THE UNIVERSITY OF ASTON IN BIRMINGHAM

DIFFERENCES BETWEEN THEORETICAL PREDICTIONS AND OPERATIONAL PERFORMANCE IN A STOCK CONTROL APPLICATION

Henry James

Submitted for the degree of Doctor of Philosophy, October 1982

THE UNIVERSITY OF ASTON IN BIRMINGHAM

DIFFERENCES BETWEEN THEORETICAL PREDICTIONS AND OPERATIONAL PERFORMANCE IN A STOCK CONTROL APPLICATION

S U M M A R Y

Over the past decade, several experienced Operational Researchers have advanced the view that the theoretical aspects of model building have raced ahead of the ability of people to use them. Consequently, the impact of Operational Research on commercial organisations and the public sector is limited, and many systems fail to achieve their anticipated benefits in full.

The primary objective of this study is to examine a complex interactive Stock Control system, and identify the reasons for the differences between the theoretical expectations and the operational performance. The methodology used is to hypothesise all the possible factors which could cause a divergence between theory and practice, and to evaluate numerically the effect each of these factors has on two main control indices - Service Level and Average Stock Value. Both analytical and empirical methods are used, and simulation is employed extensively.

The factors are divided into two main categories for analysis - theoretical imperfections in the model, and the usage of the system by Buyers. No evidence could be found in the literature of any previous attempts to place the differences between theory and practice in a system in quantitative perspective nor, more specifically, to study the effects of Buyer/computer interaction in a Stock Control system.

The study reveals that, in general, the human factors influencing performance are of a much higher order of magnitude than the theoretical factors, thus providing objective evidence to support the original premise. The most important finding is that, by judicious intervention into an automatic stock control algorithm, it is possible for Buyers to produce results which not only attain but surpass the algorithmic predictions. However, the complexity and behavioural recalcitrance of these systems are such that an innately numerate, enquiring type of Buyer needs to be inducted to realise the performance potential of the overall man/computer system.

∿ ∿ ∿ ∿ ∿ ∿ ∿ ∿ ∿

Submitted by HENRY JAMES, 1982

for the degree of DOCTOR OF PHILOSOPHY

Key Words : STOCK CONTROL; SYSTEM PERFORMANCE; IMPLEMENTATION

CONTENTS

PAGE

ACKNOWLEDGEMENTS

I wish to thank Mr. J.B. Kidd (Internal Supervisor), Professor C.D. Lewis (Associate Internal Supervisor) and Mr. G. Barnes (External Supervisor) for their help and guidance given throughout the project. I am also indebted to Mr. R.M. Hill of the University of Exeter for many stimulating discussions. Finally, I am grateful to Miss Janice Chater for typing a difficult script with so few errors, and to Mrs. Elsa Crawshaw for inserting the mathematical expressions.

List of Tables

List of Tables (Cont'd)

List of Tables (Cont'd)

List of Figures

List of Figures (Cont'd)

List of Figures (Cont'd)

List of Figures (Cont'd)

List of Figures (Cont'd)

"Call a man a machine if you want to, but don't underestimate him when you come to do experiments on him. He's a non-linear machine; a machine that's programmed with a tape you can't find; a machine that continually changes its programming without telling you; a machine that seems to be especially subject to the perturbations of random noise; a machine that thinks, has attitudes, and emotions; a machine that may try to outguess you in your attempts to find out what makes him tick, an effort in which, unfortunately, he is sometimes successful."

A. CHAPANIS (1)

CHAPTER 1

INTRODUCTION

1.1 DETERMINATE AND INDETERMINATE SYSTEMS

Computer applications which have been implemented in industrial, commercial and public organisations can be usefully classified according to their determinacy. Systems with highly structured inputs, relatively invariable procedures and predictable outputs can be said to possess a high degree of determinacy. Conversely, systems with inputs of a random or probabilistic nature, flexible procedures, and outputs which are difficult to predict can be said to be indeterminate.

The concepts of determinacy and probability are closely related. In determinate systems the probability of correctly predicting the outcome is high, whereas in indeterminate systems it is low. In determinate systems the Operators apply highly programmed actions to predictable inputs and the task is one of straightforward information handling; in indeterminate systems the Operators must be allowed to vary their actions contingent upon the input patterns and thereby exercise decision-making functions. Ashby (2) has proposed a 'Law of Requisite Variety' which asserts that in order to maintain a system in a steady state, the Operators (or, more generally, 'controlling forces') must possess a sufficient repertoire of controlling actions to counteract any environmental disturbances acting on the system. Hence if the inputs possess stochastic properties, the Operators must exercise a commensurate variety of responses to achieve the desired goals.

Indeterminate systems have both positive and negative aspects. On the one hand the system has sufficient intrinsic flexibility to adapt to unpredicted or previously unexperienced environments and competitive systems. On the other hand, the necessity to select from a large number of options requires a much higher grade of

Operator and an increased risk of Operator error. Some responses can be programmed in advance of system operation, but many are contingent upon the Operators' perception of a stream of unstructured, unpatterned and sometimes irrelevant inputs.

It is clear that indeterminate systems are, per se, complex. It is also intuitively obvious that complex systems require heightened perceptual and cognitive skills from the Operators. Indeed, the Operators may be pushed towards, and perhaps beyond, their limits of efficient performance. Although the advantages of automation are well established in the cases of psychomotor functions, they are much more dubious for cognitive functions. Certainly, at the present time there is no question of completely replacing the human cognitive function with computers. However, man is far from being an **optimal** decision-maker and computer-aiding can be a significant help.

The majority of computer systems which have been implemented to date are characterised by a high degree of determinacy. Systems which effect the basic book-keeping functions of a Company are typical examples. Here the manual system has been replaced by an automated equivalent, usually with a 'spin-off' of improved management information to justify the computer. The computer often supplants clerical operations, and the system runs in a highly programmed manner.

Some computer systems, however, have a much greater degree of indeterminacy. Typical examples are Stock Control, Production Scheduling, and Vehicle Route Planning. In those systems, environmental forces acting on the system are much more difficult to predict. In most of the applications, Operational Research techniques are invoked. Often, an algorithm is embedded in the system to produce an 'optimum' solution. However, the quest for an optimum solution often results in simplifying the model to a point where it fails to incorporate all of the decision variables, and, as a result, attaches undue weight to the remaining easily measured variables. Hence, depending on the nature and magnitude of the intangibles, the algorithm may be addressing only a sub-set of the problem, and the acceptance of the solution may be invalid.

Heuristic methods have been tried in place of optimisation techniques. Some examples will be described in Section 1.2. Here the computer adapts its processing routines to seek a solution within the confines of general decision rules. This is intuitively commendable but at the present time the potential of the method has not yet been realised in commercial software.

The present research is based upon one particular example of an indeterminate system - a complex stock control application where the buyers interact with an automatic algorithm. The demand and lead time inputs are of a probabilistic nature, and the algorithm is regularly overridden in order to gain commercial advantage from opportunistic purchasing. It is not claimed that this system is representative of all indeterminate computer systems, but it is reasonable to suppose that many of the findings aprópos man/computer interaction will be relevant to other application areas.

1.2 RESEARCH AREA

The research is aimed at identifying and placing in perspective the factors contributing to the divergence between the predicted and achieved performance in a stock control application, and to proceed from there to advance the present state of knowledge in man/computer system design in this type of indeterminate system. The long-term objective of such research is to work towards a symbiotic partnership between man and computer whereby they interact to attain a better overall performance than either could achieve independently.

The concept of man/computer symbiosis (3) is based on the assertion that man has something valuable to contribute on-line to problem-solving tasks which a computer cannot do automatically. The notion is not applicable to problems with a small well-defined easily quantifiable set of variables for which a true optimum can be calculated mathematically. It is relevant to the much more common case where the Operators have access to facts and opinions not available to the computer.

The author could find no evidence in the literature on Stock Control or Human Factors that the objectives have been researched before. Specifically, no recorded cases were found of buyers interacting with an automatic stock control algorithm. Singleton, in a private communication with the author, confirms that "there tends to be little (literature) between the two extremes of 'how to do it' management articles and the abstract approach to the use of computers for mathematical modelling". A few boundary areas of research into man/computer interaction are described briefly to put this central aspect of the present research into context.

Heuristic Methods

Heuristics are 'rules-of-thumb' which obviate the need to explore every combinatorial possibility when solving a complex problem. The principle is akin to Simon's (4) concept of 'satisficing', i.e. obtaining a 'good enough' rather than optimal result. Fuller (5) suggests that heuristic methods are useful for two general types of problem:-

i) Those which are too large to manipulate with traditional O.R. models.

ii) Those which are too ill-defined or loosely-structured to be expressed in precise mathematical terms.

A thorough search through O.R. literature written during the past decade has revealed that the most commonly recorded applications of this technique are in the areas of production scheduling, facilities siting, vehicle/travelling salesman routing, and, to a lesser extent, stock control.

A number of heuristic programs have been written for production scheduling. Haider et al (6) have compared the performance of one of the reportedly more effective routines (the 'slack/remaining operations' heuristic) with computer-assisted human schedulers. The objective of the exercise was to minimise the average tardiness of jobs processed through a job shop in a given time period. It was concluded that interactive man/computer scheduling (without an algorithm) is usually significantly superior to an automatic heuristic. It is important to recognise that the heuristic was here left to run without human intervention, i.e. no attempt was made to combine the advantages of the human and heuristic.

Facilities siting applications are typified by an experiment described by Schneider (7) and by Schneider et al (8). Here the problem was to site 5 transportation centres (described as ambulance depots in one of the papers) from 77 possible locations (nodes) such that:-

a) the aggregate travelling time to all of the source points was below a specified figure, and

b) the longest individual time was also below a specified figure.

Two separate heuristics were used to meet the respective objectives, each operating by evaluating an initial set of sitings, then shifting the sites to each adjacent node and re-evaluating the times. It was found that the heuristics provided slightly better solutions than the best human analyst operating by trial-and-error methods, and substantially better than the worst human analyst. Further, the computation time of the heuristics was significantly reduced by allowing the human analyst to select potentially fertile search areas at starting and intermediate junctures. This involved an appreciable amount of interaction between the analyst and heuristic.

The effectiveness of heuristic methods for solving the classic 'Travelling Salesman' problem has been investigated by Michie et al (9). They compared three methods of finding the shortest route between a number of call points:-

i) a human equipped with paper, pencil and rubber

ii) a computer using one of the best available heuristics

iii) a human equipped with a computer-driven display and light pen and a program for evaluating trial solutions (with permissible back-tracking).

The results indicated that the human/computer interactive method produced an average improvement of 6% over the 'paper and pencil' method, and the heuristic produced a further improvement of 2%. However, a particularly gifted scheduler consistently outperformed the heuristic. Moreover, the heuristic was estimated to be 8 times more costly to administer than the interactive method, though this could be reduced by a human scheduler guiding the search.

Heuristic methods have been used in stock control applications for determining operating policies and ordering parameter values. Naddor (10), for example, proposes heuristic decision rules for calculating the ordering parameters in an s,S policy without knowledge of the probability distribution for demand in lead time. These produce near-optimal results which often become optimal when the figures are rounded. In a more extensive study, the same author (11) produces ordering parameters for reorder level, reorder cycle and s,S policies using heuristic methods, again with near-optimal results. Schwarz (12) uses a heuristic to determine the stocking policy which minimises average system costs. Wagner et al (13) use heuristic rules derived by empirical methods to determine the ordering parameters in an s,S policy; and Freeland (14) achieves the same objective using a different method. These and other studies have consistently succeeded in obtaining near-optimal operating policies and ordering parameter values by heuristic methods. The author was unable to find any evidence, however, of attempts to produce stock control heuristics or other algorithms which interact with buyers or other human Operators.

Man-machine Interface Design

There is a plethora of literature on this area, much of which is addressed at the notion of a single Operator interacting with a machine at an ergonomic level. The typical characteristics are:-

a) The Operator senses information through displays and he reacts by activating controls.

b) The Operator is on-line to the machine.

c) The human functions employed are mainly perceptual and psychomotor.

d) The responses are near-instantaneous.

Superficially this has little in common with the integration of human cognitive functions with mathematical algorithms. However, there are certain principles which are common to both situations. This will be explored in Chapter 2.

Participative Systems Design

The notion of a socio-technical system was first developed at the Tavistock Institute. It is based on the premise that any man/machine system contains both a technology and a social structure linking the Operators with the technology and with each other. Total systems design should therefore encompass the human needs of the Operators as well as the technical and economic factors.

Participative systems design, as described by Mumford (15), is a practical application of this philosophy which involves handing over some of the design responsibilities to the employees who will eventually operate the system. This approach is centred around the philosophy of protecting human values in an age of inexorable automation. It is argued that this results in improved system performance due to improved motivation. Farrow (16), however, in a study for the European Social Fund, reports that "Evidence from the visits (to small and medium-sized Companies) was not strong enough to indicate whether or not there is a relationship between the extent of consultation and the success or failure of computer-based systems."

The aims of this approach are complementary to the present research. The present work does assume the willing and enthusiastic cooperation of the Operators - an assumption which was perceived to be generally valid.

Aids to Decision Making

Several techniques and methodologies have been developed for utilising the computer to facilitate decision-making. Most of these apply mainly to strategic decisions rather than the recurrent type of operational decision with which a buyer is normally confronted. For this reason, a thorough review of decision-making aids is not appropriate here, but a few are referred to briefly to give the flavour of the man/computer roles.

Statistical Decision Theory, described amongst others by Ackoff and Saseini (17), is a technique which evaluates the expected monetory values arising from the decision paths open to the decision-maker. The structuring of the problem and the assignment of probabilities and cost estimates is carried out by the decision-maker: the computer is used to evaluate the monetory returns from each path.

Cognitive Mapping, described by Eden, Jones and Sims (18) and based on earlier work by Kelly (19), assists individuals in the understanding of their own subjective interpretations of decision situations. The mapping process consists of modelling the subject's personal construction of the situation, first in linguistic terms then in matrix form. The computer is used for path analysis to highlight illogicalities in the subject's thought processes.

Analysis of Options, described by Radford (20), is a practical approach to the Theory of Games which was pioneered by Von Neumann and Morgenstern (21). Here the decision situation is not treated as static, but as one involving conflict between the interested parties, all of whom must be satisfied with the outcome before a stable solution is possible. The computer is used for trying out combinations of acceptable participant postures to search for stable resolutions, each of which would form the basis of a possible final solution.

It can be seen that in all cases the modus operandi is for the decision-maker to structure a complex multi-variable problem and invoke the computer for what is essentially a clarification function. The decision-maker then uses the output to form a judgement and act accordingly.

1.3 RESEARCH INTO IMPLEMENTATION

This research is directed towards the implementation and operational problems associated with a particular stock control system. Whereas it is believed that the theoretical basis of the system embodies some of the more advanced concepts found in stock control literature, no attempt is made to compare the merits of the theory employed with alternative models. It is claimed only that the system is an example of an advanced stock control system which is not badly designed nor unduly difficult to operate vis-à-vis other systems which fulfill a similar function. Hence the problems which will be uncovered are not an exaggeration of those which are likely to exist (probably undetected) in other such systems. The emphasis on operational performance means that the study begins where most studies on stock control end i.e. after the design and development stage.

Much of the work on stock control and other operational research models has been directed towards the technical and theoretical considerations. Implementation studies have not kept pace with the theoretical developments, and consequently there has been an increasing amount of disquieting evidence that OR is falling into disrepute because of what Schultz and Slevin (22) call an "implementation gap" between the designers and users. Other experienced OR practitioners have drawn attention to this problem. For example, Ackoff (23) opines that this is one of the main causes of failed projects. In an examination of 48 projects he was unable to find one which failed for technical reasons. Bonder (24) warns that "conscious efforts should be initiated to develop an empirically-based operational research.......it is time to stop the continual controversy between OR practitioners and mathematical theorists regarding the practice of OR." In a similar vein, Grayson (25) states that "Management science has grown so remote from, and unmindful of, the conditions of 'live' management that it has abdicated its usability"; Powell (26) considers that "little attention has been paid by OR/MS researchers to implementations of findings"; and Hildebrandt (27) opines that in the practical application of OR, "model construction and model implementation must be considered to be the main bottleneck."

In spite of these misgivings, the volume of literature on implementation has accelerated over the past decade. Wysocki (28) compiled a bibliography of 276 papers on OR implementation in 1979, 200 of which were written in the previous six years. Around 30 of the readily-available papers were reviewed by the present author. It was found that they almost invariably view the problems in a sociological perspective with little or no objective evidence to support the findings. The need for strong management support, user involvement in the design and implementation processes, explicit goals and objectives, etc. punctuate most of the articles.

Ginzberg (29) divides the literature into 'normative' and 'empirical' categories. The normative literature consists primarily of the reflections of researchers with experience in particular fields: the empirical literature contains, in the main, factor analyses of the variables which affect the success of implementation. From the latter body of literature, Ginzberg isolates 140 different factors which could influence the performance of operational systems. Factor analysis clearly has a potentially objective foundation (depending on how the factors are defined and measured), but Ginzberg's aggregation exercise is of doubtful value as the situation-dependent variables are not really additive.

The present research is of an empirical nature, but the numerical results are strictly contextual, i.e. they are applicable to stock control systems only. Having quantified the important factors in the implementation process, some of these will be seen to apply to other application areas, but their relative importance cannot be guaranteed to remain the same.

1.4 PERFORMANCE OF MAN/COMPUTER SYSTEMS

In the author's experience of designing and implementing computer-based systems over the past sixteen years, and of observing numerous other implementations, the incidence of outright failure is extremely rare, i.e. developed systems nearly always attain and maintain an operational state. This definition of failure is, however, overgenerous to the Designer, and Meister (30) reasonably asserts that "inability to satisfy system performance criteria adequately represents partial or complete system failure." Judged according to this criterion, Parkin (31) reports that estimates of failure are as high as 70%.

As the performance criteria of commercial systems are often couched in qualitative terms (e.g. "better information") the evaluation of success or failure must by highly subjective. Retrospective surveys of system implementations are also distorted by the tendency of participants to present a picture which minimises personal negligence. Any attempts to analyse past failed projects through questioning is therefore of doubtful objectivity.

It is reasonable to assume that the failure rate for indeterminate systems such as stock control will be higher than that for determinate systems, due to their higher levels of complexity. It can therefore be concluded that if 'reasonable' (i.e. not too stringent and not too lax) quantitative acceptance criteria were applied to an indeterminate system, the probability of 'success' would be less than 30%.

1.5 RESEARCH METHODS

Sinaiko and Buckley (32) argue that "Man-machine system evaluation is one of the hardest kinds of experimentation." This is because of the integration of human Operators, who obey the principles of biology, with machines, which obey the principles of physics. The experimental difficulties are therefore caused mainly by the unpredictability of the Operators. In the present research, some of this difficulty is removed by concentrating on the effects of Operator actions and accepting their reasons at face value. No attempt is made to probe the inner recesses of the human mind to search for subconscious or ulterior reasons for these actions.

The orientation towards actions favours experimentation as the principal research tool and a predominantly numerate approach. Historically, the 'experimental method' has been found to be the most fertile way of conducting research, particularly in the physical sciences. The main advantages over observational methods are:-

a) The Experimentor can determine relationships between variables in complex systems by fixing the state of other variables. This technique has been used extensively in formulating the laws of physics.

b) Conditions can be constructed precisely and the outcomes hypothesised and tested.

c) Extraneous factors, which are often experienced as noise in operational systems, can be controlled or eliminated. The internal dynamics of the system can then be studied free from environmental interactions.

d) Conditions can be reproduced for verification purposes.

e) The results are not subject to the reporting bias commonly experienced in opinion surveys. The objectivity is marred only by the perceptions of the Experimenter.

f) Extreme conditions which may not occur in a finite observational period can be tested.

The experimental approach is supplemented by questioning the Operators where the reasons for their actions are not apparent. This was conducted on an informal basis wherever possible so that commercial exigencies were seen to take precedence over academic studies.

The experimental work was, in the main, limited to a single system, viz. the stock control application described in Chapter 3. The reason for this limitation was one of availability - the author was intensively involved in the design and implementation of the system and he was responsible for its technical performance over most of the research period. This afforded ample opportunities for observation and research. This concentration of effort on a single system clearly had both positive and negative aspects which need no explanation. In retrospect the positive aspects are considered to predominate, as complex systems, ipso facto, require a depth of understanding which is inconsistent with a diffused study. It is considered that the system behaviour, which was sometimes counter-intuitive, would never have been adequately explained without the substantial amount of post-implementation experimental work which was carried out as part of the research - an important conclusion in itself.

The performance of the stock control system is measured by two Control Indices - Service Level and Average Stock. These were selected for the following reasons:-

a) They provide a sufficient joint criterion for evaluating stock-holding efficiency, i.e. the cost of funding a given service capability.

b) Once defined they are completely objective and they can be measured accurately without elaborate monitoring systems.

c) They are amenable to aggregation across products by conversion to monetory values.

d) They are 'close' attributes of a stock control system, i.e. there is a short chain between cause and effect. With more distant attributes such as profitability or return on investment, when the index value changes it may be difficult to ascribe that proportion of the change which is due to the stock control system and that which is due to other factors.

Wherever possible, the effects upon system performance of the phenomena investigated are expressed in quantitative terms. The view taken is that, whilst quantification is often difficult, unless it is attempted, the end product of this research will merely add to the body of 'how to do it' literature which is based upon subjective reasoning. At best this would fail to give proper perspective to the issues investigated, and at worst Forrester (33) has demonstrated that intuition can be totally misleading when analysing the behaviour of non-linear systems.

Because of the concentration of empirical work on to a single system, the thesis contains some stock control theory which has been taken further than the standard textbook treatment, in which unjustifiable simplifying assumptions are often made (e.g. that demand and lead time are uncorrelated). This is, however, only a 'spin-off' from the work - the main aim is to gain an understanding of why sound theoretical concepts do not often transfer well to the operational situation. A comparative analysis of mathematical models is not therefore appropriate. Savants of stock control theory will be well aware that several such models exist, each representing an idealised solution to a real problem. Experience has shown that a full implementation requires more than an idealised model and the actual formulations must be studied to appreciate the true nature of the system. This specificity, however, is not particularly relevant to the dynamics of the system, and there is little danger in applying the outcomes of the study to other types of stock control system.

1.6 EVOLUTION OF THE PROJECT

As with all research projects, a balance has to be struck between providing answers to preconceived questions, and taking full advantage of opportunities which present themselves during the course of the investigations. Clearly, a rigid adherence to the original conception may well stifle explorations into potentially fertile paths which are not in focus at the inception of the project. On the other hand, a deviation along every path which appears promising could result in an end product which has little relevance to the problem which instigated the research. The judgement of the researcher is here paramount.

The end product of the research was originally foreseen as a taxonomy of elements which cause differences between theory and practice, expressed in quantitative terms. Fig. 1.1 (based on a schema by Buzan (34)) illustrates these elements in a tree-like structure. It is worth pointing out that the theoretical differences identified in Fig. 1.1 do not necessarily imply design faults, e.g. if any theoretical probability distribution is chosen, however good the choice there must be some differences between it and the observed distributions.

Several of the branches in Fig. 1.1 were investigated concurrently and early results suggested that:-

a) many of the factors (especially the 'External') are too distant from the Control Indices to be evaluated quantitatively,

b) the 'Human Interactions' factors are of a much higher order of magnitude than the others, and these are strongly influenced by the indeterminate nature of the system.

c) in some instances it is possible to improve upon the

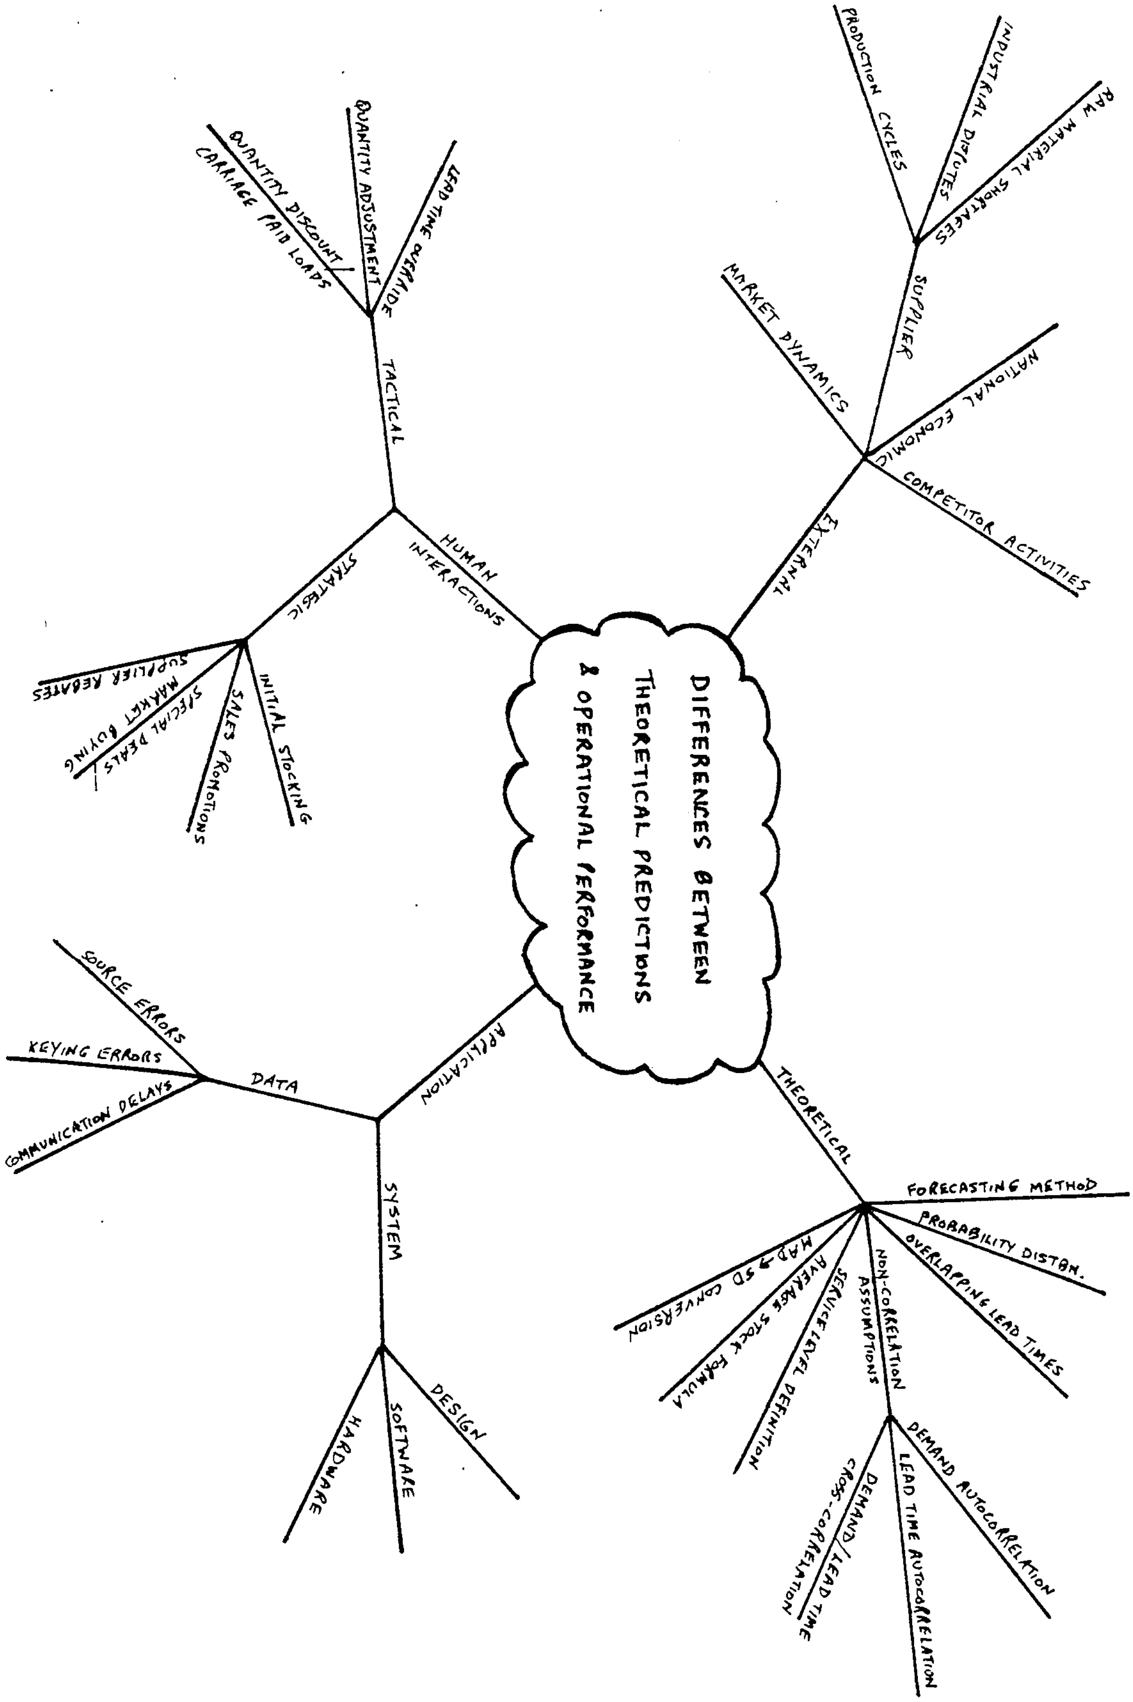

Fig. 1.1 Potential Differences Between Theoretical Predictions and Operational Performance

performance of the model when it is transferred to the 'real world' situation. Hence there will be positive elements to balance the dysfunctional elements, and

d) the influence of several of the 'Application' factors is negligible, e.g. to date there has been no instance of a hardware failure of more than 10 minutes duration (the switching time to a standby computer). In these cases there is nothing to evaluate, so they are not referred to further.

In view of these preliminary findings, it was felt justifiable to re-orient the project towards the central theme of the effects of human actions upon the performance of the system. These effects are best evaluated after all other factors have been accounted for. The revised approach is therefore to evaluate the other ('theoretical' etc.) factors first, then to evaluate the effects of the human actions given these other imperfections and sources of uncertainty.

1.7 CONTENTS OF THE THESIS

Chapter 2 considers the allocation of function between men and computers. A general historical perspective is first given, before discussing the alternative philosophies of man-dominant, computer-dominant, and balanced systems. Examples are given of balanced systems which achieve a symbiotic relationship between men and computers, and the discussion moves towards the class of systems under which the case study used in this thesis is subsumed.

Chapter 3 describes the stock control system used as the case study in some depth. Section 3.1 sketches a profile of the host organisation and its product range. The buying and stock control functions are examined, and the aggregate customer demand profile is analysed over a 5-year period. Section 3.2 explains that the design principles were influenced by the entrepreneurial traditions in the building trade with the consequent need for a high degree of freedom of action by the decision-makers. Three levels of system operation are described - the strategic functions of setting service level and stock investment targets, the operational functions of determining when and how much stock to procure, and the administrative record-keeping functions. The need to define service level in business language rather than statistical terms is defended. Section 3.3 defines the objective function of the system and derives the main mathematical formulations. An example of how the ordering mechanisms should realise the target service level is worked through at a conceptual level. Forecasting techniques are explained, and the override facilities permitted to the buyers are outlined. Section 3.4 discusses possible probability distributions to represent the pattern of demand in lead time and explains why the Gamma distribution was selected. Important properties of the Gamma distribution are explained. Section 3.5 defines the performance evaluation criteria and gives pre - and post-operational results.

Chapter 4 analyses the expected behaviour of the system. This is performed as two distinct exercises - the first is a System Dynamics study which examines the time-varying behaviour of the system; the second is a Sensitivity Analysis which examines the static effects of varying the five main variables.

Section 4.1 presents a brief resume of the System Dynamics method developed by Forrester (35). It is pointed out that the model built for this present research encompasses the Supplier/ Distribution Centre/Branch network, and all tests are carried out for a typical slow-moving product and again for an influential fast-moving product. Four disturbances are imposed upon the model - a step change, linear trend, sinusoidal signal and random noise - and the consequential effects on the main variables are charted over a two-year period.

Section 4.2 analyses the sensitivity of the two main Control Indices - Service Level and Average Stock - to the states of the Mean Demand rate, Standard Deviation of Demand, Mean Lead Time, Standard Deviation of Lead Time and Nominal Order Interval. The analyses are carried out on a single set of base data with each factor being varied in turn. The generalisability of the results to any other base data is checked as far as possible by the application of a specially developed relational logic.

Chapter 5 describes the use of simulation techniques in the research. Section 5.1 explains the reasons for using simulation in conjunction with analytical methods. An outline is given of three simulation programs which are used to represent the operational system in varying levels of complexity. Section 5.2 describes how a Gamma generator was selected, tailored for use in the simulation programs and statistically tested. Section 5.3 justifies the use of a Gamma distribution for generating both demand and lead time data to produce an acceptable approximation to a convolved Gamma distribution for demand in lead time. The service level errors arising from this approximation are evaluated.

Chapter 6 evaluates the effects of theoretical errors and assumptions on the two main Control Indices. Section 6.1 examines the suitability of the Gamma distribution to represent the pattern of demand during lead time. Section 6.2 evaluates the errors introduced by the combined effect of using numerical approximations to the Gamma functions and an inexact uniform factor for converting Mean Absolute Deviation to Standard Deviation. The system assumptions of random and independent sequences of demand and lead time variates are checked in Section 6.3; and the implicit assumption of non-overlapping lead times is tested in Section 6.4. Section 6.5 examines the consequences of setting and measuring service levels according to slightly different criteria; and discrepancies caused by using the standard textbook formula for average stock instead of an exact formulation are investigated in Section 6.6. Section 6.7 examines the effects of stock balancing vis-à-vis product independence; and Section 6.8 concludes the Chapter by examining the effects of using first-order exponential smoothing as the main forecasting technique.

Chapter 7 analyses the effects on system performance of buyer interventions into the normal homeostatic operation of the computer algorithm. Section 7.1 describes the principles of allocation of function adopted in the system, and explains the buyer override facilities. Section 7.2 presents the results of buyer training courses, and assesses the buyers' utilisation of the various types of override. Section 7.3 evaluates the effects of lead time overrides in the operational system; and Section 7.4 goes on to examine the possible application of probabilistic estimating to improve the accuracy of lead time forecasting. Section 7.5 looks at the effects of buyers changing computer recommended order quantities because of a distrust of the computer calculations; and Section 7.6 evaluates the effects of gross quantity overrides for genuine commercial reasons. Finally, Section 7.7 evaluates the consequences of data errors on system performance.

Chapter 8 draws conclusions from the study, and suggests certain design principles emanating from the findings as well as areas for further research.

As a general principle, figures and tables which are of immediate relevance are embodied within the text. A glossary of mathematical symbols, which occur throughout, are included once only as Appendix 1.

CHAPTER 2

THE ROLE OF THE HUMAN OPERATOR IN COMPLEX SYSTEMS

2. THE ROLE OF THE HUMAN OPERATOR IN COMPLEX SYSTEMS

Since the Industrial Revolution, when machinery was introduced on a scale which impinges on the quality of working life, the allocation of function between man and machine has been debated. Perhaps the most obvious solution to the allocation of function problem is to allow each component - man and machine - to do the tasks which it can best perform. A number of attempts have been made to classify the functions which are best suited to machines and those which are more efficiently performed by men. The classical attempt at such a separation of function is due to Fitts (36), the basis of which is reproduced in Table 2.1.

Property	Machine	Man
Speed	Much superior	Lag of 1 second
Power	Consistent at any level. Large constant standard forces and power available	2 h.p. for 10 secs 0.5 h.p. for a few minutes 0.2 h.p. for a day
Consistency	Ideal for routine, repetition, precision	Not reliable - should be monitored. Subject to learning and fatigue
Complex Activities	Multi-channel	Single channel. Low information throughput
Reasoning	Good deductive. Tedious to re-program	Good inductive. Easy to re-program
Computation	Fast, accurate. Poor at error correction	Slow, subject to error. Good at error correction
Input	Some outside human senses e.g. radiation	Wide range and variety of stimuli dealt with by one unit
Overload	Sudden breakdown	Graceful degradation
Intelligence	None. Incapable of goal or strategy switching	Can deal with unpredicted. Can anticipate, adapt
Manipulative Abilities	Specific	Great versatility and mobility

Table 2.1 Fitts List - Relative Advantages of Men and Machines

When machines were primitive, the limiting factors on overall system performance were clearly determined by the speed and reliability of the mechanisms, and design efforts were directed towards machine technology. As the machines became faster and more reliable, further improvements in system performance became dependent upon improved operator performance as well as machine performance. During the 20th century there is evidence of an exponential growth in technology according to several criteria, e.g. speed of transport. Singleton (37) argues that these apparently exponential curves must turn out to be the early part of Sigmoid curves, on the grounds that the human component of technological systems have a limit on learning capacity and a ceiling on knowledge. Even if these limits are not absolute, they certainly cannot keep pace with an exponential growth pattern. It follows that for further improvements in system performance to take place, attention must be directed towards the management and operation of technology.

The advent of the computer constitutes a quantum leap in the progress of technology, which calls for a re-evaluation of the traditional principles of role allocation. Some of the disadvantages of machines in the Fitts list are becoming less valid (e.g. computers are becoming much less tedious to re-program, and overload conditions may initiate a 'fail soft' procedure). Indeed, where the machine happens to be a computer, software development has profound implications on the tasks to which it is suited.

The majority of existing computer applications have been developed primarily to improve the efficiency of the base operating and administrative procedures in large organisations. In essence they were introduced to speed up numerical calculations and data handling, i.e. they exploit to the full the speed, consistency and computation attributes identified by Fitts. All of these functions are inherent in routine clerical work, and it is not surprising that early computer applications were concentrated at this level.

The computer aiding of higher management functions was, in the main, confined to the provision of voluminous reports which were conveniently produced as a by-product of the data passing through the computer. Hence, management were flooded with masses of semi-structured information which they had neither the time nor cognitive ability to convert to a form which would be instrumental in improving the running of the business. The dangers of cluttering the mind with too much information are brought out in an experiment by Baker and Goldstein (38). They showed that human assimilation of information is greatly facilitated by presenting it in a sequential fashion rather than in batches. They conclude that "the experimental data reveal no value in displaying material which may contain potential information, and indicates that much extraneous material delays problem solution by increasing display search time." The conclusion is certainly no less valid if the 'display' is a copious amount of printed output; and the sequential presentation can be parallelled by the use of interactive terminals.

Miller (39), discussing task allocation in decision-making, declares that humans are needed because they can "exercise judgement, interpret meanings, weigh values, generalise experiences, separate wheat from chaff in a problem, act on incomplete information, and bring a broader content of information to bear on an issue that is convenient to program in advance".

Shrenk (40) argues that in spite of remarkable advances in machine technology, which has made it possible to automate many functions formerly assigned to men, there remain fundamental intellectual features in decision-making, such as setting objectives, allocating resources, diagnosing environmental conditions, defining and selecting courses of action and interpreting intricate patterns of events which cannot be handled by formal algorithms. Vaughan and Mavor (41) advocate conversational man-computer systems where "man directs, via guidelines, goals and constraints; the computer generates hypotheses and alternatives; man selects out those to be developed; the computer tests; the man evaluates the results of the tests. The computer opens up and widens both the problem and the solution space; the man evaluates and narrows it down." This concept of the interspersion of activities between man and computer is one which has gained favour in recent years. This may be because the prevalent earlier concept of providing the decision-maker with all of the potentially relevant information en bloc so that he could pick and choose according to his intuition had been demonstrated to be unworkable.

The general philosophies to computer-aiding discussed above are now translated into a more usable form by examining some examples. It is useful to structure the discussion using a categorisation of computer-aiding advanced by Press (42):-

1) Man instructs; machine executes.
2) Machine instructs; man executes.
3) Balanced man-machine systems.

Nearly all of the present commercial computer systems are in the first category. A Sales Ledger system provides a typical example. The Systems Analyst, or the user he represents, decides how the ledger will be structured, whether it will be 'open-item' or 'balance brought forward', the customer credit limits, the debt collection stages, etc. The rules are presented to the computer in the form of a set of instructions, and the computer processes the input transactions in a totally pre-defined manner. The computer's role is entirely one of execution. The instruction set may be convolved into complex branches and loops which may give the impression that the computer is making choices, but this is purely illusory - it is merely following predetermined paths and taking no active part in the problem-solving process. Likewise, there is no active computer participation when parameter values are asked for by the computer, or when input errors are notified to the operator.

In these systems the man-computer partnership is totally one-sided with the man the totally dominant partner. It is the present author's view that systems were designed this way because the Designers took the most obvious and least imaginative route of replacing a clerical system by an automated equivalent. Any fear about an encroachment on human primacy did not come into it. This latter sociological perspective is discussed by Sackman (43) who asserts that computers should be used to support a continuous stream of human functions and decisions. Citing the SAGE air defense system he writes:-

"In SAGE operations, overall system response to environmental air defence stress is primarily defined, organised, and sustained by a continuous stream of human functions and decisions. System response is, by and large, secondarily determined by the automated component, at least as far as interpretation and control of the threat environment are concerned. Human actions may override and redefine computer inputs, operating parameters, and outputs in practically all significant activities affecting the course of air defense operations, particularly those concerning the conduct of the air battle."

His description of the system, however, indicates that the partnership is not nearly so one-sided as the traditional commercial systems, and some of the operations are heavily interspersed and could quite reasonably be connoted as symbiotic, e.g. man establishes tracks; computer monitors tracking; man determines if friendly; computer checks identification; man decides whether to intercept; computer provides interception guidance. From this description it seems very doubtful if this system could survive without the computer. If it cannot, then the relationship is, by definition, symbiotic.

The incidence of 'Machine instructs; man executes' systems is relatively rare even if one takes the freedom to view the system purely in an operational mode. (If the design implications are considered, it could be argued that man always has the dominant role). However, at least two application areas could reasonably be regarded as computer-dominated, viz. process control and teaching machines.

In process control, the computer is engaged on a monitoring task, and when one or more variables transgress prescribed limits an alarm system is triggered which instructs the operator to take certain actions. The computer is here utilizing its superior attributes for scanning large quantities of data, and for following detection, diagnostic or corrective algorithms. Andow and Lees (44) describe the alarm system in a nuclear power station whereby the computer performs a sequential pattern recognition function. Alarm conditions can also be based on indirect or inferred variables such as rates of change. It is probably fair to conclude that these complex threshold conditions could not be reliably interpreted by the operator and the system could be justifiably described as computer-dominated even though the operator may hold the responsibility for the ultimate action.

The development of teaching machines was pioneered by Pask (45). Here the computer presents the subject with questions, the choice of which is determined by the deficiencies in the subject's previous answers. The computer is employed as a teacher-simulator, displaying the human ability to probe into areas which cannot be predetermined but are only made manifest when a dialogue with the individual is entered into. Press (42) cites a research study into a similar field where a computer prompts a business executive into considering various factors and interrelationships when formulating strategic plans and objectives. Again the subject is gaining more from the computer than vice-versa, so it is not unreasonable to regard the system as machine-dominant.

The remainder of this Chapter will review the middle-ground between man-dominance and machine-dominance, i.e. what Press (42) refers to as 'balanced systems'. The concept of man-computer symbiosis is intellectually satisfying as it emphasises the principle of 'man-with-a-computer' rather than 'man-versus-computer' which several authors have advocated. The notion of a symbiotic relationship was first introduced by Licklider (3), who wrote:-

Whitfield (46) reviews the state of the art of man-computer symbiosis in 1975, citing examples from the whole spectrum of man-computer dominance (which is perfectly justifiable as the concepts of 'symbiosis' and 'equal balance' are by no means synonymous). The examples given here are, however, regarded as both symbiotic and balanced.

One of the best examples of man-computer symbiosis is the Probabilistic Information Processing method pioneered by Edwards (47). This is based on the premise that inherently subjective processes are just as amenable to scientific treatment as any other processes. His method is to break an intellectual task into little pieces, make separate judgements about each piece, and re-aggregate the results using a formal aggregation rule. The aggregation rule is based on Bayes' theorem, which specifies how opinions, expressed formally as numerical subjective probabilities, can be combined objectively. Men are good at judging the diagnostic meaning of a single datum (i.e. a fixed point against which a variable can be assessed). What they are not good at is combining the judgements, which is where Bayes' theorem is invoked. This implies a division of effort between men and computers which is summarised in Fig. 2.1.

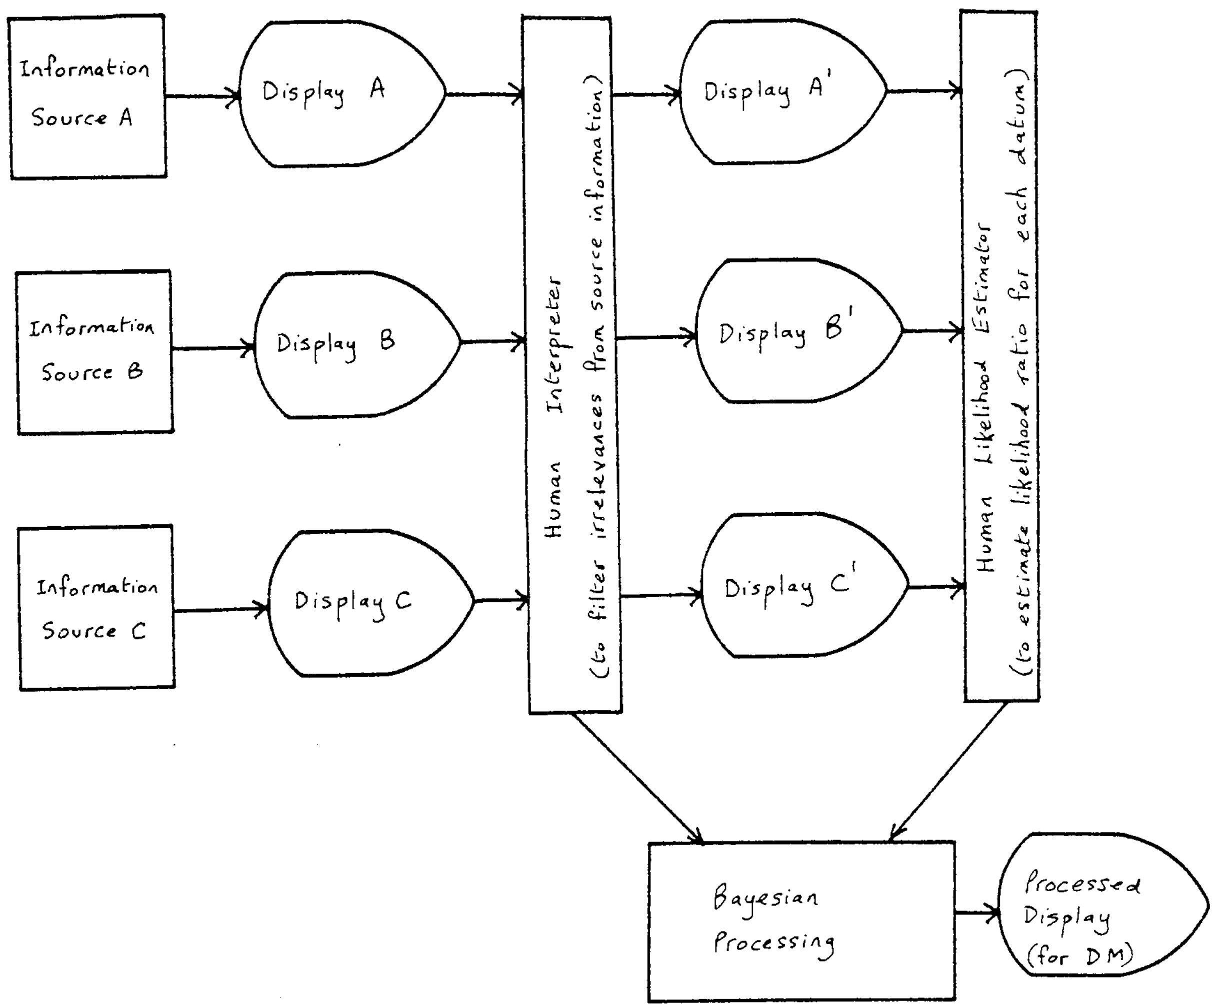

Fig. 2.1 Diagrammatic Representation of Probabilistic Information Processing

Given a payoff matrix, the decision-maker could be provided with the expected value of various options and, if there are no other considerations, the decision could be a matter of simply selecting the maximum.

The cases cited as heuristic methods in Section 1.2 are also good examples of balanced systems. It will be seen from these that the usual strategy for developing balanced systems is to extend computer-only systems to allow for interaction with, and participation by, a man who is enabled to input decisions and to obtain information which assists him in making those decisions. This strategy is useful for problems which can be automated, but where all of the information known to the investigator cannot be represented explicitly or where routines to process it economically cannot be devised. It will be seen later in this thesis that the buyer-computer system used as the case study follows these principles exactly. The computer uses an automatic algorithm for generating recommended purchase orders, and the buyer intervenes by supplying the algorithm with information which could not be feasibly or economically obtained as part of its regular source data set. The process can be represented by the flow diagram in Fig. 2.2.

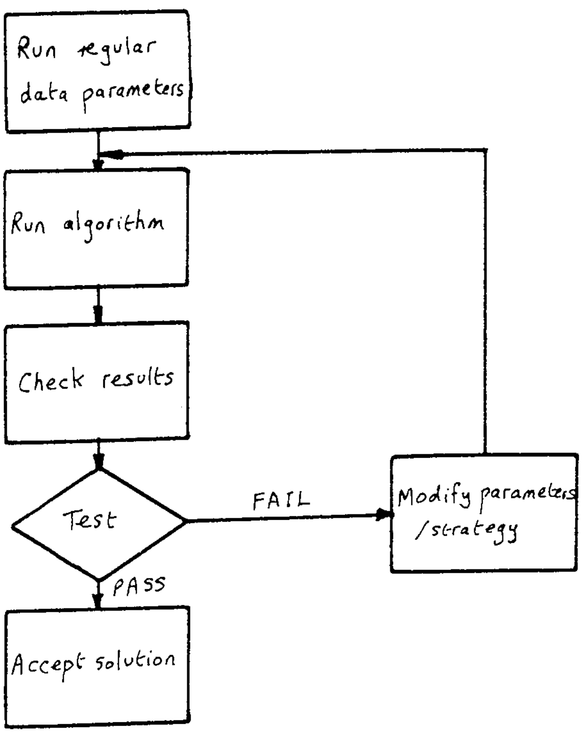

Fig. 2.2 Flowchart for Man Interacting with Computer Algorithm

In the case study, the data parameters are such factors as supplier lead times, customer demand rates etc., and the strategy could be an adherence to the target service level. Clearly some variations to this schematic might be necessary for different applications, but it is considered to be a reasonable generalised model for the type of man-computer system in which the case study is subsumed. It is also evident that with this approach the interesting parts of the task remain with the man, thus heightening his skills by concentrating his efforts on to the purely intellectual activities. Also, the amount of computer-aiding can be varied according to the individual needs of the operator, i.e. highly skilled operators with a good intuitive feel for what makes the system 'tick' would be justified in overriding the computer and supplementing the regular source data much more than less experienced individuals.

CHAPTER 3

DESCRIPTION OF CASE STUDY IN INVENTORY CONTROL

3.1 THE SETTING

The system controls inventory at four Distribution Centres for a major distributor of building materials in the U.K. The Organisation comprises 12 semi-autonomous Trading Units which provide a national coverage. A total product range of 24,000 lines is maintained.

The Buying and Distribution organisation is illustrated in outline in Fig. 3.1. This shows only the most common information and material flows - in practice many variations exist.

The Trading Units share the four Distribution Centres, which stock all of the lines except bulk products (such as sand, cement, bricks and aggregates) where transportation is costly. The major buying negotiations are conducted on a national basis by a central support function. Buying teams at the Distribution Centres then procure the goods to satisfy the Trading Unit requirements. Stock management is the responsibility of the Distribution Centre buyers, within the constraints of the deals negotiated by the central support group and the Trading Unit commercial strategies. The Distribution Centres each stock between 6,000 and 9,500 product lines and they supply most of the contract orders and the larger customer orders by direct delivery.

Each Trading Unit operates a number of Branches which service both the trade and the general public. The Branches obtain almost all of their stocks by regular replenishments from the Distribution Centres. Most Branch sales are collected by the customer, but many Branches also provide a local delivery service. A typical Branch has a range of approximately 2,500 of the more popular product lines.

Trading activity is geared very strongly to the national economic

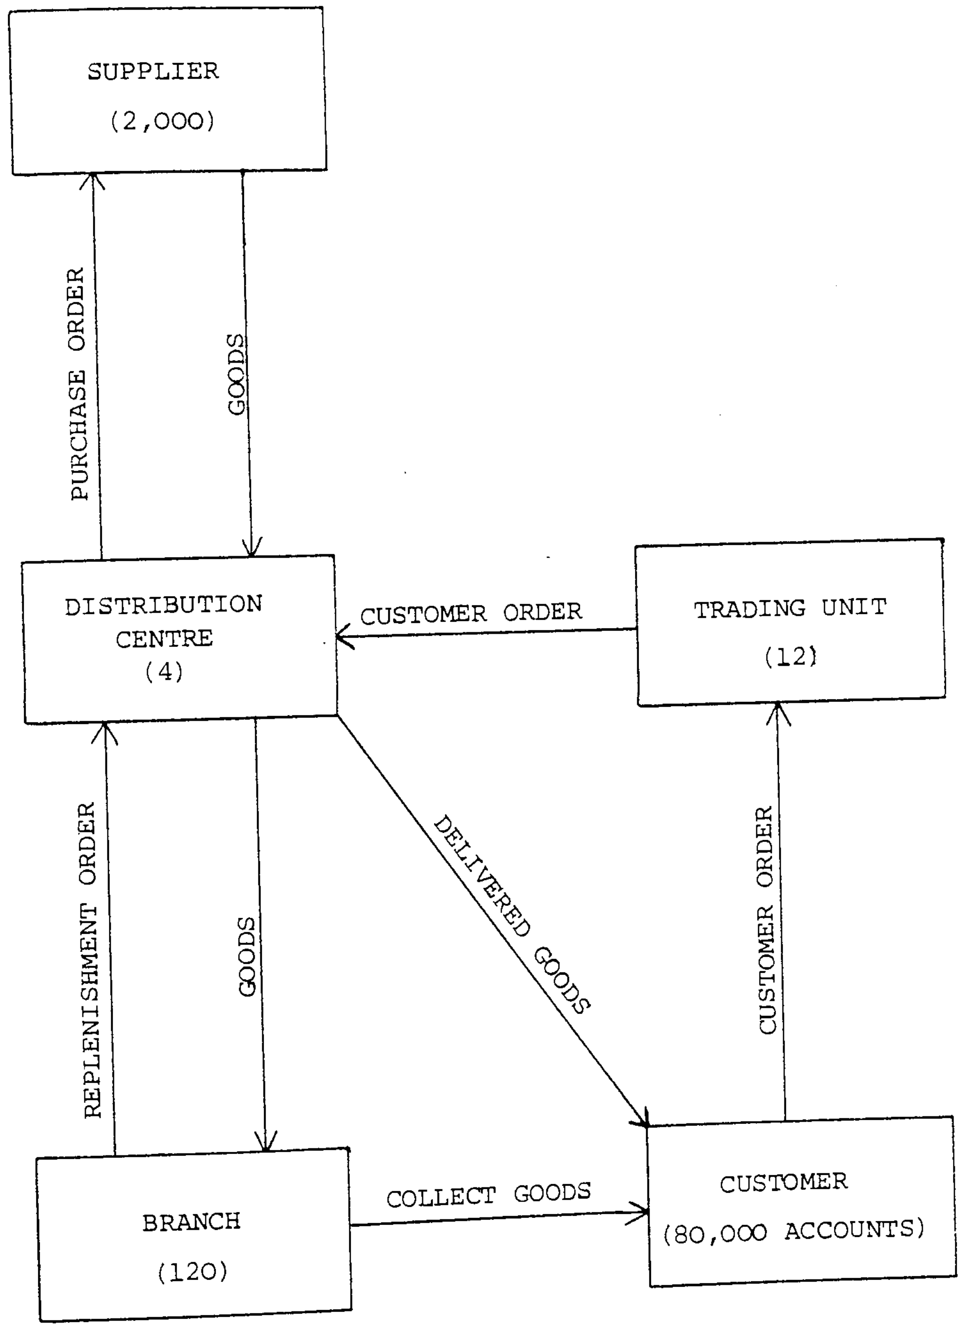

Fig. 3.1 Buying and Distribution Organisation

climate, as the building trade is particularly susceptible to financial cutbacks in the public sector, affecting housing programmes. The business activity characteristics can be illustrated to good effect by utilising the 'Classical Decomposition Method' for time series forecasting with respect to sales, as described by Wheelwright and Makridakis (48). Briefly, it is asserted that historic demand data is composed of three underlying sub-patterns as well as random influences. These are a trend factor, a cyclical factor and a seasonal factor. By decomposing the data to isolate these factors, forecasts can be formed. Fig. 3.2 shows the results of applying the analytical procedures to decompose the data for the period from March 1976 to March 1981. There were no major aquisitions or disposals of Trading Units during this period. The base data comprises the monthly sales totals. The effects of inflation were removed by dividing the figures by the Index for Construction Material Prices published by the Department of Industry for the respective months. As this Index was set to 100 to correspond to the average prices during 1975, this procedure deflates the sales values to a mid-1975 base.

The linear trend was first isolated by calculating the regression line over the period. The seasonal pattern was determined by first calculating 12-month moving averages (which are, ipso facto, deseasonalised), appropriating these to the 6th (central) month in each case, then calculating the ratio of the actual to the moving average for each month to obtain its seasonal factor. Finally the medial averages of the seasonal factors were computed for each calendar month over the 5-year period. The business cycle was identified by dividing the moving average for each month by the ordinate of the regression line for the month.

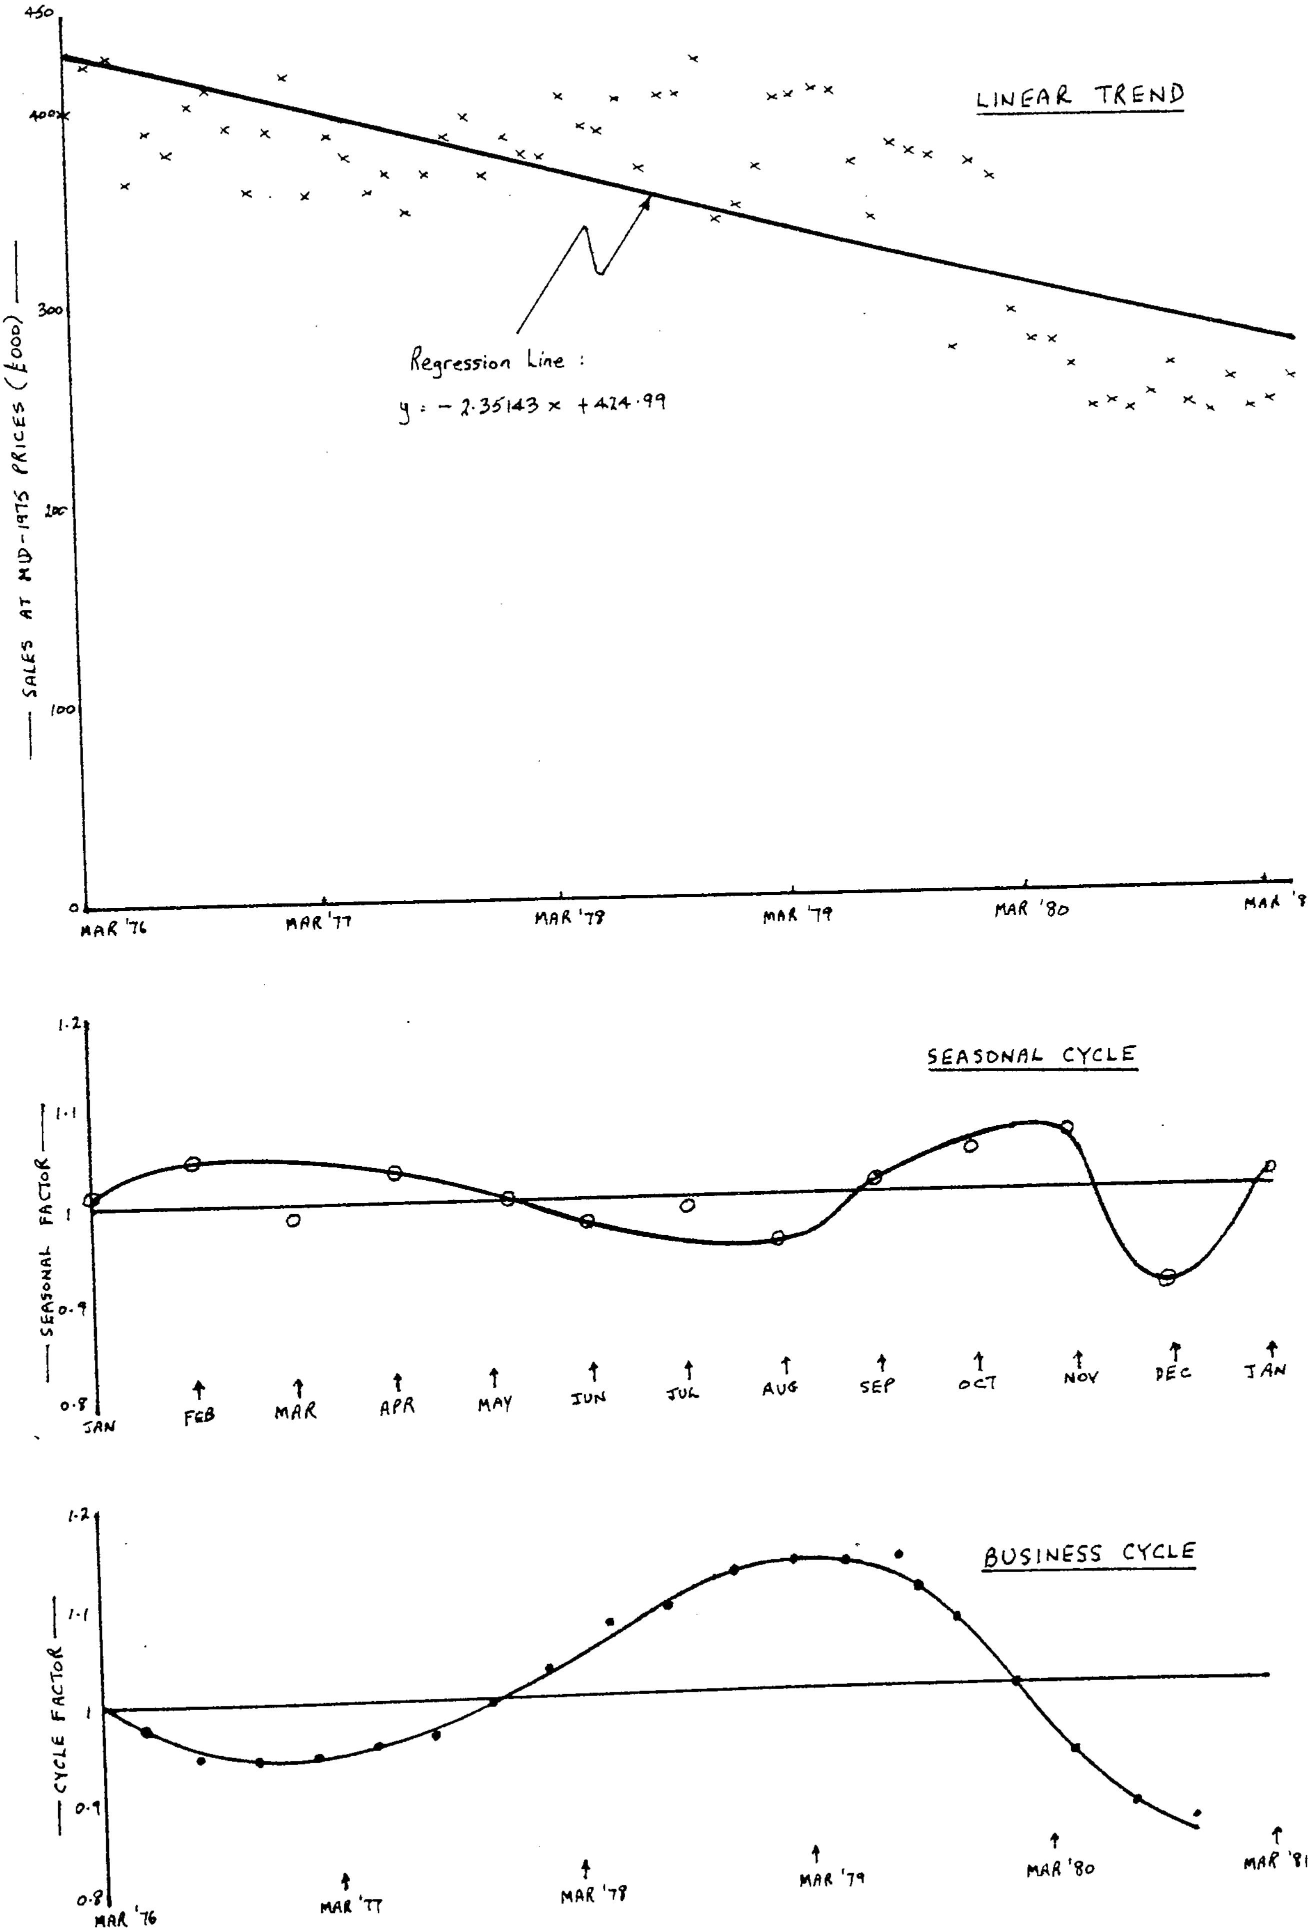

Fig. 3.2 Subpatterns in Sales Data March 1976 - May 1981

The salient points which can be elicited from Fig. 3.2 are:-

a) Trading activity was at a fairly stable level until the end of 1979 after which it declined sharply.

b) There is evidence of a business cycle with a period of around $4\frac{1}{2}$ years. The severe national economic recession of 1980/81 is clearly reflected in a sharp increase in the amplitude of the cycle over the previous low period in the Spring of 1977. (As the linear trend has been removed, this graph gives an exaggerated impression of the level of activity during the peak period in 1979).

c) There is some evidence of seasonal peaks in the Spring and the Autumn. This can be verified in practice by the high level of building activity in the Spring and early Summer, and by the installation of central heating systems and insulation prior to the Winter. The low point in December is caused by closures of selling points over the Christmas period.

The long-term cyclicity and, to a lesser extent, the seasonal patterns gives rise to instability in the level of stock investment and in the product mix. Whitin (49) gives a good account of the interrelationship between stock investment and business cycles. In periods of recession a policy of retrenchment is invariably adopted, capital being diverted away from stocks and invested in maintaining the fixed asset base to take maximum advantage of the next high-activity period. Also, depressions in house building are usually accompanied by an increase in renovation, repair, and maintenance activity. This results in a shift of the balance of stocks away from construction materials and industrial tools towards restorative materials. Some volatility of the product range is also engendered by fashion changes which are often imposed by the manufacturers.

Colour changes in sanitaryware and decorating materials are common examples which can result in a significant obsolescence factor. This also gives rise to the problem of estimating potential sales for new lines, and the attendant problem of assessing the volume required to fill the pipeline of stockholding locations with initial supplies.

The supply situation is extremely variable across the product sectors. Paint, for example, can normally be obtained from local depots within 24 hours, whereas lead times of 6-10 weeks are typical for suppliers of ironmongery and tools. Many suppliers operate on the basis of production cycles at irregular time intervals which result in highly variable and unpredictable lead times. Industrial disputes are not uncommon in some sectors of the business and alternative sources are not readily available for many vital products such as copper tube and radiators. Minimum order loads are a common practice, and these may operate on the basis of weight, volume, cost or numbers of units. Relatively small orders, even if they satisfy the minimum requirements, are often held back by the supplier pending the placement of further orders. This again gives rise to lead time variability.

The wide product range contains a high proportion of slow-moving items which support the sales of the more profitable products. Fig. 3.3 depicts the dispersion in demand rates, which follows the '80/20 rule' i.e. 20% of the products account for approximately 80% of the costed demand on a Distribution Centre.

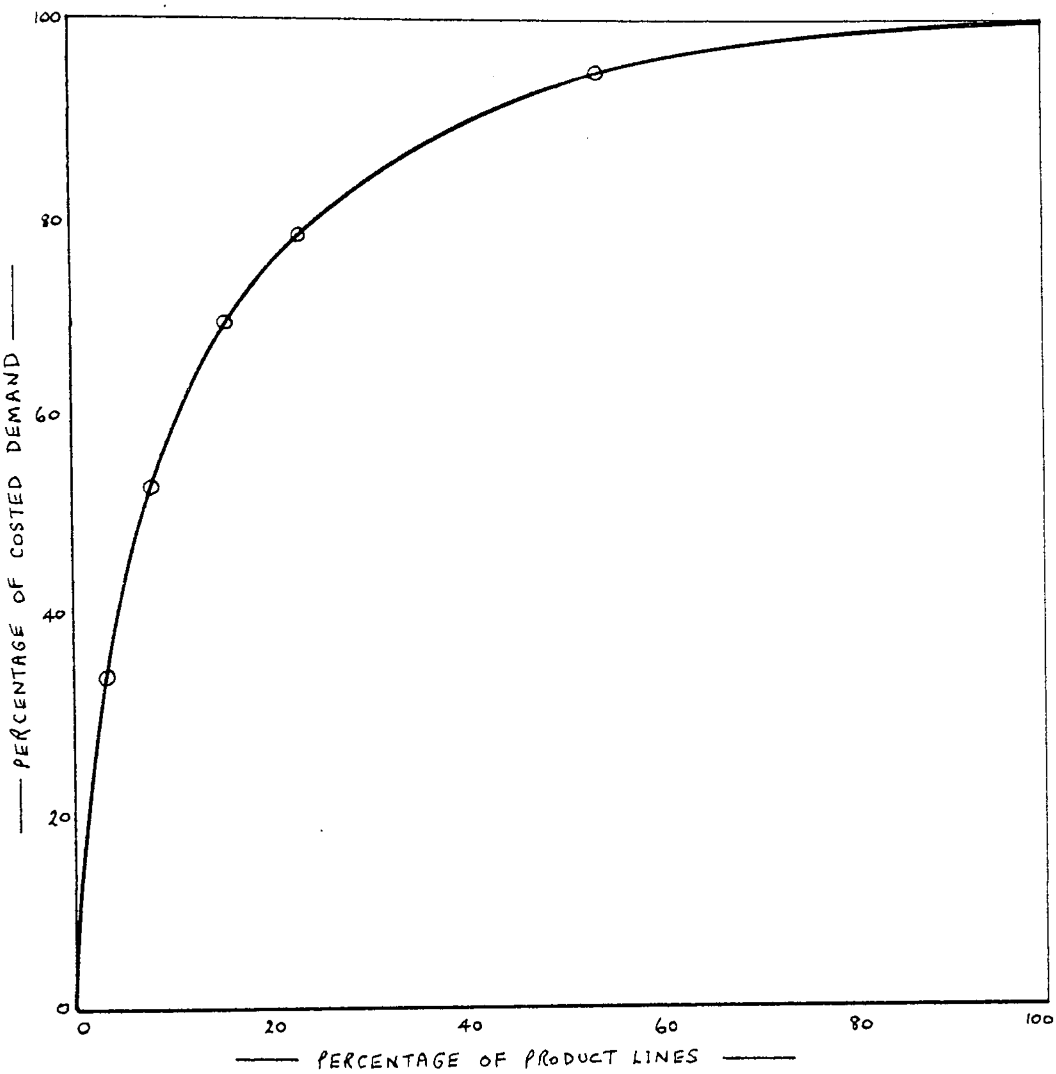

Fig. 3.3 Relationship Between Percentage Lines and Percentage Demand

At Branch level the sales volumes are of course attenuated, and approximately 75% of the lines are sold at a rate of less than one unit per week. Their very low usage rate appears to self-contradict the marketing justification for retaining them on the grounds of maintaining customer support for the more profitable products.

A small proportion of the products are assembled into kits at the Distribution Centres e.g. kitchen units, central heating packages. In these cases the demand rates for the components are strongly correlated. In all other cases demand rates between the products can be assumed to be uncorrelated as the average number of item lines on a customer order is only 2.7.

The stock levels at the Distribution Centres and Branches are controlled by two independent computer systems which are both operated remotely from a central location. The only link between the systems is a file containing data relating to Distribution Centre despatches to the Branches. This is created by the Branch system and is passed to the Distribution Centre system for the purpose of updating the stock balances and demand volumes. There is no multi-echelon control system (because of mathematical intractability), and in practice strategic redistribution of materials between stockholding locations is limited to the allocation of bulk supplies which are accepted in total at one Distribution Centre as part of a centrally-negotiated contract with the supplier.

The system which controls Distribution Centre stocks (known within the Organisation as the 'Inventory Management' system) is the subject of the case study for this thesis.

3.2 DESIGN CHARACTERISTICS

The traditions and trading environment in the Builders Merchant industry dictated to a large extent the philosophy upon which the Inventory Management system is based. The Organisation developed over a number of years by the acquisition of relatively small Companies and the Managers brought with them the traditional entrepreneurial skills of small traders. Rather than attempting to get these Managers to revolutionise their practices and conform to a rigid set of programmed rules, the system was designed to provide a basal control system for stock replenishment whilst allowing the Managers and Buyers ample scope to exploit commercial opportunities. There are many reasons for permitting market intelligence to be injected, such as sales promotions, new product introduction, fashion changes, competitors' actions, uncertainty of supplies and a particular sensitivity to the state of the economy. For these reasons there are ample facilities embedded in the system for human judgement to supervene.

The system is designed around three levels of control. The top level can be designated 'strategic' and this is concerned with stock investment decisions and the concomitant considerations of stock ranging and customer service. The top level control function is facilitated by inbuilt simulation capabilities which display the consequences of alternative courses of action and operating parameter levels. The second level is concerned basically with stock replenishment - when to procure more stock and how much to get. At this level the computer provides an advisory service and the Buyers and Product Managers must decide whether or not to accept the computer recommendations. The third level is much more deterministic in nature and it involves posting transactions to keep the stock and order balances up to date, counting stock, and the other routine operational work. The system takes over the book-keeping tasks completely, and the management reporting functions are now contained entirely within the computer system.

The concept of levels of control is not new and taxonomies of control and decision functions have been advanced by several writers such as Ansoff (50), Simon (4), Anthony (51), and Beer (52). It is, however, helpful to view the system as a hierarchy of control systems and the associated decision levels.

Because of the conceptual gap between the inbred management operating styles and formalised technological systems it is especially important to provide the human/computer interfaces in a suitable language. For this reason 'Service Levels' are defined throughout in terms of the percentage of annual demand met ex-stock, rather than in terms of stockout probabilities. The problem with using 'probability of stockout' as a measure of service is that it indicates how often a stockout is likely to occur, but it gives no idea of the amount of potential sales which could be expected to be lost given a stockout. Hence it gives no clear indication of lost business, and it is not a particularly meaningful measure of service to the customer. Lewis (53) and Brown (54) make a clear distinction between these alternative measures of customer service.

3.3 THEORETICAL FRAMEWORK

The objective function of the control system is to maximise the Gross Profit from each Product Group subject to a given level of stock investment. Gross Profit is the product of Gross Margin and Sales, Gross Margin being defined as:-

$$\frac{\text{Discounted Selling Price - Net Buying Price}}{\text{Discounted Selling Price}}$$

This objective function is fairly narrow compared to some models which have been developed (e.g. the Arrow, Harris, Marschak model (55)) as it was decided to exclude cost functions where quantification is contentious e.g. stockout penalty costs.

Optionally, the system will maximise Sales instead of Gross Profit, by the simple expedient of setting all of the Gross Margin factors to unity. To date, this option has been exercised by all Distribution Centres in order to permit management to manipulate margins for various commercial reasons without prejudice to the balance of stocks.

The system does not attempt to balance stocks <u>across</u> Product Groups nor does it determine operating parameter levels. These decisions are heavily influenced by the prevailing economic and commercial climate and they are left to management discretion.

Within the system, products are classified into three hierarchical levels, which correspond to the three levels of control already described.

i) Product Group (first level)

Each of the 25 Product Groups represents a major sector of the business e.g. Ironmongery, Electrical, Heating. This is a management reporting category used consistently by all computer systems.

ii) Buying Family (second level)

A Buying Family is a group of products which must be considered collectively when placing orders. In most cases it will constitute the complete range of products obtained from a supplier, but on occasions it is desirable to split a supplier's range into more than one Buying Family. Each Buying Family must be completely contained within one Product Group.

iii) Individual Product (third level)

Strategic control is exercised by the management setting one service level for each Product Group, using information from a 'Stock Strategy Report' (Fig. 3.4) supplied by the system. This shows the relationship at Product Group level between Service Level, Average Stock Value, Stockturn and Expected Lost Sales per annum for a range of service levels between 50% and 99.8%.

Product Group	Service Level (%)	Av.Stock (£)	Stockturn (p.a.)	ELS p.a. (£)
Sanitaryware	97.0	34263	7.5	7947
	96.0	31340	8.1	10589
	95.0	29056	8.7	13254
	94.0	27197	9.2	15913
	93.0	25632	9.6	18565
Heating	97.0	73149	4.4	9844
	96.0	67115	4.7	13145
		⋮		

Fig. 3.4 Example of a Stock Strategy Report

The computer effects daily monitoring of physical stocks, on-order balances, demand rates, lead times etc. Recommended purchase orders are generated at Buying Family level to achieve the target service levels. At the same time the stocks are balanced to maximise Sales or Gross Profit. The control mechanism is now described in detail, using the symbols defined in Appendix 1.

As stated, 'Service Level' is defined as the percentage of costed annual demand which is satisfied ex-stock. Demand received during a stockout is assumed to be lost.

'Protection Level' is defined as the probability of remaining in stock when an order just placed is received.

The probability density function is described by a Gamma distribution. The properties of the Gamma distribution are discussed at length in Section 3.4.

The 'risk period' is the time interval between the stock cover crossing the reorder point and the resulting order being received. The time interval between the stock cover crossing the reorder point and the next opportunity to order is uniformly distributed with mean $T/2$ and variance $T^2/12$. The demand during the risk period is distributed according to $f(x)$ with:

$$\mu_{LD} = \mu_D(\mu_L + T/2) \qquad (1)$$

and $$V_{LD} = (\mu_L + T/2)\sigma_D^2 + \mu_D^2(\sigma_L^2 + T^2/12) \qquad (2)$$

For a gamma distribution the important 'shape parameter',

$$k = \frac{\text{Mean}^2}{\text{Variance}} \qquad (3)$$

On the occasions where $x>R$, potential sales of $(x - R)$ will be lost.

$$\therefore \text{ ELS in risk period} = \int_R^\infty (x-R)\ f(x)\ dx$$

And, assuming 250 working days per annum, ELS p.a.= $\frac{250}{I}\int_R^{\infty}(x-R)\,f(x)\,dx$

A common approximation for average stock level is given by:-

$$R - \mu_D(\mu_L+T/2) + \mu_D\, I/2 \qquad (4)$$

Considering the n products in a Product Group:

$$\text{Total ELS p.a. (at Cost Price)} = 250\sum_{i=1}^{i=n}\frac{C_i}{I_i}\int_{R_i}^{\infty}(x-R_i)\,f(x)\,dx$$

$$\text{Total demand p.a.} = 250\sum_{i=1}^{i=n}\mu_{D_i}\,C_i$$

$$\text{SL for Product Group} = \frac{250\sum_{i=1}^{i=n}\mu_{D_i}C_i - 250\sum_{i=1}^{i=n}\frac{C_i}{I_i}\int_{R_i}^{\infty}(x-R_i)\,f(x)\,dx}{250\sum_{i=1}^{i=n}\mu_{D_i}C_i} \qquad (5)$$

$$\text{Average Stock Value} = \sum_{i=1}^{i=n}C_i(R_i - \mu_{D_i}(\mu_{L_i} + T/2) + \mu_{D_i}I_i/2) \qquad (6)$$

$$\text{Gross Profit per unit} = \frac{M_iC_i}{(1-M_i)}$$

$$\text{Expected lost Gross Profit p.a.} = 250\sum_{i=1}^{i=n}\frac{M_iC_i}{(1-M_i)I_i}\int_{R_i}^{\infty}(x-R_i)\,f(x)\,dx \qquad (7)$$

If the total stock value for the Product Group is constrained to be $\leqslant K$, the Lagrange Multiplier technique can be used to minimise (7) subject to (6) :-

$$\Phi(R_i,\lambda) = 250\sum_{i=1}^{i=n}\frac{M_iC_i}{(1-M_i)I_i}\int_{R_i}^{\infty}(x-R_i)\,f(x)\,dx$$
$$+\lambda\Big(\sum_{i=1}^{i=n}C_i(R_i-\mu_{D_i}(\mu_{L_i}+T/2) + \mu_{D_i}I_i/2)-K\Big)$$

For a minimum, $\frac{\partial\Phi}{\partial R_i} = -250\,\frac{M_iC_i}{(1-M_i)I_i}\int_{R_i}^{\infty}f(x)\,dx+\lambda C_i = 0 \qquad (8)$

and, $$\frac{\partial\Phi}{\partial\lambda} = \sum_{i=1}^{i=n} C_i(R_i - \mu_{D_i}(\mu_{L_i} + T/2) + \mu_{D_i} I_i/2) - K = 0 \qquad (9)$$

From (8), $$\int_{R_i}^{\infty} f(x)\,dx = \frac{\lambda I_i (1-M_i)}{M_i\ 250} = (1-P_i) \qquad (10)$$

Whilst the Lagrange Multiplier is introduced as an arbitrary variable for the purpose of minimising the lost profit function subject to the stock investment constraint, if the margin factors are set to unity, it, the Lagrange Multiplier, does have a meaningful interpretation.

Transposing equation (10) and omitting the factor $(1-M_i)/M_i$,

$$\lambda = (1-P_i)\ \frac{250}{I_i}$$

= Probability of stockout in order cycle x Orders per annum

= Average number of stockouts per annum

Hence the average number of stockouts per annum will be the same for all items in the Product Group. Gerson and Brown (56) reached the same conclusion for a backorders system which minimises backordered demand for a given stock investment.

As the Coefficient of Variation ($= \sigma / \mu$) is generally lower for fast-moving products than for slow-moving products, the former could be expected to lose a smaller proportion of their demand, given a stockout, than the latter. Hence the system will assign higher implicit service levels to steady fast-moving lines than to slow-moving lines with a lumpy demand pattern. In a comparison of six decision rules, Brown (57) concludes that the 'equal shortage' principle compares favourably with the other alternatives for field warehouses.

The computer routines for calculating the stable λ values and the resulting Buying Family ordering points are shown schematically in Fig. 3.5. The 'functions' referred to are numerical approximations to Gamma properties derived by Johnston (58):-

<u>Function 1</u>

Calculates R given P, k

<u>Function 2</u>

Calculates ELS given P, k

<u>Function 3</u>

Calculates ELS given y, k

The ordering routines are described by a fictitious example of a Product Group comprising a single Buying Family with 200 items. Assume that a Product Group service level of 95% has been set, the target order interval for the Buying Family is 2 weeks, and the aggregate costed demand is £200,000 p.a.

Then, ELS p.a. for Buying Family,

$$\sum_{i=1}^{i=200} ELS_i \frac{250}{I_i} = (1 - 0.95) \times £200,000$$
$$= £10,000$$

And, ELS per order for Buying Family,

$$\sum_{i=1}^{i=200} ELS_i = \frac{£10,000}{25}$$

$$= £400 \text{ which is the order point}$$

On each overnight stock status update on the computer, an ELS calculation is performed using Function 3. This represents the amount of sales which would, on average, be lost before a purchase order placed immediately is delivered. This amount increases daily as the stock falls, as illustrated in Fig. 3.6.

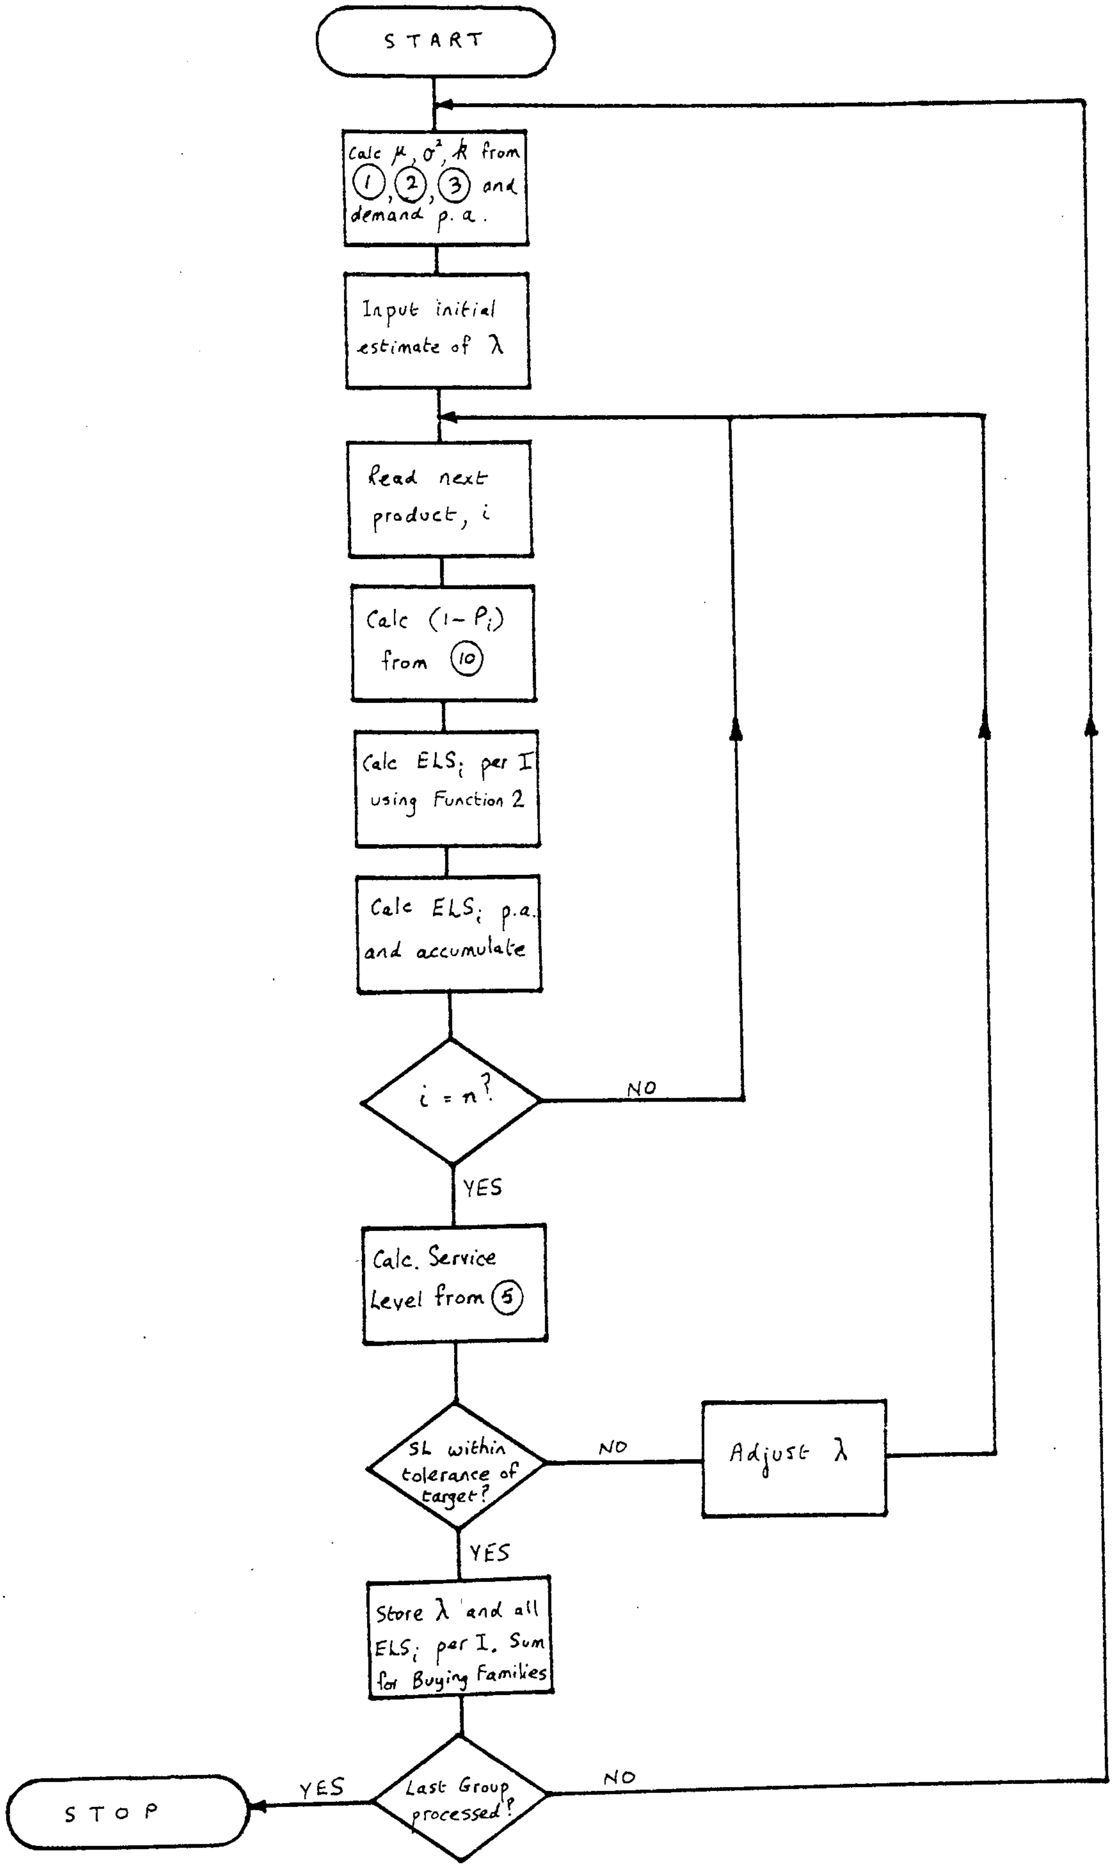

Fig. 3.5 Schematic of Routine for Calculating ELS Reorder Values

Product	ELS Day 1 (£)	ELS Day 2 (£)		ELS Day n (£)
1	£0.75	£2.06	. . .	£5.75
2	£2.63	£4.22	. . .	£7.72
3	£0.00	£0.06	. . .	£0.68
.	.	.		.
.	.	.		.
.	.	.		.
.	.	.		.
.	.	.		.
200	£5.22	£14.10	. . .	£38.70
BUYING FAMILY	£56.74	£88.20	. . .	£403.34 ← Generate order

Fig. 3.6 Example of Purchase Order Generation

After each ELS calculation the total for the Buying Family is compared with the order point. On day n, when the ELS exceeds the order point a recommended purchase order is generated and transmitted to the Distribution Centre buyers for actioning. The recommended quantities are calculated as:

$$R_i + \mu_{D_i} I_i - ELS_i - \mu_{D_i} T/2 - y_i$$

where R_i is calculated using Function 1.

