can be analysed into several components, some of which are essential and some avoidable. The supply lead time comprises:

- order review time – the intervals at which the low stock situation is reviewed (e.g. printout every day, hourly, weekly, etc.)
- order processing time purchasing and communication) – the time it takes to:
 - review the order
 - decide to buy
 - transmit to the ordering system
 - raise an order
 - gain the appropriate authority
 - inform the supplier (post, fax, e-mail)
- supplier lead time (vendors' manufacturing, buying and dispatch)
- transport time – transfer of the item from supplier to receiving bay
- receiving time – time taken for goods inwards and updating the stores records.

From this analysis, the actual supply lead time contains some elements of administration in addition to the lead time quoted by the supplier.

In stock control equations, the supply lead time is assumed to be constant. It is very important that the estimate of lead time in the system reflects the true lead time. If the supply lead changes, then the lead time in the calculation should be changed to reflect this. It is a continuous task and part of the maintenance of the stock systems. Changes can normally be made by groups of items from the same source, or by using the actual lead time to update an (exponentially weighted) standard lead time.

Keep the lead time realistic.

Setting the lead times

If the lead time on the control system is set too short, this will lead to stockouts, since the stock level will be kept too low. If the lead time is set overcautiously (long), excess stock will result. The compromise for setting the lead time is a fine judgement for the inventory controller.

Minimizing the lead time

Longer lead times make forecasting more difficult and increase the safety stock. Reducing the lead time is therefore desirable and can lead to stockless supply if the supply lead time is short enough. It is part of inventory management, therefore, to work on reducing the overall lead time.

One of the simplest ways of reducing the lead time is to specify it more precisely. 'Delivery in the month of June' as a request gives the supplier plenty of leeway. 'Delivery on 12 June at 11.18 a.m.' suggests an entirely different expectation, and is likely to gain a much better delivery performance (during the day required rather than the month). Numerous experiments attempting this have shown that a majority of suppliers will improve their delivery precision greatly, simply by being asked.

Companies also find that supplier reliability is often affected by lead time. With a lead time of one day, if the supplier of an A class item is a day late, then this can be a major issue. If the lead time is twelve weeks and the supplier is a day late, this is often expected, or even considered as on time. One day's worth of stock is held to offset this. The perception of delivery performance is as a proportion of lead time. However, from an inventory perspective, a day's worth of stock is the same value, whether it is added to an extra six weeks' worth or to one day's worth of total inventory.

From the analysis of lead time above, the total lead time should be minimized by:

- inputting information to stores and computer systems without delay (removing administration leadtime). It is important to minimize the time between:
 - the stock getting low in the stores and the supplier receiving the purchase order
 - receiving the delivery and the event being displayed on the company computer system

- partnership with major suppliers (giving schedule flexibility or kanbans).

Once these elements of lead time are small enough, (see also Chapter 8, section on purchasing procedures) the suppliers' lead times should also be examined. Supply lead time consists of actual physical transport time, manufacturing time in some instances and organizational time. Lead time should be a negotiated time, not something that happens without manipulation.

Example

A supplier quotes six-week delivery. What does this mean? It could be:

- *all the existing stock is allocated to someone else*
- *there is a stockout because the safety stock was not high enough or the reordering procedure broke down*
- *it takes six weeks to get these from their supplier.*

In the first two cases, the six weeks are caused by priorities, first, because the order is not jumping the queue and, second, because the supplier does not consider the order important enough to put in higher safety stock or improved systems. This is a matter for the supplier and the customer to negotiate. The third case could be argued as similar. If their supplier is a manufacturer, then the six weeks could reflect the manufacturing time. However, only a very little part of manufacturing time is taken up with production. The manufacturing time per unit assessed by work measurement (time being worked on) is much smaller than the production lead time taken for the work to pass through the manufacturing plant. This is usually more than ten times the measured work content. There is, therefore, potential for manufacturing lead time to be collapsed to a much smaller timescale. However, this cannot be done on a one-off basis; it has to be done gradually for repetitive products because of its effect on capacity balance and replanning. The amount of replanning manufacture has to be minimized in the short term or else the plant becomes inefficient and delivery performance suffers. Again, collaboration with suppliers may enable them to reduce their own lead time.

Because of the large numbers of items in inventory, reduction of lead time has to be carried out for selected types of items. These are initially:

- A class items
- items where a small change in lead time can make them into supply-to-order items.

Effects of lead time

Lead time does not govern stockholding. Order quantity is not affected by lead time. It is possible to order twenty weeks' worth of an item with one week's lead time and one week's worth of an item with a twenty-week lead time (as long as further orders are placed each week). True the review level is calculated using the lead time, but the review level determines only what order cover is required and not how much stock is held.

Lead time does affect the amount of safety stock to hold. The longer the lead time, the wider the absolute variation in demand and therefore the higher the safety stock for a given service level. Fortunately the amount of extra stock only increases as the square root of the lead time.

In summary, the effects of long lead time are generally:

- forecasting demand to cover the supply period is more difficult
- the reliability of the supplier may be poorer.

Errors in forecasting cause higher stocks to be held. If demand is certain over the lead time, then stocks can be very low. A short lead time means that:

- the forecast is more likely to be correct
- errors will be smaller.

Dealing with inconsistent lead times

With closer partnerships and higher expectations, inconsistent delivery is not the problem that it used to be. However, customers are now more sensitive to supply reliability, so strategies should be considered for any supplier where delivery is inconsistent and changing supplier is undesirable.

It is difficult to set reasonable stock levels or to provide good customer service if lead times are not known accurately. The policy on inconsistent lead times should be determined by the impact of the problem, and on

whether the financial risk is high or low. For high turnover value items (A class) the inconsistency has to be resolved through scheduling and working more closely with suppliers to smooth out the inconsistencies. The objective should be to smooth out the variations rather than to improve the updating, since each time the lead time is updated, the stock parameters will be changed and the order book modified.

For low turnover value items (C class) the situation is not so important unless there are very many inconsistent suppliers. The strategy for avoiding the problem for routine supplies is to use pessimistic lead time in calculations. This will cause higher stocks, which may be necessary if the item is important. However, as the purchase quantity is probably large, there are only a few occasions on which there is a chance of running out so C class items may not be worth the investment in extra stock to provide the extra availability.

In some instances lead times are cyclic. This is the result of the supplier having a fixed plan and not responding to customer demand.

Examples of this are where there is a logical sequence of production which cannot be broken such as size (steel-mill rolling plan – thickest to thinnest over several weeks), contamination (pure to impure in process industries, light to dark in injection moulding) or other cycles (food crops etc.).

In this case supply orders should ideally be placed when the lead time is minimum. Using target stock levels (described at the end of this chapter) the timing of orders can be arranged to achieve this. However, the process can only work if any changes to the supply cycle is communicated to the customer.

Effect of order frequency on safety stock

The safety stock calculated by the methods described earlier in this chapter are adequate for most practical situations. However the calculations can be improved because the risk of stockout is not constant all the time. Availability is good when there has just been a delivery. The risk of running out are greatest when the stock is lowest, i.e. while the goods are below the review level and a delivery is expected. The main risk during a year is, therefore, the risk during the lead times. This depends upon the order

frequency. The more replenishment orders for an item that are placed per year, the higher the risk and the greater should be the safety stock. This factor is usually small compared to the benefits of ordering more frequently and reducing the inventory costs and risks of obsolescence. A simple addition to the calculation above suffices for control, as shown in the following example.

More refined calculation adjusts for batch size.

Example

Assume that the historical data for an item is as follows:

- Annual demand = 10 000 (200 per week)
- MAD = 100 (weekly)
- Service level required = 94 per cent (2 MADs) (measured as number of items short per year)
- Lead time = 4 weeks
- Order frequency = 6 weeks

So the total stock shortfall allowed = 1 – (94 per cent of 10 000) = 600 per year.

If there are $50 \div 6$ orders per year (8.3 orders per year) then the allowed stock deficit per stockout is:

$$\frac{600}{50} \times 6 = 72 \text{ (assuming 50 sales weeks per year)}$$

$$\text{Usage in the lead time} = \frac{10\,000}{50} \times 4 = 800$$

The allowed stockout is therefore 72 in the lead time (LT) demand of 800 so

$$\text{Service level required} = 1 - \frac{72}{800} = 91 \text{ per cent}$$

91 per cent = 1.5 (MADs) but this is for a 4-week MAD since the period is the lead time, so

$$\text{Safety Stock (SS)} = 1.5 \times 100 \times \sqrt{4} = 300$$

and

$$\text{Review level} = 300 + 200 \times 4 = 1100$$

The average stock level is SS + 1/2 order quantity

$$= 300 + 1/2 \times 6/50 \times 10\,000$$

$$= 300 + 600 = 900$$

Now if the order frequency is changed to 10 weeks, then there are now 5 orders per year

$$\text{Allowed stock deficit per stockout} = \frac{600}{5} = 120$$

The usage in the LT = 800 (as before), but because there are less risk periods

$$\text{Service level required} = 1 - \frac{120}{800} = 85 \text{ per cent}$$

Eighty-five per cent is 1.2 MADs (i.e. 4-week MADs). Therefore:

$$\text{Safety stock} = 1.2 \times 100 \times \sqrt{4} \text{ (in weekly MADs)}$$

$$\text{Safety stock} = 240$$

$$\text{Review level} = \text{SS} + \text{usage in LT} = 240 + 800 = 1040$$

$$\text{Average stock level} = 240 + 1/2 \times 10/50 \times 10\,000 = 1240$$

The review level and safety stocks are less, but the average stock is more because of the larger order quantities. If the order frequency is changed to 4 weeks, the service level per stockout becomes 94 per cent, the same as the simple calculation. This is to be expected since there is continuous exposure to stockout risk. The safety stock becomes 400, the review level 1200 and the average stock 800.

Target stock levels

Application of target stock levels

The target stock level (TSL) is a maximum stock level which can be used to calculate the order quantities. Target stock levels marginally differ from

the review level approach to ordering: Review levels fix the order quantity and vary the order frequency, whereas TSLs fix the order frequency and vary the order quantities.

When ordering regularly from a supplier, this enables the stock controller to have a cyclic work routine, (such as always ordering from supplier X on Thursdays) and the TSL approach is therefore good to use for the regular demand items, and is recommended for A class products.

Target stock levels support efficient purchasing.

The TSL is a level to which the stock is topped up (theoretically) when the cycle time for ordering arrives. The process for setting up the ordering procedure is, first, to establish a routine for placing orders on the various suppliers by specifying the day each week or month on which the order is going to be dealt with and, second, on the allotted day, review the stock of all the goods supplied by that supplier. The amount to be ordered is given by the formula:

$$\text{Order quantity} = \text{TSL} - \text{Free stock} - \text{Supply orders outstanding}$$

Here the ‘orders outstanding’ are, of course, only those orders which are due within the current supply lead time. Orders booked further ahead are not included in the calculation.

Example

<i>TSL has been calculated as</i>	34
<i>Free stock is currently</i>	16
<i>Outstanding order due tomorrow</i>	5

Therefore

$$\begin{aligned} \text{Order quantity} &= 34 - 16 - 5 \\ &= 13 \end{aligned}$$

Calculation of target stock levels

The calculation of TSLs is very similar to that for review levels. An extra factor is included in the calculation of the TSL. The effective difference is

that stock replenishment now needs to cover the period until the review cycle comes round again. If the review cycle were a month and the supply lead time only one day, then the TSL would have to cover a month's extra usage (plus safety stock) because the next order would be placed only in a month's time. Thus the formula for the TSL is given by:

$$\text{TSL} = [\text{Usage rate} \times (\text{Lead time} + \text{Review period})] \\ + \text{Safety stock}$$

And there is a similar adjustment to increase the safety stock level:

$$\text{TSL} = [\text{Usage rate} \times (\text{Lead time} + \text{Review period})] \\ + \text{Safety stock}$$

$$\text{Safety stock} = \text{Customer service factor} \times \text{MAD} \\ \times \sqrt{(\text{Lead time} + \text{Review period})}$$

The demand 'in review period' is included because the stock level may not be checked continuously, or deliveries may only be made once per month. If we do not include the extra time in these processes then the risk of running out of stock is increased.

Example

Review period for an A class item = 1 week.

Supply lead time = 5 weeks.

Average demand rate = 4 per week.

MAD = 2.5 measured on a week basis.

*Required customer service = 90 per cent
(this gives a customer service factor of 1.6).*

Therefore,

$$\text{Safety stock} = 1.6 \times 2.5 \times \sqrt{(5 + 1)} \\ = 9.8$$

and

$$\text{TSL} = 4 \times (5 + 1) + 9.8 \\ = 33.8$$

Applying target stock levels

The TSLs are levels of stock to which the free stock never rises in a stable situation (see Figure 7.3.) The actual stock will rise to a maximum level which is

$$\text{TSL} - (\text{Usage rate} \times \text{Lead time})$$

The TSLs can be used for long review periods (e.g. one month) with short lead times or, of course, with short review periods (one week) and long supply lead times. In this latter case, there will always be one or more orders outstanding on the suppliers. This does not affect the validity of the calculation.




Figure 7.3 Target stock levels

It is important to keep on adjusting the TSLs to match the usage rates and lead times which are experienced, or else this approach can lead to last minute purchases and extra work.

It is also important to ensure that the stock does not run out without warning during the order cycle. This is why it is essential to monitor continuously to find those items where the stock is falling below the safety stock level, calculated in this chapter.

Key points

- Review levels should be continuously updated in line with demand.
- Review levels do not avoid stockouts; to do this requires expediting when physical stocks get low.
- Batch sizes are more important than lead time in keeping stocks low.
- Using target stock levels is good when buying many items from one supplier.

The changing role of purchasing

- The role of procurement
- Single sourcing
- Supply partnerships
- Vendor evaluation and rating
- Types of orders
- Calculating delivery quantities
- Effective purchasing administration

Modern supply practice

Improving the supply

The development of supply chain management changes the relationships between companies providing goods to one another. This has affected the role of the purchasing function. Purchasing has traditionally been an order-by-order process based on obtaining the best combination of price, quality and delivery (availability). Estimates are requested, best options are negotiated, orders are placed and delivery is monitored.

This approach has been eroded over the years and the time-consuming task of obtaining estimates is not now usual in routine material supply. The development of supplier relationships has introduced longer-term arrangements involving contracts, call-offs and schedules. This has changed the role of the purchasing function. Instead of individual orders, supply is negotiated less frequently. This has meant a different style of purchasing activity because low level routine negotiation is being replaced by high level contracts. The number of suppliers is also reduced.

Purchasing people are no longer order processors.

These changes have a major effect on the type of expertise required in purchasing. The traditional type of 'buying' is being superseded by a

combination of high-level contract management and supply scheduling as part of the material control activity. This means that materials management will absorb the purchasing operations as part of control and scheduling. Why is this now practical and possible? The answer lies in the current supply environment:

- total quality – supply to specification is expected
- supplier partnerships – work closely with supplier
- logistic control – consider production and transfer costs
- linear production – manufacture products through same process route
- single sourcing – only one supplier for each item
- supply chain management – maximize effectiveness and minimize total supply cost.

The new role of purchasing in a company is in negotiation – a skill which inventory management has been practising internally for a long time. Purchasing no longer requires the wheeler-dealer, buying from a favoured source, hard-hitting on price and always with a good idea of where to go for a cheap deal or a quick supply. Purchasers must now be contract managers, working in partnership and shrewd enough to balance the well-being of their company with the aspirations of their supply partners. They will have fewer suppliers and maintain multilevel collaboration within the supplying and user company. In fact, in some companies the purchasing agreements will be carried out by top executives, as they were in the early days of business enterprise. This situation arises because of the improvement in communication between companies and the better organization of the supply chain.

The purchasing environment

The environment in most businesses has changed from a supplier-led to a customer-led operation. This has resulted from the increase in international trading, multinational competition and the improvements in transport systems.

The development of the customer-led market has caused major changes in supply structures. There have been significant developments:

- Supplier relations are much closer.
- Quality management works on audit, not inspection.
- Stock control philosophy is to keep low inventory.
- Supply chain structure is fixed.

- Number and type of suppliers – there are fewer suppliers and these are comparable sized companies.
- Distribution has become reliable and a direct cost, not an overhead.
- Administration has less people with tighter targets.

There has been a gradual change of power from supplier to customer leading to changes from a product-based supplier to a service-based supply (see Table 8.1).

Table 8.1 Changing attitudes to service

Service was supplier led	Service is customer driven
Manufacturer provides products	Customer requires solutions
Supplier provides narrow range of options	Customer sources wider range from each supplier
Non-standard individual items	Items are commodities, system components
Generic service	Different service for each customer
Delivery when available	Fast reliable service
'Provide products' philosophy	Packaged items ready for customers to use or sell

This means that the requirements for suppliers and customers have changed (see Table 8.2).

Table 8.2 Effects of changing attitudes to service

	Issues for service providers	Issues for users
Supplier-led service	Supply lead times Single logistic chain High demand rates Good margins	Variety of suppliers Many orders to control
Customer-led service	Good availability Wider type of demand Stock and supply risks Commodity type supply	Convenience Good control and price Single source Supplier partnerships

*Single sourcing**Work closely with fewer suppliers.*

A major development in purchasing has been in the area of supplier relations, with a pressure to decrease the number of suppliers and to develop partnerships with them. This stems from the need for improved supply, and has led to tighter control and higher quality standards. One solution to the supply problem is seen to be a single supplier for each item, providing the supply risks can be minimized. The approach toward single sourcing can be made in the same way as any other quality improvement process:

- Improve process until a problem is encountered.
- Identify the cause of the problem.
- Find means of solving or avoiding the issue.
- Advance to the next problem.

Managers often consider that their own situation is not appropriate for single-source supply. Traditionally, companies purchased items from a variety of suppliers, choosing a supplier at any time on the grounds of cost or availability or, less often, quality. Such a formal, arm's length relationship with the customer giving an 'order' to the most fitting supplier, was more costly and time-consuming than a single-sourcing arrangement, and did not necessarily give the customer the best service.

A supplier who is not confident of the customer will:

- present a price from which to negotiate
- only meet the customer's specification and no more
- offer improvements only under duress
- provide no add-on services
- consider only the short-term profit
- fear losing the next order.

This negative attitude stops the supply chain from being co-ordinated properly. The exception to this situation is where the customer is so dominant that the supplier is treated like a part of the customer's company and is dictated to by the customer. (This is the situation in many Japanese companies.)

The normal objections to single sourcing are:

- the supplier has too much power
- inflated prices

- risk of poor delivery performance
- lack of flexibility
- supplier may cease trading
- divulging sensitive information
- risk of quality problems.

Focus on least total cost operating, not lowest unit price.

These are real fears for those not accustomed to single sourcing and present the first barriers to be crossed.

There must be compensating benefits to encourage companies to adopt single sourcing. These benefits can be considered as short-term operational advantages, and as long-term improvements. The long-term benefits are:

- supplier is responsible for the item
- closer relationship with suppliers
- joint improvement programmes to reduce cost
- buyers' leverage is enhanced
- focus on quality and efficiency
- fewer vendor contracts
- more stable schedules.

Looking at single sourcing from the supplier's point of view, the advantages are:

- higher share of business
- better opportunity to plan resources
- confidence in continuing demand
- opportunity for product and process development.

The disadvantages are:

- more dependence on one customer
- need to divulge financial information.

In the long term, the development of mutual confidence and a close working relationship will enable two partners in the supply chain to plan facilities to match the ongoing demand, focus on quality and delivery (just in time?), eliminate formalities (time and paperwork), and provide a service rather than simply a supply of items. The aim is to optimize added value throughout the supply chain to the benefits of both partners. This requires a commitment by both parties to collaboration over several years.

The short-term benefits of single sourcing can be classified as:

- reduction in buying and paperwork
- less vendor management
- easier traceability for non-conformancies
- single shipping route
- consolidated deliveries
- faster replenishment cycle.

To make single sourcing work in practice, management have to be convinced that the cost savings from these benefits far outweighs the potential costs of the risks identified.

Take the supplier point of view.

What really is single sourcing? The obvious answer is ‘only one supplier for each item’. This does not mean putting all your eggs in one basket, but it does mean relying on a supplier for a particular size of egg. This always gives the back-up that the size 3 egg supplier could be persuaded to provide size 2 if there is a big supply problem. Single sourcing also means empowering the supplier to provide the quantity, quality and delivery without fail. Strategies to support single sourcing are therefore:

- ensure supplier profitability
- ensure that demand is important to supplier (significant turnover)
- partnership relationship
- management and monitoring of performance.

Vendor appraisal

The traditional purchasing targets of good quality, price and delivery should be interpreted to mean total quality, long-term total cost and consistent or JIT supply.

Quality is the responsibility of the supplier. Quality issues at the supplier interface are concerned with the provision of usable goods, delivered safely and on time, and accurate information including the customer’s identity code, about what is sent, shipment by shipment, box by box.

It has been normal practice in some industries, to inspect goods arriving, on a 100 per cent basis in some instances, and on a sample basis in more advanced companies. It is also customary for manufacturers to inspect goods being dispatched. Items are therefore inspected on dispatch from one

company and inspected on arrival at the next, either to ensure the integrity of the transport or because there is a degree of mistrust. The integrity problem is relatively easy to solve, at least within one country. The mistrust is more difficult.

There are always stories of suppliers sending wrong quantities or wrong items, even though the consignment may have been inspected on dispatch, particularly if the supplier is also the manufacturer. The integrity of the method of supply may also be questioned, more so if it involves international or long-distance transportation. If a customer inspects a delivery and finds it defective in some way, he or she will complain to the supplier, and the questions which then arise are:

- Who paid for a full inspection and when?
- Who separates the defective parts?
- Who pays for the return transport?
- Who owns the problem?
- Who pays for a vendor rating and quality control system?

In most cases the answers to these questions is the customer. If the supplier is not worried about this aspect of customer service (especially for a single source or accredited supplier), then getting the customer to carry out the quality checks seems to be a good cost-saving idea. Only where the supplier's business is jeopardized by the quality situation is the matter resolved. The options for the customer are:

- Send the whole shipment back.
- Establish and circulate a league table for supplier rating.
- Charge the suppliers for the cost of inspecting their work.
- Communicate serious problems to their chief executive.

Make suppliers pay for cost of their quality.

The ability of a business to meet commitments to its customer depends on the quality of support from its suppliers. If that support is missing, the business must look elsewhere for better service.

Appraisal criteria

Audit is necessary as part of vendor assessment, to ensure that the excellence of supply is maintained. The purchaser's appraisal consists of three major areas:

- Technical – is the item fit for function?
- Supply – is the quality, quantity, timing, identification and cost acceptable?
- Support – is the supplier assisting in developments, queries, re-engineering?

The appraisal criteria in each of these areas can be listed and ranked according to importance – some features will be essential and others merely desirable, and suppliers can be classified by their past achievements against the set criteria. Results can then be used to rate the suppliers, as shown in Table 8.3 for the type of inspection required.

Table 8.3 Classification of suppliers

Class	Status	Inspection
Certified vendor list	First choice Close relationship	No inspection
Qualified vendor list	Select group of suppliers Business quality audited Offered expanding business	Minor audit inspection
Approved vendor list	Used on occasions Items not available elsewhere Prototypes, non-inventory items, office supplies	Guarantees or full inspection

Vendor rating systems are operated by most major companies to ensure that their suppliers meet specified criteria. The final result is usually to categorize the suppliers into these classes:

- *Certified vendor list* – these are the small band of good suppliers who need little supervision and minimal quality checks. They should include the major suppliers.
- *Qualified vendor list* – medium supplier performance, but causing some non-conformance costs. Acceptable but needs work to improve.
- *Approved vendor list* – this includes all the other suppliers. Their performance history is either low quality or unproven, so they require a quality assessment project to either upgrade their achievement or dismiss them.

Operating vendor rating

To establish and maintain vendor rating, a programme of vendor liaison has to be created, which should include not only a rating system, but also: frequency of inspections; inspection topics, focus and method; other visits and liaisons; the duration of the visits and methods of recording and relaying the results.

For those suppliers not given a top classification, the purchaser should report back on areas of concern, publicize to the suppliers the selection criteria and objectives, and agree an improvement programme.

When a rating system has been specified, it has to be actioned. It is, therefore, useful to specify in advance the conditions under which suppliers would be abandoned or supported. It may not be possible in all cases to apply these conditions because of restrictions on supply, but it is useful to have prime criteria defined for the best working partnership.

The objectives of vendor selection are primarily to ensure adequate supplier performance, and an agreement on financial rewards and penalties is very useful to make it work in practice. The supplier gets a premium if the rating is excellent and the purchaser gets a discounted price if the rating is poor. The overall effect of this can be to reduce the total quality cost between the two companies. As a spin-off from vendor selection further benefits accrue, namely:

- fewer suppliers
- certified suppliers
- objective performance data
- long-term business relationships
- better understanding and dialogue.

Supply partnerships

Organizing a partnership

The modern relationship of supplier and customer is considered in terms of:

- long-term collaboration
- open relationships
- mutual development
- supply chain management
- quality conformance.

There are also some practical considerations including geographic location and flexibility.

The supplier base consists of a wide variety of different types of suppliers providing individual blends of services. These services should provide what the purchaser requires but, of course, there is always room to improve in one way or another. Some suppliers are retained because they are essential (particularly where one part of an organization is purchasing from another); others are retained because they are excellent. There is a continuing need to develop supplier relationships, to secure improvements and, sometimes, to seek alternative sources where suppliers are unable to meet the requirements.

When selecting possible vendors, both long-term and short-term considerations should be used. What information is there on the supplier, and do they meet, or are they capable of meeting, the supply quality standards? It is important to ensure that the proposed demand level will not overload the suppliers, causing them to fail. If the style, objectives and business ethic of the supplier match the customer's, then a partnership is more likely to be successful.

Partnerships give long-term success.

The main reason for supply partnerships is to improve effectiveness for each of the participants through better inventory management. Improved mutual trust means better forecasting, less safety stock, availability of stock records to each party, shared risk and, therefore, low inventory. Successful partnerships agree on attainable targets which stretch both partners to give an equal share. The effects go outside inventory management into sharing development skills, operational and financial data, and providing support to each other.

The process for developing a new supply partner is usually:

- Select vendor.
- Assess expertise.
- Accredited.
- Agree supply type.
- Test performance.
- Set up agreement.
- Monitor performance.

