SUPPLY CHAIN DEMAND MANAGEMENT IN THE FOOD SECTOR

David Nicholas BRYANT

This Thesis is submitted for the degree of

Doctor of Philosophy in the

University of London and for the

Diploma of Imperial College

Feb. 99

Department of Electrical Engineering

Imperial College of Science, Technology & Medicine

LSE
LONDON
UNIV.

ABSTRACT

In order to meet demands for improved efficiency and customer service, higher quality and reduced environmental pollution, supply chain management and operation are increasingly being based on 'partnership' rather than confrontational principles. Partnership operation usually involves sharing of demand information and a more integrated approach to management of the entire supply chain. This research seeks to develop improved methods of demand management, which exploit partnership operation within the food sector.

Methods are developed for modelling, for daily monitoring and forecasting demand, firstly at the till and then along the entire supply chain. Also methods for detecting unexpected changes in demand level or pattern, especially during product promotions, are developed. These have been extensively tested on real data and are suited to partnership sharing initiatives and vendor management schemes.

Basic results for fixed cycle replenishment systems are developed together with a corresponding approach to multi-product warehouse replenishment. A manufacturing based decision support system is developed which integrates manufacturing scheduling and warehouse operation to obtain 'synchronous replenishment'. It also provides procedures for tuning key supply chain parameters. The system has been assessed, for a specific plant and warehouse, using real operational data.

The use of short term forecasting partnership and based approaches to distribution management is considered. This involves pulling orders forward and relaxing normal 'on time' commercial constraints. In a case study the approach yielded near 100% lorry loading, a 35% drop in deliveries and lorry trips, with minimum stock increase at the retail warehouse.

To my wife Elizabeth, my daughter Sophia, and to my Parents.

ACKNOWLEDGEMENTS

Firstly, I wish to thank my supervisor, Professor R. I. Kitney, for his encouragement, interest and guidance in completing this thesis.

Thanks are also expressed to

Robin Wells – Supply Chain Group, Tesco

Barrie Smith – Distribution Manager – Ambient, Unilever

David Mace – Supply Chain Team, Unilever

who gave generously of their time and expertise. Also many thanks to the Directors of Greycon Ltd, for their financial and technical support.

CONTENTS

		Page
Title page		1
Abstract		2
Dedication		3
Acknowledgements		4
Contents		5
CHAPTER 1		
1.0	**SUPPLY CHAINS AND THE FOOD SECTOR**	
1.1	Introduction	11
1.2	The Food Sector Supply Chain	13
1.3	Supply Chain Activities	16
	1.3.1 Raw Materials	16
	1.3.2 Manufacturing	16
	1.3.3 Manufacturing Warehousing	17
	1.3.4 Retail Regional Distribution Centres	17
	1.3.5 Retail Stores	18
1.4	Management Strategies	19
1.5	Supply Chain Performance Measures	22
	1.5.1 Non Financial Measures	22
	1.5.2 Financial Measures	23
	1.5.3 Supply Chain Benchmarking	24
1.6	Information Exchange	25

1.6.1 Customer Forecasts 25

1.6.2 Exception Reporting 26

1.6.3 Supplier 26

1.6.4 Responsiveness 26

1.7 Research Program 29

1.7.1 Forecasting 30

1.7.2 Inventory Management 31

1.7.3 Distribution 32

1.7.4 Claims of Originality 33

References 35

CHAPTER 2

2.0 DEMAND MODELLING AND WEIGHTED LEAST SQUARES APPROACHES

2.1 Introduction 37

2.2 Model Formulation 38

2.2.1 Algorithm Selection 43

2.2.2 Overdispersion 44

2.3 Least Squares Approaches 45

2.4 Singular Value Decomposition 45

2.4.1 Result 1 (Eckart and Young) 46

2.4.2 Generalised Least Squares Approaches 48

2.5 Results and Discussion 51

2.5.1 Till Data Results 51

2.5.2 Store Daily Mode 53

2.5.3 Store Seasonal Mode 55

2.5.4 RDC Data Results 55

2.5.5 RDC Daily Mode 58

2.5.6 RDC Seasonal Mode 59

References 60

CHAPTER 3

3.0 MAXIMUM LIKELIHOOD APPROACHES

3.1 Introduction 62

3.2 Off-Line Results 62

3.3 On-Line Developments 65

3.4 Event Detection 69

3.5 Results and Discussion 72

3.5.1 Till Data Results 72

3.5.2 Store Daily Mode 74

3.5.3 Store Seasonal Mode 75

3.5.4 RDC Data Results 76

3.5.5 RDC Daily Mode 78

3.5.6 RDC Seasonal Mode 79

3.5.7 Event Detection – Promotional Mode Adaption 80

3.5.8 Event Detection – Demand Level Change 81

3.6 Replenishment System Effects 83

3.7 Off-Line Modelling of Weather Mode Effects 85

3.7.1 Weather Mode Results 87

3.8 Appendix 1 Maximum Likelihood Derivations 92

3.9 Appendix 2 Adaptive Exponential Smoothing 94

3.10 Appendix 3 Weather Mode Derivations 95

References 98

CHAPTER 4

4.0 INVENTORY MANAGEMENT CONTROL

4.1 Introduction 99

4.2 Fixed Cycle Systems 99

4.2.1 Deterministic Demand 100

4.2.2 Normal Demand Distribution 102

4.2.3 Poisson Demand Distribution 103

4.2.4 Batch Flow Effects 104

4.3 Multi Product Inventory Control 108

References 112

CHAPTER 5

5.0 MANUFACTURING DECISION SUPPORT SYSTEM DESIGN

5.1 Introduction 113

5.2 System Functionality 116

5.2.1 On-Line Operation 117

	5.2.2 Off-Line Operation	118
5.3	System Modelling	120
	5.3.1 Raw Materials and Packaging	121
	5.3.2 Information Requirements	121
	5.3.3 Raw Materials Scheduling	123
	5.3.4 Manufacturing Processes	124
	5.3.5 Production Scheduling	125
	5.3.6 Product Demand Data Base	128
5.4	Synchronous Replenishment System	130
	5.4.1 Forecasting	130
	5.4.2 Stock Replenishment	131
	5.4.3 Priority Allocation	131
	5.4.4 Production Scheduling	132
5.5	User Interface Design	133
	5.5.1 Products	133
	5.5.2 Units	135
	5.5.3 Orders	136
	5.5.4 Workcentres	138
	5.5.5 Schedules	141
	5.5.6 Reactive scheduling	142
	5.5.7 Synchronous Replenishment System	145
5.6	Parameter Tuning – Results and Discussion	148
	5.6.1 Weeks Forward Cover	149
	5.6.2 Minimum Run Length	162
	5.6.3 Safety Stock and Forecast Accuracy	165

5.6.4 Results and Conclusion 165

References 167

CHAPTER 6

6.0 VEHICLE LOAD AND DISTRIBUTION

6.1 Introduction 168

6.2 OTIF Relaxation 168

6.3 Case Study 171

6.4 Retailer Benefits 174

6.5 Manufacturing Benefits 174

References 177

7.0 DISCUSSION AND CONCLUSIONS 178

1.0 SUPPLY CHAINS AND THE FOOD SECTOR

1.1 Introduction

Supply chain management has been defined as "the organisation, planning, control and execution of the goods flow from development and purchasing through production and distribution to the final customer in order to satisfy the requirements of the market at minimum cost and minimum capital use." Holmes [95]. This definition can usefully be extended to service providing activities such as health care and airlines, which also have substantial supply chains. Ballon [99].

Modern supply chains have to meet demands for high levels of customer service, a wide choice of products, fast response to market demand, low stock, improved quality and reduced environmental pollution. In attempting to meet these conflicting objectives, suppliers, manufacturers, distributors and retailers have sought to exploit technological advances such as Electronic Data Interchange (EDI) and integrated computer control. There is also an increasing practice of seeking to run the supply chain as a partnership rather than through confrontation. Management opinion of most UK leading food retailers accept that they cannot further improve their current performance without establishing effective partnership operations. Under partnership operation the aim is to run the supply chain, from raw materials, through manufacturing and retailing, to the end customer, as an integrated whole. This collaboration includes the sharing of information and the exploitation of improved communication through EDI and improved decision making .

This thesis seeks to develop concepts, analysis and technology needed for more effective partnership based approaches to supply chain management within the food sector. The food sector, part of the Fast Moving Consumer Goods (FMCG) sector, is generally accepted as one of the most advanced in its supply chain development.

The characteristics of the food supply chain differ, for instance, from more traditional "engineering" based manufacturing (cars, aircraft etc.) as shown in Fig. 1. (Source:Unilever). The food manufacturer's stock is dominated by raw materials and finished goods, whereas the engineering manufacturer typically holds far more in-process stocks in the form of components and sub-assemblies. From the cost structure chart it can be seen that the food sector spends a higher percentage on sales, marketing and distribution due to the highly competitive nature of the FMCG marketplace. The engineering manufacturer adds more value to the raw materials and bought-in items through what is generally a more complex manufacturing process. The high levels of customer service, fast response to market demands and efficient distribution required in the food sector suggest a need to concentrate resources on improving supply chain performance, particularly to the end customer.

COMPARISON OF A TYPICAL FOOD MANUFACTURER
WITH A MAJOR ENGINEERING COMPANY

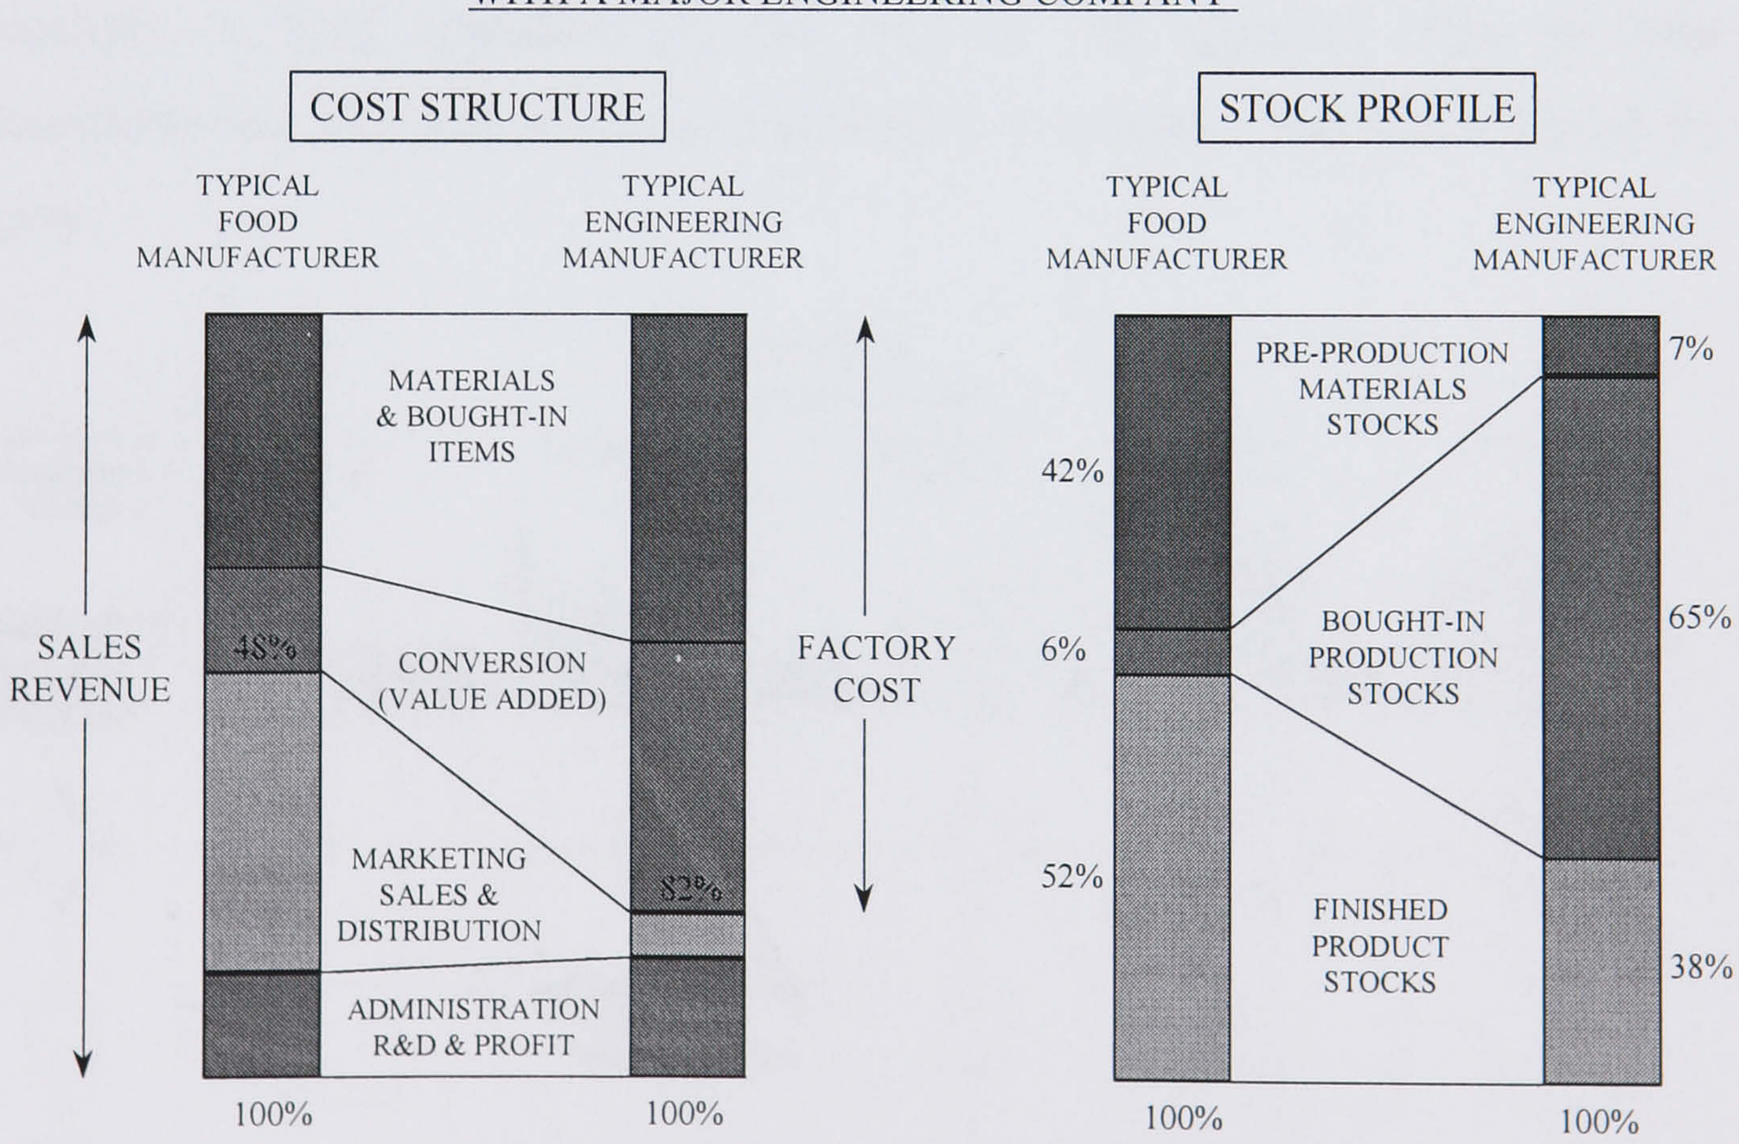


Fig. 1

1.2 The Food Sector Supply Chain

Consumer spending on food within the UK reached an estimated £37.2bn in 1995 and is estimated to reach £42.6bn by the year 2000. The number of food retailers is highly concentrated with the top five grocery retailers (Tesco, Sainsburys, ASDA, Safeway and Somerfield) sharing 66% of all retail food sales. Of the top four, the top two (Tesco and Sainsburys) represent 37% of the market. There are some 30,000 supermarkets in the UK carrying on average 30,000 product lines occupying 37 million square feet and employing 320,000 people. The distribution network required to service the UK food market carried 37.5 billion tonne kilometres in 1995 making it the largest category of road transport users.

Virtually every product we buy has been through a supply chain. A typical supply chain consists of raw material suppliers, manufacturers and retailers. Virtually all of these activities have and are part of a large supply chain. For example, a food manufacturer (see Fig. 2) will typically have its own manufacturing site and a storage/distribution warehouse which are linked by lorry.

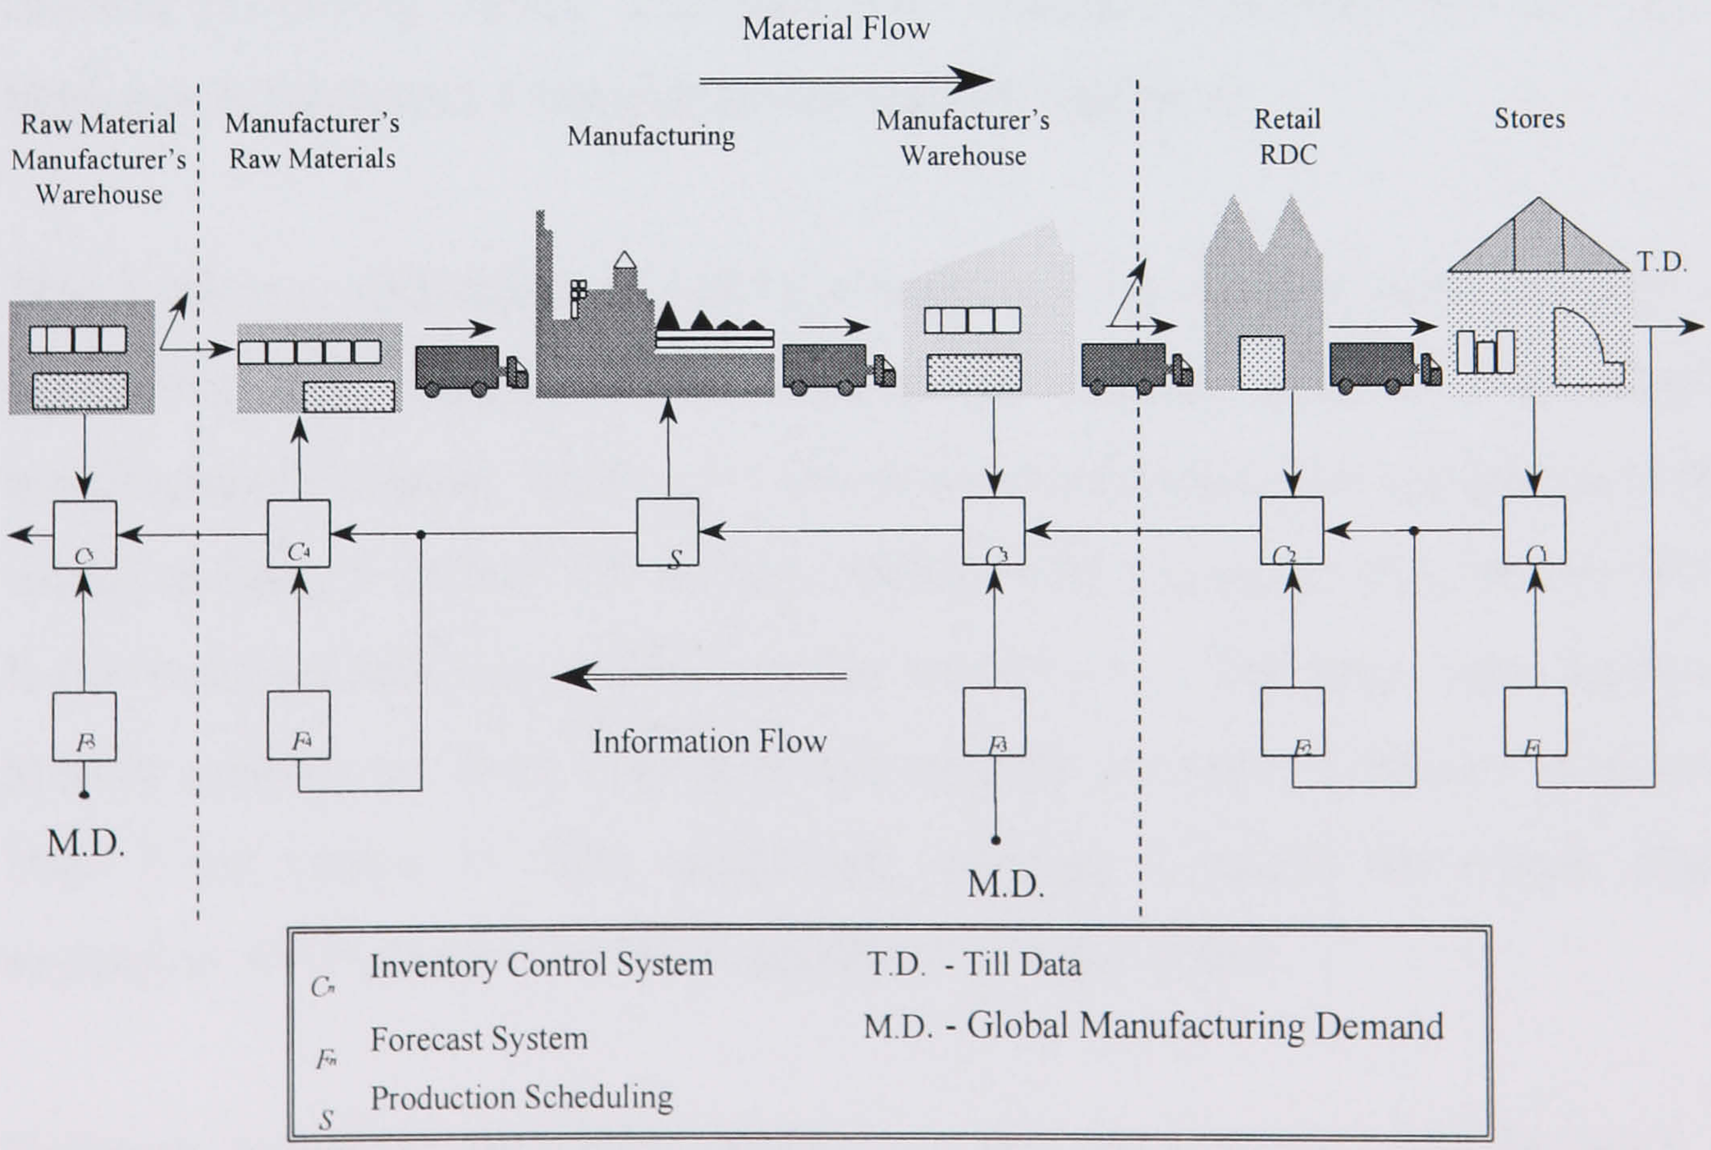


Fig. 2

This is the "internal" supply chain over which the manufacturer has complete control. Raw materials and packaging will be delivered to the manufacturer's warehouse via other supply chains. The finished product will then be transferred from the distribution warehouse to many customers e.g. retailers. The raw materials supplier and retail customer are part of the "external" supply chain over which the manufacturer has a varying degree of control depending on the extent of power they can exert in the market place. The power of "suppliers" and "customers" anywhere in a supply chain depends on several key factors. Suppliers are likely to be powerful when they are high in concentration, when there is customer dependency on the supplier`s product, when there is the possibility of supplier forward integration if sufficient margins are not obtained and when the supplier's individual customers are of little importance, due to their large number. Customers (e.g. Tesco) are likely to be more powerful when they are high in concentration, when there are alternative sources of supply and there is a threat of backward integration if satisfactory prices cannot be obtained. Johnson and Scholes [88]. Clearly, given the high concentration of grocery retailers and the increasing number of "own brand" products, the retailers are generally more powerful than the suppliers, emphasising the suppliers need to be reactive to customer needs. Prodham [93]. Similarly, manufacturers tend to be more powerful than raw material and packaging suppliers.

This work was undertaken as part of a major collaborative research program in partnership with Tesco plc, the leading UK retailer, and the Anglo-Dutch manufacturer Unilever. Unilever is one of the largest industrial companies in the world, ranking 21 in the 1995 Fortune 500 list with a turnover of £30bn in 1997. It has business interests in 88 countries world-wide. The three main areas of product category are food, detergents and personal products. Unilever owns over 1000 brand names in these categories, including Cornetto ice cream, Flora margarine, OXO, Birds Eye Walls and Persil to name a few.

Unilever suggested that their Batchelors dry soup range, manufactured at Ashford, as a candidate for supply chain research. This involved taking a "slice"

of the overall supply chain as obviously it would be an impossible task to research the entire range of products and customers. The "slice" limited consideration to a specific product group and a specific partnership. However a market-wide view had to be taken at the manufacturer's warehouse and plant, as Unilever do not "make to order" and hence customer differentiation was not possible. Total market demand is also important in warehouse operation and distribution as all warehouse products have to be considered. The Tesco Middleton Regional Distribution Centre (RDC) was selected as it had the largest dry soup demand. The supply chain was defined for this research as beginning at the interface between the Unilever plant and the raw materials supplier through to the Tesco store and the end consumer as illustrated in Fig. 2.

Tesco Plc is currently the UK's largest supermarket chain. They have approximately 580 stores with a combined annual turnover of £17.8bn in 1997-98. They have five store formats; Superstore, Supermarket, Compact, Metro and Express. The Superstores are the large out of town stores, the Supermarkets are slightly smaller and are generally situated nearer towns. The Compact is a smaller version of the Supermarket, the Metro is a smaller version of the Superstore with a greater accent on convenience products, aimed at affluent towns. The Express is a micro store contained within petrol stations. Tesco also has stores in Ireland, Hungary, Poland, the Czech Republic and Slovakia. Tesco are at the forefront of new retail initiatives such as loyalty cards, home delivery, 24hr opening and personal banking.

1.3 Supply Chain Activities

Considerable time has been spent on site studying the manufacturer and retailer current business operations. Data was collected from most activities to facilitate modelling the existing "as is" approach. A brief description of the key supply chain activities is given below.

1.3.1 Raw Materials

Unilever work on a two week re-order cycle. They nominally hold two weeks stock plus safety stock at their raw materials warehouse. The amount of stock held is measured in weeks forward cover (wfc) which represents the number of weeks, based on current demand forecasts, the existing stock held will last without replenishment. If stock of a material is insufficient to meet the forecast demand plus safety stock, an order is placed with the supplier. Lead times have to be taken into consideration as some of the more exotic spices used in dry soup have to be imported. These ingredients are stocked at higher wfc to avoid frequent reordering. There are approximately 1200 raw materials required to produce the dried soups range.

1.3.2 Manufacturing

The dry soup manufacturing activity consists of four parallel production lines. Line A produces Cup-a-Soup (CAS) and Slim-a-Soup (SAS). Line B produces Cup-a-Soup Special (CASS), Slim-a-Soup Special (SASS) and Cup-a-Soup World (CASW). Line C is an overflow for Line B (its production rate is slower) and Line D produces Variety Packs. The production process involves three main stages. Firstly, the raw materials are mixed according to the product recipe. The mix is loaded into a dispensing hopper which supplies the correct quantity to fill a single sachet. The filled sachets are conveyed to the cartoning stage where a number of sachets are placed in a box which is shrink rapped and stored on a

pallet ready for delivery to the Peterborough warehouse. A carton of dried soup holds either nine or fourteen sachets, depending on the variety. A pallet typically holds 256 cartons per layer and can hold up to 4 layers. Once the run is complete, the Bill of Materials updates the stocks of each raw material used.

1.3.3 Manufacturing Warehousing

The Unilever warehouse at Peterborough holds stocks of finished goods to service the customers' demands. Each of these products is maintained at a level of approximately 6wfc plus safety stock. Safety stock is specified as a fixed percentage of wfc, typically 35%. Unilever does not make products "to order" i.e. to fulfil a particular customer order. The factory produces stock in order to maintain the total stock at the warehouse at the wfc level, based on the most recent total demand forecasts. Replenishment stocks are delivered from the Ashford factory using a virtually continuous lorry service. Once an order is received, it is processed and picked. The order is then entered into the vehicle routing system and a delivery date established. The Unilever warehouse aims to deliver orders within four working days of receipt. In some cases this can be shortened, depending on the availability of stocks and the importance of the customer.

1.3.4 Retail Regional Distribution Centre

The Tesco RDC is a warehouse servicing the stock requirements of the 580 or so stores. Tesco has 13 RDCs and a further 6 Centres which are sub-contracted. They are sited to serve a geographical region with particular categories of product. A typical RDC covers about 300,000 square feet and handles 50 million units a year. Generally there are two RDCs holding similar product categories, one serving the northern and one the southern half of the UK. The categories represent groupings of products with similar storage requirements (ambient, chilled and frozen etc.). This research is primarily concerned with the Middleton

RDC as this is where dry soups are stored. Goods In, storage and goods out are organised by separate computer systems. Many of the resources required for unloading are used for loading (manpower, fork lifts etc.) and storing, so careful scheduling is required to manage these finite capacity resources.

1.3.5 Retail Stores

As stated earlier, Tesco have five store categories, Superstores (255), Supermarkets (220), Metro Stores (48), Compacts (42) and Express (21). The stores are serviced with stock from various RDCs providing different product categories. The stores will be replenished from a particular RDC between one and six times per week, dependent on the store size and turnover. Stock is delivered in two waves, the first contains perishable goods with short shelf life, where little or no stock is held at the store. These arrive between midnight and 8 a.m. The second wave arrives during the day. This contains items that have longer shelf lives such as frozen and non-food items; these are supplied for stocking shelves prior to the following days trading. Stock is unloaded from the lorry, unpacked and transferred directly to the shop floor. In some cases, where the delivery exceeds shelf capacity, excess stock is stored in the back of the store. For Tesco, the process of transferring stock to the shop floor represents the single largest cost in their internal supply chain, due to the high demand on labour resources.

1.4 Management Strategies

In the competitive environment of the late twentieth century, innovations in corporate management strategy have life cycles of about fifteen years. These innovations produce shifts in competitive position and corporate fortunes, and there after become accepted wisdom. Stalk and Thomas [90].

Three innovations, first developed within the Japanese automotive industries, have gained wide acceptance and have influenced management approaches and objectives. The first is that of "Just in Time" (JIT) manufacturing, the second that of "Lean" operation and the third is that of "Time Based" competition. These have been encapsulated in the Efficient Consumer Response (ECR) movement.

The JIT approach has been summarised as "produce and deliver finished goods just in time to be sold, subassemblies just in time to be assembled into finished goods, fabricated parts just in time to go into subassemblies, and purchased materials just in time to be transformed into fabricated parts" Schonberger [82]. The JIT philosophy has been translated into the language of retailing with such concepts as "synchronous flow" embodied in the ECR objectives.

The lean manufacturing philosophy Voss [92] seeks to minimise resource usage generally and to offer increased customer choice and improved service. It has been extended into the wider concept of the supply chain, inventory being regarded as waste which should be eliminated, and otherwise minimised Saunders [94].

Time based management of factors within the supply chain is now recognised as a significant source of competitive advantage and an important strategic weapon. In the US retail sector. Walmart grew three times as fast as its competitors and enjoyed twice the return on its capital. It accomplished this by reducing its supply chain inventories to a quarter of that of its competitors and by offering four times the choice of product. It replenished the stores at approximately four

times the rate of its competitors. Many of its stores replenished daily, whilst its competition replenished every two weeks. Stalk and Thomas [90]. This forward looking company has more recently pioneered the use of Vendor Managed Inventory and the sharing of detailed demand information with its many suppliers. Many of the aims of time based competitions are also implicit in the aims of lean manufacturing.

The Efficient Consumer Response (ECR) movement first began in the United States grocery industry in 1993. It is now the major policy initiative in food sector supply chain management world wide. The objective is to improve the operation of the grocery supply chain, especially through better co-operation between trading partners. In a recent survey (Institute of Grocery Distribution [97]) the improvements offered through ECR promise to deliver financial benefits worth around 5.7% of consumer price across Europe, which would represent a £4.6bn saving in the UK. It offers a framework for analysing the supply chain to identify opportunities for improvement.

ECR advocates four areas where benefits can be created:

- *Product Replenishment* - Efficient replenishment requires a supply chain that is flexible enough to react quickly to changes in demand. The replenishment of goods is a most fundamental function within the food supply chain, it is one that is not carried out with uniform efficiency. For example, in the US grocery market, the level of inventory within the supply chain often exceeds 100 days whereas in the UK the figure is less than 30 days. The efficiency of non-value-added activities (warehouse operation, physical distribution etc.) varies widely. Synchronous replenishment encapsulates the idea of replenishment "just in time" to meet customer needs.

- *New Product Introduction* - In the groceries market, it has been estimated that over 80% of new products introduced do not obtain a significant long term presence.

- *Trade and Consumer Promotions* - Lack of co-ordination and speed of response to product promotions within the supply chain can lead to mismatched inventory, misdirected incentives and unfulfilled demand.

- *Range and Assortment* - The joint planning of displays and layouts has been shown to boost sales, providing benefits to both manufacturer and retailer. The manufacturer enjoys higher product sales while the retailer has the benefit of increased direct sales and ancillary sales arising from the consumers' presence in the store. Effective forecasting and replenishment are key in this area as displays and layouts can have sudden impact on demand levels. The replenishment system dictates how much store capacity is required for stocking each product and so efficient control will maximise the number and variety of products held and hence offered to the consumer.

The latter three areas are classified as Category Management i.e. demand generation. Most of the research in this report can be related to areas within the ECR framework.

1.5 Supply Chain Performance Measures

Within both the worlds of manufacturing and retailing, a series of measures have been developed to enable a business within a particular industry to benchmark its performance against "best practices". Both financial and non-financial measures of supply chain performance have been developed. The following are relevant to this work.

1.5.1 Non-Financial Performance Measures

Forecast Accuracy - The importance of improved forecasting was emphasised by experienced managers and confirmed by many of the conclusions of this work. This quantifies the accuracy of forecast against actual demand and may, for instance, be defined in terms of the Mean Absolute Percentage Error (MAPE) Makridakis and Wheelwright [89].

On Time In Full - OTIF measures the percentage of orders which arrive on time and in full. If an item in an order arrives on time and in full it is given the value 1. If either it is not on time or not in full it is given the value zero. The scores for each item in the order are then totalled and expressed as a percentage to give an OTIF figure. OTIF is usually applied to a node of the supply chain and can be customer specific e.g. the Unilever Peterborough warehouse may be achieving an average OTIF of 95% for the Tesco Middleton RDC. The cost of stocking out to a retailer is illustrated in recent research carried out in the United States Barry [96] which found that a typical afternoon shopper in the average supermarket will face stock-outs on 8.2% of items in the categories studied. Of the shoppers studied, 20% went elsewhere or did not make the purchase and 20% delayed the purchase. When an alternative was purchased, on average there was a 6% drop in value of the product. The net effect of this was that the retailer lost 46% of the potential sales dollars that would have been spent on those out of stock items.

Lead Time - This measures the time taken between placing an order with a supplier and receiving that order at the customer's warehouse/store. Tesco expect to receive orders from Unilever nominally within 4 working days. The performance measure for lead time is typically expressed as actual against planned lead time.

Stock Holding - The stock holding within a supply chain is often measured in terms of weeks forward cover (wfc). Stock holding is often divided in to a "working stock" and a "safety stock" component. Working stock represents the expected demand over the wfc period. Safety stock represents the stock held to avoid stock out due to random variations from the expected demand. Its level is determined by the customer service level, as specified in terms of expected stock out frequency. Thus, if a 5% stock out frequency is allowed, then the probability of the demand exceeding safety stock plus working stock is .05 etc.

Lorry Loading Efficiency - The average amount of stock loaded on a lorry is divided by the maximum load possible (usually measured in cases) and expressed as a percentage. With some heavier products, loads can be weight restricted.

Pipeline Map – The pipeline map is a useful tool in analysing and diagnosing the supply chain and identifying areas for improvement. The supply chain (internal and external) "is mapped by a series of horizontal and vertical lines. The horizontal lines represent the average time spent in the major processes between stockholding points. The vertical lines represent the average time spent waiting at the stockholding points. Both are drawn on the same scale." Scott and Westbrook (1991). Two of the key measurements used in industry are "pipeline length" which is the sum of the horizontal lines. This represents the time to "pull" a product through the whole supply chain. The other measure is "pipeline volume" which is the sum of both the horizontal and vertical lines. This can be thought of as the time taken to drain the entire supply chain. The purpose of such an analysis is to identify where either the vertical lines or horizontal lines can be reduced.

These performance measures can be applied to products and product groups, individual customers or customer groups or by geographical area depending on the type of analysis being carried out. In addition to these global performance measures, many companies have their own performance measures to keep track of internal process improvements e.g. number of switch-overs per line per week, average switch over time, operational efficiency etc.

1.5.2 Financial Performance Measures

The current trend is towards Activity Based Costing (ABC). This involves breaking down and costing the processes and activities within each supply chain activity. There has been a move towards apportioning components of ABC on a customer basis in order to facilitate customer and geographically based pricing structures. Traditional cost accounting is still used to apportion overheads and other indirect costs.

1.5.3 Supply Chain Benchmarking

Assessing the supply chain involves measuring the performance of many diverse processes. In order to assist managers in this assessment, “benchmarking” approaches have been developed which aim to give "world class" gauges of what is considered "best practise" within the industry. These benchmarks cover many areas including stock levels, distribution, order fulfilment, production switch-over costs/times and customer service level. They measure levels of achievement which have been attained by the best performers in an industry, although they have no mathematical pretensions to being optimal.

1.6 Information Exchange

It is now possible to communicate large amounts of detailed data between two or more parties virtually instantaneously. The formatting of data has been standardised by the EDI (Electronic Data Interchange) protocol. EDI has removed the need for vast amounts of paperwork, duplication and reduced the levels of error in communications. Ordering, invoicing and payment can all be carried out electronically. As the transfer of data is virtually instantaneous, the customer is able to place orders based on up to the minute recorded demand. There are a number of obvious ways in which the sharing of information leads to mutual advantages. These include:

1.6.1 Customer Forecast

Suppliers traditionally have had to forecast the demand for their products based on the existing orders placed by their customers. In effect there was an "information barrier" between parties, leaving the supplier to forecast future demand without any assistance from the customer. Towill (1992). Now it is possible for the customer to communicate forecast information to the supplier, giving details of promotions, uplift estimates and estimated orders. This information allows the supplier to refine long, medium and short term forecasts and plan production accordingly, thus improving the level of service offered to the customer whilst enabling lower stocks to be held at the supplier warehouse. All of these benefits are enhanced by accurate forecasting capabilities. In a recent development, Tesco, Sainsburys and Safeway have begun offering detailed information to their key suppliers. The systems (TIE - Tesco Information Exchange, SID - Sainsburys Information Direct and SSIS - Safeway Supplier Information Service) supply daily sales data at both store and RDC level and their own forecasts of short term stock requirements. This new development enables the supplier to plan more accurately and model the customer ordering process and demand. It is expected that all of the major supermarket chains will be supplying this information within the next few years.

1.6.2 Exception Reporting

Given the most advanced forecasting techniques and reliable data, promotional events often proceed differently than expected. For example, a predicted uplift in demand for a product on promotion can often be exceeded in reality due to unforeseen factors such as weather, competitor activity or no obvious reason. The retailer needs to detect these situations as early as possible to minimise the loss of potential demand. Once a revised estimate of the retailer order requirements can be established, this information can be communicated to the supplier, who can take appropriate action (e.g. adjusting the short term manufacturing schedule to make as much of the promoted product as possible, a typical use of reactive scheduling). Typically, a Tesco promotion lasts four weeks. The existing systems take between one to two weeks to detect and react to exceptional promotional behaviour.

1.6.3 Supplier Shortages

Due to unforeseen circumstances such as ingredient rejection, machine break down or unexpected demand, a supplier may find itself unable to fulfil all of a customer's order. This information needs to be communicated to the customer so that a decision can be made on what to do with the outstanding order quantity e.g. arrange for a substitute product or cancellation, and the customer can prepare accordingly.

1.6.4 Responsiveness

A supply chain is "responsive" if it is capable of reacting quickly to changes in supply chain resources or market demand. Speedy exception detection enables revised information to be passed via EDI along the supply chain.

The role of information technology, forecasting and decision support systems in responsiveness can be illustrated by considering simple scenarios and by referring to Fig. 2. Consider, for instance, the responses that might be required to a significant under estimation of promotional uplift and market demand (T.D.). If no information is shared, the system may be responsive within the retailer's internal supply chain, but upstream response is only initiated when the effects of the increased demand work their way up through the supply chain, possibly causing stock outs. However, if forecasts and other information are shared, the following automatic responses are feasible:

(i) Rapid detection of the demand increase/decrease and automatic generation of an early management exception report.

(ii) Update of till demand forecasts (F_1) and of store replenishment system demand forecasts (C_1).

(iii) Update of RDC demand forecasts (F_2) and the RDC replenishment system demand forecasts (C_2).

(iv) Update of warehouse demand forecasts (F_3) and replenishment forecasts (C_3).

(v) Update of manufacturing schedule (S_1).

(vi) Update of manufacturing raw material forecasts (F_4) and the raw material warehouse replenishment system demand forecasts (C_4).

(vii) Update of raw material manufacturer's demand forecast (F_5) and the raw material manufacturer's replenishment forecasts (C_5).

These forecast updates are model based, and can be rapidly generated along the entire supply chain, provided each partner shares information upstream and downstream. Note that:

The use of forecast information introduces proactive responsiveness. Action can be taken ahead of a potential problem. Forecast based demands can be "smoothed" during scheduling by pulling forecast demands forward in time, thus avoiding stock out.

Each "supplier" benefits mainly from its respective downstream "customer" forecasts. Thus till data demand forecasts $\left(F_1\right)$ are of value at the RDC level, but of less direct uses at warehousing and manufacturing level, due to the effects of replenishment system dynamics.

Scenarios such as the one above may be repeated for other supply chain changes, such as over estimation of demand at any point, ingredient rejection in manufacturing or plant breakdown. These may produce complex supply chain effects as described in Chapter 5, Section 5.2.1 “Reactive Scheduling”. For instance, an ingredient rejection or a facility failure may affect the customer service levels of a range of goods.

1.7 Research Program

The problems of design, management and control of a supply chain network has attracted wide interest from researchers over the last half century. For instance, the problem of design of a supply chain network and of the efficient scheduling of fleets of vehicles within the network has been intensively studied using the concepts of combinatorial optimisation and of integer programming, Christofides (1979); more recent developments have applied the methods of Neural Networks and Genetic Algorithms to this class of problem. Pham (1999). Early in the study of supply chain dynamics, Forrester (1961) described reasons for the observed development of large swings in demand, particularly in the early stages of a supply chain. This work gave rise to a widespread study of this phenomena Towill (1992) and to the formulation of the "law of industrial dynamics".

These researches have been limited to the management control of a given supply chain network and to the development of concepts and technologies which were considered implementable and viable within a modern supply chain.

The discussion so far has been concerned with describing the main activities within a modern supply chain and with outlining key management control objectives within food sector supply chains. In order to progress further, a program of detailed discussions was undertaken, over a period of several years, with supply chain experts and managers from industry. The objectives agreed in these discussions formed the basis of the entire programme of work, which primarily aims to be of practical application. As stated earlier, it was agreed to consider the design of a system concerned with dry soup supply chain activity. This is thought to be representative of the Fast Moving Consumer Goods (FMCG) sector.

Each of the partner's supply chain computer control systems involves investments of millions of pounds. It is, of course, highly desirable that any results from this work are implementable within the environment of these

systems, and are reasonably comprehensible to managers and informed users. These "non-invasive" requirements limit the viable range of options for the programme. For instance, inventory control systems based on classical theory, such as the two-bin system or the fixed cycle system Battersby (1970) are broadly acceptable within modern computer control systems. In contrast, methods for instance based on sophisticated optimisation concepts might be more problematic.

The main themes of research are as follows:

1.7.1 Forecasting

During discussion, retailers, manufacturers and suppliers all placed high priority on the desirability of more accurate forecasts. Supply chain partners utilise both long and short term forecasts. The long-term forecasts (typically six to eighteen months) are used for strategic planning. Medium and short-term forecasts (one day to several weeks) are used in supply chain management, scheduling and everyday running of the business. The methods developed here were found to be effective for strategic, medium and short term forecasting.

Initial investigation showed that demands at all points along a supply chain contained significant "systematic" components which are dependent on various indicators such as day, season, promotional period etc. In Chapter 2, a model form was therefore proposed that takes account of all these factors. It is multiplicative in structure, which is in agreement with the language and concepts used by experienced supply chain managers. Bowersox and Closs [96]. As all supply chain data is count data, a Poisson like distribution was proposed. The resultant model does not have the linear plus normal noise structure frequently assumed in the management control literature. The model form turns out to have a number of potential advantages. For instance, the same form fits data along the entire supply chain even though the range of demand varies by many orders. The

Poisson-like distribution assumption provides simple variance models that can be exploited in the design of efficient inventory control systems.

The main sources of forecasting are illustrated in Fig. 2. Of primary interest is the Till Data (TD) demand signal at the retail outlet, which drives the entire supply chain. In the supply chain considered, there were approximately two hundred stores each generating demand signals for many thousands of products. Hence there was a need for an efficient on-line forecasting procedure. As shown in Fig. 2, the next forecast (F_2) is concerned with the total demand on the Regional Distributions Centre (RDC). This is the first example of an "internally" generated demand signal. The next main external demand comes from the Manufacturing Demand (M.D.) on the manufacturing warehouse. This is quite distinct from the RDC demand source since it represents the entire demand for the warehouse product from all retail outlets served. As shown in Fig. 2 the manufacturing activity then generates a demand for all supplier products. This is the next example of "internally" generated demand, which must be forecast for efficient running of the supply chain. These patterns then repeat down the supply chain.

In Chapter 2 we consider the off-line modelling of till data and RDC demand data using the Weighted Least Squares approach of Singular Value Decomposition (SVD). In chapter 3 we consider the development of real time algorithms for till and RDC data, based on Maximum Likelihood Estimation (MLE) approaches. Algorithms that are adaptive, efficient in terms of real-time use and able to handle the large volumes of data are evolved.

1.7.2 Inventory management control

As shown in Fig. 2, almost the entire supply chain can be viewed as a network of inventory replenishment systems. The design of such systems is considered in Chapters 4 and 5. For reasons stated earlier, the methods used are restricted to

those of classical control theory and are therefore implementable for many modern systems. In chapter 4 we consider the effects of flow constraints on inventory and use the results of Chapters 2 and 3 to develop stock control procedures. It is shown that safety stock forms a substantial component of the minimal attainable stock level and is highly dependent on forecast accuracy. Chapter 4 concludes with development of a multi-product inventory control approach which attempts to attain satisfaction of the conflicting requirements of "lean" operations and high standards of customer service. A substantial implementation of this scheme together with results and assessments of performance are described in Chapter 5. This involved the use of recorded supply chain data over a period of 6 months and of a commercial modelling and scheduling package. The results indicate that substantial reductions in stock level are possible.

1.7.3 Distribution

As shown in Fig. 2, a supply chain involves many distribution operations. This is a potential source of environmental pollution and inefficiency – it is estimated that 40% of the lorries on our road are empty and for the supply chain studied, lorry load efficiency is typically in the region of 75%. Chapter 6 consists of a case study in which short term forecasting is used to attain near 100% lorry load efficiency. The approach is of particular interest when, as is increasingly the case, retailers insist on ever decreasing stock levels and re-order quantities. This can certainly produce inefficiencies both in warehousing and in distribution. Potential inefficiencies in distribution can frequently be avoided by the use of the forecast-based approaches described. These involve relaxation of the OTIF constraints traditionally imposed on suppliers. The substantial transportation cost savings resulting from the use of the methods described are particularly easy

to realise if the partners adopt a Vendor Managed Inventory approach to supply chain management.

1.7.4 Claims to Originality

The model defined in Chapter 2 uses a multiplicative structure to describe the systematic effects and assumes a Poisson-like distribution. The explanatory variables selected were day of the week, week of the year, promotional period and weather effects. This model reflects in detail the views of experienced managers at both Tesco and Unilever and turned out to be very general, the same model structure being applicable to all nodes of the supply chain. So far as is known this model structure is original within this field of application.

The application of Singular Value Decompostion Approaches to the modelling of Supply Chain demand and its extension to Generalised Least Squares (GLS) is thought to be original.

In Chapter 3, the application of Maximum Likelihood methods for evaluation of the model parameters gives identical solutions to those obtained using GLS. The equations have relatively simple explicit solutions and turn out to provide a basis for the development of an efficient on-line forecasting algorithm. It is also to introduce adaptive control concepts in to the on-line algorithm in order to obtain rapid exception reporting and promotional monitoring. These features are important in the design of efficient Supply Chain systems. The developments are thought to be original.

In Chapter 4 the design of replenishment systems is considered. The research concentrates on fixed cycle systems since these are so widely used in practice. It is shown that the effects of batch flow constraints produce a limit cycle effect which has important implications on supply chain performance. A formula is developed for estimating the effects of key parameters on attainable minimum stock levels. This can be used as a benchmarking tool and as an aid to tuning key supply chain parameters. Finally in Chapter 4 we point out that the much discussed "Synchronous Replenishment" features aimed at in modern supply

chain management, requires that the operations of warehouse replenishment and manufacturing scheduling must be integrated and propose a scheme for realising this. The developments outlined above are each thought to be original.

In Chapter 5, the attainment of "Synchronous Replenishment" is demonstrated by the development and assessment of a Manufacturing Decision Support System (MDSS). This uses a commercial scheduling package based on discrete event simulation methods to model the manufacturing and warehousing processes. The results demonstrate that significant gains in Supply Chain performance are attainable by the correct choice of key operational parameters such as weeks forward cover, safety stock parameters, Minimum run length and Incremental quantities as described in Chapter 5. The design of the system and its assessment using six months recorded data is thought to be original.

In Chapter 6, demand forecasts and OTIF relaxation are used to pull forward orders to achieve improved lorry load efficiencies and reduced deliveries with a minimum increase in stock residency. These concepts, which are demonstrated in a case study are thought to be original.

References

Ballon, R.H. 1999. *Business Logistics.* Prentice Hall

Barry, A. 1996. *Steeling The Show.* Supply Management, April

Battersby, A. 1970. *A Guide to Stock Control.* 2nd Edition p75 – 96 Pitman Publishing.

Bowersox, D.J. and Closs, D.J. 1996. *Logistics Management.* McGraw Hill

Christofides, N. 1979. *Combinatorial Optimisation.* Wiley.

Christopher, M. 1998. *Logistics and Supply Chain Management.* Financial Times. Pitman Publishing

Forrester, J. W. 1961. *Industrial Dynamics*. MIT Press.

Holmes, G. 1995 *Supply Chain Management.* The Economist Intelligence Unit

Institute of Grocery Distribution. 1997. *Efficient Consumer Response*

Johnson, G. and Scholes, K. 1988. *Exploring Corporate Strategy.* Prentice Hall

Makridakis, S. and Wheelwright, S.C. 1989. *Forecasting Methods for Management.* John Wiley and Sons

Pham, D.T. and Karaboga, D. 1999. *Intelligent Optimisation Techniques:Genetic Algorithms, Tabu Search, Simulated Annealing and Neural Networks.* Springer-Verlag

Prodham, B. 1993. *Financial Perspectives in the Supply Chain.* Templeton College Management Research Paper

Saunders, M. 1994. *Strategic Purchasing and Supply Chain Management.* Pitman

Schonberger, R.J. 1982. *Japanese Techniques. Nine Hidden Lessons in Simplicity.* The Free Press

Scott, C. and Westbrook, R. 1991. *New Strategic Tools for Supply Chain Management*. Int. Journal of Physical Distn & Logistics Management, Vol. 21, No 1, p23-33.

Stalk, G. and Thomas, M.H. 1990. *Competing Against Time.* The Free Press

Towill, D.R. 1992. *Supply Chain Dynamics – the change engineering challenge of the mid 1990s*. Proceeding of the Institution of Mechanical Engineers. Part B. Journal of Engineering Manufacture. Vol 206. No B4. p235-238.

Voss C.A. 1992. *Manufacturing Strategy*. Chapman Hall

2.0 DEMAND MODELLING AND GENERALISED LEAST SQUARES APPROACHES

2.1 Introduction

The development of a model based forecasting procedure typically involves three stages, namely:

(i) Selection of model class.

(ii) Algorithmic selection and assessment.

(iii) Assessment of forecasting capability.

Box and Jenkins [76]. In the next two chapters we will be concerned with all three stages.

Application of model based forecasting to manufacturing and business management began in the 1950's when the number of applications grew rapidly, triggered off by two significant developments. The first was the arrival of the digital computer, and the second, the development of a range of exponential smoothing and forecasting techniques. These smoothing techniques, originally developed by operational researchers, have found increasing application in inventory control and short term demand forecasting.

The established popularity of exponential smoothing approaches can be attributed to their transparent simplicity and their computational efficiency. The need for computational efficiency is especially evident in forecasting for a modern retail supply chain, which involves data storage and daily forecasting of tens of thousands of products for hundreds of stores.

Exponential smoothing is based on the assumption that the process can be modelled as the additive sum of a fixed mean component, and a random Normally distributed component. In its simplest form, the mean is assumed equal to the exponentially weighted average of past observations. Brown [59], who pioneered the application of exponential smoothing to stock control and supply chain management, modified these assumptions and introduced an additional smoothing equation modelling trend in the mean level components Holt [57] and Winters [60] improved Browns formulation by introducing separate smoothing factors for mean, trend and season estimation. Makridakis [89].

The pioneering work of Box and Jenkins [76] provided the first systematic framework for analysis and identification of the general ARIMA class of models. This was followed by development of a range of powerful modelling and forecasting procedures including Kalman filtering. These developments allowed exponential smoothing approaches to be identified within a more formal mathematical framework. A survey of exponential smoothing and links with ARIMA modelling is given by Gardner [85].

The developments outlined above usually assume a process model consisting of the sum of a mean component plus a Normally distributed random component. The assumption of Normality may be problematic when applied to count data, which is necessarily positive and whose distribution is often well approximated by a Poisson rather than Normal distribution. The Normality assumption may be reasonable if the count is sufficiently high. Even then the assumption of constant variance is inappropriate if the mean count level varies.

2.2 Model Formulation.

Supply chain demand data consists of a series of counts and hence the assumption of a Poisson distribution is reasonable. The mean count level, however, typically displays pronounced systematic patterns or 'modes' of

variation and these are sometimes represented multiplicatively in the forecast model Bowersox and Closs [96]. These modes are usually understood to be functions of certain 'explanatory' variables. The following modes are generally perceived as important.

(i) Daily Mode (D). Customer buying frequently varies in level with day of the week, and therefore day of the week is a natural 'explanatory' variable. Variation with hour of the day etc. may certainly occur but daily patterns are usually of main practical interest.

(ii) Seasonal Mode (S). Demand usually varies in a systematic way with the yearly calendar, with peaks at Easter, Christmas and other public holidays. A convenient 'explanatory' variable for these changes is week of the year.

(iii) Promotional Mode (P). Promotional activities such as special offers, price reductions and advertising campaigns each produce a demand 'uplift' over the period of the promotion . The duration of the promotion in the food sector is typically in the order of four weeks. A convenient 'explanatory' variable is time from the start of the promotion. The time may be scaled in weeks, or in days depending on the data and the application.

(iv) Weather Modes (W). Changes in weather parameters such as temperature, rainfall or hours of sunshine are thought to systematically affect supply chain demand level, and represent candidates as 'explanatory' variables.

It should be noted that weather parameters have considerable seasonal dependence, hence we would anticipate that seasonal (S) and weather (W) modes to be highly correlated. If, as is frequently the case in practice, weather information is not to hand, then inclusion of weather related seasonal modes is not possible. As we will show, modelling without weather related modes

frequently yields very useful forecasts. It would appear that weather related modes may only be significant when prolonged extreme weather conditions develop.

As is well known, weather changes are chaotic and notoriously difficult to forecast in the UK except in the very short term. Although longer term weather forecasts are unreliable, the modelling of weather modes is still of practical value, for instance in the planning and management of contingency stock and capacity.

(v) Event Detection Modes (E). Demand level changes may occur along a supply chain without warning or known explanation. For instance, promotional uplift may differ significantly from plan or competitor promotion may cause an unexpected down turn in demand. Such events are by their natures, unpredictable. It is, however, of practical importance that significant level changes in demand are detected and estimated without delay. Supply chain conditions may then be re-adjusted to avoid high stock levels developing, especially of limited life and promotional products, or under-stocking and loss of sales and customer service.

The modes listed above are thought to reflect most of the effects of concern in supply chain planning and management. For some applications, particular modes are important. For instance, in long term strategic planning, the seasonal mode (S) is of more concern, whereas in short term supply chain management, the daily mode (D) and the promotional mode (P) are important.

All of the modes are usually present in till data. As we will demonstrate, under certain conditions, these effects may transmit 'upstream' within the supply chain with little attenuation, and hence the same model structures suffice for modelling retail regional distribution centre (RDC) demand, and even manufacturing demand.

Consider the modelling of daily demand for a single product at store level. Assuming the store is open I days per week and that sequential demand data collected for J weeks has been arranged in an $I \times J$ array Y say, where y_{ij} equals the recorded demand count on day i of week j.

If a Poisson distribution form is assumed for each demand, then

$$\Pr(y_{ij} = k) = \frac{e^{-\mu_{ij}} \mu_{ij}^{\ k}}{k!} \qquad 1.0$$

for $1 \le i \le I, 1 \le j \le J$ and where $\mu_{ij} = E(y_{ij})$ and $k = 0,1,2..$ Hastings and Peacock [75].

It is convenient to assume a multiplicative structure for the systematic modes discussed above as representative of demand behaviour. Bowersox and Closs [96]. Assuming the observations are dependently distributed, then a possible form is

$$\mu_{ij} = \lambda D(i) S(j) P(j) W(T_{ij}) \qquad 2.0$$

where the mode functions D, S, P and W are functions of the explanatory variables day (i), week (j) and weather variable T_{ij}, which might represent, for instance, the temperature on day i of week j. Additional weather effects may be introduced through allowing further functions.

In order to uniquely define the model we must define the functions D, S, P and W appropriately. One convenient procedure is to introduce norm constraints $\|D\| = \|S\| = \|P\| = \|W\| = 1$ say, where the norm, which may be selected for algorithmic convenience. The variable λ then represents the level of demand prior to modulation by the mode effects.

The weather function W, a continuous function of the explanatory weather variable T_{ij}, may be approximated by a discrete form, for algorithmic purposes. If for instance T_{ij} represented temperature on day i of week j, and (T_0, T_N) represents the feasible range of temperature, and $T_0 < T_1 < T_2 ... < T_N$, then we may define

$$W(T_{ij}) = w_k \qquad 3.0$$

if $T_{k-1} < T_{ij} \leq T_k$ say for $k = 1...N$

It will be convenient to simplify the notation by defining

$$D(i) = d_i \qquad i = 1...I$$

$$S(j) = s_j$$

$$P(j) = p_j \qquad j = 1...J$$

and re-expressing equation 2.0 as

$$\mu_{ij} = \lambda d_i s_j p_j w_k \qquad 4.0$$

where as described above w_k is determined by the value of the weather explanatory variable on day i of week j.

Equation 2.0 models the mode effects as a simple product. We can also consider a more general form.

$$\mu_{ij} = \sum_{k=1}^{r} \lambda_k D_k(i) S_k(j) P_k(j) \qquad 5.0$$

which represents the sum of r mode products. The Singular Value Decomposition approach developed later in this chapter can be used to fit this more general form.

When selecting among different models for the same data, the Akaike Information Criteria (AIC) offers a convenient and widely used measure of model effectiveness. This may be expressed in the form

$$AIC = D + 2P$$

where D is the appropriate deviance and P the number of model parameters. Lindsey [95]. The model with the smallest AIC value is selected.

2.2.1 Algorithmic Selection.

Each of the models developed above are heteroscedastic, since as is well known for a Poisson distribution

$$Var(Y_{ij}) = \mu_{ij}$$

and hence the variance varies systematically. The two most widely used approaches to fitting this class of model, are based firstly on Generalised Least Squares (GLS) methods, and secondly, the Maximum Likelihood approach, Carroll and Ruppert [88]. The GLS approach developed in this chapter employs Singular Value Decomposition (SVD) concepts. These are now widely used in signal processing. The development and assessment of SVD based approaches to supply chain demand modelling is considered later in this chapter. In the next chapter we consider the direct calculation of the Maximum Likelihood parameters for the model.

2.2.2 Overdispersion.

For the ideal Poisson distribution, the mean equals the variance. This condition is not usually satisfied in practice. McCullagh and Nelder [84]. Even in closely controlled measurements of particle arrival rate of a Geiger counter, the counts exhibit higher variance than the pure Poisson model would predict. In this case 'overdispersion' can be traced to small 'dead time' effects, present in the Geiger counter, following a particle arrival. Over dispersion is common in Poisson-like modelling, although the mechanism generating this effect is frequently not clear and may have a number of possible causes. McCullugh and Nelder [84], Lindsey [95]. Overdispersion usually occurs in supply chain demand data and hence the modelling of overdispersion is of considerable importance in stock control and in the maintenance of customer service. In the absence of any evident mechanism producing overdispersion, a frequently invoked assumption is simply that

$$Var(Y_{ij}) = \theta^2 \mu_{ij} \qquad 6.0$$

where $\theta > 1$. A moment estimator for θ^2 is

$$\hat{\theta}^2 = X^2/(IJ - P) \qquad 7.0$$

where

$$X^2 = \sum_{i=1}^{I}\sum_{j=1}^{J}(y_{ij} - \hat{\mu}_{ij})^2 \Big/ \hat{\mu}_{ij} \qquad 8.0$$

and p is the number of parameters contained in the model McCullugh and Nelder [84]. The statistic X^2 represents a weighted least squares error measure appropriate to the pure Poisson process, and equals the Pearson statistic.

2.3 Least Squares Approaches

We first develop a least squares fitting procedure for the models of equations 5.0 and 2.0 defined earlier, and based upon a singular value decomposition of the data array matrix. This least squares fitting procedure is then used within a weighted least squares routine to obtain optimal fits. As is well known, for independent observations, maximum likelihood and weighted least squares estimates are equivalent for distributions in the exponential family and in particular the Poisson distribution. Bradley [76], Carroll and Rupert [88].

2.4 Singular Value Decompostion

Singular value decomposition (SVD) based methods have been widely applied in signal processing and numerical analysis Datta [95]. An early application to consumer demand modelling is described by Farmer and Potton [66] who were concerned with forecasting demand on the UK electricity grid. A similar application was considered by Matthewman and Nicholson [68].

Consider the modelling of demand count data, recorded for I days per week and over a period of J weeks, and arranged in the form of an $I \times J$ array Y say, where y_{ij} represents the count on day i of week j. The singular value decomposition of the array may be expressed in the form

$$Y = U\begin{bmatrix} \Lambda & 0 \\ 0 & 0 \end{bmatrix} V^T \qquad 9.0$$

where Λ is diagonal with positive entries $\sigma_1 \geq \sigma_2 \ldots \geq \sigma_\ell > 0$, the singular values of Y, and U and V are orthogonal matrices of dimension I and J respectively. The matrix Y is of rank $\ell \leq Min(I, J)$, Stewart [73]. For simplicity we represent equation 9.0 as $Y = U\Lambda V^T$, it being understood that Λ may have zero diagonals, (the zero singular values). If $U_i (V_j)$ denotes the

$i^{th}(j^{th})$ columns of $U(V)$, then equation 2.0 may be conveniently expressed in the dyadic form

$$Y = \sum_{k=1}^{\ell} \sigma_K U_K V_K^T \qquad 10.0$$

The following celebrated result on dyadic approximation has found wide application

2.4.1 RESULT 1 (Eckart and Young [39]).

If the first $r \leq \ell$ dyadic terms only in equation 10.0 are retained, then the resultant matrix

$$X^r = \sum_{k=1}^{r} \sigma_k U_k V_k^T \qquad 11.0$$

is the best rank r least squares approximation to Y in the sense that

$$\sum_i \sum_j (y_{ij} - x_{ij})^2$$

is minimised over all matrices X of rank r or less.

To illustrate the application of Result 1, consider the least squares fitting of equation 2.0. For simplicity assume that no promotional effects are present in the data contained in array Y, and hence that $P(j) = 1$ for $1 \leq j \leq J$. Also that only seasonal effects are taken into account and that $W(T_{ij}) \equiv 1$. Hence the systematic mode model of equation 2.0 reduces to

$$\begin{aligned} \mu_{ij} &= \lambda\, D(i)\, S(j) \\ &= \lambda\, d_i\, s_j \end{aligned} \qquad 12.0$$

From Result 1, the optimum rank 1 least squares fit of the array Y may be expressed

$$X^1 = \sigma_1 U_1 V_1^T \qquad 13.0$$

where $\|U_1\|_2 = \|V_1\|_2 = 1$ since U and V are orthogonal. If in equation 12.0 we select the scaling of D and S so that $\|D\|_2 = \|S\|_2 = 1$ then Result 1 shows that a best least squares model fit may be obtained setting

$$\begin{aligned} d_i &= U_1(i) \qquad 1 \le i \le I \\ s_j &= V_1(j) \qquad 1 \le j \le J \qquad 14.0 \\ \textit{and} \quad \lambda &= \sigma_1 \end{aligned}$$

and hence the components of vector $U_1(V_1)$ provide optimum discrete values of function $D(S)$ and the parameter λ in equation 12.0 may be identified with the dominant singular value of the data array matrix Y.

If promotional effects are present in the data arranged in Matrix Y, then equation 5.0 becomes

$$\mu_{ij} = \lambda d_i s_j p_j \qquad 15.0$$

say, where $p_j = 1$ outside the range of any promotional effects. The procedure described above for non promotional data may still be applied, except that in equation 14.0 we make the identification

$$s_j p_j = V_1(j) \qquad 16.0$$

for all values of j influenced by promotional effects. If, as is frequently the case, the seasonal term s_j may be assumed constant through the period of the promotion, then p_j may be estimated by simple division.

Inclusion of the weather mode $W(T_{ij})$ in equation 2.0 is problematic using SVD approaches. The inclusion of this mode introduces a third independent variable and there is no SVD analysis applicable to three dimensional arrays.

The more complex model of systematic effects, given by equation 5.0 above may also be fitted in the optimal least squares sense, by appeal to Result 1. Assuming no promotional effects present in the data then follows similar reasoning to that used for the rank 1 case above, we can easily show that the best least square fit is given by identifying the parameters in equation 12.0 with the SVD terms as follows

$$\begin{aligned} d_i^{\,k} &= U_k(i) \qquad 1 \le i \le I \\ s_j^{\,k} &= V_k(j) \qquad 1 \le j \le J \\ \text{and} \quad \lambda^k &= \sigma_k \end{aligned} \tag{17.0}$$

for $k = 1...r$. Hence the components of vector $U_k(V_k)$ provide optimum sample values of the mode functions $D^k(S^k)$ and the scalars λ^k equal the kth singular value of the data array Y. If promotional effects are present these may be included by modifying equation 10.0 to

$$s_j^{\,k} p_j^{\,k} = V^k(j)$$

for all week j influenced by promotional effects.

2.4.2 Generalised Least Squares Approaches

The development so far has been concerned with least squares fitting criteria. For a weighted least squares criteria of the form

$$\chi^2 = \sum_{i=1}^{I}\sum_{j=1}^{J}\left(y_{ij} - \mu_{ij}\right)^2 / c_{ij}$$

the weighting for a Poisson model is $c_{ij} = \mu_{ij}$. Carroll and Rupert [88]. This criteria can be expressed

$$\chi^2 = \sum_{i=1}^{I}\sum_{j=1}^{J}\left(\frac{y_{ij}}{\sqrt{c_{ij}}} - \frac{\mu_{ij}}{\sqrt{c_{ij}}}\right)^2$$

$$= \sum_{i=1}^{I}\sum_{j=1}^{J}\left(\bar{y}_{ij} - \bar{x}_{ij}\right)^2$$

where $\bar{y}_{ij}$ is a weighted observation, and $\bar{x}_{ij}$ is of the form of equation 4.0. provided μ_{ij} and c_{ij} are. This suggests a generalised least squares procedure, Carroll and Rupert [88], of the form

Algorithm (conceptual)

(i) Calculate the weighted data array $\bar{Y}$ where

$$\bar{y}_{ij} = y_{ij} / \sqrt{c_{ij}^{\,o}}$$

where $c_{ij}^{\,o}$ is an initial estimate of μ_{ij} $(1 \le i \le I; 1 \le j \le J)$

(ii) Calculate the least squares approximation X say, to Y using Result 1

(iii) Set $c_{ij} = \bar{x}_{ij}\sqrt{c_{ij}^{\,o}}$

BIRL
LONDON
UNIV.

(iv) If $\left\| c_{ij} - c_{ij}^{\,o} \right\|$ is small for $1 \le i \le I . 1 \le j \le J$. set $\mu_{ij} = c_{ij}$ and stop. Otherwise set $c_{ij}^{\,o} = c_{ij}$ and return to step (i).

The solution to the un-weighted least squares problem usually provides an excellent initial estimate $c_{ij}^{\,o}$ to μ_{ij}. There is no consensus on the maximum number of iterations to be used in practice. Matloff (84) proposes only one, Williams (59) only two and Carrol and Rupert (88) at least two. We have found that two cycles usually gives adequate accuracy.

2.5 Results and Discussions

We consider results obtained from fitting both till and RDC. demand data using Weighted Least Squares based approaches and discuss model rank selection. As shown later, under certain conditions, modes present at the till demand level transmit upstream in the supply chain, with remarkably little attenuation but with increased dispersion. This effect will be demonstrated using RDC demand data.

2.5.1 Till Data Results

Fig. 1 illustrates demand level, for a 7 day week and over a period of 20 weeks. No promotion is present.

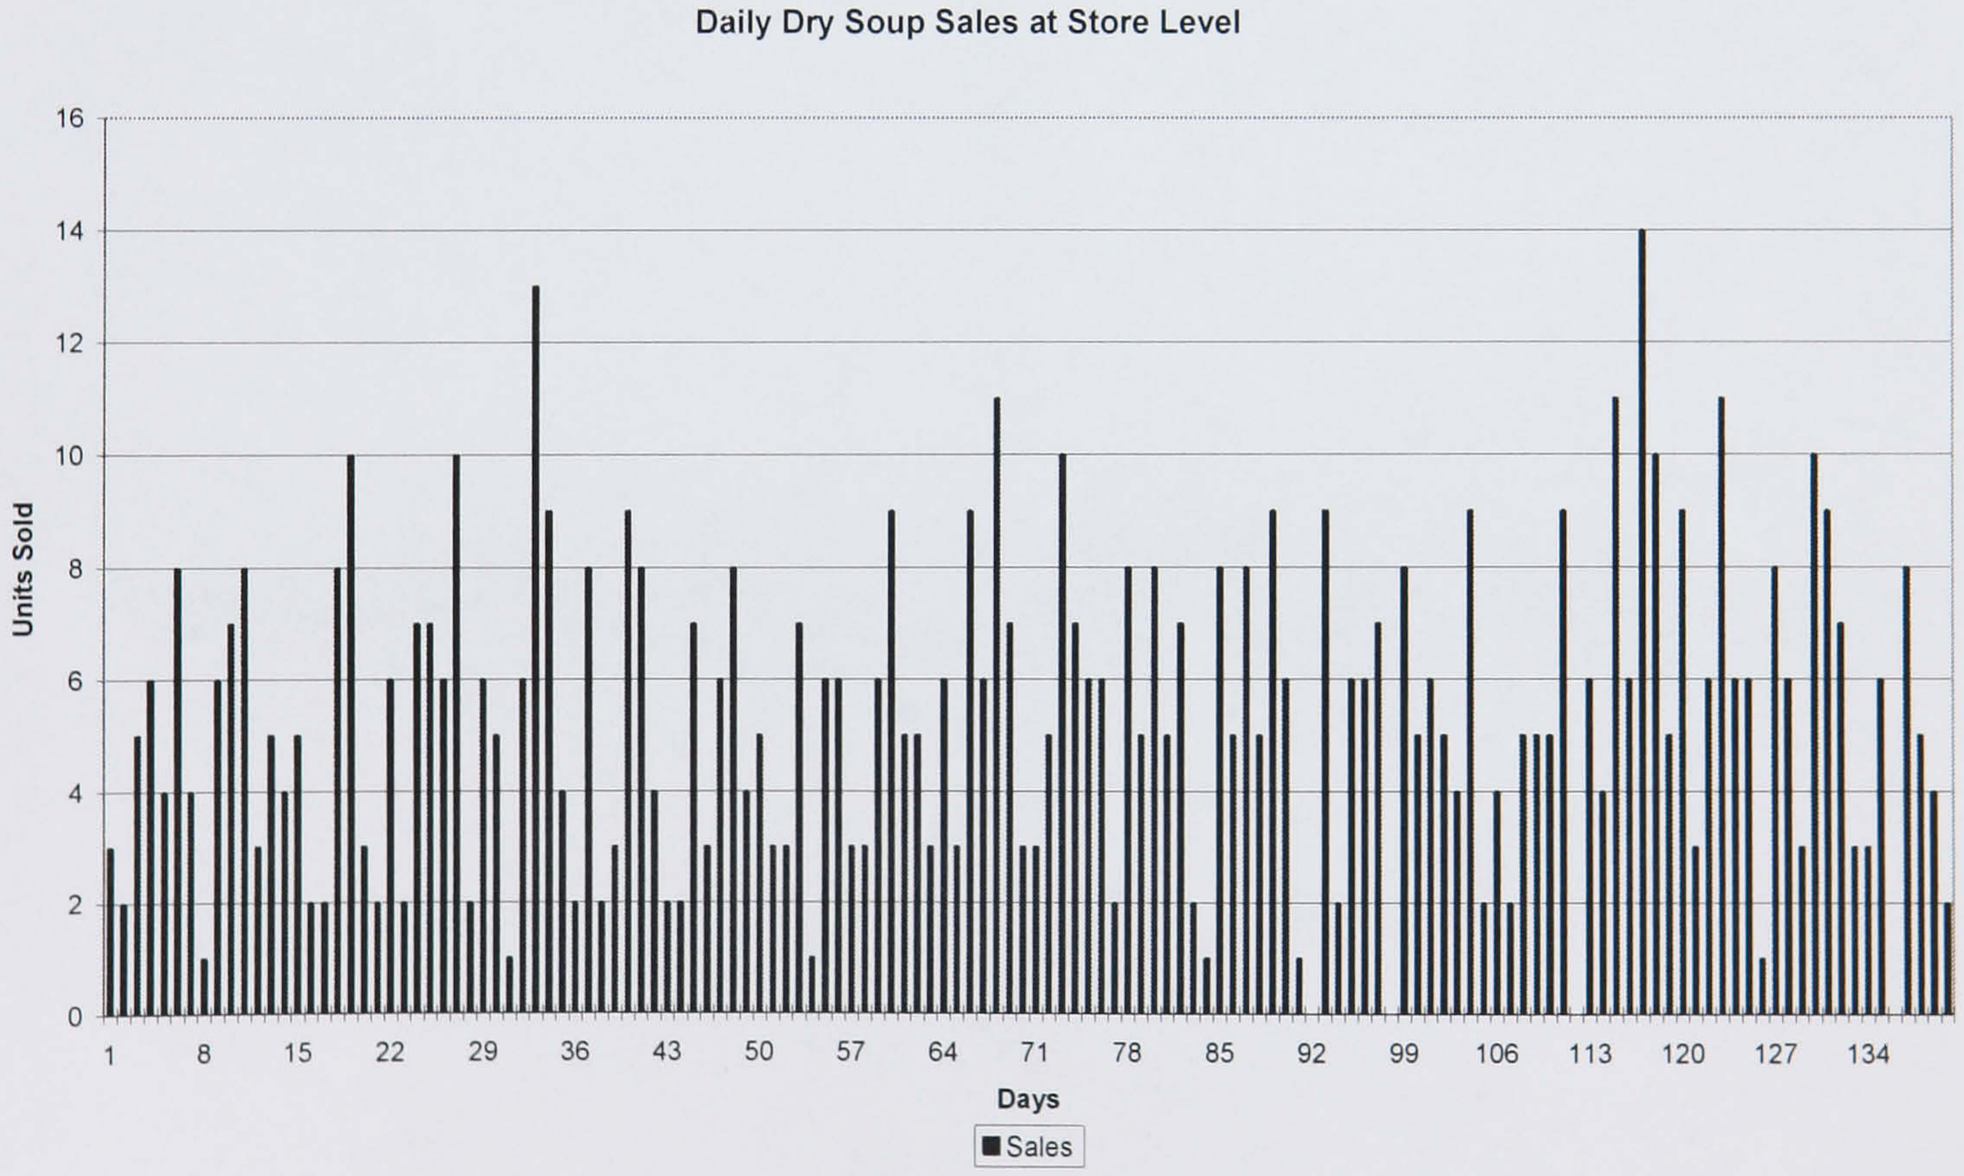


Fig. 1

A singular value analysis of the data array yielded:

$\lambda_1 =$ 66.66 $\lambda_2 =$ 13.70 $\lambda_3 =$ 12.60

$\lambda_4 =$ 11.80 $\lambda_5 =$ 9.37 $\lambda_6 =$ 7.88

$\lambda_7 =$ 6.15

The AIC values for the rank 1 (rank 2) models were 166 (194). This strongly suggests choice of the rank 1 model. The error autocorrelation for a rank 1 model is shown in Fig. 2. This suggests there is little information in the error, which is almost all within the 95% confidence interval band, and that the error is of near white noise.

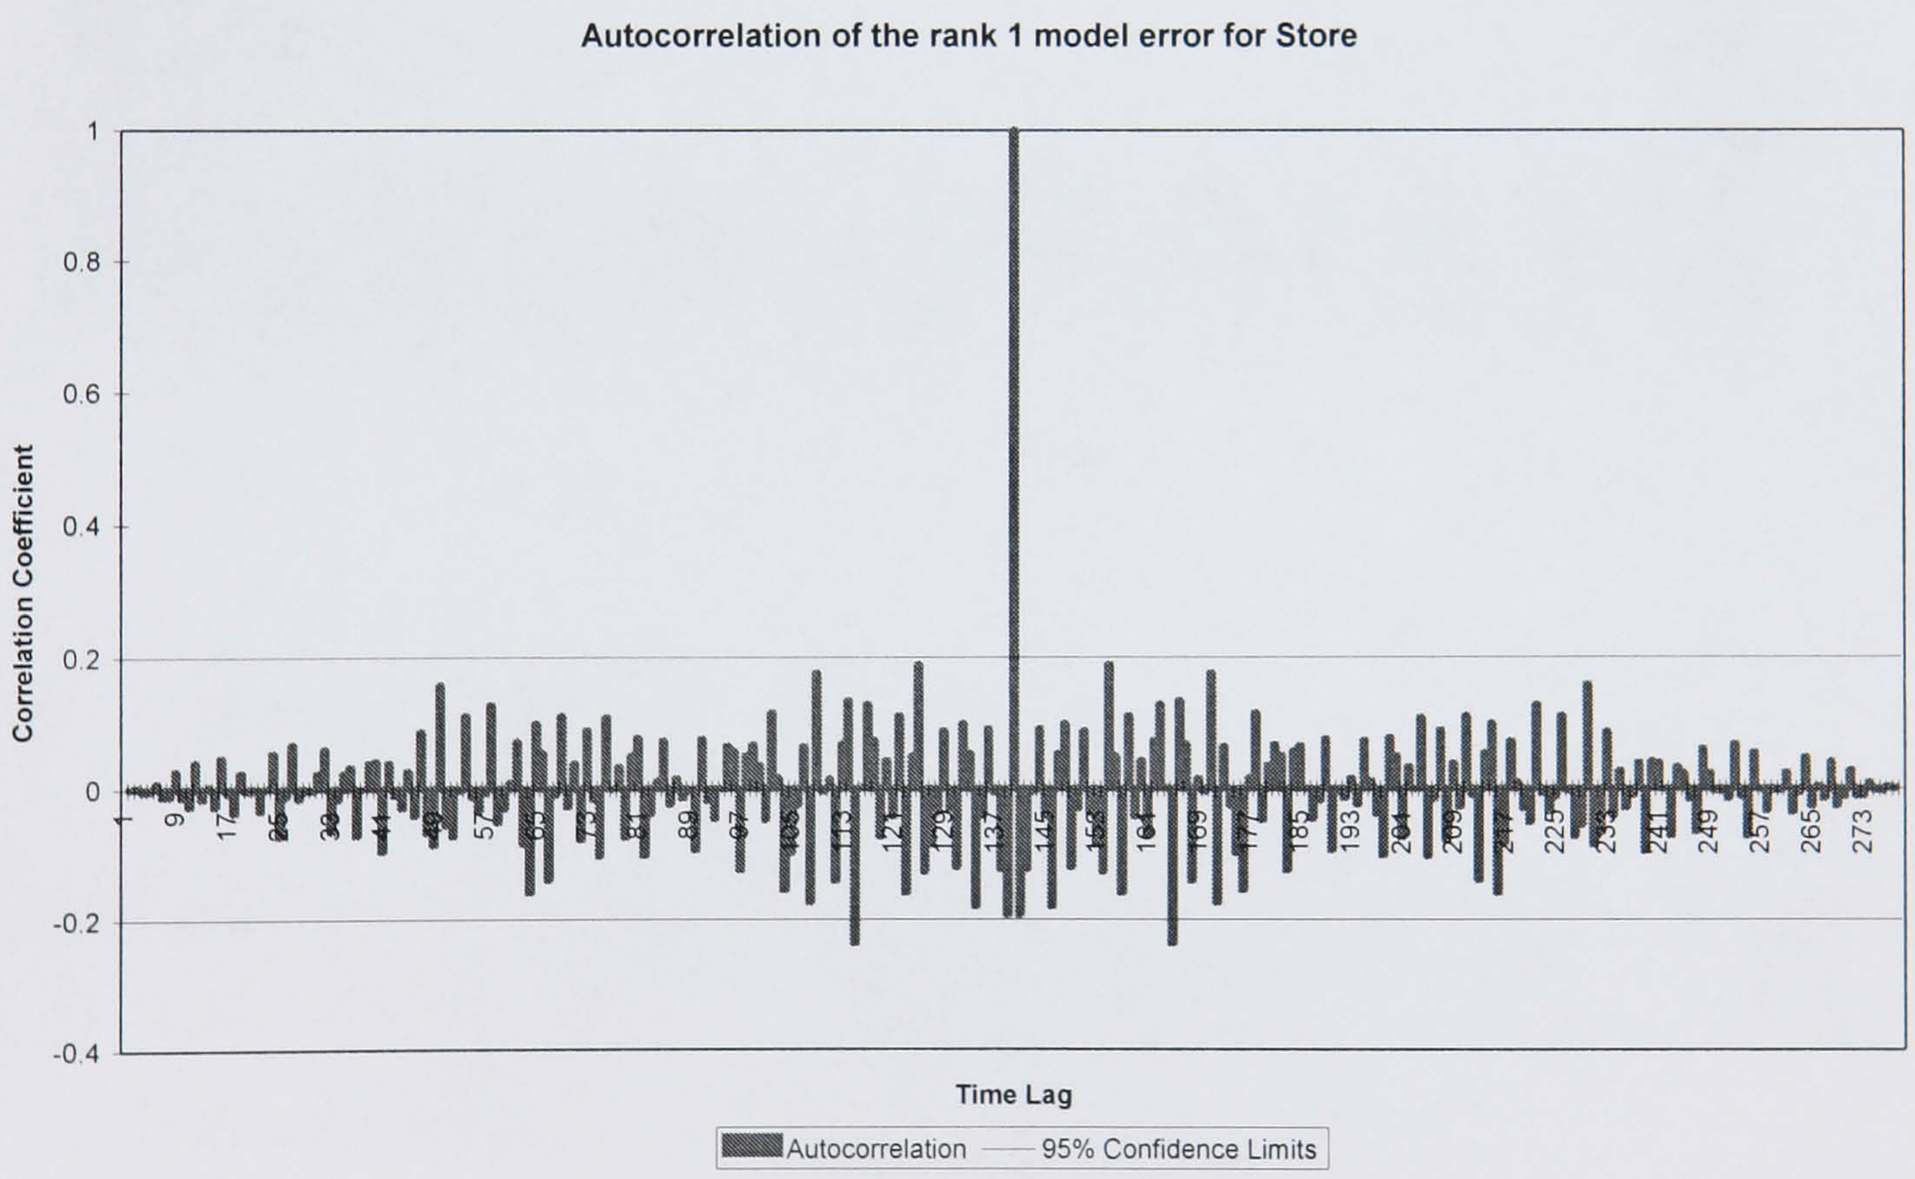


Fig. 2

There is one autocorrelation value which lies outside the 95% confidence limits. This corresponds to a lag of 25 days. Although much is known about the operations at store level (replenishment periods, product minimum order quantities, case sizes etc.), no rational explanation for this correlation has been

found. We note that its magnitude is only just outside the bounds for white noise. As can be seen from the Normal Q-Q Plot in Fig. 3, the weighted residuals appear Normally Distributed (Mills [96]), suggesting that the assumption of Poisson like error distribution is reasonable. The overdispersion parameter, fitted according to equation 6.0, was $\theta^2 = 1.27$.