On average, a particular product should have a stock cover of $R_i - \mu_{D_i} T/2$ when an order is generated. Hence the average order quantity is $\mu_{D_i} I_i - ELS_i$ (i.e. the average sales per order interval) which keeps I at a stable value of 2 weeks.

By working backwards through the logic, if £403.34 of potential sales are lost before the order is delivered (in practice this would average £400 as a correction for overshoot is incorporated) and this is repeated a further 24 times, the losses in the full year would be approximately £10,000 which would produce the target service level.

The forecasting of demand and lead times is carried out separately using simple exponential smoothing in both cases. The first two moments of demand in the risk period are calculated using equations 1 and 2. A demand filter is incorporated to avoid distortions during and immediately after stockouts. Seasonality is catered for by the inclusion of base index tables of weekly factors for a number of known seasonal patterns. The demand forecast is first deseasonalised by dividing by the factor corresponding to the average age of the data constituting the EWMA, then re-seasonalised by multiplying by the factors for the projected period.

A facility for improving forecasts for products subject to intermittent demands is incorporated in the system but this has not yet been invoked. Croston (59) has shown that a EWMA is a biassed estimator where a demand is not registered on every updating cycle. The Mean and MAD of both issue size and issue interval are recorded by the system should Croston's method be found necessary.

The buyers are allowed a high degree of freedom for incorporating subjective judgements into the computer routines. They are permitted to:-

a) introduce 'Service Level Adjustment Factors' at Product or Buying Family level where commercial considerations are perceived to outweigh the mathematical optimum deployment of stocks. These factors adjust the individual service levels by a predictable amount.

b) override lead time forecasts where market intelligence is available. Whilst overrides are in operation the statistical forecasts are updated but they take no part in the ordering calculations.

c) change recommended purchase order quantities, or suppress order generation completely.

The buyers are guided and prompted into intervention action by several computer reports. These include reports on low stocks, excess stocks, urgent orders and demand and lead time statistics.

3.4 THE GAMMA DISTRIBUTION

Although Service Level is the measure of service adopted in the system, the Protection Level is an integral part of its calculation and the probability density function of the demand during the risk period must be determined. As stock is reviewed daily in the Inventory Management system the terms 'risk period' and 'lead time' are almost synonymous. (The average risk period is actually a half day greater than the average lead time). As the term 'demand in lead time' is commonly used in Stock Control literature, it will be used in the remainder of this thesis to denote the demand during the risk period. Nevertheless, all of the relevant formulae in the system and simulation programs incorporate the review cycle.

The most commonly used statistical distributions to describe the probability density function of demand in lead time are the Normal, Negative Exponential and Poisson. This is almost certainly because they are widely used in other statistical applications and are comparatively well understood and documented. In the field of Stock Control the generalised application of any one of these distributions is unlikely to be tenable as each is suitable for specific demand characteristics, viz. the Normal for fast-selling products, the Negative Exponential for slow-movers, and the Poisson for products which sell in single units. Also, the mathematical properties of these distributions (for Poisson, $\mu = \sigma^2$; for Negative Exponential, $\mu = \sigma$; Normal extends to $\pm \infty$) are extremely unlikely to apply across a heterogeneous range of products.

A number of more versatile distributions have been propounded, such as the Gamma, Lognormal, Weibull, and Compound Poisson. The Inventory Management system uses the Gamma distribution mainly because Johnston (60) has demonstrated its suitability for a significant number of Builders Merchants. Burgin and Wild (61) and Burgin (62) have demonstrated the suitability of the Gamma distribution for Stock Control applications.

The Gamma is a 'variable shape' distribution described by a shape parameter (or modulus), k, and a scale parameter, α. The profile of the distribution for low values of k is given in Fig. 3.7. The important parameter, k, will be shown to be equal to the square of the Mean divided by the Variance, hence it is a measure of the consistency of the variates.

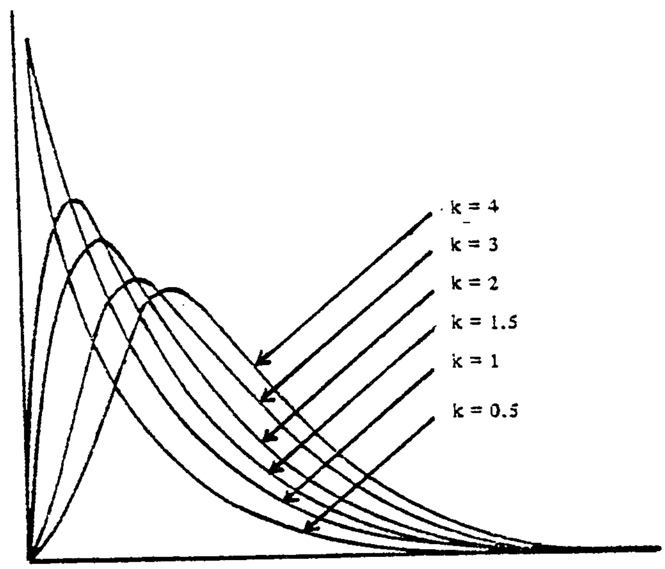

Fig. 3.7 Shape of the Gamma Distribution

For $0 < k \leq 1$ the distribution is monotonic decreasing. In the special case of k = 1 the Gamma reduces to a Negative Exponential. For $k > 1$ it is unimodal with a positive skew. As k is increased the mode centralises until as $k \to \infty$ the distribution has been shown (62) to be a very close approximation to the Normal. Generally, as the demand rate of a product increases, its k value increases, as the ratio of the Mean to the Variance becomes greater. Hence the Gamma retains the characteristics of the Normal and Negative Exponential for fast-movers and slow-movers respectively whilst avoiding the sharp dichotomy for intermediate movement rates.

A Gamma variate is defined as:

$$f(x) = \frac{\alpha^k x^{k-1} e^{-\alpha x}}{\Gamma(k)} \text{ for } 0 \leqslant x \leqslant \infty; \alpha>0; k>0 \qquad (11)$$

where $\Gamma(k)$ is the complete Gamma function:

$$\Gamma(k) = \int_0^\infty \alpha^k x^{k-1} e^{-\alpha x} dx$$

An important property of $\Gamma(k)$ is that:

$$\Gamma(k) = (k-1)\,\Gamma(k-1)$$

$$\Gamma(k+1) = k\,\Gamma(k)$$

$$\Gamma(k+2) = (k+1)\,\Gamma(k+1) = (k+1)k\,\Gamma(k)$$

$$\vdots$$

$$\Gamma(k+n) = (k+n-1)\,\Gamma(k+n-1) = (k+n-1)(k+n-2) \ldots\ldots\ldots\ldots k\,\Gamma(k)$$

As $\Gamma(1) = 1$, for integer values of k, $\Gamma(k) = (k-1)!$

By substituting k = 1 into (11),

$$f(x) = \frac{\alpha^1 x^0 e^{-\alpha x}}{\Gamma(1)} = \frac{\alpha e^{-\alpha x}}{0!} = \alpha e^{-\alpha x}$$

Hence, for k = 1, the Gamma reduces to a Negative Exponential. For all other values of k, the Gamma is a k-fold convolution of the Negative Exponential i.e. if $y = x_1 + x_2 + x_3 + \ldots\ldots\ldots x_k$ and x_i are all drawn from a Negative Exponential distribution with parameter α, then y will be drawn from a Gamma distribution with parameters k, α.

The incomplete Gamma integral, which represents the Protection Level (P) in inventory control theory, is defined as:

$$\int_o^R f(x)\,dx = \int_o^R \frac{\alpha^k x^{k-1} e^{-\alpha x}}{\Gamma(k)}\,dx \qquad (12)$$

For mathematical tractability, the following transformation is often employed:

Let $(\alpha x) = v$

Then $\frac{dv}{dx} = \alpha$ and $dv = \alpha dx$

Hence, $P = \int_o^{\alpha R} \frac{v^{k-1} e^{-v} dv}{\Gamma(k)}$

It will be shown in (17) that $\alpha R = U\sqrt{k}$ where U is a standardised reorder level expressed in standard deviations of demand in lead time, i.e.

$$U = \frac{R}{\sigma} \qquad (13)$$

$$\text{Hence, } P = \int_o^{U\sqrt{k}} \frac{v^{k-1} e^{-v} dv}{\Gamma(k)} \qquad (14)$$

The parameters k and α can be expressed in terms of the first two moments of the distribution, μ and σ, as follows:-

The mean of a distribution is the first moment about the origin,

$$\therefore \quad \mu = \int_o^\infty x \; f(x) \; dx$$

$$= \int_o^\infty \frac{x\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)}$$

But $\quad \Gamma(k) = \dfrac{\Gamma(k+1)}{k}$

$$\therefore \; \mu = \frac{k}{\alpha} \int_o^\infty \frac{\alpha^{k+1} x^k e^{-\alpha x} dx}{\Gamma(k+1)}$$

Since $\quad \displaystyle\int_o^\infty \frac{\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)} = 1$

Then by increasing the value of k by 1,

$$\int_o^\infty \frac{\alpha^{k+1} x^k e^{-\alpha x} dx}{\Gamma(k+1)} = 1$$

Hence, $\quad \mu = \dfrac{k}{\alpha}$

The variance of a distribution is the second moment about the mean

$$\therefore \quad \sigma^2 = \int_o^\infty \frac{x^2\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)} - \mu^2$$

But $\Gamma(k) = \dfrac{\Gamma(k+2)}{k(k+1)}$

$$\therefore\ \sigma^2 = \frac{k(k+1)}{\alpha^2} \int_0^\infty \frac{\alpha^{k+2} x^{k+1} e^{-\alpha x} dx}{\Gamma(k+2)} - \mu^2$$

Since $\displaystyle\int_0^\infty \frac{\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)} = 1$

Then by increasing the value of k by 2,

$$\int_0^\infty \frac{\alpha^{k+2} x^{k+1} e^{-\alpha x} dx}{\Gamma(k+2)} = 1$$

$$\therefore\ \sigma^2 = \frac{k(k+1)}{\alpha^2} - \mu^2$$

$$\therefore\ \sigma^2 = \frac{k(k+1)}{\alpha^2} - \frac{k^2}{\alpha^2} = \frac{k}{\alpha^2}$$

$$\therefore\ k = \alpha^2\sigma^2 = \frac{k^2}{\mu^2}\sigma^2$$

$$\therefore\ k = \frac{\mu^2}{\sigma^2} \qquad (15)$$

And, $\alpha = \dfrac{k}{\mu} = \dfrac{\mu^2}{\sigma^2\mu}$

$$\therefore\ \alpha = \frac{\mu}{\sigma^2} \qquad (16)$$

From (13), (15) & (16),

$$\alpha R = \frac{\mu}{\sigma^2} U\sigma = \frac{\mu}{\sigma} U = U\sqrt{k} \qquad (17)$$

An expression for calculating the Protection Level is here derived using the technique of expanding the integral proposed by Winters (63).

From (14), $P = \frac{1}{\Gamma(k)} \int_o^{U\sqrt{k}} v^{k-1} e^{-v} dv$

Introducing $e^{-U\sqrt{k}} e^{U\sqrt{k}}$ (= 1) into the expression,

$$P = \frac{e^{-u\sqrt{k}}}{\Gamma(k)} \int_o^{U\sqrt{k}} e^{U\sqrt{k}-v} v^{k-1} dv$$

But, $e^{U\sqrt{k}-v} = 1 + \frac{(U\sqrt{k}-v)}{1!} + \frac{(U\sqrt{k}-v)^2}{2!} + \ldots\ldots\ldots$

$$= \sum_{n=o}^{n=\infty} \frac{(U\sqrt{k}-v)^n}{n!}$$

$$\therefore P = \frac{e^{-U\sqrt{k}}}{\Gamma(k)} \int_o^{U\sqrt{k}} \sum_{n=o}^{n=\infty} \frac{v^{k-1}(U\sqrt{k}-v)^n dv}{n!}$$

$$= \frac{e^{\cdot U\sqrt{k}}}{\Gamma(k)} \int_o^{U\sqrt{k}} \sum_{n=o}^{n=\infty} \frac{v^{k-1}(U\sqrt{k}-v)^n (k+n) dv}{n!(k+n)}$$

$$= \frac{e^{-U\sqrt{k}}}{\Gamma(k)} \int_o^{U\sqrt{k}} \sum_{n=o}^{n=\infty} \frac{v^{k-1}(U\sqrt{k}-v)^{n-1}(U\sqrt{k}-v)(k+n) dv}{n!(k+n)}$$

$$= \frac{e^{-U\sqrt{k}}}{\Gamma(k)} \int_o^{U\sqrt{k}} \sum_{n=o}^{n=\infty} \frac{v^{k-1}(U\sqrt{k}-v)^{n-1}\left[k(U\sqrt{k}-v)+n(U\sqrt{k}-v)\right] dv}{n!(k+n)}$$

$$= \frac{e^{-U\sqrt{k}}}{\Gamma(k)} \int_o^{U\sqrt{k}} \sum_{n=o}^{n=\infty} \frac{kv^{k-1}(U\sqrt{k}-v)^n + v^{k-1}(U\sqrt{k}-v)^{n-1} nU\sqrt{k} - v^{k-1}(U\sqrt{k}-v)^{n-1} nv\, dv}{n!(k+n)}$$

$$= \frac{e^{-U\sqrt{k}}}{\Gamma(k)} \sum_{n=o}^{n=\infty} \frac{\int_o^{U\sqrt{k}} kv^{k-1}(U\sqrt{k}-v)^n dv + \int_o^{U\sqrt{k}} nU\sqrt{k} v^{k-1}(U\sqrt{k}-v)^{n-1} dv - \int_o^{U\sqrt{k}} nv^k (U\sqrt{k}-v)^{n-1} dv}{n!(k+n)}$$

Let $y = (U\sqrt{k}-v)^n$

$$\therefore \quad \frac{dy}{dv} = -n(U\sqrt{k}-v)^{n-1}$$

And, let $dz = kv^{k-1}dv$

$$\therefore \quad z = v^k$$

The first integral in the expression is $\int y\,dz$ and the third integral is $\int z\,dy$. Integrating by parts,

$$\int ydz + \int zdy = yz$$

$$\therefore P = \frac{e^{-U\sqrt{k}}}{\Gamma(k)} \sum_{n=o}^{n=\infty} \frac{v^k(U\sqrt{k}-v)^n + nU\sqrt{k}\int_o^{U\sqrt{k}} v^{k-1}(U\sqrt{k}-v)^{n-1}dv}{n!\ (k+n)}$$

Expanding the series,

$$P = \frac{e^{-U\sqrt{k}}}{\Gamma(k)} \Bigg[\frac{v^k}{k} + \frac{v^k(U\sqrt{k}-v) + U\sqrt{k}\int_o^{U\sqrt{k}} v^{k-1}dv}{(k+1)}$$

$$+ \frac{v^k(U\sqrt{k}-v)^2 + 2U\sqrt{k}\int_o^{U\sqrt{k}} v^{k-1}(U\sqrt{k}-v)\,dv}{2!\ (k+2)}$$

$$+ \frac{v^k(U\sqrt{k}-v)^3 + 3U\sqrt{k}\int_o^{U\sqrt{k}} v^{k-1}(U\sqrt{k}-v)^2dv}{3!\ (k+3)}$$

$$+ \ \ldots\ldots\ldots\ldots \Bigg]_{v=o}^{v=U\sqrt{k}}$$

$$= \frac{e^{-U\sqrt{k}}}{\Gamma(k)} \left[\frac{v^k}{k} + \frac{v^k(U\sqrt{k}-v) + U\sqrt{k}\ v^k/k}{(k+1)} \right.$$

$$+ \frac{v^k(U\sqrt{k}-v)^2 + 2U\sqrt{k}(v^k U\sqrt{k}/k - v^{k+1}/(k+1))}{2!\ (k+2)}$$

$$+ \frac{v^k(U\sqrt{k}-v)^3 + 3U\sqrt{k}((U\sqrt{k})^2 v^k/k - 2U\sqrt{k}\ v^{k+1}/(k+1) + v^{k+2}/(k+2))}{3!\ (k+3)}$$

$$\left. + \ldots\ldots\ldots\ldots \right]_{v=o}^{v=U\sqrt{k}}$$

At the upper limit, all terms containing $(U\sqrt{k}-v)$ are eliminated; at the lower limit, all terms are eliminated. The expression therefore simplifies to:

$$P = \frac{e^{-U\sqrt{k}}}{\Gamma(k)} \left[\frac{(U\sqrt{k})^k}{k} + \frac{(U\sqrt{k})^{k+1}}{k(k+1)} + \frac{(U\sqrt{k})^{k+2}}{k(k+2)} - \frac{(U\sqrt{k})^{k+2}}{(k+1)(k+2)} + \frac{(U\sqrt{k})^{k+3}}{2k(k+3)} \right.$$

$$\left. - \frac{(U\sqrt{k})^{k+3}}{(k+1)(k+3)} + \frac{(U\sqrt{k})^{k+3}}{2(k+2)(k+3)} + \ldots\ldots\ldots\ldots \right]$$

$$= \frac{e^{-U\sqrt{k}}}{\Gamma(k)}\left[\frac{(U\sqrt{k})^{k}}{k} + \frac{(U\sqrt{k})^{k+1}}{k(k+1)} + \frac{(U\sqrt{k})^{k+2}}{k(k+1)(k+2)} + \frac{(U\sqrt{k})^{k+3}}{k(k+1)(k+2)(k+3)} + \ldots\ldots\right]$$

$$= \frac{e^{-U\sqrt{k}}(U\sqrt{k})^{k-1}}{\Gamma\ (k)}\left[\frac{(U\sqrt{k})}{k} + \frac{(U\sqrt{k})^{2}}{k(k+1)} + \frac{(U\sqrt{k})^{3}}{k(k+1)(k+2)} + \frac{(U\sqrt{k})^{4}}{k(k+1)(k+2)(k+3)} + \ldots\right]$$

$$\therefore\ P = \frac{e^{-U\sqrt{k}}(U\sqrt{k})^{k-1}}{\Gamma\ (k)} \sum_{n=o}^{n=\infty} \frac{(U\sqrt{k})^{n+1}}{\prod_{i=o}^{i=n} (k+i)} \qquad (18)$$

which is essentially the same expression as that obtained by Winters.

The term Γ (k) in the denominator can be evaluated using a series expansion, several of which are available. The following, due to Hastings (64), is accurate to within ±0.0000013:

$$\Gamma(1 + x) = 1 + a_1 x + a_2 x^2 + a_3 x^3 + a_4 x^4 + a_5 x^5 + a_6 x^6 + a_7 x^7$$

for $0 \leqslant x \leqslant 1$

where, $a_1 = -0.57710166$

$a_2 = +0.98585399$

$a_3 = -0.87642182$

$a_4 = +0.83282120$

$a_5 = -0.56847290$

$a_6 = +0.25482049$

$a_7 = -0.05149930$

$\Gamma(x) = \frac{\Gamma(1 + x)}{x}$ where x is the fractional part of k.

For non-integral values of k > 1,

$\Gamma(n.x) = (n + x - 1)(n + x - 2) \ldots\ldots\ldots (x + 1)\,\Gamma(x + 1)$

where n is the integer part of k.

Computer programs were written to evaluate $\Gamma(k)$ and P which are listed in Appendix 2. It was found that for $0.1 \leq k \leq 6$ and $0.5 \leq U \leq 6$, the evaluation of P using the first 30 terms of equation 18 agrees with Pearson's Tables (65) of the Incomplete Gamma Function to within ± 0.000003 with a Mean Absolute Error of zero to the sixth significant digit.

The following derivation of an expression for Percentage Lost Sales in lead time (PLS) is based on a method by Burgin (62).

$$PLS = \frac{ELS}{\mu} = \frac{1}{\mu}\int_R^\infty (x-R)\ f(x)\ dx$$

Where μ = mean demand in lead time

Substituting from 11,

$$PLS = \frac{1}{\mu} \int_R^\infty \frac{(x-R)\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)}$$

$$= \frac{1}{\mu}\left[\int_0^\infty \frac{(x-R)\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)} - \int_0^R \frac{(x-R)\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)}\right]$$

$$= \frac{1}{\mu}\left[\int_0^\infty \frac{x\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)} - R\int_0^\infty \frac{\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)}\right.$$

$$\left. - \int_0^R \frac{x\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)} + R\int_0^R \frac{\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)}\right]$$

But, $\Gamma(k) = \dfrac{\Gamma(k+1)}{k}$

$$\therefore \quad PLS = \frac{1}{\mu}\left[\frac{k}{\alpha}\int_o^\infty \frac{\alpha^{k+1}x^k e^{-\alpha x}dx}{\Gamma(k+1)} - R\int_o^\infty \frac{\alpha^k x^{k-1}e^{-\alpha x}dx}{\Gamma(k)}\right.$$

$$\left. - \frac{k}{\alpha}\int_o^R \frac{\alpha^{k+1}x^k e^{-\alpha x}dx}{\Gamma(k+1)} + R\int_o^R \frac{\alpha^k x^{k-1}e^{-\alpha x}dx}{\Gamma(k)}\right]$$

$$= \frac{1}{\mu}\left[\frac{k}{\alpha} - R - \frac{k}{\alpha}\int_o^R \frac{\alpha^{k+1}x^k e^{-\alpha x}dx}{\Gamma\ (k+1)} + R\int_o^R \frac{\alpha^k x^{k-1}e^{-\alpha x}dx}{\Gamma\ (k)}\right]$$

Grouping the terms in the first integral to integrate by parts,

Let $U = \dfrac{\alpha^{k+1}x^k}{\Gamma(k+1)}$

$$\therefore \quad \frac{dU}{dx} = \frac{\alpha^{k+1}kx^{k-1}}{\Gamma(k+1)} = \frac{\alpha^{k+1}x^{k-1}}{\Gamma(k)}$$

And let $dv = e^{-\alpha x}dx$

$$\therefore \quad v = -\frac{e^{-\alpha x}}{\alpha}$$

$$\therefore \quad \int_o^R \frac{\alpha^{k+1}x^k e^{-\alpha x}dx}{\Gamma(k+1)} = \int_o^R Udv$$

Integrating by parts,

$$\int_o^R Udv = Uv - \int_o^R v\ dU$$

$$= \left[- \frac{\alpha^{k+1} x^k e^{-\alpha x}}{\Gamma(k+1)\alpha} \right]_{x=o}^{x=R} + \int_o^R \frac{e^{-\alpha x}\alpha^{k+1} x^{k-1} dx}{\Gamma(k)\alpha}$$

$$= - \frac{(\alpha R)^k e^{-\alpha R}}{\Gamma(k+1)} + \int_o^R \frac{\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)}$$

$$= - \frac{(\alpha R)^k e^{-\alpha R}}{\Gamma(k+1)} + P \qquad \text{(from (12))}$$

$$\therefore \ PLS = \frac{1}{\mu} \left[\frac{k}{\alpha} - R - \frac{k}{\alpha} \left(- \frac{(\alpha R)^k e^{-\alpha R}}{\Gamma(k+1)} + P \right) + RP \right]$$

Substituting from (15), (16) & (17),

$$PLS = (1-U/\sqrt{k})(1-P) + \frac{(U\sqrt{k})^k e^{-U\sqrt{k}}}{\Gamma\ (k+1)} \qquad (19)$$

A computer program was written to evaluate PLS from this expression, which is listed in Appendix 2. The resulting values correspond exactly to a table computed by Burgin and Norman (66) to the four decimal places tabulated except in a few randomly-scattered cases where there was a difference of ±0.0001.

Equations (18) and (19) satisfy the theoretical requirements of the system. However, they are extremely cumbersome to use on a routine basis, and also an inversion would be required to calculate U given P and k. Instead, the following numerical approximations, derived by Johnston (58), were employed.

Function 1 - to calculate U given P and k

$$U = A0 + A1 \log(1-P) + A2 (1-P)^2 + A3 (1-P) \log(1-P)$$

$$\text{where } A0 = 0.0106179 - 0.0156841\,k^2 + 1.66011 \log k - 0.365992 (\log k)^2 + 0.145241\,k \log k$$

$$A1 = -0.998223 - 0.00231704\,k^2 + 0.357714 \log k - 0.106577 (\log k)^2 + 0.0201662\,k \log k$$

$$A2 = -1.48338 - 0.000741918\,k^2 + 1.46426/k - 0.206282 \log k$$

$$A3 = 2.76031 - 2.72033\,k - 0.0544844\,k^2 + 3.13504 \log k + 1.04581\,k \log k$$

Function 2 - to calculate PLS given P and k

$$PLS = 0.01 \left(A1 (1 - P) + A2 (1 - P)^2\right)$$

$$\text{where } A1 = 9.4608205 + 101.30969/k - 9.5595537/k^2$$

$$A2 = 20.574471 + 9.9995001/k - 27.350124/k^2$$

Function 3 - to calculate PLS given U and k

$$PLS = 0.01 \left(A1\,e^{-U} + A2\,Ue^{-U} + A3 (e^{-U})^2\right)$$

$$\text{where } A1 = 26.684318 + 84.664245\,k - 14.77837\,k^2 + 2.3625622\,k^3 - 0.10083282\,k^4$$

$$A2 = 42.992034 - 51.061508\,k + 12.425395\,k^2 - 1.53408\,k^3 + 0.0635186\,k^4$$

$$A3 = 29.0487 - 26.4592\,k - 1.8668146\,k^2 - 1.0568659\,k^3$$

All expressions were derived for $0.5 \leqslant k \leqslant 12$; $0.5 \leqslant P \leqslant 0.999$.

All logarithms are to base e.

Johnston (58,60) reported the following levels of accuracy for the functions:

Function 1 supplies U to within 0.004, with a Mean Absolute Deviation of 0.0026 on U values from 1.2 to 7.5.

Function 2 supplies PLS to within 0.0015, with a Mean Absolute Deviation of 0.0009 (or 3.7% of the true PLS) for PLS values from 0.0002 to 0.2.

Function 3 supplies PLS to within 0.005 with a Mean Absolute Deviation of 0.0029. At values of the PLS > 0.01 the average error corresponds to 6.2% of the true PLS.

The system calculates Mean Absolute Deviation of demand in lead time rather than the Standard Deviation (for ease of computation) then converts to Standard Deviation by using a constant factor. The relationship between MAD and σ is derived below:

$$\text{MAD} = |\mu - x| \quad \text{for all values of } x$$

$$\therefore \text{MAD} = \int_0^{\mu} (\mu - x) f(x)\,dx + \int_{\mu}^{\infty} (x-\mu) f(x)\,dx$$

$$= \mu\int_0^{\mu} f(x)\,dx - \int_0^{\mu} xf(x)\,dx + \int_{\mu}^{\infty} xf(x)\,dx - \mu\int_{\mu}^{\infty} f(x)\,dx$$

$$= 2\mu\int_0^{\mu} f(x)\,dx - \mu\int_0^{\mu} f(x)\,dx - \mu\int_{\mu}^{\infty} f(x)\,dx - 2\int_0^{\mu} xf(x)\,dx$$

$$+ \int_0^{\mu} xf(x)\,dx + \int_{\mu}^{\infty} xf(x)\,dx$$

$$= 2\mu\int_0^{\mu} f(x)\,dx - \mu\int_0^{\infty} f(x)\,dx - 2\int_0^{\mu} xf(x)\,dx + \int_0^{\infty} xf(x)\,dx$$

But, $\int_0^{\infty} f(x)dx = 1$, and $\int_0^{\infty} xf(x)dx = \mu$

$$\therefore \text{ MAD } = 2\mu \int_0^{\mu} f(x)dx - \mu - 2\int_0^{\mu} xf(x)dx+\mu$$

$$= 2\mu \int_0^{\mu} \frac{\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma\ (k)} - 2\int_0^{\mu} \frac{x\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma\ (k)}$$

$$= 2\mu \int_0^{\mu} \frac{\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma\ (k)} - 2\ \frac{k}{\alpha} \int_0^{\mu} \frac{\alpha^{k+1} x^k e^{-\alpha x} dx}{\Gamma\ (k+1)}$$

$$= 2\mu \left[\int_0^{\mu} \frac{\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma\ (k)} - \int_0^{\mu} \frac{\alpha^{k+1} x^k e^{-\alpha x} dx}{\Gamma\ (k+1)} \right]$$

The second integral is seen to be of the same form as the first, with k replaced by k+1.

Integrating $\int_0^{\mu} \frac{\alpha^{k+1} x^k e^{-\alpha x} dx}{\Gamma(k+1)}$ as in the derivation of (19) to make the second integral identical to the first:

$$\text{MAD} = 2\mu \left[\int_0^{\mu} \frac{\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)} + \frac{(\alpha\mu)^k e^{-\alpha\mu}}{\Gamma(k+1)} - \int_0^{\mu} \frac{\alpha^k x^{k-1} e^{-\alpha x} dx}{\Gamma(k)} \right]$$

$$= \frac{2\mu(\alpha\mu)^k e^{-\alpha\mu}}{\Gamma\ (k+1)}$$

Substituting μ and σ from (15) and (16):

$$MAD = \frac{2k^k \sigma}{\sqrt{k}\Gamma(k)e^k}$$

$$\text{or,} \quad \frac{\sigma}{MAD} = \frac{\sqrt{k}\Gamma(k)e^k}{2k^k} \qquad (20)$$

Using the Hastings expansion for $\Gamma(k)$, the conversion factors for MAD to Standard Deviation are given in Table 3.1.

k	σ/MAD
0.1	2.0929
0.2	1.7299
0.3	1.5870
0.4	1.5097
0.5	1.4611
0.6	1.4278
0.7	1.4036
0.8	1.3852
0.9	1.3708
1.0	1.3591
1.2	1.3416
1.4	1.3290
1.6	1.3195
1.8	1.3121
2.0	1.3062
2.5	1.2956
3.0	1.2885
4.0	1.2796
5.0	1.2743
6.0	1.2708
7.0	1.2683
8.0	1.2664
9.0	1.2650
10.0	1.2640
11.0	1.2628
12.0	1.2620
∞	1.2534

c f 1.25 for Normal Distribution

Table 3.1 Conversion Factors for MAD to Standard Deviation

3.5 PERFORMANCE APPRAISAL

There are provisions in the system for monitoring performance at Product Group, Buying Family and Product level. The most important measure of total system performance is a regular comparison at Product Group level between the target or forecast values and the achieved values of the following parameters:

i) Service Level

Target service levels for each Product Group are set at the beginning of each monitoring period (normally six months) to represent the percentage of total costed demand which must be satisfied at the first attempt. The achieved service levels are calculated as the proportion of days for which each product has been in stock, weighted according to their respective costed demands.

ii) Demand

The forecast demand for each Product Group is the costed sum of exponentially-weighted moving averages of product issue data. The EWMA's are updated weekly only if the respective products have not been out of stock at any time during the week (hence issues and demand are synonymous). The achieved demand is computed as the sum of the total costed issues during the monitoring period divided by the number of days the respective products were in stock, converted to an annual figure.

iii) Average Stock Value

The forecast average stock value is calculated as in equation 6. The achieved value is the costed sum of EWMA's of the closing stock balances at the end of each week.

iv) Stockturn

The forecast stockturn is computed as:

$$\frac{\text{Forecast demand p.a. - ELS p.a.}}{\text{Forecast Average Stock Value}}$$

The achieved stockturn is taken as the sum of the costed issues per annum divided by the achieved average stock value.

The report showing the above comparisons also contains a valuation of 'Excess Stock' at Product Group level. For each product, the excess stock is defined as the difference, if positive, between the physical stock balance and the Maximum Order Cover (MOC). The MOC is the reorder level plus the expected sales per order cycle i.e. it is the point which the computer orders up to. Excess stock can be caused either by the buyer exceeding the computer recommended order quantities, or by the MOC falling due to changes in demand, lead time etc. The buyers are provided with a monthly listing of the excess stocks at product level.

The system was introduced into the first Distribution Centre on a pilot basis in March 1979. Prior to the implementation a sample of 911 products (11.4%) was monitored for a period of six months whilst under the control of a manual recording system. This yielded the following results for the sample:-

Average Service Level	=	78.2%
Total Annual Sales	=	£2.261m
Total Average Stock	=	£0.408m
Average Stockturn	=	5.5

The period from March 1979 to September 1979 was taken up by monitoring the pilot exercise and loading the remaining products on to the system. The first system predictions were made in September

1979 and these were compared with the subsequent performance at March 1980. The comparison is given in Table 3.2. It can be seen that the overall service level was within an acceptable tolerance of the target value and this represented an improvement of 7.33% over the manual system as measured by the sample. Fig. 3.8, however, depicts a falling demand situation. This could be expected to result in higher than predicted service levels, as first-order exponential smoothing produces forecasts which lag the observed values when a trend is present, and this would cause overordering. Also there are significant discrepancies between the predicted and observed service levels at Product Group level which suggests compensating errors (Mean Absolute Error = 6.0%).

The achieved average stock value is seen to be almost double the predicted value. This can be analysed as follows:-

Average Stock Value	=	£2.432m
Excess Stock	=	£0.587m
∴ 'Effective' Stock	=	£1.845m
Theoretical Optimum	=	£1.251m
∴ Penalty for Imbalance	=	£0.594m

The counted stock value at the March 1980 stocktake was £2.44m, compared with £2.23m at the September 1979 stocktake. This represents a reduction of 1.8% in real terms, but this is well within the tolerance for counting errors. The general assessment of the system in March 1980 was that it had helped to improve service levels but made no appreciable impact on stock investment. It was recognised that stock reductions were not achievable in the short term as deficiencies would be made good immediately whereas overstocks would take some time to dissipate.

Product Group Code	Predicted Performance				Achieved Performance				
	Target SL	Av.Stock	Stockturn	ELS p.a.	SL	Av.Stock	Stockturn	Demand p.a.	Excess Stock
A1	90.0	94318	8.2	85505	82.7	214027	4.2	1093029	39435
B1	85.0	9979	12.3	21568	95.7	33115	5.4	187345	11926
B2	87.5	51477	9.5	69893	87.6	89956	7.5	769925	30735
B3	87.5	44862	6.6	42401	92.0	94076	4.8	485882	20437
C1	90.0	167756	5.8	107747	96.6	209789	5.1	1098654	52370
C2	87.5	127573	12.7	231338	84.7	278457	8.0	2616518	59627
C3	90.0	82282	6.9	63540	90.0	124436	8.9	1225857	16378
D1	85.0	233666	6.6	274085	92.1	438464	5.2	2496905	99408
D2	90.0	8569	8.4	7966	62.0	19636	4.2	131979	7492
D3	95.0	1611	7.2	615	98.7	5947	3.5	20871	2270
E1	85.0	205902	8.5	308944	70.6	474292	4.7	3162364	141167
F1	90.0	54524	16.6	100459	92.2	116522	7.6	955360	27505
G1	87.5	44847	13.2	84816	89.8	101705	7.4	833479	22214
G2	90.0	20312	15.1	34019	86.2	40980	8.8	419549	13510
G3	95.0	6418	7.7	2607	98.0	13694	3.8	53261	4253
G4	85.0	1362	14.2	3407	93.1	4080	6.7	29237	1776
H1	90.0	79222	8.2	72166	88.1	150264	4.9	843368	32448
I1	95.0	16278	8.9	7633	96.6	22482	5.4	126565	3661
Overall	87.76	1250958	8.71	1518709	85.53	2431922	5.82	16550149	586612

Table 3.2 Comparison between Predicted and Achieved Performance Sept 1979 - March 1980

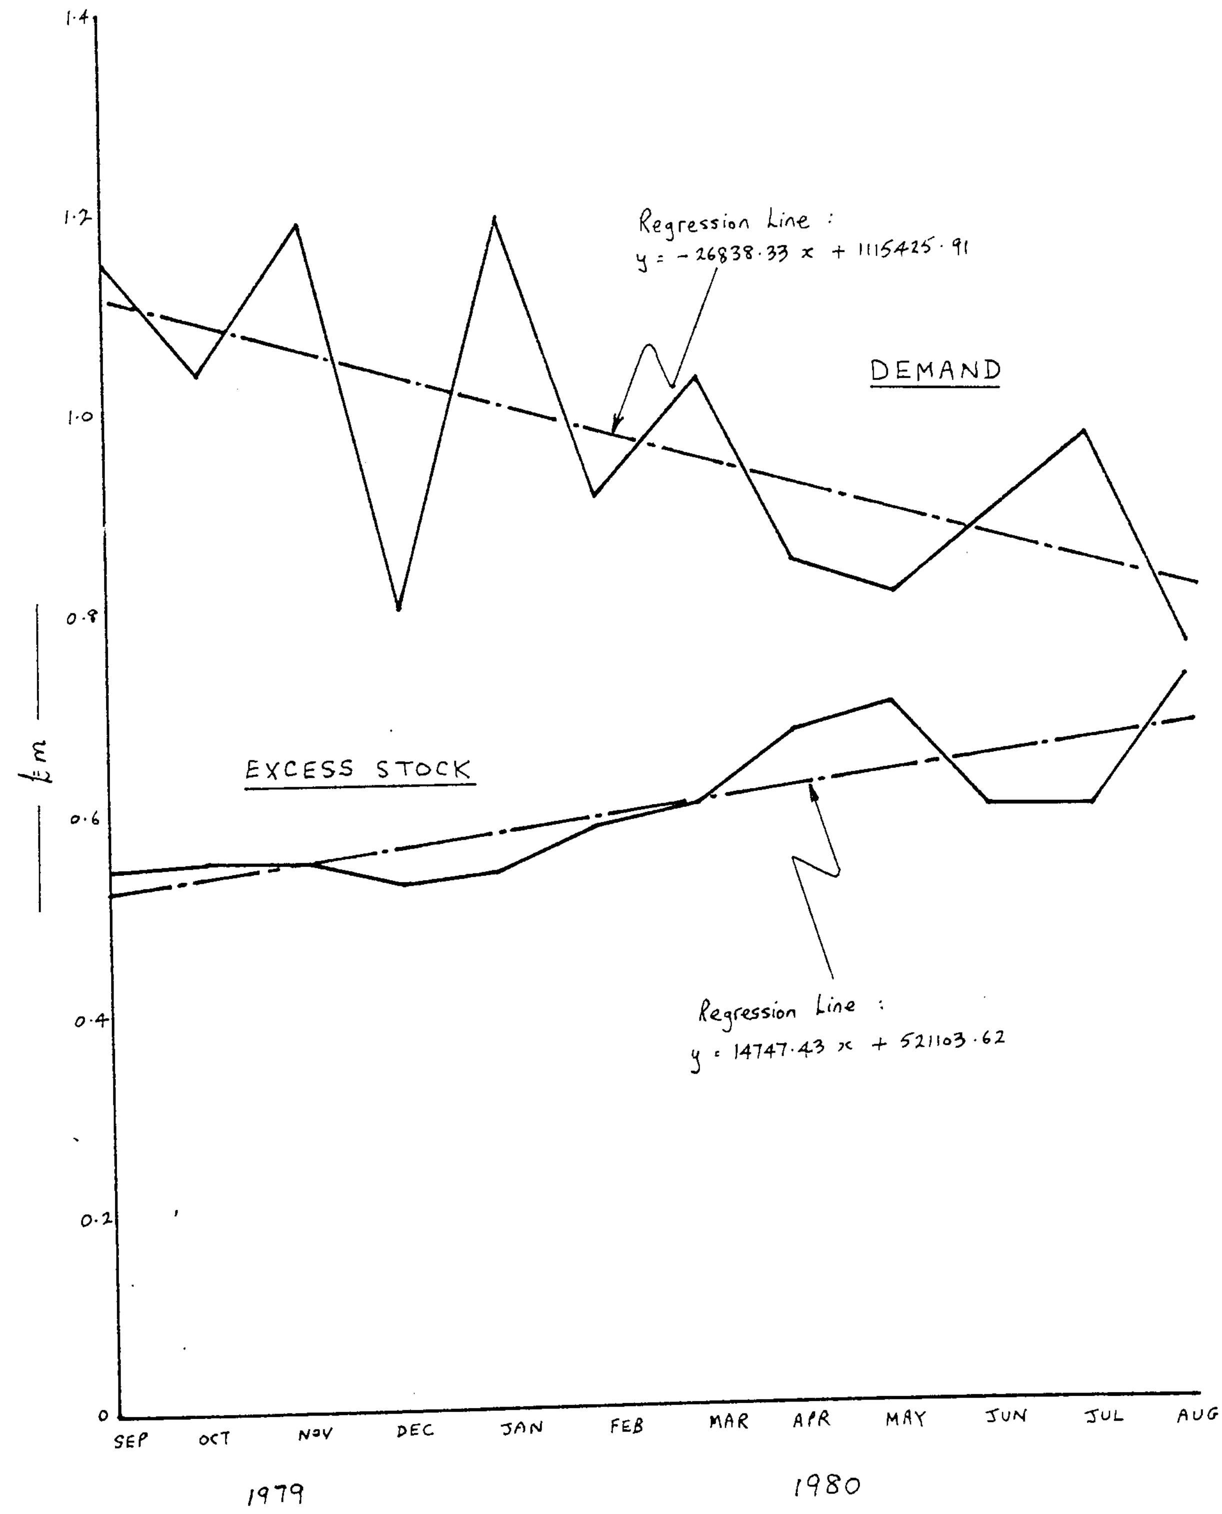

Fig. 3.8 Levels of Demand and Excess Stock Sept 1979 to Aug 1980

Two minor service level adjustments were made in March 1980 and a second comparison between predicted and achieved performance was made in September 1980 (Table 3.3). This showed a substantial improvement. The overall service level rose to 91.41% in a situation of further declining demand and the Mean Absolute Error of the Product Group service levels narrowed to 4.9%. The average stock value was 7.9% below that for the previous period (in real terms), analysed as follows:

Average Stock Value	=	£2.365m
Excess Stock	=	£0.743m
∴ 'Effective' Stock	=	£1.622m
Theoretical Optimum	=	£1.358m
∴ Penalty for Imbalance	=	£0.264m

The reduction in the imbalance allied to the closer adherence to service level targets is indicative of a much more effective degree of control. The increase in excess stock was found in a later investigation to be due to the intervention of the human element and this is explored more fully in Chapter 7.

A further performance review was carried out in March 1981 which revealed a further improvement in service level adherence and about the same level of excess and imbalanced stock as the second review. The other three Distribution Centres implemented the system in March 1980, November 1980 and September 1981 and their early results have followed a similar pattern to that experienced by the first Distribution Centre.

Product Group	Predicted Performance				Achieved Performance				
Code	Target SL	Av.Stock	Stockturn	ELS p.a.	SL	Av.Stock	Stockturn	Demand p.a.	Excess Stock
A1	90.0	130614	5.9	85308	92.9	237496	3.9	996471	55951
B1	85.0	11262	11.7	23318	93.9	29275	5.1	160343	11460
B2	87.5	48500	11.7	80892	89.1	90596	9.3	947142	13095
B3	85.0	45523	7.6	60846	96.7	84911	3.7	326817	26584
C1	90.0	77103	11.5	98267	97.3	150690	6.6	1029908	40078
C2	87.5	169735	10.5	253826	90.8	279759	7.0	2157214	71883
C3	90.0	106057	9.7	114087	91.2	100551	9.6	1062723	14494
D1	85.0	261675	7.1	327215	94.6	404672	4.0	1719772	172334
D2	90.0	7601	8.4	7089	88.1	23058	3.7	98071	4932
D3	95.0	4231	4.3	950	97.8	7663	2.5	19989	4458
E1	85.0	220102	7.6	296917	86.3	499187	4.0	2289504	221442
F1	90.0	61360	11.6	79362	93.9	96716	11.2	1153527	11778
G1	87.5	53423	11.1	84538	92.5	116796	6.4	810033	30386
G2	90.0	51336	6.3	36036	83.6	46873	9.1	507241	8108
G3	90.0	5862	7.5	4869	94.8	10673	4.7	52964	3541
G4	85.0	1162	18.3	3767	95.8	4517	6.2	29442	1672
H1	90.0	88144	6.9	67786	91.8	157606	4.8	831373	45234
I1	95.0	14274	7.2	5396	92.0	24117	6.3	166113	5518
Overall	87.61	1357964	8.49	1630469	91.41	2365156	5.55	14358647	742858

Table 3.3 Comparison between Predicted and Achieved Performance March 1980 - Sept 1980

CHAPTER 4

ANALYSIS OF SYSTEM BEHAVIOUR

4.1 SYSTEM DYNAMICS STUDY

4.1.1 PRINCIPLES OF SYSTEM DYNAMICS

In spite of the pervasiveness of systems, studies of their dynamic behaviour are of quite recent origin. This applies particularly to systems containing people, as the behaviour is influenced by the voluntary actions of the individuals. The most practical advances in developing a methodology for studying system behaviour were made at the Massachusetts Institute of Technology (M.I.T.) from 1956 onwards. The first systems to be studied were industrial corporations and a significant milestone was reached in 1961 when Professor J.W. Forrester (35) published the results of the first phase of the research programme. This provided the impetus for further research, both horizontally into other industries, and vertically into different orders of systems. The M.I.T. group progressed to urban dynamics (67), which involved the interaction of local population, job opportunities, house availability and capital expenditure; and thence to world dynamics (68) which encompassed natural resources, population, quality of life, capital investment and pollution in a global perspective. With the extension of the scope of the work the original designation of 'Industrial Dynamics' was replaced by the broader term 'System Dynamics'.

In an industrial or business context, System Dynamics is concerned with the interrelationships between the flows of men, money, materials, orders and capital expenditure, together with the information and decision network which effects the rapprochement. The approach is based upon charting these flowpaths and representing the dynamic interactions by mathematical expressions which together constitute a model. The parameters and structure of the model can then be changed to test the effect of changes on the real system without incurring substantial expenditure and commercial risk.

The model consists of two fundamental building blocks - rates and levels, and these are interspersed along the flowpaths. The levels describe the state of the system variables at a point in time. The rates describe the speed of inflow to, or outflow from the levels. If the system is brought to rest the levels will be observable, the rates will not. In a Stock Control application, inventory is a clear example of a level, and the instantaneous rates of receipts and despatches are the rates which maintain this level. The levels are therefore net accumulations of rates, and the level equations perform an integration function. Calculus is not used for the integration functions, as analytical solutions are not always possible using linear differential equations (69), and the introduction of complex mathematics would remove the equation structure from the purview of most practising Managers. Instead, the M.I.T. group use the rectangular (Euler) method of integration which consists of recalculating rates after each 'solution interval' and accumulating these to form the levels. As the solution interval is fixed and does not tend to zero, this is not a true integration, but with a small solution interval relative to the time constants in the system it is a very good approximation. The differential equations required for calculus are replaced by difference equations which reduce the mathematics to simple algebra.

An important sub-structure of the model is the feedback loop. Complex systems could be aptly described as 'packages of feedbacks' and it is these which determine the internal behaviour of the system. Four of the most important types of feedback loop have been identified by Forrester (70), and these are depicted in Fig. 4.1.

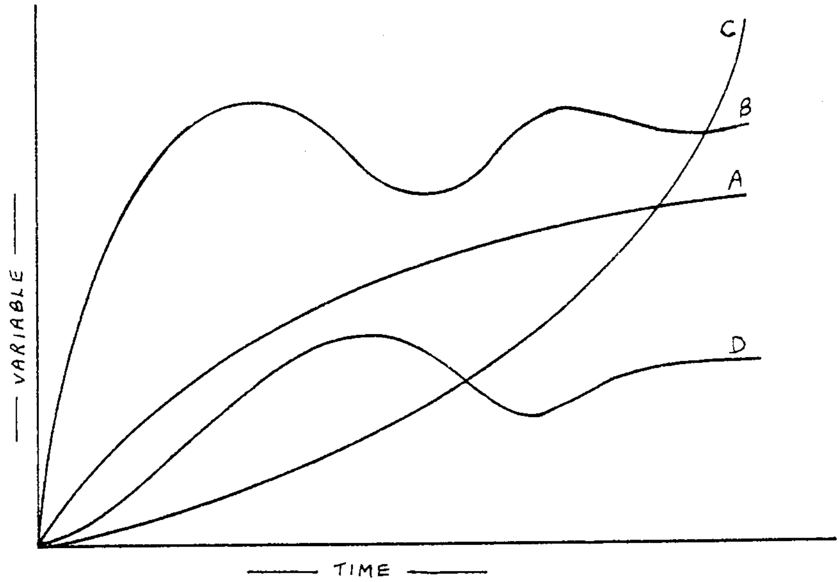

Fig. 4.1 Types of Feedback Loop

Curve A shows the time response of the variable approaching its final value without overshoot or oscillation. This is typical of a first-order negative feedback loop. 'First-order' implies a single level variable with an associated rate, whereas higher order structures contain a concatenation of levels and rates.

Curve B displays an oscillation of diminishing amplitude before the variable attains its stable value. This represents a second (or higher) - order negative feedback loop. The oscillation is caused by delaying effects introduced by other levels in the system e.g. if an inventory level begins to fall below its desired value, the ordering rate is increased and the level of unfilled orders at the supply source rises. This results in an acceleration of the supply rate which duly corrects the inventory deficit. The ordering rate is cut back and after a time lag the unfilled order pool falls and the supply rate is decelerated. During this reaction time the supply rate exceeds the ordering rate and the inventory overshoots its desired value. The sequence of actions is then repeated with the opposite polarities and the inventory level oscillates around its desired value.

Curve C is an exponential curve with an increasing gradient. This depicts a growth characteristic which is engendered by a positive feedback mechanism.

Curve D is a composite of curves C and B. This suggests two or more coupled loops. The initial time response assumes an exponential profile which suggests a dominant positive feedback loop. The early growth is constrained by some non-linear condition such as market saturation or the attainment of production capacity. A second-order negative feedback then predominates and the variable achieves a steady state. More complex profiles would be obtained with a more intricate network of feedback loops. The basic profiles are nevertheless discernible in most structures.

With a short solution interval, the independent variable, time, is incremented in small fixed steps between recomputations. The flows are therefore analagous to liquids flowing through tubes with reservoirs interposed. In essence, the model effects a continuous system simulation, and rapid changes to variables, especially discontinuities, introduce potential inaccuracies into the fixed step integration procedure. In the real system, however, the changes in the variables are fundamentally discrete - purchase orders are accumulated until 'best terms' are available, receipts arrive in lorry loads, taps are sold in pairs, wallpaper is sold in 'room lots', etc. The variables therefore experience quantum changes rather than smooth gradations. Forrester (35) argues that "real systems are more nearly continuous than is commonly supposed" on the grounds that aggregation of products and transactions smoothes the flows to the extent that a continuous process is an effective first approximation. Aggregation does, however, dilute behaviour patterns especially if the component items are in different phases of their operating cycle. Aggregation should therefore aim to achieve a balance between smoothing individual actions and preserving the innate behavioural characteristics of the operation. Childs (71), in a mathematical analysis of aggregation, states that "aggregation of products within each of these two categories (stocked or produced to order) is therefore assumed to present relatively minor estimation difficulties" . Carnell (72), in an empirical System Dynamics study, has observed that quantising variables gives a less stable result with a greater amplification of swings and a shortening of frequency.

System Dynamics produces results which could be interpreted either quantitatively or qualitatively, depending upon the nature of the application and the formulation of the equations. Models can be broadly divided into Process Flow models, where tangible operations are represented in detail, and Policy models, where more abstract concepts based heavily on the information and decision network are represented. The degree of quantification which may be imputed to the results is clearly much higher for Process Flow models, and the dynamic behaviour is usually more orderly than that encountered in Policy models. Roberts (73), for example, has modelled the Research and Development process using relationships based upon qualitative curves represented in the model by tables of assigned numbers. The results obviously reflect the nature of the input and it would clearly be inappropriate to attempt to elicit the magnitude of the variables. Meyer and Roberts (74) assert that these models are highly beneficial in helping to formalise thought processes, particularly as the qualitative curves are often constructed from subjective interpretation of the system relationships. However abstract the subject being modelled, it is of prime importance that there is an identifiable structure with deterministic causal relationships which can be expressed numerically.

Several Continuous System Simulation Languages (CSSL's) have been developed to facilitate the execution of the models on digital computers. Carnell (72) provides an account of the features and historical development of CSSL's. The M.I.T. group developed the DYNAMO (DYNAmic MOdelling) language for the purpose of running the Industrial Dynamics models. The symbolic notation given in Fig. 4.2 was developed in conjunction with DYNAMO and the following conventions are used:

The subscript 'K' to a variable name denotes 'time now', 'J' denotes the time at the previous recalculation point, and 'L' the time at the next recalculation point. All levels are subscripted with J, K or L. Rates are subscripted by 'JK' to denote the solution interval up to 'time now', or 'KL' to denote the solution interval from 'time now' to the next recalculation point. The processing sequence is:-

a) Calculate levels at K using levels at J and rates over JK.

b) Calculate auxiliaries at K using levels at K and rates over JK.

c) Calculate rates over KL using levels and auxiliaries at K.

d) Index time forward by one solution interval and repeat.

The DYNAMO (75) language contains a number of useful features including a graph-plotting facility, extensive error checking and recovery routines, the resequencing of equations into the correct calculation order, an automatic documentation facility, and a number of functions such as generating random numbers and representing relationships by tables.

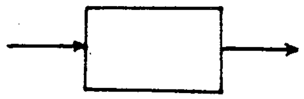

Level: Accumulation of materials or orders. The present value resulting from inflows and outflows. Describes the state of a variable at a point in time.

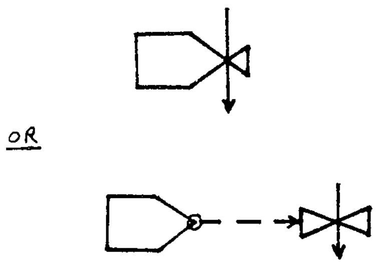

Flow Rate: Defines the instantaneous rate of flow between levels. The value is determined by Levels and, optionally, Auxiliaries.

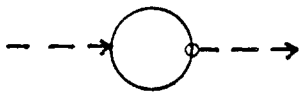

Auxiliary: A subordinate equation which is part of a rate formulation. Introduced for ease of comprehension and mathematical simplicity.

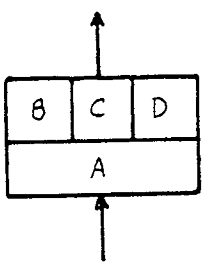

Delay: A shorthand representation of the cascaded levels and flows which constitute a delay.

A = total quantity in the delay

B = order of the delay

C = outflow rate

D = time constant

Material Flow

Order Flow

Information Flow

Source of Information

Fig. 4.2 Diagrammatic Notation

The DYNAMO CSSL was not used for the present study, but many of its features were written in to a specially developed simulation program which was coded in the standard BASIC language for operation on a microprocessor. This approach was adopted purely for computing convenience. The equations were written using the DYNAMO mnemonic conventions then translated into BASIC notation. The fundamental process of calculating difference equations iteratively is executed in exactly the same manner as the DYNAMO CSSL without any degradation in accuracy.

4.1.2 SCOPE AND DEVELOPMENT OF THE MODEL

The system delimitation encompasses the Branch/Distribution Centre/Supplier hierarchy as shown in Fig. 4.3. The supplier is represented as a single level which is tantamount to a delay in the ordering process. As the supplier is included in the system, the market is the only source of exogenous variables.

The Inventory Management system controls the stocks at Distribution Centres. The dynamics of Distribution Centre stocks are therefore of prime interest. The Distribution Centres operate independently, thus the model can be restricted to a single Distribution Centre with a complement of 20 Branches which are of the same order of size. Orders on a Distribution Centre are split about equally between Branch replenishment orders and direct customer orders. Hence the system comprises a 2-level structure superimposed onto a 3-level structure. No attempt is made in the model to segregate these structures as they co-exist in the real system. Each Branch stocks a heterogeneous range of approximately 2,500 products but their demand patterns can all be adequately represented by a Gamma probability distribution (Chapters 3 and 5). The incorporation of multiple demand characteristics into the model would introduce undue

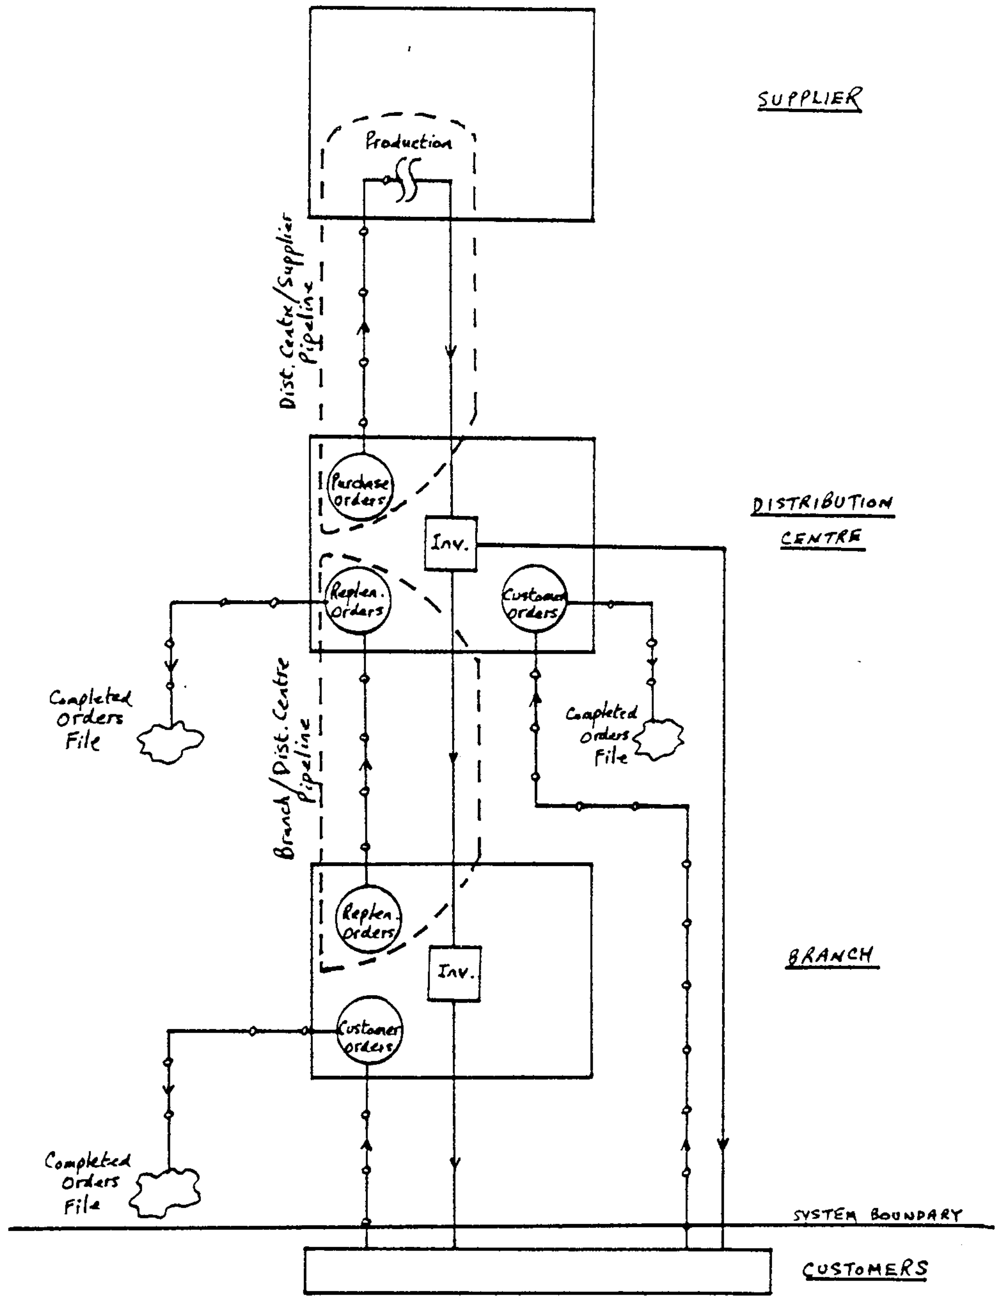

Fig. 4.3 Schematic of Order and Material Flows

complexity and probably obscure any dynamic behaviour inherent in the system. The approach is therefore to select a typical product and assume a homogeneous range. With this assumption, product aggregation is merely a scaling function with no particular advantage. The magnitude of the units input to the model is therefore at single product level but the concept is of many similar products. The selection of a 'typical' product is contentious. Intuitively, a product with a sales rate of around one unit per week per Branch would be considered typical in that this rate would be of the correct order of magnitude for a large proportion of the range (refer Fig. 3.3). On the other hand, a relatively small number of products account for a high proportion of the total revenue and they must exert a strong influence on the behaviour of the total system. The approach taken is to run all of the tests twice - once for a typical product and once for an influential product.

A simple aggregation of Branches is not possible, as the 'square root law' asserts that the total inventory in a system is proportional to the square root of the number of stockholding locations (76). This 'law' is, however, no more than a rule-of-thumb based upon assumptions which are not all valid in this application (e.g. orders must be placed according to the Economic Order Quantity principle which has a square function in its formulation). A better approach for calculating the aggregate Branch inventory is to scale down the order rate to individual Branch level, calculate the commensurate inventory, and multiply by the number of Branches.

The main purpose of the model is to examine the time-varying behaviour of the system and to check if any internally-generated cycles exist. Other strategic uses for the model are, however, envisaged (e.g. to test the effect of altering the balance between centralised and decentralised stocks), and for this reason all of the system constants are entered as parameters rather than embedded in the equations. The model represents a Process Flow and, as such, generates results which have quantitative significance. As the system inventory levels are of paramount importance, the expressions for desired inventory are developed in some detail to ensure that they are consistent with the formulae in the Inventory Management system. Forrester's (35) simple proportional relationship between stock level and demand rate (i.e. an assumed constant stockturn) is not appropriate.

The model was developed in three stages. First, the Forrester (35) Production/Distribution example (pp 137-186) was replicated in BASIC notation and run on a microprocessor. This demonstrated that:-

a) the inbuilt DYNAMO macros could be identified and written in their elemental form e.g. the third-order exponential delay function DELAY3 is equivalent to three cascaded levels and rates,

b) the time to run a full scale model on a microprocessor is not prohibitive nor is there any degradation in accuracy.

The second stage consisted of building a simplified model of the Organisation with a minimum number of equations and aggregated variables. The delays and other time constants were made fairly realistic. This demonstrated that the system is very stable and the CLIP, MAX and MIN functions (75) were unlikely to be required to constrain ill-behaved variables.

The third stage involved breaking down the simplified model into single variables and deriving relationships which matched the real system as closely as possible. This model showed the same stable behaviour as the simplified model when subjected to a simple disturbance. A comprehensive validation exercise was then carried out to compare the early results with events in the real system. This was achieved by collecting operational data (e.g. the Distribution Centre stock balances during the period after a new Branch was opened) and by asking Managers their opinion about how long changes take to work through the system. It was concluded that the model behaved in a manner which was generally representative of the real system. This validation exercise was repeated when different types of input disturbance were tested.

The program provides for four standard disturbance signals which modify the base flow rate. They are input as a percentage of the base flow rate and they have the same proportional effect on the customer orders received by the Branch and those received directly by the Distribution Centre. By simple changes to the program the disturbances can be applied in any combination, though in practice combinative disturbances were not found to be particularly instrumental in furthering the understanding of system behaviour. The basic disturbance functions in the program are:-

i) Step change - unidirectional or bidirectional steps may be effected at any point in the run.

ii) Trend - this takes the form of a linear increase or decrease.

iii) Sinusoidal - a sine wave of any specified period and amplitude may be applied.

iv) Random Noise - the program incorporates a random number generator, and a scaling routine to produce the desired magnitude relative to the base flow rate.

4.1.3 CONSTRUCTION OF THE MODEL

A complete flowchart of the model is presented in Figs. 4.4a and 4.4b. The mnemonics are defined in Appendix 3. The following equations relate to the diagrammatic symbols with a one-to-one correspondence. The Branch sector is formulated first (equations 1 to 16) followed by the Distribution Centre/Supplier Sector (equations 17 to 32). In general, the level equations represent the simple updating of a balance and are self-evident. The rate equations embrace the decision functions peculiar to the operating environment and some explanation is usually given. The solution interval is denoted 'DT' (Delta Time).

$$AUB.K = AUB.J + DT\ (CRB.JK - MDB.JK) \quad (1)$$

$$AIB.K = AIB.J + DT\ (MRB.JK - MDB.JK) \quad (2)$$

$$MDB.KL = \frac{AUB.K}{DFB.K} \quad (3)$$

The delay DFB is assumed to be variable, depending upon the stock availability at the Branch. Equation 3 can be best explained by considering a pool of unfilled orders (AUB) with the inflow suspended. The outflow rate (MDB) would be the total number of orders divided by the time taken for the last input order to traverse the delay.

$$DFB.K = DSB \times e^{-2.9957 \times \frac{AIB}{NIB}} + DHB \quad (4)$$

Equation 4 is derived with reference to Fig. 4.5.

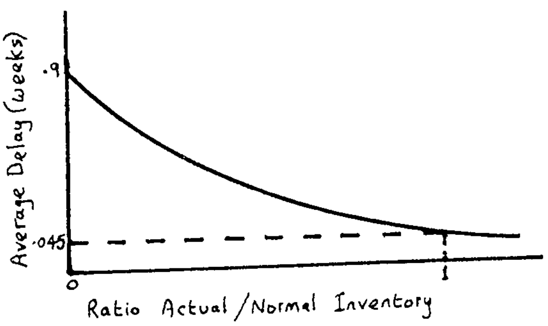

Fig. 4.5 Relationship Between Branch Inventory and Delay in Filling Customer Orders

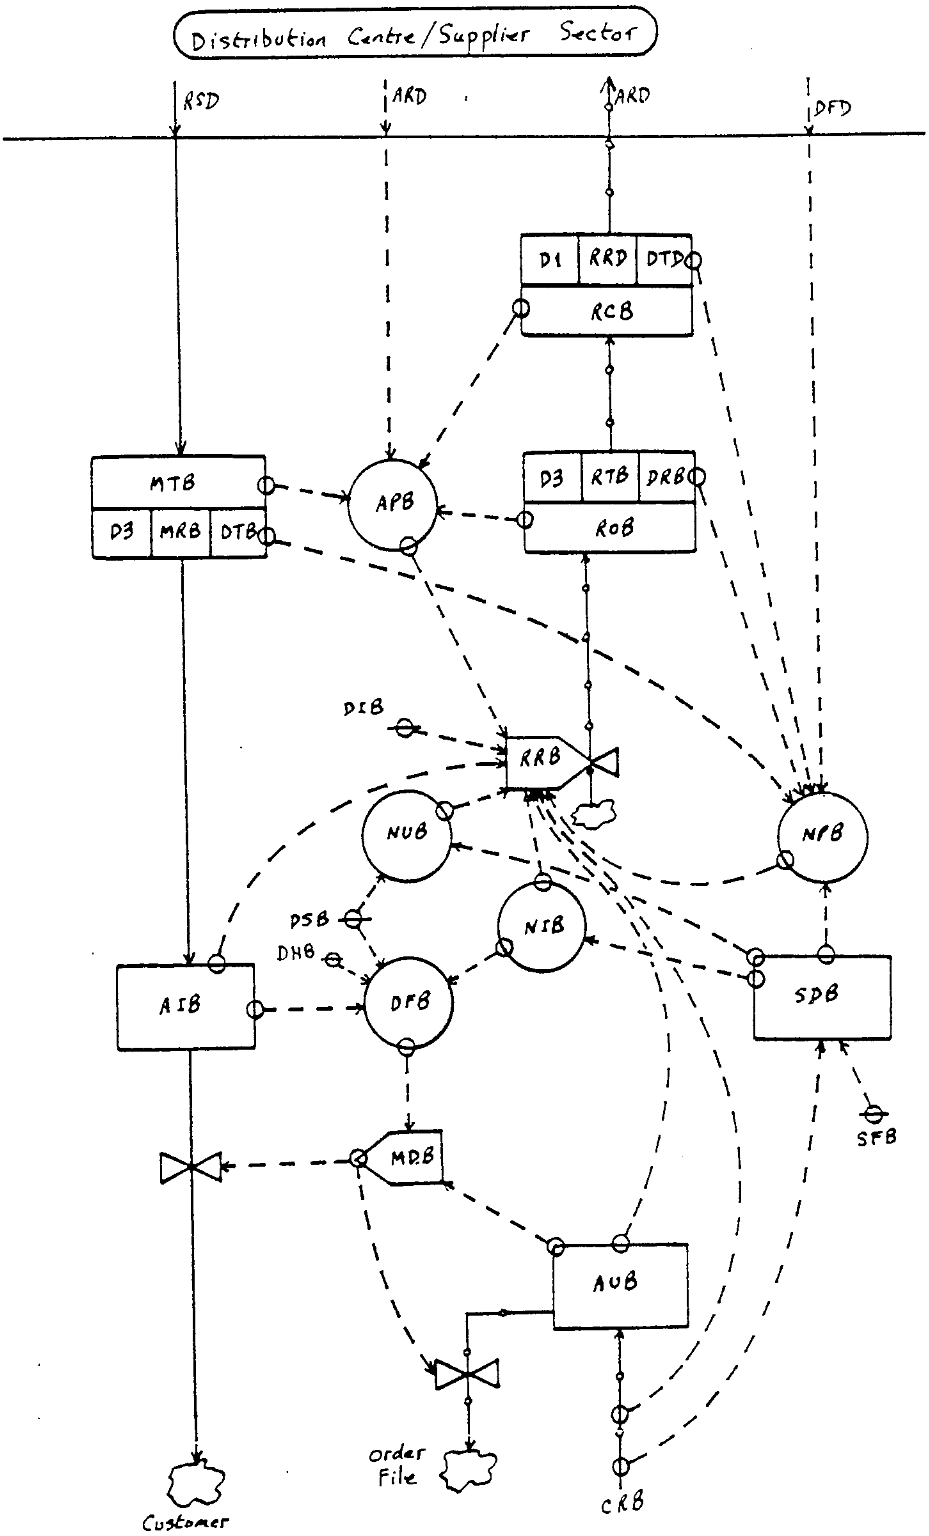

Fig. 4.4a Flow Diagram of Branch Sector

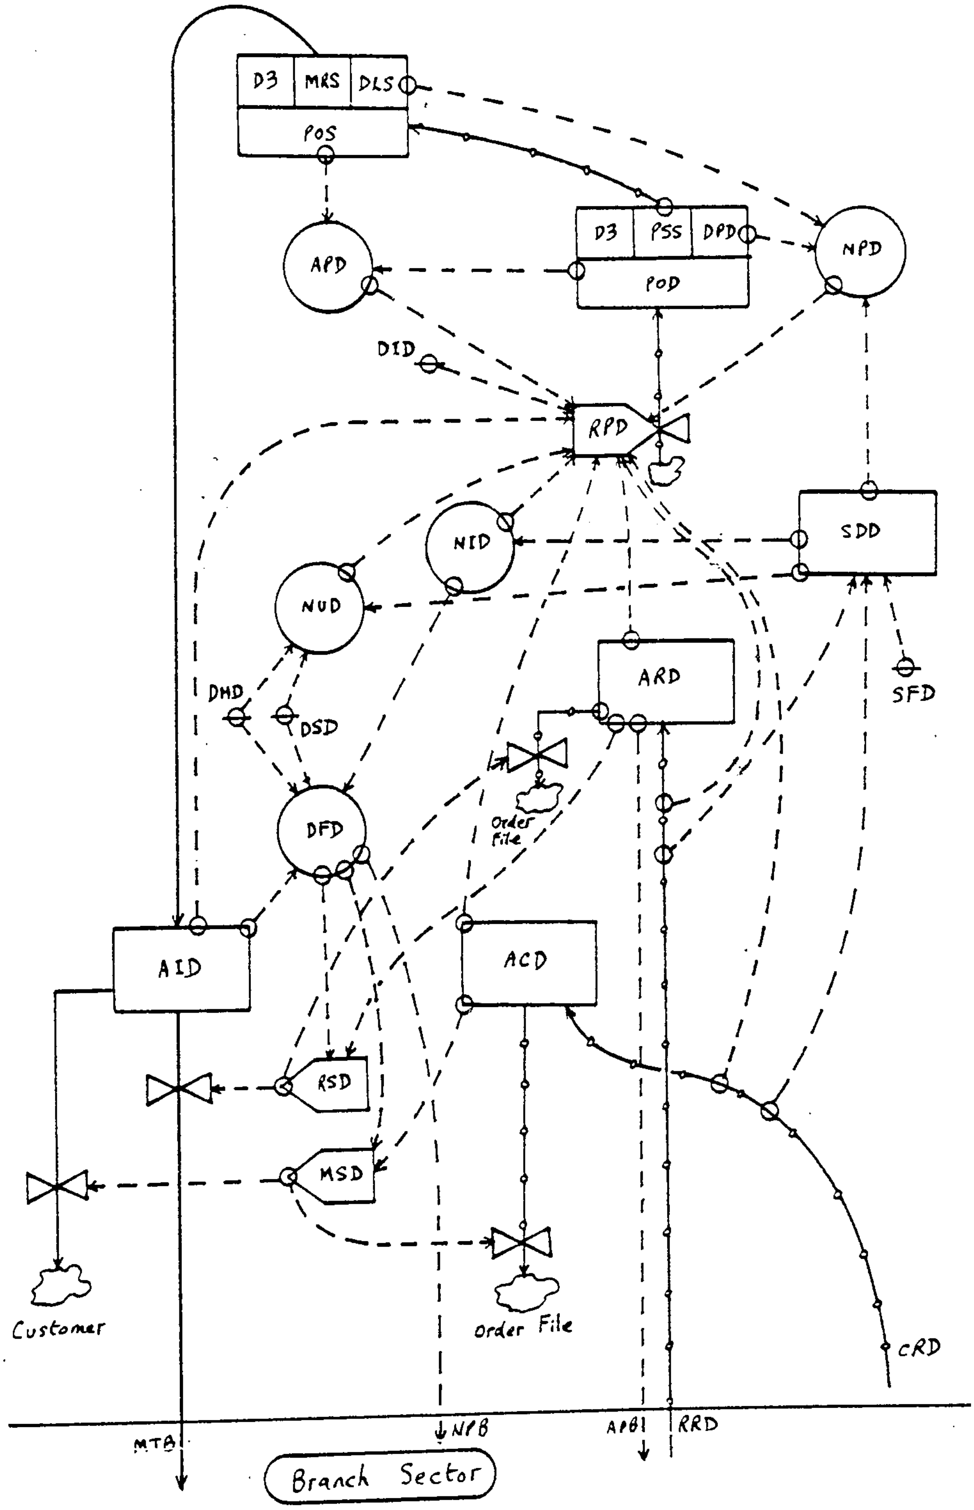

Fig. 4.4b Flow Diagram of Distribution Centre/Supplier Sector

The Branch service level is assumed to be 95% - hence approximately 5% of the products could be expected to be out of stock at any point in time. If the actual inventory is zero, all orders must be referred to the Distribution Centre, which entails an average delay (DSB) of 0.9 weeks (i.e. one half the Branch review cycle plus the replenishment lead time). If the actual inventory is at its normal level, 95% of the products will be in stock at the Branch and there will be an average delay of 0.9 weeks in acquiring the remainder from the Distribution Centre. Thus the overall average delay due to stock unavailability is 0.05 x 0.9 = 0.045 weeks. If the Branch has more stock than normal, the overall delay will be less than 0.045 weeks. At the limit, with infinite stock, there would be no delay.

These characteristics are well represented by a Negative Exponential curve of the general form $y = ae^{-bx}$. From Fig. 4.5:

For $y = 0.9$ and $x = 0$; then $a = 0.9$

For $y = 0.045$, $x = 1$, $a = 0.9$; then $b = 2.9957$

$\therefore y = 0.9e^{-2.9957x}$

or, $DFB.K = DSB \times e^{-2.9957 \times \frac{AIB}{NIB}}$

An average handling delay of 0.2 weeks is added in respect of the minority of products which are delivered to the customer from the Branch using local transport.

$$NIB.K = NB \times \left[SRB \times \left(2{\cdot}25 \times \left(\frac{SDB.K}{NB}\right)^{1{\cdot}5} + 0.0833 \times \left(\frac{SDB.K}{NB}\right)^{2} \right)^{0{\cdot}5} - 0{\cdot}4 \times \frac{SDB.K}{NB} \right] \quad (5)$$

Each product line is normally replenished from the Distribution Centre every week, as the reordering parameters are currently set to encourage small frequent replenishments. The Working Stock is therefore assumed equal to the average weekly demand.

The status of each product is reviewed weekly, therefore the interval between reaching the reorder level and the next review is uniformly distributed with Mean 0.5 weeks and Variance $1^2 \div 12 = 0.0833$. Assuming a fixed lead time of 0.4 weeks:

Mean demand in lead time, $\mu_{LD} = \mu_D \ (0.5 + 0.4) = 0.9\ \mu_D$

where μ_D = mean demand per week

And variance of demand in lead time,

$V_{LD} = 0.9\ V_D + \mu_D^2\ (0 + 0.0833)$

where V_D = variance of demand per week

The Variance Law, described by Brown (77) postulates a relationship between the Variance and Mean of demand of the form:

$\text{Variance} = a \times \text{Mean}^b$

Trials with live data have produced values of 2.5 for a and 1.5 for b using weekly demand data.

i.e.

$$V_D = 2.5\,\mu_D^{1\cdot5}$$

Substituting,

$$V_{LD} = 0.9 \times 2.5\,\mu_D^{1\cdot5} + 0.0833\,\mu_D^2$$

Gamma Modulus, $k = \dfrac{\mu_{LD}^2}{V_{LD}}$ (Chapter 3, equation 3)

$$= \frac{0.9^2\ \mu_D^2}{2.25\,\mu_D^{1\cdot5} + 0.0833\,\mu_D^2}$$

A 95% service level entails losing 5% of the demand in every order cycle (on average) and the whole of the loss is normally incurred during the lead time. Thus the Percentage Lost Sales (PLS) during the lead time $= \dfrac{1.0}{0.9} \times 0.05 = 0.0556$

Given k and PLS, the value of U (the standardised reorder level) can be obtained by an iterative application of the routines in Appendix 2.

Reorder level = U σ_{LD}

And $\sigma_{LD} = V_{LD}^{0\cdot5}$

$\therefore$ Reorder level = U $(2.25\mu_D^{1\cdot5} + 0.0833\ \mu_D^2)^{0\cdot5}$

And, Average Stock = Reorder Level - Demand in lead time

+ ½ Working Stock

$$= U(2.25\ \mu_D^{1\cdot5} + 0.0833\ \mu_D^2)^{0\cdot5} - 0.9\ \mu_D + 0.5\ \mu_D$$

$$= U(2.25\ \mu_D^{1\cdot5} + 0.0833\ \mu_D^2)^{0\cdot5} - 0.4\mu_D$$

The 'NB' factors are introduced to give the combined stock for all Branches with a joint demand rate of SDB.

The standardised reorder level is a variable depending on the demand rate μ_D. Table 4.1 illustrates the average combined stock level obtained from this expression with a range of joint demand rates, assuming 20 Branches.

Demand p.w.	k	U	Av. Stock (wks)
1	0.080	5.796	18.06 (18.06)
20	0.347	3.657	103.80 (5.19)
100	0.743	3.155	289 (2.89)
3000	3.034	2.934	3360 (1.12)

Table 4.1 Standardised Reorder Levels and Average Stocks for Varying Branch Demand Rates

The actual Branch stockturn is of the order of 10, giving an average 5.2 weeks' stock. This is consistent with the figure for the product with a demand of 1 unit per week which is generally regarded as 'typical'. The relative stocks for the other demand rates were checked against actual product stocks and the results were reasonably consistent with the Table.

The standardised reorder level is input to the model as a parameter, as the iterative application of formulae employing series expansions is extremely time consuming. Strictly, the value of the standardised reorder level should change as the disturbance signals alter the input order rate. However, trials have shown that provided the input order rate does not change by more than 20%, the change in the standardised reorder level is no more than 1.5%. As the last term in the expression is independent of the standardised reorder level, the average stock will change by less than 1.5% which is not significant. The input value is therefore retained for the duration of the run.

$$SDB.K = SDB.J + \frac{DT}{SFB} (CRB.JK - SDB.J) \qquad (6)$$

This represents a first-order exponential smoothing process with a time constant SFB. Branch demand is smoothed using a 26-week moving average, but this would be cumbersome to manipulate in the model. Brown (77) has shown that the time constant used in exponential smoothing is equivalent to $(n+1)/2$, where n is the number of periods in a moving average having the same average age of data. An exponential expression is therefore substituted for the moving average with SFB = 27/2 = 13.5

$$RRB.KL = CRB.JK + 1/DIB\ (NIB.K - AIB.K + NPB.K - APB.K + AUB.K - NUB.K) \qquad (7)$$

This formulates the stock replenishment rate required by the Branch in terms of the deficit between its stock ownership and its present commitment. The term DIB is a time constant denoting the average delay in correcting inventory and pipeline imbalances. This is set at 6.5 weeks, which is half the quarterly recalculation period for the product reordering parameters.

NPB.K = SDB.K(DRB+DTD+DFD.K+DTB) —— 8

The normal quantity of materials and replenishment orders in the pipeline between Branch and Distribution Centre is calculated as the product of the smoothed demand on the Branch and the sum of the constituent delays. The delay in placing replenishment orders is set to 0.5 weeks, which is half the Branch stock review cycle. The delay in transmitting replenishment orders to the Distribution Centre is 0.1 weeks and the transportation delay is 0.2 weeks.

APB.K = ROB.K + RCB.K + ARD.K + MTB.K —— 9

Equation 9 aggregates the actual orders and materials in the pipeline at their various stages in the circuit.

NUB.K = 0.05 x DSB x SDB.K —— 10

This assumes that 5% of the items in the Branch will be subject to a stockout delay at any given time (as equation 4).

ROB.K = ROB.J + DT(RRB.JK - RTB.JK) —— 11

RTB.KL = DELAY3(RRB.JK,DRB) —— 12

The delay in placing replenishment orders by the Branch has the characteristics of an infinite-order delay. In practice, replenishment order requirements accumulate over a weekly stock review cycle. They are then keyed into electronic data recorders and transmitted to the Distribution Centre via the Computer Centre. Branch stock reviews are carried out on a rota basis, one-fifth of the range being reviewed daily. Each product is not reviewed on the same day by all Branches. Hence the load on all Sections of the Distribution Centre are evenly spread over the week. Any replenishment occasioned by a Branch sale is therefore input to a delay of up to five days with an instantaneous discharge. This should be properly represented with a 'boxcar' or 'pulse' (75) function but these were found merely to introduce minor irregularities into the curves without having any appreciable effect on either the amplitude or period produced by a third-order exponential delay. Consequently the third-order function was preferred.

$$RCB.K = RCB.J + DT(RTB.JK - RRD.JK) \qquad (13)$$

$$RRD.KL = DELAY1(RTB.JK, DTD) \qquad (14)$$

Replenishment orders are transmitted to the Computer Centre and stored on computer files whilst awaiting 'call off' by the Distribution Centre. Calls are made in accordance with warehouse activity schedules and the picking and loading activity tends to be highest during the early morning and to decline gradually throughout the day. A first-order exponential delay therefore fits this routine.

$$MTB.K = MTB.J + DT\ (RSD.JK - MRB.JK) \qquad (15)$$

$$MRB.KL = DELAY3(RSD.JK, DTB) \qquad (16)$$

A third-order exponential delay is a good qualitative fit for the transportation delay. The output response is initially zero; it rises slowly as deliveries are made to Branches in the immediate vicinity of the Distribution Centre; the bulk of the deliveries are made within 4-8 hours of despatch; and some of the more remote Branches are serviced via Transit Depots where the goods are offloaded and transferred to local vehicles the following day.

$$ARD.K = ARD.J + DT(RRD.JK - RSD.JK) \qquad (17)$$

$$ACD.K = ACD.J + DT(CRD.JK - MSD.JK) \qquad (18)$$

$$AID.K = AID.J + DT(MRS.JK - MSD.JK - RSD.JK) \qquad (19)$$

$$RSD.KL = \frac{ARD.K}{DFD.K} \qquad (20)$$

$$MSD.KL = \frac{ACD.K}{DFD.K} \qquad (21)$$

The stock unavailability delay, DFD, is assumed to be variable depending upon the general inventory level at the Distribution Centre.

$$DFD.K = DSD \times e^{-2.9957 \times \frac{AID}{NID}} + DHD \qquad (22)$$

This expression was derived in the same manner as equation (4) with a similar assumption that 95% of the products are in stock when inventory is at a normal level and the remainder must be ordered. The delay when stock is unavailable, DSD, is set at 6.0 weeks, comprising 2.0 weeks for half the average time for the computer to generate a recommended Purchase Order, and 4.0 weeks for the average supplier lead time. The handling delay for processing replenishment or customer orders is set at 0.2 weeks.

$$NID.K = SRD \times (10{\cdot}25 \times SDD.K^{1{\cdot}5} + 4{\cdot}0033 \times SDD.K^2)^{0{\cdot}5} - 2{\cdot}1 \times SDD.K \qquad (23)$$

The average order cycle at a Distribution Centre is 4.0 weeks, which produces a Working Stock of approximately 4 weeks' demand.

The status of each product is reviewed daily, therefore the interval between reaching the reorder level and the next review is uniformly distributed with mean 0.1 weeks and variance $0.2^2/12 = 0.0033$. (This is true on average even though the joint status of a family of products is the order generation criterion). Assuming a mean lead time of four weeks with a variance of 4:

Mean demand in lead time, $\mu_{LD} = \mu_D \ (4.0 + 0.1) = 4.1\,\mu_D$

where μ_D = Mean demand per week

Substituting from the Variance Law,

$$V_{LD} = 4.1 \times 2.5\ \mu_D^{1\cdot5} + 4.0033\ \mu_D^2$$

$$= 10.25\ \mu_D^{1\cdot5} + 4.0033\,\mu_D^2$$

$$\text{Gamma modulus, } k = \frac{\mu_{LD}^{2}}{V_{LD}}$$

$$= \frac{4.1^{2}\ \mu_D^{2}}{10.25\mu_D^{1.5} + 4.0033\ \mu_D^{2}}$$

$$\text{PLS during the lead time} = \frac{4.0}{4.1} \times 0.05 = 0.0488$$

The standardised reorder level (U) for k, PLS can be obtained as for equation 5, thus:

$$\text{Reorder level} = U\ (10.25\ \mu_D^{1.5} + 4.0033\ \mu_D^{2})^{0.5}$$

$$\text{And, Average Stock} = \text{Reorder Level} - 4.1\mu_D + 2.0\ \mu_D$$

$$= U\ (10.25\ \mu_D^{1.5} + 4.0033\ \mu_D^{2})^{0.5} - 2.1\ \mu_D$$

The standardised reorder levels and average stock values for varying demand rates are given in table 4.2. These were also verified against the actual stocks in the Distribution Centre.

Demand p.w.	k	U	Av. Stock (wks)
2	1.494	2.945	15.56 (7.78)
40	2.989	3.015	202.00 (5.05)
200	3.555	3.078	918.57 (4.59)
6000	4.065	3.177	26164.94 (4.36)

Table 4.2 Standardised Reorder Levels and Average Stocks for Varying Dist. Centre Demand Rates

$$SDD.K = SDD.J + \frac{DT}{SFD}\ (RRD.JK + CRD.JK - SDD.J) \qquad 24$$

Demand at the Distribution Centre is smoothed using first-order exponential smoothing with a time constant of 7.

Demand for a period consists of both Branch replenishment orders and direct customer orders.

$$RPD.K = RRD.JK + CRD.JK + 1/DID(NID.K - AID.K + NPD.K - APD.K + ARD.K + ACD.K - NUD.K) \quad (25)$$

This formulates the purchasing rate required by the Distribution Centre in terms of the deficit between its owned and committed stock. The average adjustment delay, DID, is 2.0 weeks, which is half the average time between recalculating Lagrange Multipliers in the Inventory Management system. The reordering parameters are recalculated every time an order is placed, using the Lagrange Multipliers.

$$NPD.K = SDD.K(DLS + DPD) \quad (26)$$

The normal quantity of materials and purchase orders in the pipeline between Distribution Centre and supplier is computed as the product of the smoothed demand on the Distribution Centre and the sum of the delays. The average lead time delay is 4.0 weeks and the average delay in placing purchase orders is 2.2 weeks, comprising 2.0 weeks for half the order generation cycle and 0.2 weeks for vetting and mailing the orders.

$$APD.K = POD.K + POS.K \quad (27)$$

$$NUD.K = SDD.K(DSD \times 0.5 + DHD) \quad (28)$$

$$POD.K = POD.J + DT(RPD.JK - PSS.JK) \quad (29)$$

$$PSS.KL = DELAY3(RPD.JK,DPD) \quad (30)$$

The delay in placing purchase orders on the supplier has the same transient response characteristics as the delay in placing replenishment orders on the Distribution Centre. A third-order exponential function is used for reasons already discussed.

$$POS.K = POS.J + DT(PSS.JK - MRS.JK) \quad (31)$$

$$MRS.KL = DELAY3(PSS.JK,DLS) \quad (32)$$

The transient response characteristics of the supplier lead time delay would depend upon the nature of the production processes, such as batch or flow line, length of runs, ease of re-tooling, etc. Forrester (35) has suggested that a third-order exponential delay is generally suitable for production delays.

The system operation is assumed to commence in an equilibrial state. This is effected by a set of initialising equations which:-

a) set all flow rates in the Branch sector and the smoothed Branch demand equal to CRB,

b) set all flow rates in the Distribution Centre/Supplier sector and the smoothed Distribution Centre demand equal to CRB + CRD,

c) set the Branch and Distribution Centre inventories according to equations 5 and 23,

d) set all other levels as the product of their delay constants and initial flow rates.