This monitoring is then fed back to the supplier and maintained on a continuous basis.

Pricing method

Because of the huge variation in market and trading conditions, the purchase price can be fixed in a number of ways:

- market price – variable from day to day
- fixed price – optional supply channel, mandatory supply channel
- market adjusted price – using an agreed formula
- fixed term contract – with discounts.

Each of these has advantages for supplier and customer in different markets, and the choice should be one which both partners consider to be fair.

The ordering process*Planning orders*

There are two separate aspects of procurement to be considered, namely:

- purchasing
- reordering.

A large proportion of inventory control in a mature business will be the reordering of supplies from an agreed vendor, and begins with a communication that a further delivery is required. This situation is more about material control and scheduling than traditional purchasing.

If the purchasers can separate out those items where there are routine orders from those which are not likely to be repeated, then a different focus can be made in the initial negotiations.

The strategy for supply should take into consideration the overall commercial, logistic, inventory and added-value services which are to be provided by the supply chain, and a variety of purchasing options can be considered, as appropriate. Supply can be obtained from:

- the open market
- a manufacturer
- a dedicated provider (specialist or distributor).

Each of these supply sources has different characteristics which should be examined to ensure their use is right for the items to which they are applied (see section on order types below). Current practice is to set up a supply pipeline which goes through as few stages as possible. However, the nearer the original source (manufacturer) of the item, the less flexibility there is likely to be, the longer the lead times and the higher the supply volumes.

Order types

With the increased pressure for effective supply, simple one-off purchasing processes are largely irrelevant. No longer does a buyer act as a dealer, seeking a number of suppliers, negotiating price and delivery, and ordering a batch. This is far too time-consuming and ineffective for the bulk of stock items.

Scheduling suppliers is part of inventory control.

Ordering practices separate the purchasing aspects from the reordering aspects. Purchasing is strategic and carried out at intervals to set the structure, volume and price of supply. Reordering is carried out as a result of stock control and forecasting processes as necessary for managing replenishment. The two aspects require different skills and can be carried out by different individuals. Supply orders can be broadly classified into:

- one-off orders
- bulk purchases
- annual contracts
- split delivery purchases
- rolling schedules.

An analysis of these types of orders and their application is shown in Table 8.4. The more organized companies use JIT or scheduled supplies. Unfortunately, much of the software used in purchasing systems assume one-off orders will be placed, and so schedulers are forced to carry out extra work and make schedules look like split delivery orders. Rolling schedules and JIT supply are the only options which meet the real needs of both the supplier and the customer.

Table 8.4 Analysis and application of supply orders

Order type	Advantages	Risks	Application
One-off orders	No commitment	No continuity of supply	Capital purchase
Bulk purchases	Traditional method for low prices	High inventory; high risk	C class items
Annual contracts	Assured supply from awkward supplier	Need to forecast accurately one year ahead	Seasonal purchases (food)
Split delivery purchases	Lower inventory	Large stock cover or commitment	'Order' based supplier; C or B class items
Rolling schedules	Assured supply, low inventory, low commitment	Non-adherence to schedule, poor forecast adjustment	A class items, delivery more frequent than fortnightly
Just in time	'Zero' inventory pull system; short effective lead time	Delivery failure; unreliable forecast of demand	High volume limited range; responsible suppliers

Order quantities

If the aim is to minimize overall cost, by keeping order quantities low, efficient stock monitoring and goods received systems are required. If this necessitates more frequent ordering, then the purchasing workload is increased. However, an increase in inventory actually results from the delivery, not from the ordering, and the prime task must be to reduce the delivery quantity. Sometimes this is easy where reducing the order quantity is more difficult. There are several approaches, including:

- supply scheduling (see next section)
- multiple item deliveries
- source from suppliers who sell in small quantities.

A reduction in delivery quantities can be achieved in two ways, namely:

- increase the variety of items per delivery
- increase the delivery frequency.

The stock value is decreased in direct proportion to the number of lines delivered.

Consider the simplistic situation of a weekly delivery from one supplier who delivers one item one week and another the alternate week. They deliver a fortnight's worth of each item each time, which causes an average of one week's stock of each. Instead, it is agreed that half the quantity of

each will be delivered every week. The delivery quantity is now enough for one week and the average stock is then half a week's worth.

Go for less orders and more frequent deliveries.

The stock is halved by simply splitting the order quantity between the items supplied.

As paperwork would be increased by the smaller deliveries, it is important to minimize the effect of the extra time spent by limiting this approach to A class items.

Pareto-based order quantities

The key decisions in stock control are when to order (discussed in Chapter 7) and how much to order because the larger the order quantity, the higher the risk of excess stockholding. The order quantity depends on the value and usage rate of the item, and purchase order quantities should be related to the predicted requirements. The use of the Pareto technique described in Chapter 3 provides the best balance between usage rate, costs and ordering effort.

The simple logic is to order A class items frequently to avoid stock value and order C class infrequently to avoid too many orders and deliveries.

A typical order pattern would be:

- A class: order enough for one week.
- B class: order enough for one month.
- C class: order enough for ten weeks.

This, of course, is not dependent on the lead time taken by the supplier. Deliveries can be arranged every day even if the lead time is four weeks, it just means that there are about twenty supply orders outstanding all the time. The lead time does affect the safety stocks and therefore the overall stock level.

Get A class items delivered weekly or daily.

It is a gross assumption that a week's worth of stock is the appropriate order quantity for A class, but it appears to be generally correct across a wide range of businesses and markets. The main exceptions are fast moving goods where there is daily delivery and the equivalent balance in days is used as the order level. Companies where A items are delivered less frequently than fortnightly have large stores and unbalanced stocks, or they have several months' lead times for slow moving A items. They have the wrong idea.

Table 8.5 Pareto order patterns for 1000 stock lines

Class	Moving lines (%)	Order cover (weeks)	Turnover (%)	Average weeks of value	Orders per year ¹
A	10	1	65	0.325	5200
B	20	4	25	0.500	2600
C	70	10	10	0.500	3640
Total	100		100	1.325	11 440

Note: ¹For 1000 moving stock lines.

The effects of ordering pattern on the number of line orders placed per year has already been discussed in Chapter 3, a reduction from 12 000 orders per year to 11 440 per year. If the above order pattern is applied to the example provided in Table 3.8, then the values shown in Table 8.5 result.

The result, compared with monthly purchase for everything, is a reduction of stock from 2 weeks' worth to 1.325 weeks' worth (plus safety stock in each case) and a decrease in line orders of just under 5 per cent. This gives more orders but less stock than shown in Table 3.8, and is better balanced.

The ABC order quantities reflect the natural ordering practice for many inventory controllers, but the Pareto technique provides rules for controlling all the stock in an appropriate manner and makes it work better.

Scheduling supply

Rolling schedules reserve supply and reduce risks.

The best arrangement for ensuring reliable supply and low stocks is the rolling schedule. It is commonly thought that parts which take a long time to arrive should be ordered in big batches. This is not necessary. Large savings in inventory cost can be made by receiving regular small amounts of the high turnover value items which have long lead times. For example an imported part may take two months from the order to receipt into stores. If there are 420 used per month, then an order could be placed every month for about 420. This would lead to a stock of 210 on average (plus safety stock). Alternatively, an order can be placed each week for delivery in two months' time from the date of each order. This would mean a delivery quantity of 100 each time and average batch stock of fifty. However it might be possible to place orders every day, in which case the batch stock would be reduced to an average of ten. (Distribution methods have to be altered

to deal with more frequent transfers without incurring extra costs.) This is a significant reduction in stock and cost without any effect on customer service. The number of orders per year may increase, but the workload can be reduced instead by using supply schedules rather than single orders.

Rolling schedules are not fixed on an annual basis because this may lead to a problem at the end of the year where there is either a shortfall or a large unwanted amount still to be delivered. The best schedule is an agreement between supplier and customer to the supply of an agreed quantity over a number of weeks. Outside this period the customer indicates the expected requirement so that the supplier can organize capacity and materials to cover it. Although the supplier would prefer a long fixed schedule, while the customer would like very short commitment, the scheduling technique has many advantages when operated properly with the full commitment of both supplier and customer to keep to the schedule. They include:

- less risk for both customer and supplier
- reduced stocks for the customer
- reduced stocks for the supplier
- the opportunity to plan supplies
- regular flow of materials.

A further advantage of scheduling arises from a regular flow of orders where, instead of placing an order or a new schedule, the old schedule can be used and extended. This enables a continuous supply of items at the average rate without the need for any purchasing action. There should be a simple agreement between supplier and customer that the schedule quantity rolls forward unless there is specific instruction otherwise, so that supplies are ensured. The only reason for changing schedules should be to adjust rather than create.

Not only does the rolling schedule embody the principle of 'no change – no action' (the classic management by exception principle) but it improves the stability of supply. Even if the supplier's lead time varies, the scheduled quantity is not affected, nor should the regularity of deliveries be affected either. A schedule is essentially a sequence of delivery times, not order placement times, so the equivalent of lead time for schedules is the inflexible front end of the schedule. The only change should be in the length of the fixed part of the schedule. In consequence, as the lead time increases, the customer has a longer time ahead before the programme can be changed. Schedules solve many of the problems of ordering and minimize the risks.

To use schedules successfully, both supplier and customer have to be committed to the quantities agreed. There should also be no sudden changes

in schedule outside the fixed part of the schedule. These are caused by late or overcorrection to demand or stock levels and, if not avoided, will give a stop-go ordering pattern. When usage changes the schedules should be modified gradually. Minor action early prevents a major problem later on.

Accurate systems are required to control scheduled orders. Running orders and continual deliveries are included on the schedule, and the cumulative delivery quantities have to be correct with a check to show which scheduled delivery is actually being received. A computerized technique is an advantage where there are many items to schedule.

The economic order quantity approximation

Where there are no other guidelines to determine the order quantity a theoretical model can be used, the 'economic order quantity' (EOQ). This model considers batch size, Q , which results from balancing the cost of holding stock and the cost of ordering it. All costs are assumed either to vary with the batch quantity or to be invariant with the order cost. For an item with a cost of P and annual sales of A , the cost of holding stock is given by $\frac{1}{2} \times r \times P \times Q$, where r is the financing factor. The half arises through averaging the usage rate to get stock levels. r is usually 20 per cent to 30 per cent, because it takes into account the cost of borrowing the money for the stock and the running cost of the warehouse. The order cost, S , contains the cost of running the order department, goods receiving, and supply progressing and can be typically US\$100 per order.

Consider an item with a usage of 420 per month (5040 per year) and a cost of \$35. Using typical values of \$100 for the cost(s) of raising an order and 25 per cent for the return expected on the stock investment, the ordering options are shown in Table 8.6. The costs for various order patterns show a minimum for fifteen orders per year, so this is said to be the EOQ.

EOQ does not work in reality.

At the bottom of Table 8.6 is the actual minimum value, fifteen orders per year. This has been calculated using the formula:

$$Q = \sqrt{\frac{2 \times A \times S}{r \times P}}$$

Substituting the values in this equation gives the economic order quantity, Q , of 339, which suggests that the item should be ordered fifteen times per year.

Table 8.6 Balancing order and stockholding costs

Orders per year (N)	Weeks between orders	Order quantity (OQ) (A/N)	Ordering cost (OC) ($N \times S$)	Stock cost (SC) ($0.5 \times OQ \times P$)	Total cost (OC + SC)
1	52.0	5040	100	22 050	22 150
2	26.0	2520	200	11 050	11 225
4	13.0	1260	400	5 513	5 913
15	4.0	388	1 300	1 696	2 096
26	2.0	194	2 600	848	3 448
35	1.5	144	3 500	630	4 130
52	1.0	97	5 200	424	5 624
104	0.5	48	10 400	212	10 612
Balance					
15	3.5	339	1485	1 485	2 970

Although the EOQ theory gives simple answers, it generally gives too high a stock level and should only be used where there is no alternative. There are many reasons why EOQ gives the wrong answer in practice. These include:

- Order and stockholding costs are assumed to be fixed, but these costs should vary according to the stock and ordering situation.
- The valuation will give different results when estimating by full or marginal costs.
- The demand is assumed to be regular, and ignores batching and timing of issues which have a major impact on order costs in reality.
- Split deliveries and schedules do not fit well into the cost equations
- The EOQ ignores the balancing of stocks which is the most important factor in manufacturing and many other cases.
- There is no allowance for co-ordination of orders for similar items where the order cost can be shared.

There are other more minor problems which arise from the impractical assumptions made by the model. The application of this EOQ technique is therefore in the quick assessment of the stock and ordering situation. It is not appropriate for use in manufacturing, nor should it be used as part of a stock control system.

As can be seen in the example, the total cost is insensitive to the number of orders per year near the minimum cost. An increase of order frequency from fifteen times per year to twenty times per year only increases the total cost by 4 per cent but decreases the inventory by 26 per cent. Even ordering

every two weeks rather than every 3.5 weeks will increase the total cost by 14 per cent but decrease the inventory by 43 per cent as evaluated from the table.

Good techniques can reduce the order quantities below those given by the economic order calculation without increasing the costs. In many situations, batch sizes in the region of 60 per cent of EOQ minimize the actual costs. The best answers are given by considering the practical situation and keeping the delivery quantities as small as can be sensibly achieved.

Other order size considerations

There are situations where the supplier imposes a minimum batch quantity, or a penalty for ordering small quantities. This suggests that there is a mismatch between the supplier's perception of a 'reasonable amount' and that of the customer's requirements, and is a sign that the supplier is not interested in the amount of business which is available. Good liaison between purchaser and supplier could set up a rolling schedule to assure the supplier of continued custom, albeit for small quantities.

Minimum order quantities, as set by suppliers, are generally arbitrary round numbers. It is possible to erode them by a few per cent in many instances, and gradually move to the required quantity.

Minimum order constraints on A class items should be studied energetically to avoid expensive stockholding. Where the minima are imposed on C class, then the solution may be to accept the constraint as long as goods are likely to be used in under a year. The minimum order quantity problem can be resolved in a number of ways offering different degrees of flexibility as shown in Table 8.7.

Table 8.7 Minimum order constraints

Constraint	Initial policy	Long-term changes
Minimum order value	Place order for phased delivery	Negotiate on volume of annual business
Minimum shipment size	Co-ordinate orders for widest mix of purchases	Change method of shipment
Minimum order quantity per line	Order for phased delivery	Schedule, or arrange payment when stock is used
Minimum delivery quantity per line	Share item with others (locations or competitors)	Renegotiate or look for alternative supply

Where the supplier is a sole source, then more effort has to be made to resolve the situation:

- Customers may be persuaded to accept an alternative.
- A second source may be found.
- The item can be sourced through a third party (avoiding significant cost increase).
- The item could be destocked and the minimum quantity supplied direct to the customer.

Negotiation and an understanding of the situation can sometimes lead to a sensible accommodation on both sides.

Apart from minimum order constraints there are other situations requiring specific policies. In many situations, the forecast will be considered to be acceptably accurate for a number of weeks ahead, up to a limit, commonly called the 'planning horizon' illustrated here for each. (*Note: planning horizon is the length of forward planning. In Table 12.4 it is ten weeks.*)

Inventory should flow through the business.

Different situations give rise to different order policies:

- *Single warehouse* – forward cover to horizon fixed by Pareto class.
- *Standard product* – cover to horizon against forecast.
- *Distributed stockholding* – co-ordinate supply to all stores.
- *Manufacture to order* – buy and make required quantity plus scrap allowance.
- *Parts for assemblies* – make balanced sets.

In manufacturing, batch quantities should be the same for purchase and production where possible. The actual batch sizes are often determined by the practical limits of capacity on a process, the overall process throughput time and the physical need to make a batch a more easily handled quantity. By co-ordinating the manufacture of like products the batch size of each can be reduced because the set-up costs are reduced. Thus a slow-moving item should be produced in small batches along with similar fast-moving items.

Purchasing processes

The stages of purchasing for a reorder are usually as follows:

- 1 Usage causes requirement.
- 2 System identifies need to purchase.

- 3 Need is assessed and actioned if necessary.
- 4 Purchase is initiated.
- 5 Order is produced.
- 6 Order is authorized (if high value).
- 7 Order is sent to supplier.
- 8 Supplier initiates supply.

If the workload in purchasing is too high to allow proper vendor development and liaison, companies must find ways to reduce the ordering workload by attending to the eight stages and rationalizing them, as follows:

- 1 *Usage causes requirement.* This is the initial step in the reorder process and the whole supply process is dependent on the information being accurate. The first issue is to ensure that the stock and demand information is correct.
- 2 *System identifies need to purchase.* Much time can be wasted searching through printouts, screens or even stores, to see what is really required. Systems should identify by exception reporting those items considered to be ready for reordering. The reporting of this should err on the safe side by selecting all items which need ordering plus extra items whose demand patterns may lead to a high usage, a decision on which can then be taken by the inventory controller.

It is always useful to have a communication from stores which reports what they consider to be at risk. This avoids orders getting overlooked and adds weight to the supplier-expediting process.

- 3 *Need is assessed and actioned if necessary.* The information provided by the system should initiate orders for a minimum of 85 per cent of the items where orders are recommended. If this is not the case, then the stock control system is inadequate and requires improving. The development of good stock control systems through forecasting or MRP is not a complex computer problem. Solutions are available for most situations, including sporadic demand and seasonal sales.
- 4 *Purchase is initiated.* The actual purchasing process can be a simple confirmation at the press of a computer key, or a time-consuming discussion with one or more suppliers. Some companies even require a quotation for routine supplies! For fast-moving items, the ordering process should be done in a flash once the order quantity has been decided. (For most items the order quantities should be determined using the formulae discussed previously.)

- 5 *Order is produced.* Order production is part of the normal computerized purchasing system and is printed automatically. Manual orders are to be avoided since they complicate the system.

Organize routines minimizing purchasing workload.

There is a good argument for a petty cash purchase system where over-the-counter purchases of low value are carried out by an authorized person outside the normal purchasing system. This will avoid a large bulk of formal purchases and immediate access to the items required (often maintenance parts). Control over these purchases should be through the monitoring of monthly purchase values.

- 6 *Order is authorized (if high value).* For capital goods and non-routine purchases, there has to be executive involvement in the commitment of company money. For routine supplies, this process only serves to delay the purchase, and is a great disadvantage. The signing of reams of purchase orders by a senior person wastes their time and usually the orders are not investigated during this process. There should therefore be generous spend limits for buyers. Executives need to control the purchases by:
- (a) monitoring the value of order placed and outstanding order commitment weekly or monthly
 - (b) signing the major value orders (say twenty per week)
 - (c) auditing the general order level by random checks each week
 - (d) setting targets for reduction of value of purchases relative to demand.
- 7 *Order is sent to supplier.* By post an order can take two days to reach the supplier, longer for overseas suppliers. The use of fax or more modern methods enable the supplier to receive the information immediately. The use of e-mail, wide area networks (WANs) or the data superhighway is important for ensuring fast ordering. This saves stockholding as well as improving communication. The well-organized purchaser will not have to use the telephone for normal routine orders. The telephone is time-consuming, and should be reserved for supplier liaison and negotiation.
- 8 *Supplier initiates supply.* The information sent to the supplier should then be processed into their supply system. Any delivery problems should be relayed to the customer without them having to prompt. In situations where the customer is apt to alter the demand, there is a tendency not to report difficulties.

Key points

- Buying consists of creating strategic contracts and supply scheduling.
- Close co-operation with suppliers pays dividends.
- Single sourcing is best, but risks have to be identified and ameliorated.
- Quantitative assessment of suppliers will identify potential problems.
- Organize ordering to minimize delivery quantities and effort.
- Use Pareto ABC as a basis.
- Avoid the use of EOQ.
- Review purchasing control and administration to reduce work.

Forecasting demand

- The role of forecasting
- Types of forecasts
- Avoiding forecast data error
- Managing data for forecasts
- Benefits of group forecasting

Options for assessing demand

Buying safety stock, or purchasing a batch quantity, means making a prediction that the items will eventually be sold. This is a crude type of forecast. In order to improve the stockholding (increase availability and reduce stock) the quality of the forecast has to be improved. If forecasts were excellent and the supplier could predict when the customers would want items, they could be obtained and delivered just when they were needed: Good forecasting means low stock. Poor forecasting means high stock.

Forecasting should be based on data which is accurate and appropriate for the purpose. This is often a problem, and a forecast frequently has to rely more on balanced estimates of future sales rather than history. The data used to calculate the forecast include established demand patterns but ignore irregular demands for which stock is not normally held, such as scheduled customer orders (for which specific purchases are required to cover the demands). For example:

- one-off promotions
- sales campaigns
- one-off very large orders (generally orders where the demand is more than four MADs).

Good forecasting results in low stocks.

The selection of the forecasting technique is a difficult decision and either the inventory controller selects the most appropriate one or employs a focus

forecasting technique. In general, it pays to use the most sophisticated technique available because it is a better model for the demand pattern and gives better results. Forecasting techniques can be divided in many ways. One simple classification is between forecasts which use prediction and those which use history, another is between qualitative and quantitative figures, and another between intrinsic forecasting and extrinsic forecasting.

Considering these concepts in turn, historical forecasts rely on looking back to predict the future. Historical forecasting forms the basis of most inventory forecasting because it has sufficient detail, and demand for products usually changes continuously. However, it is obvious that looking back to go forward could lead to disaster and this is where predictive forecasting plays a role. People will predict an event based on understanding of future changes, rather than extending what has already happened. The problem with predictive forecasting is that it is often qualitative. A sales campaign will increase demand, but by how much? The qualitative prediction is marginally useful, but for stock control a quantitative assessment is necessary.

The difference between intrinsic and extrinsic is that the intrinsic forecasting is based on information which is available within the system. Extrinsic forecasting uses information which comes from external sources such as trade statistics or market information. It is convenient to split extrinsic information into leading indicators, concurrent indicators and lagging indicators, according to whether the information collected shows future, current or past activity. For instance, government statistics on housing starts is likely to be a useful leading indicator to a national roof tile manufacturer. The same information would be a concurrent indicator, showing what is happening to the general demand this month, to a company hiring digging equipment for house foundations, and for architects it would explain the level of activity which they had already experienced. Obviously, leading indicators are the most useful, even though the information provided is often qualitative.

The manipulation of forecasts is part of the inventory manager's art, since computerized systems are unable to provide the right interpretations when conditions change.

Types of inventory forecasts

Some types of forecasting are used only for specialized purposes. Of those used regularly by inventory managers, the following five approaches are the most useful and form the basis of stock control:

- 1 *Market research.* There is nothing like knowing for sure what customers are going to buy. Firm orders are the ideal! Sometimes it is possible to obtain firm order commitment, but more often it is possible to gain only estimates from customers. Good liaison with the biggest customers, or the biggest users of A class items, can have a major impact on forecasting accuracy. If customers trust a supplier, they will provide accurate information, since they should be in the best position to know their market and their likely order requirements.

An alternative would be to carry out surveys to discover trends in the market. These can be particularly useful in one-off investigations for new products or for determining the effect of different sales or promotional policies. It is a skilled and complex job to devise market surveys which give factual results for calculating the demand level. A market survey is only cost-effective as an ongoing exercise for a narrow item range where the volume is high.

Another method of assessing customer requirements is to ask the sales people. The risk with this is that sales staff sometimes confuse the difference between maximum potential sales and expected average sales. Consequently, their estimates often have to be interpreted before use!

- 2 *Market demand models.* These are generally based on knowledge of the customer's markets. If the major factors which cause the demand can be identified, then a model can be created. The influence of each factor can be assessed and put into an equation, or the factors can be input into a simulation which can develop its own parameters.

Choose the right forecasting technique for each product.

Models are very effective where there are leading indicators, and where there are simple deterministic relationships causing demand. Some models can be financial (expenditure on promotion versus increased turnover), others technical (item needs replacing every \times years) or based on commercial factors.

- 3 *Historical techniques.* Forecasts should be based on historical data, and then modified by the influences. The only exceptions to this rule are for new products, and those which have a causal link to another forecast (via a model or MRP).

There is a wide variety of historical techniques to choose from. Some are simple and give adequate results, others are much more sophisticated and give excellent results in the right circumstances. In general, the more sophisticated the technique, the more data is required to realize its potential, so established products can be forecast best using these. Also,

a weighted average technique will be better than a simple average. There are many approaches to historical forecasting, favoured by a variety of authorities. The basic models for these will be discussed in the following chapters.

- 4 *Minimum stock levels.* Many systems, both manual and computer have a 'minimum stock level' field displayed. This field is usually used for the 'reorder level' instead of for the minimum stock, and is either a user input field or fixed by some useless rule (such as cover for 'x' weeks).

With a large range of stock items, it is a very big task to maintain stock levels up to date. In practice, where the minimum is too low, the stock controller will increase the minimum level, but where the stock is too high, there is no availability problem and so fewer adjustments are made. As a result, the stockholding can easily increase gradually.

Minimum stock levels should be avoided and proper safety stock levels should be recalculated on a regular bases by a computer system. An area of development of stock control has been the improvement in forecasting techniques used by the professional stock controller. These should continue to be used to the maximum.

- 5 *Demand patterns.* The demand for stock items often follows a pattern which results from a variety of customer requirements for each item. There are two general approaches predicting demand patterns:
- (a) treat the usage as a pattern and analyse the demand profile (an intrinsic approach)
 - (b) investigate the causes of the demand and forecast the requirements separately (an extrinsic approach).

The first alternative has the advantage that it is based on accessible information and is quick and easy to apply. This type of modelling does not normally predict rapid changes to the usage rate because it is essentially based on the historic demand pattern.

If a change in demand is understood, predictions can be made of likely usage and total likely requirements. This method should provide a more accurate answer, but requires thorough communication, knowledge and management.

Predicting for a group is more accurate than for each item.

In practice option 1 based on historical data is used for the majority of items. High-value items and critical items are usually treated by the second method, often through an informal system, although such systems should be avoided if possible in favour of extrinsic forecasting which will manage stock more efficiently.

Product life cycle

For a company's marketing and manufacturing departments, a product has a life cycle of three phases – design, development and manufacture. The end of manufacture is determined by marketing strategy and the introduction of a new more competitive product. This product life cycle is shown in Figure 9.1. It should be noted that product life cycles are getting much shorter with new technology. The product life cycle for the mechanical typewriter was more than twenty-five years; for the modern word processor it is frequently less than one year.




