Aircraft Rotable Inventory Optimisation

Development of a new solution

Michael MacDonnell

Doctor of Philosophy

Aston University

October 2008

Aston University

Aircraft Rotable Inventory Optimisation

Development of a new solution

Michael MacDonnell

PhD 2008

Analysis of the use of ICT in the aerospace industry has prompted the detailed investigation of an inventory-planning problem. There is a special class of inventory, consisting of expensive repairable spares for use in support of aircraft operations. These items, called rotables, are not well served by conventional theory and systems for inventory management.

The context of the problem, the aircraft maintenance industry sector, is described in order to convey some of its special characteristics in the context of operations management.

A literature review is carried out to seek existing theory that can be applied to rotable inventory and to identify a potential gap into which newly developed theory could contribute.

Current techniques for rotable planning are identified in industry and the literature: these methods are modelled and tested using inventory and operational data obtained in the field.

In the expectation that current practice leaves much scope for improvement, several new models are proposed. These are developed and tested on the field data for comparison with current practice.

The new models are revised following testing to give improved versions. The best model developed and tested here comprises a linear programming optimisation, which finds an optimal level of inventory for multiple test cases, reflecting changing operating conditions.

The new model offers an inventory plan that is up to 40% less expensive than that determined by current practice, while maintaining required performance.

Keywords: repairable spares management, operations research

Acknowledgements

Thanks to Professor Tom Allen of Massachusetts Institute of Technology and the National Institute of Technology Management, University College Dublin for providing the initial direction and collaboration for this research agenda.

Enterprise Ireland provided funding for the research that forms the background to this project, grant PRP00-AMT-03.

FLS Aerospace provided the data set and problem expertise.

Colleagues in the Product Support Technical Committee of the American Institute of Aeronautics and Astronautics, comprising specialists from Airbus, Boeing and other major aerospace organisations, gave valuable industry knowledge, validation and feedback.

For their expert guidance and collaboration, thanks are due to Dr Ben Clegg, Dr Doug Love and the late Dr Nelson Tang of Aston Business School.

Thanks to Professor Andrew Deegan and colleagues of the Management Information Systems Department, University College Dublin for their collegiality and support.

Above all, thanks to Mum, Dad and Carolyn for many years of loyalty, patience and encouragement.

Contents

Figures

Tables

Glossary

term	*description*
"go"	a category of inventory that can fail without preventing an aircraft from failing, e.g., a coffee maker – essentiality code (Ess Code) 3
"go-if"	a category of inventory that can conditionally fail without preventing an aircraft from operating, e.g., a cockpit instrument can fail providing two backups are working – essentiality code (Ess Code) 2
"no-go"	a category of inventory whose failure prevents an aircraft from operating, e.g., control surfact actuator – essentiality code (Ess Code) 1
Δfill	change in the number of satisfied requests for spares resulting from increasing the level of a part i from quantity j-1 to j
actual holding	the total stock held by the operator of a given part number
average part value	average value of an item of stock from a solution = total cost / total inventory count; an indicator of solution efficiency
binary variable	a decision variable with 0 and 1 as the only allowed values
case	a use case based on observed or recommended practice, e.g., different SL values; the basis for comparison of the performance of different models against each other
CLP	Current List Price: an input cost value obtained from the part supplier as the latest full price for a new item
constraint	a mathematical bound on a solution to enforce a condition, e.g., target Service Level must be met
count	quantity of a given part number from a model solution
CWSH	Cost-Wise Skewed Holding, a heuristic model where line items are grouped into cost bands and given different target Service Levels
decision variable	an element of a problem formulation to be assigned a value in the solution, e.g., X23e represents the quantity 5 (from the letter e) of line 23
demand event	an arising, or part removal, leading to a request for a spare part
deviation metric	a ratio measuring the degree to which a solution corresponds to the best solution for a given test case
Ess Code	Essentiality Code: a classification of parts into three levels with values 1, 2 and 3, corresponding to "no-go", "go-if" and "go" respectively
fill list	the inventory pick list, the quantities of each line item prescribed by a solution to reach the objective function or target Service Level
fill rate	the expected number of satisfied demand events for a given quantity of a particular line item
formulation	a model with a set of parameters defined for solution
Gauss	a normal probability distribution, which may be used in forecasting satisfaction rates of demand events
GBV	Gross Book Value, the accounting value of a spare part; does not allow for depreciation, holding or disposal costs
global SL	the required demand satisfaction rate for a combined inventory set

heuristic	practical algorithm to give an easy solution to an optimisation problem; not expected to be theoretically optimal
holding cost	the cost of given quantity (holding) of a part; not the cost of holding inventory
holding quantity	the quantity of a part prescribed by a solution
i	represents a part number in a model formulation: in the data set tested, i can have values from 1 to 300
index number	the number in the sequence of part numbers or line items held = i; used in data storage, problem formulation and solution sorting
j	represents a quantity of a part in a model formulation: for the data tested, j typically ranges from 1 to 15, but in some cases up to 30 or 60
line	index number or part number
Line replaceable unit, LRU	a rotable spare that can be changed on an aircraft "on the line", without removing the aircraft from service
linear program	an algebraic model consisting of an expression to be maximised or minimised and a set of problem constraints
linear program solver	a software tool used to solve an LP formulation by finding the best fit to the objective function while maintaining constraints
LP	linear program, also a model consisting of a linear program formulation of the full inventory set (all part numbers)
LP3	a linear program model where the inventory set is divided by essentiality code and each of three sub-problems is formulated and solved separately
LP-combined	linear program formulation of the full inventory set
LP-split	same as LP3
MA	Marginal Analysis, a model aiming to allocate inventory quantities in the order that they contribute the best value to operations
MC	Marginal Contribution, the incremental fill rate provided by increasing a part quantity, divided by the part's cost
MLP	Manufacturer's List Price, the full price of an original part
model	a mathematical method of solution, which may be run for many cases with different formulations
MTBR	Mean Time Between Removals, the mean statistical number of hours after which a part is expected to be removed from service; the part has not necessarily failed but requires replacing
MTBRIP	Mean Time Between Removals, Initial Provisioning, the predicted number of hours after which a part is expected to be removed from service, as recommended by the manufacturer and agreed by the customer at the time that the purchase of spares stock is negotiated
objective function	a mathematical statement used to drive a linear program, e.g., minimise cost
owned stock	the actual stock held by an operator, obtained from operational data, not used in model formulation but used for comparison
part number	a numerical designation given to a common part type, treated as interchangeable in inventory