The delay constants are written into the computer program immediately before the initialising equations and assigned numerical values. As internal data is communicated electronically between the Branches and Distribution Centre, some of the delays are of very short duration compared with those used by Forrester who presumably assumed paper flows in all cases. Because of this, a very small solution interval (DT) of 0.05 weeks was used. This was selected by trial and error as the maximum solution interval which produced completely smooth curves. The 2,080 iterations necessary for a 2-year simulation can be accomplished in approximately 20 minutes on a microprocessor.

4.1.4 TEST RESULTS

4.1.4.1 STEP INCREASE

Fig. 4.6 shows the result of applying a 20% increase in demand to a typical product having an initial demand of one unit per week per Branch plus 20 direct orders on the Distribution Centre per week. A unidirectional step increase is encountered when, for example, a new Branch is opened or a new range of products is added. A summary of the maximum changes to the variables is presented in Table 4.3.

Variable	Mnemonic	Max % Change	Time Period
Flow Rates			
Customer Orders (Total)	CRB/CRD	+20.0	-
Total Orders on Distribution Centre	TRD	+23.1	7
Purchase Orders sent to Supplier	PSS	+40.7	7
Materials Received from Supplier	MRS	+37.8	12
Inventory Levels			
Branch	AIB	-6.3	5
Distribution Centre	AID	-12.1	7
Unfilled Order Levels			
On Branch	AUB	+28.8	7
On Distribution Centre	ATD	+89.8	8

Table 4.3 Maximum Changes to Variables after 20% Step Increase

As approximately 95% of the Customer Orders on the Branch are satisfied immediately, the Branch inventory begins to fall as soon as the increase is applied. This causes an increase in the Branch

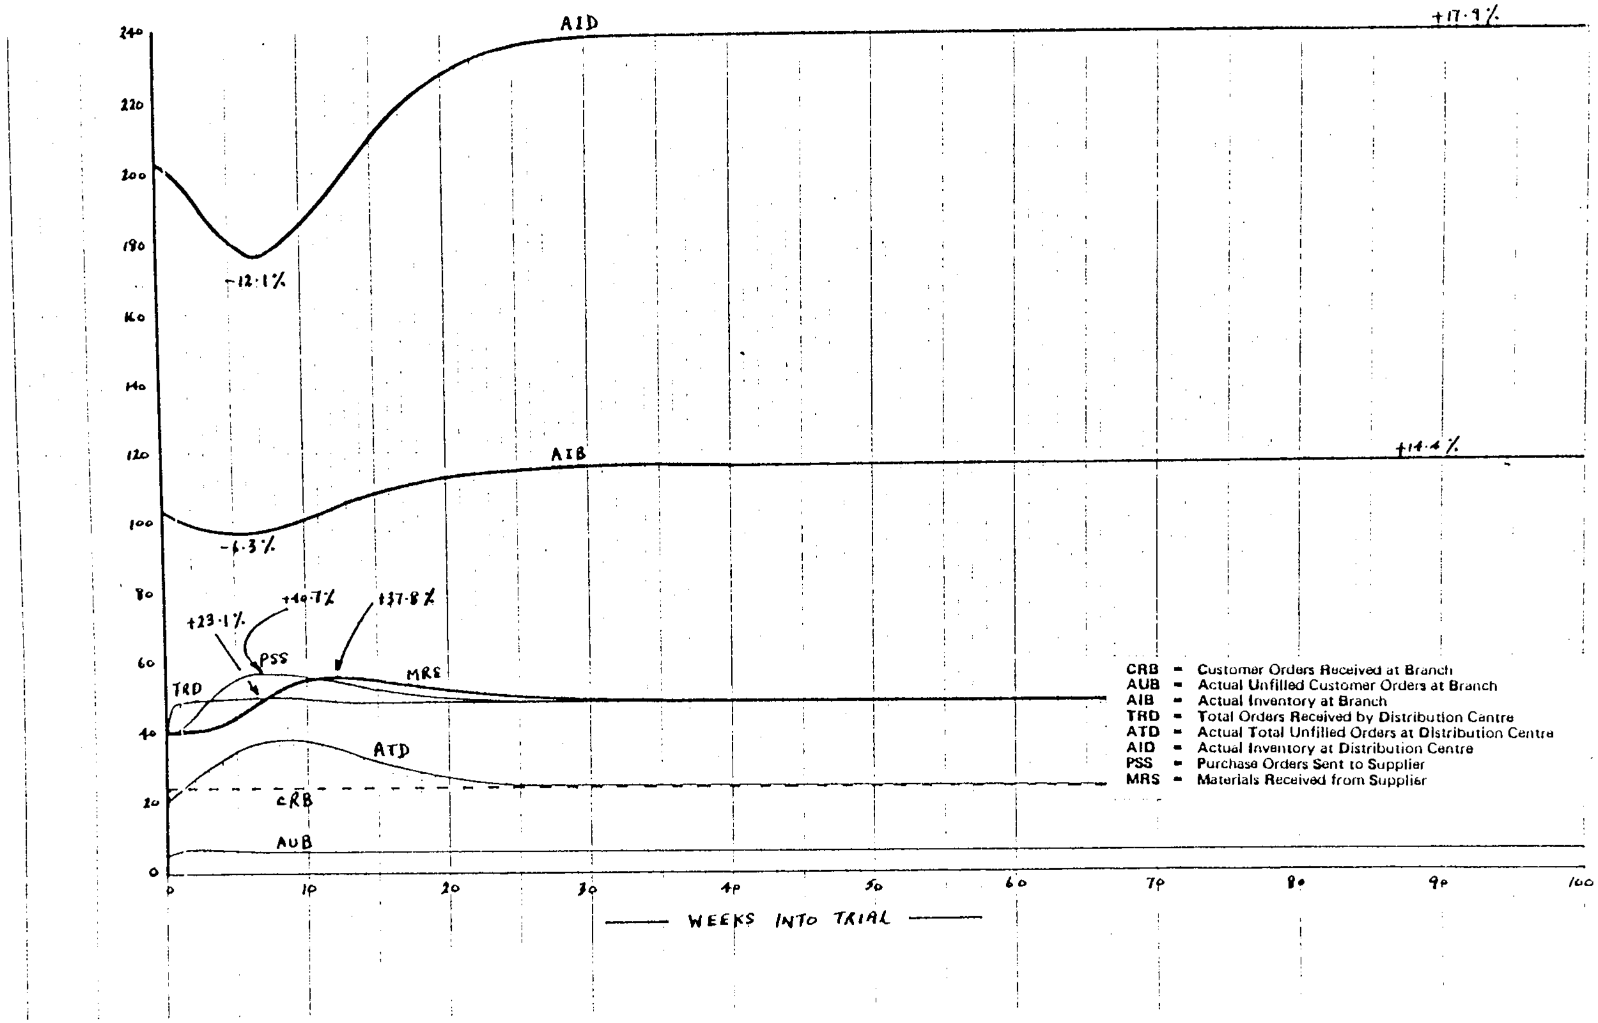

Fig. 4.6 Step Increase of 20%

orders on the Distribution Centre:

a) to replace the materials sold,

b) to increase the Branch inventory to a level commensurate with the perceived new order rate,

c) to increase the volume of orders and materials in the pipeline between the Branch and Distribution Centre.

The peak increase in the Branch ordering rate is seen to be limited to 23.1%, i.e. only 3.1% above the new customer ordering level. The constraint is caused by:-

a) the generally low levels of Branch stocks resulting from a frequent 'stock-topping' operation with short, fixed lead times,

b) the short duration of pipeline delays with electronic data communications,

c) the dampening effect of the customer orders received directly by the Distribution Centre.

The increased ordering rate on the Distribution Centre is seen to cause a maximum depletion of 12.1% seven weeks after the step increase. The effect on the Distribution Centre is almost twice as severe as that on the Branch as:-

a) the amplification of ordering in the Branch is passed on to the Distribution Centre,

b) the lead time for the Distribution Centre is much greater than for the Branch, hence the corrective action takes longer to effect.

The peak ordering rate on the supplier is more than double the magnitude of the source increase. This is due to:-

a) the increased ordering rate on the Distribution Centre,

b) the requirement to raise Distribution Centre stocks to a level commensurate with the perceived demand,

c) the need to increase the volume of orders and materials in the pipeline between Distribution Centre and supplier. This effect is much greater than the counterpart at the Branch as the total pipeline delay includes the supplier lead time and is almost 5 times as great in steady state conditions.

Not surprisingly the peak Distribution Centre ordering rate coincides with the minimum inventory, and the peak supply rate occurs with a lag of 5 weeks and some attenuation.

The dynamics of the Distribution Centre inventory and its effect on service level is of central interest. This is examined in detail by reference to Table 4.4.

Time Period	Orders Recd (TRD)	Actual Inv (AID)	Inv for 95% Service Level	Effective Std. ROL	Effective Serv Lvl%
0	40.00	202.00	202.00	3.015	95.0
2	48.38	194.73	239.77	2.617	91.6
4	48.89	183.91	242.05	2.506	90.6
6	49.19	177.89	243.40	2.445	89.7
7	49.24	177.61	243.62	2.441	89.7
8	49.23	179.02	243.58	2.454	89.9
10	49.07	185.62	242.86	2.516	90.7
12	48.84	194.94	241.83	2.604	91.4
15	48.52	209.56	240.40	2.743	92.6
20	48.26	227.31	239.23	2.909	94.0
25	48.18	235.27	238.87	2.983	94.7
104 (Stable)	48.00	238.07	238.07	3.015	95.0

Table 4.4 Effective Service Levels over Unstable Period

The inventory necessary to provide an average service level of 95% is calculated by applying the order rate (TRD) to equation (23). The effective standardised reorder level is also calculated from equation (23) using the order rate (TRD) and substituting the actual inventory (AID) for the normal inventory (NID). The corresponsing PLS value can be obtained from the formulae in Appendix 2, and the service level as:

$$1 - \left(\text{PLS} \times \frac{\text{Lead Time}}{\text{Order Interval}} \right)$$

The results show that a step increase in demand would cause a marked reduction in service level for at least six months. As there is no appreciable overshoot (In fact there is an overshoot which peaks at 0.34% of the final stable value in the 45th week, which is not discernible on the graph) there is no compensating increase in service level. The absence of overshoot and the slow response is contributed to by the heavy damping effect of the demand smoothing factors at both the Distribution Centre and Branch. In practice the inventory level is a factor in the Expected Lost Sales calculation which generates recommended purchase orders. The problem would therefore be alleviated by the computer advancing purchase orders and shortening the order cycle. The extent of the alleviation would depend upon the constitution of the buying families involved, which is outside the scope of a continuous system simulation model.

4.1.4.2 STEP INCREASE (FAST-MOVING PRODUCT)

Fig. 4.7 depicts the results of repeating the previous test with a fast-moving product (15mm copper tube). The initial demand rate is 150 lengths per week per Branch plus 3,000 lengths per week taken as direct orders on the Distribution Centre. The maximum changes to the variables are shown in Table 4.5.

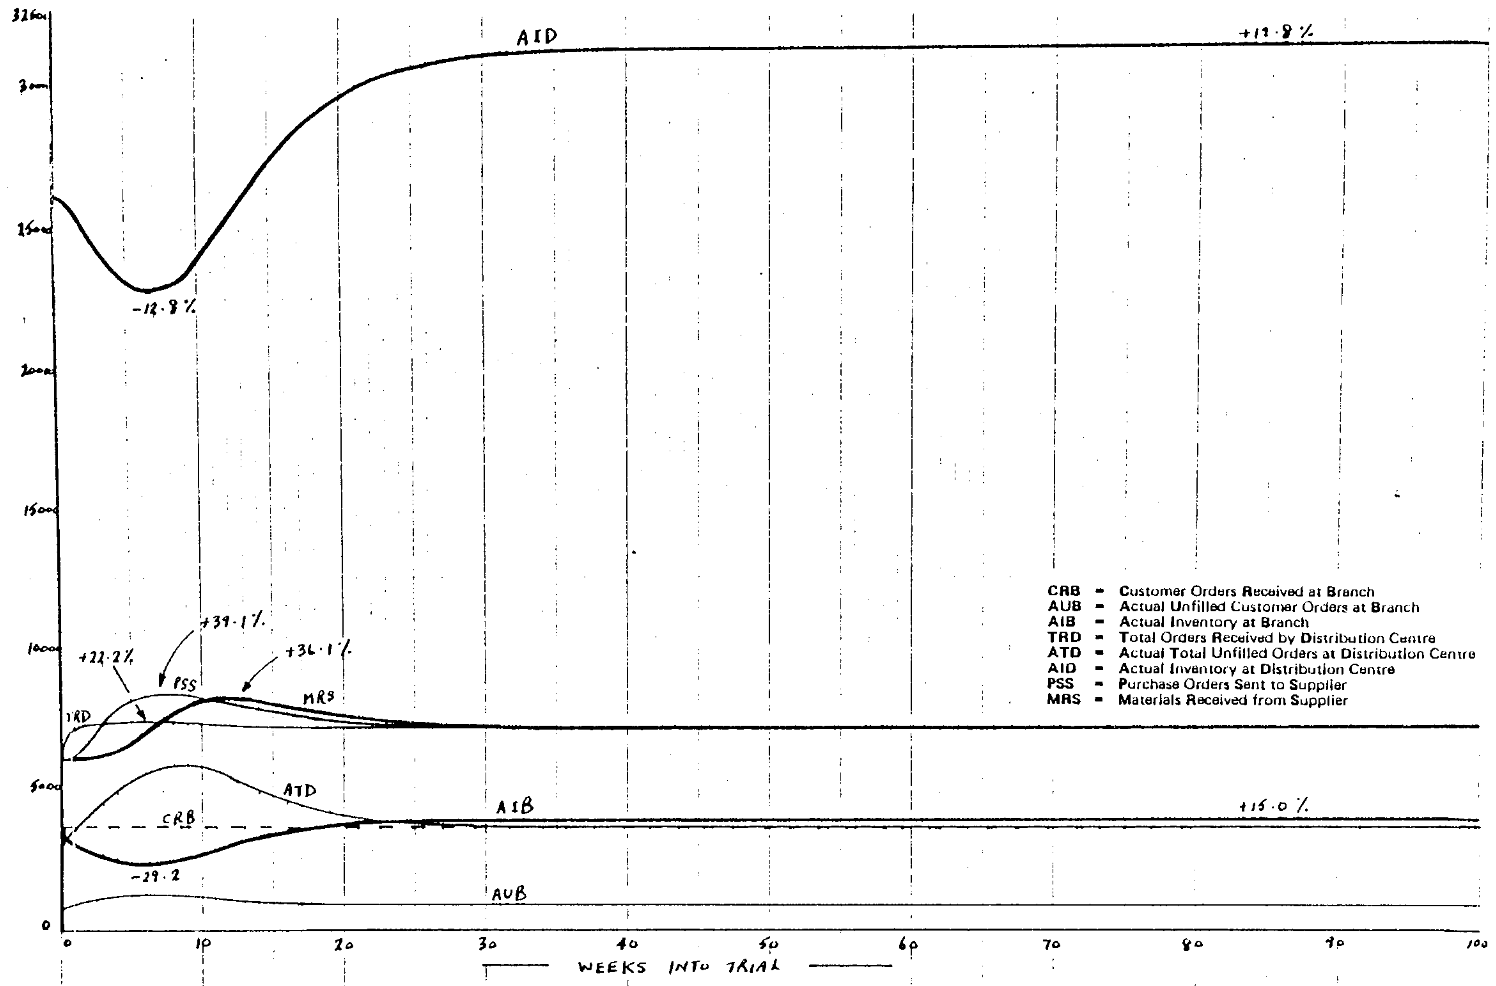

Fig. 4.7 Step Increase of 20% with Fast-Moving Product

Variable	Mnemonic	Max % Change	Time Period
<u>Flow Rates</u>			
Customer Orders (Total)	CRB/CRD	+20.0	-
Total Orders on Distribution Centre	TRD	+22.2	7
Purchase Orders Sent to Supplier	PSS	+39.1	7
Materials Received from Supplier	MRS	+36.1	12
<u>Inventory Levels</u>			
Branch	AIB	-29.2	6
Distribution Centre	AID	-12.8	7
<u>Unfilled Order Levels</u>			
On Branch	AUB	+36.4	7
On Distribution Centre	ATD	+93.5	8

Table 4.5 Maximum Changes to Variables after 20% Increase for Fast-Moving Product

The correspondence between this table and Table 4.3 is extremely close except for the depletion of the Branch inventory. This is because of the very low Buffer Stocks held at Branches for fast-moving products.

The Distribution Centre inventory stabilises after approximately 30 weeks, which concurs with that for the typical product. The minimum inventory level of the fast-moving product provides an effective service level of 89.4% compared with 89.7% for the typical product. These results suggest that the dynamics of the inventory at the Distribution Centre are not significantly influenced by the movement rates of the products. This conclusion was reinforced by further trials, and all other results here reported refer to the typical product.

4.1.4.3 STEP DECREASE

A step decrease rarely occurs in practice but a recent case was encountered when the servicing of a group of Branches was transferred from one Distribution Centre to another. Hence one experienced a step increase whilst the other experienced a step decrease. The outcome of a step decrease of 20% is given in Fig. 4.8. This is not a mirror-image of Fig. 4.6, as some of the relationships are non-linear and the circular sequence of calculations results in all of the variables assuming different values.

The maximum inventory level at the Distribution Centre provides an effective service level of 98.6% and the inventory again takes approximately 30 weeks to stabilise.

4.1.4.4 REVERSED STEP INCREASE

This test consists of a step increase of 20% followed by a step decrease of 20% after 8 weeks. This simulates a sales promotion which is a regular trading practice in the industry. Promotions usually take the form of price reductions, and market intelligence suggests that normal sales levels appertain as soon as the promotion is withdrawn. Legal restrictions to the manner in which promotions are advertised militate against long-term promotions and encourage the switching of promotional items at short intervals. This is thought to have a major de-stabilising effect on the system and it has frequently been observed to be a major determinant of excess stocks.

The outcome of a single promotion is shown in Fig. 4.9, and Table 4.6 summarises the peak effects.

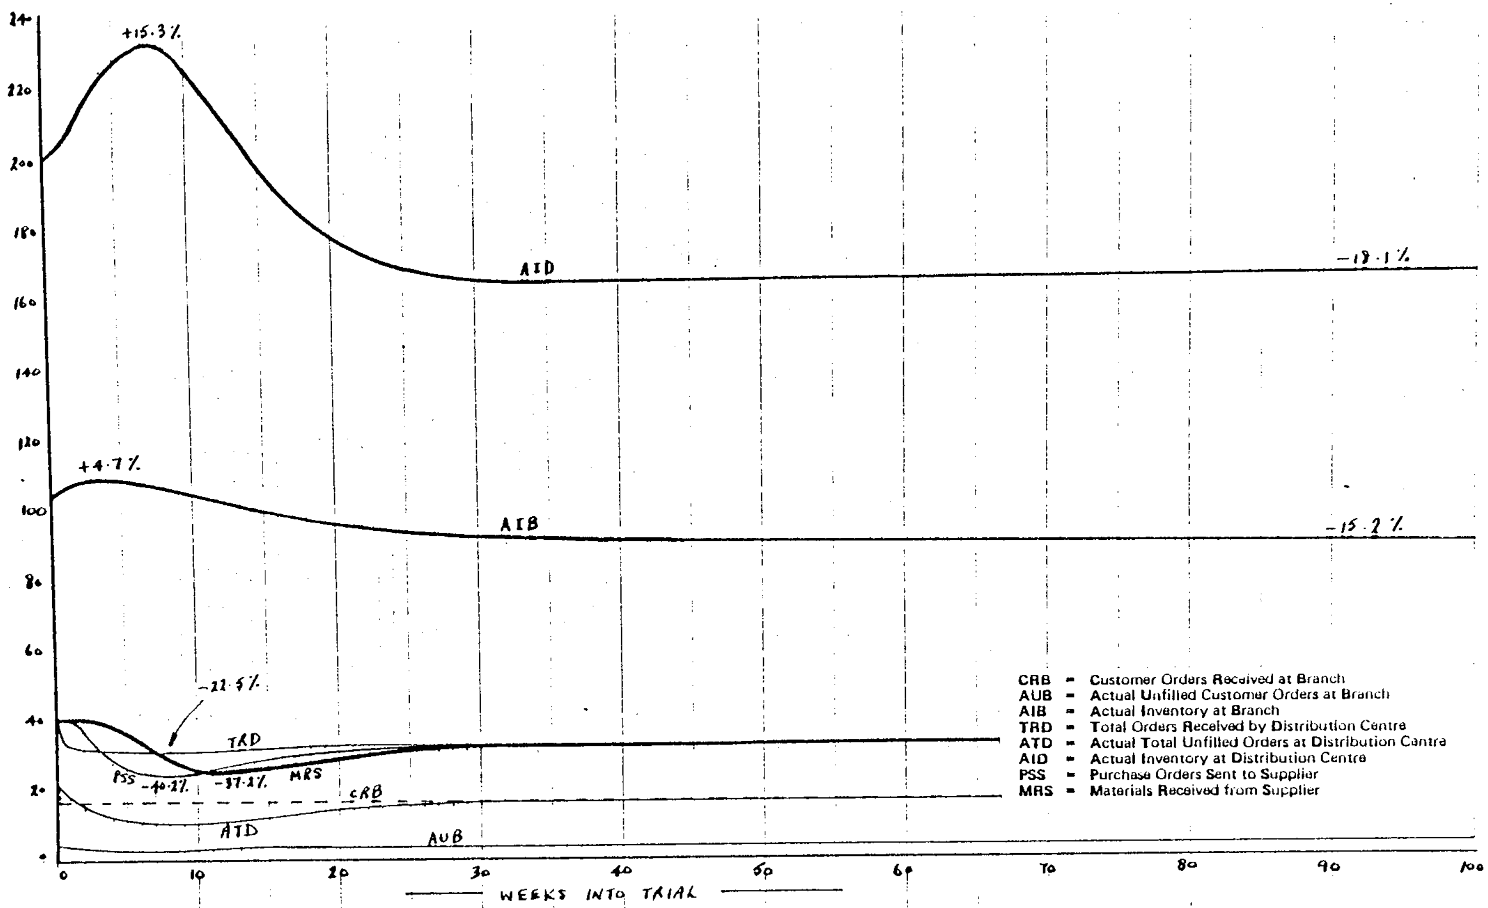

Fig. 4.8 Step Decrease of 20%

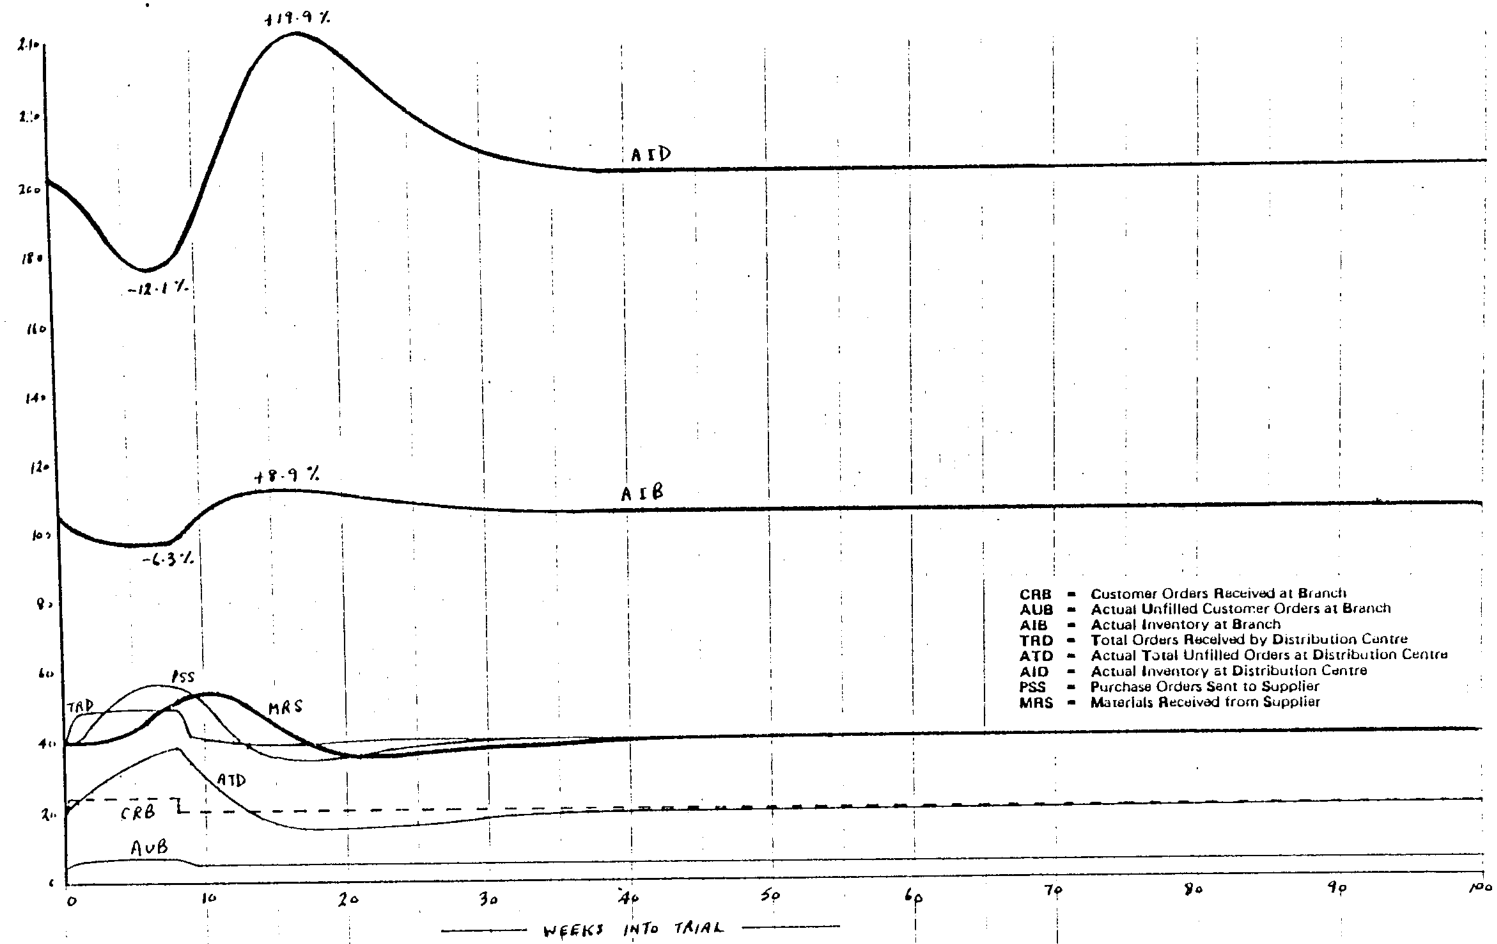

Fig. 4.9 Reversed Step Increase of 20%

Variable	Mnemonic	Max % Change (Overshoot)	Time Period
Flow Rates			
Customer Orders (Total)	CRB/CRD	+20.0	1-8
Total Orders on Distribution Centre	TRD	+23.1 (-1.5)	7 16
Purchase Orders Sent to Supplier	PSS	+40.7 (-12.8)	7 17
Materials Received from Supplier	MRS	+35.8 (-10.7)	11 22
Inventory Levels			
Branch	AIB	-6.3 (+8.9)	5 15
Distribution Centre	AID	-12.1 (+19.9)	7 17
Unfilled Order Levels			
On Branch	AUB	+28.8 (-2.4)	7 16
On Distribution Centre	ATD	+89.8 (-24.2)	8 19

Table 4.6 Maximum Changes with 20% Reversed Step Increase

When the Branch sales level falls, its receipts exceed its issues and the inventory rises. At the same time the forecast requirements for stocking the Branch and the Branch/Distribution Centre pipeline are reduced. Consequently the Branch orders on the Distribution Centre are reduced and the Distribution Centre inventory rises steeply (cf Fig. 4.6 where the increased demand is sustained). After an average delay of two weeks for the computer ordering cycle, the purchase ordering rate is reduced; and after a further four weeks for the lead time, the supply to the Distribution Centre is decreased commensurately. During the whole of the system reaction time the supply into the Distribution Centre exceeds the issues and the inventory level substantially overshoots its final value. All of the other variables experience some degree of overshoot before attaining their stable values.

Because of the central importance of the dynamics of the Distribution Centre inventory and its effect on service level, the trial was repeated with a 50% sales increase, which is not uncommon during a promotion. The actual inventory (AID) and the effective service level this provides are given in Fig. 4.10. The model was modified to calculate the service level every solution interval using the method explained in Sect. 4.1.4.1. The average service level over the first 40 weeks was calculated as:

$$\frac{\sum_{DT=1}^{DT=800} (\text{TRD x Service Level})}{\sum_{DT=1}^{DT=800} \text{TRD}}$$

and the average inventory as:

$$\frac{1}{800} \sum_{DT=1}^{DT=800} \text{AID}$$

The result indicates that whereas a 95% service level can be achieved with 202.00 units of stock in a steady-state condition, the sort of disturbance induced by a promotion requires an average of 225.33 units (+11.5%) to achieve a service level of 93.0 (an increase of 40% in Expected Lost Sales). As promotions are applied as a

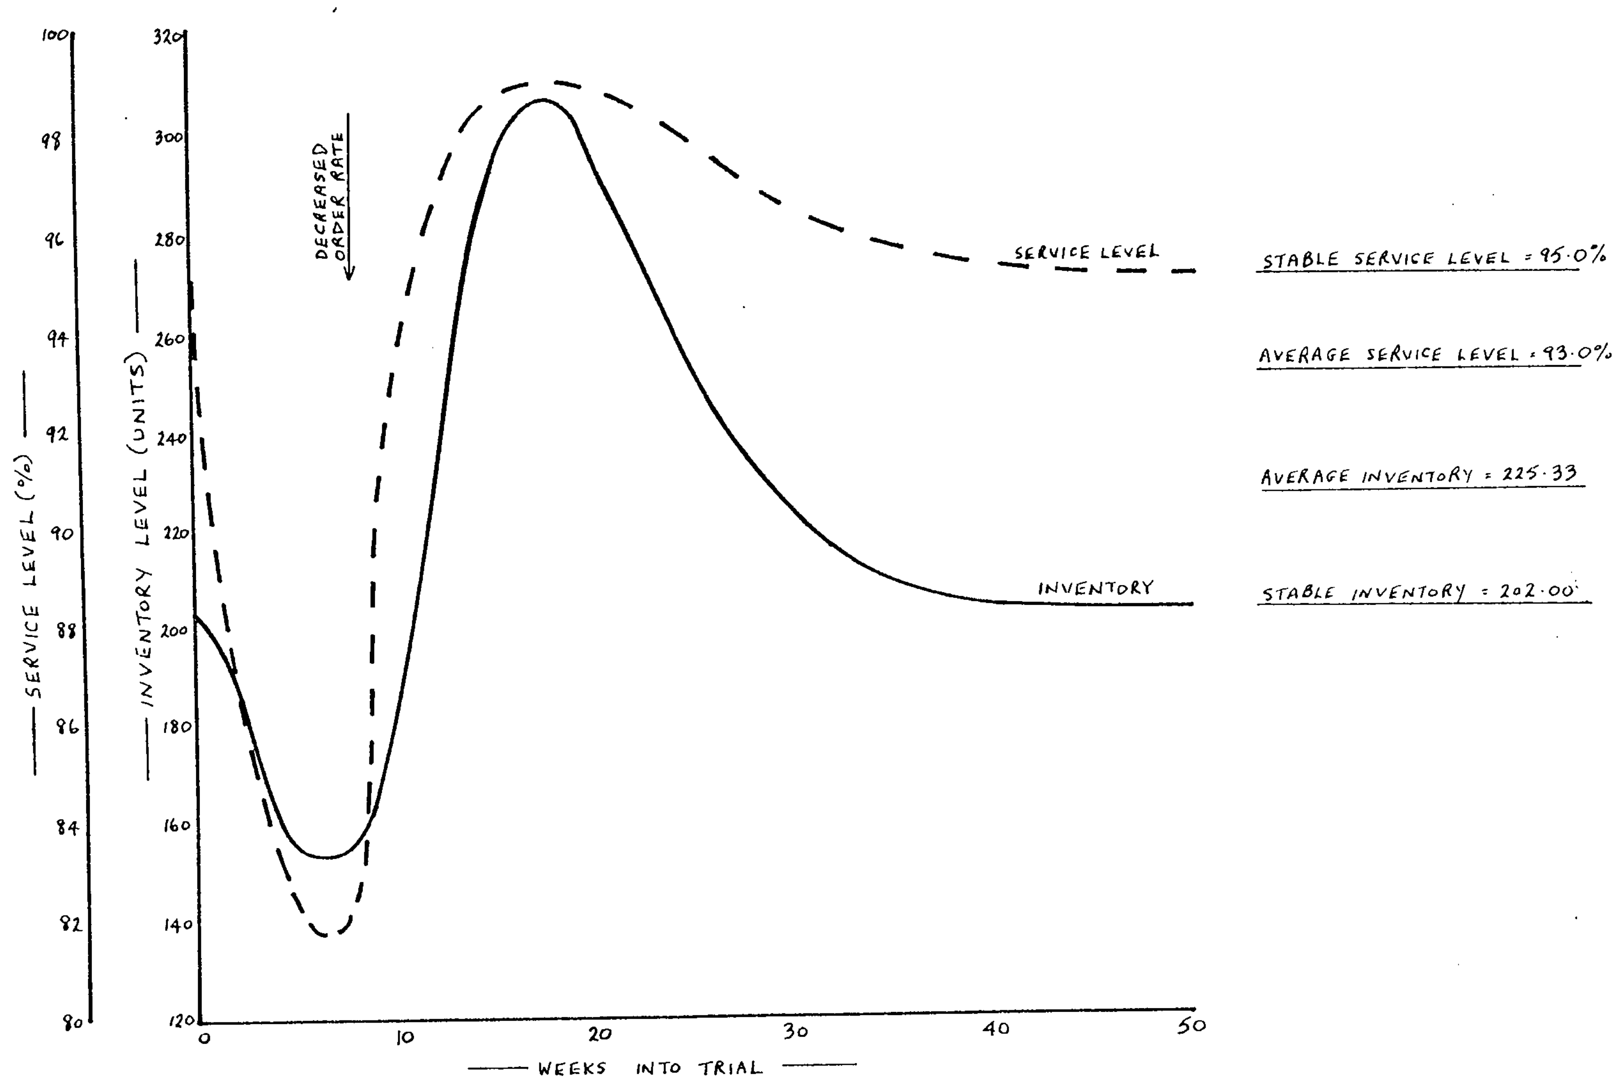

Fig. 4.10 Reversed Step Increase of 50% - Comparison of Stable and Average Inventory and Service Level

regular feature of trading, this degradation in performance is extremely significant.

The simulation implicitly assumes that there is no foreknowledge of changes to exogenous variables. In practice, promotions are usually planned several weeks in advance and hence the buyers have the opportunity to predict the extra stock needed and acquire it in advance. The trial was repeated a further three times - first by increasing the initial inventories by exactly the additional amount sold in the promotion; then by increasing the inventories by 50% of the additional amount; and finally increasing the inventories by 150% of the additional amount. This simulates a perfect prediction by the Buyer, an underestimate, and an overestimate respectively. The results are shown in Table 4.7.

Prediction	Service Level	Inventory
Perfect	95.3%	227.58
50% Underestimate	94.3%	225.84
50% Overestimate	95.9%	230.41

Table 4.7 Effect of Buyer Predictions on Performance during Promotions

This suggests that a good prediction can obviate the service level degradation with only a small increase in stock above the disturbance level. However, the practice of imposing step disturbances on the system, albeit to increase sales seems certain to be a contributory factor in the general overstocking problem experienced within the Inventory Management system.

4.1.4.5 LINEAR TREND (RAMP FUNCTION)

The Inventory Management system was implemented in the first

Distribution Centre in March 1979. This coincided with the start of a severe economic recession and, in real terms, sales fell steadily over the next two years (Fig. 3.2). The decline over this period was equivalent to approximately 0.5% per week. The effect on the system variables is shown in Fig. 4.11.

The onset of the decline in demand first manifests itself as a small increase in Branch stocks. The Branch reordering rate is consequently cut back, which is reflected in an increase in the Distribution Centre inventory. The sustained decline, however, results in a reduction in the demand forecasts, and the desired inventories and pipeline requirements are reduced commensurately. The ordering rates are therefore decreased, which eventually results in a fall in the inventories. After approximately 30 weeks all of the variables exhibit a steady declivity. Thereafter, the actual inventories (AIB, AID) are always greater than the normal inventories (NIB, NID). This is because the normal inventories influence the procurement rates, which in turn influence the actual inventories. Thus the effective service levels exceed the target values. The average Distribution Centre service level over the 2-year simulation is calculated as 96.5% and the average actual inventory as 172.51 (-14.6%). This offers a general explanation for the achieved service level exceeding the target value between March 1980 and September 1980 (Table 3.3) when the demand declined very steeply (Fig. 3.2); and the achieved and target values being comparable between September 1979 and March 1980 (Table 3.2) when the demand was fairly even. It does, however, suggest that the overstocking problem is worse than that indicated on the later report (Table 3.3).

4.1.4.6 SINUSOIDAL FLUCTUATION

A sinusoidal input with a one-year period simulates the trading

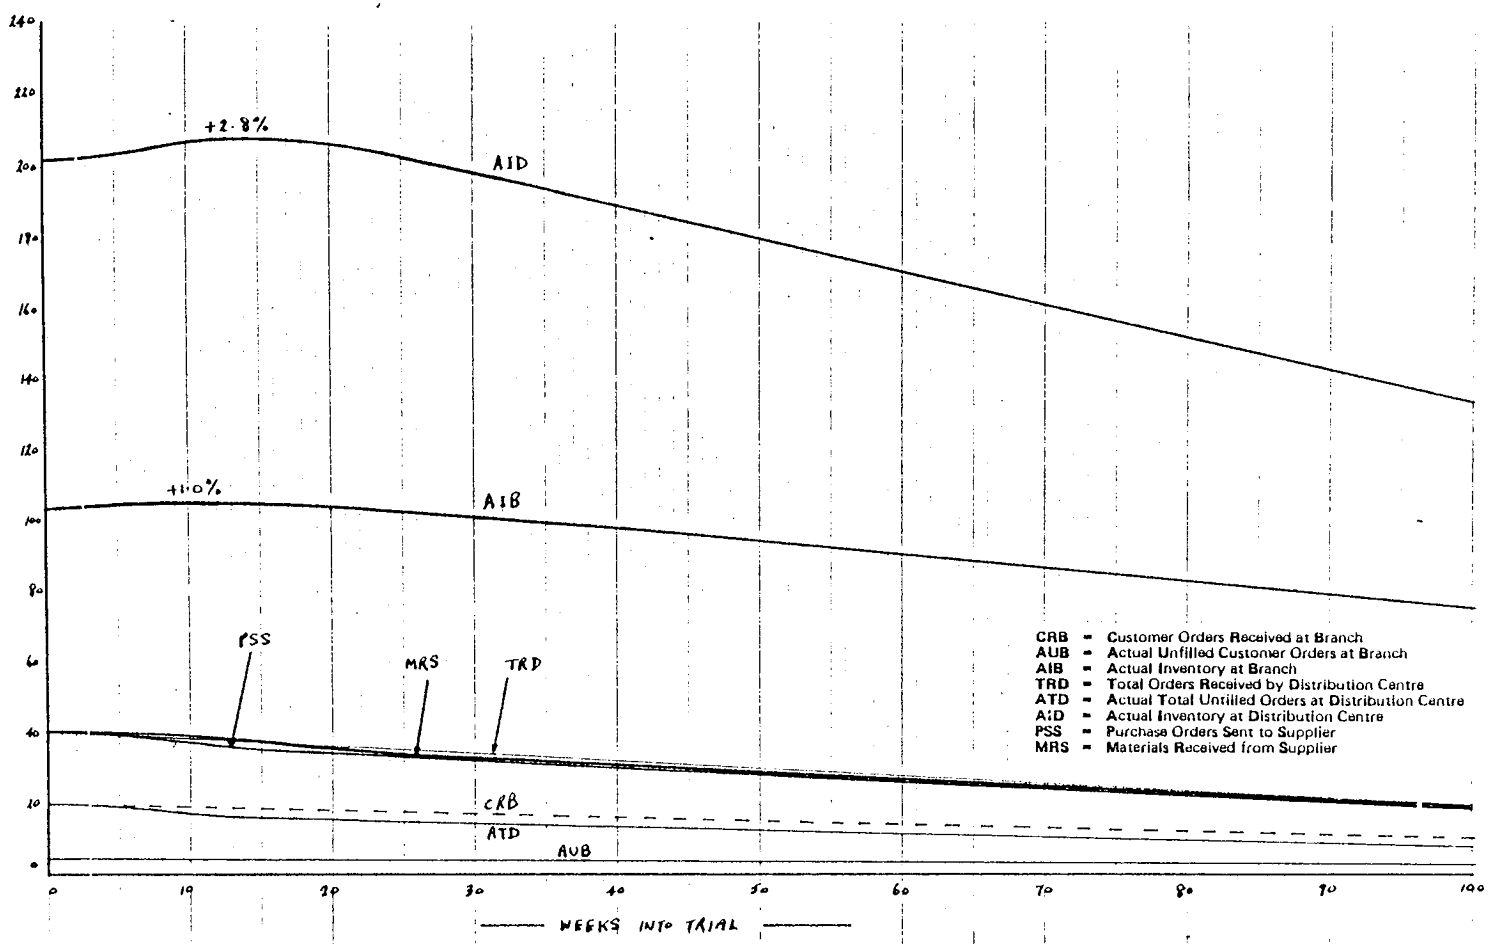

Fig. 4.11 Trend Decrease of 0.5% Per Week

pattern for seasonal products with a single peak. Fig. 4.12 displays the results of applying a sine wave of 10% amplitude to the customer order rate. The maximum and minimum points for the variables are presented in Table 4.8, together with the phase shift from the input. The tabulated figures refer to the second upper and lower swings which appear to be unaffected by the initial transient conditions. The upper and lower swings are seen to be unsymmetrical. Forrester (35) has observed the same characteristic when applying a sinusoidal input (pp 175-177) and he attributes this to system non-linearities.

Variable	Mnemonic	Max % Amplitude (Minimum)	Weeks Lag
Flow Rates			
Customer Orders (Total)	CRB/CRD	+10.0 (-10.0)	-
Total Orders on Distribution Centre	TRD	+11.5 (-11.5)	0
Purchase Orders Sent to Supplier	PSS	+20.9 (-20.6)	0
Materials Received from Supplier	MRS	+20.1 (-19.8)	3
Inventory Levels			
Branch	AIB	+6.8 (-6.7)	17
Distribution Centre	AID	+14.7 (-13.7)	20
Unfilled Order Levels			
On Branch	AUB	+14.1 (-12.4)	0
On Distribution Centre	ATD	+50.1 (-33.6)	-2

Table 4.8 Maximum Changes with 10% Sinusoidal Input

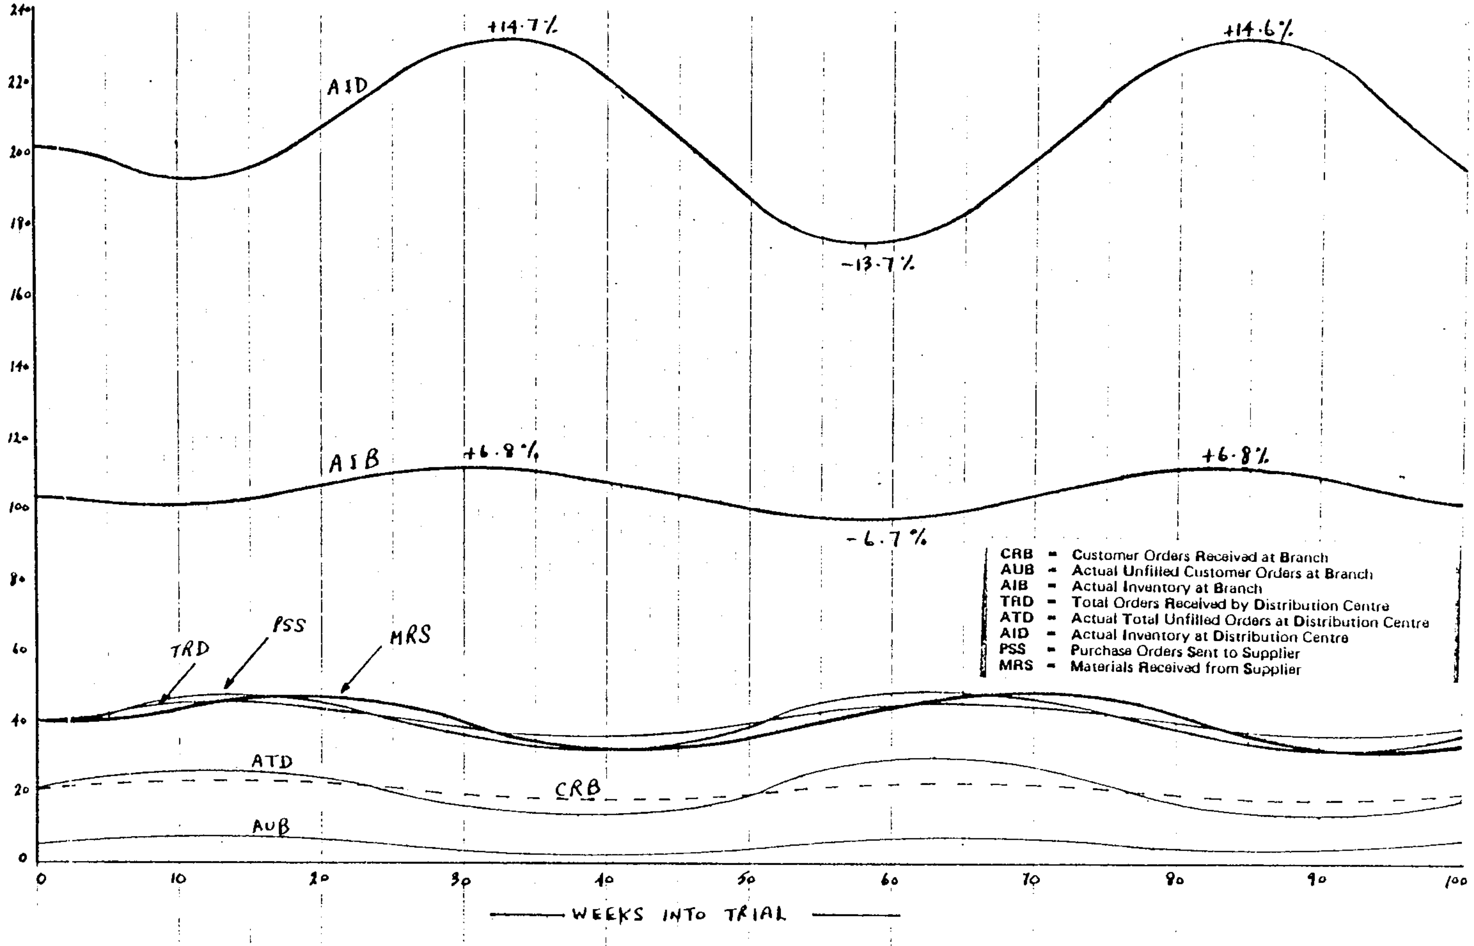

Fig. 4.12 Sinusoid with Annual Period (10% Amplitude)

The fact that the variable ATD leads the input signal is an example of counter-intuitive behaviour (33). Forrester (35) shows that in a system containing amplification, at some frequencies it is possible for a dependent variable to lead the independent variable (pp 415-417).

Of greater potential significance is the 20-week phase shift between the total orders on the Distribution Centre and its inventory level. The reasons for this can be elicited from Fig. 4.12. The inventory level turning points will occur when the receipts equate with the issues i.e. when the receipt and issue curves intersect. For reasons of clarity, Distribution Centre issues (MSD + MTB) have not been plotted, but it can be seen from equation (22) that the delay in filling orders (DFD) in steady state conditions is only 0.5 weeks. Hence the issue curve will be roughly in phase with the order curve (TRD). It can be seen that TRD attains its second peak in week 65, and it is not intersected by the receipts curve (MRS) until week 85, when AID is at a maximum.

The conclusion that the inventory level is at a minimum when it is most needed is not as serious as might be expected. The average inventory over the duration of the trial is 202.15, though this is not strictly comparable with the steady state value of 202.00 as the run is terminated at a low point in the inventory cycle. The average service level is 94.8%, an increase in ELS of 4%. This strongly suggests that abrupt shocks to the system have a much more deleterious effect than smooth wave inputs, and seasonality per se is not a major reason for performance degradation.

4.1.4.7 COMPOUND SINUSOIDAL FLUCTUATION

Fig. 3.2 identifies a business cycle of approximately 4.5 years. Some experimentation with simple sinusoidal inputs have revealed that, for equal amplitude, increasing the period has a beneficial effect on system performance i.e. the lower the frequency the more stable is the system. A sine wave with 4.5-year period and 20% amplitude produces an extremely stable result with negligible performance degradation. The combination of the business cycle with a seasonal trading pattern does, however, cause significant degradation. Fig. 4.13 depicts the state of the main variables (some are omitted for clarity) when subjected to a one-year sinusoidal input superimposed onto a 4.5-year sinusoidal input, with amplitudes of 10% and 20% respectively. As with the simple sinusoidal input, the Distribution Centre inventory is seen to be almost 180° out of phase with the ordering rate. The average inventory level is 203.6 (0.8% above the steady-state level), and the average effective service level is 94.3% (an increase of 14% in ELS). This result is not particularly meaningful as the averaging process over 4.5 years has no practical significance. The system results are normally monitored every six months, so the effect of the business cycle is likely to be perceived as a positive or negative ramp function.

4.1.4.8 RANDOM FLUCTUATIONS

In practice all of the input functions will be subject to random noise. Some experimentation involving superimposing random noise on to the basic signals already described has revealed that the variables display the same basic characteristics but with some distortion. The pursuit of this approach is therefore unlikely to prove instructive. Instead, the application of random deviations to a steady input flow is examined (Fig. 4.14). Forrester (35) has found that this can reinforce a system's intrinsic periodicity and initiate any tendencies to oscillate (pp 177-180). A pseudo-random number generator was used to produce random numbers between 0 and 1 which were scaled to give a

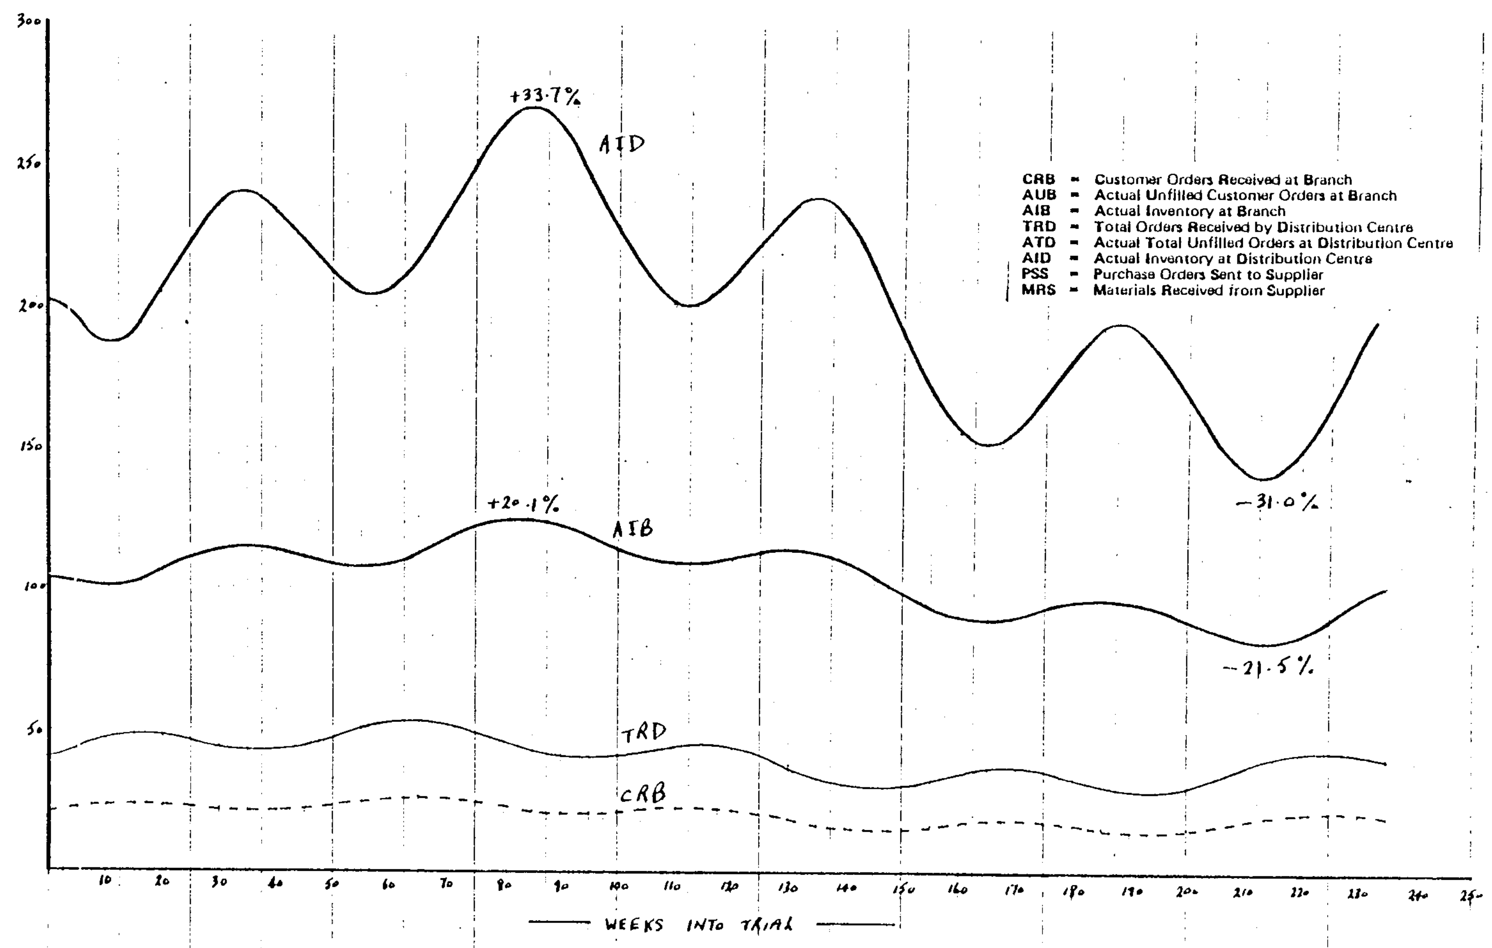

Fig. 4.13 Compound Sinusoid (Amplitudes 10%/20%)

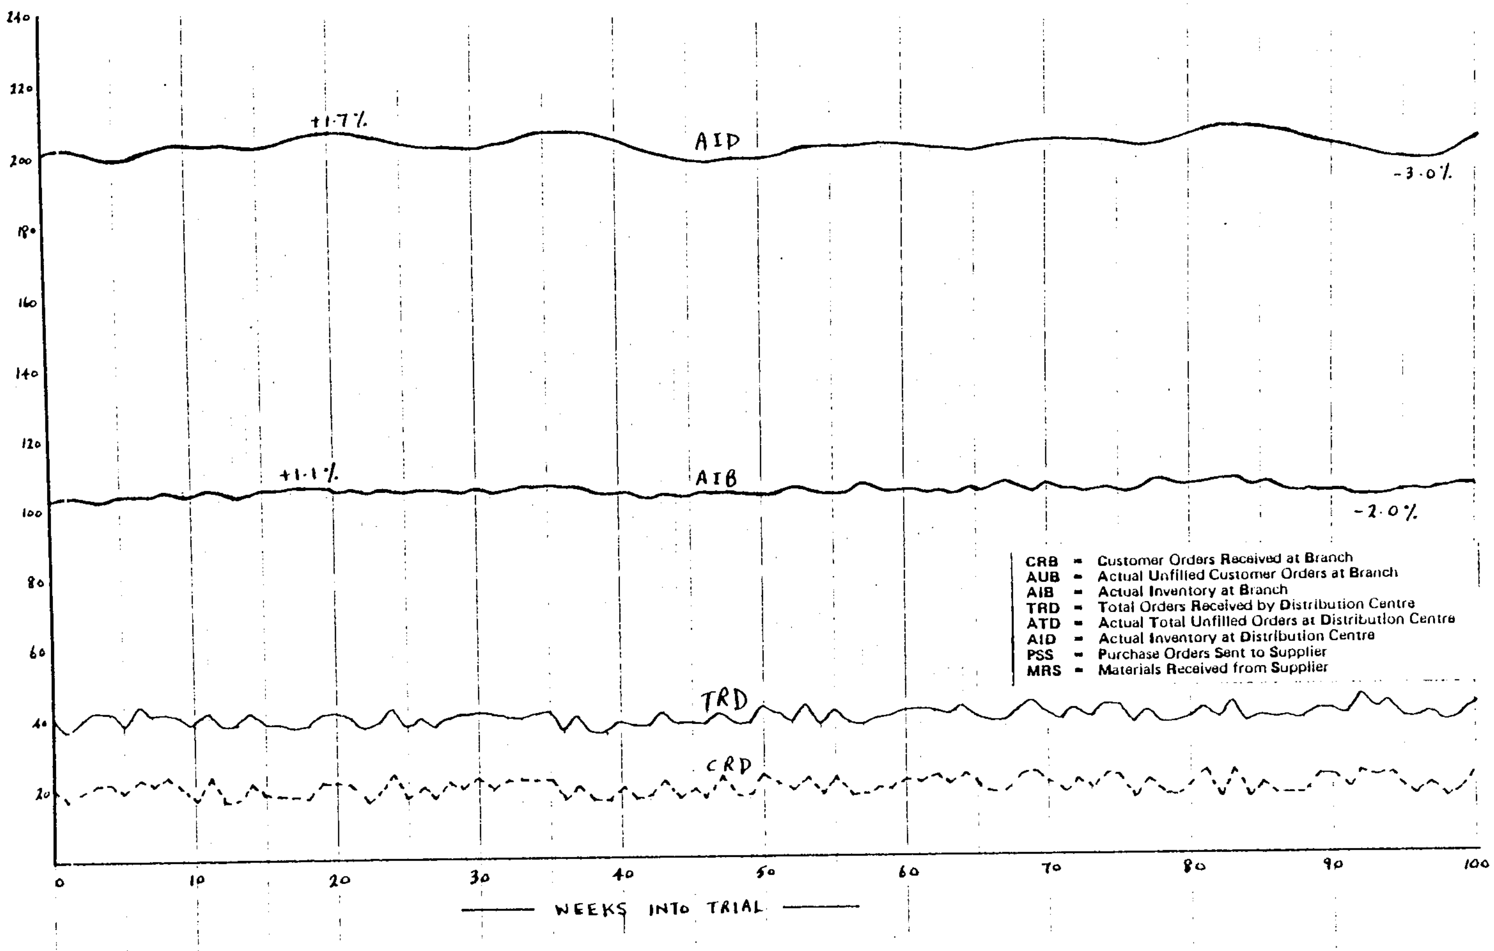

Fig. 4.14 Random Fluctuations with 20% Amplitude

maximum absolute deviation of 20% on to the base flow rates. The Distribution Centre inventory exhibits a wave form, but no regular period is discernible. The test was repeated three times with different random number seeds and an exaggerated amplitude of 50% (Fig. 4.15). It is concluded that the system is not influenced by any inherent natural cycles.

The trial with the 20% deviation produced an average Distribution Centre inventory level of 201.50 and an average effective service level of 94.9%. The inventory level is marginally less than the steady-state level, but this is due to the average order level (TRD) being 39.8 units per week, which is slightly less than the steady-state rate. The evidence suggests that the basic characteristics and values of the variables are not significantly altered by normal week-to-week order fluctuations.

4.1.5 CONCLUSIONS FROM SYSTEM DYNAMICS STUDY

1) The Branch/Distribution Centre/Supplier system is extremely stable with very little tendency to oscillate. This is due primarily to the heavy damping factors in the forecasting systems at both Branch and Distribution Centre. For the same reason, the inventories are slow to regain their steady-state values after being disturbed. In reality they would be continually homing towards a steady state without ever achieving it.
2) There is a considerable amount of amplification of the variables as the effects of disturbances work through the system, but not as much as other observers have found with similar structures. The limitation is due to the electronic communication of data, fast delivery service to the Branches, and circumvention of the Branches by 50% of the customer orders.
3) The imposition of sales promotions at frequent intervals acts as

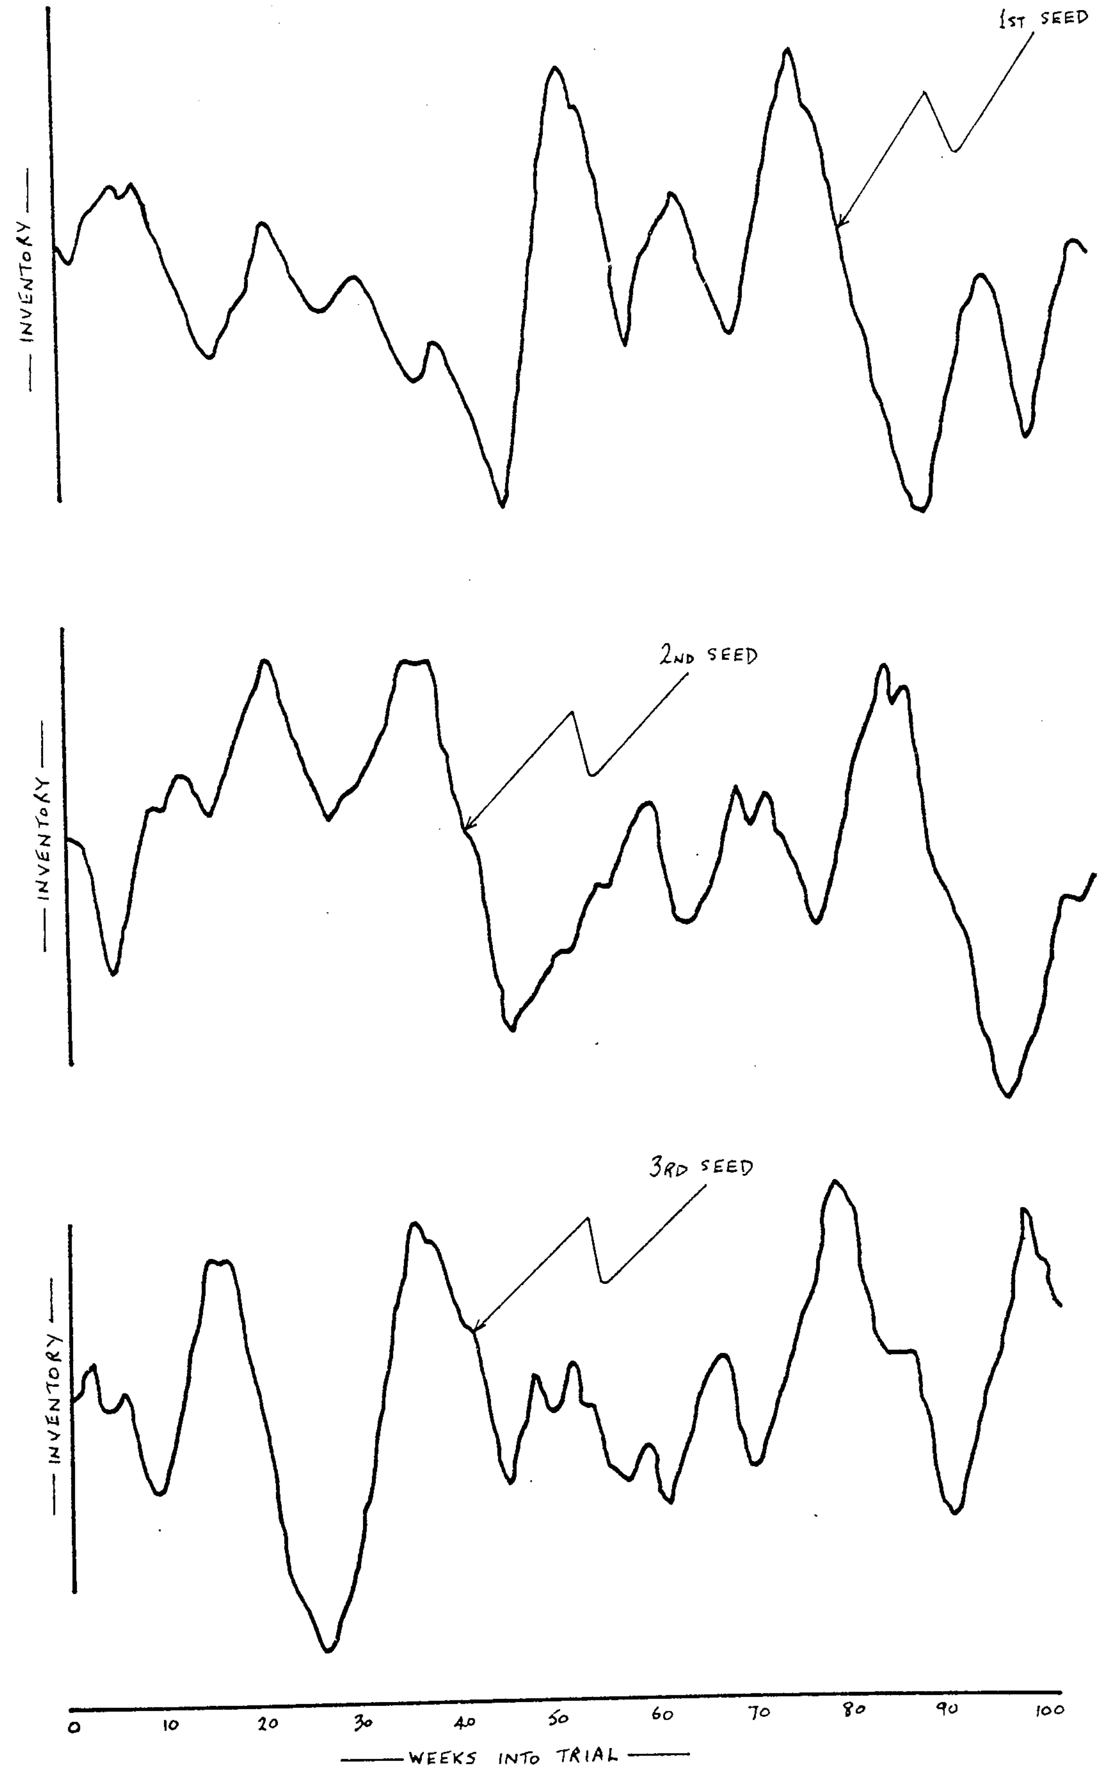

Fig. 4.15 Effect on Inventory at Distribution Centre of Random Deviations of $\pm 50\%$ on Weekly Orders

a series of shocks to the system which has a severe adverse effect on both service level and inventory levels. This can be mitigated by judicious advance buying.

4) Seasonality in trading activity has little effect on the dynamic behaviour of the system, assuming it takes a smooth wave form. It does produce a phase shift of almost 180° between the inventory and order peaks, but this in itself has little numerical effect on the average values of inventory levels and effective service levels.

5) The business cycle can be regarded as a succession of ramps with varying inclinations and declinations. In periods of volume contraction, the expected service levels will be higher than the target values and the inventory will be lower than the predicted level: the opposite will occur in periods of expansion.

6) There is no evidence of a natural cycle in the system. The internal feedback mechanisms produce amplification but little or no oscillation.

7) The effect of random noise is minimal. It is likely to be subordinated to the effect of quantising the variables (such as rounding to the nearest pack size). This will be explored in Chapter 5 with a model which is more appropriate for handling discrete quantities.

4.2 SENSITIVITY ANALYSIS

Section 4.1 examined the dynamic behaviour of the Branch/ Distribution Centre/Supplier system when subjected to various types of disturbance. This Section investigates the static effects on the service level and average stock level at a Distribution Centre by altering the values of the five main variables i.e. Mean Demand per unit time (μ_D), Standard Deviation of Demand per unit time (σ_D), Mean Lead Time (μ_L), Standard Deviation of Lead Time (σ_L), and Nominal Order Interval (I). The tests are conducted on a single product, the base values of the variables being:

$$\left.\begin{array}{lll} \mu_D & - & 40 \text{ units per week} \\ \sigma_D & - & 25.1488 \text{ units per week} \\ \mu_L & - & 4 \text{ weeks} \\ \sigma_L & - & 2 \text{ weeks} \end{array}\right\} \quad k = 2.989$$

I - 4 weeks

In the operational system, sales maximisation is achieved using a Lagrange Multiplier (λ) for each Product Group. Strictly, these should be recalculated when any variable changes its value. In practice, this would be prohibitive in computer time, and stock balancing is confined to once per month. Hence changes to variables have two distinct effects - an immediate effect and a long-term effect after the Lagrange Multipliers have been recalculated. The recalculation restores the target service levels to their correct values. Figs. 4.16 to 4.20 show the immediate effects of the change, and Figs. 4.21 to 4.25 show the long-term effects. In the latter set of graphs the Lagrange Multiplier values are shown in place of service levels. All results are based on initial target service levels of 90% and 95%, and each variable is deviated independently within the limits of 50% to 200% of its base value.

In order to generalise the results for all base data, the relationships in Johnston's Function 1 and Function 2 (Section 3.4) between Gamma shape factor (k), standardised reorder level (U) and probability of stockout (1-P) have been analysed. These are plotted in Figs. 4.26 to 4.31 and they are independent of base data values.

The basic equations required for the analysis are re-stated here for convenience. The minor effects of the review cycle of 1 day have been omitted from the analysis to simplify the algebra, but they are incorporated into the graphical results.

$$SL = 1 - \left(PLS \times \frac{\mu_L}{I} \right) \qquad (33)$$

$$AV = U\,\sigma_{DL} - \mu_{DL} + \mu_D\, I/2$$

$$\text{But } k = \frac{\mu_{DL}^2}{\sigma_{DL}^2}$$

$$\therefore AV = \mu_{DL}(U/\sqrt{k} - 1) + \mu_D\, I/2 \qquad (34)$$

$$k = \frac{\mu_D^2 \mu_L^2}{\mu_L \sigma_D^2 + \mu_D^2 \sigma_L^2} \qquad (35)$$

$$\lambda = (1-P)\ 250/I \qquad (36)$$

The relationships between U, k and (1-P) are first examined. As many of the sub-relationships act in opposition, the following relational operators have been devised to identify the dominant sub-relationships:

$x \ \alpha \ y$: x varies in the same direction as y by an equal factor

$x \ \overset{>}{\alpha} \ y$: x varies in the same direction as y by a larger factor

$x \ \overset{<}{\alpha} \ y$: x varies in the same direction as y by a smaller factor

$x \ \overset{\lessgtr}{\alpha} \ y$: x varies in the same direction as y by a larger, smaller or equal factor depending on data values

$x \ \alpha \ \frac{1}{y}$: x varies in the opposite direction to y by the same factor

$x \ \overset{>}{\alpha} \ \frac{1}{y}$: x varies in the opposite direction to y by a larger factor

$x \ \overset{<}{\alpha} \ \frac{1}{y}$: x varies in the opposite direction to y by a smaller factor

$x \ \overset{\lessgtr}{\alpha} \ \frac{1}{y}$: x varies in the opposite direction to y by a larger, smaller or equal factor depending on data values

$x \ \overset{\lessgtr}{\alpha} \ y$ or $\frac{1}{y}$: x varies in either direction relative to y depending on data values

$x \ \hat{\alpha} \ y$: x is not related to y

Johnston's Function 2, for a constant (1-P) reduces to the quadratic,

$PLS = a + \frac{b}{k} - \frac{c}{k^2}$ where a, b and c are positive constants.

Differentiating w.r.t. k and equating to zero:

$$\frac{d_{PLS}}{dk} = \frac{-b}{k^2} + \frac{2c}{k^3} = 0 \qquad \therefore k = \frac{2c}{b}$$

And $$\frac{d_2 PLS}{dk^2} = \frac{2b}{k^3} - \frac{6c}{k^4}$$

Substituting 2c/b for k: $$\frac{d_2 PLS}{dk^2} = \frac{2b(b^3)}{(2c)^3} - \frac{6c(b^4)}{(2c)^4} = \frac{-b^4}{8c^3}$$

The PLS therefore has a single turning point at $k = 2c/b$ and this is a maximum. Evaluating the constants reveals that the maximum PLS for $(1-P) \leqslant 0.5$ occurs for $k = 0.44$, i.e. outside the working region of the curves (Fig. 4.26). In the working region the graphical evidence indicates that:

$$PLS \overset{<}{\alpha} \frac{1}{k} \tag{A}$$

The relationships between PLS and (1-P) for $0.5 \leqslant k \leqslant 6$ are given in Fig. 4.27.

For $k = 1$ the gamma distribution reduces to a negative exponential, a characteristic of which is:

$$PLS = (1-P) \tag{B}$$

For $k < 1$, from graphical evidence,

$$PLS \overset{<}{\alpha} (1-P) \tag{C}$$

For $k > 1$, from graphical evidence,

$$PLS \overset{>}{\alpha} (1-P) \tag{D}$$

Fig. 4.29 shows that, except for very high values of (1-P),

$$U > \sqrt{k} \tag{E}$$

Fig. 4.30 indicates that, except for very high values of (1-P),

$$(U/\sqrt{k} - 1) \overset{<}{\alpha} \frac{1}{k} \tag{F}$$

Fig. 4.31 shows that:

$$(U/\sqrt{k} - 1) \overset{<}{\alpha} \frac{1}{(1-P)} \tag{G}$$

It can be seen by inspection of equation (35) that k is directly related to μ_D and μ_L and inversely related to σ_D and σ_L.

The degree of change in k relative to the independent variables is now investigated.

From equation (35), a factor change, x, to μ_L will produce a factor change x^2 in the numerator and a factor change $< x$ in the denominator.

$$\therefore \; k \overset{>}{\alpha} \mu_L \qquad \text{(H)}$$

If a factor change, x, is applied to μ_D, k could change by a bigger or smaller factor depending on the values of the constants. Assuming the factor change to k is equal to x :

$$\text{then, } xk = \frac{x^2\mu_D{}^2\mu_L{}^2}{\mu_L\,\sigma_D{}^2+x^2\mu_D{}^2\sigma_L{}^2}$$

$$\therefore \; \frac{x\mu_D{}^2\mu_L{}^2}{\mu_L\sigma_D{}^2+\mu_D{}^2\sigma_L{}^2} = \frac{x^2\mu_D{}^2\mu_L{}^2}{\mu_L\sigma_D{}^2+x^2\mu_D{}^2\sigma_L{}^2}$$

$$\therefore \; x(\mu_L\sigma_D{}^2+\mu_D{}^2\sigma_L{}^2) = \mu_L\sigma_D{}^2+x^2\mu_D{}^2\sigma_L{}^2$$

$$\therefore \; (x-1)\mu_L\sigma_D{}^2 = x(x-1)\mu_D{}^2\sigma_L{}^2$$

$$\therefore \; x = \frac{\mu_L\sigma_D{}^2}{\mu_D{}^2\sigma_L{}^2}$$

Therefore, if $x > \mu_L\sigma_D^2/\mu_D^2\sigma_L^2$ the factor change to k will be $<x$; and if $x < \mu_L\sigma_D^2/\mu_D^2\sigma_L^2$ the factor change to k will be $>x$ .With operational data, $\mu_L\sigma_D^2/\mu_D^2\sigma_L^2$ usually evaluates <0.5 (it is 0.395 for the sample base data). Therefore, in the general case, within the limits of variation $0.5 \leqslant x \leqslant 2.0$:

$$k \overset{<}{\alpha} \mu_D \qquad \text{(I)}$$

If a factor change, x, is applied to σ_D which results in a factor change $\frac{1}{x}$ to k:

$$\text{then,} \quad \frac{\mu_D^2\mu_L^2}{x(\mu_L\sigma_D^2+\mu_D^2\sigma_L^2)} = \frac{\mu_D^2\mu_L^2}{\mu_L x^2\sigma_D^2+\mu_D^2\sigma_L^2}$$

$$\therefore \quad x(\mu_L\sigma_D^2+\mu_D^2\sigma_L^2) = \mu_L x^2\sigma_D^2+\mu_D^2\sigma_L^2$$

$$\therefore \quad x(1-x)\ \mu_L\sigma_D^2 = (1-x)\ \mu_D^2\sigma_L^2$$

$$\therefore \quad x = \frac{\mu_D^2\sigma_L^2}{\mu_L\sigma_D^2}$$

which is the reciprocal of the expression obtained above when μ_D is varied, and usually evaluates > 2.0

$$\therefore k \overset{<}{\alpha} \frac{1}{\sigma_D} \qquad \text{(J)}$$

If a factor change, x, is applied to σ_L which results in a factor change $\frac{1}{x}$ to k:

then, $$\frac{\mu_D^2\mu_L^2}{x(\mu_L\sigma_D^2+\mu_D^2\sigma_L^2)} = \frac{\mu_D^2\ \mu_L^2}{\mu_L\sigma_D^2+\mu_D^2x^2\sigma_L^2}$$

$$\therefore \quad x(\mu_L\sigma_D^2+\mu_D^2\sigma_L^2) = \mu_L\sigma_D^2+\mu_D^2x^2\sigma_L^2$$

$$\therefore \quad (x-1)\mu_L\sigma_D^2 = x(x-1)\mu_D^2\sigma_L^2$$

$$\therefore \quad x = \frac{\mu_L\sigma_D^2}{\mu_D^2\sigma_L^2}$$

which is the same expression as for varying μ_D and usually evaluates < 0.5

$$\therefore \quad k \overset{>}{\alpha} \frac{1}{\sigma_L} \qquad \text{———— K}$$

Having derived the sub-relationships, these are now used to analyse the effects on service level and average stock of changing the values of the variables.

a) <u>Varying μ_D</u>

$k \overset{<}{\alpha} \mu_D$ (I)

$PLS \overset{<}{\alpha} \frac{1}{k}$ (A)

$$\therefore \quad PLS \overset{<}{\alpha} \frac{1}{\mu_D}$$

$$\frac{\mu_L}{I} \quad \hat{\alpha} \quad \mu_D$$

$$\therefore \text{ PLS } \frac{\mu_L}{I} \ \overset{<}{\alpha} \ \frac{1}{\mu_D}$$

$$SL \ \overset{<}{\alpha} \ \frac{1}{(PLS \ \mu_L/I)} \quad \text{for } SL > 0.5 \quad (33)$$

$$\therefore \quad \underline{\underline{SL \ \overset{<}{\alpha} \ \mu_D}}$$

$$k \ \overset{<}{\alpha} \ \mu_D \qquad (I)$$

$$(U/\sqrt{k}-1) \ \overset{<}{\alpha} \ \frac{1}{k} \qquad (F)$$

$$\therefore \ (U/\sqrt{k}-1) \ \overset{<}{\alpha} \ \frac{1}{\mu_D}$$

$$\mu_{DL} \ \alpha \ \mu_D$$

$\therefore \quad \mu_{DL}(U/\sqrt{k}-1) \ \overset{<}{\alpha} \ \mu_D$ (as the direct sub-relationship dominates).

$$\mu_D \ I/2 \ \alpha \ \mu_D$$

$$\therefore \quad \underline{\underline{AV \ \overset{<}{\alpha} \ \mu_D}}$$

The empirical results (Fig. 4.16) are consistent with these general relationships. The service level is seen to be quite sensitive to falling demand. After recalculating the Lagrange Multipliers (Fig. 4.21) the re-establishment of the service levels is seen to moderate the stock increases. The relationship between the recalculated Lagrange Multipliers is confirmed as follows:

For a constant (re-established) SL, PLS is constant (33)

$k \overset{<}{\alpha} \mu_D$ (I)

For a constant PLS, $(1-P) \overset{<}{\alpha} k$ (Fig. 4.27)

$\therefore (1-P) \overset{<}{\alpha} \mu_D$

$\lambda \; \alpha \; (1-P)$ (36)

$\therefore \underline{\underline{\lambda \overset{<}{\alpha} \mu_D}}$

This is entirely consistent with gamma distribution theory. A higher demand rate increases k, which produces a more symmetrical probability distribution with a shorter 'tail'. Hence the ELS given a stockout is less, so there must be more stockouts to compensate (for the same service level). Therefore λ (i.e. the number of stockouts per year) is increased.

b) <u>Varying σ_D</u>

$k \overset{<}{\alpha} \frac{1}{\sigma_D}$ (J)

$PLS \overset{<}{\alpha} \frac{1}{k}$ (A)

$\therefore PLS \overset{<}{\alpha} \sigma_D$

$$\frac{\mu_L}{I} \;\hat{\alpha}\; \sigma_D$$

$$\therefore \quad PLS \times \frac{\mu_L}{I} \;\overset{<}{\alpha}\; \sigma_D$$

$$SL \;\overset{<}{\alpha}\; \frac{1}{(PLS\ \mu_L/I)} \quad \text{for } SL > 0.5 \ (33)$$

$$\therefore \quad \underline{\underline{SL \;\overset{<}{\alpha}\; \frac{1}{\sigma_D}}}$$

$$k \;\overset{<}{\alpha}\; \frac{1}{\sigma_D} \qquad (J)$$

$$(U/\sqrt{k}-1) \;\overset{<}{\alpha}\; \frac{1}{k} \qquad (F)$$

$$\therefore \quad (U/\sqrt{k}-1) \;\overset{<}{\alpha}\; \sigma_D$$

$$\mu_{DL} \;\hat{\alpha}\; \sigma_D$$

$$\mu_D\ I/2 \;\hat{\alpha}\; \sigma_D$$

$$\therefore \quad \underline{\underline{AV \;\overset{<}{\alpha}\; \sigma_D}}$$

Fig. 4.17 confirms the above relationships. The immediate effect of an increase in σ_D is to provide a lower service level with slightly more stock. The increased stock level is expected, as σ_D has a direct effect on the buffer stock. The reduction in the service level is caused by σ_D decreasing k without the concomitant adjustment to (1-P) via λ. Hence the higher level of ELS per stockout arising from the more skewed probability distribution is not counteracted by a decrease in the expected frequency of stockouts. Fig. 4.22 shows that at the recalculation the number of stockouts per annum (λ) falls to restore the service level to its original target value. This gives rise to a further increase in the average stock to support the restored service level.

c) Varying μ_L

$k \overset{>}{\alpha} \mu_L$ (H)

$PLS \overset{<}{\alpha} \frac{1}{k}$ (A)

$\therefore\ PLS \overset{<}{\underset{\alpha}{>}} \frac{1}{\mu_L}$

$\frac{\mu_L}{I} \alpha\ \mu_L$

$\therefore\ PLS \quad \frac{\mu_L}{I} \overset{<}{\underset{\alpha}{>}} \mu_L \text{ or } \frac{1}{\mu_L}$

$SL \overset{<}{\alpha} \frac{1}{(PLS\ \mu_L/I)}$ for SL >0.5 (33)

$\therefore\ SL \overset{<}{\underset{\alpha}{>}} \mu_L \text{ or } \frac{1}{\mu_L}$

$k \overset{>}{\alpha} \mu_L$ (H)

$(U/\sqrt{k}-1) \overset{<}{\alpha} \frac{1}{k}$ (F)

$\therefore\ (U/\sqrt{k}-1) \overset{<}{\underset{\alpha}{>}} \frac{1}{\mu_L}$

$$\mu_{DL} \ \alpha \ \mu_L$$

$$\therefore \quad \mu_{DL}(U/\sqrt{k}-1) \ \overset{\lessgtr}{\alpha} \ \mu_L \text{ or } \frac{1}{\mu_L}$$

$$\mu_D \ I/2 \ \hat{\alpha} \ \mu_L$$

$$\therefore \ \underline{\underline{AV \ \overset{\lessgtr}{\alpha} \ \mu_L \text{ or } \frac{1}{\mu_L}}}$$

Figs. 4.18 and 4.23 reflect the immediate and long-term effects of the interplay between the opposing sub-relationships and the uncertainty of the outcome. Both sets of results contain counter-intuitive elements. The immediate effect of improving the lead time is a deterioration in the service level; and in the long term there is no significant reduction in stock - in fact under certain circumstances a marginal increase ensues. The inconclusiveness in the final relationships implies that the outcome of varying the mean lead time is contingent upon data values. This uncertain system behaviour could easily evoke inappropriate buyer action.

d) <u>Varying σ_L</u>

$$k \ \overset{>}{\alpha} \ \frac{1}{\sigma_L} \qquad (K)$$

$$PLS \ \overset{<}{\alpha} \ \frac{1}{k} \qquad (A)$$

$$\therefore \quad PLS \ \overset{\lessgtr}{\alpha} \ \sigma_L$$

$$\frac{\mu_L}{I} \hat{\alpha} \sigma_D$$

$$\therefore \text{PLS} \cdot \frac{\mu_L}{I} \underset{\alpha}{\lessgtr} \sigma_L$$

$$\text{SL} \overset{<}{\alpha} \frac{1}{(\text{PLS } \mu_L/I)} \text{ for SL>0.5} \qquad (\ 33\ )$$

$$\therefore \underline{\underline{\text{SL} \underset{\alpha}{\lessgtr} \frac{1}{\sigma_L}}}$$

$$k \overset{>}{\alpha} \frac{1}{\sigma_L} \qquad (\ K\ )$$

$$(U/\sqrt{k}-1) \overset{<}{\alpha} \frac{1}{k} \qquad (\ F\ )$$

$$\therefore (U/\sqrt{k}-1) \underset{\alpha}{\lessgtr} \sigma_L$$

$$\mu_{DL} \hat{\alpha} \sigma_L$$

$$\mu_D \; I/2 \; \hat{\alpha} \; \sigma_L$$

$$\therefore \underline{\underline{\text{AV} \underset{\alpha}{\lessgtr} \sigma_L}}$$

Figs. 4.19 and 4.24 demonstrate that the system is extremely sensitive to changes in σ_L . The immediate effect of increasing σ_L is a substantial deterioration in service level, and the longer term effect is an exponential escalation of the stock level. There are two main reasons for the large increase in stock:

i) σ_L is a potent factor in determining the Buffer Stock. The standard deviation of demand in lead time is an important factor, which is calculated as:

$$\sigma_{LD} = \sqrt{\mu_L \sigma_D^{\,2} + \mu_D^{\,2} \sigma_L^{\,2}}$$

(This assumes no cross-correlations or auto-correlations of demand and lead time - refer Sect. 6.3)

$\mu_D^{\,2}\,\sigma_L^{\,2}$ (and thus σ_L) exerts a dominant influence as μ_L is the only variable which is not squared.

ii) There is a strong 'backlash' effect of the service level restoration. Doubling σ_L causes the 95% target service level to fall to 88%. Stock Level (plotted vertically)/ Service Level curves rise very steeply in this region.

e) <u>Varying I</u>

(i) <u>For k<1</u>

$(1-P)\ \alpha\ I$ (36)

$PLS\ \overset{<}{\alpha}\ (1-P)$ (B)

$\therefore\ PLS\ \overset{<}{\alpha}\ I$

$\dfrac{\mu_L}{I}\ \alpha\ \dfrac{1}{I}$

$\therefore\ PLS\ \dfrac{\mu_L}{I}\ \overset{<}{\alpha}\ \dfrac{1}{I}$

$SL\ \overset{<}{\alpha}\ \dfrac{1}{(PLS\ \mu_L/I)}$ (33)

$\therefore$ <u>$SL\ \overset{<}{\alpha}\ I$</u>

(ii) <u>For k=1</u>

$(1-P)\ \alpha\ I$ (36)

$PLS = (1-P)$ (C)

$\therefore\ PLS\ \alpha\ I$

$\dfrac{\mu_L}{I}\ \alpha\ \dfrac{1}{I}$

$\therefore\ PLS\ \dfrac{\mu_L}{I}$ is constant

$SL = 1-$(a constant)

$\therefore$ <u>SL is constant.</u>

(iii) <u>For k>1</u>

$(1-P)\ \alpha\ I$ (36)

$PLS\ \overset{>}{\alpha}\ (1-P)$ (D)

$\therefore\ PLS\ \overset{>}{\alpha}\ I$

$\dfrac{\mu_L}{I}\ \alpha\ \dfrac{1}{I}$

$\therefore\ PLS\ \dfrac{\mu_L}{I}\ \overset{<}{\alpha}\ I$

$SL\ \overset{<}{\alpha}\ \dfrac{1}{(PLS\ \mu_L/I)}$ (33)

$\therefore$ <u>$SL\ \overset{<}{\alpha}\ \dfrac{1}{I}$</u>

$(1-P) \; \alpha \; I \qquad (36)$

$(U/\sqrt{k}-1) \overset{<}{\alpha} \frac{1}{(1-P)} \qquad (G)$

$\therefore \quad (U/\sqrt{k}-1) \overset{<}{\alpha} \frac{1}{I}$

$\mu_{DL} \; \hat{\alpha} \; I$

$\therefore \; \mu_{DL} \; (U/\sqrt{k}-1) \overset{<}{\alpha} \frac{1}{I}$

$\mu_D \; I/2 \; \alpha \; I$

$\therefore \; \underline{\underline{AV \underset{\alpha}{\overset{<}{>}} I \text{ or } \frac{1}{I}}}$

Fig. 4.20 confirms the immediate effects i.e. increasing the order interval reduces the service level (as $k > 1$) and there are clearly counteracting influences on the average stock. The long-term effects (Fig. 4.25) show little change as the service level correction is relatively small. As the order interval does not affect k, its influence on system behaviour is limited and it may be varied by the System Controllers with relative impunity. (The consequences of reducing it below the lead time are significant and are dealt with in Section 6.4).