Figure 9.1 Product life cycle

For a service organization, the need for stockholding may extend beyond the product manufacturing cycle since customers will retain products for many years after the end of manufacture and still expect spares to be available to support a service for their equipment. For some companies, products manufactured more than twenty years ago are still in use and being repaired as required. In some cases, service organizations have announced stop dates for service in conjunction with sales initiatives to sell customers new products. However, this has frequently failed since third party maintenance organizations have seen an opportunity and taken over a segment of the business which is very lucrative.

The product life cycle for the service organization is different from the marketing and manufacturing life cycle and can be divided into three phases:

- introductory phase
- supply phase
- obsolescence phase.

Introductory phase

During this period, which should coincide with the design and development phases, the inventory planner will be working with the product project team, when the customer service strategy for the product should be decided. To determine the stock requirement, the inventory planner will want to know the forecast demand and the possible geographical distribution of customers, which will enable the planner to determine the requirement and place time-phased orders for the items necessary to support the market strategy for the product ahead of the launch date.

Product launch is a critical period in the life of the item and it is essential to have sufficient to support the first customers. However, if too much is bought initially, the profit can be seriously impacted because much of the initial purchase may never be used if too much 'just in case' inventory is bought. The planner must monitor carefully early life usage as the sales build up and adjust his ordering accordingly.

Supply phase

Established products have a steady or slowly varying demand pattern, and during this part of the product life cycle major demand of stock items occurs. Inventory control is concerned at this time with maintaining availability. The ongoing requirement can be determined from historical usage data. Standard stock control systems with provisioning facilities can be used for reordering replenishment stock. At some point during this period, manufacturing may cease which will create a 'last time buy' opportunity for distributors. Discretion is needed to ensure excessive 'just in case' inventory is not bought at this time.

Obsolescence phase

It is important that the inventory planner is involved as the requirement for the product declines. The task now is to prevent a build-up of surplus stock, and the stock control system should order fewer replenishment parts as the usage falls. However, a large proportion of inventory will still be in place to meet the requirements of customer service. When the product demand ceases then all the inventory should be withdrawn promptly. Companies find this a valuable exercise since it provides an opportunity to redeploy

inventory in other countries where the product life cycle is in its introductory or supply phase, and not its obsolescence phase.

During this obsolescence phase, consideration should be given to the relative importance of customer support and the high risk of having to write off unsold stock.

Causes of forecasting inaccuracy

People sometimes say ‘you can’t forecast this’ – a statement which is untrue. Everything is capable of being forecast, but some things can be forecast well and others are very difficult to forecast. To improve the quality of forecasts, and therefore of stock control, it is useful to review what risks there are and what causes forecast inaccuracy. Typical pitfalls for the forecaster are as follows:

- inaccurate data
- sales information rather than demand statistics
- bias
- speed of response to change
- poor assessment of supply capability
- inclusion of extra demand in the forecast
- shortage of data.

Inaccurate data

The whole of inventory management depends upon knowing how much stock is available so that the appropriate controls are maintained and actions taken. Inaccurate data leads to poor availability, unreliable customer service and excess stocks. It is therefore imperative for stores to maintain records accurately, and not just to maintain records by a weekly or monthly physical stock check. Records should be maintained accurately through a variety of techniques including:

- automatic recording and checking
- batch control
- stores empowerment
- perpetual inventory checking
- bar coding.

If records are not accurate, then it is not worth improving inventory management processes until inaccuracies are small enough to be a minor influence on customer service.

Sales information rather than demand statistics

When considering inventory statistics, it is normal to discuss 'demand'. However, most companies' records focus on 'sales'. The difference is only important where the actual amount issued does not correspond with the amount ordered. This could happen in a number of ways, including sending boxed quantities, or when there is a stockout. In the case of a stockout, the lack of sales could lead to the conclusion that no stock is required if no notice were taken of the actual demand.

Professional stock records should therefore monitor 'demand' (order requirements at their delivery dates) rather than just sales.

Example

Taking an extreme example: if an item has a demand in three months of thirty, forty-five and forty, but there is no stock of the item in the first two months and then a delivery during the next month, the sales records would give monthly sales of 0, 0 and 115. The actual fluctuations in demand are only about ten whereas the sales statistics would give a much higher value to month 3.

Bias

Sales tend to use a forecast either as a motivator to the sales team by exaggerated optimism, or to ensure that stock is available. As a consequence, projections from sales departments are notoriously high. Stock controllers have to resolve this basic issue either by interpreting sales data during forecasting, or by making sales (or marketing) departments more responsible for the accuracy of their forecasts. This can be done by apportioning financial responsibility for excesses.

Get the data right, then you can forecast.

Speed of response to change

It is difficult to get the right compromise between overreaction to events and lack of response. One or two periods of high demand could signify an increase in market demand or simply a statistical variation. Either assumption will cause risk. Similarly when considering what time intervals

to use for collecting data, if the time period is long (e.g. a month), there could have been major changes in customer demand profile within that time or, for short periods, demand can be erratic and make it difficult to interpret trends.

Poor assessment of supply capability

Fulfilment of increasing demand may be restricted by availability from suppliers. This can happen even with good suppliers who are used increasingly until they become overloaded and cease to supply successfully. This will result in forecasts which are not achieved. Supplier restrictions have to be understood by the forecaster. If a supplier cannot cope with peaks in demand, then there is no point in including them in forecasts for peak rates.

Inclusion of extra demands in the forecast

Special offers, campaigns and fixed scheduled orders should not be included in the routine forecasts. They add variation to the demand rate which is not true random demand. Instead, there should be planned supply to meet the requirement which is assessed separately. For example, a demand pattern may be:

20 510 30 5 25 520 15 35 25 515

and the reason for this may be one major customer who requires 500 every four weeks plus a variety of customers who are taking about fifteen per week. A sensible policy is to stock for the many small demands and to order in specially for the 500 to arrive just before it is required. The 500 would therefore not be included in the data for forecasting but as a separate type of demand in the item stock record.

The ability to forecast closely to the actual demand rate will depend on the selection of the method of forecasting. This can be done automatically by the computer from the historical models, but staff may also possess additional information which could be significant in forecasting future demand.

One-off abnormal demands can be excluded from forecasts by using a filter set at, say four MADs so that demands outside this range are treated as non-stock sales.

Shortage of data

For a 'good' historical forecast, sufficient data is required. It would be sensible to go back into history until the demand became different but it is often difficult to assess this from the data. For seasonal sales, it is essential to have data for more than one year.

The smaller the amount of information available, the less accurate will be the forecast. For example, the statement 'last two months' sales were 500' gives very little information on which to base stock levels. If the information available is 'the sales were 300 for month 1 and 200 for month 2' then this starts to give an indication of the demand pattern. If weekly demand data were available, then it would be possible to see whether the demand was:

- only two customer orders or sporadic demand
- regular orders each week or each day
- decreasing demand.

This suggests that it is better to use more smaller time periods rather than large ones, unless the demands in the time periods become non-standard (e.g. if demands on Fridays are different from the rest of the week, daily time periods would give cyclic demand with a peak each Friday). How much data is needed to make a decision? The answer is indicated in Figure 9.2. The size of the forecast error depends on how erratic the demand is –




Figure 9.2 Effect of sample size on accuracy of mean

as indicated by the vertical scale. Figure 9.2 suggests that there is a reasonable degree of accuracy afforded by twelve periods of demand history.

Work with the data available – it's the best you have!

Methods of improving forecasting

Feedback and monitoring

The simplest improvement to forecasting comes from ensuring that the forecasts are modified to reflect actual demand. This simply requires the re-evaluation of forecasts at the end of each period. New forecasts have then to be applied to inventory calculations to optimize the stockholding. Problem items can easily be identified in each period using a stock cover exception report (see Chapter 3).

Feedback from the forecast error can show where the forecast technique is lacking and can identify when the forecasting method is not coping and needs changing.

Data aggregation

Statistically there is a major benefit from using as much data as possible by going back further in time or spreading wider geographically, locally and globally.

The obvious risks of using old information are:

- undetected population changes have altered the usage rates
- demand patterns have changed through modifications in user practice
- items have been substituted by alternatives or modified
- contracts or customer base have changed significantly.

If a wider geographic population is used, then the risks are:

- usage patterns differ regionally
- different demand patterns (or modification mix) exist locally
- data from different sources may not be compatible.

It is therefore a matter of assessment and balance as to the range of data which can be used for individual items. In general, the wider the information base, the more reliable the forecast. The accuracy improves as the square root of the amount of data included.

Comparative items

For a high usage item, random fluctuations are small compared with the forecast demand. Where demand is small, fluctuations can be large compared to the forecast. Alternatives to aggregation are, first, to find a similar item which is expected to have the same demand profile and, second, to vary the forecast for a minor item by scaling the forecast for a major item.

Example

Consider two items in the same product range, one with a demand of 100 per month and the other with a demand of 2 per month. The forecast for the major item shows that the forecast has drifted from 100 per month to 110 per month. Assuming that the profiles are similar, the demand for the minor item should increase by 10 per cent as well – to 2.2. It is much more difficult to detect confidently a change from 2 to 2.2 than from 100 to 110. It does not matter whether the two items change in exactly the same way. This technique works if the minor item is more likely to follow the major item than to remain static.

If there are a large number of items in a product range with low usages, comparison with the major items can be a very good way of allowing for changes in the underlying level of demand. The forecast for each item can be based on individual historical demand, and the resultant forecast multiplied by the relative change factor used on the major item.

Key points

- Better forecasting gives low stock levels.
- There are many techniques and options so several should be used.
- Forecast can be predictive, based on product or sales information, or on historical data.
- Record accuracy is paramount for realistic forecasting.
- An ongoing effort should be made to keep the records accurate.
- Extending the amount of data (by aggregation or copying profiles) increases the reliability of the forecasts.

Historical forecasting techniques

- Historical averages
- Better forecasts with weighted averages
- Choosing the weighting factors
- Choosing the best forecast
- Warning when forecasts are bad

Basic forecasting techniques

Forecasting

Minimum stocks are possible when a good knowledge of demand exists. If demand were known accurately enough and far enough in advance, then no stock would be necessary. On the other hand, if there is very little market information, high stocks are needed to ensure reasonable customer service. Forecasting is therefore one of the important aspects of stock control. It is a complex art which embraces many factors.

There are two basic approaches to forecasting:

- assessing future market requirements
- using demand history.

Assessment of demand

Assessment of demand requires the knowledge of customers, products and background conditions. That comes from discussion and evaluation with users further down the supply chain and through techniques such as market research, questionnaires and customer surveys. The strength of these techniques lies in providing general information on the levels of business and on the likely uptake of a new product, or the phase-out of an old one. This information is invaluable where large changes in demand level are likely to occur.

These techniques need to be better managed to make them more acceptable as forecasting tools for each line item, and inventory management need to work with sales departments to ensure that they do not draw optimistic conclusions – higher demand – which then has to be second guessed for planning stocks. However, assessment of the market demand is an essential tool for the professional stock controller, and the technique should be used for A class items and those at the start and end of the product life cycles.

One method of determining customer demand is to assess potential sales. Many companies produce estimates, enquiry responses or have contracts and formal agreements. Using a proven conversion factor on this information to calculate an expected average demand by item number can be the best way of estimating demand. Of course, a customer order could be considered as a market prediction, since there is always the possibility of them changing the date or quantity.

Demand history

The second approach to forecasting, using demand history, provides a convenient basis for predicting demand. For most items it is:

- readily available
- detailed by item number
- generally reliable
- easy to use.

However, using demand history is rather like rowing a boat. To go forward means sitting looking backwards and working hard. This is fine if there is no obstruction in the way, but it is advisable to look round from time to time to see what is in front. In a similar manner, inventory controllers use history to determine the future. They have the problem that they are steering many hundreds of items at the same time, and it is not practical to look round for each one to see if it is heading for collision. What is required is a general warning mechanism to detect impending disaster, and a forward look for the more sensitive and important items.

Historical forecasting works for the vast majority of items and is therefore a basic tool of inventory control. The more sophisticated the forecast, the better the results, but there is always the need to keep abreast of real changes in the market so that forecasts can be overridden if necessary.

Base forecasts on history.

In practice, a variety of simplified models, based on major features of the demand pattern, can be used to meet stock control requirements for standard items. More sophisticated forecasting methods are applied to the non-standard items.

Historical forecasting methods are based on the mathematical manipulation of historical data. This approach is fine for most items, and can lead to excellent forecasts where the demand pattern is consistent. It is preferable to off-the-cuff estimates because it is:

- consistent
- rapidly calculated
- frequently updated
- based on facts not reaction to pressures.

However it should be stressed again that it is desirable to modify forecasts in the light of additional information.

Moving average

The most straightforward method of forecasting the next period's sales is to take the average of sales for each preceding period, add them up and divide by the number of months. The moving average technique is a formal way of calculating this for each period in turn, taking into account a fixed number of preceding periods (usually weeks or months). Financial accounts have a fixed annual cycle in which history builds up from the start of the year. This data is unsuitable for stock control, since a fixed number of periods of history is required. The moving average can be calculated from:

$$\text{Average demand} = \frac{\text{Demand for N periods}}{N}$$

where N is the number of periods over which the average is taken. The selection of N is of prime importance in getting the best from the moving average technique. It is a compromise between the inaccuracy caused either by too little data (see Figure 9.2) or by using old irrelevant data. Often the average is taken over twelve months (or fifty-two weeks) to avoid complications due to holidays or seasonality. However, it has to be remembered that by averaging fifty-two weeks of history, the data used is on average twenty-six weeks old: a lot may have changed since then! Therefore it could be safer to use a shorter moving average.

Example

Take, for example, the demand history for three items over a twelve-week period, as shown in Figure 10.1. The average demand for all of them is 20.5 but the reliability of the figure 20.5 as a forecast for the next week could be different in each case. In Figure 10.1A the demand is in between 11 and 28, and 20.5 is quite a reliable forecast. For Figure 10.1B demand is more sporadic and 20.5 is not such a good estimate – perhaps more data would be useful. In Figure 10.1C there is a gradual increase in demand and 20.5 is rather low, so taking a shorter average would give a better answer in this case (i.e. a smaller number of periods).

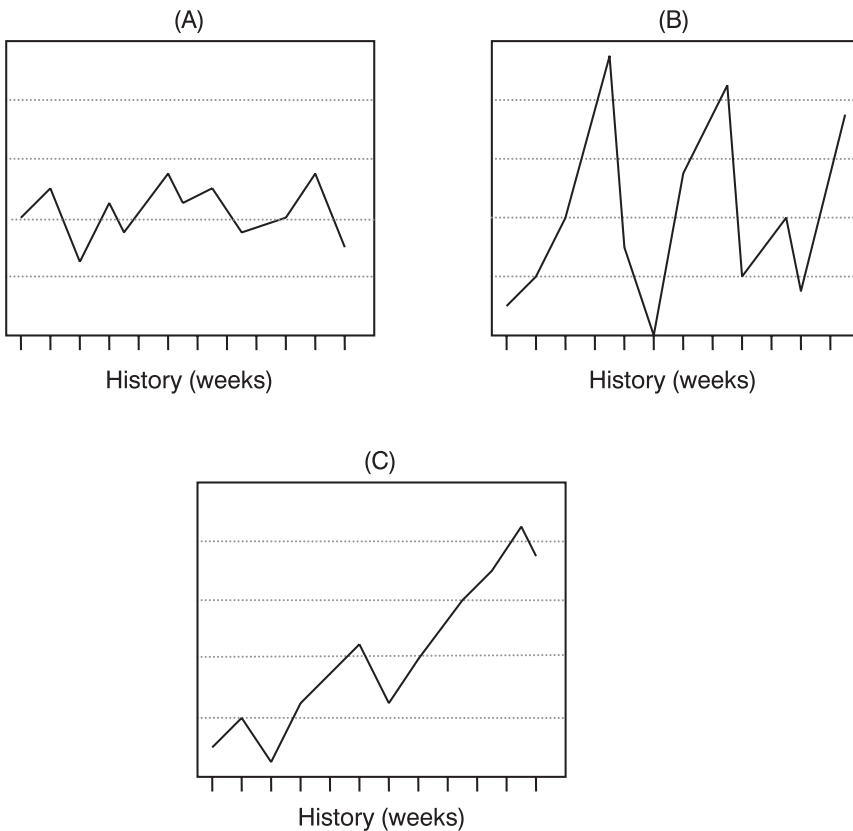


Figure 10.1 Historical demand patterns

In general a 'suitable' number of periods is chosen for a whole range of stock items. This is not the best option but a practical solution. Calculations including many time periods are best where underlying demand is static, where there are large variations between demand in each period (Figure 10.1B), where stock level is not too critical and the forecasting is allowed to run itself without much supervision. Reducing the number of periods in the average enables the forecast to cope faster with a changing demand level (Figure 10.1C). This is useful where the demand is smooth or where the sales in a period are dependent upon those in the previous period (autocorrelation).

Each inventory line needs its own averaging technique.

Taking an average is really a method of segregating trends and random fluctuations. If the average includes data for a long period then any spurious effects are damped out, but this will also make the result insensitive to trends. Conversely a short moving average with N small, gives good response to changes in market but is affected by fluctuation.

Example

The histories plotted in Figure 10.1 are provided with six-week moving averages whereas three-week moving averages are used in Figure 10.2. Notice that the smoothest line is the six-week average. The three-week average is more variable, but it also responds better to the trend in Figure 10.2C. In practice a three-week moving average contains rather a small amount of data, so if it provides the best forecast, then there is a major trend and a more advanced forecasting method would normally be used or, if not, then the three weeks could be split into days or half weeks to get a reasonable amount of data.

Alternative calculation methods for a moving average

Calculating the moving average by summation month by month can be long winded, especially when dealing with twelve-month averages. There is a shorter way of generating the new moving averages.

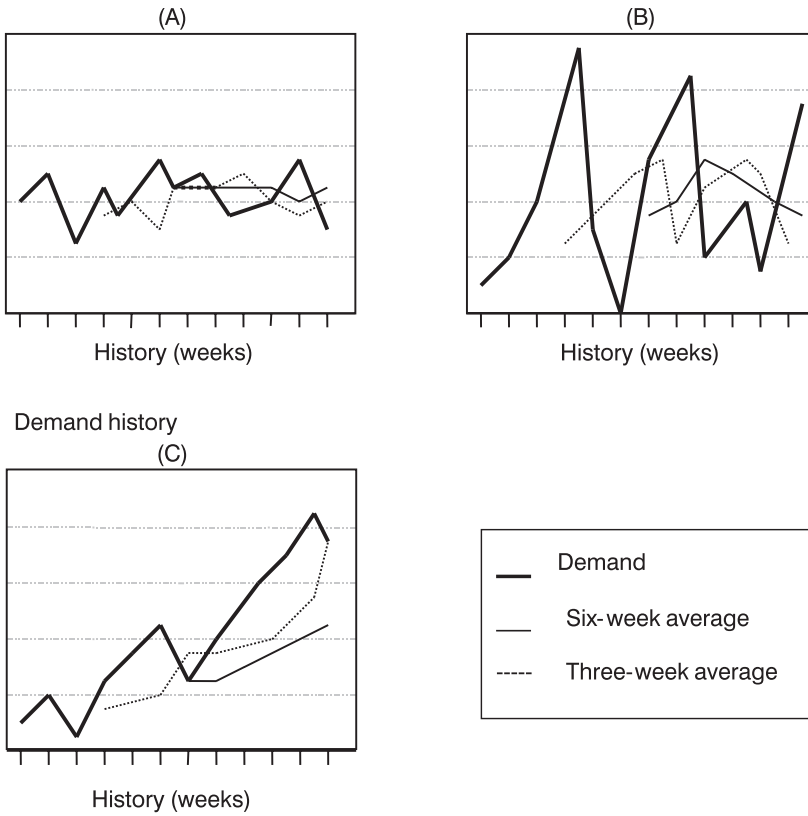


Figure 10.2 Moving averages

Example

Consider that the demand for an item over the previous seven months was:

Month	1	2	3	4	5	6	7
Demand	17	11	15	12	14	15	?

The average demand, $A_o = 14$ for the first six months. In the following month the sales were 11. The new six-month average demand A_N is therefore 13.

To calculate the old moving average, A_o the total six-month demand of 84 was divided by 6. The demand for 17 in month 1 is not now included in the calculation, since it is older than six months. However a simple way of calculating that A_N is 78 was $84 - 17 + 11 = 78$.

Therefore the new average can be calculated by working out:

$$\text{New average} = \frac{\text{Total demand} + \text{New demand} - \text{Oldest demand}}{N}$$

Even easier, and more practical, is to find the difference between the old demand data (to be discarded) and the new period data, and divide the difference by the number of periods in the moving average. Then this can be added to the old average to give the new average, or:

$$\text{New average} = \text{Old average} + \frac{(\text{New demand} - \text{Oldest demand})}{N}$$

Moving average is the most common forecasting method because it is simple to apply and easy to understand. However, it is not the most favoured forecasting method because it has to be calculated from data extending back over the month by month period, and equal weights are given to all periods. It is a matter of judgement whether the most recent information is more important in a particular instance. Often it is, and better techniques are required.

Moving averages for C class items.

Moving average is most successful where the demand fluctuates widely, because more sophisticated methods cannot identify demand trends and profiles in the presence of sporadic demand. Moving average is often used for C class items because the demand for these is often very variable compared with the mean, and it is more important to have a forecast which is generally reliable rather than one which requires some management.

Weighted averages

Benefits of weighted averages

The moving average technique suffers from two difficulties. First, it does not respond well to changes, as it gives equal importance to all periods and, second, a large amount of information is required to recalculate the average each time.

An average which takes more notice of recent history and less notice of older data has major benefits. The use of a weighted average improves historical forecasting. There are many options for weighting the averages

but there is one option which is better than the others and is used almost exclusively.

The technique of exponential smoothing avoids both the problems outlined above. It produces a weighted average which is based on all the information available (see Figure 10.4). The calculation of a new average only requires the old average, the new demand and a weighting factor.

Exponential smoothing

The simplest proper forecasting technique is exponential smoothing.

The name 'exponential smoothing' comes from the fact that contributions from history for an exponential curve go backwards in time. There is theoretically a contribution from many years of data, although this will become increasingly insignificant as it ages.

Forecasting with exponential weighting is the simplest of the professional forecasting tools. A new forecast demand, or historical average, is obtained by mixing a portion of the old average with a portion of the new demand.

$$\begin{aligned} \text{New forecast} &= \alpha \times \text{Demand in Last Period} + (1 - \alpha) \\ &\quad \times \text{Forecast for last period} \end{aligned}$$

The mixing portion, α , is the smoothing factor.

A new forecast has to be calculated as soon as the data for the previous week or month (i.e. last period) has been collected. The new forecast then applies to the week or month starting at that time.

The weighted forecast is easier to use when written down mathematically:

$$A = \alpha D + (1 - \alpha)B$$

where

A is new forecast

B is the old forecast

D is the demand in the period just completed

α is the smoothing factor.

Table 10.1 Exponential forecasting

Month	Sales	$\alpha = 0.1$		$\alpha = 0.3$		Cumulative Error
		Forecast	Error	Forecast	Error	
1	10	10	0.00	0.00	10	0.00
2	14	10.00	4.00	4.00	10.00	4.00
3	6	10.40	-4.40	4.40	11.20	-5.20
4	8	9.96	-1.96	1.96	9.64	-1.64
5	16	9.76	6.24	6.24	9.15	6.85
6	12	10.39	1.61	1.61	11.20	0.80
7	7	10.55	-3.55	3.55	11.44	-4.44
8	11	10.19	0.81	0.81	10.11	0.89
9	13	10.27	2.73	2.73	10.38	2.62
10	5	10.55	-5.55	5.55	11.16	-6.16
11	10	9.99	0.01	0.01	9.31	0.69
12	8	9.99	-1.99	1.99	9.52	-1.52
		MAD = 2.74		MAD = 2.90		

'A' can be described as the 'average', which it is, although it is being used as a 'forecast'. The value of α is between 0.1 and 0.4 as discussed below.

The smoothing effect of exponential smoothing is illustrated in Table 10.1 and Figure 10.3. Like moving averages, this calculation simply analyses the demand to give an average plus the remainder which is treated as statistical fluctuation. The smoothing factor, α , defines the amount of notice which is taken of each period's demand.

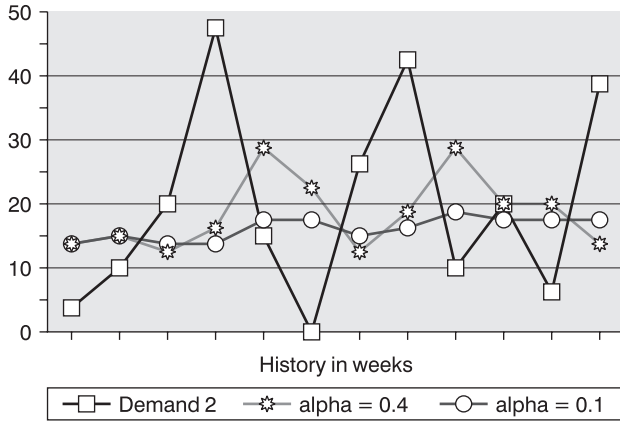
Selection of the smoothing constant, α

The nature of the forecast can be changed just as for moving averages, by altering the value of the smoothing constant, α . Low values of α make the forecast consistent; high values make it reactive to change. Making it too high would render the result unreliable because it would depend too heavily upon demand in the last period. For effective exponential forecasting, the range of values which can be chosen for α are 0.1 to 0.4.

If the demand varies slowly then a low value of α , say 0.1, is suitable (see Figure 10.3). If there are changes in average demand level to which the forecast should adapt then a high value of α , say 0.4, is chosen.