Poisson	a probability distribution for discrete events with a low frequency, considered the best fit for the failure of aircraft parts
REMS	removals, the number of parts of a given line removed from service during the planning period
rotable	an inventory classification for parts that are maintained and returned to stock following failure, typically with a life span equal to that of the parent fleet; inventory that is not consumed or discarded; a float of rotables is held to allow line changes without disruption to aircraft operations; also called Line Replaceable Units; may be a composite entity, e.g., an engine containing rotables
sequence number	same as line
SL	Service Level, the required or achieved average probability of demand events being satisfied. An exponential function with respect to inventory levels, full satisfaction or 100% SL is not attainable in practice
SL scaling	applying weights to parts with Ess Code 2 or 3 in a combined model, to reduce the selection of these parts relative to Ess Code 1 items
SLAirbus	Set of SL values prescribed by Airbus: 95, 89 and 75% for Ess Code 1, 2 and 3
SLFLS	Set of SL values used by FLS and in line with Boeing recommendations: 95, 93 and 90% for Ess Code 1, 2 and 3
solution	the outcome of a formulation, i.e., the set of decision variables produced by a formulation of a model
Stk	same as owned stock
stochastic	a process where events occur around a probability distribution, such as the removal of aircraft rotables
subject to	linear program model term to introduce constraints after the declaration of the objective function
target SL	SL to be met by a solution
TAT	Turn Around Time, the elapsed time, in days, from a rotable being removed from service to it becoming available in stock following diagnosis, routing, repair and receipt; spares holding quantities are directly proportional to TAT
TCH	Total Component Hours, the number of hours of operation encountered by all installed parts in the planning period, equal to the number of aircraft in operation multiplied by the number of components of that part number on an aircraft multiplied by the mean number of hours flown by that aircraft type
total cost	the sum of the quantity of a part number prescribed by a solution multiplied by the cost of the part
Total holding cost	the sum of all total costs of all parts prescribed by a solution
total inventory count	the sum of the quantities of all parts prescribed by a solution
UAF	Un-Availability Factor: the portion of the planning period during which a part is in the repair cycle; equal to TAT divided by the planning period (normally 365 days)
value	cost data used in problem formulation, based on CLP or GBV

weighted cost	cost of an Ess Code 2 or 3 part scaled up by an SL scaling value to make it relatively less attractive to a solution than an equivalent Ess Code 1 part
weighted demand	demand for an Ess Code 2 or 3 part scaled down by an SL scaling value to make it relatively less attractive to a solution than an equivalent Ess Code 1 part

Chapter 1: Introduction

A substantial decision problem has been identified for development: the theory emerging from this work comprises a novel application of standard mathematical techniques in a well-defined operational setting. While the benefits of this work apply to the airline industry, there is potential to extend the models developed into other arenas.

This study looks at the problem of planning optimal levels of spare parts that are used in daily operations of commercial aircraft. These items, generally called **rotables**, and specifically termed **Line Replaceable Units**, are characterised as follows:

- they can be replaced on an aircraft "on the line", meaning without significant disruption to service, if any;
- they are generic, such that a given type of spare can be fitted on any aircraft of suitable type;
- they are valuable and worth repairing – rotables are so called since they rotate through stock, being used and repaired as required, and usually surviving the lifetime of the parent aircraft fleet; rotables are part of an airline's assets, they are not considered consumable;
- while inventory levels will fluctuate in the short term, due to failures and replacements, and will change in the long term due to fleet changes, there is no net change in inventory in the medium term, which is the inventory-planning horizon;
- they are not well represented by production-oriented inventory management systems, since they are not consumed;
- they may be bought and sold to other users and are fully interchangeable among airlines (subject to airlines' engineering policies).

In certain cases the term rotables may be used in a more general sense (for example, engine parts in the overhaul cycle), but in this study rotables are defined as LRUs.

The world rotable inventory holding is estimated at over $80bn and it is thought that much of this is excessive, due to poor planning, and never used. Despite advances in ICT and optimisation sciences, this pool continues to grow. Since rotables are a major source of revenue for the aircraft Original Equipment Manufacturers (OEMs), there is a conflict of interest when OEMs like Airbus, Boeing and the engine manufacturers advise airlines on initial provisioning.

This study looks at the literature on inventory planning in general and rotable inventory planning in particular and includes industry sources for reference on practice and trends.

In the expectation that there is scope to develop better solutions than those currently in use and in the literature, a set of models is built and tested. A common data set, obtained from a Maintenance, Repair and Overhaul provider (MRO) is used to evaluate both known practice and the new models.

The emphasis in this work is on the empirical development and testing of new models, which show large gains in efficiency and cost reduction for the data set tested. The new models are conclusively shown to offer significant advantage over current practice by a thorough assessment of the results.

The thesis concludes with recommendations for improved practice and claims for contribution to theory in the area of operations management, specifically in the planning of rotable inventory levels.

The layout and content of this thesis are presented here to show the progression from literature search to the intended contribution to theory. The logical flow between sections is illustrated in Figure 1 below, where numbers refer to the chapters of this dissertation.

The chapters that follow, and their aims, are listed below.

2. Industry Background – the characteristics of the operational area being studied: size and type of organisation, the supply chain, typical business processes, application of ICT solutions, trends and driving

forces, outlook for the industry taking into account environmental changes.

3. Rotable Inventory Optimisation literature review – assess published theory and experience of the identified problem; consider how general inventory theory applies to this special case; demonstrate the gap in the literature to be addressed by the present work.

4. Methodology – a prescriptive schedule to perform an in-depth study with theoretical substance and value, presenting new techniques to build on current practice and existing theory.

5. Model Formulation and Implementation – design specifications for each of the solution methods proposed herein and describe their implementation, together with the preparation of the required input data.

6. Results – obtain the solution values and performance measures for the different solutions.

7. Analysis – compare solution outputs for comparison and review and make recommendations for the best use of the solution methods developed and tested here. A further exercise in scenario analysis is introduced here: the effect on demand, inventory levels and cost of varying airline fleet size.