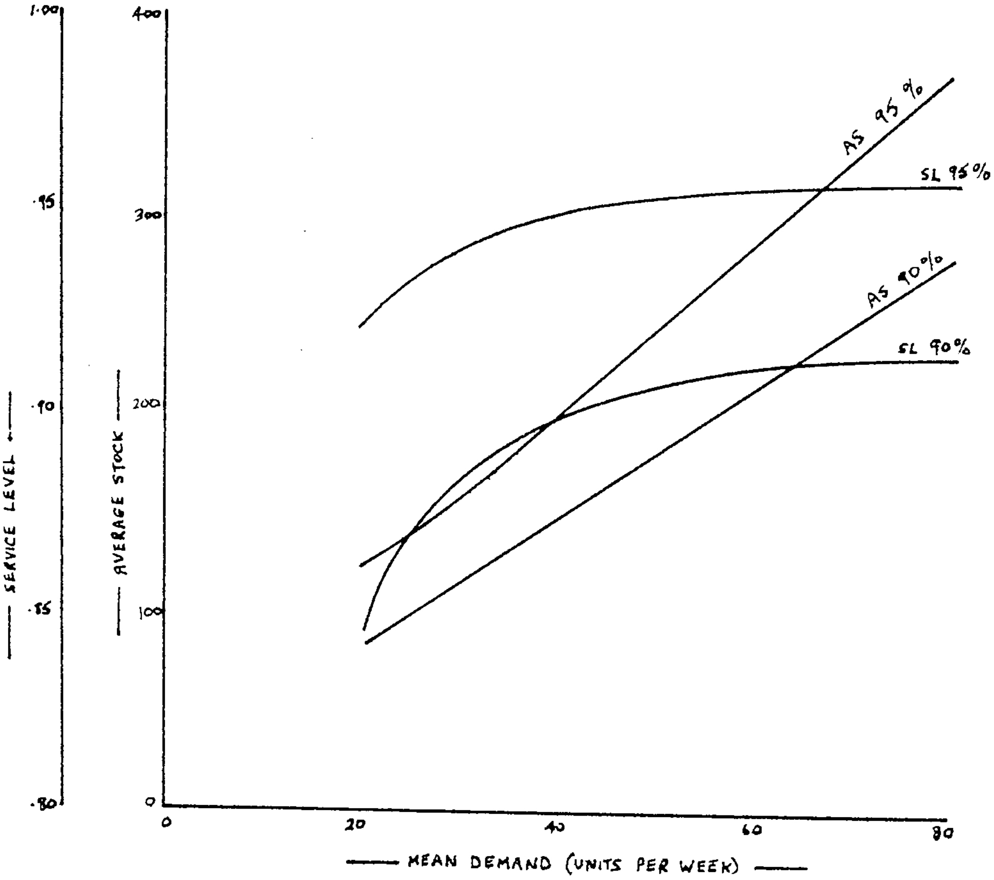

Fig. 4.16 Sensitivity of Service Level and Average Stock to Mean Demand Changes without Recalculation of λ

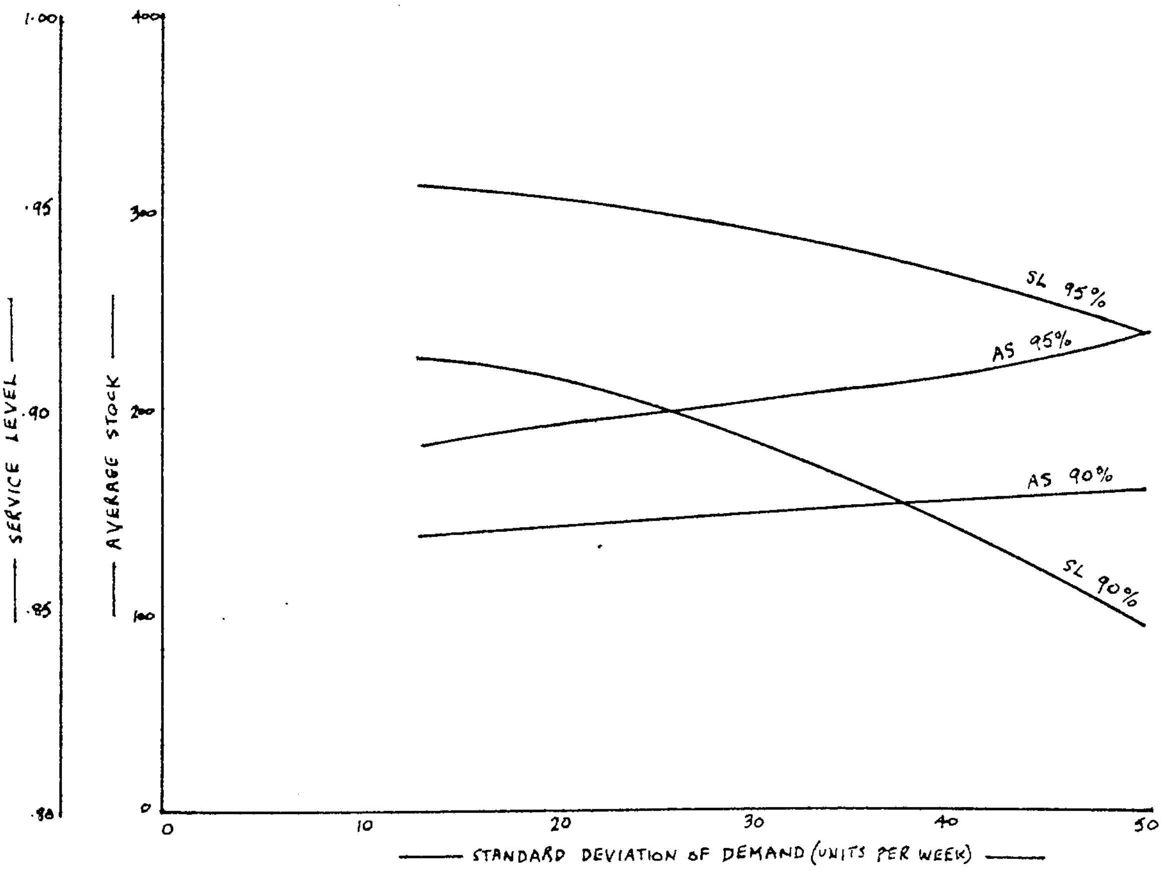

Fig. 4.17 Sensitivity of Service Level and Average Stock to Std. Dev. of Demand Changes without Recalculation of λ

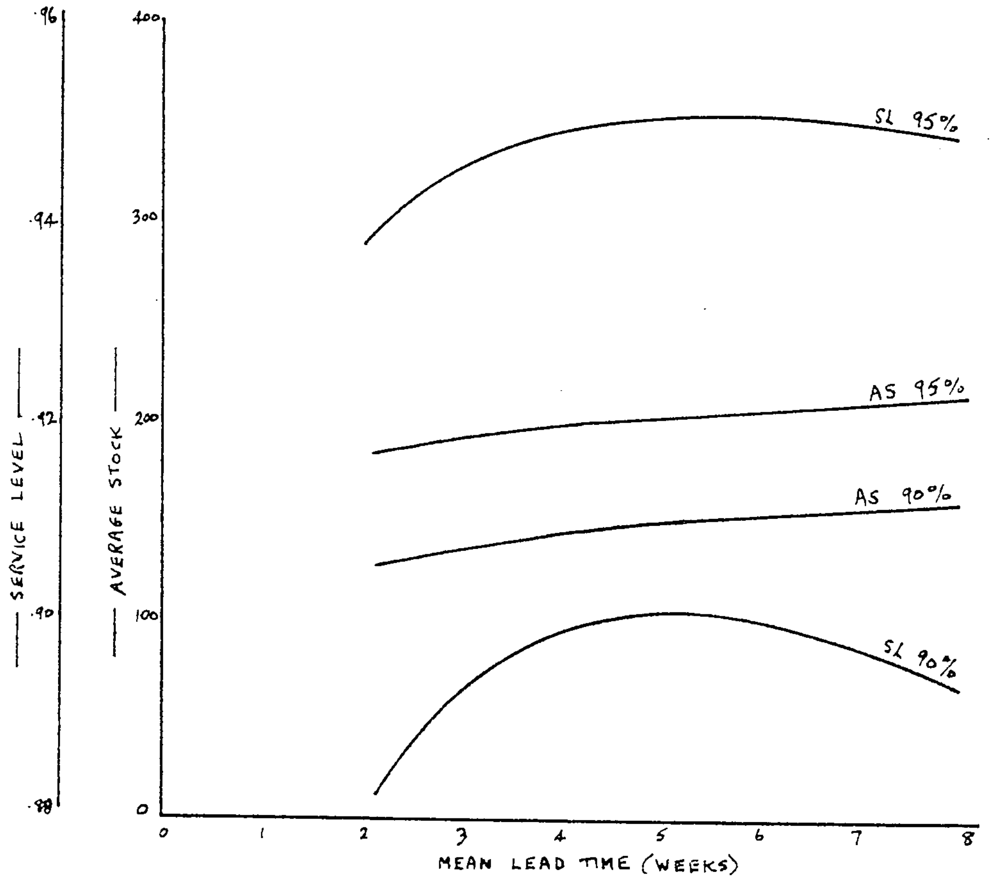

Fig. 4.18 Sensitivity of Service Level and Average Stock to Mean Lead Time Changes without Recalculation of λ

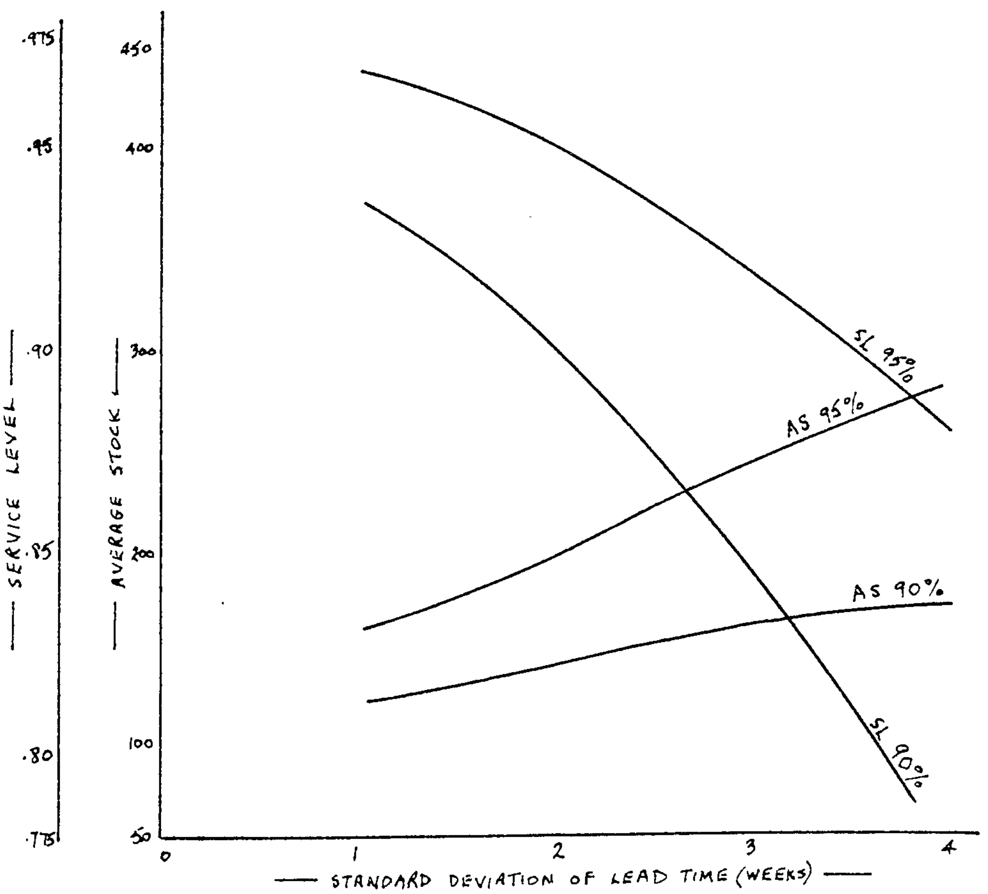

Fig. 4.19 Sensitivity of Service Level and Average Stock to Std. Dev. of Lead Time Changes without Recalculation of λ

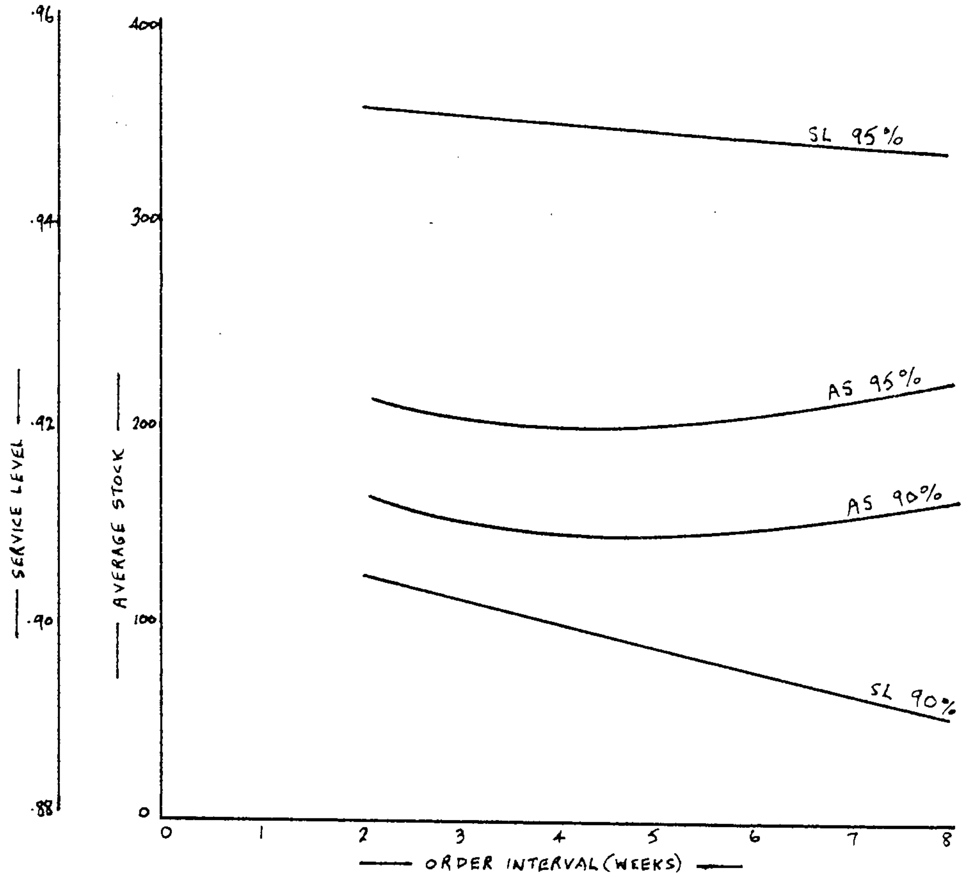

Fig. 4.20 Sensitivity of Service Level and Average Stock to Order Interval Changes without Recalculation of λ

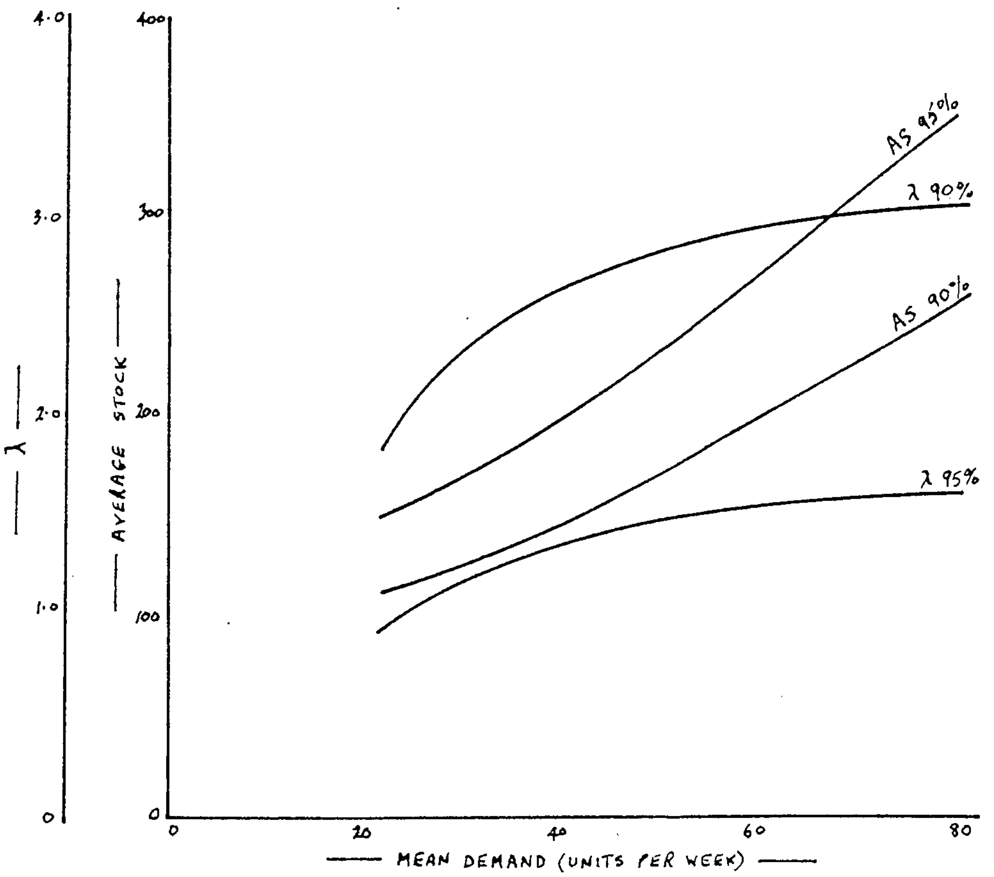

Fig. 4.21 Sensitivity of Average Stock and λ to Mean Demand Changes with Recalculation of λ

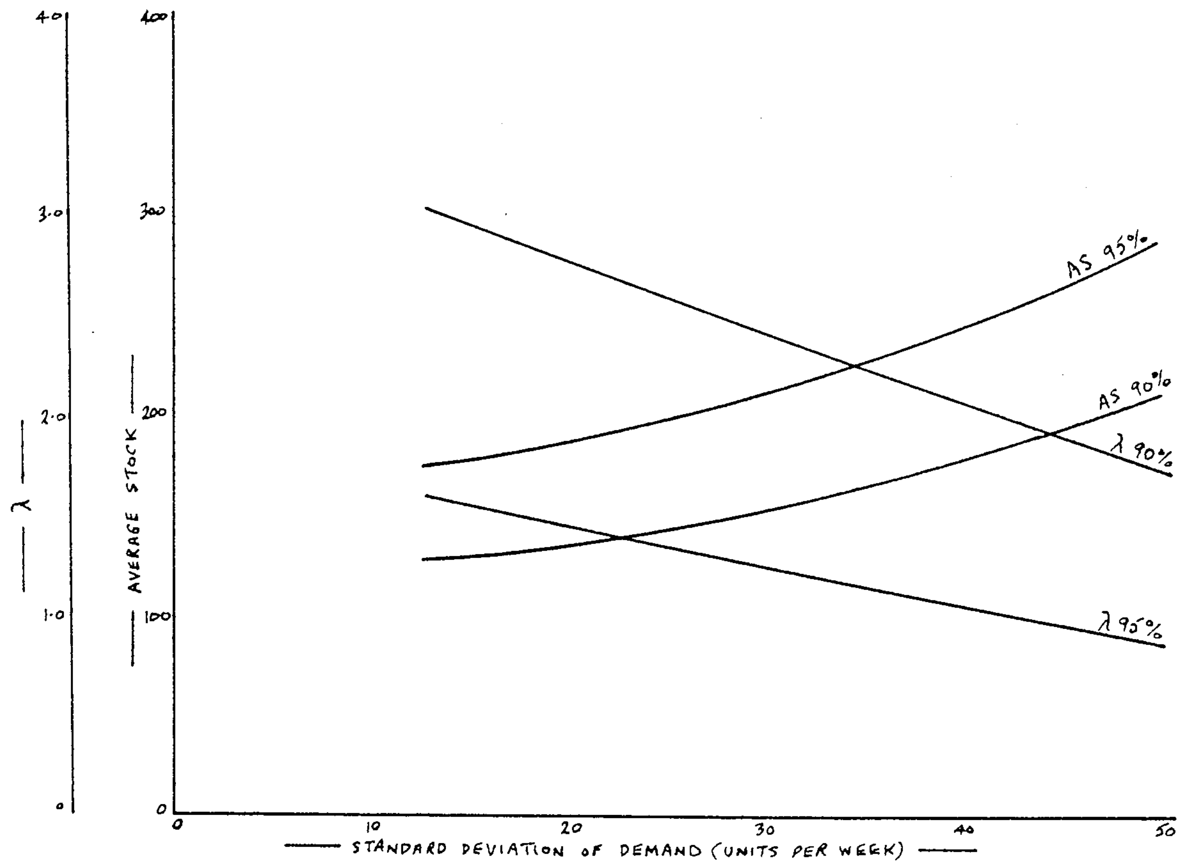

Fig. 4.22 Sensitivity of Average Stock and λ to Std. Dev. of Demand Changes with Recalculation of λ

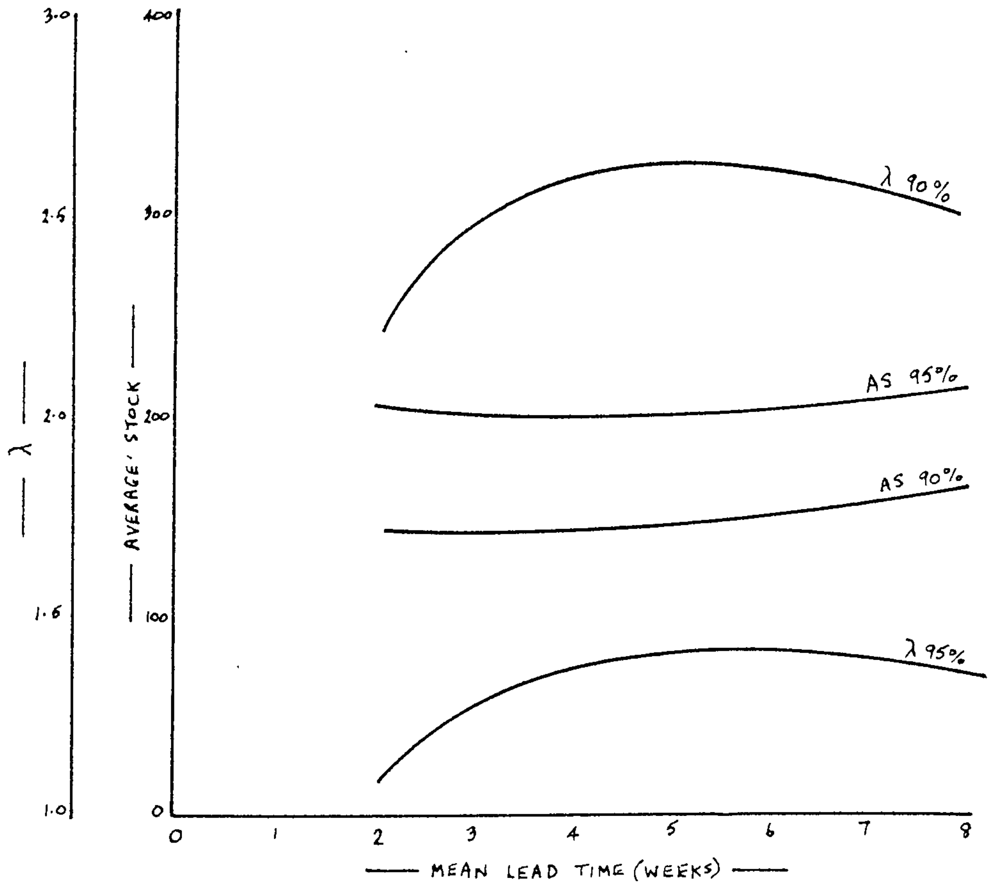

Fig. 4.23 Sensitivity of Average Stock and λ to Mean Lead Time Changes with Recalculation of λ

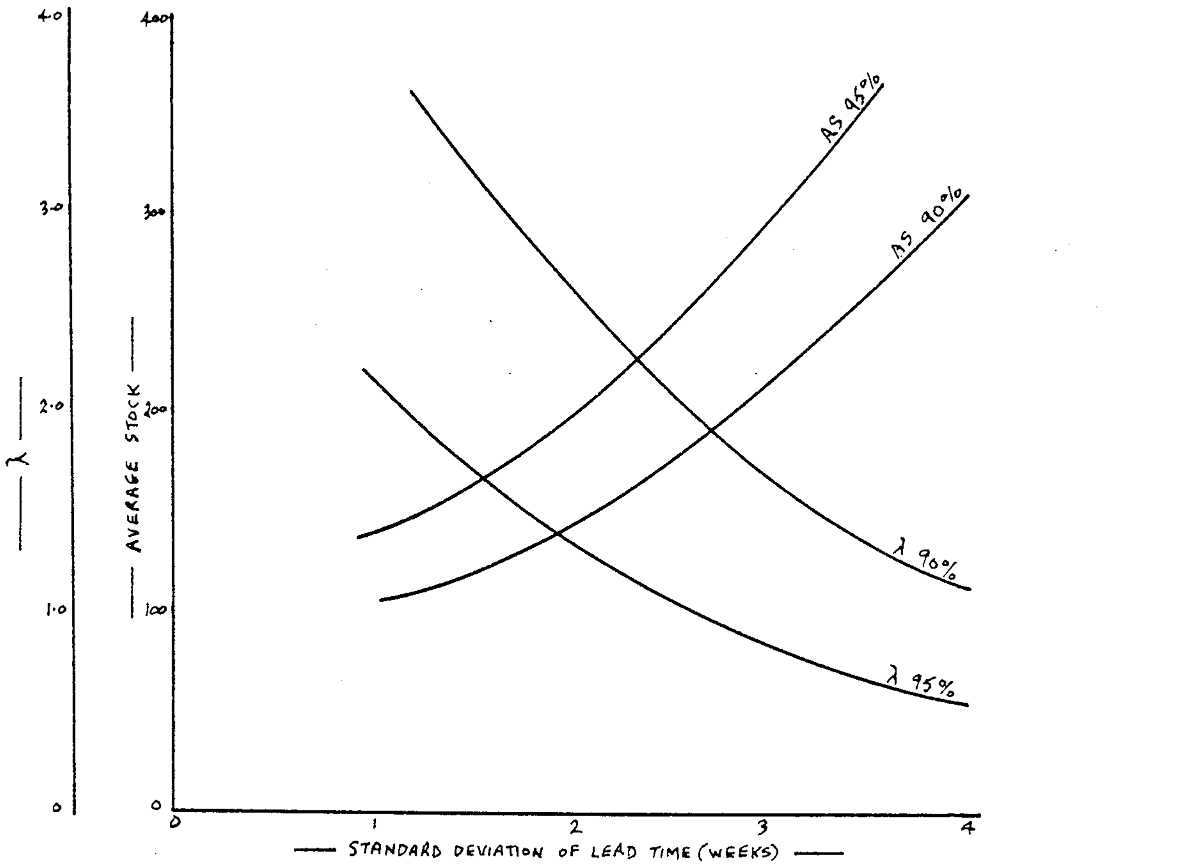

Fig. 4.24 Sensitivity of Average Stock and λ to Std. Dev. of Lead Time Changes with Recalculation of λ

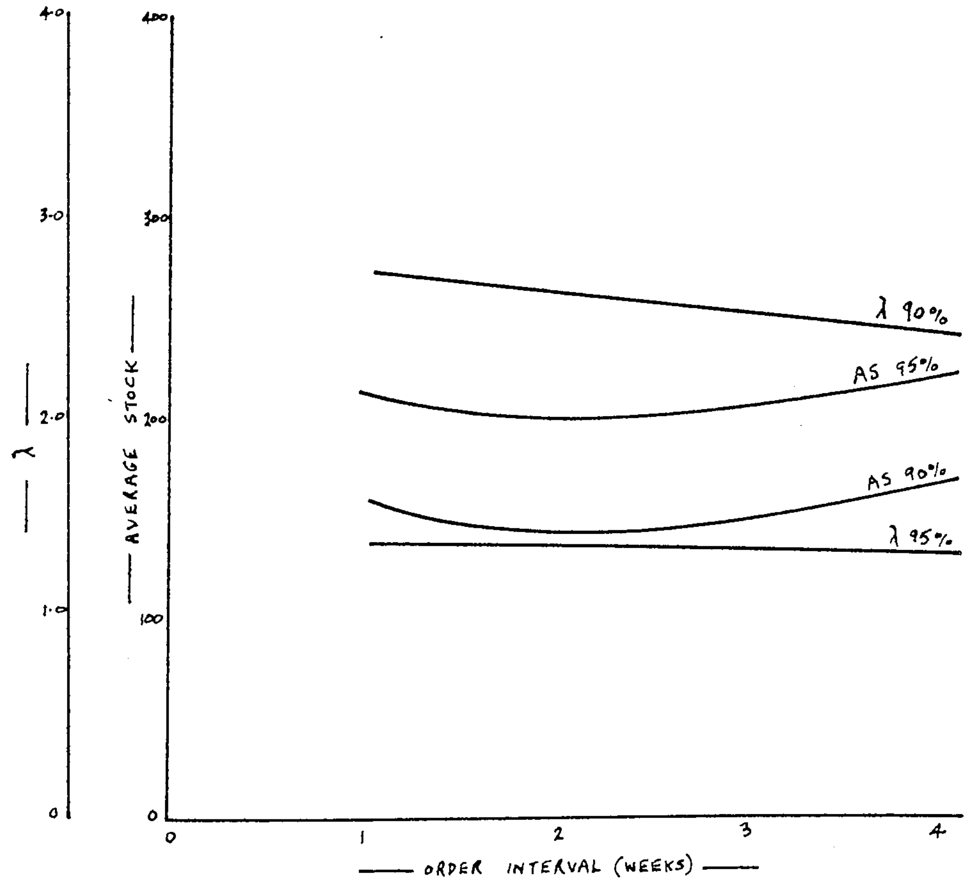

Fig. 4.25 Sensitivity of Average Stock and λ to Order Interval Changes with Recalculation of λ

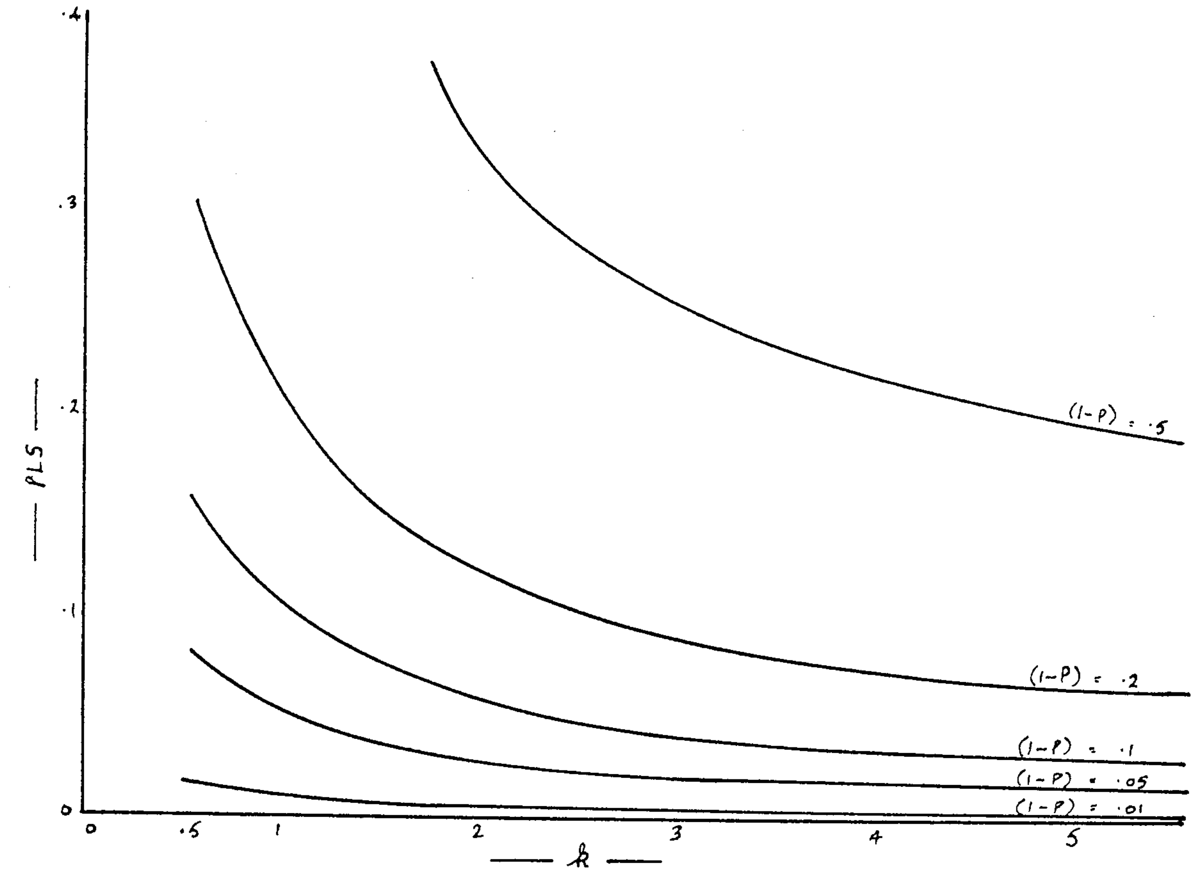

Fig. 4.26 Relationship Between PLS and k for Constant Probability of Stockout

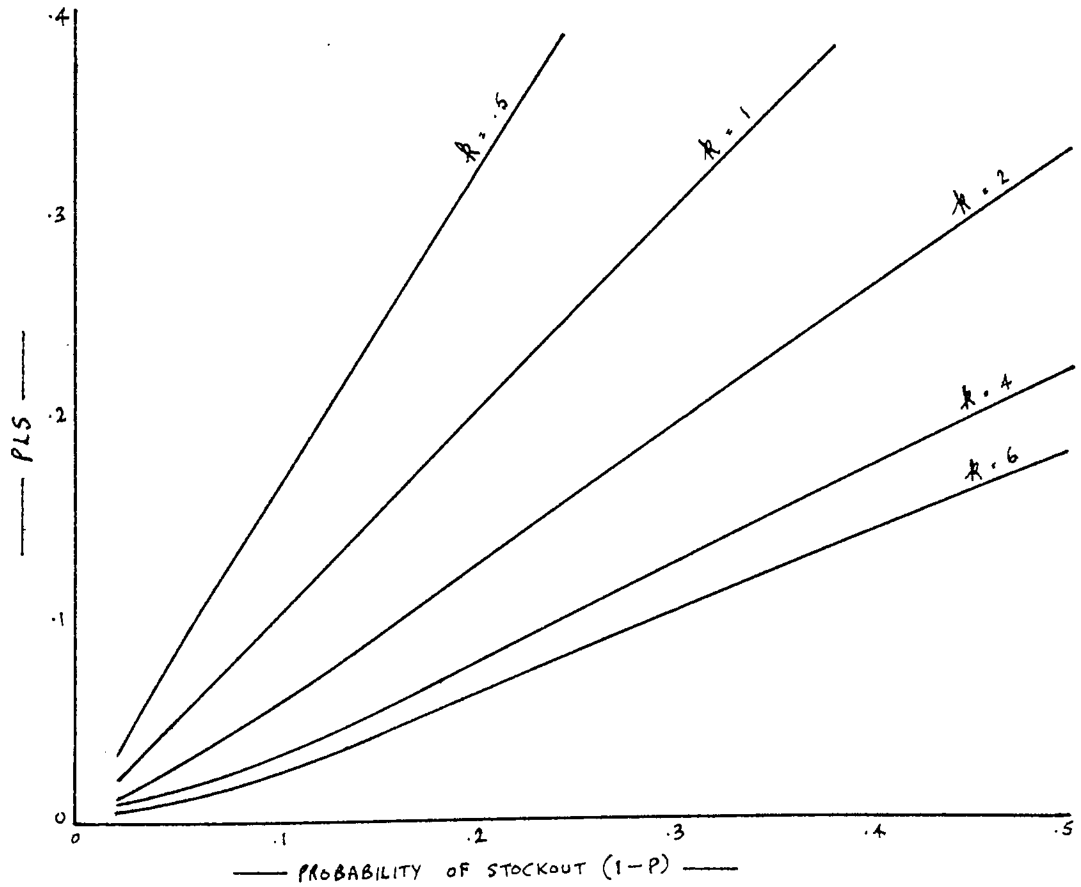

Fig. 4.27 Relationship Between PLS and Probability of Stockout for Constant k

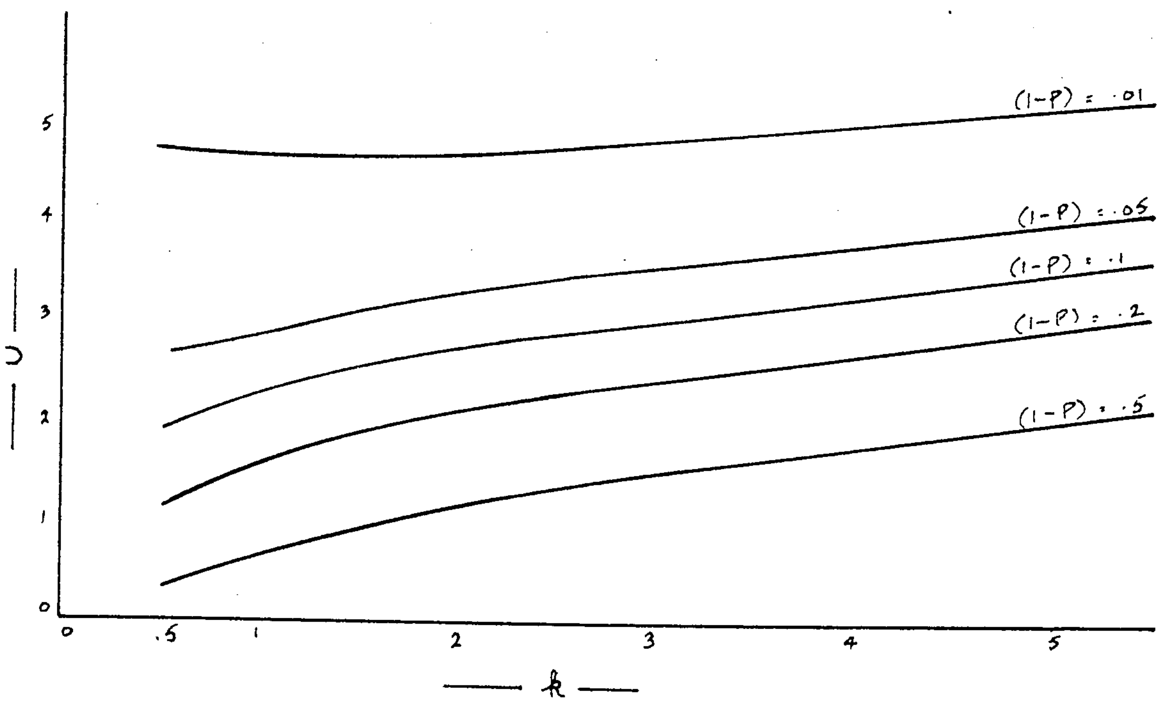

Fig. 4.28 Relationship Between Standardised Reorder Level (U) and k for Constant Probability of Stockout

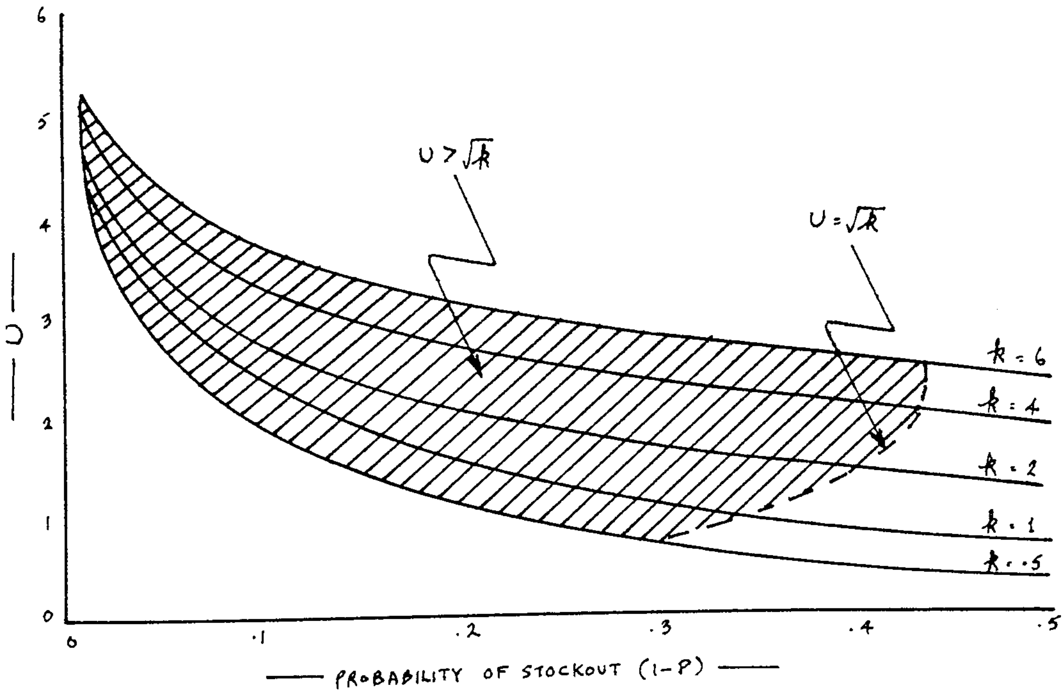

Fig. 4.29 Relationship Between Standardised Reorder Level (U) and Probability of Stockout for Constant k

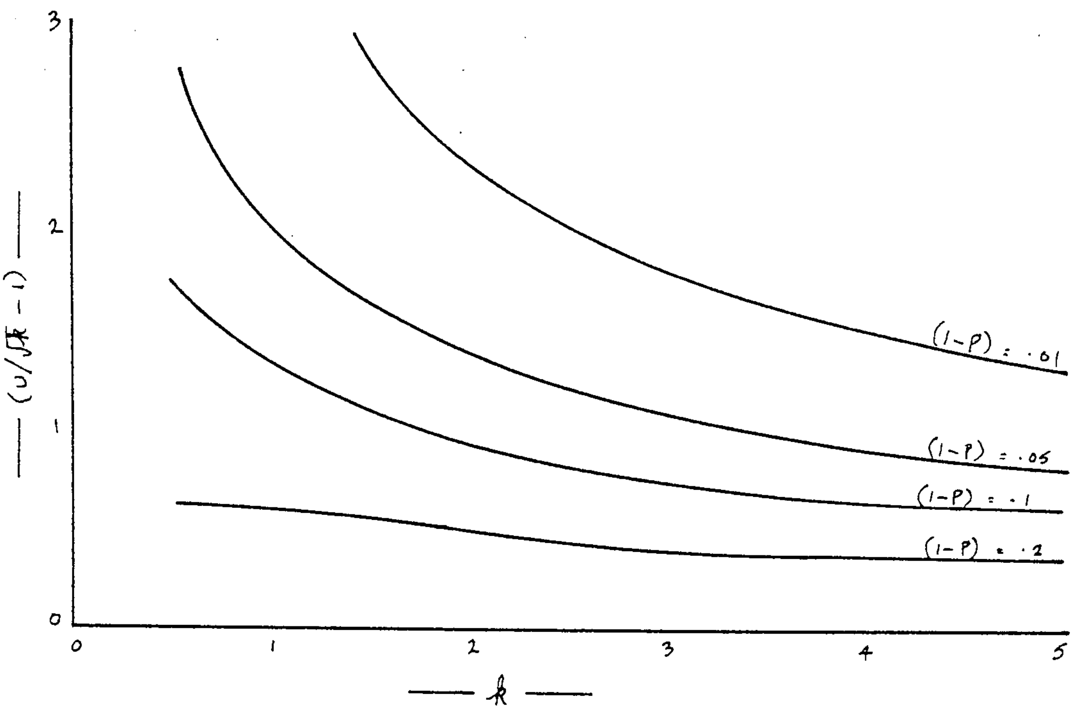

Fig. 4.30 Relationship Between (U/ $\sqrt{k}$-1) and k for Constant Probability of Stockout

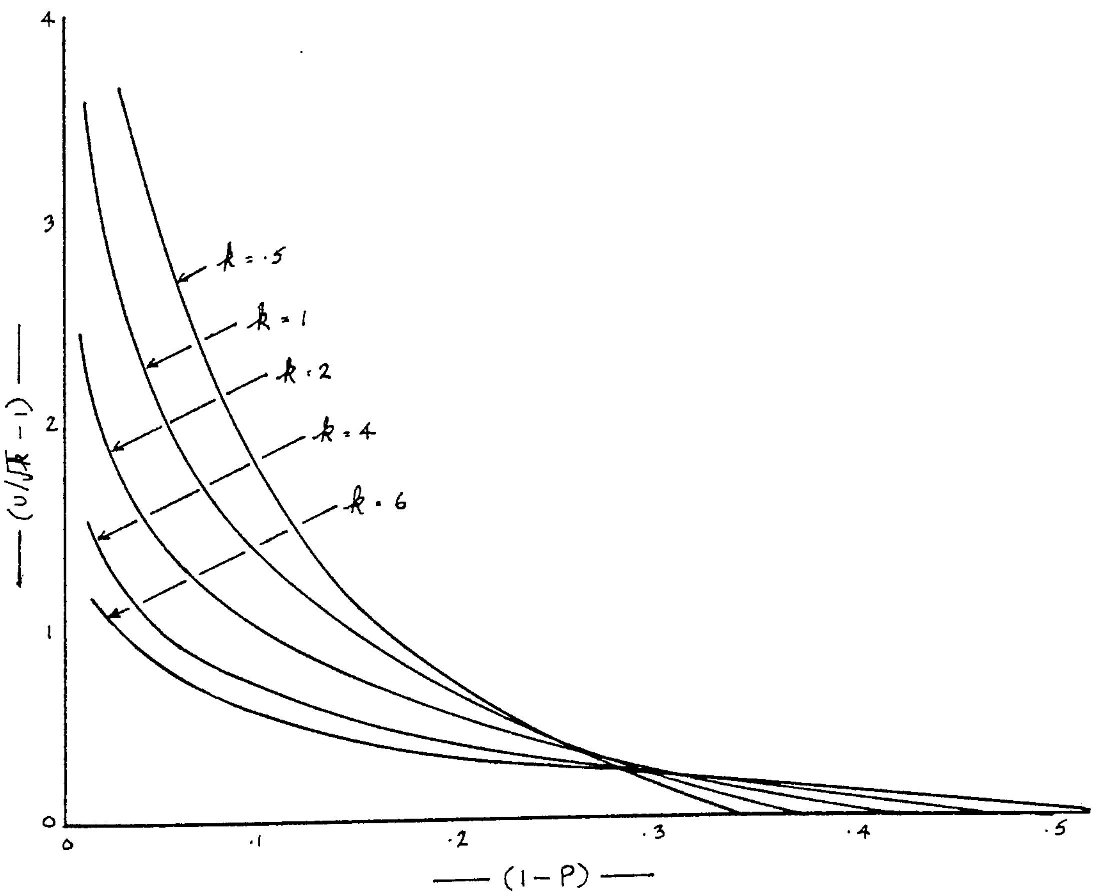

Fig. 4.31 Relationship Between (U/ $\sqrt{k}$ -1) and Probability of Stockout for Constant k

Summary

1. In most cases, the long-term effects of changes to variables are significantly different from the immediate effects. The long-term effects on average stock are usually more severe than the immediate effects.

   However, in the operational system the value of a Lagrange Multiplier is the result of the state of the variables for all products in a Product Group. A change to the state of a variable for an individual product will therefore have only a marginal effect on the Lagrange Multiplier, but this marginal effect will be proliferated across all of the products in the Group. The effect on the individual product will be similar to the short-term effect but with the service level restored.

2. The short-term effects on service level are quite serious in some cases. To the extent that the performance of the system can be perceived as a time-integration of the short-term effects, this could be a significant factor in the overall discrepancy between target and achieved service levels. Stock levels would also be affected by the delayed effect of the recalculation. Changes tending to alter stock levels generally have a muted effect before the recalculation takes place.

3. There are two cases where the system exhibits counter-intuitive behaviour even in the long-term:

   i) Improved lead times, if they are not accompanied by reduced standard deviations, can under some circumstances cause a marginal increase in stock.

   ii) Reducing order intervals can cause stocks to increase.

4. The standard deviation of lead time is by far the most potent

variable in changing the state of the service level and average stock. The attainment of consistent lead times is therefore a far more worthy objective than the attainment of short lead times.

5. There is an optimum value of the order interval for achieving stock minimisation which is not incorporated into the stock balancing optimisation function. However, the Average Stock/ Order Interval curve is quite flat around the minimum point, so the consequences of operating on either side of the optimum are not likely to be serious.

CHAPTER 5

THE USE OF SIMULATION TECHNIQUES IN THE STUDY

5.1 REQUIREMENT FOR SIMULATION AND MODELS USED

Simulation techniques have been used extensively to solve stock control problems over the past two decades. A comprehensive coverage of the use of simulation is given by Tocher (78). Most early applications used differential analysers (e.g. Lewis (79)), but much of the later simulation work has been carried out on the more versatile and readily available digital computers.

Simulation has been used in this study:-

a) to verify that the mathematical formulae produce the theoretical results where the formulae are exact,

b) to check the magnitude of errors produced by inexact numerical methods where an analytical solution is not possible,

c) to observe the dynamic behaviour of the system,

d) to test the effects of proposed system changes without incurring commercial risk,

e) to test the long-run effects of the system.

A System Dynamics simulation model has already been described in Chapter 4. Three other simulation programs of varying complexity were also constructed for more general use:-

1. SIMSIMPLE. This program is written in BASIC for operation on a microprocessor. It handles a single product only. The order trigger takes the form of a simple fixed reorder level in units which must be calculated 'off-line' from the desired Percentage Lost Sales per lead time (PLS) and Gamma modulus (k) using equations 18 and 19 (Sect. 3.4) in iterative mode. The order size is calculated as the reorder level plus the average sales per order cycle less the available stock. Demand and lead

time variates are generated separately from a routine which calculates Gamma random variates assuming a stationary distribution.

2. IMSIM. This program is also written in BASIC for operation on a microprocessor. It is also confined to single-product operation, but it reorders using the three Johnston functions in a manner identical to the operational system. The forecasting mechanism is also replicated from the live system. The Lagrange Multiplier must be calculated 'off line' and input. A Gamma generator is incorporated for producing demand and lead time data. The program has printing facilities at full, abridged and summary levels.

   A companion program, IMSIMB, operates in an identical fashion but allows the following Buyer overrides to be applied in interactive mode:-

   i) Cancel order
   ii) Change order quantities
   iii) Raise manual order
   iv) Change lead time forecast

3. FT01. This program is written in FORTRAN for mainframe operation. It replicates all of the main facilities in the operational system, including stock balancing for up to five products. Order generation and forecasting routines are replicated from the operational system. Demand and lead time variates may be generated from a Gamma, Normal or Observed distribution. All of the buyer overrides permitted in the operational system are available. Graph plotting facilities are incorporated as well as reporting at full and summary levels.

5.2 SELECTION OF A GAMMA GENERATOR

Several algorithms exist for generating Gamma variates. They have been compared by Atkinson (80) and by Atkinson and Pearce (81). In this work Atkinson's algorithm is used for $0 < k \leqslant 1$ and Ahren's (82) GS algorithm for $k > 1$. These were the most efficient algorithms for the ranges of k most frequently encountered in the system. The two selected algorithms were integrated into the computer routine given in Appendix 4. This is suitable for all values of k but is most efficient for $0 < k \leqslant 4$.

The random numbers generated from the microprocessor software were found to contain an unacceptable bias. This was overcome by the simple expedient of removing 0, 1, 2 and 3 significant digits from consecutive random numbers. The mainframe random number generator was free from bias and the output did not require scrambling.

The output from the Gamma generator is displayed in Fig. 5.1 for k values of 0.2, 0.5, 1.0, 2.0, 4.0 and 8.0. The same data is subjected to χ^2 tests (Table 5.1) over the apposite region of the curves ($P > 0.5$). An acceptably good fit is indicated for all values of k, and, most importantly for reorder level calculations, there is no evidence of lengthening or truncation of the tails compared with the theoretical curves. It is concluded that the algorithm would not introduce any significant errors into the simulation runs.

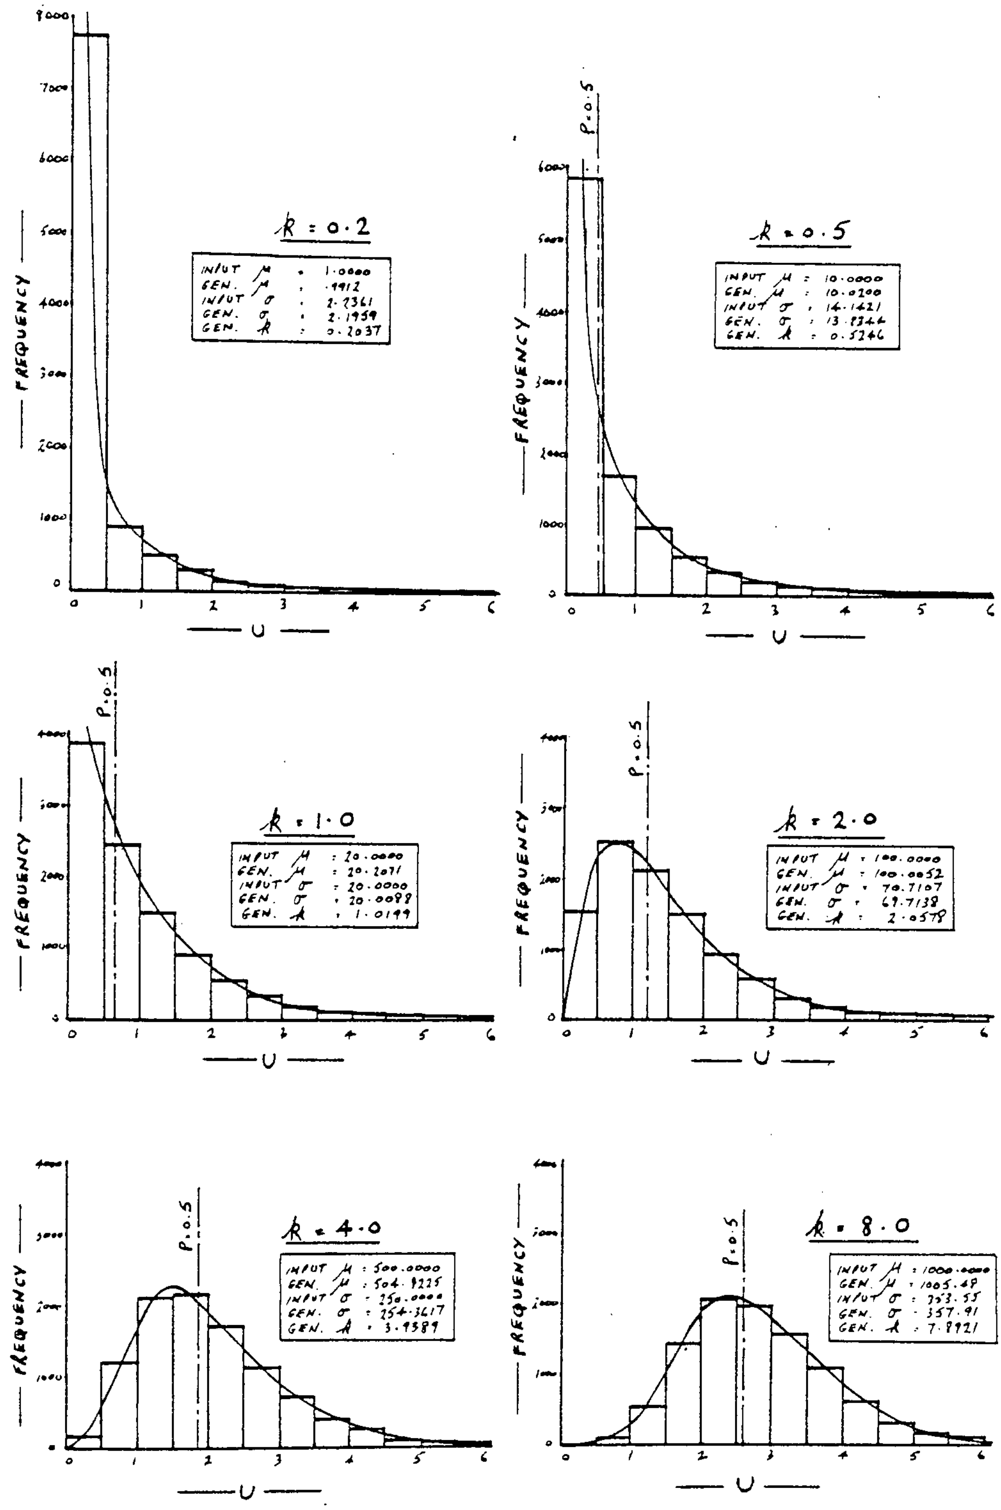

Fig. 5.1 Comparison Between Output of Gamma Generator and Theoretical Gamma Distribution (Histogram = Generator, Graph = Theoretical)

k	Obs.(O)/ Exp.(E)	Class Interval (= U/2)																		χ^2 (D.F.)
		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17+	
0.2	O	7781	885	476	272	168	112	79	62	34	27	24	25	11	10	8	7	5	14	11.35
	E	7788	869	452	274	179	121	85	60	44	32	24	17	14	9	8	5	5	14	(13)
	$(O-E)^2/E$	.01	.29	1.27	.01	.68	.67	.42	.27	2.27	.78	0	3.76	.64	.11	0		.17		
0.5	O	5922	1688	941	536	336	201	128	81	50	43	21	18	14	11	4	4	1	1	11.73
	E	5996	1660	891	527	325	207	133	87	58	38	25	17	12	7	6	3	3	5	(13)
	$(O-E)^2/E$	.91	.47	2.81	.15	.37	.17	.19	.41	1.10	.66	.64	.06	.33	2.29	.67		.50		
1.0	O	3845	2419	1482	896	543	323	192	109	67	50	36	12	13	5	2	2	2	3	9.83
	E	3935	2386	1448	878	532	323	196	119	72	44	26	16	10	6	3	3	1	2	(11)
	$(O-E)^2/E$	-	.46	.80	.22	.23	0	.08	.84	.35	.82	3.85	1.00	.90	.17			.11		
2.0	O	1561	2528	2134	1512	961	580	306	183	110	64	36	12	9	3	0	1	0	0	9.92
	E	1583	2548	2127	1479	941	569	331	189	106	58	32	17	10	5	2	1	1	1	(9)
	$(O-E)^2/E$	-	-	.02	.74	.43	.21	1.89	.19	.15	.62	.50	1.47	.10			3.60			
4.0	O	181	1196	2104	2144	1681	1144	696	398	240	97	63	21	18	6	6	2	3	0	13.86
	E	190	1239	2099	2137	1685	1138	694	394	212	109	54	26	12	6	3	1	1	0	(8)
	$(O-E)^2/E$	-	-	-	.02	.01	.03	.01	.04	3.70	1.32	1.50	.96	3.00			3.27			
8.0	O	0	75	558	1424	2058	1969	1572	1069	630	326	173	87	32	11	11	2	1	2	9.92
	E	1	85	583	1434	2016	2006	1581	1053	619	330	162	75	32	14	6	2	1	0	(7)
	$(O-E)^2/E$	-	-	-	-	-	.68	.05	.19	.20	.05	.75	1.92	0	.64		5.44			

(Note: Relevance of fit is confined to the right of the heavy line)

Table 5.1 χ^2 Tests for Generated Gamma Variates

5.3 STATISTICAL CONVOLUTIONS TO GENERATE A GAMMA DISTRIBUTION OF DEMAND IN LEAD TIME

The control system is based upon the theory that the demand whilst awaiting replenishment assumes a Gamma probability distribution. The corroboration of this assumption is treated in Section 6.1. For simulation purposes it is obviously necessary to generate demand and lead time data separately. The objective, therefore, is to convolve two probability distributions to produce a Gamma output, and for the two input distributions to represent adequately demand and lead time data respectively.

Burgin (62) has indicated that it is not possible to obtain an explicit expression for a combined integral for a Gamma-distributed demand coupled with a Rectangular, Normal, Gamma or Log Normally-distributed lead time. He has obtained (83) an explicit expression for a Normally distributed demand coupled with a Gamma-distributed lead time, but the demand assumption is not appropriate for a product range with a preponderance of slow-moving lines.

The probability distributions of demand and lead time are first examined separately. The data collection and aggregation methods are explained in full in Section 6.1. The essential results are presented in Fig. 5.2/Table 5.2 for demand, and Fig. 5.3/Table 5.3 for lead time. In both cases the observed data is tested against a Gamma hypothesis.

Prima facie, it appears that the Gamma distribution represents the demand distribution well at high k values, but is an inadequate fit at low k values. It is evident that the histogram for $k = 0.5$ shows serious irregularities, and there are appreciable irregularities in the histogram for $k = 1$. A close inspection of

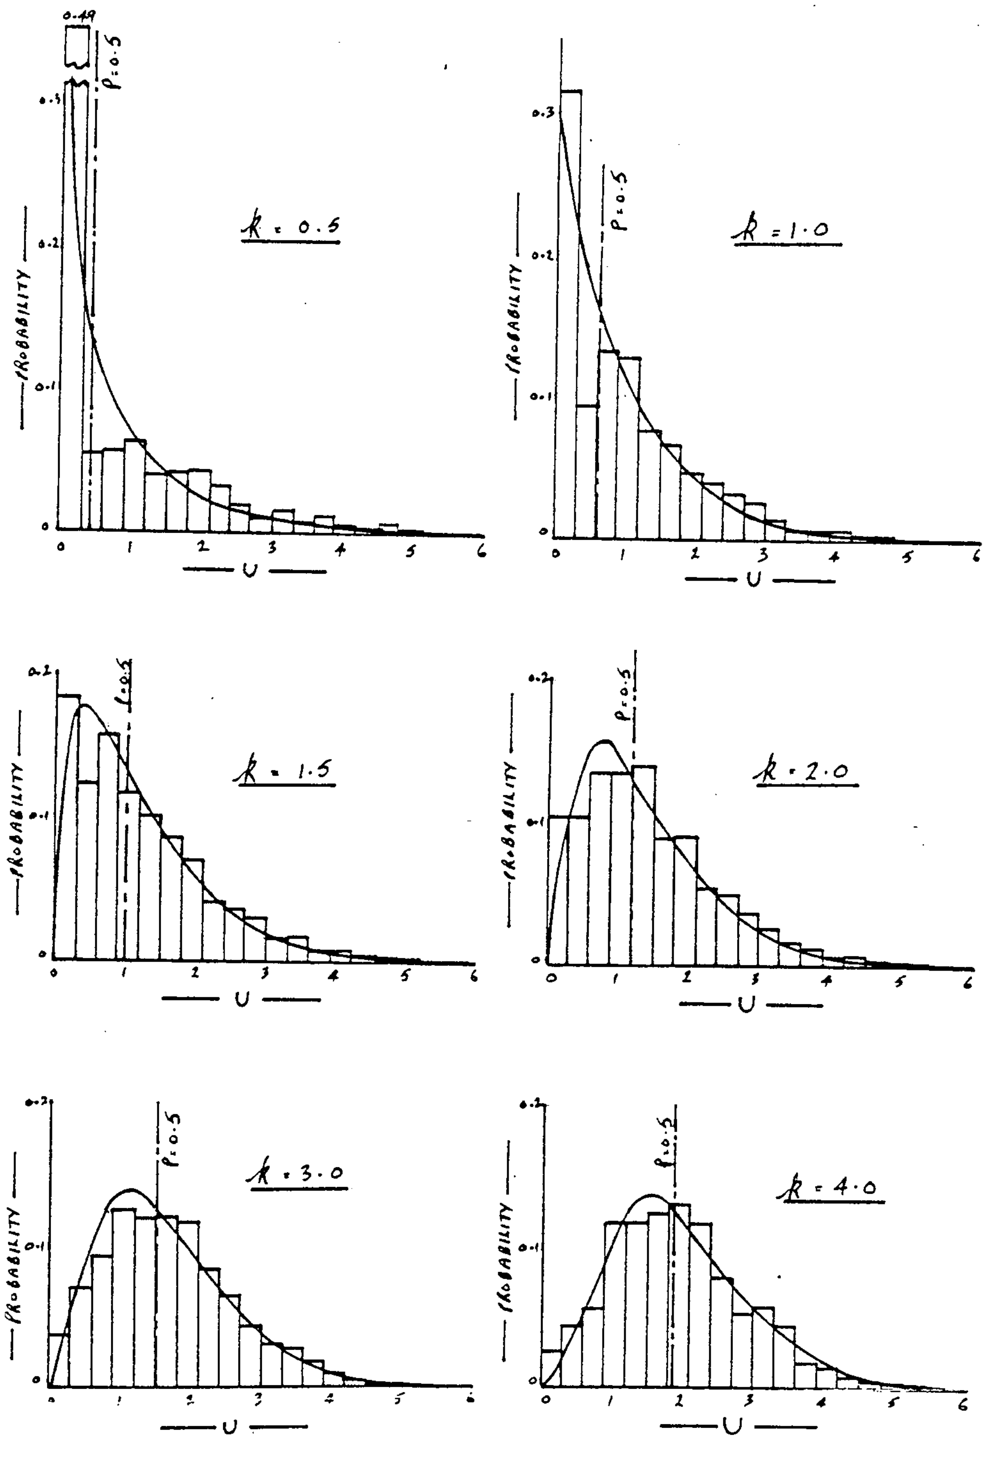

Fig. 5.2 Observed and Theoretical Gamma Distributions for Demand per Unit Time (Histogram = Observed, Graph = Theoretical)

k (mid-pt)	Obs. (O) / Exp. (E)	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17+	χ^2 (D.F.)
0.5	O	13309	1217	1294	1531	936	950	996	767	436	203	417	142	285	92	44	142	33	99	3228.8
	E	11108	3613	2237	1525	1085	794	591	444	337	259	199	153	119	92	73	57	43	164	(14)
	$(O-E)^2/E$	-	1588.9	397.5	0	20.5	30.6	277.5	235.0	29.1	12.1	238.8	0.8	231.6	0	11.5	126.8	2.3	25.8	
1.0	O	2528	765	1065	1013	591	507	376	324	239	192	117	104	50	43	36	12	9	11	148.1
	E	2069	1533	1135	841	623	461	342	253	188	139	103	77	57	42	31	23	17	42	(13)
	$(O-E)^2/E$	-	-	4.3	35.2	1.6	4.6	3.4	19.9	13.8	20.2	1.9	9.5	0.9	0	0.8	5.3	3.8	22.9	
1.5	O	640	428	542	401	357	297	246	148	128	102	55	55	29	29	18	9	6	7	38.4
	E	472	614	553	455	357	274	206	154	113	83	60	44	31	23	16	12	8	21	(12)
	$(O-E)^2/E$	-	-	-	6.4	0	1.9	7.8	0.2	2.0	4.3	0.4	2.8	0.1	1.6	0.3	0.8	0.5	9.3	
2.0	O	247	244	317	310	333	214	213	126	118	81	60	40	30	10	11	8	3	1	42.6
	E	161	333	366	337	284	227	176	133	99	72	52	37	26	19	13	9	5	16	(10)
	$(O-E)^2/E$	-	-	-	-	8.5	0.7	7.8	0.4	3.6	1.1	1.2	0.2	0.6	4.3	0.3	0.1	13.8		
3.0	O	58	104	140	190	177	179	175	124	97	69	48	43	29	15	10	4	4	3	14.2
	E	23	105	174	203	201	179	149	117	90	67	49	35	24	17	12	8	5	9	(9)
	$(O-E)^2/E$	-	-	-	-	-	0	4.5	0.4	0.5	0.1	0	1.8	0.9	0.2	0.3	2.0	3.5		
4.0	O	16	28	38	82	83	87	92	82	56	36	40	30	13	11	8	5	6	2	7.8
	E	2	22	54	81	94	94	86	72	58	44	33	24	17	12	8	5	3	7	(7)
	$(O-E)^2/E$	-	-	-	-	-	-	0.4	1.4	0.1	1.5	1.5	1.5	0.9	0.1	0		0.4		

(Note: Relevance of fit is confined to the right of the heavy line)

Table 5.2 χ^2 Tests for Demand per Unit Time

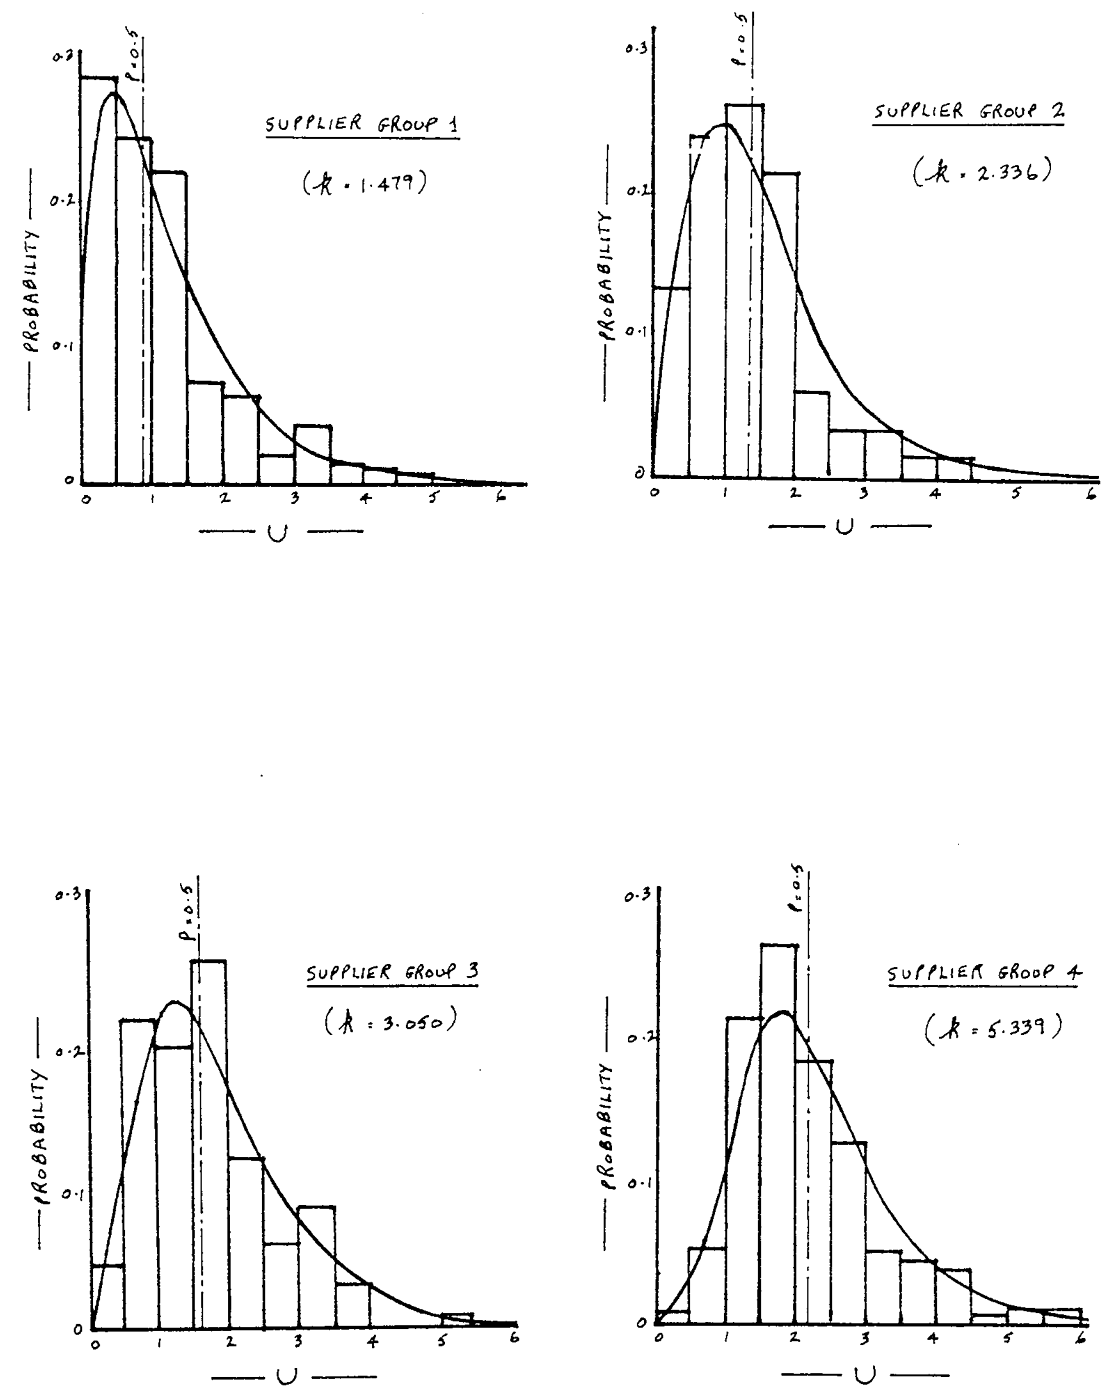

Fig. 5.3 Observed and Theoretical Gamma Distributions for Lead Time (Histogram = Observed, Graph = Theoretical)

k	Obs. (O) / Exp. (E)	Class Interval (= U/2)												χ^2
		0	1	2	3	4	5	6	7	8	9	10	11+	(DF)
1.479	O	64	54	50	16	14	5	10	4	3	2	0	0	
	E	57	58	41	26	16	10	6	3	2	1	1	1	12.41
	$(O-E)^2/E$	-	0.28	1.98	3.85	1.00	2.50	2.67			0.13			(4)
2.336	O	20	36	39	32	9	4	5	2	2	0	0	0	
	E	17	36	33	25	16	10	6	3	2	1	0	0	10.55
	$(O-E)^2/E$	-	-	1.09	1.96	3.06	3.60	0.17			0.67			(3)
3.050	O	6	29	27	35	16	8	11	4	0	0	1	0	
	E	7	26	31	26	19	12	7	4	2	1	1	1	8.99
	$(O-E)^2/E$	-	-	-	3.12	0.47	1.33	2.29			1.78			(2)
5.339	O	1	11	47	58	40	27	11	10	8	1	2	2	
	E	1	12	34	46	44	33	22	13	7	4	2	0	8.22
	$(O-E)^2/E$	-	-	-	-	0.63	1.09	5.50	0.69	0.14		0.17		(3)

(Note: Relevance of fit is confined to the right of the heavy line)

Table 5.3 χ^2 Tests for Supplier Lead Times the raw data reveals that these are due to the combined effect of:-

a) a very small number of movements during the sampling period, and

b) the tendency of a great many building products to sell in particular multiples (e.g. taps in pairs, door hinges in threes, wall tiles in packs of 9).

A computer tabulation of the most frequent customer purchasing units of slow-moving products reconciles well with the lumpiness of the histograms. A much better fit would certainly be obtained if products with small numbers of movements or large unit purchase sizes relative to the mean were removed from the sample, but there is no justification for this as a significant proportion of the range displays these selling characteristics.

χ^2 tests for the demand distribution were also carried out against Normal and Poisson hypotheses. Both of these fared significantly worse than Gamma. It is concluded that Gamma represents well the demand distribution of products with a fast or medium movement rate. There are serious discrepancies in the fit for slow-moving products, but due to the quantisation problems it is unlikely that any other continuous distribution would provide a significantly better fit. Johnston (60) has also measured the demand distribution of 784 building materials, each with a k value around 3. He does not present numerical results of fit tests, but his visual evidence is consistent with Fig. 5.2.

The results for the lead time tests are much more consistent than the demand tests in spite of the much smaller amount of data. The χ^2 tests do not refute the null hypothesis and there is a good correspondence in the all-important extreme right hand tails of the curves which heavily influence the expected lost sales calculations. χ^2 tests were repeated with Normal and Poisson hypotheses but the results were generally inferior to those obtained with Gamma.

There are very few references to lead time distributions in stock control literature even though in many cases they are at least as important as demand distributions. This is presumably because of measurement difficulties and the contentious assumption that they are drawn from a stationary distribution over perhaps several years. In order to compare the relative merits of the Gamma distribution with other probability distributions, the lead time data detailed in Table 6.2 was subjected to χ^2 and Kolmogorov-Smirnov goodness of fit tests against six probability distributions using the DISFIT (84) package developed at The University of Aston. The results are given in Table 5.4. A tick indicates acceptance at the 95% significance level (i.e. the null hypothesis cannot be rejected with 95% confidence), and a cross indicates rejection at the same level.

The χ^2 test appears from the results to be the more stringent. Also, the Kolmogorov-Smirnov test utilises a maximum absolute deviation between the theoretical and observed cumulative distributions as the test statistic, which could easily occur in the main body of the distribution whilst the tail is of paramount interest. For these reasons, greater emphasis is placed on the χ^2 results, which favour Gamma to the other distributions tested.

As the mathematics of convolving a Gamma demand distribution with a Gamma lead time distribution are intractable, an empirical convolution is now carried out. The method used is to generate and store 10,000 lead times using the Gamma generator, recall each in

Supplier Code	Probability Distribution											
	Normal		Gamma		Neg.Exp.		Uniform		Lognormal		Poisson	
	χ^2	K-S	χ^2	K-S	χ^2	K-S	χ^2	K-S	χ^2	K-S	χ^2	K-S
B	x	✓	✓	✓	x	x	x	x	✓	✓	x	✓
CN	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	–	–
CO	x	✓	✓	✓	x	x	x	x	✓	✓	–	–
D	x	✓	✓	✓	✓	✓	x	x	✓	✓	–	–
E	x	✓	✓	✓	x	x	x	x	✓	✓	x	✓
F	x	✓	x	✓	x	x	x	x	x	✓	✓	✓
G	x	✓	x	✓	✓	✓	x	x	x	✓	–	–
NT	x	✓	✓	✓	x	✓	✓	✓	x	✓	–	–
PE	x	✓	✓	✓	✓	✓	x	x	x	✓	–	–
PO	x	x	x	x	x	x	x	x	x	x	x	x
U	x	✓	✓	✓	x	x	x	x	✓	✓	x	✓
WD	✓	✓	✓	✓	x	x	x	x	✓	✓	x	✓
WP	x	✓	x	✓	x	✓	✓	x	x	✓	–	–
Totals	2	12	9	12	4	6	3	2	7	12	–	–

Notes: χ^2 = Chi-Square Test
K-S = Kolmogorov - Smirnov Test
✓ = Acceptance of fit
x = Rejection of fit
– = Test inappropriate (Mean > 20)

Table 5.4 Tests of Alternative Distributions for Supplier Lead Times turn, and generate a daily demand for each constituent day in the lead time, again using the Gamma generator. The summation of these over a lead time constitutes a demand in lead time variate. The exercise was carried out six times with the first two moments of demand and lead time selected to produce convolved k values of 0.5, 1.0, 1.5, 2.0, 3.0 and 4.0 respectively. The results, including χ^2 tests, are given in Figs. 5.4 to 5.9. The 'U' ordinates corresponding to percentage lost sales values of 0.15, 0.10, 0.05 and 0.02 are also shown, as the goodness of fit gets progressively more important towards the end of the tail. In each case the χ^2 test has been applied to the region $0.5 < P < 1.0$.

The evidence suggests that the resultant convolved distribution is an extremely close approximation to a Gamma distribution. However, it is unlikely to be an exact Gamma distribution for the following reasons:-

a) The χ^2 values are consistently higher than those obtained with the simple application of the generator.

b) There are clearly discernible discrepancies at low 'U' values in the region of the mode.

c) In at least some cases the convolved distribution appears to truncate more quickly than Gamma.

The last reason is potentially extremely important, as an inadequate incidence of very high variates could give rise to overprotection by the calculated reorder levels. This is now investigated by first working through an illustrative example for $k = 1$, target service level = 0.95, lead time = order interval. The frequencies shown in the χ^2 table in Fig. 5.5 are given in greater detail in Table 5.5.