If the variability of the demand is high from month to month, then a low α is required to show up a stable average. Where the demand is smoother a larger value of α can be used. This responds more rapidly to trends in the average. If little is known about the item, or a value of α is required for a range of items, then a median value of $\alpha = 0.2$ can be used initially. The

(A) Variable demand



(B) Rising demand

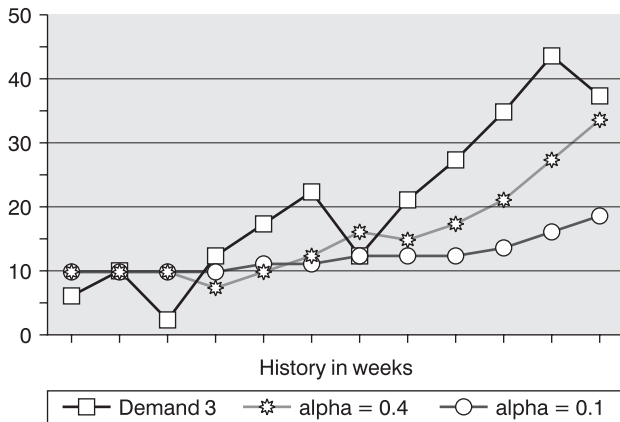


Figure 10.3 Effect of different values of α

value of α is then reviewed at regular intervals and altered where necessary.

In Figure 10.1 a sales pattern has been predicted each month using exponential smoothing with $\alpha = 0.1$ and $\alpha = 0.3$. The forecast for the next month has been worked out using the old average and the error. The success of the forecasting has been gauged by calculating the MAD in each case. The result indicates that $\alpha = 0.1$ gives better answers in this particular case.

Alternative formula

This formula has an alternative, more convenient, method of calculation using the error on the forecast instead of the actual demand, which differs from the original form as follows:

$$\begin{aligned} \text{New forecast} &= \alpha \times \text{Actual demand} + (1 - \alpha) \times \text{Old forecast} \\ &= \alpha \times \text{Actual demand} - \alpha \times \text{Old forecast} + \text{Old forecast} \\ &= \text{Old forecast} + \alpha (\text{Actual demand} - \text{Old forecast}) \end{aligned}$$

so

$$\text{New forecast} = \text{Old forecast} + \alpha (\text{Error in old forecast})$$

This formula is used as frequently as the original but it is important to ensure that any negative signs are taken into account when using this version.

In exponential smoothing (and moving average), the forecast for subsequent periods is the same as that for the next period. It is possible to project this forecast further ahead, but its accuracy will be poorer since it will not predict a trend in demand (up or down).

Let the computer do the calculations and keep them updated.

Contributions of history

The moving average takes data at full value until it becomes too old and falls off the end of the data table. With exponential smoothing, history never gets discarded, it just becomes less significant. The forecast for this week, with $\alpha = 0.2$ is 20 per cent of last week's demand and 80 per cent of history. But that 80 per cent was calculated from 20 per cent of the demand of the week before and 80 per cent of the previous history, i.e. last week is 20 per cent of the forecast, the week before 16 per cent (20 per cent of the 80 per cent) and previous forecasts are 64 per cent (80 per cent of 80 per cent). Taking it back further, the contribution from three weeks ago is 12.8 per cent and from four weeks ago is 10.24 per cent.

The contribution from each period to the forecast is shown in Figure 10.4 going back twelve periods for $\alpha = 0.4$, $\alpha = 0.2$ and $\alpha = 0.1$. The forecast using $\alpha = 0.4$ takes little heed of demands which are more than




Figure 10.4 Contribution from each period to the average

about six periods old. The forecast using $\alpha = 0.1$ still has significant contributions from twelve-month-old data, and demands from as long ago as eighteen months can affect the answer significantly.

Initializing the forecast

The fact that exponential smoothing depends on old history gives a problem of when to start the forecast. What should be taken as the ‘old forecast’ in the equation? There are two ways of solving this, namely, estimate and history.

For a new product, or for the addition of a new item to the stock range, it is best to make an estimate of expected demand to start off the forecasting. This can be done by:

- inspired estimate
- model or replicate a similar item
- an arithmetic average of such data as is available
- historical forecast.

The effect of this estimate will become reduced at the rate shown in Figure 10.4. It is not a good practice to start with zero as the ‘old forecast’ since it takes too long to settle down, especially with low values of α .

Where there is some usage information, a large number of items to forecast or inexperienced staff, then a historical forecast is better. If the

item already has eighteen months' worth of usage history, then the technique is to work out a six-month average for the demand up to one year before. This can then be used as the 'old forecast' to start twelve months of calculation using the exponential formulae to bring the forecast up to the present day. Once the initial forecast has been established, the exponential average is self-perpetuating.

Improved values for mean absolute deviation

Any tabulated information can be weighted using exponential smoothing. It is important that the forecast is weighted towards the most recent data, and that the variability data incorporates the most reliable up-to-date information to set the safety stocks. The calculation of MAD or SD should therefore be exponentially weighted to allow their wide application in assessing forecasts and calculating safety stocks with exponentially weighted MAD and forecast error functions. Where standard deviations are used, the weighting has to be applied to the variance (SD)². Updating of the MAD can be done most conveniently using the second version of the exponential smoothing formula:

$$\text{New MAD} = \text{old MAD} + \beta (\text{Error in MAD})$$

where

$$\text{Error in MAD} = \text{Actual absolute deviation} - \text{Old MAD}$$

and β is the MAD smoothing constant and is exactly the same as α but using β shows that it is the MAD that is being smoothed rather than the forecast.

It is normal to take $\beta = 0.1$ when smoothing the MAD. The reason for this is that the MAD reflects the characteristics of the demand pattern, and this would only be expected to alter gradually. If there is a sudden change in the variability of demand it is likely to be detected first through a change in demand level and forecast bias.

If the forecast for period 1 in the example (Table 10.2) is 36.2 and the MAD 8.5 from a previous assessment, the new MAD can be calculated either by taking the errors from the last six periods and averaging them to get a simple average, or by taking an exponentially weighted average. The last two columns in Table 10.2 show the difference between the simple and exponentially weighted average.

Table 10.2 Exponentially weighted MAD

Period (wk/mth)	Demand	Exponential forecast ($\alpha = 0.2$)	Error	Simple MAD (6 period)	Smoothed MAD ($\beta = 0.1$)
1	28	36.20	8.20	8.50	8.50
2	45	34.56	10.44		8.47
3	33	36.65	3.65		8.67
4	39	35.92	3.08		8.17
5	19	36.53	17.53		7.66
6	27	33.03	6.03	8.57	8.64
7	22	31.82	9.82	8.16	8.38
8	48	29.86	18.14	8.43	8.53
9	36	33.49	2.51	9.71	9.49
10	25	33.99	8.99	9.52	8.79
11	18	32.19	14.19	10.50	8.81
12	41	29.35	11.65	9.95	9.35
Forecast for period 13		31.68		10.88	9.58

Examples

1 MAD Forecast = 4, actual absolute deviation = 6

$$\text{New MAD} = 4 + 0.1 \times (6 - 4) = 4.2$$

2 MAD forecast (old MAD) = 4

Demand forecast = 17

Actual demand for period = 15

so actual absolute deviation = 2

$$\text{new MAD} = 4 + 0.1 (2 - 4) = 3.8$$

This calculation would be the same if the actual demand was 19. Although MAD always has a positive value, the error in the MAD can be positive or negative.

Exponentially weighed MADs, which give a better assessment of current variability, should be used wherever possible. An example using exponentially weighted averages and MADs is shown in Figure 10.5 together with the exponentially weighted errors of the mean. The exponentially weighted standard deviation is insensitive to variances during each time period because a low value of β is used. This is a good feature for ensuring a stable calculation of safety stock, but it does make the task of setting the initial SD more critical. The initial estimate used for MAD or SD will affect the

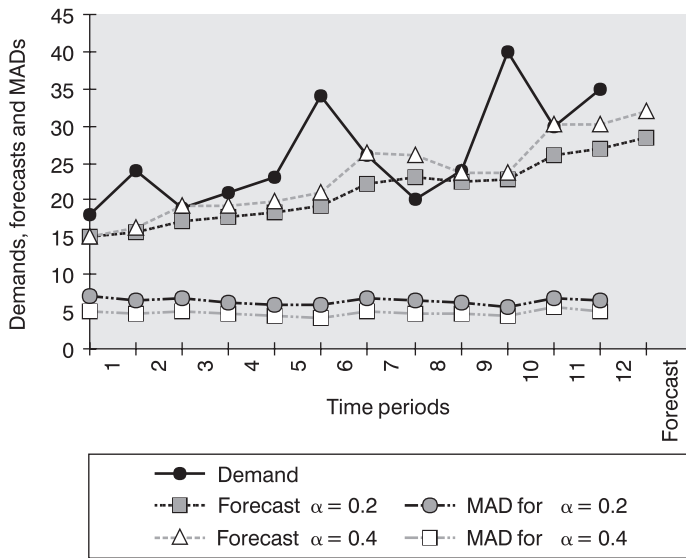


Figure 10.5 Exponential forecasts and MADs

calculation for many cycles, and can give misleading results if not chosen carefully. (An initial assumption of zero for MAD or SD is not good enough.)

In order to compensate for a change in variance or to reduce the effect of initial conditions, the weighting factor can be increased for a few periods but β should be kept below 0.2.

Choosing the best forecast

Application of forecast types

The forecast models discussed previously give a variety of estimates for the future, some of which will turn out to be better than others. The problem is how to choose the best option. At one time, it was necessary to decide in advance what the characteristics of the demand pattern were, and to choose the best model to use. Nowadays, technology can choose the best forecast, so long as the rules are defined.

The general principles which have been used to choose between the simple forecasting models are shown in Table 10.3. (The forecasting models are described in this chapter and the next.)

The exponentially weighted models have the advantage that they are more sensitive to recent events and, therefore, are generally to be preferred

Table 10.3 Uses of simple forecasting models

Forecasting model	Demand characteristics
Moving average	Static demand level, irregular pattern
Regression analysis	Continuous trend, irregular pattern
Exponential smoothing	Variable demand, gradual changing level, and slow moving spares items
Base series	Seasonal demand, mature products with low variability
Double exponential	Varying trend in demand, low variability, product phase-in and obsolescence
Fourier	Product with long history, cyclic demand
More sophisticated techniques	Much history often needed, for longer-term forecasting
Specific models	Based on market knowledge, potentially most accurate

over the simple averaging techniques such as moving average, regression, or trend measurement.

Table 10.3 can be taken as an indication of which method is likely to give best results in any circumstances, or it could be used to set the model for each item for forecasting. With the large number of lines in many stores, some gross assumptions have to be applied, and rules can be suggested such as those shown in Table 10.4.

Focus forecasting

With the technology now available it is not necessary to make these assumptions. The approach is to calculate a range of options and then to select the best average.

An assumption in making this selection is that the forecast which has been most successful recently will be the best one for the period being forecast. This assumption is open to debate, but is the best working option.

Table 10.4 Simplistic selection of forecasting models

Pareto class	Model	Method of use
A	Double exponential	Monitor and use knowledge
B	Single exponential	Exemption report and control
C	Moving average	Accept forecast
Other	Product specific	Supply to order

In inventory control the forecasts are usually for a short period ahead, making this concept more acceptable.

Do many forecasts and choose the smallest MAD.

The question still remains as to which forecast can be considered the best. The best forecast is the most accurate one, of course. The most accurate forecast is the one which has the least errors (errors are the difference between the actual demand and the forecast). Therefore, the choice of the best forecast can be made by finding the one which recently has demonstrably given the least errors.

Now the basis of a safety stock calculation is that stock is required to cover a divergence between the forecast and the actual demand (perfect forecasts mean no safety stock). The calculation of the safety stock is based on finding an average divergence (SD or MAD) as one does in working out the accuracy of the forecast. The MAD is a measure of the forecast errors – this is focus forecasting.

If there are, say, three different forecasts for the same item, created in different ways, then the way to choose the best one is to calculate the MAD for each (see Table 10.5). The best forecast is the one with the smallest MAD.

The use of a simple MAD is reasonable but, as in the case of safety stocks, a better measure is readily available. This is because a forecast which was brilliant a few months ago, could now be rather poor and it would be preferable to choose a forecast which has given reasonable results in the last couple of periods. For this reason a weighted MAD is used – in fact an exponentially weighted MAD (see the section on improved values for mean

Table 10.5 Three different forecasts for the same item

Month	Demand		Model	Forecast demand	Error (MAD)
1	28		12 period		
2	45	Moving	6 period	26.8	5.9
3	33	averages	$\alpha = 0.1$	22.0	7.4
4	39	Exponential	$\alpha = 0.2$	28.3	9.5
5	17	smoothing	$\alpha = 0.3$	24.7	6.2
6	27		$\alpha = 0.4$	23.8	5.7
7	22	Double	$\alpha = 0.2$	23.9	5.2
8	24	exponential	$\alpha = 0.3$	21.9	5.4
9	15	smoothing	$\alpha = 0.4$	21.8	5.0
10	18			21.4	5.7
11	28				
12	25				

absolute deviation, earlier in this chapter). This provides a more reliable method of assessing the potential of alternative forecasts.

Monitoring forecasts

Forecast tracking using cumulative sum of errors

The forecasting methods discussed will cope with most demands, but they need adjustment from time to time when the method is inappropriate to the conditions. We therefore need a monitoring process to enable prompt action to be taken.

Tracking signals provide a monitor of how well the forecast is performing. When the errors fall outside acceptable limits, the tracking signal identifies the problem or modifies the controls.

Example

The demand pattern for an item over the past twelve periods is shown in the two left-hand columns of Table 10.5. As a result the forecasts in the three right-hand columns were calculated. By looking at the MAD column, it can be seen that the best forecast is double exponential smoothing with $\alpha = 0.3$, and the worst forecast is exponential smoothing with $\alpha = 0.1$. Therefore the forecast to choose is 21.8. However at the next period end a different forecast might well prove to be even more accurate.

The Cusum technique provides a warning signal when there is a significant asymmetric shift of actual results from the forecast. It is very simply obtained by cumulating the errors in forecast over successive periods (see the last column of Table 10.1). When the forecast is correct, the actual demand is distributed equally above and below the forecast and cumulative error is small. Once the forecast diverts from the average demand, the errors build up period by period.

If the Cusum reaches an unacceptably high level, an error signal can be triggered and the forecast adjusted. The size of the error in the mean depends on the variability of the demand pattern and so the limit on acceptable forecast errors is set in terms of MAD – six times the MAD. However, for important items, values above four times the MAD should be investigated.

When the Cusum has signalled a problem by a high error value, the Cusum should be reset to zero. Otherwise the controlled forecast will give a level which is consistently high.

Monitor forecast performance.

Improving control – Trigg’s tracking signal

When the Cusum indicates that there is a significant discrepancy, the forecast needs to be changed rapidly. For exponentially weighted averages this means increasing the smoothing constant, α (see section on selection of the smoothing constant, α). Where there is a large variety of items this will result in many changes to be made. Trigg’s method is a way of feeding back the Cusum information automatically and using it to modify α . Instead of generating the cumulative SUM of error, the tracking signal first works out the average error. This is done using the exponential smoothing formula on the forecast errors and is called the ‘smoothed error’, calculated from the actual errors, in the normal way:

$$S_t = \delta e_t + (1 - \delta)S_{t-1}$$

Where S is the standard error, e is the actual error and δ is the smoothing factor for the errors. The subscripts on the standard errors, t and $t-1$, stand for this period just completed and last period. This is similar to the usual exponential smoothing of the MAD:

$$MAD_t = \beta e_{t-1} + (1 - \beta)MAD_{t-1}$$

δ and β are set to the same value (0.1 or 0.15).

The tracking signal provides a figure of merit for the quality of forecast, and is calculated by the ratio

$$\text{Tracking signal} = \frac{S_t}{MAD_t}$$

The tracking signal shows whether the forecast is good or not. It is considered good if it is not biased (i.e. not consistently high or consistently low) where S and MAD are the exponentially smoothed values. The difference between the calculation of S and MAD is simply that S takes into account whether the forecast errors are up or down, and the MAD only uses the size (absolute value) of the forecast errors. Hence the biggest value

that the ratio, the tracking signal, can have is 1, when all the errors are either positive or negative. The smallest value is zero where the smoothed errors compensate and are very small compared with the individual fluctuations in demand.

Trigg used this to give direct feedback into the α factor in the forecast. The argument is that an accurate forecast gives a low value for tracking signal, and it needs little modification, so a low value of α is best. If the tracking signal is large, there is a significant error in the forecast so a high α factor is required to alter the forecast. The tracking signal can therefore be used as the smoothing constant for the forecast, and this will enable an exponentially weighted forecast to be more self-compensating each period. For this purpose, the tracking signal has to be a positive number and so any negative tracking signals (TS) are changed to positive.

The forecast F_t is then determined by:

$$F_{t+1} = F_t + |TS_t| \times e_t$$

where e is the forecast error for the latest period, and the value of TS is always positive. The tracking signal can be used for feedback into exponential smoothing or double exponential smoothing. It is also possible, but less beneficial, to calculate a tracking signal using normal averages. To make the tracking signal work properly the value of δ has to be the same as the value of β . Since both of these are usually set to 0.1 this does not present a problem.

There is a danger when using Trigg's method that the α factors become very high (above 0.4) when it tends to overcompensate for small changes. In the discussion of exponential forecasting, it was stated that the value of α should be below 0.5. It is therefore advisable in practice to modify the formula so that values above 0.5 are not created.

The tracking signal is very useful and allows a computer-calculated forecast to monitor itself and warn the inventory manager when it is producing bad forecasts. The tracking signal can be used simply as a warning flag. If the tracking signal is high, say above 0.4, then a manual adjustment to the forecast is required, and an exception report can be triggered.

Trigg's forecasts can be a more accurate method.

The tracking signal can be combined with focus forecasting, the optimum forecast being selected by choosing the one with the smallest smoothed MAD as normal. If this forecast has a tracking signal below, say 0.4, then

the forecast is acceptable. If the tracking signal is greater than the unreliability level (0.4), then the forecast with the next smallest MAD can be chosen. This is a simple, mechanical and effective way of getting good forecasts across the range of stock items.

Example

In Table 10.6 the characteristics of the demand pattern change after period 12.

The forecast gradually drifts up and the MAD increases. The smoothed error which was very small up to period 12 starts to increase. This is because the change in forecast lags behind the actual demand level. Consequently the tracking signal increases. It reaches 0.5 in period 11, which is quite quick. At that stage the forecast can be readjusted or an alternative model used, or it continues to give poor results for a while (periods 12 to 14).

Table 10.7 shows what happens when the tracking signal is fed back to change the α as proposed by Trigg. The α factors are much smaller at the beginning. For the first few periods this is a result of the initial condition that the smoothed error is set to zero. The tracking signal stabilizes and responds less to the low demand in period 7, which is sensible given the consistently larger demand in the previous periods. When the demand level changes, the tracking signal increases and enables the forecast to rise faster to cope with the problem. However, this does not fully compensate and the tracking signal still rises above 0.4. The forecast rose quickly to the new level as the α factor increased. It did this without the inventory controller having to intervene. The effect is shown graphically in Figure 10.6, which contrasts the demand, the forecast with $\alpha = 0.2$ and the forecast using Trigg's method. The data are taken from Tables 10.6 and 10.7 and represents a typical demand level change. The only problem in period 19 is that the tracking signal is still rather large. As the smoothed error continues to decrease because the forecast is again accurate, this tracking signal will reduce gradually.

Although it is not the ideal answer, since it lags behind the demand pattern significantly in a period of change, it is an improvement where there are many stock items and little time to manage them individually. However, it would be expedient to limit the maximum value of tracking signal fed back to 0.4 and to manually readjust when it exceeds this value.

Table 10.6 Forecast tracking

Period	Demand	Exponential forecast ($\alpha = 0.2$)	Error in forecast	Smoothed MAD $\beta = 0.1$	Smoothed error $\delta = 0.1$	Tracking signal
1	13	16.00	-3.00	4.00	0.00	0.00
2	1	15.40	-14.40	5.04	-1.44	-0.29
3	10	12.52	-2.52	4.79	-1.55	-0.32
4	17	12.02	4.98	4.81	-0.89	-0.19
5	11	13.01	-2.01	4.53	-1.01	-0.22
6	19	12.61	6.39	4.71	-0.27	-0.06
7	16	13.89	2.11	4.45	-0.03	-0.01
8	15	14.31	0.69	4.08	0.04	0.01
9	26	14.45	11.55	4.82	1.19	0.25
10	38	16.76	21.24	6.47	3.20	0.49
11	31	21.01	9.99	6.82	3.88	0.57
12	40	23.01	16.99	7.84	5.19	0.66
13	35	26.40	8.60	7.91	5.53	0.70
14	32	28.12	3.88	7.51	5.36	0.71
15	Forecast	28.90				

Table 10.7 Forecast using Trigg's tracking signal

Period	Demand	Trigg's exponential forecast	Error in forecast	Smoothed MAD $\beta = 0.1$	Smoothed error $\delta = 0.1$	Tracking signal
1	13	16.00	-3.00	4.00	0.00	0.00
2	1	16.00	-15.00	5.10	-1.50	-0.29
3	10	11.59	-1.59	4.75	-1.51	-0.32
4	17	11.08	5.92	4.87	-0.77	-0.16
5	11	12.02	-1.02	4.48	-0.79	-0.18
6	19	11.84	7.16	4.75	0.00	0.00
7	16	11.84	4.16	4.69	0.42	0.09
8	15	12.21	2.79	4.50	0.66	0.15
9	26	12.62	13.38	5.39	1.93	0.36
10	38	17.41	20.59	6.91	3.79	0.55
11	31	28.72	2.28	6.44	3.64	0.57
12	40	30.01	9.99	6.80	4.28	0.63
13	35	36.30	-1.30	6.25	3.72	0.60
14	32	35.52	-3.52	5.98	3.00	0.50
15	Forecast	33.76				

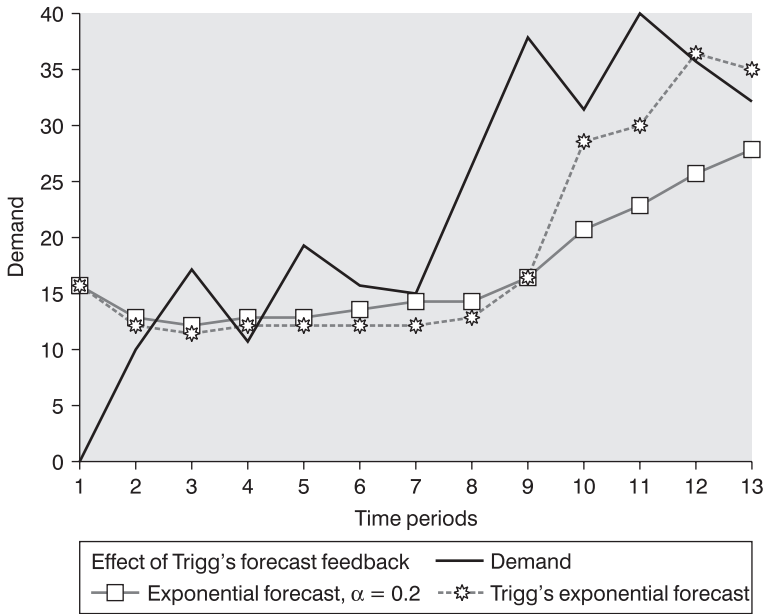


Figure 10.6 Effect of Trigg's forecast feedback

Key points

- Historical forecasting is essential for basic stock control.
- It is necessary to add information about future events.
- A rolling average gives a basic forecast.
- This can be greatly improved using exponential smoothing.
- Carry out several forecasts, then select that with the smallest MAD.
- A tracking signal shows when the forecast is wrong.

Advanced forecasting methods

- Better forecasting for rising or falling demand
- Coping with seasonal demand
- Other techniques

More forecasting tools

Improved forecasting techniques

Exponential smoothing is the basic model, which can, of course, be improved. It splits demand into two components:

- average, which is recalculated each period
- random fluctuations about the averages.

The next stage of complexity is to add either:

- a constant increase or decrease in demand, using double exponential smoothing, or
- a seasonal variation of demand, using base series forecasting.

Alternatively, entirely different approaches to intrinsic forecasts can be used giving:

- unweighted average trends (regression analysis)
- profitability of changes in demand (Bayesian forecasting)
- curve fitting (Fourier analysis)
- or more complex methods. (These are noted but not discussed as exponential methods are a better option.)

The mathematics can become quite complex and are usually left to the computer. In general, the more complex the model the more history is

required which is why simple forecasting has been acceptable until recently. Intrinsic forecasts are best for data collected under similar conditions, even when the demand is erratic. Then it is quite possible that exponential forecasting with $\alpha = 0.1$, or even long moving averages are best. If a change from single to double exponential smoothing gives a significant improvement in the forecast for an item, then perhaps more sophisticated methods would work better, but there are fewer and fewer items benefiting from the better forecasting techniques. As a general rule forecasting can benefit more by improving the customer communications to detect changes in the market than by going to highly sophisticated models.

Double exponential smoothing

As stated in the last section, single exponential smoothing splits the demand into 'average + fluctuations'. This is improved in double exponential smoothing by analysis into 'average + trend + fluctuations.' In Figure 11.1




Figure 11.1 Improved forecasting through double exponential smoothing

the difference in forecast can be seen. The exponential forecast for the next period will show a weighted average, whereas by adding the average trend, the predicted rate of demand will be higher, and the forecast will be better.

To understand double exponential smoothing it is best to go back to simple mathematics. Figure 11.2 shows a normal straight-line graph given by the formula $y = bx + a$, with slope m and cutting the y axis at point c . The double exponential model is equivalent to this straight line, with y being the demand and x being time. Single exponential smoothing is like the line ' a ' above the x axis.




Figure 11.2 Double exponential model

In double exponential smoothing the forecast of demand, F_t , is therefore composed of two factors. $F_{t+1} = a_t + b_t$, with one factor arising from the movement of the average due to random fluctuations, similar to single exponential smoothing. The second factor calculates the trend of variations and adds a ‘slope’ factor. This has important advantages because the extrapolation to future periods continues, the increase or decrease being observed rather than assuming a static demand level.

In addition to this, double exponential smoothing can also compensate for lag in the forecast. It is obvious in graph (C) of Figure 10.1 that the forecasts are lagging behind an increasing demand level (and will signal a tracking error). Double exponential smoothing can bring the forecast in line with the demand once a linear trend has been established. This can be very useful when there is a gradually increasing or decreasing average demand.

The lag in an exponentially weighted average is simply given by the expression

$$\text{Lag} = \left(\frac{1 - \alpha}{\alpha} \right) \times \text{Trend}$$

where α is the smoothing factor.

From Figure 11.2 it can be seen that the double exponential forecast for one time period ahead is $a + b$ (demand + trend). The forecast for two periods ahead will be $a + 2b$ (demand + twice trend).