8. Discussion – compare the findings of the experimental research with the literature; consider the implications of the research for the subject field and other fields; argue contributions to the theoretical fields of knowledge explored here, namely operations management, inventory management and operations research

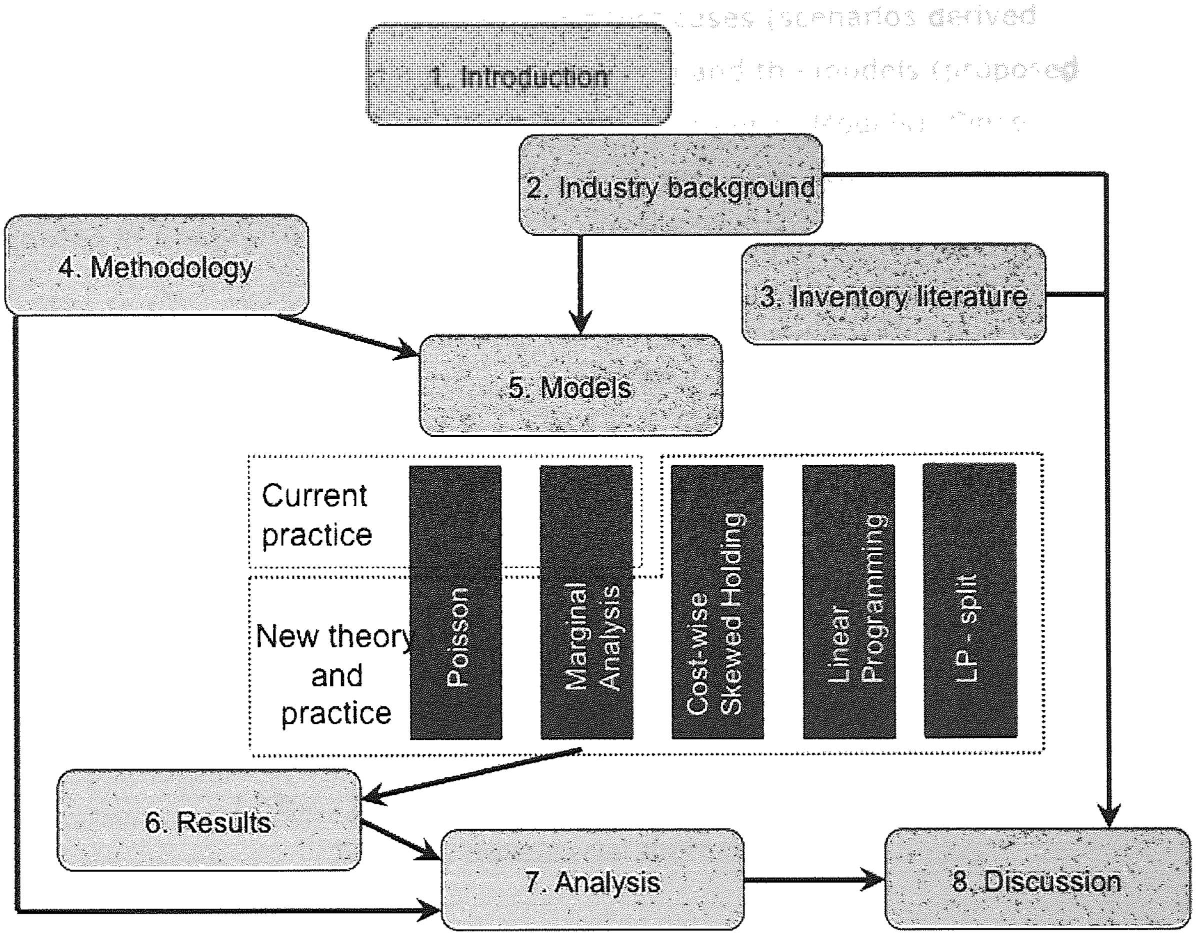

Figure 1.1: project schema

The core of this work is a set of rotable inventory planning models, which are tested on a common data set for a range of scenarios.

The five models shown in Figure 1.1 represent implementations of known practice (Poisson and Marginal Analysis) and new models. The LP-split model gives the best results and forms the basis of recommendations for improved practice.

The testing of multiple scenarios on the same data set across all of the models gives a sensitivity analysis and quality control of the models, showing the effect of changes in scale and other operating conditions. An interesting outcome is that, for an aircraft fleet of moderate size, the economy of scale available from using the same spares pool to support a growing fleet, while significant, is not as great as suggested by the literature.

Table 1.1 shows the structure of the experimental work to be performed, which comprises the relationship between test cases (scenarios derived from the literature and presented in Chapter 3) and the models (proposed in Chapter 4, Methodology and developed in Chapter 5, Models). Since there are 5 cases and 5 models, there are 25 test runs, which are recorded in Chapter 6 (Results) and interpreted in Chapter 7 (Analysis).

<table>
<tr><td colspan="2"></td><td colspan="5">Cases – Chapter 3</td></tr>
<tr><td colspan="2"></td><td>Case 1
base</td><td>Case 2
fewer</td><td>Case 3
faster</td><td>Case 4
bigger</td><td>Case 5
best</td></tr>
<tr><td rowspan="5">Models – Chapters 4, 5</td><td>Poisson</td><td colspan="5" rowspan="5">25 test runs – Chapter 6</td></tr>
<tr><td>Marginal Analysis</td></tr>
<tr><td>Cost-wise skewed</td></tr>
<tr><td>LP</td></tr>
<tr><td>LP3</td></tr>
</table>

Table 1.1: experimental structure

Chapter 2: Industry Background

This chapter presents background descriptive information on the problem area, drawing chiefly on first-hand observation and experience. It is intended to communicate the nature of the problem area and stops short of addressing the theory and related literature to be addressed by the empirical work that forms the core of this study.

The aircraft maintenance industry sector has been chosen for study here, as it offers interesting opportunities for applied research, with scope for theoretical developments therein.

The nature of the products involved – aircraft and engines – is explored in some detail in order to illustrate the characteristics of flight operations and maintenance and to convey the richness of information required to manage these operations.

This chapter takes a broad look at an industry, considers usage of ICT within the industry and within constituent organisations, and finally describes an operational problem, which is the core research in this study.

2.1 Problem context – industry composition

The aircraft maintenance industry presents interesting opportunities for the application of management science techniques for improved decision-making, process support and cost saving. The current work follows from a research project funded by the Irish government (Enterprise Ireland grant reference PRP00-AMT-03), which studied the Irish aviation industry for potential benefits of improved uses of Information and Communications Technology and management science tools. Several useful prototypes were developed in the course of this project: it is proposed to further develop and validate these solutions with the collaboration of firms not previously involved in the research.