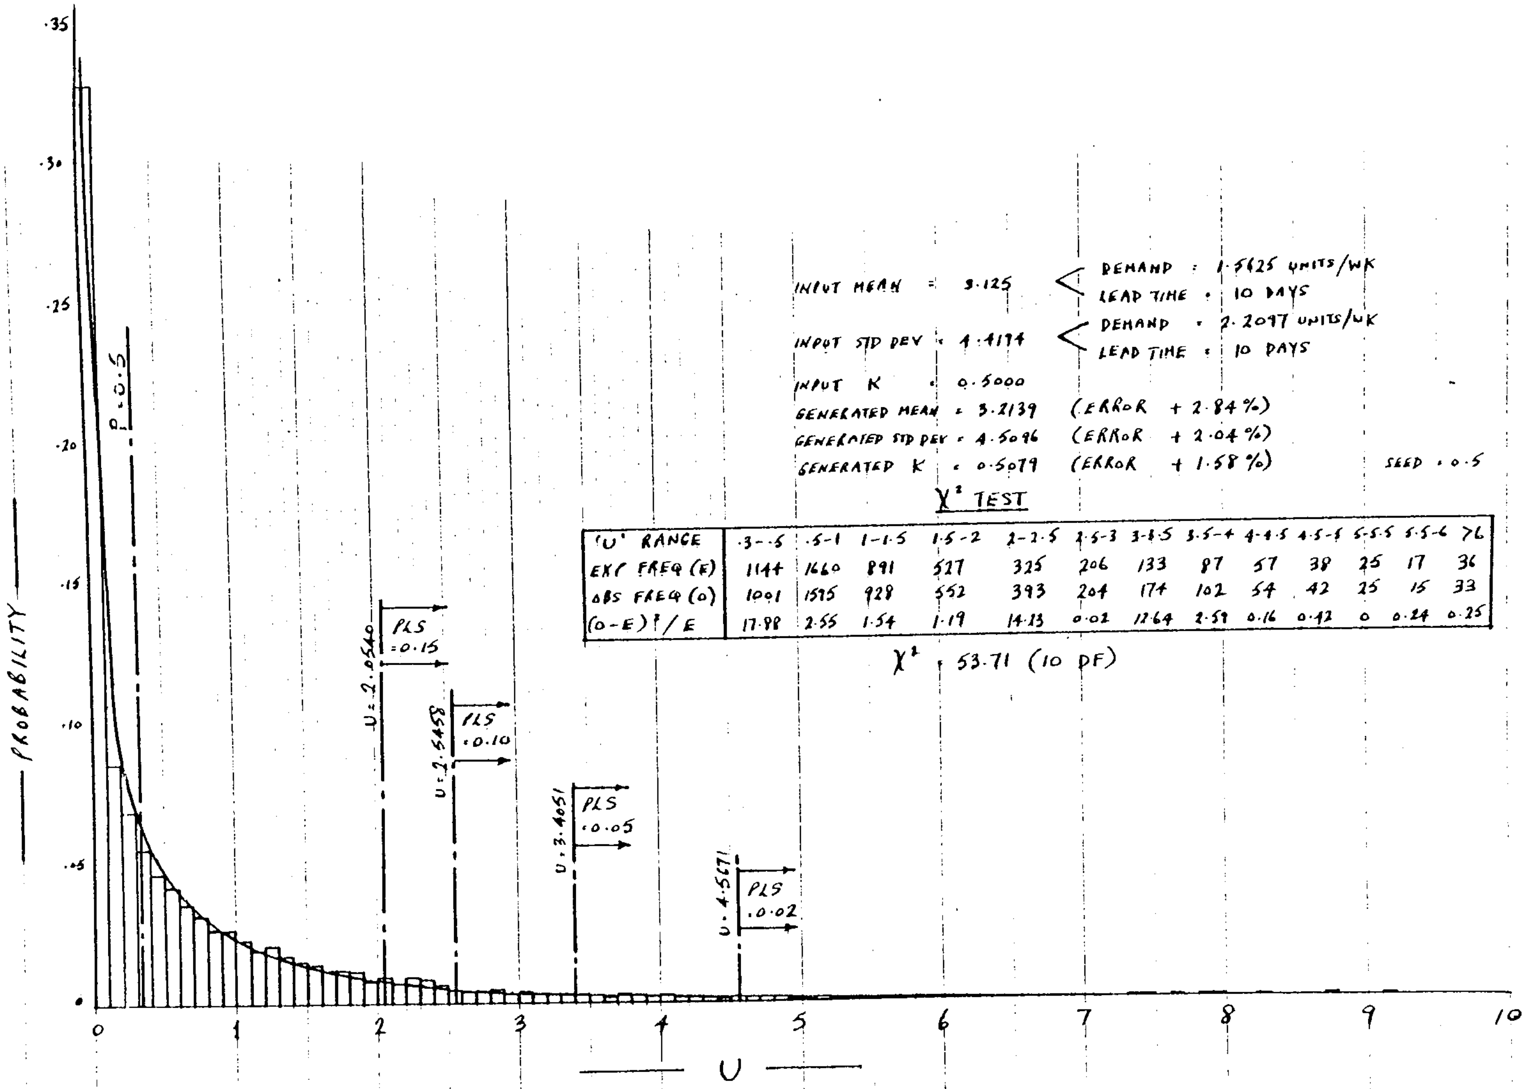

Fig. 5.4 Probability Distribution for Convolved Gamma/Gamma with k = 0.5

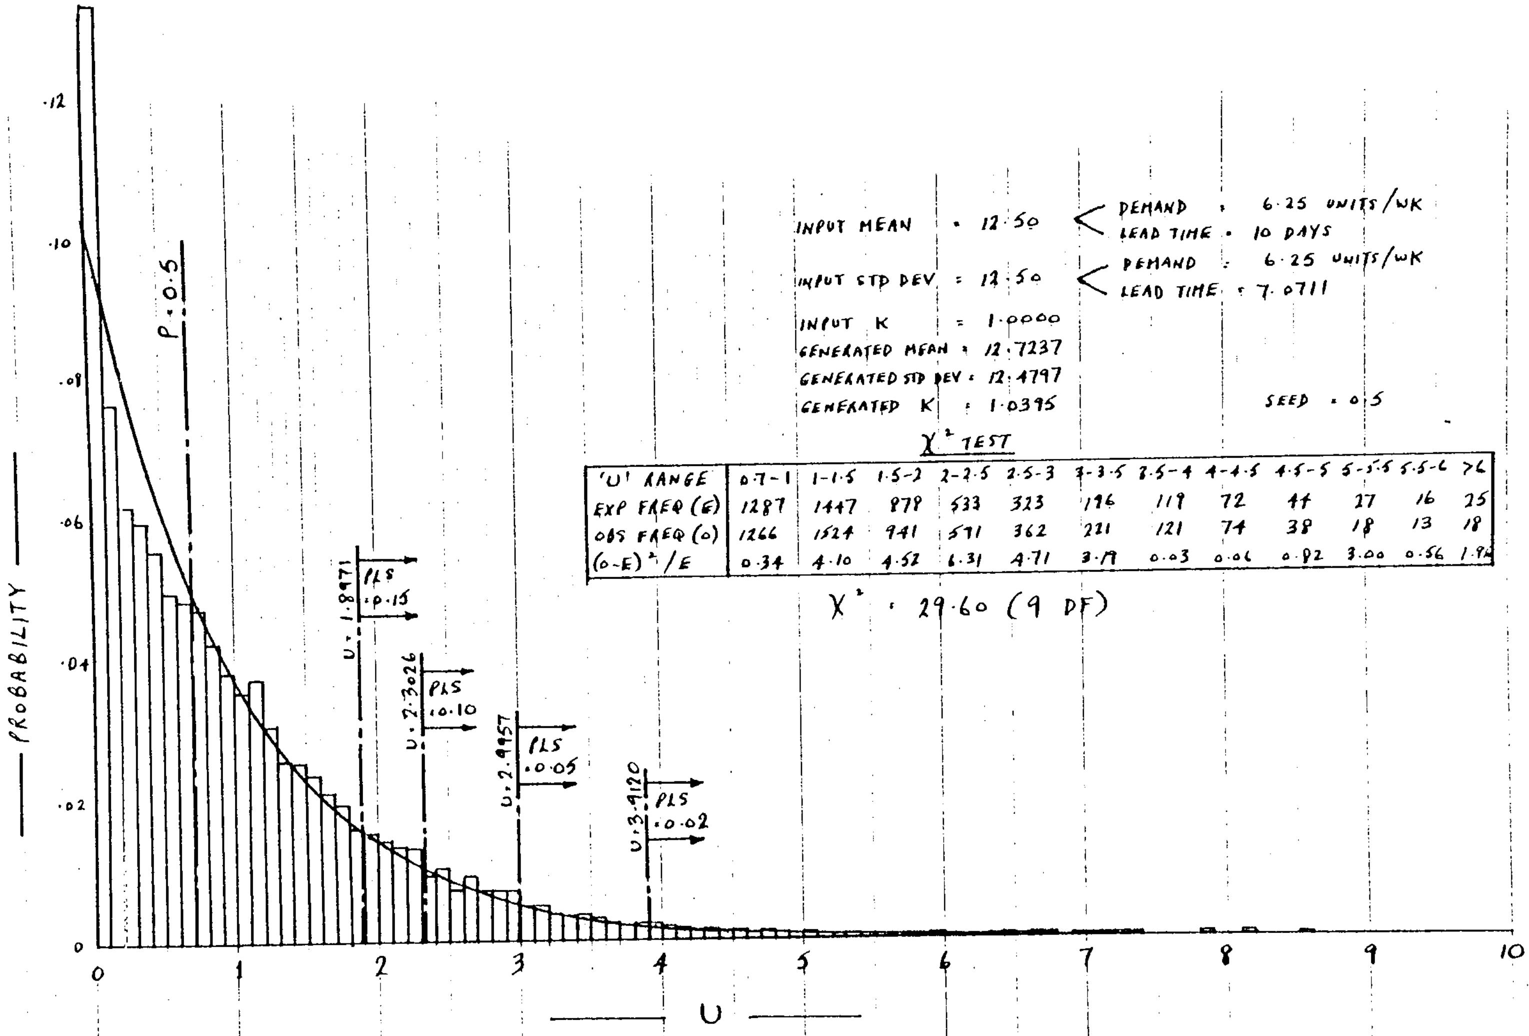

'U' RANGE	0·7-1	1-1·5	1·5-2	2-2·5	2·5-3	3-3·5	3·5-4	4-4·5	4·5-5	5-5·5	5·5-6	>6
EXP FREQ (E)	1287	1447	878	533	323	196	119	72	44	27	16	25
OBS FREQ (O)	1266	1524	941	571	362	221	121	74	38	18	13	18
$(O-E)^2/E$	0·34	4·10	4·52	6·31	4·71	3·19	0·03	0·06	0·82	3·00	0·56	1·96

Fig. 5.5 Probability Distribution for Convolved Gamma/Gamma with k = 1

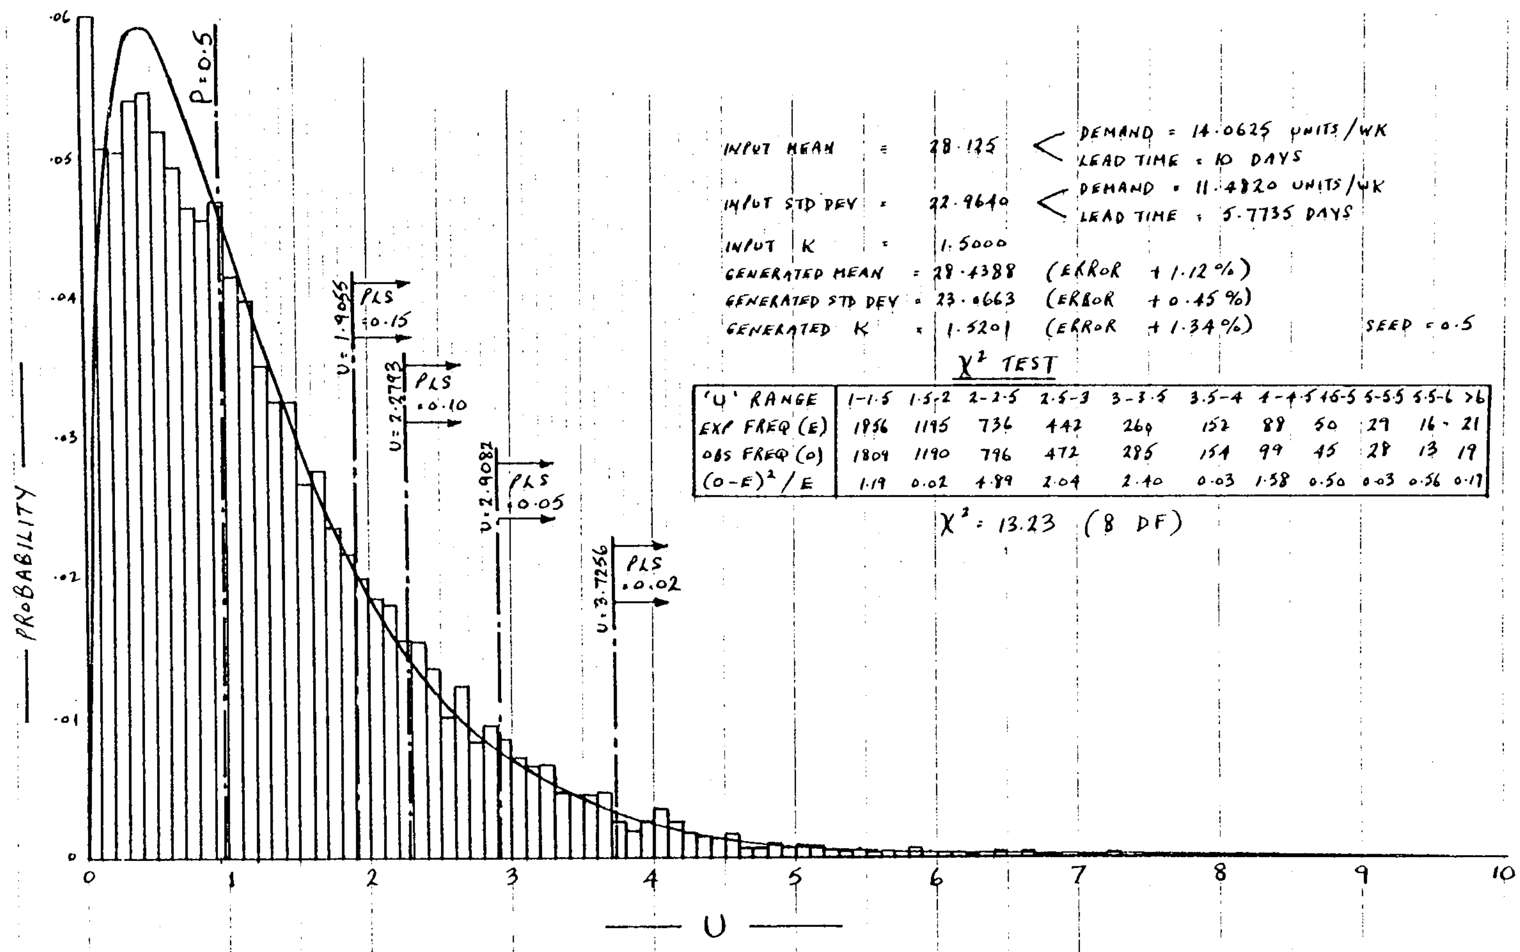

Fig. 5.6 Probability Distribution for Convolved Gamma/Gamma with k = 1.5

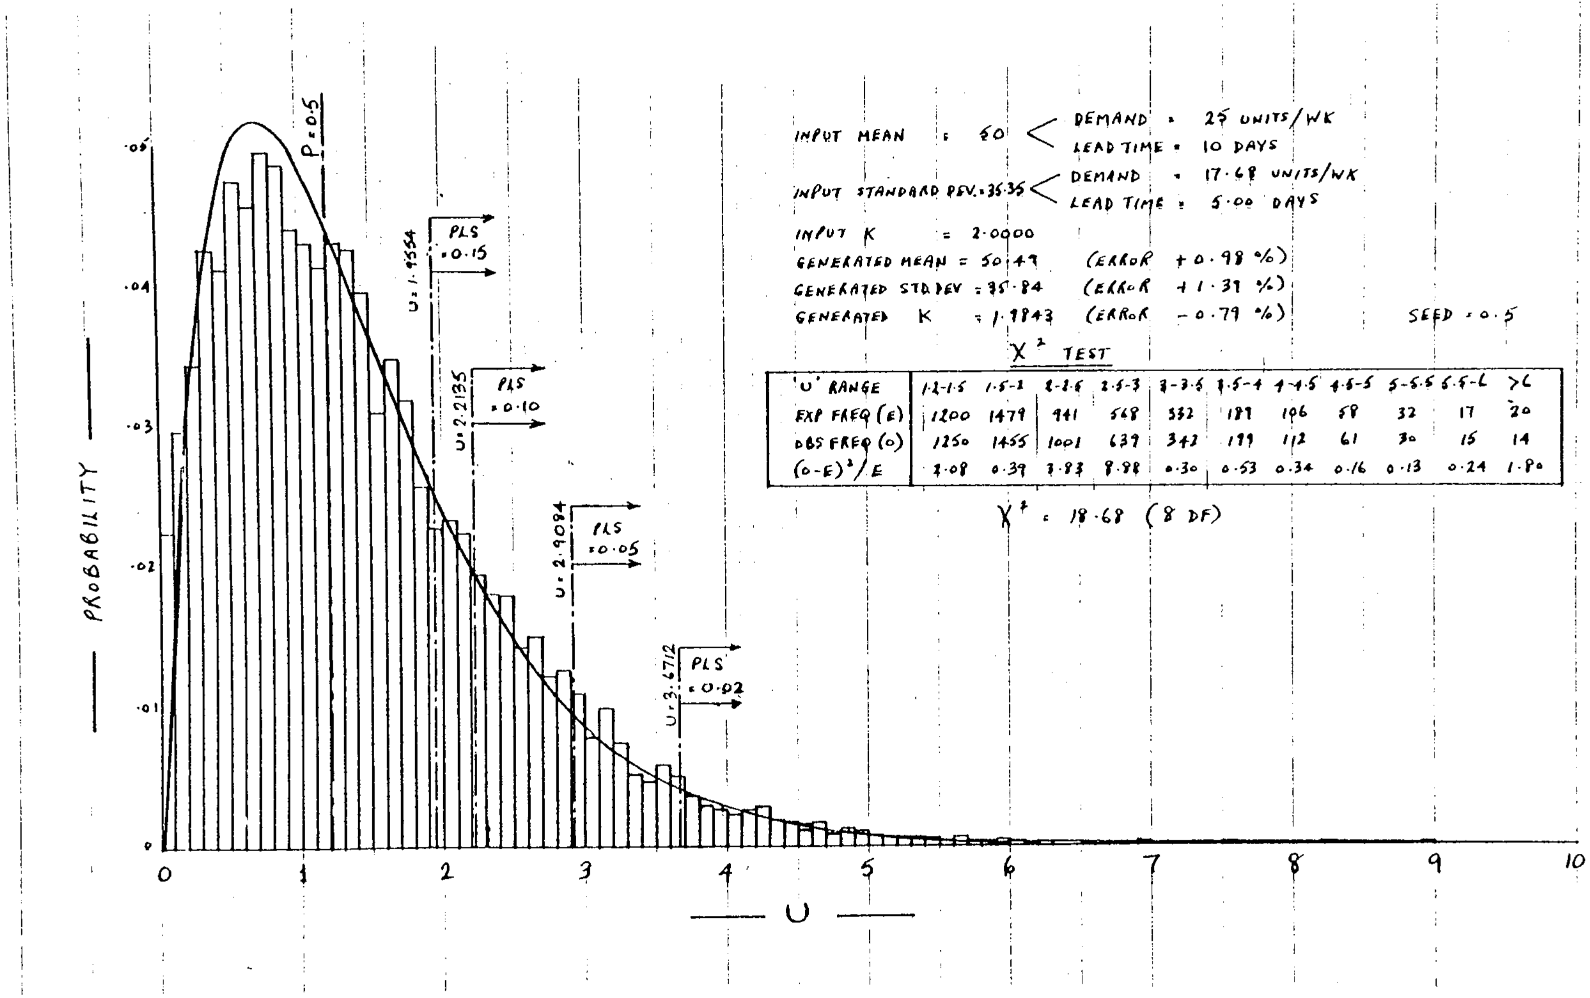

'U' RANGE	1.2-1.5	1.5-2	2-2.5	2.5-3	3-3.5	3.5-4	4-4.5	4.5-5	5-5.5	5.5-6	>6
EXP FREQ (E)	1200	1479	941	568	332	189	106	58	32	17	20
OBS FREQ (O)	1250	1455	1001	639	342	199	112	61	30	15	14
$(O-E)^2/E$	2.08	0.39	3.83	8.88	0.30	0.53	0.34	0.16	0.13	0.24	1.80

Fig. 5.7 Probability Distribution for Convolved Gamma/Gamma with k = 2

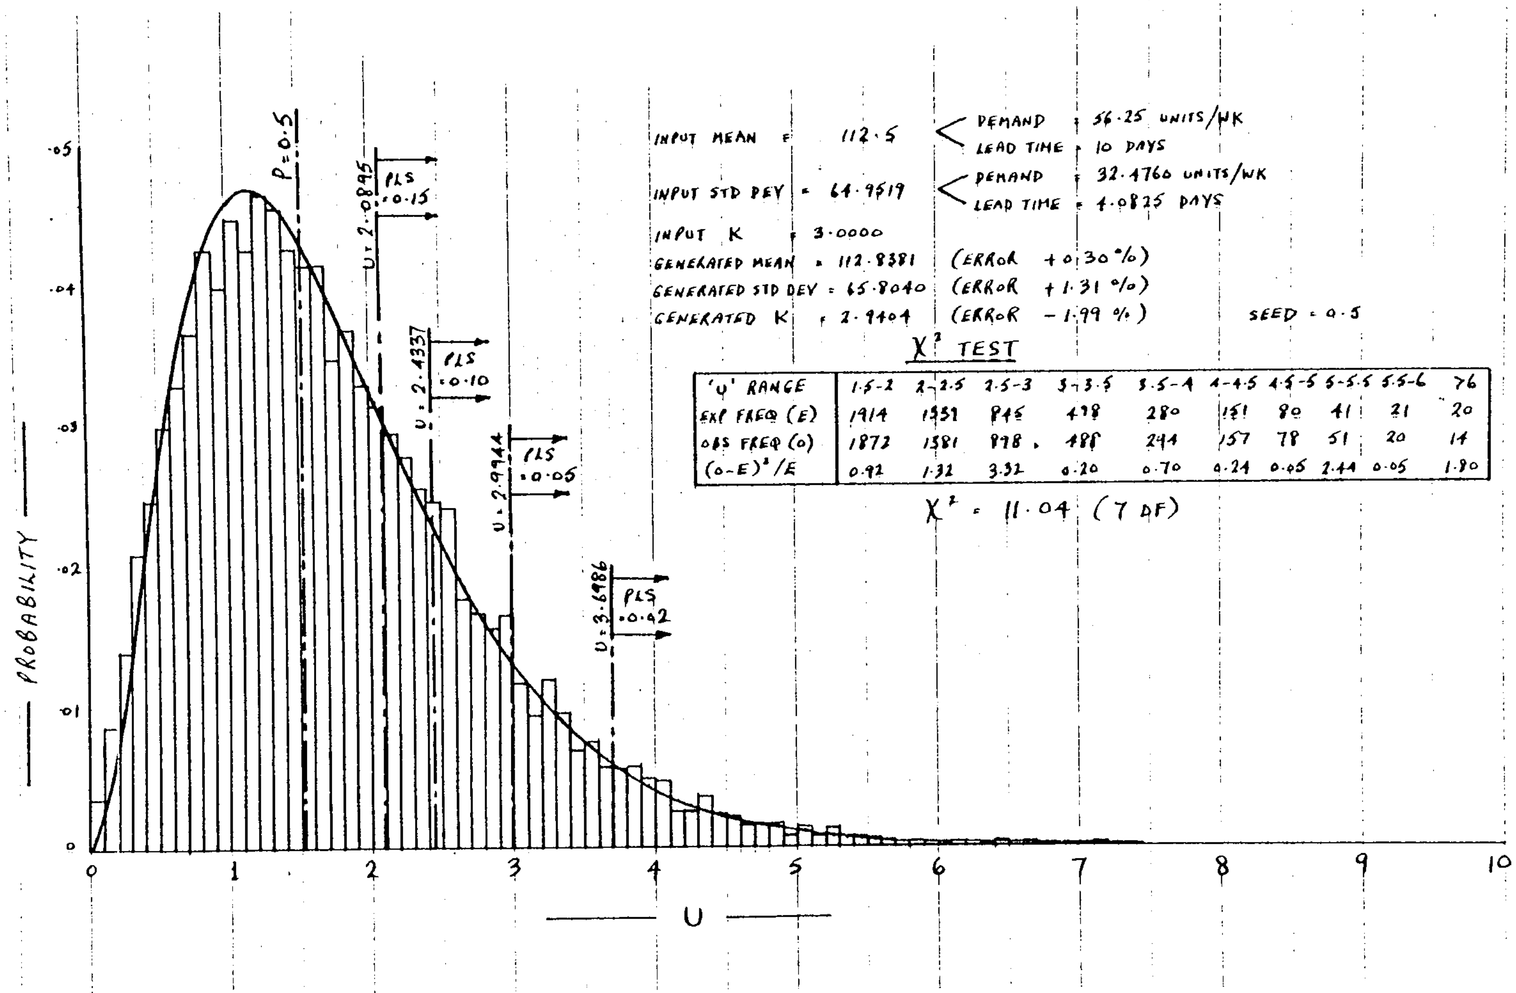

'U' RANGE	1.5-2	2-2.5	2.5-3	3-3.5	3.5-4	4-4.5	4.5-5	5-5.5	5.5-6	>6
EXP FREQ (E)	1914	1551	845	478	280	151	80	41	21	20
OBS FREQ (O)	1872	1581	898	488	244	157	78	51	20	14
$(O-E)^2/E$	0.92	1.32	3.32	0.20	0.70	0.24	0.05	2.44	0.05	1.80

Fig. 5.8 Probability Distribution for Convolved Gamma/Gamma with k = 3

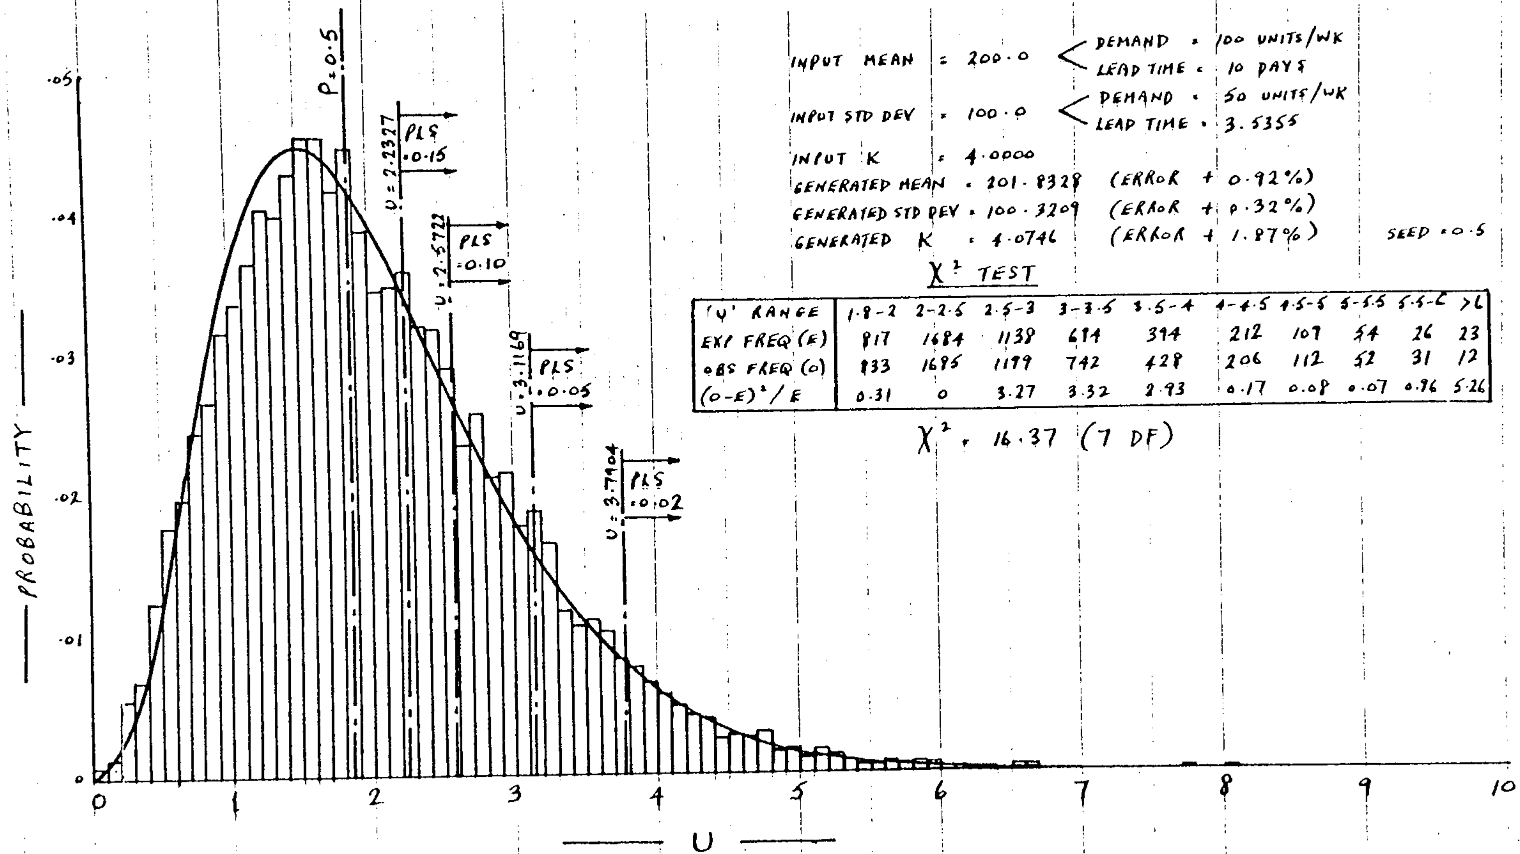

χ^2 TEST

'U' RANGE	1.8-2	2-2.5	2.5-3	3-3.5	3.5-4	4-4.5	4.5-5	5-5.5	5.5-6	>6
EXP FREQ (E)	817	1684	1138	614	394	212	107	54	26	23
OBS FREQ (O)	833	1695	1199	742	428	206	112	52	31	12
$(O-E)^2/E$	0.31	0	3.27	3.32	2.93	0.17	0.08	0.07	0.96	5.26

$\chi^2 = 16.37$ (7 DF)

Fig. 5.9 Probability Distribution for Convolved Gamma/Gamma with k = 4

		U (Fractional Part-designated by Upper Limit)									
		.1	.2	.3	.4	.5	.6	.7	.8	.9	0
U (Integer Part)	0	1315	761	617	595	551	491	483	468	419	379
	1	351	368	302	254	249	234	205	192	158	152
	2	140	132	129	88	102	69	89	67	68	69
	3	56	57	38	34	36	30	23	20	24	24
	4	19	16	13	15	11	14	5	10	6	3
	5	6	3	2	4	3	2	3	2	2	4
	6	1	0	0	1	2	1	2	2	0	1
	7	1	1	1	1	0	0	0	0	2	0
	8	0	1	0	0	0	1	0	0	0	0
	9	0	0	0	0	0	0	0	0	0	0

Table 5.5 Generated Frequencies for U-Bands of 0.1 for k = 1

For target SL of 0.95, PLS per order = 0.05

As lead time = order interval, PLS per lead time = 0.05

'U' value corresponding to PLS of 0.05 = 2.9957 (Fig. 5.5)

$R = U\ \sigma_{LD} = 2.9957 \times 12.50$

$= 37.4463$

(The fractional part is retained to avoid quantising errors)

Potential sales will be lost whenever a demand in lead time variate (x) exceeds 37.4463, and the amount of lost sales will be (x - 37.4463). Assuming each variate is at the mid-point of its U-Band, the general expression,

$$ELS = \int_{R}^{\infty} (x-R)\ f\ (x)\ dx$$

may be approximated by

$$\frac{ELS}{\sigma_{LD}} = \sum_{U_M=U}^{U_M=\infty} (U_M - U) \times \text{Prob} \quad (U_M \pm \frac{\Delta U}{2})$$

where U_M is the mid-point of any U-Band

and ΔU is the width of the U-Bands

$$\frac{ELS}{\sigma_{LD}} = \left[\frac{(2.9957+3.0)}{2} - 2.9957\right] \times \frac{69}{10000} \times \frac{(3-2.9957)}{0.1}$$

(Proportion of band-width for band containing U)

$$+ \left[\frac{(3.0 + 3.1)}{2} - 2.9957\right] \times \frac{56}{10000}$$

$$+ \left[\frac{(3.1+3.2)}{2} - 2.9957\right] \times \frac{57}{10000}$$

$$\vdots$$

$$+ \left[\frac{(8.5+8.6)}{2} - 2.9957\right] \times \frac{1}{10000}$$

$\therefore$ ELS = 0.0445 σ_{LD}

For lead time = order interval and k = 1,

demand per order interval = $\mu_{LD} = \sigma_{LD}$

$$\text{Service Level} = \frac{\text{Demand per order - ELS}}{\text{Demand per order}}$$

$$= \frac{\sigma_{LD} - 0.0445\,\sigma_{LD}}{\sigma_{LD}} = \underline{0.9555}$$

This result confirms that the truncation of the distribution observed in the χ^2 table leads to an appreciable amount of over-protection. As simulation is used extensively in the study the above calculation was repeated for all combinations of:

k = 0.5, 1.0, 1.5, 2.0, 3.0, 4.0

Target SL = 0.85, 0.90, 0.95, 0.98

Random Number Seed (for generator) = 0.5, 0.6, 0.7, 0.8, 0.9.

The full set of results is given in Table 5.6 and the average expected service levels are shown graphically in Fig. 5.10.

Assuming a Normal distribution of the sample errors, the null hypothesis that the sample results do not have the input mean (i.e. target service level) is refuted in only 3 of the 24 cases at the 90% confidence level. However, the graphical evidence convincingly demonstrates that at high service levels the reorder levels would overprovide and at low service levels they would underprovide. The overprovision errors are greatest for low k values and the underprovision errors are greatest for high k values.

In summary, the convolution of Gamma demands with Gamma lead times gives an acceptably close approximation to a Gamma distribution of demand in lead time. Unless otherwise stated all of the following simulation trials use the Gamma generator for both demand and lead time variates. Where necessary minor corrections are made with reference to Table 5.6.

k	Target SL	SL Run 1 Seed = 0.5	SL Run 2 Seed = 0.6	SL Run 3 Seed = 0.7	SL Run 4 Seed = 0.8	SL Run 5 Seed = 0.9	Average SL (Std. Dev.)
0.5	.85	.8416	.8474	.8443	.8592	.8546	.8494(.0073)
	.90	.8971	.9003	.8983	.9093	.9088	.9028(.0059)
	.95	.9524	.9528	.9513	.9570	.9618	.9551(.0043)
	.98	.9825	.9819	.9816	.9832	.9888	.9836(.0030)
1.0	.85	.8491	.8469	.8427	.8446	.8499	.8466(.0030)
	.90	.9027	.9012	.8963	.8998	.9041	.9008(.0030)
	.95	.9555	.9533	.9499	.9536	.9552	.9535(.0022)
	.98	.9842	.9818	.9810	.9837	.9830	.9827(.0013)
1.5	.85	.8445	.8460	.8506	.8425	.8416	.8450(.0036)
	.90	.8979	.8981	.9028	.8955	.8945	.8978(.0032)
	.95	.9510	.9507	.9535	.9494	.9477	.9505(.0021)
	.98	.9816	.9831	.9818	.9804	.9799	.9814(.0013)
2.0	.85	.8437	.8505	.8502	.8409	.8510	.8473(.0046)
	.90	.8972	.9018	.9019	.8931	.9026	.8993(.0041)
	.95	.9507	.9531	.9532	.9454	.9526	.9510(.0033)
	.98	.9812	.9828	.9822	.9772	.9814	.9810(.0022)
3.0	.85	.8459	.8455	.8470	.8404	.8480	.8454(.0029)
	.90	.8979	.8976	.8985	.8938	.8988	.8973(.0020)
	.95	.9493	.9505	.9511	.9472	.9497	.9496(.0015)
	.98	.9801	.9816	.9815	.9798	.9802	.9806(.0008)
4.0	.85	.8454	.8468	.8406	.8415	.8412	.8431(.0028)
	.90	.8980	.8987	.8932	.8946	.8944	.8958(.0024)
	.95	.9507	.9509	.9470	.9481	.9482	.9490(.0017)
	.98	.9813	.9826	.9794	.9797	.9801	.9806(.0013)

Table 5.6 Service Level Errors Caused by Convolving Gamma Demands with Gamma Lead Times

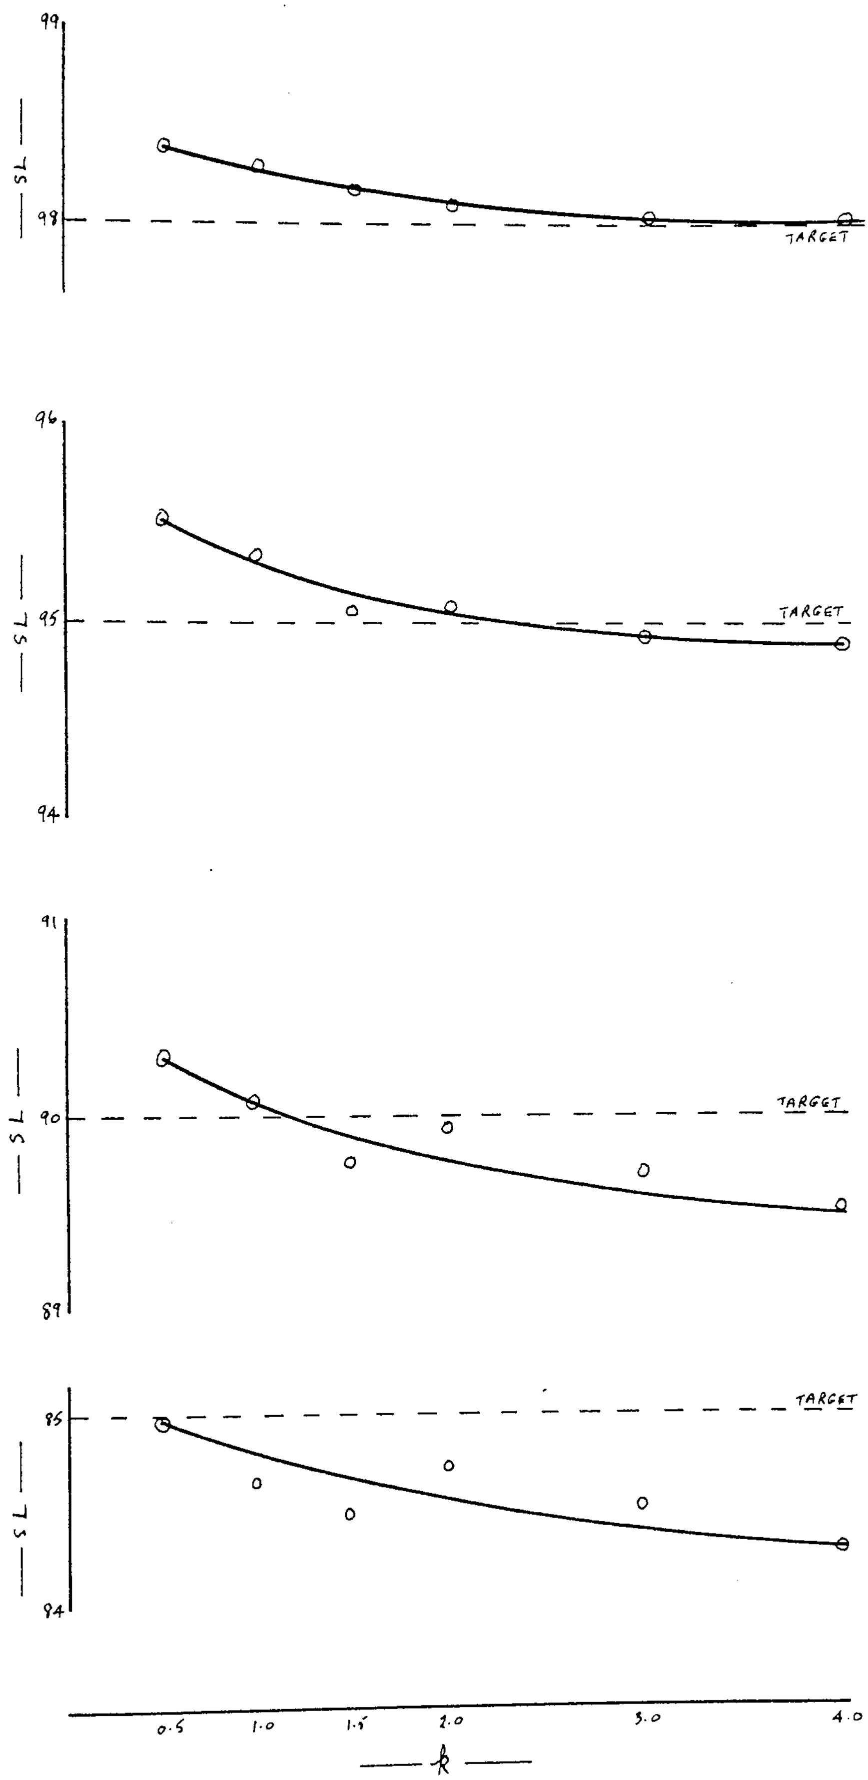

Fig. 5.10 Service Level Errors from Gamma/Gamma Convolutions

CHAPTER 6

ANALYSIS OF THEORETICAL ERRORS

6. ANALYSIS OF THEORETICAL ERRORS

This chapter attempts to quantify the effects on the main Control Indices - Service Level and Average Stock Value - of errors and inaccuracies in the theory. In some cases numerical approximations were used to facilitate computation even though exact analytical solutions were known. Here the errors introduced can be quantified quite accurately. In other cases the 'pure' theory was found to be mathematically intractable and here it is necessary to invoke the simulation methods described in Chapter 5 to assess the magnitude of the errors. Simulation methods invariably produce results of a statistical nature and, unless the causative error-producing factors are known, the results may be specific to the test data. In some instances as much as 200 years' data were processed to attempt to obtain statistically-significant evaluations.

It is important to recognise that errors in the application of the theory are deferred until Chapter 7. Also the merits and demerits of the stock control concepts around which the system is built are not compared with alternative philosophies. The theoretical imperfections are evaluated using this specific model as a base.

In most instances the error sources are investigated individually with the tacit assumption that they can be summated. The possible synergistic effects of combining all permutations of the error conditions are too complex to evaluate.

6.1 ERRORS IN THE PROBABILITY DISTRIBUTION

For reasons explained in Section 3.4, a Gamma distribution was selected to represent the pattern of occurrences of demand during the lead time. As with any theoretical distribution, it does not fit the data perfectly, nor is it equally appropriate for all products/suppliers. Hence errors will be introduced by employing the Gamma distribution rather than the observed distributions.

In order to assess the effects of individual errors on total system performance, sampling and aggregation techniques are used. Also the length of time required to obtain sufficient observations of demand during each lead time is prohibitive and the only practical approach is to combine current demand data with the most recent lead time data available.

Sales data were collected over a three-month period (77 working days) for a retail outlet adjoining a Distribution Centre. As customer demands which cannot be satisfied from the outlet's shelf stocks are obtained immediately from the Distribution Centre, 'sales' and 'expressed demand' are synonymous except when both sources are out of stock. In the event, sales were recorded for 5961 products which is approximately 60% of the range of the Distribution Centre. The data were keyed into a computer and the weekly mean, standard deviation and k values were calculated for each product. The products were grouped into six k_D-bands as k determines the profile of the distribution. (Note - subscript D denotes weekly demand).

The selected bands were:

k_D = 0.25 - 0.75 (mid-point 0.5)

k_D = 0.75 - 1.25 (mid-point 1.0)

k_D = 1.25 - 1.75 (mid-point 1.5)

$k_D = 1.75 - 2.5$ (mid-point nominally 2.0)

$k_D = 2.5 - 3.5$ (mid-point 3.0)

$k_D = 3.5 - 4.5$ (mid-point 4.0)

Products with a k_D value < 0.25 or > 4.5 were ignored. The daily demands for all of the remaining products were normalised according to the calculated demand rates of fictitious products which obey the Variance Law for the respective mid-point k values e.g. for mid-point $k_D = 0.5$

$$V_D = 2.5\ \mu_D^{1.5}$$

$$k_D = 0.5 = \frac{\mu_D^2}{V_D} = \frac{\mu_D^2}{2.5\mu_D^{1.5}} = 0.4\sqrt{\mu_D}$$

from which $\mu_D = 1.5625$

If μ_D for a particular product in the 0.25-0.75 band was calculated as, say, 4.36, then each daily demand would be divided by 4.36/1.5625 = 2.7904. This adjustment brings all of the products within a k_D band on to the same scale. The frequencies of occurrence of the modified daily demands were then aggregated for each k_D band and converted to probabilities (Table 6.1).

The collection of lead time data has inherent difficulties. With a monthly order interval, a period in excess of four years would be required to collect 50 observations. Even if order placement and receipt dates are available over the previous four years, the statistical validity of assuming that the data are drawn from the same distribution is questionable. The approach taken was to select a sample of typical suppliers upon whom orders are placed at fairly frequent intervals, and to derive each lead time from recorded order and receipt dates over the previous two years. The sample of suppliers covered all of the major Product Groups, and some suppliers who are known to manufacture to order were included.

k_D (mid-pt)	Normalised Daily Demand (Units)																		
	0	1	2	3	4	5	6	7	8	9	10	11 - 15	16 - 20	21 - 30	31 - 40	41 - 50	51 - 75	76 - 100	101 - 150	151 - 200
0.5	.9210	.0229	.0229	.0103	.0088	.0078	.0029	.0012	.0015	.0004	.0003	-	-	-	-	-	-	-	-	-
1.0	.8089	.0254	.0294	.0220	.0243	.0235	.0198	.0105	.0122	.0043	.0054	.0109	.0031	.0004	-	-	-	-	-	-
1.5	.7120	.0157	.0251	.0216	.0277	.0289	.0309	.0230	.0221	.0159	.0167	.0336	.0164	.0092	.0014	-	-	-	-	-
2.0	.6120	.0093	.0238	.0261	.0234	.0268	.0289	.0302	.0270	.0234	.0233	.0756	.0312	.0280	.0086	.0021	.0003	-	-	-
3.0	.4921	.0034	.0083	.0106	.0153	.0259	.0156	.0208	.0175	.0176	.0220	.1176	.0822	.0780	.0392	.0185	.0131	.0018	.0006	-
4.0	.3967	.0014	.0014	.0045	.0052	.0073	.0144	.0130	.0099	.0177	.0068	.1110	.0779	.1570	.0579	.0477	.0529	.0130	.0040	.0002

Table 6.1 Probabilities of Occurrence of Normalised Daily Demands for k-bands

The lead times are given in full in chronological sequence in Table 6.2.

In most cases the number of observations is inadequate for statistical analysis, so data aggregation was again carried out. In this case, four groups of suppliers were formed, the members of each group having broadly similar k_L values. Again the observations were normalised to a common mean value, as follows:

	Group	Normalised Mean	No. of Observations
1.	PE, D, PO	40 days	222
2.	G, WP, F	30 days	149
3.	CN, CO, NT	50 days	137
4.	U, B, E, WD	15 days	218

A correlation test was carried out for a sample of 100 products between k_D and k_L. This produced a Coefficient of Correlation of -0.1067 which suggests that they are not correlated. This is to be expected as products with a steady selling pattern and those which sell spasmodically are usually purchased on the same orders and delivered together. All of the sample suppliers provide a range of fast-moving and slow-moving products, therefore the selling characteristics tend to be independent of the lead times at a given point in time. Any k_D may therefore be convolved with any supplier group with equal validity. The following pairings were convolved:

i) k_D = 0.5/Supplier group 1

ii) k_D = 1.0/Supplier group 1

iii) k_D = 1.5/Supplier group 2

iv) k_D = 2.0/Supplier group 2

v) k_D = 3.0/Supplier group 3

vi) k_D = 4.0/Supplier group 4

Supplier	μ_L (Days)	σ_L (Days)	k_L	Recorded Lead Times
B	15.62	7.66	4.158	10,15,10,14,20,13,12,14,13,1,13,13,8,15,10,14,29,18,17,16,16,9,9,23,12,15,32,27,39,27,20, 10,24,9,13,8,10
CN	43.19	25.69	2.826	8,25,5,10,42,33,58,17,17,77,51,21,14,99,99,30,51,63,20,31,56,42,52,39,63,32,79,79,35,49, 66,19
CO	39.70	22.60	3.087	64,73,80,144,36,118,30,33,37,69,20,25,26,27,44,61,38,44,45,25,31,32,35,43,36,38,46,38,45, 68,45,50,56,48,49,57,41,42,51,21,22,45,55,77,13,16,31,32,18,39,44,56,29,35,22,24,34,16,18, 21,24,21,24,45,31,1,19,36,22,17,21
D	45.68	38.35	1.419	30,30,53,75,3,25,49,48,89,55,89,8,12,53,49,53,30,10,71,124,16,77,31,43,77,149,161,7,13,47, 11,45,10,22,8,24,24,15
E	14.78	6.09	5.890	20,25,22,16,8,17,13,33,9,10,11,10,22,27,24,36,26,15,14,14,26,15,13,12,16,10,13,15,24,18,8, 16,8,13,11,11,8,16,13,14,15,8,11,14,9,12,14,19,16,13,13,12,8,27,7,8,12,18,13,23,10,14,12, 13,18,11,13,6,13,9,10,12,12,17
F	14.56	9.45	2.374	11,15,16,11,15,15,8,12,7,12,17,12,6,17,12,16,29,25,64,34,14,11,10,18,15,10,4,10,8,6,18,14, 10,11,21,17,13,11,15,9,13,12,10,10,31,40,6,6,10,15,16,18,8,6,15,10,8,15,11
G	33.07	21.99	2.263	10,14,10,23,31,47,59,38,57,71,37,40,41,33,20,7,13,14,19,22,13,37,44,29,15,24,43,76,16,89, 25,45,25,14,77,11,16,59,67,15,10
NT	65.97	36.96	3.186	49,88,72,62,54,53,26,38,13,37,79,119,69,67,28,70,80,104,119,45,29,101,117,59,114,33,13,22, 22,122,110,59,140,30
PE	38.86	37.40	1.079	6,15,46,54,3,13,33,36,80,92,7,22,105,117,6,16,56,32,5,70,73,31,24,32,7,14,42,29,78,6,13, 132,17,29,111,18,12,149,126,3,6,108,13,85,20,20,41,3,6,61,34,11,74,72,10,11,43,9,5,27,7,6, 12,33,49
PO	10.87	8.29	1.718	11,30,19,6,9,6,11,9,4,18,6,6,3,6,5,5,8,12,16,12,3,37,4,9,9,6,3,6,5,4,14,7,6,11,16,27,41, 10,12,22,3,4,5,10,12,6,8,15,21,8,11,27,9,10,3,9,3,1,4,10,7,8,8,8,3,11,10,14,2,9,5,3,0,3,7, 2,8,15,8,28,12,11,5,11,4,12,4,27,12,12,5,6,12,27,3,39,10,19,27,18,10,27,19,5,1,7,4,14,11, 11,34,16,18,8,13,10,10,3,5
U	12.22	6.22	3.864	7,7,9,6,9,11,7,6,14,11,7,8,5,11,13,8,9,6,31, 16,15,19,35,20,19,15,20,16,15,9,9,16,8,17, 10,18,17,24,15,15,8,16,9,5,20,10,13,10,8,8,6,10,5,7,8,6,18,9
WD	14.67	5.50	7.114	6,8,29,12,23,19,11,8,14,10,12,12,21,16,18,32,13,14,17,18,21,12,16,12,27,16,11,10,9,11,10, 22,13,21,13,13,13,8,12,13,12,13,7,15,13,15,19,14,15
WP	33.96	22.15	2.351	10,19,16,7,12,11,51,77,41,17,42,67,23,36,58,74,11,8,29,42,69,4,25,12,23,40,54,31,67,38,67, 65,70,12,44,24,25,51,71,14,27,46,24,12,19,13,9,13

Table 6.2 Sample of Recorded Lead Times

The convolution process consisted of:

a) setting up Table 6.1 as an array in the computer,

b) reading each normalised lead time in turn,

c) generating a random number between 0 and 1 for each day in the lead time,

d) extracting the normalised daily demand corresponding to the random number e.g. for k = 0.5 any random number < 0.9210 selects zero; a random number between 0.9210 and (0.9210+0.0229) selects 1; etc. Where a probability encompasses a range of daily demands, the most common issue size is taken e.g. '12' for the range 11-15,

e) accumulating the daily demands over the lead time to produce a demand in lead time variate.

For each convolution, the generated mean, standard deviation, k_{LD} and frequency distribution are presented in Table 6.3. The theoretical frequencies for Gamma distributions with the same respective k_{LD} values are also given, together with the calculations and results of χ^2 'goodness of fit' tests. The χ^2 values do not refute the null hypothesis at a 90% confidence level in any of the six cases, and they are sufficiently small to suggest that a Gamma hypothesis is appropriate. Certainly the fit is superior to that obtained by Beckman and Muth (85) for the demand for machine parts using a Negative Exponential (a subclass of Gamma) which they considered a good simple approximation. The same data are shown graphically in Fig. 6.1 but with the frequencies converted to probabilities. In each case the fit is relevant only to the tail of the curve beyond P = 0.75. Table 6.4 shows that with a target service level of 0.95, 99.47% of the annual costed demand (upon which all group service levels are based) is accounted for by

k_{LD} / μ_{LD} / σ_{LD}	Obs.(O) / Exp.(E)	Class Interval (= U/3) 0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18+	χ^2 (D.F.)
0.924	O	79	31	20	31	21	10	5	9	5	5	1	1	1	1	1	1	0	0	0	
9.018	E	69	45	31	22	16	11	8	6	4	3	2	2	1	1	1	0	0	0	0	4.73
9.384	$(O-E)^2/E$	-	-	-	-	1.56	0.09	1.13	1.50	0.25	0.20		0								(4)
1.198	O	56	40	24	31	19	14	13	7	5	5	2	3	1	0	1	1	0	0	0	
39.104	E	50	45	35	26	18	14	10	7	5	3	3	2	1	1	0	0	0	0	0	1.60
35.733	$(O-E)^2/E$	-	-	-	-	0.06	0	0.90	0	0.50		0.14									(3)
1.497	O	17	31	28	14	26	9	7	7	5	2	1	0	0	0	0	0	0	1	1	
67.477	E	23	29	25	20	14	11	8	6	4	3	2	1	1	1	0	0	1	0	0	2.16
55.144	$(O-E)^2/E$	-	-	-	-	-	0.36	0.13	0.17	0		1.50									(2)
1.743	O	15	27	24	25	19	11	7	6	7	4	1	2	0	0	0	0	0	0	1	
123.676	E	17	27	25	21	17	12	9	6	5	3	2	2	1	1	0	0	1	0	0	2.12
93.667	$(O-E)^2/E$	-	-	-	-	-	0.08	0.44	0	0.80	0		0.80								(3)
3.069	O	4	8	21	16	25	16	15	8	7	8	2	4	1	0	0	2	0	0	0	
404.277	E	3	11	19	21	20	17	13	10	8	5	3	3	1	1	1	0	1	0	0	1.34
230.779	$(O-E)^2/E$	-	-	-	-	-	-	0.31	0.40	0.13	0.50		0								(2)
3.413	O	0	14	32	32	30	34	18	15	13	8	8	7	2	1	3	0	0	1	0	
257.404	E	3	14	27	32	33	28	23	18	13	9	6	4	3	2	1	1	0	1	0	1.72
139.338	$(O-E)^2/E$	-	-	-	-	-	-	-	0.50	0	0.11	0.67	0.44		0						(3)

(Note: Relevance of fit is confined to the right of the heavy line)

Table 6.3 χ^2 Tests for Demand in Lead Time Data

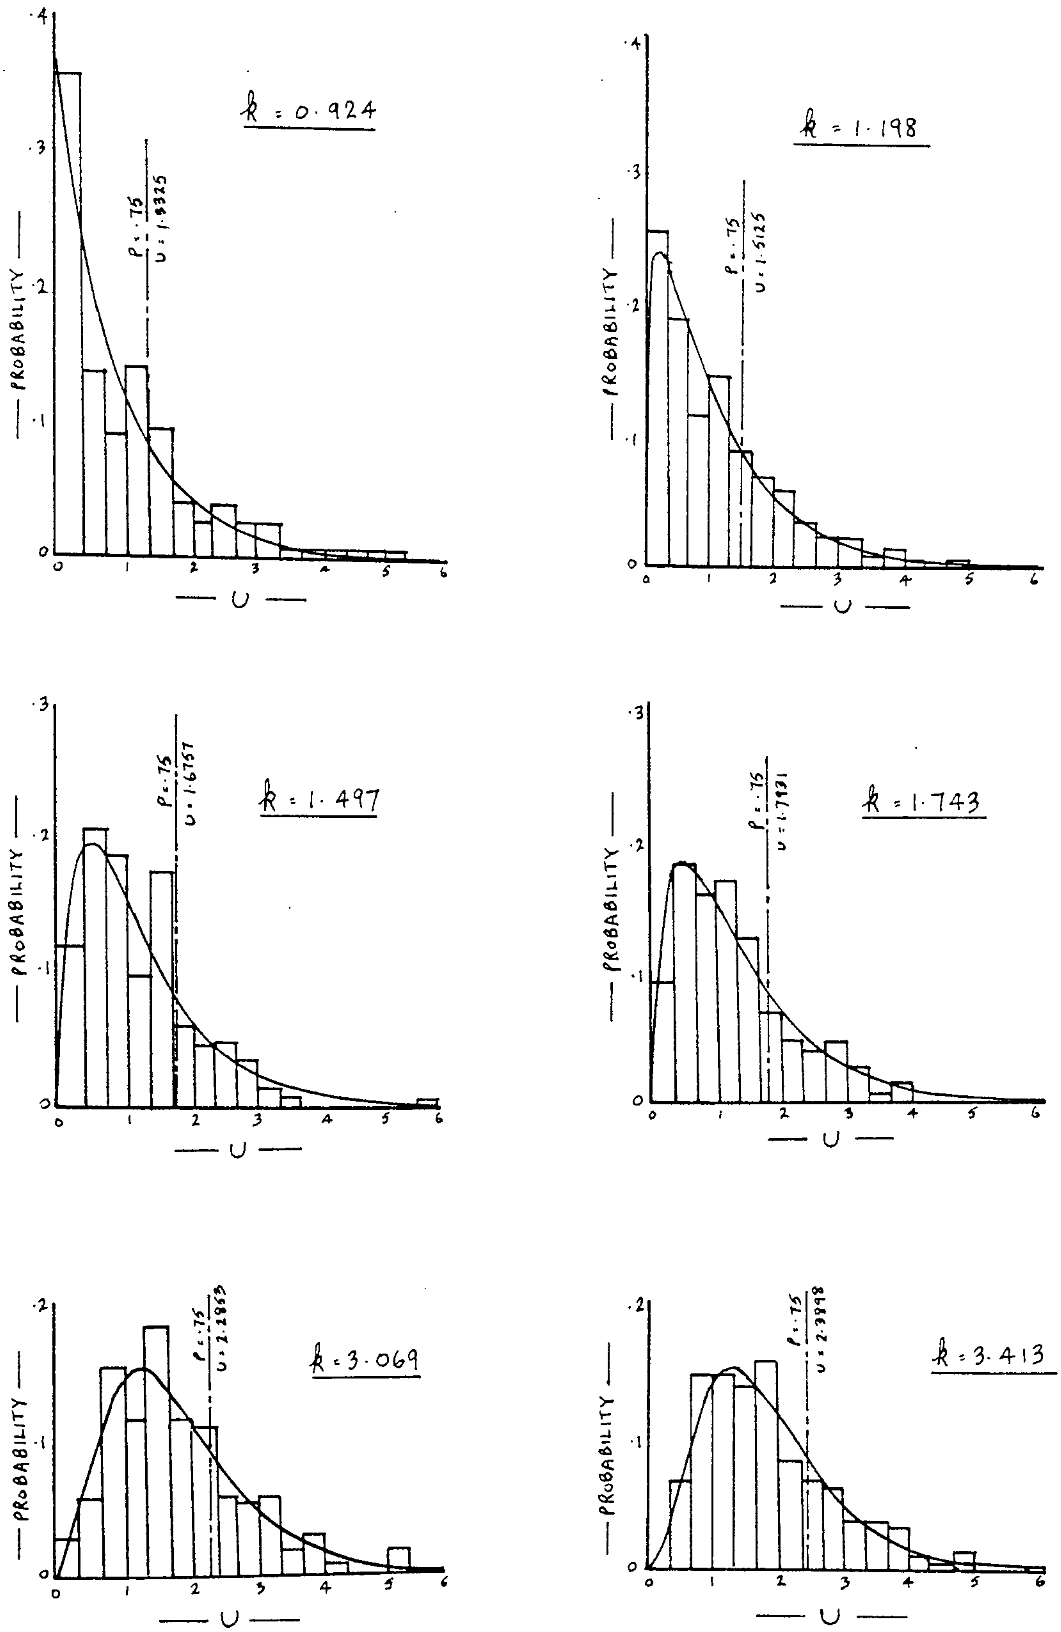

Fig. 6.1 Observed and Theoretical Gamma Distributions for Demand in Lead Time (Histogram = Observed, Graph = Theoretical)

products with P > 0.75. The same cut-off is applied in the χ^2 tests.

Protection Level (P) Range	Annual Costed Demand		
	(£)	(%)	(% cumulative)
< 0.70	13341	0.10	0.10
0.70 - 0.75	57367	0.43	0.53
0.75 - 0.80	276160	2.07	2.60
0.80 - 0.85	170766	1.28	3.88
0.85 - 0.90	3912934	29.33	33.21
0.90 - 0.95	2209280	16.56	49.77
0.95 - 0.99	6617166	49.60	99.37
> 0.99	84049	0.63	100.00
	13341063	100.00	

Table 6.4 Analysis of Protection Level Ranges

The visual evidence suggests a reasonably good fit in all cases, assuming any lumpiness is caused by the limited sample sizes and the customer purchasing multiples discussed in Sect. 5.3. Neither the statistical tests nor the visual inspection, however, indicate the expected effects on system performance. This is best achieved by evaluating the expected lost sales in the generated data for demand in lead time. An example is worked through for convolution (i) assuming the lead time is one half of the order interval and the target service level is 0.95.

$$PLS = 0.05 \times \frac{I}{\mu_L}$$

$$= 0.05 \times 2 = 0.10$$

From an iterative application of equations (18) and (19) (Sect. 3.4), the value of U corresponding to $k = 0.924$, PLS = 0.10 is 2.3170.

$$R = U\sigma_{LD}$$

$$= 2.3170 \times 9.384 = 21.743$$

Assuming a static reorder level, lost sales are experienced whenever the demand in lead time exceeds it. This occurs in the generated data in the following instances:

Demand in Lead Time	Lost Sales
30	8.257
27	5.257
24	2.257
43	21.257
26	4.257
35	13.257
26	4.257
31	9.257
23	1.257
29	7.257
32	10.257
23	1.257
50	28.257
22	0.257
23	1.257
23	1.257
22	0.257
30	8.257
25	3.257
45	23.257
38	16.257
29	7.257
28	6.257
25	3.257
26	4.257
Total	191.425

Number of observations = 222

Total demand during generated lead times $= 222\,\mu_{LD}$

$= 2001.996$

Total demand during elapsed period $= 2001.996 \times \frac{I}{\mu_L}$

$= 2001.996 \times 2 = 4003.992$

$$\text{Service Level} = \frac{\text{Demand} - \text{Lost Sales}}{\text{Demand}}$$

$$= \frac{4003.992 - 191.425}{4003.992}$$

$$= \underline{\underline{0.9522}}$$

and % error in ELS = $(.0478 - .05)/.05 = \underline{\underline{-4.4\%}}$

This procedure is repeated for all convolutions with target service levels of 0.90, 0.95, 0.98 and Order Interval/Lead Time ratios of 2.0, 1.0 and 0.5. The results are shown in Table 6.5. As convolutions (i) and (ii) and convolutions (iii) and (iv) used the same lead time data, the average results are tabulated for each pair, as it would be incorrect to include them separately for aggregation purposes.

It can be seen that at low k values the errors in the ELS are predominantly negative. The very large positive errors at the highest service level were caused by the incidence of a single variate of 7.6 σ_{LD}. The preponderance of negative ELS errors indicates that the system tends to overprovide at low k values. At higher k values the polarity of the errors is about even, taking into account their magnitude. The visual evidence in Fig. 6.1 supports these evaluations.

Table 6.5 contains a weighting factor for each cell, which represents the proportion of the total annual costed demand which is accounted for by products within bandwidths of the cell

k_{LD}	OI/LT Ratio →	2.0			1.0			0.5		
	Target SL →	.90	.95	.98	.90	.95	.98	.90	.95	.98
	Target PLS →	.20	.10	.04	.10	.05	.02	.05	.025	.01
0.924/1.198	U →	1.6115/ 1.6137	2.3170/ 2.2818	3.2537/ 3.1558	2.3170/ 2.2818	3.0252/ 2.9437	3.9642/ 3.8125	3.0252/ 2.9437	3.7353/ 3.6014	4.6760/ 4.4662
	Expected SL →	.9012	.9533	.9849	.9065	.9602	.9871	.9204	.9647	.9914
	% Error ELS →	-1.2	-6.6	-24.5	-6.5	-20.4	-35.5	-20.4	-29.4	-57.0
	Weighting Fact. →	.0011	.0342	.0232	.0426	.1079	.0081	.0521	.0592	.0005
1.497/1.743	U →	1.6364/ 1.6631	2.2792/ 2.2919	3.1088/ 3.0953	2.2792/ 2.2919	2.9083/ 2.9019	3.7262/ 3.6893	2.9083/ 2.9019	3.5282/ 3.4991	4.3376/ 4.2748
	Expected SL →	.9108	.9554	.9806	.9109	.9550	.9724	.9091	.9394	.9547
	% Error ELS →	-10.8	-10.8	-3.0	-10.9	-10.0	+38.0	-9.1	+21.2	+126.5
	Weighting Fact. →	.0001	.0125	.0301	.0005	.0699	.0505	.0410	.1364	.0285
3.069	U →	1.8452	2.4432	3.1763	2.4432	3.0021	3.7036	3.0022	3.5359	4.2146
	Expected SL →	.8963	.9467	.9792	.8933	.9473	.9815	.8946	.9501	.9824
	% Error ELS →	+3.7	+6.6	+4.0	+6.7	+5.4	-7.5	+5.4	-0.2	-12.0
	Weighting Fact. →	0	.0027	.0101	0	.0112	.0441	.0021	.0521	.0111
3.413	U →	1.8948	2.4902	3.2144	2.4902	3.0428	3.7323	3.0428	3.5678	4.2324
	Expected SL →	.9047	.9480	.9795	.8960	.9478	.9820	.8957	.9525	.9826
	% Error ELS →	-4.7	+4.0	+2.5	+4.0	+4.4	-10.0	+4.3	-5.0	-13.0
	Weighting Fact. →	.0015	.0007	.0048	.0002	.0041	.0652	.0008	.0297	.0608

Table 6.5 Expected Service Levels for Varying k Values, Target Service Levels and Order Interval/Lead Time Ratios

identifying parameters. The weighting factors are derived from Table 6.6 which also delineates the bands. This table was obtained by reading through the computer files and aggregating annual costed demand and stock value for 45 mutually exclusive parameter selections e.g. 0.02% of the annual costed demand is accounted for by products with k < 0.75 and SL < 0.9250 and OI/LT > 1.499. These are applied to Table 6.5 assuming $k_{LD} = 0.924/1.198$ represents both of the bands <0.75 and $0.75 - 1.5$; $k_{LD} = 1.497/1.743$ represents $1.5 - 2.5$; $k_{LD} = 3.069$ represents $2.5 - 3.5$; and $k_{LD} = 3.413$ represents > 3.5. The bands represented by the service level and Order Interval/Lead Time cell identifiers are self-explanatory. In the cells representing two k bands the weighting factors are added together. The target service levels refer to the values allocated to individual products by the system to accomplish the optimisation function based on Product Group settings of 0.95 for all groups. The weighting factors do not run in smooth gradations partly because the bandwidths are unequal and partly because of data clustering caused by the influence of Buying Families.

An overall system error in service level due to imperfections in the distribution fit can now be calculated. A total system service level is first calculated assuming all products in each cell had the characteristics of the cell identifying parameters. Assuming the total demand on the Distribution Centre = T, then demand on cell 1,1 (using x, y convention) = 0.0011T, and ELS for cell 1,1 = 0.1 x 0.0011T.

$$\therefore \text{Total Target ELS} = \sum_{1,1}^{4,9} (1 - \text{Target SL}) \times \text{Weighting Factor} \times T$$

$$= 0.1 \times 0.0011 \times T$$

$$+ \; 0.05 \times 0.0342 \times T$$

$$+$$

.

.

.

k-band	SL band	OI/LT band	Weighting Factor based on Costed Demand	Weighting Factor based on Stock Value
< 0.75	< .9250	> 1.499	.0002	.0010
	.9250-.9649	> 1.499	.0054	.0055
	⩾ .9650	> 1.499	.0008	.0010
	< .9250	.75-1.499	.0026	.0065
	.9250-.9649	.75-1.499	.0022	.0034
	⩾ .9650	.75-1.499	.0007	.0005
	< .9250	< .75	.0004	.0017
	.9250-.9649	< .75	.0003	.0006
	⩾ .9650	< .75	0	.0001
0.75-1.5	< .9250	> 1.499	.0009	.0011
	.9250-.9649	> 1.499	.0288	.0308
	⩾ .9650	> 1.499	.0224	.0220
	< .9250	.75-1.499	.0400	.0575
	.9250-.9649	.75-1.499	.1057	.1435
	⩾ .9650	.75-1.499	.0074	.0067
	< .9250	<.75	.0517	.0985
	.9250-.9649	<.75	.0589	.0570
	⩾ .9650	<.75	.0005	.0017
1.5-2.5	< .9250	> 1.499	.0001	.0003
	.9250-.9649	> 1.499	.0125	.0083
	⩾ .9650	> 1.499	.0301	.0247
	< .9250	.75-1.499	.0005	.0007
	.9250-.9649	.75-1.499	.0699	.0720
	⩾ .9650	.75-1.499	.0505	.0383
	< .9250	<.75	.0410	.0429
	.9250-.9649	<.75	.1364	.1165
	⩾ .9650	<.75	.0285	.0143
2.5-3.5	< .9250	> 1.499	0	0
	.9250-.9649	> 1.499	.0027	.0020
	⩾ .9650	> 1.499	.0101	.0051
	< .9250	.75-1.499	0	0
	.9250-.9649	.75-1.499	.0112	.0091
	⩾ .9650	.75-1.499	.0441	.0370
	< .9250	<.75	.0021	.0035
	.9250-.9649	<.75	.0521	.0475
	⩾ .9650	<.75	.0111	.0076
> 3.5	< .9250	> 1.499	.0015	.0009
	.9250-.9649	> 1.499	.0007	.0009
	⩾ .9650	> 1.499	.0048	.0037
	< .9250	.75-1.499	.0002	.0004
	.9250-.9649	.75-1.499	.0041	.0022
	⩾ .9650	.75-1.499	.0652	.0503
	< .9250	<.75	.0008	.0014
	.9250-.9649	<.75	.0297	.0328
	⩾ .9650	<.75	.0608	.0386

Table 6.6 Weighting Factors for k/Service Level/Order Interval-to-Lead Time Ratio Categories

.

.

$$+ 0.02 \times 0.0608 \times T$$

$$= 0.0470\ T$$

$$\text{Overall Target SL} = \frac{T - 0.0470\ T}{T}$$

$$= \underline{\underline{0.9530}}$$

The error of 0.0030 is clearly due to the assignment of uniform parameters to all products in each cell.

The expected service level is now calculated by the same method.

$$\text{Total Expected ELS} = \sum_{1,1}^{4,9} (1 - \text{Exp.SL}) \times \text{Weighting Factor} \times T$$

$$= 0.0988 \times 0.0011 \times T$$

$$+ 0.0467 \times 0.0342 \times T$$

+

.

.

.

.

$$+ 0.0174 \times 0.0608 \times T$$

$$= 0.0447\ T$$

$$\text{Overall Expected SL} = \frac{T - 0.0447\ T}{T}$$

$$= \underline{\underline{0.9553}}$$

As the 'boxing' technique introduces an error of 0.0030,

$$\text{the true expected SL} = \frac{0.9553}{0.9530} \times 0.95$$

$$= \underline{\underline{0.9523}}$$

$$\text{or \% error in ELS} = \underline{\underline{-\ 4.6\%}}$$

In summary, the observed distribution truncates rather more quickly than the Gamma distribution. This engenders a slightly higher service level than the target setting. The overprovisioning is concentrated in the slow-moving products with low k values. The average stock value should not be affected by this error as the reorder level is calculated assuming a Gamma distribution for reporting as well as order-generating purposes. Hence there will be no difference between the target and expected values.

6.2 ERRORS IN NUMERICAL APPROXIMATIONS TO GAMMA TABLES AND IN CONVERSION FROM MAD TO STANDARD DEVIATION

Fig. 6.2 illustrates the utilisation of the three Johnston functions defined in Section 3.4. It can be seen that they constitute the foundation of the control system around which the system operates.

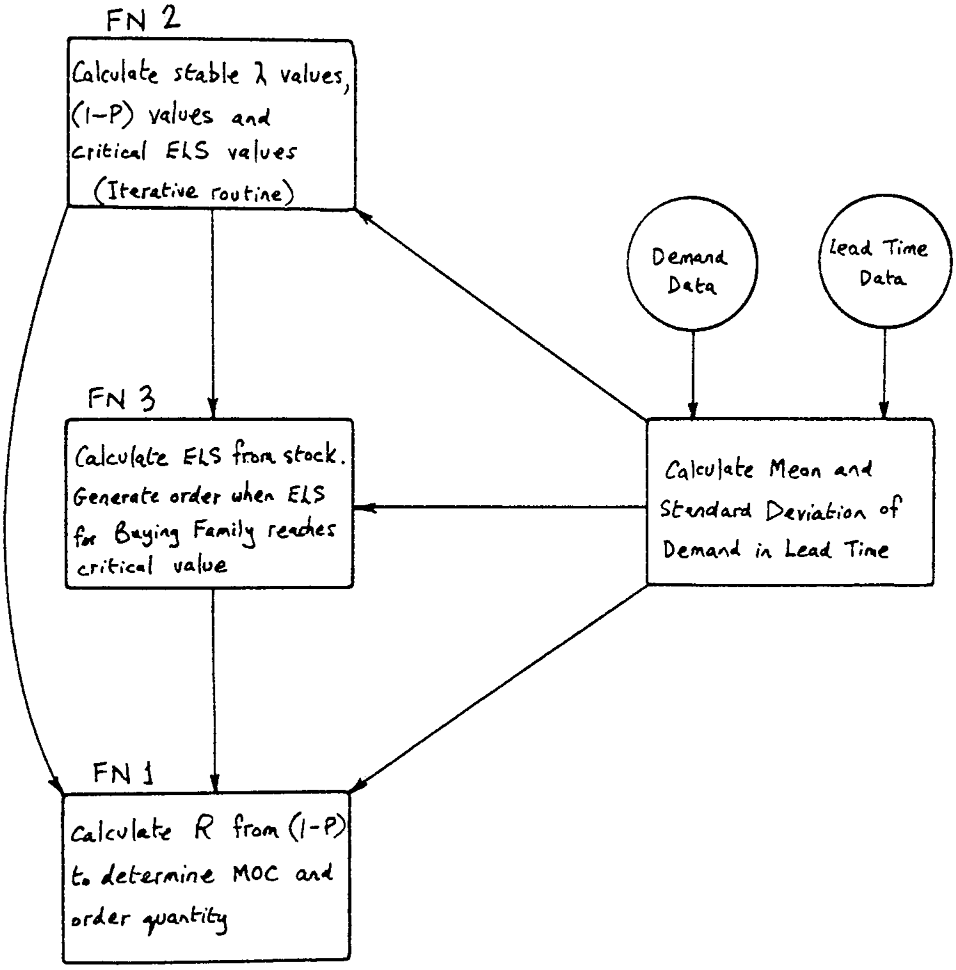

Fig. 6.2 Utilisation of Johnston Functions in Control System

The accuracy of the functions has a potentially important bearing upon the performance of the system. As the region of interest in this particular application is concentrated into a relatively small part of the range over which the expressions were developed, more selective accuracy checks are here carried out.

Also, the functions all use the Gamma modulus, k, which is calculated from the mean and standard deviation of the demand in the lead time. Errors in the calculation of the standard deviation are therefore inextricably linked to errors in the functions and they are evaluated jointly.

The regions of interest are identified by analysing the data for a Distribution Centre with respect to k and P (Tables 6.7 and 6.4 respectively).

k_{LD} Range	Products	Annual Costed Demand		Average Stock Value	
		%	% cumulative	%	% cumulative
< 0.25	17	.0002	.0002	.0005	.0005
0.25 - 0.50	43	.0014	.0016	.0025	.0030
0.5 - 0.75	274	.0111	.0127	.0172	.0202
0.75 - 1.0	902	.0574	.0701	.0913	.1115
1.0 - 1.25	1494	.1641	.2342	.2183	.3298
1.25 - 1.5	896	.0949	.3291	.1091	.4389
1.5 - 1.75	734	.1124	.4415	.0988	.5377
1.75 - 2.0	599	.0928	.5343	.0837	.6214
2.0 - 2.5	902	.1646	.6989	.1355	.7569
2.5 - 3.0	513	.0817	.7806	.0722	.8291
3.0 - 3.5	292	.0517	.8329	.0397	.8688
3.5 - 4.0	231	.0522	.8845	.0406	.9094
4.0 - 4.5	126	.0240	.9085	.0185	.9279
4.5 - 6.0	244	.0545	.9630	.0456	.9735
6.0 - 12.0	139	.0328	.9958	.0224	.9959
> 12.0	28	.0044	1.0002	.0040	.9999
Totals	7434	1.0002	-	0.9999	-

Table 6.7 Analysis of k Ranges

A random sample of 102 products produced a Correlation Coefficient between k and P of -0.0471 which indicates that there is no correlation. Hence products with a k value $\leqslant$ 6 and a P value of $\leqslant$ 0.99 account for 0.9630 x 0.9937 = 0.9569 of the annual costed demand. These limits always result in $U < 6$, therefore the accuracy tests are delimited by $k \leqslant 6$, $U \leqslant 6$, $P \leqslant 0.99$. Within the overall confines of these limits, it can be deduced from Tables 6.7 and 6.4 that 0.9502 of the annual costed demand is accounted for by products with $0.5 \leqslant k \leqslant 6$ and $0.75 \leqslant P \leqslant 0.99$; and 0.8431 by products with $0.5 \leqslant k \leqslant 4$ and $0.85 \leqslant P \leqslant 0.99$. Maximum interest is therefore focussed into these smaller regions.

Values of P and PLS were calculated for U, k combinations for $0.5 \leqslant U \leqslant 6$, step 0.5; $0.1 \leqslant k \leqslant 6$, step 0.1 using equations 18 and 19 in Section 3.4. The error in these computations has already been demonstrated to be negligible (Section 3.4), thus the results are taken as the theoretical base against which the errors in the Johnston functions are compared. The following values were computed from the Johnston functions:

a) U from Function 1, with (1-P) and k as input.

b) PLS from Function 2, with (1-P) and k as input.

c) PLS from Function 3, with U and k as input.

The percentage errors from the functions were calculated and these are presented in grid form in Figs. 6.3, 6.4 and 6.5 for Functions 1, 2 and 3 respectively. In each case the smaller curved quadrilateral delineates $0.5 \leqslant k \leqslant 4.0$, $0.85 \leqslant P \leqslant 0.99$ and the larger one delineates $0.5 \leqslant k \leqslant 6.0$, $0.75 \leqslant P \leqslant 0.99$.

The errors in Function 1 were found to be extremely small (average 0.07% of the true U) in the region of maximum interest and they compare closely with Johnston's published figures for wider limits (refer Sect. 3.4).

The errors in Function 2 (average 1.61% of the true PLS) compare

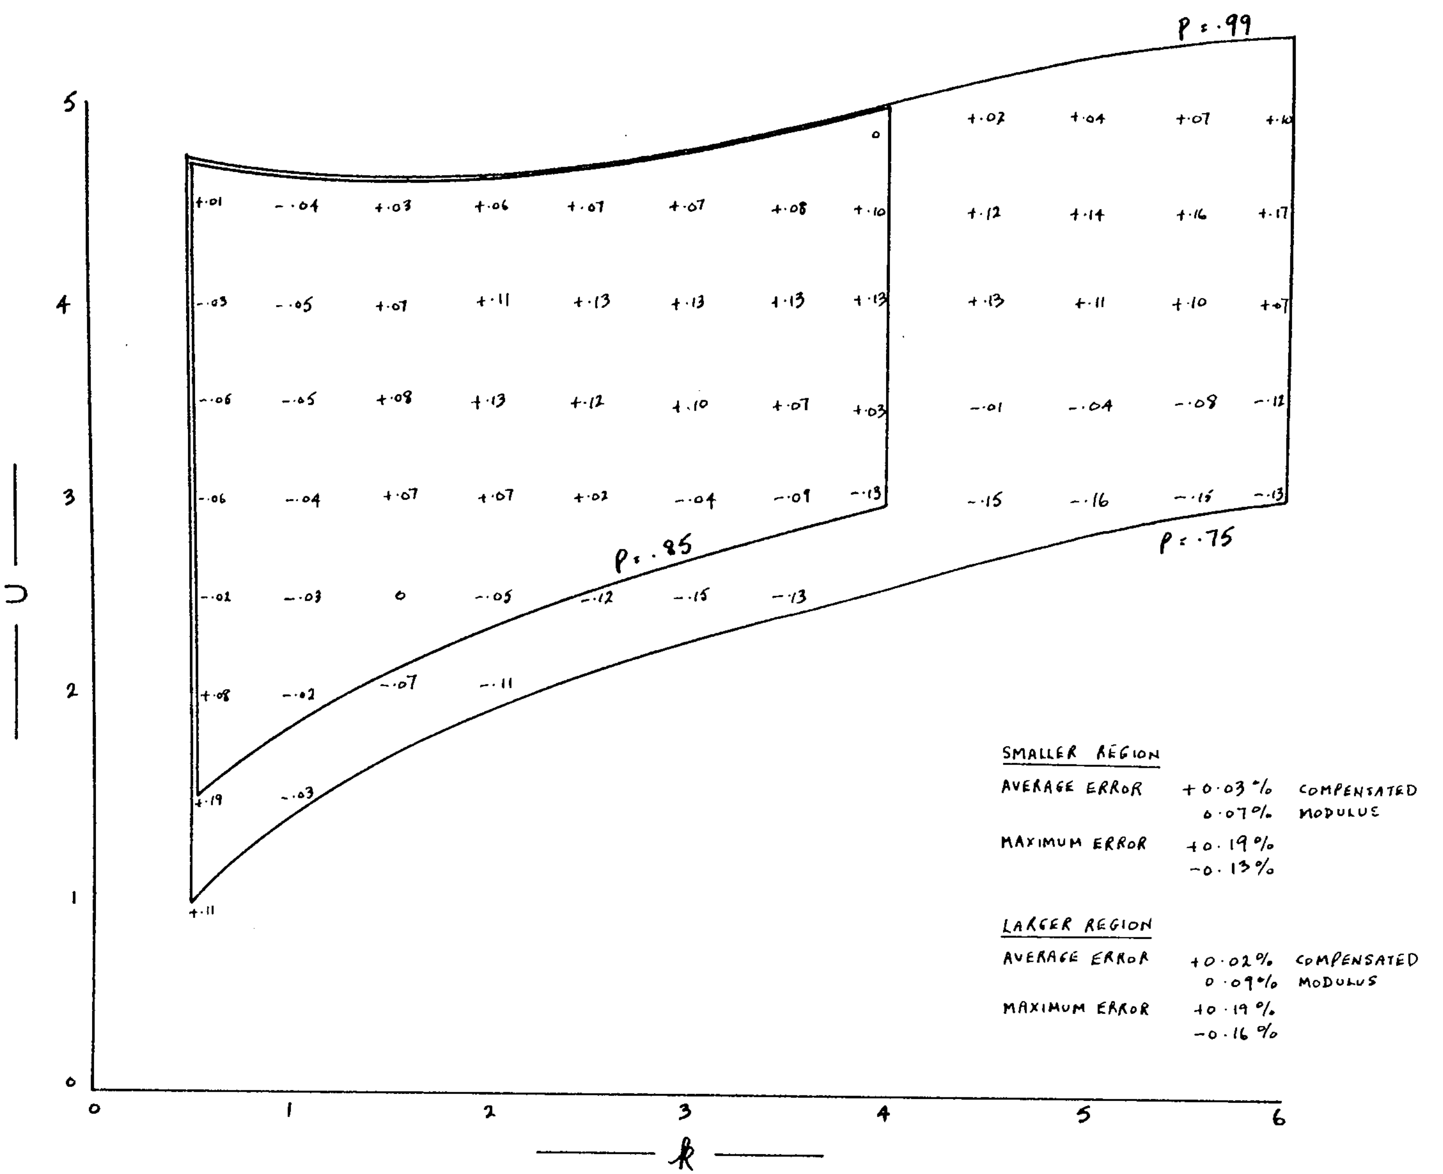

Fig. 6.3 Errors in Function 1

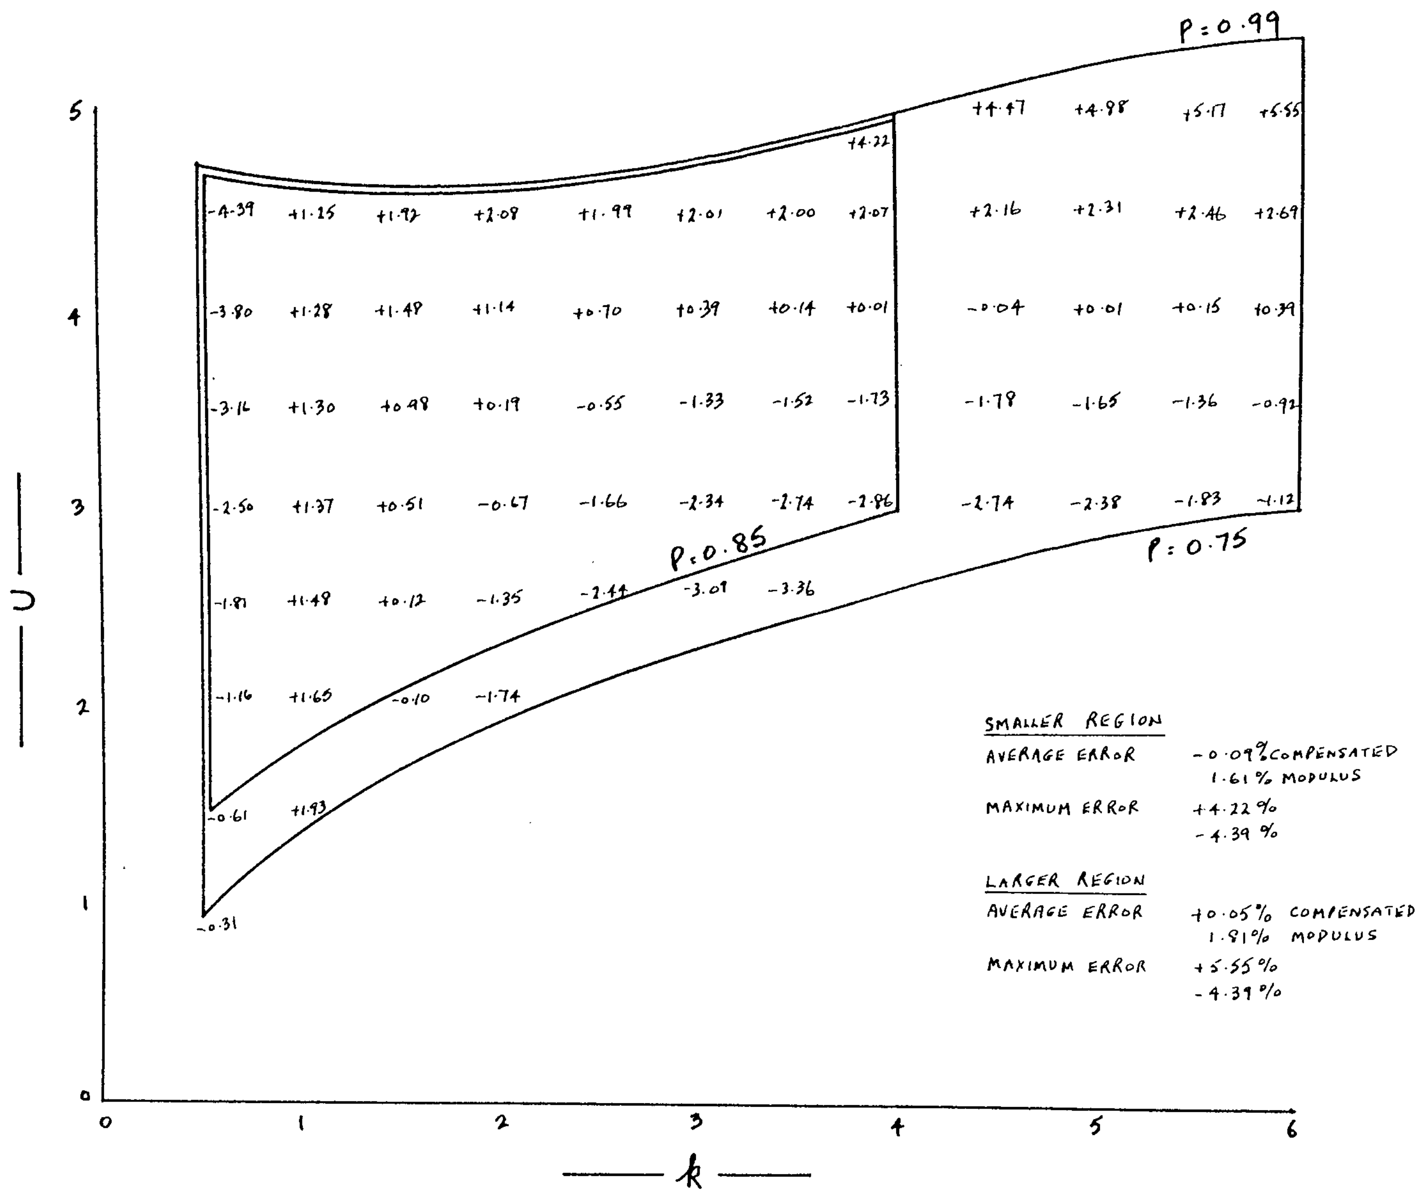

Fig. 6.4 Errors in Function 2

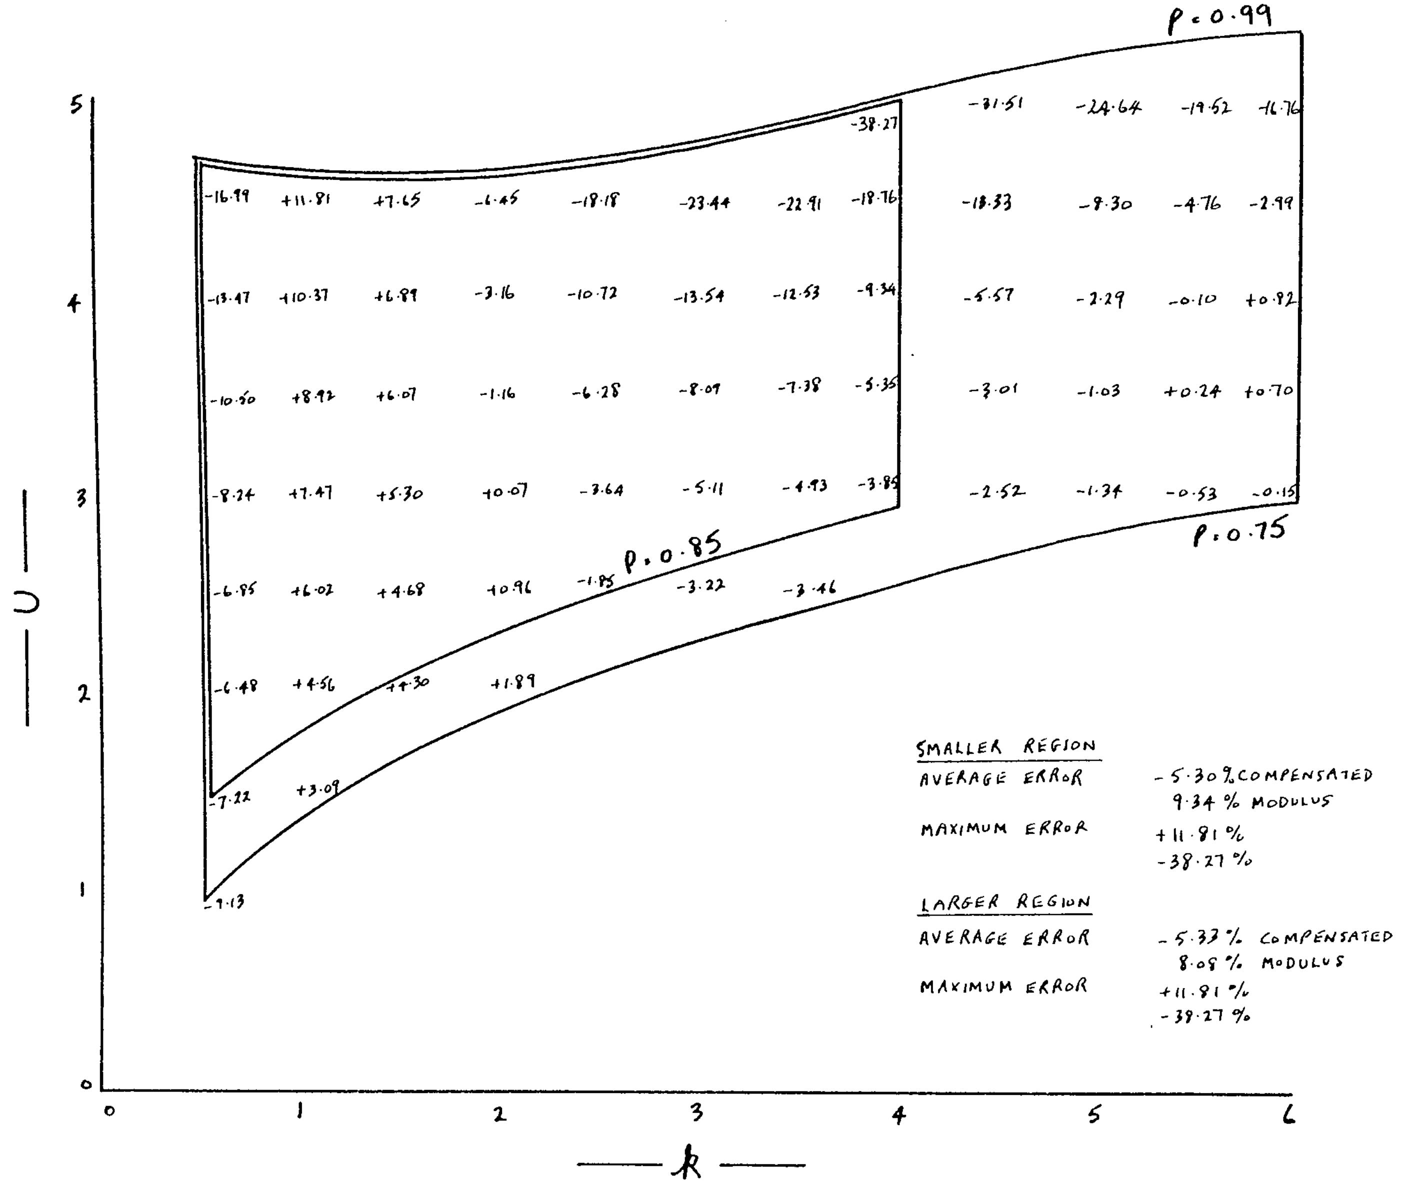

Fig. 6.5 Errors in Function 3

favourably with Johnston's equivalent figure (3.7%) for a different region.

The errors in Function 3 (average 9.34% of the true PLS) are larger than those reported by Johnston (cf 6.2%) and the errors in some areas are unacceptably high. An inspection of Tables 6.7 and 6.4 reveals that the greatest concentration of data is around the point k = 1.125, P = 0.97. Fig. 6.5 suggests that errors of approximately 10% would be experienced in this region.

Because of the cumbersome calculation of $\Gamma(k)$ in the expression which relates MAD to Standard Deviation (Section 3.4, equation 20), a uniform factor of 1.33 is used in the system for conversion. The true conversion factors are evaluated in Table 3.1 and the inherent errors in using 1.33 are given in Table 6.8.

k Range (Mid-point)	% Error
0.50	+9.86
0.75	+4.84
1.25	+0.64
1.75	-1.21
2.25	-2.19
2.75	-2.85
3.25	-3.29
3.75	-3.62
4.25	-3.89
4.75	-4.09
5.25	-4.25
5.75	-4.39
9.00	-4.89

Table 6.8 Errors in Converting MAD to Standard Deviation

The effects on service level and average stock of the combined conversion and function errors are examined by first working through an example for a single product.

Assume σ_{LD} = True standard deviation of demand in lead time

and let true $k = 0.5$,

target service level = 90%,

Nominal Order Interval = 2 x Mean Lead Time

From Table 3.1, true MAD_{LD}/σ_{LD} conversion factor = 1.4611

$$\therefore \quad 0.5 = \frac{\mu_{LD}^2}{1.4611^2 \times MAD_{LD}^2}$$

$$\therefore \quad \frac{\mu_{LD}^2}{MAD_{LD}^2} = 0.5 \times 1.4611^2$$

$$\text{Calculated } k = \frac{\mu_{LD}^2}{1.33^2 \times MAD_{LD}^2}$$

$$= \frac{0.5 \times 1.4611^2}{1.33^2}$$

$$= \underline{\underline{0.60343}}$$

For a 90% service level, an average of 10% of the demand must be lost in each order cycle.

$\therefore$ PLS = 0.1 x 2 = 0.2

From the values of P and PLS computed from U, k combinations, the value of U corresponding to PLS = 0.2, k = 0.5 is 1.7391. Hence, true reorder level,

$$R = 1.7391 \times \sigma_{LD}$$

$$= 1.7391 \times 1.4611 \times MAD_{LD}$$

$$= \underline{2.5410 \times MAD_{LD}}$$

When the iterative application of Function 2 (Fig. 6.2) culminates in a stable condition of the variables, the PLS must be 0.2, as the attainment of the target service level is the acceptance criterion. Using a k of 0.60343 this yields a (1-P) of 0.13709.

An order is generated when the available stock produces a PLS of 0.2 from Function 3. For a single product, the available stock and reorder level are then coincident. With a k value of 0.60343, Function 3 requires a U value of 1.62883 to produce a PLS of 0.2. U is itself calculated using the standard MAD_{LD} / σ_{LD} conversion factor of 1.33.

$$\therefore \text{ R from Fn 3 } = 1.62883 \times \sigma_{LD}$$
$$= 1.62883 \times 1.33 \times MAD_{LD}$$
$$= \underline{2.1663 \times MAD_{LD}}$$

At this reorder level, where the system actually triggers orders, the true value of U = 2.1663/1.4611 = 1.48265 which corresponds to a true PLS of 0.24305.

The order size is determined by subtracting the available stock from the Maximum Order Cover (MOC). The MOC is obtained by adding the average sales per nominal order cycle to the reorder level calculated from Function 1, using k = 0.60343, (1-P) = 0.13709. This produces a U value of 1.67780,

$$\therefore \text{ R from Fn 1 } = 1.67780 \times \sigma_{LD}$$
$$= 1.67780 \times 1.33 \times MAD_{LD}$$
$$= \underline{2.2315 \times MAD_{LD}}$$

The three reorder levels are shown pictorially in Fig. 6.6.