There are several versions of the double exponential formula. The simpler version is Brown’s model. A more comprehensive model is Holt’s

model, but this requires two smoothing constants (one for the base forecast 'a', and one for the trend 'b'). Holt's method does allow optimization of the smoothing factors for selected items by iterative programming.

The simpler formula is adequate normally and uses only one smoothing factor, which can be denoted as α since it behaves like the single exponential smoothing factor. The method of calculating values of 'a' and 'b' and the forecast is shown in Figure 11.3.

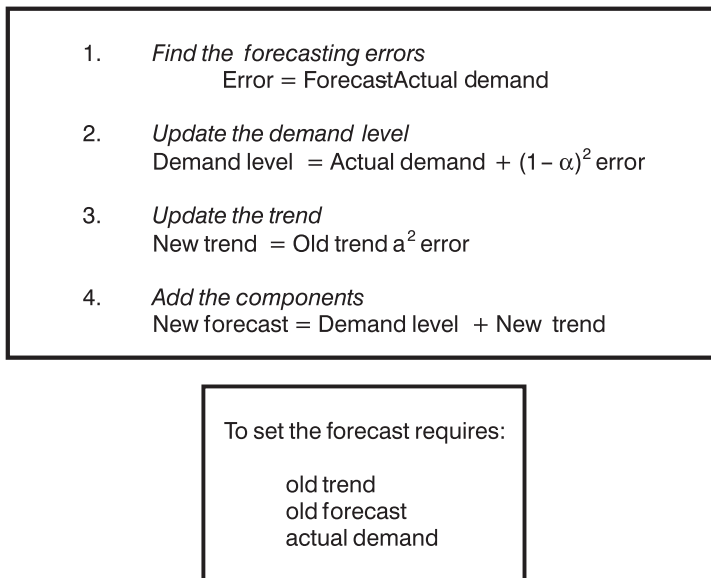


Figure 11.3 Double exponential smoothing

The general forecast for n periods ahead is:

$$F_{t+n} = 'a' + 'b'n$$

By this linear extrapolation of trend, a much more accurate forecast of demand in future periods can be made, if a linearly increasing or decreasing sales pattern is expected.

For items with widely varying demand patterns, the estimated trend can be misleading. A high or low value for the previous month's demand reflects too heavily on the calculated trend. If α is reduced to compensate, then the model becomes unresponsive. In this case it is better to take a longer sample period to even out the statistical variations, and this is often done by

Table 11.1 Double exponential forecast

Time period	Forecast F	Demand D	Error e	Demand level a	Trend b
				65	-5.00
1	66.0	55	5.0	58.2	-5.20
2	53.0	50	3.0	51.9	-5.32
3	46.6	45	1.6	46.0	-5.38
4	40.6	55	-14.4	45.8	-4.81
5	41.0	80	-39.0	55.0	-3.25
6	51.8	105	-53.2	70.9	-1.12
7	69.8	130	-60.2	91.5	1.29
8	92.8	155	-62.2	115.2	3.77
9	118.9	180	-61.1	140.9	6.22
10	147.1	205	-57.9	168.0	8.53
11	176.5	230	-53.5	195.8	10.67
12	206.4	255	-48.6	223.9	12.61
13	236.5				

Note: $\alpha = 0.2$

calculating the average demand over the last three periods and using the averaged error to modify ‘a’ and ‘b’.

The trend calculation is smoothed through the demands over preceding periods. Therefore, the greater the smoothing the slower is the trend to respond to change in trend. This causes lag – the amount of lag is shown in Table 11.1.

Changing forecast methods to double exponential smoothing means that the initial values of ‘a’ and b_{t-1} have to be estimated. The smoothing constant behaves like the smoothing constant for single exponential smoothing. The larger the value, the faster it responds to change, but the more unreliable is the value.

As the constant is squared for double exponential smoothing, selection of the correct value is more important. Lower values of α are chosen. In Figure 11.4 the equivalent future demand is shown for single and double exponential smoothing.

The method of estimating ‘a’ and ‘b’ depends upon how much information is available. If the demand for one period only is known or an estimate of it, then there is little alternative to setting ‘a’ equal to that value, ‘b’ to zero, and using a high value of smoothing constant so that the forecast will respond rapidly to the actual trend. Thereafter, the smoothing constant can be reduced to a more stable value.

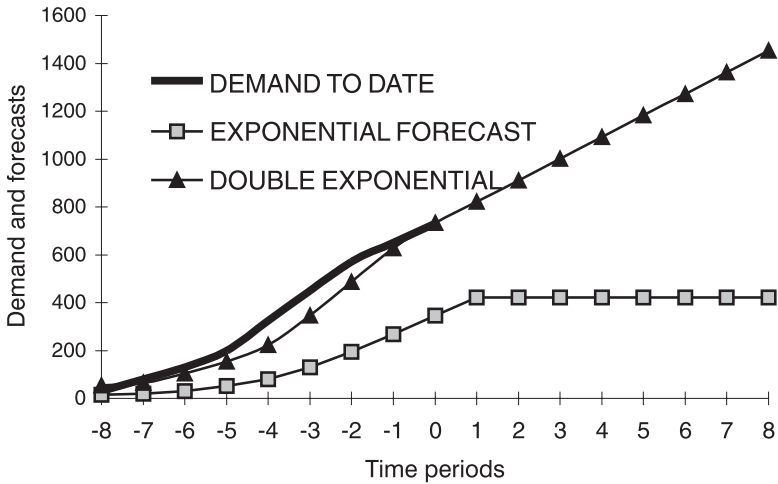


Figure 11.4 Forecasting several periods ahead

For an established product, the history is known and the values of ‘a’ and ‘b’ can be estimated accurately. In this case it is best to plot a graph and estimate the value of demand, F_t , and a trend ‘ b_t .’ Then

$$‘a_t’ = F_t - b_t$$

These values of ‘a’ and ‘b’ can then be used as a starting condition for double exponential forecasting.

An example of double exponential smoothing is given in Table 11.1 and the effect of single period forecasting in Figures 11.4 and 11.5.

Forecasting for seasonal sales

Forecasts do not cope with seasonal demand without using base series.

For many products seasonal or cyclic demand patterns exist. This can be caused by natural seasonal factors, for example temperature, rainfall or the probability of thunder storms, or by other seasonal factors like holiday patterns, Christmas or other religious festivals, or financial year ends. These factors can be the overriding cause of demands, or a small reason for fine tuning the forecast. Seasonal demands can be masked by the sporadic

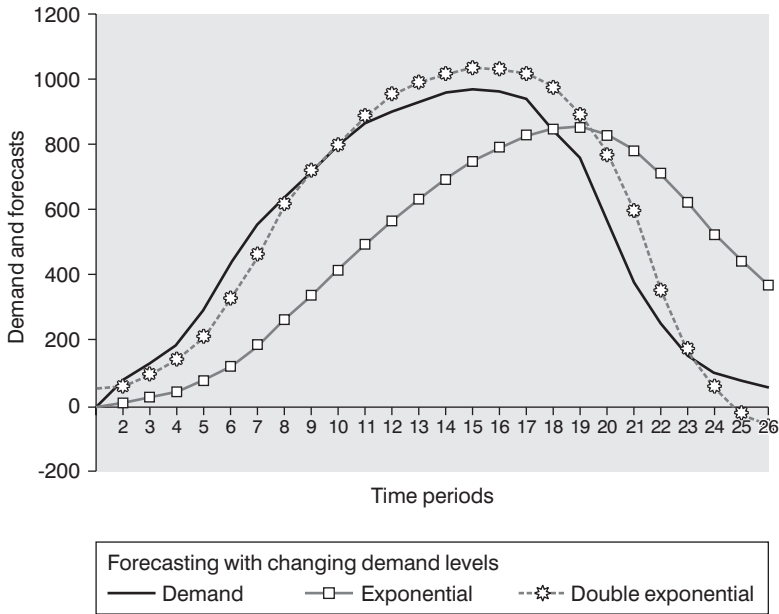


Figure 11.5 Forecasting with changing demand levels

nature of the demand for an item. Where this happens, as is often the case for C class items, it may well be expedient to ignore the seasonal variation and treat this as a random fluctuation too. Here a twelve-month moving average is a simple solution, although the safety stock will, of course, inflate the service outside the faster moving season and cause more stockouts during the season.

Historical data

To use history to forecast seasonal demand requires at least a year’s worth of data! In fact the more years for which data are available, the more reliable the forecast, unless there has been a change in the pattern of requirements for the item. This is usually a matter of interpretation, since the data is often distorted by specific, non-repeating factors to do with the market, price changes at irregular intervals or large contracts. The historical data often has to be interpreted to get the best forecast.

Generally seasonal data is collected in monthly buckets, although where Christmas dominates the sales, weekly information around November and December is often needed. The average demand for each corresponding month for three years or more is averaged so that an

average monthly demand pattern emerges. It is useful to use weighted averages of demand to bias the data in favour of the most recent years. Where history is shorter, seasonal data can be used but there is less confidence in the results.

Base series technique

The classic technique for calculating seasonal demand is base series forecasting. This is a method of producing a seasonal demand forecast as an extension to exponential smoothing, and makes the following assumptions: that the seasonal variation is known (from the data) and that forecast errors are due to changes in the average demand level or to random fluctuations.

The process for updating the forecast is:

- 1 Take the new month's (or week's) demand.
- 2 De-seasonalize it.
- 3 Compare with the forecast.
- 4 Update the forecast for the next months (exponentially).
- 5 Re-seasonalize the forecast.

Before the forecasting can start, seasonal demand factors have to be calculated, and an initial forecast made for the average demand. The processes 1 to 5 are then carried out.

Example

Demand for July is forecast as thirty-six. This is because the average monthly demand is thirty and history has shown that the demand next month is typically 20 per cent above average. At the end of the month, the records showed that thirty-three were sold. What is the likely demand for August, which has typically proved to be 10 per cent above average?

The demand in July should have been $30 \times 1.20 = 36$, but instead it was 33, which is 27.5×1.2 . (The assumption is that the seasonal factor is correct.) The de-seasonalized forecast should have been 30 but was 27.5. Therefore the forecast was too high and can be exponentially smoothed downwards (assuming α is 0.2 here) as follows:

$$\begin{aligned}
 \text{New de-seasonalized forecast} &= \alpha \times \text{de-seasonalized} \\
 &\quad \text{demand} \\
 &+ (1 - \alpha) \times \text{old} \\
 &\quad \text{de-seasonalized forecast} \\
 &= 0.2 \times 27.5 + (1 - 0.2) \times 30 \\
 &+ 29.5
 \end{aligned}$$

If the de-seasonalized forecast is 29.5, and the demand for August is 10 per cent above average, then the new expected sales for August are 29.5 + 10 per cent of 29.5, i.e. 32.45.

The de-seasonalized demand forecast of 29.5 can be used to forecast all future months by multiplying by their respective seasonal factors.

Other methods

From the wide variety of forecasting methods available, only the simplest have been described. Others such as regression analysis, Fourier analysis and Bayesian forecasting are based on different assumptions about the demand pattern.

More comprehensive approaches such as Holt Winters and Wagner Within, which extend the methods to greater sophistication, can be used to give more accurate forecasts, but require more historical information.

As the quality of sales information is often inconsistent, the opportunity to use refined techniques is limited but the development of improved forecasting will make it easier for the stock controller to use better techniques in future. Meanwhile, application of exponential techniques is adequate for many purposes although changes will be required in other areas including:

- the measurement of demand not sales
- the segregation of stock from non-stock demand for an item
- time bucket control.

The application of refined forecasting will then provide useful benefits.

Key points

- Double exponential smoothing gives tangible benefits where there are
 - significant trends in demand
 - long lead times for A class.
- Double exponential works best where demand is consistent.
- Seasonal demand has to be forecast using base series.
- More sophisticated techniques need better data and may not make a significant improvement in the forecasts.

Material requirements planning – an alternative to forecasting

- Identifying dependent demand
- The MRP process
- Benefits of using MRP
- Structured demand and bills of materials
- Triggering orders – projected available
- Meeting demand – available to promise

Avoiding uncertainty

When there is a clear picture of the demand for a particular item then stock can be minimized. No stockholding would be needed if demand were known precisely and far enough in advance because supply could be matched exactly with demand.

For items outside the range stocked, the customer orders and receives the supply after a lead time. There may be some transient stock because the supplier is asked to deliver in advance of the time that the customer requires it. The greater the confidence in supply, the shorter time items can spend in stores.

Treat demand as dependent to reduce inventory.

The situation with non-stock supply is mirrored for raw materials in manufacturing and for production components in assembly. Stock is required to feed the process, and firm plans exist for usage rate just as for scheduled customer orders and non-stock supplies. The approach to these types of inventory is to treat them as dependent demand rather than the independent demand as discussed in the previous chapters. The contrasts between dependent and independent demand are shown in Table 12.1. The essential feature of dependent demand is that it is calculated from the

Table 12.1 Dependent and independent demand

Independent	Dependent
Inventory-level system	MRP system
Forecast	Calculated
Keeps stock	Supply as required
All lines separate	Lines co-ordinated
Reactive	Proactive
Good customer service	Very good customer service
High inventory	Low inventory

demand of the next item up the supply chain; the aim is to have stock when it is to be used and no stock the rest of the time.

The inventory profile for dependent demand is shown in Figure 12.1. When the need is identified an order is placed. After the supply lead time, the delivery arrives. The items are then transferred or transformed for the customer and, when this is complete (which takes the process lead time), then the item can be sent out. The stock is therefore only on site during the process lead time and the investment in stock is very low. There is potentially some safety stock included in case the customer order quantity turns out to be greater than that originally agreed. The safety stock here can be calculated from the forecast errors; in this case the ‘forecast’ is the originally agreed delivery quantity. Ideally the safety stock should be zero.

A basic concept used in planning dependent demand is backward scheduling. This is simply the process of taking the delivery date required by

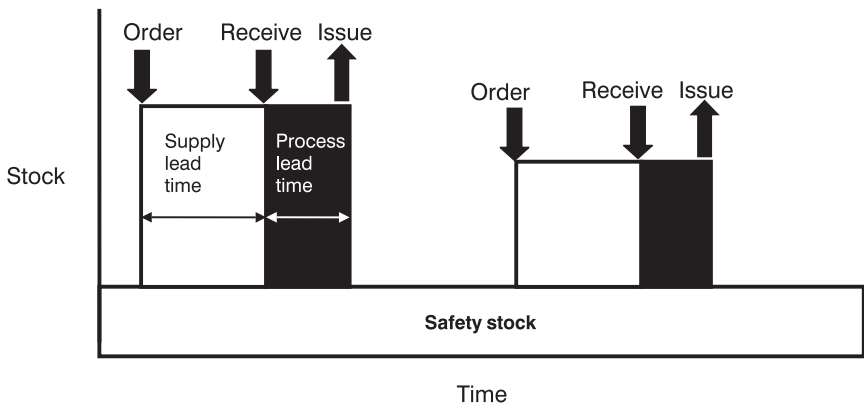


Figure 12.1 Inventory levels for dependent demand

the customer, and working back to identify when the item needs ordering from the supplier. 'Offsetting by the lead time' is the jargon used for this method. In Figure 12.1 the 'issues' should be made at the time requested by the customer. Moving back by the 'process lead time' gives the time to start on that order. If the system is working properly, this should happen just as the supplier delivered (marked 'receive' in the Figure 12.1). Going back further by the supply lead time identifies when to place the purchase order.

If this calculation results in a date which has already passed, the options are either to find a way of reducing the lead time or to forward schedule from today's date. This forward schedule will then estimate a later customer delivery date, so that the customer can be informed and the plan delayed.

Material requirements planning

The structure

There is a well-developed technique for planning dependent demand called material requirements planning. This is a generally applicable technique for all types of dependent demand. The basic concept is to have stock when it is needed and to have none the rest of the time. With dependent demand, the size and timing of the requirements are known from the next level down the supply chain. Material requirements planning can therefore give good control. It is used extensively in manufacturing because there can be several levels of dependent demand in a product (assemblies made from sub-assemblies made from components made from purchased materials.) Here MRP is well developed and complex. In many inventory situations MRP is a much simpler process.

The MRP process

Material requirements planning is a useful planning tool outside manufacturing. Normal stock management treats each item as completely independent, unaffected by usage rates of any other item and the demand is assumed to happen at random. The MRP approach is entirely different, based on:

- interdependent usage rates between stock lines
- acquiring stock for when it will be used.

Many customers tend to use items together, so demand is not truly independent. For example, there must be a connection between the

demand for hinges and doors, paint and paint brushes, nuts and bolts, although, in general business, there is no certainty that the sale of one item will be matched by the sale of the linked item. If a customer buys 100 bolts, there will be no automatic order for 100 nuts. They may already have a stock of nuts, or use some in threaded holes or, even, buy from elsewhere. The situation will become clear after a while. For example, it may work out in the long term that for every 100 bolts bought by a customer, they require forty nuts of one type and twenty of another. This relationship could be consistent, because the customer's use is consistent, sometimes without identifying the reason. The demand for these items could therefore be interlinked, and could vary as their level of business changes. It is important, therefore, for the items to be balanced in stock.

Material requirements planning links in the inventory systems with the sales and financial plans of the company. It is a professional technique for a modern company and it requires adequate information, particularly:

- a master schedule of planned supply to customers for each product, projected far enough ahead to permit ordering of bought-out items, raw materials and components. This 'planning horizon' must be long enough to cover the procurement lead time for bought-out items plus the total manufacturing time
- a well defined 'bill of materials' showing purchased items and manufactured components to identify what to make and buy
- lead times for the purchase or manufacture of all parts, realistic times taken
- accurate records of stock, work in progress and on-order parts.

Material requirements planning enables parts to be scheduled for the day that they are required. The advantages over the traditional inventory level systems are shown in Table 12.2.

Table 12.2 Advantages of MRP over traditional inventory levels

Inventory-level system	Material requirements planning
Treats each part individually	Deals with structure
Depends on demand history	Looks at future plans
Assumes average	Handles erratic demands
Aims to keep stock levels up	Holds stock only to cover demand
Priorities inflexible	Sensitive to priority changes
Mainly runs itself	Needs managing

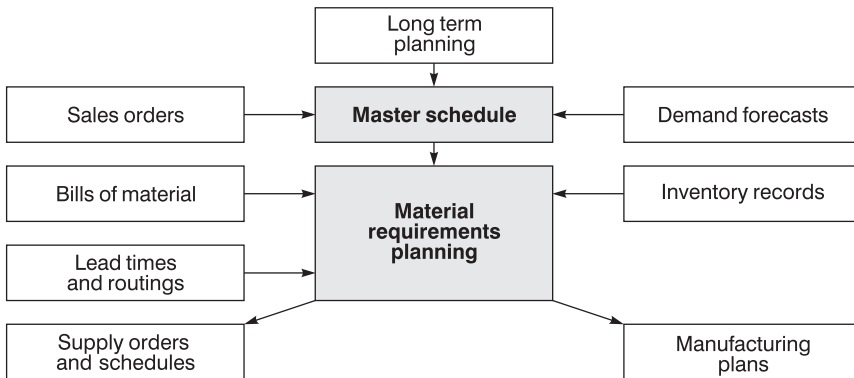


Figure 12.2 Material requirements planning

The two main attributes of MRP, namely ‘time-phasing’ and ‘structure’ are now being applied in a wide variety of stocking situations outside manufacturing. If it is known, or can be forecast, when major customers want their supplies, any business can create a master schedule, with the consequential reduction in stocks.

Material requirements planning is a business planning tool as well as a material supply calculation. The planning of MRP is shown in Figure 12.2. The process, known as manufacturing resource planning (or MRP II), starts with a long-term plan, which identifies what range of items are offered, how fast, with what market focus, and all the policy decisions to enable inventory control to operate successfully. As a result of this long-term process, actions should have been taken to ensure that the resources (people, plant, systems, logistics, organization) are available to make the policies work successfully. This long-term plan stretches ahead at least six months and, typically, by two to three years, depending how fast-changing the market is. Long-term planning is usually carried out by product group rather than by product. A master schedule is generated from the long-term plan, and is simply a list of what the business is going to send out, allowing for the real constraints (see section below).

Material requirements planning logic

The MRP backward scheduling process starts by identifying the customer requirement in the ‘master schedule’. Once the output to customers has been decided, then the focus is on achieving it. If there is already stock, how far ahead on the schedule will it last? If there is no inventory, then what has to be purchased – the item itself or are there components which are needed?

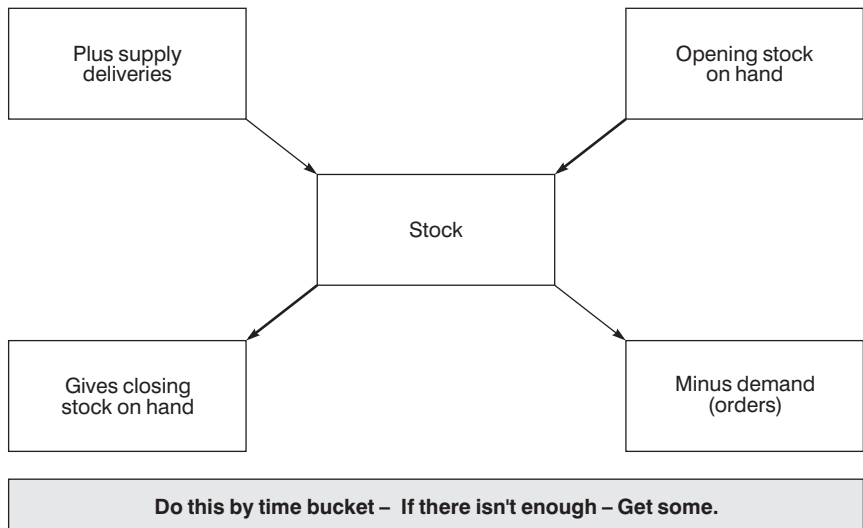


Figure 12.3 MRP logic

A bill of materials is required to cover this. At the same time, information on how long the supply will take and where it will be done – the lead times and routings – is needed. All this information is fed into the MRP calculation.

The actual MRP calculation is very simple. First it batches up all the requirements within each time bucket (week or day) so that there is no differentiation between different customer orders. Inventory is deducted from this total requirement. If the inventory is sufficient to cover the demand, then the same calculation is carried out for the next period. If the inventory is insufficient then delivery (order receipt) is required.

The process is simply shown in Figure 12.3.

$$\text{Initial stock} + \text{Receipts} - \text{Demand} = \text{Closing stock}$$

If closing stock is above zero (or safety stock) repeat calculation for next 'time bucket'. If closing stock is below zero (or safety stock) then create more receipts by ordering more. This then causes extra supply for orders and demand for that component.

This netting off process is carried out for all periods of the final process first, and then it is repeated for the supplying level (components or assemblies) if there is one.

The result is a list (or a series of lists) of what has to be done by when; a list of when to place and receive purchase orders and, where appropriate,

a manufacturing ‘work to’ programme for each production area (called ‘manufacturing plans’ in Figure 12.2).

Bill of materials

Classic stock control deals with each item individually, independent of all others. However, for most inventory, the usage of one item is linked to that of others, e.g. printer cartridges and paper. Inventory controllers need a way to acknowledge this link and to save the work of two forecasts where one will suffice (see Table 12.1). This can be done using an inventory bill of materials.

Traditionally, a bill of materials (BOM) is the list of parts, ingredients or materials needed to produce or assemble the required end product. The bill of materials for a product is not simply a parts list. It contains more information. Most products are packaged and the box, packaging material, pallet and even the documentation is part of the full BOM. The BOM can also include essential tooling for production. This is a useful way of planning for all the resources necessary for the process.

In manufacturing, an assembly is made up of components. If they are all bought out and given to an operator to assemble and pack, then the BOM




Figure 12.4 Balanced demand – bill of materials

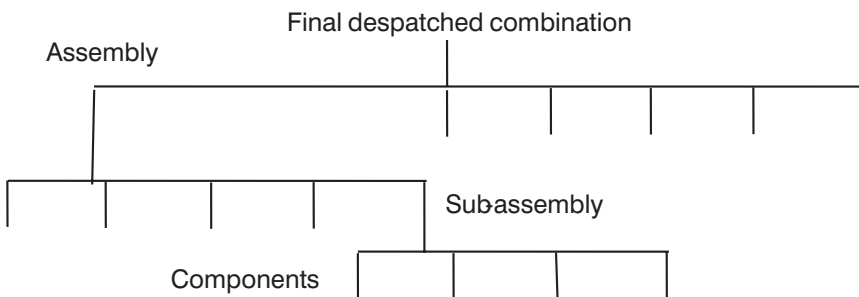


Figure 12.5 Structured demand – bill of materials

is as shown in Figure 12.4. This is a flat BOM. If the assembly and packing are separate, then the BOM is as shown in Figure 12.5. In a supply chain this would correspond to different suppliers of items which are to be sold together in the end.

Time delays between processes require recording and therefore add levels on to the BOM.

Planning the structure of supply

Example

In many situations there is a balance between the usage of different items. For example, DIY stores know that someone fitting a door hinge will probably need screws. They will either have the screws added to the pack, or remind the purchaser that they are needed.

Many activities result in the use of a variety of items, and if this usage is a fixed ratio, or even if the ratio of usages is consistent on average, then the ratio can be expressed in a loose 'bill of materials'. This could be used by a supplier providing a range of items to major customers. If the causes of use are understood, a planning (statistical) BOM can be created even though there is no manufacture involved. The use of this tool cuts down the amount of forecasting, since one forecast will provide the information for all the items included in the BOM.

Find similarity in demand, then use it.

The inventory control BOM usually consists of two levels, the same as a JIT manufacturing bill. The top level is the demand profile for the collection of stock items. The level below is the dependent demand for each of the item codes. Since the relative usage rates are not entirely fixed, the bill has to be monitored and corrected according to the usage changes of the ratios. This situation is very similar to a manufacturer where there is major variation in process yield or there is product scrap.

The demand for each item within the statistical BOM will vary for individual customer orders. This means that actual safety stock of each item will consist of two components:

- safety stock against the error in the overall forecast
- individual line safety stock against the variances in the demand BOM.

In practice, the statistical BOM is an excellent and underutilized method for distributors to give enhanced service where customers buy a range of items.