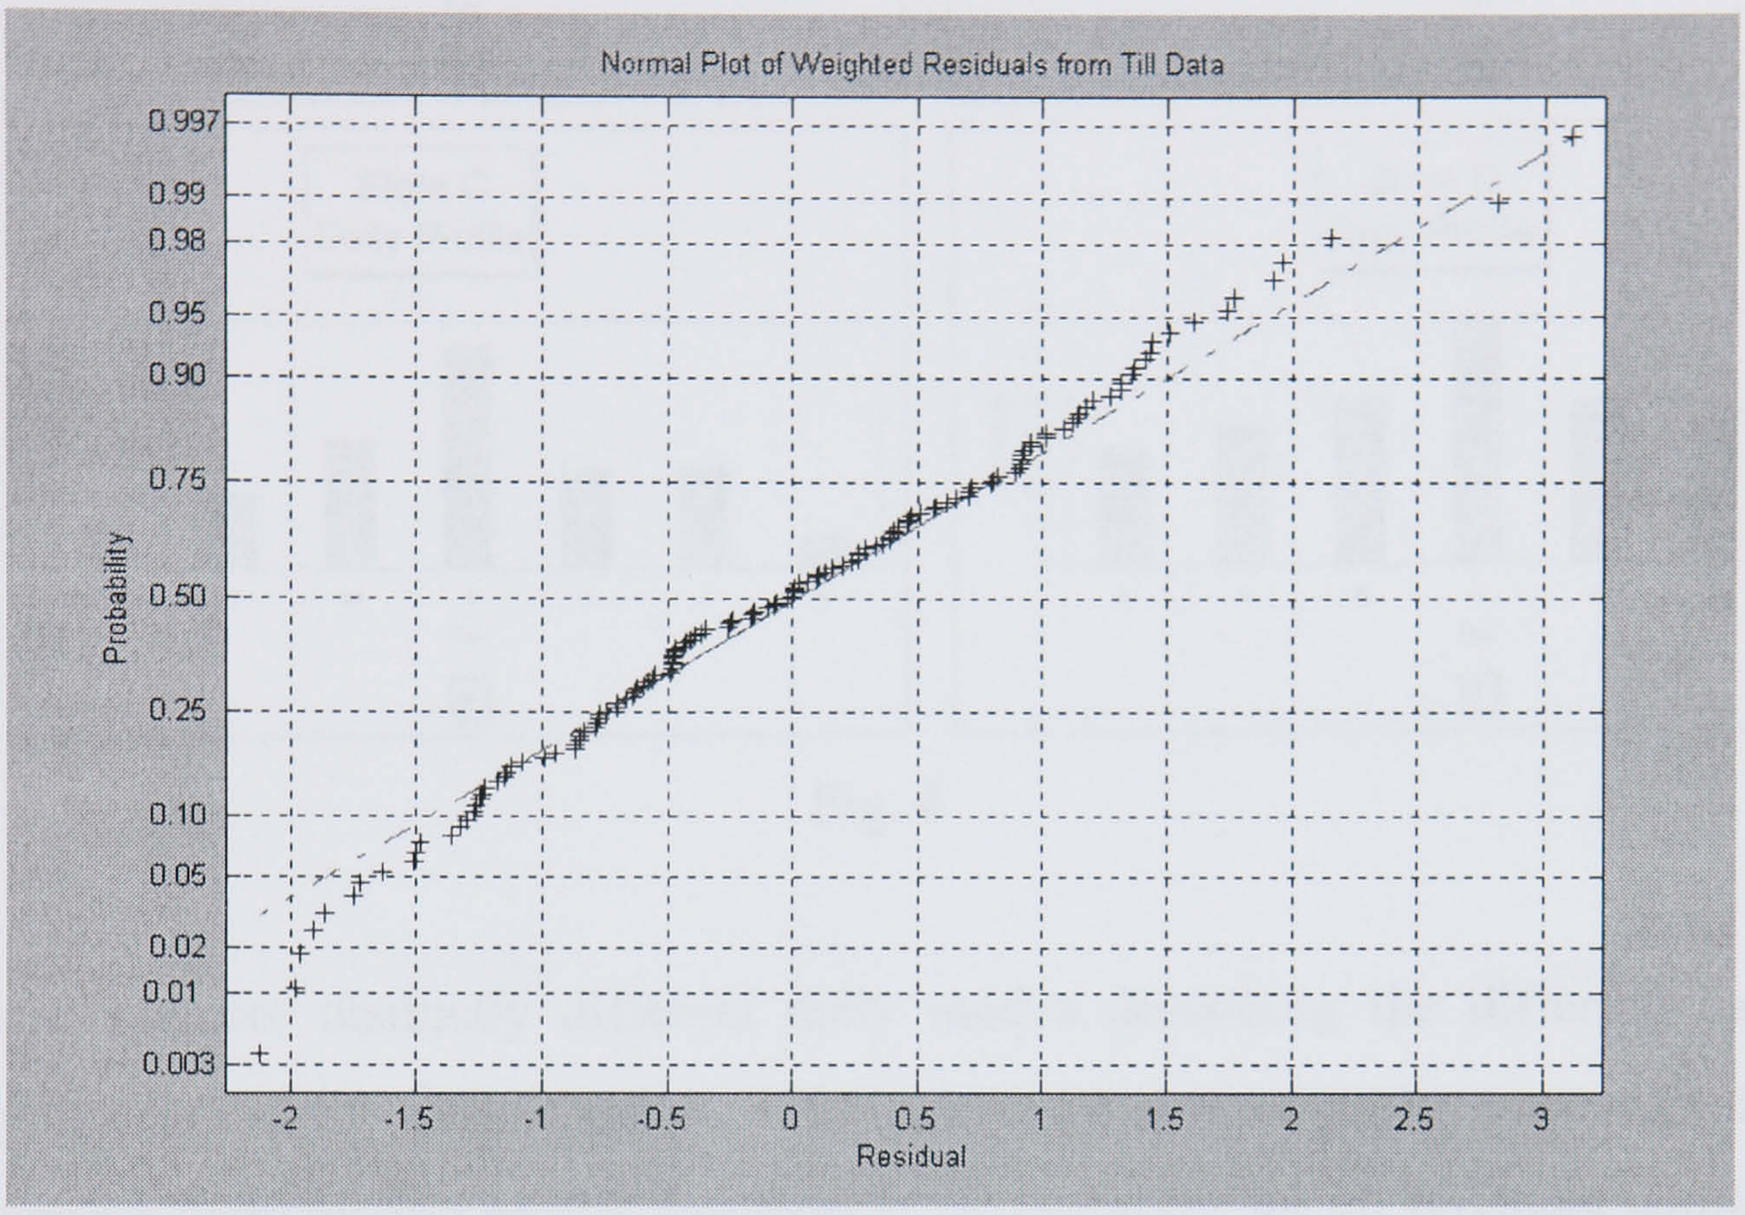


Fig. 3

2.5.2 Store Daily Mode

The daily mode values, as shown below, corresponds to the components of vector U_1 in the rank 1 model. Fig. 4 shows four different forms obtained for the factor $D(i)$ and for fitting data from four stores for the same product.

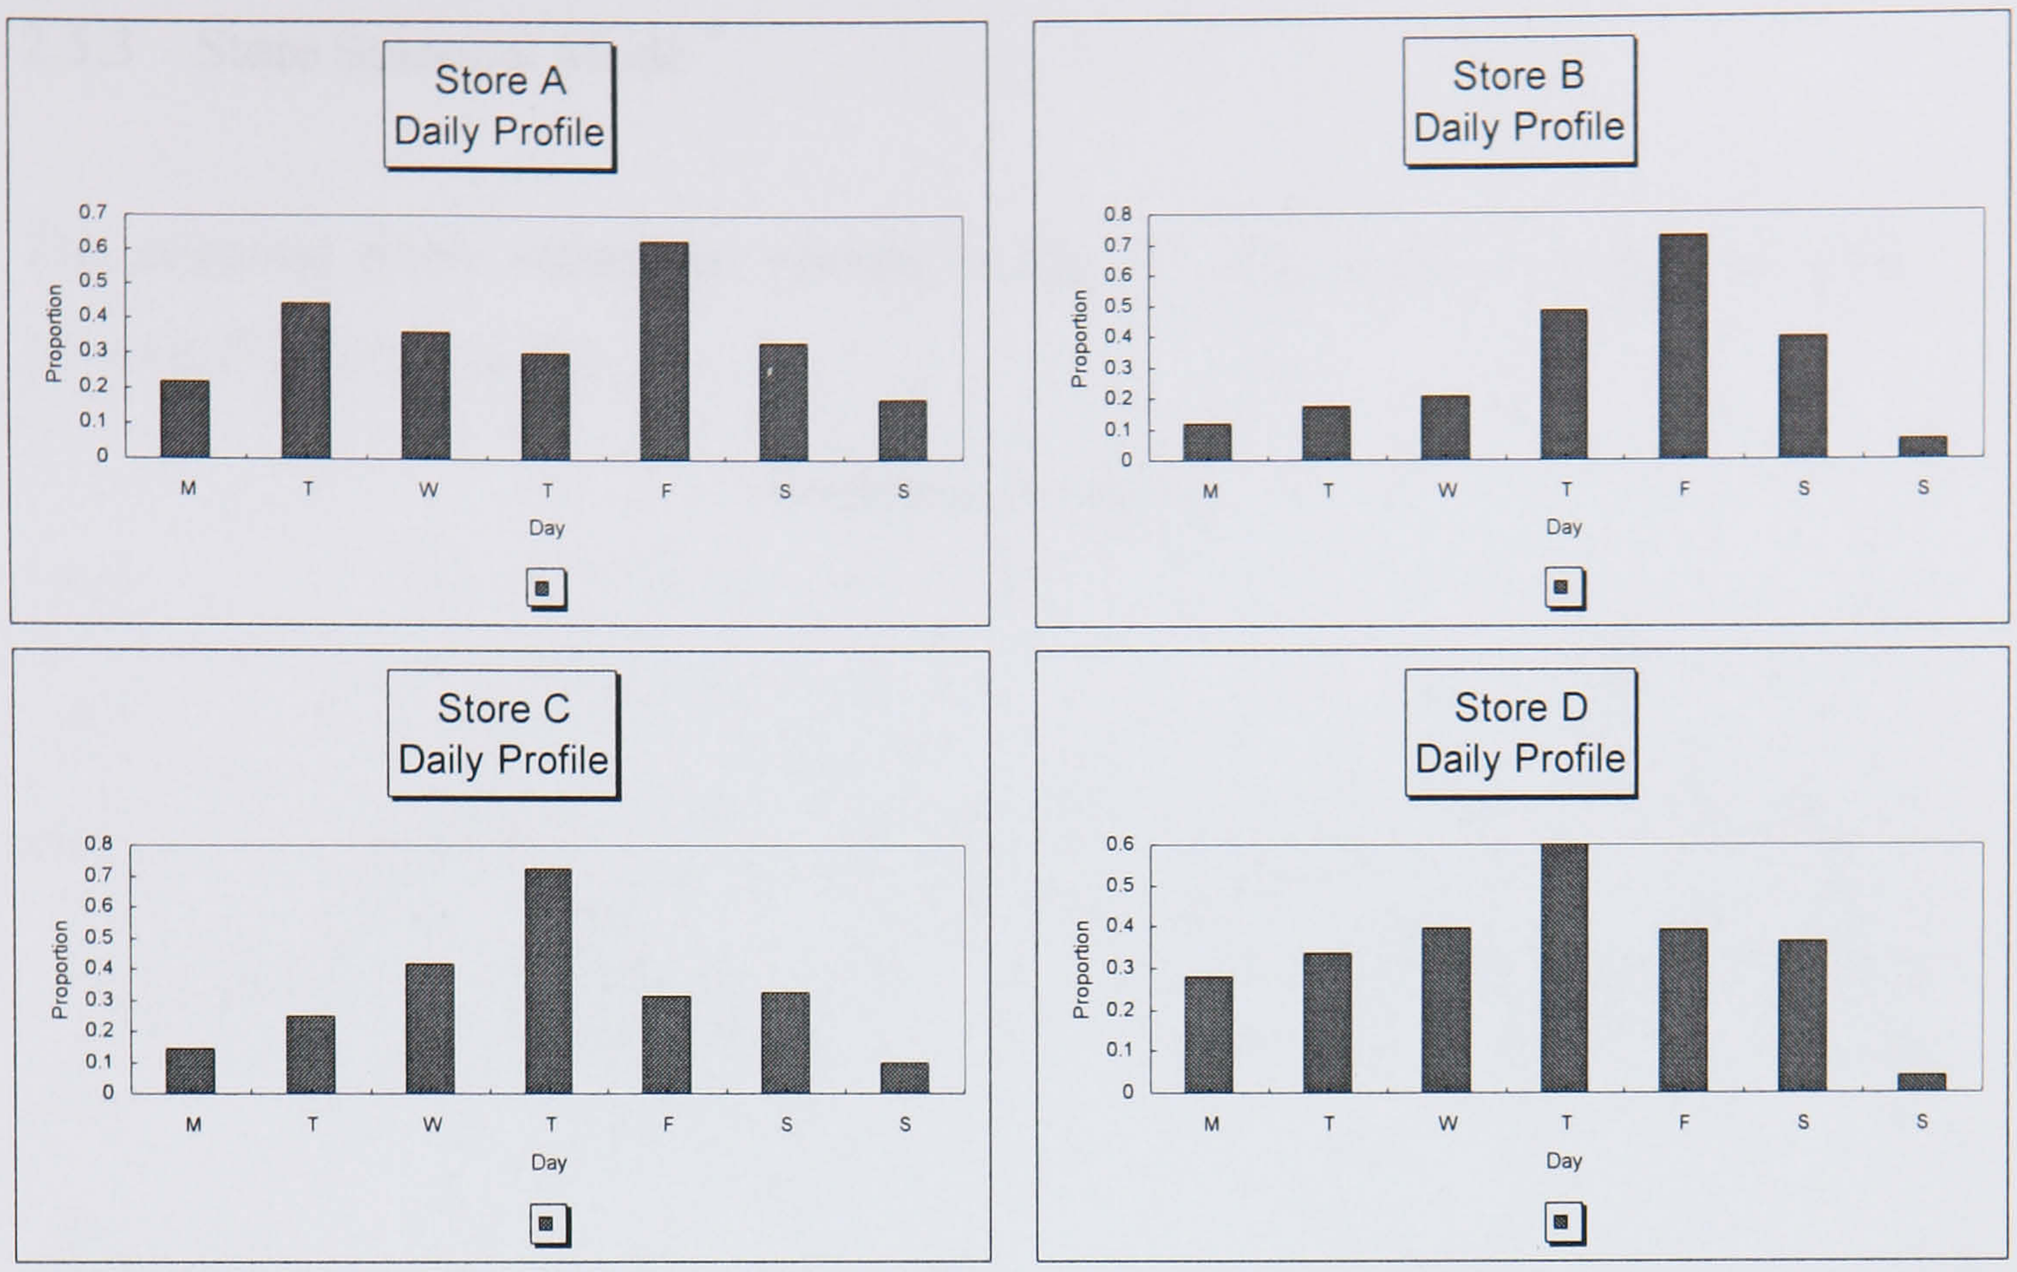


Fig. 4

These indicate distinctly different daily modes describing the differing daily buying patterns of the customers. This indicates that it would be non-optimal to assume a daily buying profile common to all stores (or for all products within a store).

In an attempt to find an explanation for the different buying profiles, a number of stores factors were investigated, including the sociological grouping of the customers who frequent the store, the size of the car park, the size of the store, and the distance out of town. The profiles shown in Fig. 4 were all for stores whose customers were mainly within the B, C1, C2 social groups, and as is evident from the figure, this sociological characteristic does not seem to explain the difference in the buying patterns. This investigation is continuing. The practical implication is that optimum forecasting for stores will only be obtained if stores and products are individually forecast. This is straightforward to implement using the real time algorithm developed later.

2.5.3 Store Seasonal Mode

The seasonal mode values are shown in Fig. 5. The results are for the store demand illustrated in Fig. 1.

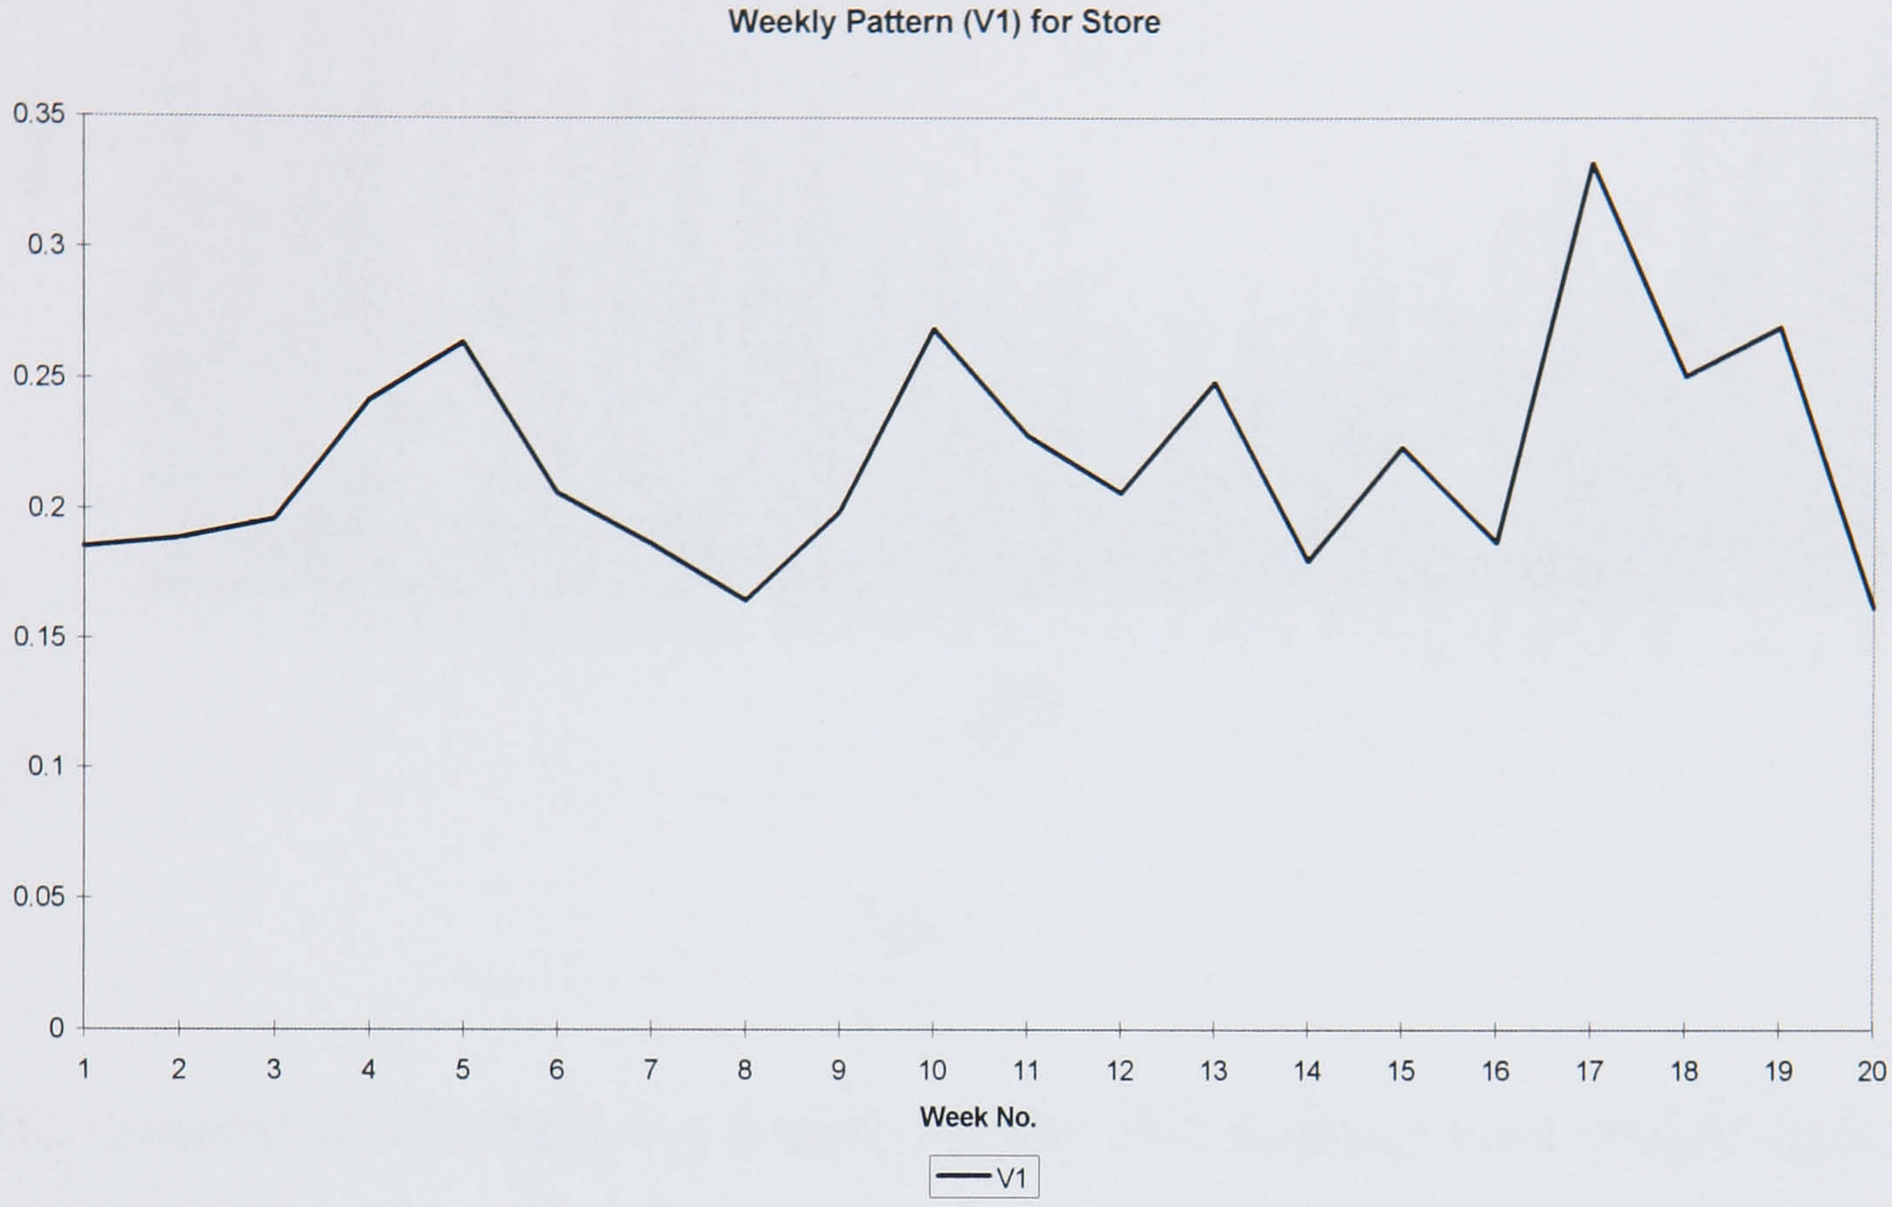


Fig. 5

As is suggested from visual examination of the store demand (Fig. 1), there is little suggestion of a seasonal pattern in the data.

2.5.4 RDC Data Results

Fig. 6 illustrates a daily demand sample for a dry soup at an RDC as recorded for a 7 day week over a period of 28 weeks.

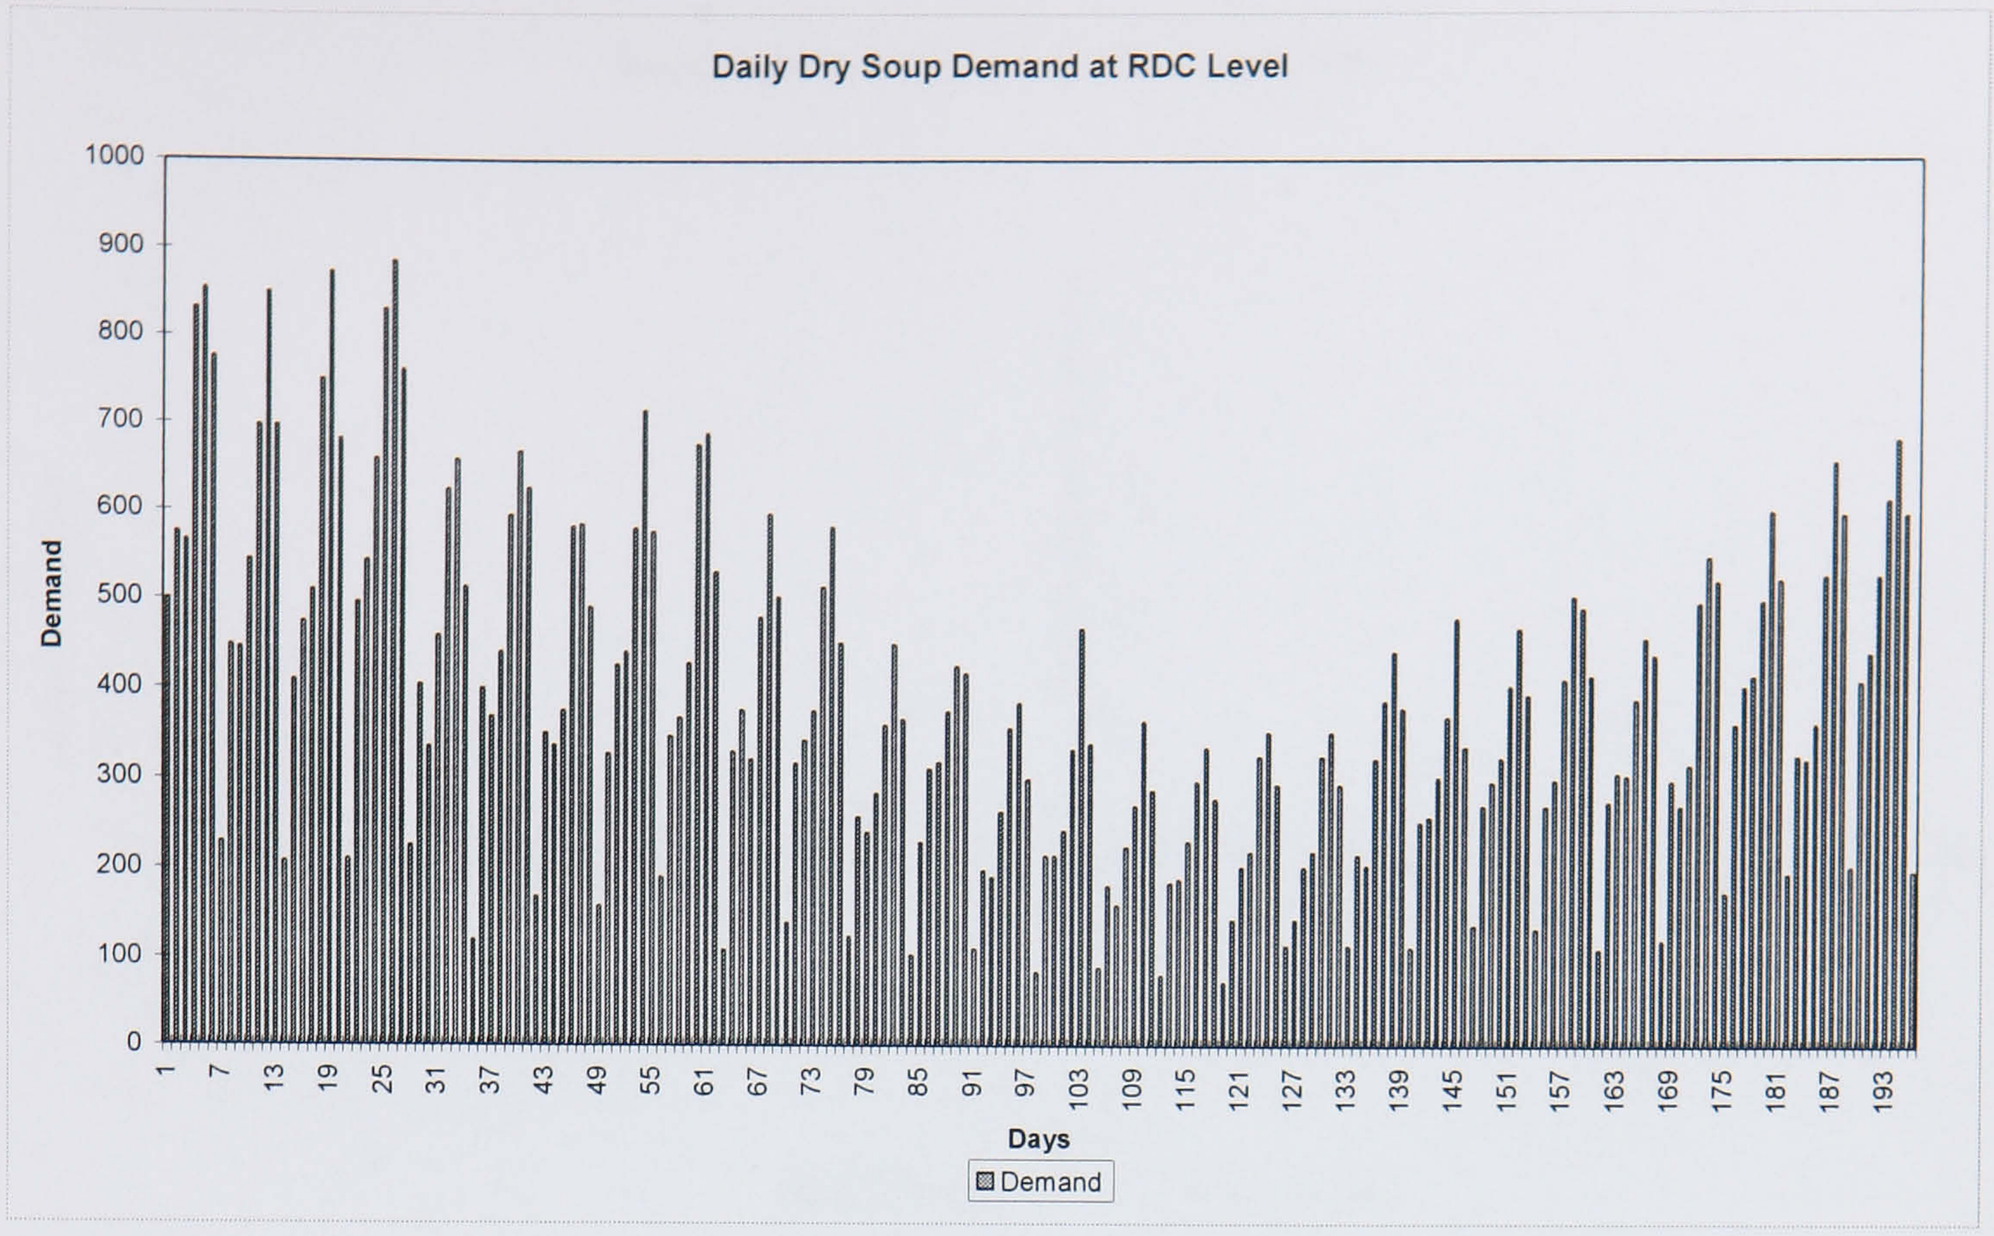


Fig. 6

The presence of both daily and seasonal modes in this data is more evident than in the till data of Fig. 1. If the distribution was pure Poisson, then the signal to noise ratio for each sample is $(\text{mean})^2/\text{variance}$ which equals μ_{ij}. Since the RDC demand approximates 200 times the store demand, the improved signal to noise ratio is reasonable. A singular value analysis of the corresponding data array Y yielded

$\lambda_1 =$ 5947.9 $\lambda_2 =$ 220.5 $\lambda_3 =$ 195.2

$\lambda_4 =$ 142.1 $\lambda_5 =$ 138.0 $\lambda_6 =$ 125.4

$\lambda_7 =$ 80.4

The AIC values for the rank 1 (rank 2) models were 230 (266). This strongly suggests choice of the rank 1 model. . The autocorrelation of the rank 1 model error is shown in fig 7.

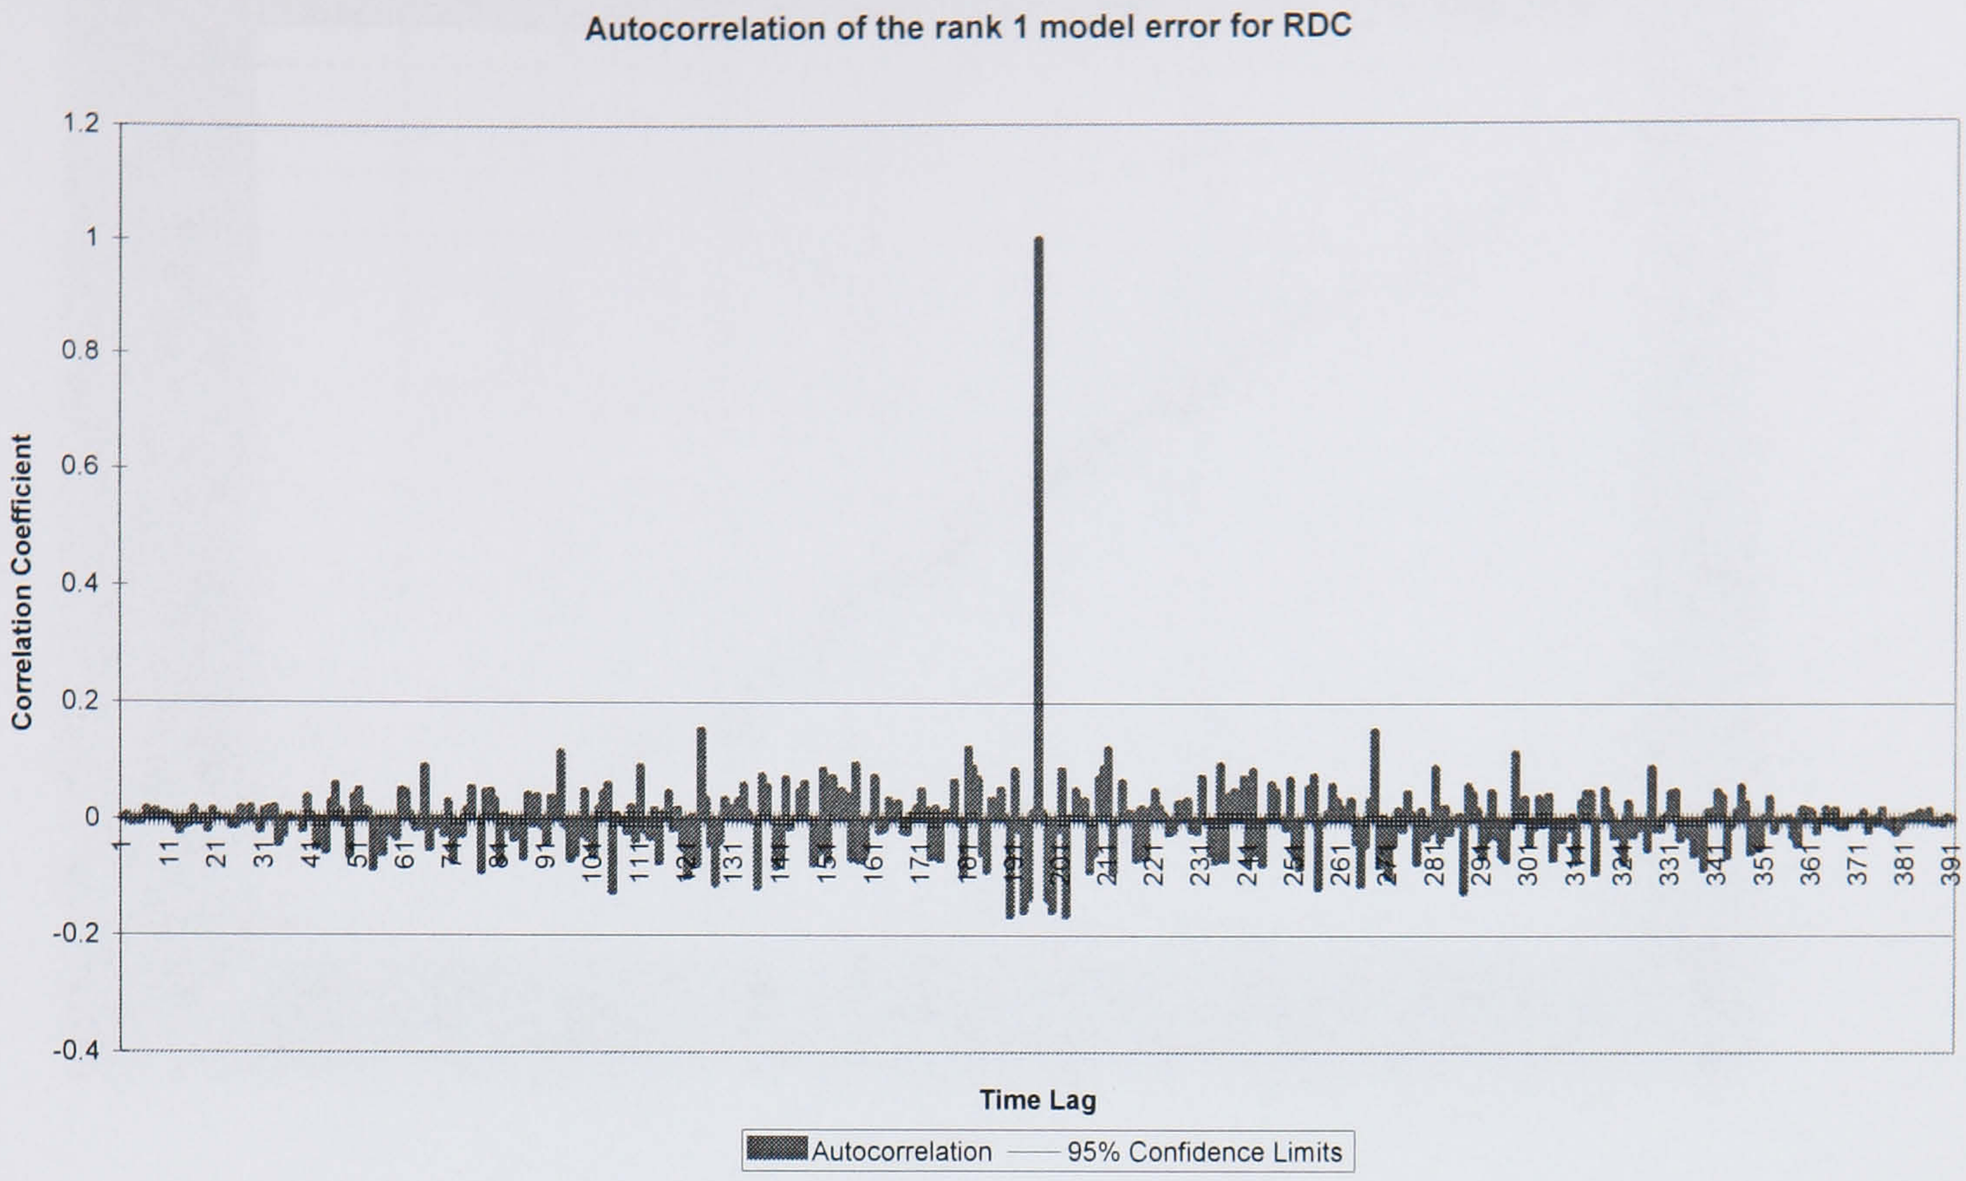


Fig. 7

The coefficients all fall within the usual 95% confidence intervals, this suggests that there is little or no model information in the model error.

The weighted residuals shown in the Q-Q below (Fig. 8) appear Normally Distributed, confirming the assumption of an independent Poisson like error distribution at RDC level is reasonable. The evaluation of the dispersion parameter, according to equation 6.0, gives $\theta^2 = 5.03$. This is significantly higher than the dispersion parameter for the stores ($\theta^2 = 1.27$). As shown later, the increased dispersion is reasonable, since the RDC demand is the sum of about 200 stores' replenishment systems. These systems introduce extra variation through the effects of batch flow constraints. (Also the individual store replenishment profile differs significantly from the till data profile).

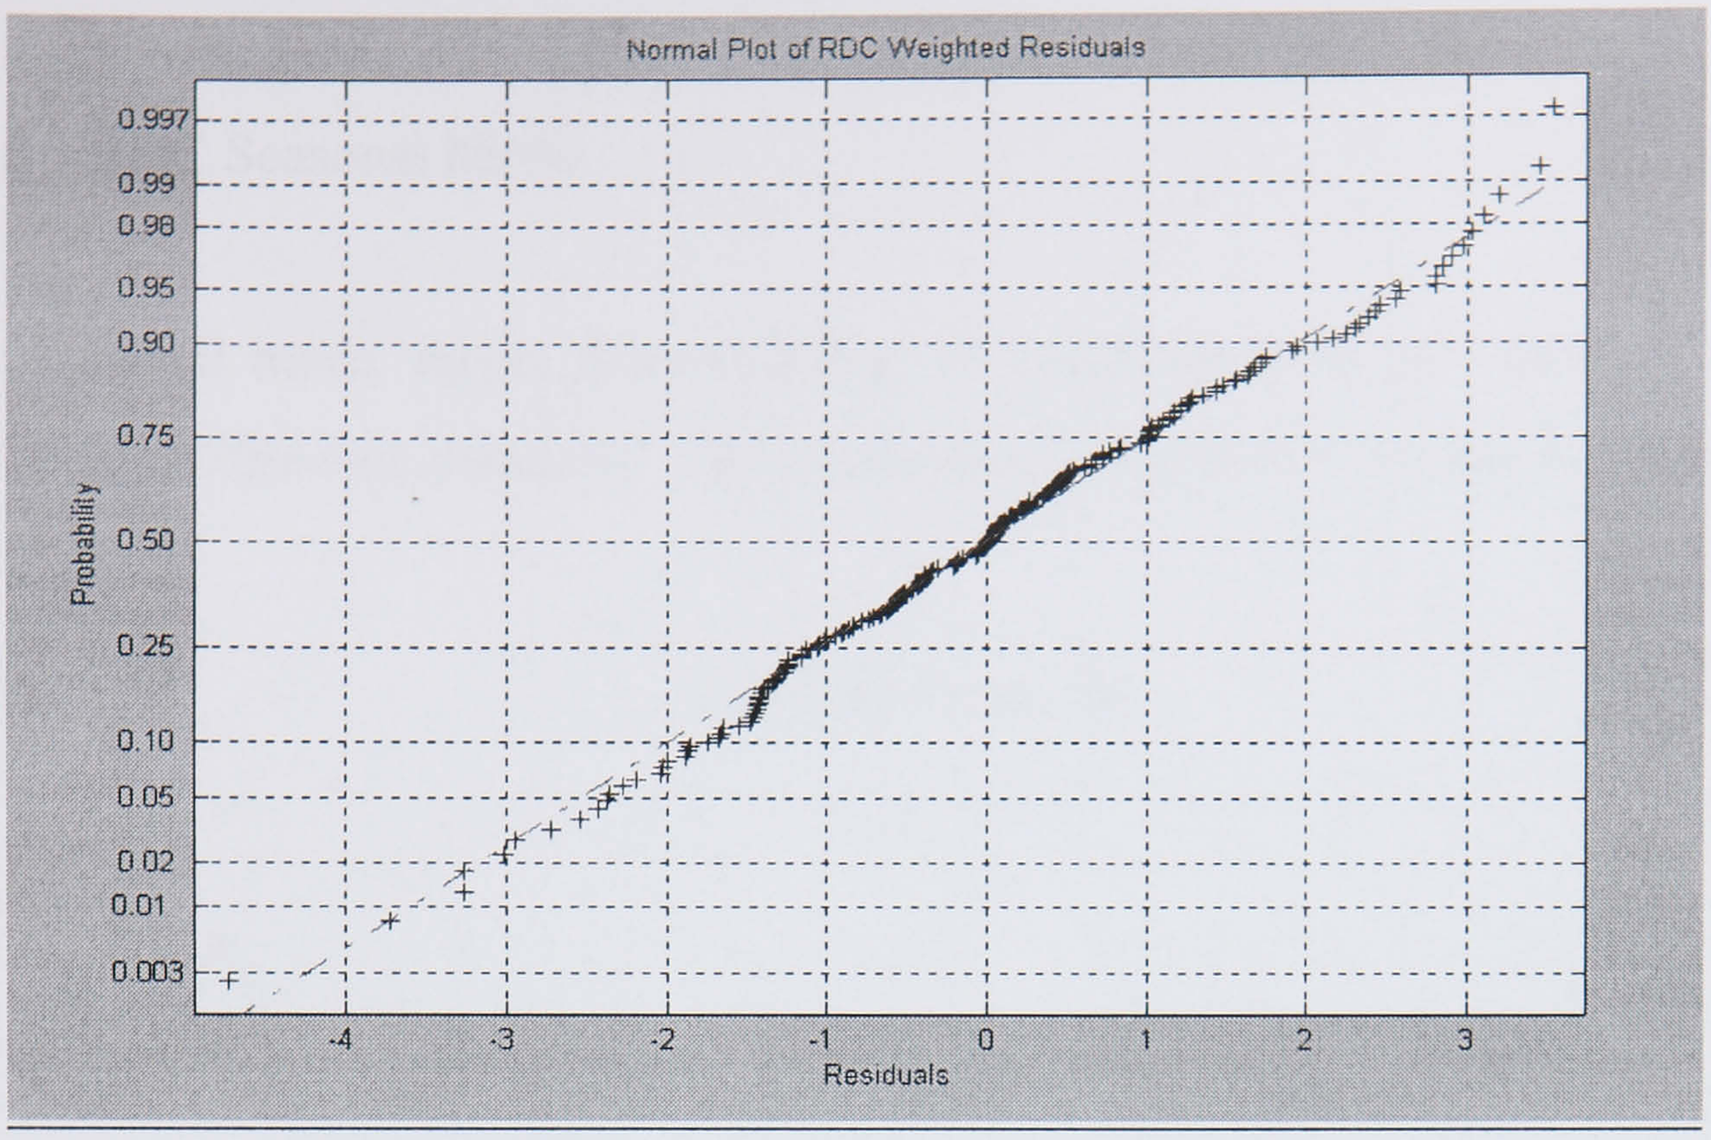


Fig. 8

2.5.5 RDC Daily Mode

Fig. 9 shows a typical daily mode for an RDC. As mentioned earlier in this chapter, the RDC has an improved signal to noise ratio compared with a store. This results in a stable daily mode at RDC level for a given product.

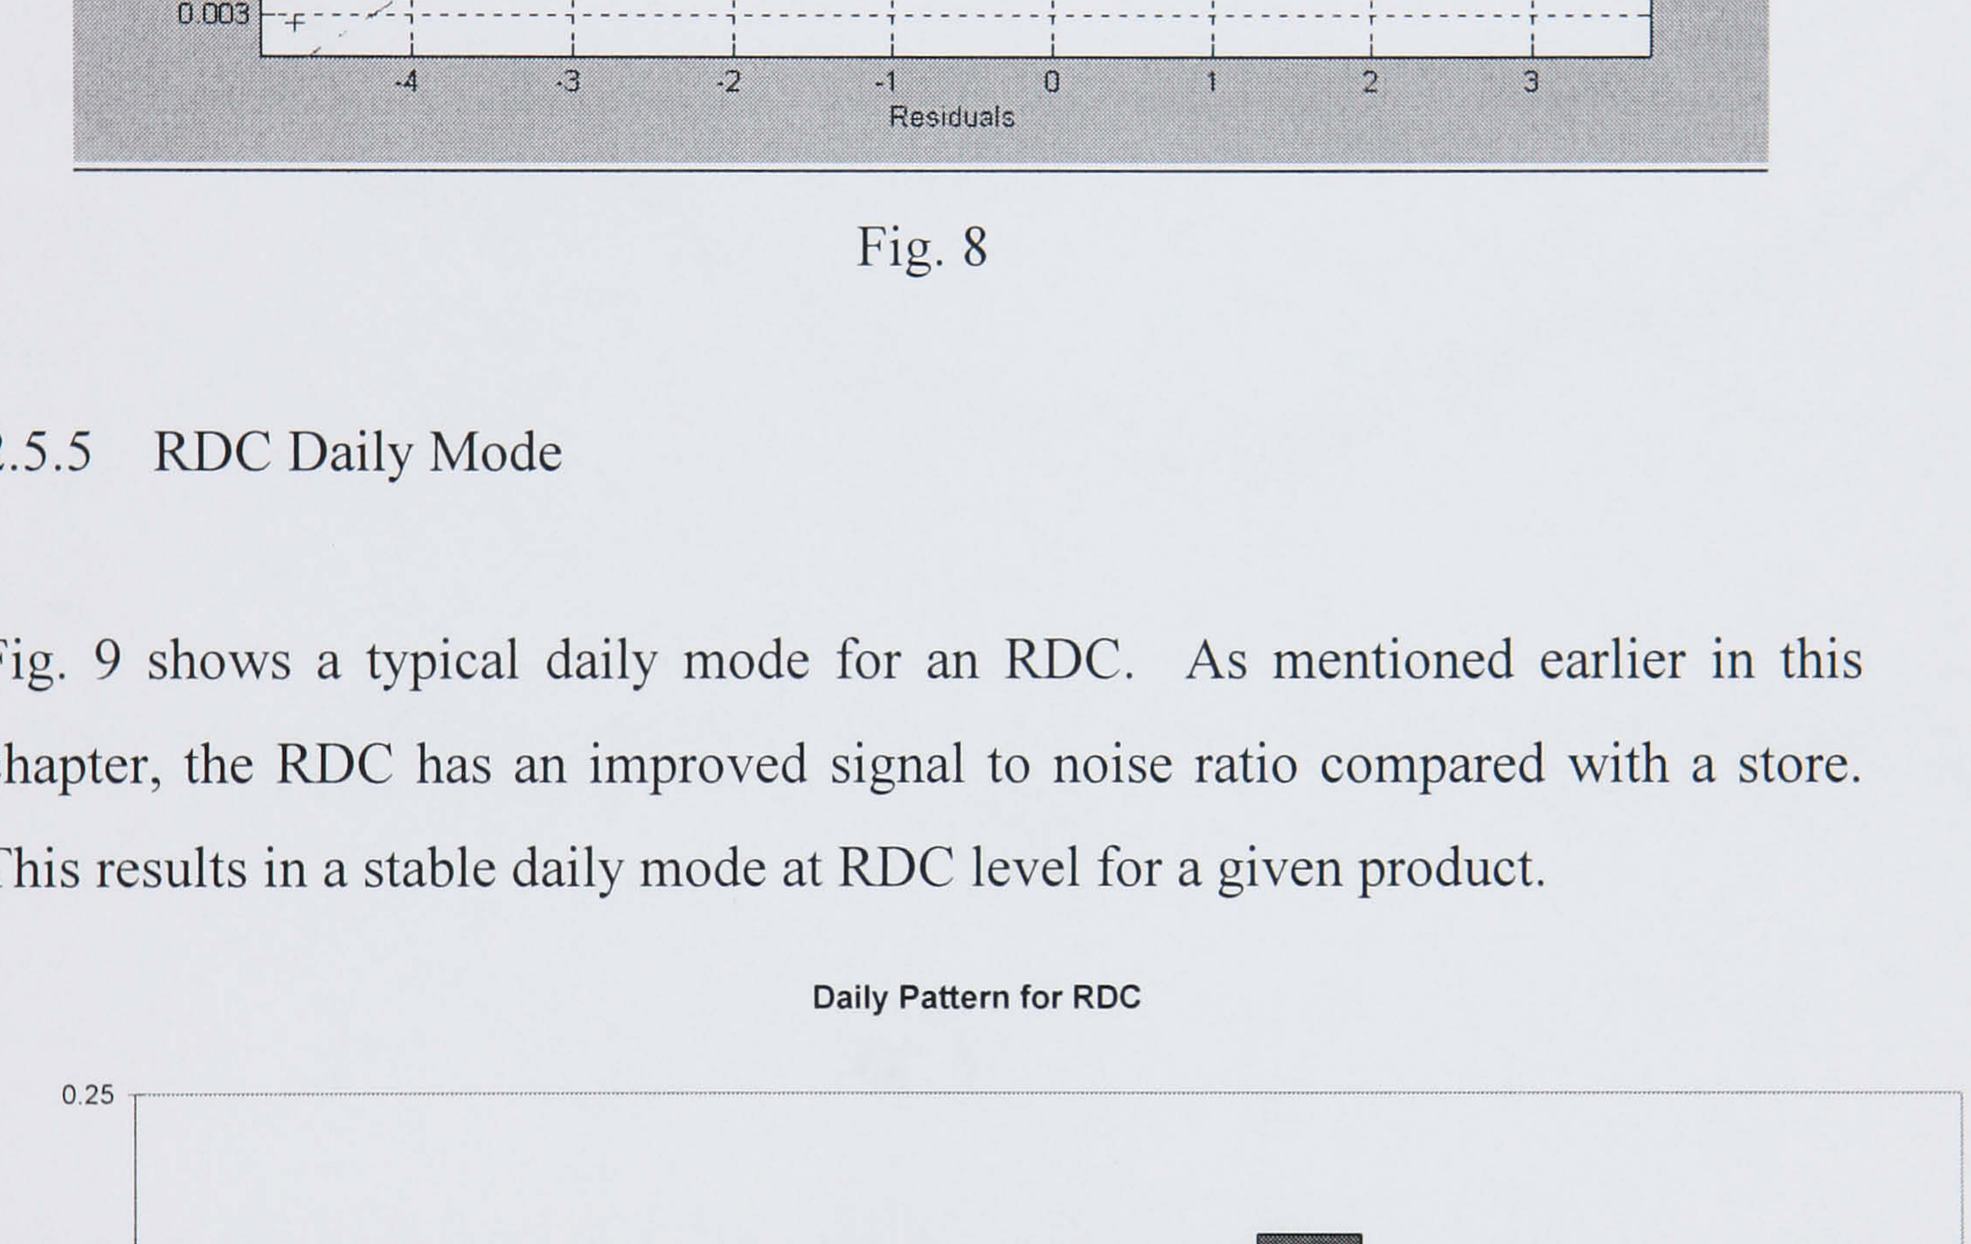


Fig. 9

2.5.6 RDC Seasonal Mode

The seasonal mode values shown in Fig. 10 corresponds to the component of vector V_1, of the rank 1 model. The results are for the R.D.C. demand illustrated in Fig. 6

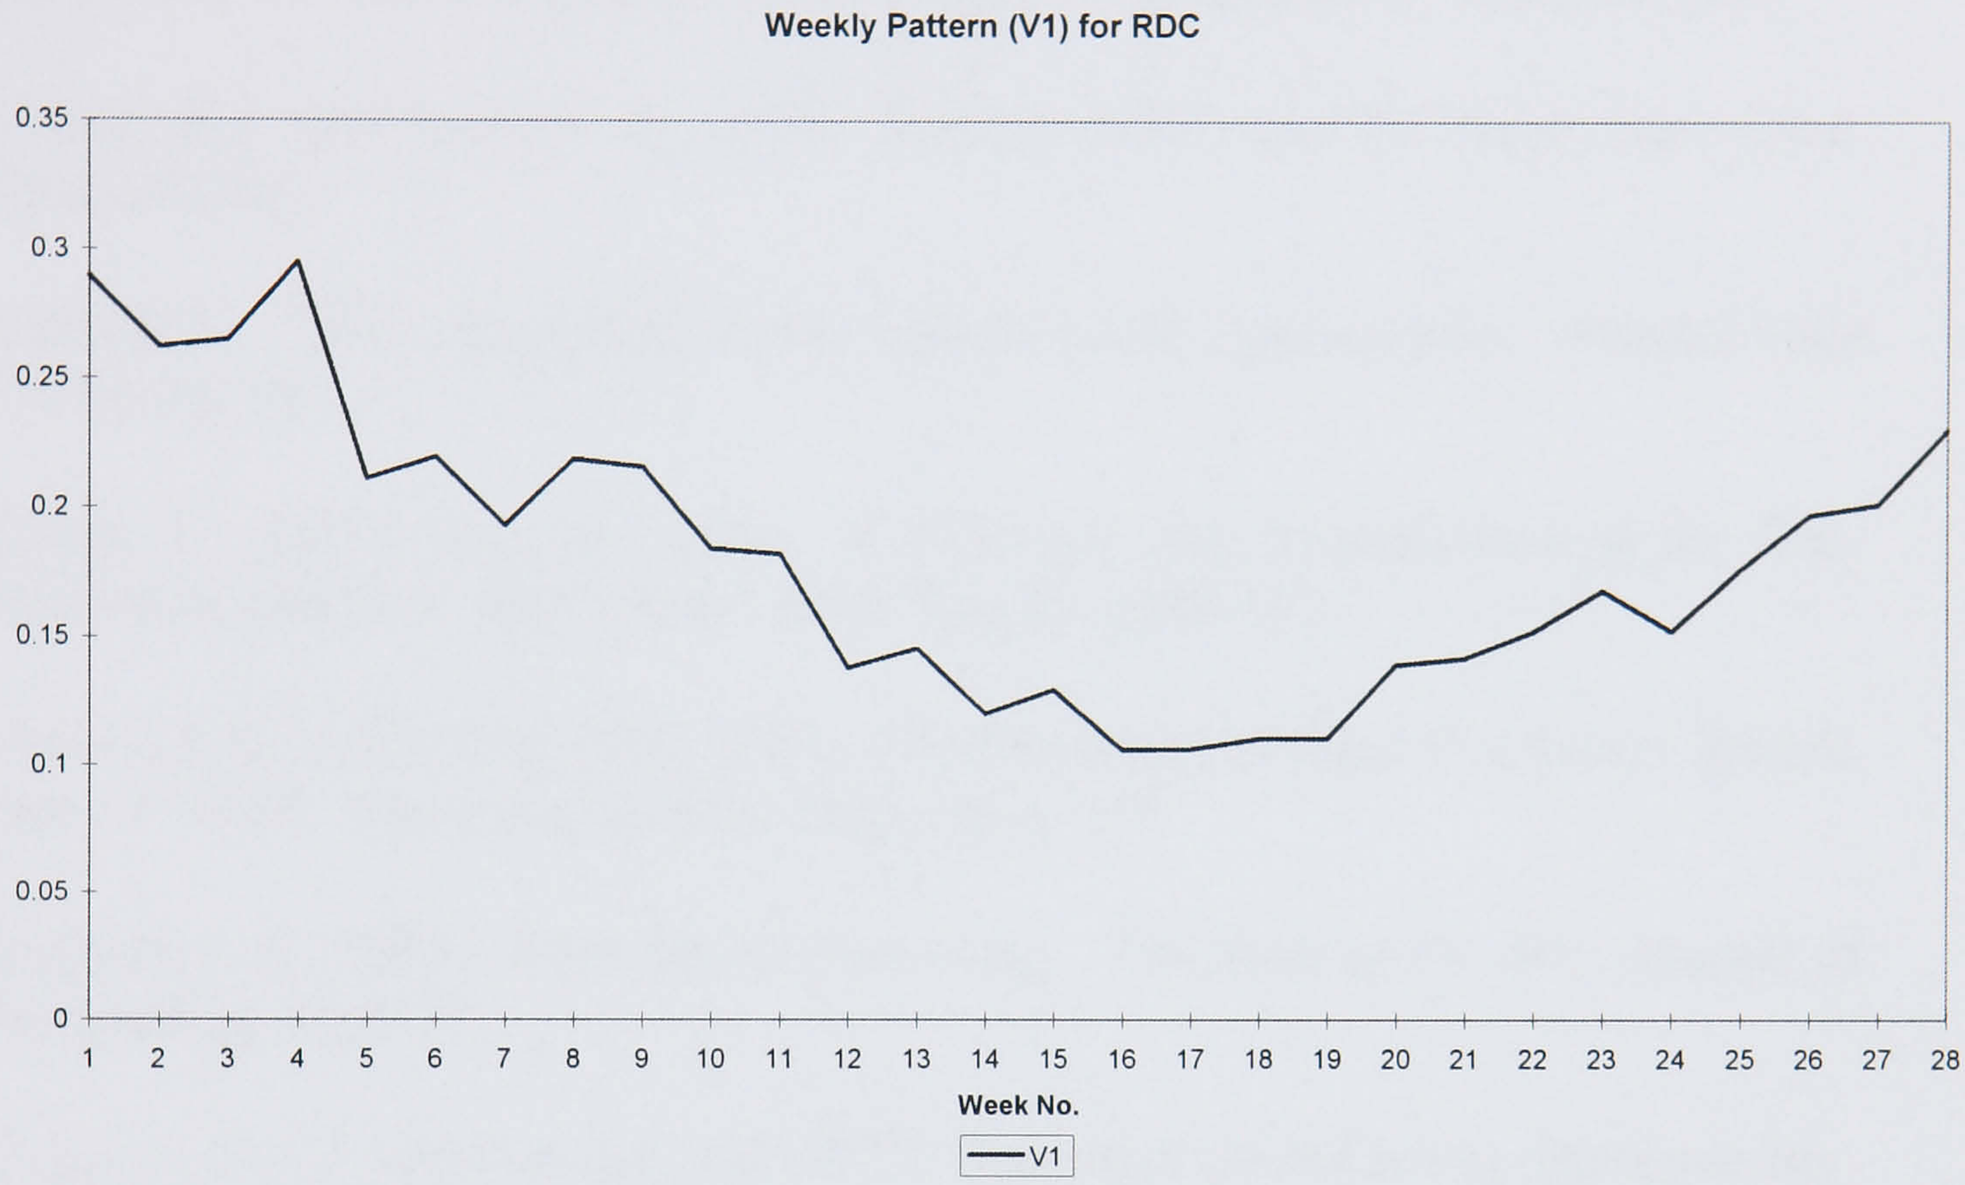


Fig. 10

This seasonal component could be fitted with say, a quadratic, however for short term forecasting and for stock control purposes, it usually suffices to approximate the seasonal variations locally by a constant or constant slope approximations.

References

Box, G.F.P. and Jenkins, G.M. 1976. *Time Series Analysis: Forecasting and Control.* Hodder Day

Brown, R.G. 1959. *Statistical Forecasting for Inventory Control.* 2nd Edition. McGraw Hill

Bowersox, D.J. and Closs, D.J. 1996. *Logistics Management.* McGraw Hill

Carroll, R.J. and Ruppert, D. 1988. *Transformation and Weighted Regression.* Chapman Hall

Datta B.N. 1995. *Numerical Linear Algebra and Applications.* Brooks Cole Publishing Co

Eckart, C. and Young, G. 1939. *A Principal Axis Transformation for Non-Hermitian Matrices.* Bull. Amer. Math. Soc.45. p118-121.

Farmer, E.D. and Potton, M.J. 1966. *The Prediction of Load in a Power System.* Proc. 3rd IFAC Congress, London 1966, paper 21F.

Gardner, E.S. 1985. *Exponential Smoothing: The State of the Art.* Journal of Forecasting, 4, p1-27.

Hastings, N.A.J. and Peacock, J.B. 1975. *Statistical Distributions.* Butterworths

Holt, C.C. 1957. *Forecasting Seasonal and Trends by Exponentially Weighted Moving Average.* Office of Naval Research, Research Memorandum 52

Lindsey, J.K. 1995. *Introductory Statistics.* Oxford University Press

Makridakis, S. and Wheelwright, S.C. 1989. *Forecasting Methods for Management.* John Wiley and Sons.

Matloff, N., Rose, R. and Tai, R. 1984. *A Comparison of Two Methods for Estimating Weights in Regression Analysis.* Jnl. Stat. Computational Simulation 19, p265-74

Matthewman, P.D. and Nicholson, H. 1968. *Techniques for Load Prediction in the Electricity Supply Industry.* Proc. Inst. of Elec. Eng. Vol 115, No 10, 1451-58.

McCullagh, P. and Nelder, J.A. 1984. *Generalised Linear Methods.* Chapman Hall.

Mills, T.C. 1996. *Time Series Techniques for Economists.* Cambridge University Press.

Stewart, G.F. 1973 *Introduction to Matrix Computations*. John Wiley.

Williams, E. 1959. *Regression Analysis*. Wiley.

Winters, P.R. 1960. *Forecasting Sales by Exponentially Weighted Moving Average*. Management Science.

3.0 MAXIMUM LIKELIHOOD APPROACHES

3.1 Introduction

In the previous chapter we developed a generalised least squares (GLS) approach to fitting the supply chain demand model. The GLS solution, if convergence is achieved, is equal to the maximum likelihood equation solution for the Poisson distribution. Carroll and Ruppert (1988). As these equations are easily solved and can be adapted for use on-line, we consider the alternative approach of directly solving the maximum likelihood equations for the model. The first section of the chapter considers the fitting of a model in which the weather modes are ignored, as in the previous chapter. For this model we also develop on-line forecasting and jump detection procedures. Finally we consider results from off-line fitting of a model which includes weather modes.

3.2 Off-Line Development

As in the previous chapter, consider again an $I \times J$ array Y in which entry y_{ij} represents the demand on day i of week j. Assuming the distribution for y_{ij} is Poisson, then as is well known, Hastings and Peacock [75], the probability function may be expressed

$$f(y_{ij};\mu_{ij}) = \Pr(y_{ij} = k)$$

$$= \frac{\mu_{ij}^{\,y_{ij}} e^{-\mu_{ij}}}{y_{ij}!}$$

$$= \exp\{y_{ij}\log\mu_{ij} - \mu_{ij} - \log y_{ij}!\} \qquad 1.0$$

and the log likelihood for observation y_{ij} is therefore

$$\ell(\mu_{ij};y_{ij}) = \{y_{ij}\log\mu_{ij} - \mu_{ij} - \log y_{ij}!\} \qquad 2.0$$

for $1 \le i \le I$, $1 \le j \le I$, where $\mu_{ij} = E(Y_{ij})$, and ! indicates factorial.

Assuming the observations are independent, the joint log likelihood function for the array Y is

$$\ell(\mu_{ij}; Y) = \sum_{i=1}^{I} \sum_{j=1}^{J} \ell(\mu_{ij}; y_{ij}) \qquad 3.0$$

From equation 12.0 of the previous chapter, the simplified rank 1 model, assuming only daily and seasonal effects is

$$\mu_{ij} = \lambda d_i s_j \qquad \begin{matrix} 1 \le i \le I \\ 1 \le j \le J \end{matrix} \qquad 4.0$$

and as in the previous chapter, to uniquely specify the model, we introduce constraints, this time of the form

$$\sum_{i=1}^{I} |d_i| = \sum_{j=1}^{J} |s_j| = 1 \qquad 5.0$$

Substituting equation 4.0 into equation 3.0 we obtain

$$\ell(\lambda, d_i, s_j; Y) = \sum_{i=1}^{I} \sum_{j=1}^{J} \{ y_{ij} \log \lambda d_i s_j - \lambda d_i s_j - \log y_{ij}! \} \qquad 6.0$$

In Appendix 1 we show

RESULT 2

The maximum likelihood estimate of μ_{ij} may be expressed

$$\hat{\mu}_{ij} = \hat{\lambda}\hat{d}_i\hat{s}_j$$

$$\text{where} \quad \hat{d}_i = S_i^r \Big/ \left\| S^r \right\|$$

$$(\hat{\lambda}s_j) = S_j^c$$

$$\hat{\lambda} = \left\| S^c \right\| = \sum_{i=1}^{I}\sum_{j=1}^{J} y_{ij}$$

where $\|\cdot\|$ denotes the vector 1-norm, Barnett [90] and we define

$$S_i^r = \frac{1}{J}\sum_{j=1}^{J} y_{ij} \qquad i = 1\ldots I$$

$$S_j^c = \sum_{i=1}^{I} y_{ij} \qquad j = 1\ldots J$$

which are sufficient statistics.

To include the promotional mode, equation 4.0 may be modified to the form

$$\mu_{ij} = \lambda d_i s_j p_j \qquad 7.0$$

where $p_{ij} = 1$ outside the promotional period. For this case, as in the previous chapter, Result 2 may be directly applied to give an estimate of the product $\left(\lambda s_j p_j\right)$. Results based on Result 2 are not presented since they were near identical to those for the weighted least square approach of the last chapter.

3.3 On Line Model Developments

It is desirable that models used in on-line application should react quickly to changes in demand, and should be recursive, efficient in storage, robust and easy to maintain. The off-line estimation procedures of Result 2 were developed for processing weekly blocks of data and do not meet these requirements. For instance, statistics such as S_i^r used in Result 2 clearly become unresponsive when the number of weeks is large. This is obvious from expressing the statistic in the recursive form Gelb [74],

$$S_i^r(J) = S_i^r(J-1) + \frac{1}{J}[y_{iJ} - S_i^r(J-1)] \qquad 8.0$$

$i = 1 \ldots I$, where $(\cdot)$ indicates the number of weeks used in calculating S_i^r. Clearly the statistic $S_i^r(J)$ will react slowly to changes in demand level once J becomes large. S^r may be viewed as an unnormalised estimate of the daily response vector D and whose components are assumed to be constant, as assumed in the model. If the demand sequence $Y_{i\bullet}$ is assumed Poisson then an approximate model, assuming constant mean μ_i is

$$y_{ij} = \mu_i + \varepsilon_{ij} \qquad 9.0$$

where $\varepsilon_{ij} \in N(0, \mu_i)$ Kalbfleisch [85] and it is easily shown that $S_i^r(J)$ is the best least square estimates of μ_i minimising $\sum_{j=1}^{J} \varepsilon_{ij}^2$. Gelb [74]. A convenient and widely used approach, to developing a more responsive statistic is based on the use of a weighted least square criteria

$$E_i = \sum_{j=1}^{J} \beta^{J-j} e_{ij}^2 \qquad 10.0$$

where errors j units ago we are weighted by $\beta^{J-j}\left(j=1...J\right)$. This criteria 'remembers' error information in an assumption memory length of

$$N = 1/\left(1-\beta\right) \qquad 11.0$$

the weighting of error information dying away with a time constant of N intervals, Clarke and Gawthrop [75].

If a process parameter changes and the error level increases, it is desirable that β and N are reduced, to allow the model to adapt parameters to the new condition. Various procedures for adapting the forgetting factor β have been developed, Kershenbaum [88]. In Appendix II we consider a proposal by Fortesque et al [81] who suggested the forgetting factor be adjusted to keep an "information" measure Σ_0 say, constant. As shown in Appendix II this leads to the adaptive exponential smoothing procedure.

Algorithm 1. (Conceptual) For $i=1...I, j=1...J$.

$$\text{(i)} \qquad \varepsilon_{ij} = y_{ij} - S_i^r\left(j-1\right)$$

$$\text{(ii)} \qquad \beta = \Sigma_0^i / \Sigma_0^I + \varepsilon_{ij}^2$$

$$\text{(iii)} \qquad S_i^r\left(j\right) = S_i^r\left(j-1\right) + \left(1-\beta_{ij}\right)\varepsilon_{ij} \qquad 12.0$$

The parameter Σ_0^I is related to the sum of squares of the errors over the memory length, Fortesque et al [81], and one possible guideline for choosing Σ_0^i is therefore

$$\Sigma_0 = \sigma_0^2 N_0$$

where σ_0^2 is the expected process noise variance. If a nominal value of demand $\hat{\mu}_i$ of μ_i is available, then a possible guideline to choice of parameter Σ_0^I is

$$\Sigma_0^i = N_0 \hat{\mu}_i \hat{\theta}^2 \qquad 13.0$$

$i = 1 \ldots I$ where N_0 controls the memory length and $\hat{\theta}^2$ is the estimated dispersion coefficient.

Algorithm 1 provides a procedure for updating the statistics S_i^r used in result 2. We now consider a procedure for also obtaining daily estimates of the statistic S_j^c used in Result 2. On day i of week j the data array does not generally form a complete $I \times J$ array. A complete block may be formed by deleting the first i days of data from an array, and renumbering the remaining array elements, as indicated in Fig 1.

Sequential Array Statistics S^r and S^c

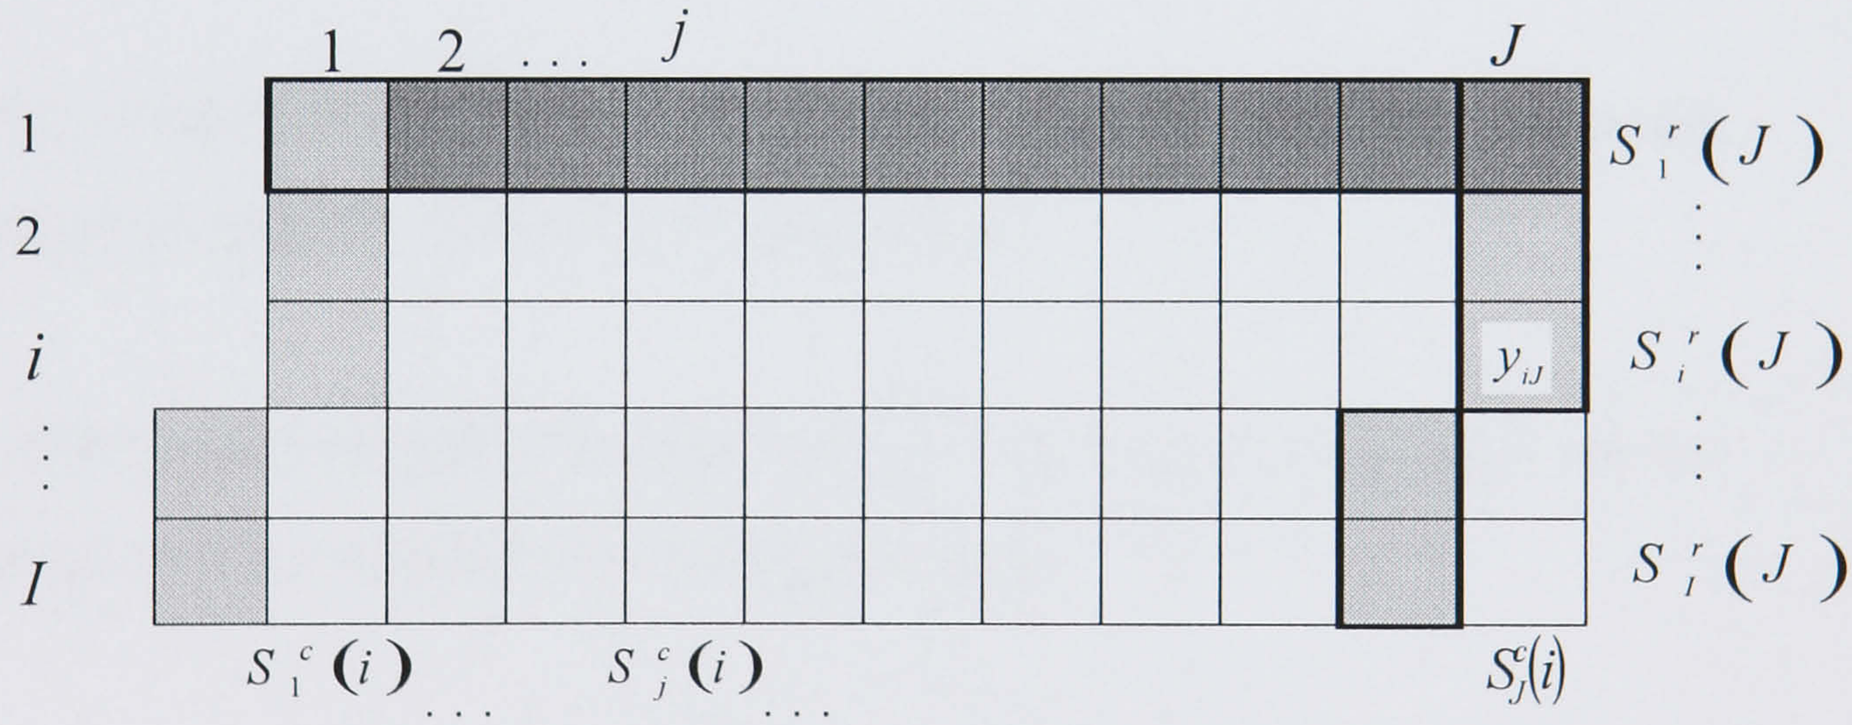


Fig. 1

The vector S^c for this modified array will, for instance, have its last element $S_J^c(i)$ equal to the sum of the last I data points. This statistic may be expressed in the sequential form

$$S_J^c(i) = S_J^c(i-1) + \left(y_{iJ} - y_{i(J-1)}\right) \qquad 14.0$$

a similar relationship holding for all components of S^c.

Result 2 may now also be expressed in a sequential form as follows:

$$\hat{\mu}_{iJ} = \left(\hat{\lambda s}\right)_{iJ} \hat{d}_{iJ}$$

$$\text{where } \left(\hat{\lambda s}\right)_{iJ} = S_J^c(i) \qquad 15.0$$

$$\hat{d}_{iJ} = S_i^r(J) \Big/ \left\| S^r \right\|$$

where the components $S_J^c(i)$ and $S_i^r(J)$ are updated using the sequential relationships of equation 14.0 and 12.0 respectively.

If an r step ahead prediction is required on day i of week j then, assuming the model parameters remain sensibly constant, we have

$$\hat{\mu}_{(i+r)J} = \left(\hat{\lambda s}\right)_{iJ} \hat{d}_{(i+r)J}$$

where the estimates on the right have been evaluated on day i of week J and to avoid tedious notation, we assume $(i+r) < I$. If the forecast is required over a period of months, say, then the assumption of constant parameters is not usually valid, and the seasonal characteristics of the model must be modelled.

3.4 Event Detection

Sudden jumps in demand may occur in practice without warning or known explanation. It is essential to detect the onset of these changes and estimate there magnitudes as quickly as possible in order to allow time for re-adjustment to manufacturing and operating conditions along the supply chain in order to avoid stock out or overstocking.

Significant differences between planned and actual response do occur during promotional periods due usually to incorrect estimation of promotional uplift. We consider the detection and estimation of promotional uplift error first. The model for the expected demand on day i of week j, from equation 7.0, now including daily uplift factor, can be expressed

$$\mu_{ij} = \left(\lambda s_{ij}\right) d_{ij} p_{ij}$$

where p_{ij} represents the uplift on day i of week j during the promotional period and assumes the value unity outside the promotion period.