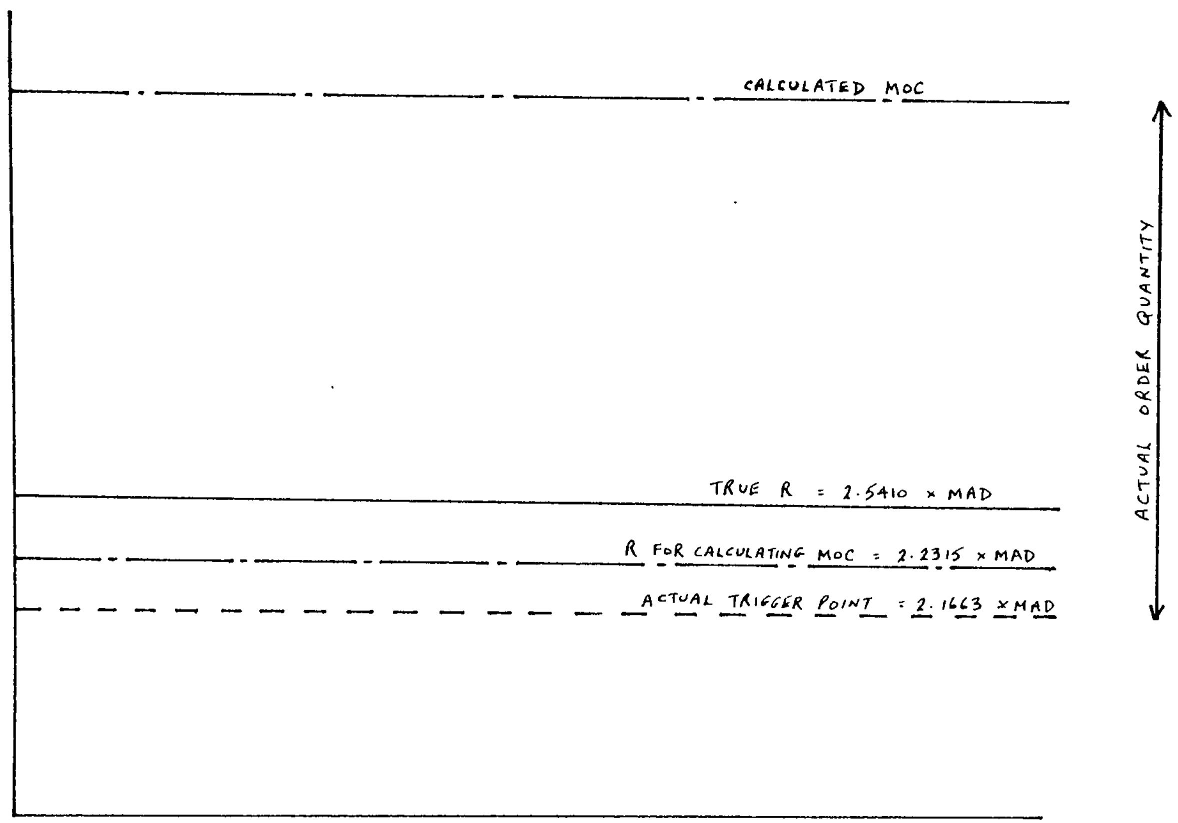

Fig. 6.6 System Reorder Levels

For a stable ordering pattern the average order quantity must equate with the average sales per order cycle. From Fig. 6.6,

Order Qty = MOC - Available stock at placement

The MOC is calculated as the reorder level (using Fn 1) plus the demand per nominal order cycle less the target ELS per order.

$\therefore$ Order quantity $= R\ (\text{Fn } 1) + 2 \times \mu_{LD} - 0.2 \times \mu_{LD} - R(\text{Fn } 3)$

$$= 2.2315 \times MAD_{LD} + 1.8 \times \mu_{LD} - 2.1663 \times MAD_{LD}$$

For $k = 0.5$, $\mu_{LD} = \sqrt{0.5} \times 1.4611 \times MAD_{LD}$

$\therefore$ Order quantity $= MAD_{LD}\ (2.2315 + (1.8 \times \sqrt{0.5} \times 1.4611) - 2.1663)$

$$= 1.9249 \times MAD_{LD}$$

$$\text{Demand per expected order cycle} = \frac{\text{Sales per expected order cycle}}{\text{Expected SL}}$$

$$= \frac{1.9249 \times MAD_{LD}}{\text{Expected SL}}$$

$$\text{Expected SL} = \frac{\text{Demand per expected order cycle - revised ELS per order}}{\text{Demand per expected order cycle}}$$

$$= \frac{1.9249 \times MAD_{LD}/\text{Exp. SL} - 0.24305 \times \mu_{LD}}{1.9249 \times MAD_{LD}/\text{Exp. SL}}$$

Substituting $\sqrt{0.5} \times 1.4611 \times MAD_{LD}$ for μ_{LD} and transposing,

Expected SL = $\underline{\underline{0.8846}}$

or, error in ELS = $\underline{\underline{+15.4\%}}$

The Average stock level is calculated in the system as:

$AS = R\ (\text{Fn } 1) - \mu_{LD} + \frac{1}{2}$ Nominal Order Interval demand

$= R\ (\text{fn } 1) - \mu_{LD} + \mu_{LD}$

$= 2.2315 \times MAD_{LD}$

The expected AS = R (Fn 3) − μ_{LD} + ½ Order quantity

$$= 2.1663 \times MAD_{LD} - \sqrt{0.5} \times 1.4611 \times MAD_{LD} + 0.5 \times 1.9249 \times MAD_{LD}$$

$$= 2.0956 \times MAD_{LD}$$

$$\text{Error in AS} = \frac{2.2315 - 2.0956}{2.0956} = \underline{\underline{+6.49\%}}$$

The calculations were repeated for all permutations of k_{LD} = 0.5, 1.0, 2.0, 3.0, 4.0; Target Service Level = 0.90, 0.95, 0.98; Order Interval/Lead Time Ratio = 2.0, 1.0, 0.5. The results are given in Table 6.9 and the expected service levels are shown graphically in Fig. 6.7. As the service level errors are small when the MAD/SD conversion errors are zero (at k ≃ 1.4) and the errors are at a maximum when the conversion errors are highest (at k = 0.5), this strongly suggests that the conversion errors dominate the function errors. The errors in average stock are seen to be predominantly positive which indicates that the levels of stock required by the system are in general less than the predicted levels.

The weighting factors in Table 6.9 are taken from Table 6.6. As there are errors in both the service levels and average stock values, two sets of weighting factors are used based on costed demand and stock value respectively.

Overall system errors were calculated by applying the weighting factors exactly as in Section 6.1. The expected service level (allowing for the 'boxing' error) evaluates as 0.9500, which means that, fortuitously, the positive and negative errors cancel each other exactly and there is no overall system error. The overall average stock value error is +0.94% which also implies a favourable degree of cancellation.

In summary, in spite of sizeable errors in some individual cases for

k_{LD}	OI/LT Ratio →	2.0			1.0			0.5		
	Target SL →	.90	.95	.98	.90	.95	.98	.90	.95	.98
	Target PLS →	.20	.10	.04	.10	.05	.02	.05	.025	.01
0.5	Expected SL →	.8846	.9362	.9705	.8758	.9304	.9683	.8665	.9273	.9689
	% Error ELS →	+15.4	+27.6	+47.5	+24.2	+39.2	+58.5	+33.5	+45.4	+55.5
	Wt. Fact. ELS →	.0002	.0054	.0008	.0026	.0022	.0007	.0004	.0003	0
	% Error AS →	+6.49	+2.05	+0.75	+2.45	+0.98	+0.82	+1.05	+0.81	+1.12
	Wt. Fact. AS →	.0010	.0055	.0010	.0065	.0034	.0005	.0017	.0006	.0001
1.0	Expected SL →	.8991	.9489	.9792	.8956	.9465	.9779	.8856	.9399	.9724
	% Error ELS →	+0.9	+2.2	+4.0	+4.4	+7.0	+10.5	+14.4	+20.2	+38.0
	Wt. Fact. ELS →	.0009	.0288	.0224	.0400	.1057	.0074	.0517	.0589	.0005
	% Error AS →	+5.99	+1.33	-0.43	+1.71	-0.22	-1.00	-0.25	-0.96	-1.27
	Wt. Fact. AS →	.0011	.0308	.0220	.0575	.1435	.0067	.0985	.0570	.0017
2.0	Expected SL →	.9032	.9522	.9813	.9031	.9528	.9817	.9044	.9541	.9825
	% Error ELS →	-3.2	-4.4	-6.5	-3.1	-5.6	-8.5	-4.4	-8.2	-12.5
	Wt. Fact. ELS →	.0001	.0125	.0301	.0005	.0699	.0505	.0410	.1364	.0285
	% Error AS →	+7.61	+2.46	+0.82	+3.56	+1.33	+0.73	+1.57	+0.88	+0.84
	Wt. Fact. AS →	.0003	.0083	.0247	.0007	.0720	.0383	.0429	.1165	.0143
3.0	Expected SL →	.9022	.9517	.9811	.9029	.9527	.9817	.9063	.9553	.9829
	% Error ELS →	-2.2	-3.4	-5.5	-2.9	-5.4	-8.5	-6.3	-10.6	-14.5
	Wt. Fact. ELS →	0	.0027	.0101	0	.0112	.0441	.0021	.0521	.0111
	% Error AS →	+9.54	+3.48	+1.56	+5.47	+2.47	+1.87	+3.10	+2.23	+2.51
	Wt. Fact. AS →	0	.0020	.0051	0	.0091	.0370	.0035	.0475	.0076
4.0	Expected SL →	.9016	.9516	.9814	.9030	.9531	.9820	.9065	.9556	.9831
	% Error ELS →	-1.6	-3.2	-7.0	-3.0	-6.2	-10.0	-6.5	-11.2	-15.5
	Wt. Fact. ELS →	.0015	.0007	.0048	.0002	.0041	.0652	.0008	.0297	.0608
	% Error AS →	+10.79	+3.94	+1.55	+6.55	+2.68	+1.78	+3.51	+2.29	+2.46
	Wt. Fact. AS →	.0009	.0009	.0037	.0004	.0022	.0503	.0014	.0328	.0386

Table 6.9 Errors in Service Level and Average Stock Value resulting from Functions and MAD/SD Conversions

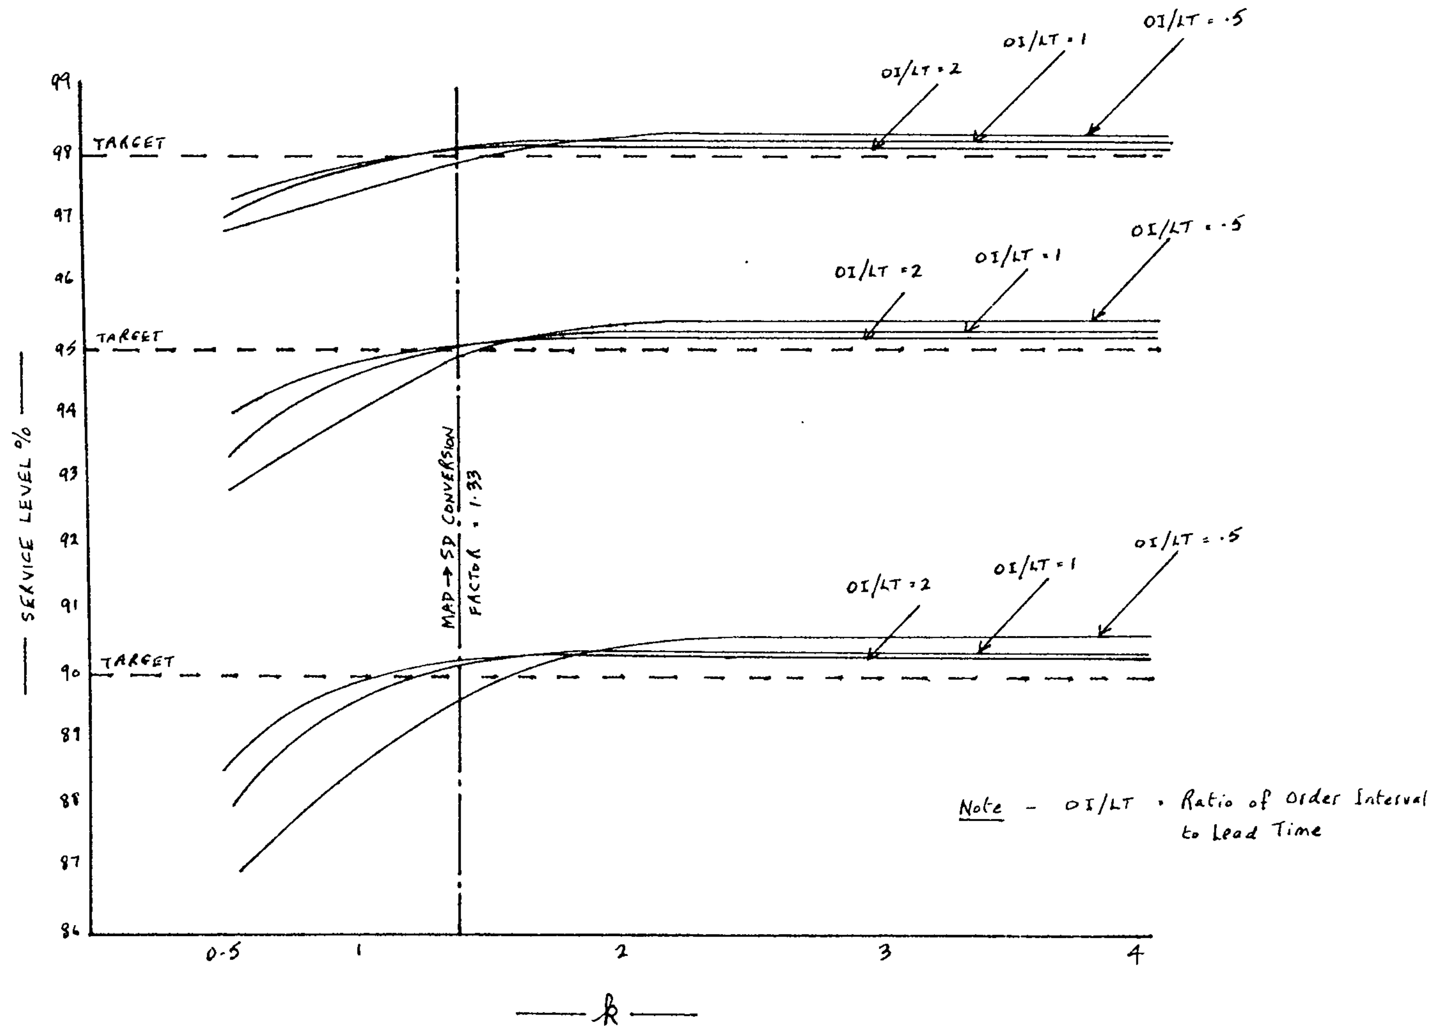

Fig. 6.7 Effect on Service Level of Function Errors and MAD/SD Conversion Errors

both the functions and the MAD/SD conversions, the overall system errors in the service level and average stock value are negligible. This is due to:-

a) a judicious, if somewhat fortunate, choice of the uniform MAD/SD conversion factor of 1.33, and

b) counteraction of errors due to system interactions.

6.3 ERRORS DUE TO DEMAND AND LEAD TIME CORRELATION AND AUTOCORRELATIONS

The common assumption in stock control theory that demand and lead times are independent random variables implies that:-

a) demand and lead times are not cross-correlated,

b) lead times are not autocorrelated, and

c) demand is not autocorrelated.

These assumptions are examined in turn and the system effects of breaches are assessed.

6.3.1 CROSS-CORRELATION BETWEEN DEMAND AND LEAD TIMES

A positive correlation between demand and lead time suggests that high demand periods would be accompanied by long lead times and vice versa. This infers that the variance of demand in lead time would be greater than if the peaks and troughs occurred independently. Buffer stocks calculated on the assumption of independence would therefore be inadequate to provide the intended service levels. Conversely, a negative cross-correlation would result in over-provisioning.

The mode of reasoning employed in the System Dynamics study in Section 4.1 suggests that a positive cross-correlation is a likely outcome of market dynamics. A period of high consumer demand would deplete retail stocks which would result in an increased ordering rate on the warehouse to:-

a) satisfy the higher level of retail demand,

b) raise the retail stock levels commensurately, and

c) increase the volume of pipeline orders and materials to sustain the higher throughput.

The raised ordering rate would result in a depletion of warehouse stocks, and an analagous process would cause a further amplification in the ordering rate on the supplier or factory. The consequence would be an increased loading on the supplying source and an extension of lead times. This process would take some time to work through the retail-distribution-supplier chain, so there would be a time lag between peak demand and maximum lead times.

The study was initiated by tabulating the average lead times extant in each month between February 1978 and August 1980 for each of the 13 suppliers described in Section 6.1. Monthly demand (strictly 'sales') data were collected for the Product Groups pertaining to the suppliers over the same period. The demand data were deflated by dividing by the Index for Construction Material Prices for the respective months, then deseasonalised by dividing by seasonal factors elicited using the 'Classical Decomposition Method' of time series forecasting described in Section 3.1 (Note - deseasonalising is justified on the grounds that the demand data processed in the Inventory Management system has already been deseasonalised by the forecasting routines. Hence any autocorrelation caused by seasonal influences has been removed before demand and lead time are convolved).

Standard correlation tests between the pre-processed demand and lead time data were carried out for:-

a) synchronised demand and lead time,

b) a lead time lag of one month, and

c) a lead time lag of three months.

As the demand was measured at the Distribution Centre (i.e. the retail stage was excluded), a 3-month time lag was considered adequate for the lead time to react to demand changes. The results are given in Table 6.10.

Supplier Code	Number of Observations	Lead Time Lag			5%* Significance Level
		Zero	1 Month	3 Months	
B	22	-.0222	+.0665	+.1294	.42
CN	15	-.4099	-.5939	-.2675	.51
CO	16	+.4007	+.2434	+.4365	.50
D	12	+.5316	+.5101	+.6855	.58
E	24	+.3913	+.2425	+.1077	.41
F	25	+.0491	-.1073	+.1751	.40
NT	18	+.4437	+.3852	+.3906	.47
PO	14	-.3683	-.1344	+.2761	.53
WD	26	-.0117	+.1264	+.1920	.39
G	20	-.1390	-.0538	-.1274	.44
PE	26	-.0210	+.0837	-.3388	.39
U	31	-.3311	-.1009	-.0824	.35
WP	23	+.2200	+.3630	+.3681	.41
	Average	+.0564	+.0793	+.1496	

*Values refer to a 2-tail significance test (86)

Table 6.10 Correlation Coefficients for Demand/Lead Time

There is some evidence of a predominance of positive correlations with a 3-month lag, but a fairly even balance of positive and negative correlations with a zero or 1-month lag. This is intuitively supportable as most of the major suppliers carry 2-3 months' shelf stocks which would have to be depleted or severely eroded before lead times would be generally affected. However, the exponential smoothing factors for demand and lead time are set to the same value, which implies synchronisation when the demand and lead time forecasts are convolved. The zero time lag correlations are therefore more directly relevant to buffer stock calculations, and hence service level considerations, than the lagged correlations attributable to market dynamics. With zero time lag, none of the Correlation Coefficients attain the 5% significance level and there is no definite polarity bias. It is therefore concluded that cross-correlation between demand and lead times cannot be substantiated as a contributory cause of service level errors.

6.3.2 LEAD TIME AUTOCORRELATION

On considering autocorrelated lead times, a distinction will first be drawn between the simple case where lead times are temporally independent, and the compound case where lead times are subject to possible overlap.

If the order interval is sufficiently large relative to the lead times to preclude the possibility of overlapping lead times, it is here asserted that the assumption of a random sequence of lead times may be vitiated without affecting the service level, provided that:-

a) the lead time distribution is stationary, and

b) the order point is calculated using the population mean and variance and is not updated by recent observations.

An autocorrelated lead time sequence would certainly engender an autocorrelated demand in lead time sequence, but this should not affect the variance and hence the buffer stock. Thus the desired service level should be obtained if measured over a sufficiently long period.

When lead times overlap, demand is shared between two or more lead times, and an autocorrelated sequence of demand in lead time variates arises even if the lead times are not themselves autocorrelated. If lead times are autocorrelated at low-order time lags, then long lead times tend to follow each other consecutively and the amount of shared demand increases. Stockouts during periods of shared demand will be demonstrated to be a serious problem in Section 6.4. As this condition is exacerbated by autocorrelated lead times, then the problems due to autocorrelated lead times can be subsumed under the wider problem area of overlapping lead times. This wider problem is treated by simulation methods in Section 6.4. As the simulation data must reflect the lead time autocorrelations with an acceptable degree of fidelity, the magnitude of the autocorrelations are first assessed here.

Lead time autocorrelations were investigated by utilising the SIBYL-RUNNER package of forecasting programs developed by Makridakis and Wheelwright (87). The data for the 13 suppliers recorded in Table 6.2 were input to the package and subjected to autocorrelation analysis for time lags 1-24 at zero and first levels of differencing. The analysis was performed by the SIBYL segment of the package purely for the purpose of computing the Autocorrelation Coefficients and associated information (The RUNNER segment enables optional forecasting routines to be run after the data has been analysed). The data for all suppliers were found to be stationary, which justifies the assumption made earlier in this Section. At the zero level of differencing (i.e. the simple analysis of raw data) a typical pattern of autocorrelations over the time lags was evident, an example of which is given in Fig. 6.8.

```
TIME LAGS                                                   AUTOCORRELATIONS
 24                    :    *     :                          0.01496
 23                    :   *I     :                         -0.03821
 22                    :   *I     :                         -0.06179
 21                    :    I*    :                          0.03049
 20                    :    I*    :                          0.05661
 19                    :  * I     :                         -0.09795
 18                    :   *I     :                         -0.04552
 17                    :    *     :                         -0.02199
 16                    :   *I     :                         -0.07425
 15                    :    *     :                          0.00874
 14                    :    I*    :                          0.06074
 13                    :    I *   :                          0.08545
 12                    :    I *   :                          0.10314
 11                    :    *     :                          0.01210
 10                    :    *     :                          0.00470
  9                    :    I*    :                          0.05520
  8                    :    I*    :                          0.07164
  7                    :    *     :                          0.01461
  6                    :    I   * :                          0.13576
  5                    :    I*    :                          0.07291
  4                    :    I    *:                          0.19980
  3                    :    I    *:                          0.20992
  2                    :    I     :  *                       0.40107
  1                    :    I     :*                         0.31200
          I.I.I.I.I.I.I.I.I.I.I.I.I.I.I.I.I.I.I.I.I
         -1                   0                   +1
```

Fig. 6.8 Typical Example of Lead Time Autocorrelations

The pattern is characterised by the incidence of apparently meaningful positive autocorrelations for the lower order time lags and very small autocorrelations with no polarity bias for the remaining time lags. This can be supported on intuitive grounds, as the time lags represent the intervals between placing orders, which can vary considerably. Hence the only autocorrelations which would be expected would pertain to the lower order time lags.

Table 6.11 shows the Autocorrelation Coefficients for Time Lag 1 for all suppliers.

Supplier Code	Number of Observations	Autocorrelation Coefficient	5%* Significance Level
B	37	+.3934	.32
CN	32	+.1406	.35
CO	71	+.3120	.23
D	38	+.2505	.32
E	74	+.2446	.23
F	59	+.4179	.26
NT	34	+.1805	.33
PO	119	+.1472	.18
WD	49	+.0349	.28
G	41	+.0992	.31
PE	65	+.0134	.24
U	58	+.3331	.26
WP	48	+.2624	.28
	Average	+.2177	

*Values refer to a 2-tail significance test (86)

Table 6.11 Lead Time Autocorrelations for Time Lag 1

It can be seen that, not only are all of the Autocorrelation Coefficients positive, but 5 of the 13 exceed the 5% significance value. It is concluded that successive lead times are consistently positively autocorrelated with an average Autocorrelation Coefficient of approximately +0.22.

6.3.3 DEMAND AUTOCORRELATION

Positive demand autocorrelation indicates that particular lead times could encompass runs of high demand or low demand. Hence the extremes of demand in lead time would be wider than if the demands occurred randomly i.e. the variance would be greater. Buffer stocks computed on the premise of random demand sequences would therefore be inadequate to realise the set service levels.

Weekly demand (strictly 'sales') data were collected for each of 18 Product Groups over the period January 1978 to January 1980. The readings were first deflated by dividing by the Index of Construction Material Prices, then deseasonalised by dividing by seasonal factors. An Autocorrelation Coefficient for Time Lag 1 was calculated for each Product Group (Table 6.12).

Product Group	Autocorrelation Coefficient	Weighting Factor		Weighted Coefficient
C3	-.2171	.1066	.1942	-.130
C1	-.0237	.0876		
C2	+.0605	.1573	.3383	+.066
B3	+.0665	.0363		
E1	+.0726	.1447		
B2	+.2921	.0510	.1127	+.297
B1	+.2967	.0168		
H1	+.3025	.0449		
D1	+.4120	.0936	.2194	+.440
G2	+.4227	.0402		
D2	+.4343	.0051		
F1	+.4817	.0805		
D3	+.5203	.0012	.1354	+.590
A1	+.5556	.0638		
G1	+.6051	.0516		
G4	+.6396	.0031		
I1	+.6673	.0127		
G3	+.7338	.0030		

Note: the 5% significance value of autocorrelations for 24 observations (22 D.F.) is 0.41

Table 6.12 Demand Autocorrelations for Time Lag 1

Weighting factors were assigned to the Product Groups according to their annual costed demand. The Autocorrelation Coefficients were grouped into five categories for ease of aggregation, e.g. for the first group the joint Autocorrelation Coefficient was calculated as:

$$\frac{.1066}{.1942} \times (-.2171) + \frac{.0876}{.1942} \times (-.0237) = \underline{\underline{-.130}}$$

It can be seen that there is a strong predominance of positive autocorrelations and this is significant at the 5% level for 10 of the 18 Product Groups. It is concluded that weekly demands are subject to a strong positive autocorrelation.

Ray (88) has produced a formula for the variance of demand in lead time when demands are autocorrelated in a first-order auto-regressive sequence:

$$\sigma_{LD}{}^2 = \mu_D{}^2 \sigma_L{}^2 + \mu_L \frac{\sigma_D{}^2}{(1-a)^2} - \frac{2a\ \sigma_D{}^2}{(1-a)^2 (1-a^{\ 2})} \ (1 - pgf_L(a)\ )$$

where 'a' denotes the demand Autocorrelation Coefficient and pgf_L denotes the probability generating function for lead times. The derivation of the formula suggests that it is only appropriate for discrete lead time probability distributions, and the three tabulated examples in the paper - for Geometric, Uniform and truncated Poisson - are unsuitable as approximations to Gamma. Wagner (89), however, observes that "the Gamma distribution is the continuous analog of the Negative Binomial". To test this, the probability density functions for Gamma and Negative Binomial were compared for equal mean and standard deviation parameters (corresponding to a Gamma modulus, k, of 2.25). The result is shown in Fig. 6.9.

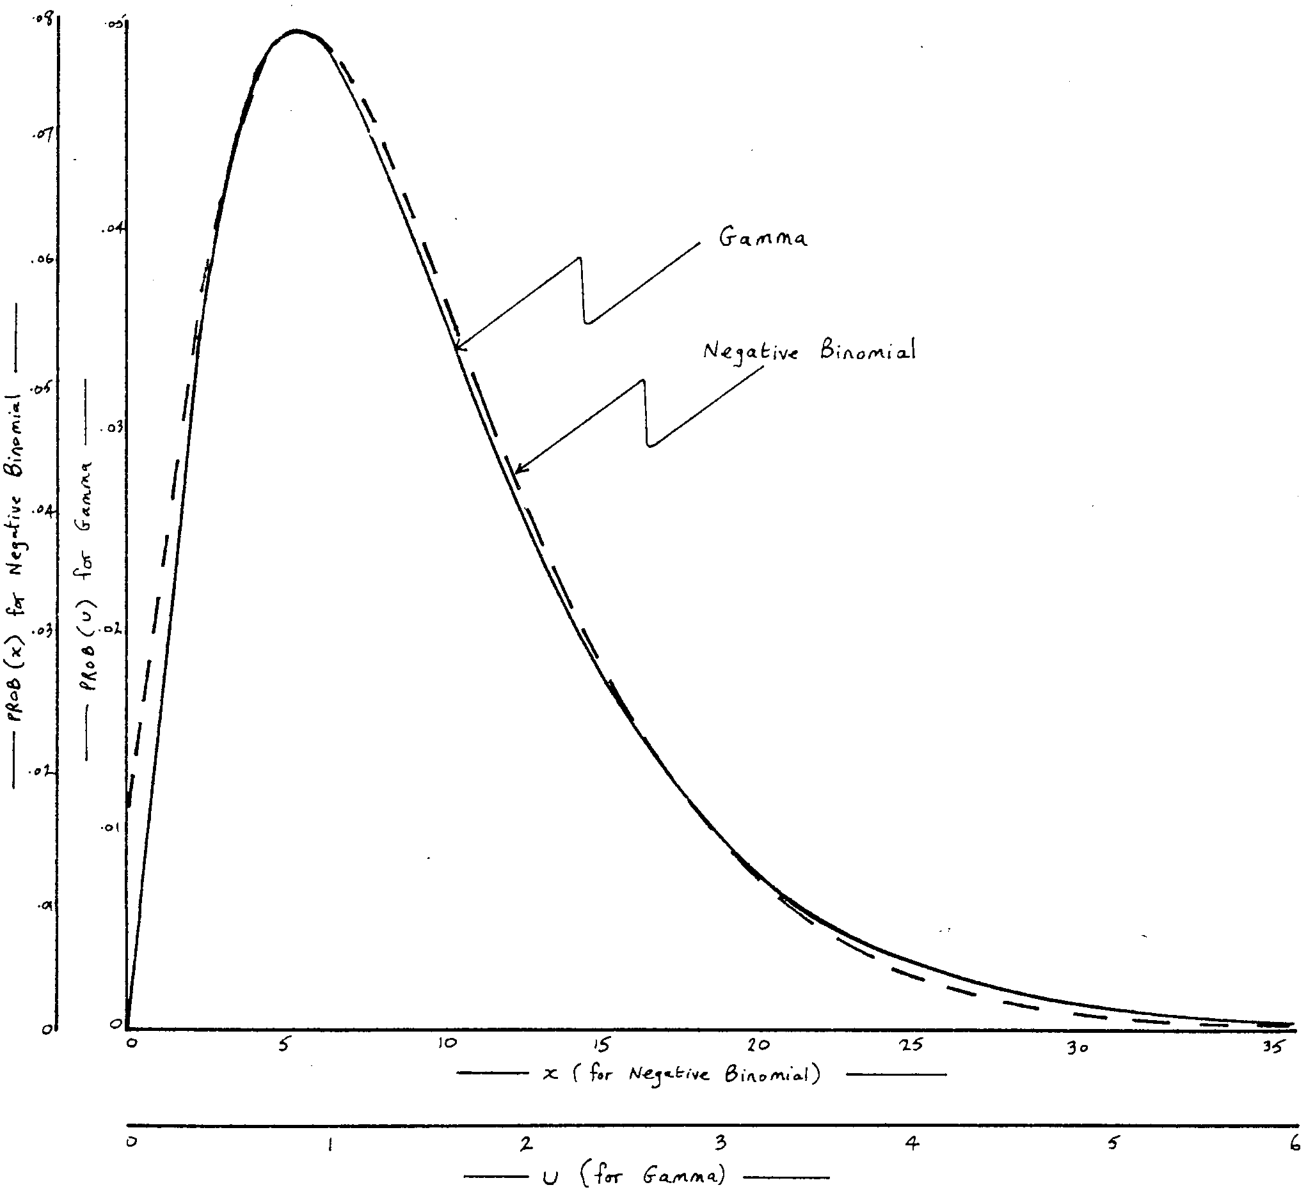

Fig. 6.9 Comparison Between Gamma and Negative Binomial for k = 2.25

It can be seen that the correspondence is extremely close. The test was repeated for other values of k with similar results except for very low values of k which are not appropriate to lead time parameters. It is concluded that the Negative Binomial could be used in place of Gamma for evaluating demand autocorrelation errors without significantly affecting the results.

For the negative binomial, the pgf, p(z), is defined as:

$$p(z) = (p/(1-qz))^r$$

where z is any random variable 0, 1, 2 ....∞

r is any integral value $= \dfrac{\mu^2}{(\sigma^2-\mu)}$

$$p = \frac{\mu}{\sigma^2}$$

and $q = (1-p) = \dfrac{\sigma^2-\mu}{\sigma^2}$

In the Ray equation the Autocorrelation Coefficient, a, takes the place of the random variable z.

Before discussing the aggregation requirements to obtain a total system effect, a particular example is worked through using the data values:

Target service level (SL) = 95%

μ_D = 1 unit/week

σ_D = 1.5811 units/week

μ_L = 2 weeks

σ_L = 1.4226 weeks

I = 4 weeks

a = -.130

From the normal expression for σ_{LD} (Section 3.3 eqn. 2),
but ignoring the marginal effect of the review cycle, T, for simplicity,

$$\sigma_{LD}^{\ 2} = \sigma_D^{\ 2}\mu_L + \mu_D^{\ 2}\sigma_L^{\ 2}$$

$$= (1.5811^2 \times 2) + (1^2 \times 1.4226^2)$$

$$= 7.0235$$

and $\sigma_{LD} = 2.6502$

$$k = \frac{(\mu_L \mu_D)^2}{\sigma_{LD}^2}$$

$$= \frac{(2 \times 1)^2}{7.0235}$$

$$= 0.5695$$

$$PLS = (1 - SL) \times \frac{I}{\mu_L}$$

$$= (1 - 0.95) \times \frac{4}{2}$$

$$= 0.1$$

Standardised reorder level, from an iterative application of equations (18) and (19) using k and PLS,

$$U = 2.4794$$

$$R = U \ \sigma_{LD}$$

$$= 2.4794 \times 2.6502$$

$$= 6.5709$$

This is the effective reorder level calculated by the system without regard to the autocorrelation of demand.

To obtain the expected service level taking into account the autocorrelation of demand, the negative binomial parameters are first calculated:

$$p = \frac{\mu_L}{\sigma_L^2}$$

$$= \frac{2}{1.4226^2}$$

$$= 0.9882$$

$q = 1 - p$

$= 1 - .9882$

$= .0118$

$$r = \frac{\mu_L^2}{(\sigma_L^2 - \mu_L)}$$

$$= \frac{2^2}{(1.4226^2 - 2)}$$

$= 168.13$ (168 to the nearest integer)

$pgf_L(a) = (p/(1-qa))^r$

$= (.9882/(1 + .0118 \times .130))^{168}$

$= 0.10522$

Variance of demand in lead time with demand autocorrelation,

$$\sigma_{LD}^{2'} = \mu_D^2 \sigma_L^2 + \frac{\mu_L \sigma_D^2}{(1-a)^2} - \frac{2a\sigma_D^2(1-pgf_L(a))}{(1-a)^2(1-a^2)}$$

$$= 1^2 \times 1.4226^2 + \frac{2 \times 1.5811^2}{(1 + .130)^2} - \frac{2 \times (-.130) \times 1.5811^2 (1-.10522)}{(1 + .130)^2 \quad (1 - .130^2)}$$

$= 6.4024$

and $\sigma_{LD}' = 2.5303$

Value of k with demand autocorrelation,

$$k' = \frac{(\mu_L \mu_D)^2}{\sigma_{LD}^{2'}}$$

$$= \frac{(2 \times 1)^2}{6.4024}$$

$= 0.6247$

Actual standardised reorder level provided,

$$U' = \frac{\text{Reorder Level calculated by system}}{\sigma_{LD}'}$$

$$= \frac{6.5709}{2.5303}$$

$$= 2.5969$$

PLS provided from k' and U' (using equations (18) and (19)),

$PLS' = 0.0870$

$$\text{SL provided} = \left(1 - PLS' \times \frac{\mu_L}{I}\right)$$

$$= (1 - .0870/2)$$

$$= \underline{0.9565 \text{ (or } 95.65\%)}$$

In order to assess the system effect of demand autocorrelation across a Distribution Centre it is necessary to carry out these calculations for numerous selected parameter values, assign weighting factors and aggregate the results. The variables which could directly affect the results are: μ_D, μ_L, σ_D, σ_L, a, the ratio of μ_L/I, and the target service level. It is clearly not feasible to compute expected service levels for all combinations of these, so the following principles were adopted:-

a) The primary variable for the purpose of this exercise is the Autocorrelation Coefficient, a. The five group values in Table 6.12 were used.

b) μ_D and μ_L exert a different influence on the variance of demand in lead time, and the Ray formula corrects only the portion of the standard expression containing μ_L (The other portion, $\mu_D^2 \sigma_L^2$, is unaffected). Also, the sensitivity analysis shows that service level is influenced in quite different ways by varying μ_D and μ_L (Figs. 4.16 and 4.18). It was therefore considered essential to vary μ_D and μ_L independently.

c) σ_D values were determined according to the Variance Law: $\sigma_D^2 = 2.5\ \mu_D^{1 \cdot 5}$. This is justified in Section 4.1.3.

d) σ_L values were determined according to the Variance Law: $\sigma_L^2 = \mu_L^{1\cdot5}$. It is shown in Section 6.4 that this relationship represents the central third of all lead time variances.

e) The target service level was assumed to be 95% for all products. Table 6.13 shows that with all Product Group settings at 95%, 76% of the total annual costed demand is accounted for by products with allocated service levels between 93% and 98%. The spread around the Group setting is therefore not unacceptably great.

Service Level %	% Total Annual Costed Demand
< 80	1.25
80-90	4.02
90-91	2.34
91-92	4.04
92-93	5.32
93-94	10.74
94-95	13.42
95-96	15.45
96-97	18.36
97-98	18.03
98-99	6.66
> 99	0.35

Table 6.13 Spread of Product Service Levels with Product Groups set to 95%

The full results are given in Table 6.14. The weighting factors were derived by first carrying out a multiple enquiry on the stock file to ascertain the proportion of the total annual costed demand accounted for by each μ_D / μ_L combination, then multiplying these by

μ_D (units/week)	μ_L (weeks)	Autocorrelation Coefficient				
		-0.130	+0.066	+0.297	+0.440	+0.590
1	2	.9565 .0091	.9458 .0158	.9238 .0053	.9016 .0103	.8662 .0063
1	4	.9604 .0102	.9430 .0178	.9036 .0059	.8602 .0116	.7856 .0071
1	6	.9624 .0046	.9415 .0081	.8910 .0027	.8307 .0052	.7203 .0032
1	8	.9630 .0014	.9408 .0025	.8838 .0008	.8118 .0016	.6722 .0010
5	2	.9544 .0119	.9472 .0207	.9322 .0069	.9168 .0134	.8911 .0083
5	4	.9570 .0122	.9453 .0212	.9195 .0071	.8906 .0137	.8387 .0085
5	6	.9583 .0118	.9444 .0206	.9120 .0069	.8731 .0134	.7990 .0082
5	8	.9586 .0013	.9440 .0023	.9086 .0008	.8638 .0015	.7734 .0009
10	2	.9535 .0246	.9477 .0429	.9356 .0143	.9230 .0278	.9016 .0172
10	4	.9557 .0209	.9462 .0364	.9255 .0121	.9022 .0236	.8597 .0146
10	6	.9567 .0073	.9456 .0127	.9198 .0042	.8889 .0082	.8287 .0051
10	8	.9568 .0031	.9454 .0055	.9174 .0018	.8824 .0036	.8102 .0022
50	2	.9520 .0233	.9488 .0406	.9420 .0135	.9348 .0263	.9220 .0163
50	4	.9531 .0114	.9479 .0199	.9365 .0066	.9235 .0129	.8990 .0079
50	6	.9537 .0039	.9475 .0068	.9336 .0023	.9168 .0044	.8830 .0027
50	8	.9538 .0017	.9479 .0029	.9328 .0010	.9142 .0019	.8752 .0012
100	2	.9515 .0251	.9491 .0438	.9440 .0146	.8385 .0284	.9287 .0175
100	4	.9524 .0080	.9484 .0139	.9399 .0046	.9301 .0090	.9114 .0056
100	6	.9528 .0012	.9481 .0021	.9378 .0007	.9252 .0014	.8998 .0009
100	8	.9528 .0010	.9482 .0017	.9372 .0006	.9236 .0011	.8940 .0007

Expected Service Level

Weighting Factor

Table 6.14 Expected Service Levels with Demand Autocorrelation the proportion of the total annual costed demand attributable to each Autocorrelation Coefficient (refer Table 6.12) e.g. for cell $\mu_D = 1/\mu_L = 2/a = -.130$, .0468 of the total annual costed demand is accounted for by products with $\mu_D < 2.5$ and $\mu_L < 3$ (the approximate mid points of the calculation points). The weighting factor for a = -.130 is .1942, therefore the composite weighting factor for the cell is .0468 x .1942 = .0091. The multiplicative process to obtain the composite factor is reasonable as there are wide spreads of lead times and demand rates in every Product Group. They can therefore be assumed to be independent of the autocorrelation weightings.

The aggregation exercise was completed by calculating:

$$\sum^{\text{all cells}} (\text{Service Level x Weighting Factor})$$

which evaluates as .9250, or <u>92.50%</u>

The service level degradation due to demand autocorrelation appears to be very significant. It is particularly severe on items with long lead times - a conclusion also reached by Ray. A second finding concurred with Ray is that the impact of positive demand autocorrelation on the variance of lead time demand is very much greater than the opposite influence of negative demand autocorrelation. An independent finding which can be elicited from this study is that the service levels for fast-moving items are much less sensitive to demand autocorrelation that those for slow-moving items.

It the system formulations were refined to take account of demand autocorrelation the result would be a substantial increase in stock levels to counteract the service level degradation caused by ignoring it. The increase in stock investment could be evaluated with a reasonable degree of accuracy by undertaking an exercise similar to the one here carried out for service levels. It is, however, outside the scope of this study as the purpose here is to evaluate the differences between theory and practice for a given system. In this respect order points are set according to a certain theoretical formulation and these order points are themselves the main determinants of the stock levels which ensue. Measured stock levels should not therefore differ from predictions.

6.4 OVERLAPPING LEAD TIME ERRORS

This error came to light only after several months of live operation of the system. It was reported by one of the Stock Controllers as an overstocking condition associated with long lead times. Some preliminary investigations proved that long lead times per se were not the cause of the problem, but rather a high ratio of lead time to order interval i.e. when more than one purchase order is outstanding at the same time.

The nature of the problem and the measures taken to overcome it are described here in depth as it is believed that this is a very common problem which has been inadequately treated in the literature. Indeed the only attempt at a mathematical solution was found in an unpublished paper (90), and this addressed itself to very specific conditions. Most standard textbooks on stock control ignore the problem, though it is relevant to most, if not all, stock control models whenever a buyer has more than one outstanding order on a particular supplier in his file. In the Inventory Management system, this condition will be shown to be the rule rather than the exception.

After some deterministic mathematical solutions had proved inadequate and inappropriate, the approach taken was to derive a correction routine by extensive simulation trials. The following considerations were taken into account when constructing test data and executing the trials:

a) The basic simulation program SIMSIMPLE was selected so that the results would not be clouded by other system complexities. The use of this simplified reorder level model also proves the generality of the problem. As the program operates with a single product, a duration of 200 years was used for each trial with a nominal order interval of 16 days. This produces an average of 3250 orders with around 400 stockouts. A total of over 200 trials were carried out on an IBM 3032 computer.

b) An examination of operational lead time data revealed that orders on a particular supplier regularly overlap but they very rarely overtake. Most commonly, orders placed in different time periods are delivered together. The marked degree of autocorrelation in Table 6.11 supports these findings. The simulation program was modified to prevent orders overtaking by transposing generated lead times where necessary. This process tends to produce runs of long lead times followed by runs of short lead times, which quite accurately resembles the sequence of measured lead times given in Table 6.2. The re-sequencing process has a negligible effect on the mean and standard deviation of the generated lead times, as a lead time would be rejected only if it exceeds the next largest lead time in the available array by more than the order generation interval.

c) The tendency for orders to overlap is clearly greatest when the lead time variability is large. For instance, if the mean lead time is 80% of the order interval there may be no overlap with a fairly constant lead time, whereas a large standard deviation could cause a substantial degree of overlap.

It was found by experimenting with the Variance Law on a sample of 137 suppliers that 34.3% of the variances are well represented by the relationship $\sigma_L^2 = 0.5\mu_L^{1.5}$;

34.3% by $\sigma_L^2 = \mu_L^{1 \cdot 5}$; and the remaining 31.4% by $\sigma_L^2 = 2.5\ \mu_L^{1 \cdot 5}$. For a set of typical lead times this produces the standard deviations shown in Table 6.15.

Mean lead time (μ_L) (days)	Standard Deviation of lead time (σ_L)		
	*C = 0.5	C = 1.0	C = 2.5
3.50	1.81	2.56	4.05
4.83'	2.31	3.26	5.15
7.50	3.20	4.53	7.17
10.16'	4.03	5.69	9.00
15.50	5.52	7.81	12.35
31.50	9.40	13.30	21.02
47.50	12.49	18.09	28.61
63.50	15.91	22.49	35.57

*C = Coefficient of Variance Law

Table 6.15 Sets of Lead Time Variabilities used in Trials (The mean lead time values were chosen so that when 0.5 days are added for the mean review cycle, Risk Period/Order Interval ratios of ¼, ⅓, ½, ⅔, 1, 2, 3, 4 are produced).

d) The following parameters were maintained at constant values throughout, though the correction formula was subsequently checked with independent data:

μ_D = 10 units per week

σ_D = 5 units per week

Target SL = 95%

e) Fixed reorder levels were calculated for each run, as follows:-

i) Calculate PLS from the target SL.

ii) Calculate K_{LD} from the demand and lead time parameters.

iii) Calculate U by an iterative application of equations (18) and (19) (Sect. 3.4) to give the desired PLS for the k_{LD} value.

iv) Convert U to a reorder level by multiplying by σ_{LD}.

The statistical accuracy of the results was first checked by running the same data with six different random number seeds. This produced service levels of 95.33%, 95.28%, 95.38%, 95.57%, 95.03% and 94.94%. These figures have a mean value of 95.26% and a standard deviation of 0.23%. Assuming a Normal distribution of sampling errors, it is 95% certain that any trial result will be within $\pm$ 1.96 x 0.23% = 0.45% of the mean (where 1.96 is the value of the standard Normal deviate at the 95% level).

The first set of trials was carried out to ascertain the magnitude of the problem. The results are presented in Table 6.16 and the achieved service levels are depicted in Fig. 6.10. Before discussing these results, the precise nature of the problem is clarified by examining in depth an abbreviated run of 10 years' duration which was carried out on a microprocessor for expository purposes. The example has a Lead Time/Order Interval ratio of 1 and a variance coefficient of 2.5. The salient facts are given in Table 6.17. Fig. 6.11 examines one of the overlap occurrences which affect the result.

According to the theory, reorder levels (or their ELS equivalents) are set such that there is a fixed probability of their being exceeded by the demand in any lead time. When they are exceeded, the 'correct' amount of potential sales must be lost to provide the desired service level. Thus any demand in lead time

Variance Coefficient		Ratio of Lead Time to Order Interval							
		¼	⅓	½	⅔	1	2	3	4
0.5	Average SL →	.9445*	.9428*	.9445*	.9485	.9622	.9688	.9747	.9798
	% Error ELS →	+11.0	+14.4	+11.0	+3.0	-24.4	-37.6	-49.4	-59.6
	Wt. Factor ELS →	.0008	.0073	.0071	.0398	.1252	.1088	.0395	.0145
	% Error AS →	-	-	-	0	+9.10	+13.28	+24.28	+30.18
	Wt. Factor AS →	.0008	.0046	.0081	.0363	.1338	.1096	.0345	.0153
1.0	Average SL →	.9442*	.9450*	.9491*	.9539	.9644	.9758	.9748	.9810
	% Error ELS →	+11.6	+10.0	+1.8	-7.8	-28.8	-51.6	-49.6	-62.0
	Wt. Factor ELS →	.0008	.0073	.0071	.0398	.1252	.1088	.0395	.0145
	% Error AS →	-	-	-	+2.51	+9.49	+25.95	+31.56	+38.56
	Wt. Factor AS →	.0008	.0046	.0081	.0363	.1338	.1096	.0345	.0153
2.5	Average SL →	.9515*	.9500*	.9619*	.9670	.9713	.9765	.9834	.9853
	% Error ELS →	-3.0	0	-23.8	-34.0	-42.6	-53.0	-66.8	-70.6
	Wt. Factor ELS →	.0007	.0067	.0065	.0365	.1146	.0996	.0361	.0133
	% Error AS →	-	-	-	+11.35	+25.56	+43.44	+69.90	+66.18
	Wt. Factor AS →	.0008	.0042	.0074	.0333	.1225	.1004	.0316	.0140

Note - dashes indicate no correction necessary
* - corrected for overshoot

Table 6.16 Errors in Service Level and Average Stock Value resulting from Overlapping Lead Times (Uncorrected)

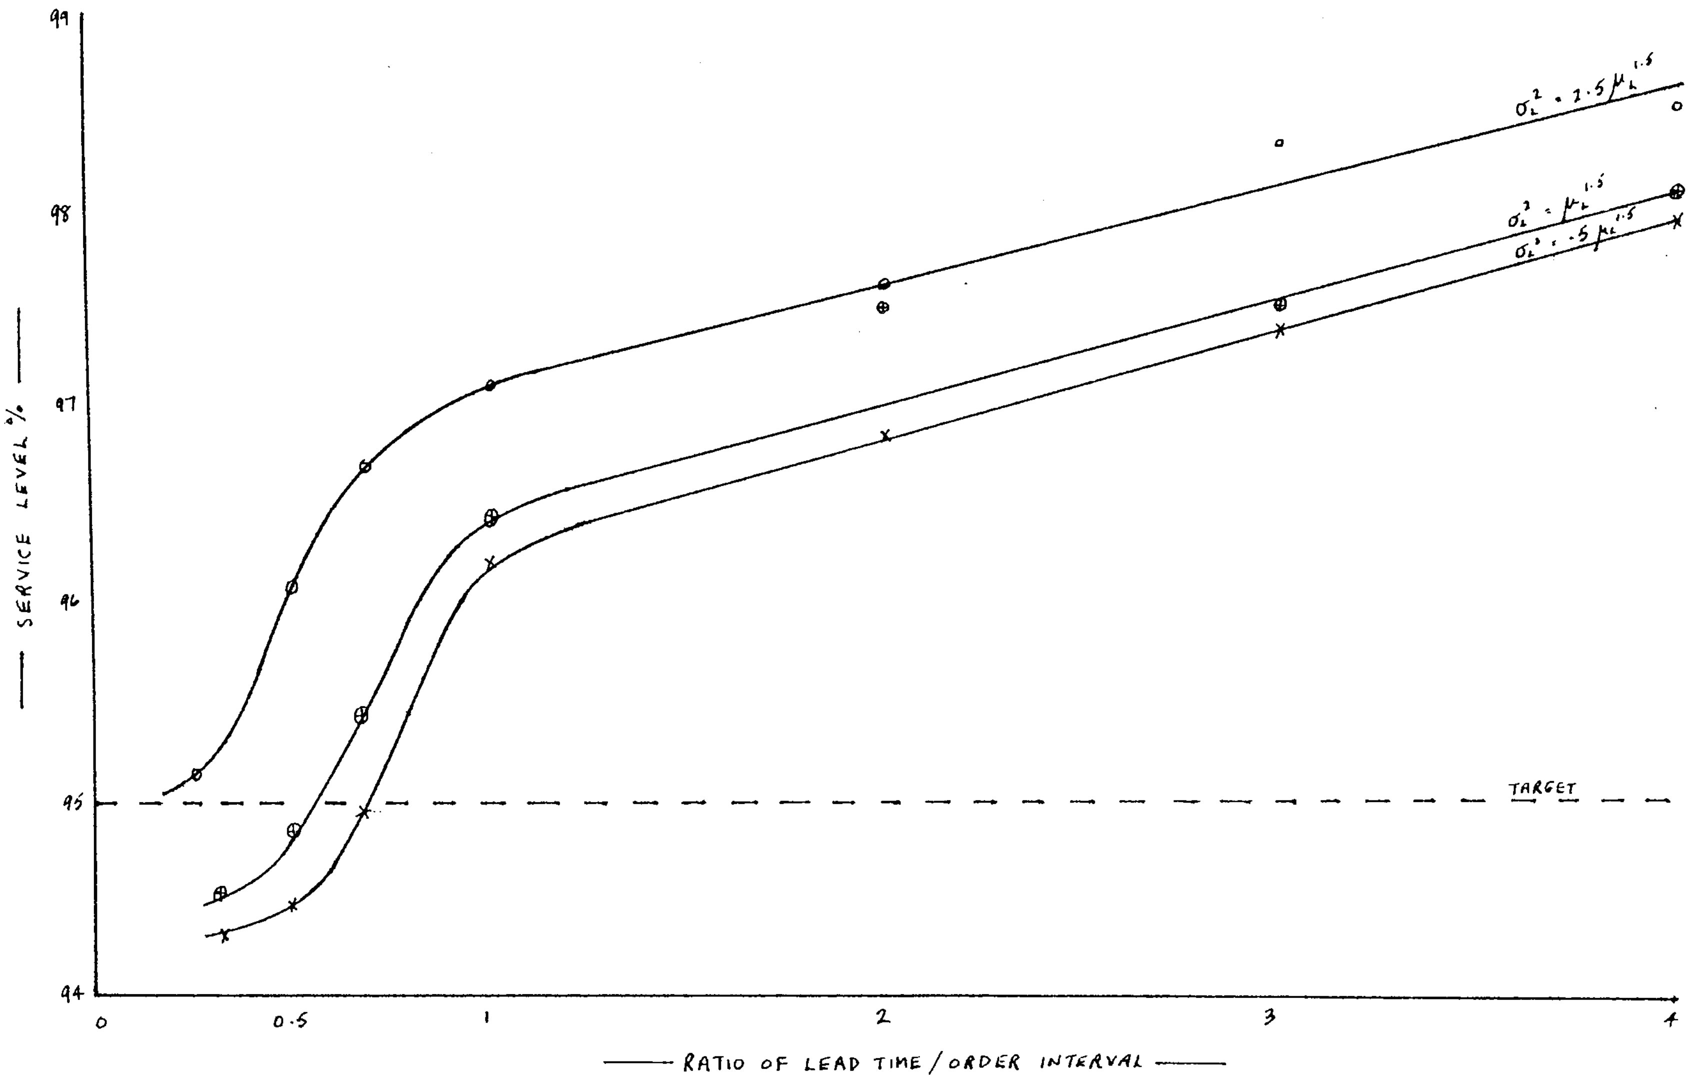

Fig. 6.10 Achieved Service Levels with Varying Lead Time/Order Interval Ratios (Uncorrected)

Order No.	Lead Time (days)	Stock-on-Hand at Placement(a)	Demand in Lead Time (b)	(b)-(a) if (b)>(a)	Actual Lost Sales
1	16	71	39	-	-
2	35	76	56	-	-
6	36	72	79	7	7
20	44	75	96	21	21
38	31	72	85	13	13
39	33	76	82	6	
41	64	76	133	57	57
42	53	69	113	44	
43	44	73	94	21	
53	29	75	85	10	10
66	30	75	80	5	7
67	34	75	82	7	
127	29	76	77	1	1
149	43	75	110	35	39
150	50	76	115	39	
151	48	74	105	31	
161	9	76	8	-	-
			Totals	297	155

Total Demand

Theoretical SL = (5441-297)/5441 = 0.9454

Expected SL = (5441-155)/5441 = 0.9715

Table 6.17 Comparison of Theoretical and Actual Lost Sales with Overlapping Lead Times

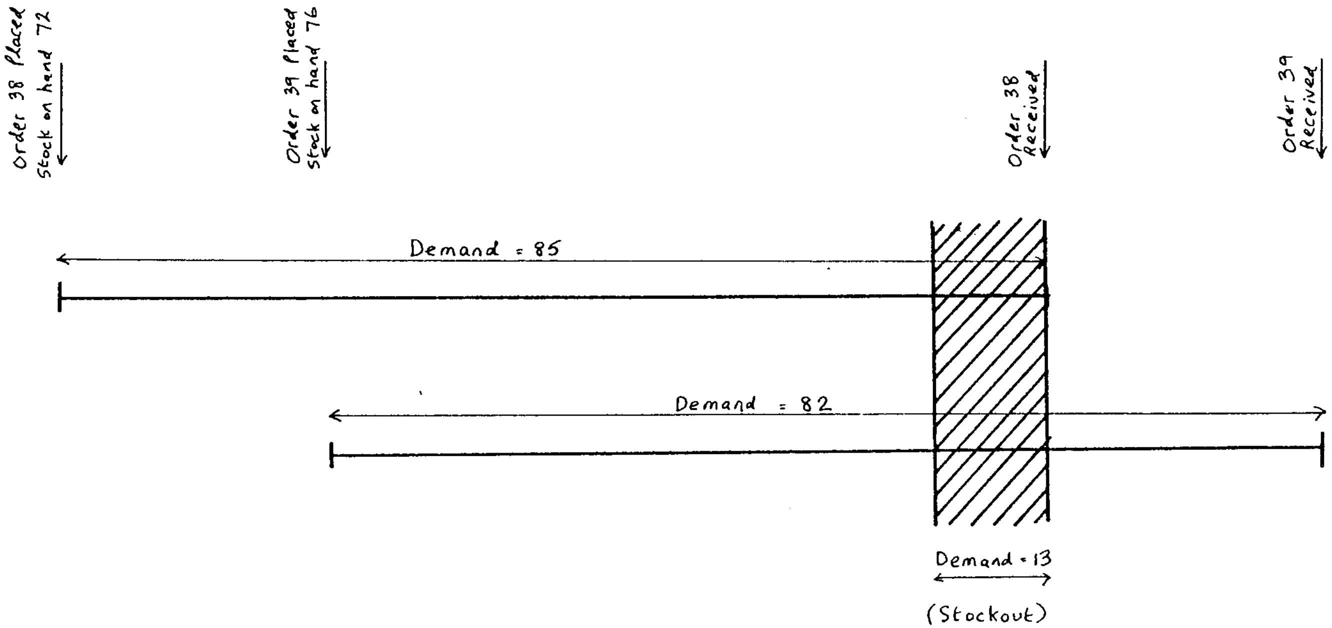

Theoretical Lost Sales = (85 − 72) + (82 − 76) = 19

Actual Lost Sales = 13

Fig. 6.11 Discrepancy Between Theoretical and 'Expected Actual' Lost Sales with Overlapping Lead Times which exceeds the reorder level should incur its full quota of lost sales. Table 6.17 shows four instances where two or more long lead times overlap, and they do not incur sufficient lost sales. The first case is detailed in Fig. 6.11. When order 38 is placed, the stock-on-hand is 72. The demand during its lead time is 85, thus 13 units are unsatisfied. When order 39 is placed, the stock-on-hand is 76. The demand during its lead time is 82, thus 6 units should be unsatisfied. They are not lost, however, as the demand is effectively reduced by the stockout caused by order 38. Thus only 13 units are lost in total instead of the theoretical 19 units. The overall effect of the four instances of overlapping lead times is that only 155 units are lost instead of the theoretical 297. The achieved service level is therefore 97.15% compared with the 94.54% which would have been achieved if the long lead times had each achieved their full share of lost sales.

An important concomitant of the overlapping lead time problem is the incidence of necessarily autocorrelated lead time demands, as successive lead times share some of the same demand. This vitiates the common assumption of a random sequence of lead time demand variates.

Fig. 6.10 indicates the magnitude of service level errors with increasing degrees of overlap. The high variability curve shows that the problem exists to some degree even when the lead time is very much smaller than the order interval. With a ratio of unity there is a severe problem - this illustrates why a deterministic mathematical solution would not be effective, viz. in the deterministic case there would be no overlap and hence no correction would be applied.

Table 6.16 presents the service level and average stock errors in a form suitable for system aggregation. The following principles were applied in constructing the table:

a) A correction for overshoot was applied where the reorder levels are low relative to the maximum daily demand rates. This is necessary as replenishment orders are rarely placed when the stock-on-hand is exactly equal to the reorder level. Large individual demands can deplete the stock to well below the reorder level when an order is generated thus greatly increasing the stockout risk in that cycle. Overshoot is not a problem in the operational system as orders are generated according to the predicted expected lost sales for a <u>family</u> of products.

   The correction in the simulation is based on an expression for average overshoot ($\bar{K}$) given by Lampkin (91):

$$\bar{K} = \tfrac{1}{2}(\mu_D - 1 + \sigma_D^2 / \mu_D)$$

   where μ_D and σ_D are strictly the parameters of individual demand order sizes. However, they apply reasonably well to daily demands, as, for the majority of products in the range there is rarely more than one order in a single day.

   This has the effect of increasing the reorder level by $\bar{K}$. The correction adds approximately 0.5% to the result in most cases (decreasing as the reorder level increases).

b) The weighting factors for ELS represent the proportion of the annual costed demand for the Distribution Centre accounted for by the cell identifying parameters, e.g. 0.0008 of the annual costed demand is accounted for by products with a variance coefficient < 0.75 and a Lead

Time/Order Interval ratio < 0.29 (i.e. halfway between the adjacent parameter values). The weighting factors for average stock are formed in the same way but using costed stock value as the apportionment base.

c) The percentage error in average stock is calculated as:

$$\frac{AS - AS'}{AS'} \times 100\%$$

where AS is the average stock obtained during the trial

and AS' is the average stock which would have been obtained with the corrected reorder level to give a 95% service level (explained later).

The system service level, obtained from

$$\sum^{\text{all cells}} \text{Average SL} \times \text{Weighting Factor for ELS},$$ is 96.84%. This is equivalent to an error of -36.8% in the expected lost sales, which puts the magnitude of the error in better perspective.

The system error in terms of the average stock value is easily derived. It will be shown that the correction for overlapping lead times is uniformly effective. The weighting factors therefore apply to the correct stock levels and not to the levels inflated by the error.

Let total stock value at Distribution Centre = S

$$\text{Then total stock value given error} = \sum^{\text{all cells}} \text{Wt. Factor} \times S \times \left(\frac{100 - \%\ \text{Error}}{100}\right)$$

$$\text{And, system error} = \frac{\sum^{\text{all cells}} \left[\text{Wt. Factor} \times S \times \left(\frac{100 - \%\text{Error}}{100}\right)\right] - S}{S}$$

$$= \sum^{\text{all cells}} \left[\text{Wt. Factor} \times \left(\frac{100-\%\text{Error}}{100} \right) \right] - 1$$

$$= \sum^{\text{all cells}} \text{Wt. Factor} \times \% \text{ Error}$$

This evaluates as +20.86%, which is a measure of the additional stock required to achieve the inflated service level caused by the error. It is worth pointing out that with a perfectly balanced stock, increasing the service level from 95% to 96.84% requires approximately 15% more stock. The increase of almost 21% reflects the fact the additional stock is not being employed to the best financial advantage of the Organisation.

The correction to the overlapping lead time errors was derived as follows:-

a) The reorder levels for each point on Fig. 6.10 were adjusted by trial and error until the achieved service levels were as close as possible to 95% i.e. a further adjustment of 1 unit would cause a larger error in the opposite direction. These individual reorder levels (R') were regarded as effecting the 'perfect' correction insofar as this is possible within the constraints of the exercise.

b) Corresponding standardised reorder levels were computed :

$$U = \frac{R'}{\sigma_{LD}}$$

c) The related PLS values (PLS') for U' and k_{LD} were obtained from equations (18) and (19).

d) The ratio PLS'/PLS was plotted against the ratio of Lead Time/Order Interval for each of the variance coefficients (C) (Fig. 6.12(a)). A linear approximation appeared

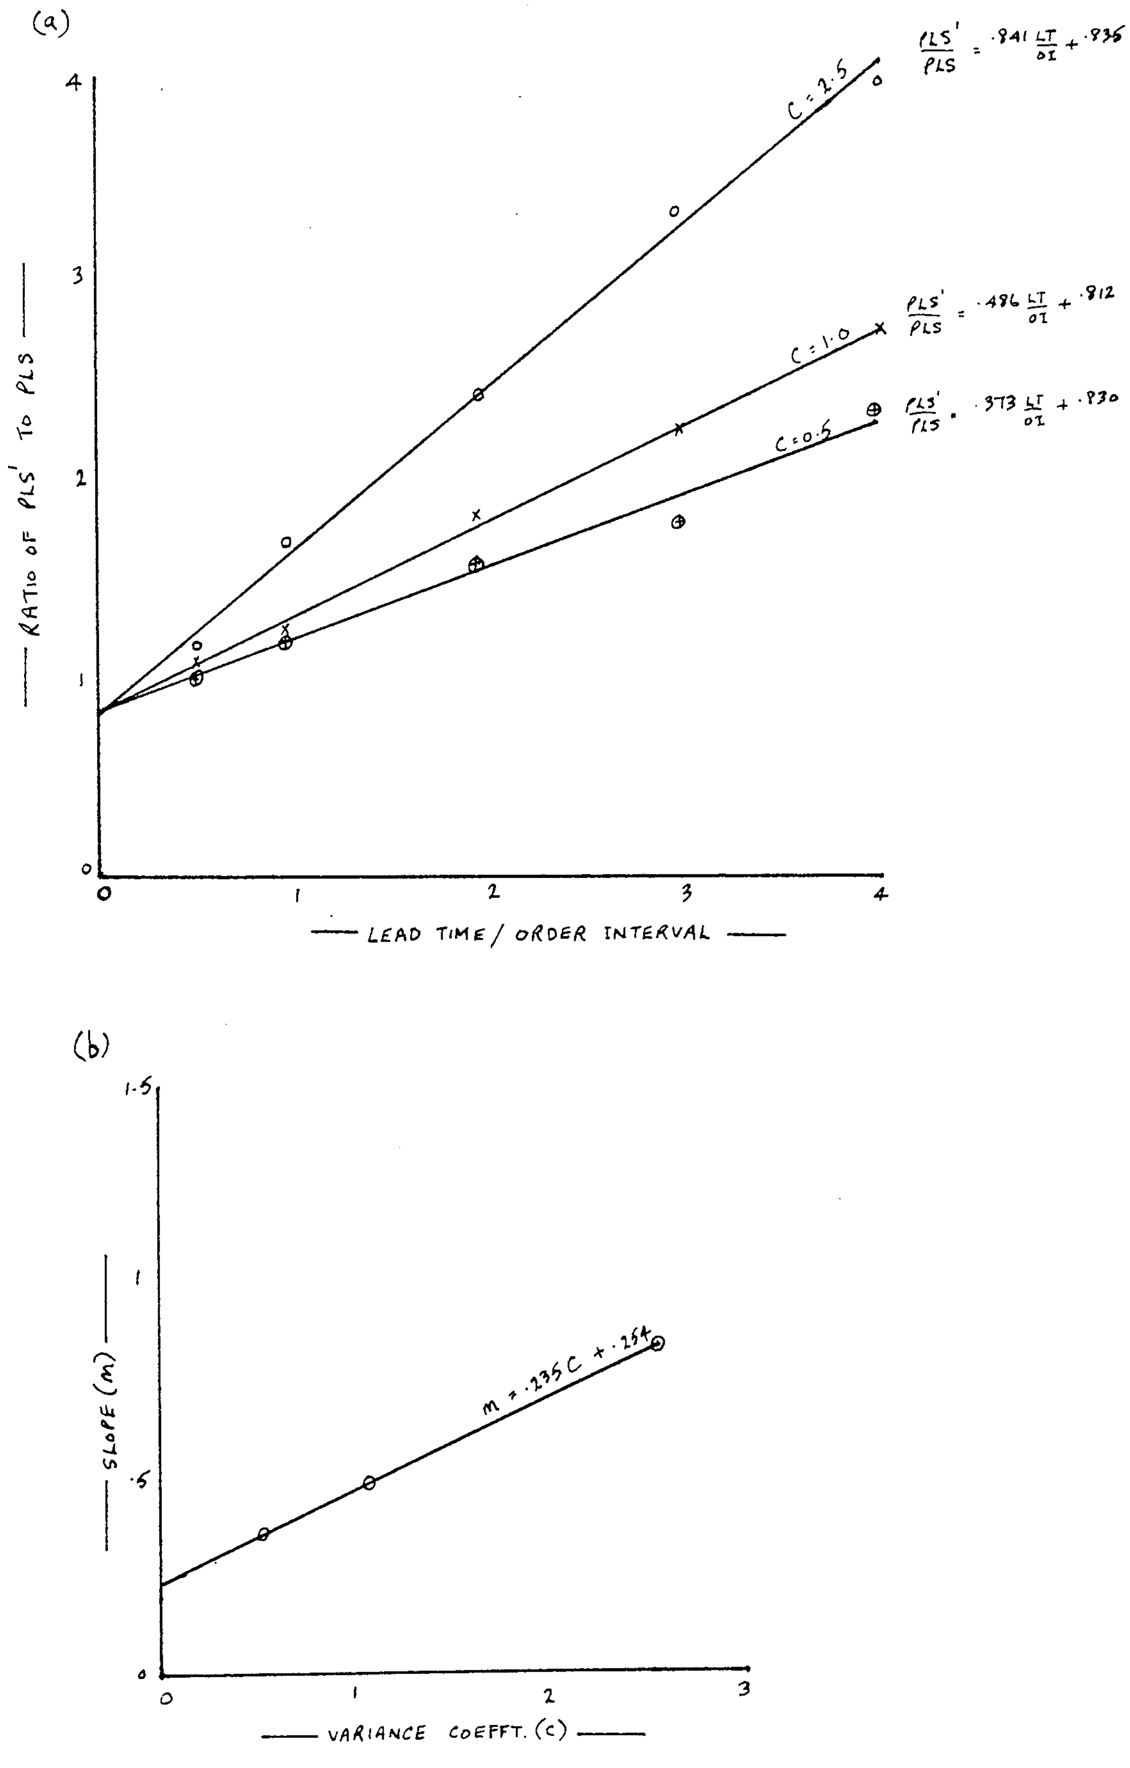

Fig. 6.12 Graphs to Derive Correction for PLS

reasonable in all cases, therefore a linear regression was carried out to obtain the best straight-line fits indicated. The intercepts on the $PLS^{'}/PLS$ axis were found to be very consistent and the average value of 0.826 was taken to apply to all three.

The slopes of the lines (m) were plotted against their variance coefficients (Fig. 6.12(b)). A linear approximation again seemed justified and a linear regression obtained the fit,

$$m = 0.235C + 0.254$$

But, $$\frac{PLS^{'}}{PLS} = m\,\frac{(\mu_L + 0.5)}{I} + 0.826$$

and $$\sigma_L^{2} = C\,\mu_L^{1.5}$$

Substituting for m and C and transposing,

$$PLS^{''} = PLS\left[\left(0.235\,\frac{\sigma_L^{2}}{\mu_L^{1.5}} + 0.254\right)\frac{(\mu_L + 0.5)}{I} + 0.826\right]$$

where $PLS^{''}$ is the PLS value produced by the formula, which may be slightly different to the $PLS^{'}$ value used to derive the formula.

In the operational system the PLS values used in the computation of the order points are replaced by $PLS^{''}$. In the simulation the $PLS^{''}$ values were converted back to reorder levels ($R^{''}$) which in many cases were the same as $R^{'}$ to the nearest integer. The results are given in Table 6.18 and the average service levels are shown graphically in Fig. 6.13.

By repeating the procedures illustrated earlier to assess the total system effects, with the correction incorporated the expected system service level evaluates as 95.08% (ELS error -1.6%), and the

Variance Coefficient		Ratio of Lead Time to Order Interval							
		¼	⅓	½	⅔	1	2	3	4
0.5	Average SL →	.9445*	.9428*	.9445*	.9493	.9470	.9526	.9488	.9533
	% Error ELS →	+11.0	+14.4	+11.0	+1.4	+6.0	-5.2	+2.4	-6.6
	Wt. Factor ELS →	.0008	.0073	.0071	.0398	.1252	.1088	.0395	.0145
	% Error AS →	-	-	-	0	0	0	+2.23	+0.99
	Wt. Factor AS →	.0008	.0046	.0081	.0363	.1338	.1096	.0345	.0153
1.0	Average SL →	.9442*	.9450*	.9491*	.9477	.9515	.9553	.9559	.9528
	% Error ELS →	+11.6	+10.0	+1.8	+4.6	-3.0	-10.6	-11.8	-5.6
	Wt. Factor ELS →	.0008	.0073	.0071	.0398	.1252	.1088	.0395	.0145
	% Error AS →	-	-	-	0	+0.30	+3.44	+3.13	+1.32
	Wt. Factor AS →	.0008	.0046	.0081	.0363	.1338	.1096	.0345	.0153
2.5	Average SL →	.9515*	.9500*	.9619*	.9527	.9491	.9520	.9473	.9447
	% Error ELS →	-3.0	0	-23.8	-5.4	+1.8	-4.0	+5.4	+10.6
	Wt. Factor ELS →	.0007	.0067	.0065	.0365	.1146	.0996	.0361	.0133
	% Error AS →	-	-	-	-0.52	+0.16	0	+3.96	-3.16
	Wt. Factor AS →	.0008	.0042	.0074	.0333	.1225	.1004	.0316	.0140

Note - dashes indicate no correction necessary

* - corrected for overshoot

Table 6.18 Errors in Service Level and Average Stock Value resulting from Overlapping Lead Times (Corrected)

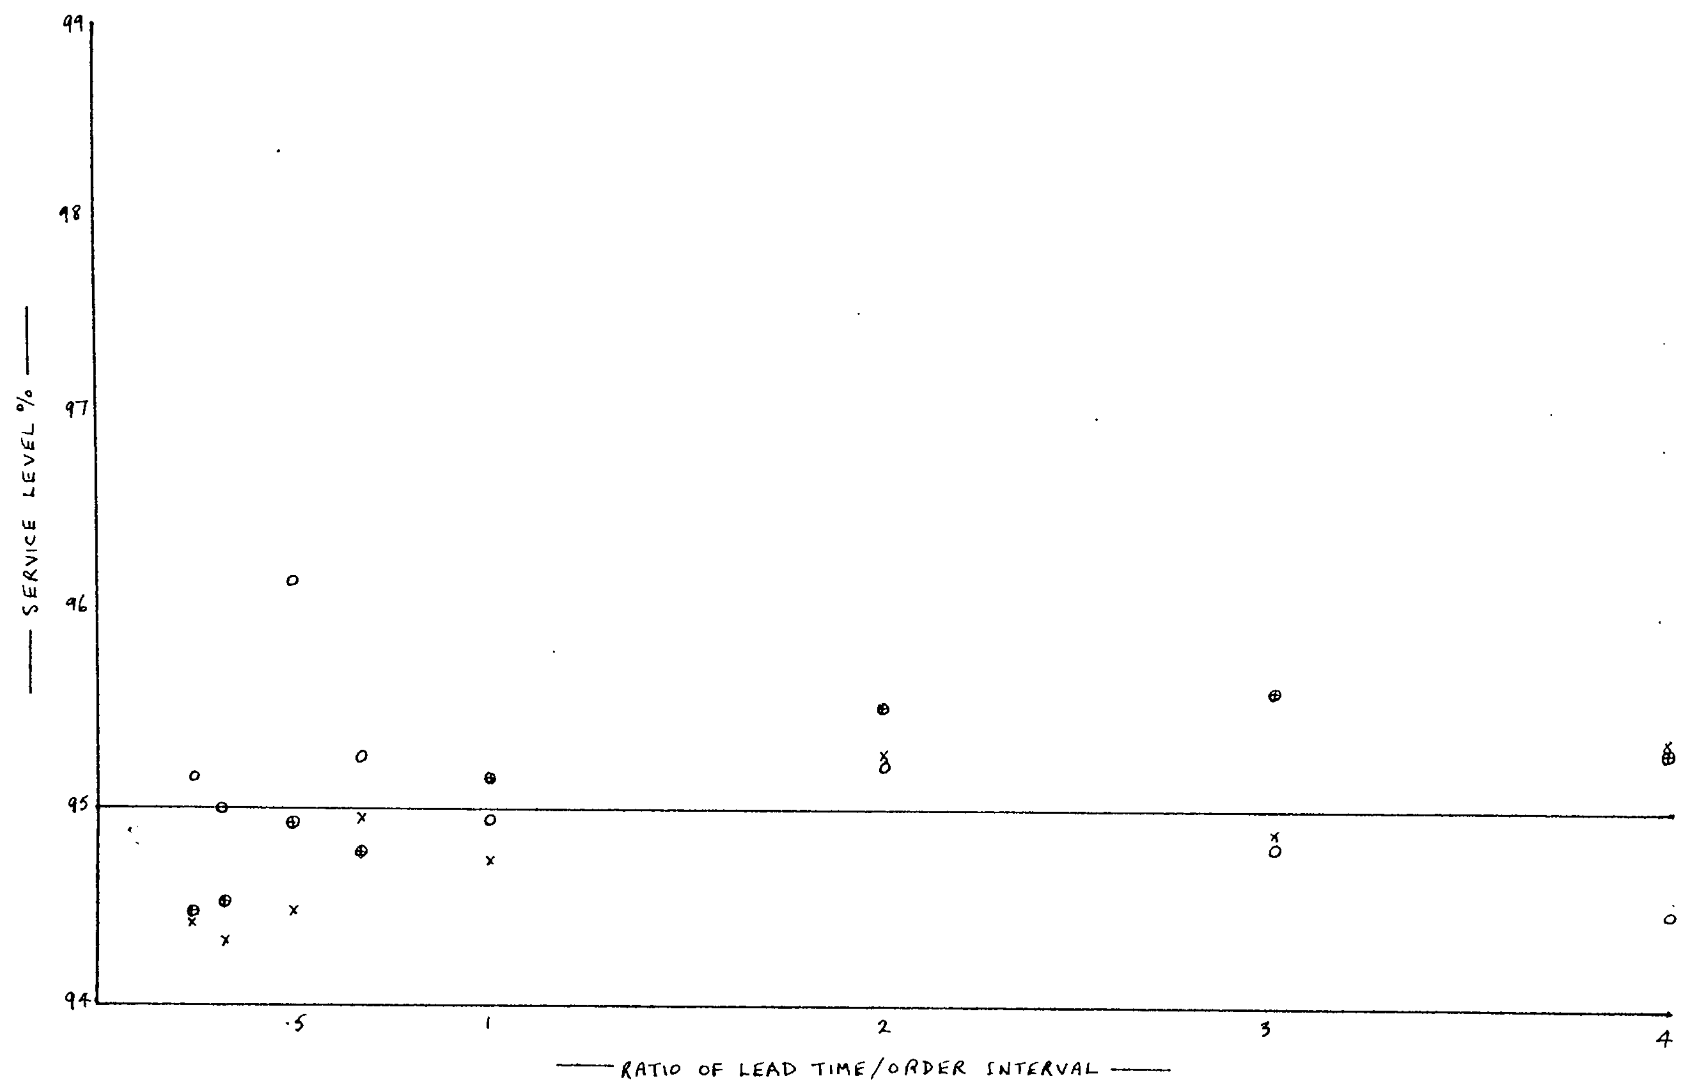

Fig. 6.13 Achieved Service Levels with Varying Lead Time/Order Interval Ratios (Corrected)

increase in average stock is 0.72%.

The correction was the result of several months' work after the fault had been reported. A temporary less effective correction was incorporated whilst the simulation work described was undertaken. For the purpose of this study, it should be assumed that the fault was present in the results reported in Tables 3.2 and 3.3, and the reported service levels and average stock values are inflated by the error.

6.5 INCONSISTENCY IN SERVICE LEVEL DEFINITIONS

As stated in Section 3.3, 'Service Level' is defined as "the percentage of costed annual demand which is satisfied ex-stock". The expression used in the theoretical formulations (eqn. 5) is:

$$SL = \frac{\text{Costed Forecast Demand p.a. - Expected Lost Sales p.a.}}{\text{Costed Forecast Demand p.a.}}$$

As demand which is not satisfied ex-stock is assumed to be lost, this expression is consistent with the verbal definition. (It is assumed throughout that there is no unexpressed demand i.e. all demand results in orders).

It is not feasible to measure achieved service level the same way as it is set i.e. by expressing satisfied costed orders as a percentage of total costed orders. The problem is due to the Branch replenishment philosophy. This system is based on a Reorder Cycle policy where the stock is reviewed weekly on a rota basis (one-fifth of the range is reviewed each day), and topped up to a Maximum Stock Level. If the Distribution Centre is out of stock of an item, there is no backordering procedure - it is simply reordered the following week. Therefore, assuming no replenishment:

$$\text{Deficiency at end of week } n = \text{Deficiency at end of week } n-1 + \text{Sales for week } n$$

It can be seen that the deficiency at the end of week n-1 is ordered twice - once at the end of week n-1 and again as part of the order at the end of week n. Hence, service level measured on this basis would be understated.

Achieved service level is actually measured as:

$$\sum^{\text{all products}} \text{Item Cost} \times \left(\frac{\text{Days in Stock}}{\text{Days in Stock + Days out of Stock}} \right)$$

The two measures are clearly not identical for the following reasons:-

a) If an item is out of stock but there are no orders for it, the measured service level will fall progressively each day, whereas the defined service level will not be affected.

b) As the state of the bins is only reviewed daily, some inaccuracy is introduced e.g. an item could go out of stock during a day, fail to satisfy demands, and be replenished the same day without the computer being aware of the stockout.

Notwithstanding these differences, if the assumption is made that the demand for an item is the same when it is out of stock as when it is in stock, it is reasonable to suppose that the proportion of demands satisfied will be the same as the proportion of time in stock. The measures are therefore equivalent apart from the precision of measurement and the chance factor. Lampkin, in a personal communication with the author, agrees with this conclusion. He writes:-

> "As far as I can see, under most reasonable assumptions, the expected proportion of time 'in stock' is equal to the expected proportion of demands met 'ex-stock'. It follows that an unbiassed estimate of the one is an unbiassed estimate of the other."

In order to check this conclusion empirically, 15 sets of results from simulation program FT01 were analysed by counting the days in and out of stock and comparing the service level so derived with the reported service level (which is calculated as the percentage of costed demand satisfied ex-stock). The results are shown in Table 6.19, the target service level being set at 90% in all cases.

Run	SL from Lost Sales	SL from Days In/Out
1	88.0	87.8
2	89.6	91.7
3	89.6	88.1
4	86.8	89.5
5	92.0	88.2
6	88.8	88.6
7	92.4	92.3
8	86.4	89.7
9	88.0	90.1
10	91.2	89.8
11	92.8	92.0
12	90.0	91.3
13	92.0	90.7
14	92.8	92.3
15	88.4	89.3
AVE	89.9	90.1

Table 6.19 Comparison of Service Levels Using Alternative Definitions

It can be seen that:-

a) the average service levels are quite consistent, and

b) the direction of the error is about equal (9 in one direction; 6 in the other).

It is therefore concluded that the errors introduced into the operational system due to an inconsistency in the way service levels are set and measured are negligible.

6.6 ERRORS DUE TO APPROXIMATION IN AVERAGE STOCK FORMULA

Ignoring the negligible effect of the review cycle, the expression used in the system for average stock is:

$R - \mu_D \mu_L + \mu_D I/2$ (refer Section 3.3, equation 4)

This is the standard textbook formula which is derived with reference to Fig. 6.14 (a).

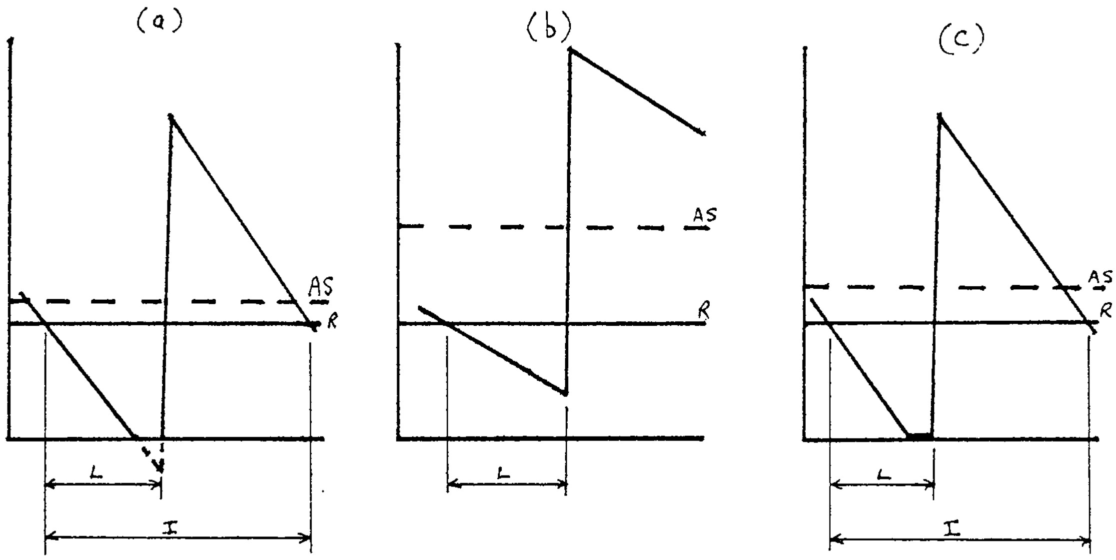

Fig. 6.14 Alternative Formulations for Average Stock

Stock at order placement = R

Let demand in lead time = x with p.d.f. f(x)

Stock just before order arrives = R - x

Stock just after order arrives = $R - x + \mu_D I$

$$\therefore \text{Average Stock} = \int_0^\infty \tfrac{1}{2}\left[(R-x)+(R-x+\mu_D I)\right] f(x)\,dx$$

$$\text{But,} \quad \int_0^\infty x\, f(x)\, dx = \mu_D \mu_L$$

$$\therefore \text{Average Stock} = R - \mu_D\mu_L + \mu_D I/2$$

This assumes that when a stockout occurs the physical stock is negative, therefore it is an underestimate of the true average stock.

Fig. 6.14 (b) represents only that part of the probability distribution where $x \leq R$.

Average stock just before order arrives

$$= \int_0^R (R-x)\ f(x)\ dx$$

$$= \int_0^\infty (R-x)\ f(x)\ dx - \int_R^\infty (R-x)\ f(x)\ dx$$

$$= \int_0^\infty R\ f(x)\ dx - \int_0^\infty x\ f(x)\ dx - \int_R^\infty (R-x)\ f(x)\ dx$$

$$= R - \mu_D\mu_L + \int_R^\infty (x-R)\ f(x)\ dx$$

But, $\int_R^\infty (x-R)\ f(x)\ dx$ = ELS per order cycle

∴ Average stock just before order arrives = $R - \mu_D\mu_L$ + ELS

And, average stock just after order arrives

$$= R - \mu_D\mu_L + ELS + \mu_D\, I$$

∴ Average Stock over cycle = $R - \mu_D\mu_L + \mu_D\, I/2$ + ELS

As this encompasses only the incidences of lead time demand which are insufficient to cause a stockout, this expression for average stock is an overestimate.

The true value of average stock would derive from the profile represented in Fig. 6.14 (c). Clearly the assumption used so far that the average over the cycle is 'half the sum of the expected physical just before arrival and the expected physical just after arrival' is no longer valid as the rate of depletion is not constant. This makes an analytical solution of the true average mathematically intractable unless simplifying assumptions are made. Barrington Taylor and Oke (92) had to make such simplifying assumptions when dealing with the backorders case which is certainly simpler than the lost sales case. Their approach of calculating the upper and lower bounds of the average stock is followed here. For the upper bound, the expression described as an overestimate is used, and for the lower bound, the underestimate. An example is first worked through using the data values used for the example in Section 6.3.3:

μ_D = 1 unit/week

σ_D = 1.5811 units/week

μ_L = 2 weeks

σ_L = 1.4226 weeks

I = 4 weeks

SL = 95%

The reorder level (R) is calculated as 6.5709 (refer Section 6.3.3). The ELS per order cycle is 5% of the mean demand per order cycle $= 0.05\ \mu_D\ I$.

$\therefore$ Lower bound of Average Stock $= R - \mu_D\mu_L + \mu_D\ I/2$

$= 6.57 - (1 \times 2) + (1 \times 4)/2$

$= 6.57$ units

And, Upper bound = Lower bound + ELS

$= 6.57 + (0.05 \times 1 \times 4)$

$= 6.77$ units

% Increase over Lower bound $= \dfrac{(6.77 - 6.57)}{6.57} \times 100$

$= \underline{3.04\%}$

The aggregation technique is similar to that described in Section 6.3.3. In this case, the weighting factors represent the proportion of the total costed average stock attributable to the respective μ_D / μ_L combinations. The results are presented in Table 6.20.

The percentage by which the upper bound exceeds the lower bound for the Distribution Centre as a whole is calculated as:

$$(3.04 \times 0.582) + (2.20 \times .0914) + \ldots\ldots\ldots(2.79 \times .0042)$$

$$= \underline{3.55\%}$$

It is therefore concluded that the formula used for calculating average stock in the system understates the true value by $< 3.55\%$.

μ_D (units/week)	μ_L (weeks)	Av. Stock A (underestimate)	Av. Stock B (overestimate)	$\frac{(B - A)}{A} \times 100$	Weighting Factor
1	2	6.57	6.77	3.04	.0582
1	4	9.11	9.31	2.20	.0914
1	6	11.12	11.32	1.80	.0336
1	8	13.14	13.34	1.52	.0297
5	2	23.88	24.88	4.19	.0563
5	4	32.54	33.54	3.07	.0735
5	6	39.85	40.85	2.51	.0467
5	8	47.79	48.79	2.09	.0175
10	2	43.46	45.46	4.60	.0856
10	4	58.75	60.75	3.40	.1021
10	6	71.95	73.95	2.80	.0545
10	8	86.77	88.77	2.30	.0270
50	2	188.26	198.26	5.31	.0699
50	4	250.65	260.65	3.99	.0635
50	6	306.72	316.72	3.26	.0256
50	8	373.39	383.39	2.68	.0150
100	2	362.62	382.62	5.52	.1010
100	4	480.49	500.49	4.16	.0375
100	6	587.72	607.72	3.40	.0072
100	8	717.45	737.45	2.79	.0042

Table 6.20 Upper and Lower Bounds of Average Stock

6.7 STOCK BALANCING CONSIDERATIONS

The principle of stock balancing to optimise the objective function has been described in Section 3.3. It can be seen from Fig. 3.5 that when the Lagrange Multiplier (λ) achieves a stable value, each constituent product in the Group will have been assigned an implicit service level. These individual service levels are strongly related to the k_{LD} values of the products. As an example, the group of five products used on the buyers' training course (Section 7.2) are examined. Table 6.21 shows the assigned service levels.

Product Code	μ_D (units/week)	MAD_D (units/week)	k_{LD}	Theoretical SL (%)	Net Cost (£)
A	8.00	3.44	2.02	90.6	3.959
B	130.00	32.54	2.54	92.1	1.950
C	38.50	11.56	2.40	91.7	1.173
D	74.75	19.97	2.49	91.9	1.497
E	7.00	5.88	1.08	85.0	23.754

Table 6.21 Relationship Between Assigned Service Levels and k Values

The salient points are:-

a) In general, fast-moving lines are assigned the highest service levels.

b) Product E, which has a particularly erratic selling pattern, is heavily penalised.

These results are entirely consistent with the theoretical objectives of the system, in that products with a steady selling pattern would require smaller buffer stocks (expressed as weeks' demand) to sustain a given service level. Tautologically, they would generate more sales from a given value of stock - hence they are favoured by the allocation process.