A top-down approach has been applied to documenting current practice in the industry, which has been shown to be highly generic and predictable through a series of detailed process analysis exercises. The first output, a business process reference model, serves to document common processes governing aircraft maintenance activity.

2.1.1 Supply chain overview – member firm generic types

The types of organisation comprising the aircraft maintenance supply chain are described here. The scale of the industry and its business environment are discussed later in this chapter.

1. **Airline operator**: all maintenance activity is triggered by airline operations, scheduled or unscheduled. The vast majority of maintenance activity is scheduled by elapsed time (e.g., corrosion inspection) or hours and number of flights accrued (most maintenance). However, there is a trend with modern aircraft design to have increasing levels of "on-condition" maintenance, where aircraft systems measure their own performance and report deterioration. This is now standard practice with jet engines, so the removal time of an engine is not planned but occurs when performance has dropped to a defined level. Thus engine changes do not coincide with maintenance activity for the rest of the aircraft, referred to as the airframe.

Unscheduled maintenance results from system malfunctions, and in many cases the aircraft can be returned to service without significant downtime by the exchange of the failed system component. Components (or sub-systems) with the ability to be replaced 'on the line' (without removing the aircraft from service) are referred to as Line Replaceable Units. In the case that an LRU can be maintained and returned to spares inventory, it will be categorised as a rotable item (since the component rotates through inventory, i.e., it may be removed from inventory and returned to inventory after use or maintenance many times throughout its life).

2. **Base station:** an airline will need a *line maintenance* capability, whereby troublesome components can be diagnosed and often replaced on the line (during aircraft stops at airports): this function will often be performed by the airline, but will sometimes be outsourced, especially at remote stations. The airline also needs a *technical services* function, an engineering capability to schedule maintenance, make decisions on unscheduled maintenance action and consult maintenance providers and OEMs. This engineering function may be performed by the airline, but can also be outsourced, often to the primary maintenance provider.

3. **Primary maintenance provider**, usually referred to as a Maintenance, Repair and Overhaul provider, or **MRO**, provides (i) technical services, (ii) materials management and (iii) skilled labour to airline customers.

(i) *Technical services* may include maintenance scheduling, OEM instructions and upgrades and production and archiving of technical records for operational and regulatory compliance. Where an MRO provides technical support to an airline, they will receive operational data from the airline (hours and flights operated by each aircraft, component failures and removals) and use this to plan maintenance, bearing in mind the need to maintain sufficient fleet in service to keep the airline in business. Seasonal demand may be taken into account, so that some major maintenance will be scheduled early during quieter periods. Another important function is trend monitoring: when an aircraft exhibits exceptional behaviour, such as elevated fuel consumption or excessive vibration, decisions will need to be taken regarding major maintenance. Engine maintenance is mostly conducted by separate firms to airframe (everything except engines) maintenance and is scheduled separately. Thus engines will be changed overnight on aircraft that will continue in service until airframe work is needed. Therefore technical services for engine management may be provided by an engine MRO, as well as being provided by the airframe MRO.

(ii) *Materials management* by an MRO involves provisioning of parts needing replacement during maintenance, and may also include rotable support for aircraft line operations and selling of surplus materials. Where an MRO has several customers with common rotable requirements (similar aircraft types), there is huge potential for cost saving by pooling the requirements of several airlines and supporting them with a common pool of rotables. However, this is limited by airline policies, where the airline may not accept parts previously used by another airline. Some airlines will have their spare inventory held by an MRO – this is referred to as consignment stock, which may be governed by complex rules of ownership and exchange with other stock owned by other customers or the MRO itself.

(iii) The largest element of activity in MRO is *labour* supply (although the largest contributor to profit is usually materials management). Licenced technicians will specialise in aircraft types, skills (such as painting or welding) and component repair (such as radios). Maintenance contracts are generally negotiated such that a prescribed maintenance check (job) requires a fixed budget for labour (e.g. 4,000 manhours) at a negotiated hourly rate for a combination of mechanics, supervisors and inspectors. Managing labour productivity (hours billed for a job over actual hours consumed) and labour utilisation (total hours billed in a period over total hours available and paid) is a critical task and most surviving MROs are an amalgamation of earlier defaults. Engine MROs tend to be more profitable than airframe MROs, mainly due to the high materials portion of the job, so that margins on material trading generate significant income.

Engines and airframe are typically dealt with by separate organisations according to separate maintenance schedules. Thus when an aircraft is removed from service for major airframe maintenance, the engines may be left in place without maintenance, or they may be removed for spares and replaced when the airframe nears return to service. Equally, when engines require maintenance, they are removed from an active aircraft and replaced without disruption to service – this is usually carried out over night and takes around 6 hours. Thus, while engine and airframe MROs are similar types of organisation from a supply chain perspective, they are usually different entities for a given airline and will be subject of separate commercial arrangements.

An airline's choice of MRO depends mainly on commercial negotiation and takes into account elements like mark-up on replacement parts, turn-around times, transport costs, spares support and past delivery and quality performance. An airline may split maintenance contracts between MROs to get the best terms, although this has disadvantages in coordination and reduction of scale. In February 2008, Aer Lingus divided its airframe maintenance into four areas (wheel and brake, line support, overhaul and components) for separate negotiation by tender, citing their objective as finding the best terms (RTE 2008). However, this was a strategy to break away from their captive MRO and former subsidiary SR Technics in Dublin, where there were political problems. SR Technics was awarded one contract, line maintenance, which by its nature is locally based and could not be supported by another company. It can be expected that an airline will receive better support and commercial terms from an MRO getting all of its business, so it is unlikely that splitting contracts is a viable long-term strategy.

4. Repair vendor – repair agent or specialist subcontractor. These are typically firms of several hundred employees with a customer base that is fairly small in number but globally dispersed. They tend to concentrate on skills and capabilities that are expensive to develop and for which there is a small market, for example, precision machining of structural engine cases, deposition of high-temperature ceramic coatings on engine airfoils and overhaul of engine bearings. These firms are typically independently controlled and occupy a small, stable market, with good profitability but high ongoing research, qualification and selling costs. Given their profitability, however, there has been a trend for OEMs to buy up increasing amounts of repair capacity.

5. Parts trader – intermediaries buying excess inventory from airlines, OEMs, MROs and subcontractor and selling back to the same constituency. Some will specialise in aircraft types or product types, such as engine parts and they may invest in obsolete parts to have them upgraded by a subcontractor and made saleable. The parts trading industry sector exists to a large extent because of the intransparency of information: if airlines and MROs had full visibility of the availability of spares, they would not need to trade with traders. In other words, if a buyer could easily see relevant offerings from sellers, they could trade directly without the need for an intermediary.