If $\hat{p}_{ij}$ represents the a priori estimate for p_{ij}, and we assume all substantial model errors are associated with this mode, then the forecast demand is

$$\hat{\mu}_{ij} = \lambda d_i s_j \hat{p}_{ij}$$

Based on this assumption, the ratio of the actual demand y_{ij} to the forecast demand $\hat{\mu}_{ij}$ has the expected value

$$\begin{aligned} \alpha_{ij} &= E\left(y_{ij} / \hat{\mu}_{ij}\right) \\ &= \mu_{ij} / \hat{\mu}_{ij} \\ &= p_{ij} / \hat{p}_{ij} \end{aligned} \qquad 16.0$$

and variance

$$E\left[\left(y_{ij}/\hat{\mu}_{ij}\right)-\alpha_{ij}\right]^2$$

$$= E\left[\frac{y_{ij}}{\hat{\mu}_{ij}} - \frac{\mu_{ij}}{\hat{\mu}_{ij}}\right]^2$$

$$= E\left(y_{ij}-\mu_{ij}\right)^2 \Big/ \hat{\mu}_{ij}{}^2$$
$$= \theta^2 \mu_{ij} \Big/ \hat{\mu}_{ij}{}^2 \qquad 17.0$$
$$= \theta^2 \alpha_{ij} \Big/ \hat{\mu}_{ij}$$

The adaptive exponential smoothing algorithm developed in Appendix 2.0 may be used to estimate α_{ij} and to update $\hat{p}_{ij}$ as follows

Algorithm 2 (Conceptual)

$$\text{(i)}\ \ \varepsilon_{ij} = \left(y_{ij}/\hat{\mu}_{ij}\right) - \alpha_{i(j-1)}$$

$$\text{(ii)}\ \ \lambda_{ij} = \Sigma_0^{ij} \Big/ \left(\Sigma_0^{ij} - \varepsilon_{ij}^2\right)$$

$$\text{(iii)}\ \ \alpha_{ij} = \alpha_{i(j-1)} + \left(1-\lambda_{ij}\right)\varepsilon_{ij}$$

$$\text{(iv)}\ \ \hat{p}_{(i+1)j} = \hat{p}_{ij}\alpha_{ij} \qquad 18.0$$

The parameter Σ_0 is related to the sum of squares of the process errors and from equation 17.0 a possible choice is

$$\Sigma_0^{ij} = N_0\theta^2\alpha_{(i-1)j} \Big/ \hat{\mu}_{ij}$$

The discussion so far has been concerned with estimation of errors attributed to the promotional uplift parameter p_{ij}. Sudden and unexplained changes in demand level also occur during non-promotional periods. The analysis and approach above can be modified to treat such changes, assuming the daily (d_i) and seasonal (s_j) parameters remain sensibly constant.

The daily model during non-promotional periods, from equation 4.0, may be expressed in the form

$$\mu_{ij} = \lambda_{ij}\, d_{ij}\, s_{ij}$$

a change in demand level being represented as a change in parameter λ . If $\lambda_{ij}\left(\hat{\lambda}_{ij}\right)$ are the actual (estimated) values on day i of week j then we may define α_{ij} as above, and following the same reasoning, obtain

$$\alpha_{ij} = E\left(y_{ij}/\hat{\mu}_{ij}\right) = \mu_{ij}/\hat{\mu}_{ij} = \lambda_{ij}/\hat{\lambda}_{ij}$$

Algorithm 2 may therefore be applied to the detection of level changes, except that step (iv) is modified to

$$\hat{\lambda}_{(i+1)j} = \hat{\lambda}_{ij}\alpha_{ij}$$

3.5 Results and Discussion

The results obtained using the 'on-line' MLE approach are presented below. To allow comparison, the same demand data is used as with the off-line SVD based approach (see Figs. 1 & 6 of Chapter 2). Results are produced for both till data and RDC data. As the MLE 'on-line' model is concerned with forecasting demand, the model errors shown below are for a one day ahead demand forecast against actual demand and therefore differ from the data fitting error used to assess the SVD model.

3.5.1 Till Data Results

The autocorrelation of the forecast error shown in Fig. 2 suggests that there is little information in the error. The autocorrelation coefficients lie within the 95% confidence interval bands, suggesting the model error is near random.

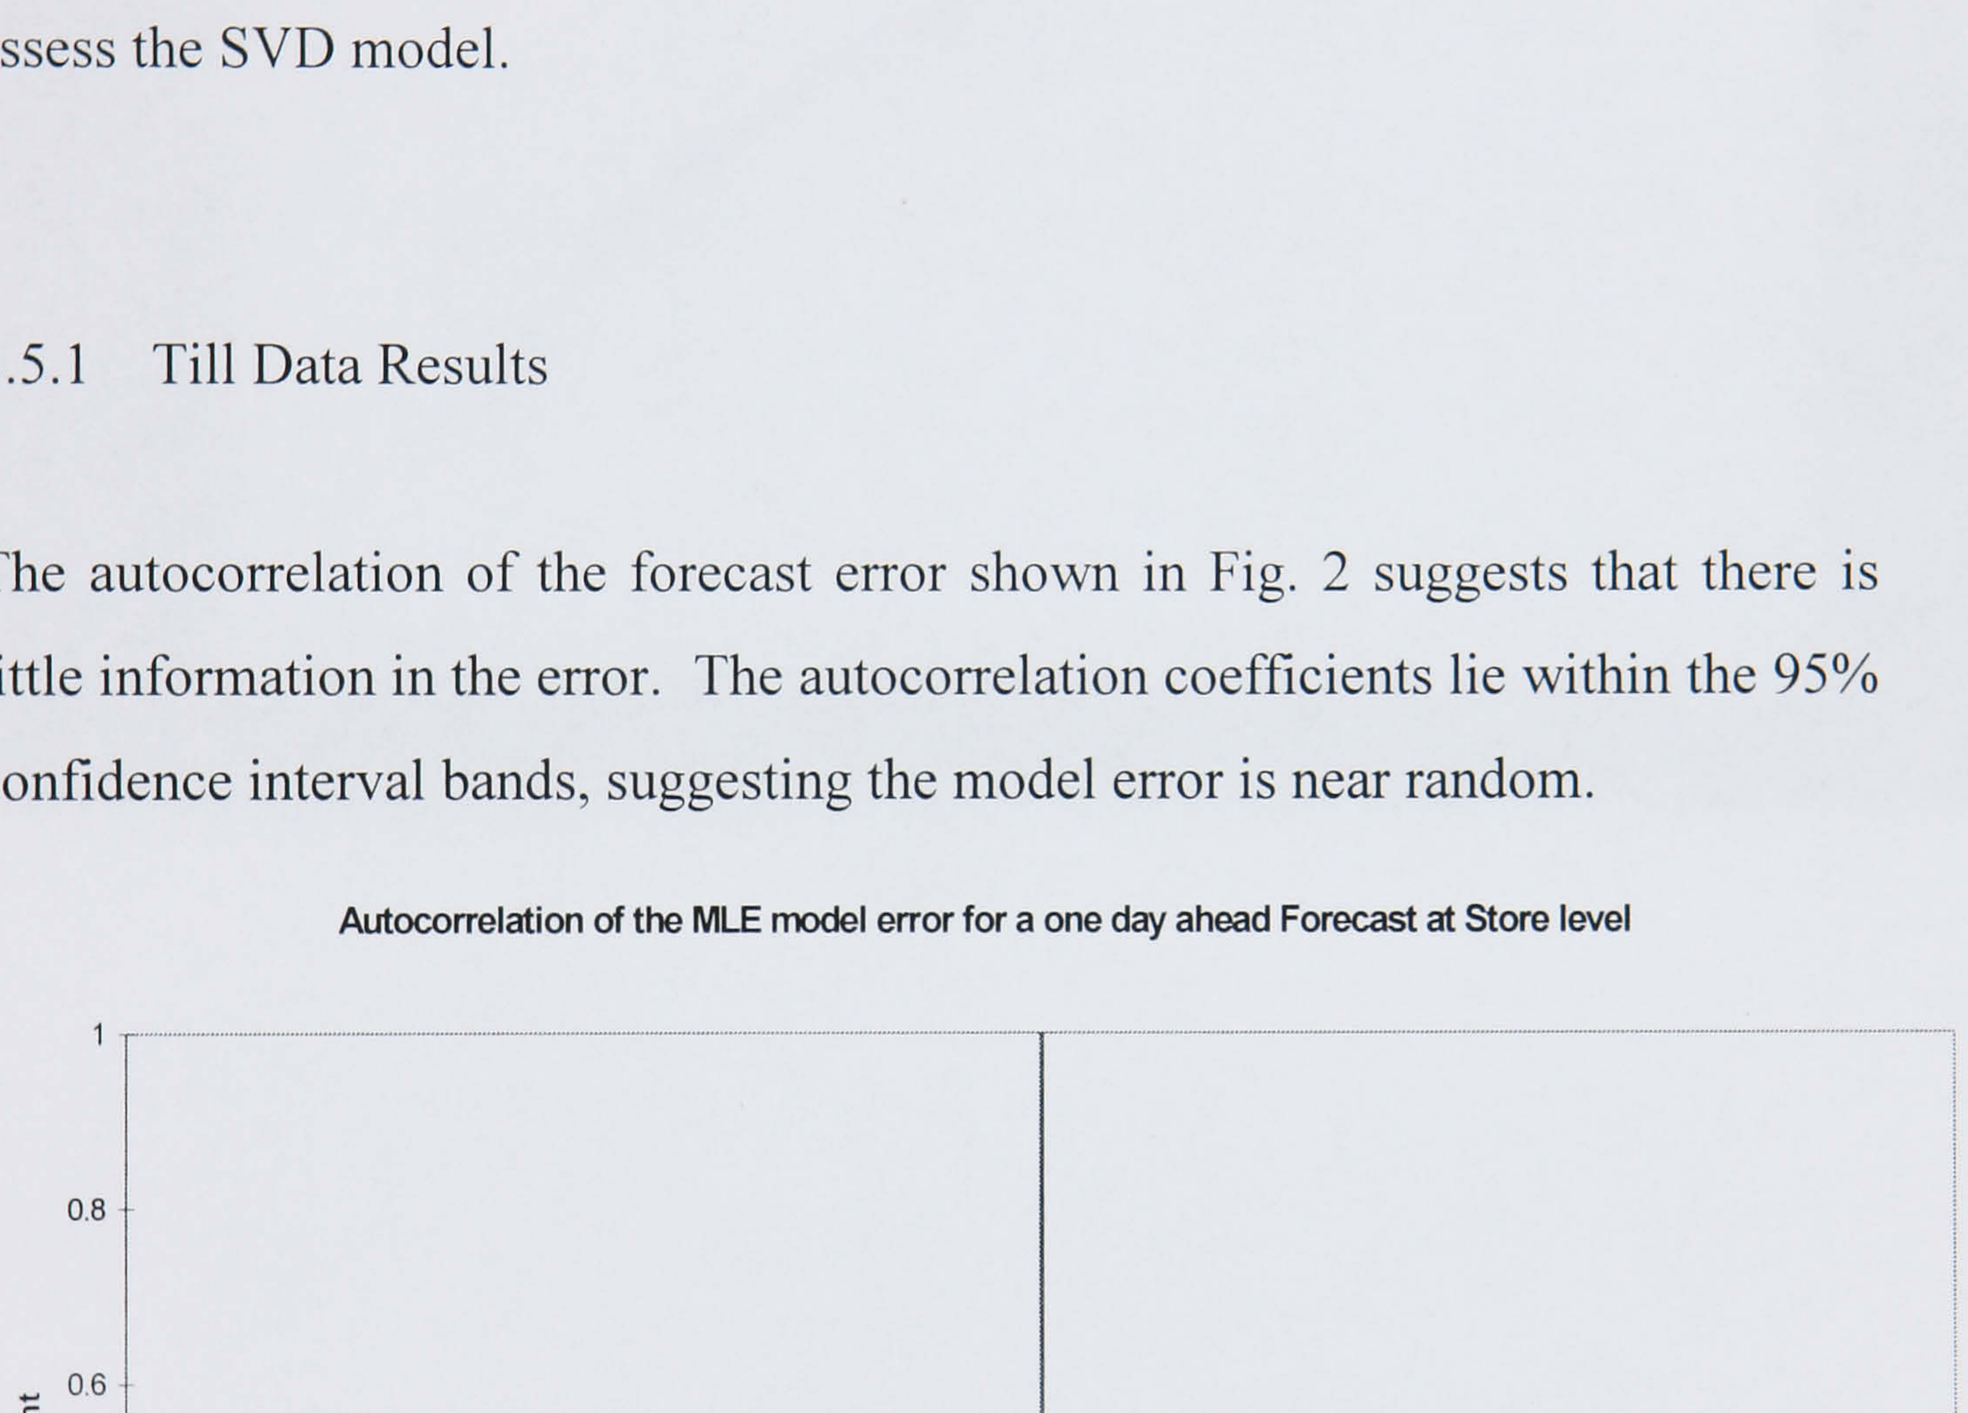


Fig. 2

The Normal Q-Q plot below (Fig. 3), of the weighted residuals suggest that the assumption of a Poisson like error distribution is reasonable.

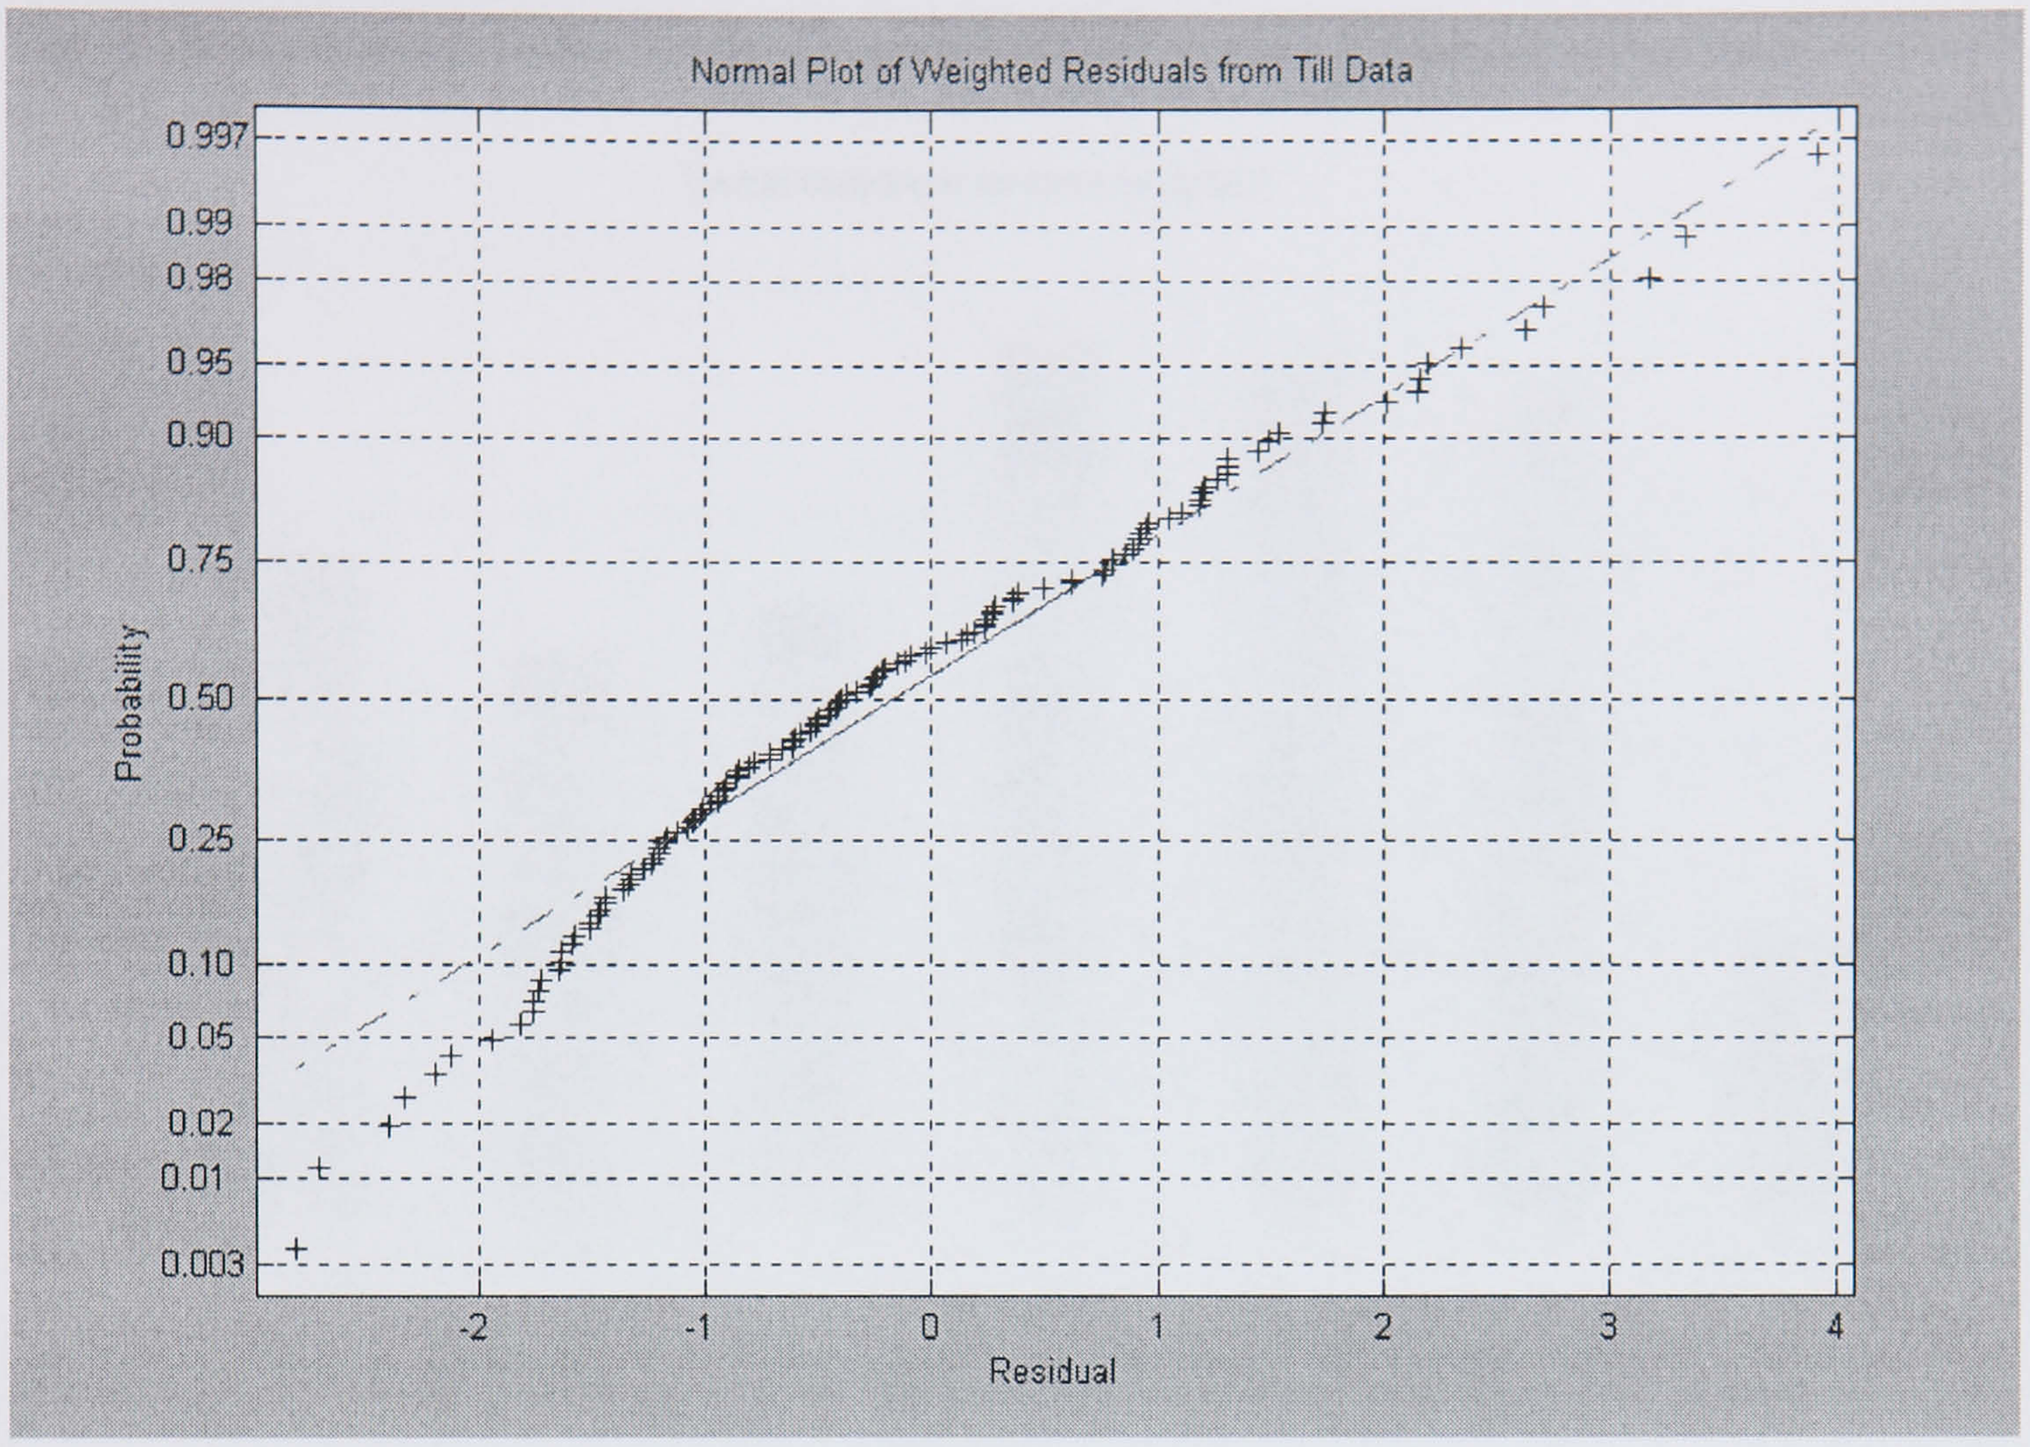


Fig. 3

3.5.2 Store Daily Mode

Fig. 4 illustrates a typical "Daily Fingerprint" for a week's sales at store level

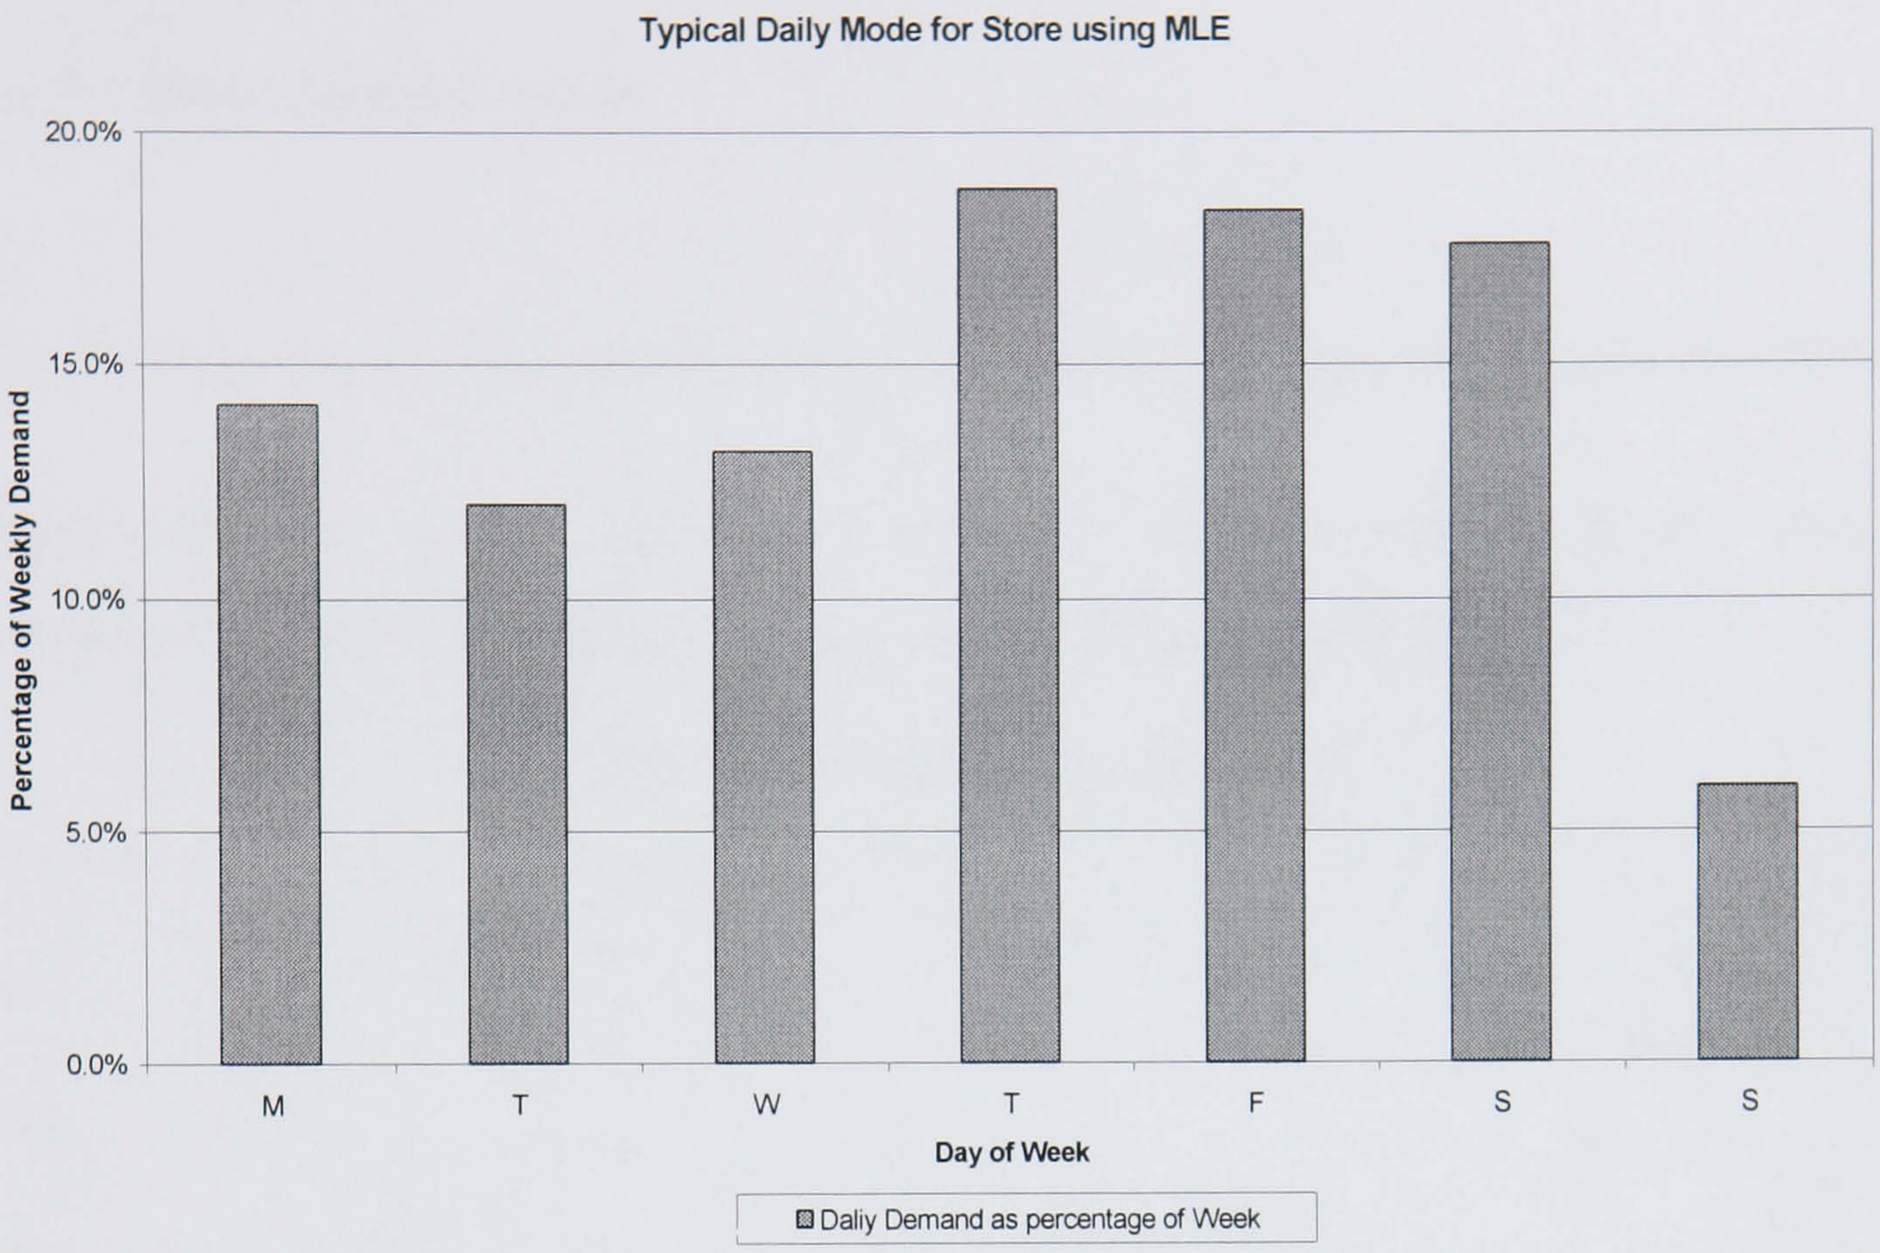


Fig. 4

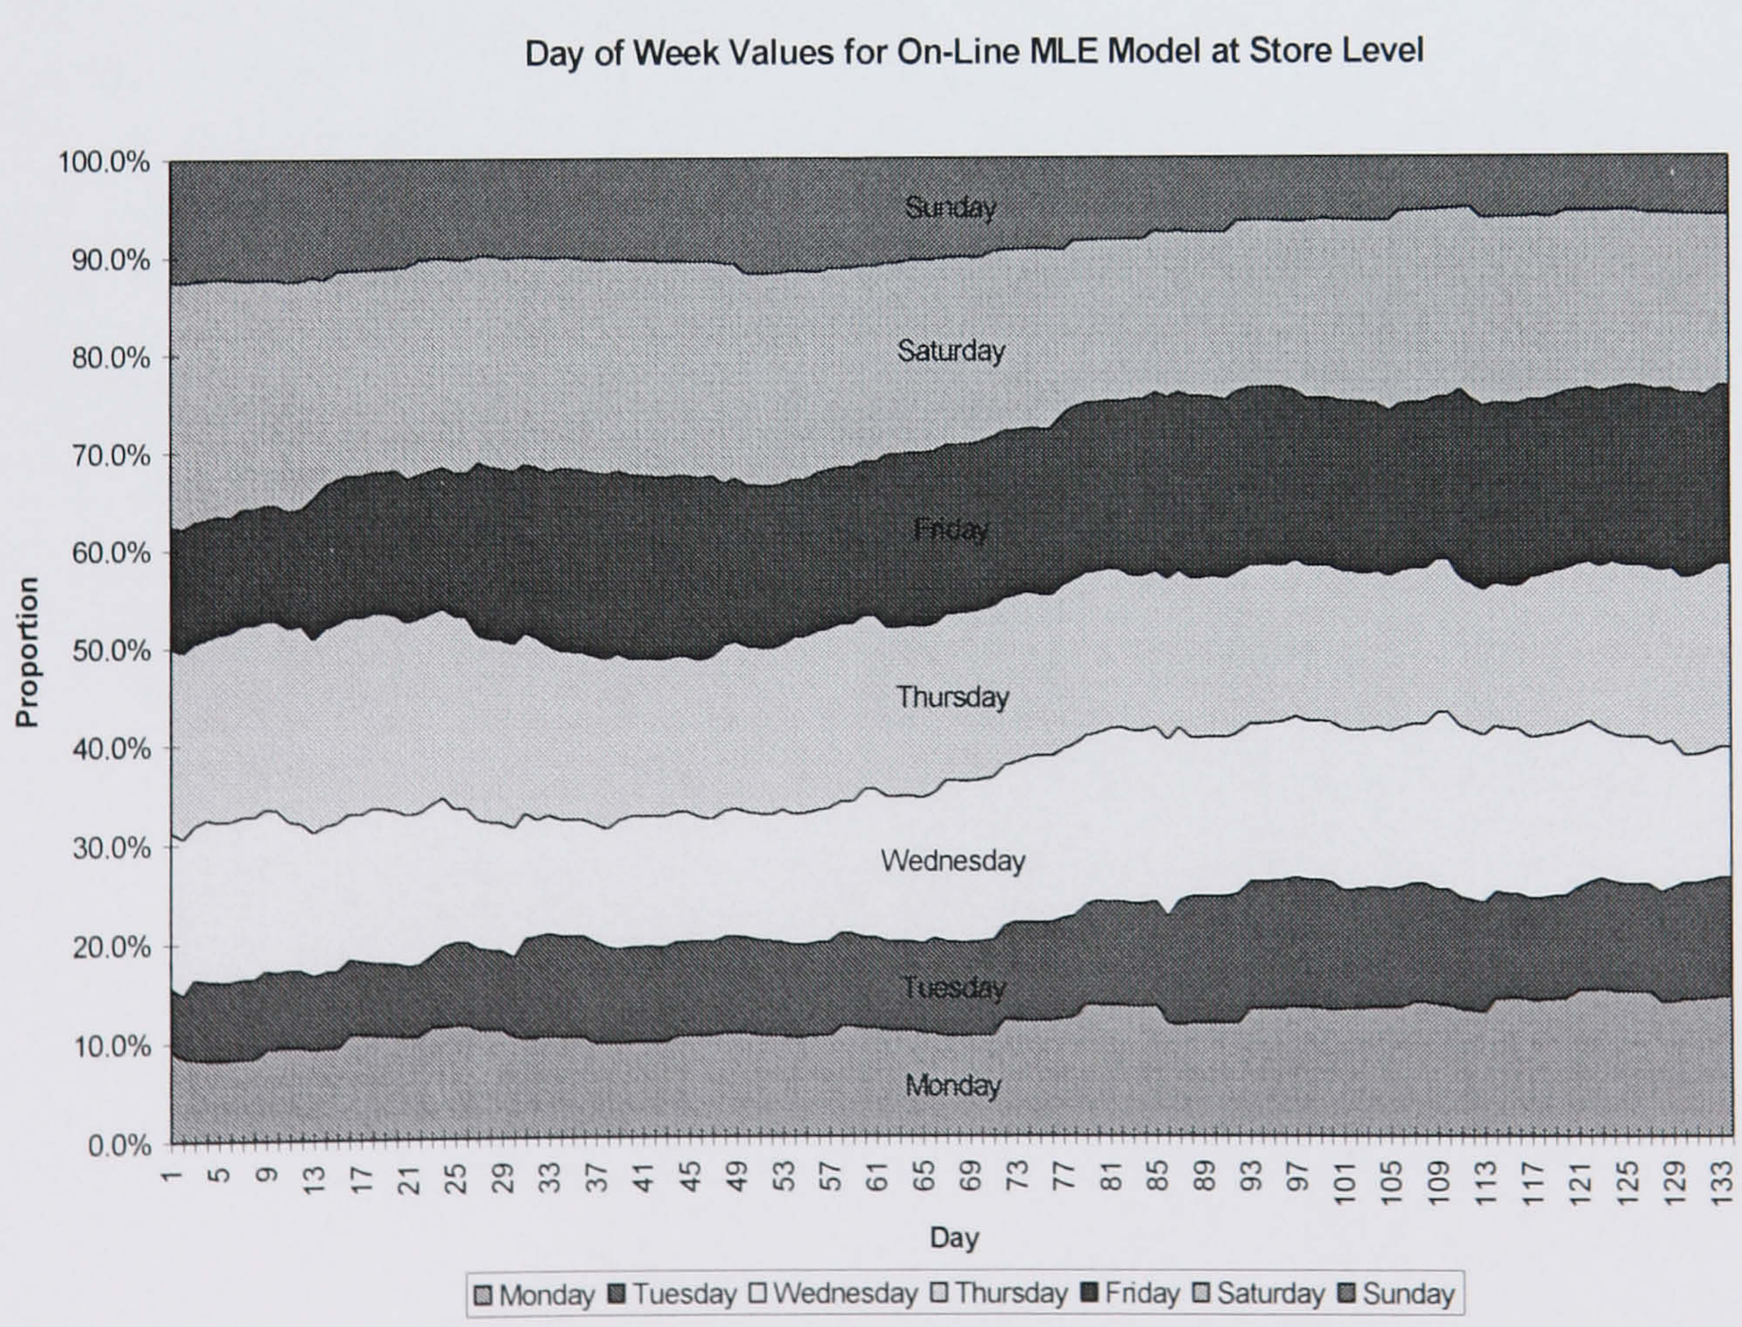


Fig. 5

These daily values are reasonably stable as can be seen in Fig. 5, which shows the fingerprint on a daily basis throughout the period tested.

3.5.3 Store Seasonal Mode

The seasonal mode $\left(\hat{\lambda}s\right)$ shown in Fig. 6 is for the dry soup used in the previous chapter at store level. The results obtained appear similar in shape when compared to the SVD seasonal mode in the previous chapter (Fig. 5).

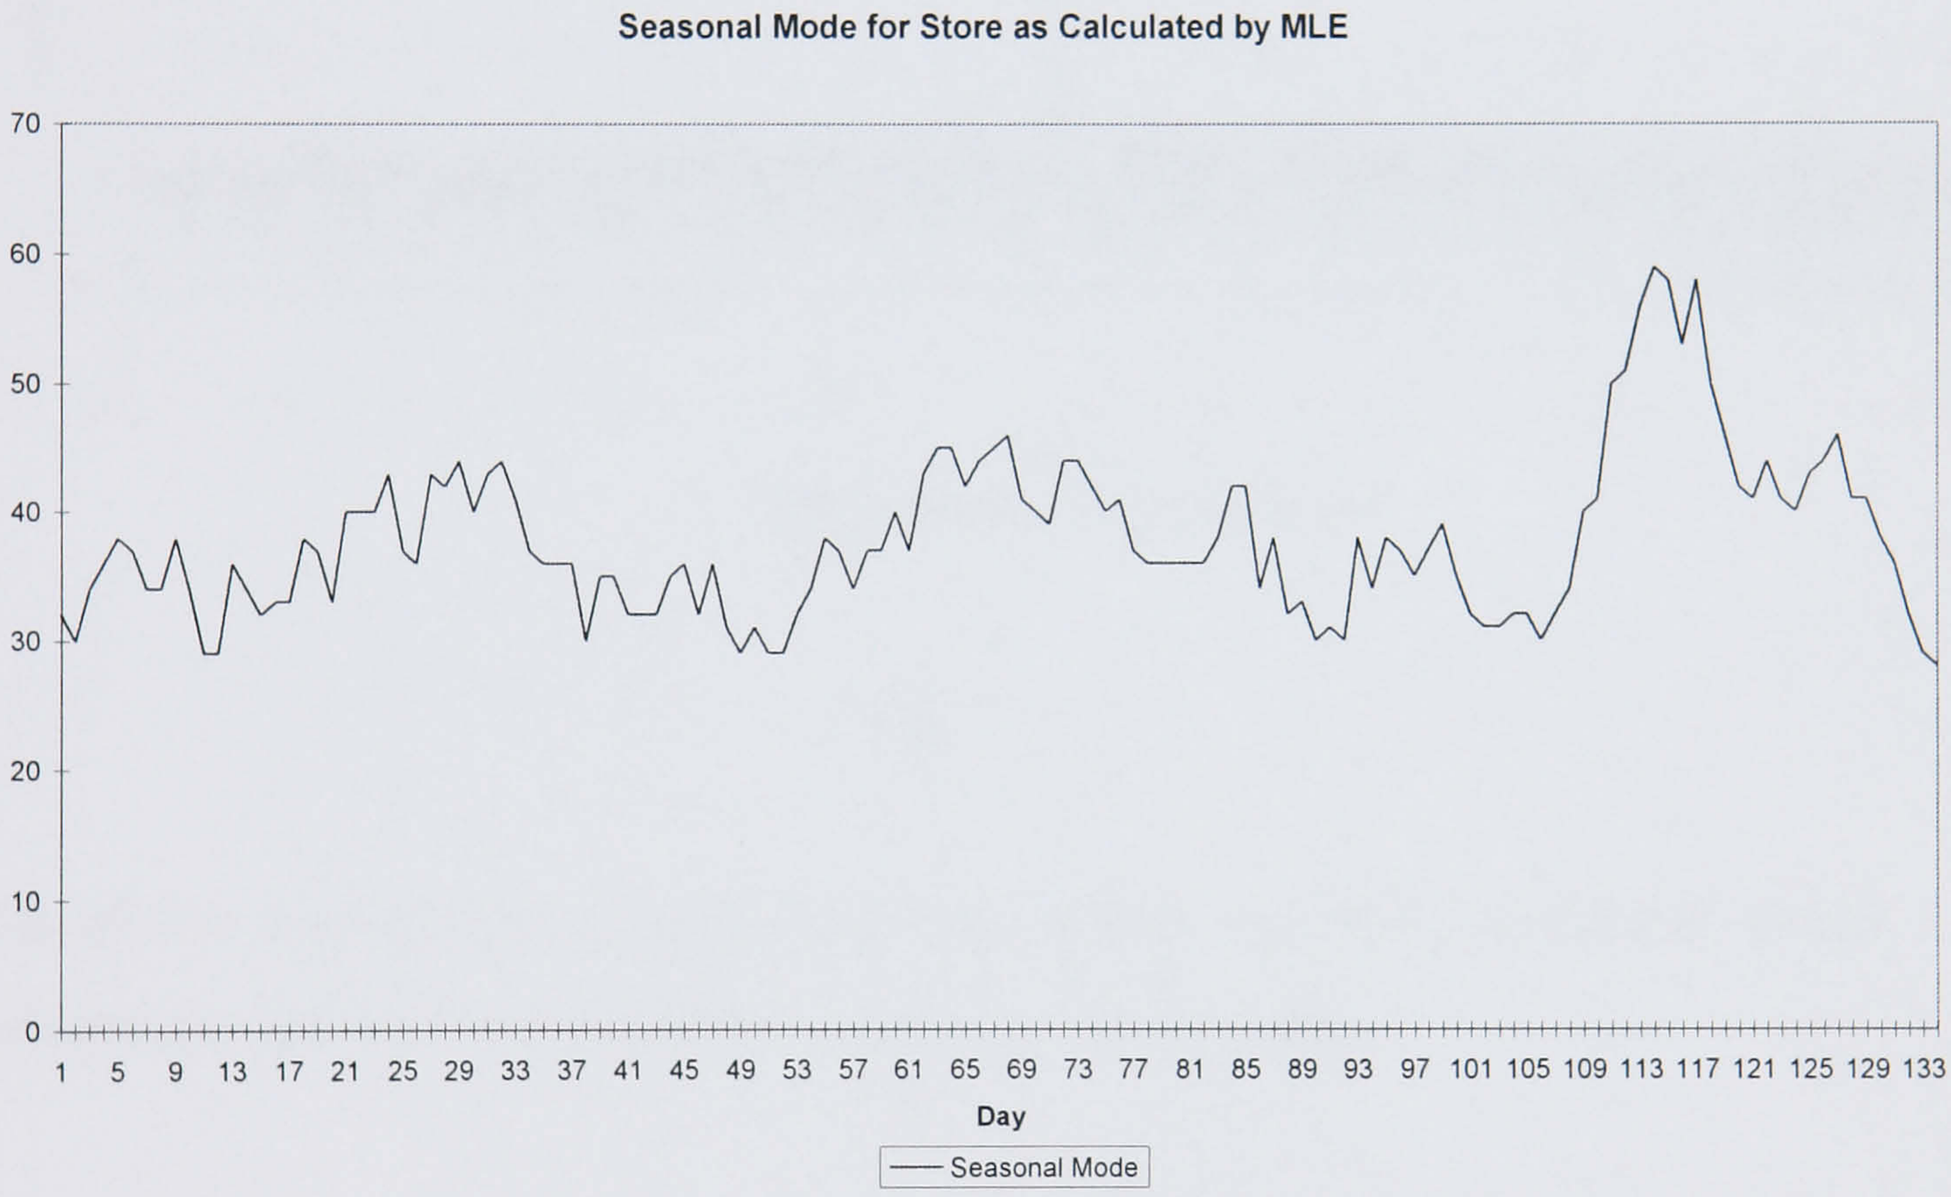


Fig. 6

3.5.4 RDC Data Results

The autocorrelation of the MLE model error is shown in Fig. 7.

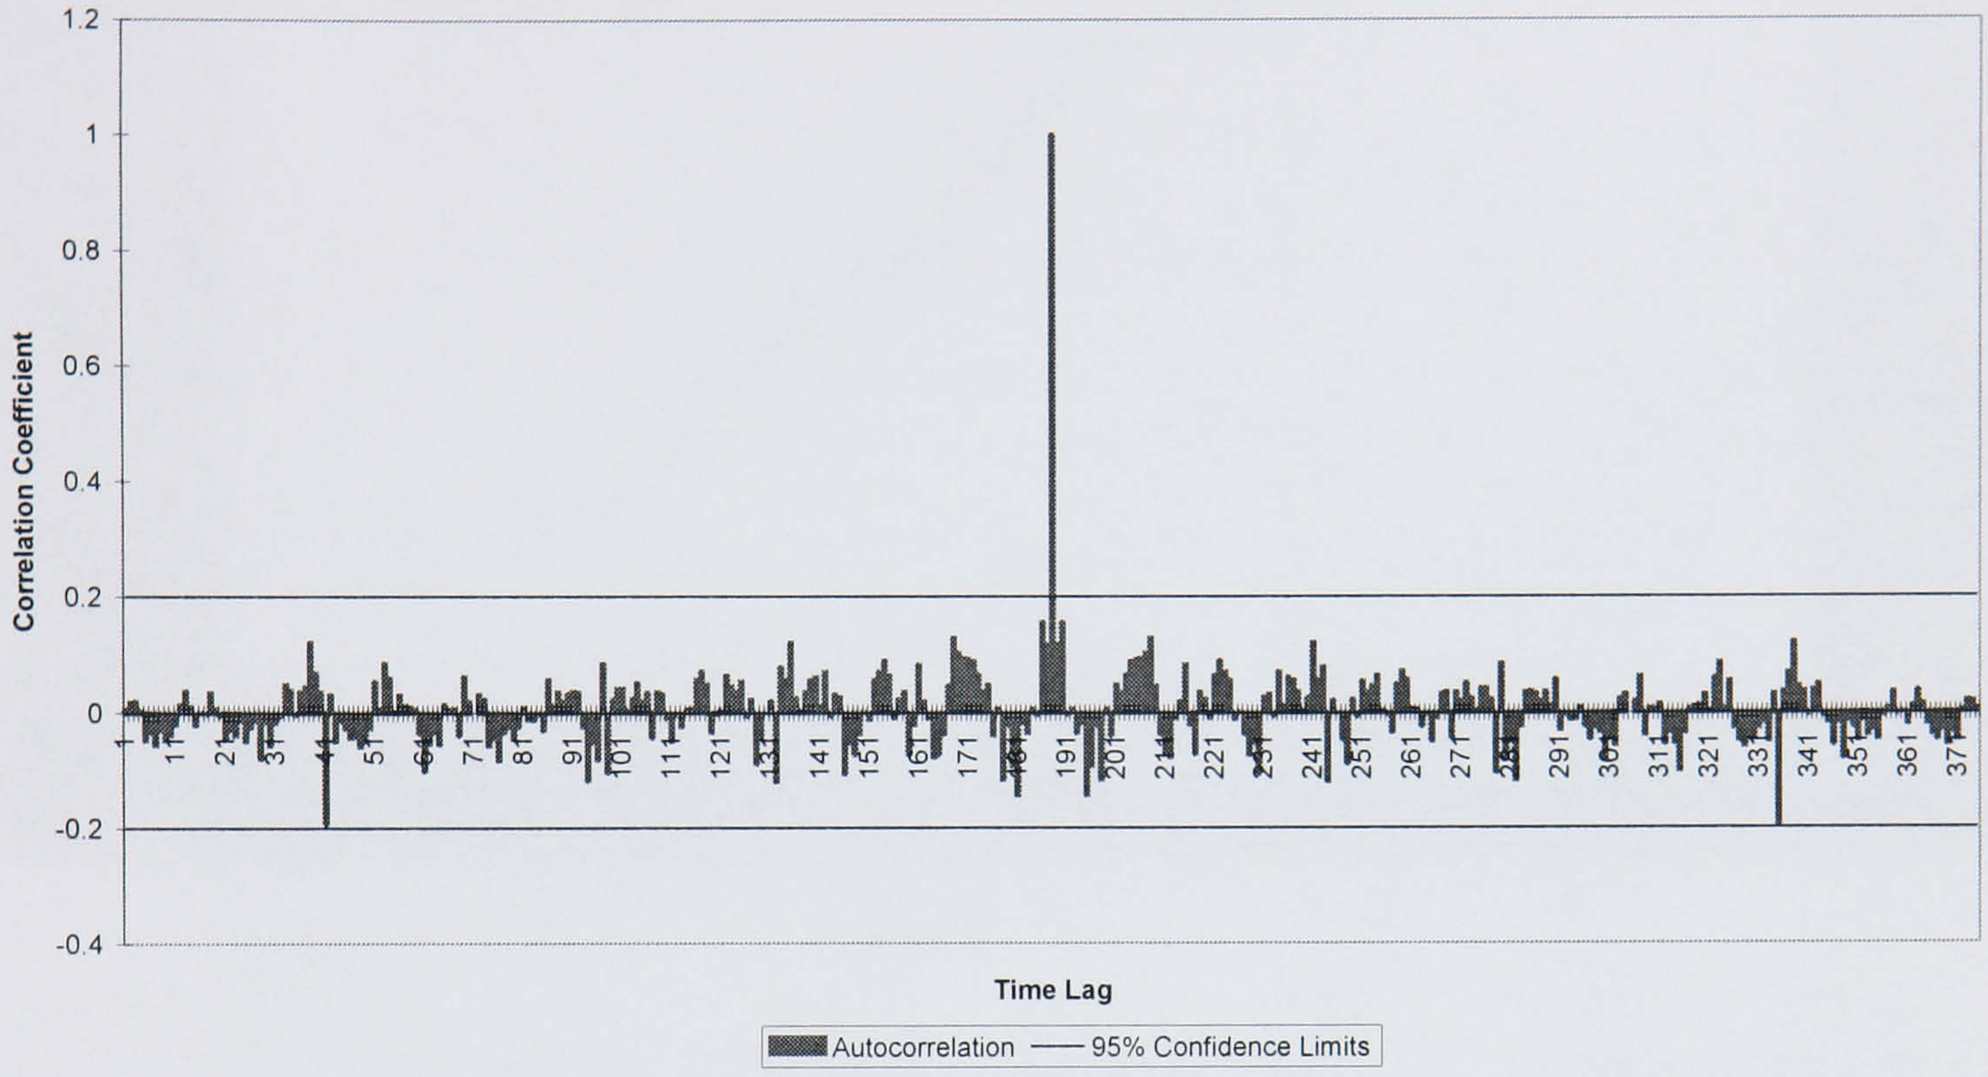


Fig. 7

All of the autocorrelation coefficients are within the 95% confidence limits suggesting very little model information in the model error.

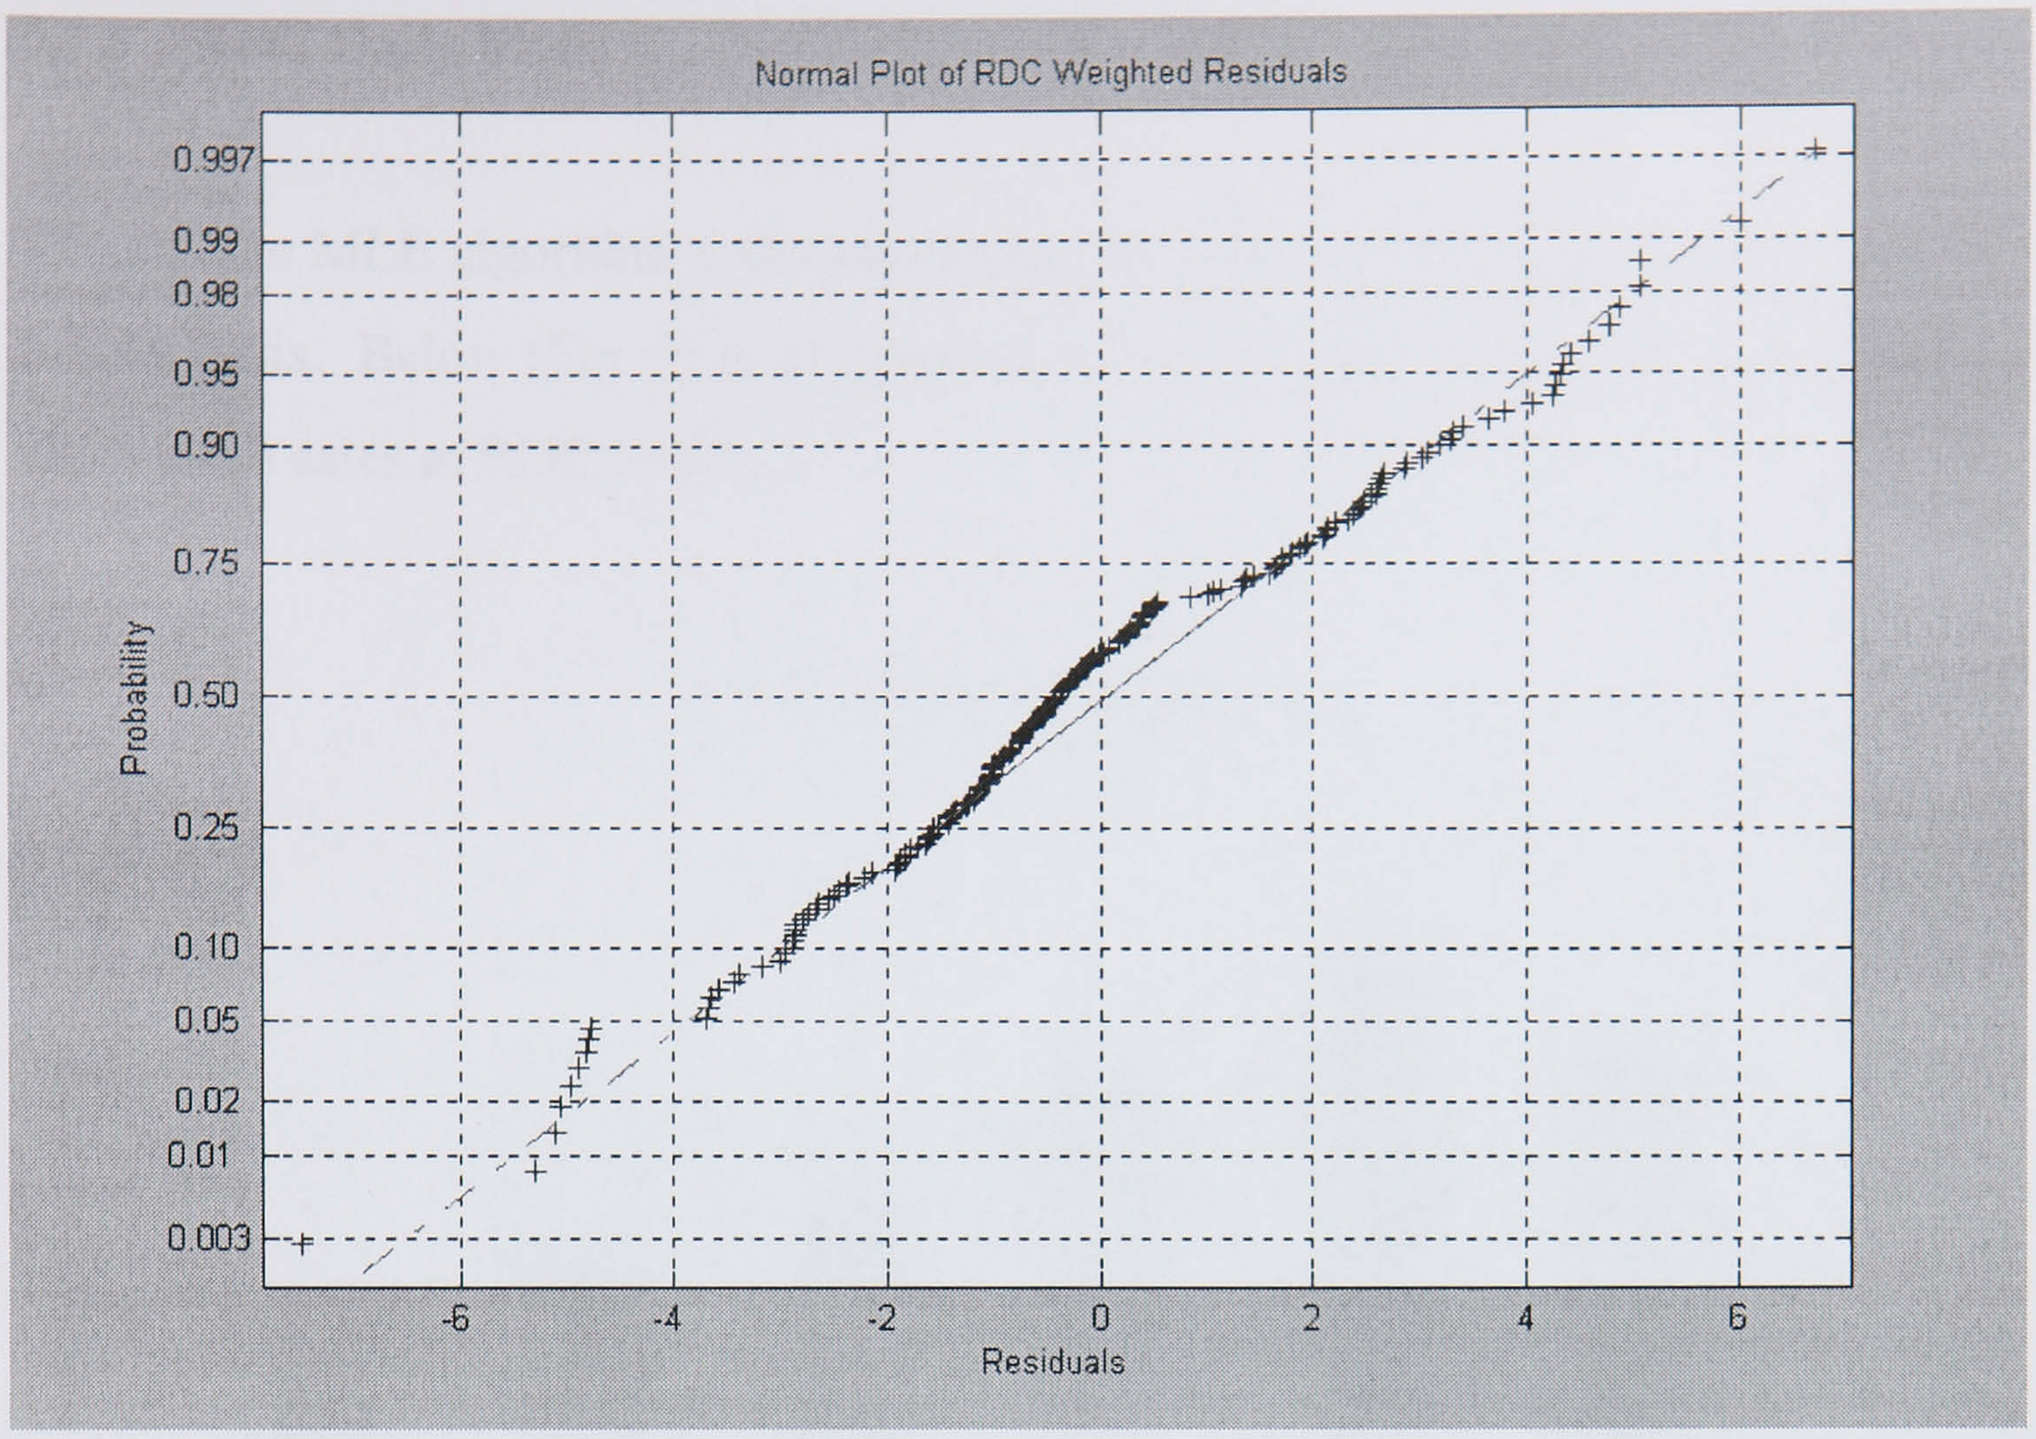


Fig. 8

The Normal Q-Q plot of the weighted residuals (Fig. 8), suggests that the assumption of a Poisson like error distribution is still reasonable.

3.5.5 RDC Daily Mode

The on-line MLE algorithm calculates values for the daily and weekly modes on a daily basis. Below (Fig. 9) is a typical set of values for the "Daily Fingerprint" of a week's sales at RDC level.

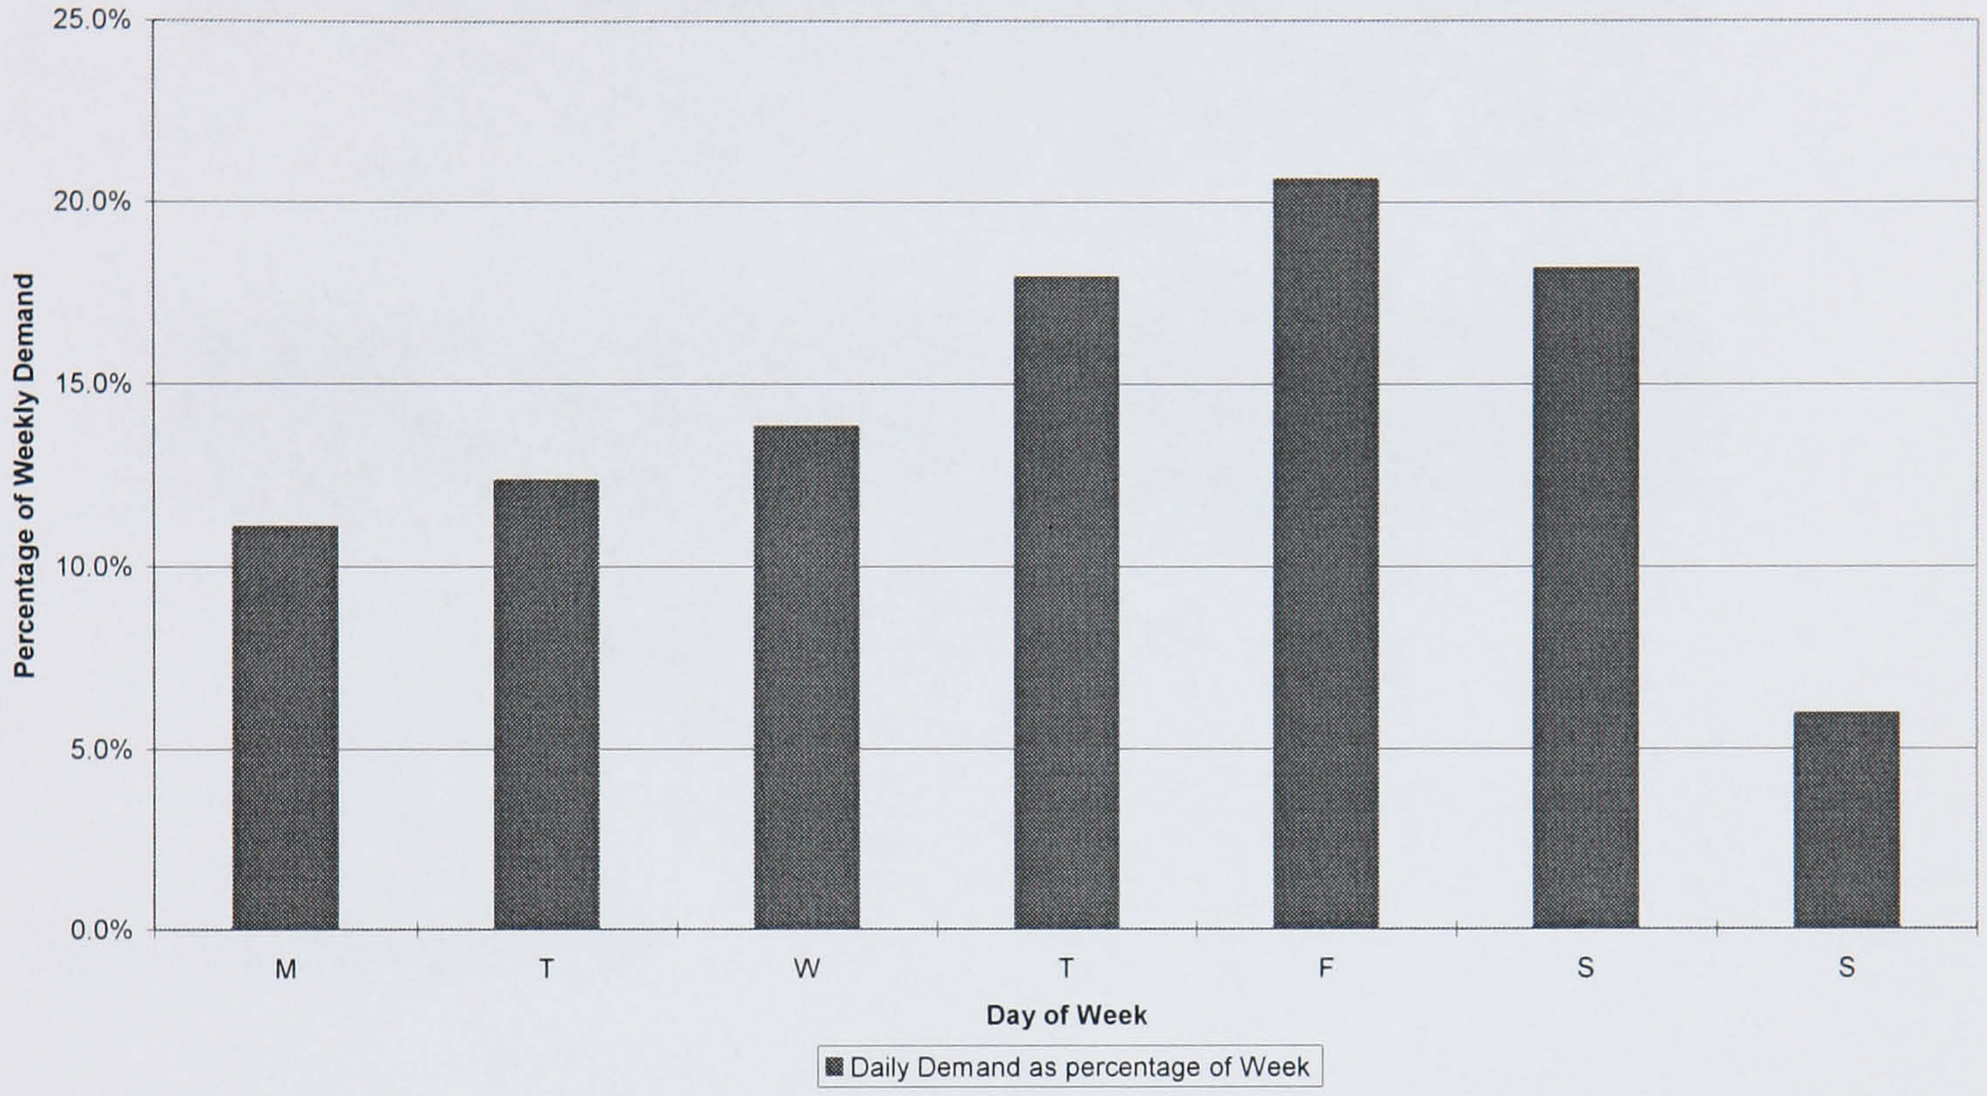


Fig. 9

These daily values are reasonably stable for this example as illustrated by Fig. 10 which shows the daily values through out the period tested.

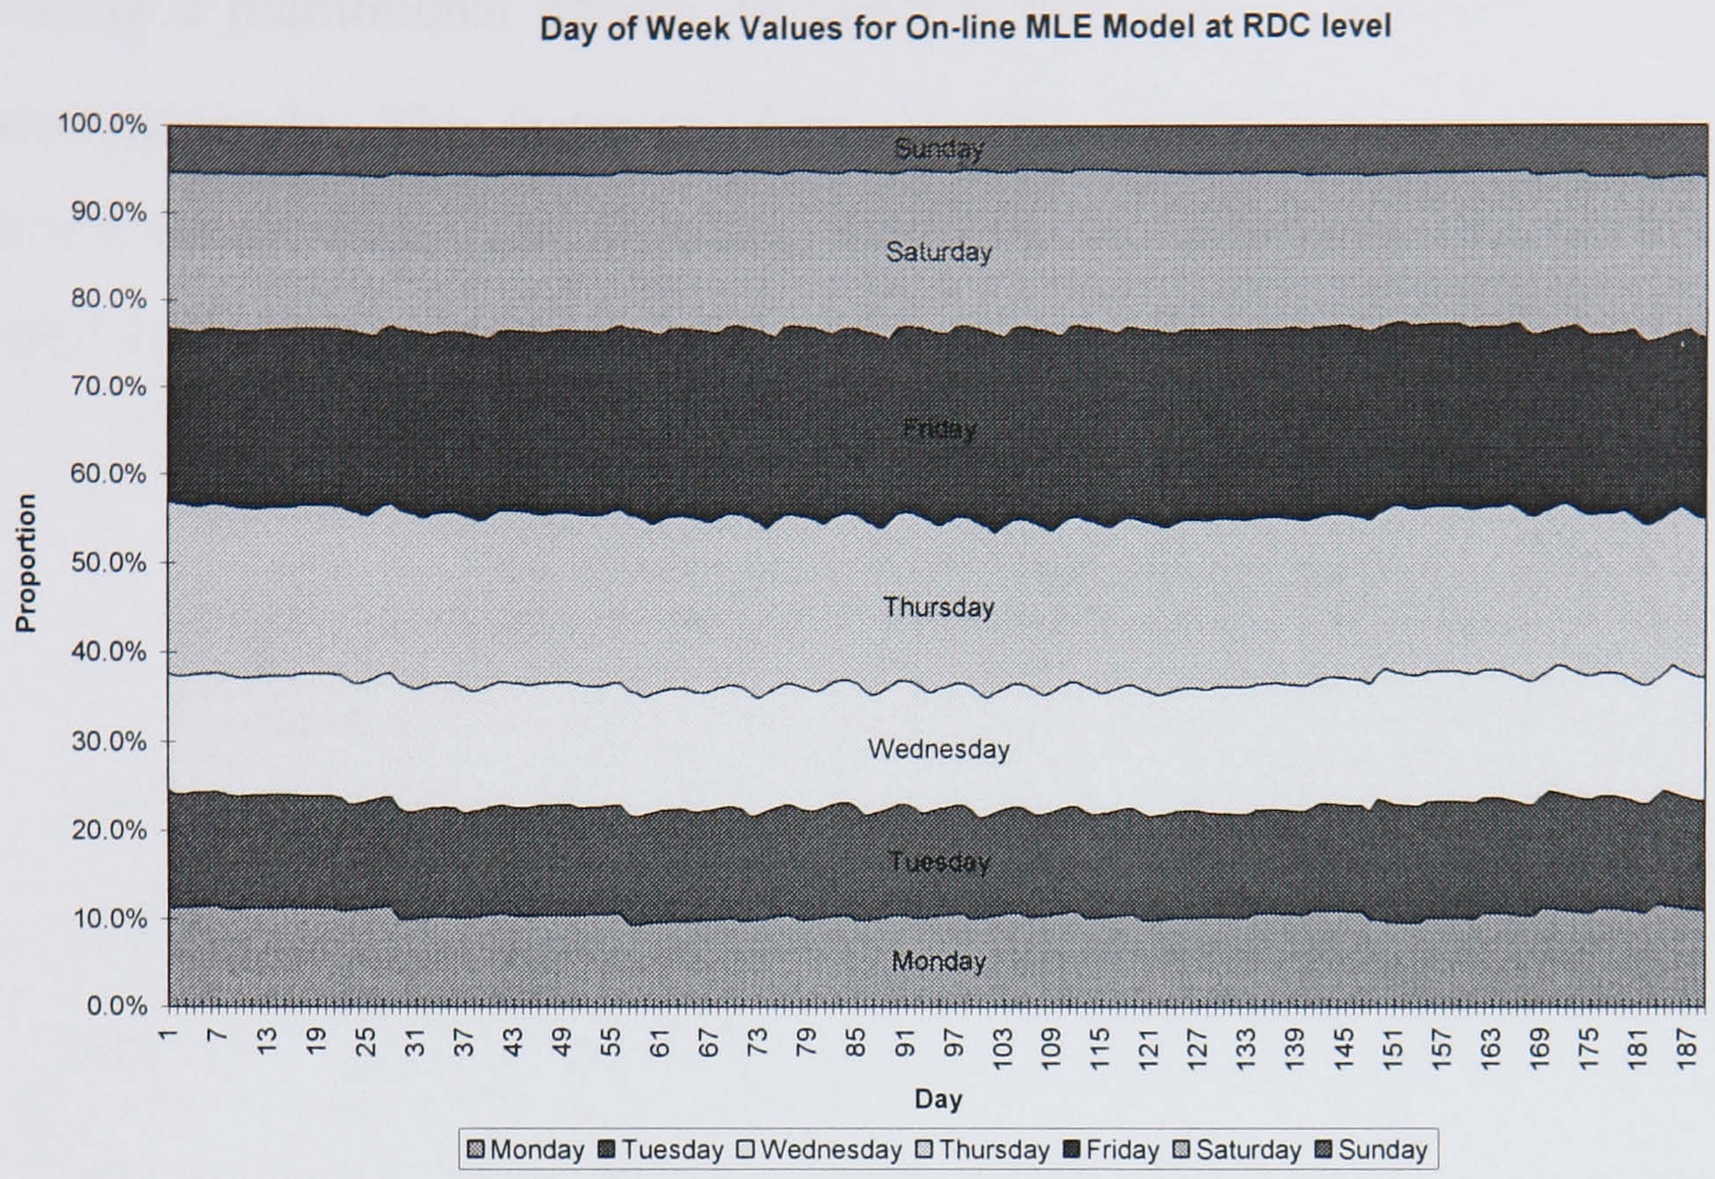


Fig. 10

3.5.6 RDC Seasonal Mode

Fig. 11 is a plot of the Seasonal Mode of the demand data. The shape is similar to Fig. 10 of chapter 2. The results are smoother due to the effects of daily updating.

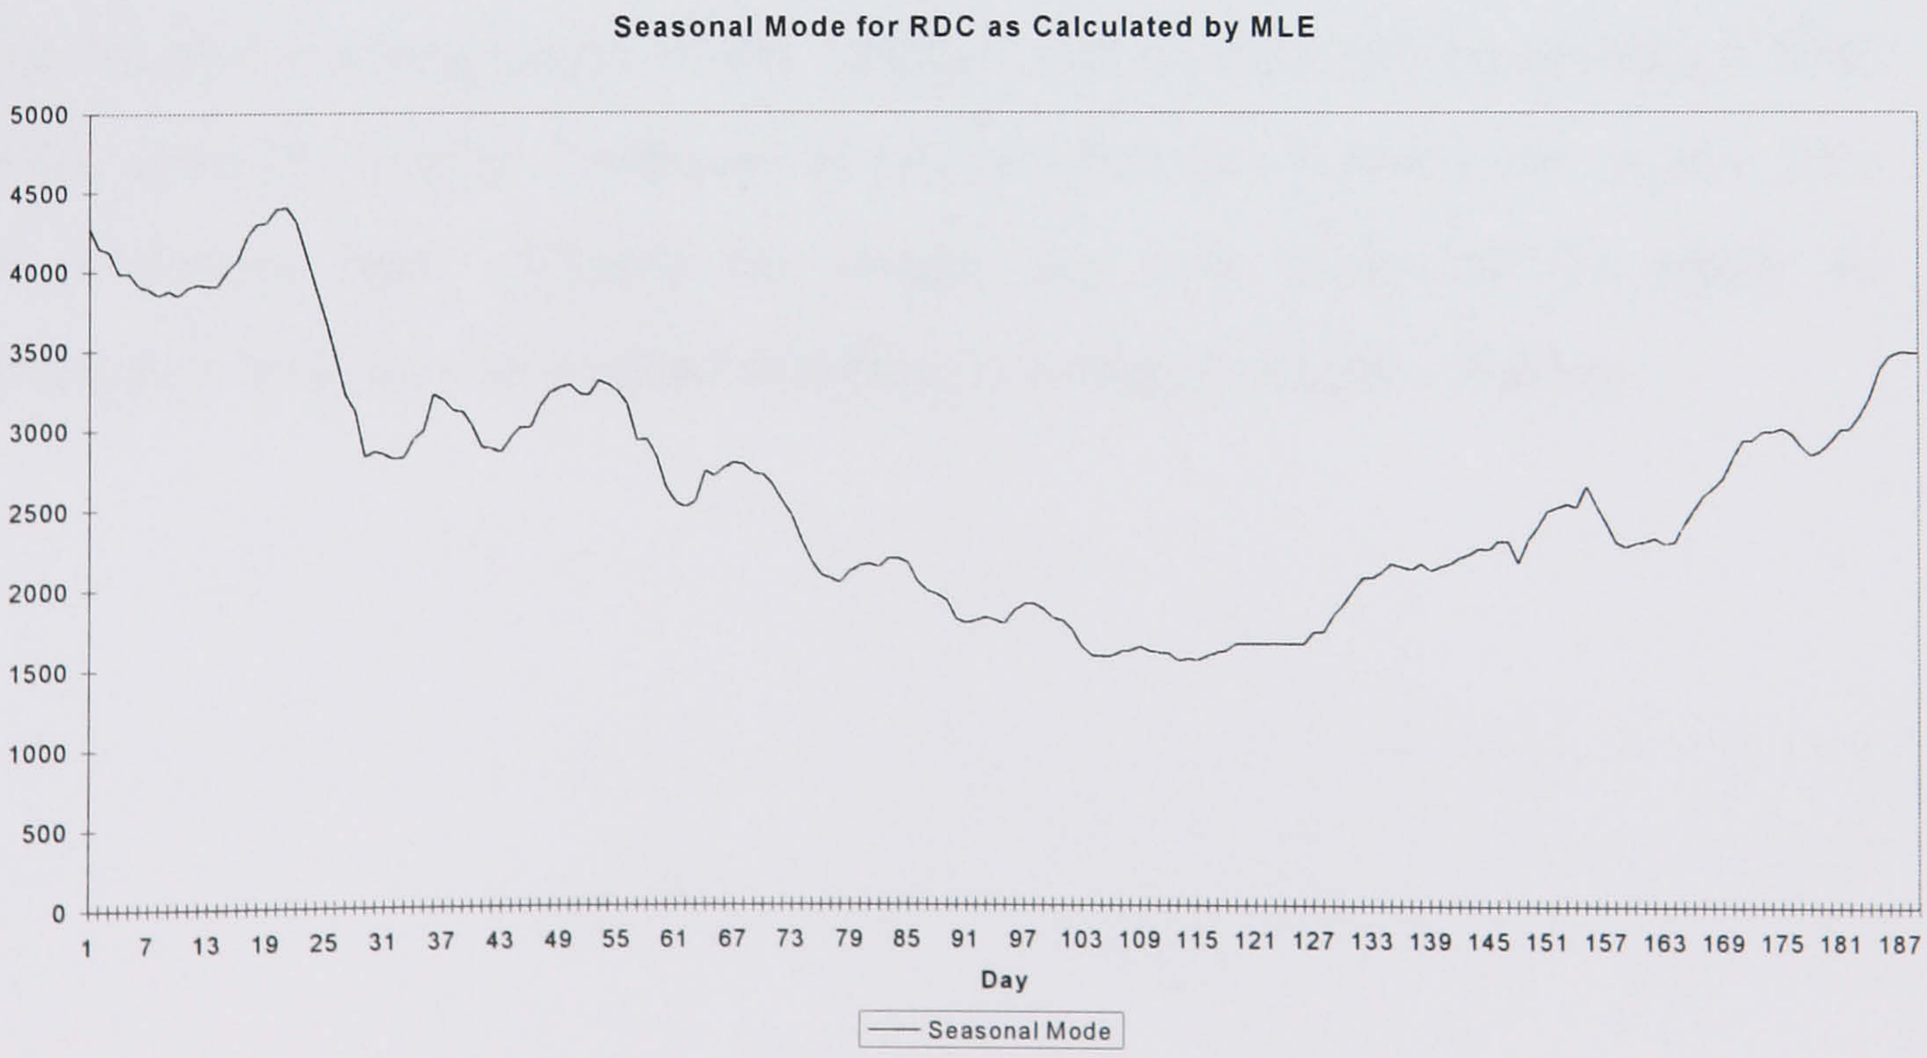


Fig. 11

3.5.7 Event Detection - Promotional Mode Adaption

During a promotional period, an uplift factor $(\hat{p})$ is introduced to boost the seasonal mode. This factor is a forecast of the expected demand response to the promotional offer. The following sequence of graphs show the effect of the uplift factor on the forecast.