Arguments are often advanced that such algorithms fail to account for important value judgements. For example, whatever service level is assigned to bath plugs will certainly not be high enough - it would be unthinkable to have an available supply of baths but no plugs. The interaction of buyers with the algorithm is the central theme of this work, which will be examined in depth in Chapter 7. For the present, the effects of product interdependence on the theory are investigated without the complication of buyer interference.

Program FT01 (Section 5.1) was used to run 100-week simulation trials on the five products listed in Table 6.21. The operating parameters were:

SL - 90% (for the Group as a whole)

μ_L - 10.64 days

σ_L - 6.22 days

I - 10 days

A summary of ten sets of results using different random number seeds to generate the demand and lead time data is presented in Table 6.22.

Run	Product A	Product B	Product C	Product D	Product E	Group
1	84.36	93.12	89.95	89.67	88.96	90.65
2	89.57	93.34	90.75	89.85	88.32	90.89
3	88.56	92.30	89.70	90.11	86.85	89.79
4	89.90	94.65	91.01	93.34	85.91	91.29
5	94.02	92.64	90.77	88.53	88.64	90.62
6	92.42	93.52	93.87	90.14	84.20	90.05
7	90.52	93.92	93.98	91.61	86.91	91.23
8	87.77	91.71	88.12	89.53	88.21	89.76
9	91.25	93.86	91.80	90.48	85.43	90.42
10	95.26	93.85	92.07	91.16	83.64	90.70
MEAN	90.36	93.29	91.20	90.44	86.71	90.54
S.D.	3.15	0.87	1.82	1.33	1.88	0.54
EXPECTED	90.6	92.1	91.7	91.9	85.0	90.0

Table 6.22 Individual Service Levels Obtaining from Simulation Trials

The main findings are:-

1) The service level allocations are broadly in line with the theoretical expectations.

2) The group results are much more consistent than those for the individual products. There are two reasons for this:

   a) The cancellation effect of aggregating + and - individual product errors.

   b) Replenishment orders are triggered according to the group status. Hence individual products will be above or below their implicit reorder levels when orders are generated. Their individual results can therefore be expected to vary much more.

3) As a corollary of 2 b), the dominant product in the group (B, which accounts for over 40% of the group costed demand) is subject to less variation than the other products. This is because it has a greater influence on order generation.

It is concluded that all products will achieve their theoretical service levels but the non-influential products (low value, slow-moving) will experience much larger swings than if they were controlled independently. Such products can experience prolonged stockouts whilst they are awaiting order generation by the influential products; or, conversely, they can be amply stocked for substantial periods if the influential products exceed their forecast sales.

The average group costed stock value from the above trials is £1,621.86 (standard deviation £51.28). The theoretical average costed stock using equation (4) in Section 3.3 to support a service level of 90.54% with this group is £1,558.89. The surplus of 4.0% is barely outside the upper bound of average stock identified in Section 6.6. It is therefore concluded that the process of stock balancing does not have any unexpected effects on the average stock calculations.

6.8 FORECASTING ERRORS

During the system design stage, an analysis of demand over a 13-week period indicated that the pattern was reasonably horizontal, and random influences were the main source of variation. With no trend or data pattern and random incidences of variates, it can be shown that the historical Mean is the best estimate of the next occurrence. It is also clear that an exponentially-weighted moving average with a low smoothing factor constitutes a good approximation of the historical Mean. Therefore, the best evidence available at the time suggested that first-order exponential smoothing would be a suitable technique for demand forecasting. This also had the advantages that it is easy to comprehend, it is inexpensive to compute, it requires very little data storage, it can be initiated without an extensive data collection exercise, and the sensitivity to recent events can be altered by a simple parameter change without the necessity to re-compile computer programs.

The main forecasting technique is supplemented by a de-seasonalising/re-seasonalising procedure using tables of seasonal factors only for products which are subject to seasonal influences. Demand filtering of extreme values is also incorporated.

Exponential smoothing was also selected for forecasting lead times, though with less reasoned justification. In this case the manual records were of questionable reliability (especially apropos part-deliveries) and the method was selected more by default through having insufficient information than for its positive suitability. Certainly there is strong evidence arising from this study which favours a different approach.

With the benefit of hindsight, it seems likely that more elaborate techniques would have produced better forecasts for both demand and lead time, particularly after the end of 1979 when the economic recession vitiated the demand stationarity assumption. However, it is considered that whilst more accurate forecasting would have reduced uncertainty and thus enabled the same service level to have been obtained with a smaller stock investment, the forecasting errors do not contribute to the discrepancies between predicted and achieved performance, provided the parameters of the probability distribution are stationary and the demand in lead time variates occur in a random sequence. (If the parameters of the distribution are not stationary, the effects have already been evaluated in Section 4.1.4.5; if the variates do not occur randomly, the effects have been evaluated in Section 6.3). This is because the forecasting errors (MADs) determine the reorder levels, which in turn determine the subsequent level of actual stocks. With the same level of forecasting errors, the actual stocks could therefore be expected to be consistent with the predictions and the target service levels should result.

It is therefore concluded that inaccuracies in the forecasting method did not contribute to the performance variations other than through the factors already evaluated.

CHAPTER 7

ANALYSIS OF PERFORMANCE OF HUMAN OPERATIVES

7.1 ALLOCATION OF FUNCTION IN THE INVENTORY MANAGEMENT SYSTEM

The design principles in the Inventory Management system are strongly orientated towards buyer freedom of action, with the computer taking over the repetitive high volume computational tasks. The main considerations were:-

a) Personal flair is the dominant characteristic of trading activity in the industry. Sales promotions are instigated at relatively short notice at local or national level according to such factors as market buoyancy, competitor activity and local events, which are not normally subject to prediction by automatic methods.

b) Buyers have a close and regular contact with suppliers and they are often in a position to influence lead times (e.g. by switching to an alternative supplier). In certain cases they are certainly better equipped to estimate lead times than even the best forecasting system.

c) Many building materials are subject to fluctuating prices (e.g. copper-based products react to the international copper market). There is therefore sound reason for buying when the price is advantageous. Similarly, suppliers often offer special deals according to their own trading strategies.

d) Impending shortages can sometimes be anticipated by overseeing the general industrial scene. For instance, the manufacture of radiators is extremely vulnerable to shortages of oxygen for welding. Reported industrial disputes which may affect the delivery of oxygen therefore make it prudent to obtain extra supplies of radiators. As with many building products, radiator shortages prejudice the sales of associated products such as boilers and copper tube - hence it is extremely important to take prompt action to prevent or mitigate the effects of long-term shortages.

In most cases buyer initiatives should be self-financing. For example, if large stocks are purchased for a sales promotion, the resultant additional profit should outweigh the extra stockholding costs (including a realistic allowance for excess stocks of products which do not meet the promotional expectations). In other cases, buyer initiatives should reduce the detrimental effects of unavoidable shortages. For these reasons it was considered prudent not to impose system constraints on instinctive purchasing decisions.

The system philosophy is that the computer effects the routine regulation of stocks by calculating order points from all of the pertinent variables according to service level targets set by management. When orders are generated they are printed on remote terminals in the form of 'recommended orders'. Together with each recommended order are:-

i) an 'alternative order' which is made up to a 'best terms' value input by the buyer only if the recommended order falls short of that value, and

ii) the state of all of the variables used to calculate the order quantities.

On receipt of the recommended order, the buyer has the following options:-

a) accept the recommended order, in which case it is placed, or

b) accept the alternative order, in which case it is placed, or

c) change any quantities on the recommended order, or

d) ignore the recommendations altogether.

If the recommendations are ignored, then the order will reappear every day with increasing quantities. The buyer then has the option of suppressing printing until he elects to re-introduce it. If the buyer reduces order quantities, this will advance the next order, and vice versa. At any time during the operation of the system the buyer has the following additional options. He can:-

e) request an order without it being due,

f) override the forecast lead time with his own estimate. If this is done the computer will use the override for all calculations until an accompanying 'horizon' input by the Buyer has expired - then it will revert to the forecast (updated if there have been any receipts during the horizon),

g) change the service levels for individual products from those allocated by the stock balancing function. The same facility is available at Buying Family level.

It can be seen that buyers can use the computer to varying degrees depending upon their individual preferences and their management's guidelines. In practice, approximately 60% of the recommended orders are placed unchanged, and in a further 10% of the cases the alternative orders are placed. Approximately 30% of the orders are therefore manually adjusted. It follows that as nearly all of the routine tasks of recording orders, matching receipts, calculating demand rates, calculating lead times, calculating order points, etc., are now performed by the computer, much more of the buyers' time is now spent on more cognitive functions. More effort can now be devoted to product ranging, progress chasing, supplier liaison, monitoring of excess stocks, etc.

7.2 BUYER TRAINING COURSES

Before the system was introduced into each of the four Distribution Centres, the respective buyers attended a 2-day training course at the computer centre. This comprised a short introductory session to explain the general principles of the system, followed by a participatory simulation exercise using the FT01 simulation program (refer Section 5.1).

The buyers were formed into teams, each team comprising two or three buyers. In total, eleven teams have been through the exercise. Each team was asked not to divulge the nature of the exercise to following teams, therefore it is assumed that each team started on an equal footing.

The product data consisted of 'live' data for five typical products which are purchased from the same supplier. The Product Data sheet given to each team member is summarised in Table 7.1.

Product	μ_D (units/week)	MAD_D (units/week)	Net Cost per Unit (£)	Opening Stock (units)
A	8.00	3.44 (F.Steady)	3.959	68
B	130.00	32.54 (Steady)	1.950	1002
C	38.50	11.56 (Steady)	1.173	192
D	74.75	19.97 (Steady)	1.497	374
E	7.00	5.88 (Erratic)	23.754	75

μ_L (all products) = 10.64 days

MAD_L = 4.68 days (fairly steady)

I = 10 days

Table 7.1 Product Data Sheet for Buyers' Training Course

The opening stock figures refer to the counted stocks of the products at one of the Distribution Centres immediately prior to the commencement of the first course. All other figures were derived from clerical records for the same Distribution Centre.

Each team member was also presented with an instruction sheet explaining the objectives of the exercise and the rules. An abridged version of the sheet is given in Table 7.2. This was fully explained and the differences between the simulation program and the operational system were pointed out, viz.:-

a) The simulation handles only five products, whereas Buying Families contain up to 200 products.

b) Part-deliveries are not catered for in the simulation.

c) Alternative orders are not included in the simulation.

d) Service level adjustments are not possible in the simulation.

The override options and the procedures for effecting them were explained in detail. It was stressed that when an override is made the program is run through to the end of the 100 weeks and the repercussions are not completely predictable. As the performance measurements are taken over the whole period, chance plays some part in the outcome.

A target service level of 90% was used throughout. In practice, when the program is allowed to run through without interruption the achieved service levels are between 89.5% and 92.0% in over 90% of the trials. It is now known that the positive error is mainly attributable to overlapping lead times (refer Fig. 6.9 with a Lead Time/Order Interval ratio of c 1.0 and a low variability lead time). Each course was initiated by an uninterrupted run through the 100 weeks, and the outcome results were taken as the datum for the team.

INVENTORY MANAGEMENT SIMULATION EXERCISE

Introduction

The exercise is intended to give you some practice in operating the computerised Inventory Management system which will be introduced shortly. It takes the form of a computer programme which faithfully 'simulates' the new system. This means that you can practice using the new facilities before you are required to do so in real life (much like an aircraft flight simulator). Please read this instruction sheet carefully before starting.

Overview

We will simulate the purchasing, receipts and issues of five products over a 100-week period. The products are all obtained from the same supplier and the details are given on the attached Product Data Sheet which we will explain to you. Throughout the exercise you will be provided with three regular pieces of information:-

1. Recommended Purchase Orders with order quantities on the day they are generated by the computer.
2. Notification of receipts.
3. A weekly stock report giving the latest demand and lead time forecasts as well as the physical stock and on-order balances.

You can assume that the data has been fed into the computer accurately, so the information given to you is error-free.

A mock 'supplier' will be available at all times to provide you with lead time estimates. The lead times are rearranged from their generated sequence by the computer so that they come in 'runs'. The supplier can anticipate the rearrangement fairly well, but he himself cannot be absolutely certain of particular lead times, and his uncertainty is progressively greater for lead times further in advance. This should be a realistic representation of the reliability of the information you would get from a real supplier.

The demand data is known in advance and you will be warned of periods of high demand. You can regard these as sales promotions and they will almost certainly cause stockouts if you take no action.

Objective

Your objective is to obtain a better result than if you merely follow the computer recommendations all of the time. This can be achieved in two ways (or a combination of them):-

1. Achieve a higher service level without increasing the average stock value.
2. Achieve the same service level with a lower average stock value.

Overrides

You can intervene in four possible ways:-

1. Change recommended order quantities.
2. Raise a manual order.
3. Suppress the generation of recommended orders.
4. Override the computer's lead time forecast.

You are allowed a total of six overrides in the 100 weeks and you can use any combination of the four types. After you apply an override the computer will project the effects through the remaining part of the 100 weeks - so chance plays some part in the outcome.

Table 7.2 Instructions for Buyers' Training Course

In all, four different random number seeds were used to generate the data, hence the datum levels are not quite the same for all teams. In order to assess the profitability implications of different service level/average stock value results, an additional Control Index, denoted 'Net Profit p.a.' was computed as:

$$\sum^{\text{All Products}} (\text{Demand p.a.} \times \text{Unit Cost} \times \text{Service Level} \times \text{Gross Margin}) - (\text{Average Stock} \times \text{Unit Cost} \times \text{Stockholding Fraction})$$

where Gross Margin (as defined in Section 3.3) = 0.30

and Stockholding Fraction (i.e. the annual cost of holding £1 stock capital) = 0.20

After each override was effected the following measurements were recorded for the five products combined:

Achieved service level

Average stock value

Average stockturn

Additional net profit p.a. (over starting value)

The full results for all teams are given in Table 7.3. It should be noted that:-

a) the result for a given override reflects the compound effect of that override and all previous overrides, and

b) the objectives (Table 7.2) imply a policy constraint which has a bearing on the selection of overrides e.g. it is not permissible to double the stock investment even if this is justifiable on the grounds of improved profitability.

c) the result of each override was conveyed back to the team and discussed with them. Hence the learning process was reinforced by continuous feedback. After each override, the Controllers of the exercise (the author and two

Team	Override No.	Service Level (%)	Stock Value (£)	Stockturn	Additional 'Net Profit p.a.' (£)
1	-	90.1	1985	14.2	-
	1 ↓	89.7	1949	14.4	(29.31)
	2 *	90.6	2009	14.1	40.84
	3 ↓	89.9	1876	15.0	3.54
	4	90.2	1886	15.0	28.93
	5	90.4	1862	15.2	51.98
2	-	90.1	1985	14.2	-
	1	90.5	2020	14.1	29.51
	2 *	92.0	2062	14.0	158.02
	3 ↓	91.9	1964	14.7	168.50
	4 ↑	94.8	2157	13.8	394.60
	5 ↓	93.0	1981	14.7	265.50
3	-	90.1	1985	14.2	-
	1	90.4	1967	14.4	30.98
	2 ↓	88.9	1857	15.0	(83.93)
	3 *	89.5	1861	15.1	(29.97)
	4 ↓	89.0	1857	15.0	(74.81)
	5 ↑	93.7	2167	13.6	292.19
	6	92.7	1943	15.0	245.72
4	-	90.1	1985	14.2	-
	1	90.1	1995	14.2	(2.00)
	2 ↓	91.0	1786	16.0	121.95
	3 *	90.4	1817	15.6	60.98
	4 ↑	93.0	1945	15.0	272.70
	5 ↓	92.8	1847	15.8	274.05
5	-	90.1	1985	14.2	-
	1	90.1	1983	14.2	0.40
	2 ↓	89.2	1774	15.8	(39.95)
	3 *	90.9	1767	16.1	116.62
	4 ↑	93.9	2095	14.1	324.85
6	-	90.1	1985	14.2	-
	1	90.9	1978	14.4	74.42
	2 ↓	89.5	1778	15.8	(13.37)
	3 *	91.3	1817	15.8	143.13
	4 ↑	94.4	1924	15.4	404.69
	5 ↓	92.7	1792	16.2	275.92
7	-	90.4	1646	17.3	-
	1	89.5	1632	17.3	(79.35)
	2 ↑	90.5	1775	16.1	(16.67)
	3 ↓	90.7	1746	16.4	7.38
	4	92.3	1788	16.2	145.03
	5	92.5	1782	16.3	164.48
	6	92.8	1763	16.6	195.67
	7 ↓	91.9	1724	16.8	121.32
	8 ↓	92.2	1697	17.1	154.10

Team	Override No.	Service Level (%)	Stock Value (£)	Stockturn	Additional 'Net Profit p.a.' (£)
8	-	90.4	1646	17.3	-
	1	90.6	1648	17.3	17.86
	2	90.1	1681	16.9	(34.38)
	3	91.2	1696	16.9	63.02
	4	91.8	1789	16.2	99.19
	5	91.9	1835	15.8	99.12
	6	93.3	1796	16.4	234.70
	7	92.7	1769	16.5	185.34
	8 ↓	92.1	1733	16.7	137.77
9	-	90.9	1570	18.2	-
	1	91.8	1609	17.9	74.35
	2 ↓	91.3	1551	18.5	40.31
	3 *	92.9	1558	18.7	184.95
	4 ↑	92.1	1590	18.2	105.53
	5	92.3	1624	17.9	116.99
	6 *	92.6	1528	19.0	163.57
	7 ↑	93.2	1688	17.3	186.34
10	-	89.8	1640	17.6	-
	1	88.8	1578	18.1	(78.88)
	2 *	90.8	1577	18.5	103.87
	3	93.3	1621	18.5	323.26
	4 *	94.5	1672	18.2	422.60
	5 ↓	93.8	1544	19.6	384.30
	6	93.9	1557	19.4	390.83
	7 ↓	93.5	1523	19.8	361.12
	8	94.0	1531	19.8	405.16
11	-	89.8	1640	17.6	-
	1	89.8	1640	17.6	0
	2 *	90.7	1588	18.4	92.54
	3 *	94.5	1692	18.0	418.60
	4 ↑	94.9	1710	17.9	451.51
	5	94.4	1736	17.5	400.67
	6 ↓	94.4	1564	19.4	435.07
	7	94.7	1583	19.3	458.65

* Extra quantities purchased for sales promotion

↓ Lead time forecast decreased

↑ Lead time forecast increased

Table 7.3 Performance Results from Buyers' Training Course colleagues who were involved in the design of the system) ran off a fresh audit report for the whole of the 100-week period. Appendix 5 comprises an extract of this report from the end of Week 2 to the end of Week 6. At the end of Week 2 the forecasts (EWMAs) of demand and MAD of demand are revised and the critical ELS (i.e. the order point) is re-calculated. At the end of Week 3 Day 2, the ELS calculated from the status of the physical stocks exceeds the order point ELS. This generates Order No. 1 which is placed the following day. On Week 5 Day 5 the ELS based on the physical stock plus on-order status exceeds the critical value, therefore Order No. 2 is placed on Week 6 Day 1. The generated lead time for Order No. 1 is 16 days, therefore the receipt is recorded at the start of business (by convention) on Week 6 Day 4. The quantities for Order No. 1 are transferred into the physical, and Order No. 2 is shifted into the first order slot. It can be seen that this report provides sufficient detail to trace through the repercussions of any override to give a full explanation to the team.

The performance of Team 4, which is selected for typicality and expository convenience, is examined in detail before discussing the results in general.

Override 1 involved increasing the quantities of two of the products on Order No. 1 because it was considered that the computer recommendations were insufficient to meet their preconceived notion of buffer stock requirements. The original audit print revealed that these products would have remained in stock without the override. As the computer compensated for the manual intervention by decreasing the quantities on the next order by the same amounts, the average stock over the period was affected only marginally. The override was therefore unnecessary but not particularly detrimental.

Override 2 comprised the application of a 'supplier' lead time estimate of 5 days in place of the computer EWMA of 11 days. The override was maintained over a period of 6 weeks, and in the event the estimate proved fairly accurate. The order quantities were reduced substantially and a considerable reduction in the average stock value ensued. The service level improvement was the chance outcome of changing the time phasing of the order generation relative to the demands.

Override 3 consisted of substantially increasing a recommended order quantity of product E (a costly slow-mover) in anticipation of a sales promotion. Inadequate communication with the 'supplier', however, resulted in the consignment arriving too late to prevent or mitigate the effects of the stockout. Excess stock was carried for 8 weeks which added significantly to the average stock value. The decline in the service level was due to a reversal of the phase change caused by Override 2.

Override 4 involved applying a 'supplier' lead time estimate of 15 days over a 12-week period when the computer EWMA was initially 7 days. This proved well-founded and a number of potential stockouts were averted. This increased the service level by 2.6% for a modest increase in the average stock value.

Override 5 consisted of another lead time change. On this occasion the computer EWMA of 13 days was replaced by a 'supplier' estimate of 6 days. In the event the average lead time proved to be 9 days and the team were somewhat fortunate in that the lost sales increased only marginally. The average stock value, as expected, fell significantly.

The first override tried by nearly all of the teams was to adjust quantities on the first recommended order. The buyers from all of the Distribution Centres were of the uniform opinion that the computer was under-ordering the fast-moving products and over-ordering the slow-moving product with the erratic selling pattern (product E). This attitude was conditioned by the standard industry practice of ordering up to a set number of weeks' cover. The stock cover calculated by the system is, however, quite sensitive to the standard deviation of demand (Fig. 4.22), and the number of weeks' sales this represents is far from constant across the product range. Hence the buffer stock expressed as a number of weeks' sales is much greater for product E than for the other products.

The early overrides were also characterised by minor quantity increases or the placement of small manual orders to try to prevent or alleviate impending stockouts. In the introductory session, buyers showed a marked reluctance to accept the concept of expected lost sales. The notion of planning to lose sales was considered by most buyers to be an anathema - hence they tended to act to prevent every threatened stockout. Most teams recognised early in the exercise that minor changes to order quantities were having a negligible effect on system performance, and that they were unlikely to make significant gains by this means. Some teams, however, persisted with this type of override (e.g. Team 8 used their first five overrides to change order quantities), and were very reluctant to allow the computer to exercise homeostatic control in normal circumstances. The limitation on the number of allowable overrides eventually forced the teams to try more adventurous methods, and it was strongly emphasised by the Controllers that tampering with order quantities in routine circumstances would be an unproductive utilisation of their time when using the operational system.

The tactic of buying in extra stocks for the purpose of servicing notified sales promotions was used to good effect by most of the teams. The overrides marked with an asterisk on Table 7.3 were all utilised for this purpose. It can be seen that additional purchasing to meet anticipated needs which are known to the buyer but not to the computer is a very effective instrument for raising service levels above the target values. This can be achieved with a minimal stock investment penalty provided the timing of the additional purchasing is correct and the stocks are consumed in the volumes anticipated.

In the latter stages of the exercise most of the overrides took the form of lead time adjustments (denoted by directional arrows in Table 7.3). It is clearly evident that by reducing the amount of uncertainty in the forecasts appreciable performance improvements can be made. It can be seen that when the lead time forecasts are increased into line with the expected levels, substantial service level improvements can be made at the expense of a very modest increase in stock investment; whereas when the forecasts are decreased into line with the expected levels, substantial cuts in stock investment are possible for only a marginal sacrifice in service level.

In general, all of the teams performed quite creditably and completed the exercise on a better performance level than the unaided computer (though Teams 5, 7, 8 and 9 did not finish within the specified guidelines). Some teams quickly learned that routine homeostatic regulation is very adequately performed by the computer, and manual intervention is only necessary where the buyer has access to more information than the computer. Other teams were noticeably slower to grasp this and showed a marked reluctance to allow the computer to take over even the most repetitive tasks.

The training task was considered to be successfully accomplished in a limited timescale. Fig. 7.1 depicts the improvement in team performance as the exercise progressed. It is very important to recognise that the apparent decline in performance on the last override by Teams 2, 3, 6 and 8 was due in every case to an attempt to drop the stock investment artifically to a level within the specified guidelines. Less damaging measures would undoubtedly have been taken if the exercise had not been terminated at that point.

An analysis of Table 7.3 reveals that in terms of profitability, as defined, the most beneficial override is to increase the lead time forecast when the EWMA is below the best estimate. The nine applications of this produced an average increase in 'Net Profit' of £145.93. Conversely, the 16 applications of a downward lead time adjustment produced an average reduction in 'Net Profit' of £18.99. This measure is, however, unlikely to reduce profits in the operational system as:-

a) the service level was set at only 90% in the simulation (in order to concentrate more stockout events into the exercise). At this level improvements in service tend to benefit profitability more than reducing stock investment

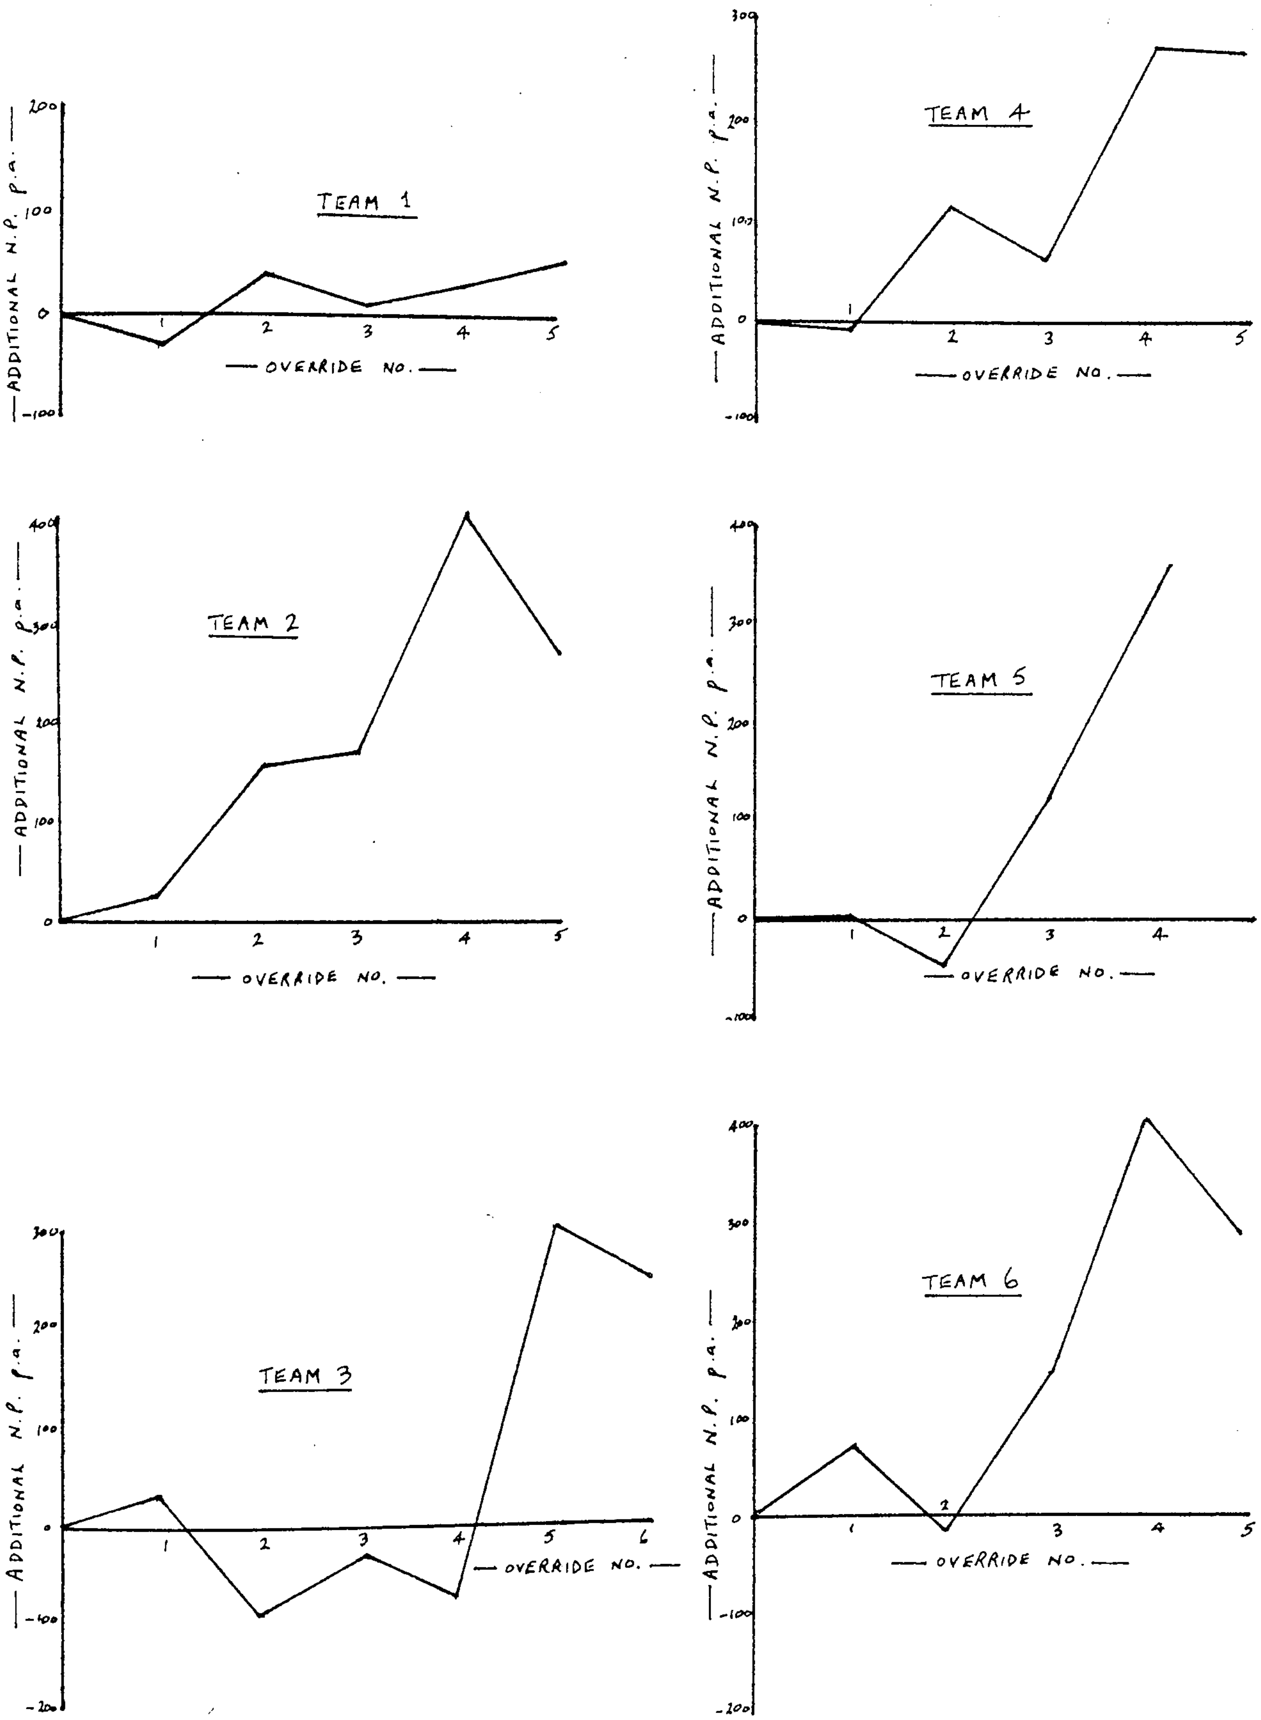

Fig. 7.1 Profitability Index after Successive Overrides

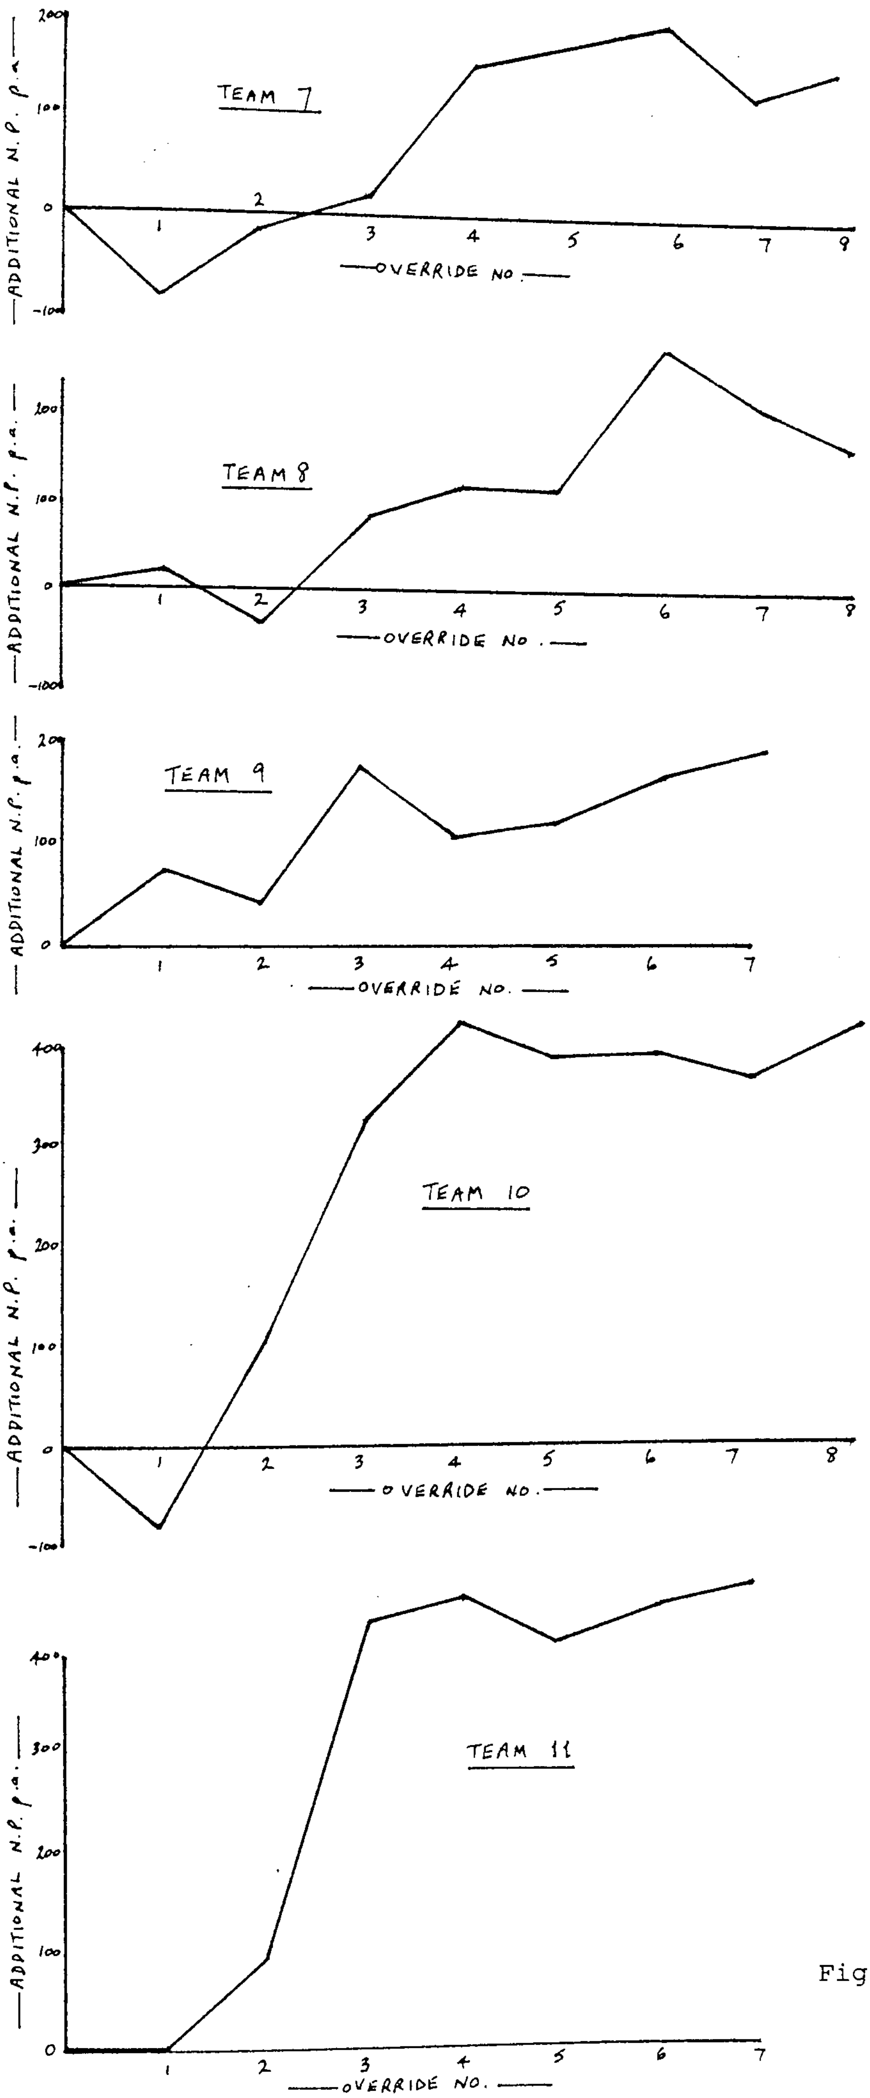
TEAM 7
ADDITIONAL N.P. p.a.
200
100
0
-100
1 2 3 4 5 6 7 8
OVERRIDE NO.
TEAM 8
ADDITIONAL N.P. p.a.
200
100
0
-100
1 2 3 4 5 6 7 8
OVERRIDE NO.
TEAM 9
ADDITIONAL N.P. p.a.
200
100
0
1 2 3 4 5 6 7
OVERRIDE NO.
TEAM 10
ADDITIONAL N.P. p.a.
400
300
200
100
0
-100
1 2 3 4 5 6 7 8
OVERRIDE NO.
TEAM 11
ADDITIONAL N.P. p.a.
400
300
200
100
0
1 2 3 4 5 6 7
OVERRIDE NO.

Fig. 7.1 (Cont'd)

does, and

b) the products selected for the exercise have high stock-turns relative to most other products in the range of a Distribution Centre - hence the numeric value of the average stock is small compared with the annual demand. The 'Net Profit' formula presented earlier in this Section clearly shows that under these conditions small movements in service level can outweigh much larger movements in stock investment.

The Controllers also observed a non-linear relationship between the strength of lead time overrides and their effects. For instance, if the computer EWMA is 10 days and the best estimate is 20 days, an intermediate override of 15 days produces far more than half of the potential benefits. This was observed to apply to overrides in either direction - hence quite conservative actions are likely to produce disproportionately large rewards. This phenomenon could not be quantified, but it was checked on a number of occasions after the original observation had been made.

Table 7.3 also reveals that the 12 overrides for promotional buying produced an average increase in 'Net Profit' of £116.39. Considering the teams were notified of only one or two promotions (depending on the chance clustering of demand data), this is also a highly effective means of improving profitability. The remaining 27 overrides which were applied for miscellaneous reasons produced an average gain in 'Net Profit' of £26.88. This reinforces the earlier conclusion that small gains can be made by making minor adjustments, but when 7,000 or more products are being monitored, the time could in all probability be employed more advantageously elsewhere.

The buyers themselves felt that, on the whole, the exercise simulated their working environment with a high degree of fidelity. In particular, they considered that the demand generation and lead time generation and re-sequencing routines produced realistic data. They were, however, of the opinion that the lead time information was more reliable than they would expect to get in practice, and that promotional sales rarely materialise as planned. At the conclusion of the course all of the buyers recognised the advantages of computer assistance in the buying and stock control functions, and there was no perceived inbuilt resistance to the introduction of the system.

7.3 EFFECTS ON PERFORMANCE OF LEAD TIME OVERRIDES

The buyers' training courses provided consistent objective evidence that the computer performance predictions could be improved upon in practice by the judicious use of market intelligence. The computer predictions are based upon the forecasting errors (MADs) which have occurred in the past, with the implicit assumption that they will continue at the same level in the future. If the buyers can provide better estimates than the computer EWMAs, then the forecasting errors which ensue will be of a lower order than those provided for. Hence the buffer stocks will overprovide, and the resulting service levels will exceed the target levels. Tautologically, the target service levels could be obtained with reduced buffer stocks.

Table 7.3 indicates that improving lead time forecasts is a particularly effective way of improving performance. Substantial service level gains can be made for modest increases in stock, or, conversely, sizeable stock reductions can be achieved at the expense of only marginal service level deterioration. These observations are now studied by analytical methods to remove the random effects inherent in the simulation.

To avoid compound effects, a continuous review, fixed reorder level, fixed quantity system will be used for a single independent product with a stationary lead time distribution. Assume the last 12 lead times for supplier PO in Table 6.2 are drawn from a population with the same parameters:

.........14,11,11,34,16,18,8,13,10,10,3,5.......

Then, μ_L = 12.75 days; σ_L = 7.58 days

And assume μ_D = 5 units/week; σ_D = 3 units/week;

Target SL = 95%; I = 30 days

Then, $PLS = (1-0.95) \times \dfrac{30}{12.75}$

$= 0.117647$

$$\sigma_{LD} = \sqrt{\frac{12.75}{5} \times 3^2 + 5^2 \times \frac{7.58^2}{5^2}}$$

$= 8.9670$

$$k_{LD} = \frac{12.75^2 \times 5^2}{5^2 \times 8.9670^2}$$

$= 2.0217$

U_{LD} for k_{LD} = 2.0217 and PLS = 0.117647 is 2.1734

$R = U_{LD} \quad \sigma_{LD}$ (Sect. 3.4 equation 13)

$= 2.1734 \times 8.9670$

$= 19.49$ units

AS (Sect. 3.3 equation 4 with T = 0)

$$= 19.49 - \frac{12.75 \times 5}{5} + \frac{30 \times 5}{2 \times 5}$$

$= 21.74$ units

The performance expectation in each individual order cycle is now examined by calculating posterior service levels and average stock expected values given the occurrence of the respective lead times. This is equivalent to assuming a fixed lead time. Weatherburn (93) proves that for a fixed lead time, L, and a Gamma-distributed demand with modulus k, the demand in lead time assumes a Gamma distribution with modulus L k.

Considering the first lead time in the set (of 14 days):

$$k_D = \frac{\mu_D^2}{\sigma_D^2}$$

$$= \frac{5^2}{3^2}$$

$$= 2.7778$$

From Weatherburn, $k_{LD} = k_D L$

$$= 2.7778 \times \frac{14}{5}$$

$$= 7.7778$$

$$\sigma_{LD} = \sqrt{\frac{14}{5} \times 3^2 + 0}$$

$$= 5.0200$$

Given the calculated R of 19.49, $U_{LD} = \frac{19.49}{5.0200}$

$$= 3.8825$$

PLS for $k_{LD} = 7.7778$ and $U_{LD} = 3.8825$ is 0.034048

$$SL = 1 - \left(0.034048 \times \frac{14}{30}\right)$$

$$= 0.9841 \text{ (or } 98.41\%)$$

$$AS = 19.49 - \frac{14}{5} \times 5 + \frac{30 \times 5}{2 \times 5}$$

$$= 20.49 \text{ units}$$

Hence the posterior expectation of service level during the order cycle containing the 14-day lead time is 98.41% and the average stock expectation is 20.49 units. By repeating these calculations for the other lead times in the set, the results presented in Table 7.4 are obtained.

Lead Time (days)	Expected SL (%)	Expected Av. Stock (units)
14	98.41	20.49
11	99.56	23.49
11	99.56	23.49
34	51.54	0.49*
16	96.87	18.49
18	94.52	16.49
8	99.92	26.49
13	98.92	21.49
10	99.74	24.49
10	99.74	24.49
3	100.00	31.49
5	99.99	29.49
Av.	94.90	21.74

*The formula used for calculating average stock is unsuitable for low service levels

Table 7.4 Posterior Calculations of Expected Service Level and Average Stock for Individual Lead Times

The process of treating the order cycles as mutually independent introduces some minor inaccuracies e.g. the 34-day lead time would set back the generation of the subsequent order by 4 days. However the effects of these inaccuracies are negligible compared to the service level differences between the order cycles.

Suppose now that the buyer had forewarning of the 34-day lead time, and an override of the exact duration was applied at the correct time. (The effects of imperfect overrides will be examined later). The computer would recalculate the reorder level to give a

95% service level during the period of the override. When such an override is applied, the computer calculates an associated standard deviation to be used with it, in accordance with the Variance Law with a coefficient derived from the historical standard deviation. In this case:

$$\text{Coefficient (C)} = \frac{\sigma_L^{\ 2}}{\mu_L^{\ 1\cdot 5}}$$

$$= \frac{7.58^2}{12.75^{1\cdot 5}}$$

$$= 1.262$$

$$\text{Standard deviation to be used} = \sqrt{1.262 \times 34^{1\cdot 5}}$$

$$= 15.82 \text{ days}$$

$$\sigma_{LD} = \sqrt{\frac{34}{5} \times 3^2 + 5^2 \times \frac{15.82^2}{5^2}}$$

$$= 17.6486$$

$$k_{LD} = \frac{34^2 \times 5^2}{5^2 \times 17.6486^2}$$

$$= 3.7114$$

$$\text{PLS for 95\% SL during override period} = (1-0.95) \times \frac{30}{34}$$

$$= 0.044118$$

U_{LD} for k_{LD} = 3.7114 and PLS = 0.044118 is 3.1755

$$R = 3.1755 \times 17.6486$$

$$= 56.04 \text{ units}$$

A reorder level of 56.04 units would therefore give a prior service level expectation of 95%. Given the incidence of the 34-day lead time with this reorder level the posterior service level and average stock expected values are now calculated.

$$\sigma_{LD} = \sqrt{\frac{34}{5} \times 3^2 + 0}$$

$$= 7.8230$$

$$k_{LD} = k_D \; L$$

$$= 2.7778 \times \frac{34}{5}$$

$$= 18.8890$$

$$U_{LD} = \frac{56.04}{7.8230}$$

$$= 7.1635$$

PLS for $k_{LD} = 18.8890$ and $U_{LD} = 7.1635$ is 0.000853

$$SL = 1 - \left(0.000853 \times \frac{34}{30}\right)$$

$$= 0.9990 \text{ (or } 99.90\%)$$

$$AS = 56.04 - \frac{34}{5} \times 5 + \frac{30 \times 5}{2 \times 5}$$

$$= 37.04$$

Substituting these values into Table 7.4, the average expected service level over the complete period becomes 98.93% and the average stock = 24.79 units. It can be seen that the expected lost sales has been cut to 21.4% of the target value for an increase in stock investment of 14.0%. This can be compared with the stock investment required to give a 98.93% service level without an override facility. This is depicted by the curve in Fig. 7.2, which shows that an average stock of 33.5 units would be required i.e. an increase of 54.1% over the 95% base level. It is clear from Fig. 7.2 that any override which produces a coordinate on the convex side of the curve represents an improvement over the unaided computer. The result of the override is denoted by 'IE' and this clearly represents a substantial improvement. However, repeated overrides

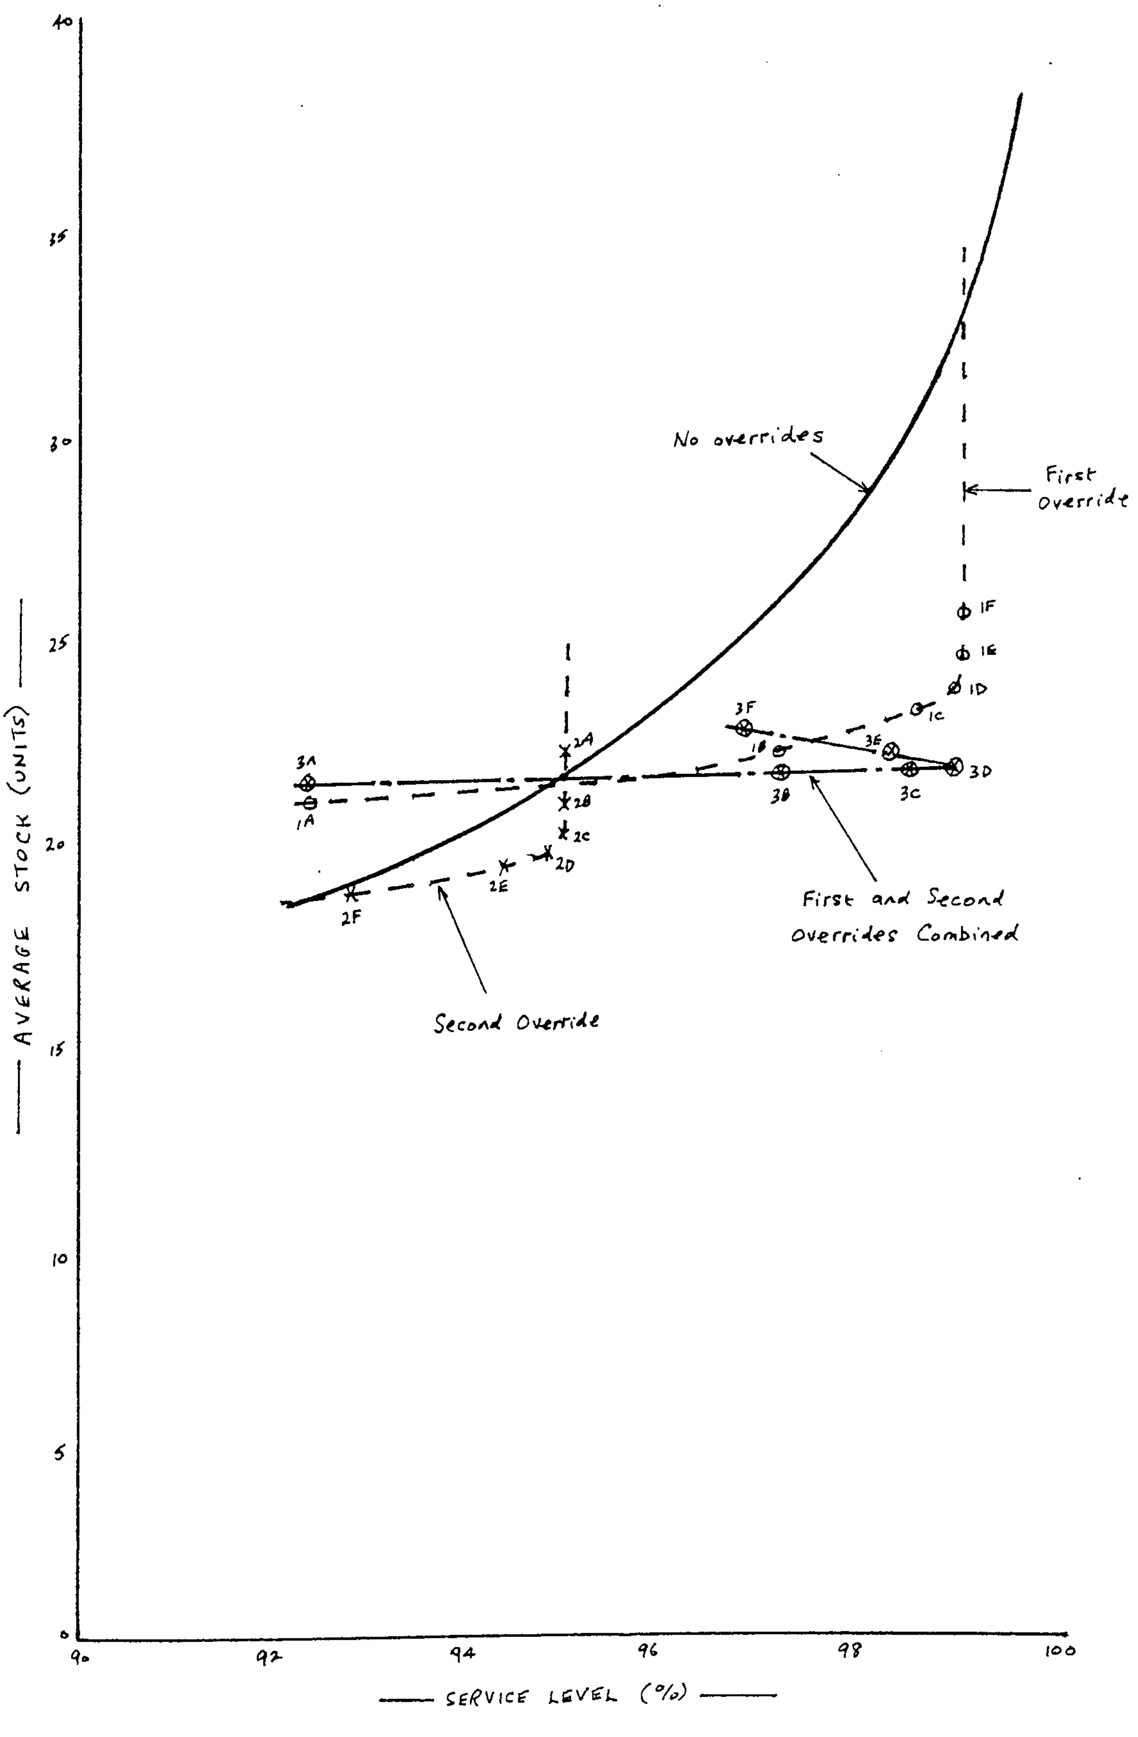

Fig. 7.2 Effects of Overrides of Varying Strengths on Service Level/
Average Stock Level Relationship

of this type without counterbalancing measures may well be unacceptable from a stock investment viewpoint.

Suppose now the buyer had reliable advance warning of a fall in lead times for the last two orders, and an override of 4 days was applied before the penultimate order was placed. The reorder level would be recalculated as follows:

$$\text{Standard deviation to be used} = \sqrt{1.262 \times 4^{1.5}}$$

$$= 3.17 \text{ days}$$

$$\sigma_{LD} = \sqrt{\frac{4}{5} \times 3^2 + 5^2 \times \frac{3.17^2}{5^2}}$$

$$= 4.1532$$

$$k_{LD} = \frac{4^2 \times 5^2}{5^2 \times 4.1532^2}$$

$$= 0.9276$$

$$\text{PLS for 95\% SL during override period} = (1-0.95) \times \frac{30}{4}$$

$$= 0.375$$

U_{LD} for k_{LD} = 0.9276 and PLS = 0.375 is 0.9756

$$R = 0.9756 \times 4.1532$$

$$= 4.05 \text{ units}$$

Given the incidence of the 3-day lead time with this reorder level, the posterior service level and average stock expectations are now calculated:

$$\sigma_{LD} = \sqrt{\frac{3}{5} \times 3^2 + 0}$$

$$= 2.3238$$

$$k_{LD} = 2.7778 \times \frac{3}{5}$$

$$= 1.6667$$

$$U_{LD} = \frac{4.05}{2.3238}$$

$$= 1.7428$$

PLS for $k_{LD} = 1.6667$ and $U_{LD} = 1.7428$ is 0.181789

$$SL = 1 - \left(0.181789 \times \frac{3}{30}\right)$$

$$= 0.9818 \text{ (or } 98.18\%)$$

$$AS = 4.05 - \frac{3}{5} \times 5 + \frac{30 \times 5}{2 \times 5}$$

$$= 16.05$$

A similar calculation for the 5-day lead time with the same reorder level gives a service level over its order cycle of 94.59% with an average stock of 14.05 units. If Table 7.4 is adjusted for these changes without adjusting for the first override, the expected service level over the whole period becomes 94.30% and the average stock 19.17 units. This result is denoted by point '2E' on Fig. 7.2. If both overrides are applied the expected service level over the period becomes 98.33% and the average stock 22.21 units. Thus the joint effect of the two counterbalancing overrides is to produce a very substantial service level increase for virtually no increase in stock investment (point '3E' on Fig. 7.2).

The foregoing results are possible with perfect (or near-perfect) overrides. The more practical case is now considered with overrides subject to varying degrees of error. The first override is first considered where the estimate is in the right direction but only 50% of the correct strength i.e. midway between the computer forecast and the ensuing actual.

Then, override $= 12.75 + \left(\frac{34 - 12.75}{2}\right)$

$$= 23.38 \text{ days}$$

$$\sigma_L = \sqrt{1.262 \times 23.38^{1.5}}$$

$$= 11.94 \text{ days}$$

$$\sigma_{LD} = \sqrt{\frac{23.38}{5} \times 3^2 + 5^2 \times \frac{11.94^2}{5^2}}$$

$$= 13.5885$$

$$k_{LD} = \frac{23.38^2 \times 5^2}{5^2 \times 13.5885^2}$$

$$= 2.9604$$

PLS for prior SL expectation of 95% = $(1-0.95) \times \frac{30}{23.38}$

$= 0.064157$

U_{LD} for $k_{LD} = 2.9604$ and PLS = 0.064157 is 2.7913

R based on override = 2.7913×13.5885

$= 37.93$

It has been shown that with an effectively fixed lead time of 34 days, $\sigma_{LD} = 7.8230$ and $k_{LD} = 18.8890$. Therefore, for the posterior condition, $U_{LD} = \frac{37.93}{7.8230}$

$= 4.8485$

PLS for $k_{LD} = 18.8890$ and $U_{LD} = 4.8485$ is 0.048030

$$SL = 1 - \left(0.04803 \times \frac{34}{30}\right)$$

$$= 0.9456 \text{ (or } 94.56\%\text{)}$$

$$AS = 37.93 - \frac{34}{5} \times 5 + \frac{30 \times 5}{2 \times 5}$$

$$= 18.93 \text{ units}$$

If these figures are substituted in Table 7.4 for those relating to the 34-day lead time (51.54% and 0.49 units), then the expected value of the service level over the complete period becomes 98.48% and the average stock becomes 23.28 units.

Repeating these calculations for overrides of -25% (i.e. 25% in the wrong direction), 25%, 75% and 150% of the correct strength gives the results shown in Table 7.5, and graphically by points 1A to 1F on Fig. 7.2.

Override LT(days)	% Correct Strength	Period of Override		Full Period		Designation on Fig. 7.2
		SL (%)	AS (units)	SL (%)	AS (units)	
7.44	-25	20.43	-8.89*	92.31	20.96	1A
-	0	51.54	0.49	94.90	21.74	Base
18.06	25	79.06	9.75	97.19	22.51	1B
23.38	50	94.56	18.93	98.48	23.28	1C
28.69	75	98.83	26.60	98.84	23.92	1D
34.00	100	99.99	37.04	98.94	24.79	1E
44.63	150	100.00	51.05	98.94	25.95	1F

*The Average Stock formula can produce negative values with very low service levels

Table 7.5 Posterior Expected Values of Service Level and Average Stock (First Override)

The same exercise for the second override gives the results shown in Table 7.6, and by points 2A to 2F on Fig. 2.

Override LT(days)	% Correct Strength	Period of Override		Full Period		Designation on Fig. 7.2
		SL (%)	AS (units)	SL (%)	AS (units)	
14.94	-25	100.00	34.33	94.90	22.38	2A
-	0	99.99	30.49	94.90	21.74	Base
10.56	25	99.97	26.64	94.89	21.10	2B
8.38	50	99.86	24.79	94.88	20.46	2C
6.19	75	99.24	18.92	94.77	19.81	2D
4.00	100	96.40	15.06	94.30	19.17	2E
1.81	150	88.05	11.41	92.90	18.56	2F

Table 7.6 Posterior Expected Values of Service Level and Average Stock (Second Override)

Table 7.7 and points 3A to 3F indicate the combined results of applying both overrides with the same strength.

% Correct Strength	Full Period		Designation on Fig. 7.2
	SL (%)	AS	
-25	92.31	21.28	3A
0	94.90	21.74	Base
25	97.19	21.87	3B
50	98.46	21.99	3C
75	98.71	21.99	3D
100	98.34	22.21	3E
150	96.94	22.77	3F

Table 7.7 Posterior Expected Values of Service Level and Average Stock (First and Second Overrides Combined)

In order to check that the essential characteristics of Fig. 7.2 are not unduly influenced by the demand data, the complete exercise was repeated with a typical slow-moving product ($\mu_D = 1$, $\sigma_D = 1.4$). The result is shown in Fig. 7.3. A further verification exercise was performed with a typical fast-moving product ($\mu_D = 50$, $\sigma_D = 25$). The result is not included as the profile of the graph is virtually identical to that in Fig. 7.2.

The analytical results are now summarised:-

a) The balanced effect of two opposing overrides is to reduce the expected lost sales substantially whilst maintaining the stock at around the target level.

b) In all three exercises, the best results are obtained by applying overrides which are somewhat conservative. Overrides of 75% strength (i.e. 75% of the way between the previous average and ensuing actual) appear best in all cases in either direction. The results are not particularly sensitive to the accuracy of the overrides - anywhere between half strength (points C) and full strength (points E) are almost equally beneficial. This conclusion substantiates the observations of the Controllers of the buyers' course that there is a non-linear relationship between the strength of overrides and their effect.

c) Over-strength overrides are always counter-productive compared with full-strength overrides. This is because full-strength overrides produce posterior expected values of service level approaching 100% during the override period, and stronger measures merely add stock for no

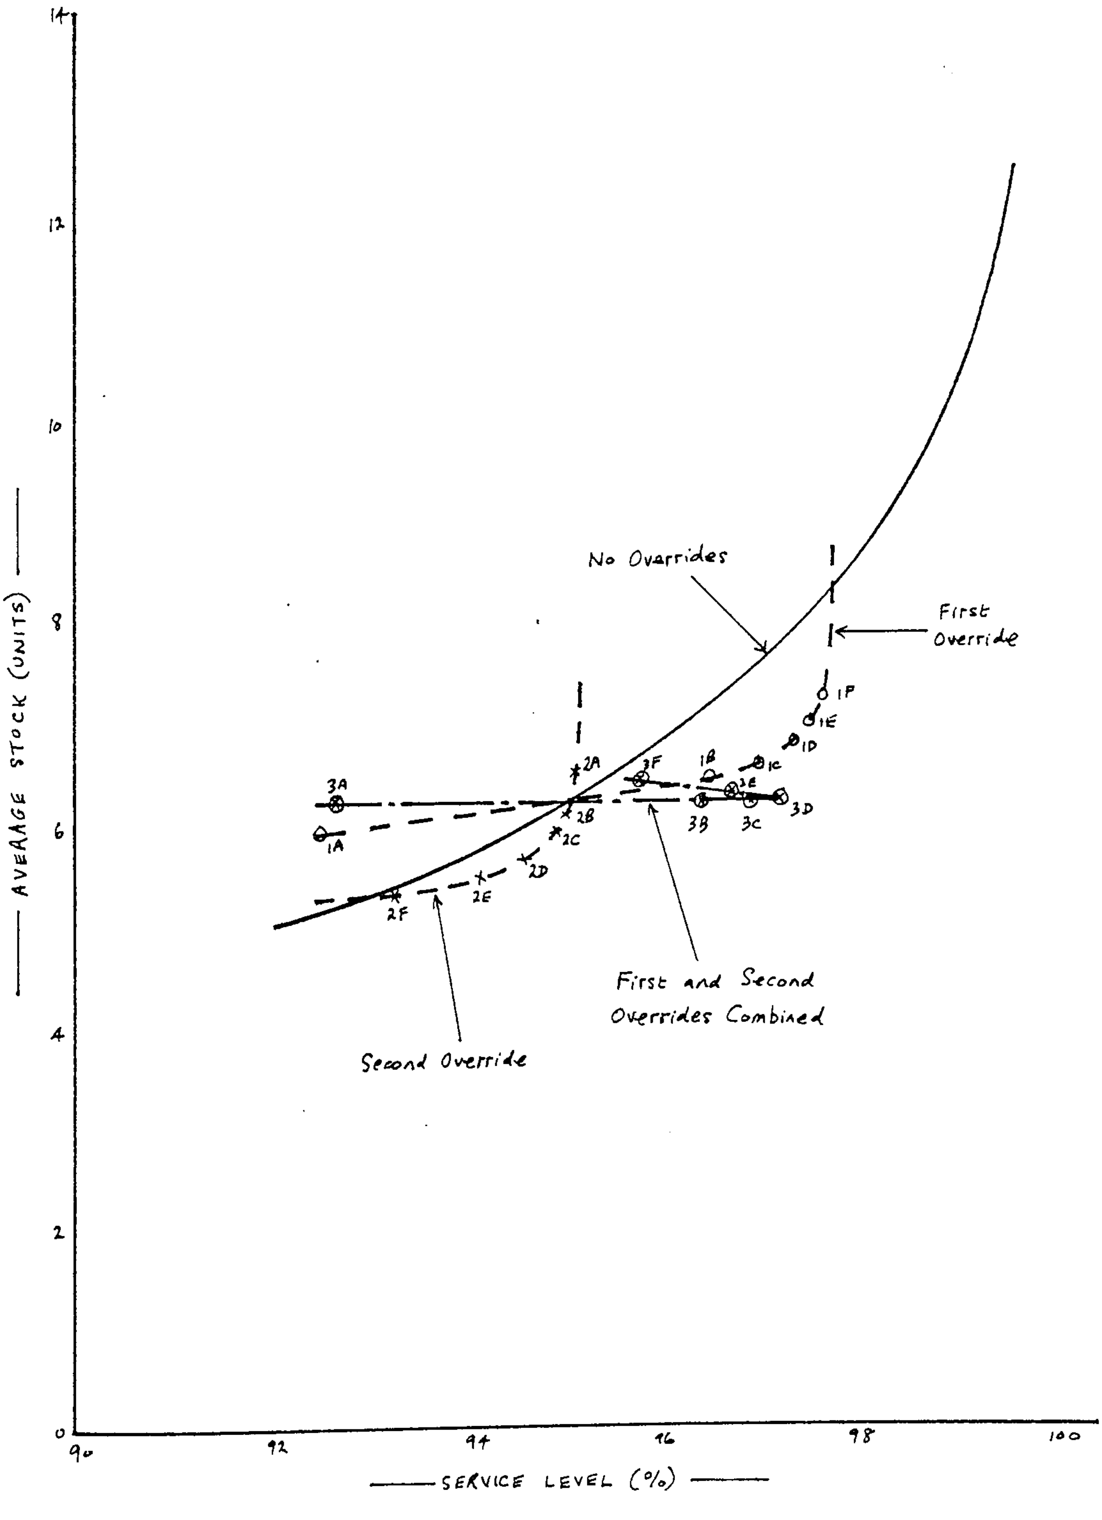

Fig. 7.3 Effects of Overrides of Varying Strengths on Service Level/ Average Stock Level Relationship (Slow-Moving Product)

service benefit.

d) As expected, overrides applied in the wrong direction (points A) are always counter-productive. The effect of a 25% strength override in the wrong direction is seen to be extremely deleterious when the actual is greater than the computer forecast.

e) The greatest improvements are brought about by averting, or at least mitigating, the potential lost sales due to very large lead time demands (usually caused by very long lead times). With the original data set, the protection level corresponding to a 95% service level is 81.73%, i.e. on average only 1 order in every 5.5 results in a stockout. If this stockout can be anticipated and averted by a timely increase in stock over a limited period, the overall service level attained benefits greatly for a marginal overall increase in stock investment. The counterbalancing effects of reducing lead time estimates when the actual lead times are low are far less significant but they are necessary to keep stock investment levels within budgetary constraints.

f) The underlying reason for the performance improvements is the reduction in uncertainty in the lead time forecasts. This reduces buffer stocks for a given service level which produces a more efficient stockholding operation.

g) It is significant that the best results achieved on the buyers' course compare closely with the analytical results (Team 11 succeeded in almost halving the expected lost sales without increasing the stock investment). Whilst the conditions are not exactly comparable, there is a strong suggestion that results which are quite close to the theoretical optimum can be obtained by buyers in a controlled environment. Also, the added complexity of multiple products being ordered collectively does not appear to have an appreciable effect on the outcome.

Lead time overrides have been monitored in the operational system on a number of occasions, by comparing the override forecast errors with the EWMA forecast errors. The most comprehensive check was carried out in March 1982 when all of the overrides in operation at that time were examined. For the four Distribution Centres combined, overrides were applied to 137 Buying Families out of a total of 1795 (7.6%). Of these, 72 had been applied by buyers for the Distribution Centre who had pioneered the development of the system and who had subsequently shown the greatest affinity to it. This represents 14.7% of their Buying Families. On the other hand, one of the other Distribution Centres who had elected to accept only token site training did not have a single override in operation.

Of the 137 overrides, a sample of 46 was selected for analysis. The results are given in Table 7.8. An analysis of the forecast errors reveals that:

$$\sum \left|(\text{Actual} - \text{EWMA})\right| = 391$$

$$\sum \left|(\text{Actual} - \text{Override})\right| = 260$$

$$\sum (\text{Actual} - \text{EWMA})^2 = 4793$$

$$\sum (\text{Actual} - \text{Override})^2 = 2144$$

The Mean Absolute Errors with the overrides are therefore 33.5%

Over-ride No.	Lead Time EWMA (days)	Lead Time Override (days)	Lead Time Actual (days)	Prior Exp. of AS (units)	Posterior Exp.of AS Without O/R(units)	Posterior Exp.of AS With O/R (units)	Posterior Exp.of SL Without O/R (%)	Posterior Exp.of SL With O/R (%)
1	6	4	2	15.92	19.92	15.94	99.44	98.67
2	20	10	12	31.27	39.27	19.39	99.79	95.68
3	23	12	9	33.93	47.93	26.53	99.97	98.98
4	14	10	20	25.57	19.57	11.39	92.96	82.19
5	14	10	6	25.57	33.57	25.39	99.83	99.28
6	14	7	6	25.57	33.57	18.92	99.83	97.75
7	18	10	8	29.42	39.42	23.39	99.90	98.54
8	16	20	12	27.53	31.53	39.27	99.29	99.79
9	10	15	3	21.39	28.39	38.55	99.83	99.98
10	11	15	17	22.47	16.47	24.55	91.16	96.86
11	14	20	11	25.57	28.57	40.27	99.02	99.85
12	20	15	9	31.27	42.27	32.55	99.92	99.62
13	26	5	8	36.52	54.52	12.34	99.99	91.26
14	32	20	20	41.49	53.49	31.27	99.93	98.38
15	3	12	1	12.40	14.40	34.53	99.02	99.99
16	26	20	7	36.52	55.52	44.27	100.00	99.97
17	11	15	5	22.47	28.47	36.55	99.68	99.93
18	22	15	8	33.06	47.06	33.55	99.97	99.73
19	20	15	6	31.27	45.27	35.55	99.98	99.88
20	17	15	12	28.48	33.48	29.55	99.48	99.04
21	39	20	19	47.03	67.03	12.74	99.99	98.69
22	30	15	17	39.87	52.87	24.55	99.95	96.86
23	25	20	23	35.67	37.67	28.27	99.12	97.10
24	29	20	17	39.05	51.05	34.27	99.93	99.17
25	28	15	30	38.21	36.21	11.55	98.33	77.81
26	32	18	20	41.49	53.49	27.42	99.93	97.34
27	28	15	12	38.21	54.21	29.55	99.98	99.04
28	25	20	20	35.67	40.67	31.27	99.54	98.38
29	20	15	23	31.27	28.27	18.55	97.10	90.96
30	15	8	18	26.55	23.55	9.13	96.17	78.56
31	32	15	8	41.49	65.49	33.55	100.00	99.73
32	31	25	25	40.69	46.69	35.67	99.68	98.68
33	26	15	25	36.52	37.52	16.55	98.96	87.95
34	30	20	14	39.87	55.87	37.27	99.98	99.62
35	24	20	15	34.80	43.80	36.27	99.84	99.50
36	13	15	8	24.57	29.57	33.55	99.48	99.73
37	21	10	22	32.17	31.17	9.39	98.11	76.97
38	23	15	17	33.93	39.93	24.55	99.63	96.86
39	29	15	22	39.05	46.05	19.02	99.74	91.77
40	31	17	23	40.69	48.69	22.48	99.80	94.22
41	20	15	9	31.27	42.27	32.55	99.92	99.62
42	23	15	24	33.93	32.93	17.55	98.27	89.53
43	22	18	11	33.06	44.06	36.42	99.92	99.71
44	22	10	8	33.06	47.06	23.39	99.97	98.54
45	14	10	16	25.57	23.57	15.39	96.68	90.39
46	16	10	14	27.53	29.53	17.39	98.79	93.38
				1478.91	1821.91	1212.04	99.08	95.77
				Totals	Totals	Totals	Averages	Averages

Table 7.8 Effects of Lead Time Overrides Applied in the Operational System

less than those with the EWMAs, and the squared errors with the overrides are 55.3% less. Also the absolute error with the override was smaller than that with the EWMA in 29 cases out of the 46 (63.0%). It is therefore strongly indicated that the overall effect on system performance should be clearly beneficial. In order to quantify this benefit the prior and posterior expected values of service level and average stock were calculated in the manner described earlier. The results are given in Table 7.8, columns 5-9. The following assumptions were made:-

i) Each order is assumed to be for a single product with a mean demand rate of 5 units per week and a standard deviation which conforms with the Variance Law. The use of a single product enables stocks to be added without the complication of varying unit costs. In order to check that the results are not unduly influenced by the selected product parameters, the exercise was repeated with two other products with mean demand rates of 1 unit per week and 50 units per week respectively. The results were not fundamentally different to those presented here.

ii) A nominal order interval of 20 days is used throughout. This gives a fixed relationship between PLS and service level which makes posterior service level variations easier to comprehend.

iii) The target service level is 95% throughout. This means that the 'Prior Expected Value of Service Level' is 95% in all cases and hence it is not tabulated.

Before examining the results in detail, it should be appreciated that they are strongly influenced by the trading climate at the time. During the latter half of 1981 and the early months of 1982 lead times were generally improving. The System Dynamics study in Chapter 4 has shown that under these circumstances the system substantially overprovides. Table 7.8 shows that the lead time EWMAs exceed the actual lead times in 38 cases (82.6%). The posterior expected values of service level can therefore be expected to exceed substantially the prior expected values. This tendency is reinforced by the absence of any very high lead times which would have caused very low service levels each of which would counterbalance several high ones. Another consequence of the EWMAs exceeding the actual lead times is that the posterior expectations of average stock will be greater than the prior expectations. This is because the mean demand in lead time (which is subtracted from the reorder level - refer equation 4) is numerically smaller in the posterior condition.

In the light of the analytical study leading to Fig. 7.3, the overrides are examined in four groups.

Group 1

This consists of overrides 1, 3, 5, 6, 7, 10, 14, 18, 21, 24, 27, 28, 31, 32, 34 and 44. All of these were applied in the correct direction and they were between half and full strength. According to the analytical study their effect should be near-optimal. For this group the overall results are:

Posterior service level without override - 99.32%

Posterior service level with override - 98.75%

Posterior aggregate stock without override - 719.50 units

Posterior aggregate stock with override - 437.25 units

Thus, the stock is cut by 39.2% for a very marginal deterioration in Service level.

Group 2

This comprises overrides 12, 16, 19, 20, 35, 41 and 43, all of which were applied in the correct direction but below half of the correct strength. According to the previous analysis they should produce weaker benefits than Group 1. The overall results are:

Posterior service level without override - 99.87%

Posterior service level with override - 99.62%

Posterior aggregate stock without override - 306.67 units

Posterior aggregate stock with override - 247.16 units

As with Group 1, a marginal service level deterioration is evident, but here the stock reduction is only 19.4%.

Group 3

Overrides 2, 3, 22, 23, 26, 33, 38, 39, 40 and 46 were all applied in the correct direction but too strongly. This condition is represented by the 'F' points in Fig. 7.2. Though they are depicted there as beneficial, if the strength had been increased further they would have crossed the 'Normal Relationship' curve and become counterproductive. Hence the effect of an overstrength override is indeterminate. The overall results are:

Posterior service level without override - 99.57%

Posterior service level with override - 94.24%

Posterior aggregate stock without override - 439.54 units

Posterior aggregate stock with override - 211.96 units

In terms of stockholding efficiency this result is broadly neutral. Fig. 7.2 and many similar graphs with different data (not included here) indicate that approximately twice as much stock is required to raise the service level from 95% to 99½%.

Group 4

This consists of overrides 4, 8, 9, 11, 15, 17, 25, 29, 30, 36, 37, 42 and 45, all of which were applied in the wrong direction. The analytical study suggests that the effect of these must be detrimental. The overall results are:

Posterior service level without override	- 98.00%
Posterior service level with override	- 91.21%
Posterior aggregate stock without override	- 356.20 units
Posterior aggregate stock with override	- 315.67 units

The modest reduction of 11.4% in the stockholding in no way compensates for the severe degradation in service level even if this action constitutes an acceptable strategy shift.

The application of overrides in the operational system is now summarised:-

a) As there was a general improvement in lead times during the period, it is certain that only a small fraction of the total number of justifiable overrides were applied. The reason most commonly given for this is an alleged unreliability of suppliers in quoting realistic lead times.

b) The effect on total system performance of this reluctance to override lead times would be to produce results which are more akin to the posterior expectations without overrides i.e. service levels and average stocks which are well above the prior calculations.

c) When overrides were applied a definite net advantage ensued, and the portion of the system affected was steered to a condition much nearer the strategic objectives.

d) The quality of the overrides fell far short of that achieved on the training courses. This is considered to be due to three main reasons:-

i) The information given to buyers by suppliers is less reliable than that given by the course Controllers. In the operational situation, over a quarter of the overrides were applied in the wrong direction, whereas on the courses this occurred only in a few instances of mistiming.

ii) On the courses, the buyers were given feedback on the effects of every override, together with a reinforcement of the principles which should be adopted. In the operational system there is no regular feedback of forecasting errors.

iii) On the courses, the compression of 2 years' data into 2 days forced a concentration of effort into applying overrides. In the operational system, this activity is interspersed with many other disparate activities. Hence lead time movements may be unnoticed.

e) The system effects of the overrides are completely in line with the analytical expectations. Properly applied overrides can be seen to produce a much more efficient stockholding operation than a sole reliance on historical data for forecasting.

f) It is extremely significant that the human manipulation of lead time forecasts produces performance changes of a much higher order than any of the theoretical factors analysed in Chapter 6. It is therefore of paramount importance to raise the achievement of human performance in this area of activity.

7.4 PROBABILISTIC ESTIMATING FOR LEAD TIME OVERRIDING

It has been demonstrated in Section 6.1 that the pattern of lead times approximates a Gamma distribution. This deduction was arrived at by measuring lead times over a long period of time before subjecting the data to statistical analysis. The form of the distribution and its parameter values therefore represent the amalgam of circumstances which existed over the whole of the data collection period. At any point in time, the full variety in the supply circumstances would not be expected to exist. It should therefore be possible to obtain a more accurate forecast by obtaining information which would identify the subset of circumstances extant at that particular time. As previously stated, a reduction in uncertainty would permit buffer stocks to be reduced for the same level of service.

The essential role of the buyer here is to reduce uncertainty by providing the system with lead time information which is specific to the prevailing circumstances and which cannot be elicited from the historical data. In view of the importance of lead time overrides on system performance, this Section examines possible ways of improving them. The discussion should be regarded as exploratory, and the experimental results are intended only to provide direction for further research.

In the present system, the buyer applies an override in the form of a point estimate, this being his judgement of the most likely eventuality (i.e. his perceived mode). This is used in place of the computer EWMA as the mean value of the forecast. Buyers have unanimously indicated that they do not take account of a spread of probabilities when submitting overrides - they simply use the most likely eventuality. The possibility of using probabilistic estimates is now examined. This would mean that the buyer would submit his subjective assessment of the likelihood distribution of the next lead time falling into mutually exclusive ranges. Such an assessment has two distinct elements compounded:-

1) The Estimator's familiarity with the statistical concepts of probability.

2) The Estimator's contextual knowledge of the domain in which the judgement is required.

Winkler and Murphy (94) refer to these as the 'normative' and 'substantive' standards of goodness respectively.

The usual technique for judging probability assessments is to compare the stated probabilities with the subsequent event using 'Scoring Rules' (refer Kidd (95)). A number of Scoring Rules have been proposed, a common one being the 'Quadratic Rule'. An experiment conducted by Staël Von Holstein (96) will be described to illustrate the method.

The estimation task involved the movement of buying prices on the Stockholm Stock Exchange. The subjects were asked to assign probabilities (summing to unity) for five classes of outcome which were considered to be more or less equi-probable according to past records:

i) Price decreases > 3%

ii) Price decreases > 1% $\leqslant$ 3%

iii) Price changes $\leqslant$ 1% in either direction

iv) Price increases > 1% $\leqslant$ 3%

v) Price increases > 3%

The Quadratic Rule is expressed as:

$$S_j(r) = 2r_j - \sum_{i=1}^{i=n} r_i^2$$

where r is the expressed probability of event class i occurring

j is the event class which actually occurs

n is the number of event classes

S_j (r) is the score received for the assessment

e.g. if a subject assigns probabilities of .05; .25; .50; .20; 0 to five classes and the second outcome occurs, then:

$$S_j(r) = (2 \times .25) - (.05^2 + .25^2 + .50^2 + .20^2 + 0^2)$$

$$= +.145, \text{ which is the score for the assessment.}$$

Two characteristics of the Quadratic Rule are important apropos the present study:-

1) The maximum score possible is +1, which occurs when a probability of 1 is assigned to a single class and the assessment is correct. Then,

$$S_j(r) = (2 \times 1) - (1^2 + 0^2 + 0^2 + 0^2 + 0^2) = +1$$

The minimum score possible is -1, which occurs when a probability of 1 is assigned to a single class and the assessment is wrong. Then,

$$S_j(r) = (2 \times 0) - (1^2 + 0^2 + 0^2 + 0^2 + 0^2) = -1$$

2) If equal estimates of 0.2 are assigned to each of the five classes,

$$S_j(r) = (2 \times .2) - (.2^2 + .2^2 + .2^2 + .2^2 + .2^2) = +0.2$$

More generally, for n classes,

$$S_j(r) = \left(2 \times \frac{1}{n}\right) - \left[\left(\frac{1}{n}\right)^2 + \left(\frac{1}{n}\right)^2 + \left(\frac{1}{n}\right)^2 + \ldots \left(\frac{1}{n}\right)^2\right] = \frac{1}{n}$$

Therefore, a neutral assessment is awarded a positive score, and that is high with a small number of classes.

It follows from these characteristics that the Quadratic Rule favours an even spread of estimation (i.e. 'hedging one's bets'). In an extreme case of dichotomous estimating, assigning a probability of .5 to both classes would ensure a score of +0.5, whereas assigning a probability of 1 to either class would give an expected score of ½ (+1, -1) = 0.

In the present study, each of six buyers was asked to make a probabilistic lead time assessment for each supplier upon whom an order was about to be placed. Seven classes were used, which were not equi-probable though they were designed to give a reasonably even spread of lead time incidences. In all, 29 assessments were made which are shown as the middle vector for each supplier in Table 7.9. Buyer A submitted the assessments for suppliers 1-5, B for 6-8, C for 9-13, D for 14-18, E for 19-23, and F for 24-29. The four sets of scores - S^1, S^2, S^3 and S^4 - are now explained.

The S^1 scores were calculated on the basis that the buyer made a point estimate at the modal value. e.g. for Supplier 1,

$$S_j(r) = (2 \times 1) - (0^2 + 0^2 + 1^2 + 0^2 + 0^2 + 0^2 + 0^2) = \underline{+1.00}$$

The S^2 scores used the full probabilistic assessment, e.g.

$$S_j(r) = (2 \times .90) - (0^2 + .05^2 + .90^2 + .05^2 + 0^2 + 0^2 + 0^2) = \underline{+0.99}$$

The S^3 scores are based on the prior lead time distributions shown as the top vector. These were obtained by calculating the Gamma modulus (k) from the prior mean and standard deviation, then using equation (18) in Section 3.4 to calculate Gamma probabilities. This can be assumed to represent the algorithmic forecast unmodified by the buyer.

e.g. $$S_j(r) = (2 \times .18) - (.09^2 + .17^2 + .18^2 + .15^2 + .21^2 + .16^2 + .04^2) = \underline{+.20}$$

Supplier	Probabilities for Lead Times (Weeks)*							Prior Mean LT	Next Actual LT	S^1	S^2	S^3	S^4
	0-1	1-2	2-3	3-4	4-6	6-10	>10						
1	.09	.17	.18	.15	.21	.16	.04	4.00	2.40	+1.00	+.99	+.20	+.99
	0	.05	.90	.05	0	0	0						
	0	.05	.91	.04	0	0	0						
2	.25	.27	.19	.12	.11	.05	.01	2.40	3.60	-1.00	-.91	.04	-.87
	0	.05	.95	0	0	0	0						
	0	.07	.93	0	0	0	0						
3	.13	.21	.19	.15	.18	.12	.02	3.40	5.20	+1.00	+.98	+.19	+.99
	0	0	0	.05	.85	.10	0						
	0	0	0	.04	.89	.07	0						
4	.56	.26	.11	.05	.02	0	0	1.20	0.80	-1.00	-.81	.72	-.62
	.05	.95	0	0	0	0	0						
	.10	.90	0	0	0	0	0						
5	.04	.10	.13	.14	.23	.25	.11	5.47	5.00	-1.00	-.46	.28	-.26
	0	0	.80	.10	.10	0	0						
	0	0	.74	.10	.16	0	0						
6	.01	.04	.08	.10	.21	.32	.24	7.60	7.20	-1.00	+.04	+.42	+.27
	0	0	0	.40	0	.20	.40						
	0	0	0	.20	0	.32	.48						
7	.02	.09	.11	.12	.23	.28	.15	6.16	4.00	+1.00	+.50	+.27	+.77
	0	0	0	.50	.50	0	0						
	0	0	0	.34	.66	0	0						
8	.01	.04	.06	.09	.20	.32	.28	8.05	8.00	+1.00	+.78	+.41	+.85
	0	0	0	.10	.10	.60	.20						
	0	0	0	.03	.07	.69	.21						
9	.37	.29	.16	.09	.07	.02	0	1.80	1.60	-1.00	-1.00	+.32	-1.00
	0	0	1.00	0	0	0	0						
	0	0	1.00	0	0	0	0						
10	.56	.26	.11	.05	.02	0	0	1.20	1.00	+1.00	+1.00	+.72	+1.00
	0	1.00	0	0	0	0	0						
	0	1.00	0	0	0	0	0						
11	.09	.17	.18	.15	.21	.16	.04	4.00	3.30	-1.00	-.81	+.14	-.84
	0	0	.95	.05	0	0	0						
	0	0	.96	.04	0	0	0						
12	.03	.08	.12	.13	.23	.27	.14	6.00	7.80	-1.00	-.62	+.36	-.55
	0	0	0	0	.90	.10	0						
	0	0	0	0	.88	.12	0						
13	.21	.25	.19	.13	.14	.07	.01	2.70	3.60	+1.00	+1.00	+.08	+1.00
	0	0	.01	.99	0	0	0						
	0	0	.01	.99	0	0	0						
14	.01	.04	.06	.09	.20	.32	.28	8.00	8.00	-1.00	-.28	+.41	-.01
	0	0	0	0	.80	.20	0						
	0	0	0	0	.71	.29	0						
15	.01	.03	.07	.09	.20	.32	.28	7.90	2.90	-1.00	-.50	-.09	-.57
	0	0	0	.50	.50	0	0						
	0	0	0	.31	.69	0	0						
16	.25	.27	.19	.12	.11	.05	.01	2.40	4.60	+1.00	+.50	+.02	+.46
	0	0	0	.50	.50	0	0						
	0	0	0	.52	.48	0	0						

Table 7.9 Scores for Buyer and Algorithm Estimates

Cont'd/...

Supplier	Probabilities for Lead Times (Weeks) *							Prior Mean LT	Next Actual LT	S^1	S^2	S^3	S^4
	0-1	1-2	2-3	3-4	4-6	6-10	>10						
17	.21	:26	.19	.13	.13	.06	.02	2.65	2.20	-1.00	-.63	+.19	-.77
	0	0	0	0	.75	.25	0						
	0	0	0	0	.87	.13	0						
15	.09	.17	.18	.15	.21	.16	.04	4.00	0.80	-1.00	-1.00	+ .02	-1.00
	0	0	0	1.00	0	0	0						
	0	0	0	1.00	0	0	0						
19	0	.01	.02	.04	.12	.29	.52	11.40	6.40	+1.00	+.97	+.21	+.97
	0	0	0	0	.10	.85	.05						
	0	0	0	0	.04	.87	.09						
20	.03	.09	.12	.13	.23	.27	.13	5.90	5.50	+1.00	+.86	+.28	+.86
	0	0	0	.10	.70	.20	0						
	0	0	0	.06	.71	.23	0						
21	.09	.17	.18	.15	.21	.16	.04	4.00	1.90	-1.00	-1.00	+.18	-1.00
	0	0	1.00	0	0	0	0						
	0	0	1.00	0	0	0	0						
22	.02	.08	.11	.12	.23	.29	.15	6.20	3.60	-1.00	-.82	+.05	-.79
	0	0	0	0	.90	.10	0						
	0	0	0	0	.88	.12	0						
23	.17	.24	.19	.14	.16	.09	.01	3.00	2.60	-1.00	-.60	+.20	-.80
	0	0	0	0	.75	.15	.10						
	0	0	0	0	.89	.10	.01						
24	.12	.20	.19	.15	.19	.12	.03	3.60	2.40	+1.00	+.98	+.22	+.99
	0	0	.90	.08	.02	0	0						
	0	0	.92	.06	.02	0	0						
25	.04	.12	.14	.14	.23	.24	.09	5.12	4.30	-1.00	-.78	+.29	-.74
	0	0	.90	.08	.02	0	0						
	0	0	.89	.08	.03	0	0						
26	.03	.08	.12	.13	.23	.27	.14	6.00	5.00	-1.00	-.26	+.28	-.06
	0	0	.60	.30	.10	0	0						
	0	0	.54	.29	.17	0	0						
27	.09	.18	.18	.15	.21	.15	.04	3.98	1.40	-1.00	+.14	+.20	+.15
	0	.30	.60	.10	0	0	0						
	0	.31	.61	.08	0	0	0						
28	.03	.10	.13	.13	.23	.26	.12	5.60	5.50	-1.00	+.45	+.28	+.38
	0	0	0	0	.45	.50	.05						
	0	0	0	0	.43	.54	.03						
29	.02	.06	.09	.11	.22	.30	.20	6.93	4.20	+1.00	+.81	+.24	+.73
	0	0	0	.05	.65	.27	.03						
	0	0	0	.02	.61	.34	.03						
									Average	-.24	-.02	+.25	+.02

S^1 = Buyer point estimate
S^2 = Buyer probabilistic estimate
S^3 = Algorithm estimate
S^4 = Joint Buyer/Algorithm estimate

* Top figure in each cell = Prior Probability
Middle figure in each cell = Likelihood Distribution
Bottom Figure in each cell = Posterior Probability

Table 7.9 (Cont'd)

The S^4 scores are calculated using a Bayesian combination of the prior probabilities and the buyer's likelihood distribution, following a method by Kidd (97). The posterior probabilities, shown as the bottom vector, are computed as:

$$Po_i = \frac{Pr_i \times L_i}{\sum_{i=1}^{i=7} Pr_i \times L_i}$$

where Po_i is the posterior probability for event class i

Pr_i is the prior probability for event class i

L_i is the buyer's assessed probability for event class i

e.g. for the 2-3 week class for Supplier 1,

$$Po_i = \frac{(.18x.90)}{\left[(.09x0)+(.17x.05)+(.18x.90)+(.15x.05)+(.21x0)+(.16x0)+(.04x0)\right]}$$

$$= .91$$

And, $S_j\ (r) = (2 \times .91)-(0^2 + .05^2 + .91^2 + .04^2 + 0^2 + 0^2 + 0^2) = \underline{+.99}$

Bayesian forecasting methods have been the subject of a considerable amount of research over the past decade with differing conclusions. Bunn (98), and Green and Harrison (99) have reported promising results, whereas Dikenson (100) has concluded that the statistical combination of forecasts is not beneficial. In principle, the Bayesian method uses historic data (which is normally the sole basis of computer forecasting algorithms), and combines this with a further datum to form posterior probabilities upon which the forecast is made. Insofar as the datum usually takes the form of a subjective likelihood distribution, the principle is entirely consistent with the notion of man-computer symbiosis discussed in Chapter 2. It is hypothesised that this offers a sound basis for working towards better lead time forecasting methods.