The other value brought by the parts trader, as well as technical expertise, is the resources to take risks: traders will buy surplus inventory or the inventory of a failed airline or retired aircraft type, with a view to adding some value and re-selling the inventory. This may comprise the purchase of parts to be overhauled and upgraded by a repair subcontractor. For example, a trader bought five old-generation Boeing 747 engines at the liquidation auction of PanAm's spares in the early 1990s. The company then agreed a deal with an engine MRO whereby the MRO provided three complete overhauled engines at no material cost to the trader (who paid for some of the labour required). The trader then sold the engines to an African airline, which still operated this older engine type, making a significant profit. The MRO added the surplus parts obtained to its inventory, having overhauled them.

Aircraft parts trading, which can be very profitable, is a very specialised sector. There have been several attempts to launch independent e-commerce exchanges for aircraft spares – for example www.componentcontrol.com (June 2008). However, rather than operating transactional exchanges, these sites typically provide listings with the result that buyers and sellers are put in contact and proceed with person-to-person negotiation. Thus the business remains highly manual, which reflects price sensitivity and the complex nature of the parts. Rather than being true commodities, aircraft parts require certification, service history and traceability, to the extent that in most cases, parts without full paperwork have no market value.

6. **Original Equipment Manufacturer** – aircraft manufacturers such as Boeing and Airbus, and providers of cockpit, cabin and ancillary equipment and spares. The OEMs are responsible for initial certification of equipment and will communicate with airlines and MROs with updates to parts for modernisation and compliance with safety requirements. The main OEMs will often have representatives permanently based with large MROs and airlines to deal with queries and upgrades – this may be part of the support package negotiated with the sale of aircraft.

7. **Other service provider** – engineering services, finance, IT, regulatory bodies. There are firms providing insurance and lease finance, business process consulting and software tools to the airline industry – some of this activity may involve the maintenance supply chain. For example, an aircraft lease must address maintenance – which party (owner or operator) pays for unscheduled maintenance and plans and pays for scheduled maintenance. The aircraft owner will have "hand-back criteria", i.e., aircraft and engines will have to have a minimum service time available from the return by the customer, so that they can be sent on a new lease without needing major overhaul at the leasing company's expense.

2.1.2 Aircraft Maintenance Supply Chain Reference Model

Business processes are highly generic in the aircraft maintenance industry, since they are highly regulated, operate in a very technical environment, involve many metrics and are based around common products and procedures originating from a small number of manufacturers. Practices in the airline industry have tended to originate from American military practice, which continues to be the main driver of product research in aerospace. Flight operations and maintenance procedures are prescribed and overseen by regulatory bodies, airline operators and aircraft manufacturers. In maintenance operations, processes are almost universally based around manufacturers' maintenance manuals. Therefore there is a very high degree of commonality across all firms performing maintenance on equipment from the same manufacturers. Further, as practices become established, it is common for maintenance tasks to be standardised across brands. Thus a D-check (heaviest level of scheduled maintenance shop visit for an aircraft) for an Airbus A320 will be similar in nature, content and processes to a D-check on a Boeing B737. Given the standard nature of processes throughout the industry, it is therefore reasonable to build a representative set of business processes to reflect the typical operations employed by an airline and its suppliers in managing the technical side of its business, i.e., aircraft operations and maintenance. This work focuses on maintenance activities more than airline operations.

Business processes are structured, repeatable operations carried out by organisations of interest, which may occur inside an organisation or between organisations, for example, sending an item for repair to a subcontractor. These processes are typically shown in a graphical form to illustrate the flow of actions, material and information, decision points and connections to other processes.

The objective of a reference model is to compile and present a representative set of process maps to gain consensus among users about (a) how processes are currently performed and (b) how they might be improved. The reference model can then be the vehicle for systems implementation and organisational change. A high-level view of the Aircraft Maintenance Supply Chain Reference Model is shown in Figure 2.1. This model has been based on processes mapped at the following entities:

1. airline = Aer Arann Express

2. base station = Aer Arann Express

3. primary maintenance provider (MRO) = Shannon MRO

4. specialist subcontractor = PWAI

5. parts trader = Magellan

6. OEM (not mapped)

7. Other service provider (not mapped)

Entities 6 and 7 above were not mapped since they are peripheral to the aircraft maintenance industry sector.

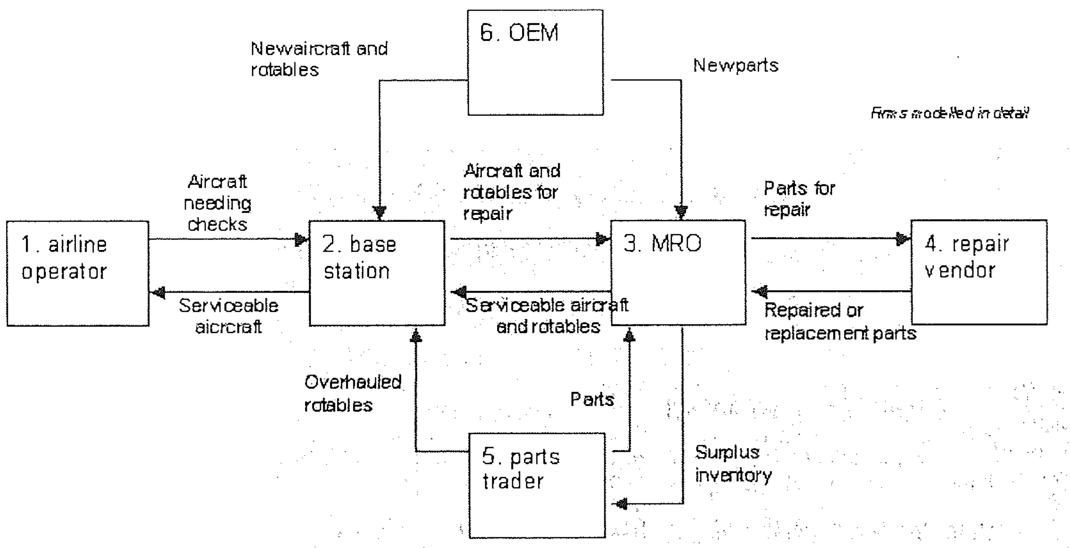

Figure 2.1: the aircraft maintenance bi-directional supply chain

The process mapping work has been validated with the subject firms named above, together with expert groups including the American Institute for Aeronautics and Astronautics and Massachusetts Institute of Technology's Lean Aircraft Initiative (MacDonnell 2004).