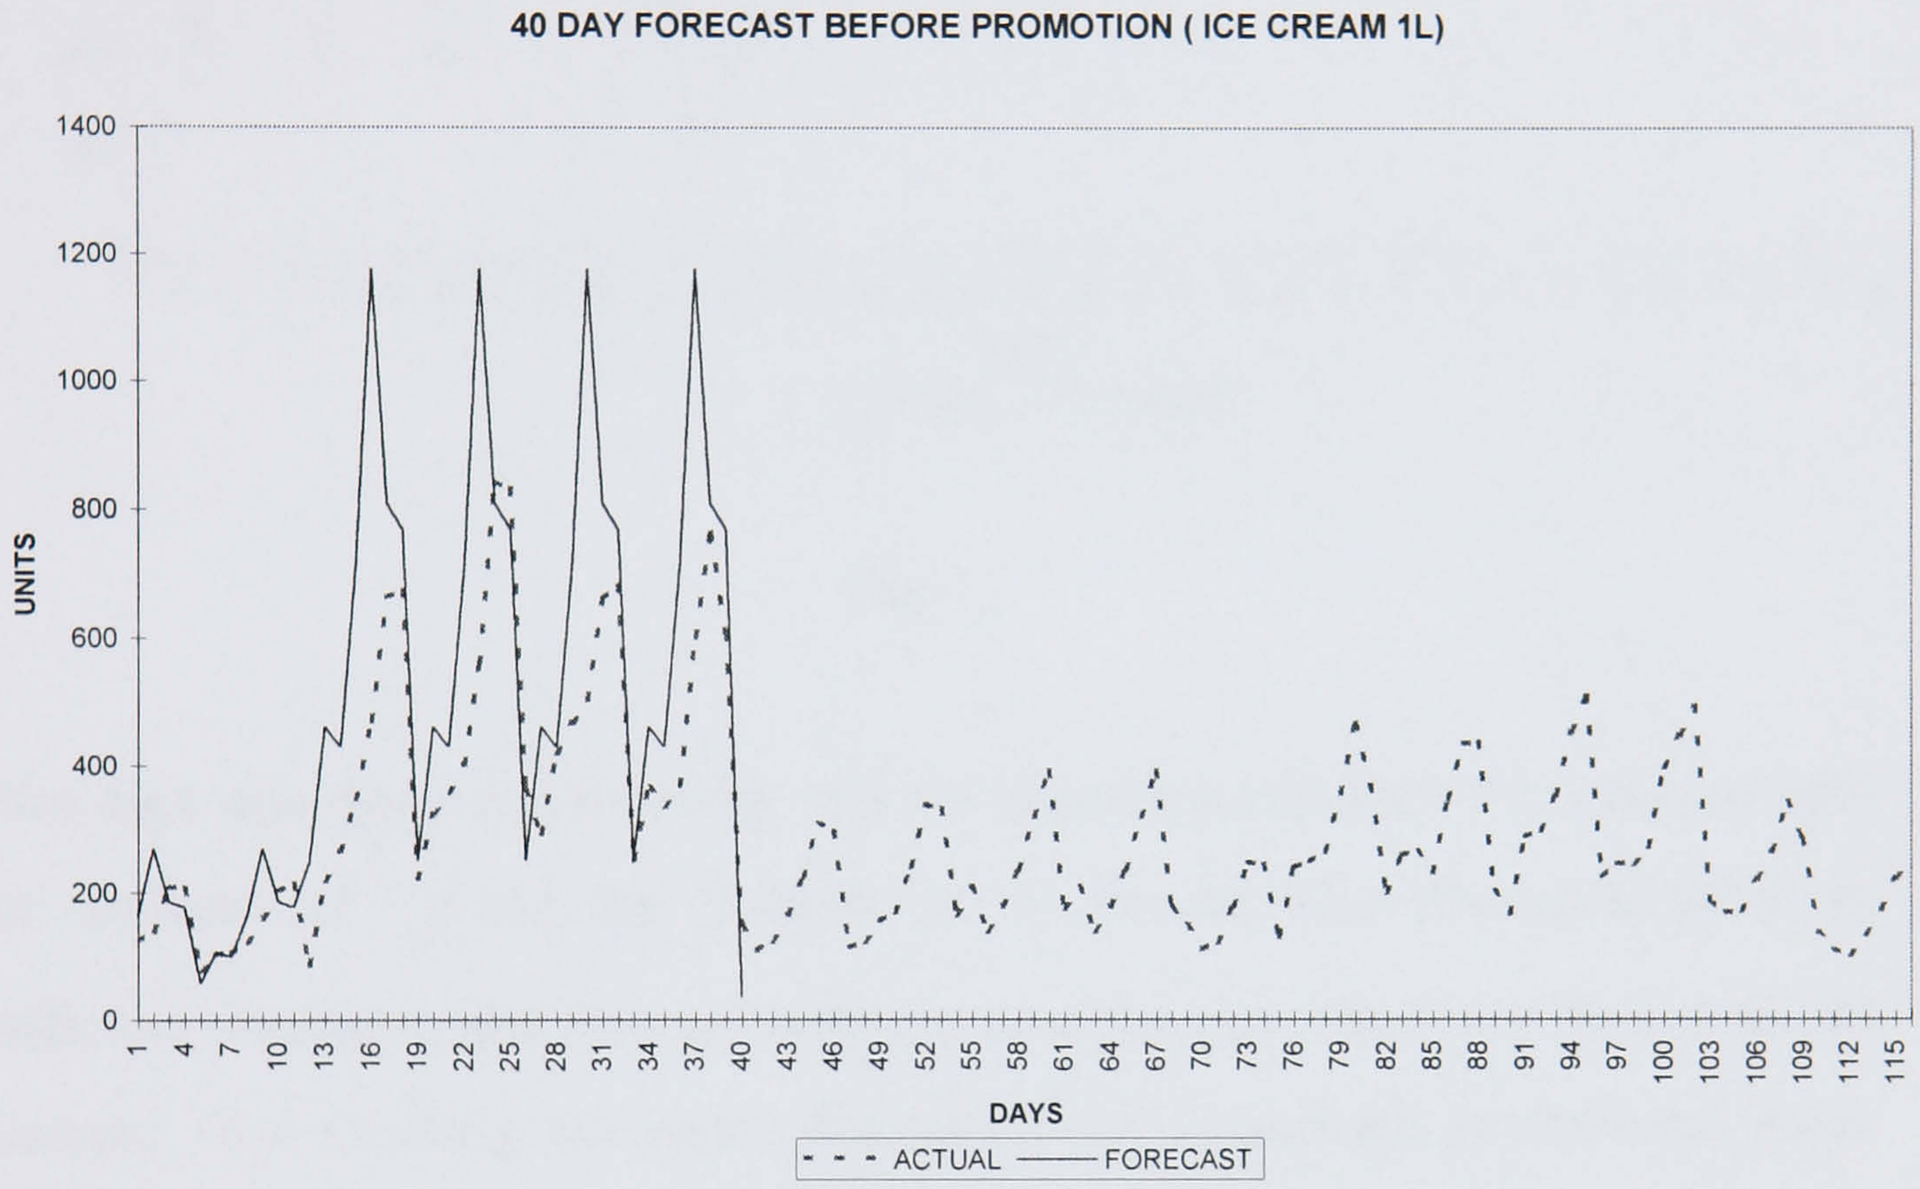


Fig. 12

Fig. 12 shows a long range 40 day forecast over a four week promotion at RDC level, using the retailer's estimate of promotional uplift, nearly two weeks prior to promotion start. Clearly the retailer has over estimated the effect the promotion will have on demand and there is a danger of over-stocking.

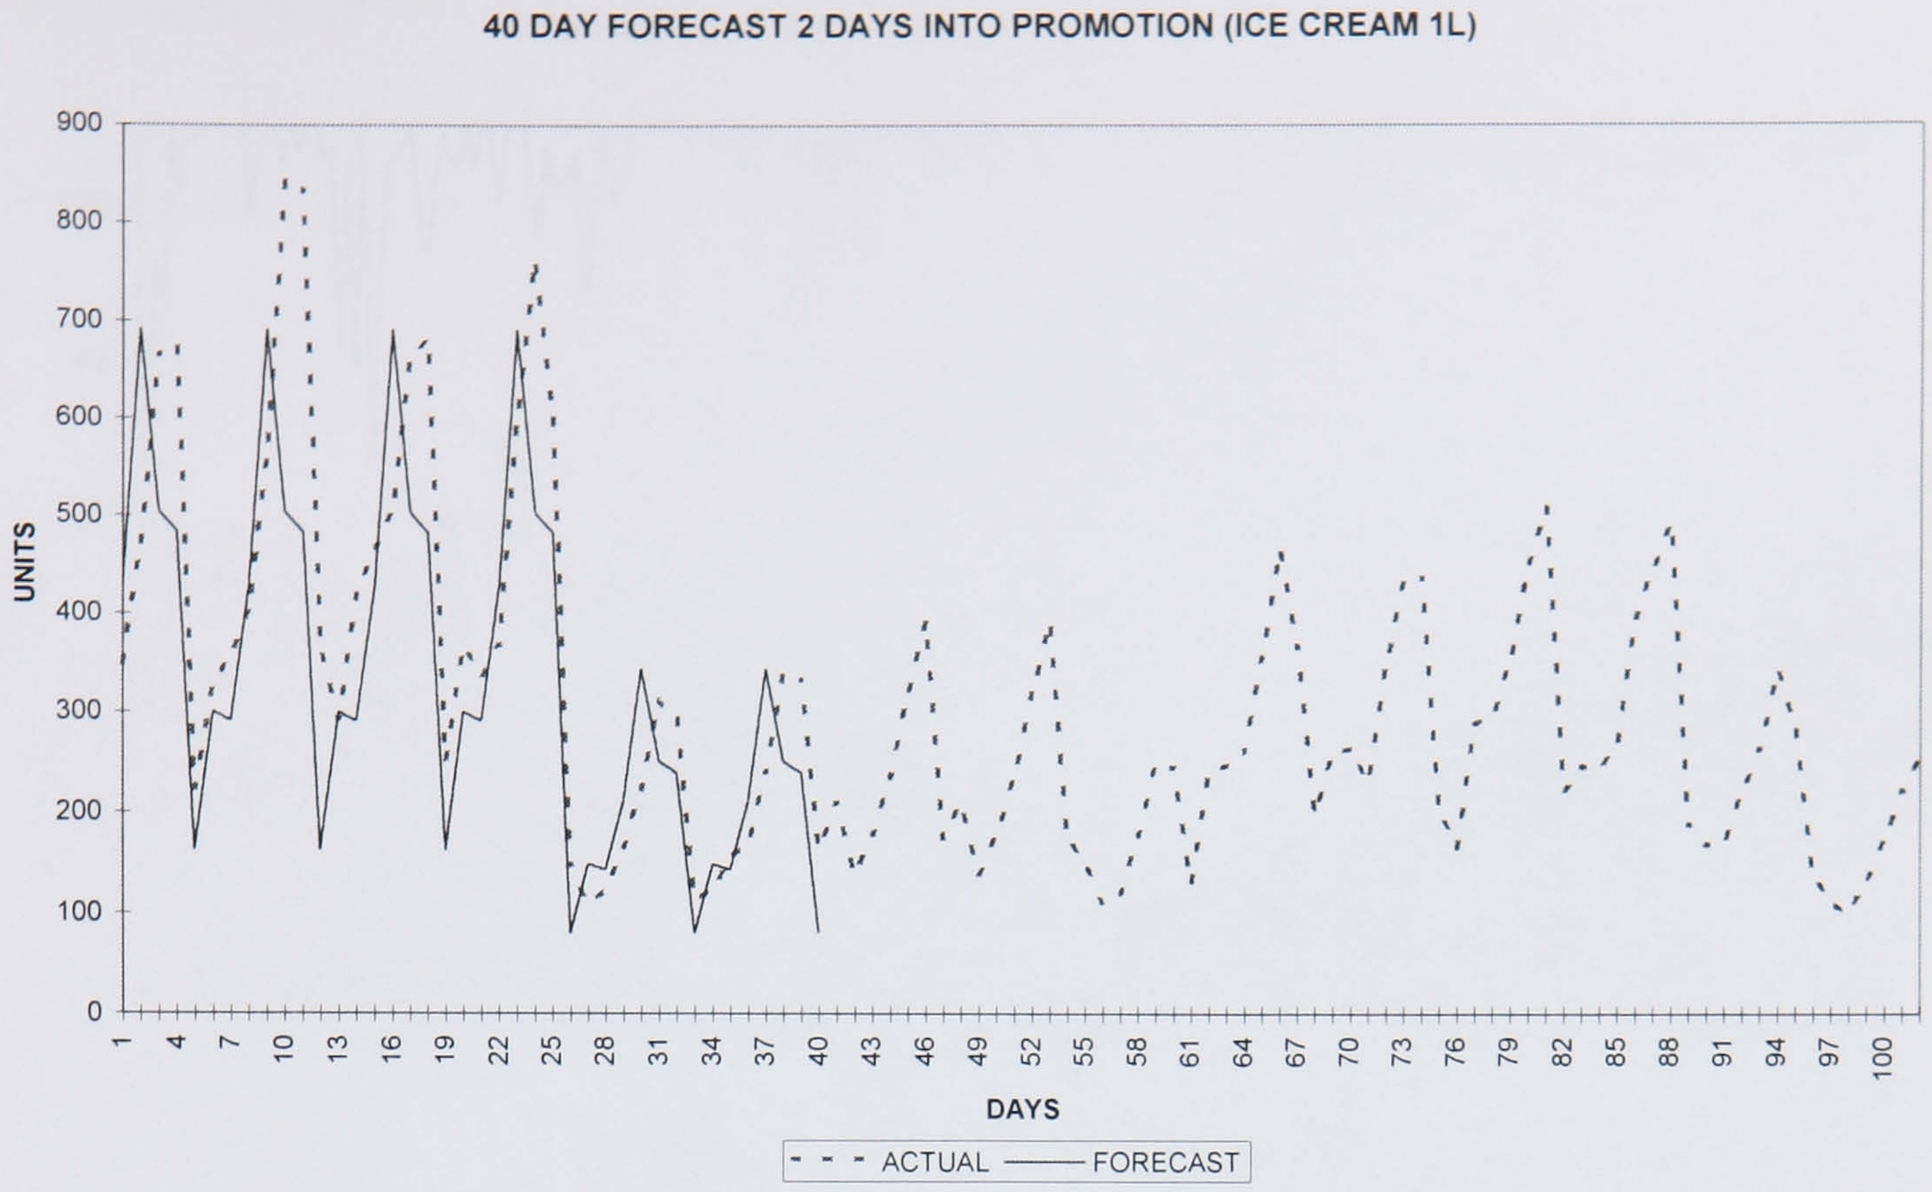


Fig. 13

Two days into the promotion (Fig. 13), the algorithm (equation 18.0) has revised the estimate of p and has "locked on" to the actual promotional uplift in sufficient time to amend forward orders to accommodate the lower than expected demand, thus avoiding the losses associated with redundant promotional stock which is frequently specially packaged. The time to detect a promotional mismatch is clearly a significant improvement on the current reaction period of approximately 14 days experienced by Tesco. The end of the four week promotion can be seen on day 26. Forecasts after this date have the uplift factor removed.

3.5.8 Event Detection – Demand Level Change

Fig. 14 shows the demand for a soup at RDC level and the adaptive smoothing factor α of Algorithm 2, which reacts to unexpected changes in demand level.

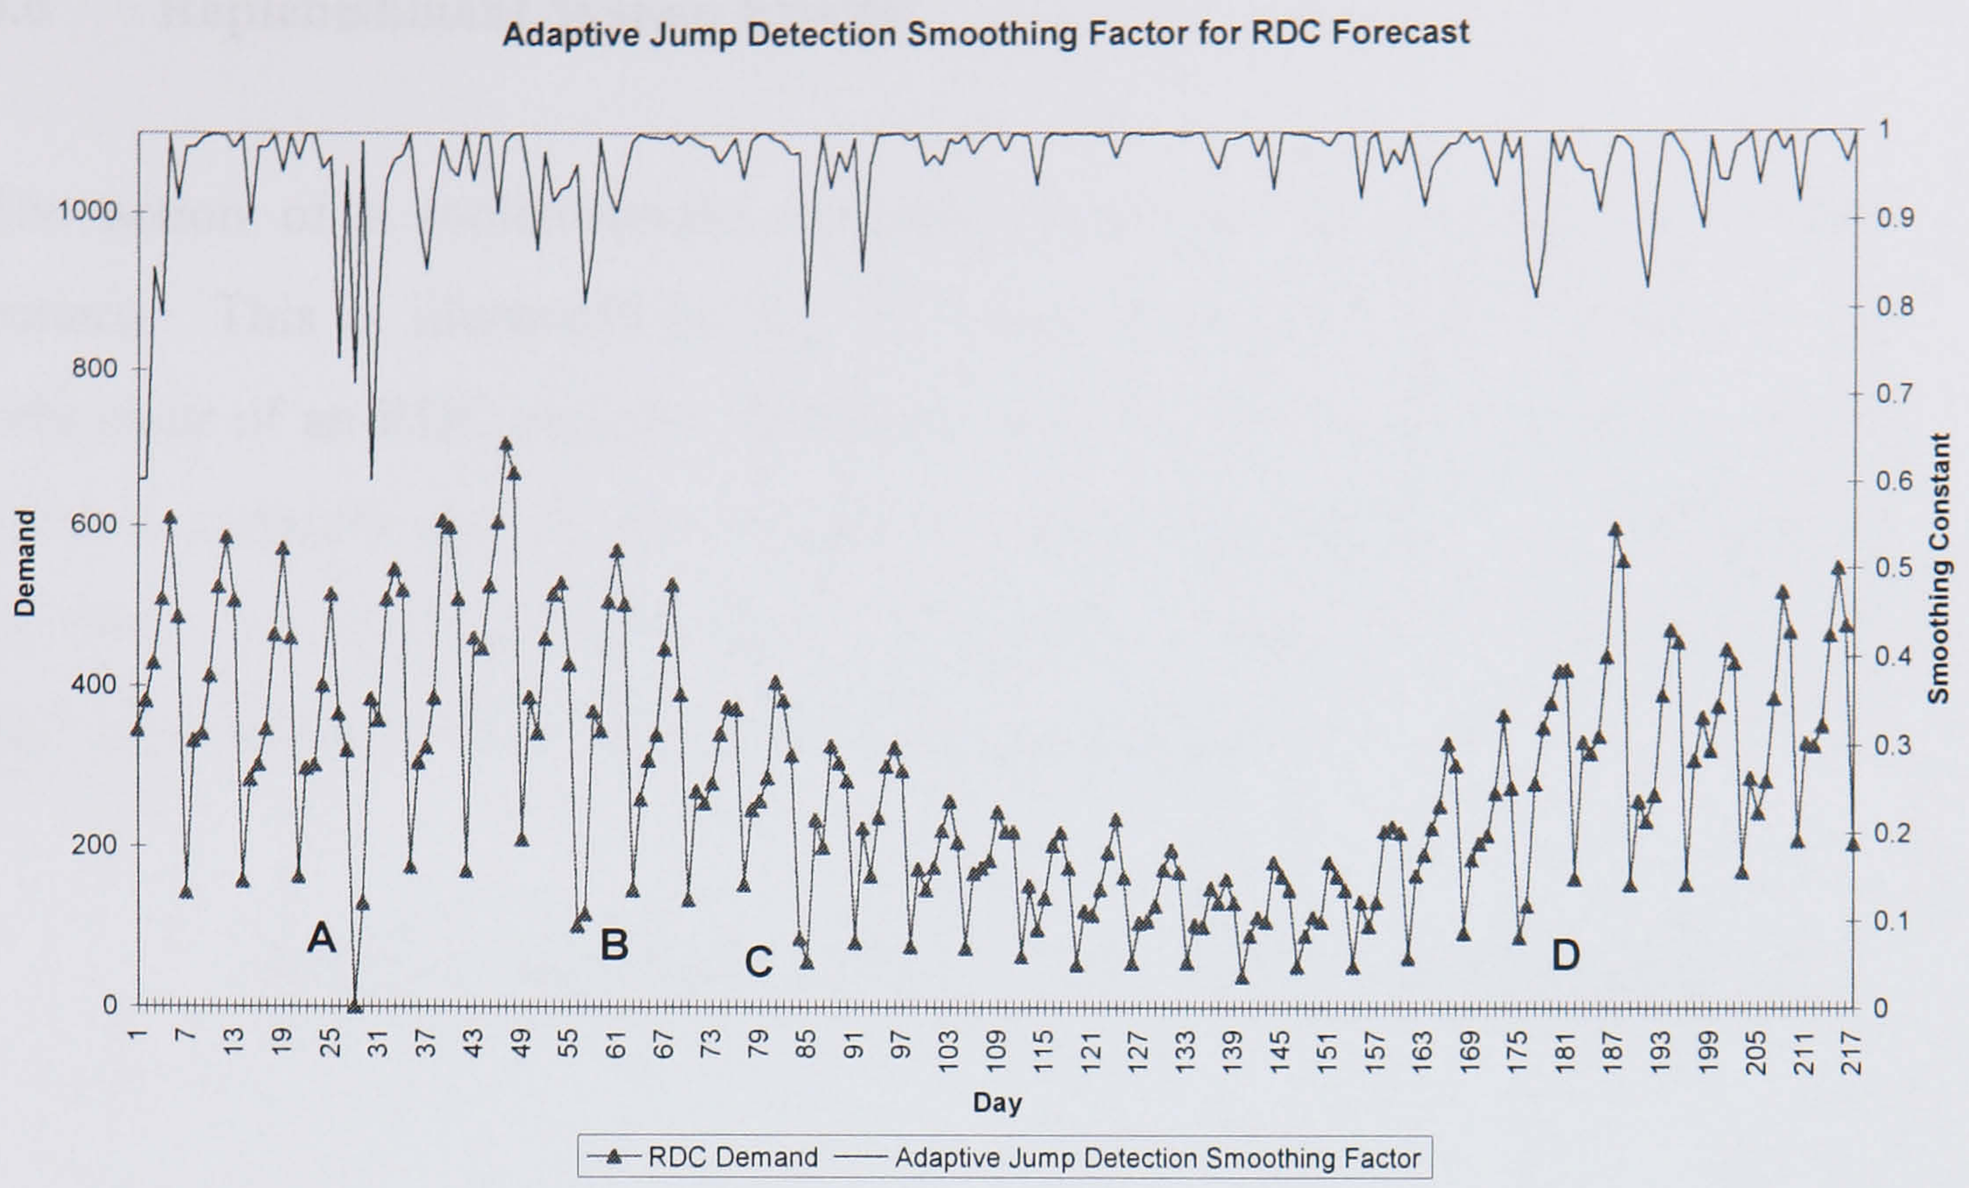


Fig. 14

In this example $N = 10$ and $\theta^2 = 5.2$. Although this product has no promotional activity logged during the period shown, changes in mean level are present. In such situations, it can be seen that the smoothing constant α is reduced in order to give priority to current information so that the system adapts quickly. Closer examination of the demand data at points A, B, C and D in Fig. 14, where α reacts sharply, revealed unusual discrepancies in the data. At point A the demand fell to zero, the probability of this occurring is very small and investigation revealed that the zero demand occurred on Easter Sunday and hence all stores were closed. At point B the demand fell by 70% compared with the demand on the same day for the previous week and similarly, at point C, a 77% drop was experienced. At point D, the daily demand pattern altered for one week and then returned to normal. These unusual daily demand levels and patterns are not easily seen by eye. The adaptive jump detection can be used to rapidly adjust to unexpected changes in the mean level of demand and highlight demand data which is unusual. The latter function would be useful when exception reporting, as it would highlight any behaviour which may warrant further investigation (e.g. manual intervention in data collection, stores failing to place products on promotion.).

3.6 Replenishment System Effects

The action of a replenishment systems can completely change the demand pattern. This is illustrated in Fig. 15 which shows the actual deliveries and behaviour of an RDC during a Christmas period. The system is close to failure (almost certainly due to the effects of manual intervention and operational problems which often accompany the Christmas demand) even though all deliveries made by the supplier were on time and in full.

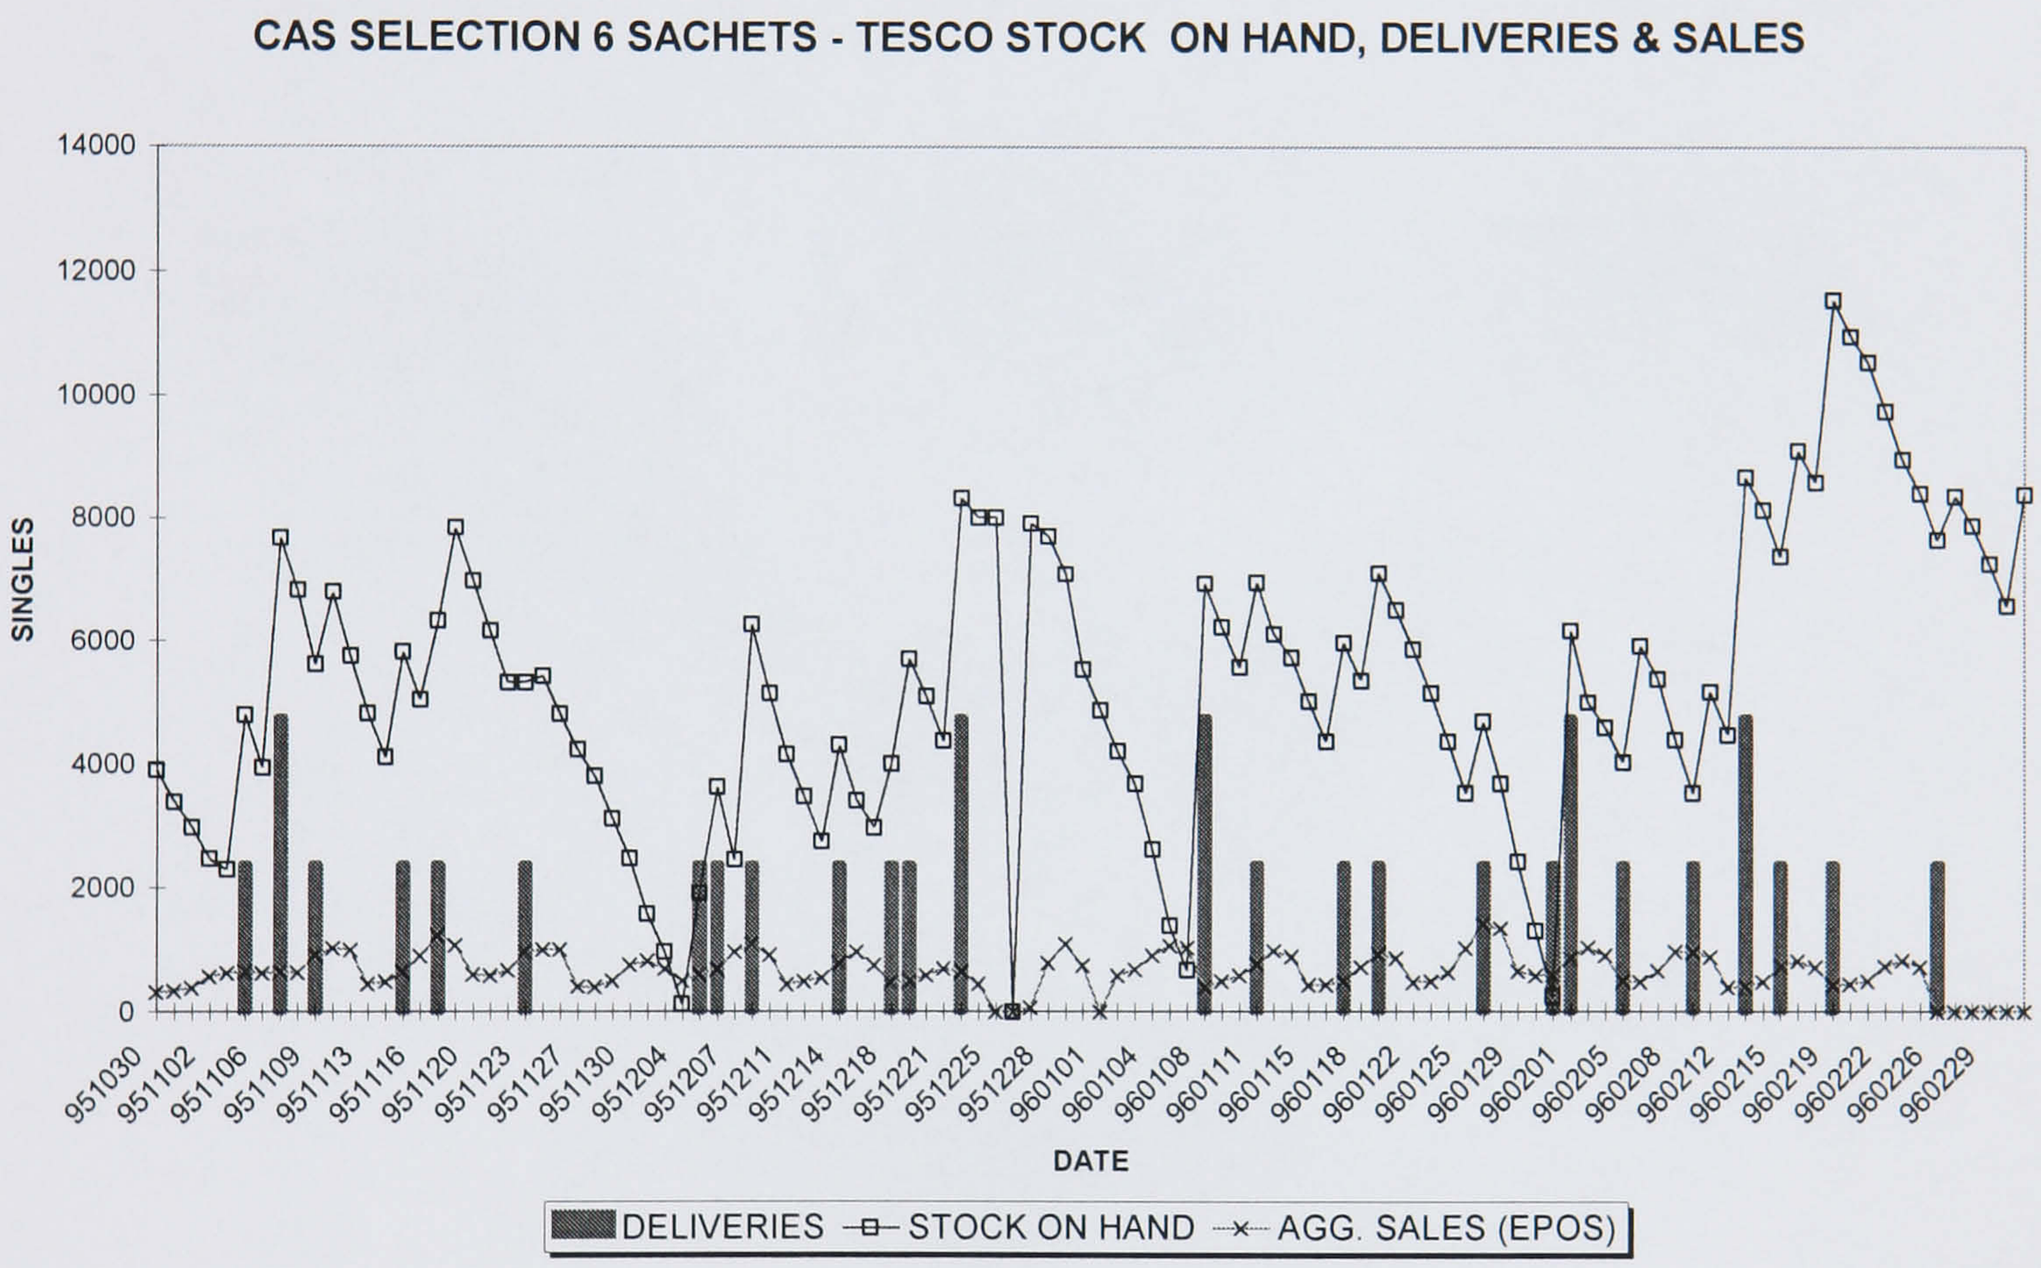


Fig. 15

Notice that the RDC replenishment orders are pulse like, whereas the demand signal on the RDC is smooth in character. In the operation of Vendor Managed Inventory systems and information sharing systems such as TIE and SID, the information of interest to the supplier is the pulse like replenishment orders, and not the smooth demand on the RDC. In order to model the replenishment orders it is necessary to feed the warehouse demand through a model of the replenishment system.

Similar considerations apply to the store demand on an RDC which is also pulse like. It is therefore noteworthy that the pulse like orders, which include round up effects from 200 or so store replenishment systems, when aggregated together, effectively average out much of the noise generated by the individual store orders (Fig. 16). Tests showed that about twenty or more store's demands were averaged, then a close approximation to the demand pattern of the RDC was obtained.

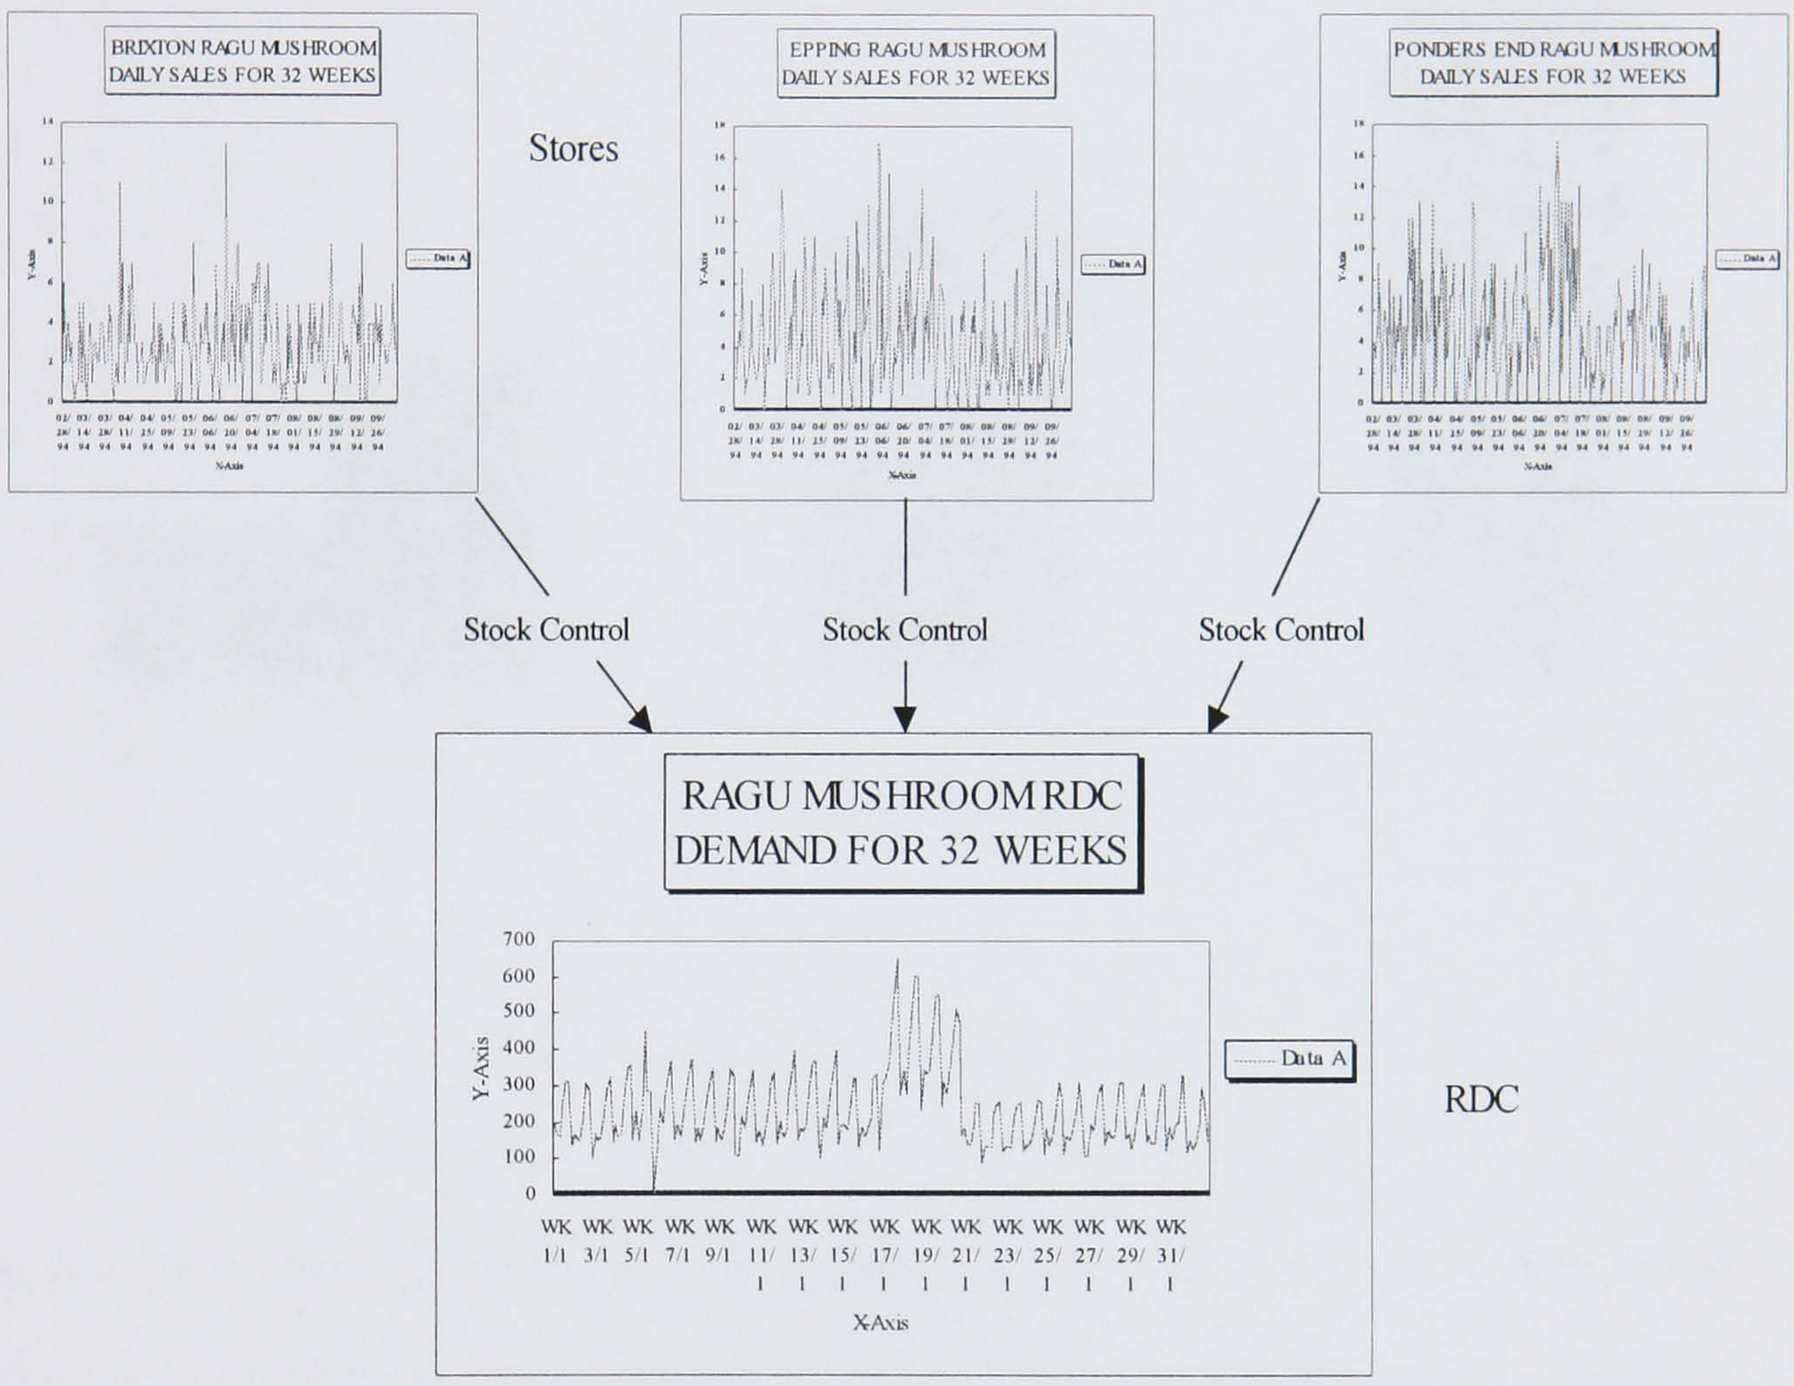


Fig. 16

3.7 Off-Line Modelling of Weather Mode Effects

The discrete model of a weather mode proposed in Chapter 2, is illustrated in Fig. 17. The mode value is w_k if the explanatory variable in day i of week j has the value T_{ij} and $T_{k-1} < T_{ij} \leq T_k$ for $(1 \leq k \leq N)$.

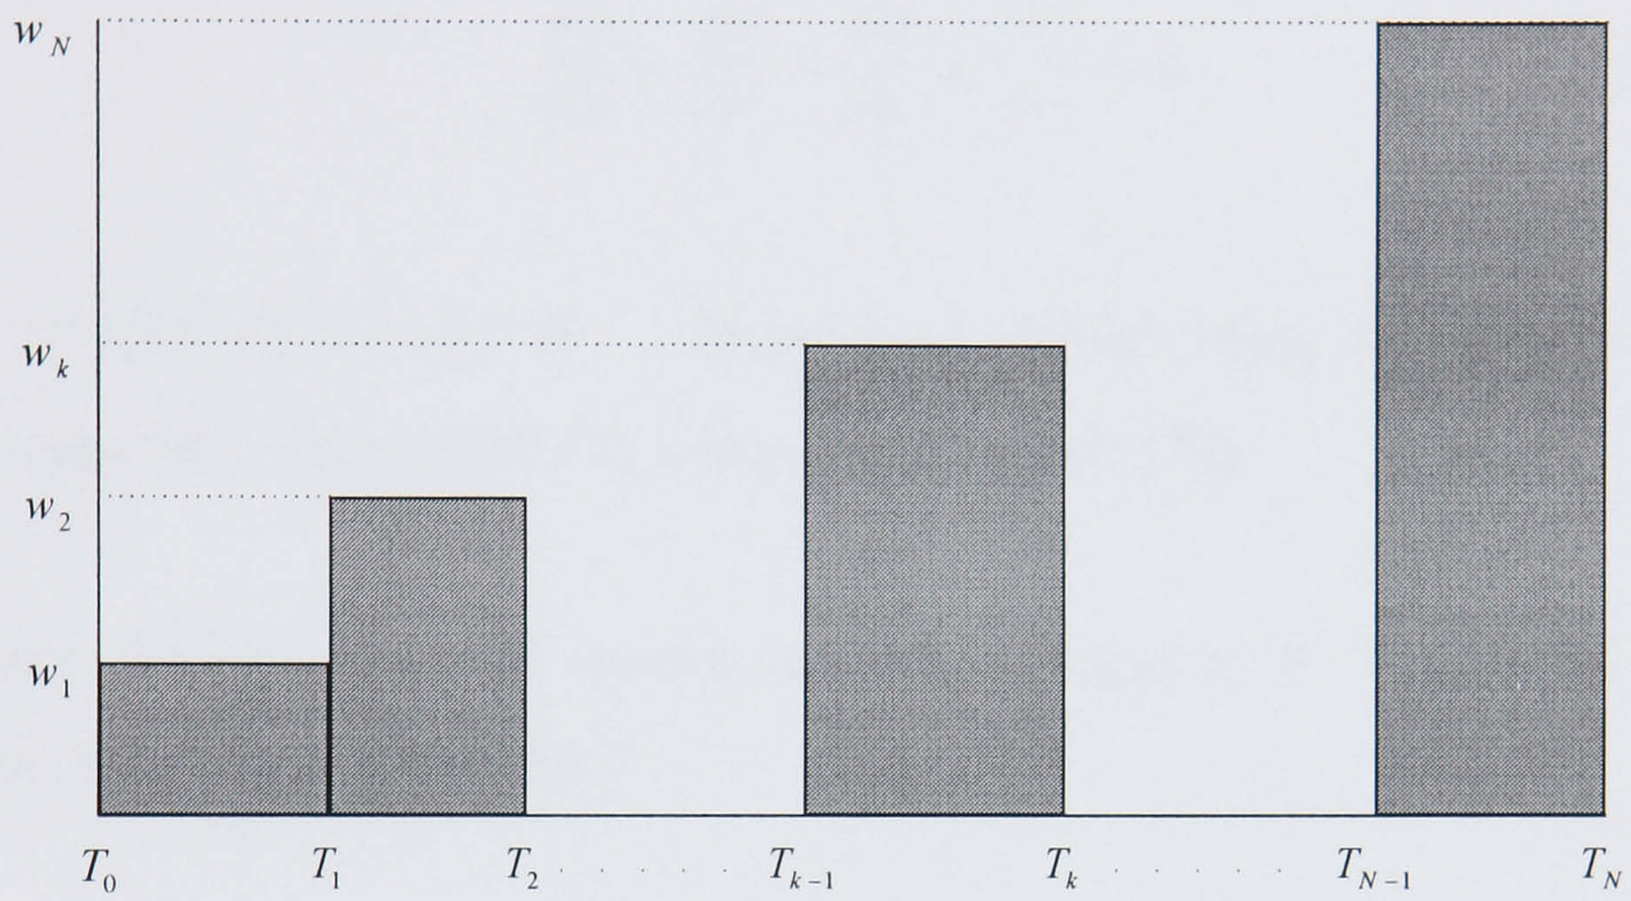


Fig. X Discrete Weather Mode Form

Fig. 17

The discrete intervals T_k need not be equally spaced (in the results discussed later, they were chosen so that roughly equal numbers of explanatory variables occurred in each interval). The parameter w_k $(1 \leq k \leq N)$ has to be determined by model fitting.

The expected demand on day i of week j, proposed in Chapter 2 was of the form

$$\mu_{ij} = \lambda d_i s_j w_k \qquad 19.0$$

Substituting this expression into the likelihood function of equation 3.0 for the $I \times J$ array Y, we obtain

$$l(\lambda, d_i, s_j, w_k; Y) = \sum_{i=1}^{I}\sum_{j=1}^{J}\left\{y_{ij}\log\lambda d_i s_j w_k - \lambda d_i s_j w_k - \log y_{ij}!\right\} \qquad 20.0$$

The maximum likelihood values of the model parameters satisfy the $(I + J + N + 1)$ equation

$$\frac{\partial l}{\partial \lambda} = \frac{\partial l}{\partial d_i} = \frac{\partial l}{\partial s_j} = \frac{\partial l}{\partial w_k} = 0 \qquad 21.0$$

for $1 \le i \le I, 1 \le j \le J, 1 \le k \le N$. These were solved using the generalised Newton approach, Greenspan [74], Ortega and Reinbolt [72].

Expression for the first and second derivatives required to implement this procedure are given in Appendix 3.

In order to fit the model parameters off line, explanatory variables were collected over a period of 12 months, together with RDC demand data. The RDC at Strood was chosen. Daily demand for all dry soups and several varieties of ice cream were collected. The Meteorological Office, Bracknell provided data from East Malling weather station, which gave a sensible coverage for the region supplied by the Strood RDC. The potential explanatory daily variables collected were:

- Maximum Temperature (°C)
- Mean Temperature (°C)
- Rain Fall (mm)
- Sunshine (hrs)

Results from fitting these variables indicated that the weather effects were not significant. Sample results are discussed below. A possible explanation for this negative result is that weather effects only become significant when prolonged extreme conditions develop, and that these extremes did not develop during the test period (only four days in the year fell below +5 °C and the maximum temperature reached was +24.5 °C).

3.7.1 Weather Mode Results

A graph of daily hours of sunshine, rainfall and mean temperature in the region of East Malling, obtained from the Met. Office, is shown in Fig. 18, and the demand for the soup and ice cream to

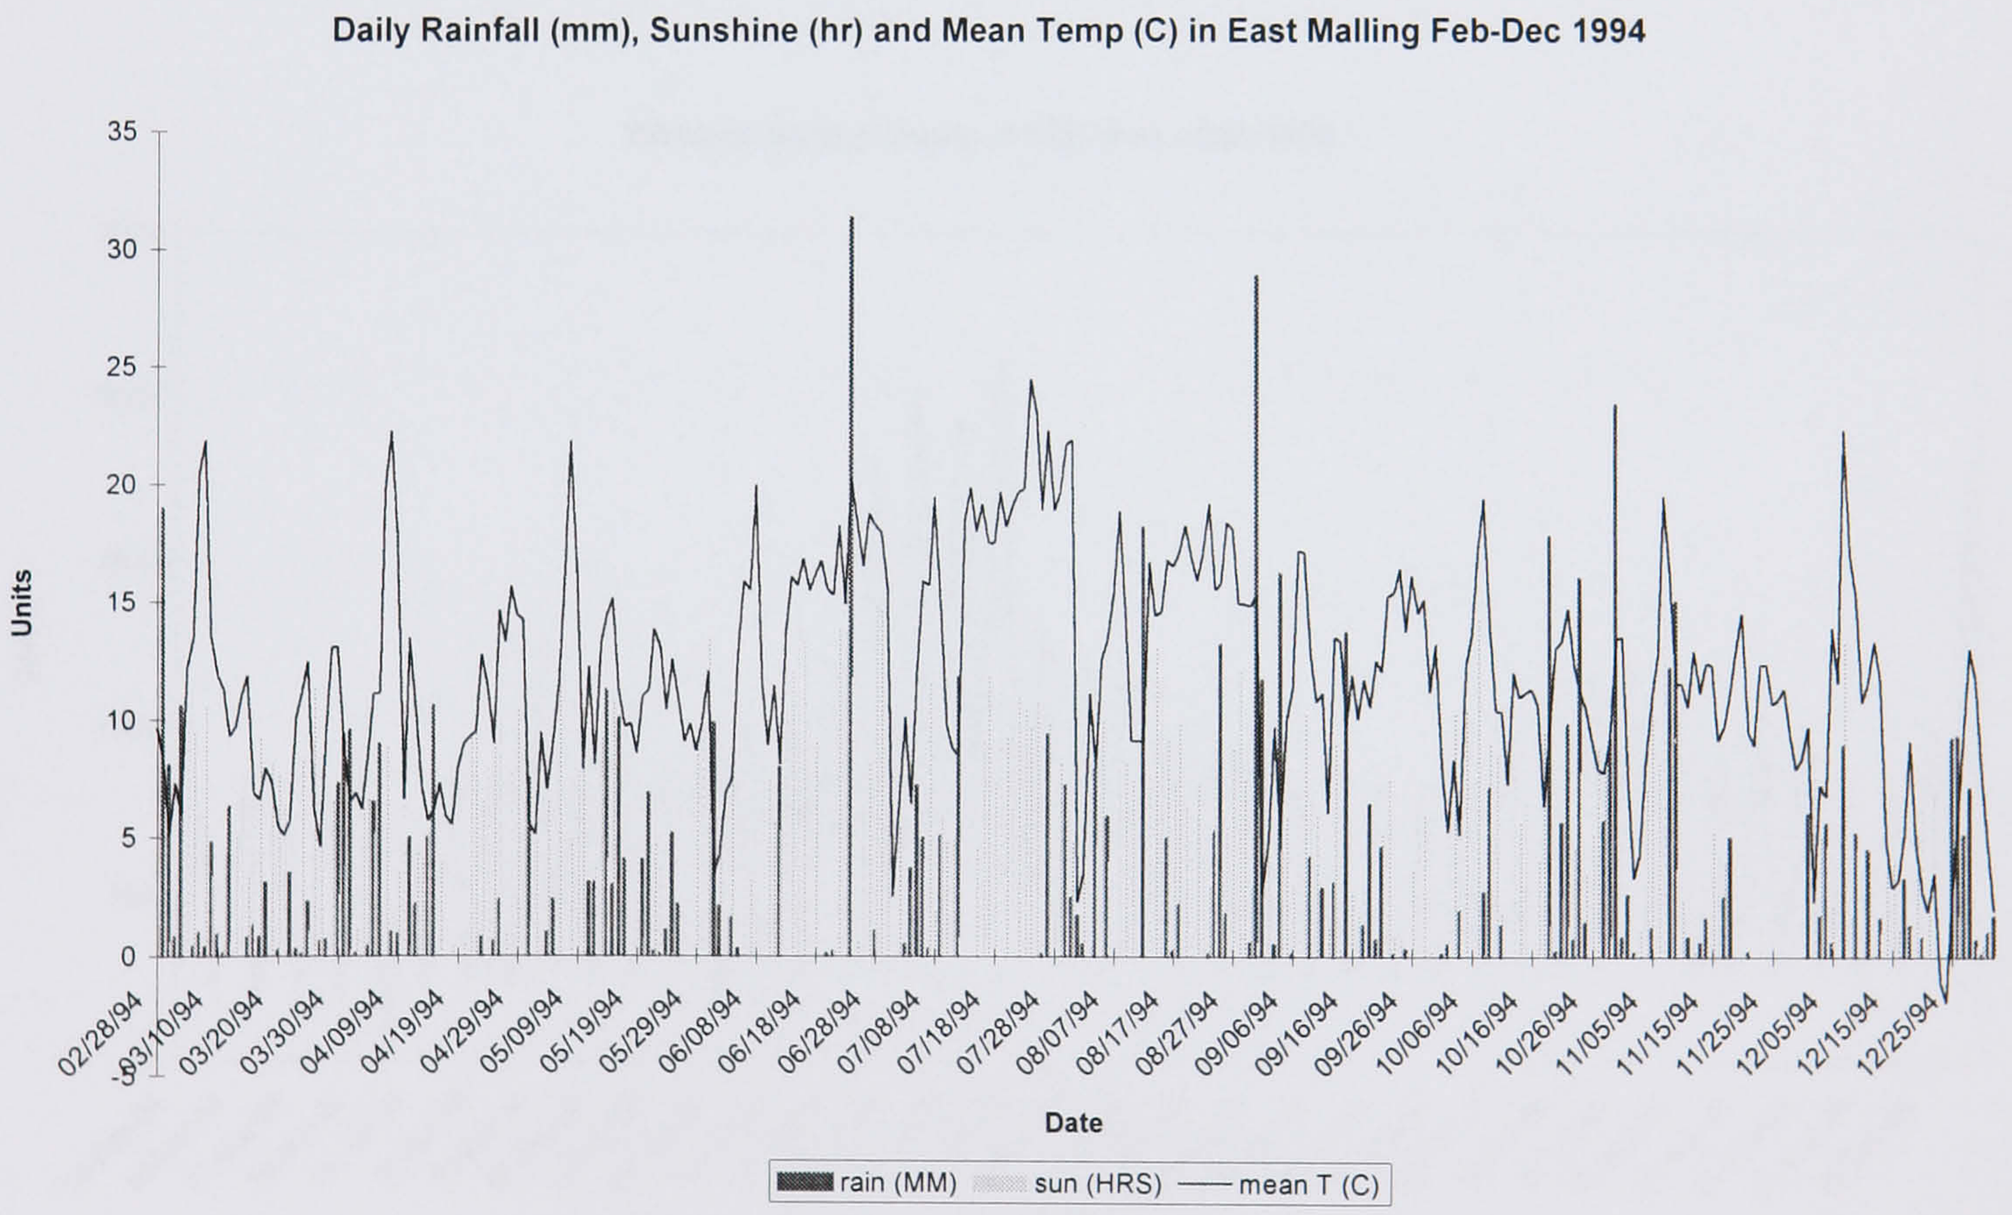


Fig. 18

be discussed are shown in Fig. 19 and Fig. 20.

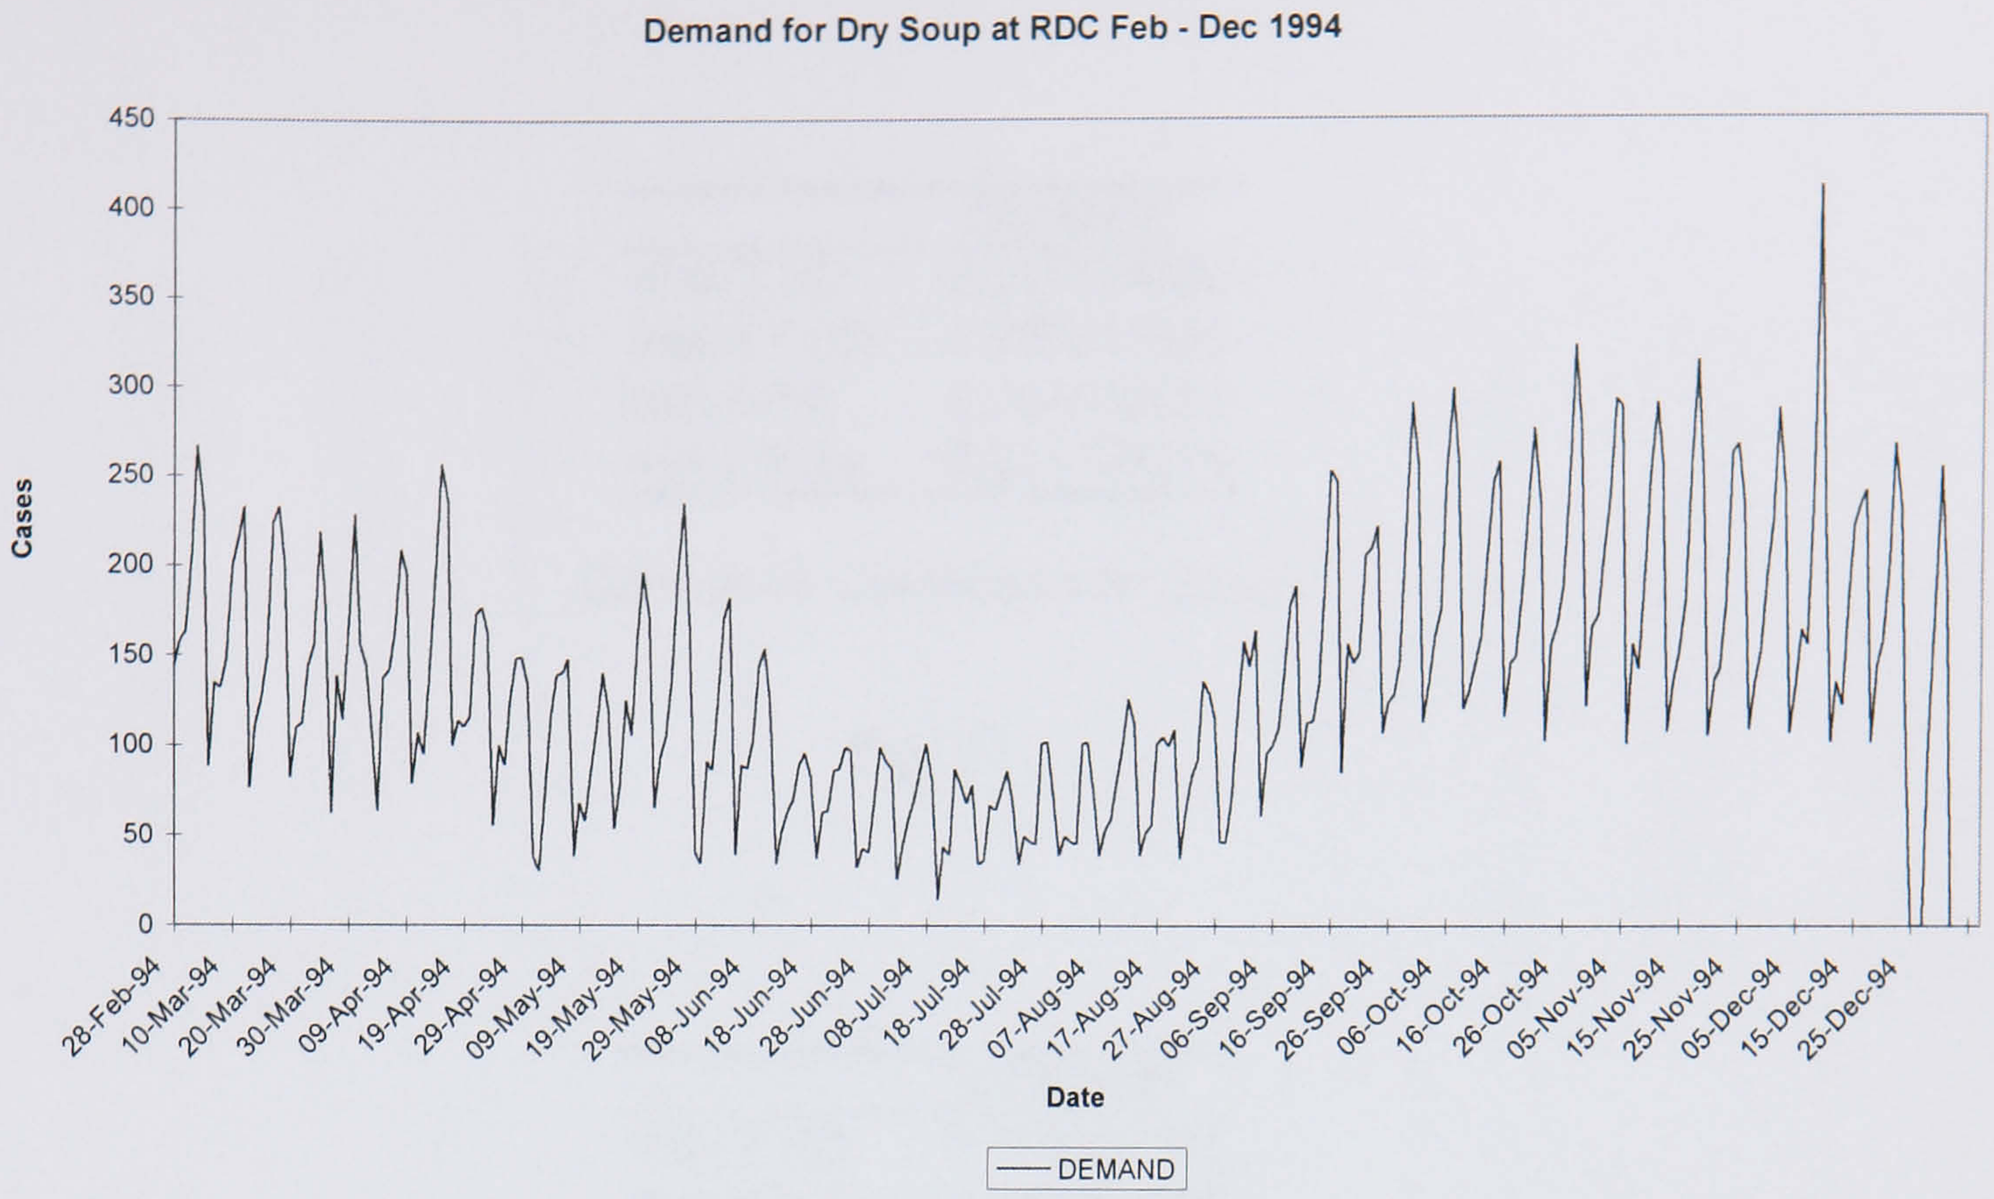


Fig. 19

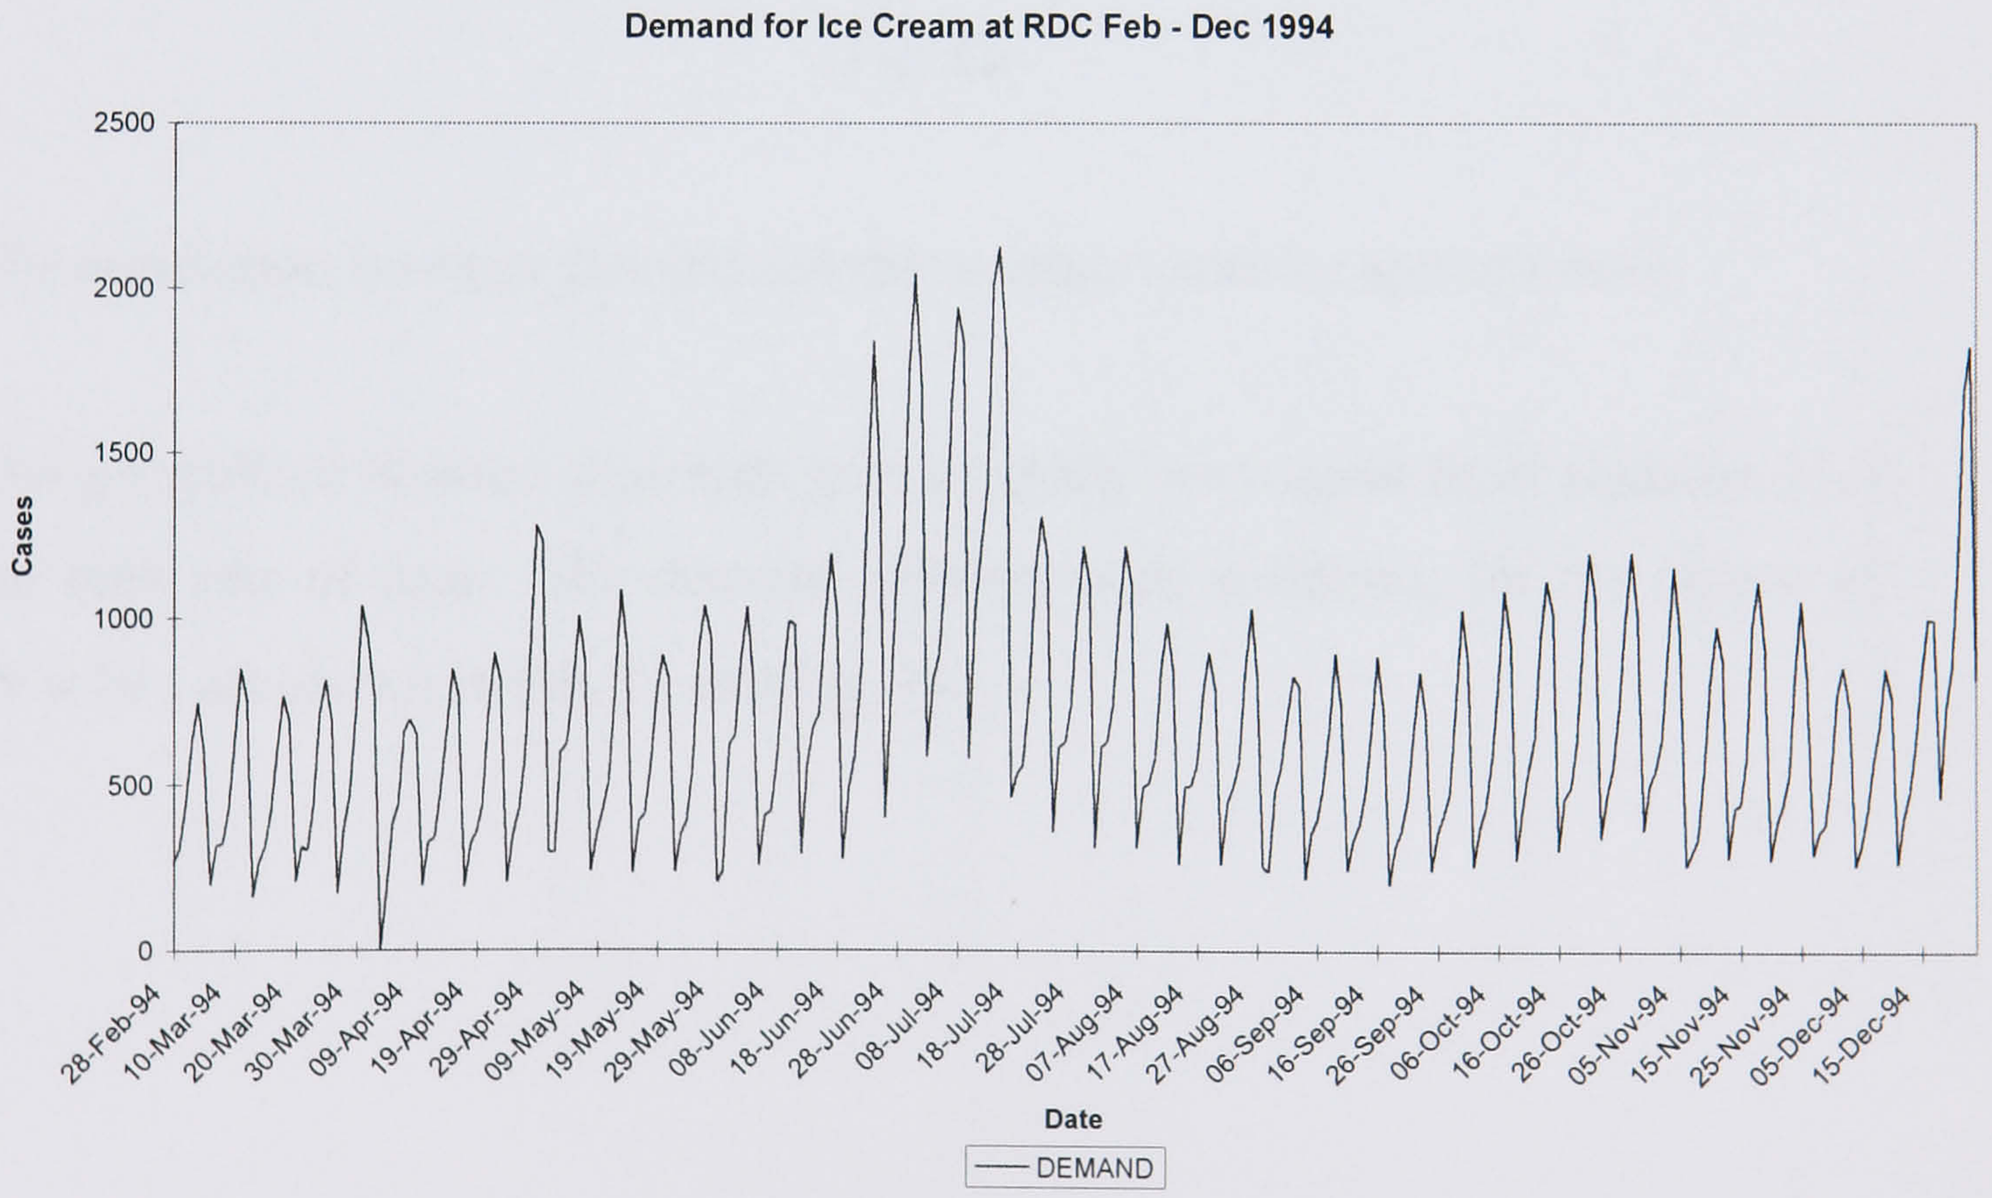


Fig. 20

The correlation coefficients were taken to give an initial indication of the effects of the weather variables on demand for the two products. These are shown in Fig. 21 and Fig. 22.

	DEMAND
max T (C)	-0.311394605
mean T (C)	-0.280295995
rain (MM)	0.054402096
sun (HRS)	-0.254195319

Correlation Coefficients for Soup

Fig. 21

	DEMAND
max T (C)	0.119691405
mean T (C)	0.097294897
rain (MM)	-0.071511164
sun (HRS)	0.147050501

Correlation Coefficients for Ice Cream

Fig. 22

The correlation between demand and the weather variables appears weak.

The generalised Newton algorithm gave a rapidly convergent fit of equation 21.0 for both sets of data. The discrete weather mode solutions, for the choice of $N = 10$, are shown in Fig. 23 and Fig. 24.

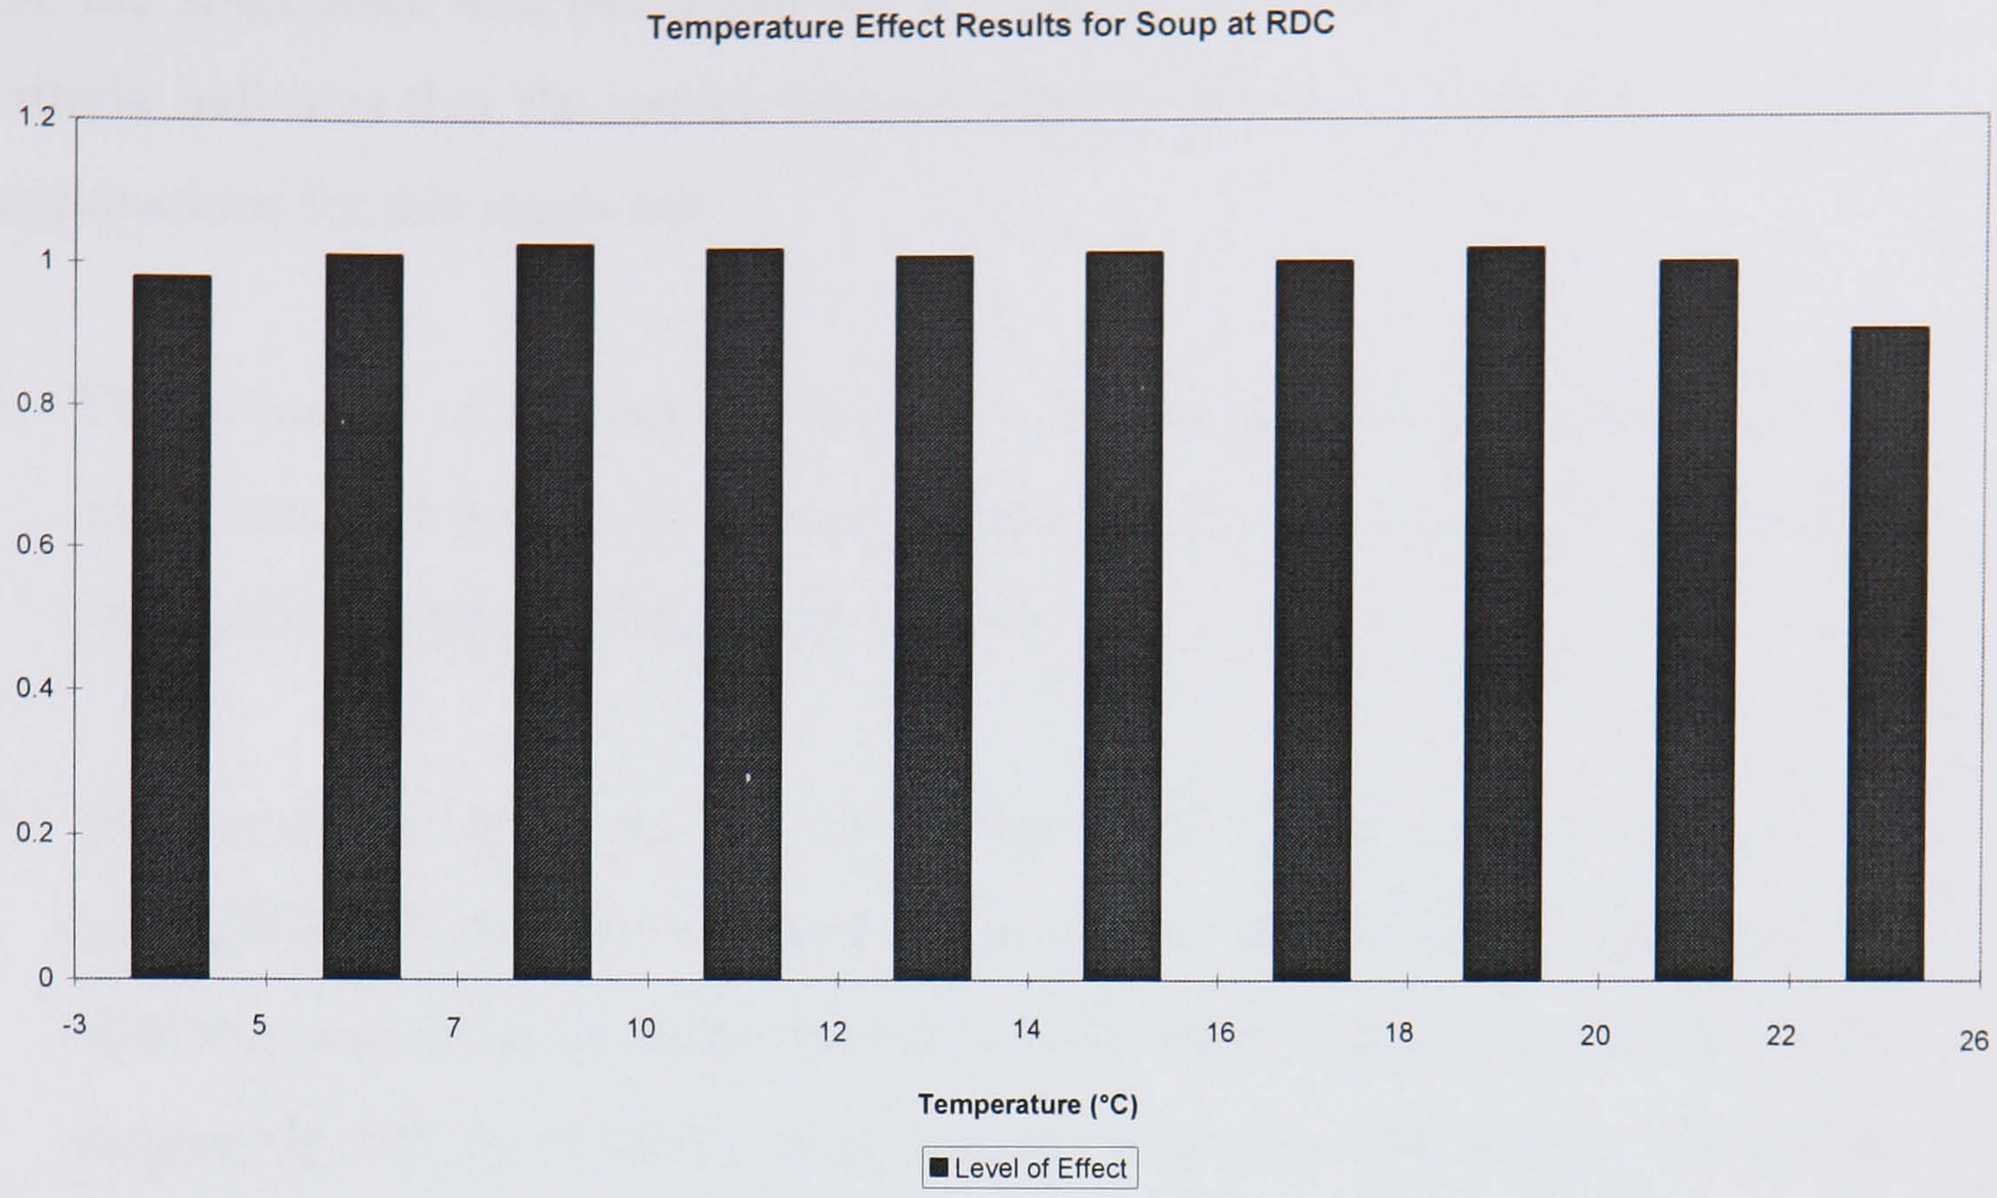


Fig. 23

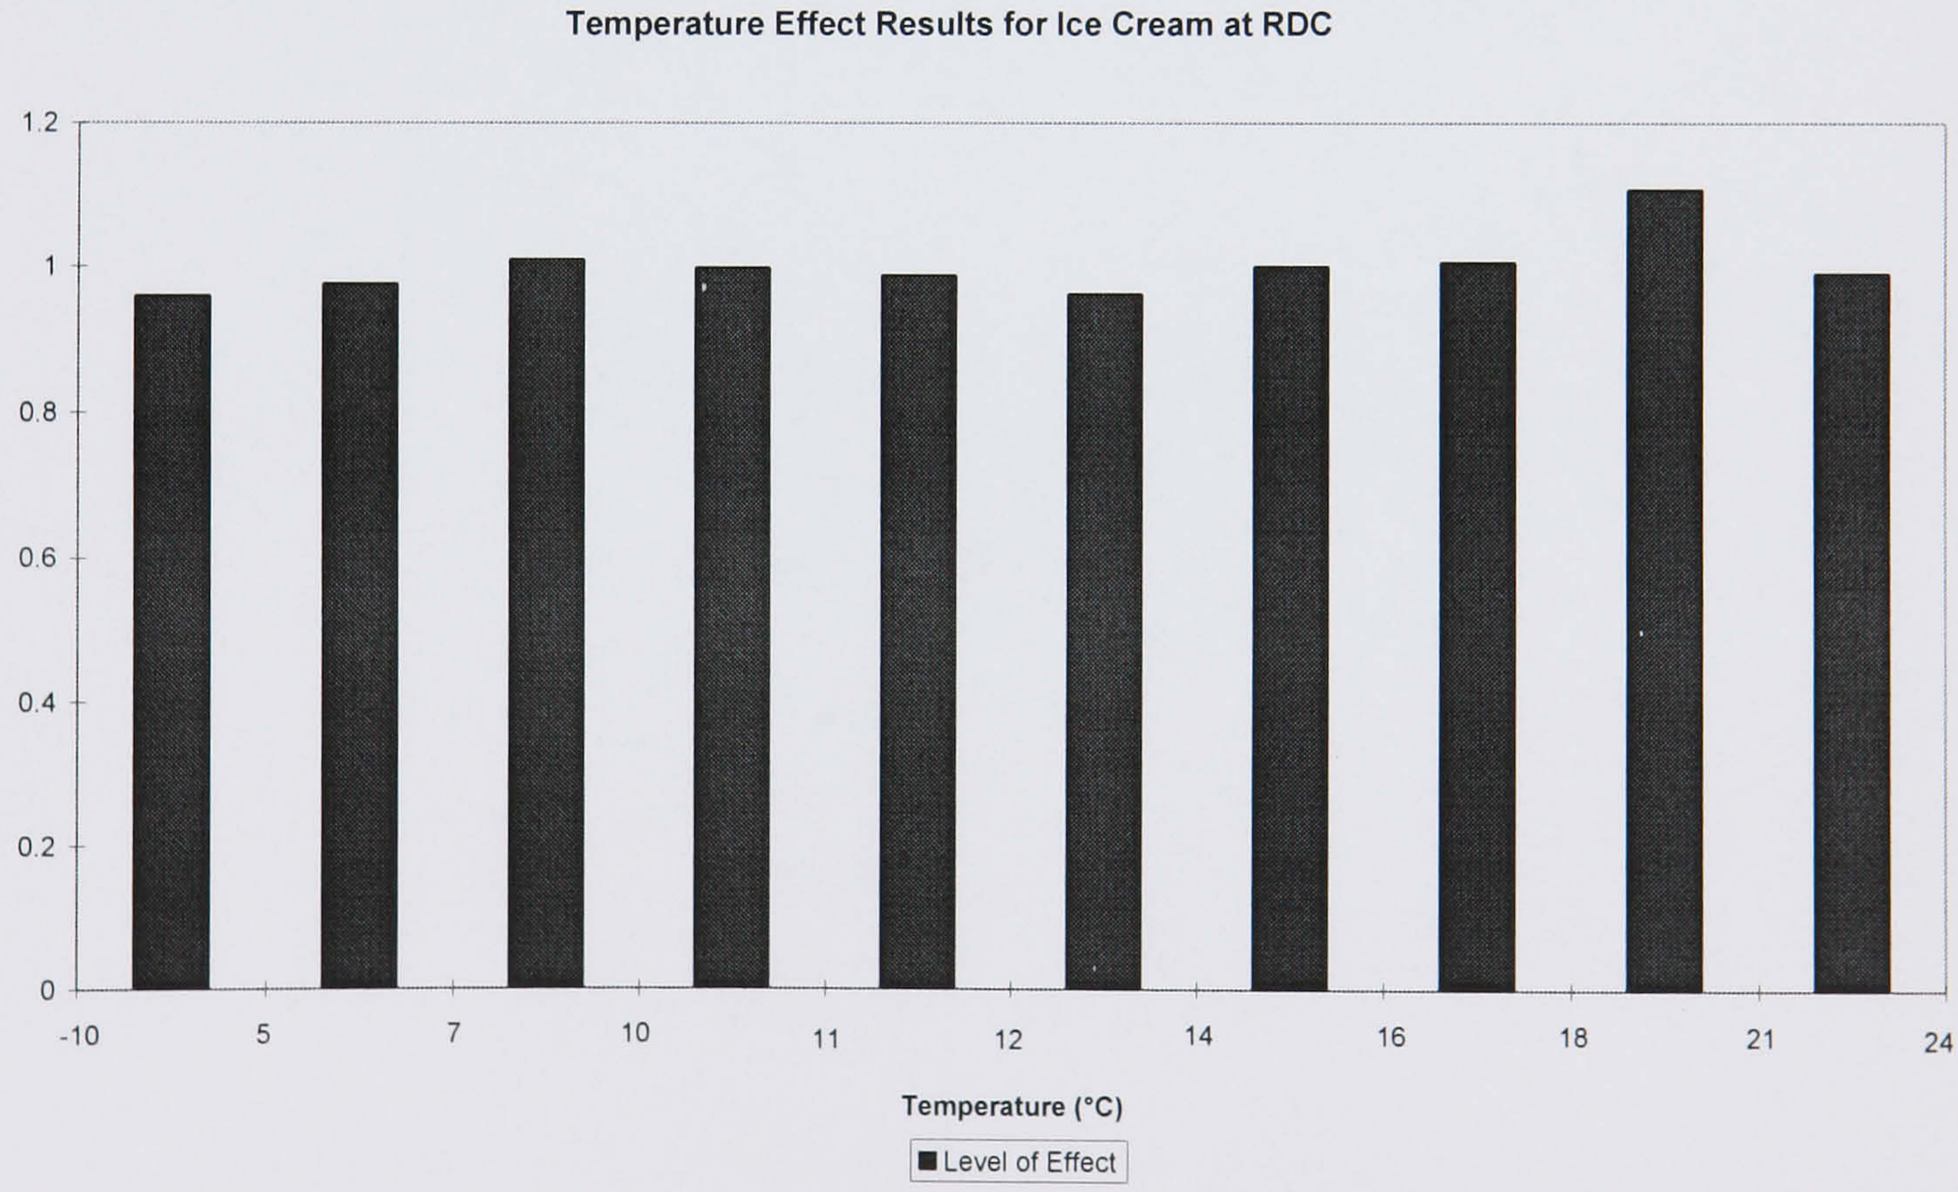


Fig. 24

It would seem that none of the temperatures bands has any significant effect on the level of demand on soup or ice cream. The Akaike Information Criterion (AIC), Lindsey [95] was evaluated with (without) the weather mode. The AICs

for the soup were 432 (413) and for ice cream 524 (518) In both cases the criteria indicates that the model without weather effects is preferred. Possible explanations for this result are:

1. The influence of weather on demand is negligible until prolonged extreme conditions, not present in the data, develop. This appears to be a plausible explanation, from examination of the data.