In the present study it has to be concluded that the results presented in Table 7.9 are on the whole disappointing. The reasons for this are considered to be:-

a) Considering that the buyers had the opportunity to use every means at their disposal to obtain their estimates (including sight of the computer EWMAs), their forecasts were not very good. The Mean Absolute Error of the mid-point of the range into which they assigned their modal estimates was 1.35 weeks, compared with 1.43 weeks for the EWMA.

b) The buyers used very little spread for assigning probabilities. Fig. 7.4 depicts the average of the assessments for each Buyer as 'personal p.d.f.'s' with the modal estimates overlaid. It can be seen that Buyers A, C and E, in particular, placed a great deal of confidence in their modal estimates, and Buyer D never used more than two classes. In most cases this confidence was misplaced. The S^1 scores show that the modal values did not occur in 18 of the 29 cases. Also, of the four instances where the Buyer assigned 100% probability to one class, three of these proved to be wrong. In the more extensive experiment already referred to, Staël Von Holstein (96) found that of 40 assessments concentrated into one class, only 12 proved correct. Both of these results substantiate the widely held view that humans ascribe far too much certainty to subjective forecasts.

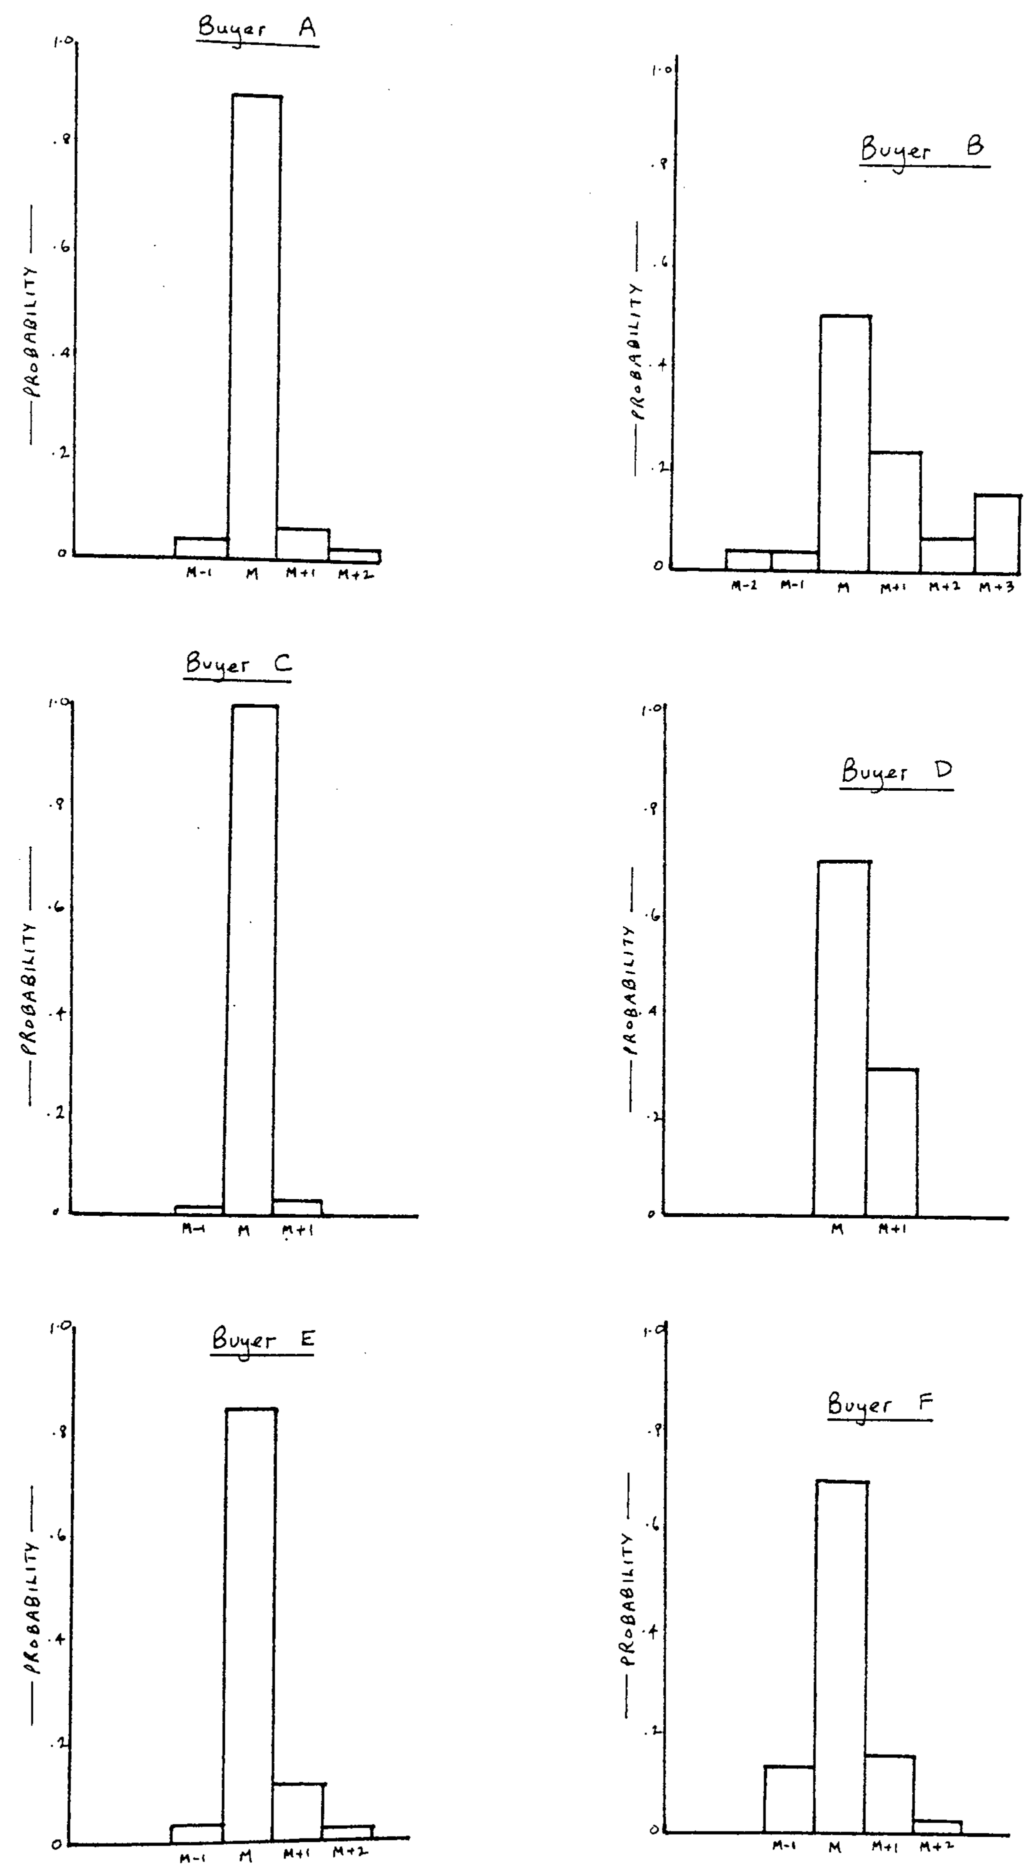

Note - M = Mode

Fig. 7.4 Buyer Personal Probability Density Functions

c) The prior probabilities and likelihood distribution are combined multiplicatively to form the posterior distribution. It follows that the spread of the posterior distribution is equal to or less than the smaller of the two component spreads. The buyers' tendency to ascribe too high a degree of certainty to their favoured outcome is therefore passed through to the combined forecast.

d) As previously indicated, the Quadratic Rule favours dispersed estimating. The fact that only 3 of Staël Von Holstein's 72 subjects recorded a final score which exceeded that which would have been obtained by consistently making neutral assessments is, in the opinion of the present author, due more to this reason than Staël Von Holstein's deduction that the task was too difficult.

In spite of the fact that the Bayesian combination appeared to be outperformed by the algorithm alone, there are some indications that with training in probabilistic estimating, the buyers could contribute towards combined forecasts which would consistently outperform the algorithm. The results for Supplier 8 provide a good example. In this case both the buyer and the algorithm are correct in their modal estimate. The buyer is, however, more certain in his assessment (60% vs. 32%), though he has allowed a reasonable dispersion around the mode. The effect is that the buyer and algorithmic forecasts reinforce one another to give a joint modal probability of 0.69. It is important to observe that the posterior distribution assumes a similar form to the prior distribution. This should therefore provide better cover in the regions where the lead time is most likely to occur.

If buyers were trained in probabilistic estimating, it seems likely that the joint forecast should outperform the algorithm in most cases. The buyer's modal probability could be expected to exceed the algorithm's modal probability (otherwise he would not be reducing uncertainty by applying his special knowledge), but he should allow a sufficient spread for the algorithm to impose weightings on the other classes which would reduce the probable error if the modal value is not correct or if the buyer and algorithmic modes do not coincide. Here the notion of symbiosis is clearly apparent - the buyer provides the algorithm with specialist knowledge of the current circumstances, and the algorithm provides the fine statistical weightings.

It should be noted that the results of this experiment do not contradict the results obtained on the buyers' course (Section 7.2). In this experiment, the buyers were asked to make forecasts on all orders placed in a short time period. On the course, they applied overrides only in special circumstances where they had strong reason to believe that the historic records would be inappropriate for forecasting.

7.5 EFFECTS ON PERFORMANCE OF ORDER QUANTITY OVERRIDES FOR NON-COMMERCIAL REASONS

An analysis of order placements over a 6-week period in May/June 1982 revealed that at one of the Distribution Centres 37% of the normal computer orders were either pre-empted by buyer ordering action or else the recommended quantities were changed before the order was placed. There were two main reasons for these actions:-

1) A buyer distrust of the computer calculations.

2) Bulk purchasing for commercial reasons.

This Section deals with the first reason. The second reason is treated in Section 7.6.

Prior to the introduction of the computer system, the buyers operated on a doctrine of maintaining buffer stocks equal to a set number of weeks' supply. This number of weeks was normally constant from product to product except for those of special commercial importance, e.g. white paint.

It was immediately evident on the buyers' training course that this practice had been strongly ingrained, and there was a strong tendency to change any computer recommendations which deviated far from it. Because of time pressures they were unable to follow their instincts on every occasion, and at the end of the course they invariably had to accept that the computer had provided sufficient cover to give the service intended. Nevertheless, it was suspected that they might have a tendency to revert to their former practice when they returned to the 'real world'. This proved to be the case and a number of refresher sessions were held. Even so, it is an agreed philosophy that the buyers should have the authority to override the computer whenever they are unhappy with the recommendations, and they carry the ultimate responsibility for the consequences.

After the order placement analysis, an exercise was undertaken to check if the buyers' conception of buffer stock requirements had changed after a lengthy period of operation of the computer system. Seven buyers were asked to estimate the buffer stock for each set of parameters given in Table 7.10 assuming a fixed order interval of 3 weeks. The main elicitations from Table 7.10 are:-

1) Apart from an isolated aberration by Buyer 5, all buyers recognise that higher service levels require additional buffer stocks. However, the rate of increase in buffer stocks to support the higher service levels is not nearly large enough. This is surprising as the asymptotic profile of the Service Level/Average Stock graph had been impressed upon them repeatedly.

2) All buyers except Buyer 3, and possibly Buyer 4 whose results are inconsistent, recognise that high variability of demand and lead time require increased buffer stocks.

3) Only four buyers achieved the correct ranking of the four Groups (3, 1, 4, 2 ascending). This is not surprising, as the correct ranking cannot be ascertained with certainty by inspection alone.

It is concluded that all of the buyers have abandoned their previous notion of a fixed number of weeks' supply as a suitable criterion for determining the magnitude of buffer stocks. With one or two exceptions, they are aware of the direction of the relationships between buffer stock and service level, and between buffer stock and data variability. The operation of the system at various formal service level settings has clearly helped to

Parameter Set	μ_D (units/wk)	MAD_D (units/wk)	μ_L (days)	MAD_L (days)	Target SL (%)	Buyer Estimate of Buffer Stock (Weeks' Supply)								Computer Buffer
						1	2	3	4	5	6	7	Av.	
1	5	1.5	15	3	90	3.6	1.2	4.5	1.4	1.2	1.2	4.4	2.5	0.3
2	5	1.5	15	3	95	4.8	1.4	4.8	1.6	1.4	1.5	4.8	2.9	0.9
3	5	1.5	15	3	99	6.0	1.8	5.0	1.8	2.0	1.6	5.0	3.3	2.0
4	5	4.5	15	12	90	4.4	6.2	4.5	4.0	5.0	6.5	10.0	5.8	4.4
5	5	4.5	15	12	95	6.8	6.8	4.8	4.4	3.6	6.9	11.0	6.3	6.7
6	5	4.5	15	12	99	8.6	7.2	5.0	4.8	6.0	7.2	12.0	7.3	12.0
7	100	15.0	15	3	90	3.5	0.7	4.5	6.5	0.8	1.0	4.5	3.1	0.1
8	100	15.0	15	3	95	4.6	0.9	4.8	6.8	0.9	1.1	4.8	3.4	0.6
9	100	15.0	15	3	99	6.5	1.1	5.0	7.2	1.1	1.1	6.0	4.0	1.5
10	100	37.6	15	12	90	4.0	3.7	4.5	5.0	3.5	4.0	7.2	4.6	2.8
11	100	37.6	15	12	95	4.6	4.1	4.8	5.2	4.0	4.2	7.6	4.9	4.4
12	100	37.6	15	12	99	7.1	4.4	5.0	5.5	5.0	4.4	8.0	5.6	8.2

Table 7.10 Estimation of Buffer Stocks by Buyers

eradicate the concept of a fixed number of weeks - for they would have been forced to accept that if Management wished to aim for a higher service level they must permit the buffer stocks to rise. The buyers nevertheless have a poor appreciation of buffer stock magnitudes.

A dissatisfaction with buffer stock levels can result in two types of override being applied. The buyer can either:-

1) raise an order before the computer order is generated to protect products which he feels are vulnerable, or
2) adjust the computer recommended order quantities downwards where he considers the products are over-protected.

In the first case the buyer is discouraged from placing an order for just the vulnerable items, but instead to request a computer recommended order even though one is not due. On receipt of such a request, the computer ignores the ELS trigger point and produces a balanced order with reduced quantities. If perpetuated, this principle of generating an order for a family of products when the first one falls to its reorder point is akin to one described by Brown (57) as a 'can-order' point method. The result is that the target service level can then be regarded as a minimum and the expected service level is indeterminate. Simulation trials using program FT01 have revealed that the average service level obtained on a 95% setting is approximately 97.2% and the average stock is 38% above the normal value. As only around 6% of all orders are pre-empted in this way, it is estimated that the practice adds 0.13% to the achieved overall service level (on a setting of 95%) and 2.3% to the average stock investment.

In the second case, the reduced order quantities of the products which the buyer feels are over-protected will in fact result in their under-protection. There is also a second more insidious effect, in that the reduced quantities will advance the next order, thus reducing the order interval. As the ELS trigger point is calculated from the target service level on the basis of a planned number of risk points per annum, by increasing the number of risk points per annum and keeping the ELS unchanged, the service level for the whole family will be prejudiced. This is another case where a sound understanding of the system concepts is essential before the normal homeostatic functioning may be interfered with.

Around 8% of all orders are adjusted downwards. It has been observed that this reduces the order interval by an average of around 20%, which produces an increase in ELS of 25% p.a. On a service level base of 95%, the expected service level will be 100 - (1.25 x 5) = 93.75%. The overall service level then becomes (0.92 x 95) + (0.08 x 93.75) = 94.90%. The overall average stock could be expected to decrease by around 1.5%.

It is concluded that the counterbalanced effect of adjusting order quantities because of a distrust of the computer buffer stock calculations is small overall, but it could assume significant proportions if the practice became widespread.

7.6 EFFECTS ON PERFORMANCE OF ORDER QUANTITY OVERRIDES FOR COMMERCIAL REASONS

An analysis of the constitution of the physical stock at the two largest Distribution Centres was carried out in June 1981 and September 1981 respectively. Distribution Centre 'A' had a physical stock value of £4.07m of which £1.05m was classified as 'excess' (i.e. above the computer calculated Maximum Order Cover). Distribution Centre B had a physical stock of £2.65m of which £0.63m was excess. There were a total of 9,180 products at 'A' and 7,510 at 'B' giving a joint total of 16,690.

In each case an 'Excess Stock' computer report was produced which listed all of the items where the physical stock exceeded the MOC. This contained 3,674 items (40% of the range) for Distribution Centre 'A' and 3,402 items (45%) for Distribution Centre 'B'. As the computer orders up to the MOC, there had clearly been a substantial amount of overbuying. Every item on these reports was scrutinised by the relevant buyer who assigned a reason for the overstock wherever possible. As this exercise was dependent mainly on memory, a high degree of accuracy was not expected. The reasons were categorised by the author, and some of these are reported in Table 7.11.

Reason for Excess	Dist.Centre 'A' Value (£)	Dist.Centre 'B' Value (£)	Dist.Centres 'A' + 'B' Value(£)	%	Items
Sales Promotion	280422	94385	374807	22.2	1468
Special Deal	86284	29041	115325	6.8	524
Supplier Rebate	64713	21781	86494	5.1	466
Initial Stocking	164342	156012	320354	19.0	1297
Improved Discount	160832	87092	247924	14.7	976
Total (Commercial)	756593	388311	1144904	67.7	4731
Total (Non-Commercial)	290578	254929	545507	32.3	2345
Grand Total	1047171	643240	1690411	100.0	7076

Table 7.11 Classification of Excess Stocks

The non-commercial reasons comprise:-

a) System dynamics, including dynamic recalculation of the MOC (in a declining demand climate, the MOC often falls below the physical).

b) Buyer adjustments as described in Section 7.7.

c) Customer and Branch returns.

d) Unidentified reasons.

The commercial reasons are examined separately.

Sales Promotion

As previously stated, sales promotions are a regular trading characteristic in the industry. Expressed as a percentage of the effective (i.e. non-excess) stock, the increase due to sales promotions is:

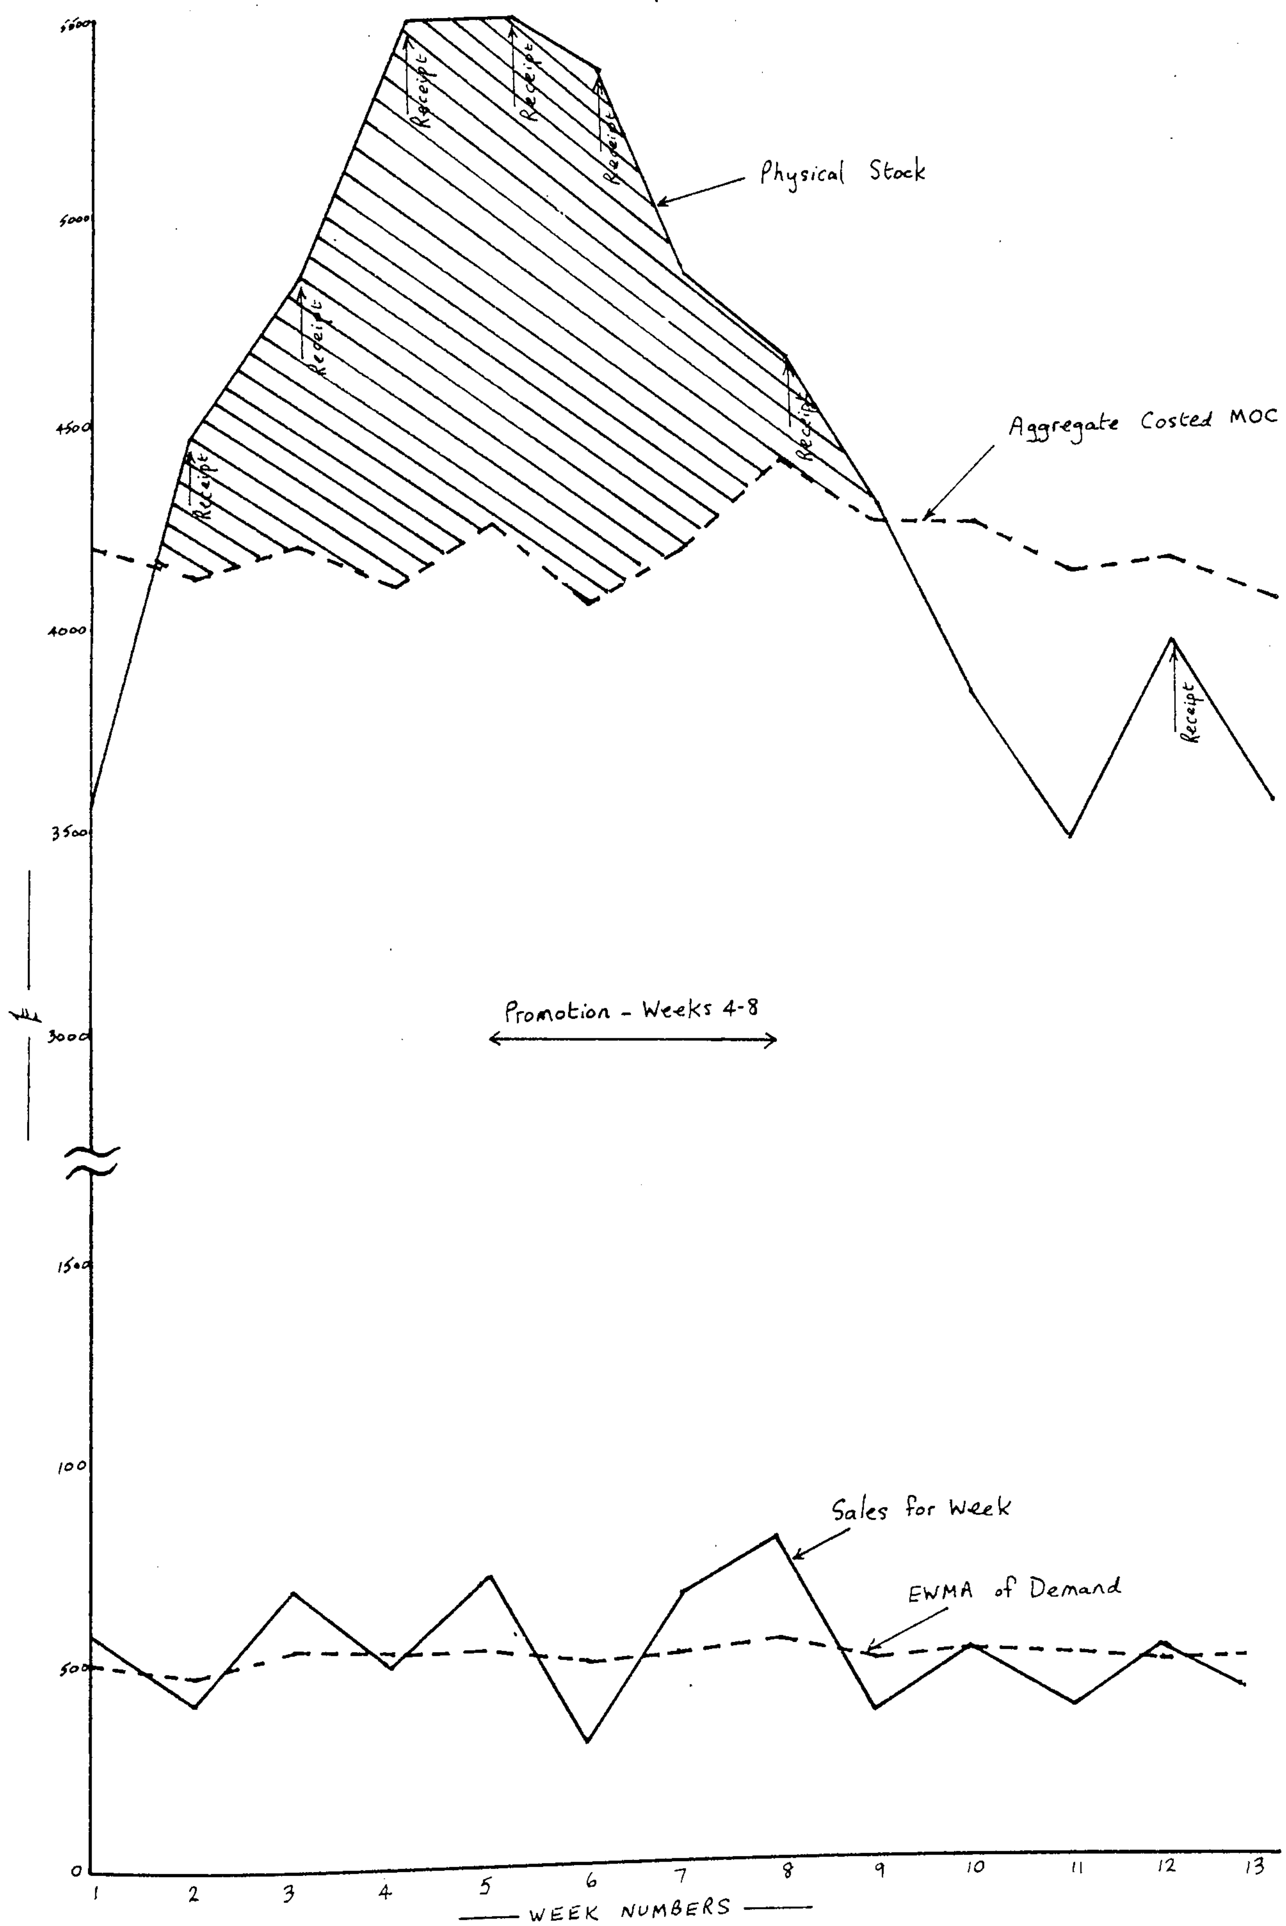

Fig. 7.5 Example of Sales Promotion

$$\frac{374807}{5040000} \times 100$$

$$= \underline{7.44\%}$$

If it is assumed that the excess stock provides 100% service level for the appropriate products, and the products without an excess achieve the overall target of 95%, then the expected service level excluding the products which are in excess for other reasons is:

$$\frac{(100 \times 1468) + (95 \times (16690 - (7076 - 1468)))}{(1468 + (16690 - (7076 - 1468)))}$$

$= \underline{95.58\%}$ (N.B. For strict accuracy the items should be costed)

An example of a sales promotion for a family of 21 products is given in Fig. 7.5. This clearly shows that whilst the buyer managed the procurements extremely well by obtaining regular receipts just before and during the promotional period then curtailing them immediately afterwards, a considerable overstock was carried over a period of eight weeks. This is because the promotion was almost totally ineffective. The buyer had no option but to procure sufficient supplies to meet the marketing estimates. Commercial realities dictate that this is the price paid for an attempt to improve the market share. As promotions are undertaken rotationally, a constant excess stock level of the order indicated is quite plausible without impugning the buyers in any way.

Special Deals

Special terms are sometimes offered by suppliers in respect of their own sales promotions, or perhaps because of their own over-stock problems. Here, excess stocks are justified only if they are self-financing over the overstock period, i.e.

Additional discount + extra profit arising from effective 100% service level + reduction in ordering costs > increase in stockholding costs

Given the financial parameters, this is a straightforward calculation, and it is not proposed to investigate the justification for the excess stock in this category. The effects on average stock and service level are, however, quantified on the same basis as for the sales promotions:

$$\text{Increase in AS} = \frac{115325}{5040000} \times 100$$

$$= \underline{2.29\%}$$

$$\text{Overall SL} = \frac{(100 \times 524) + (95 \times (16690 - (7076 - 524)))}{(524 + (16690 - (7076 - 524)))}$$

$$= \underline{95.25\%}$$

Supplier Rebate

Supplier rebates are normally obtained annually if all Trading Units combined purchase an agreed goods value. Towards the end of the year, if the purchases to date fall short of the agreed value, Distribution Centres are often directed to make bulk purchases so that the Organisation obtains the rebate. Again, such parameters should be self-financing on the basis of a calculation similar to that for special deals. The effects on average stock and service level are:

$$\text{Increase in AS} = \frac{86494}{5040000} \times 100$$

$$= \underline{1.72\%}$$

$$\text{Overall SL} = \frac{(100 \times 466) + (95 \times (16690 - (7076 - 466)))}{(466 + (16690 - (7076 - 466)))}$$

$$= \underline{95.22\%}$$

Initial Stocking

When a new product line is introduced into the range of a Distribution Centre, sufficient supplies are procured for the Distribution Centre and all of its Branches. It is, however, tested in a few Branches before being generally introduced in order to minimise shelf re-merchandising activity. The initial stocks can therefore reside in the Distribution Centre for a considerable period before being dispersed. Also, if the buyer overestimates the sales potential, the Branches will reduce their shelf stock allocations accordingly and the residue will remain at the Distribution Centre. During the build-up of sales over the first few months, the computer will have little or no sales history, therefore most of the stocks will appear as excess. The effects on average stock and service level are:

$$\text{Increase in AS} = \frac{320354}{5040000} \times 100$$

$$= \underline{6.36\%}$$

$$\text{Overall SL} = \frac{(100 \times 1297) + (95 \times (16690 - (7076 - 1297)))}{(1297 + (16690 - (7076 - 1297)))}$$

$$= \underline{95.53\%}$$

Improved Discount

The minimum order value necessary to receive the most favourable supplier terms is the basis for calculating the computer 'Alternative Order'. Even though this is a computer order, the product order quantities added to the physical stock plus on-order balances will exceed the normal MOCs. Here the system informs the buyer that he must obtain n.nn% additional discount to justify moving from the recommended to the alternative order i.e. the self-financing formula is included in the order generation program. As approximately 10% of all recommended orders are accompanied by an alternative order and nearly all of these can be justified by a small additional discount, the excess value for this category is consistent with the buyers following the computer guidelines. The effects on average stock and service level are:

$$\text{Increase in AS} = \frac{247924}{5040000} \times 100$$

$$= \underline{4.92\%}$$

$$\text{Overall SL} = \frac{(100 \times 976) + (95 \times (16690 - (7076 - 976)))}{(976 + (16690 - (7076 - 976)))}$$

$$= \underline{95.42\%}$$

The implications on system performance of commercially-motivated bulk overrides are summarised in Table 7.12.

Reason for Excess	Increase in AS (%)	Expected SL (%)
Sales Promotion	7.44	95.58
Special Deal	2.29	95.25
Supplier Rebate	1.72	95.22
Initial Stocking	6.36	95.53
Improved Discount	4.92	95.42

Table 7.12 Summary of Effects of Commercial Actions on Service Level and Average Stock

As the effects of these measures are mutually exclusive, they constitute a major reason for the differences between the expected and achieved system performance results.

7.7 EFFECTS ON PERFORMANCE OF DATA INACCURACY

System performance is clearly directly dependent upon the timing of the placement of purchase orders. If the computer stock records, in particular, are inaccurate, then orders will be placed at the wrong time, and discrepancies between the predicted and achieved performance will inevitably result. This Section evaluates the likely magnitude of these errors. For clarity, the effects of joint ordering are removed by examining the products independently, and by using reorder levels in units instead of the ELS equivalents.

A sample of 104 products was selected for monitoring the accuracy of the computer stock balances at intervals of approximately one week. Unfortunately, several of these products were either removed from the stock range or else they moved so slowly that continued monitoring was not worthwhile. Consequently, the following analysis is based on a reduced sample of 51. The results of one check are given in Table 7.13, from which it can be seen that the computer records were correct in only 7 cases (13.7%). Several such checks revealed that the discrepancies tend to remain at about the same level, which suggests that the transaction data is delayed rather than lost completely. The reasons for the delays have been investigated but they are not reported here.

Table 7.13 shows the effects of the data inaccuracies on the expected service level and average stock. An example, using product A, is worked through to illustrate the method.

μ_D = 37.80 units/week

σ_D = 24.68 units/week

μ_L = 13.00 days

σ_L = 5.32 days

Product	μ_D (units/ week)	Phys. Stock (units)	Comp. Stock (units)	Phys. - Comp. (units)	Expected SL (%)	% Inc. in AS over theoretical
A	37.80	115	117	-2	94.77	-1.52
B	49.38	275	275	0	95.00	0
C	4.88	28	29	-1	94.28	-4.38
D	14.88	147	144	3	95.66	4.53
E	20.05	240	265	-25	88.20	-27.87
F	27.91	345	329	16	96.72	13.45
G	7.24	195	223	-28	35.44	-86.83
H	7.09	22	21	1	95.53	3.44
I	1.72	6	4	2	97.22	18.71
J	0.63	23	22	1	97.66	25.05
K	81.20	1157	1237	-80	88.46	-27.48
L	21.04	191	300	109	99.95	136.48
M	9.72	130	171	-41	0.00	-
N	32.85	377	432	-55	84.22	-37.61
O	3.34	42	65	-23	0.00	-
P	0.93	50	50	0	95.00	0
Q	4.98	146	127	19	99.51	73.28
R	44.47	88	147	-59	84.97	-36.41
S	7.75	14	14	0	95.00	0
T	5.69	42	41	1	95.66	4.46
U	4.30	33	34	-1	94.10	-5.46
V	26.19	23	80	-57	70.26	-60.28
W	24.35	102	146	-44	78.55	-48.58
X	1.21	28	25	3	97.64	29.64
Y	0.46	22	24	-2	83.24	-48.43
Z	0.51	13	10	3	99.01	57.20
a	0.39	40	36	4	99.81	127.39
b	0.40	24	24	0	95.00	0
c	1.55	26	28	-2	92.04	-18.28
d	0.92	19	19	0	95.00	0
e	3.96	6	4	2	96.49	12.48
f	1.42	9	5	4	97.86	33.42
g	1.93	32	50	-18	0.00	-
h	2.40	0	-19	19	99.83	119.57
i	1.52	0	0	0	95.00	0
j	1.62	0	-2	2	96.42	13.35
k	0.81	32	34	-2	89.21	-30.53
l	6.65	27	22	5	96.92	17.04
m	7.44	19	34	-15	84.85	-41.18
n	6.70	41	53	-12	88.03	-32.14
o	1.85	65	73	-8	73.69	-62.56
p	5.30	100	96	4	96.76	15.29
q	6.10	123	135	-12	86.01	-38.08
r	0.74	15	15	0	95.00	0
s	0.64	0	2	-2	85.76	-38.80
t	3.10	71	73	-2	93.44	-9.81
u	11.30	54	55	-1	94.73	-1.85
v	6.27	37	44	-7	90.11	-24.72
w	5.34	139	141	-2	94.02	-6.49
x	3.35	145	153	-8	89.31	-25.18
y	3.12	67	63	4	97.67	26.45

Table 7.13 Effects of Physical Stock Discrepancies

$I = 5$ days

Target SL $= 95\%$

Then, $\sigma_{LD} = \sqrt{\mu_L \sigma_D^2 + \mu_D^2 \sigma_L^2}$ (ignoring the review cycle)

$$= \sqrt{\frac{13}{5} \times 24.68^2 + 37.80^2 \times \frac{5.32^2}{5^2}}$$

$$= 56.5796 \text{ units}$$

$$k_{LD} = \frac{\mu_D^2 \mu_L^2}{\sigma_{LD}^2}$$

$$= \frac{37.80^2 \times 13^2}{56.5796^2 \times 5^2}$$

$$= 3.0172$$

$$PLS = (1 - SL) \frac{I}{\mu_L}$$

$$= (1 - 0.95) \times \frac{5}{13}$$

$$= 0.0192$$

U_{LD} for $k_{LD} = 3.0172$ and $PLS = 0.0192$ is 3.7303

And, $R = U_{LD} \sigma_{LD}$

$$= 3.7303 \times 56.5796$$

$$= 211.06$$

AS (Sect. 3.3 equation 4 with T = 0)

$$= 211.06 - 37.80 \times \frac{13}{5} + \frac{37.80 \times 5}{2 \times 5}$$

$$= 131.68$$

If it is assumed that the stock discrepancy of -2 units is still present when the stock falls to the reorder level (this is the best estimate available), then the physical stock (R') will be 209.06 (actually 209, but the fraction is retained to avoid quantising errors) when the order is placed.

$$\text{Corresponding } U'_{LD} = \frac{R'}{\sigma_{LD}}$$

$$= \frac{209.06}{56.5796}$$

$$= 3.6950$$

And, PLS for $k_{LD} = 3.0172$ and $U'_{LD} = 3.6950$ is 0.0201

$$SL = 1 - \left(0.0201 \times \frac{13}{5}\right)$$

$$= 0.9477 \text{ or } \underline{94.77\%}$$

$$\text{And } AS' = 209.06 - 37.80 \times \frac{13}{5} + \frac{37.80 \times 5}{2 \times 5}$$

$$= 129.68 \text{ units}$$

$$\therefore \text{ Increase in AS} = \left(\frac{129.68 - 131.68}{131.68}\right) \times 100$$

$$= \underline{-1.52\%}$$

In this particular case the errors are small, but it can be seen from Table 7.13 that they can be very substantial where the stock discrepancy is large relative to the demand rate.

The costed expected service level for the sample,

$$\sum_{i=1}^{i=51} \frac{C_i \; SL_i \; \mu_{D_i}}{C_i \; \mu_{D_i}} \qquad \text{evaluates as } 90.20\%,$$

and the overall percentage increase in costed average stock is -5.23%. This indicates that with this particular sample, approximately double the predicted amount of lost sales could be expected to occur without a commensurate reduction in stock. Whilst there is some bias towards negative discrepancies (26 negative, 18 positive, 7 zero), it is nevertheless apparent that a degraded service level could be expected for the Distribution Centre as a whole. This is because a positive error can produce at best a 100% service level, which gives 5.3% (i.e. 5/95) extra sales, whereas a negative error can result in the whole of the predicted sales being lost. It is also clear that positive and negative discrepancies of an equal magnitude will not produce fully-compensated results. The positive error will give rise to a large stock increase relative to the improvement in service level, whereas the negative error will tend to cause a large service level deterioration for a relatively small stock saving. The combined result is a degradation in stockholding efficiency.

CHAPTER 8

CONCLUSIONS

8. CONCLUSIONS

8.1 The system has been well received by the management and buyers, but a measure of disappointment is felt that the performance predictions were never fully realised. The overall service level was raised from a totally inadequate level to one which has settled to within what the management regard as a reasonable tolerance of a commercially-acceptable target. The stock investment, however, has shown only a relatively small movement towards the greatly reduced level predicted.

Numerous reasons for performance variations have been demonstrated, some of which cause positive service level errors and some negative. It is important to recognise, however, that fully-compensated service level errors carry a stock overhead. (To take a simple example, if two identical products have a 95% service level target and they achieve 99% and 91% respectively, this will require substantially more stock than if both achieve the target). Also, errors of any sort usually result in a deterioration in stockholding efficiency, i.e. a service level achieved in error will almost certainly result in more stock than if the same service level was targetted and achieved without error. It is therefore quite feasible that the achieved service levels adhere fairly closely to the targets without a commensurate adherence by the stock investment.

8.2 The performance variations due to human actions were demonstrated to be of a much higher order of magnitude than those due to technical factors. It was shown in Section 7.2 that when the system is allowed to operate in a reasonably stable manner, it is possible with buyer intervention to halve the lost sales from the theoretical ELS without affecting the stock level. Similarly, regular bulk buying practices for various reasons were shown in Section 7.6 to cause large accumulations of excess stock. In contrast, with two exceptions (which are discussed in Section 8.11), the effects on performance of theoretical factors are small, and numerical approximations are shown in Section 6.2 to be inconsequential. These findings provide objective support for the views expressed by the proponents of redirecting more Operational Research efforts towards implementation activities, as discussed in Section 1.3.

8.3 The study produced strong evidence that a buyer/computer combination is potentially capable of substantially outperforming either when acting independently. The computer, operating without human intervention, produces results which approximate the predictions, but with a sacrifice of the commercial advantages of opportunistic buying. The buyer on his own produces the sort of results indicated in Section 3.5 before the computer system was introduced. The potential of the buyer/computer team is evaluated with respect to lead time forecasting in Section 7.3. Near-optimal results should be achievable in practice, as they are shown to be fairly insensitive to the accuracy of the overrides, provided the buyers err on the side of conservatism. Judicious advance purchasing to meet known periods of high demand is seen in Section 7.2 to have a similarly beneficial effect.

The task allocation is clear: the buyer provides an intuitive estimation of factors which have no known history and therefore cannot be analysed by statistical techniques, and the computer provides from a wide base of data projections of future probabilities whose complexity is such as to defy mental analysis. This division of responsibility accords with the principles advocated by several authors, as discussed in Chapter 2.

The modus operandi is that the buyer interrupts the homeostatic regulation exercised by the algorithm only when he has information about particular circumstances which is not available to the computer. The injection of this information into the algorithm reduces uncertainty, which enables the stockholding efficiency to be improved. The perpetuation of this mode of action should result in an achievement which surpasses that of the algorithmic prediction which is itself based upon the full range of uncertainty experienced in the past.

8.4 The use of probabilistic buyer estimates could be expected to further raise the level of achievement indicated in Section 8.3 by refining the buyer's input. Whilst the results of the exploratory exercise (Section 7.4) were somewhat disappointing, there were glimpses of some very promising possibilities which could be developed for practical application. The failure to obtain favourable results is attributed mainly to the Buyers' lack of a sound understanding of probabilities, and their consequent tendency to assign undue certainty to their estimates. Much better results would be expected from graduates in a numerate discipline with equivalent contextual knowledge.

In the best case it was seen that a reasonable spread of the Buyer estimate reinforces a good algorithmic estimate, whereas it also has a correcting influence on a bad algorithmic estimate. At the same time the algorithmic estimate provides cover for the Buyer estimate in the regions dictated by the historical pattern of incidences (i.e. the probability distribution). This principle of 'Buyer helps computer' at the same time as 'computer helps Buyer' gives the true flavour of a symbiotic partnership. This area, however, needs further research before it could be considered for practical application.

8.5 The failure to transfer the satisfactory course results to the operational system is considered to be due to two main causes:-

i) The generous amount of feedback given to the buyers on the course was not continued at the same level in the operational system. On the course, the direct and indirect effects of every action were examined in depth and the possible effects of alternative actions were also contemplated. Hence there was a prompt reinforcement of the actions which produced successful outcomes. In the operational system, a good balance of reports is provided but they do not in general indicate the direct effects of specific actions. In retrospect, more feedback should have been built into the system, e.g. buyer and computer lead time forecasts could have been compared with actuals, promotional sales and stocks could have been monitored, stockouts resulting from cancelling or reducing recommended orders could have been reported, etc. All of these additional features obviously carry a computing overhead, and consequently there are commercial limits to the amount of feedback provided. However, in this instance the provision of more direct feedback is justified.

ii) The operational system was subjected to more 'shocks' than the simulated system. Whilst sales promotions and supply failures were included in the simulation, the extent of these and other disruptive forces were not fully representative of the real situation. Certainly, opportunistic buying cannot be suppressed, as that would blunt any competitive advantage. Yet the practice of bulk buying was shown to disrupt the control system and gave rise to large performance variations.

The problem of reconciling commercial buying with a self-regulating algorithm is not easy to resolve. One fascinating area of research might be to draw a cybernetic analogue with the way the human body reconciles the exterofective activities of the upper regions of the brain with the basal homeostasis which is governed by the autonomic nervous system and endocrine system. Certainly, in nature, the homeostasis is not destroyed. A more prosaic solution would be to extend the education process beyond the staff who are directly involved in using the system. Many of the commercial issues are handled by central marketing and negotiating teams who have no direct responsibility for the stock-holding consequences. A greater understanding of, and commitment to, the system by these groups could only be beneficial.

8.6 The system was found to possess profound behavioural characteristics even when allowed to operate under the control of the algorithm without human intervention. Many of these characteristics are a function of the dynamics of the system and its interaction with the internal and external environment. The following are examples of system behaviour which could not have been foreseen with confidence without a System Dynamics exercise:-

a) A moderate step increase in demand of 20% of the base value produces an Expected Lost Sales peak of more than double the original value, and the service level takes around six months to stabilise (Sect. 4.1.4.1).

b) Step reversals of demand induce 'backlash' effects which result in greater deviations in service level and stock level than unidirectional steps of the same magnitude (Sect. 4.1.4.4).

c) The onset of a linear downward trend causes stock to increase over the first three months before it starts to decline (Sect. 4.1.4.5).

d) The continuation of a linear downward trend gives a stockholding performance which is permanently better than the predictions (Sect. 4.1.4.5).

e) With a cyclic demand pattern, the stock levels may be completely out of phase with the demand levels. Hence stock is most readily available when it is least needed (Sect. 4.1.4.6).

Insofar as the performance predictions are concerned, the dynamic effects would be incorporated into the EWMAs as a hidden component. This may cause a discrepancy between prediction and actual, depending upon whether the causal factors recur with sufficient regularity to be absorbed in the aggregation process. Whether or not they recur, they are identifiable elsewhere in the Organisation before they are recorded as orders on the Distribution Centre. Hence, they could conceivably be incorporated into the forecasts by employing a 'feedforward' process in the forecasting mechanism. This should reduce uncertainty and thus improve stockholding efficiency. This is another area which would benefit from further investigation.

8.7 The conceptual framework, mathematics and dynamic behaviour of the system render it an indisputably complex package to comprehend. The numerous facets which have been examined in this work underline the different order of com-lexity between this type of indeterminate system and the determinate systems which comprised most early computer applications. Moreover, the dynamic effects are specific to the organisation - its structure, internal and external environment, operating procedures, and communication media. It follows that a full understanding of the behaviour of the system can only be gained in situ. Learning solely from the experiences of others working in different organisations and circumstances will not be adequate. This intellectual challenge to the Organisation has two extremely important implications - the Organisation's attitude to Research and Development expen-diture, and the Personnel Specification for the buyers - which will be discussed in Sections 8.8 and 8.9 respectively.

8.8 Most commercial establishments are reluctant to spend much time and effort on R & D work unless it is directly related to product development. A common attitude to the introduction of computer systems is to purchase a 'package' from a computer manufacturer or software house, perhaps add a few custom-built management reports, and install it according to the adage that "it will meet 90% of the requirements for minimum cost." This may well be a sound strategy for determinate systems, but all of the indications in this study are that it is a false economy for complex indeterminate systems.

There is no doubt that the system studied produced immediate benefits over the previous manual system, but when it was first introduced none of the problems uncovered by this study were recognised, and operational performance was well below its potential. As the project followed a typical course, the indications are that development effort is usually cut off too soon, and that post-operational development might be expected to be of the same order of magnitude as pre-operational development. Certainly a sound in-house knowledge of the system in situ is essential, and this should be experientially supported.

8.9 A higher grade of buyer is required to fulfill the demands of the system. Without implying any derogation, it is evident that the mental constructs of the buyers are unsuitable for managing the system efficiently. Though basically numerate, they were stretched to understand the basic statistical concepts (which is not surprising as they were originally employed to replenish stock according to some very simple rules). Moreover, this study has amply demonstrated that an understanding of the basic concepts is not enough. There are phenomena which sometimes run counter to intuitive expectations, as exemplified below, and it requires some flexibility of thought beyond a straightforward application of the concepts to explain them:-

a) An increase in mean lead time can reduce the average stock (Sect. 4.2, Fig. 4.21).

b) An increase in the interval between placing orders can reduce the average stock (Sect. 4.2, Fig. 4.25).

c) Consistent lead times are generally more important than short lead times (Sect. 4.2, Figs. 4.21 & 4.22).

Numerous similar examples were referred back to the design team for investigation. The answers to most of the queries were derivable by logical reasoning, but some required empirical verification. Quite often, the buyers were unconvinced by the explanations. For example, some buyers still override the computer because they do not understand the probability theory which underlies buffer stock calculations (Sect. 7.5).

It is not considered feasible to expect the existing buyers to acquire the specialised mental attributes now demanded of them. These attributes are partly innate and partly developed by an education in certain numerate disciplines. It is therefore unrealistic to expect these qualities to be acquired by a re-training programme, though this must be employed as a necessary short-term expedient. Training is not an adequate substitute for education.

In the longer term, persons with the mental attributes outlined above must be inducted and trained in buying skills. This may well mean employing graduates or equivalent in this capacity. This raises the moral issue - not debated here - of introducing automated systems which require élitist operating staff, with the resultant deterioration in job opportunities for the remainder of the population. Certainly the findings from this study do not support the popular conception of automatic systems having a de-skilling effect on jobs. The cognitive elements of the original buying function are distilled by elimination of the mundane stock replenishment tasks and the introduction of taxing intellectual challenges.

8.10 The traditional methods of training were found to be inappropriate for this type of system. Much of the expertise needed to manage the system efficiently was gained post-operationally. In the circumstances the problems were referred back to the design team for investigation, and any lessons learned were disseminated by memorandum or verbal communication. It would be much better for buyers to be given simulation tools to carry out their own investigations. Periodic experience-sharing sessions could then be convened which would be much more valuable than the traditional 'classroom' techniques where the Trainer is expected to provide all of the solutions. In indeterminate systems, it is not possible to provide a set of definitive instructions which cover all eventualities. As circumstances change, operator actions must adapt, and a much greater degree of self-reliance is called for.

8.11 It is significant that the only two theoretical factors which gave rise to serious performance variations - overlapping lead times and demand autocorrelation - are nearly always treated as assumptions in standard textbooks. Further research is required to provide generalised solutions to these problems. (In this work solutions were devised which are specific to the system). In the meantime it is recommended that proper recognition is given to the importance of these problems even though satisfactory solutions do not currently exist. There is a tendency for students, in particular, to regard assumptions as inconsequential and use the results regardless.

8.12 Performance variations due to computer hardware, software and operations were found to be virtually non-existent.

Since the inception of the system there has not been a single mainframe failure resulting in a downtime of more than 10 minutes. Occasional failures in remote terminals did occur, but they were too infrequent to have had any tangible effect on system performance. A log of system design and programming errors revealed that after the initial 'settling-in' period when the system was loaded up, there were very few errors which could conceivably have caused performance variations. Computer operational errors did sometimes cause delays in printing recommended orders, and occasionally a day was missed out altogether. The consequential delay in sending out purchase orders may have had a marginal effect on stockouts, but again the occurrences were of insufficient frequency to produce a measurable service level degradation.

8.13 The data errors evaluated in Section 7.7 were found to produce potentially serious performance variations, with a strong tendency to degrade service levels. These errors are caused primarily by the time lags associated with a batch updating system, and by physical stock counting errors. Neither of these problems are easy to rectify: at least 30 data collection terminals would be required in each Distribution Centre to provide a real-time updating facility; and the layout of the storage areas makes counting irregularly-shaped products extremely difficult from ground level. It is also perhaps relevant that the Operatives responsible for the submission of most of the data are not the main beneficiaries from the system.

These rather prosaic problems are faced by most data processing installations. The evaluations in Section 7.7 confirm that they are a major cause of system performance degradation, and as such they warrant constant attention.

8.14 Finally, an observation made in Section 1.5 is returned to - namely that the true behaviour of the system would never have been fully understood without the degree of analytical and experimental work here undertaken. It would be reasonable to question the justification for the inbuilt complexity. The type of heuristics described in Section 1.2 may well have produced similar results with a much simpler mechanism.

The author is of the opinion that, with some qualification (mainly in respect of the need for a refined optimisation function), the complexity is a feature of the requirement, and a complex requirement must be met with a complex solution. Simplifying assumptions bring their own problems. The complex situation has to be understood to determine if the simplifying assumptions are valid. If they are not valid, then the wrong problem is being addressed. Even if they are valid, they may not remain so indefinitely. In the system studied, the states of the operating variables were shown to be both time-dependent and situation-dependent.

It has been demonstrated that much of the difficulty in comprehension is caused by the way the system behaves in situ, and this would not be greatly changed by substituting it with a simple reorder level system. There is a need for both the designers and buyers to gain sufficient exposure to the way the system reacts to different conditions, and to understand what 'makes it tick'. This cannot be gained from any Manual. It is highly desirable for buyers to become what Singleton (37) has termed 'Concept Operators'. It is seen that limited success is being achieved by buyers acting as 'Programmed Operators', but, by definition, a full set of predetermined responses cannot be specified for indeterminate systems. In order to get the best out of the total system, the buyers must fully understand the concepts underlying the algorithm with which they interact. Only then could the positive outcome from the interventions be expected to outweigh the multiplicity of dysfunctional factors which the system encounters. Until that state is attained, the predictions will not be matched by the achievements.

<u>APPENDICES</u>

Glossary of Mathematical Symbols

α	Gamma shape factor
AS	Average stock level
C	Item cost per unit
ELS	Expected Lost Sales (units, unless otherwise stated)
EWMA	Exponentially-weighted Moving Average
f(x)	Probability density function of demand in lead time
I	Order Interval
k	Gamma modulus
k_D	Gamma modulus for demand per unit time
k_L	Gamma modulus for lead time
k_{LD}	Gamma modulus for demand in lead time
λ	Lagrange multiplier
M	Gross Margin
μ_D	Mean demand per unit time
μ_L	Mean lead time
μ_{LD}	Mean demand in lead time
MAD	Mean Absolute Deviation
MAD_D	Mean Absolute Deviation of demand per unit time
MAD_L	Mean Absolute Deviation of lead time
MAD_{LD}	Mean Absolute Deviation of demand in lead time
MOC	Maximum Order Cover
MSL	Maximum Stock Level
P	Protection Level
PLS	Percentage Lost Sales (in lead time unless otherwise stated)
R	Reorder level

σ_D	Standard deviation of demand per unit time
σ_L	Standard deviation of lead time
σ_{LD}	Standard deviation of demand in lead time
SL	Service Level
T	Review cycle
U	Standardised reorder level
V_D	Variance of demand per unit time
V_L	Variance of lead time
V_{LD}	Variance of demand in lead time
x	Demand in lead time variate (unless otherwise stated)
y	Stock cover (physical plus on-order)

Computer Routines to Calculate Gamma Functions

(BASIC Notation)

1. To calculate Γ (k), given k:

```
10 INPUT "K";K
20 A1 = -.57710166: A2 = .98585399: A3 = -.87642182: A4 = .8328212
30 A5 = -.5684729: A6 = .25482049: A7 = -.0514993
40 IF K = 1 THEN GK = 1: GØTØ 190
50 KI = INT(K)
60 KF = K-KI
70 IF = KF = 0 THEN 130
80 GP = 1 + A1*KF + A2*KF↑2 + A3*KF↑3 + A4*KF↑4 + A5*KF↑5
90 GP = GP + A6*KF↑6 + A7*KF↑7
100 GF = GP/KF
110 IF KI = 0 THEN GK = GF: GØTØ 190
120 IF KI = 1 THEN GK = GP: GØTØ 190
130 GK = 1
140 FØR I = 1 TØ KI - 1
150 GK = GK*(K - I)
160 NEXT I
170 IF KF = 0 THEN 190
180 GK = GK*GP
190 PRINT GK
200 END
```

2. To calculate P given U, k and Γ (k):

```
10 INPUT "K";K
20 INPUT "GAMMA K"; GK
30 INPUT "U"; U
```

```
40  INPUT "TERMS IN SERIES"; M
50  RK = SQR(K)
60  A = (RK*U)↑(K - 1)*EXP(-RK*U)/GK
70  SX = O
80  FØR N = O TØ M
90  PI =
100 FØR R = O TØ N
110 PI = PI*(K + R)
120 NEXT R
130 X = (RK*U) ↑(N + 1)/PI
140 SX = SX + X
150 NEXT N
160 P = A*SX: PRINT P
170 END
```

3. To calculate PLS given U, k, $\Gamma(k)$ and P:

```
10  INPUT "K";K
20  INPUT "GAMMA K"; GK
30  INPUT "U"; U
40  INPUT "PROTECTION LEVEL"; P
50  RK = SQR(K)
60  PLS = (1-U/RK)*(1-P)+(U*RK)↑K*EXP(-U*RK)/K/GK: PRINT PLS
70  END
```

List of Mnemonics Used in System Dynamics

ACD Actual unfilled Customer orders at Distribution Centre

AIB Actual Inventory at Branch

AID Actual Inventory at Distribution Centre

APB Actual level of Pipeline orders at Branch

APD Actual level of Pipeline orders at Distribution Centre

ARD Actual unfilled Replenishment orders at Distribution Centre

ATD Actual total (Branch + Cust.) unfilled orders at Distribution Centre

AUB Actual Unfilled customer orders at Branch

CRB Customer orders Received at Branch

CRD Customer orders Received at Distribution Centre

DFB Delay in Filling orders at Branch

DFD Delay in Filling orders at Distribution Centre

DHB Delay in Handling orders at Branch

DHD Delay in Handling orders at Distribution Centre

DIB Delay in adjusting Inventory at Branch

DID Delay in adjusting Inventory at Distribution Centre

DLS Delay due to Lead Time at Supplier

DPD Delay in placing Purchase orders at Distribution Centre

DRB Delay in placing Replenishment orders at Branch

DSB Delay due to Stockouts at Branch

DSD Delay due to Stockouts at Distribution Centre

DTB Delay in Transporting materials to Branch

DTD Delay in Transmitting replenishment orders to Distribution Centre

MDB Materials Despatched from Branch

MRB Materials Received by Branch

MRS Materials Received from Supplier

MSD Materials Sent to customers from Distribution Centre

MTB Materials in Transit to Branch

NB Number of Branches per Distribution Centre

NIB Normal Inventory at Branch

NID Normal Inventory at Distribution Centre

NPB Normal Pipeline levels Branch/Distribution Centre

NPD Normal Pipeline levels Distribution Centre/Supplier

NUB Normal Unfilled orders at Branch

NUD Normal Unfilled orders at Distribution Centre

POD Purchase Orders awaiting placement at Distribution Centre

POS Purchase Orders in process at Supplier

PSS Purchase Orders Sent to Supplier

RCB Replenishment orders in storage on Computer from Branch

ROB Replenishment orders awaiting placement at Branch

RPD Required Purchasing rate by Distribution Centre

RRB Required Replenishment rate by Branch

RRD Replenishment orders Received by Distribution Centre

RSD Replenishment materials Sent by Distribution Centre

RTB Replenishment orders Transmitted from Branch

SDB Smoothed Demand at Branch

SDD Smoothed Demand at Distribution Centre

SFB Smoothing Factor for Branch demand

SFD Smoothing Factor for Distribution Centre demand

SRB Standardised Reorder Level for Branch

SRD Standardised Reorder Level for Distribution Centre

TRD Total orders Received by Distribution Centre

Computer Routine for Generating Gamma Random Variates

(BASIC Notation)

```
010  K = MEAN*MEAN/VAR
020  T = K-1
030  LAMBDA = MEAN/VAR
040  IF T < = 0 THEN 200
050  MU = (-1 + SQR(1 + 4 * T))/2/T
060  R = (T *(1-MU)↑ T)/(T *(1-MU)↑ T + 1/MU * EXP(-MU*T))
070  R1 = RND(RN)
080  R2 = RND(RN)
090  E = -LØG(R2)
100  IF R1 >R THEN 140
110  X = R1*T/R
120  IF T * LØG(T/X) - T + X <= E THEN 180
130  GØTØ 070
140  R3 = RND(RN)
150  X = -1/MU * LØG ((1-R3) * EXP (-MU*T))
160  IF E >= T * LØG (T/(1-MU)/X) + (1-MU) * X - T THEN 180
170  GØTØ 070
180  GAMMAVAR = X/LAMBDA
190  GØTØ 320
200  R1 = RND (RN)
210  R2 = RND (RN)
220  E = -LØG (R2)
230  P = R1 * (K + EXP(1))/EXP(1)
240  IF P > 1 THEN 280
250  X = P ↑ (1/K)
260  IF X <= E THEN 310
270  GØTØ 200
280  X = -LOG ((K + EXP(1) - P * EXP(1))/K/EXP(1))
290  IF (1-K) * LØG (X) <= E THEN 310
300  GØTØ 200
310  GAMMAVAR = X/LAMBDA
320  END
```

Notes

1. 'RN' is a random number generated directly from the seed

2. Statements 050 - 180 execute Atkinson's algorithm: statements 200 - 310 execute Ahren's algorithm.

IM WEEKLY ROUTINES

WEEK NO 2

UPDATE EWMA OF DEMAND

```
A     0.15 *  14   +  0.85 *   8.30  =    9.15
B     0.15 * 111   +  0.85 * 126.55  =  124.22
C     0.15 *  46   +  0.85 *  36.48  =   37.90
D     0.15 *  91   +  0.85 *  78.84  =   80.66
E     0.15 *   5   +  0.85 *   6.25  =    6.06
```

UPDATE EWMA OF MAD OF DEMAND

```
A     0.15 *   5.70  +  0.85 *   3.22  =    3.60
B     0.15 *  15.55  +  0.85 *  31.11  =   28.78
C     0.15 *   9.52  +  0.85 *  11.85  =   11.50
D     0.15 *  12.16  +  0.85 *  21.06  =   19.73
E     0.15 *   1.25  +  0.85 *   5.75  =    5.07
```

UPDATE AVERAGE STOCK LEVEL

```
A     0.15 *  44   +  0.85 *  66.50  =   63.12
B     0.15 * 784   +  0.85 * 985.95  =  955.66
C     0.15 * 121   +  0.85 * 188.25  =  178.16
D     0.15 * 181   +  0.85 * 358.70  =  332.04
E     0.15 *  68   +  0.85 *  74.70  =   73.70
```

RECALCULATE CRITICAL LOST SALES VALUE

```
A              £   6.51
B              £  37.54
C              £   7.31
D              £  18.90
E              £  43.73
                -------
        TOTAL  £114.00
                -------
```

IM DAILY ORDER REVIEW

WEEK NO 3

DAY	PRODUCT	PHYS. STOCK	OUTSTANDING ORDERS					EXPECTED LOST SALES
1	A	44	0	0	0	0		£ 2.74
1	B	781	0	0	0	0		£ 1.93
1	C	95	0	0	0	0		£ 20.73
1	D	175	0	0	0	0		£ 77.51
1	E	68	0	0	0	0		£ 1.52
							TOTAL	£104.43
2	A	43	0	0	0	0		£ 3.01
2	B	752	0	0	0	0		£ 2.66
2	C	86	0	0	0	0		£ 26.03
2	D	158	0	0	0	0		£ 97.29
2	E	67	0	0	0	0		£ 1.65
							TOTAL	£130.64 > PLACE ORDER NO 1
3	A	38	8	0	0	0		£ 2.25
3	B	723	0	0	0	0		£ 3.59
3	C	73	120	0	0	0		£ 1.93
3	D	129	280	0	0	0		£ 4.49
3	E	67	0	0	0	0		£ 1.65
							TOTAL	£ 13.91
4	A	38	8	0	0	0		£ 2.25
4	B	668	0	0	0	0		£ 6.13
4	C	64	120	0	0	0		£ 2.45
4	D	99	280	0	0	0		£ 6.68
4	E	67	0	0	0	0		£ 1.65
							TOTAL	£ 19.16
5	A	38	8	0	0	0		£ 2.25
5	B	665	0	0	0	0		£ 6.30
5	C	54	120	0	0	0		£ 3.18
5	D	99	280	0	0	0		£ 6.68
5	E	67	0	0	0	0		£ 1.65
							TOTAL	£ 20.06

IM WEEKLY ROUTINES

WEEK NO 3

UPDATE EWMA OF DEMAND

A	0.15	*	6	+	0.85	*	9.15	=	8.68
B	0.15	*	119	+	0.85	*	124.22	=	123.43
C	0.15	*	67	+	0.85	*	37.90	=	42.27
D	0.15	*	82	+	0.85	*	80.66	=	80.86
E	0.15	*	1	+	0.85	*	6.06	=	5.30

UPDATE EWMA OF MAD OF DEMAND

A	0.15	*	3.15	+	0.85	*	3.60	=	3.53
B	0.15	*	5.22	+	0.85	*	28.78	=	25.24
C	0.15	*	29.10	+	0.85	*	11.50	=	14.14
D	0.15	*	1.34	+	0.85	*	19.73	=	16.97
E	0.15	*	5.06	+	0.85	*	5.07	=	5.07

UPDATE AVERAGE STOCK LEVEL

A	0.15	*	38	+	0.85	*	63.12	=	59.36
B	0.15	*	665	+	0.85	*	955.66	=	912.06
C	0.15	*	54	+	0.85	*	178.16	=	159.54
D	0.15	*	99	+	0.85	*	332.04	=	297.09
E	0.15	*	67	+	0.85	*	73.70	=	72.69

RECALCULATE CRITICAL LOST SALES VALUE

A	£ 6.27
B	£ 36.60
C	£ 8.40
D	£ 18.47
E	£ 43.83
TOTAL	£113.56

IM DAILY ORDER REVIEW

WEEK NO 4

DAY	PRODUCT	PHYS. STOCK	OUTSTANDING ORDERS					EXPECTED LOST SALES
1	A	36	8	0	0	0		£ 2.15
1	B	655	0	0	0	0		£ 6.11
1	C	28	120	0	0	0		£ 10.29
1	D	94	280	0	0	0		£ 6.69
1	E	66	0	0	0	0		£ 1.49
							TOTAL	£ 26.73
2	A	36	8	0	0	0		£ 2.15
2	B	634	0	0	0	0		£ 7.45
2	C	27	120	0	0	0		£ 10.50
2	D	82	280	0	0	0		£ 7.84
2	E	51	0	0	0	0		£ 4.96
							TOTAL	£ 32.90
3	A	34	8	0	0	0		£ 2.63
3	B	613	0	0	0	0		£ 9.04
3	C	25	120	0	0	0		£ 10.94
3	D	64	280	0	0	0		£ 9.89
3	E	51	0	0	0	0		£ 4.96
							TOTAL	£ 37.47
4	A	34	8	0	0	0		£ 2.63
4	B	610	0	0	0	0		£ 9.29
4	C	16	120	0	0	0		£ 13.14
4	D	42	280	0	0	0		£ 13.05
4	E	51	0	0	0	0		£ 4.96
							TOTAL	£ 43.08
5	A	34	8	0	0	0		£ 2.63
5	B	594	0	0	0	0		£ 10.73
5	C	5	120	0	0	0		£ 16.37
5	D	5	280	0	0	0		£ 20.50
5	E	51	0	0	0	0		£ 4.96
							TOTAL	£ 55.20

IM WEEKLY ROUTINES

WEEK NO 4

UPDATE EWMA OF DEMAND

A	0.15	*	4	+	0.85	*	8.68	=	7.98
B	0.15	*	71	+	0.85	*	123.43	=	115.57
C	0.15	*	49	+	0.85	*	42.27	=	43.28
D	0.15	*	94	+	0.85	*	80.86	=	82.83
E	0.15	*	16	+	0.85	*	5.30	=	6.91

UPDATE EWMA OF MAD OF DEMAND

A	0.15	*	4.68	+	0.85	*	3.53	=	3.70
B	0.15	*	52.43	+	0.85	*	25.24	=	29.32
C	0.15	*	6.73	+	0.85	*	14.14	=	13.03
D	0.15	*	13.14	+	0.85	*	16.97	=	16.39
E	0.15	*	10.70	+	0.85	*	5.07	=	5.92

UPDATE AVERAGE STOCK LEVEL

A	0.15	*	34	+	0.85	*	59.36	=	55.55
B	0.15	*	594	+	0.85	*	912.06	=	864.35
C	0.15	*	5	+	0.85	*	159.54	=	136.36
D	0.15	*	5	+	0.85	*	297.09	=	253.28
E	0.15	*	51	+	0.85	*	72.69	=	69.44

RECALCULATE CRITICAL LOST SALES VALUE

A		£ 6.14
B		£ 35.53
C		£ 8.33
D		£ 18.77
E		£ 50.98
	TOTAL	£119.75

IM DAILY ORDER REVIEW

WEEK NO 5

DAY	PRODUCT	PHYS. STOCK	OUTSTANDING ORDERS					EXPECTED LOST SALES
1	A	34	8	0	0	0		£ 2.05
1	B	506	0	0	0	0		£ 17.67
1	C	0	120	0	0	0		£ 18.88
1	D	0	280	0	0	0		£ 23.80
1	E	51	0	0	0	0		£ 12.22
							TOTAL	£ 74.62
2	A	29	8	0	0	0		£ 3.40
2	B	485	0	0	0	0		£ 21.18
2	C	0	120	0	0	0		£ 18.88
2	D	0	280	0	0	0		£ 23.80
2	E	51	0	0	0	0		£ 12.22
							TOTAL	£ 79.48
3	A	28	8	0	0	0		£ 3.76
3	B	397	0	0	0	0		£ 43.95
3	C	0	120	0	0	0		£ 18.88
3	D	0	280	0	0	0		£ 23.80
3	E	51	0	0	0	0		£ 12.22
							TOTAL	£102.60
4	A	27	8	0	0	0		£ 4.15
4	B	393	0	0	0	0		£ 45.39
4	C	0	120	0	0	0		£ 18.88
4	D	0	280	0	0	0		£ 23.80
4	E	51	0	0	0	0		£ 12.22
							TOTAL	£104.43
5	A	24	8	0	0	0		£ 5.56
5	B	362	0	0	0	0		£ 58.05
5	C	0	120	0	0	0		£ 18.88
5	D	0	280	0	0	0		£ 23.80
5	E	36	0	0	0	0		£ 34.47
							TOTAL	£140.76 > PLACE ORDER NO 2

IM WEEKLY ROUTINES

WEEK NO 5

UPDATE EWMA OF DEMAND

A	0.15	*	10	+	0.85	*	7.98	=	8.28
B	0.15	*	232	+	0.85	*	115.57	=	133.03
C	0.15	*	45	+	0.85	*	43.28	=	43.54
D	0.15	*	49	+	0.85	*	82.83	=	77.76
E	0.15	*	15	+	0.85	*	6.91	=	8.12

UPDATE EWMA OF MAD OF DEMAND

A	0.15	*	2.02	+	0.85	*	3.70	=	3.45
B	0.15	*	116.43	+	0.85	*	29.32	=	42.39
C	0.15	*	1.72	+	0.85	*	13.03	=	11.33
D	0.15	*	33.83	+	0.85	*	16.39	=	19.01
E	0.15	*	8.09	+	0.85	*	5.92	=	6.24

UPDATE AVERAGE STOCK LEVEL

A	0.15	*	24	+	0.85	*	55.55	=	50.82
B	0.15	*	362	+	0.85	*	864.35	=	789.00
C	0.15	*	0	+	0.85	*	136.36	=	115.90
D	0.15	*	0	+	0.85	*	253.28	=	215.28
E	0.15	*	36	+	0.85	*	69.44	=	64.42

RECALCULATE CRITICAL LOST SALES VALUE

A	£ 6.04
B	£ 43.27
C	£ 8.09
D	£ 18.22
E	£ 54.03
TOTAL	£129.66

IM DAILY ORDER REVIEW

WEEK NO 6

DAY	PRODUCT	PHYS. STOCK	OUTSTANDING ORDERS					EXPECTED LOST SALES
1	A	24	8	13	0	0		£ 1.54
1	B	316	0	270	0	0		£ 23.01
1	C	0	120	120	0	0		£ 1.22
1	D	0	280	170	0	0		£ 1.89
1	E	35	0	6	0	0		£ 36.16
							TOTAL	£ 63.83
2	A	22	8	13	0	0		£ 1.91
2	B	261	0	270	0	0		£ 33.77
2	C	0	120	120	0	0		£ 1.22
2	D	0	280	170	0	0		£ 1.89
2	E	35	0	6	0	0		£ 36.16
							TOTAL	£ 74.95
3	A	22	8	13	0	0		£ 1.91
3	B	241	0	270	0	0		£ 38.72
3	C	0	120	120	0	0		£ 1.22
3	D	0	280	170	0	0		£ 1.89
3	E	34	0	6	0	0		£ 38.60
							TOTAL	£ 82.34

RECT OF ORD NO 1 : UPDATED LT = 0.15 * 16 + 0.85 * 10.64 = 11.44
UPDATED MAD OF LT = 0.15 * 5.36 + 0.85 * 4.68 = 4.78

DAY	PRODUCT	PHYS. STOCK	OUTSTANDING ORDERS					EXPECTED LOST SALES
4	A	30	13	0	0	0		£ 2.29
4	B	237	270	0	0	0		£ 47.65
4	C	118	120	0	0	0		£ 1.62
4	D	259	170	0	0	0		£ 3.32
4	E	34	6	0	0	0		£ 44.17
							TOTAL	£ 99.04
5	A	25	13	0	0	0		£ 3.82
5	B	233	270	0	0	0		£ 48.93
5	C	92	120	0	0	0		£ 3.05
5	D	237	170	0	0	0		£ 4.54
5	E	34	6	0	0	0		£ 44.17
							TOTAL	£104.50

IM WEEKLY ROUTINES

WEEK NO 6

UPDATE EWMA OF DEMAND

```
A    0.15 *   7  + 0.85 *   8.28 =   8.09
B    0.15 * 129  + 0.85 * 133.03 = 132.43
C    0.15 *  53  + 0.85 *  43.54 =  44.96
D    0.15 * 110  + 0.85 *  77.76 =  82.59
E    0.15 *   2  + 0.85 *   8.12 =   7.20
```

UPDATE EWMA OF MAD OF DEMAND

```
A    0.15 *  1.28 + 0.85 *  3.45 =  3.12
B    0.15 *  4.03 + 0.85 * 42.39 = 36.63
C    0.15 *  9.46 + 0.85 * 11.33 = 11.05
D    0.15 * 32.24 + 0.85 * 19.01 = 20.99
E    0.15 *  6.12 + 0.85 *  6.24 =  6.22
```

UPDATE AVERAGE STOCK LEVEL

```
A    0.15 *  25 + 0.85 *  50.82 =  46.95
B    0.15 * 233 + 0.85 * 789.00 = 705.60
C    0.15 *  92 + 0.85 * 115.90 = 112.32
D    0.15 * 237 + 0.85 * 215.28 = 218.54
E    0.15 *  34 + 0.85 *  64.42 =  59.86
```

RECALCULATE CRITICAL LOST SALES VALUE

```
A          £  5.76
B          £ 41.79
C          £  8.32
D          £ 19.64
E          £ 54.20
           -------
TOTAL      £129.72
           -------
```

REFERENCES

1. A. CHAPANIS (1956) "The Design and Conduct of Human Engineering Studies", San Diego State College Foundation.
2. W.R. ASHBY (1956) "An Introduction to Cybernetics", published Chapman & Hall.
3. J.R.C. LICKLIDER (1961) "Man Computer Symbiosis", IRE Trans. on Human Factors in Electronics Vol. 1.
4. H.A. SIMON (1976) "Administrative Behaviour: A Study of Decision Making Process in Administrative Organisation 3rd Ed.", published Collier Macmillan, Canada.
5. J.A. FULLER (1978) "Optimal Solutions Versus 'Good' Solutions: An Analysis of Heuristic Decision Making", OMEGA 6, 6 pp 479-484
6. S.W. HAIDER, C.L. MOODIE & J.R. BUCK (1976) "An Investigation of Man/Computer Interactive Versus an Automatic Heuristic in Scheduling Job Shops", Report No. 87 Purdue Laboratory for Applied Industrial Control.
7. J.B. SCHNEIDER (1971) "Solving Urban Location Problems: Human Intuition Versus the Computer", J.Am.Inst. of Planners 1971 pp 95-99
8. J.B. SCHNEIDER, J.G. SYMONS Jr. & M. GOLDMAN (1972) "Planning Transportation Terminal Systems in Urban Regions: A Man-Machine Interactive Problem Solving Approach", Trans. Res. 6, 3 pp 257-274
9. D. MICHIE, J.G. FLEMING & J.V. OLDFIELD (1968) "A Comparison of Heuristic, Interactive, and Unaided Methods of Solving a Shortest-route Problem" in D. MICHIE (Ed) "Machine Intelligence 3" pp 245-255, published Edinburgh Univ. Press.
10. E. NADDOR (1975) "Optimal Heuristic Decisions for the s,S Inventory Policy", Management Science 21, 9 pp 1071-1073.

11. E. NADDOR (1975) "Optimal and Heuristic Decisions in Single - and Multi-item Inventory Systems", Management Science 21, 11 pp 1234-1249.

12. L.B. SCHWARZ (1973) "A Simple Continuous Review Deterministic One-Warehouse N-retailer Inventory Problem", Management Science, 19, 5 pp 555-566.

13. H. WAGNER, M. O'HAGAN & B. LUNDH (1965) "An Empirical Study of Exactly and Approximately Optimal Inventory Policies", Management Science, 11 pp 690-723.

14. J.R. FREELAND (1980) "Evaluating the Effectiveness of a New Method for Computing Approximately Optimal (s,S) Inventory Policies", Operations Research 28, 2 pp 353-364

15. E. MUMFORD (1981) "Participative Systems Design: Structure and Method", Systems, Objectives, Solutions No. 1 Vol. 1.

16. H.F. FARROW (1977) "The Effects of Computer-based Systems on Employee Attitudes and Working Skills in Small and Medium-sized Companies", for the European Social Fund.

17. R.L. ACKOFF & M.W. SASEINI (1968) "Fundamentals of Operations Research", published John Wiley.

18. C. EDEN, S. JONES & D. SIMS (1979) "Thinking in Organisations", published Macmillan.

19. G.A. KELLY (1955) "The Psychology of Personal Constructs", published Norton, New York.

20. K.S. RADFORD (1977) "Complex Decision Rules", published Reston, Virginia.

21. J. VON NEUMANN & O. MORGENSTERN (1953) "Theory of Games and Economic Behaviour", published Princeton Univ. Press.

22. R.L. SCHULTZ & D.P. SLEVIN (1975) "Implementation and Management Innovation" in R.L. SCHULTZ & D.P. SLEVIN (Eds.)

"Implementing Operations Research/Management Science", published Elsevier.

23. R.L. ACKOFF (1960) "Unsuccessful Case Studies and Why", Operations Research 8, 2 pp 259-263.

24. S. BONDER (1979) "Changing the Future of Operations Research", Operations Research 27, 2 pp 209-224.

25. C.J. GRAYSON Jr. (1973) "Management Science and Business Practice", Harvard Business Review Vol. 51 pp 41-48.

26. G.N. POWELL (1976) "Implementation of OR/MS in Government and Industry: A Behavioural Science Perspective", Interfaces 6, 4 pp 83-89.

27. S. HILDEBRANDT (1980) "Implementation - the Bottleneck of Operations Research: the State of the Art", Euro. J. of Op. Res. 6 pp 4-12.

28. R.K. WYSOCKI (1979) "OR/MS Implementation Research: A Bibliography", Interfaces 9, 2 Pt. 1 pp 37-41.

29. M.J. GINZBERG (1978) "Steps Towards More Effective Implementation of MS and MIS", Interfaces 8, 3 pp 57-63.

30. D. MEISTER (1976) "Behavioural Foundations of System Development", published John Wiley.

31. A. PARKIN (1980) "Systems Management", published Edward Arnold.

32. H.W. SINAIKO & E.P. BUCKLEY (1961) "Human Factors in the Design of Systems" in H.W. SINAIKO (Ed.) "Selected Papers on Human Factors in the Design and Use of Control Systems", published Dover Publications Inc., New York.

33. J.W. FORRESTER (1972) "Understanding the Counter-intuitive Behaviour of Social Systems" in J. BEISHON & G. PETERS (Eds.) "Systems Behaviour", published Open Univ. Press

34. A. BUZAN (1974) "Mind Your Head", published BBC.

35. J.W. FORRESTER (1961) "Industrial Dynamics", published MIT Press.

36. P.M. FITTS (1951) "Engineering Psychology and Equipment Design" in S.S. STEVENS (Ed.) "Handbook of Exp. Psychol.", published John Wiley.

37. W.T. SINGLETON (1974) "Man-Machine Systems", published Penguin.

38. J.D. BAKER & I. GOLDSTEIN (1966) "Batch vs. Sequential Displays: Effects on Human Problem Solving", Human Factors vol. 8 pp 225-235.

39. R.B. MILLER (1969) "Archetypes in Man-Computer Problem Solving", Ergonomics 12, 4 pp 559-581.

40. L.P. SHRENK (1969) "Aiding the Decision Maker - A Decision Process Model", Ergonomics 12, 4 pp 543-557.

41. W.S. VAUGHAN Jr. & A.S. MAVOR (1972) "Behavioural Characteristics of Men in the Performance of Some Decision-Making Task Components", Ergonomics 15, 3 pp 267-277.

42. L. PRESS (1971) "Toward Balanced Man-Machine Systems", Int. J. Man-Machine Studies 3 pp 61-73.

43. H. SACKMAN (1967) "Computers, System Science, and Evolving Society", published John Wiley.

44. P.K. ANDOW & F.P. LEES (1974) "Process Plant Alarm Systems: General Considerations" in C.H. BUSCHMANN (Ed.) "Loss Prevention and Safety Promotion in the Process Industries" pp 299-307, published Elsevier.

45. A.G.S. PASK (1958) "Teaching Machines", Proc. of 2nd International Conference on Cybernetics.

46. D. WHITFIELD (1975) "Man-Computer Symbiosis: A 1975 Review", A.P. Report 57, Univ. of Aston.

47. W. EDWARDS (1973) "Divide and Conquer: How to Use Likelihood and Value Judgements in Decision Making" in R.F. MILES Jr. (Ed.) "Systems Concepts" pp 87-110 published Wiley-Interscience.

48. S.C. WHEELWRIGHT & S. MAKRIDAKIS (1980) "Forecasting Methods for Management, 3rd Ed.", published Wiley-Interscience.

49. T.M. WHITIN (1957) "Theory of Inventory Management", published Princeton Univ. Press.

50. H.I. ANSOFF (1968) "Corporate Strategy", published Penguin.

51. R.N. ANTHONY (1965) "Planning and Control Systems - A Framework for Analysis", published Harvard Univ.

52. S. BEER (1966) "Decision and Control", published John Wiley.

53. C.D. LEWIS (1970) "Scientific Inventory Control" published Butterworths.

54. R.G. BROWN (1967) "Decision Rules for Inventory Management", published Holt, Rinehart & Winston, New York.

55. K.J. ARROW, S. KARLIN & H. SCARF (1958) "Studies in the Mathematical Theory of Inventory and Production", published Stanford Univ. Press.

56. G. GERSON & R.G. BROWN (1970) "Decision Rules for Equal Shortage Policies", Naval Research Logistics Quarterly 17, 3 pp 351-358.

57. R.G. BROWN (1977) "Materials Management Systems", published Wiley-Interscience.

58. F.R. JOHNSTON (1975) "Approximations Useful for Stock Control', Working Paper 25, School of Industrial and Business Studies, Univ. of Warwick.

59. J.D. CROSTON (1972) "Forecasting and Stock Control for Intermittent Demands", Opl. Res. Q. Vol. 23 No. 3 pp 289-303.

60. F.R. JOHNSTON (1980) "An Interactive Stock Control System with a Strategic Management Role", J. Opl. Res. Soc. Vol. 31 pp 1069-1084.

61. T.A. BURGIN & A.R. WILD (1967) "Stock Control - Experience and Usable Theory", Opl. Res. Q. Vol. 18 No. 1 pp 35-52.

62. T.A. BURGIN (1975) "The Gamma Distribution and Inventory Control", Opl. Res. Q. Vol. 26 No. 3, i pp 507-525.

63. P.R. WINTERS (1957) "Calculation of the Incomplete Gamma Function Suitable for Electronic Computer Application", Appx. B ONR Research Memorandum No. 48 Contract N-onr-760/01/Project NR 047011.

64. C. HASTINGS Jr. (1955) "Approximations for Digital Computers", Princeton Univ. Press.

65. K. PEARSON (1957) "Tables of the Incomplete Gamma Function", published Cambridge Univ. Press.

66. T.A. BURGIN & J.M. NORMAN (1976) "A Table for Determining the Probability of Stockout and Potential Lost Sales for a Gamma Distributed Demand", Opl. Res. Q. Vol. 27 pp 621-631.

67. J.W. FORRESTER (1969) "Urban Dynamics", published MIT Press.

68. J.W. FORRESTER (1971) "World Dynamics", published Wright-Allen Press, Cambridge, Mass.

69. F.L. WESTWATER & W.A. WADDELL (1961) "An Introduction to Servomechanisms", published English Universities Press.

70. J.W. FORRESTER (1968) "Principles of Systems", published Wright-Allen Press, Cambridge, Mass.

71. G.L. CHILDS (1967) "Unfilled Orders and Inventories - a Structural Analysis", published North-Holland, Amsterdam.

72. P.E. CARNELL (1974) "Studies of the Use of Industrial Dynamics Models in the British Steel Corporation", Ph. D. Thesis, The

University of Aston.

73. E.B. ROBERTS (1964) "The Dynamics of Research and Development", published Harper & Row, New York, Evanston & London.

74. R.J. MEYER & A.R. ROBERTS (1969) "Views of Industrial Dynamics and the Honeywell Experience with Implementation", Proc. Conf. on Applications of Continuous System Simulation Languages, ACM/IEEE/SHARE/SCI, San Francisco.

75. A.L. PUGH III (1970) "Dynamo II User's Manual", published MIT Press.

76. D.H. MAISTER "Centralisation of Inventories and the Square Root Law", IJPD 6, 3 pp 124-134.

77. R.G. BROWN (1959) "Statistical Forecasting for Inventory Control", published McGraw-Hill.

78. K.D. TOCHER (1963) "The Art of Simulation", published English Universities Press.

79. C.D. LEWIS (1967) "An Iterative Analogue Inventory Policy Cost Model", Int. J. Prod. Res. Vol. 6 No. 2 pp 135-151.

80. A.C. ATKINSON (1977) "An Easily Programmed Algorithm for Generating Gamma Random Variables", J.R. Statistic. Soc. A, 140 Part 2 pp 232-234.

81. A.C. ATKINSON & M.C. PEARCE (1976) "The Computer Generation of Beta, Gamma and Normal Random Variables", J.R. Statistic. Soc. A, 139 pp 431-461.

82. J.H. AHRENS & U. DIETER (1974) "Computer Methods for Sampling from Gamma, Beta, Poisson and Binomial Distributions", Computing (Vienna), 12 pp 223-246.

83. T.A. BURGIN (1972) "Inventory Control with Normal Demand and Gamma Lead Times", Opl. Res. Q. 23 pp 73-80.

84. "DISFIT - An Interactive 'Goodness of Fit' Package", The

University of Aston Management Centre (1978).

85. M.J. BECKMAN & R.F. MUTH (1958) "On the Two Bin Inventory Policy: An Application of the Arrow-Harris-Marschak Model" in K.J. ARROW, S. KARLIN & H. SCARF "Studies in the Mathematical Theory of Inventory and Production", published Stanford Univ. Press.

86. J. MURDOCH & J.A. BARNES (1977) "Statistical Tables for Science, Engineering, Management and Business Studies" published Macmillan.

87. S. MAKRIDAKIS & S.C. WHEELWRIGHT (1974) "Interactive Forecasting" published Scientific Press, Palo Alto, Calif.

88. W.D. RAY (1980) "The Significance of Correlated Demands and Variable Lead Times for Stock Control Policies", J. Opl. Res. Soc. Vol. 31 pp 187-190.

89. H.M. WAGNER (1962) "Statistical Management of Inventory Systems", published John Wiley.

90. D. BURNELL, J. KENNEDY & R. HARRISON (1979) "Health Service Stock Control: An Approximation to the Gamma Distribution Gives Some Explicit Results", NCB internal paper.

91. W. LAMPKIN (1967) "A Review of Inventory Control Theory", The Production Engineer, Vol. 46 No. 2 pp 57-66.

92. P.B. TAYLOR & K.H. OKE (1976) "Tables for Stock Control - Problems of Formulation and Computation", Opl. Res. Q. Vol. 7, 3 pp 747-758.

93. C.E. WEATHERBURN (1962) "A First Course in Mathematical Statistics", published Cambridge Univ. Press.

94. R.L. WINKLER & A.H. MURPHY (1968) "Good Probability Assessors", Journal of Applied Meteorology Vol. 5, 7 pp 751-758.

95. J.B. KIDD (1975) "Scoring Rules for Subjective Assessments", Opl. Res. Q. Vol. 26, 1, ii pp 183-195.

96. C.S. STAËL VON HOLSTEIN (1969) "The Assessment of Discrete Subjective Probability Distributions - An Experimental Study", Inst. of Mathematical Statistics, Univ. Stockholm, Research Report No. 41.

97. J.B. KIDD (1975) "The Bayesian Combination of a Manager's Opinion and an Algorithmically Derived Forecast", University of Aston Management Centre Working Paper Series No. 44.

98. D.W. BUNN (1975) "A Bayesian Approach to the Linear Combination of Forecasts", Opl. Res. Q. Vol. 26, 2, i pp 325.

99. M. GREEN & P.J. HARRISON (1973) "Fashion Forecasting for a Mail Order Company using a Bayesian Approach", Opl. Res. Q. Vol. 24, 2 pp 207.

100. J.P. DIKINSON (1973) "Some Statistical Results in the Combination of Forecasts", Opl. Res. Q. Vol. 24, 2 pp 253.