The aircraft maintenance sector can be viewed as the firms in the shaded area in Figure 2.1 (line station, MRO, subcontractor and parts trader). The airline and OEM interact with the sector but can be regarded as acting in separate sectors, namely airline operations and aircraft manufacture respectively.

With respect to supply chain conventions, the bi-directional nature of the supply chain in the maintenance sector is a key feature as it differs from production-oriented models, which are the staple of supply chain research (Beamon 1999, Burgess 2006). While the idea of reverse logistics is used to describe defective product returns and remanufacturing (Aras 2006), the maintenance supply chain can be considered a closed loop along the central axis in Figure 2.1. Unserviceable items (parts not fit for use) flow from airline to base station to MRO and subcontractor. This flow of products occurs "up" the supply chain, namely from customer to supplier, which is contrary to the norm. Repaired items return to the customers – in some cases the items are replaced with new, however it is more usual for items to be refurbished, perhaps containing new components.

It is useful to consider the flow of value around the supply chain, in the first instance between firms, but also within firms. The practice of value stream mapping (Manos 2006) is closely related to business process mapping and is helpful when considering flows of physical goods, flows of information and the strategic approach to business models. Value can be thought of as flowing in the opposite direction to payment: wherever payment takes place, a commensurate value must be provided in exchange.

Referring to Figure 2.1 again, the flow of value into the maintenance industry begins with the airline, which pays a line maintenance provider to maximise the operational availability of aircraft.

Aircraft maintenance events, where the aircraft is removed from service, are typically scheduled to occur following a set number of flights, hours or a combination of both. The rules for maintenance intervals depend on the nature of the system in question and are usually set by the manufacturer, although an airline may choose to exceed maintenance requirements.

2.2 Maintenance events and the role of rotables

Some systems are affected primarily by the number of flights accrued: landing gear, wheels, tyres and brakes come under most stress during landing, so the number of flights is more significant than the hours flown. Similarly, engines are run at their highest setting during take-off, so the main rotating parts (disk assemblies) will have a cycle (take off and landing) limit as well as a time limit. The restrictions on these life-limited parts will contribute to the maintenance requirement for the parent system, in this case an engine.

Apart from the systems mentioned, most other aircraft systems are monitored and scheduled for maintenance based on the hours flown. Thus for example an inspection for fatigue of structural components or corrosion of the aircraft skin will be scheduled after several thousand hours of flying.

Maintenance activities result from either planned routines or the failure of a component or system.

Maintenance arisings are the trigger for all maintenance activity and can be separated into airframe and engine events. An overview of the airframe and engine operating and maintenance processes is given in Appendix 3.

From a maintenance process perspective, rotable spares may be seen as items that can be exchanged on an operational aircraft in reaction to a failure, thereby avoiding the failure of the parent aircraft. This requires the provision of line spares, or Line Replaceable Units. The problem of planning the appropriate numbers of these spares is the focus of this project.

Thus the process triggers that determine demand for rotables are unscheduled failures arising during aircraft operations, leading to stochastic demand for spares.

2.3 The aircraft maintenance market

The main players in the aircraft maintenance business have traditionally come from the maintenance wing of large "flag-carrier" airlines, such as British Airways, Lufthansa and the surviving large USA companies, such as Delta and American Airlines. These firms continue to dominate and have continued consolidating. The majority of maintenance capability in Europe is now owned by Lufthansa Technik and SR Technics (which came from the failed Swiss Air). The OEMs, namely Airbus and Boeing, continue trying to move forward into the maintenance market with full-service contract packages for aircraft sales, which include rotable support and heavy maintenance. On the engine side of the industry, GE controls much of the maintenance market, with Rolls-Royce and CFMI moving to service-based contracts as well.

The trend toward low-cost carriers continues, with airlines favouring simple business models with single aircraft types, selling their own tickets on the Internet, and avoiding inter-line arrangements (partnerships with other airlines for ticketing and route sharing). With oil at $144 per barrel (markets.ft.com, 4th July 2008), a growing recession in 2008 and many disincentives to air travel (stealth charges, airport congestion, excessive security measures and general delays), it is conceivable that growth in air travel will decline, especially considering that much of the growth in recent years can be termed discretionary and price-sensitive.

The low-cost carrier model has led airlines to operate as "virtually" as they can, with minimal assets and outsourcing of as many functions as possible. It can be seen from the electronics industry that outsourcing is good in a growing market but equates to giving away value in the medium term, so it remains to be seen whether those airlines who outsource key activities will do well in the future.

In order to present the context for the detailed inventory modelling work that follows, this chapter has given a view of the types of organisation comprising the aircraft maintenance industry sector and the supply chain that connects them. The events leading to demand for the class of inventory considered - rotables - are stochastic operational failures. The functional purpose of rotables is to allow for the exchange of failed sub-systems without removing an aircraft from service.

Against a backdrop of rising costs and falling demand, coupled with observed poor practice in the planning of airline inventory levels, it is intended to investigate whether new models can offer reduction in the cost of inventory without compromising performance.

In order to introduce new models for improving the performance of inventory systems in the context identified, it is first necessary to review both current practice in the existing problem domain and also the relevant literature for inventory planning applied to this and other operational situations. This range of literature is evaluated in the next chapter.

Chapter 3: Literature Review

This chapter examines the literature on rotable inventory optimisation for aircraft operations support, drawing on theoretical and empirical sources from a wide range of perspectives.

The first section consists of broad background literature describing the industry, uses of information systems and supply chain characteristics.

The analysis is then approached in stages, becoming more specific and starting with descriptive literature on repairable inventory, finally concentrating on quantitative solutions to the aircraft rotable problem.

The optimisation problem is essentially a quantitative one, but different operating conditions and supply chain configurations need to be considered. Further, in the expectation of introducing improvements to the current inventory planning process, it is useful to assess the driving forces in the industry in order to derive scenarios to be modelled and tested.

Since rotable inventory is a special case – the inventory is not consumed but maintained as a set of spares in support of continuing operations – this review is confined to the management of inventory of this type. This differs from mainstream production-oriented practice, where inventory is consumed, which is a broader field of research.