2. Weather mode effects are already accounted for by the seasonal mode. The correlation between the seasonal mode effects for dry soup and hours of sunshine was however only -0.15 and with mean daily temperature, -0.29, suggesting that the seasonal variations are determined by factors other than weather.

3.8 APPENDIX 1

Maximum Likelihood Derivation

The likelihood function ℓ for the rank 1 model is given by equation 6 of chapter 3. The first order conditions for a maximum of ℓ are

$$\frac{\partial \ell}{\partial d_i} = \sum_{j=1}^{J} y_{ij} / d_i - \lambda = 0$$

$$\frac{\partial \ell}{\partial s_j} = \sum_{j=1}^{J} y_{ij} / s_j - \lambda = 0$$

$$\frac{\partial \ell}{\partial \lambda} = \sum_{i=1}^{I} \sum_{j=1}^{J} y_{ij} - \lambda = 0$$

Defining vector S^r and S^c by $S_i^r = \sum_{j=1}^{J} y_{ij}$, $S_j^c = \sum_{i=1}^{I} y_{ij}$ and $S_j = s_j$ then the above equations reduce to

$$S_i^r = \lambda d_i \; i = 1 \ldots I$$

$$S_j^c = \lambda S_j \; j = 1 \ldots J$$

$$\text{and } \|S^r\| = \|S^c\| = \lambda$$

where $\|\cdot\|$ denotes the 1 norm Barnett [90]. Provided at least one element of Y is strictly positive then $\lambda = \sum_{i=1}^{I} \sum_{j=1}^{J} y_{ij} > 0$. Therefore setting

$$D = S^r / \|S^r\|$$

$$S = S^c / \|S^c\| \qquad \text{A.1.2}$$

$$\lambda = \|S^r\|$$

we obtain a solution which satisfies the constraints of equation 5.0. With increasing J, these solutions become ill conditioned since $\|S^r\|$, $\|S^c\|$ and λ all increase without bound.

It is convenient to re-express the solution in the alternative form

$$D = \bar{S}^r / \|\bar{S}^r\|$$

$$\lambda S = S^c$$

$$\text{where } \bar{S}_i^r = \frac{1}{J}\sum_{j=1}^{J} y_{ij} \qquad i = 1 \ldots I$$

is the vector of row averages for array Y. Since $\bar{S}^r$ is a scaled version of S^r, the solution of D is unaffected. Although $\|S^c\|$ becomes unbounded, this is of no practical concern for on-line control since only the component S_j^c is of interest in real time application. In stating Result 2 of chapter 3 the bar notation for S^r has been dropped.

3.9 APPENDIX 2

Adaptive Exponential Smoothing

Consider the constant mean model,

$$y_k = \mu + \varepsilon_k \qquad \text{A.2.1}$$

where $\varepsilon_k \in N(0, \mu)$. The minimum weighted least squares estimate $\hat{\mu}_k$ which minimises

$$\Sigma_k = \sum_{i=1}^{k} \lambda^{k-i} e_i \qquad \text{A.2.2}$$

where $e_i = (y_i - \hat{\mu}_k)$ can be expressed in the recursive form

$$\hat{\mu}_k = \hat{\mu}_{k-1} + K_k e_k \qquad \text{A.2.3}$$

where

$$e'_k = (y_k - \mu_{k-1})$$

$$K_k = \left(A_k \Lambda_k A_k^T\right)^{-1}$$

A_k is a k dimensional vector of ones, and $\Lambda_k = diag[\lambda^{k-1}, \lambda^{k-1}, \ldots 1]$ Aström and Wittenmark [89]. It is easily shown that

$$K_k = \frac{1-\lambda}{1-\lambda^k} \qquad \text{A.2.4}$$

and therefore K_k converges in the limit to

$$K_k = 1 - \lambda \qquad \text{A.2.5}$$

provided $1 < \lambda < 0$. Therefore, from equation A.2.3 the optimum estimator converges to the form of an exponential smoother. If the process parameters change then it is desirable that the forgetting factor λ is reduced. Various procedures have been developed Kershenbaum[88] and we here consider the approach of Fortesque et al [81] who defines an information measure Σ in terms of the weighted sum of square of the a posteriori errors and suggested that λ be chosen to maintain this constant at Σ_0 say. They show that, for this problem, Fortesque [77]

$$\Sigma_k = \lambda\Sigma_{k-1} + (1 - K_k)e_k^2$$

from which it follows that λ is implicitly defined by

$$1 - \lambda_k = (1 - K_k)e_k^2 / \Sigma_0 \qquad \text{A.2.6}$$

From equations A.2.5. and A.2.3. we update the parameter estimate according to

$$\hat{\mu}_k = \hat{\mu}_{k-1} + (1 - \lambda_{k-1})e_k$$

and then the smoothing factor from A.2.6, so that

$$\lambda_k = \frac{\Sigma_0}{\Sigma_0 + e_k^2}$$

where, following Fortesque et al [81] we have avoided an implicit calculation by updating λ_k after parameter update and have constrained $\lambda > \lambda_{\min} > 0$ to avoid small values if very large errors develop.

3.10 APPENDIX 3

Weather Mode Derivations

By straight forward differentiation, equations 21.0 of the text may be expressed:

$$\frac{\partial l}{\partial d_i} = \sum_{j=1}^{J} y_{ij} / d_i - \lambda \sum_{j=1}^{J} s_j w_k = 0$$

$$for \quad 1 \le i \le I,$$

$$\frac{\partial l}{\partial s_j} = \sum_{i=1}^{I} y_{ij} / s_j - \lambda \sum_{i=1}^{I} d_i w_k = 0$$

$$for \quad 1 \le j \le J,$$

$$\frac{\partial l}{\partial w_k} = \sum_{M_k} y_{ij} / w_k - \lambda \sum_{M_k} d_i s_j = 0$$

$$for \quad 1 \le k \le N, \text{ and}$$

where M_k is defined as

$$M_k = \left[ij;\, T_{k-1} < T_{ij} \le T_k \right]$$

In order to implement the generalised Newton method we require

$$\frac{\partial^2 l}{\partial d_i^2} = -\sum_{j=1}^{J} y_{ij} / d_i^2 \qquad 1 \le i \le I$$

$$\frac{\partial^2 l}{\partial s_j^2} = \sum_{i=1}^{I} y_{ij} / s_j^2 \qquad 1 \le j \le J$$

$$\frac{\partial^2 l}{\partial w_k^2} = \sum_{M_k} y_{ij} / w_k^2 \qquad 1 \le k \le N$$

$$\frac{\partial^2 l}{\partial \lambda^2} = \sum_{i=1}^{I}\sum_{j=1}^{J} y_{ij} / \lambda^2$$

If x denotes the M dimension vector of model parameters

$$x = \left(\lambda, d_1 \ldots d_I, s_1 \ldots s_J, w_1 \ldots w_N\right)$$

then the generalised Newton method employs the iteration

$$x_i^{n+1} = x_i^n - \left[f_i\left(x^{ni}\right) \Big/ \frac{\partial f_i}{\partial x_i}\left(x^{ni}\right) \right]$$

where $x^{ni} = \left(x_1^{n+1}, \ldots x_{i-1}^{n+1}, x_i^n, \ldots x_M^n\right)$ and $f_i = \dfrac{\partial l}{\partial x_i}, \quad i = 1 \ldots M$, Greenspan[74]

References

Aström, K.J., Wittenmark, B. 1989. *Adaptive Control.* Addison Wesley

Barnett, S. 1990. *Matrices, Methods and Applications.* Oxford University Press

Carroll, R.J. and Ruppert, D. 1988. *Transformation and Weighted Regression.* Chapman Hall.

Clarke, D. and Gawthrop, P.J. 1975. *Self Tuning Controller.* Proc. Inst. Elect. Eng.,122.

Fortesque, T.R. 1977. *Work on Aström's self tuning regulator.* Department of Chemical Engineering Report, Imperial College, p20-28.

Fortesque, T.R., Kershenbaum L.S. and Ydstie B.E. 1981. *Implementation of Self Tuning Regulators with Variable Forgetting Factors.* Automation Vol. 17, p831-835.

Gelb, A. 1974. *Applied Optimal Estimation.* M.I.T. Press

Greenspan, D. 1974. *Discrete Numerical Methods in Physics and Engineering.* Academic Press

Hastings, N.A.J. and Peacock, J.B. 1975. *Statistical Distributions.* Butterworths

Kalbfleisch, J.G. 1985. *Probability and Statistical Inference.* Vol 2. Springer Verlag.

Kershenbaum, L.S. 1988. *Recursive Algorithms, Variable Forgetting Factors.* Systems and Control Encyclopedia Pergamon.

Lindsey, J.K. 1995. *Introductory Statistics.* Oxford University Press.

Ortega and Reinbolt. 1972. *Iterative Solution of Non-Linear Equations in Several Varaiables.* Academic Press.

4.0 INVENTORY MANAGEMENT CONTROL

4.1 Introduction.

The classical theories of inventory control were developed in the sixties and seventies and are described for instance in Maggee [58], Lewis [70], Battersby [70] and Brown [80]. These theories still represent the science based art employed in the design and operation of most supply chain systems in the retail sector. The replenishment procedures of classical theory fall into two main classes. Those in which the procedure is based solely on inventory level [S-s system] and those in which it is based on inventory level at fixed cycle time. Lewis [70]. Only fixed cycle systems are considered here, since modern computer based systems in the retail sector mainly operate on the fixed cycle principle. We develop some basic results for these systems and for use later.

4.2 Fixed Cycle Systems

Fixed cycle systems operate on the simple principle of balancing forecast demand against book stock and stock in transit, over a time window (T) usually set equal to the review period plus lead time. We will assume, as is usually the case, that the lead time (L) is a multiple of the review period and for convenience, set the review period equal to unity. Therefore the time window is usually $T = (L+1)$.

If at review point n, the forecast demand over window T exceeds the available inventory, then the deficit is made up by the current replenishment so that

$$O_n = F_n - \left(B_n + L_n\right) \qquad 1.0$$

where we define the variables

O_n Current replenishment

F_n Forecast demand over window T

B_n Current bookstock

L_n Current inventory in transit

and we note that since L_n represents the inventory arriving at review point n, plus that arriving in the next lead time period, $\left(B_n + L_n\right)$ represents the total inventory in the system.

The forecast demand is usually estimated in the form

$$F_n = W_n + S_n \qquad 2.0$$

where W_n is the expected demand over window T (the working stock) and S_n (the safety stock) is chosen to avoid unacceptable levels of stock out, during periods of higher than expected demand. The specific form of F_n depends on the nature of the inventory demand process. We consider a few cases of interest.

4.2.1 Deterministic Demand.

If the demand is constant at level D and deterministic $\left(S_n = 0\right)$ then equation 2.0 can be expressed

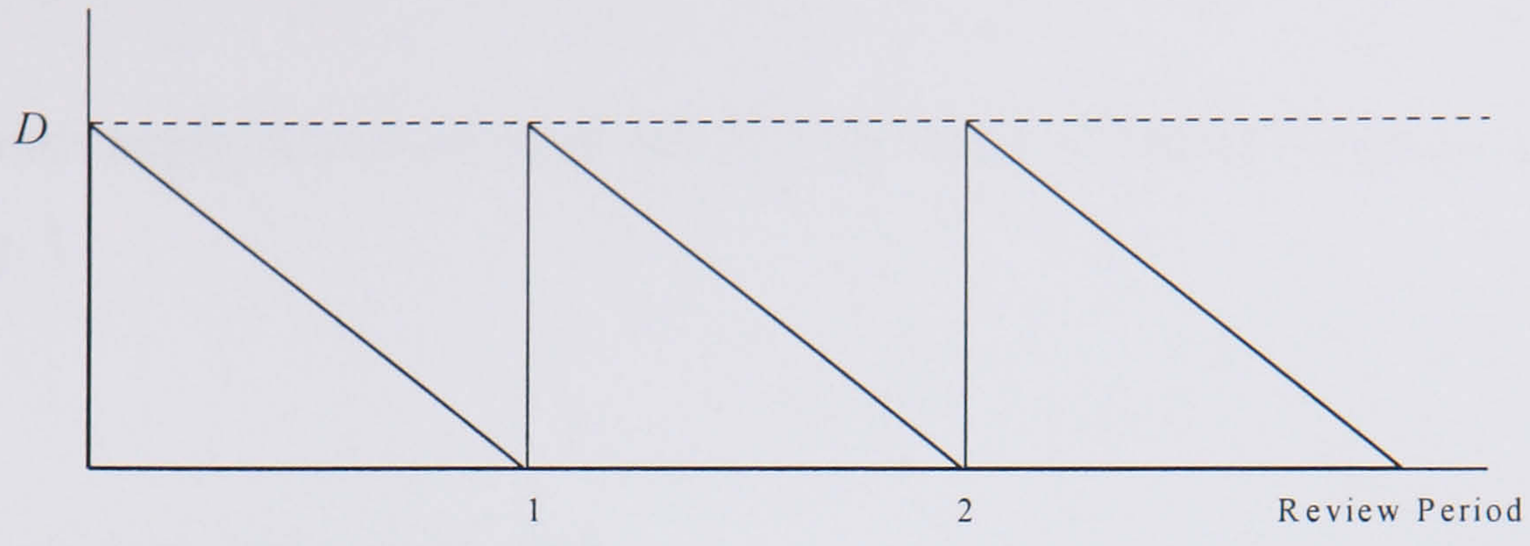
$$F_n = W_n = TD$$

Setting $O_n = D$ and $L_n = LD$, then from equation 1.0

$$\begin{aligned} B_n &= F_n - O_n - L_n \\ &= \left[T - (L+1)\right]D \end{aligned}$$

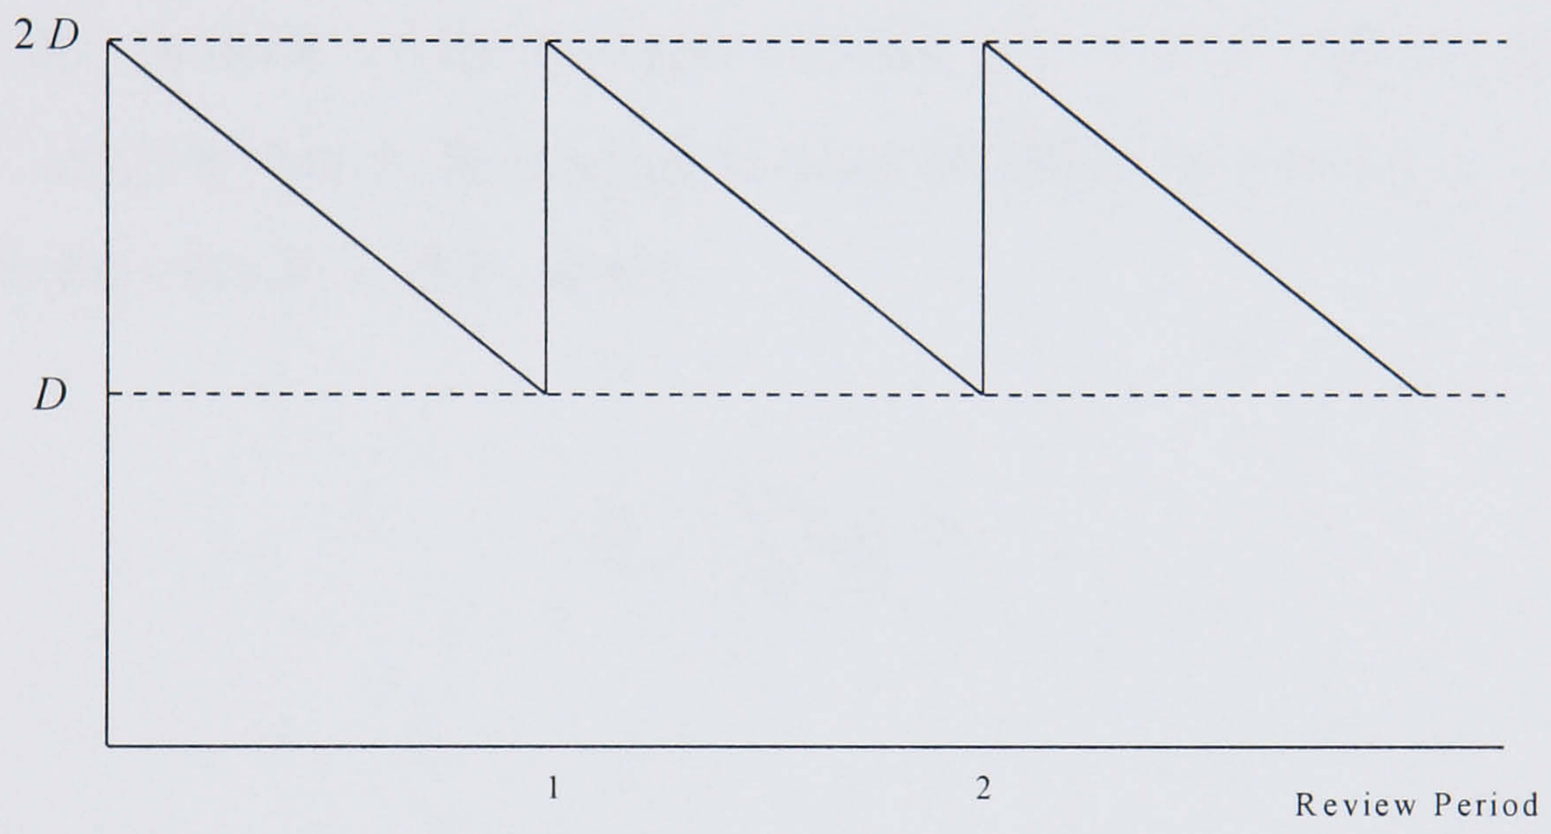
and if as is usual $T=(L+1)$, then $B_n=0$, and the solution is as illustrated in Fig. 1. The average stock level for this solution is $I=D/2$. We note that if parameter T is increased by unity to $T=(L+2)$ say, then $B_n=D$. For this solution, illustrated in Fig. 2, then average stock is $I=3D/2$.

As is reasonable, each unit increase in T above the minimum value of $(L+1)$ results in a corresponding increase in the average stock level by D. For a weekly review period, the wfc shown in Fig. 1 is one and the wfc for Fig. 2 is two. This illustrates that each unit increase in T results in a unit increase of wfc for the system. Parameter T will be used to control wfc in chapter 5. Note that the replenishment quantities are the same for both solutions. This effect will also be evident in the tuning of a realistic system as described in chapter 5.



$T = L + 1$ $\qquad\qquad$ $I = D/2$

Fig. 1



$T = L + 2$ $\qquad\qquad$ $I = 3D/2$

Fig. 2

4.2.2 Normal Demand Distribution.

If the demand is normally distributed with mean D and variance σ_d^2, per review period, then assuming independence, the average demand over T is $W_n = DT$ and the variance of W_n is $\sigma_d^2 T$. Equation 2.0 may therefore be expressed

$$F_n = DT + k\sigma_d\sqrt{T} \qquad 3.0$$

where k is chosen to guarantee acceptable stock out probability Lewis [70 - equation 5.3]. A well known estimate of average stock level, Maggee [58], is

$$I = D/2 + k\sigma_d\sqrt{T} \qquad 4.0$$

which assumes negligible stock outs and that the stock variations average out to those of Fig. 1

4.2.3 Poisson Demand Distribution.

Assuming as in the previous chapter, that the demand varies systematically and that if D_i denotes the forecast demand over review period i, then $\theta^2 D_i$ represents the variance of the demand estimate D_i, and θ^2 the dispersion parameter. At cycle time n, the average forecast demand over window T can be expressed in the form $W_n = \overline{D}_n T$, where

$$\overline{D}_n = \left(\sum_{i=n}^{n+T} D_i\right) \Big/ T$$

and the variance of W_n in the form

$$\theta^2 \sum_{i=n}^{n+T} D_i = \theta^2 \overline{D}_n T$$

Therefore equation 2.0 may be expressed in the form

$$F_n = \overline{D}_n T + k\theta\sqrt{\overline{D}_n T} \qquad 5.0$$

where k is chosen to ensure acceptable stock out probabilities. If the demand has mean $\overline{D}$ then a plausible benchmark estimate for the average stock is

$$I = \overline{D}/2 + k\theta\sqrt{\overline{D}\,T} \qquad 6.0$$

4.2.4 Batch Flow Effects.

In almost all modern retail sector supply chains, replenishment is in batches of specified size. The classical replenishment rule of equation 1.0 is therefore effectively modified. Consider the frequently occurring case, in which the flow is batched into multiples of some specified product dependant batch size M. Equation 1.0 is then effectively modified to

$$O_n = \left[F_n - \left(B_n + L_n \right) \right]_M \qquad 7.0$$

where if $X = cM + r,\ 0 \le r < M$, then

$$\begin{aligned} [X]_M &= cM & r &= 0 \\ &= (c+1)M & r &> 0 \end{aligned}$$

In order to develop some insight with the effects of batch flow on inventory levels, the simple ease of constant deterministic demand at level D was analysed and the following results obtained.[1]

RESULT 3

If the inventory demand level D is deterministic and both M and D are prime, where M is the batch size then for a fixed cycle system

(i) The average stock is $W = (M + D - 1)/2$.

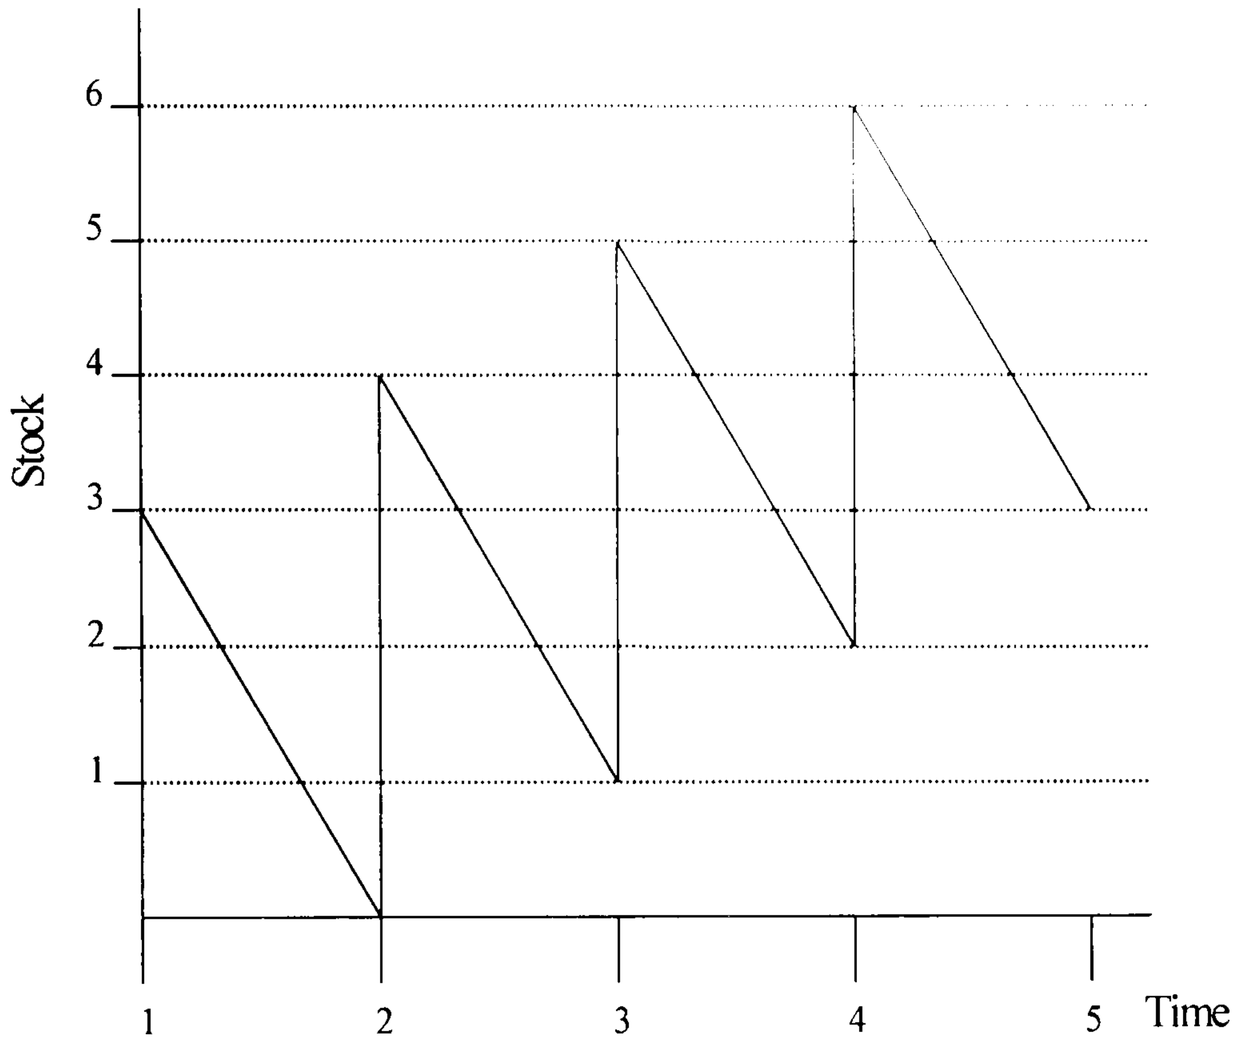
(ii) The stock level varies cyclically with period M.

(iii) During each period of length M,

[1] The results originally proposed by the author were proved by Dr. B. James in Bryant and James [96] together with more general results.

$$
\begin{aligned}
O_n &= (c+1)M \quad \text{for } r \text{ cycles} \\
&= cM \quad \text{for } (M-r) \text{ cycles} \\
\text{where } D &= cM + r, \quad 0 \le r < M
\end{aligned}
$$

The results are illustrated in Fig. 3 for the ease $M = 4, D = 3, L = 1$. Here the average stock is 3 agreeing with (i) and the replenishments satisfy (iii) with $r = 3$



n	1	2	3	4	5
B	3	0	1	2	3
L	0	4	4	4	0
O	4	4	4	0	4

Fig. 3 Batch Replenishment

Solution $(M = 4, D = 3, L = 1)$

Average Stock = 3.0

Substituting the average working stock estimate of result (i) into equation 5.0 we obtain the average stock estimate

$$I = \left(M + \overline{D} - 1\right)/2 + k\theta\sqrt{\overline{D}T} \qquad 8.0$$

This equation can be used to estimate of the effects of changing supply chain parameters on inventory level. For instance, Fig. 4 shows the estimated effects of varying lead time and dispersion coefficient on RDC warehouse stock level.

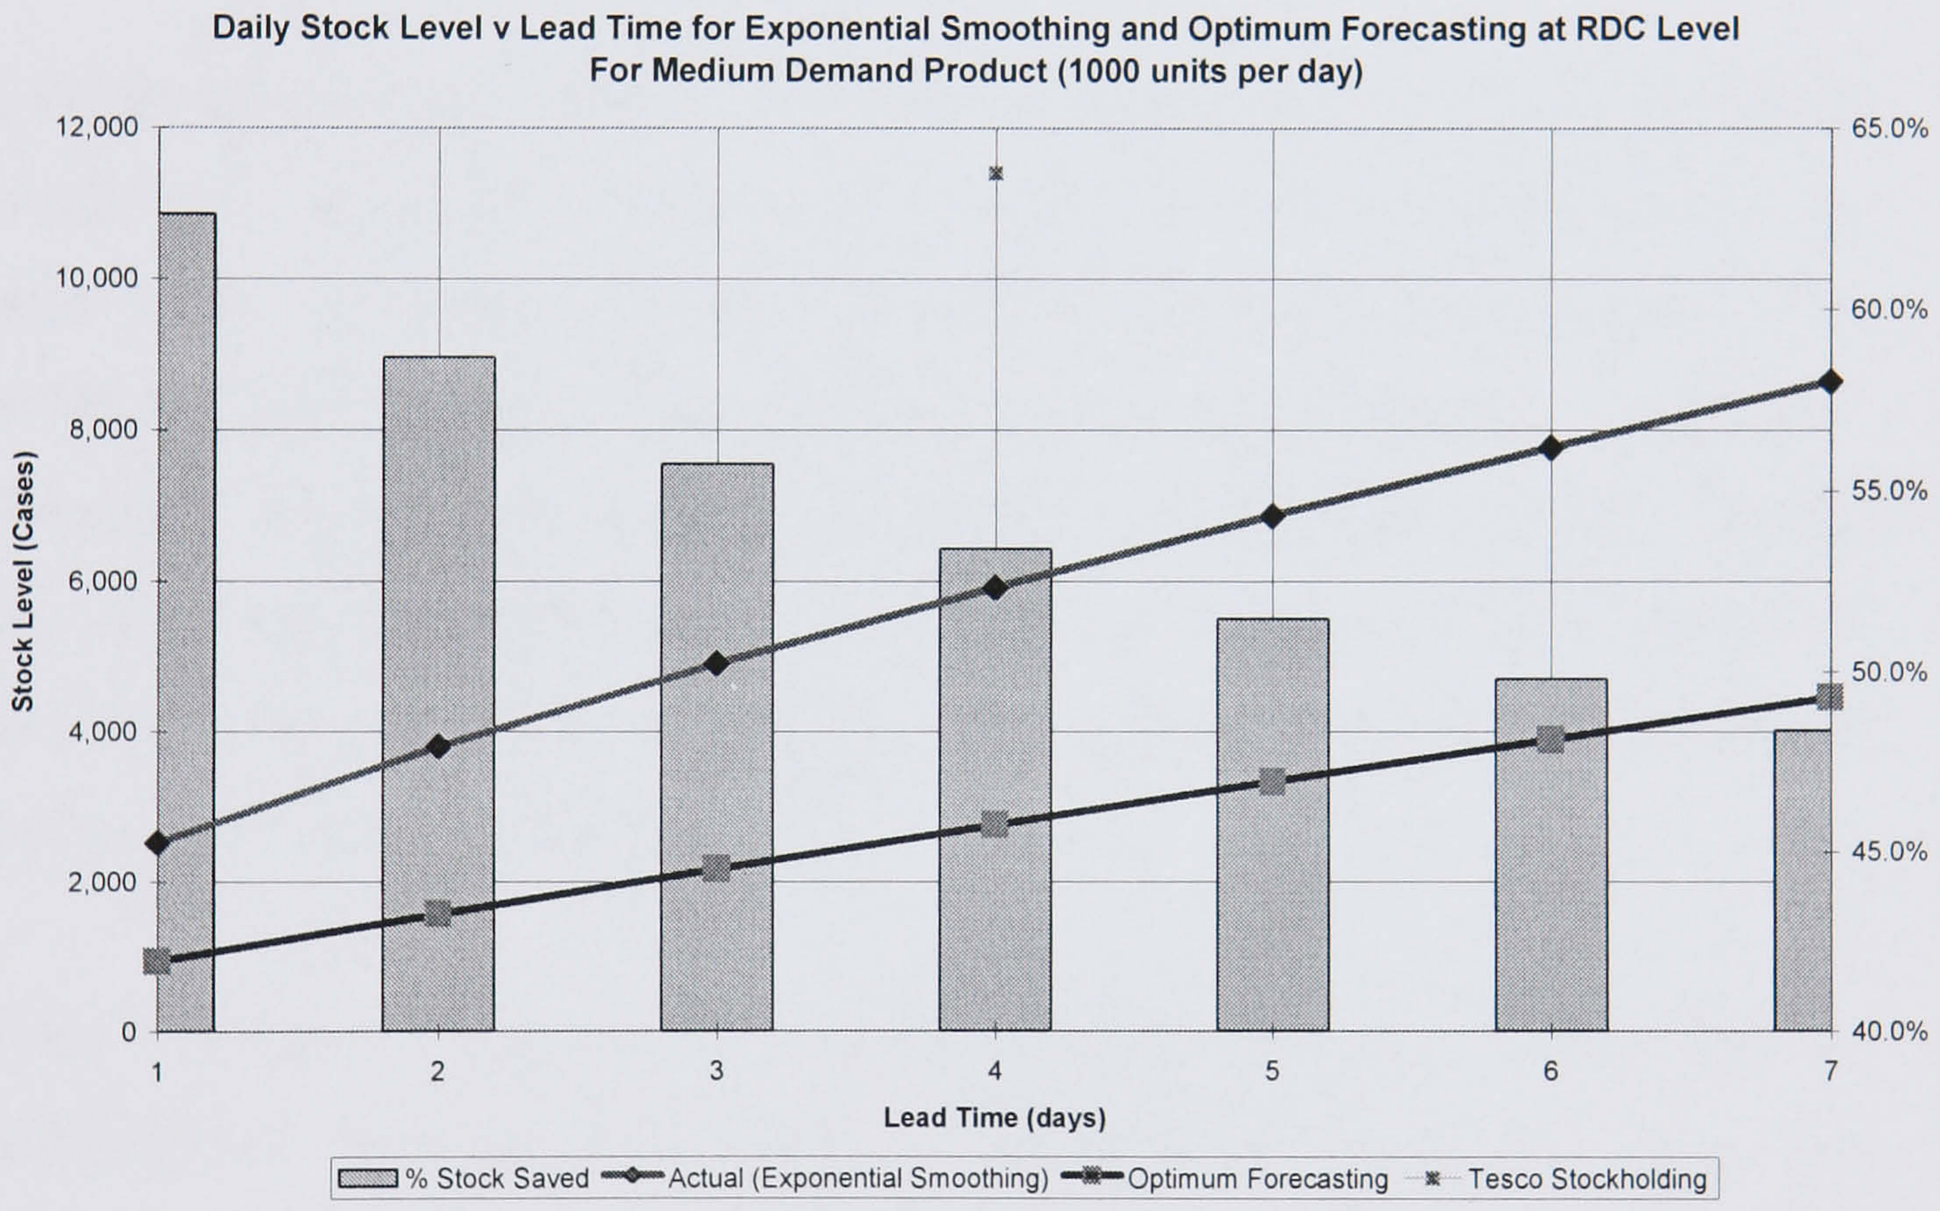


Fig. 4

The two dispersion coefficients were obtained using the procedure described in Chapter 2. The first was based on forecast errors for the procedures of Chapter 3 $(\theta = 2.0)$, the second on forecast errors for a simple exponential smoothing $(\theta = 8.0)$. As shown in Fig. 4, the actual stock levels were somewhat higher than the estimate of equation 8.0, which may be regarded as a benchmark. The dispersion parameter estimates for this example, would suggest that the choice of forecast procedure could, typically, make almost a fourfold difference in safety

stock level. Equation 8.0 will be used in the next chapter to estimate the effects of parameter changes.

Result 4 (iii) indicates that batch effects introduce cyclic changes into the replenishment characteristics. This might partly explain why the fitted dispersion coefficient at RDC level obtained in Chapter 2 $(2.5 \leq \theta \leq 5.0)$ was typically higher than the dispersion coefficient fitted at store level $(1.2 \leq \theta \leq 1.8)$.

4.3 Multi-product Inventory Control.

A single manufacturing facility will usually service a whole family of warehouse inventories. As the inventories are not serviced from independent sources, the nature of the shared resource must be taken into account in the operation of the inventory control system. Frequently the manufacturing facilities do not have the capacity to service all inventories in a single manufacturing schedule. In this case, the most critical inventories should be serviced first and this implies the scheduled items should be selected in terms of priority, the inventories nearest stock-out being given highest priority.

If the warehouse inventories are to be 'synchronously' replenished, then the replenishment quantities should arrive from manufacturing 'just in time' to meet demand. Since the demand cannot be perfectly forecast, the best that can be realised is that the inventory replenishments arrive 'just in time' to meet forecast demand. This requires that manufacturing schedules are determined by the requirements of the inventory control system as indicated in Fig. 5. It is obvious that if the orders placed on the warehouse (O) differ from this manufacture supply (Q) then the difference accumulates in the warehouse inventory.

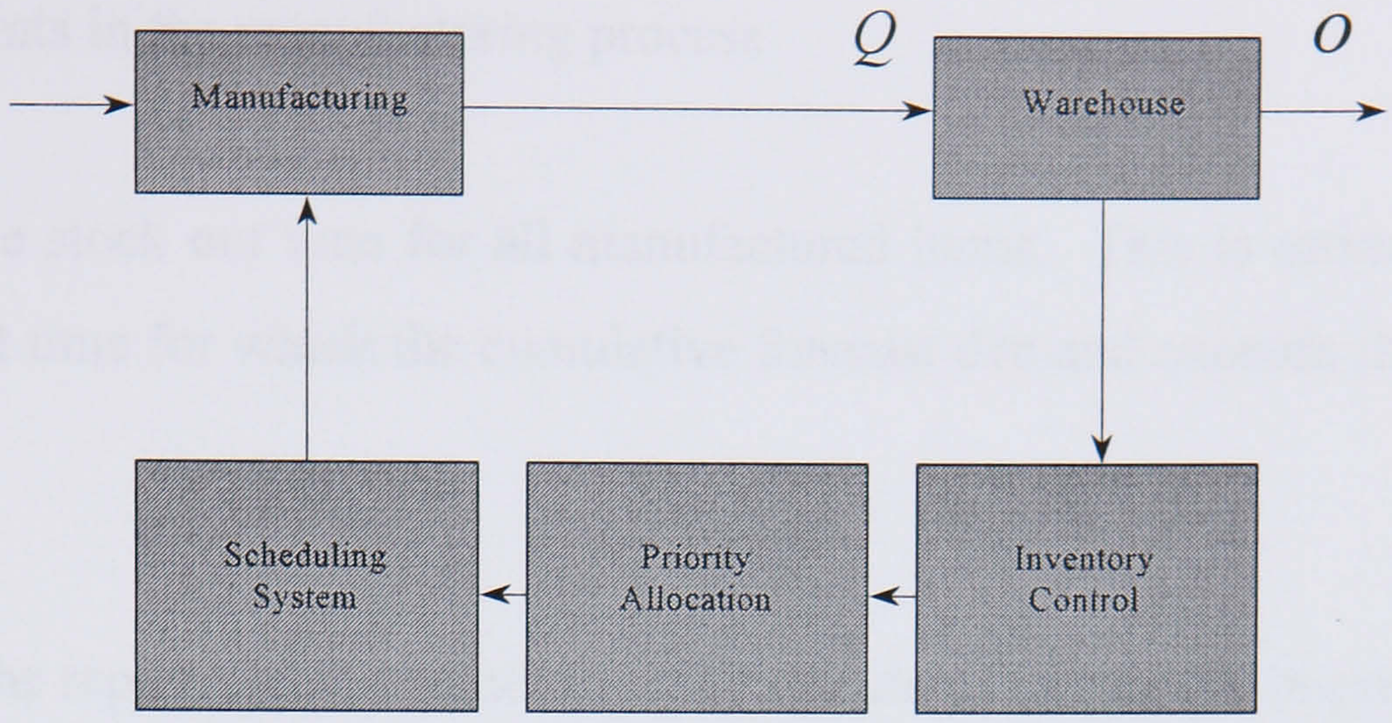


Fig. 5

If the manufacturing scheduling system and the warehouse replenishment system are to operate in 'synchronism' then the fixed cycle warehouse replenishment system parameters should be appropriately chosen - for instance, if the manufacturing facility is scheduled on a fixed cycle basis, then it is usually sensible to operate the stock control system on the same cycle (usually weekly in the retail sector). If the manufacturing schedules are, as is usually the case, subject to batch constraints arising, for instance, from material handling restrictions, then these batch constraints should be reflected in the fixed cycle inventory control system. As discussed earlier, the time window parameter (T) used in the inventory replenishment system, offers a convenient parameter for setting the inventory wfc parameter. This may be increased for reasons such as stock build up prior to an anticipated period of high demand. In the next chapter, these general concepts are applied to an industrial system design and the results assessed using industrial data. In this system, the following procedure was used for allocating schedule priorities, based on the estimated stock out time for each item manufactured.

(i) Calculate replenishment quantity (O_i) for all manufactured warehouse items, based on the fixed cycle stock control algorithm and taking into account batch constraints in the manufacturing process

(ii) Estimate stock out time for all manufactured items. This is estimated as the smallest time for which the cumulative forecast demand exceeds the available stock.

(iii) Rank the replenishment quantities (O_i) calculated in step (i), in priority order, those with smallest time until stock out being given highest priority

(iv) Schedule the replenishment quantities (O_i) in priority order until all the manufacturing capacity is used or until all items are scheduled

If an item is not scheduled in step (iv) then it re-emerges in the next cycle with new priority. It may therefore work its way through the priority list and actual replenishment may only occur every few cycles. This is likely when the manufacturing batch constraint is significantly higher than the demand over a fixed cycle. The use of stock out estimates to determine production priority is intuitively reasonable. It has been suggested in the classical literature in connection with the planning of production. Magee [58].

The combined system of Fig. 5, operates on a fixed cycle basis, balancing inventory supply against forecast demand, as is required by 'synchronous replenishment' approaches. The average manufacturing run length adjusts to maintain this balance. For instance, if demand falls then inventory replenishment and manufacturing run lengths falls (subject to minimum order constraints) since the fixed order cycle system controls to a specified wfc. Switch-over times therefore increase and the inventory supply decreases. This process continues until inventory supply and demand balance. Should the average run length become, for instance, unacceptably small then correction involves decreasing

manufacturing capacity in the system, usually in terms of available man-machine hours.

References

Battersby, A. 1970. *A Guide to Stock Control.* Pitman Publishing.

Brown, R.G. 1982. *Advanced Service Parts Inventory Control.* Materials Management Systems.

Bryant, D.N. and James, B. 1996. *Dynamic Analysis and Design.* Deliverable D3.3C Report, EAGLE D1.2B/96.1.

Lewis, C.D. 1970. *Scientific Inventory Control.* Butterworths.

Maggee, J.F. 1958. *Production Planning and Inventory Control.* McGraw-Hill.

5.0 MANUFACTURING DECISION SUPPORT SYSTEM

5.1 Introduction

This chapter describes a Manufacturing Decision Support System (MDSS) prototype, intended for use as an on-line management tool, an off-line tuning and parameter optimisation tool, and as a "what if" tool. The MDSS prototype is concerned with instants soup production and distribution. This operation involves about 50 Store Keeping Units (SKU), roughly 900 raw materials and spices ingredients and about 300 packaging materials.

The prototype MDSS system integrates the operation of manufacturing, scheduling and warehouse replenishment within one decision support system, as required for synchronous replenishment operation. The manufacturing scheduling system was programmed on an established commercial scheduling package, S-Plan, Greycon (1997) capable of providing the necessary speed of response and 'what if' capability needed. The system uses a discrete event simulation of the manufacturing process and contain a host management tools for decision making, such as Activity Based Costing, Gantt Charts, and resource usage, order status and expected delivery dates. A sample of the available displays is shown in the following text. The system is widely used in various sectors of industry world-wide. The system was programmed to full commercial standards and contained full and realistic details of the entire manufacturing process. The system was then extended to include the warehouse replenishment. It provided on-line decision making facilities and off-line parameter tuning facilities needed within the MDSS system. The whole package could be driven by off-line data recorded over a 6 month test period.

The layout of the Chapter is as follows. The functionality of the MDSS is described in the Section 5.2. The functionality covers both off-line and on-line use of the system, and include reactive scheduling and 'what if' facilities, knowledge based scheduling, machine and manpower features, and parameters tuning.

The system is model based and Section 5.3 covers the modelling of raw materials, packaging, manufacturing, and scheduling. This section also treats key displays developed within the system related to aspects of modelling.

Section 5.4 describes the "Synchronous Replenishment System" (SRS) which seeks to integrate the dynamic operation of the suppliers, manufacturing, scheduling, and warehousing operations. This is important in the attainment of reduced stock, and improving system responsiveness.

The system contains a substantial selection of graphical user interface (GUI) displays. These are used for editing the system and to produce detailed operational performance results. A selection of the key GUI displays is given in section 5.5.

For the purposes of this research, the system was used in an off-line mode, concerned with parameter tuning of the supply chain operation. In practice, this involves the tuning of key "parameters" within the supply chain systems, as described in Section 5.2.2 Off-line Operation, "Parameter Tuning", to ensure efficient operation of the overall supply chain.

Historic demand data and logged information on the "as is" supply operation are used. This data was collected over a six months period, and has been used to drive dynamic models of this supply chain. The data is processed sequentially by the system, as if it had been received in real time, and the result assessed for the test run for each parameter tuning and "what if" operation.

It is assumed in final implementation that both on and off-line modes would be available to managers. The system would have regular (usually daily) access to supply chain data, and would generate model based forecasts and provide "what if" facilities for use by management in decision making. It would also store past demand data to drive the system "off-line" in order to tune parameters and for answering 'what if' issues for instance, arising from proposed changes in manpower, shift and plant changes, and for the assessment of new product

launches and demand changes on this supply chain. The system will also supply model based forecasts for both retailers and suppliers. As will be apparent, most of these modes are straightforward to use and have established graphical user interface facilities.

5.2 System Functionality

The basic components within the MDSS system are illustrated in Fig. 1. The system has two distinct modes. The first mode is concerned with "on-line" operation and the provision of short term management tools. In this mode, the system is driven by current warehouse demand which is used to generate a short term forecast covering typically a few weeks duration. This forecasting procedure is built within the system and in the "on-line" mode the forecasts are based on current and past demand information for each SKU. The "off-line" mode of operation is concerned with the optimisation of the system operating conditions and parameters. In this mode, the system is driven by historic demand information. It may be desirable, for various reasons, to edit this information for "what if" purposes. Also as shown, the MDSS system employs an off-line procedure for calculating minimum safety stocks parameters for each SKU, taking into account uncertainties in the forecasting capability.

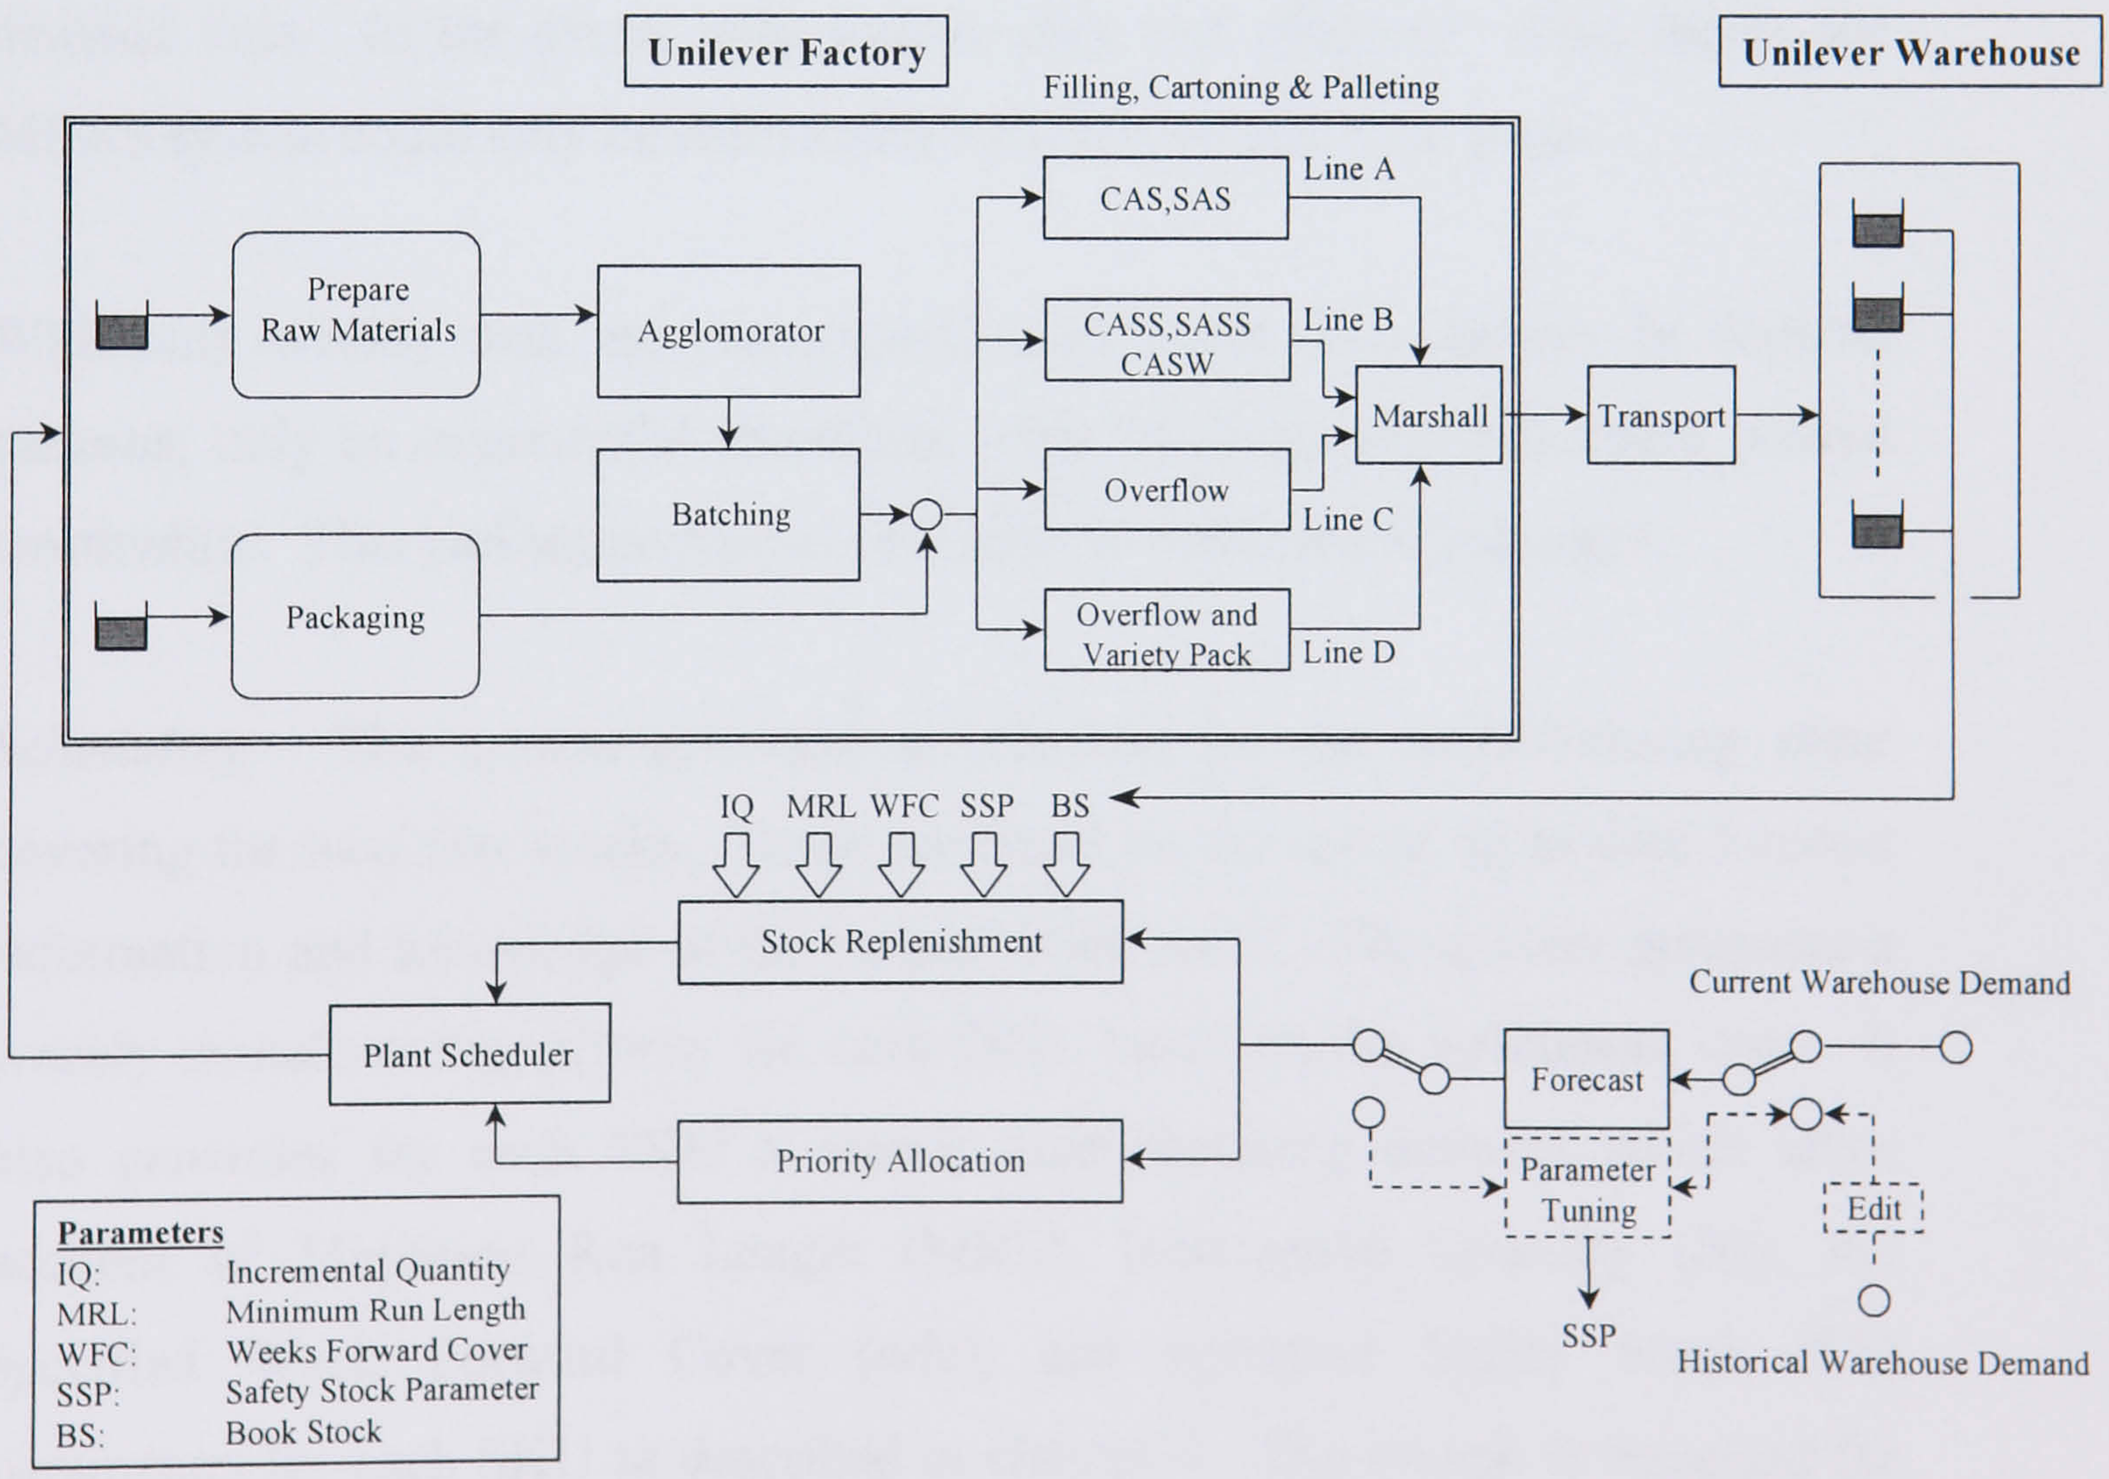


Fig. 1 MDSS/SRS System

Under operation the system provides fast response, normally under 30 sec. The on-line and off-line functionality are now described.

5.2.1 On-line Operation

Synchronous Replenishment System (SRS). The SRS system is driven by the real time warehouse demand signal as indicated in Fig. 1 and the system provides model based short term forecasts of supply chain operation. In addition, the system has a "what if" capability.

The test data used to generate these 'on-line' forecasts was recorded for each SKU for the period October 95 to April 96. It had been planned to record daily demand data. In the event, only weekly data was obtained. This meant the MDSS system could only be realistically run on a weekly time scale.

With only weekly data, no promotional information and changes in demand patterns, only an exponential smoothing class of forecasting procedure proved worthwhile. This yielded performance similar to commercial packages.

Scheduling. The system generates a schedule for the manufacturing plant covering the next few weeks. These are based on the use of up to date forecast information and knowledge of the current book stock. The system generates a weekly manufacturing priority for each SKU, based on the warehouse state. It also generates for each SKU a weekly manufacturing demand which takes account of Minimum Run Length (MRL), Incremental Quantity (IQ), the specified Week Forward Cover (wfc), and optimum Safety Stock (SS) parameters for each SKU as described in chapter 4. The system is designed for fast reaction to changes in demand at the warehouse levels. The manufacturing scheduling capability takes into account set-up, switch-over and labour

constraints and seeks to optimise the schedules, using knowledge based approaches.

Reactive Scheduling. Three most common disturbances are ingredient rejection due to quality defects, unexpected demand from a major account customer, and plant breakdown. Once information on such unplanned disturbances is fed into the MDSS system, it will forecast their effect on customer service and warehouse stock availability.

It may also be used to forecast when any failure in customer service or stock out can be resolved. Bearing in mind that the stock out will often be forecast rather than actual, this should allow the management time to consider possible actions and to inform a customer of a problem ahead of time.

In the event of raw material or packaging shortage being the cause of the stock out, the system will automatically forecast when the late order will be available, based on each supplier replenishment information. If as is often the case, a manager negotiates early replenishment, this can also be fed manually into the MDSS system.

The solution to a given short fall in supply usually requires a background knowledge of many business issues and hence the use of "what if" facility is essential. Possible causes of action may include changes in raw material supply, lead time, manpower, shift rosters, machine maintenance schedules, product manufacturing path and ingredient substitution. These and other changes may in principle be introduced on a manual basis, and their effects again forecast.

5.2.2 Off-line Operation

As shown in Fig. 1, when operated in off-line mode, the system is driven by historic warehouse demand data. The data is fed sequentially into the forecasting procedure which effectively generates "real-time" forecasts on a continuous basis. The current historical data relates to the period October 95 to March 96.

An important system function is that of parameter tuning. Here we attempt to supply best operating parameters for use in the SRS system. The off-line system may also be used in a "what if" mode to assess immediate and strategic options. Also, the off-line system has built in performance measures and display facilities for use in optimisation and tuning. We review the functionality below.

Parameter Tuning. Each of the parameters indicated in Fig. 1 (IQ, MRL, WFC, SSP) may be tuned for every manufactured SKU. Each of these parameters may be tuned to obtain minimum stock level within the warehouse, subject to the constraints of customer service and OTIF. The safety stock parameter (SSP) must be optimised within the limitations of the forecasting procedure. As discussed above, the forecasting procedure used is far from optimal. In order to deal with this a parameter tuning procedure is used. This involves the feeding of historical data through the forecasting procedure and the modelling of the error characteristics to obtain a safety stock policy constructed to meet the customer service level specifications. The tuning procedure effectively gives a minimum safety stock for each SKU, consistent with the forecasting errors and the customer service level specification. The SSP parameters used to obtain the MDSS results reported later were determined from off-line tuning.

"What if" Operation. "What if" operations are of interest in off-line operation and concerned with tactical and strategic, rather than with real-time and short-term issues. As shown in Fig. 1, the recorded demand for each or any product may be "edited" as required. An example of this occurs in the work reported below, involving the effects of a cold spell. To simulate this, all of the demands were increased by a prescribed amount, calculated from an assumed temperature drop, and the effects of this, over a prescribed period assessed in terms of stock out etc., using the edited demand.

The system can also be used to answer certain "what if" questions associated with strategic issues. For instance, the introduction of a new product can be assessed by introducing an assumed uptake characteristic and by introducing the new manufacturing paths into the system. This is a relatively easy exercise.

Similarly the system may also be used to asses the effect of changes in manufacturing capability. For instance, the introduction of new production facility can be introduced by simply changing the manufacturing paths, switch over costs, machine efficiencies, capacities and manpower requirements and then driving the modified models with the appropriate demand. Such a procedure has the advantage of presenting information in a way which is very familiar to the manufacturing staff and allowing assessment of changes under real life production conditions.

5.3 System Modelling

As shown in Fig. 2, the stages covered by the MDSS system include raw materials and packaging, manufacturing and warehousing. Each physical component is modelled together with its replenishment system and logic. In this section, we review the modelling of each of the processes and review some of the graphic user interfaces developed for model management and update. These are designed to be easily accessible to management and planning staff.

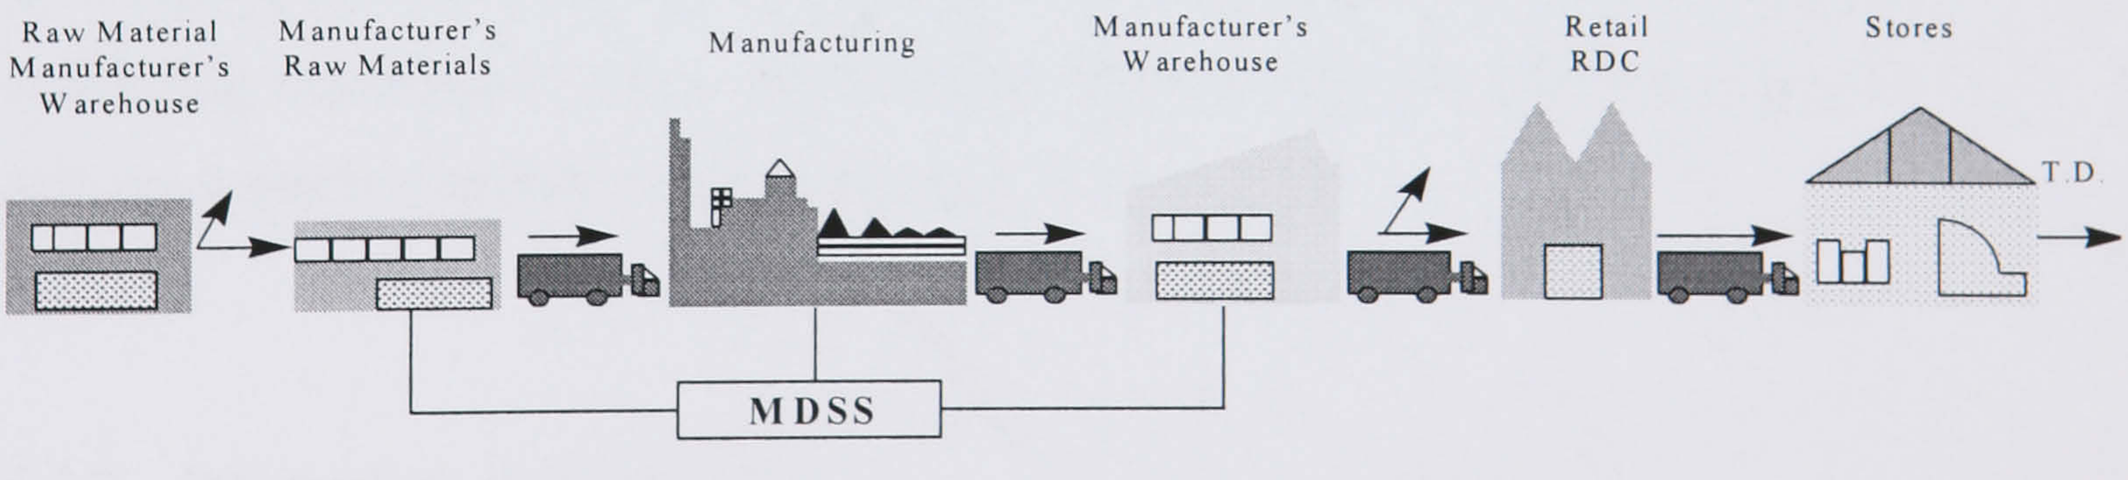


Fig. 2 The MDSS Supply Chain

5.3.1 Raw materials and packaging

As indicated in Fig. 1, the dried soup manufacturing process requires that a number of packaging and raw materials be available before the manufacturing process commences. Raw materials in the form of dried ingredients are required in strict measures in order to produce the required quantity of the final product. In addition, various spices in given ratios (the recipe) are required to make the desired quantity of soup. These are relatively expensive and are prepared ahead of manufacturing in the 'spice kitchen'. These give the products their distinctive flavour. Similarly, the dried vegetables are weighed in prescribed quantities and this affects the minimum run length and incremental constraints of the manufacturing process. Clearly, once materials are committed, there is a potential loss of reactive capability and hence operating the preparatory operations on a 'just in time' basis is advantageous. There will also be packaging material requirements further down the line at the filling stage. There is normally a limited quantity of packaging and raw materials available on site for processing and, in addition, some materials on order . The raw materials suppliers themselves have to schedule production and transport, therefore they impose a lead time and order quantity constraints, on deliveries. These are reflected in the MDSS system.

The MDSS takes into account these material requirements and constraints when planning operations. The MDSS can predict when materials will run out, suggest reordering materials to cover the shortfall and take into account the delay in starting a specific operation due to material unavailability.

5.3.2 Information Requirements

The MDSS uses the following information regarding materials:

Materials

For each material:

1. the current stock position

2. the minimum safety stock required
3. the lead time
4. minimum order quantity (MOQ).
5. Incremental Order Quantity. The minimum incremental increase in production quantity allowed once the MOQ is exceeded.

For materials which are in transit the MDSS must also have:

- quantity in transit
- estimated time of arrival

Bill of Materials Each MDSS operation can have its own bill of materials, identified by a unique identifier. For instance, operation mixing could have bill of material BOM18769 which defines that for each unit of product at this operation (e.g. kg) we need to have available the following materials:

1. kg dried tomato
2. kg sugar
3. kg salt
4. kg pepper

The MDSS ensures these material requirements are available from one of three possible sources (or a combination there of):

- directly from available stock
- materials in transit
- Issued replenishment orders

Depending on the sources used, the MDSS will delay the start of the operation until the materials (are estimated to) become available.

5.3.3 Raw Material Scheduling

The raw material suppliers operate to an agreed lead time, typically of 2 weeks duration and supply amounts with an agreed Minimum Order Quantity (MOQ). The MDSS system takes these into account when scheduling and planning operations. If a material replenishment is forecast, the MDSS will compute the correct re-order quantity and assume the agreed lead time. If a material stock out is forecast, the MDSS system computes a re-order quantity to cover the short fall, and delays the relevant operations associated with the short fall, recommencing production when the material becomes available. These replenishment operations are seen as advisory and would normally be transmitted to the main system only after confirmation by management. If a short fall is forecast together with a corresponding operational delay and customer service failure, then the management may negotiate with the supplier material delivery on a shorter lead time, in which case the agreed new delivery may be added, and the system re run to obtain updated forecasts on short fall etc.

Substitute materials - The prototype version of the MDSS does not deal (at least automatically) with the issue of substitute materials. The user can only substitute the material requirement for a particular operation manually if this is required. Two refinements are planned for real implementations. These are:

- a table with a list of substitute materials indicating whether the change represents an upgrade or downgrade and the financial implications.

- an automatic search in the list of substitutes if material deficit is predicted, aimed at minimising production delays.

Alternative suppliers - At present alternative suppliers with different lead times and prices are not considered by the MDSS. A refinement could search for alternative supplies and lead times if the order is critical in terms of delivery time.

5.3.4 Manufacturing Processes

The main activities here are indicated schematically in Fig. 1. The manufacturing activity consists of four lines (ABCD respectively) and these are all housed in Hall 1 of the factory. As indicated, the lines undertake the filling, cartoning and palletising activities.

Line A produces Cup A Soup (CAS) Slimmer Soup (SAS). Line B produces Cup A Soup Special (CASS) Slimmer Soup Special (SASS) Cup A Soup World-wide (CASSW). Line C is mainly dedicated to overflows from Line B. Line D also may handle overflow activities from Line B and is otherwise dedicated to the production of variety packets. These contain selections from the main classes of soup listed above. One of these may be produced at the time of assembly of the variety pack. The remainder in any given selection have to be produced prior to assembly of the pack. Hence, manufacture of variety packs couples the scheduling of Line D to the scheduling of Lines A, B and C.

The manufacture of the soup product involves seven steps :

Step 1: Creates spice mix in the spice kitchen . These are relatively high value products.

Step 2: Pick the various material ingredients and weigh them according to the Bill of Material [BOM] listing.

Step 3: Feed the selected spices and the weighed materials into the agglomeration process. Move agglomerated material to the batching operation. (the agglomeration process converts powder into a granular form now associated with instant soups and coffee).

Step 4: In parallel with batching operation prepare the packaging materials for each of the products being scheduled in Hall 1.

Step 5: Set up the line and initiate the filling, cartoning and palletising operations. Excluding set up, this operation will frequently involve

cleaning operations which are dependant on sequence of products being produced. There can be significant switch over costs involved and this must be taken into account in scheduling.

Step 6: When the current palletising operation is complete, enter the order into the system.

Step 7: Update BOM to revise stock level records.

The timing and efficiency of the manufacturing activities are carefully monitored. Statistics exist for the values and variability of the key processing times involved in modelling the manufacturing processes.

5.3.5 Production Scheduling

In order that the MDSS model forecasts with reasonable accuracy, account must be taken of the constraints of the real production process. These include man power and machine capacity constraints, switch over costs, run length constraints, machine efficiency and planned maintenance programs. The important factors here, include the following:

Switch over costs - There are no significant switch over costs in evidence within the preparation phase, the main switch-overs occurring in production. The switch over costs on these lines essentially dominate the stock control replenishment operations in the warehouse and hence are of considerable importance. The switch over costs may be expressed firstly in terms of lost production time, and secondly in terms of material wastage, both components being product dependent. In order to express scheduling approaches used in practice, a rule based approach was employed. The rules were used hierarchically and may be summarised as follows.

- *Meat to vegetable* - A product change such as this involves a high cost in terms of switch over. For religious and diet reasons some customers will be very concerned, for instance, that a vegetable based soup does not

contain traces of meat (the Vegetarian Society inspect all processes for food products they endorse). To ensure this, a very thorough machine cleaning operation is involved, with a consequent rise in down time and material wastage. The down time for switching between such products can be as long as half a day.

- *Light to dark* - In order to ensure high quality of product the schedules are usually "zigzagged" on a light-dark-light basis. The time for a switch over change from light to dark is typically 30 to 45 minutes. This policy is built into the system.

- *BM number* - Within a given run of products within similar group, it is desirable to sequence the runs keeping the BM numbers as close as possible. This ensures that closely related products are sequenced together, thus ensuring minimum down-time and change over costs.

- *Despatch time* - In the Synchronous Replenishment System (SRS), each run has an associated warehouse latest delivery time, computed to take accounts of priorities within the warehouse. For reasons discussed above, it is currently only possible to both forecast and schedule the warehouse on a weekly basis and hence despatch times are of less importance. However, in an order based system, or when the warehouse forecasting and the SRS system are operated on a daily basis this facility would be more important in ensuring maximum reactivity to warehouse and customer needs.

The current scheduling logic involves sequencing products by line, and then in an order determined by the hierarchical considerations discussed above. In other words, products on a given line are scheduled first of all in terms of meat/vegetable considerations and then light-to-dark considerations etc.

Emergency rules - These are not used in the current studies in view of the weekly constraints on the scheduling operation. In a daily or hourly based time scale of

operation it may be worthwhile to have the facility to over-ride the hierarchical rules described above, should the situation arise. The system has facilities for a number of these emergency rules but these will not be used in obtaining the results described here.

Start Up - Start up costs are associated mainly with the manufacturing production lines. It is well known that these lines cannot generally be started automatically. An operator's skill is required in order to get the lines up to speed. Standard times have been agreed for this phase of the schedule.

Minimum Run Length - Minimum Run Length (MRL) constraints are imposed on all products and the "as is" (MRL) constraints have been adopted here. This avoids schedules having sequences of tedious and small run lengths which would be inefficient. Thus for low demand products, minimum run lengths are imposed which are usually in excess of the requirements for stock replenishment purposes. This produces high stocks of low demand products, having the advantage of allowing production to occur in low activity times thus freeing manufacturing processes for high demand products during high demand periods. The constraints used by experienced management appear to produce well balanced schedules as will be illustrated later.

Capacity Constraints - The manufacturing process works on a three eight hour daily shifts, five day roster. Overtime is possible at weekends and is paid on a time and a third basis. During manufacturing, machines do not always work at 100% efficiency, and hence agreed machine efficiency factors have been built into the models. Furthermore, machines have to be maintained to a high level and a maintenance shift program is allowed within the models. The agglomerator is not capacity constrained.

As described below, the SRS system generates a schedule associated with each SKU in the warehouse, together with a schedule priority allocation. These are then scheduled using the approaches described above, until they are all scheduled, or the available resources for the week are all used. The manufactured products are then fed through the warehouse and the appropriate

book stocks' levels are updated. For reasons discussed above, the book stocks update can only occur on a weekly basis.

5.3.6 Product Demand Database

The model is driven by total weekly demand data, by product, on the Peterborough warehouse. This information was recorded for all of the dried soups. Actual weekly product stock level was also recorded in order to benchmark the performance of the model. Below is a list of database tables used to run the simulation:

Product Demand - This contains the weekly demand sample used to drive the simulation, for each product. The period covered is week 40 of 1995 to week 23 of 1996.

Field	***Type***	***Length***	***Description***
WEEK_NO	Character	6	Year and Week Number
CODE	Numeric	7	Product Code
QUANT	Numeric		Product Demand

Product Details - This contains the stock control and forecasting parameters for each product.

Field	***Type***	***Length***	***Description***
CODE	Numeric	7	Product Code
DESC	Character		Product Description
MOQ	Numeric		Minimum Order Quantity
IQ	Numeric		Incremental Quantity
WFC	Numeric	2	Weeks Forward Cover
FP	Numeric		Parameter for Forecasting
SSP	Numeric		Parameter for Safety Stock Calculation

Product Stock - This contains the weekly product stock levels. It is a temporary table, holding information generated by the simulation for subsequent analysis.

Field	***Type***	***Length***	***Description***
WEEK_NO	Character	6	Year and Week Number
CODE	Numeric	7	Product Code
QUANT	Numeric		Product Stock

Actual Weekly Product Stock - This contains the actual weekly product stock for each product.

Field	***Type***	***Length***	***Description***
WEEK_NO	Character	6	Year and Week Number
CODE	Numeric	7	Product Code
QUANT	Numeric		Product Stock

5.4 Synchronous Replenishment System

The key components of the Synchronous Replenishment System (SRS) are illustrated in Fig. 1. The system aims to dynamically couple the warehouse replenishment operation with that of the manufacturing plant scheduler, in order to minimise stock levels in the warehouse, whilst at the same time maintaining specified customers service levels at the warehouse. The SRS system is designed for use both in the real time mode, in which the system is used for management decision making, and in the off-line mode in which the system seeks to tune the key parameters in order to set operating conditions.

The stock levels in the system have to be maintained in order to satisfy a number of differing and conflicting objectives. First and foremost, the designated level of customer service and the maintenance of the required OTIF statistics are of paramount importance. Thus stock levels throughout the system must be adjusted to ensure that disturbances can be absorbed within the system through the provision of suitable safety stock levels. Potential sources of disturbance include unexpected uplift during promotions, a sustained major drop in temperature which may produce a rise in demand for soups, stock out of raw materials, a quality control failure or plant machinery failure. In principle, all of these can be included in the model and their affects forecast and assessed.

The system is concerned with the manufacture and supply of almost 50 SKUs to the warehouse level. This in turn involves the maintenance of supplies of approximately 900 ingredients and 300 packaging and cartoning material entities. The remainder of this section is concerned with describing the principles adopted in the SRS design and each of the functions indicated in Fig. 1.

5.4.1 Forecasting

Forecasts based on exponential smoothing are generated sequentially for each SKU. In "real-time" operation, the forecasts are generated from past and current warehouse demand information. In "off-line" operation, the forecasts are

sequentially generated using the historical data, recorded over a period of six months. The forecast procedure may also be used for tuning the safety stock parameter (SSP) as described below.

5.4.2 Stock Replenishment

The replenishment algorithms were based on the fixed cycle principle discussed earlier. The lead time and cycle time parameters were set to one week. Uplift in demand, due to seasonal and weather effects, may result in demands which are in excess of the plant capacity. It was therefore customary to control the stock levels in order to maintain a specified number of wfc to act as a buffer against unexpected surges in demand, expected during high demand periods. The stock required to maintain the forecast weeks forward cover in winter were manufactured during the off-peak demand period and maintained through the high demand period to ensure that customer service was maintained. The tuning of the wfc parameter for each SKU is important in tuning the system to obtain minimum stock level. The improvements obtained from improved tuning are described in the next section.

5.4.3 Priority Allocation

It will generally be impossible to manufacture all of the product quantities calculated in the stock replenishment algorithm above during one week. The priority allocation algorithm described in the previous chapter, was used to sort the replenishment quantities calculated by the stock replenishment algorithm, into priority order, according to the dynamic state of the warehouse.

5.4.4 Production Scheduling

Firstly, this calculation sorts the entire list of stock replenishment quantities according to the rules in section 5.3.5 . Within each line allocation the quantities are then sorted according to designated priority allocation. The products are scheduled on each line, using the normal scheduling rules employed in manufacture, until the available plant capacity for the week is fully utilised. Replenishment quantities which cannot be manufactured within the week are discarded from the list.

As is apparent from the description above, the generated schedules seek to take account of all of the manufacturing characteristics including switch over times, planned maintenance, current machine efficiency, the MRL and the IQ run length constraints, and the risk of stock out in the warehouse, in order to compute the current production schedules.

In a system implementation based on daily operation, the system may be used to regularly update the schedule to take account of any unexpected changes in manufacturing resources or warehouse demand. In this mode of usage, the line schedules may be changed at short notice. This facility would be especially useful if order based scheduling is introduced. A schedule change has to take account of the current state of the plant and cannot be changed immediately.

5.5 User Interface Design

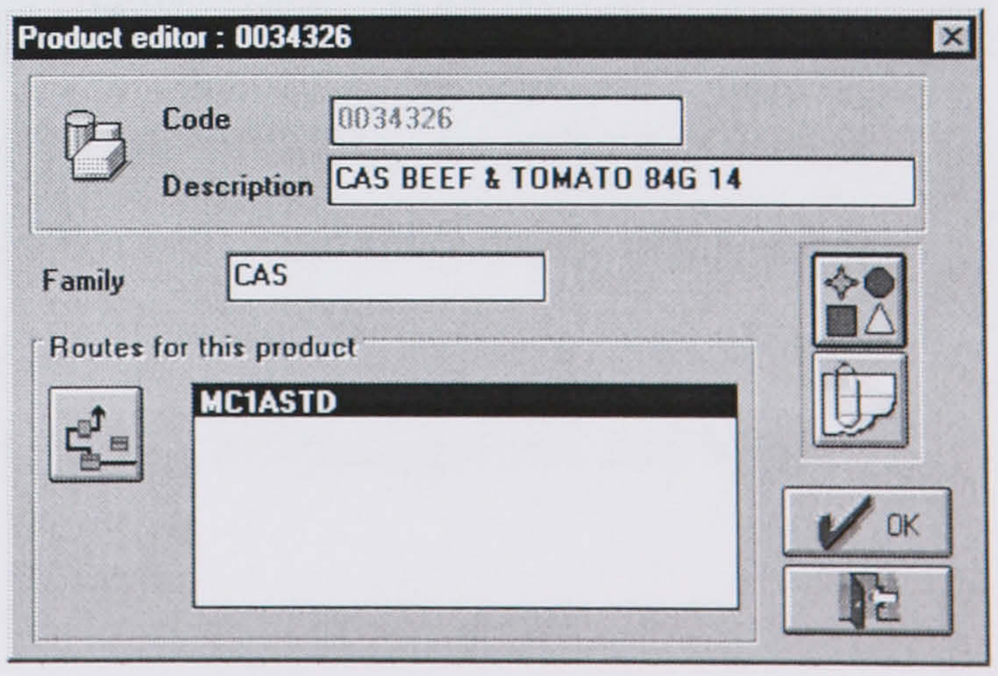
5.5.1 Products

The browser snapshot below shows a number of the 50 products that were produced together with some of their details and specifications. The user is able to drill down to a specific product by double clicking with the mouse on one of the rows.