Master planning

Master scheduling

Improvements over classic stock control can also be made by considering exactly when the demand will occur.

Example

If there is a known demand for fifty this week, a hundred next week and sixty the week after, a knowledgeable traditional controller could work out the average demand as seventy per week with an MAD of twenty and then decide on the appropriate safety stock. A different approach would be to identify the demand in each week as fifty, a hundred and sixty and arrange supply accordingly.

To do this requires a new type of record. Instead of the normal record of what has been issued each week, the requirement is for a list of what will be issued each week, showing individual weeks (or days) into the future. This is a master schedule. The result of using a master schedule instead of a review level is that the inventory is kept lower and large fluctuations in demand can be accommodated. There is still the option to keep some safety stock, calculated from the inaccuracy of the schedule each week.

The master scheduling process for receiving orders converts the existing demand forecast into firm customer demand by placing the order on the master schedule. This is a management function, not a clerical function. The master scheduler is responsible for arranging customer demands into an achievable plan. Normally, particularly outside the demand lead times, the forecast quantity will be greater than the total quantity ordered for each time period (a day or a week). However, if orders are received which exceed the forecast, decisions are required to deal with this.

Example

If the forecast is a hundred per week and the first order is for seventy in the first week, then it is no problem. If a second order is received for thirty for the first week it can also be supplied. However, if a third order is received for fifty for that week, this will have to be met from the forecast in the second week unless it is convenient to change the schedule.

The logic of allocation is, therefore, that the forecast is filled up by real customer orders to the forecast level each week. However, if the customers require more, then the master scheduler has to decide whether to:

- provide later delivery or part shipment to the customer
- reallocate goods already promised to another customer
- overload the forecast and cause panic buying and inefficiency.

The last option is commonly used but should be avoided since it upsets supplier relations. There is sometimes the option of increasing capacity by overtime, but the material supply chain may not be able to react to short-term schedule changes.

The master plan in its simplest form is a list of how many of each item the company is intending to supply to the customers each week (or each day for JIT business) – see Table 12.3.

Master planning takes the forecast and plans to supply parts to meet this demand at the right time. In the more detailed master plan (Table 12.4) the scheduled receipts are the delivery date of current firm outstanding supply orders, and the stock on hand is the inventory which results if no action is taken. The first three columns show the current situation. The remaining

Table 12.3 Master schedule

Item	Week 1	Week 2	Week 3	Week 4	Week 5	Week 6
Product 1	1000	500	2300	600	1300	1300
Product 2	300	150	450	300	350	600
Product 3	3000	2000	3000	3000	3000	3000
Product 4	1500	~	~	~	~	~
Product 5	~	~	~	~	~	~
Product 6	~	~	~	~	~	~

Table 12.4 Planned supply – planned order release

Week	Forecast	Scheduled receipts (a)	Stock on hand	Planned order receipts	Projected available	Planned order release (b)
Current			500		500	0
1	3000	5000	2500		2500	0
2	4000	5000	3500		3500	5000
3	2000		1500		1500	0
4	2000	5000	4500		4500	5000
5	4000		500		500	5000
6	3000		-2500	5000	2500	0
7	1000		-3500		1500	0
8	3000		-6500	5000	3500	0
9	4000		-10 500	5000	4500	0
10	2000		-12 500		2500	

Note: Assumes a fixed batch size of 5000 and a supply lead time of four weeks.

three show what has to be done to meet the forecast demand (i.e. plan in new supply).

Set up supply using projected available.

The planned order receipts are calculated to minimize stock and, at the same time, give the information necessary to supply forecast. The resultant stock level is shown in the next column headed ‘Projected available’, i.e. the stock levels resulting from these actions. With a four-week supply lead time, the order release (last column) has to precede the order receipt by four weeks. If a planned order release occurs in the current week, then an order has to be placed. Alternatively the planned order release column is used as the time phased gross demand for the items on the next level down on the BOM, or as a supply schedule for purchasing.

Placing the order results in the order quantity being deducted from the planned order release and the ‘planned order receipt’ being transferred to a scheduled receipt. If a safety stock is appropriate instead of letting stock go down to zero, then the stock on hand can trigger a planned order receipt at the safety stock level instead of negative stock. A safety stock of 1000 would mean that stock would be too low by the end of week 5 (Table 12.4). The planned order receipt and release would therefore be required a week earlier.

Apply it to customers using available to promise.

Delivery promising

The master plan shows what is available for customers and when.

When the customer orders are received, the forecast quantities should be available. The calculation of projected available and planned order release enables the supply to be planned. This calculation accepts the forecast and works out what supply is required to meet it. The amount of the item remaining to fulfil outstanding demand can be calculated in a similar way

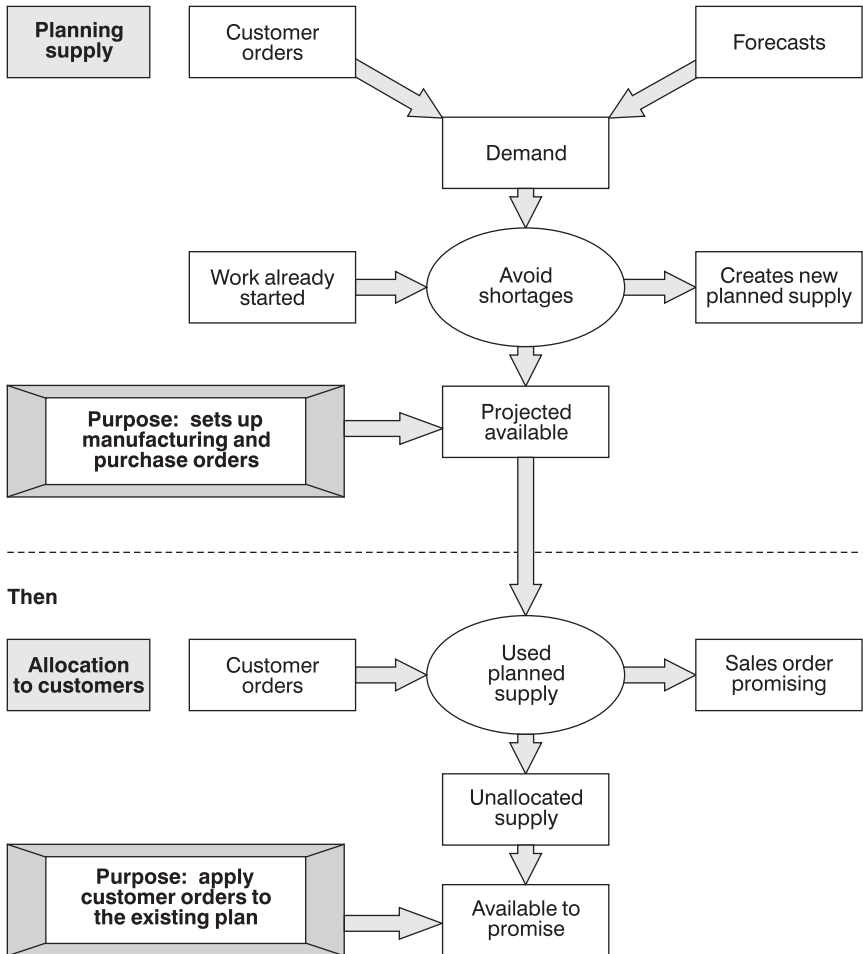


Figure 12.6 Control of supply and demand

Table 12.5 Available to promise

Delivery promising	Customer orders	Scheduled receipts	Stock on hand	Planned order receipts	Available to promise	Planned order release
Current stock			500			
1	2600	5000	2900		2900	
2	6600	5000	1300		1300	
3	1200		100		100	
4	3200	5000	1900		1900	5000
5	1200		700		700	5000
6	600		100	5000	5100	
7	4000		-3900		1100	
8	100		-4000	5000	6000	
9			-4000	5000	11 000	
10			-4000		11 000	

by looking at the actual orders instead of the forecast. The allocation of the projected available stock to real orders enables customer delivery times to be estimated using the MRP logic. This process is shown in Figure 12.6.

Creating an order book according to Table 12.4 gives the availability shown in Table 12.5. This shows the actual customer demand at the present moment but orders are constantly being received, so the situation may well change within a few minutes. The scheduled receipts and the planned order receipt are the values calculated from the projected available stock shown in Figure 12.6. The stock on hand and the available to promise columns are calculated from the real customer demand column.

Example

In week 1 there are 5000 received and 2600 issued to customers. The stock which was originally 500 is therefore 2900 by the end of that week, and these could be used to fulfil customer demand. However in week 3 the available to promise is down to 100, and the supply lead time is four weeks. If more than 100 out of the 2900 are issued in week 1, then there will be a shortage in week 3. The delivery promise for quantities up to 100 can therefore be in the current week; the delivery promise for up to 1900 can be week 4, and greater quantities can be promised either up to a maximum of the forecast or in excess of the cumulative forecast, as long as the forecast is changed to match the demand.

This change has to take account of the fact that the large demand is probably part of the forecast and therefore the forecast for other weeks will need to be reduced.

The two sectors – projected available and available to promise – ensure that customer demand is managed in line with supply. It avoids difficulties which result from accepting orders on a fixed or short lead time which then cause emergency sourcing or failing to provide delivery on time to other customers.

Batch sizes

The master schedule is an excellent tool for managing and monitoring stocks, especially where there is variable demand which can be estimated. It enables lower stock to be held than can be achieved by statistical stock-centred methods. It also links in very well with supplier scheduling. Getting the best out of master planning does require more effort because of the regular reviews and reforecasting. However, in a pure stock environment, it is only necessary to use MRP on the A class items and items where the demand is understandable and variable. The triggering of MRP by this schedule enables companies to forecast for whole ranges of stock with a single evaluation and give a balance of stock for customers who require co-ordinated supply.

(In classic MRP, a number of batch rules have been developed as many of these are based on ‘economic batch sizes’, which produce batch quantities which are unbalanced and too large for customer order requirements.)

The same logic can be used for MRP order quantities as for normal stock control. The three main alternatives are:

- lot for lot
- fixed batch size
- period cover.

Lot for lot is the first choice because it leads to a minimum of inventory.

Example

The rule is to bring in the balance required in each time period. In Table 12.4 the scheduled receipts for weeks 1, 2 and 3 would have been 2500, (3000 forecast demand less the 500 stock) 4000 and 2000, and the projected on hand would always be zero.

This gives a variable supply quantity which is inconvenient in some circumstances.

Fixed batch sizes are used widely to illustrate MRP. They are overused in practice since changes in demand level should be reflected in changes in order quantities. Fixed batch sizes are convenient to use but are only essential where there is a restriction in vessel size (distribution or manufacturing). Batch sizes based on price breaks are the result of too simplistic an approach to costing.

Choose batch sizes to minimize work and inventory.

Period cover is a way of taking the demand for the next few periods and batching it up. It is the same as lot for lot but with several time periods added together. Again the order quantity is variable, and the stock level is significantly higher. This technique reduces the number of orders and increases their size at the expense of higher inventory. The application is therefore to C class items where this is the most beneficial option.

Key points

- Look for structure in demand patterns.
- Working with dependent demand keeps inventory low.
- MRP is a simple way to trigger future supply orders – especially where order levels vary.
- Using 'available to promise' gives more accurate delivery dates.
- Batch sizes have to be defined using rules which minimize total costs and workload.

The future – inventory and logistics

- Viewing the full inventory and distribution picture
- Continuing development – the lean approach
- The basis – getting the records right
- Structuring inventory at each level
- The role of consignment stocks
- Structuring the distribution chain
- Location and depth of inventory
- Replenishment systems – fair shares & DRP
- Summary

The basis of the lean supply chain

Extending inventory management

The science of inventory management has gradually developed, not because the theories have leapt forward, but because requirements change. No longer is inventory management a single problem but it is integrated into the whole of logistics of supply. The manager is now co-ordinating

- warehouse operations
- inventory control
- distribution
- suppliers
- customer operations

and attempting to optimize the profit for the supply chain rather than just keep delivery on time up and inventory down.

However, the targets are still the same but in a wider context:

- better customer service
- lower inventory
- decreased operating costs,

but now the focus is on the whole supply being lean, not on the individual echelon.

The development in information technology (IT) communications enables information to be gathered and exchanged immediately and accurately, and makes the information available to everyone. However, the challenge is now:

- collecting data accurately
- being willing to share it along the supply chain
- understanding how to use it most beneficially
- working for the good of the supply chain and not for political gain.

The techniques discussed in previous chapters are the basic methods to be used in each step of this chain. They then have to be put together to make the chain effective and competitive.

Lean supply

The concept of lean supply concentrates on the very essence of good inventory management:

- customer focus – exceeding the customers’ expectations
- aim for perfection – not just better than the competition
- added value – only carry out activities which directly benefit the customer
- flow – so that inventory does not stop, it continues from one stage to the next
- unit movement – small batches (e.g. kanbans) moved frequently
- simplicity – in operation so that things are obvious and do not go wrong
- individual responsibility – at operator level for developing more effective operation
- integrated design approach – with additional features for the customer and ease of manufacture at low cost
- quality – always create the correct quality first time.

Use lean supply ideas to see how to develop inventory control and logistics.

Although there are many principles here, they are all interlinked and working on one is supported by working on another. Improvement is a circular process of tightening up one area then moving on to the next and eventually back again to the original as it has become the weakest area

again. When companies change to a lean attitude there are a lot of minor changes which add up to major developments. These improvements are made continuously, so the operation gets better and better.

Lean supply involves the whole business from purchasing through to dispatch. The achievement of good flow of supply requires this full participation. Purchasing has to ensure that goods are delivered on time, they are fit for purpose and labelled and packaged correctly and consistently. A design department has to ensure that a small number of products satisfy a wide range of customers and are composed of as small a number of components as possible. They should also ensure that components and modules satisfy the widest range of products. What is more the lean approach causes a continuous reassessment of the design, taking into account the change in customers' values and the needs to rationalize production processes and components.

The quality department does not inspect; it sets up the processes which either avoid mistakes occurring or enable the work to be checked as it is carried out. Purchasing is not concerned with lowest unit price of purchases, but is involved with minimizing the total cost to the customer and maximizing the value to the customer. Manufacturing need to focus on supporting customer requirement and not on throughput.

The break with traditional concepts is that improvement is part of normal work for everyone. The gradual development of the operations is led by the people carrying out the processes – the management play a supportive and guidance role rather than being the promoters of change.

The foundations of lean supply lie in the development of JIT processes described earlier in this book. Lean supply takes this concept wider so that greater benefits are gradually accrued.

Information accuracy

More accurate records will reduce inventory holding.

All the work on inventory management will come to nothing if the base information is not correct. There are many companies where the quality of the information on the system is very poor. If the same quality standards were applied to products, then these organizations would quickly be out of business. Part of the development of better inventory control is therefore to ensure that the information on the formal system is

- real time
- accurate.

Computer systems are accurate, so it is a matter of making failsafe the information systems and making people believe that it can be correct. Inaccurate records have a high cost arising from continuously checking to see if items are there and holding extra inventory to avoid stockouts, not to mention the obvious cost when customers cannot have the items promised to them. It is another challenge to get stock records right. And this goes hand in hand with the development of inventory control.

Control methods

The achievement of tighter targets has largely been the result of better communication and simple controls. Customer pressure has improved delivery performance, and financial pressure has decreased operating costs.

The key tools have been the exception reports of ‘these stocks are too low’ and ‘these stocks are too high’, which keep the stock broadly correct so long as orders are placed at approximately the right time. The use of ABC analysis, dynamic review levels, forecasting techniques and JIT supply enables inventories to be improved significantly. Effective implementation of these techniques has been slow because conditions did not encourage it. In the past, supply lead time was unreliable so that review level calculations could not be relied upon. Now that suppliers are expected to guarantee reliable delivery the calculations give the correct results and there is no need to add a bit extra (lead time, inventory, etc). The techniques described in previous chapters have proved invaluable in meeting the demands of the new market environment. Once a company is operating these successfully it is ready to contribute to the logistic effectiveness of the supply chain.

There are other techniques, built into some advanced software, which are waiting to be used more widely, and more advanced inventory control techniques will show significant advantages for the user. (It has been known for decades that Newton’s Law of gravity is not quite right, and Einstein’s theory of relativity is nearer the truth, but Newton’s Law is good enough for those of us who are not space travellers or who move below the speed of light. Similarly in inventory control, the review level calculations are not exactly correct and delivery frequency does affect the safety stock.)

Aim to get it nearly correct now rather than perfect sometime in the future.

Better formulae are readily available, but from observation, stocks are so far out from the ideal that a small amendment to the safety stock level is not

going to alter the real inventory. Similarly Gaussian statistics, simple availability calculations and independent demands are all concepts which are questionable theoretically, but which work in practice.

More advanced forecasting methods provide better answers when the future can be relied upon to mirror history (see Chapter 11). Crude tools like ABC analysis could be replaced by tools which take into account the potential different stock policies for high-value slow movers and low-value fast movers. Or ABC could be developed into a differential value analysis, which gives an individual class for each item code and optimizes every inventory line. This evaluation will take place when companies have exhausted the current methods and realize that they can benefit.

Consignment stock (consignment inventory)

A popular technique in the supply chain is the creation of consignment inventory. This can be defined as 'Inventory owned by a supplier at the customer's site'. (There are a variety of similar arrangements to this, but the essential feature is that the inventory is near the customer and owned by the supplier.) On the face of it there seems to be great benefit. Sales people like it because:

- there is effectively a captured market
- the customer should not have shortages
- supply can be slow and regular to stock instead of short lead time to customer order.

Customers like consignment stock because they can have:

- immediate availability
- supply without problems
- ability to change requirements at will
- less administration when ordering.

However, consignment stocks are a bad thing. By reviewing the situation from a distance, it can be seen that a consignment stock will:

- require the supplier to have an extra stores stock
- have inventory in the stores which is relatively slow moving
- take the pressure off efficient supply
- transfer the responsibility for the inventory from customer to supplier.

Looking at these points in turn, the first arises because the consignment is stock in a remote location and, therefore, not available to other customers. There may be duplication of stock between the supplier's own stores feeding the consignment and the consignment inventory itself.

The second point is that a smaller stores has throughput because the MAD for small throughput is relatively large (decreases by the square root of the volume). Therefore the consignment will have a surfeit of stock, even if it is controlled very well. (The chances are that the inventory level is either set very high by the supplier or agreed with the customer at an inflated level.)

If the supply chain with the consignment is in competition with one without, then the one without will have lower cost, and therefore a competitive advantage. The consignment stock is a palliative to poor planning: it may however be a short term (few months) solution to cover a specific problem.

There are a few simple rules that should be used when applying consignments:

- Do not set stock levels; set availability targets instead.
- Review variety of inventory and levels regularly (every few weeks).
- Use consignments only for a few fast-moving item types.
- Arrange who is charged with inventory left over through change of design, demand pattern and any other non-moving inventory.
- Allow consignment to be available for other users if necessary.
- Have a formal agreement for a limited period of time.

Consignment is only a temporary fix.

Logistics

Objectives of logistics

Logistics management presents new challenges to the inventory manager, who has to maintain and balance distributed stocks, additional inventories, manufacturers' work in progress, supply partnerships, and other new dimensions, including distribution costs. The need for this integration results from the changes in market structure given in Figure 13.1.

The development of supply chain strategy means that stock controllers from suppliers and customers should be sharing information and inventory in order to optimize the operation of the supply chain.

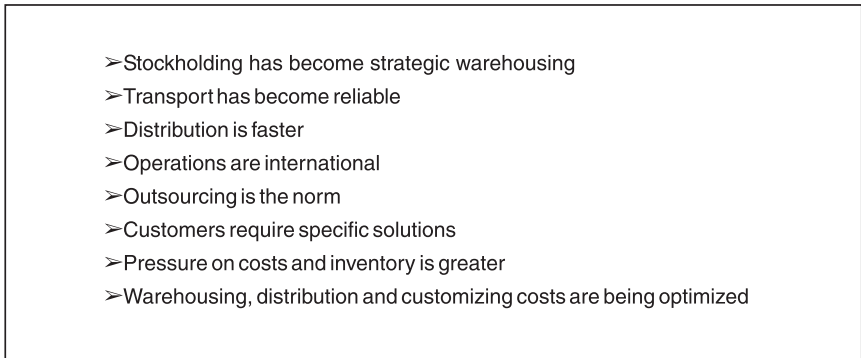


Figure 13.1 Changes in the distribution environment

Management of logistics is a simple extension of inventory management, characterized by the same basic underlying themes. The most successful supply chains are:

- integrated
- consistent
- fast moving (short lead time)
- service oriented
- customer specific.

Optimize the complete inventory structure, not each stage.

Supply chain structure

The management of a supply chain requires inventory to be held in different locations, even in different operating companies. The inventory controller can then optimize the stock across many locations. The structure of the supply chain can be altered to meet the needs of the market, and the overall balance between availability, stock investment and operating cost can be optimized by selecting the number of stages (echelons) in the supply chain and the number of stock sites in each echelon.

This can be established most simply by modelling the cost balance: generally it is more economical to hold inventory as near to the start of the supply chain as can be allowed by the demand lead time. Once the distribution structure has been decided, the proportion of stock held at each inventory level has to be calculated, as shown in Table 13.1.

Table 13.1 Supply chain inventory options

All stocks in a single central warehouse	Stock all held by end user or local stores
Minimum investment in stock	Items are available more rapidly
Accurate forecasting	Sensitive to local conditions
Higher demand rate	Duplication of stockholding
Reduced safety stocks	Poorer control
Low warehousing costs	Higher inventory investment
Better formal physical control	Greater risk of obsolescence.

The compensating factor in deciding where to hold stock is the distribution cost. In general, it is considered that:

- fast delivery costs more
- the greater the distance, the greater the cost
- the smaller the load, the higher, relatively, is the distribution cost.

However, there are other considerations to make, including the following:

- Regular shipments cost less than one-offs.
- Distribution costs are largely loading and unloading, mainly independent of distance.
- Consolidation saves cost and is easier for small, regular deliveries.
- Small loads need smaller transport and facilities which are cheaper.

The complexity arising from these latter considerations makes transport costing difficult. When considering JIT supply (Chapter 4) one of the principles discussed was continuous improvement (kaizen). This approach has to be used when reducing distribution costs and, therefore, inventory. However, to have a significant effect on the cost balance, bold leaps have to be made in distribution structure. The size and frequency of shipment have to be altered dramatically. This usually means a change in means of transport to provide smaller units faster. The savings in inventory and warehousing costs should be greater than any marginal increase in distribution costs.

The correct balance of stockholding and the location of the stock are different for each stock line, depending on their usage rates and values. Low-value, high-usage items are better nearer the point of use. High-value items with occasional usage can be held centrally.

The whole inventory system can be considered as a pipeline for supplying the fluid 'stock' to the point of sale (or use). The input to the pipeline is a continuous (scheduled) flow. For the high-usage items the pipeline can continue to the point of use. For low-usage items the movement includes a reservoir, a central stores from which the items are dispensed as required.

This analogy highlights some interesting aspects of the supply chain, and poses the questions:

- Should the system required be a stock control system or should it be a supply flow system?
- Which parts of the supply chain are best served by scheduled supply and which by batch ordering?
- If stocks are held at one echelon, do they need to be held at others as well?

For fast-moving goods, the central warehouse can act as a reservoir replenishing the local storage tanks. It is better if the central store is used purely as a transshipment point for these items. The pipeline is then connected to the distributed stocks. The central function becomes stock management co-ordinating ordering and supply, and ensuring consistent delivery.

Slow-moving items can all be held at the central stores, as long as there is the means to withdraw them and transfer them to the final customer within an acceptable time. This reduces the overall stockholding of these items. The resultant stock is structured with each stock-keeping unit (SKU) having a specified stock level based on its usage and value.

Multi-location supply chains are used commonly in some industries. Using an integrated stock information system, duplication and proliferation of stocks can be avoided. Distribution warehouses in different localities should be treated as a single level where all stock is considered as holdings of a single store apart from the safety stocks. The aim is to keep the appropriate amount at each location and if supply from the originator or manufacturer is rapid, then a central stock location can be avoided, with a corresponding reduction in stock. Careful supervision of delivery time is necessary to avoid duplication of stockholding at several locations.

The degree to which a company can take advantage of these techniques is limited by their organization and systems. Often it is necessary to make improvements as a two-stage process:

- 1 Introduce the 'stock to stock' system to get the control.
- 2 When control is good, reduce the stock at each stage in the system by improving the inventory control techniques.

Organizational control

For control of inventory in the distribution chain, the costs have to be minimized and service maximized. Distribution systems are neither ‘pull’ nor ‘push’ methods of stock control. Such methods imply that either distribution (pull) or manufacturing (push) are entirely responsible for the organization of distributed inventory. Modern techniques centre around an

- accessible database of information
- realistic schedules
- replenishment of distribution sites based on requirements, lead times and item availability.

Start from minimizing the total cost, then work out the nearest practical solution.

If the supply chain is to be operated most efficiently, then there has to be some co-ordination and organization for the whole chain. This will generally require:

- good information fed back from all the distributed stores (it helps to have integrated systems)
- central decision-making of the allocation of stock.

This is a statistical reality. Assessment of demand is more accurate for the total demand than for individual warehouses, since the larger volume has a smaller MAD. (The MAD only increases as the square root of the volume.) This means that historical forecasting is most accurate when carried out by a central planner. However, if a demand pattern changes (new large customer, loss of contract, etc.) then the local knowledge of these events needs to be included in the forecast (see Chapter 9).

The techniques commonly used for managing the distribution channel are basically:

- fair shares (for allocating available stock)
- distribution requirements planning (DRP) (for planning supply).

Both the fair shares and DRP techniques are concerned with the quantity of stock to be shipped through the distribution network.

Distributed inventory

For an inventory planner controlling stock in several regions, the issues to resolve are:

- 1 Where the stock is to be held.
- 2 What level of stock to hold in each region.
- 3 How to maintain supply at the lowest cost.

The first question is answered by working out the general distribution and inventory cost balance from the factors just discussed in the first section in this chapter. The second issue is easily resolved using the classic inventory management approach described in Chapter 7. The third issue is to ensure that the transfers are carried out most economically within the framework arranged by (1) and (2). This is an optimization process. The cost of providing items at a remote warehouse depends largely on two factors:

- the item cost
- the distribution cost.