The nature and focus of the papers reviewed vary in both their specific fits to the aircraft rotable inventory problem, and the level of detail in the analysis and modelling of the problem. Some articles entail a broad discussion of the problem and its operating environment, while others address the problem to be solved in a theoretical manner. A general classification of the literature is shown in Table 1, taking account of both the specific fit to the problem and the level of modelling detail in each paper. The top-most row in the table shows general literature, which is of interest to the problem are but does not address rotable inventory management.

While the majority of papers discussed here are from academic sources, there are also selected items from the industry press, which give useful opinion and data from the user base.

Nature	General repairable inventory, reverse logistics, aviation (Problem fit)	Specific to aircraft rotable inventory (Problem fit)
3.1 general	Aviation Week 2000, 2003 c, 2004 a, b, c, 2005 a, b, 2007 a Cassady 1998 Cheung 2005 a, b Flight International 2003 Garg 2006 Graves 2003 Jain 2001 Michaelides 2006	Nurmilaakso 2002 Piplani 2005 Porter 2008 Samaranayake 2002 Savaskan 2004 Smith 2001 Soh 2006 Spengler 2003
Quantitative	3.3 Baldenius 2005 Bashyam 1998 Brown 1984 Debo 2005 deCroix 2005 Depuy 2007 El Hayek 2005 Fortuin 1999 Fung 2001 Giri 2005 Guide 1997 Jung 1993 Keizers 2003 Kim 2007 Scarf 2002 Teunter 2008 Thonemann 2002 Wong 2005 Zhao 2005 Zorn 1999	3.5 Adams 2004 Airbus 1997 Airbus 1998 Armac 2007 Computer World 2005 Friend 2001 GE Engine Services 2002 Ghobbar 2003 a Ghobbar 2003 b Ghobbar 2004 Haas 1997 Kilpi 2004 Lee 2007 Logistechs 2006 Lye 2007 MacDonnell 2007 Shrebrooke 1968 Sherbrooke 1986
Descriptive	3.2 Aviation Week 2003 a Aviation Week 2006 Aviation Week 2007 b Bailey 2007 Buxey 2006 Fleischmann 2000 Kennedy 2002 Kranenburg 2007 Krupp 2002 Lapre 2004 Liberopoulos 2005 McKone 2002 Pati 2008 Sherwin 2000 Singh 1989 Srivastava 2008 Tedone 1989 Templemeier 2007 Yang 2000	3.4 Airbus 2001 Aircraft Technology Engineering and Maintenance 2001 Aircraft Technology Engineering and Maintenance 2007 Airline Fleet and Asset Management 2004 Aviation Week 2003 b Flight International 2004 Flight International 2005 Jackson 2003 LMI consulting 2006 Mabini 2002 Overhaul and Maintenance 2007 SAP 2007 a SAP 2007 b Weckman 2001

Table 3.1: classification of literature by nature and fit

In Table 3.1, articles are classified as listed below and in the sections that follow.

3.1 General: of background interest to the industry, trends in uses of ICT, changing business models, supply chain management, e-commerce.

3.2 Articles in the lower left quadrant are of background significance to the particular problem area and do not contain detailed mathematical models;

3.3 Articles in the upper left quadrant describe mathematical inventory models but are not an exact match to the problem area;

3.4 Articles in the lower right quadrant give useful qualitative information relating to the present inventory problem;

3.5 Articles in the upper right quadrant address the rotable management problem in a quantitative manner.

Each of these literature groups is discussed in turn, and in an order going from a general review of related discussion of the repairable inventory problem, to specific modelling solutions for aircraft rotable inventory.

The general group from Table 3.1 is addressed first, then the four classes of increasingly narrow focus above are addressed.

3.1 General

The airline industry is the least profitable of all major industries in the long term, with an average return on investment of 5.9% from 1992 to 2006 (Porter 2008). Thus there is clearly scope for new strategic approaches to managing businesses in this industry.

Considering the potential of ICT for improvement in business processes, it seems intuitive that e-commerce should thrive in the aircraft maintenance industry, where there are many interactions between buyers and suppliers. However, e-commerce exchanges launched by the major OEMs for use by their suppliers were met with little enthusiasm due to a lack of perceived benefit (Aviation Week 2000).

By 2003, two generations of e-commerce exchanges for aircraft spares had come and gone (Aviation Week 2003 a) – there are continuing initiatives for reverse auctions and the disposal of surplus stock, but in general efforts by both small and large suppliers have had limited success. A dozen exchanges were set up in 2000, with only Aeroxchange surviving in 2003 (Flight International 2003). Aeroxchange is owned by a consortium of 13 airlines and carries their purchase transactions to suppliers, with 3 million transactions carried out in 2006. However, a downside to e-commerce in the maintenance industry (in addition to a lack of tangible benefit) is the problem of price transparency (Soh 2006): it is not in a seller's interest to publish prices if this causes downward pressure on prices. Thus prices may not be published, making an exchange ineffective for automated quotation and buying.

The OEMs have been modestly successful on the selling side for taking orders for spares from airlines and MROs. Meanwhile, specialist organisations like Lufthansa Systems continue to develop web-based systems for maintenance management between operators and MROs (Aviation Week 2005 a). There are ongoing efforts in the manufacturing supply chain to adopt e-commerce, such as Exostar, which transacted

$23bn in orders among major contractors (principally in the defence sector) in 2005 (Aviation Week 2007 a).

e-commerce has been successful on the operations side of airlines, for selling to consumers (Smith 2001). Applying operations research methods to airline operations helps with the planning of aircraft capacity, which may become more flexible as selling models become more dynamic. Since aircraft utilisation determines maintenance activity, it is appropriate for airline operators to develop their use of capacity planning models to assist in maintenance forecasting.

The OEMs see their growth potential in the aftermarket (Aviation Week 2004 a) and continue to develop partnerships and new supply chain solutions with the MRO operators and companies like IBM are developing new business practices such as "Service Lifecycle Management" as they attempt to increase their presence in the technical aviation business (Aviation Week 2004 b). Major firms like Rolls-Royce are pursuing innovative supply chain relationships in an attempt to be more competitive in the maintenance market, but change is very slow (Aviation Week 2004 c).

There is clearly scope for applying new supply chain models to aircraft maintenance, for instance looking at flexibility in supply chains in manufacturing (Graves 2003), there is potential to view flexibility as uncertainty associated with the varying lead times for major maintenance items like engines. Also, there is scope for modelling and simulation of supply chain processes with a view to recommending improvements (Jain 2001), as well as the application of techniques like Systems Dynamics to modelling a complex supply chain (Spengler 2003).