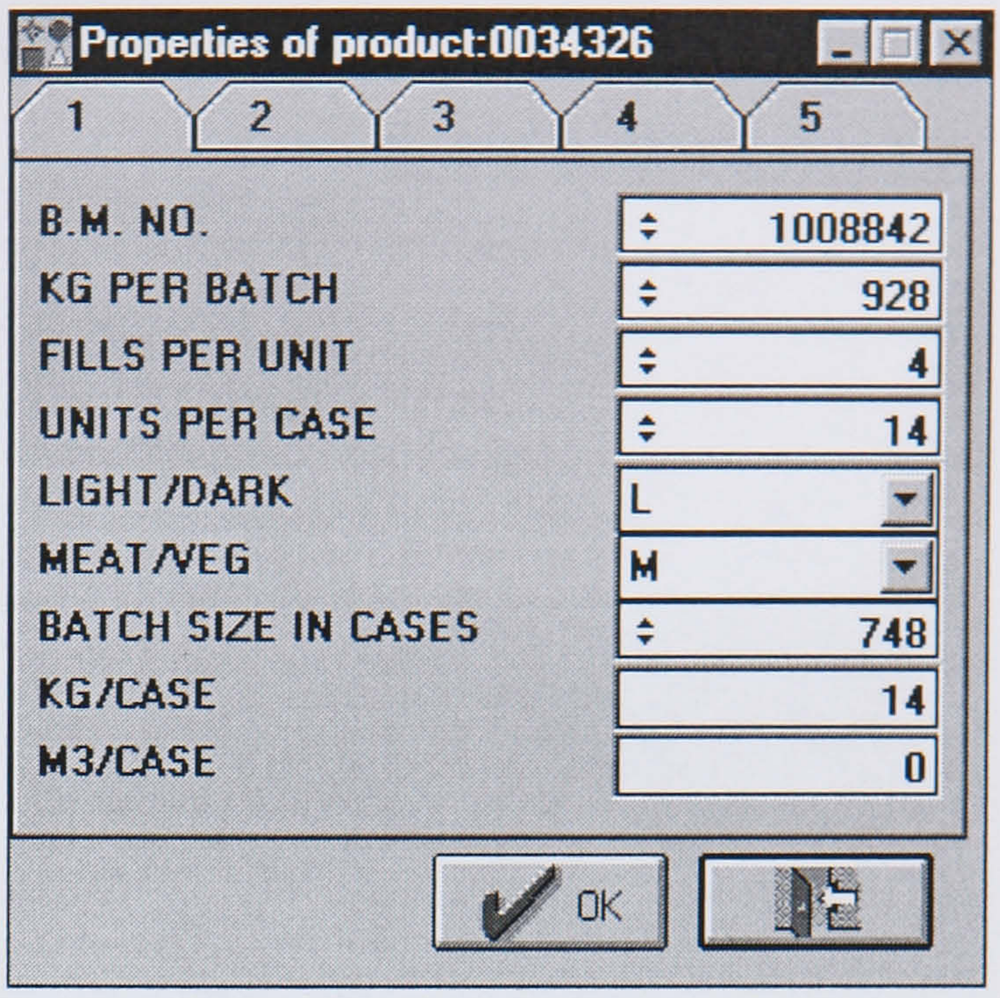
Product browser (Total: 102)

Product code	Product family code	Description	B.M. NO.	KG PER BATCH	FILLS PER UNIT	UNITS PER CASE	LIGHT/DARK	MEAT/VEG	BATCH SIZE IN CASES
0034326	CAS	CAS BEEF & TOMATO 84G 14	1008842	928	4	14	L	M	748
0034338	CAS	CAS CHICKEN 84G 14	1008526	769	4	14	L	M	606
0034347	CAS	CAS GOLDEN VEG. 80G 14	33997	900	4	14	L	V	759
0034348	CAS	CAS OXTAIL 82G 14	1008844	841	4	14	D	M	693
0034349	CAS	CAS TOMATO 92G 14	34054	970	4	14	D	V	716
0034399	CAS	CASWR FR. CHK & ASPAR 92G 9	1007919	847	4	9	L	M	953
0034401	CAS	CASWR MEX SPICY TOM 100G 9	1003027	640	4	9	D	V	679
0034402	CAS	CASWR CHIN S'C & CHK 70G 9	34598	592	4	9	L	M	882
0034403	CAS	CASWR FR. LEEK & POT 108G 9	34600	783	4	9	L	V	771
0034404	CAS	CASWR D.CAULI & BROC. 98G 9	34602	1056	4	9	L	V	1099
0035562	CAS	CAS CHICKEN & LEEK 78G 14	1008529	825	4	14	L	M	694
0036164	CAS	MCD CAS CHICK (621) 14XPKT	1008526	769	4	14	L	M	604
0036168	CAS	MCD CAS G.VEG (622) 14XPKT	33997	900	4	14	L	V	757
0036169	CAS	MCD CAS OXTAIL (620) 14XPKT	1008844	841	4	14	D	M	691
0036729	CAS	MCD CASWR MEX. TOM (631) X9	1003027	640	4	9	D	V	676
0036730	CAS	MCD CASWR CAULI/BR (632) X9	34602	1056	4	9	L	V	1095
0036731	CAS	MCD CASWR FR.LK/POT (633) X9	34600	783	4	9	L	V	769
0036732	CAS	MCD CASWR FR.CHK/AS (634) X9	1007919	847	4	9	L	M	950
1001567	SNAS	SNAS BACON & TOMATO 42G 12	1001418	596	1	12	D	M	1069
1001568	SNAS	SNAS CHILLI BF&TOM 44G 12	1008838	630	1	12	D	M	1071
1001605	SNAS	SNAS CHICKEN&MUSH 42G 12	1001418	561	1	12	D	M	1006
1001606	SNAS	SNAS SPICY VEG. 44G 12	1003010	5044	1	12	D	V	857
1001758	CAS	CASWR VARIETY PACK 8	1003027	640	5	8	D	V	1523
1001767	CAS	CAS VARIETY PACK 9	1008526	0769	6	9	L	M	1880

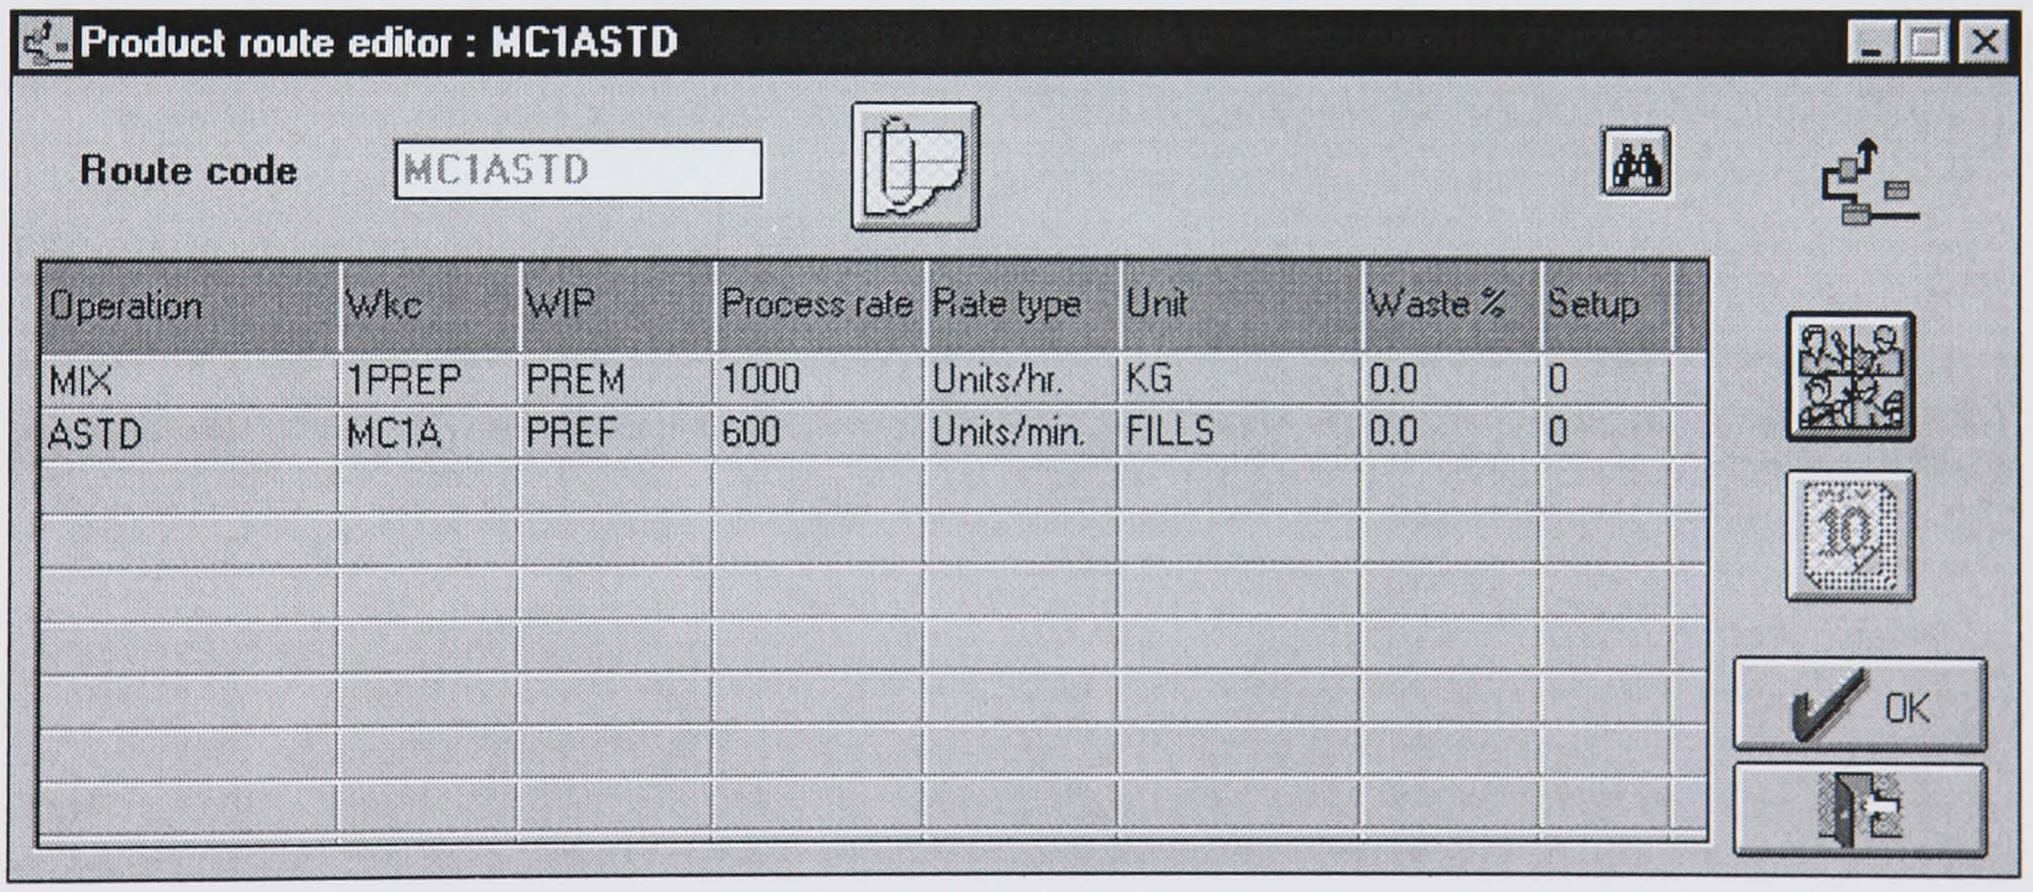
On the right hand side the product editor is shown for product 0034326 with associated details. One of the buttons leads to the product specification as shown in the icon below.



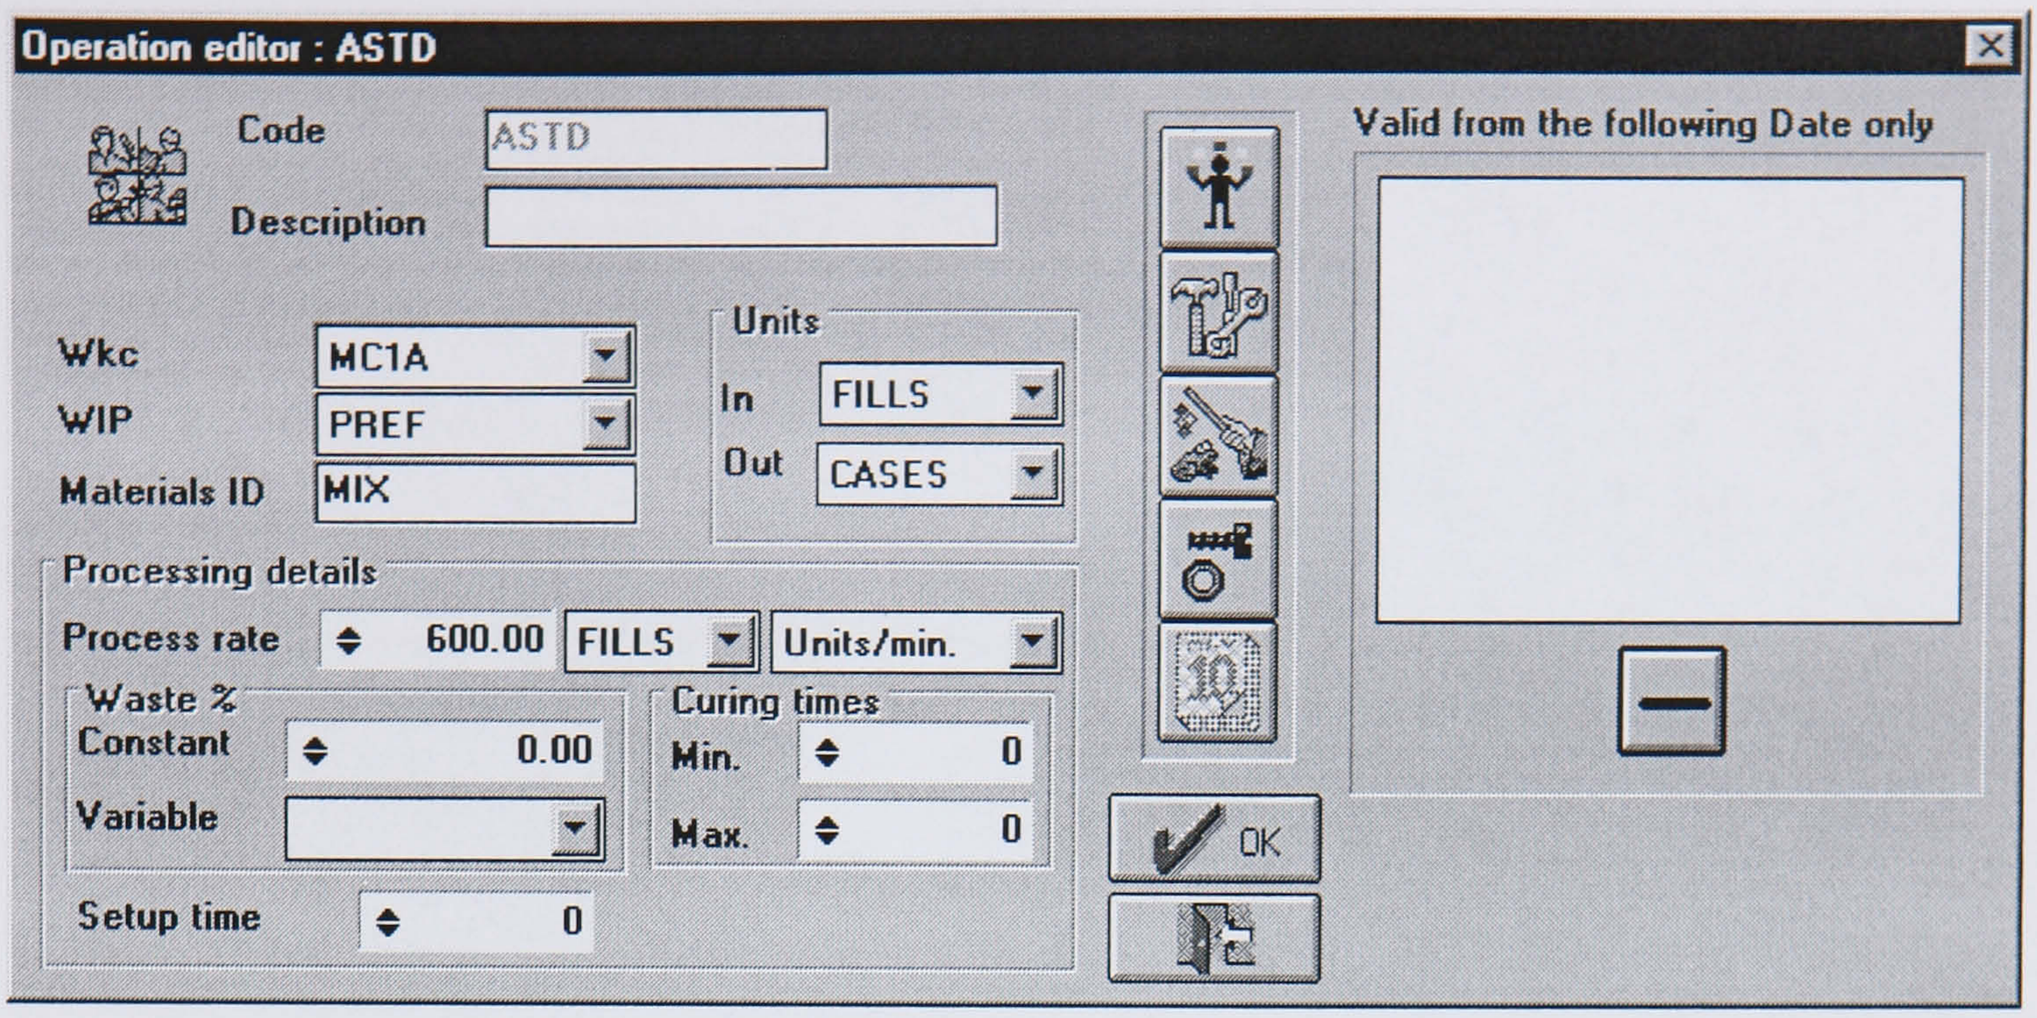
The Specifications defined here are used throughout the system either for unit conversions e.g. how many fills to produce in order to generate the required number of cases, or for sequencing the production so that switch over times (set-up times that depend on what was produced previously) are reduced.



The screen below shows the route of production for a particular soup. It includes two processing operations, a preparatory mixing operation (MIX) and a filling operation on line A. In this case, the product goes specifically to this line, but it is possible to specify that the product is produced on an alternative line if capacity is exceeded.

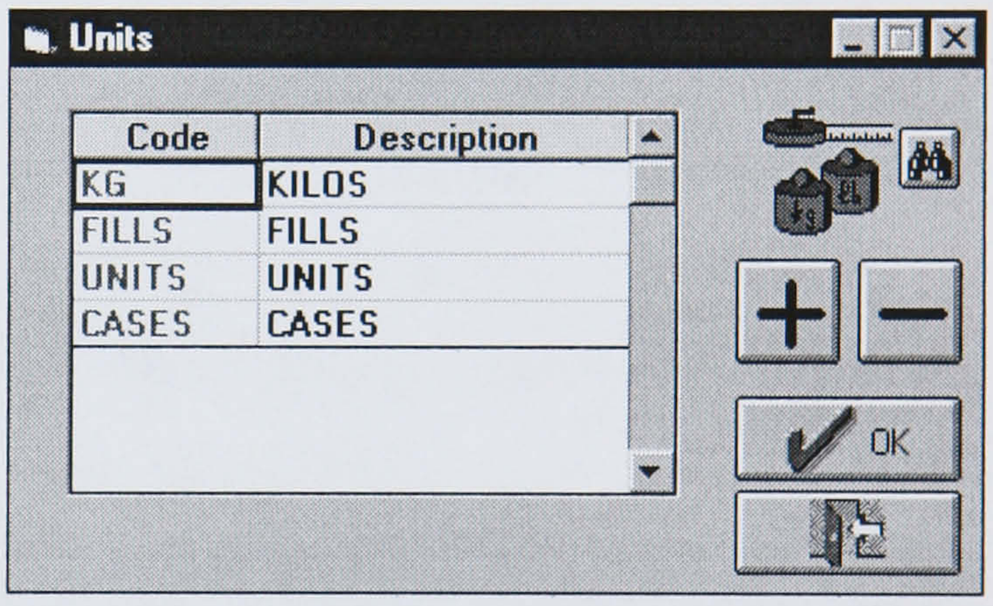


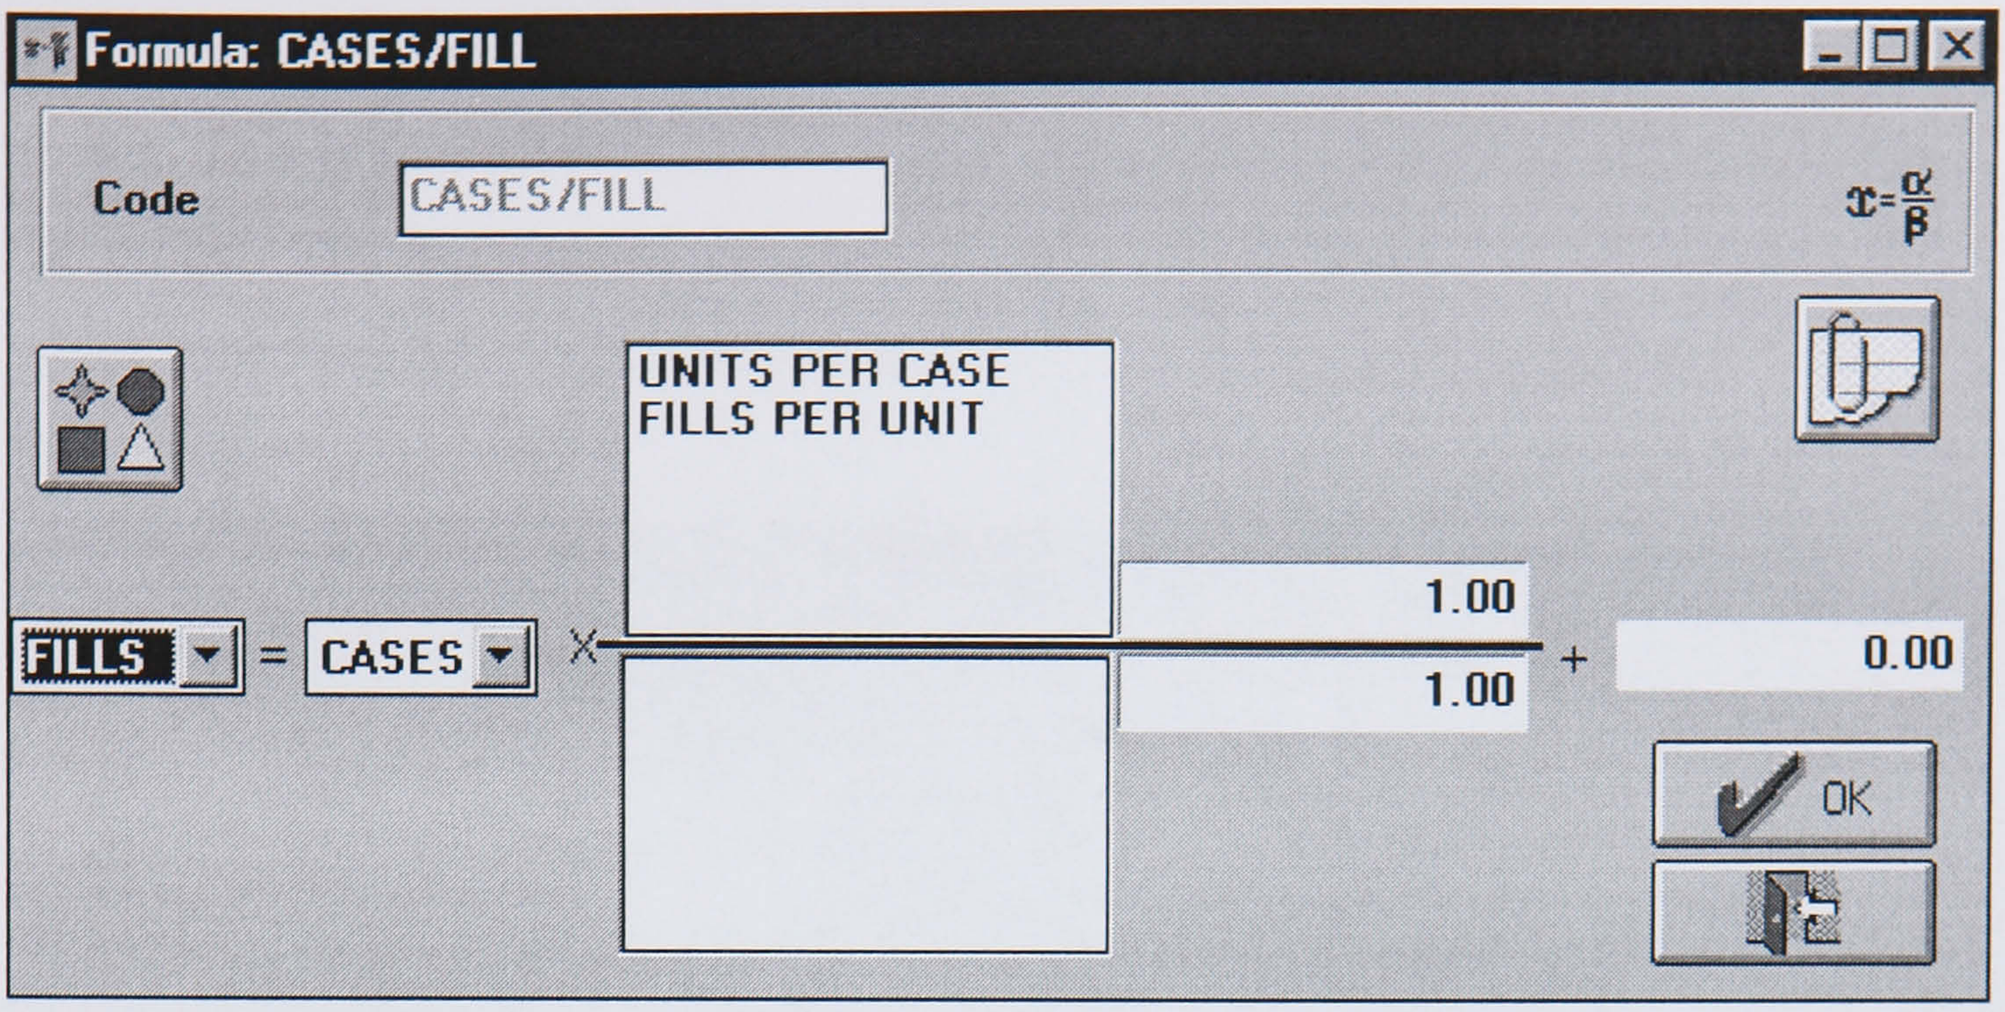
The filling operation can be viewed in detail as shown below by selecting one of the rows from the above screen. It defines the Line where the processing takes place, the WIP area where the material is coming from (in this case the mixing area), the input/output units, the processing rate and any curing times and loses.



5.5.2 Units

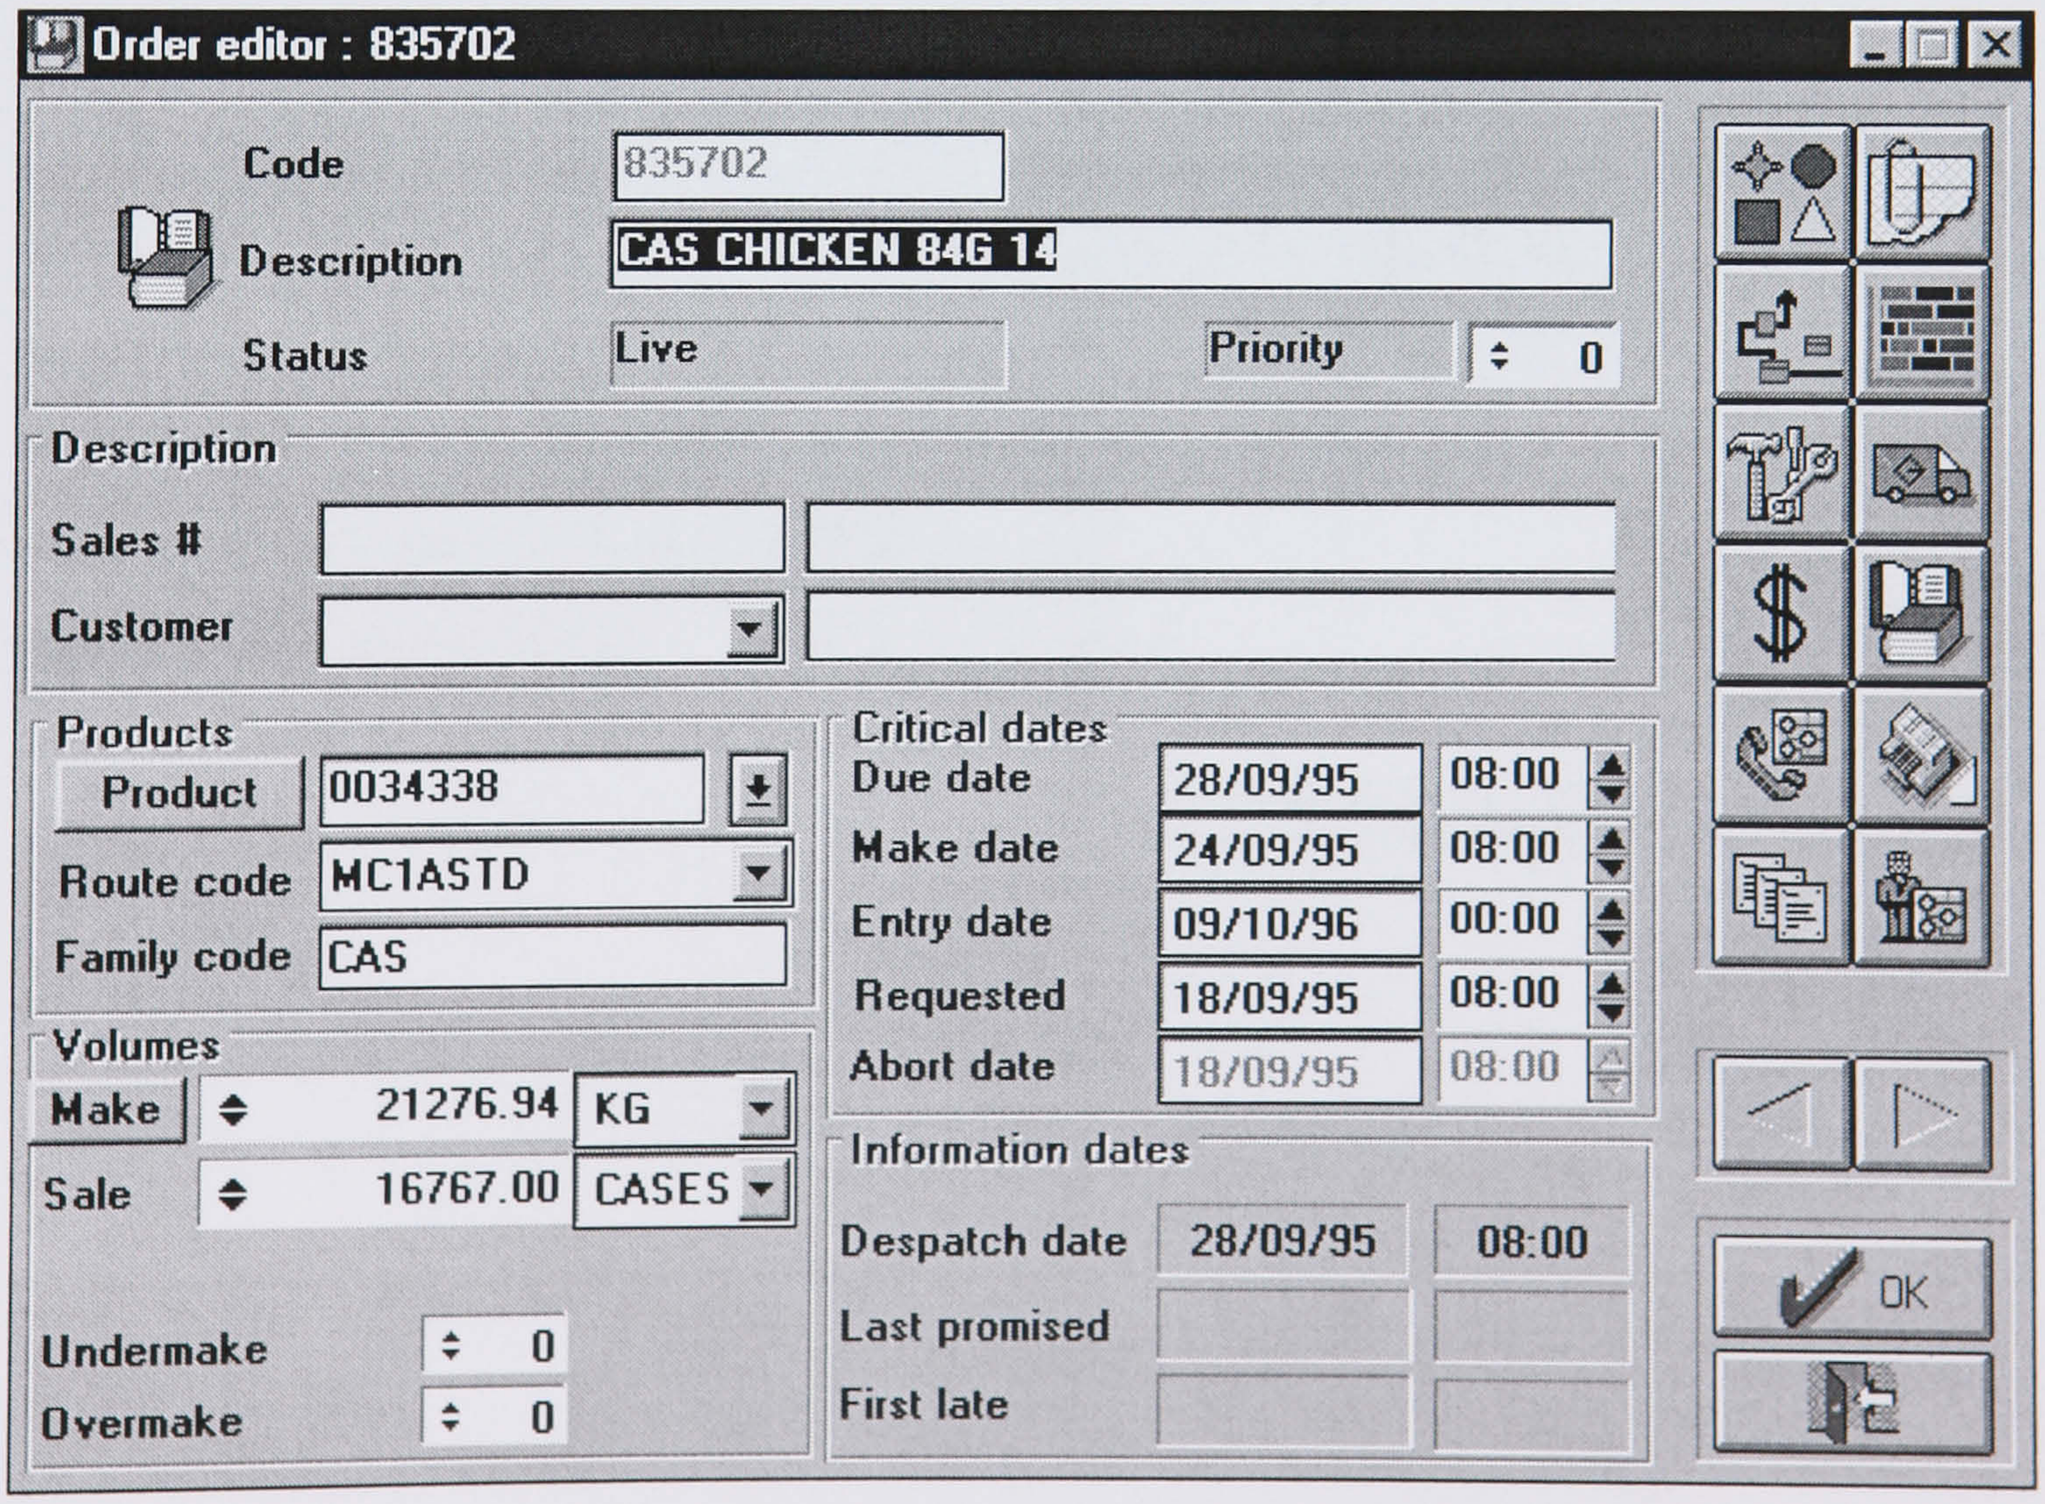
In the snapshot to the right, the various units defined for model are shown. The system is capable of converting between different units as long as the appropriate formulae have been defined which in turn are based on product specifications. In the case below, the number of fills is estimated on the basis of Units/Case and Fills/Unit.



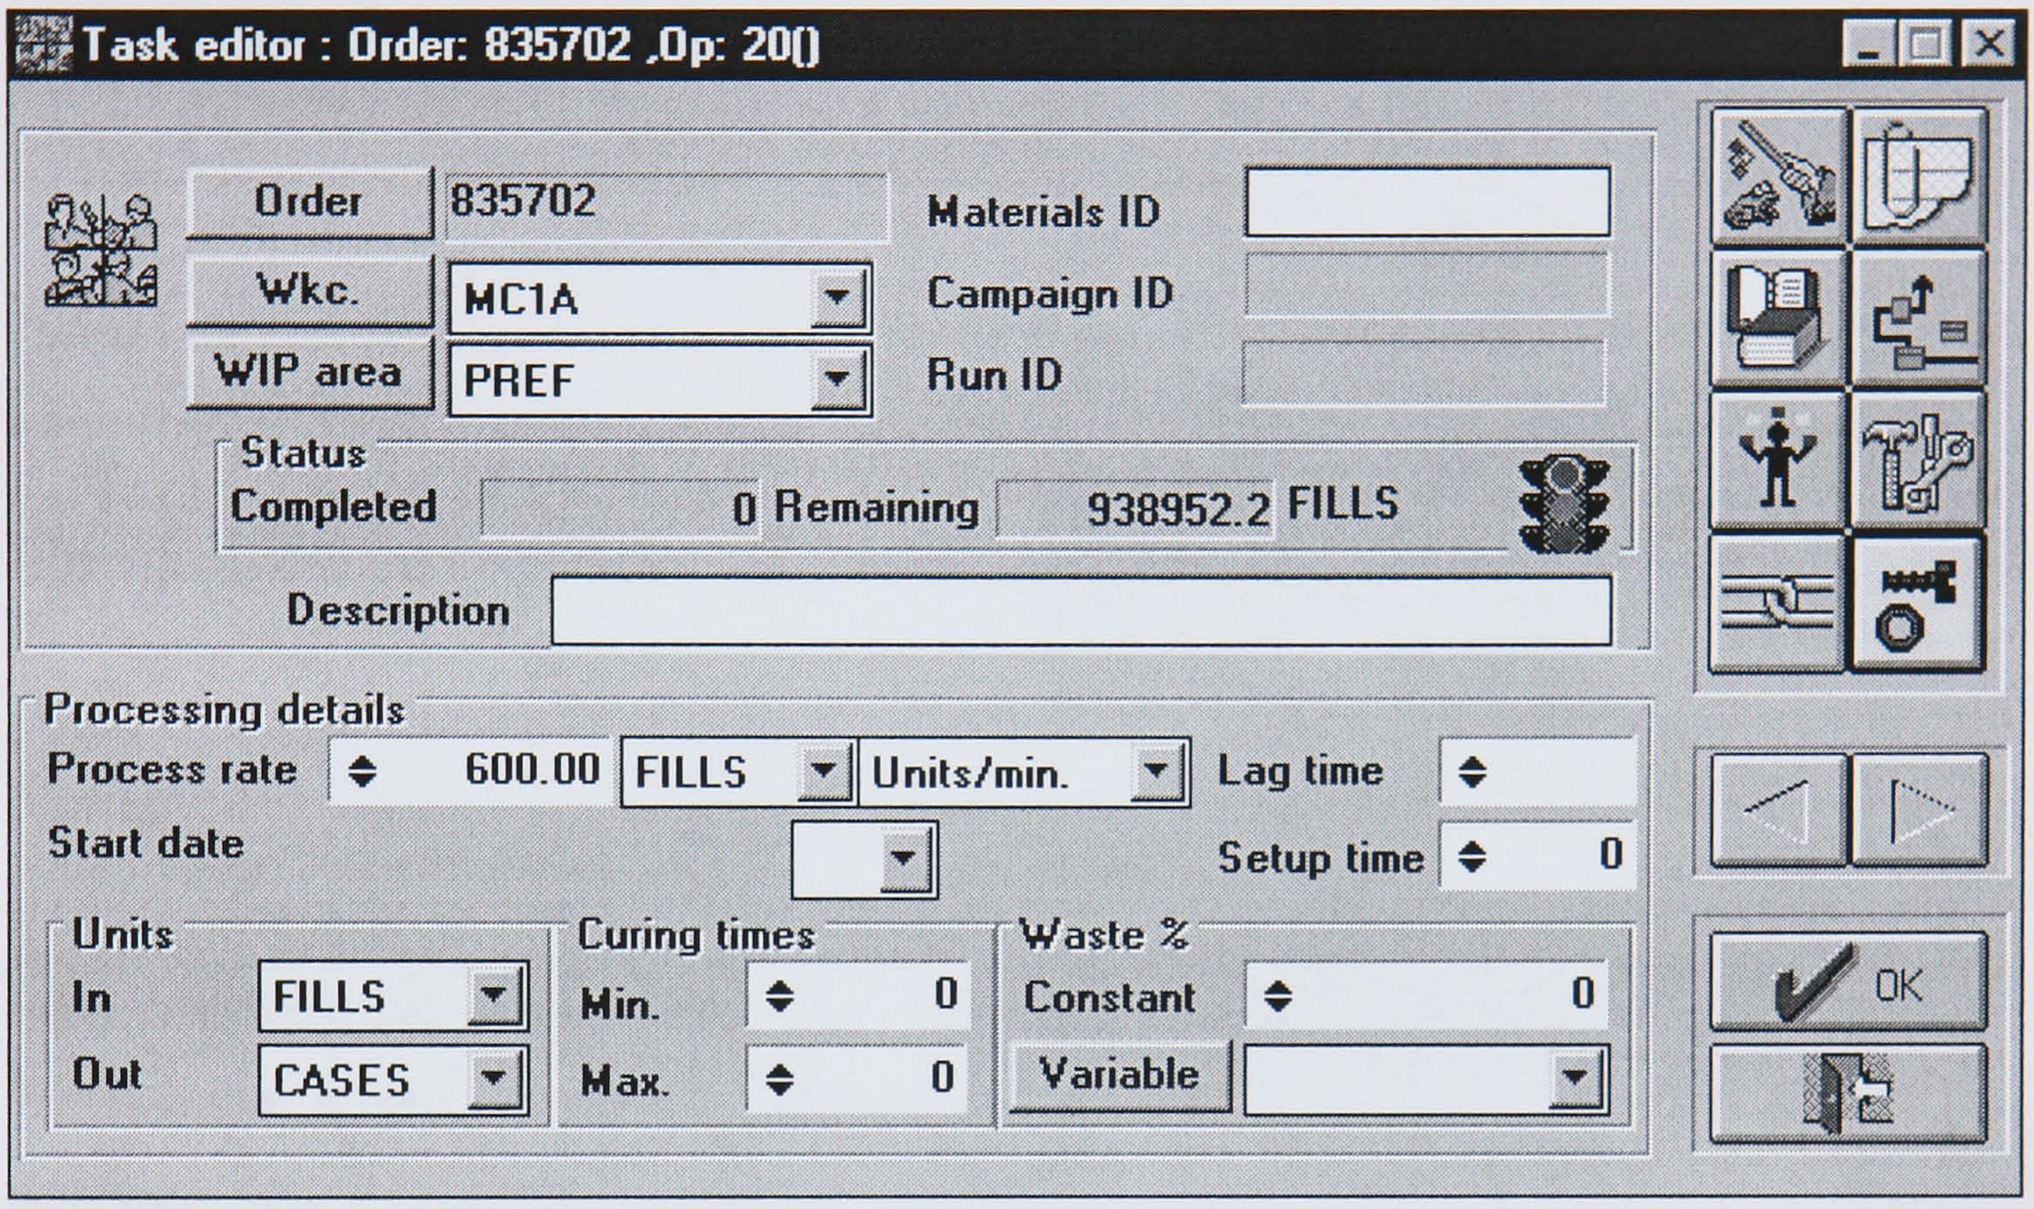


5.5.3 Orders

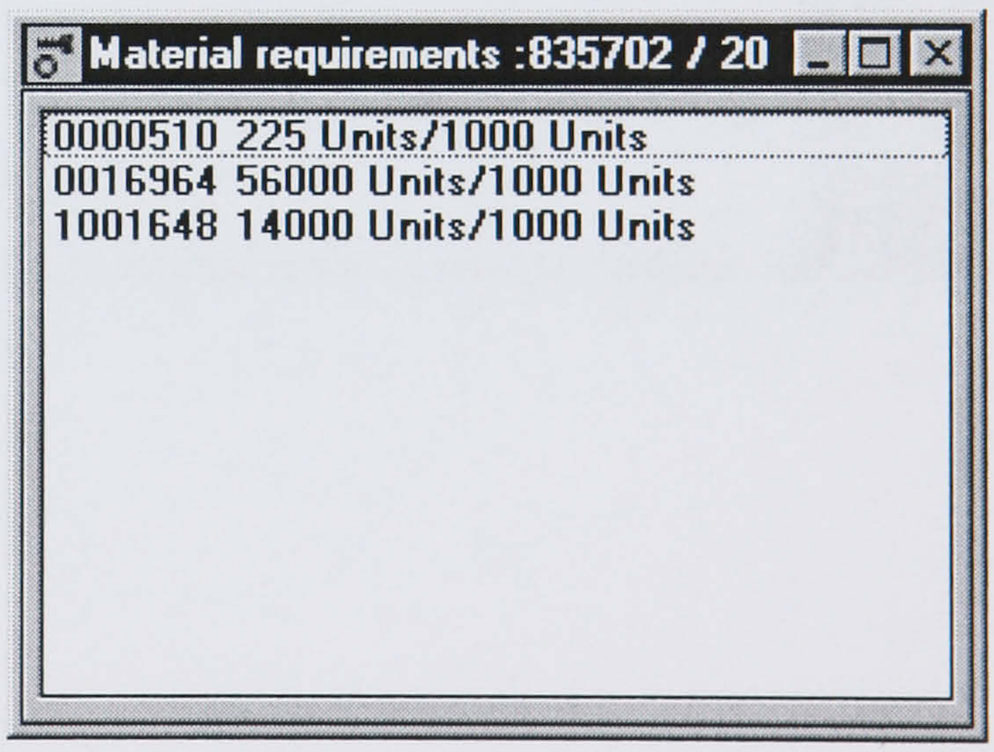
An order is shown in the section below with associated details. Some of the details are attached to the order directly from the product details, e.g. specification, production route, description, material requirements etc. Other details are specific to the order, such as due date, and quantity.



The information below defines one of the operations for producing the order above. This is an extension of the product route, again including extra details for the order. The highlighted materials button (nuts & bolts), is used to specify which processes need to be completed before this operation can commence.

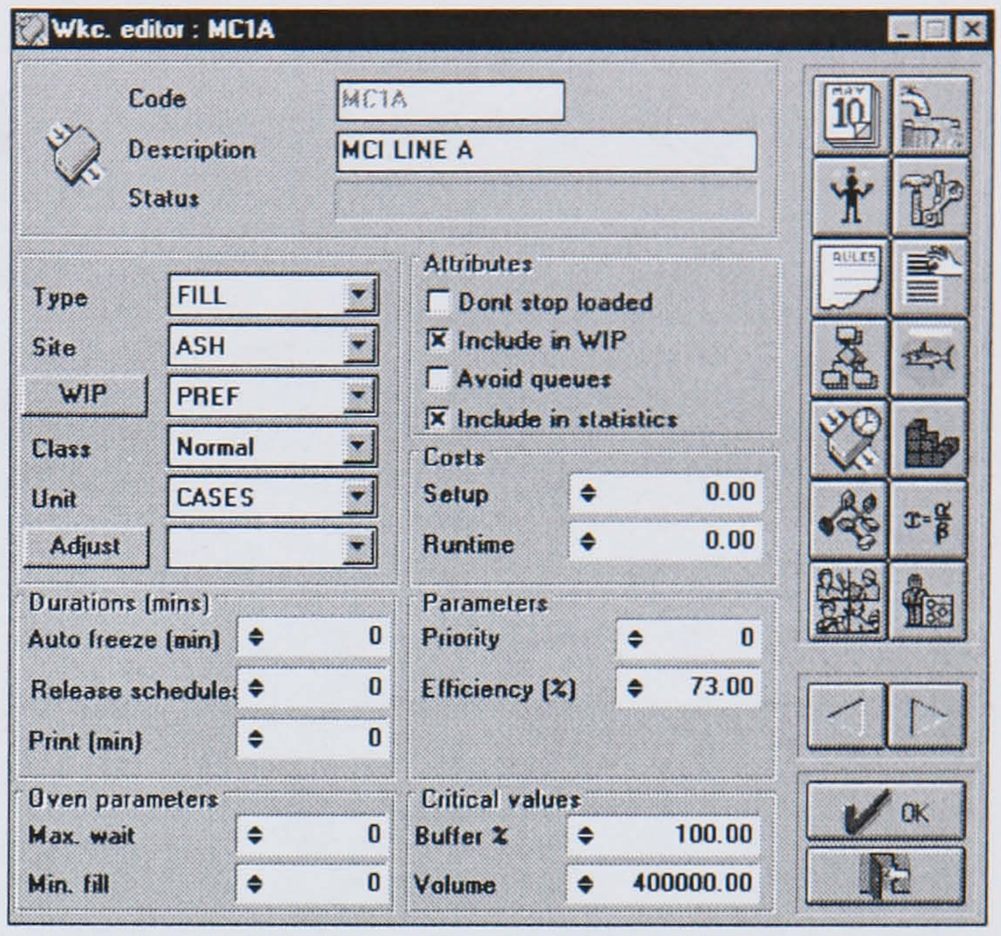


In the window to the right the material codes and their rate of usage are displayed. In this study, for reasons of confidentiality, only a small subset of these materials have been used, mainly to reflect packaging materials.

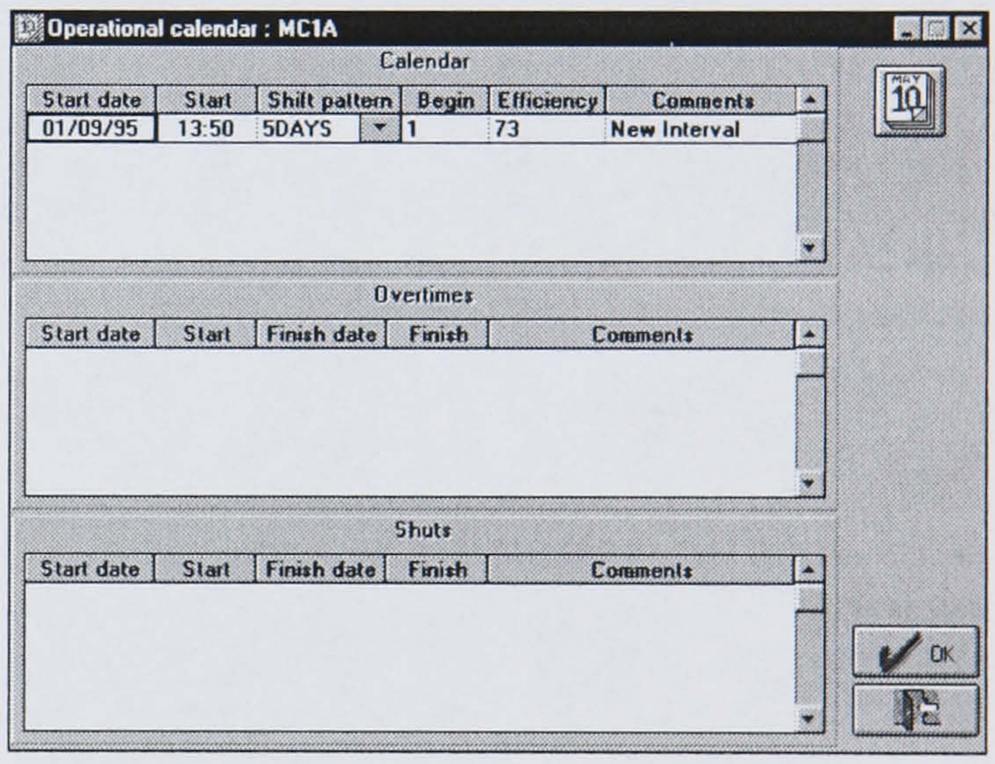


5.5.4 Workcentres

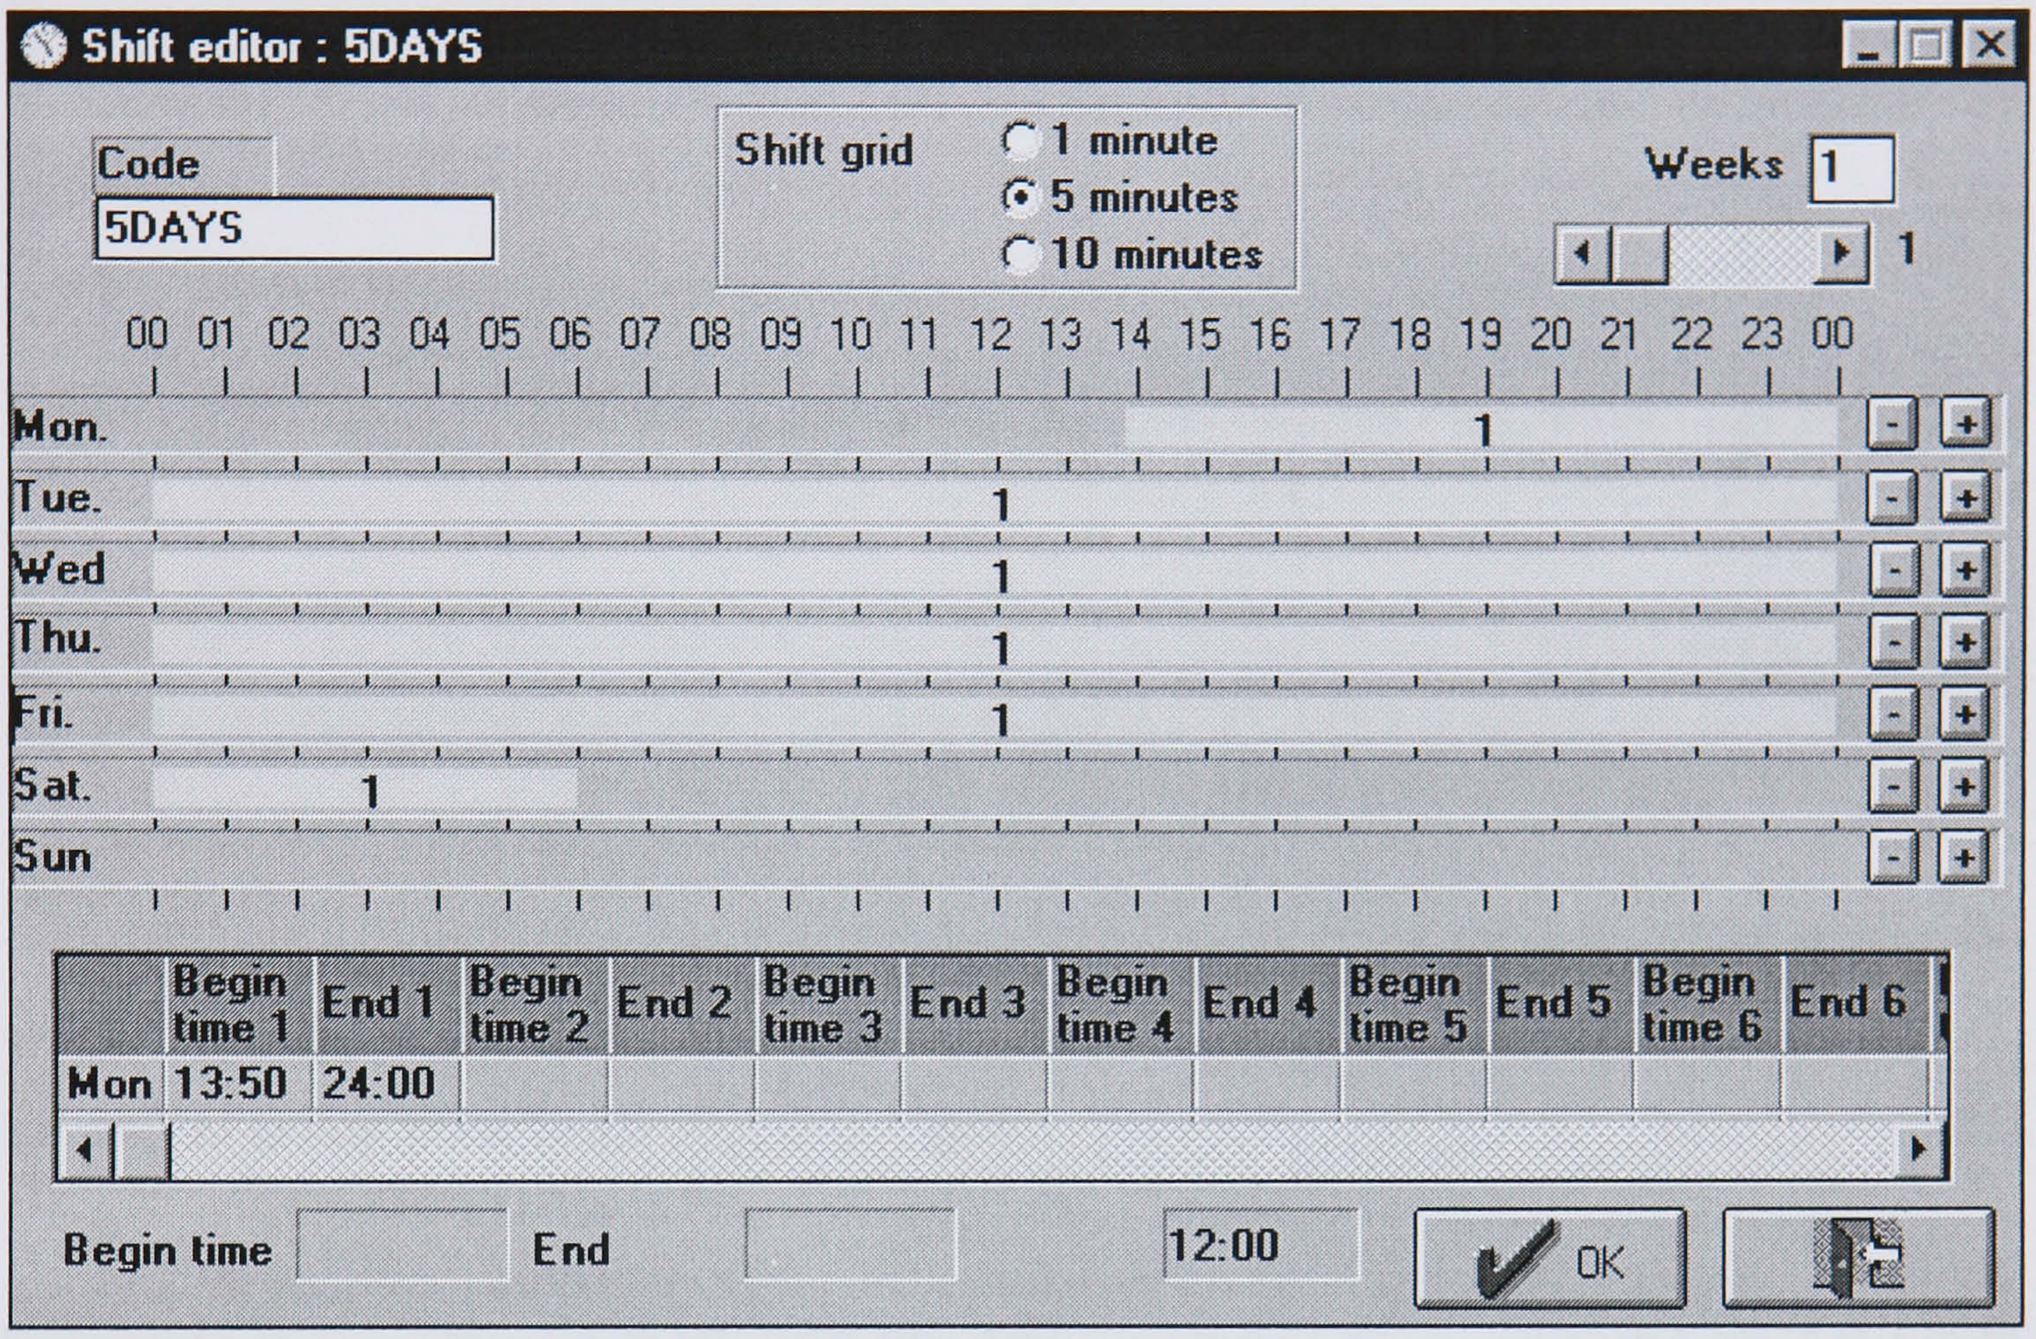
This is a snapshot of the details of one of the Workcentres.



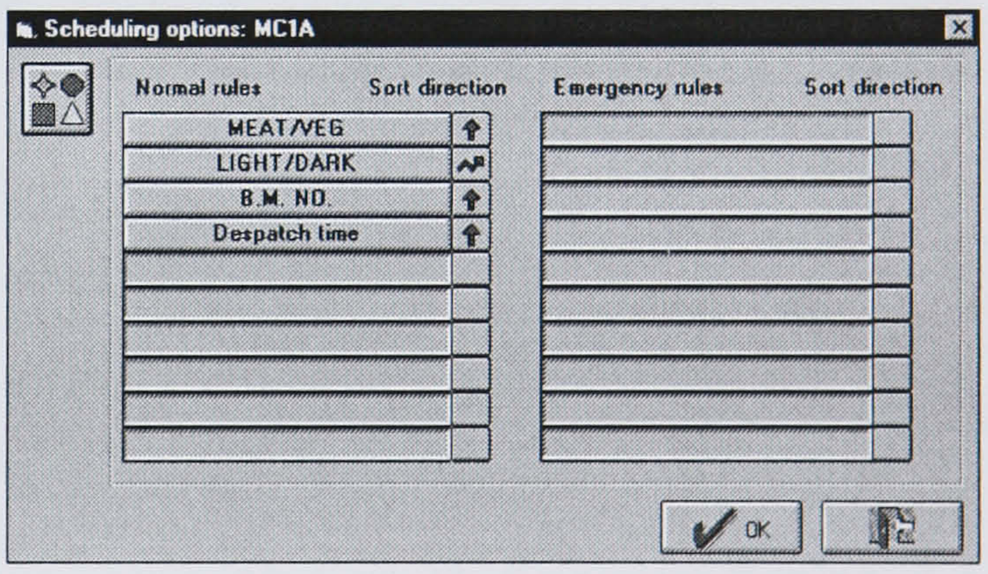
The operational calendar of the Workcentre above is shown here. The operator can introduce different shift patterns to operate at different times, maintenance or breakdown shuts, and periods of overtime.



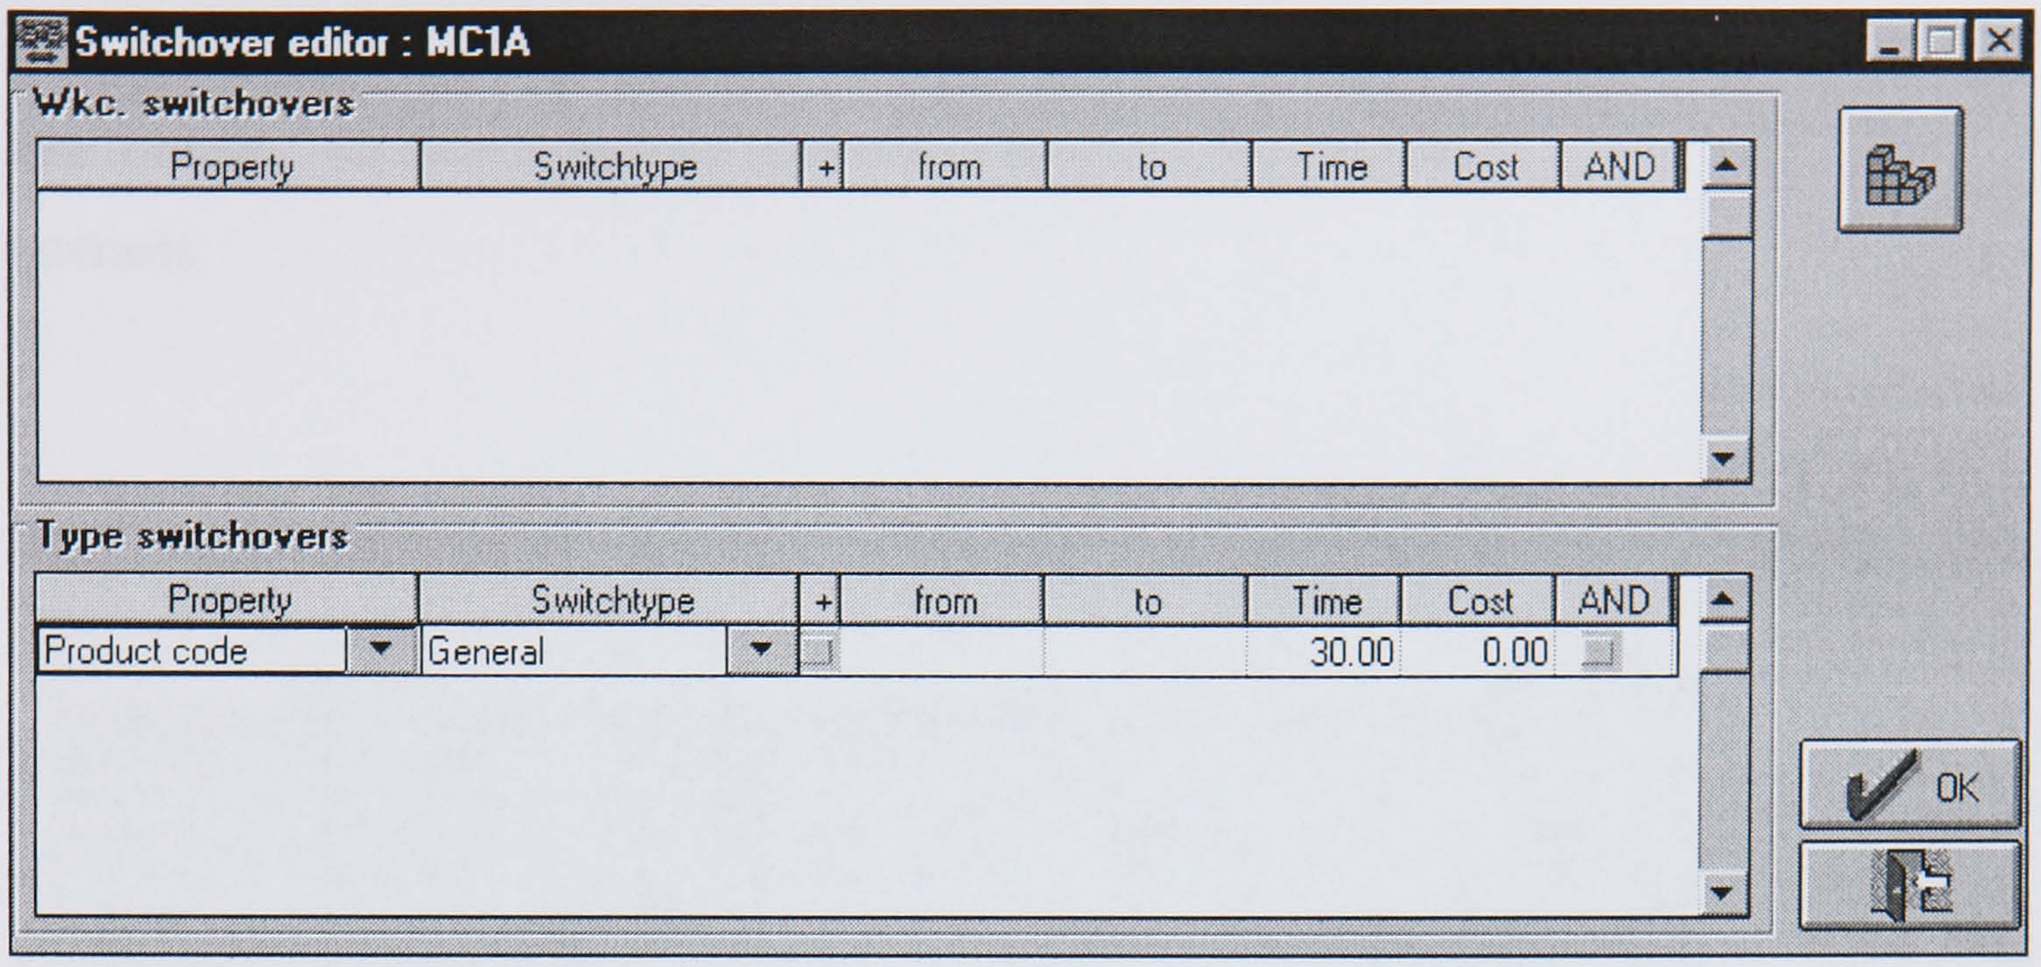
The following snapshot of the shift editor shows the operational availability for a Workcentre. Line A, illustrated below, operates on a '5DAYS' shift pattern.



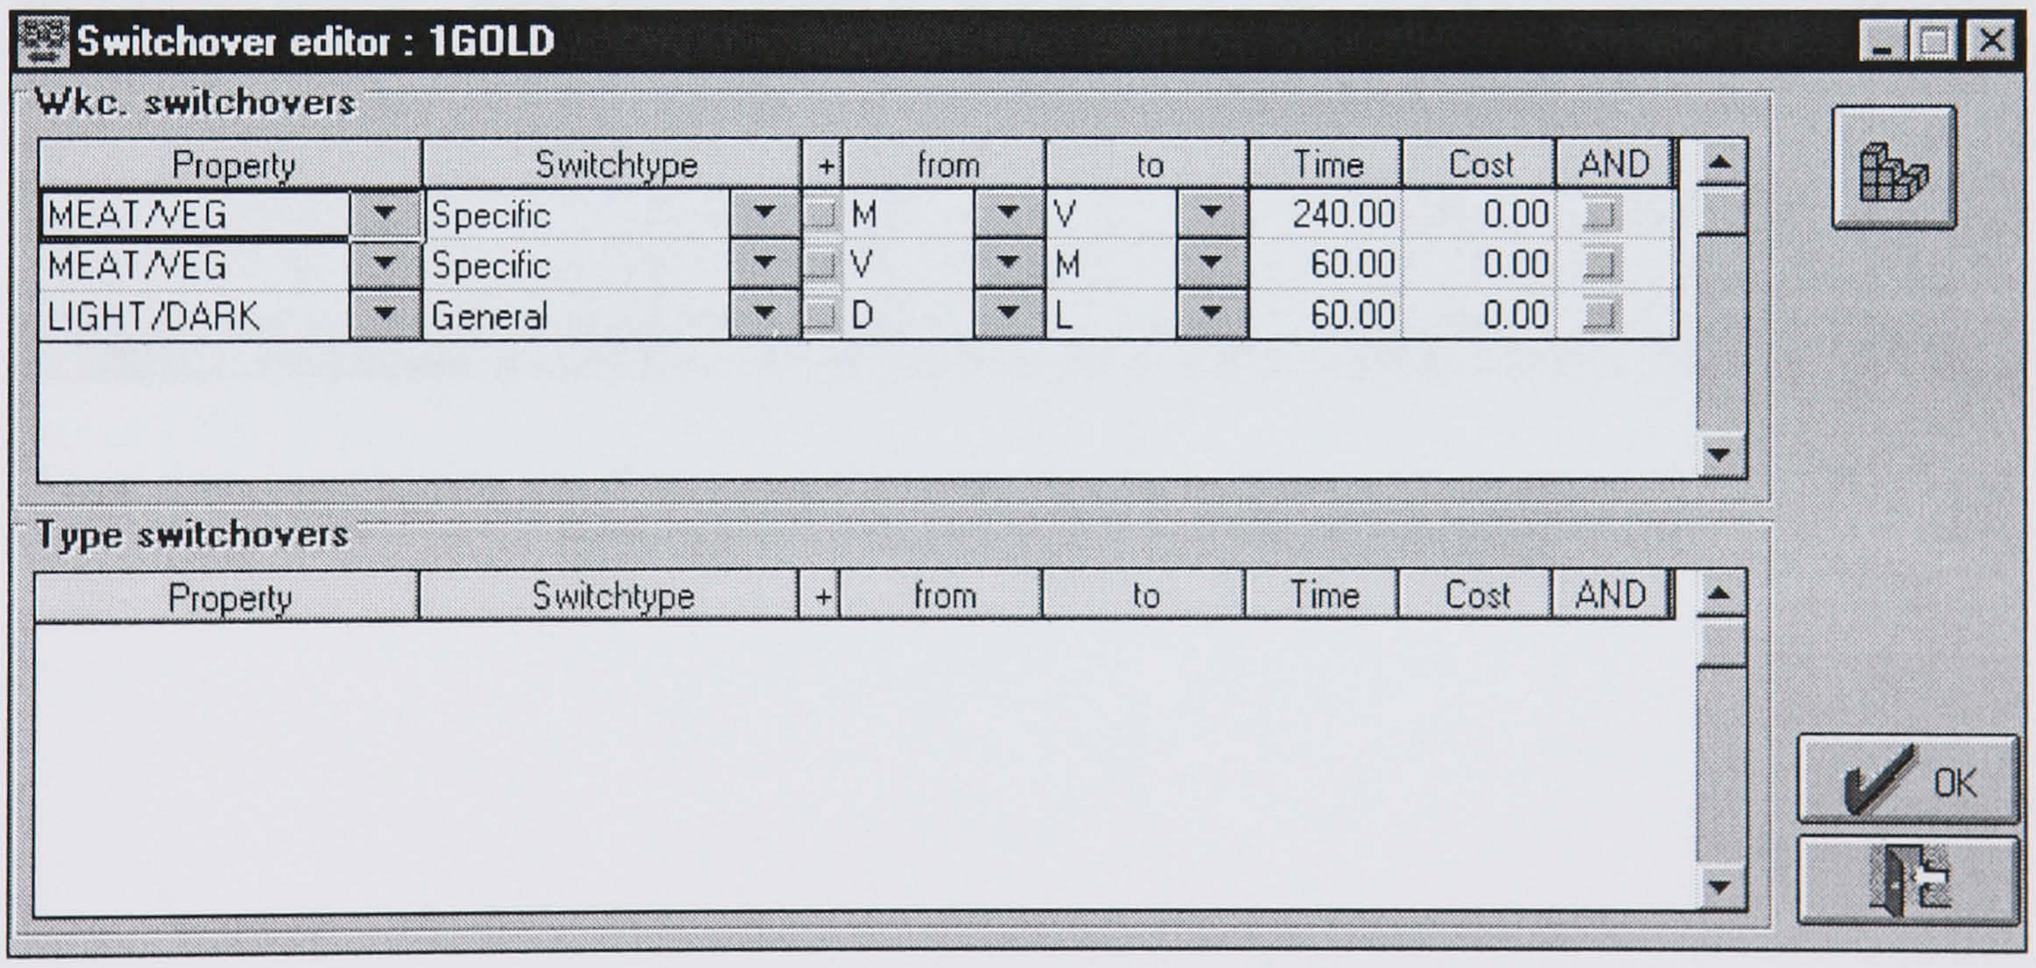
The rules used for sequencing line A are shown on the right hand side. They have been set so that we separate Meat and Vegetable, go from Light to Dark and Dark to Light soups, sequence on the bill of materials (i.e. by like product) and finally on despatch time. The overall objective of this sequence of scheduling rules is biased toward set-up time minimisation



The following snapshot shows the switch over times of line A, which depends on product change.



The rules for the Gold line below, are more complex, and include the set-up time for switching from Meat to Vegetable and vice versa in an asymmetric fashion, as well as going from Dark to Light soups.



5.5.5 Schedules

The listing below shows a typical schedule for line A. Lighter grey text indicate late orders. In this scenario, material unavailability has been responsible for lateness.