There is a unique cost for each item at each site, and SKUs provide the information required to evaluate distribution needs. The stock in some distributed warehouses is likely to be larger than necessary, particularly if they are the delivery points for outside supply, while others require supplying with stock to meet demand. Stock therefore has to be redistributed at the least cost from one location to another at the same echelon of the distribution network. If the costs of distribution between each of the stores is known, then the total costs can be optimized. This is particularly important where transport is expensive, or where there are significant differences in transport costs, as found in a multinational supply operation.

Fair shares

Fair shares analysis approaches the regional stock issue as an inventory control problem, assuming that distribution costs are not the prime concern.

The purpose of fair shares analysis is to distribute stock through the supply chain to alternative distributed storage. Stock is not held at a central warehouse but distributed to regional stores nearer to the customer demand. An implicit assumption is that the supply to the network can be

accomplished by short lead time orders. If the lead time is longer, the planning aspects of DRP have to be used.

When stock is delivered into a distributed stores network, it has to be reallocated to the various stores in the most equitable manner. The sensible approach is to cover priorities first and then share out the rest to meet future demand. Fair shares analysis is a method for achieving this using a standard procedure. The logical process is to identify the safety stocks in the supplying warehouse and in the other warehouses, and then to allocate the total available stock to meet the requirements in proportion to the demand rate.

The fair shares allocation of stock normally distributes stock in the following sequence:

- 1 Customer back orders unsatisfied.
- 2 Expected demand within supply lead time.
- 3 Safety stock at distributed stores.
- 4 Safety stock if required at supplying warehouse.
- 5 Allocation of remaining stock in proportion to demand rate.
- 6 Redistribution where a stores stock exceeds its share.

If stock is insufficient to meet the full requirement at any stage, it is allocated logic. The actual stock transfers used to achieve this distribution can be optimized through the transport algorithm. A simple example of the use of the fair shares logic is shown in Tables 13.2, 13.3 and 13.4.

A manufacturer sells items in batches of 400 to a central supply warehouse. This warehouse holds some safety stock, but distributes the items out to three regional stores. Estimates for the usage at the stores are given in Table 13.2 together with the safety stocks. If a delivery of 400 arrives from the manufacturer just at the start of week 1, the total stock available is 400 plus 200 central stock plus 280 at the regional stores (= 880). Table 13.3 shows the lead times, and develops the supply chain

Table 13.2 Stock and forecast demand

	Safety stock	Current stock	Week 1	Week 2	Week 3	Week 4	Week 5	Week 6
North stores	20	70	50	40	70	20	50	50
South stores	46	200	100	140	60	110	150	150
West stores	12	10	20	30	0	50	20	20
Total	78	280	170	130	130	180	220	220
Supply warehouse	37	200						

Table 13.3 Fair shares priority allocation

Location	Average demand	Lead time (wks)	Demand in lead time
North stores	46	1.2	56
South stores	118	1.0	119
West stores	23	0.4	10
Total	187	4.0	185

	Stock	Balance
Current stock	280	
Supply warehouse stock	200	
Current delivery	400	
Total	880	880
Shortages	0	
Demand in lead time	185	695
Stores safety stock	78	617
Warehouse safety stock	37	580

Table 13.4 Fair shares apportionment of stock¹

Stores	3.1 weeks' demand	Demand in lead time	Safety stock	Current stock	Delivery required	Total stock
North stores	143	56	20	70	149	219
South stores	366	119	46	200	331	531
West stores	71	10	12	10	83	93
Total	580	185	78	280	563	843
Supply warehouse			37	600	-563	

Note: ¹Number of weeks' cover = $580/187 = 3.10$ weeks.

calculation. There are no back orders, so the first allocation of stock is to cover regional demand in the lead time (185), then regional stores safety stocks (seventy-eight), then supply warehouse safety stocks (thirty-seven). Deducting these from the total available leaves 580 items to be split up across the regional stores. This can be done either by netting off the individual forecasts until the stock is exhausted (where these are accurate and contain real orders). Alternatively the average week's cover can be calculated. The total demand is 187 per week, so 580 will last 3.10 weeks on average. Multiplying the demand by 3.10 gives the fair share for each store. This is shown in Table 13.4.

Each store is then given a fair share consisting of:

$$\begin{aligned} & \text{Any customer back orders} + \text{Demand in the lead time} \\ & + \text{Safety stock} + 3.10 \text{ weeks' demand} - \text{Current stock.} \end{aligned}$$

This is also shown in Table 13.4. All the stock and delivery to the supply warehouse has been allocated, apart from the safety stock (thirty-seven). If one regional store already had had sufficient, then either the excess can be transferred, or that regional store (and its stock) can be excluded from the analysis and a fair share worked out among the others.

Distribution requirements planning

Distribution requirements planning provides planners with a complete forward schedule for all dispatches, arranged by date order by warehouse. Each store at the demand end of the supply chain generates a forecast for demand which is used to identify the orders they are likely to place on their supply warehouse (supply). The forecast is for demand well in advance of the normal replenishment lead times. Since the supply is planned to meet the requirement well ahead, the demand rate may have changed so the resultant replenishment delivery may be too early or late by the time demand is met.

The delivery requirements calculated by the netting logic of DRP show the latest shipping dates associated with each warehouse requirement (see Figure 13.1). Distribution requirements planning is an extension of the MRP process for distributed warehouses. The netting logic used for DRP is exactly the same as for MRP (see Chapter 12, section on MRP logic):

- The future requirement is split into time periods of a week or a day.
- The expected demand each time period is identified.
- Starting with the current period, the demand is netted off against the available stock.
- If the remaining stock is above the safety stock then this quantity is carried forward to the next time period and the calculation is repeated.
- If the available stock is less than the safety stock, then a receipt is required. The quantity requested depends on the batch rules being used.
- The request generates a supply order a lead time earlier. (The lead time corresponds to the picking, packing, transfer and receiving time.)

By using DRP, supply planners can receive a complete forward schedule for all dispatches, arranged by date and by warehouse. Transport must be made

to fit the schedule. Distribution requirements planning builds on that master scheduling, and looks at demand far past the normal lead time. The purpose of DRP is to produce a time-phased shipping requirement for the suppliers.

Distribution requirements planning is a method for planning the replenishment for the distribution network. The logic closely follows that of MRP. Like MRP the supply chain is planned by backward scheduling from the customer. This requires the establishment of fixed supply lead times and batch quantities. If this is organized by the manufacturer, the batch size can be large and the intervals of distribution correspondingly long. This, of course, adds stock throughout the supply chain and so the use of a just in time approach with very short lead times is a great advantage.

An example is shown in Table 13.5. In this table the items are distributed from a supply warehouse to three regional stores. The stores have given their estimates of demand (shown as 'forecast') and from these the stock and demand pattern needed to keep the 'projected on hand' positive are given in the 'Stock' and 'Supply' rows. The item is distributed in batches of 100 in this case. (This is the same calculation as for 'projected on hand' and 'planned receipts' for MRP in Chapter 12.)

There is a lead time of one week for delivery to the north and south stores from the supply warehouse. This is caused by the order processing, picking and transport: delivery is very fast to the west stores. The dispatches from the supply warehouse required to meet this are given in the lower half of the table, taking into account the week offset in the two cases. If the supply

Table 13.5 Example of distribution requirements planning

		Current	Wk1	Wk2	Wk3	Wk4	Wk5	Wk6
North stores	Forecast		50	40	70	20	50	50
	Stock	70	20	80	10	90	40	90
	Supply			100		100		100
South stores	Forecast		100	140	60	110	150	150
	Stock	200	100	60	0	90	40	90
	Supply			100		200	100	200
West stores	Forecast		20	30	0	50	20	20
	Stock	10	90	60	60	10	90	70
	Supply		100				100	
Supply warehouse								
Demand upon supply warehouse	North stores		100	0	100	0	100	?
	South stores		100	0	200	100	200	?
	West stores		100	0	0	0	100	0
	Total		300	0	300	100	400	0
	Stock balance	200	300	300	0	300	300	300
	Deliveries		400			400	400	

warehouse receives the goods in batches of 400, the required delivery schedule can be worked out as in the ‘Deliveries’ line. The resultant stock at the central warehouse is also shown. Once demand at this supply echelon has been determined, the demand on manufacturing can be placed into the master schedule. Again a supply lead time has to be allowed for the distribution process.

Distribution requirements planning is far simpler than fair shares and is easier to manage. An important practical difference between DRP and fair shares is that DRP relies on lot-sized quantities with pre-scheduled dispatch times. In DRP there is normally a resultant stock caused by the batch sizes, as illustrated in Table 13.5.

Some pitfalls which have to be avoided when using DRP rather than fair shares are:

- DRP is less flexible.
- It focuses on each site rather than the total system.
- Its sole concern is the balance of forecast and available stock, with little reference to customer service.
- It fixes and depends on the supply performance.

The technique is relatively rigid, since the quantities must be sent on the days scheduled. In this sense it is limited in value for transportation planners (transport must be made to fit the schedule). With DRP, quantities shown are planned, with the implication that stock is not usually currently available. Consequently, changes to shipment schedules have to be long term rather than immediate. There is little possibility of responding to shortages since the quantities shown are planned to be available, not currently available. Distribution requirements planning is best suited to situations of relatively steady demand: it is most successful where fixed delivery quantities are physically necessary (such as container loads).

Fair shares can be made to work in conjunction with DRP by making the best allocation of the shipments arriving into the network.

Wherever possible, the manufacturing and distribution systems should be synchronized, thus reducing the amount of inventory carried throughout the system. This could be achieved in the example shown in Table 13.5.

The interface between the computed lot-size requirements and the master schedule is the gross requirement from the supply warehouse, taking into account any finished goods stock held by the manufacturer. In practice, the DRP gross requirement can be exceeded by the master schedule and products can be planned in advance of the DRP schedule dates if it is beneficial for manufacturing planning and efficiency.

The distribution quantities calculated by the forecasting techniques and apportioned by fair shares often have to be adjusted because of other influences. The total shipment may have to be phased because of space limitation in a particular stores. The distribution quantities could well be rounded to the nearest pack size, or altered to take advantage of a distributor's cost structure.

Traditionally forecast demand would be made at the product level. Distribution requirements planning allows the forecast to be conducted at the SKU level. A consequence of this is the need to aggregate the local forecasts to generate forecast demand for each product and, ultimately, for product families.

It will have been observed from the example that DRP results in considerable unevenness of the gross requirements. This is caused by the fixed lot sizes. If the lot size were smaller, this effect would be minimized, but the number of deliveries would be increased.

Review

The development of inventory management is a dynamic activity, with new approaches being made, techniques being refined and new challenges being met. Successful inventory management is a balance between the use of basic techniques which are easy to use and give reasonable results, and sophisticated techniques which can give better control if they are used properly. Modern inventory managers have the opportunity to use better and more effective controls.

It is imperative that the challenges are met, since continuing competitiveness is based on ever-improving customer service and ever-reducing inventory costs. The approaches discussed in this book give companies major improvements in their efficiency across all types of business. Understanding the full implications and applications is a gradual process, but worth the effort because of the tremendous benefits.

The science of inventory management as developed in this book has to be carried out by an organized, aware and dedicated artist. The manager has to focus on the simple basic requirements of *creating the most profit for the organization through the best customer service*. This means:

- excellent forecasting
- low inventory value
- small operating costs.

The direction is towards perfect service, no costs and no inventory. To manage this it is essential to have operating controls and targets – monitoring

of current performance and tighter targets (which are achievable and agreed with those required to reach them).

There is also the balance between long-term and short-term requirements. Paying extra delivery costs to help the customer may be good for relations. As a one-off it is fine, but as a frequent activity it can damage the viability of the business. The inventory controller has always to look simultaneously at the long-term and short-term detailed picture and overall plan. By steering a consistent course, based on a set of operating principles, work is reduced and effectiveness increased.

As better control is introduced and the performance improves in one area, so it can be integrated with other activities and new challenges adopted.

Take a simple example of a manufacturer that used to hold raw materials inventory and then supply the finished items into a warehouse from where it was distributed to regional stores. Each stage of the process was struggling to gain good control in its own separate way. (Some people are still at base 1 and wrestling with these challenges.) There is a natural process of evolution from this setup:

- 1 Gain control over manufacturing (using MRP for example).
- 2 Add the finished warehouse stock to the manufacturing control process (form materials management activity).
- 3 Include distribution costs in the cost balance.
- 4 Form partnerships with suppliers to avoid duplicating the raw material stockholding.
- 5 Introduce flow through the business and work on the supply pipeline to eliminate work in progress and stores inventory.
- 6 Integrate the distributed regional stores into the inventory management system.
- 7 Optimize inventory along the supply chain.
- 8 Appoint supply chain management to optimize supply chain costs.
- 9 Extend the integrated management throughout all stages of the supply pipeline.

If the company is using the same methods as five years ago, it has a problem.

This may be a large step for companies where the supply chain members are not all fixed and where there are many companies involved. However, it is necessary if businesses are to remain competitive. Other companies are developing and improving, so maintaining the present situation is not a viable option.

The key to success is simplicity of operation. The concepts can be complicated, just as reality is, and the inventory management system has to reflect this. It does not mean that the actual day-to-day operation should be difficult. In fact, the more sophisticated the system, the more reliable it should be and the less effort to operate. However, a good system requires setting up to operate properly and this requires good inventory management expertise, using the professional techniques described in this book. Using the forecasting described here, along with ABC analysis and stock level calculations, the author has saved millions of pounds worth of inventory in many companies. Most businesses want to generate more profits: many have not recognized the key role of inventory management in increasing turnover and reducing costs.

Inventory management is a collection of techniques developed by practical people and proven in the real world. It is no use using general good management practice in this area where there are specific inventory control methods required. These have taken many years to develop and refine.

There are refinements beyond those described in this book, but in practice the methods discussed give such good benefits that other areas become the focus of improvements before these are required. In these pages there is an attempt to illustrate the application of the methods as well as the techniques themselves. The reader should be able to gain significant improvements in the stockholding simply from implementing the concepts described – and that is the next challenge. Understanding the full implications and benefits is a gradual process, as is the implementation of the techniques. Taking up the challenge is definitely worth it because of the tremendous impact on the business.

Key points

- When the inventory in a stores has been improved, consider the supply chain as a whole.
- Minimize the logistic costs whilst exceeding customer expectations.
- Decide upon the best places to situate the stock and the amount of stock to hold in each, using inventory management principles.
- Employ the techniques described in this book for the best results, then add expertise to enhance those items which are important, new or exceptional.
- Set up your systems to take the work out of inventory control.

Questions and answers

Questions

- 1 What are the three key objectives of inventory control?
- 2 Which departments put pressure on inventory managers, and what are their objectives?
- 3 'More stock gives better availability.' Challenge this conventional view.
- 4 Two companies distribute fast moving consumer goods. Their results for last year were:

Annual accounts data	Ace Distributors (£ million)	Deuce Associates (£ million)
Sales turnover	8	26
Fixed assets	2	5
Cost of sales	6.2	21
Stock value	4	8

You have been offered the job of inventory manager by both companies. Which one would you take and, based on the data provided, explain why.

- 5 The aim for Ace Distributors (in question 4) is to reduce the stock to £1.5 million. What will be the return on assets then?
- 6 What change in the return on capital (as a percentage) will there be when a company has reduced its stock value by £500 000, given the following data:

$$\text{Return on capital} = \frac{\text{Annual profit}}{\text{Total capital employed}}$$

where

capital employed = stock + capital plant

sales turnover = £3 million

profit = £400 000

capital plant = £1.1 million

stock before reduction = £1.5 million

7 Discuss the ways in which the JIT approach differs from conventional inventory management.

8 How does a company give good customer service? A company is distributing a range of fast-moving items to stockists. The following is a list of dispatches for one day. What is their customer service level, assuming this to be a normal day?

Order No 1A

	Quantity requested by customer	Quantity dispatched from stores
Item no. 1	20	15
Item no. 2	100	100
Item no. 3	50	20

Order No. 2B

	Quantity requested by customer	Quantity dispatched from stores
Item no. 1	60	40
Item no. 2	150	100
Item no. 3	20	20

Order No. 3C

	Quantity requested by customer	Quantity dispatched from stores
Item no. 1	10	10
Item no. 2	70	70
Item no. 3	50	20
Item no. 4	25	25

9 Which items in the following list are in the A, B and C classes? How do you determine this?

Item code	Weekly demand	Unit price (£)
A901	2	3.00
B662	2	8.00
C355	20	4.00
D523	40	0.50
E191	10	65.00
F807	1	9.00
G010	3	1.00
H244	1	170.00
J459	10	3.50
K488	22	0.50

10 If the current stock of the items in question 9 were as shown below:

Item	Stock quantity
A901	6
B662	10
C355	50
D523	60
E191	50
F807	50
G010	12
H244	5
J459	20
K488	5

- What is the stock cover for each of the items?
- What would be the first task resulting from this analysis?
- What is the stock cover value for the total stock?

11 A store is offering ex stock delivery to a wide variety of customers. Item 35721 is a non-seasonal item which has the following usage characteristics:

Weekly usage: 200.
 MAD: 50.
 Supply lead time: four weeks.

- If deliveries are brought in every week, what is the target stock level if 85 per cent availability ex stock is required?
- If the supplier is persuaded to deliver daily, what would be the new target stock level?
- How would the change affect the value of stockholding if the items cost £10 each?

12 The statistics on the stock of a range of items is as follows:

Item no.	Unit cost (£)	Annual usage
2A32	25.00	12
2B44	0.50	360
3D10	0.10	120
3E82	40.00	1
3F66	2.00	40
4G19	5.00	520
4H95	4.00	7
5J53	20.00	5
7N78	0.30	200
9P21	12.00	50

The stores controller is under pressure from the management to reduce the stock level to below half a month's usage. The stores personnel are adamant that they are not going to do any extra work. Show what the stock policy should be in these conditions and how the target is achieved. (Assume that safety stock is not necessary in this case.)

13 Stock levels are as follows:

- A items: £28 000, 130 items
- B items: £6500, 350 items
- C items: £3500, 600 items
- Obsolete items: £2000, 200 items.

Within the budget for next year, the estimated usage rate for each is:

- A items: £133 000 per month
- B items: £16 000 per month
- C items: £2000 per month
- Obsolete items: nil.

What is the best way to reduce the stock by £2000?

14 Do you believe that JIT simply passes the burden of holding stock on to the supplier?

15 Your managing director has asked for a short report on ways of reducing work in progress and throughput time. At present, batch sizes are about fifteen, setups average three hours, processing time per unit is two minutes, and WIP stands at thirteen weeks. Create a short report outlining in simple language the measures you would take.

- 16 (a) What are the financial benefits to a company through setting up a linear flow kanban system?
- (b) A company is intending to set up a kanban system for production. Comment on the changes which will be required in:
- responsibility
 - layout
 - operator training
 - quality management.

17 What is the role of safety stock, and why is it often not required in a JIT environment?

18 'JIT is a good theoretical idea but cannot work in our company' said the production director. 'We have customers who change their minds, a variety of processes, many products, and real people who cannot always be relied upon.' Set out the arguments you would make to persuade this director that JIT is a good idea for the company.

19 The data below is the stock record for item 345.

Date	Receipts	Unit cost	Issues	Balance
14/3	10	11		10
12/4			3	7
22/5			2	5
10/6	5	15		10
30/6			3	7
17/7			2	5

- (a) What is the value of the last issue (17/7) of item 345, based on:
- FIFO costing
 - average costing
 - LIFO costing.

(b) If you needed to reduce the stock value in your stores by £500 000, what items would you look at first and what techniques would you use to reduce the stock value?

20 Why is safety stock held?

21 What are the three factors which determine safety stock?

22 The usage for two items over the past five weeks has been:

Item	Week 1	Week 2	Week 3	Week 4	Week 5
D523	60	30	20	40	50
K488	30	15	0	65	0

- (a) What are the values of average demand and variability for each?
- (b) The items cost 50p each and the lead time from the supplier is two weeks. The inventory controller is deciding whether to try to provide a service level of 90 per cent or 95 per cent. What is the additional cost of the better service?

23 ABC Distributors have the UK dealership for a well known manufacturer of pre-packaged confectionery. Their main activity is the supply of goods to the local retailers ex stock. At the recent stocktake, ABC Distributors had 1280 different stock lines with a total value of £40 000. They order all their items during the week they fall below the reorder level.

The system works well except for the fact that the stores manager and his assistant are overloaded with work – in fact they are currently three weeks behind with the ordering. As they are ordering one month’s requirement at a time, there are often times when the stock runs out. The suppliers take two months to deliver, and the shortfalls cannot therefore be rectified immediately.

What are the changes in operating practice which you would suggest to assist them? What controls should they use which will bring the changes into effect?

24 'Re-ordering systems are not designed to stop you running out of stock.' Give two reasons why this statement can be justified.

25 What factors do you consider when setting stock levels for an independent demand item?

26 The demand history of item XX1 is as follows:

Week	Demand
------	--------

The items are supplied weekly in boxes of 100. Work out:

- (a) the review level (reorder point)
- (b) the average stockholding.

(You can assume a safety factor of 2 for 94 per cent customer service.)

27 Explain why, for some A items, we can expect the review level to always be higher than the physical stock level?

28 A company is putting expensive fountain pens into boxes at a rate of about 250 per week, together with ink cartridges. The stock controller has decided that the pens they buy are class A, while the cartridges and boxes are class C. Suggest order and delivery patterns and order quantities for the three types of item, and discuss why you have chosen them.

29 (a) It has been remarked that 'stock level for an item does not depend on its lead time'. Defend this remark using examples, and discuss why the statement is not precisely true.

- (b) Calculate the reorder level for the supply of umbrellas given the following data:
 Lead time = 2 months.
 Customer service level expected is 90 per cent.
 Demand over the last six months has been 11, 20, 8, 15, 10 and 14.
- (c) Discuss the assumptions which you have made in using this calculation.

30 A company is proposing to schedule in supplies of twenty key items each week instead of ordering them monthly.

- (a) What preparatory work do they need to do with the suppliers?
- (b) If the purchase value of these items is £1 million per year, by how much will the average inventory be reduced?
- (c) What are the benefits and extra jobs caused by scheduling?

31 What factors determine the delivery quantity?

32 Discuss ways in which stock can be reduced by suppliers and customers managing a supply chain together. What are the prerequisites for successful supply partnerships?

33 What is the difference between a target stock level system and a review level system?

34 Weekly sales of a product are:

Week	Demand
1	12
2	14
3	8
4	16
5	13
6	18

- (a) Discuss what forecasting techniques can be used to determine the demand for the next week. Indicate which you would prefer to use and the reasons for your choice.

- (b) If the initial forecast for week 1 was ten, work out the forecast demand for each week using exponential smoothing, and give your prediction for week 7. (You can use an α factor of 0.2 for this calculation).

35 What items of data (inputs) are required for the MRP calculation? For each, state the level and type of accuracy necessary to ensure reasonable results.

36 A watch manufacturer has a schedule for assembly covering ten weeks. The schedule is:

Watch X	Manufacturing orders for 1000 watches due in weeks 2, 4, 6, 8 and 10. Assembly lead time per batch is 1 week.
Watch Y	Manufacturing orders for 500 watches due in weeks 3, 6 and 9. Assembly lead time per batch is 2 weeks due to a special processing on the watch cover.
Watch Z	Manufacturing orders for 100 watches due every week. Assembly lead time is half a week. The parts for the batch of week 1 have already been issued.

Watches X and Y use the same module A and batteries (two per watch). Watch Z uses a different module B and the same batteries but only one per watch. Batteries and modules come in boxes of 500. There are seven boxes of batteries and three boxes of modules in stock. The buyer has outstanding purchasing orders for a total of 20 000 batteries and 6000 of each module, half quantities to arrive in weeks 4 and 10.

- (a) What would you expect your MRP package to recommend for purchasing? (Assume that the package is not capable of recommending rescheduling.)
- (b) What would you recommend should be done?

37 A product is manufactured in two stages. The material for the process consists mainly of one basic, bought out ingredient. Expected sales of the product are as follows:

Lot size for manufacturing store is 1000.

Current stock for manufacturing store is 1650.

Replenished lead time for all warehouses is two periods, internal sales are replenished immediately, manufacturing lead time is three weeks.

40 What are the advantages and disadvantages of keeping stock in retail stores rather than in supply depots?

41 The following ten items are the range of stock held in a small stores.

Item no.	Unit cost (£)	Annual usage
1	150	2
2	125	100
3	85	
4	50	
5	40	
6	26	
7	20	
8	10	
9	7	300
10	6	20

- Identify which are the A, B and C class items.
- Discuss how ABC is used in purchase ordering.
- Review another use of this Pareto classification.

42 The following data is available for the demand for an item:

Week	1	2	3	4	5
Demand	10	20	5	10	15

What would you estimate to be the average quantity held in stock to meet the demand for this item, treating it as:

- (a) Independent demand (assuming that a 99.2 per cent customer service level is required).
- (b) Dependent MRP demand (level 2) (assuming a one week lead time).
- (c) JIT supply (assuming a one-day delivery interval).

Specify any assumptions you make.

Answers

1

Availability (delivery on time).

Minimum inventory cost.

Minimum operating cost.

See text, Chapter 1.

2

Purchasing, if they try to buy large batches to reduce unit cost.

Sales if they want large stocks to service possible demand.

Accounts who require lowest investment in inventory to improve cash flow.

In manufacturing, production traditionally want to make large batches and to have buffer stock in case of problems.

Also pressure comes from elsewhere – designers who specify items too late and management who keep changing the targets.

See text, Chapter 1.

3 High stocks result from poor control, and poor service results from poor management of supply and forecasts. The better the inventory management, the better the availability and the less stock cover is required.

4

Annual accounts data	Ace Distributors (£ million)	Deuce Associates (£ million)
Annual profit ¹	1.8	5
Return on sales	<u>1.8</u>	<u>5</u>
=	22.5%	19.2%
Assets are	2 + 4	2 + 4
Return on assets	<u>1.8</u>	<u>5</u>
=	30.0%	38.5%

Note: ¹Sales turnover – cost of sales.

Ace has better margin and more potential for improvement in stock control.

5 51.4 per cent.