S-Plan 4 Reports

Code	Operation number	Product code	Description	Due date	Finish run date	Finish run time
946796	020	0034326	CAS BEEF & TOMATO 84G 14	22/09/95	19/09/95	18:31
006307	020	0035562	CAS CHICKEN & LEEK 78G 14	20/09/95	20/09/95	03:33
969798	020	1001953	CAS TOMATO WIP	19/09/95	20/09/95	04:29
015650	020	1008702	SAS GRAN CHK & LEEK 42G 14	23/09/95	23/09/95	00:57
003746	020	0036164	MCD CAS CHICK (621) 14XPKT	28/09/95	25/09/95	17:00
835702	020	0034338	CAS CHICKEN 84G 14	28/09/95	27/09/95	05:14
029054	020	0034326	CAS BEEF & TOMATO 84G 14	30/09/95	27/09/95	23:15
982176	020	1001953	CAS TOMATO WIP	26/09/95	28/09/95	01:29
033583	020	0034347	CAS GOLDEN VEG. 80G 14	22/09/95	04/10/95	01:58
029399	020	0034349	CAS TOMATO 92G 14	22/09/95	05/10/95	12:02

10

MC1A

Route

1PREP
MC1A

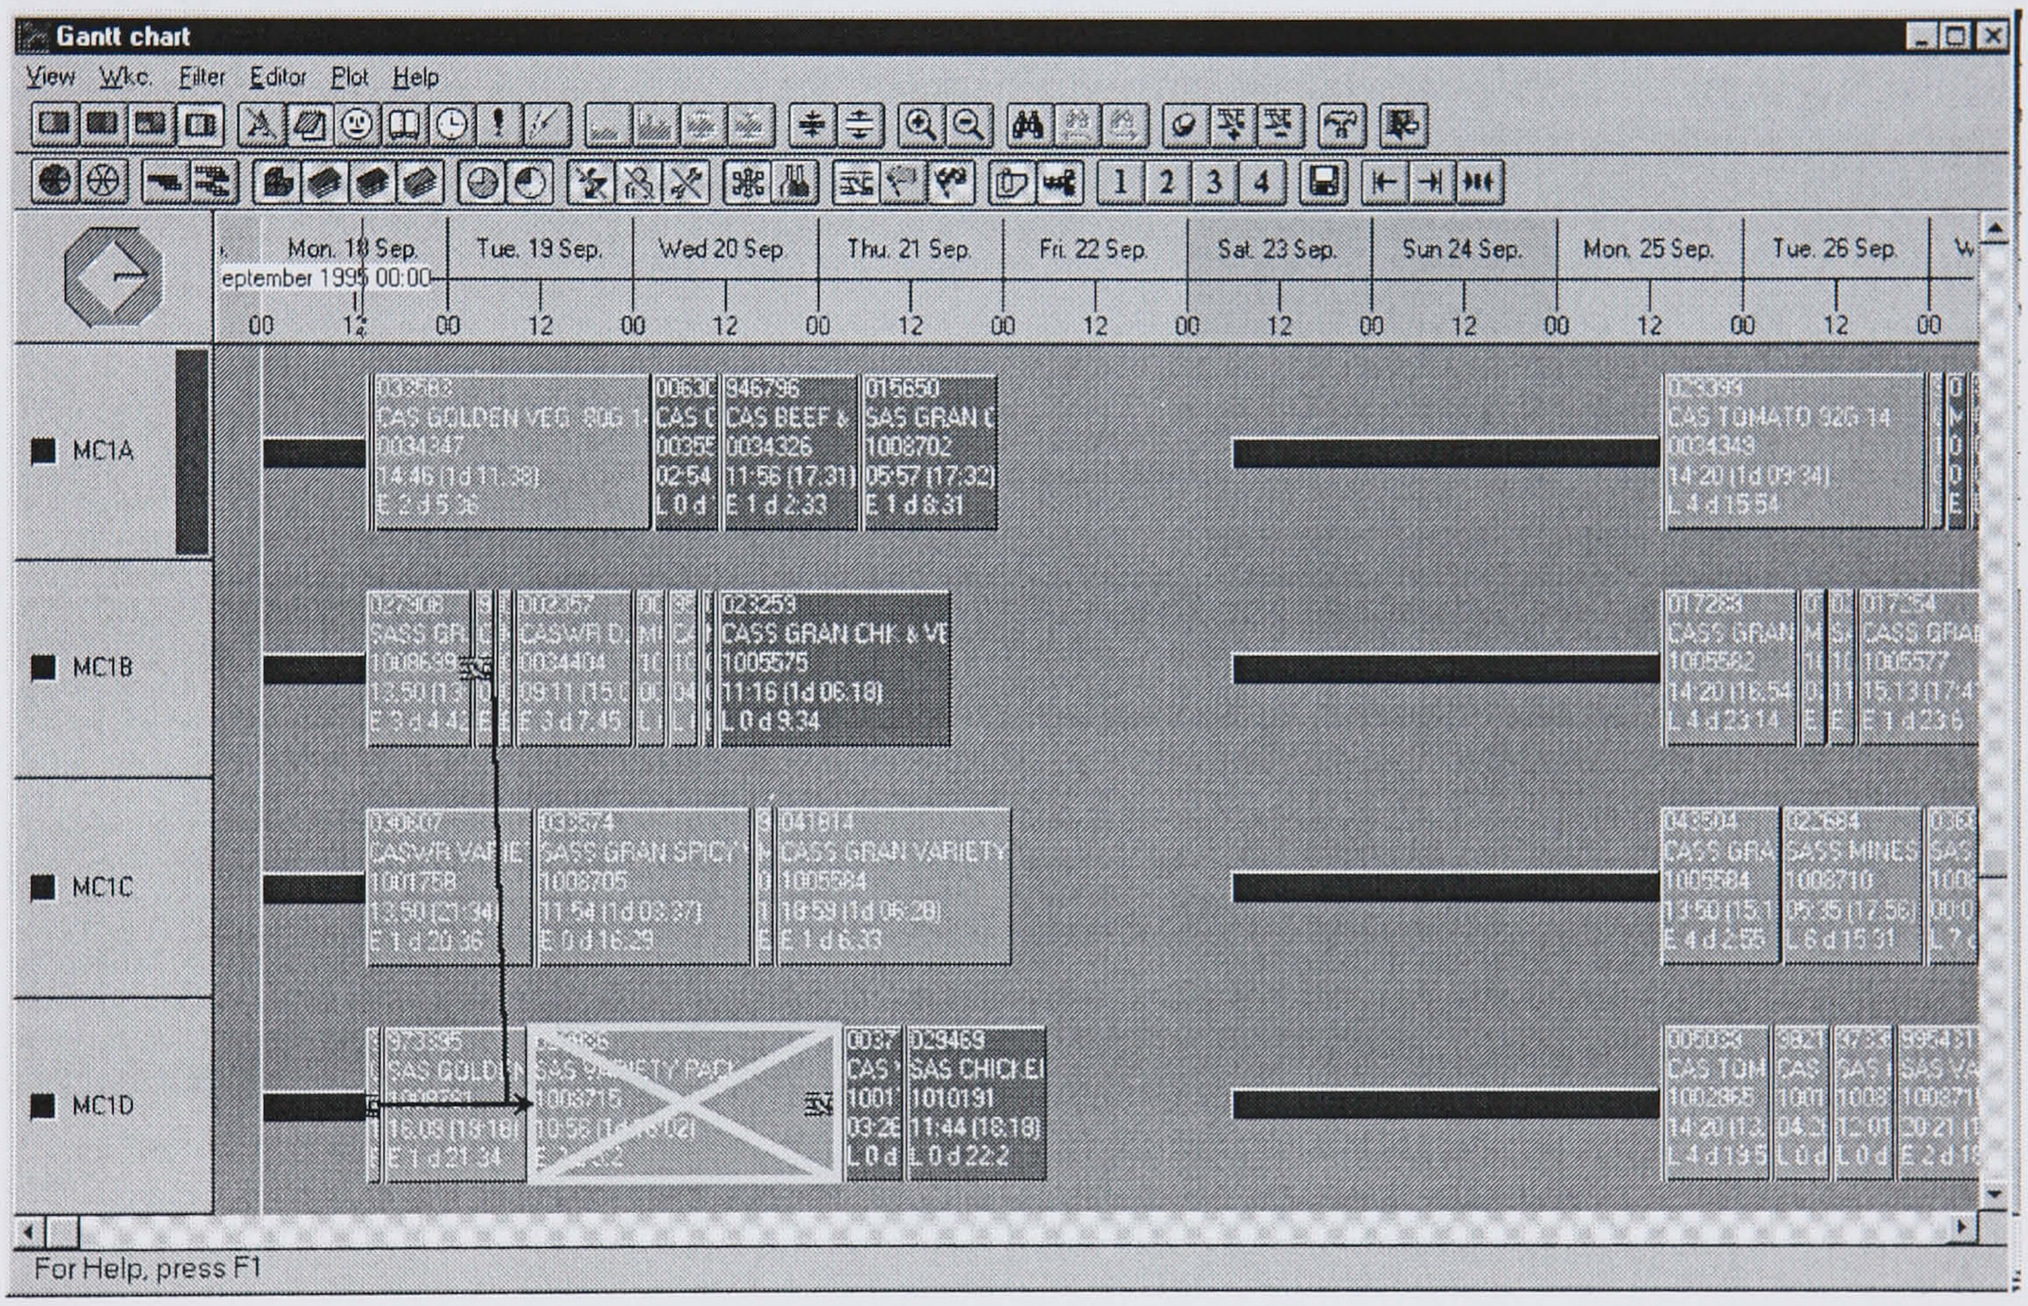
Show interrupts

Synchronise scroll

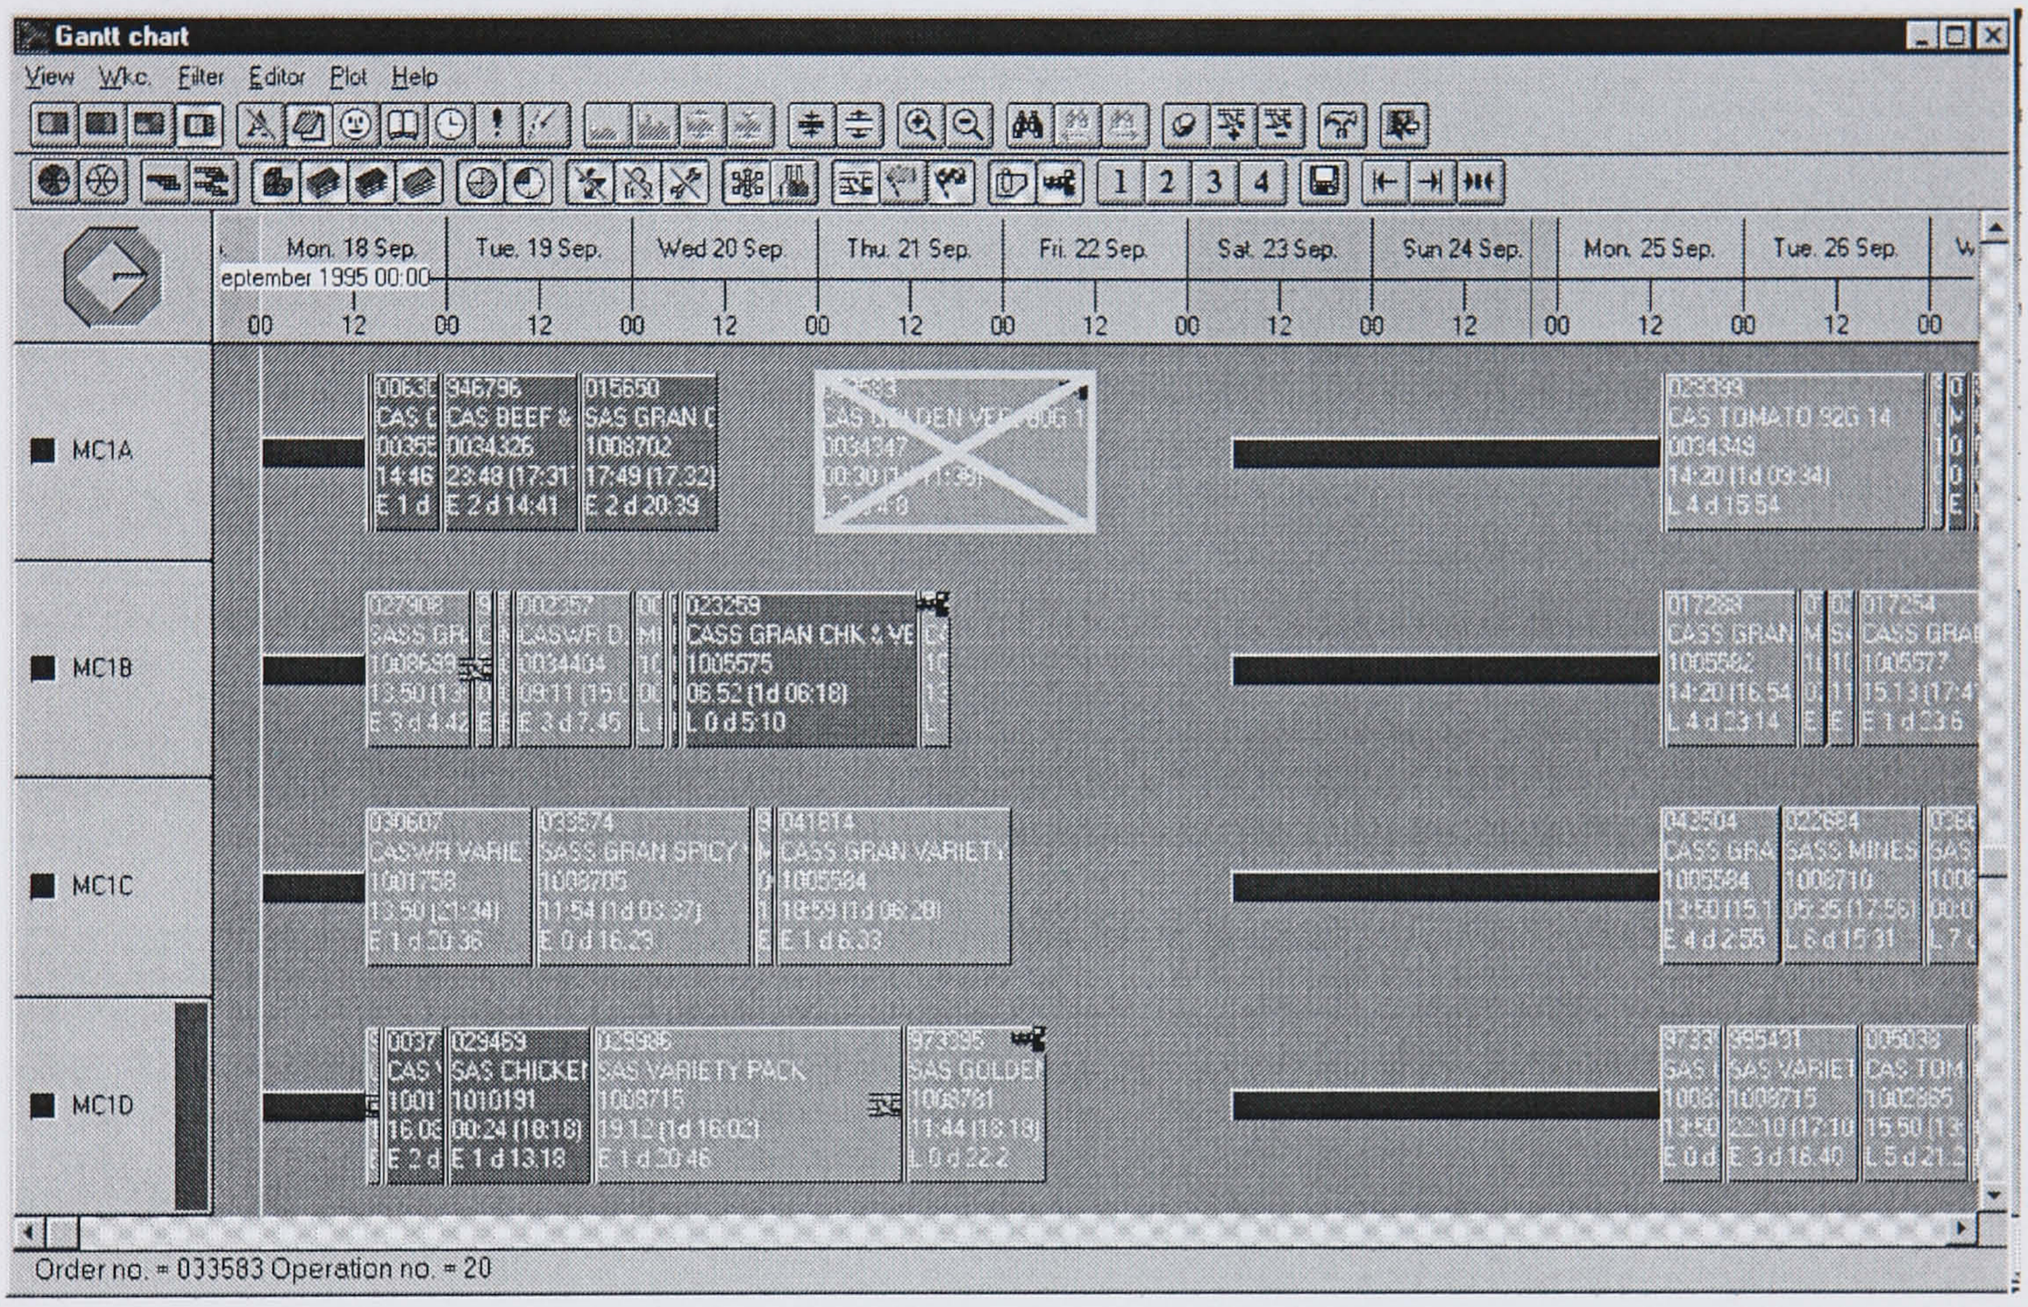
Version 1.12.1F

5.5.6 Reactive Scheduling

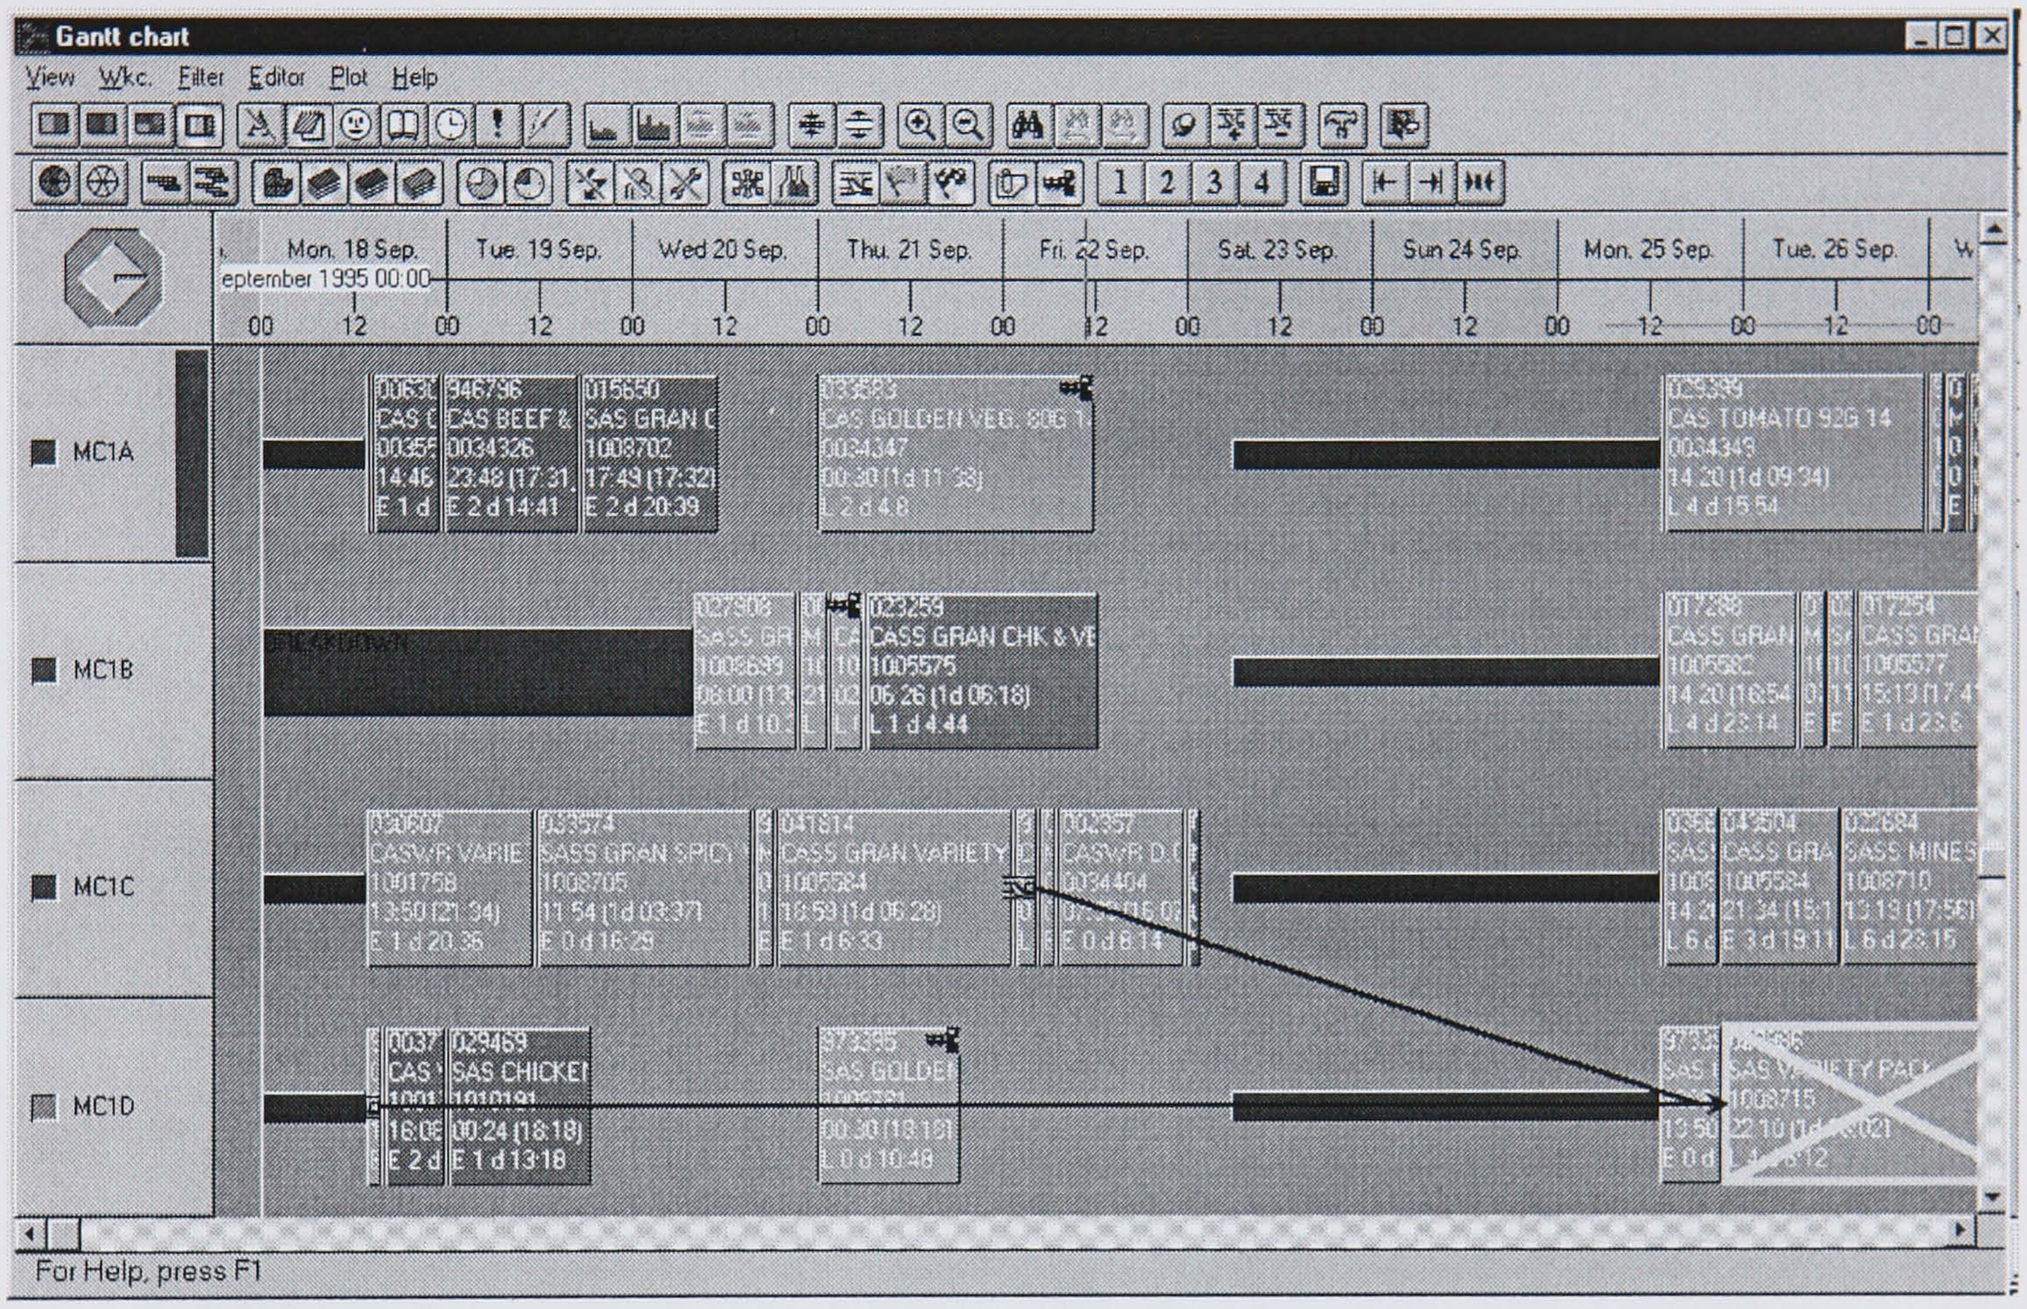
The following display shows the production at Ashford with no Material shortages in evidence. In this snapshot of the Gantt chart the darker colour pieces indicate meat, the lighter indicate vegetable soups. Also, note that the crossed piece on line D represents a variety SKU order, waiting for fill components to arrive from line B and line D (indicated by lines with arrows) before processing the final pack operation.



In the next snapshot we have introduced some material shortages to indicate the effect on the line production schedule. For instance, the crossed order cannot start processing because of shortage on packaging material '0016976' and therefore it has to wait for the material to be delivered towards the end of the week. This order becomes late. Similarly, the other orders which have been visually marked with a screw indicate the same problem.

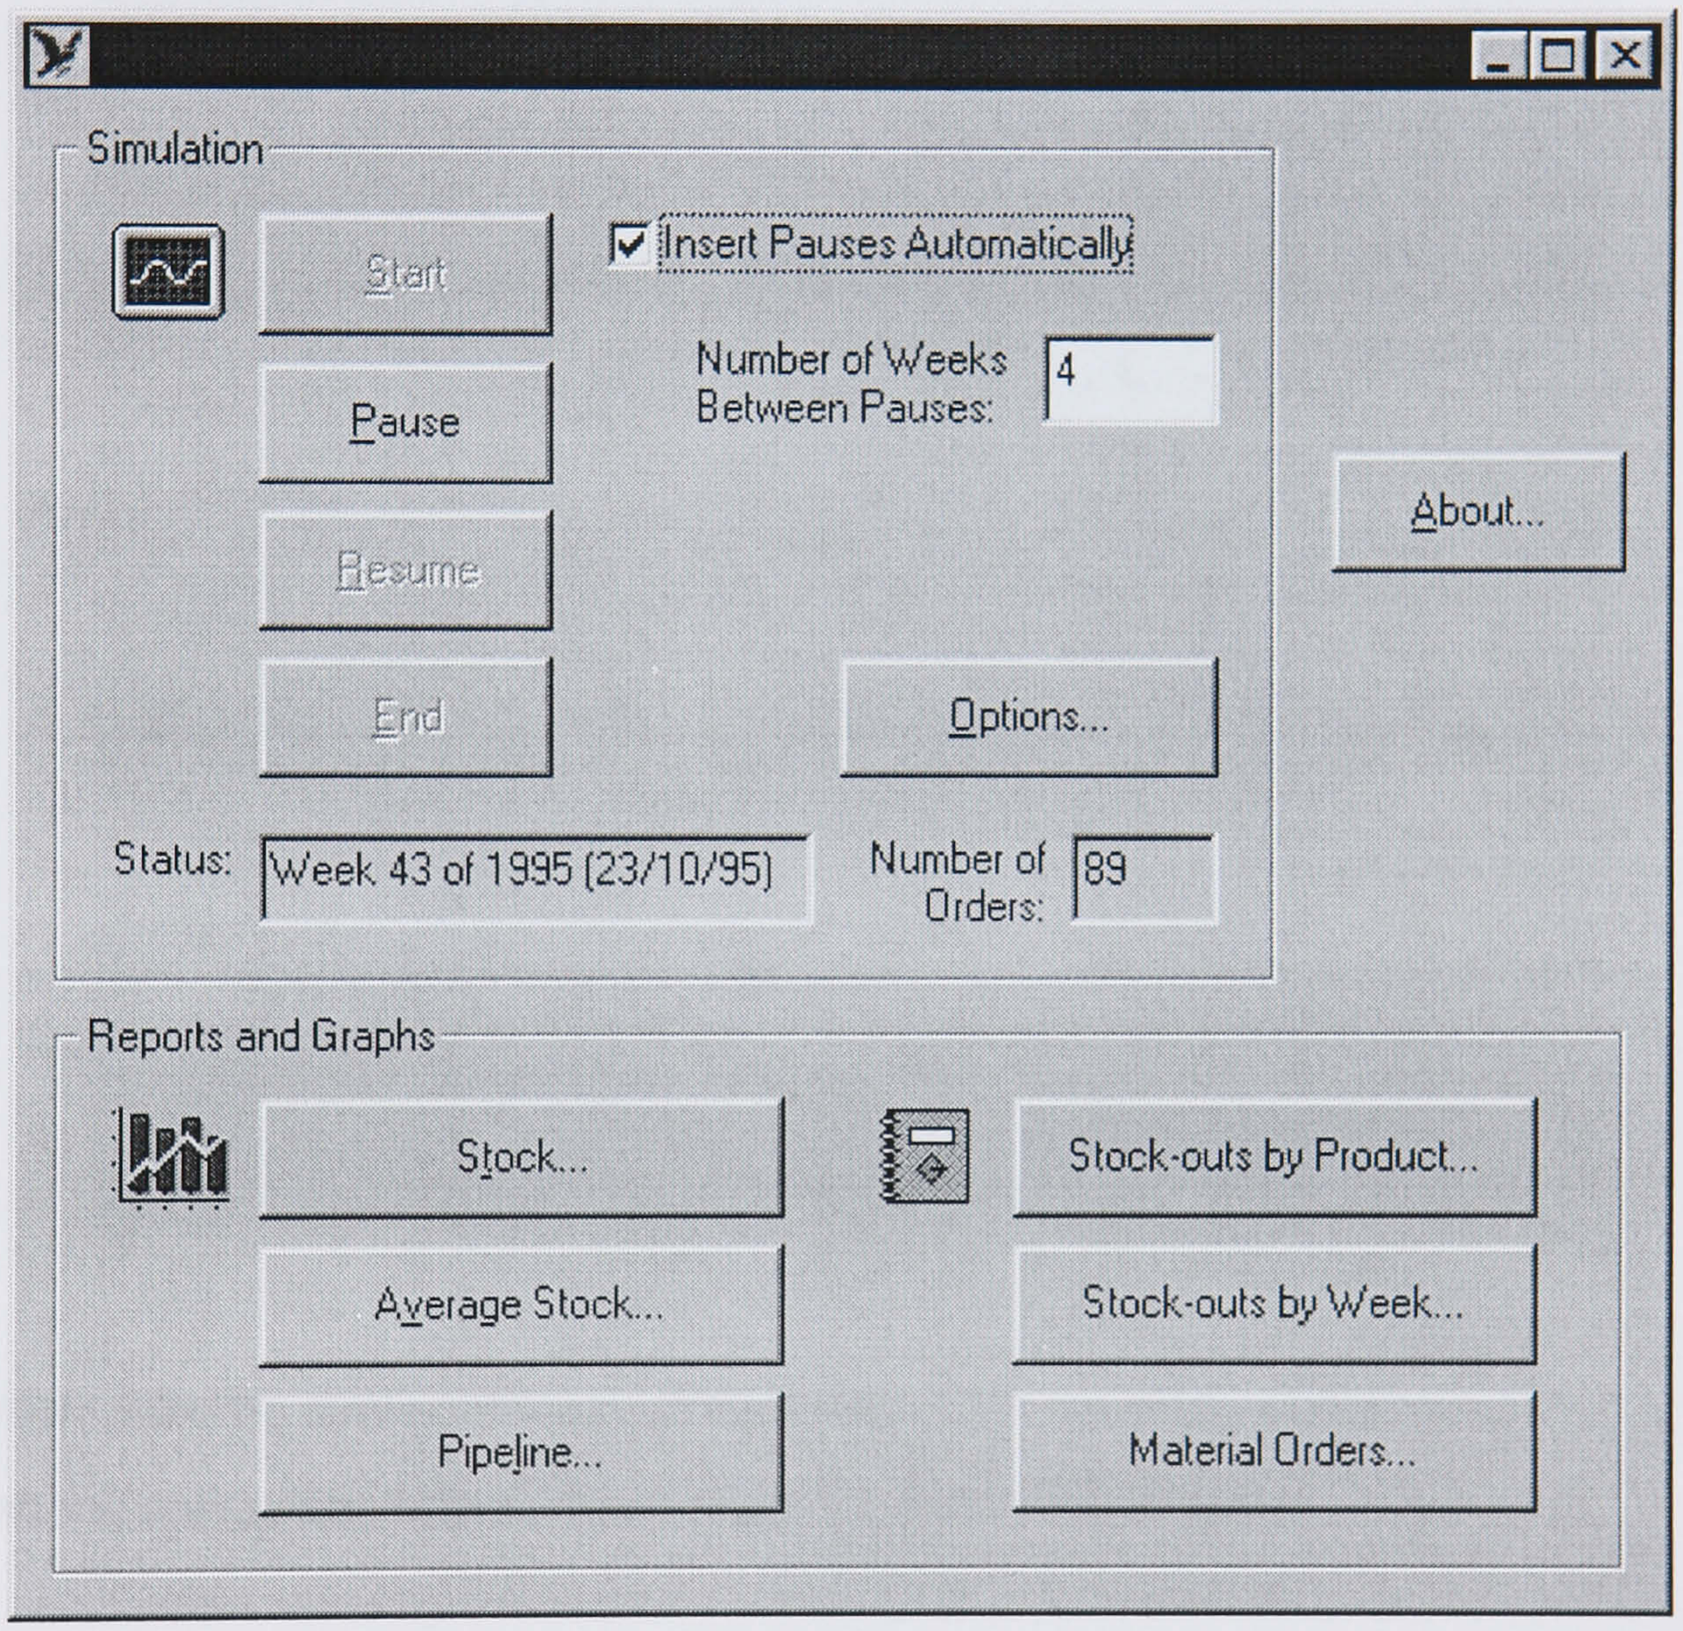


In the following snapshot the introduction of downtime causes production problems on line B. Some of the orders originally going through line B have been automatically re-routed over to line C. In particular, the variety pack order can no longer start before the weekend, and therefore the assembly processing starts the week after.

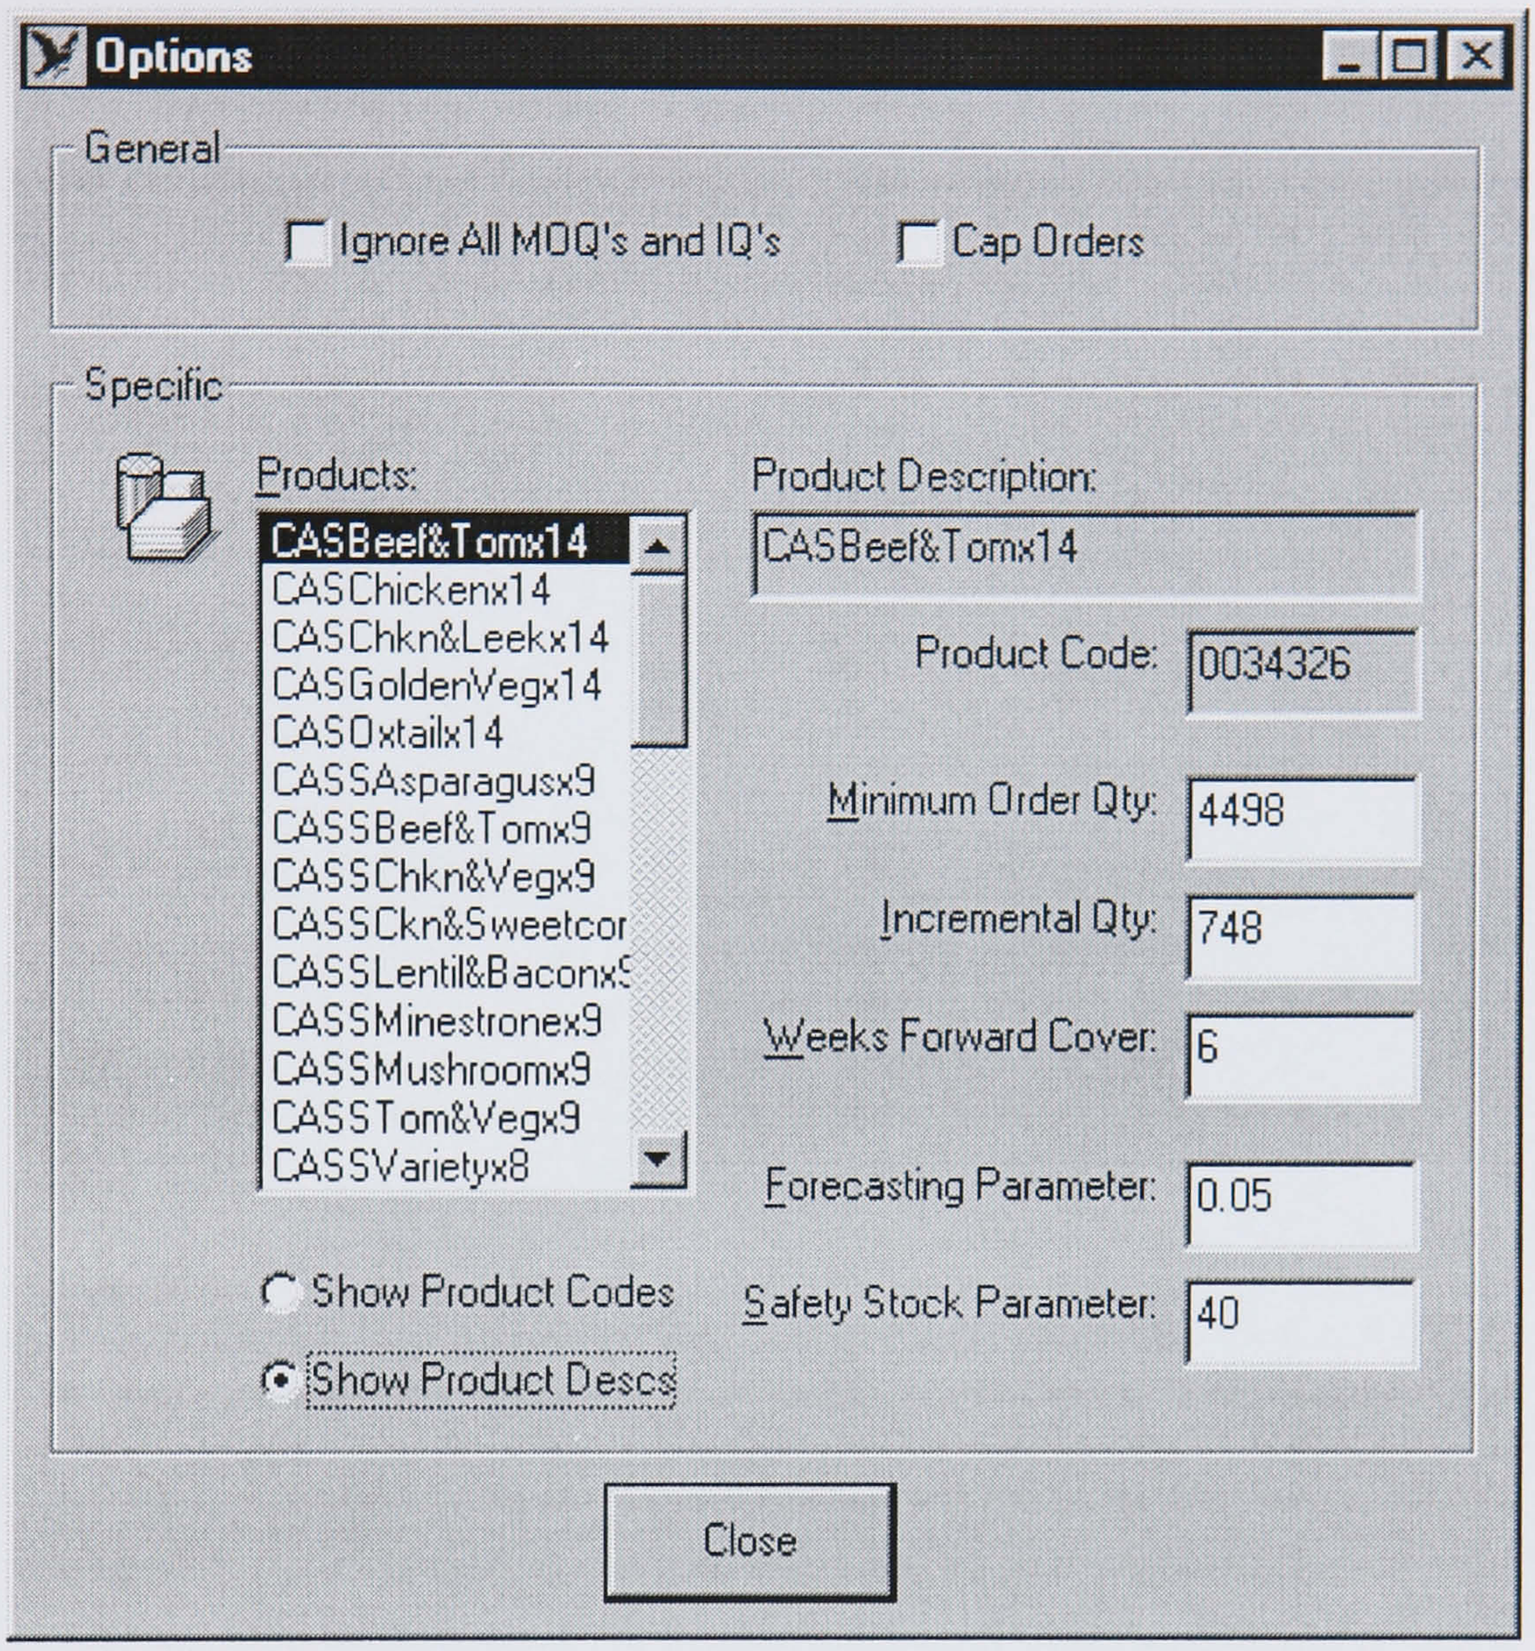


5.5.7 Synchronous Replenishment System

The front screen of the SRS is shown below. The Simulation box enables the user to specify the number of weeks the simulation is to run before pausing. Alternatively, the simulation can be interrupted manually by the user using the Pause button. The Status box indicates how far the simulation has progressed and the Number of Orders box displays the number of orders scheduled so far.



The Options screen is shown below. This enables the user to set parameters prior to simulation. The user can scroll through the database of products and view/edit parameters accordingly.



Products can be displayed by Description or by Code. There is also a global option to remove MOQ and IQ constraints from all products.

Once the simulation is complete, there are several reports and graphs displaying key results. These are accessed through the front screen. The Stock graph shows how a product's stock varies over the period of simulation. Also displayed are the demand, actual stock level experienced and all of the parameters associated with the simulation. Fig. 8 illustrates examples of this graph. The Average Stock graph shows the simulated average stock against actual average stock for all SKUs within a soup family (e.g. Cup-a-Soup) as illustrated in Fig. 3.

The Stock Out reports show any stock outs occurring over the period of simulation. These can be listed by week or by product. The Material Orders report gives example orders of raw materials required from the Supplier (this list is not exhaustive due to manufacturer confidentiality over soup recipes).

5.6 Parameter Tuning – Results and Discussion

The MDSS system has many parameters which affect the overall performance of the supply chain. The effects of many of these on key performance measures are frequently non-linear and hence the use of a model based approach.

In this section, we report on the approaches and results obtained using the MDSS system to tune seven different classes of parameters. The total number of parameters involved in these examples exceeds 200. Those selected, for discussion are as follows:

- Weeks Forward Cover (wfc). This plays a key role in the performance of the SRS System and in the minimisation of average stock levels. As explained in the previous chapter, it may be controlled by adjusting the time window (T) parameter of the individual fixed cycle system.
- Minimum Run Length (MRL). This parameter affects stock levels and manufacturing scheduling. It represents a batch constraint within the warehouse fixed cycle replenishment system.
- Safety Stock (SS) and Forecast Accuracy. The tuning and optimisation of safety stock parameters are key to the minimisation of stock level and attainment of customer service. For the procedure used here, these correspond to the dispersion parameter of equation 5.0 of chapter 4.0.

The tuning of each of these parameter classes is described in a separate section below. Each section has a similar layout. The performance measures affected by the parameter are first listed and briefly discussed. The tuning procedure is then discussed and the results documented. Finally, the results are assessed and discussed. Assessment involves quantification of the results in terms of financial and non-financial performance measures, and a discussion of the methodology.

5.6.1 Weeks Forward Cover (wfc)

Cost and performance measures. These parameters potentially influenced the following costs and performance factors:

- Average stock
- Warehouse overheads
- Manpower and overtime costs
- Opportunity costs associated with warehousing.

Tuning procedures - The SRS system seeks to maintain the stock levels for all SKUs within the warehouse at prescribed levels, described in terms of weeks forward cover, calculated from the current mean level demands. Thus the stock level in absolute terms is varied according to seasonal variations in mean demand. In the case of instant soups, this results in higher stock levels in absolute terms, during winter periods.

As described earlier, the fixed cycle inventory control system calculates the replenishment needed to bring every SKU up to the weeks forward cover (wfc) level specified within the system, taking into account the normal safety stock component and according to equation 4.0 of chapter 4, the safety stock at each wfc level is calculated using the parameter tuned safety stock rules described below.

The replenishment quantity calculated for any particular SKU must be rounded up if it is less than the prescribed minimum run length. The minimum run length parameters used here to obtained these results are the "as is" values used by management in the current system. As explained above, these are primarily introduced to avoid unexpectedly short run lengths and high switch over rates being included in the manufacturing schedules in addition to raw material lot size constraints. The tuning of the MRL parameter is discussed in the separate section below.

The procedure used to tune the wfc was as follows. Firstly, the wfc is set to "as is" figures of 8 weeks and the models are run. Comparisons are made between the stock levels computed using the SRS system and those obtained from the data logging exercise undertaken during the test period. This first test checks that the model based stock levels are in reasonable agreement with those recorded in practice.

The wfc levels were then progressively reduced uniformly for all SKUs. The SRS response to the recorded warehouse demand data over the 6 months test period was then used to compute key measures such as average stock level, switch overs, number of stock outs, overtime etc. Where appropriate, the results were investigated and explained from a practical and theoretical stand point.

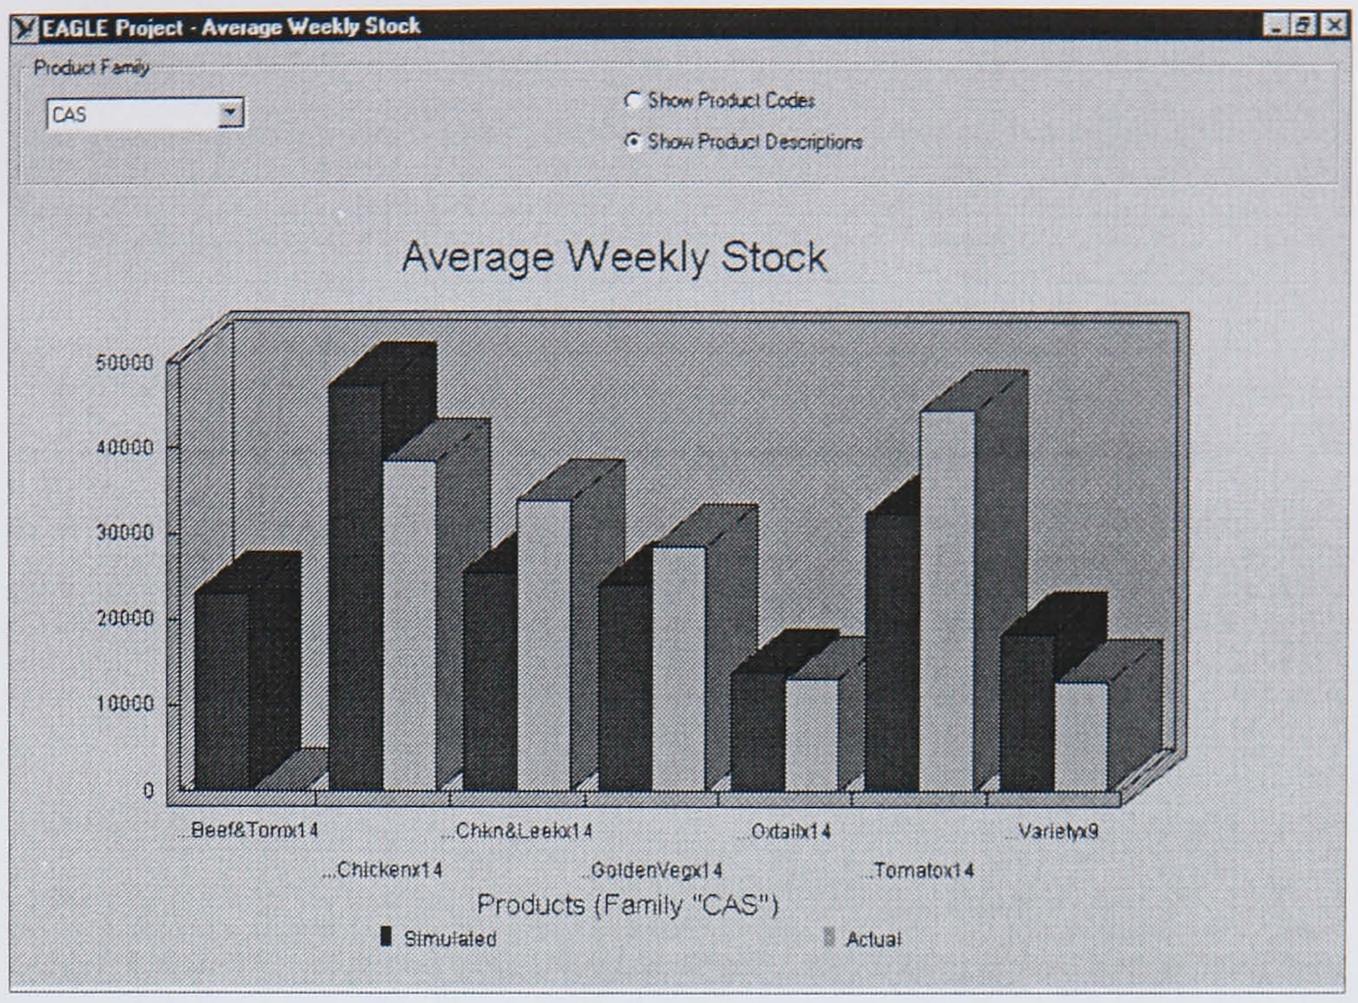
8 wfc result. The average stock results are displayed in Fig. 3. The displays are for lines selected in the high, medium and low demand range and are representative of the 50 or so products. We note that there are no "as is" results for beef and tomato in the CAS range, due to a fault in the data logging exercise. Otherwise all of the products illustrate the comparison between the actual and SRS response. The display shows the results obtained with the standard minimum run length constraint. The performance of the SRS system with the 8 wfc setting is similar but uniformly lower than the "as is" stock level recorded for the corresponding period. The stock levels also give reasonable agreement with equation 8.0 of Chapter 4.

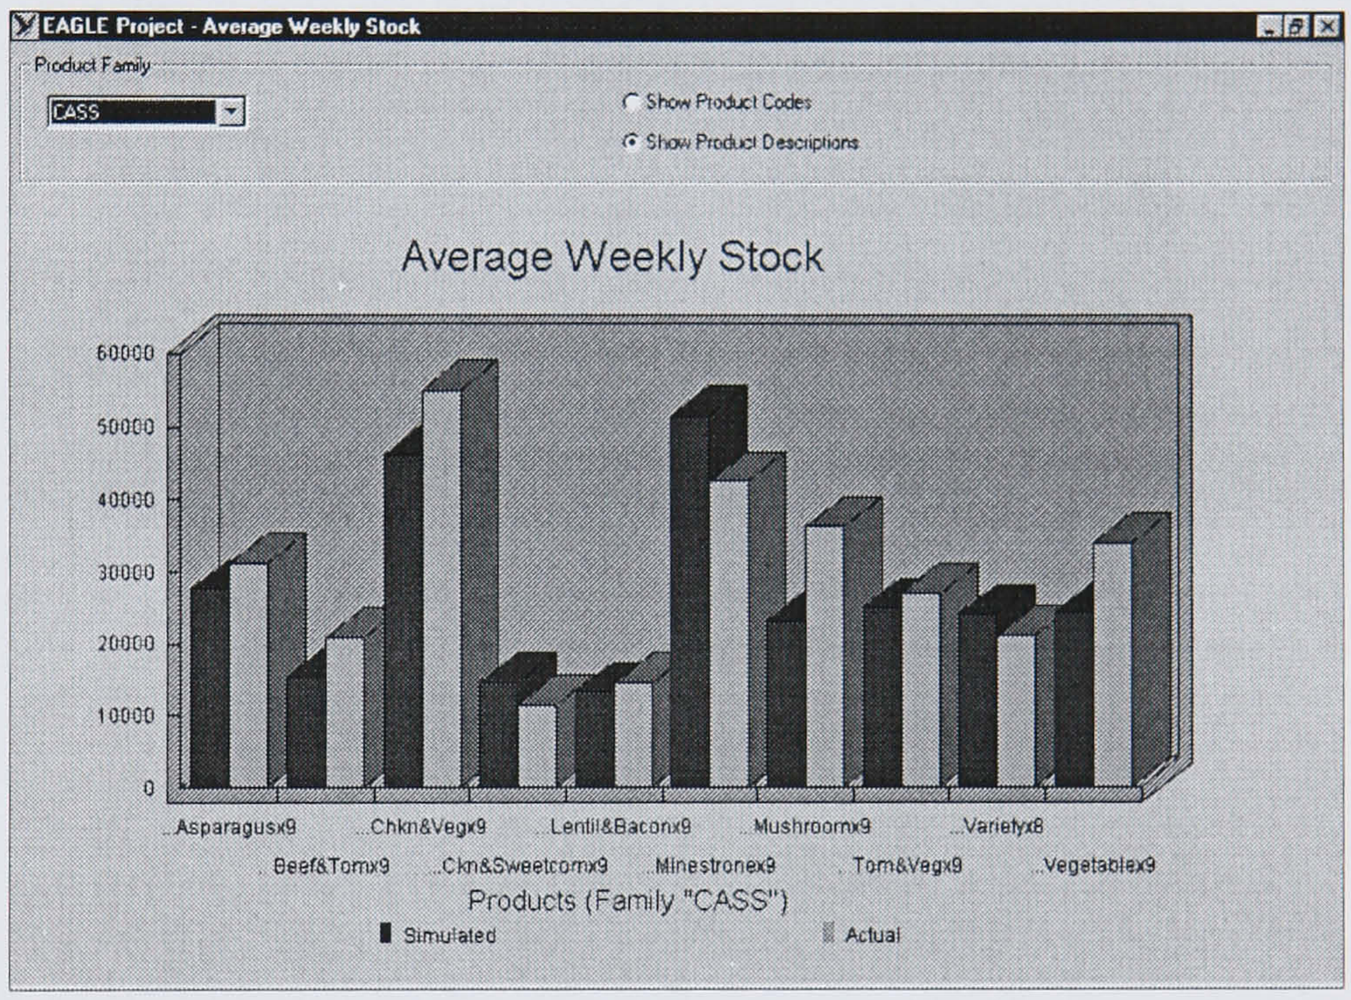
A more comprehensive assessment of the behaviour of the SRS system can be obtained from Fig. 4 which contains average stock level over the test period for every product. These are summarised in terms of product families in Fig. 5. Simple overall statistics of interest include the average week forward cover for all products and the total number of product runs undertaken on all four production lines, during the test period. For the six wfc parameter setting the grand overall averages were as follows:

Overall average stock (as is)	10.4 wfc
Overall SRS average stock	8 wfc

Total SRS replenishment runs	595

The last total refers to the total number of replenishment runs over all four lines for the total test period.


EAGLE Project - Average Weekly Stock
Product Family
CAS
Show Product Codes
Show Product Descriptions
Average Weekly Stock
50000
40000
30000
20000
10000
0
..Beef&Tomx14
..Chickenx14
..Chkn&Leekx14
..GoldenVegx14
..Oxtailx14
..Tomatox14
..Varietyx9
Products (Family "CAS")
Simulated
Actual


EAGLE Project - Average Weekly Stock
Product Family
CASS
Show Product Codes
Show Product Descriptions
Average Weekly Stock
60000
50000
40000
30000
20000
10000
0
..Asparagusx9
..Beef&Tomx9
..Chkn&Vegx9
..Chkn&Sweetcornx9
..Lentil&Baconx9
..Minestronex9
..Mushroomx9
..Tom&Vegx9
..Varietyx8
..Vegetablex9
Products (Family "CASS")
Simulated
Actual

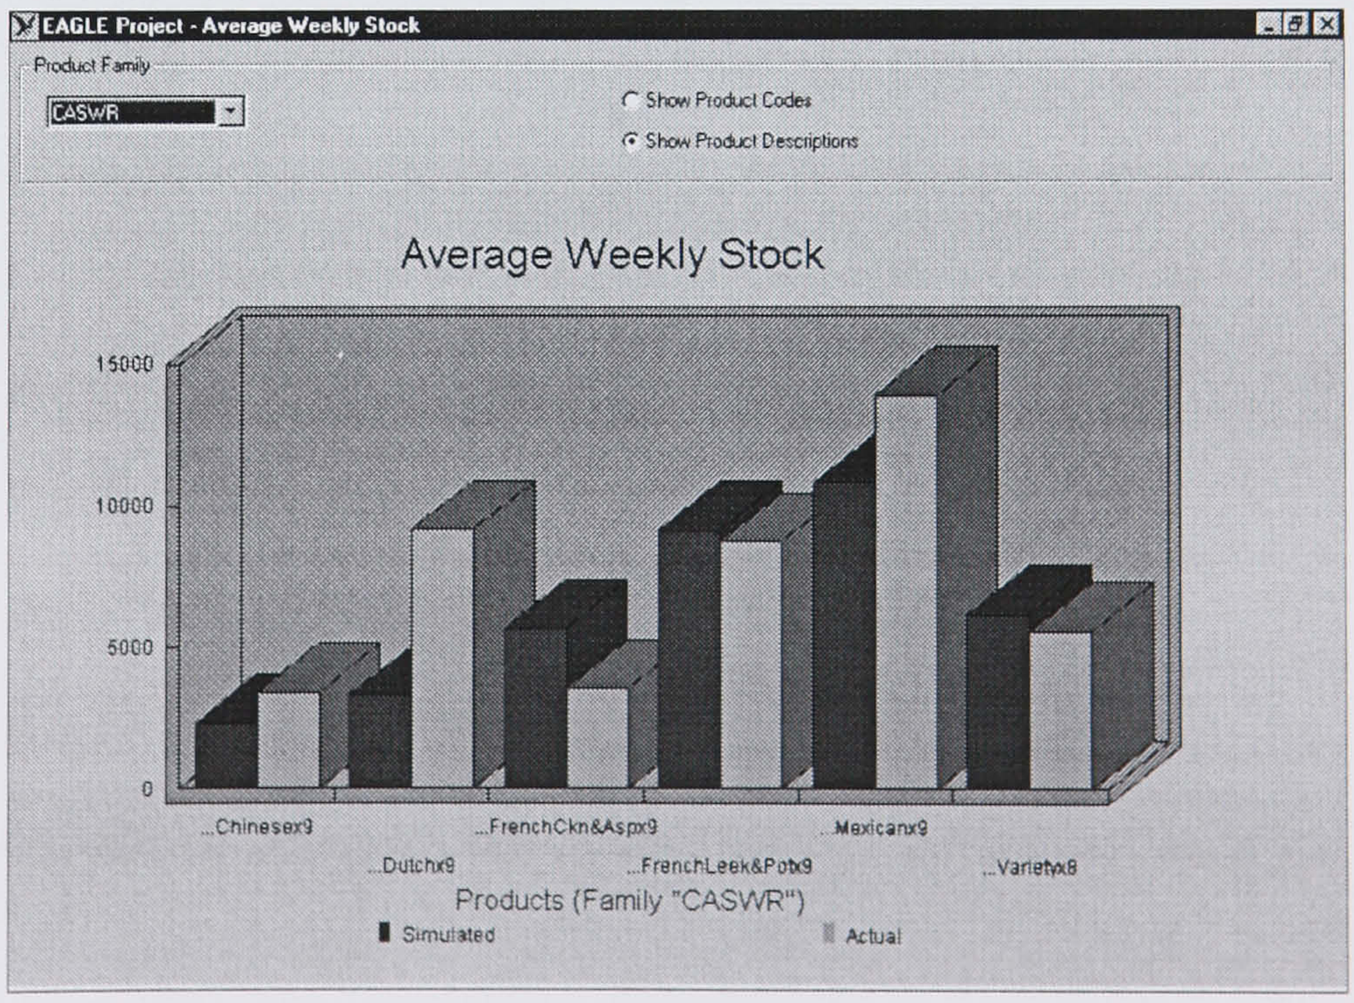

EAGLE Project - Average Weekly Stock
Product Family
CASWR
Show Product Codes
Show Product Descriptions
Average Weekly Stock
15000
10000
5000
0
...Chinesex9
...Dutchx9
...FrenchChkn&Aspx9
...FrenchLeek&Potx9
...Mexicanx9
...Varietyx8
Products (Family "CASWR")
Simulated
Actual

Fig. 3

PRODUCT	CODE	Average Demand	MRL	IQ	ACTUAL AV. STOCK	wfc	SIMULATED AV.WKLY STK 4WFC	% SAVED	SIMULATED AV.WKLY STK 2WFC	% SAVED
CAS Beef & Tom x14	0034326	3428	4498	748	n/a	n/a	15772	n/a	9560	n/a
CAS Chicken x14	0034338	8441	4242	606	38430.5	4.6	30626	20.3%	16363	57.4%
CAS Chkn & Leek x14	0035562	4354	4164	694	33796.2	7.8	16927	49.9%	10131	70.0%
CAS Golden Veg x14	0034347	3948	4554	759	28519.7	7.2	16546	42.0%	10094	64.6%
CAS Oxtail x14	0034348	1905	4158	693	13167.1	6.9	9456	28.2%	6684	49.2%
CAS Tomato x14	0034349	5323	4296	715	44387.9	8.3	23221	47.7%	13205	70.3%
CAS Variety x9	1001767	2543	0	1	12750.1	5.0	12167	4.6%	7198	43.5%
AVERAGE		4278	3702	602	28508.6	6.7	17816	32.1%	10462.1	59.2%
CASS Asparagus x9	1005576	3579	4531	197	31296.4	8.7	18263	41.6%	11401	63.6%
CASS Beef & Tom x9	1005581	1984	4530	302	20987.6	10.6	10936	47.9%	7539	64.1%
CASS Chkn & Veg x9	1005575	7347	4510	205	55212.0	7.5	29762	46.1%	15993	71.0%
CASS Ckn & Sweetcorn x9	1005578	1459	4560	285	11485.0	7.9	10620	7.5%	7580	34.0%
CASS Lentil & Bacon x9	1005580	1687	4644	258	14525.5	8.6	9619	33.8%	6716	53.8%
CASS Minestrone x9	1005582	8208	4630	463	42551.9	5.2	33755	20.7%	17569	58.7%
CASS Mushroom x9	1005574	3763	4540	227	36218.2	9.6	18180	49.8%	10944	69.8%
CASS Tom & Veg x9	1005583	3546	4572	254	26960.0	7.6	18052	33.0%	10368	61.5%
CASS Variety x8	1005584	4080	2084	2084	21145.8	5.2	16788	20.6%	9755	53.9%
CASS Vegetable x9	1005577	3899	4440	222	33965.3	8.7	18099	46.7%	10847	68.1%
AVERAGE		3955	4304	450	29434.8	7.4	18407	34.8%	10871.2	59.8%
CASWR Chinese x9	0034402	96	1764	882	3418.1	35.6	1875	45.1%	1801	47.3%
CASWR Dutch x9	0034404	274	2198	1099	9168.4	33.4	3231	64.8%	2773	69.8%
CASWR French Ckn & Asp x9	0034399	379	3812	953	3563.2	9.4	4238	-18.9%	3761	-5.6%
CASWR French Leek & Pot x9	0034403	877	4626	771	8763.4	10.0	6644	24.2%	5102	41.8%
CASWR Mexican x9	0034401	1354	4763	679	13907.9	10.3	8221	40.9%	6434	53.7%
CASWR Variety x8	1001758	553	3046	1523	5579.4	10.1	4785	14.2%	3770	32.4%
AVERAGE		589	3368	985	7400.1	12.6	4832	28.4%	3940.2	39.9%
SAS Asparagus x14	1008704	495	4191	381	4395.4	8.9	5367	-22.1%	3621	17.6%
SAS Chicken x14	1008700	605	4839	483	5349.1	8.8	5816	-8.7%	4808	10.1%
SAS Chkn & Leek x14	1008702	1411	4420	442	26028.3	18.5	8291	68.1%	5860	77.5%
SAS Golden Veg x14	1008701	1573	4210	421	18994.0	12.1	9307	51.0%	6500	65.8%
SAS Variety x9	1008715	1581	0	1	11073.1	7.0	7930	28.4%	5010	54.8%
SAS Veg & Beef x14	1008703	966	4220	422	13768.8	14.3	6799	50.6%	4865	64.7%
AVERAGE		1105	3647	358	13268.1	12.0	7252	27.9%	5110.7	48.4%
SASS Beef & Tomato x9	1008708	418	4712	589	3195.5	7.7	4771	-49.3%	4182	-30.9%
SASS Broc & Caul x9	1008713	2119	4820	482	14124.8	6.7	11586	18.0%	7971	43.6%
SASS Cajun Spicy Veg x9	1008705	1102	4850	485	5839.8	5.3	7553	-29.3%	4388	24.9%
SASS Chicken & Mushroom x9	1008707	2068	4587	417	24084.1	11.6	12320	48.8%	8116	66.3%
SASS Chicken & Sweetcorn x9	1008711	1875	5420	542	30764.2	16.4	8756	71.5%	6678	78.3%
SASS Chkn & Veg x9	1008709	1477	4272	712	13729.2	9.3	8367	39.1%	6409	53.3%
SASS French Mush & Herbs x9	1008712	589	4653	517	6496.7	11.0	4138	36.3%	3750	42.3%
SASS Leek & Potato x9	1008706	1032	4750	475	10397.7	10.1	7553	27.4%	5772	44.5%
SASS Mediterranean Tomato x14	1008699	1510	4464	496	16286.6	10.8	7275	55.3%	5456	66.5%
SASS Minestrone x9	1008710	3283	7020	702	24031.9	7.3	14916	37.9%	9826	59.1%
SASS Tom & Lentil x9	1008714	302	1305	435	1311.0	4.3	2284	-74.2%	2012	-53.5%
SASS Variety x8	1008716	2012	3204	1602	12923.4	6.4	13209	-2.2%	8403	35.0%
AVERAGE		1482	4505	621	13598.7	9.2	8561	14.9%	6080.3	35.8%
SNAS Chilli Beef&Tom Scht x1	1001568	225	1071	1071	3374.0	15.0	1949	42.2%	1726	48.8%
SNAS Chkn & Mushroom Sach x1	1001605	982	1006	1006	7491.6	7.6	6967	7.0%	4661	37.8%
SNAS Smoky Bacon Sach x12	1001567	57	1069	1069	889.5	15.5	1210	-36.0%	1165	-31.0%
SNAS Spicy Veg Sach x12	1001606	330	857	857	5186.4	15.7	3028	41.6%	2350	54.7%
AVERAGE		399	1001	1001	4235.4	10.6	3289	13.7%	2475.5	27.6%
GRAND AVERAGE		2201	3758	627	17489.4	8.0	11048	37.4%	7091.5	59.8%
TOTAL		95609			769534.6	8.0	481403	37.4%	309557.0	59.8%

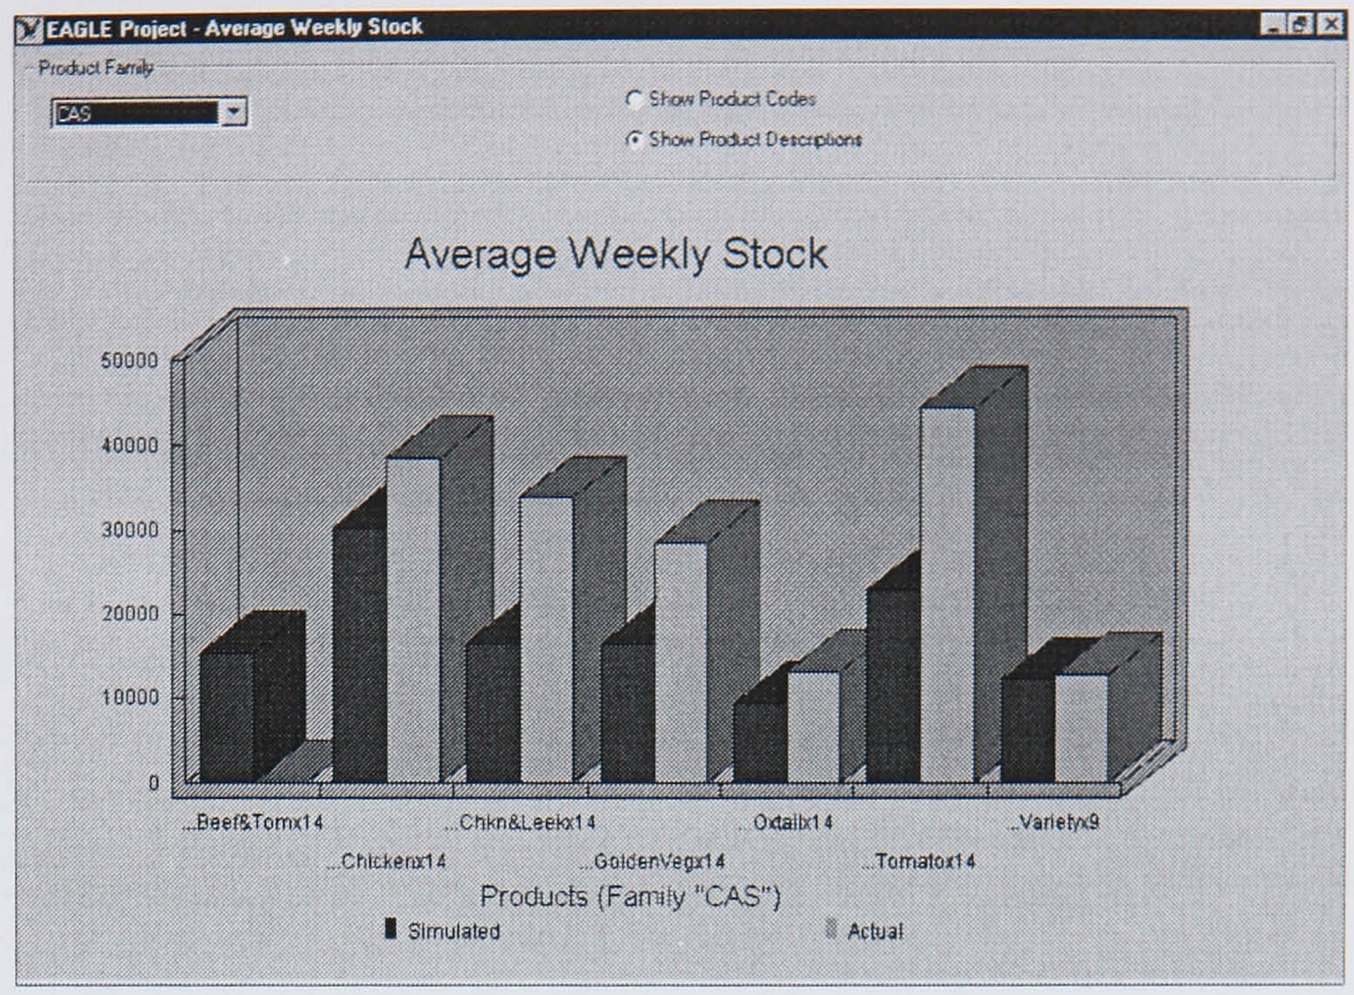
Fig. 4

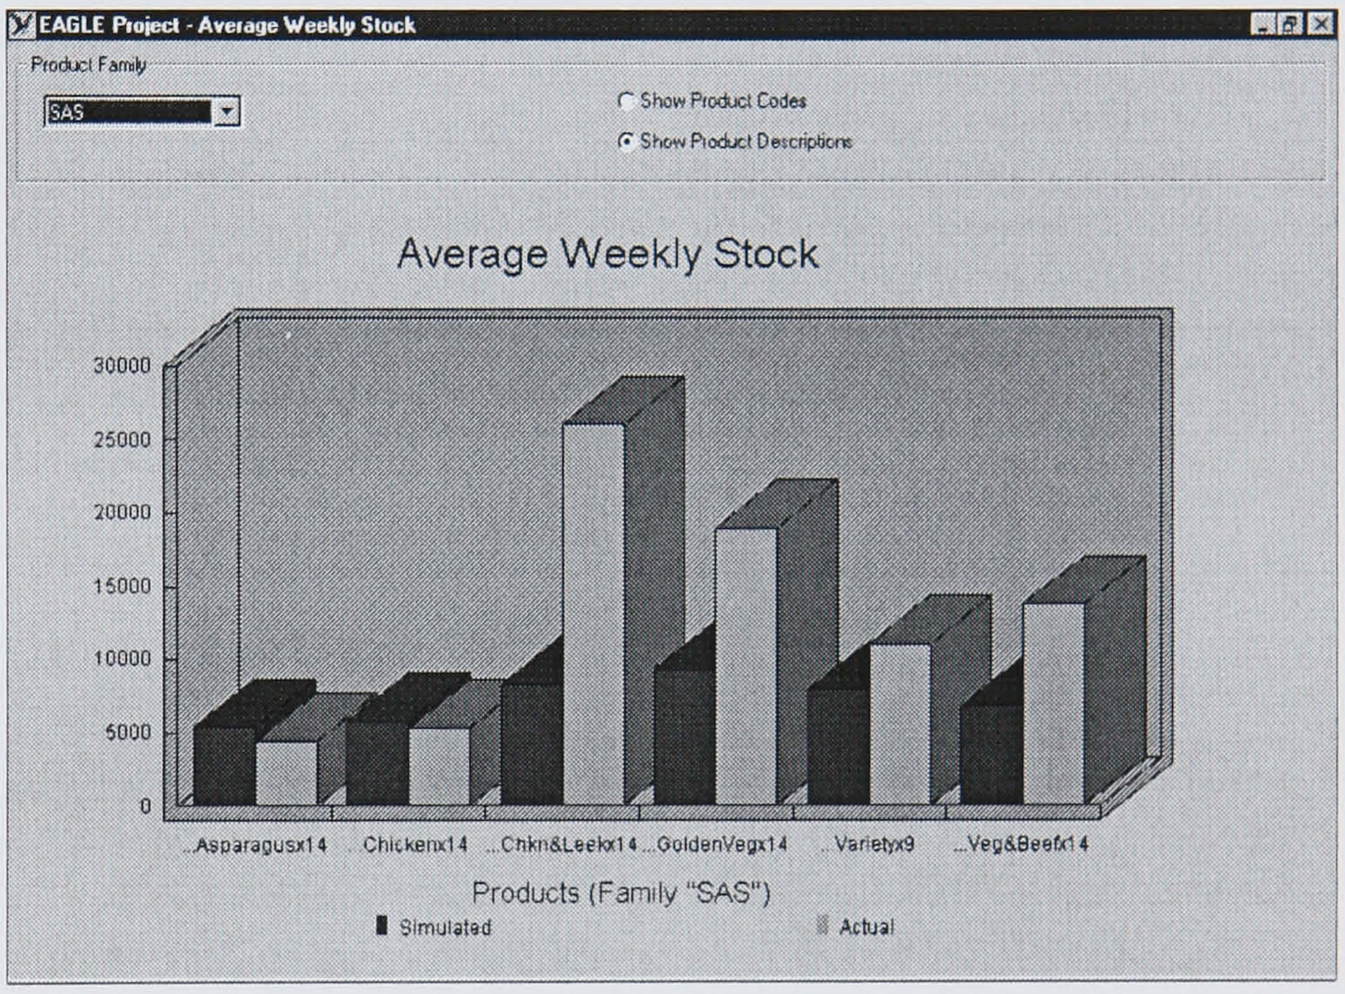
FAMILY	'AS IS'	4 WFC	2 WFC
CAS	5.7	4.2	2.4
CASS	7.4	4.7	2.7
CASWR	12.6	8.2	6.7
SAS	12.0	6.6	4.6
SASS	9.2	5.8	4.1
SNAS	10.6	8.2	6.2
GROUP	8.0	5.0	3.2

Fig. 5

4 wfc results - The same sample results as illustrated in the previous section are reproduced in Fig. 6 with the wfc set to four weeks. The high demand product (such as CAS tomato), exhibits significant reductions in stock level. In contrast to low demand products (such as CASWR) at the bottom of Fig. 6 exhibit relatively little change in stock level. This is due to the overriding effect of the

minimum run length and Incremental Quantity setting overriding the computed 4 wfc and hence holding a high stock level. The imposition of an MRL constraint increases the stock level and decreases the replenishment frequency.


EAGLE Project - Average Weekly Stock
Product Family
CAS
Show Product Codes
Show Product Descriptions
Average Weekly Stock
50000
40000
30000
20000
10000
0
...Beef&Tomx14
...Chkn&Leekx14
...Oxtailx14
...Varietyx9
...Chickenx14
...GoldenVegx14
...Tomatox14
Products (Family "CAS")
Simulated
Actual


EAGLE Project - Average Weekly Stock
Product Family
SAS
Show Product Codes
Show Product Descriptions
Average Weekly Stock
30000
25000
20000
15000
10000
5000
0
...Asparagusx14
...Chickenx14
...Chkn&Leekx14
...GoldenVegx14
...Varietyx9
...Veg&Beefx14
Products (Family "SAS")
Simulated
Actual

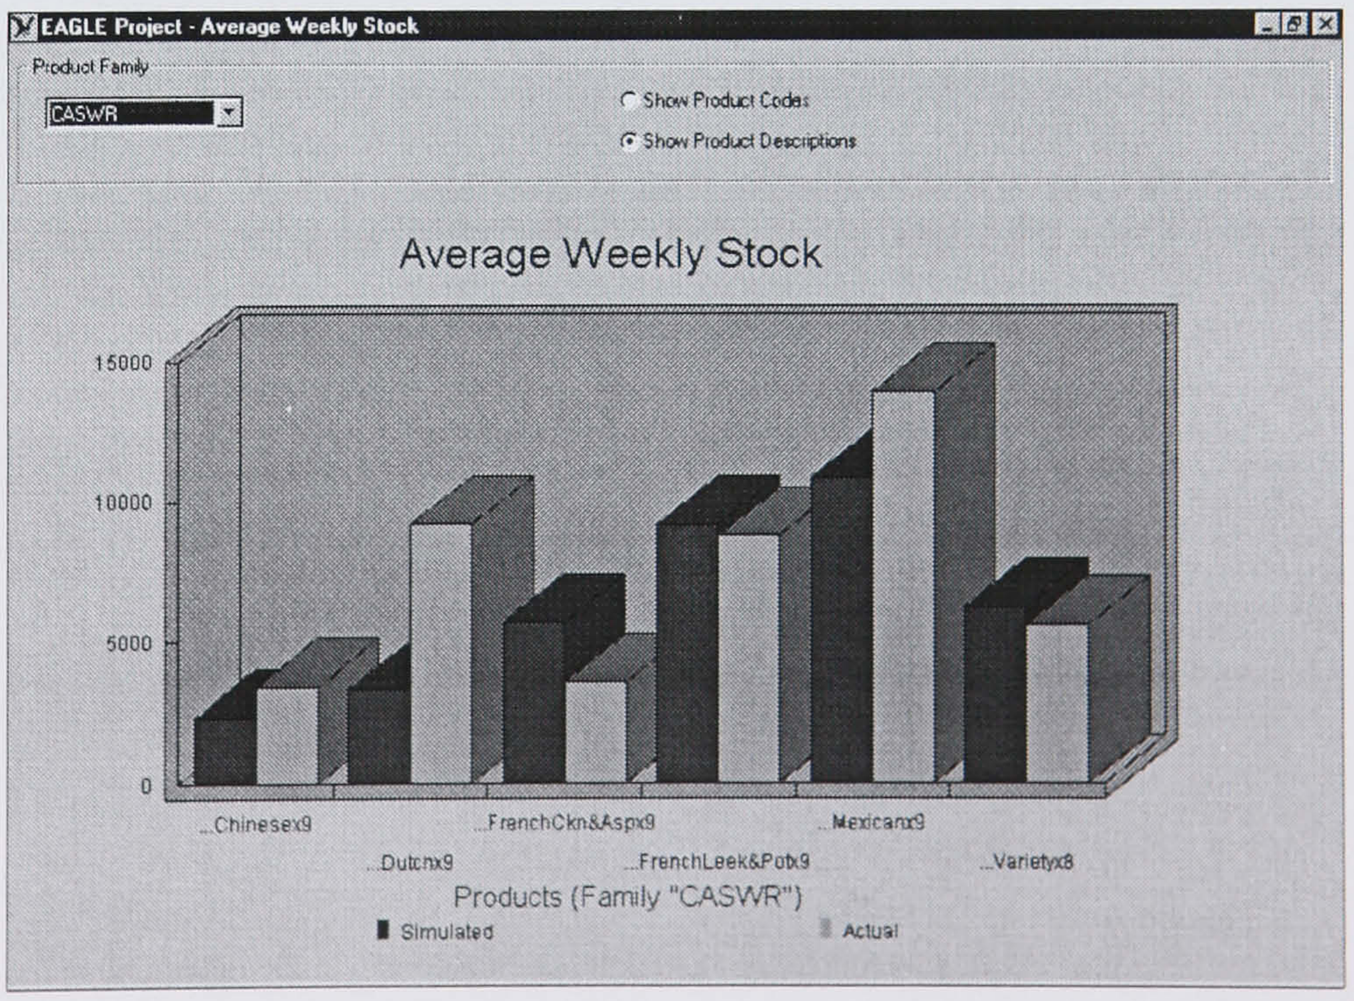

EAGLE Project - Average Weekly Stock
Product Family
CASWR
Show Product Codes
Show Product Descriptions
Average Weekly Stock
15000
10000
5000
0
...Chinesex9
...Dutchx9
...FrenchChkn&Aspx9
...FrenchLeek&Potx9
...Mexicanx9
...Varietyx8
Products (Family "CASWR")
Simulated
Actual

Fig. 6

Again, the detailed result in terms of average stock, for every product are given in Fig. 4 and the results summarised in terms of average for every product class in Fig. 5. The key overall average statistics for the four wfc are:

Overall SRS average stock	5 wfc
Total SRS replenishment runs	588

Comparing this result with that of the previous section we note that a decrease in the wfc parameter has in fact resulted in a decrease in the total number of replenishment runs, and hence switch-overs, throughout the test period. This apparent anomaly is be explained in the discussion below.

2 wfc results - In view of the success of the previous result, we now attempt an even lower wfc settings and repeat the test procedures, the results are shown in Fig. 7. Again these sample results significantly reduce stock levels and here even some of the lower demand products are exhibiting a reduction in stock level, whilst others are unaffected due to the MRL effects discussed above. As in the previous two runs, we experienced no change in the stock out statistics for any of the products (i.e. no stock outs).

The average stock levels recorded from the SRS results are again shown in Fig. 4. Details by product group are shown in Fig. 5. The overall statistics for this run were:

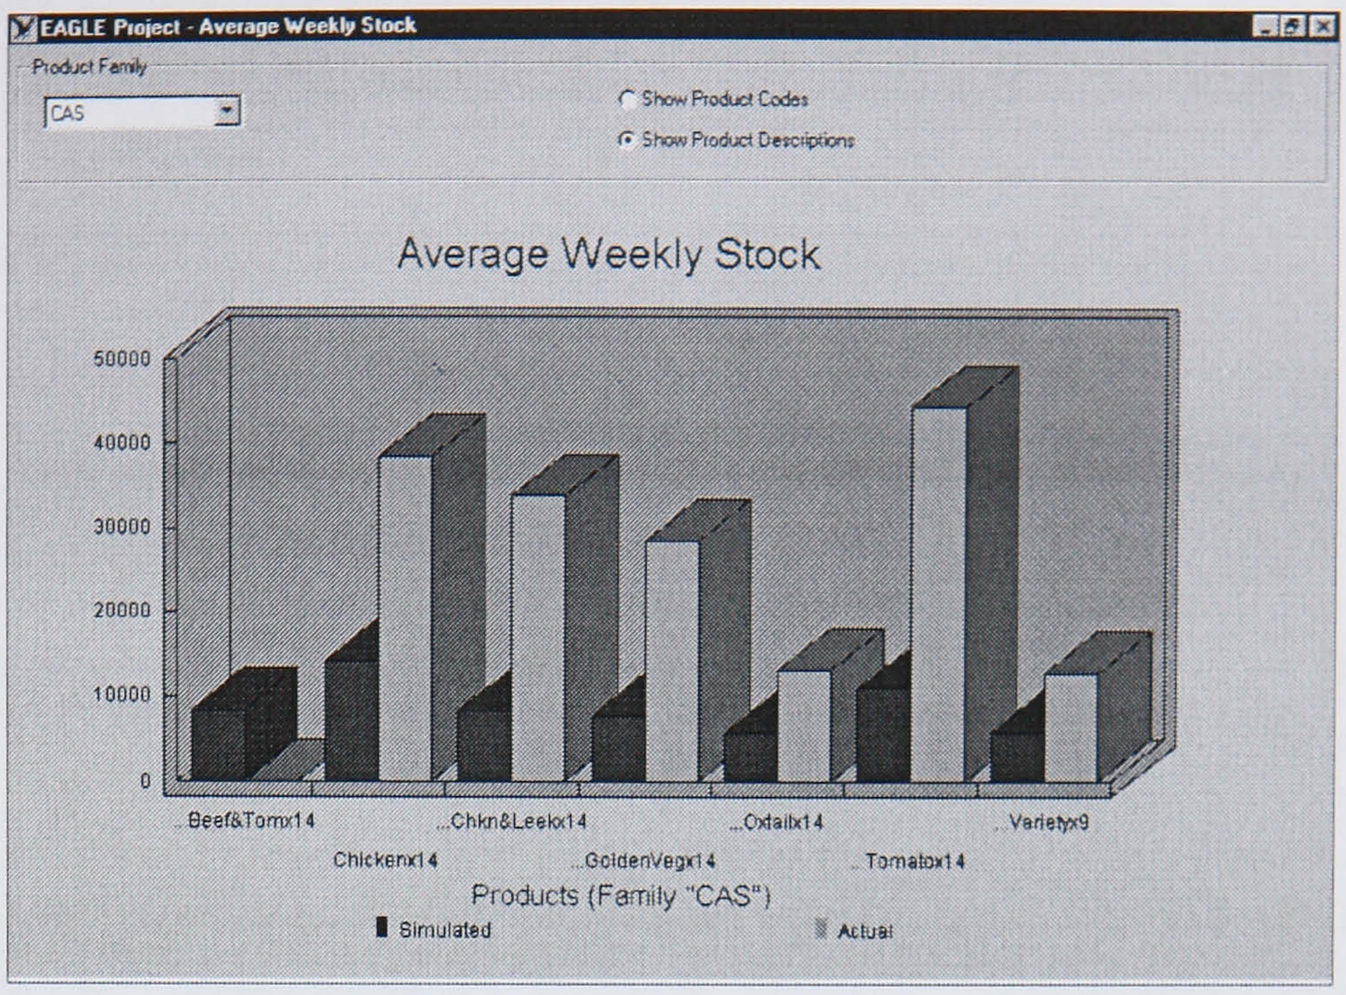
Overall SRS average stock	3.2 wfc
Total SRS replenishment runs	567

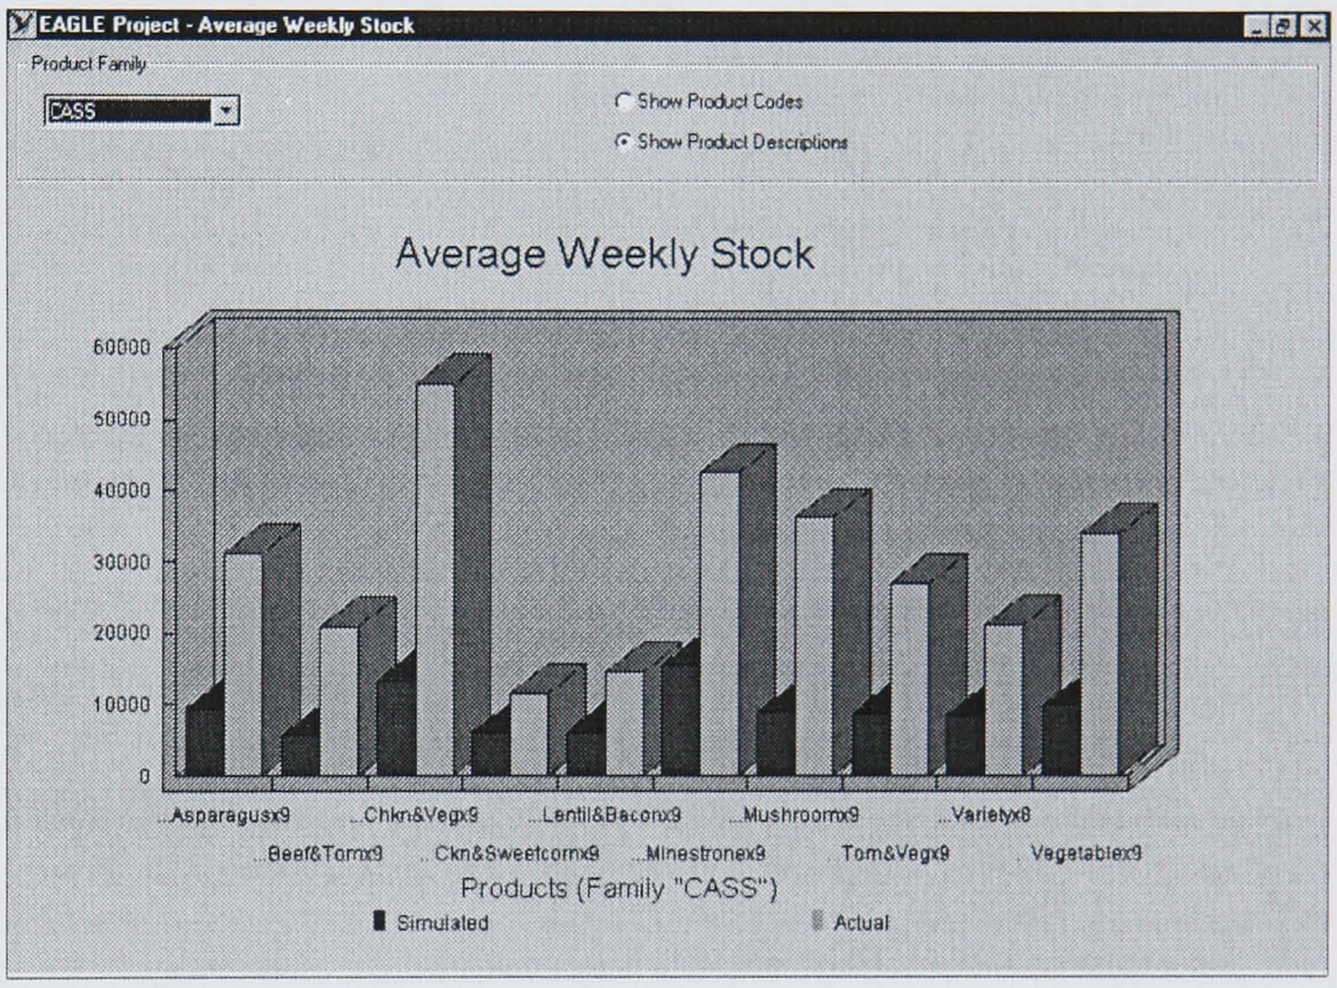
Again the results exhibit the unexpected drop in replenishment runs and switch overs.

Since only weekly demand data was collected, it is not possible to consider wfc parameters by a fraction of a week. Decreasing the wfc parameter to one week is not reasonable since the minimum wfc setting and lead time plus cycle time equals two weeks for this system. Hence, it would seem that a setting of two

weeks for the wfc parameter is optimum. The average stock figures are higher that the wfc setting. This is due to the inclusion of safety stock in the replenishment calculation.

Temperature effect - A 20% increase in demand over the four weeks prior to Christmas and after Christmas was introduced. Using the editing facility indicated in Fig. 1, the entire demand levels for all products were increased by 20% for the appropriate period and the whole tuning procedure repeated with a two week forward cover parameter setting. The system coped with this increase through the use of overtime. In particular lines, A and B were working at weekends for 6 of the total 8 week periods. Given this level of overtime, this system could cope with the simulated temperature drop without loss of customer service and stock build up. The cost of overtime can be balanced against the cost of stock increase to decide best policy and account taken of the probability of extreme weather.


EAGLE Project - Average Weekly Stock
Product Family
CAS
Show Product Codes
Show Product Descriptions
Average Weekly Stock
50000
40000
30000
20000
10000
0
...Beef&Tomx14
...Chkn&Leekx14
...Oxtailx14
...Varietyx9
Chickenx14
...GoldenVegx14
...Tomatox14
Products (Family "CAS")
Simulated
Actual


EAGLE Project - Average Weekly Stock
Product Family
CASS
Show Product Codes
Show Product Descriptions
Average Weekly Stock
60000
50000
40000
30000
20000
10000
0
...Asparagusx9
...Chkn&Vegx9
...Lentil&Baconx9
...Mushroomx9
...Varietyx8
...Beef&Tomx9
...Ckn&Sweetcornx9
...Minestronex9
...Tom&Vegx9
...Vegetablex9
Products (Family "CASS")
Simulated
Actual

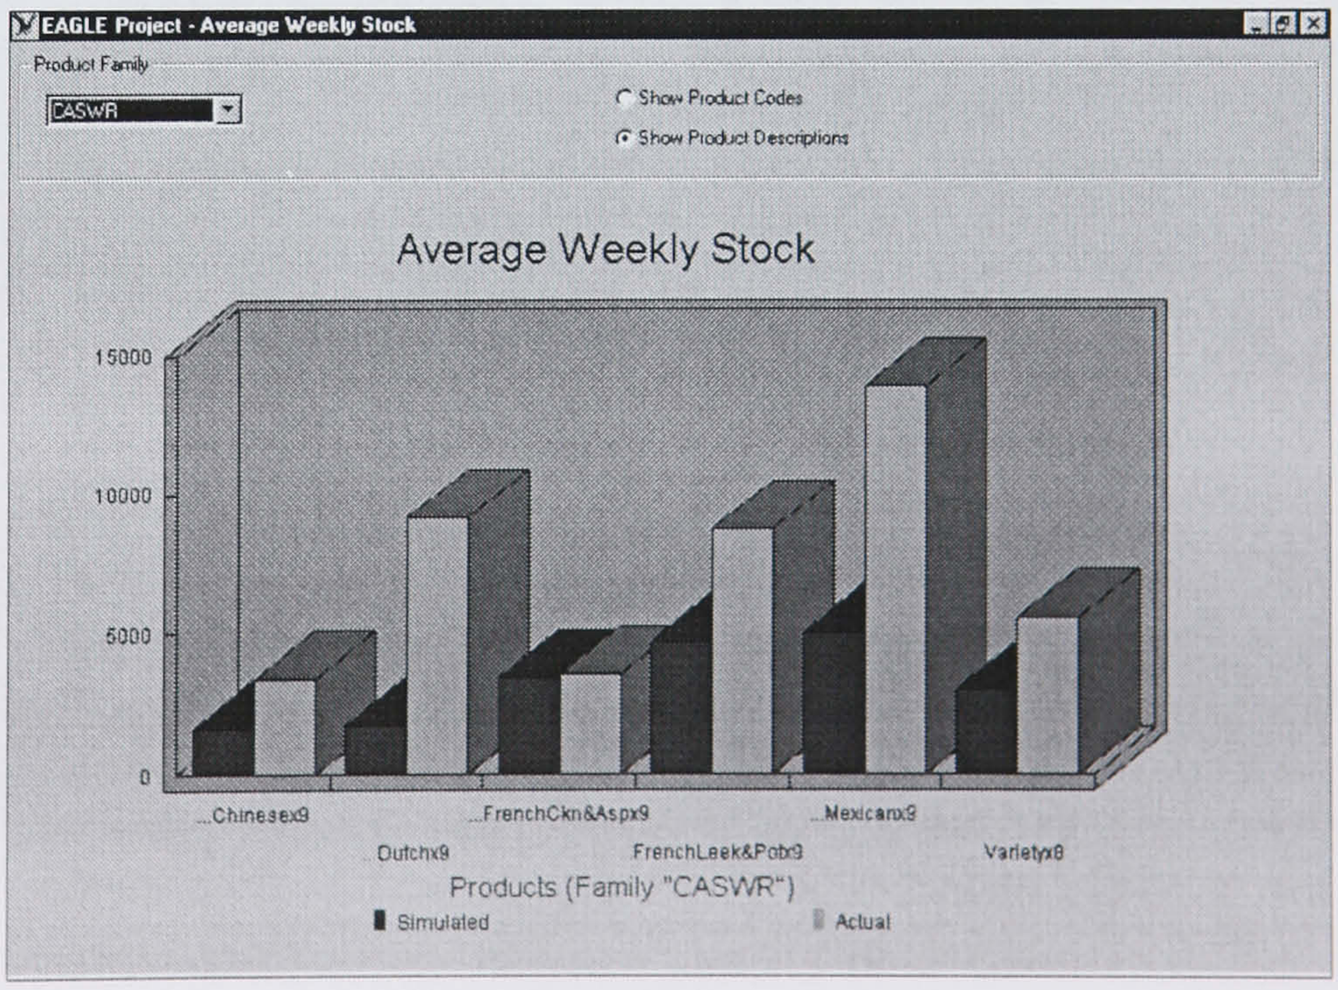

EAGLE Project - Average Weekly Stock
Product Family
CASWR
Show Product Codes
Show Product Descriptions
Average Weekly Stock
15000
10000
5000
0
...Chinesex9
...FrenchCkn&Aspx9
...Mexicanx9
...Dutchx9
...FrenchLeek&Potx9
...Varietyx8
Products (Family "CASWR")
Simulated
Actual

Fig. 7

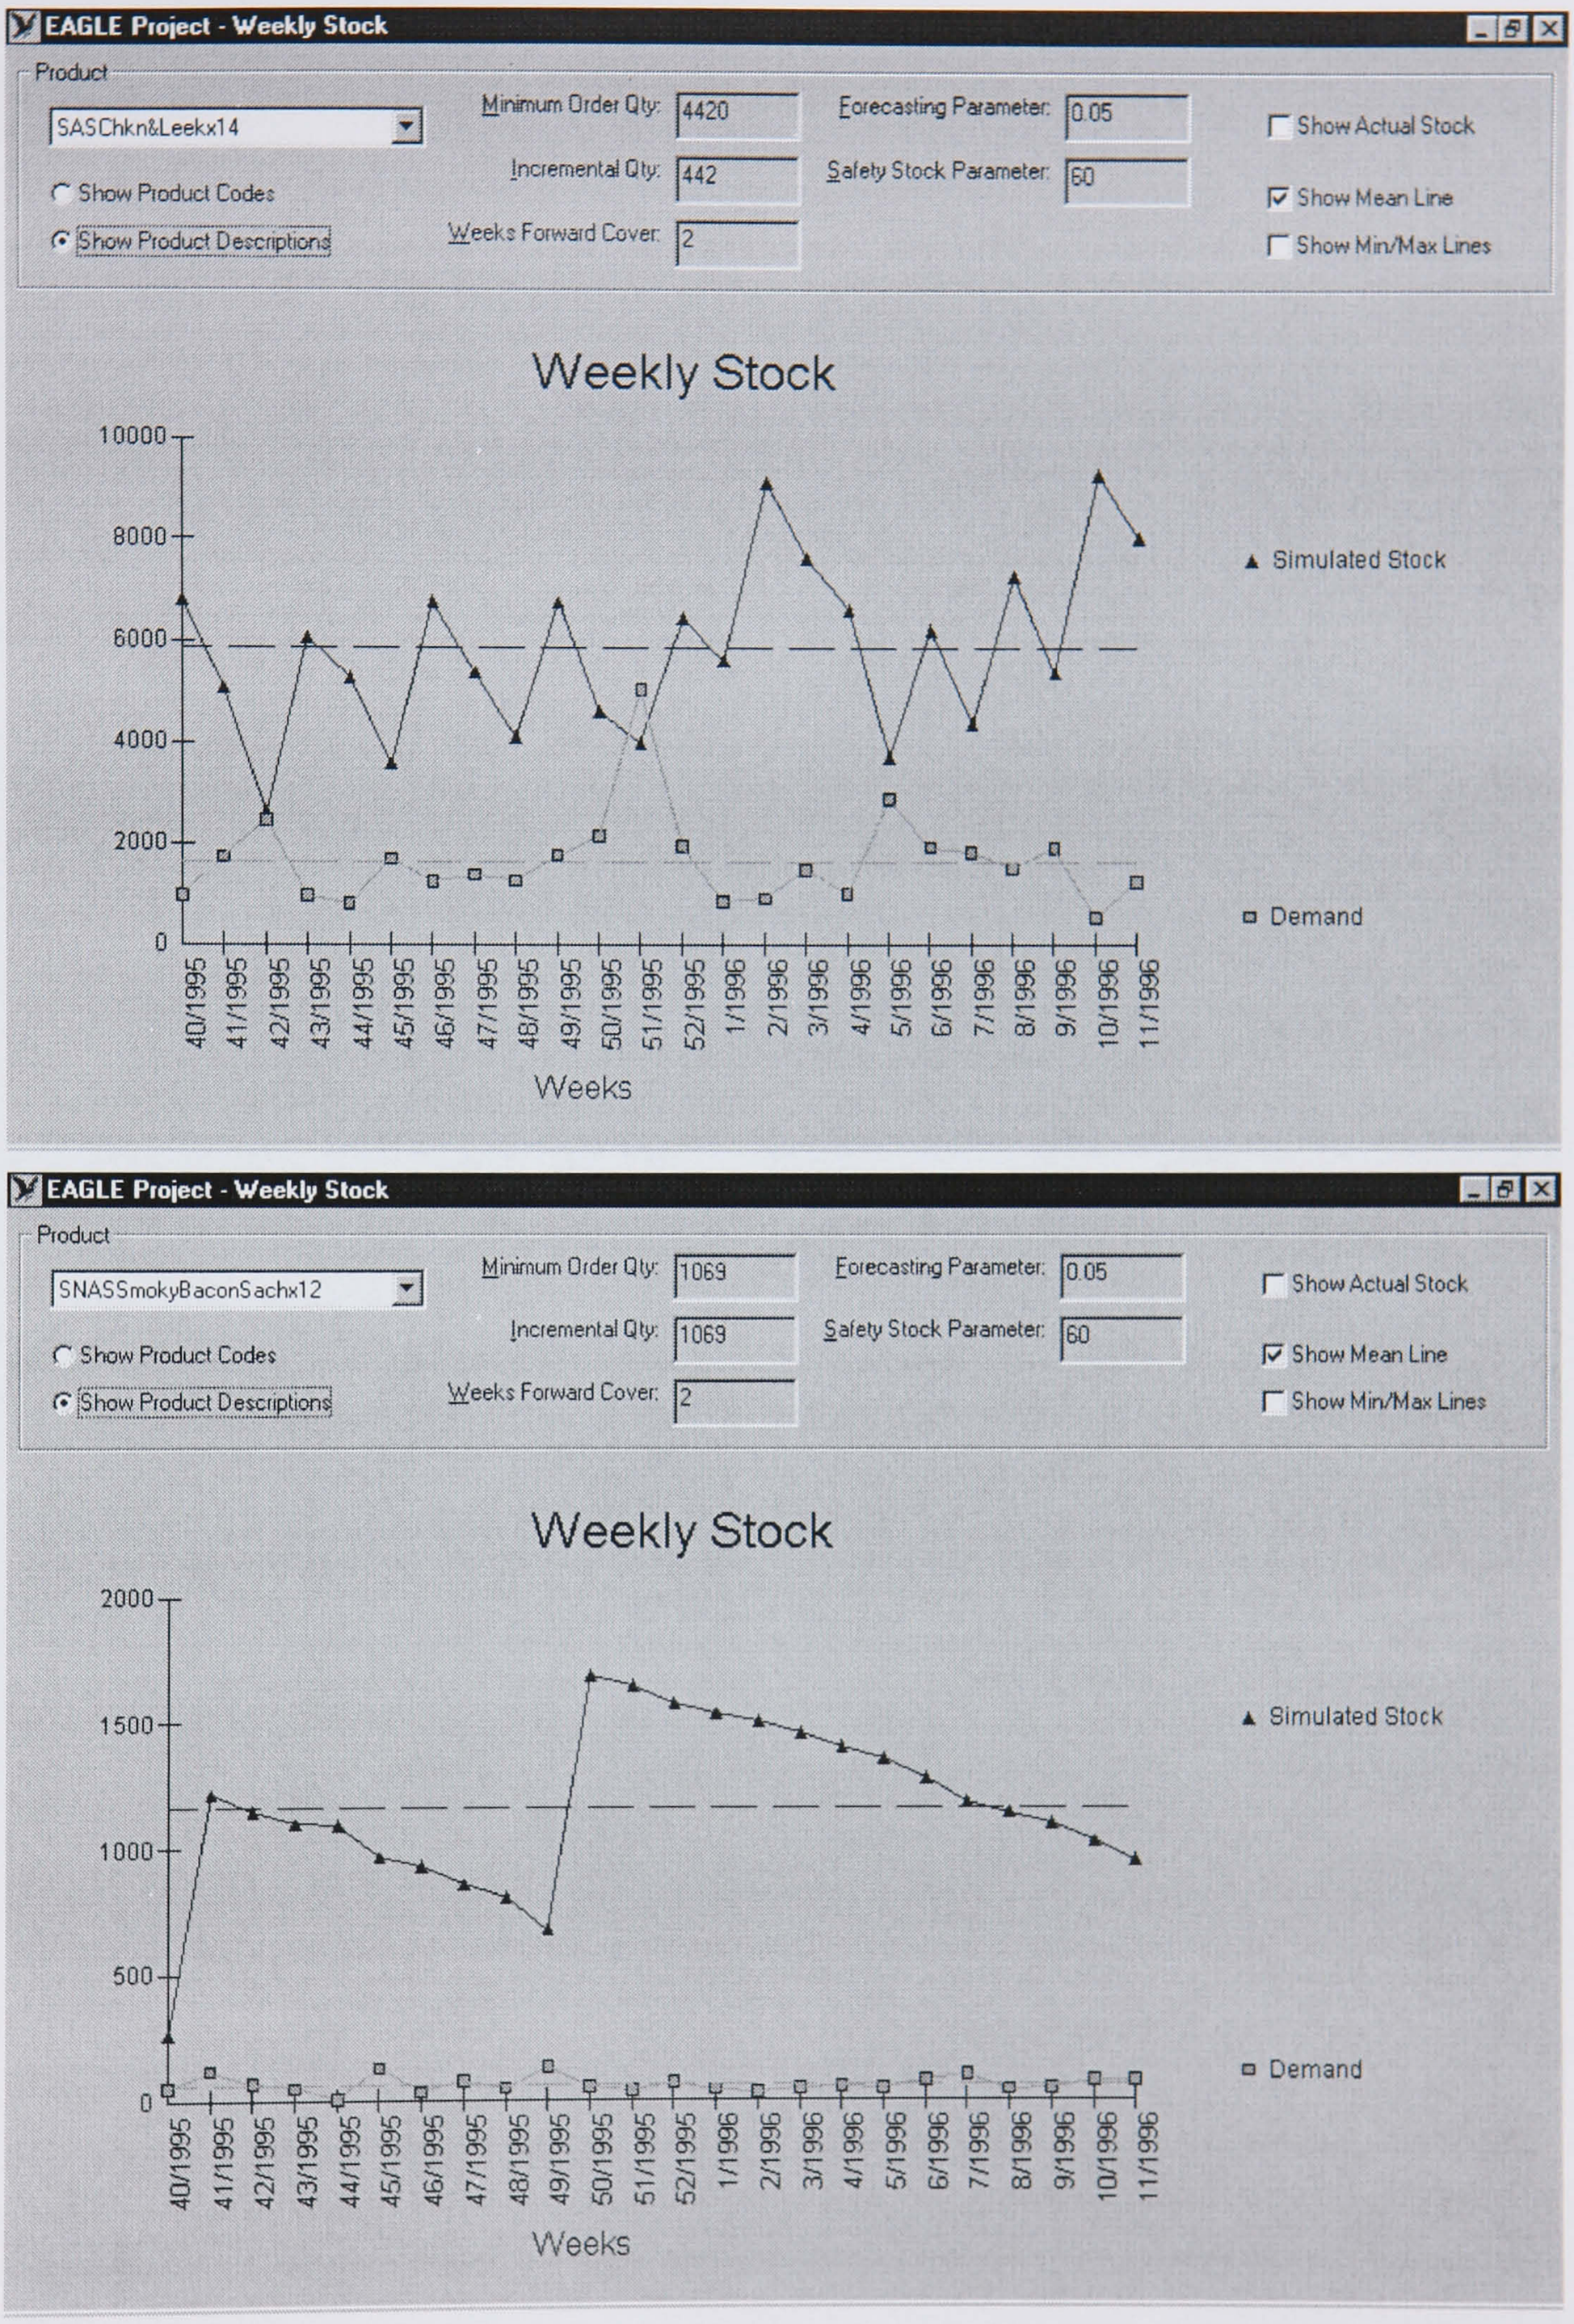

EAGLE Project - Weekly Stock
Product
SASChkn&Leekx14
Show Product Codes
Show Product Descriptions
Minimum Order Qty: 4420
Incremental Qty: 442
Weeks Forward Cover: 2
Forecasting Parameter: 0.05
Safety Stock Parameter: 60
Show Actual Stock
Show Mean Line
Show Min/Max Lines
Weekly Stock
10000
8000
6000
4000
2000
0
40/1995
41/1995
42/1995
43/1995
44/1995
45/1995
46/1995
47/1995
48/1995
49/1995
50/1995
51/1995
52/1995
1/1996
2/1996
3/1996
4/1996
5/1996
6/1996
7/1996
8/1996
9/1996
10/1996
11/1996
Weeks
Simulated Stock
Demand
EAGLE Project - Weekly Stock
Product
SNASSmokyBaconSachx12
Show Product Codes
Show Product Descriptions
Minimum Order Qty: 1069
Incremental Qty: 1069
Weeks Forward Cover: 2
Forecasting Parameter: 0.05
Safety Stock Parameter: 60
Show Actual Stock
Show Mean Line
Show Min/Max Lines
Weekly Stock
2000
1500
1000
500
0
40/1995
41/1995
42/1995
43/1995
44/1995
45/1995
46/1995
47/1995
48/1995
49/1995
50/1995
51/1995
52/1995
1/1996
2/1996
3/1996
4/1996
5/1996
6/1996
7/1996
8/1996
9/1996
10/1996
11/1996
Weeks
Simulated Stock
Demand

Fig. 8

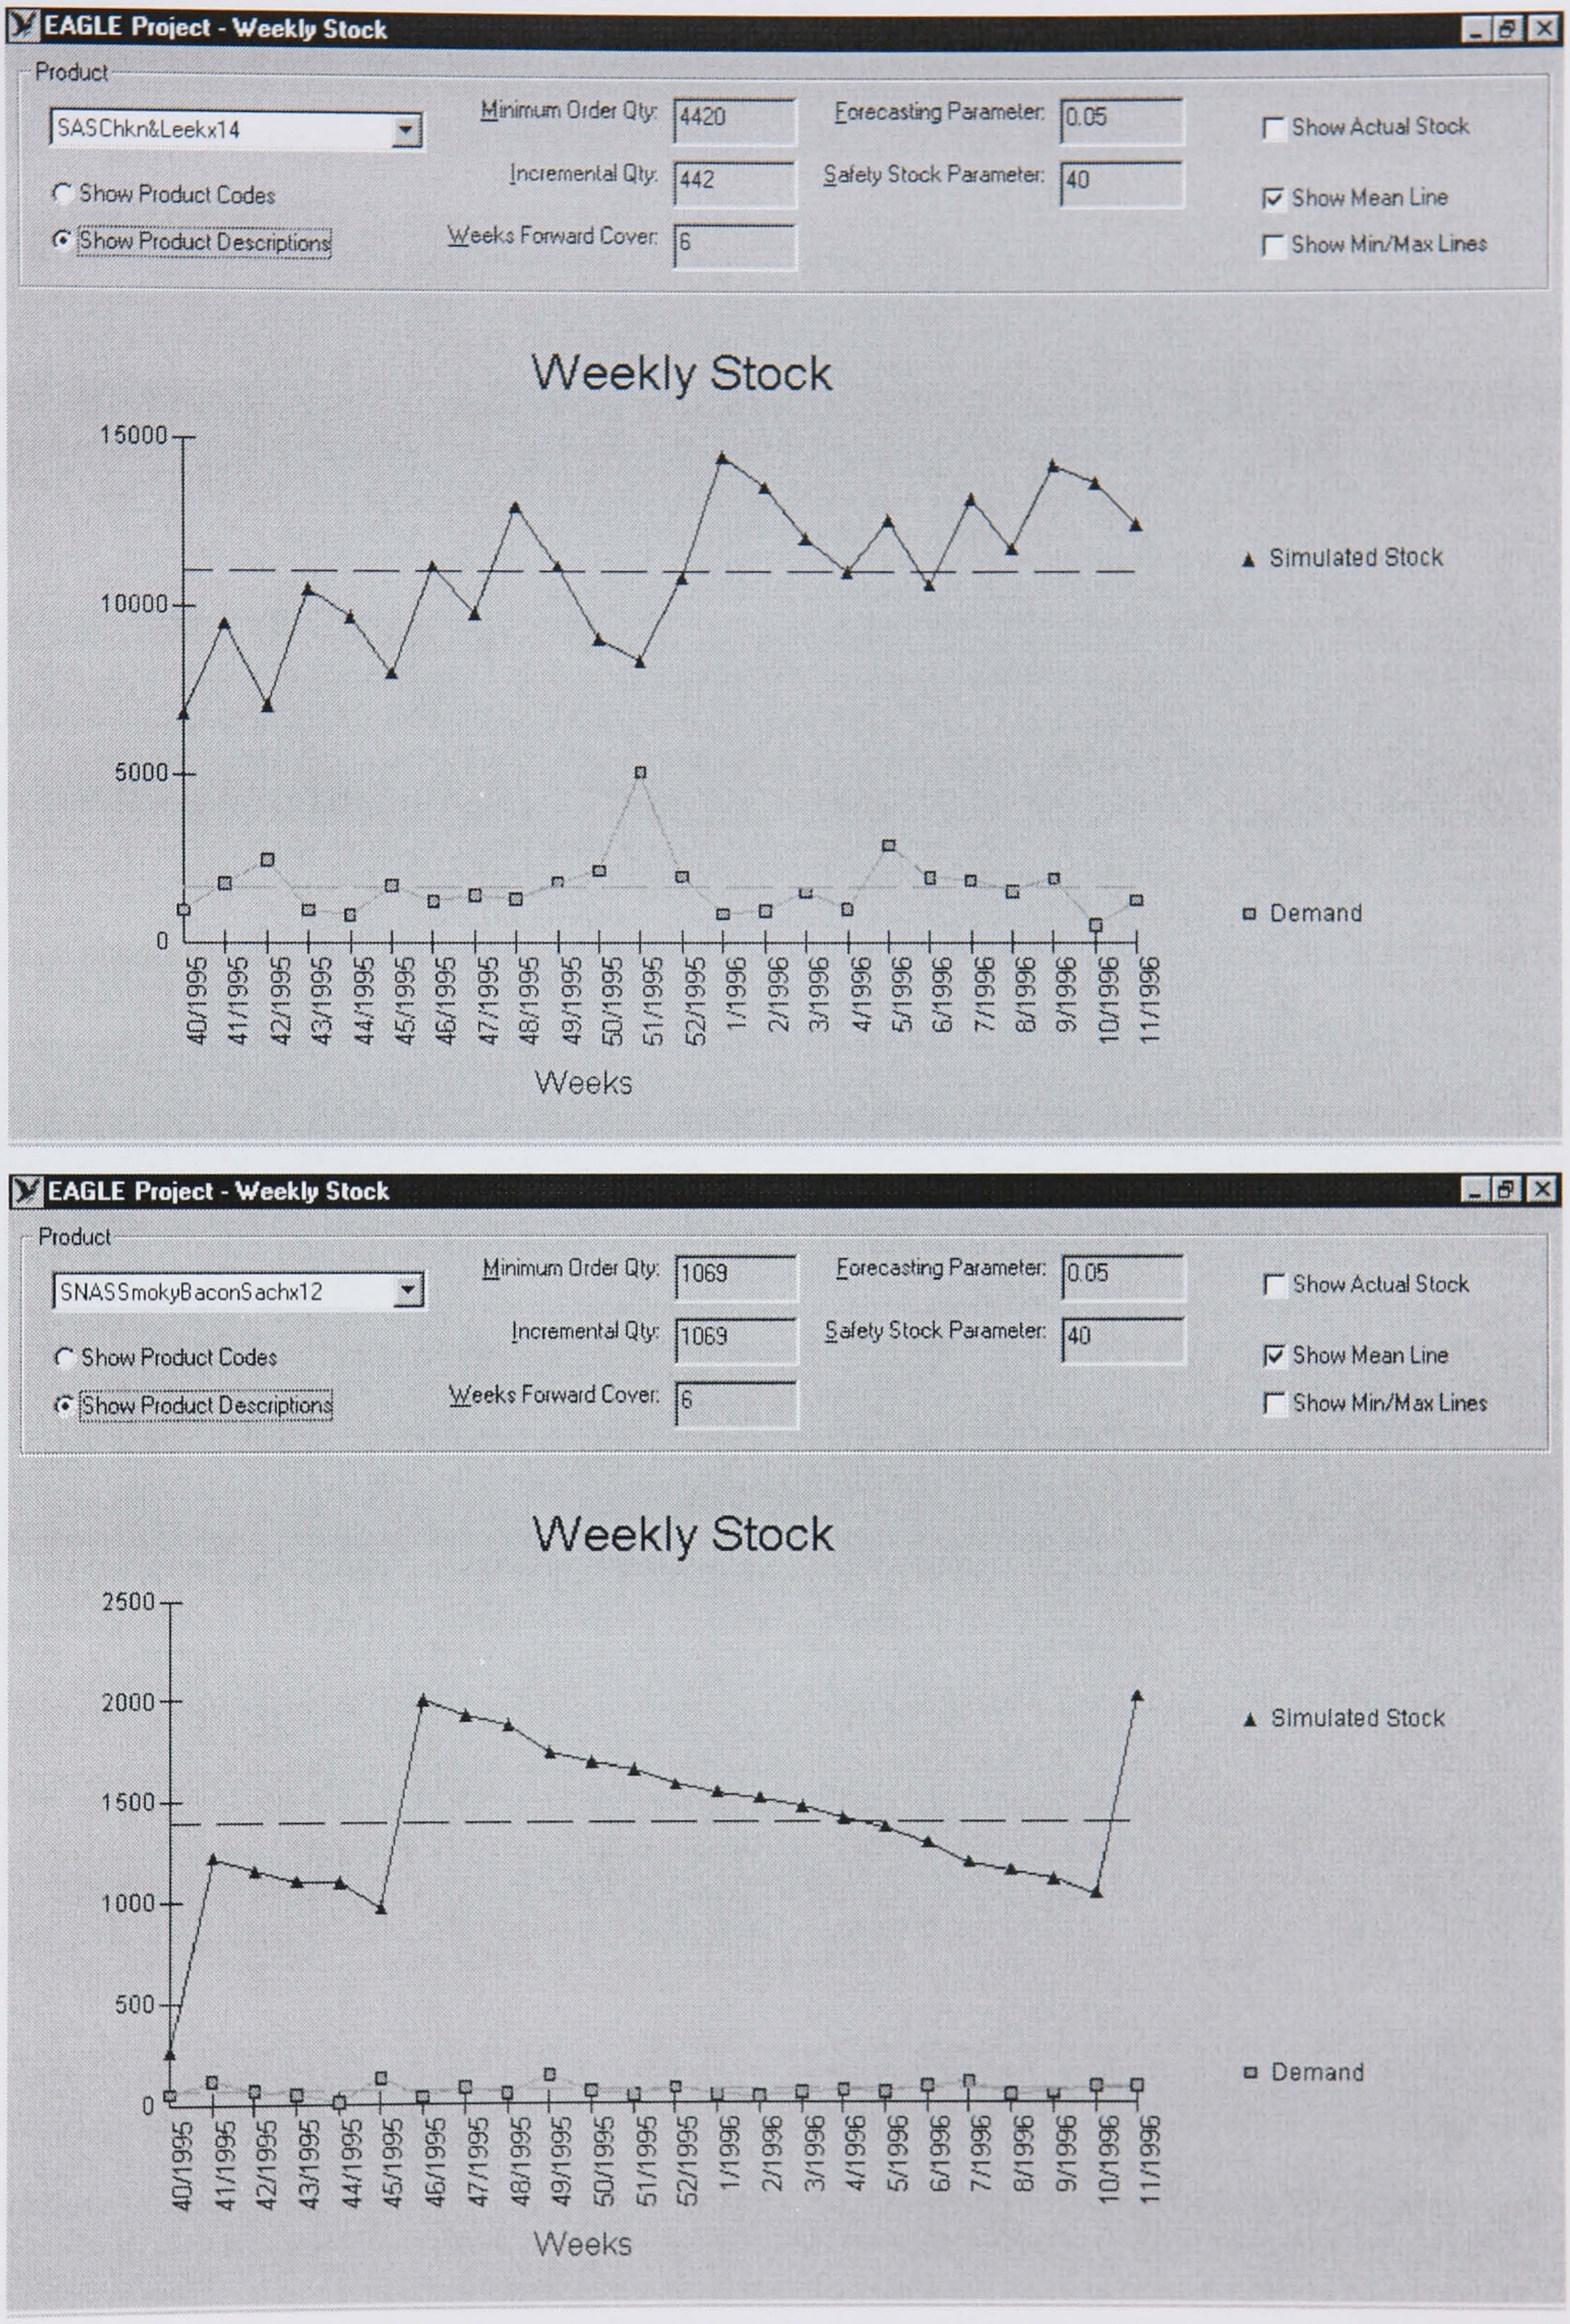

EAGLE Project - Weekly Stock
Product
SASChkn&Leekx14
Show Product Codes
Show Product Descriptions
Minimum Order Qty: 4420
Incremental Qty: 442
Weeks Forward Cover: 6
Forecasting Parameter: 0.05
Safety Stock Parameter: 40
Show Actual Stock
Show Mean Line
Show Min/Max Lines
Weekly Stock
15000
10000
5000
0
Simulated Stock
Demand
40/1995
41/1995
42/1995
43/1995
44/1995
45/1995
46/1995
47/1995
48/1995
49/1995
50/1995
51/1995
52/1995
1/1996
2/1996
3/1996
4/1996
5/1996
6/1996
7/1996
8/1996
9/1996
10/1996
11/1996
Weeks
EAGLE Project - Weekly Stock
Product
SNASSmokyBaconSachx12
Show Product Codes
Show Product Descriptions
Minimum Order Qty: 1069
Incremental Qty: 1069
Weeks Forward Cover: 6
Forecasting Parameter: 0.05
Safety Stock Parameter: 40
Show Actual Stock
Show Mean Line
Show Min/Max Lines
Weekly Stock
2500
2000
1500
1000
500
0
Simulated Stock
Demand
40/1995
41/1995
42/1995
43/1995
44/1995
45/1995
46/1995
47/1995
48/1995
49/1995
50/1995
51/1995
52/1995
1/1996
2/1996
3/1996
4/1996
5/1996
6/1996
7/1996
8/1996
9/1996
10/1996
11/1996
Weeks

Fig. 9

Discussion and conclusion - The tuning results discussed above become clearer when individual stock replenishment behaviour is considered. Fig. 8 give samples of weekly stock behaviour for high level and low level demand products for 2wfc, Fig. 9 corresponds to 6 wfc. Note that at the beginning of the high demand response of Fig. 8 the system does not replenish until the second week. This is because the chosen initial conditions used throughout the test were set to 6wfc and hence avoid an early replenishment run. In contrast, for the 6 wfc result of Fig. 9 an immediate replenishment run is requested since the initial condition is now equal to the wfc setting. This effect accounts for the anomalous result of reduction in number of switch-overs with wfc parameter settings. The result is small and entirely due to the initial condition dynamic which favours the lower wfc parameter settings.

The low demand responses for both the 2 and 6 weeks wfc parameter settings each illustrate the overriding effect of MRL which, as required, significantly reduces the replenishment frequency involved with low demand products at the cost of increased stock level. Notice that the stock level remains unchanged whether the parameter setting is set at either 2 or 6. This result is in agreement with equation 8.0 of Chapter 4.

Changes in the wfc parameter from 6 to 2 weeks resulted in the following changes.

1.	Reduction in average stock	55%
2.	Changes in production switch-over	negligible
3.	Change in overtime	0

The average number of switch-overs is close to the optimum number indicated by an internal audit optimisation study and to the number used in practice by the management team (6 per week per line).

The value of the stock level changes are estimated to be as follows:

Average value of existing dry soup stock holding	£2.725m
Value of 55% reduction in stock	£1.5m

This foregone sum would realise the internal ROCE if released for use on other projects. Assuming at ROCE of 20% this would result in an additional return of £2.66m (present value) over a five year period. Hence, under these assumptions the value of stock improvements in total is £4.16m at a plant with an annual turnover of approx. £25m.

There were no changes in predicted manpower and overtime usage until a colder winter was simulated. This is in agreement with the observation of Chapter 4, (that increase in window parameter (T) used in the fixed cycle algorithm does not affect the replenishment characteristics). The changes of manpower and overtime costs calculated in accordance with company procedures were, for the colder winter:

Line A:	10 weekends
Line B:	7 weekends
Line C:	8 weekends

Line D and the Agglomeration process did not require any overtime.

Weekends are paid on a time and a third agreement currently operating with the company. The manpower requirements shown above were predicted by the manufacturing model.

5.6.2 Minimum Run Length (MRL)

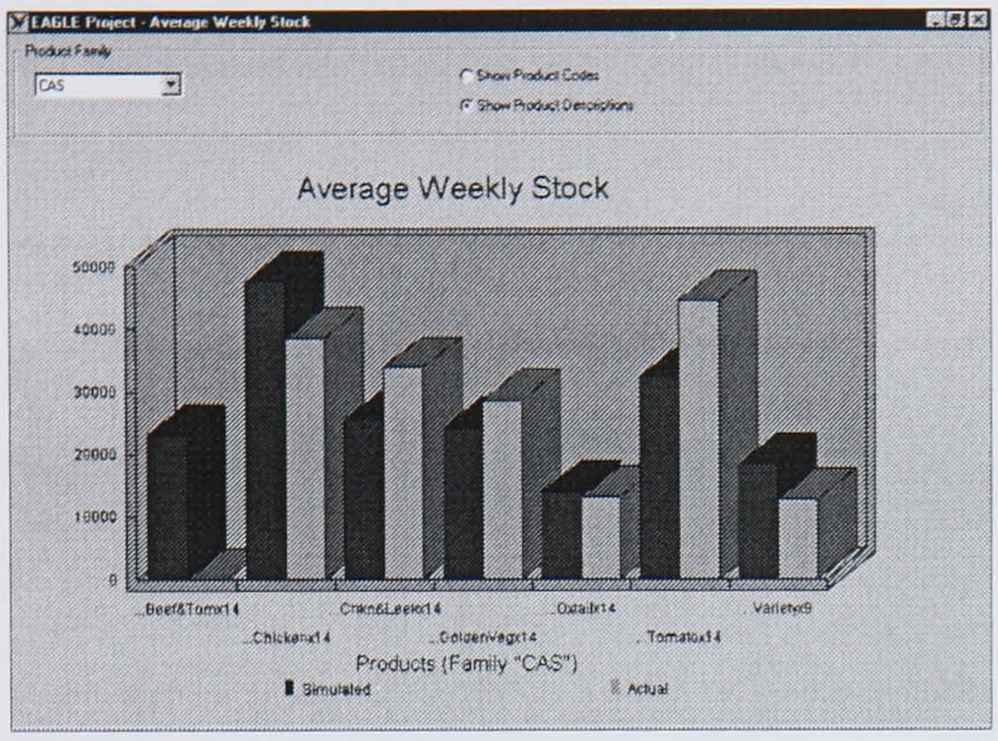
Costs and performance measure - The MRL parameter affects stock level and the performance factors listed above for the wfc parameter.

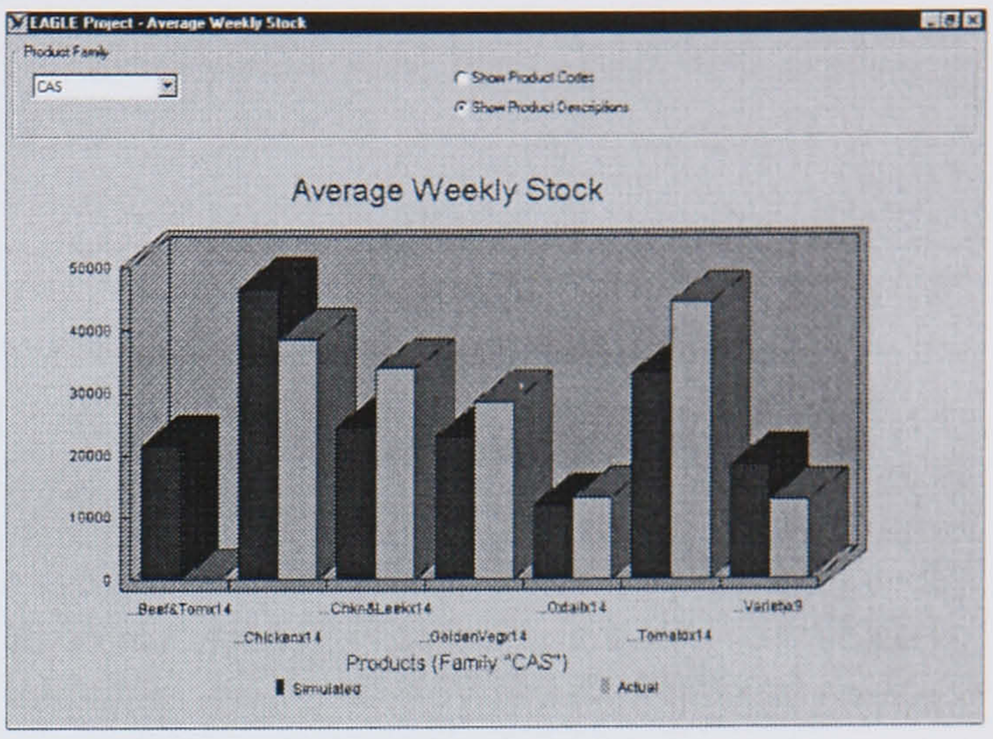
Tuning procedure - The results above were all obtained using the "as is" setting for the MRL parameter. Fig. 10 demonstrates the effect of relaxing the MRL constraint on 6wfc stock levels. The average stock level for all medium and high level demand products were virtually unchanged. This is reasonable since demand was substantially higher than the MRL parameter for these products. The stock level on low demand products did fall as expected. The switch-overs and short production run lengths did increase. The number of switch-overs increased by 70%.

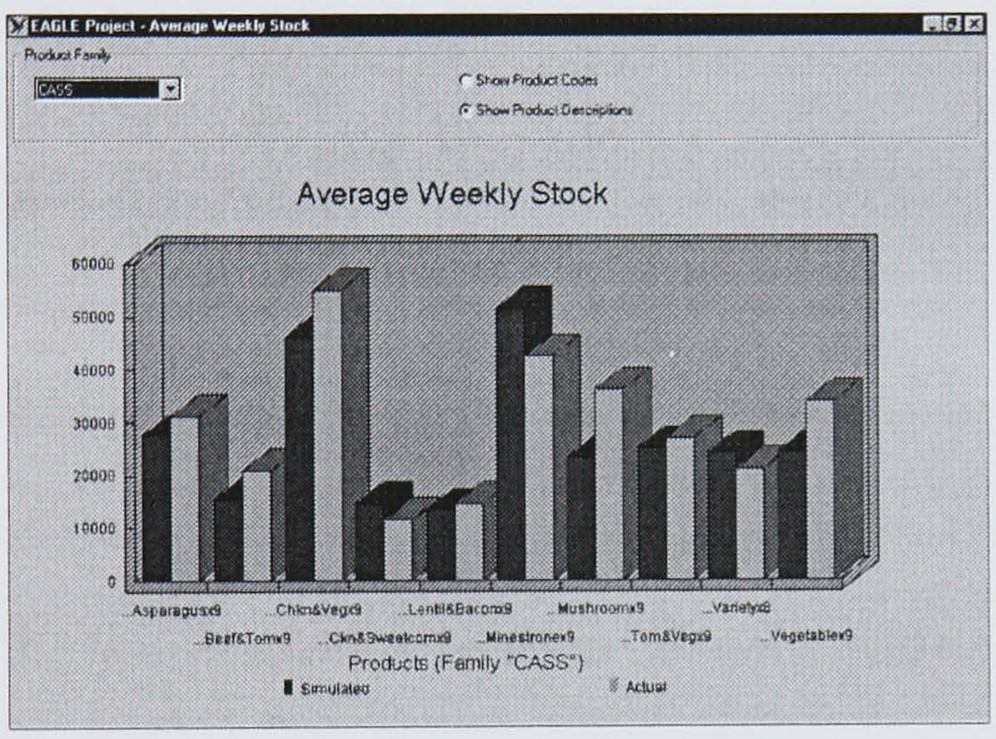
Discussions and Conclusions - The current MRL parameter settings seem reasonable. They remove small run lengths from the schedule at only a modest increase in stock level. They have little effect on key performance indicators.

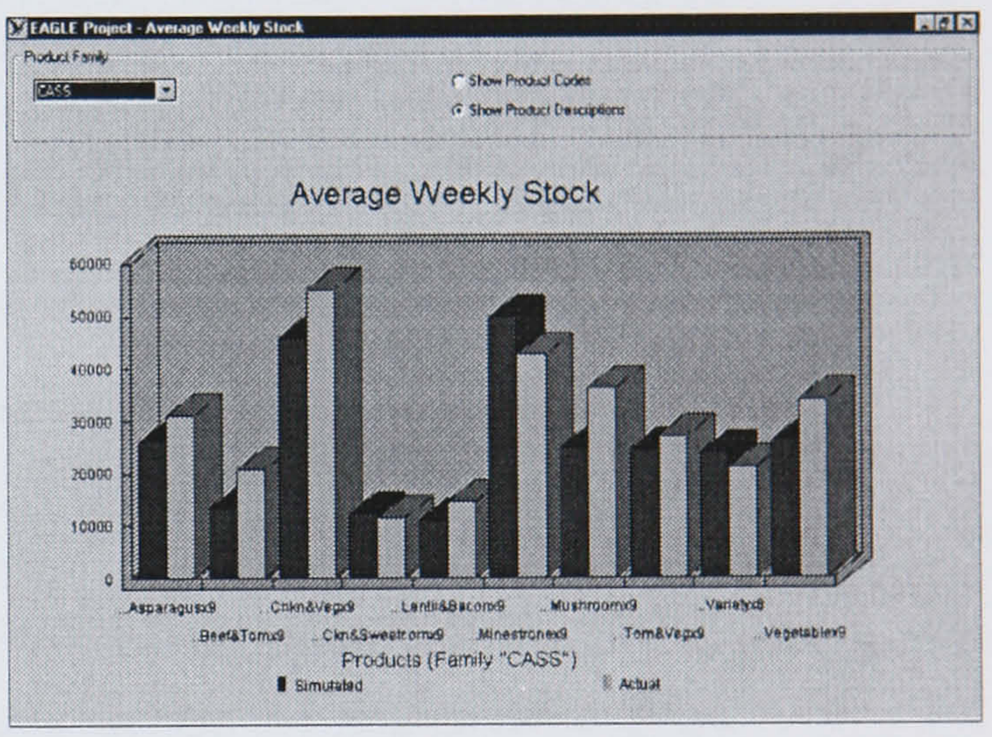
6wfc with MOQ

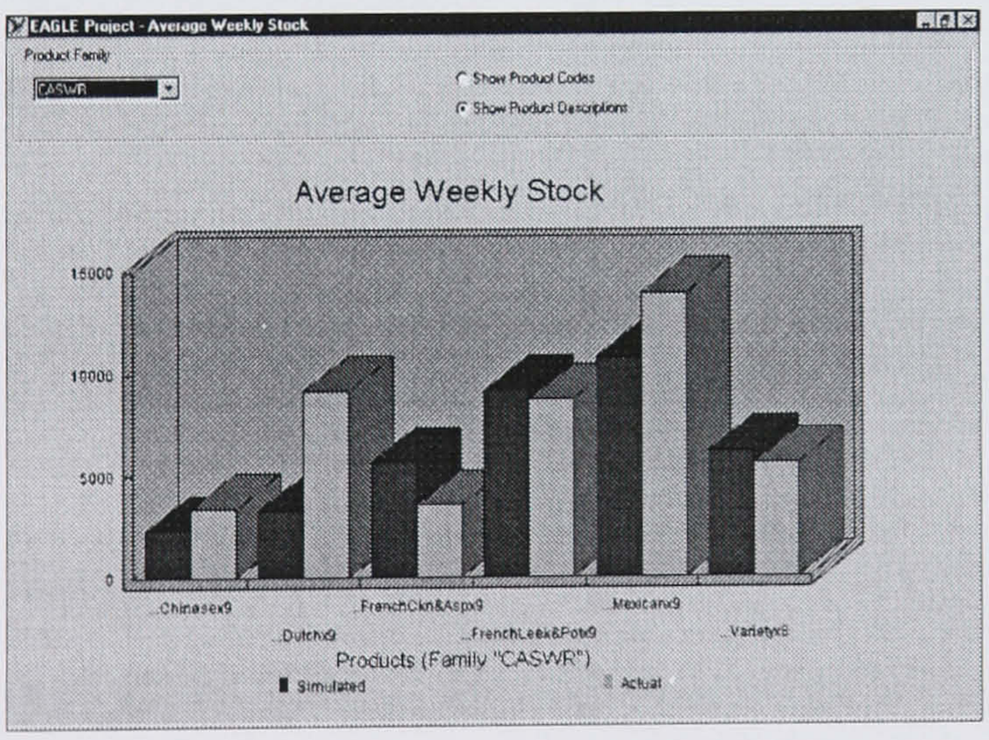
6wfc no MOQ











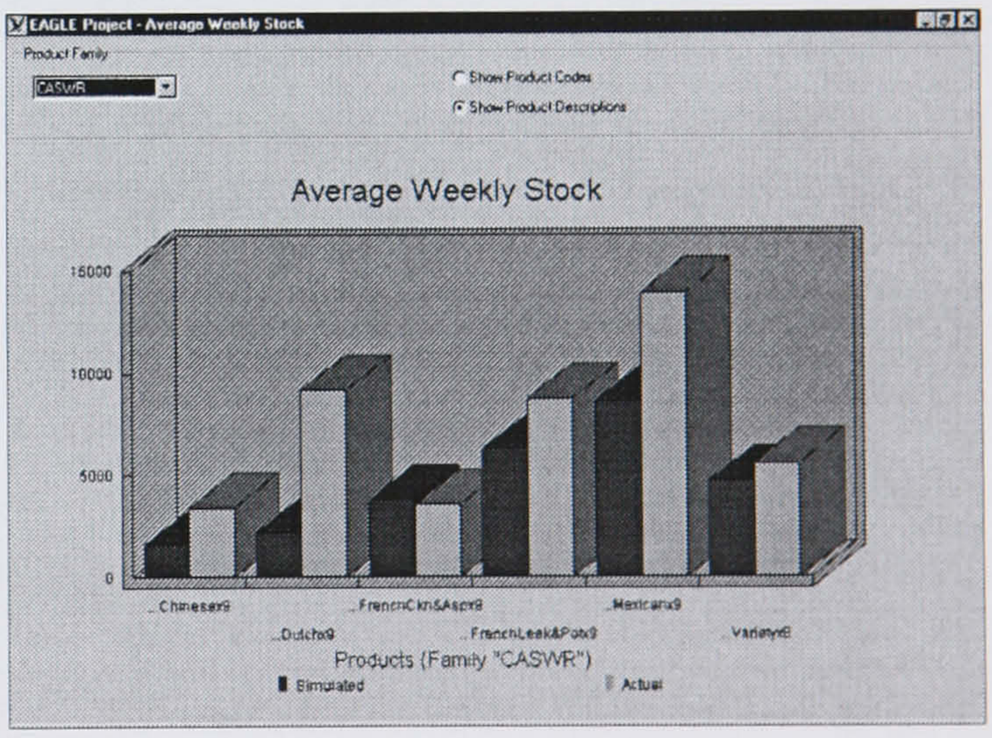


Fig. 10

5.6.3 Safety Stock and Forecast Accuracy

Costs and performance measures - These include all of the factors listed under the wfc parameter tuning section and further consideration of customer service level. Clearly, the safety stock parameters must be chosen for each separate SKU, in order to ensure that customer service is maintained at the warehouse.

Tuning procedure - Since neither daily data nor promotional information was recorded, it was not possible to develop a forecasting procedure along the lines followed for the retailer in Chapter 2, and hence a simple forecasting procedure, based on exponential smoothing has been used. An examination of the forecasting errors suggested that the variance of the demand was proportional to the estimated mean. We therefore used the procedure described in chapter 2 in order to estimate the constant of proportionality in the form of a dispersion coefficient θ. The proportionality of forecast error to demand was also observed by Brown [82] who reports that the proportionality remains constant in a manufacturing concern over 30 years, but varied from manufacturer to manufacturer. Values in the region $\theta = 40 - 50$ were obtained over all SKUs. These were used in the replenishment stock calculation of equation 5.0 of Chapter 4.

5.6.4 Results and Conclusion

The safety stock levels varied with wfc and demand level. Expressed both in absolute terms and in percentage terms, they were:

8 wfc (Safety Stock)

(i) High Demand Products

Cases(%) - 8000 (30%)

(ii) Low Demand Products

Cases(%) - 4700 (53%)

2 wfc (Safety Stock)

(i) High Demand Products

Cases(%) - 4400 (55%)

(ii) Low Demand Products

Cases(%) - 2300 (100%)

- The percentage figures are not linear due to the form of the calculation of Safety Stock.

- Reduction of the wfc figure from 8 to 2 result in about a 50% reduction in SS for all products. At a setting of 6 wfc about 25% of the total stock maybe attributed to SS. At 2 wfc about 35% is attributable to SS. If an optimum forecast could be developed whose accuracy compares to that obtained at the retail RDC, then from equation 8.0 of Chapter 4, the safety stock factor could be reduced by an order.

References

Greycon Ltd, Battersea, London, UK. 1997. *User Guide S-Plan 4 for Windows V 1.21*

Brown, R.G. 1982. *Advanced Service Parts Inventory.* Materials Management Systems, Inc.

6.0 VEHICLE LOADING AND DISTRIBUTION

6.1 Introduction

It is estimated that almost half the lorries on UK roads are empty and carrying nothing but fresh air. The largest category of lorries on our roads are associated with the food sector, who in 1995 accounted for 37.5 billion tonne kilometres [The Department of Transport]. Many bulk food sector lorries carry goods packed in pallets. The cost of delivery of a single pallet in the food sector typically breaks down to distribution (30-40%), stock holding (15-25%), warehouse overheads (15-20%) and warehouse handling (15-25%) [Unilever]. There are, therefore, both commercial and environmental reasons for attempting to improve vehicle distribution efficiencies.

This chapter demonstrates by means of case studies how short term forecasting can be exploited to obtain improved distribution efficiency. The case studies are based on distribution data for a manufacturing warehouse delivering to a retail RDC. In this system, orders from many customers were received daily within the warehouse planning system and delivery made within a few days. Orders from RDCs usually contained a range of products. Precise delivery times at the RDC were agreed by negotiation. The resultant order OTIF constraints had to be observed on delivery and, therefore, viable routes and lorry loads were also constrained. In the case study a modern route planning package was employed and lorry load efficiencies of approximately 75% were attained.

6.2 OTIF Relaxation

Lorry load efficiencies are usually more difficult to maintain under 'continuous replenishment' policies since these result in decreased order size, more frequent deliveries and reduced lead times. Some improvement in lorry load efficiencies are possible if OTIF constraints are relaxed. These relaxations usually take the

form of allowing agreed flexibility in order size (say ±5%). Such 'round up' policies represent a relaxation of the 'In-Full' half of the OTIF constraints.

An alternative option is that of relaxing the 'On-Time' (OT) half of the OTIF constraint. We propose a Forecast Based Order Restructuring (FBOR) approach in which orders are 'pulled forward' in time for early delivery at the RDC. Items pulled forward spend longer in stock and, therefore, increase the average stock level. It is proposed that items from the next RDC lorry schedule are pulled forward to improve the current RDC lorry load efficiency. This policy minimises the increase in stock residence time and average stock level.

The increase in stock residence time under FBOR is equal to time to the next RDC delivery. Items pulled forward which are not in the next RDC delivery will be associated with later deliveries and will therefore experience larger increases in stock residence time and will therefore result in an increase in average stock level.

Fig. 1 illustrates the increase in average stock incurred under FBOR policies and round up policies.

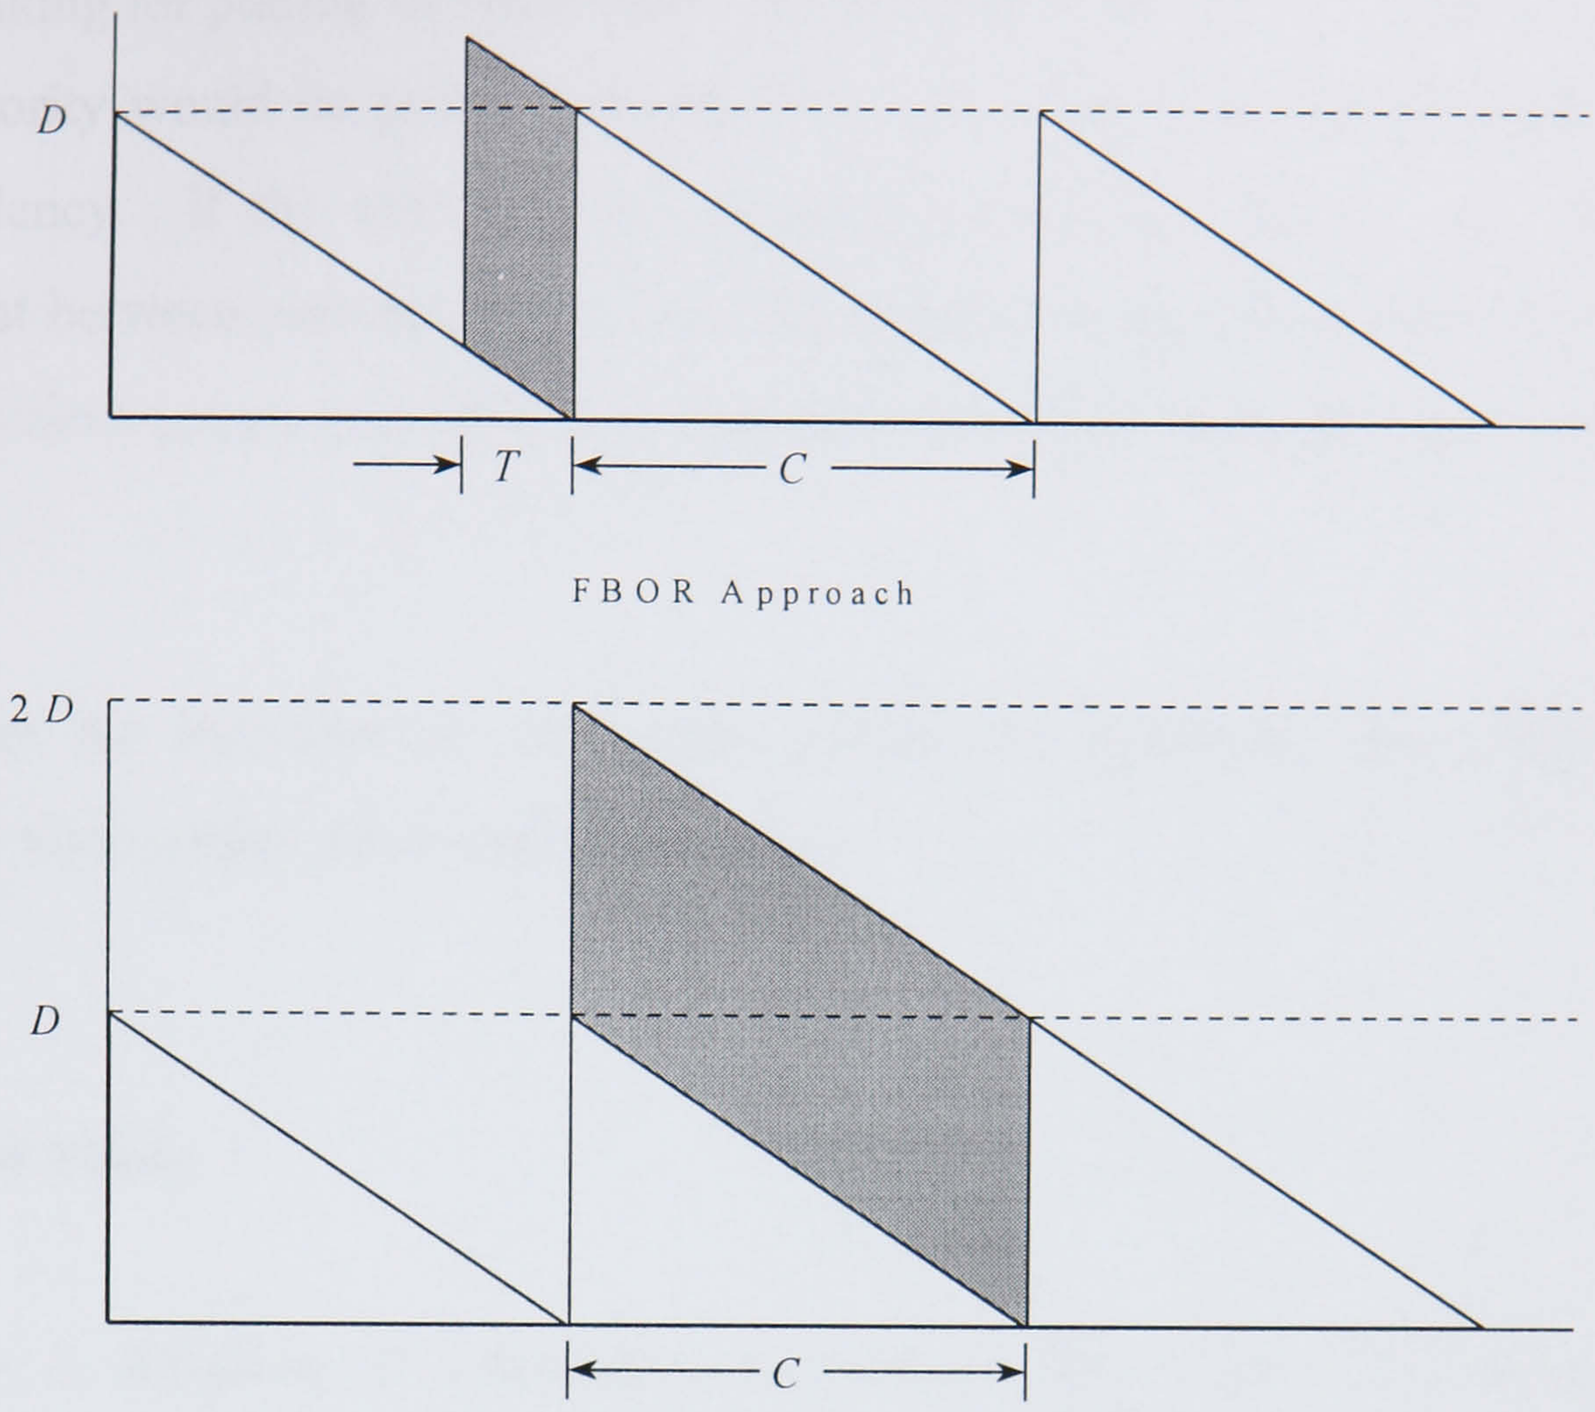


Fig. 1

The increased residence time for a round up policy equals the replenishment cycle time (C). This is always larger than the time between deliveries (T). For the case study T was slightly in excess of 1 day whereas C could typically range from 5-20 days, depending on demand and batch constraints.

The lead time for the case study system was nominally 4 days. The time between RDC deliveries was slightly in excess of 1 day. Therefore, the schedule for the next RDC lorry load was resident within the manufacturing system and no load forecast was necessary. Under Vendor Managed Inventory (VMI), short term forecasts can be used to generate orders for pulling forward. The method of allocating replenishment priorities in the Synchronous Replenishment System

described in the previous chapter and chapter 4 can equally be used to produce a priority ranking for pulling forward orders for delivery to an RDC. The orders of highest priority would be pulled forward first, thus minimising the increase in stock residency. If the next schedule is not available and there is no VMI arrangement between partners, then, if the necessary information is shared (TIE. SID etc.), orders can be generated from current book-stock levels and short term forecasts.

To illustrate the increases in lorry load efficiencies attainable from FBOR approaches we consider some case study results.

6.3 Case Study

Database - A database of information relating to three weeks of journeys between the manufacturer's warehouse and retail RDC were recorded. The following fields were recorded for each departing lorry:

1. Date
2. Route Number
3. Destination(s)
4. Customer(s)
5. Number of stops
6. Order Number
7. Product Description
8. Number of Cases
9. Cost of Journey
10. Total Pallets on Lorry

Every journey made within the chosen period was recorded.

A sample of successive orders from the supplier to the retailer's RDC were analysed. In the sample period the average drop rate was approximately daily. This is approximately the drop rate of dry soups, which over a 6 month period had an average drop rate of 1.9 days. None of the lorry loads were constrained by weight limitations. The bulk of the loads for this sample consisted of retail RDC orders with only 7% overall being concerned with independents (i.e. other smaller retail outlets). The summary information for a 3 week (15 day) sample period was:

Pallet delivered to RDC	274
No RDC Drops	14
Average Lorry Efficiency	75%

FBOR Schedule - If the orders in the system for several days ahead are allowed to be delivered ahead of time then we can maintain high lorry load efficiencies by pulling forward orders. The following results were attained.

	Solution 1	Solution 2
Pallets delivered to RDC	234	255
No RDC Drops	9	10
Lorry Efficiency	100%	98%
Reduction in RDC Drops	35%	28.5%
Average Pull Forward (days)	1.5	1.5
Maximum Pull Forward (days)	4.0	4.0

Solution 1 assumes that a delivery of 21 pallets required in the last week of the period could be combined with an order pulled forward from the week following the test period, as had been allowed for the previous two weeks.

Solution 2 does not assume this and hence a partially full lorry is inevitable in week 3. The 4 day pull forward was as a result of the fact that the retail RDC closes at weekends and some orders were pulled forward across the weekend.

As there may be reasons for clearing weekly orders and hence not allowing pull forward across a weekend, we consider two sample weeks (A and B1-B2) with this constraint. The summary information of actual results for each week was:

Week	A	B
Pallets Delivered to RDC	121	80
No Lorry Drops	7	4
Lorry Efficiency (with Independents)	80%	84%

Using FBOR the following results were achieved

Week	A	B1	B2
Pallets Delivered to RDC	121	80	80
No RDC Drops	5	4	3
Lorry Efficiency	93%	77%	100%
Reduction in RDC Drops	28%	0%	25%
Average Pull Forward (days)	0.6	1.5	1.5
Maximum Pull Forward (days)	1	2.0	2.0

In schedule A only 9 Pallets were contained in the fifth drop (without these, 100% efficiency and a drop of reduction of 37.5% is possible). Schedule B is divided in to two sets of results, in schedule B1 only two pallets were contained in the fourth drop. From a detailed look at the contents of the mixed pallets in the delivery, it seemed these two pallets could be absorbed into other partially filled pallets, hence Solution B2 is obtained.

Comment: In these solutions the pull forward periods are very constrained, due to elimination of pull forward across weekends. The results suggest that in this

case the use of independents would be advantageous to avoid near empty lorries and that the introduction of strict weekly constraints may be counter productive.

Discussions and Conclusions - Implementation of the approach would clearly involve changes in commercial practices between partners. The number of days pulled forward, of course depends on the individual product demand rates, and delivery frequency.

6.4 Retailer Benefits

The benefits to the retailer of the FBOR approach include the following:

- A significant reduction in lorries arriving at the retail RDC. At peak times, this can represent a serious bottle neck. The case study quoted above would indicate a reduction of up to 35% in lorries arriving at the RDC.

- A reduction in transportation and delivery costs.

- The RDC stock level increase is minimal

6.5 Manufacturer Benefits

- Reduced distribution costs.

- Fewer orders per week. In this example the manufacturer's warehouse handles just over a thousand orders per week. Top priority is given to the retailer's RDC orders. If the number of RDC orders were reduced by the order of 30% there would be reduction in administration.

- Improvement in quality of service and product. With multiple drop loads there is always the hazard of misplaced pallets. Mistakes of this class are extremely tedious and expensive to rectify and any commercial procedure which reduced the possibility of pallet misplacement would be of advantage.

- Increasing lorry loading efficiencies from 75% to 100% would represent a significant reduction in the number of lorries miles with consequential benefits to road congestion, pollution and the environment.

- Under Vendor Managed Inventory and pay as you use arrangements, residence time of the increased stock is minimal.

Round Up Heuristics - Lorry load efficiencies may be improved if the 'in full' constraint of OTIF is relaxed. If for instance ± 5% is allowed on any order then more efficient lorry loads might be expected. However we note:

- Constrained round up heuristics may not allow full lorry loads to be attained. For the data above, if only 5% round up is allowed then less than 4% of the 25% empty lorry space can be filled. If round up is only allowed for 'high demand' products then the improvement may be significantly less than 4%.

- If the round up percentages are increased, the potential increase in average stock at the retail RDC is correspondingly increased. If for this data, 30% round up is allowed then the full lorry loads are attained only by a 30% increase in stock.

- The book keeping associated with round up order changes (decided by the supplier after the order has been placed) are a form of overhead, to the retailer.

Pulling forward orders in the system does not appear to have the above disadvantages. For instance, the FBOR pulls forward by a time interval that is typically small compared with the stock replenishment cycle and hence involves much smoother average stock increase. The concepts and approaches used in FBOR can in principle be extended to other distribution operations including back-hauling. More basically they might be integrated with the vehicle routing procedures.

References

The Department of Transport. 1996. Transport Statistics Report – The Transport of Goods by Road in Great Britain 1995. Government Statistical Service. HMSO.

Unilever. 1997. Internal Management Report. Not Published.

7.0 DISCUSSION AND CONCLUSIONS

This research and development programme has been motivated by developments within the food sector which is generally accepted as one of the most advanced sectors in terms of supply chain development. The concepts which have been developed are thought to be applicable to supply chain development within other sectors.

We have reviewed the strategies and objectives which are widely accepted within the food sector and encapsulated within the Efficient Consumer Response movement which began in the United States. This gives high priority to the development of effective partnership based approaches to supply chain management. This work has sought to develop concepts and technology needed in realising effective partnership.

Almost the entire section of a supply chain under the control of collaborating partners can be viewed in the network of inventory replenishment systems. The replenishment loops however, may contain manufacturing, warehousing and transport distribution systems.

Within all sections of the supply chain network, forecasting play an important role and this has formed a central theme of these researches. In chapters two and three we develop effective models for forecasting at key points within a supply chain partnership network. Of first interest is the demand at the till data points since this drives the whole supply chain network. The methods developed in Chapter 3 offer very practical tools for extraction of key information from till data. The algorithms are fast and robust and capable of providing daily monitoring and forecasting for many thousands of products. They will also provide efficient exception reporting capabilities. A key idea in these developments is that of exploiting the "high frequency" data present in the daily demand patterns. These provide the fast response required in decision making within supply chains. It would seem likely that the exploitation of daily patterns would be of importance in other supply chain activities. For instance it is likely that the demand patterns say, within the Health Service or

other service sectors would be likely to exhibit very similar characteristics to those encountered at the till demand of a major retailer.

It is of regret that in spite of extensive effort, no convincing modes of weather response to information were established. There is no doubt however that some product responses change dramatically under prolonged extreme weather conditions. It would seem that in the test period selected, extreme weather conditions did not occur. The estimation of weather responses is certainly an area which is worthy of further investigation and analysis.

The forecast models seem effective along the supply chain even though the levels of demand differ by several orders. The forecasting in the manufacturing studies of Chapter 5 were rather simplistic since only weekly data was available. This short coming has now been re-addressed and recently extensive daily manufacturing demand data has been acquired. It is intended that the methods used so effectively to monitor and forecast demand at till and RDC level will be developed and applied to important area of manufacturing demand.

In Chapters 4 and 5 we considered the design of stock replenishment systems and limited our consideration to fixed cycle system, on the basis that these were mainly used in industry. The presence of significance batch constraints modified the stock level characteristic in an important way and simple formula for estimating these effects were developed. These formulae are simple to use and provide plausible estimates of the minimum stock level which can be attained in practice.

We have argued that the attainment of the ECR objectives and in particular of "synchronous replenishment" requires that the activity of warehousing and manufacturing be dynamically integrated and forecast driven. We have suggested a practical procedures for realising synchronous replenishment and have implemented a system based on this procedure as described in Chapter 5. The detailed evaluation undertaken there would indicate that there are no substantial savings with this type of integration. We have also indicated how a number of key parameters may be tuned given such a system. The essential idea is that of storing historic information in order

to run the system in the "tuning" mode, and to assess the effects of key parameter changes. From this work it would seem that the effects of these parameters can be easily understood and tuned in practice.

If then demand distribution is Poisson-like then the dispersion coefficient directly reflects the efficiency of the forecasting procedure and determines the minimum achievable level of safety stock. The tuning procedure used to evaluate the dispersion coefficient, is applicable to the whole of the supply chain, and takes in to account the efficiency or otherwise of the forecasting procedures. It is believed that such procedures would provide a valuable general tool use of which could lead to major reductions in stock level along the supply chain. Furthermore it would seem that the entire tuning routine could be built to operate in an on-line mode. This development would seem to have both commercial and research interest and be worth further investigation.

In the last chapter we have considered procedures for improving current levels of lorry loading efficiency. These were shown to yield benefits, both in terms of efficiency of distribution, and of environmental improvements. In order to realise this or similar approaches it is necessary to relax the OTIF constraint which is so widely imposed within the food sector supply chain. This seems essential as the objectives of ECR "continuous" replenishment are realised and replenishment order sizes reduce and frequent delivery increase reported in this work on research and investigations which were undertaken within the area of warehouse control. This is an area which will justify future work and the increasing importance as supply chain operations become more and more demanding.

A modern warehouse is relatively well served in terms of such areas as vehicle routing and automation of the internal operations of the warehouse. The dynamic day-to-day management of a large manufacturing warehouse or an RDC increasingly requires sophisticated scheduling capability. Such an operation involves scheduling on a day-to-day basis the departure and arrival of large number of lorries and all of the internal dynamic processing involved within the warehouse.

The key findings of this work in terms of importance to the food sector supply chain may be summarised as follows:

- *Till Data Analysis.* In any "pull-through" (demand led) system, monitoring of customer demand is a priority. The on-line procedures developed here are capable of efficiently forecasting the millions of product demands involved on a daily basis. The results would indicate that the forecasts are near optimal in the sense that the forecast error is near white noise. These techniques would seem to have considerable application within the area of Category Management and in the assessment of promotional performance. The approaches are capable of providing daily patterns of consumer demand for every product at each individual store.

- *RDC Demand Analysis.* The model developed can efficiently be used to forecast daily demand placed on the RDC with near optimum accuracy. It is also capable of rapidly detecting (1-2 days), quantifying and re-estimating demand changes due to unforeseen behaviour in, for instance, promotional uplift and unexplained changes in patterns and levels of demand. This capability is clearly important in obtaining responsive supply chain performance and maintaining high levels of customer service in a highly competitive marketplace. The responsiveness described above is due directly to the new model structure developed for this work, and the inclusion of daily response modes.

- *Weather Effects.* We have developed techniques capable of estimating the effects of numerous weather variables on demand. They are capable of analysing the effects of weather variables on the demand for individual products.

- *Safety Stock Tuning.* Procedures have been developed for modelling forecast error for any forecasting procedure and for deriving minimum stock levels required to meet a specified level of customer service. The methods are applicable to either the routine forecasting procedures widely used in industry, or to the near optimum procedures developed in this work. We have also demonstrated the significant impact of forecast accuracy on stock levels throughout the supply chain.

- *Replenishment System Analysis.* Fixed cycle system are used almost universally in the food sector. We have demonstrated that the batch flow constraints result in the development of an unavoidable limit cycles. This cycle imposes a minimum on the attainable stock level and storage capacity and therefore on the number of products that can be held in a given node of the supply chain.

- *Benchmarking.* Benchmarking is used to compare an individual company`s performance with the alleged world class standard. We have developed a formula (Chapter 4, equation 6.0) which is simple to use in practice which provides an estimate of the best attainable stock levels taking into account the characteristics of the supply chain processes including minimum order constraints, forecast accuracy, lead time and the customer service level required. The formula allows managers to investigate the effects of the various parameters concerning replenishment systems on the overall performance of the supply chain and how far their current performance is away from optimum estimates.

- *Multi-Product Synchronous Replenishment.* An approach has been developed which integrates the requirements of the of the manufacturer's warehouse with the scheduling of the manufacturing plant. This system ensures that priority products are dynamically allocated preference in the manufacturing schedule through the use of forecast times to stock out.

- *Manufacturing Decision Support System Design.* We have developed and assessed a realistic system which embodies the principles of synchronous replenishment and allows the tuning of the key supply chain parameters. The system demonstrated that gains of £2.725m were achievable from a plant with a annual turnover of approximately £25m.

- *Vehicle Loading Optimisation.* The approach developed relaxes the precious OTIF constraint in order to allow the pull-forward of orders. The resulting case study showed that near 100% lorry load efficiency is obtainable together with a 35% drop in the number of journeys made. As 30-40% of the costs of delivering a

single pallet are associated with vehicle distribution costs, these improvements represent a sizeable saving.

All of the above developments are important in enhancing partnership based operation of the supply chain.

The author has carried out modelling and assessment studies concerning the many operations within a warehouse and these would indicate that the development of dynamic planning tool is feasible. It seems that this area of work would be of considerable importance to future supply chain development.

Lastly, mention should be made of the rapidly increasing practise amongst retailers of allowing their many thousands of suppliers access to their stock and demand information. The method developed within the thesis are tailor made for exploitation of these new developments. The methods developed in Chapter 3 would allow the precise forecasting of demands for supplier products and a detailed analysis of the sales of this product within the hundreds of stores involved.

Although the developments reported here were all concerned with the food sector, it should be emphasised that, the techniques would seem to transfer usefully to supply chain operations within other industrial sectors and within some service sectors.

BIBL
LONDON
UNIV.