Annual accounts data	Ace Distributors (£ million)
Annual profit	1.8
Return on sales	$1.8/8 = 22.5\%$
New value of assets are	$2+1.5 = 3.5$
Return on assets	$1.8/3.5 = 51.4\%$

6

Capital employed before reduction = Stock + Capital plant
= £1.5 million + £1.1 million
= £ 2.6 million

Capital employed after reduction = £ 2.6 million – £500 000
= £ 2.1 million

Return on capital before reduction = £400 000/£2 600 000
= 15 per cent

Return on capital after reduction = £400 000/£2 100 000
= 19 per cent

7 There are two key ways in which JIT is a different approach:

- The level of inventory is not a simple calculation – it is a negotiated quantity.
- JIT does not set a permanent inventory situation – it empowers the people to improve and reduce the levels required.

This means that those involved with managing the supply chain always have to see ways of improving reliability of performance, reducing timescales and costs and, therefore, removing the need for inventory.

The timescales of JIT operations are measured in hours, under the umbrella of a long-term relationship and smooth supply. Conventional inventory management works with days or weeks of lead time and inventory levels and any demand pattern.

8 Solving customer requirements by providing quality goods service, information and ongoing support, and by delivery on time (see text).

Demand satisfied on time:	Dispatched	Requested	%
By order	0	3	0
By order line	5	10	50
By item code:			
Item 1	65	90	72
Item 2	270	320	84
Item 3	60	120	50
Item 4	25	25	50
Total	420	555	
Average line availability			77
Average item availability			76

Different answers are used in different situations (see text).

9

Item code	Weekly demand	Unit price (£)	Turnover (£)	Rank	Class
A901	2	3.00	6	9	C
B662	2	8.00	16	6	C
C355	20	4.00	80	3	B
D523	40	0.50	20	5	C
E191	10	65.00	650	1	A
F807	1	9.00	9	8	C
G010	3	1.00	3	10	C
H244	1	170.00	170	2	B
J459	10	3.50	35	4	C
K488	22	0.50	11	7	C

See also text, Chapter 3.

10 (a)

Item	Stock	Usage	Cover
A901	6	2	3.0
B662	10	2	5.0
C355	50	20	2.5
D523	60	40	1.5
E191	50	10	5.0
F807	50	1	50.0
G010	12	3	4.0
H244	5	1	5.0
J459	20	10	2.0
K488	5	22	0.23

(b) Ensure that supplies of item K488 are arriving within one day.

(c) This is calculated on a financial basis.

	Stock	Usage	Unit value (£)	Stock value (£)	Turnover (£)
A901	6	2	3.00	18	6
B662	10	2	8.00	80	16
C355	50	20	4.00	200	80
D523	60	40	0.50	30	20
E191	50	10	65.00	3250	650
F807	50	1	9.00	450	9
G010	12	3	1.00	12	3
H244	5	1	170.00	850	170
J459	20	10	3.50	70	35
K488	5	22	0.50	3	11
				4963	1000

$$\begin{aligned} \text{Stock cover for all stock} &= \text{£}4963/\text{£}1000 \\ &= 4.96 \text{ weeks} \end{aligned}$$

11

(a) TSL = Demand in lead time and review period plus safety stock in IT and RP

$$\begin{aligned} &= 200*(4 + 1) + 1.3 * 50 * (\sqrt{(4 + 1)}) \\ &= 1000 + 145.3 \\ &= 1145.3 \end{aligned}$$

(b) Assuming that the lead time is still four weeks and orders are placed daily:

$$\begin{aligned} \text{TSL} &= 200*(4 + 0.2) + 1.3 * 50 * (\sqrt{(4 + 0.2)}) \\ &= 840 + 133.2 \\ &= 973.2 \end{aligned}$$

(c) This would be a reduction of 172.1 in TSL. The average stock is TSL – half usage in (LT + RP). So reduction would be from 645.3 to 553.2. Reduction 92.1 units at £10 each = £921.

12 There are several ways of getting to this answer. The easiest is to perform Pareto Analysis and then apply batch sizes. Say A = every week, B = every fortnight, C = every four months. If every item is ordered each month, then there are 120 purchase events

Item no.	Unit cost (£)	Annual usage	Usage value per year (£)	Class	No of orders per year	Average stock value (£)	Percentage of value
2A32	25.00	12	300	B	24	6.25	7.5
2B44	0.50	360	180	C	3	30.00	4.5
3D10	0.10	120	12	C	3	2.00	0.3
3E82	40.00	1	40	C	1	20.00	1.0
3F66	2.00	40	80	C	3	13.33	2.0
4G19	5.00	520	2600	A	52	25.00	65.0
4H95	4.00	7	28	C	3	4.67	0.7
5J53	20.00	5	100	C	3	16.67	2.5
7N78	0.30	200	60	C	3	10.00	1.5
9P21	12.00	50	600	B	24	12.50	15.0
Total			4000		119	140.42	

Stock cover = 42.

From the table, the number of order events is now 119 (so no more work) and the stock is 0.42 of a month instead of 0.5. (See Chapter 3.)

13

Class	Stock (£)	Budget (£)	Stock cover (weeks)	% reduction required for £2000
A	28 000	133 000	0.91	7
B	6 500	6 500	1.76	31
C	3 500	2 000	7.58	57
	2 000	0	Infinite	100

Although the stock cover for A class is already under a week, it should be easier to gain £2000 by reducing this because it represents the majority of stock value. This can be achieved by scheduling or kanban. In practice some inventory reduction may also come from the other classes. (See Chapter 3.)

14 It does not pass on inventory, it eliminates it. However suppliers have to be able to respond rapidly to supplying to the smooth demand pattern required. (See Chapter 4.)

15 From the figures, the average batch only takes thirty minutes to produce but takes three hours to change over. This is inefficient. There needs to be a concerted effort on reducing the setup time. At the same time product design could be reviewed to see whether common parts can be created, so that set-ups can be avoided. There is also a problem with high WIP. Looking at the wasted time in between operations will show where this wasted time can be eliminated. The use of a kanban pull system would cause the WIP to decrease dramatically. (See Chapter 4.)

16 (a) The costs of manufacture and stockholding are reduced. Kanban create an environment where the WIP of the A class items is less than a day's worth between each process. The processes are linked together so that there is a flow of work. This enables efficient manufacture, and a reduction in waste, which contribute to lower operating costs.

The improvement in speed through the process means that any reject work is caught earlier and therefore can be rectified before too much material and capacity has been used up. There is also the direct reduction in inventory which removes working capital costs.

(b) *Responsibility.* Shop floor responsible for triggering supply of next item/batch without fail. Supplying operation responsible for replenishing in time.

Layout. Processes arranged in sequence with minimum distance between the processes.

Operator training. Operator responsible for their area and processes within it. Teamworking skills required as well as multiskilling. Operators need to be able to transfer from one process to the next in case there is a bottleneck.

Quality management. Excellent quality is a precursor to introducing kanban. Quality should be built into the process so that it will not go wrong. Quality is part of the manufacturing process, not a separate activity.

17 Safety stock decouples one process from the next. In JIT, the aim is to couple all the processes together to minimize timescales and inventory. Where JIT systems operate, the demand is smoothed to give similar requirements every day; consequently there is not the need for significant safety stock. Safety stock will be required where demand exceeds the flexibility of the manufacturing process.

18 Just in time throughput time is small, therefore there is not the need for detailed planning far ahead. This gives flexibility. By not having to forecast far ahead, then the forecasting process is more accurate. Manufacturing smaller batches means continuity of work and simpler planning, therefore less problems and less cost.

Many companies have been using JIT successfully for many years. They create the correct working environment, and applying a lean approach. They have great advantages over their competition because the costs of storage, interprocess stocks and non-value-added work are reduced.

19

Date	Receipts	Unit cost	Issues	Balance	No. left 14/3	From 10/6	Average cost
14/3	10	11		10	10		11
12/4			3	7	7		11
22/5			2	5	5	5	11
10/6	5	15		10	5	5	13
30/6			3	7	2	5	13
17/7			2	5	0		13

So, value of issue of the two items was:

(a) FIFO 22; average 26; LIFO 30.

(b) Highest stock value lines independent of usage rate. Reduce supply. See text, Chapter 5.

20 To absorb the different rates of supply and usage in two processes. The calculation of safety stocks normally only takes into account the variations in demand. The variation in supply is avoidable and should be controlled properly so that it is insignificant. The variation in demand is caused by the

‘customer’ and is therefore to be coped with and negotiated. The better the forecast of demand, the smaller the safety stock will have to be. (See Chapter 6.)

21 The calculation of safety stock depends on:

- The level of service being offered (acceptable risk of running out).
 - The predictability of the demand (MAD as a measure of forecast accuracy).
 - The replenishment lead time (but only as the square root of lead time).
- (See Chapter 6.)

22 (a)

Item	Wk 1	Wk 2	Wk 3	Wk 4	Wk 5	Average		
D523	60	30	20	40	50	40.0		
K488	30	15	0	65	0	22.0		
							MAD	MADP
D523	20	10	20	0	10	12.0		30%
D488	8	7	22	43	22	20.4		93%

(b) Lead time two weeks:

- CSF 90 per cent 1.6
- CSF 95 per cent 2.06

	MAD	SS 90%	SS 95%	Increase	Extra cost	% of turnover
D523	12.0	27.15	34.96	7.81	3.90	20
K488	20.4	46.16	59.43	13.27	6.64	60
					10.54	

Answer is £10.54.

23 Apply Pareto's Law, then organize the supply to save work (see also question 12). Change the order pattern for the 128 A class parts to weekly and schedule in to avoid paperwork and time. Order the B class every four weeks as they fall below the review level, and order the 896 C class items every eight weeks to reduce the workload.

Then work out the better safety stocks, and flexibility of supply for the A class. It is important to measure customer service level and inventory levels while the change is happening, and also to make sure that the number of orders are reducing significantly. (See Chapter 7.)

24 The reordering system creates replenishment orders, which open the supply. However demand may not be satisfied because:

- *From the theory.* The review level is calculated assuming a customer service factor which gives a small risk of stockouts (zero risk is impossible).
- *In practice.* Supply may be later than expected and the stock is insufficient to cope with the longer lead time.

25 To set the safety stock requires demand. To set the average stock level or to order requires in addition:

delivery quantity variability of demand (MAD or SD)
 level of availability required (customer service factor)
 supply lead time.

To set the review level requires in addition, forecast.

Underlying the forecast and MAD calculations are therefore an assessment of demand pattern using historical information backed up by understanding of how future demand may vary from this.

26

Average 65: safety stock = 16.67.

MAD 8.33: review level = 81.67.

average stock = 131.67.

27 The reorder level is used to trigger new replenishment orders when the stock + supply orders outstanding become low. If the delivery frequency is shorter than the lead time then the stock is low, delivery quantities are small and there are always supply orders outstanding. Hence, for long lead time A class items, ordering is frequent and stock is always less than the review level.

28

A class – order in daily, if possible on Kanban.

C class – order in weekly.

See Chapter 7.

- 29 (a) Stock level depends mainly on supply delivery quantities, customer service and demand pattern. Lead time variation is only as the square root. (However, delivery performance may be lead time dependent also.)

$$\text{Average stock} = \text{half delivery quantity} + \text{customer service factor} \times \text{MAD} \times \sqrt{\text{lead time.}}$$

So changing the lead time has the least effect on stock levels.

- (b) Average = 13.
 MAD = 3.33.
 Customer service factor is 1.6.
 Safety stock = $1.6 \times 3.33 \times \sqrt{2} = 7.53$
 Reorder level = $13 \times 2 + 7.53 = 33.53$

- (c) Continuous demand.
 Continuation of the same demand pattern.
 Normal distribution and non-seasonal demand.
 Supplier delivers on time.

- 30 (a) Work out with the supplier:

Benefits to supplier as well as customer (continuous demand, earlier payment, more consistent demand).
 Delivery arrangements (transport, packaging).
 Dispatch and receiving reorganization.
 Information processing (schedules and invoicing).

- (b) Reduction in inventory:

Average stock = half delivery quantity + safety stock.
 Assume the safety stock remains the same.
 Delivery quantity goes from four weeks' worth to one week's worth, so stock saving is two weeks to half a week.
 At £1 million per year 1.5 weeks is £28 846.

(c) Benefits:

- Lower inventory investment – better cash flow.
- Less stores space required.
- Easier movement within stores.
- Improved information system (paperwork?) efficiency.
- Extra employment in the long term.
- More frequent (but smaller deliveries).
- More receiving document processing.
- More monitoring of inventory and co-ordination with suppliers.

31 The delivery quantity primarily depends on how frequently the supplier is prepared to deliver without significant increase in cost. It will depend upon:

- The number and variety of other items bought from the same supplier.
- The type of items (e.g. handlable, slow-moving).
- The quantity in which they are to be consumed.
- The turnover value of the item.

32

- Better forecasting.
- Treating demand as dependent and avoiding safety stock effects.
- Measuring end user usage, structuring supply source.
- Integrating systems
- Square root of warehouses effect.

33 Target stock levels: variable order quantities fixed time interval between deliveries. Ideal where supply of many B and C class items are sourced from a single supplier. Organize purchasing to look at suppliers in rotation.

Review level system: interval between ordering is variable, order quantities are fixed by an ordering rule. Normal way to trigger orders. Good for A class items and slow-moving and one-product suppliers.

(See Chapter 8.)

34 (a) Extrinsic forecasts to give causes of demand.

- Market analysis and surveys.
- Historical techniques.

Moving average: OK as demand is random and varying much.

Exponential smoothing: better since there could be a slight trend.

Double exponential: will be overreactive until a trend is seen.

Base series: not appropriate as seasonality not established.

Regression: no advantage over other methods.

More sophisticated techniques: insufficient data for these.

(For details, see Chapter 10.)

(b) Exponential smoothing with $\alpha = 0.2$.

Week	Demand	Forecast
1	12	10.00
2	14	10.40
3	8	11.12
4	16	10.50
5	13	11.60
6	18	11.88
7	?	13.10

35 MRP requires:

master schedule: 95 per cent accurate.

bills of materials: 98 per cent accurate

inventory records: 95 per cent accurate

and

lead times (for suppliers and any manufacturing processes).

manufacturing routings (where appropriate).

36 *Components*

Watch	Modules	Battery quantity
X	A	2
Y	A	2
Z	B	1

Manufacturing orders

Watch type	Week 1	Week 2	Week 3	Week 4	Week 5	Week 6	Week 7	Week 8	Week 9	Week 10
X		1000		1000		1000		1000		1000
Y			500			500			500	
Z	100	100	100	100	100	100	100	100	100	100

Component demand

Module A

Gross requirement		1500	0	1000	500	1000	0	1500	0	1000	0
Stock on hand	1500	0	0	-1000	-1500	-2500	-2500	-4000	-4000	-5000	-5000
Purchases outstanding					3000						3000
Projected on hand		0	0	-1000	2500	1500	1500	0	0	-1000	2000

New orders for module A of 1000 for week 3 and 1000 for week 9.

Module B: excessive stock of module B is building up.

Battery

Gross requirement		3100	100	2100	1100	2100	100	3100	100	2100	100
Stock on hand	3500	400	300	-1800	-2900	-5000	-5100	-8200	-8300	-10400	-10500
Purchases outstanding					6000						6000
Projected on hand		400	300	-1800	3100	1000	900	-2200	-2300	-4400	1500
Planned receipts				2000				2500			
Projected available		400	300	200	5100	3000	2900	2300	2200	100	6000

New orders for the battery of 2000 for completion in week 3 and 2500 for week 7.

(b) Recommend that schedule is smoothed. There are many ways of doing this. One option is:

Module A

Module A	Week	1	2	3	4	5	6	7	8	9	10
Gross requirement		1500	0	1000	500	1000	0	1500	0	1000	0
Stock on hand	1500	0	0	-1000	-1500	-2500	-2500	-4000	-4000	-5000	-5000
Reschedule receipts				1500	500	1000	500	1000	500	1000	500
Projected on hand		0	0	500	500	500	1000	500	1000	1000	1500

Try to reschedule the module supply as the schedule above.

Battery

Battery	Week	1	2	3	4	5	6	7	8	9	10
Gross requirement		3100	100	2100	1100	2100	100	3100	100	2100	100
Stock on hand	3500	400	300	-1800	-2900	-5000	-5100	-8200	-8300	-10400	-10500
Reschedule receipts				1400	1400	1400	1400	1400	1400	1400	1400
Projected on hand		400	1700	1000	1300	600	1900	200	1500	800	2100

Try to reschedule the battery supply to 1400 per week

37 (a)

Week	Stock	1	2	3	4	5	6	7	8	9	10
Sales		210	300	140	65	200	260	150	100	190	220
FG	650	440	140	0	435	235	475	325	225	35	315
Order receipt					500		500				500
Order release		0	500	0	500	0	0	0	500		
Semi-finished	220	220	220	220	220	220	220	220	220	220	220
Order receipt			500		500				500		
Order release		0	500	0	500	0	0	0	500		
Raw material	500	0	0	0	0	0	0	0	0	0	0
Purchase receipt				500				500			
Purchase ordering		500				500					

(b)

Week	Stock	1	2	3	4	5	6	7	8	9	10
Sales		210	300	140	65	200	260	150	100	190	220
FG	650	440	140	0	0	0	0	0	0	0	0
Order receipt					65	200	260	150	100	190	220
Order release		0	65	200	260	150	100	190	220		
Semi-finished	220	220	155	0	0	0	0	0	0		
Order receipt				45	260	150	100	190	220		
Order release		0	45	260	150	100	190	220	0		
Raw material	500	500	455	195	45	0	0	0	0		
Purchase receipt						55	190	220			
Purchase ordering		0	0	55	190	220					

38 (a) Distribution requirements planning takes the forecast from warehouse and aggregates it into the master schedule. It therefore depends upon the warehouses to forecast in the first place.

Fair shares can manufacture against a forecast and then allocate the inventory against the warehouse requirement.

(b) The logic is:

- 1 Any unfulfilled demand from the warehouses.
- 2 Forecast usage which will be unfulfilled within the supply lead time.
- 3 Top up warehouse safety stock.
- 4 Allocate to central stock safety level if appropriate.
- 5 Distribute remainder in proportion to the forecast warehouse usage rates.
- 6 Transfer excesses to other warehouses if appropriate.

39

Week	Stock	1	2	3	4	5	6	7
London sales		150	320	370	190	150	320	379
Current stock	600	450	130	360	170	20	300	521
Order receipt				600			600	600
Order release		600	0	0	600	600	0	0
Manchester sales		70	70	70	70	70	70	70
Current stock	230	160	90	270	200	130	60	240
Order receipt				250				250
Order release		250	0	0	0	250	0	0
Newcastle sales		90	150	150	90	90	150	150
Current stock	310	220	70	10	10	10	40	70
Order receipt				90	90	90	180	180
Order release		90	90	90	90	180	180	
Internal sales		30	30	30	30	30	30	30
Total		970	120	120	720	1060	210	30
Current stock	1650	680	560	440	720	660	450	420
Order receipt					1000	1000		
Order release		1000	1000	0	0			

Manufacturing orders required in weeks 1 and 2.

40 Advantages:

Retail stores satisfy customer demand and, therefore, stock is available immediately to meet demand.

Any distribution problems do not affect availability.

No central stock is necessary.

Slow distribution method can be used to the retail stores.

Disadvantages:

More stock is required because there are more stock points.

Any change of products requires emptying the whole supply chain.

Redistribution of excess or supply to meet an individual shortage is more difficult.

41

Item no.	Unit cost	Annual usage	Turnover	% turnover	Cumulative % turnover
2	125	100	12 500	62.5	62.5
4	50	70	3 500	17.5	80.0
9	7	300	2 100	10.5	90.5
3	85	10	850	4.3	94.8
6	26	20	520	2.6	97.4
1	150	2	300	1.5	98.9
10	6	20	120	0.6	99.5
7	20	3	60	0.3	99.8
5	40	1	40	0.2	100.0
8	10	1	10	0.1	100.0
Total		527	20 000		

- (a) A class: item 2.
 B class: items 4, 9.
 C class: items 3, 6, 1, 10, 7, 5, 8.
- (b) ABC determines delivery quantity and reduces purchase workload (see Chapter 3).
- (c) Discussion of optimizing profits, lead time reduction or cycle counting.

42 (a) Independent demand

Assume that a 99.2 per cent customer service level is required.

Average demand = 12.

MAD = 4.4.

Customer service factor = 3.

Safety stock = $3 \times 4.4 \times \sqrt{(1)}$
 = 13.4

Delivery quantity = 12.

Average stock = Half delivery quantity + safety stock
 = $6 + 13.4$
 = 19.4

(b) Dependent MRP demand (level 2), assuming no safety stock is required, stock would be:

Week	0	1	2	3	4	5
Demand	10	20	5	10	15	

Average stock = 12.

(c) JIT supply

Assume that demand is the same every day.

(Alternative assumptions are that there should be a safety stock calculated as (a) or as one-day lead time.)

Week	0	1	2	3	4	5
Demand	10	20	5	10	15	
Daily demand	2	4	1	2	3	

Assuming that demand is dispatched every day:

Average stock = half day's demand:

1, 2, 0.5, 1, 1.5

So average stock during the period is 1.2 (assuming a one-day delivery interval).

This Page Intentionally Left Blank

Index

- ABC Analysis, 36–42
- Accurate records, 200, 214
- Aggregate inventory, 81
- Aggregate stock valuation, 91
- Allocation of stock, 79, 221
- Assessment of demand, 164
- Availability, measurement of, 18–21
 - of stock, 100
- Average value, 88

- Backorder, 22–4
- Backward scheduling, 198
- Bar coding, 158
- Base series technique, 194
- Batch, 6, 11, 52, 60, 68, 158, 191, 202
 - sizes, 210
- Baysian forecasting, 187
- Bias, 159
- Bill of material (BOM), 200, 203

- Collecting historical information, 96
- Company strategy, 15
- Consignment stock, 216
- Consuming forecast demand, 27
- Contributions of history, 174
- Control methods, 215
- Customer feedback, 27
 - relations, 17
 - service, 4, 7, 16, 17–32, 52, 82, 100, 112, 198, 212
 - service factors, 105
 - support, 71
- Cusum technique, 181

- Data aggregation, 162
- Delivery promising, 208
 - times estimation, 26, 29
- Demand history, 165
 - management, 25–31
 - patterns, 155
- Dependent and independent demand, 198
- Depreciation, 91
- Despatch information, 26
- Deviation measurement technique, 103
- Distributed inventory, 222
- Distribution requirements planning (DRP), 225
- Double exponential smoothing, 188

- Economic order quantity (EOQ), 73, 145
- Electronic data interchange (EDI), 26
- Excess and obsolete items, 109
- Expediting, 95, 117, 149
- Exponential smoothing, 171–4

- Fair shares, 222
- First in, first out (FIFO), 87
- Fixed classification, 42
- Forecasting advanced, 187–96
 - choosing the best types, 178
 - demand, 152–4
 - focus, 179
 - historical, 164–86, 193
 - inaccuracy, 158

- Forecasting advanced – *continued*
 methods for improvement, 162
 monitoring, 162, 181
 seasonal, 192
 tracking, 182
 Fourier analysis, 187
- Gaussian, 98
- Historical data analysis, 96
 forecasting, 164–86
- Information accuracy, 214
 Initializing the forecast, 175
 Inventory control objectives, 7
 forecast types, 153
 investment, 48, 62, 72, 82
 management, 4–16, 33–58
 valuation, 87
- Just in time management, 59–77
- Kanban, 12, 76
- Lag, 189
 Last in, first out (LIFO), 88
 Layout and location of stores, 14
 Lead time, effects of, 121
 inconsistent, 121
 management, 117–20
 reduction, 53–8, 74, 119
 Lean supply chain, 212
 Logistics, 7, 59, 78, 129, 217
 Lot for lot, 210
- Manufacturing lead time, 119
 Market change, 13
 Master planning, 205–11
 scheduling, 205
 Material requirements planning,
 197–211
- Mean absolute deviation (MAD), 99,
 102, 152, 176, 217
 Minimum stock levels, 155, 164
 order constraints, 147
 Moving average, 166
- Non-movers, 42
 Non-stock items, 43
 Normal demand patterns, 100
- Obsolescence, 13, 109, 157
 Order processing, 26, 118
 Order quantities, 141
 Ordering, 139
 Organization and management, 78
 Overheads, 71, 74, 88
- Pareto based order quantities, 142
 Pareto stock balance, 48
 Pareto's Law, 33–42, 109
 Perpetual inventory, 39, 62, 158
 Physical distribution, 222
 Physical inventory, 158
 Planned inventories, 14
 Planning the structure of supply, 204
 Poisson's Law, 108
 Pricing method, 139
 Product life cycle, 156
 Profit through inventory management,
 9
 Projected available, 207
 Pull (system), 60, 76, 221
 Purchase commitment, 85
 order, 15, 50, 57, 72, 95, 119,
 139–41
 order quantity, 11
 Purchasing, 129–51
 Push (system), 60, 76–7, 221
- Quality management, 69
 Queue, 84
- Reasons for stock, 11
 Record accuracy, 70, 214

- Reducing stockholding, 51–8
- Regression analysis, 187
- Reorder point, 139–40
 - quantity, 149
- Repetitive processes, 65
- Replacement value, 88
- Replenishment systems, 213, 225–8
- Review levels, 112–16
- Review of techniques, 212
- Review period, 126
- Risk measurement, 104

- Safety stock, 12, 47, 50, 52, 81,
96–111, 112, 152, 197, 204
- Scheduled supply, 143
- Setting stock levels, 112
 - stock targets, 46
- Single sourcing, 132
- Skills and systems, 92
- Slow moving items, 108
- Slow systems, 16
- Smoothing constant, 172
- Standard cost, 88
- Standard deviation, 100, 114
- Stock availability, 100
 - cover, 43
 - distribution, 222
- Stockkeeping unit (SKU), 89, 220
- Stores layout, 14

- Supplier lead time, 95, 118, 143
- Supply chain, 129, 218
- Supply practice, 129
 - partnerships, 137
 - practice, 129
 - structure, 204
- Systems and control, 15

- Target stock levels, 124–8
- Total quality management (TQM), 70
- Tracking signal, 182
- Training, 93
- Transport algorithm, 223
- Trigg's tracking signal, 182
- Turnover of stock, 43

- Unit cost of stock lines, 87

- Variance, 103
- Vendor appraisal, 134–7

- Weighted averages, 170
- Work in process (WIP), 200

- Zero inventories, 59

This Page Intentionally Left Blank