The maintenance supply chain is a special case, as it is generally a closed loop (Savaskan 2004).

Clearly there is also potential for inventory reduction in the supply chain through the design and adoption of formal frameworks (Piplani 2005) and better use of systems generally: *replacing inventory with information* (Aviation Week 2005 b).

However, the maintenance planning process is complex and it falls to the operator to make the most efficient use of their resources (Cassady 1998), so attempts to outsource should not reduce the importance given to maintenance and fleet planning. The maintenance planning process should be seen as a core competence and a source of competitive advantage to operators, who will benefit from the use of expert systems and emerging techniques like genetic algorithms to process more complex planning models (Cheung 2005 a, b). Samaranayake (2002) presents a structured approach to maintenance planning developed with Qantas and combining project management techniques with specialist expert systems for coordinating complex tasks: clearly there is scope for continued detailed modelling of internal processes, over and above e-commerce initiatives for handling buying and selling processes.

Maintenance management systems in general are not well addressed in the literature (Garg 2006) and there is an identified need for new initiatives in the aircraft maintenance business (Michaelides 2006). There are also continuing developments in e-commerce technology and standards (Nurmilaakso 2002) that could be applied to the industry.

3.2 The repairable inventory problem

Before looking at the literature on repairables and spare parts generally, it is important to define the special inventory class that is the subject of this study. Rotable inventory (MacDonnell 2007) addresses the need for aircraft to have line replaceable units (LRUs), that is replacement items that can be provided and fitted 'on the line', meaning without removing an aircraft from service. Removed parts are repaired and returned to stock, so stock levels do not change and items do not enter or leave inventory (the supply chain) other than by purchase or disposal.

Sherwin (2000) presents structured frameworks for maintenance management and identifies the category of rotable spares, although their planning is not dealt with in detail.

From an inventory policy perspective, the special case for rotables can be simply defined as an inventory system where there is no change in inventory in the medium term. In the short term there will be spares removed, then the items that they replaced are re-stocked following repair. In the long term, there will be changes in fleet size and operating conditions, which will change the requirements for rotable levels.

In reviewing the literature, the constant question is the extent to which policies, models and recommendations fit the rotable case.

This section looks at spare parts inventory policy and practice and supply chain management and ICT issues in this area.

Kennedy (2002) points out that spare parts inventories need to be planned in a different way to inventories of work in progress (WIP) and finished products in a manufacturing environment, since spares are used to keep operations going in the face of unexpected failure. The motivation behind the paper is the perception that the inventory literature is not well developed in respect of spare parts and few papers address the problem extensively. Some ways in which spare parts requirements differ from production inventory are listed below, together with comments on how these observations relate to aircraft spares.

1. Maintenance policies, rather than customer orders, drive demand. Parts may be repaired or replaced. In the present work (aircraft rotables) it is always assumed that parts are repaired, but there are more general situations, such as machinery used in manufacturing.

2. Reliability information is not completely predictive, so the time at which demand will occur is a matter of (a) experience and (b) probability.

3. Part failures are often dependent – there may be systems comprising several inventory items, and the cause of a failure may lead to other failures. This will be ignored in the work on aircraft rotables since it is an intractable problem; further, this effect should be captured in the reliability data for separate items.

4. Demand is sometimes met through cannibalisation and loans. Again, this is ruled out for aircraft spares since, while it may be acceptable in an MRO situation, it is not acceptable for aircraft operations, the subject of the present work.

5. The costs of being out of a part are hard to quantify. This is refuted in the case of aircraft, since aircraft-on-ground (AOG) events can be costed – Ghobbar (2003 a) gives a value of $50,000 per hour – and the problem is addressed with a service level requirement.

6. Spares may become obsolete. In the case of aircraft spares, there is a formal approach to configuration management designed to address this issue.

7. Large assemblies may not be held if they are expensive, so a manufacturing plant may hold component parts, requiring repair work when the large assembly is required. This is not envisaged for aircraft spares since rotables must be available in serviceable condition to be stocked.

Further factors arise in relation to maintenance spares: for instance, planned and unplanned maintenance, multiple locations and levels of locations (referred to as echelons), which can supply each other, reliability changing over time and obsolescence. For the modelling work proposed later, demand is always unplanned, assumes a single inventory location (the observed norm); and behaviour over time is taken to be stable.

While there is undoubtedly scope to improve reliability by better failure analysis and improved preventive maintenance practices (McKone 2002), it is assumed from here on that all maintenance events of interest are unplanned (those requiring rotable spares to be supplied in response to stochastic demand).

The combined global cost of spare aircraft engines was estimated at $11bn in 1996 (Kennedy 2002) and leasing spares is suggested – however outsourcing needs to be paid for and engine spares are simply a function of repair time and will always be necessary. Total aircraft spares were estimated at $45bn in 1995 (Kennedy 2002) and measures are proposed to reduce this, including leasing and pooling of inventories and shortening repair times. The airline industry is considered slow to change and it is thought to be behind other industries in improving practice in inventory management.

In general, there are many models, most of them derived from observed areas of application rather than developed from a theoretical standpoint. However, many of these models continue to focus on order levels, such as multi-item service constrained models, which does not exactly correspond with the concept of rotables. Some models capture complex supply chain situations, where there are multiple echelons (inventory locations supplying each other, or distribution centres) and there may be indentured assemblies (an inventory item comprises other inventory items). These models aim to minimise the cost of spares for a given level of performance and are driven by the number of expected back orders, i.e., the proportion of failed requests. Economic order quantity (EOQ) models are adapted to spares planning and some further reference is made to repairable items. However, it has to be said that overall the literature gravitates towards spares being treated as replacement spares, with repairs being unexpected and not the norm.

In "Companies Get Creative In Their Inventory Management Solutions" (Aviation Week 2003 c), 82% of senior executives from a range of industries say that inventory reduction is a major concern, and the presence of large inventories is symptomatic of bad management practice, poor planning and poor cost control. They are relying on improved models for forecasting, coupled with better ICT systems to help reduce inventory with the aim of improving profitability. Most airlines use home-made inventory planning applications due to a lack of suitable products in the market, but the airline's own solutions are thought deficient. Aircraft downtime is quantified at between $23,000 per hour and $50,000 per hour (Ghobbar 2003 a) and maintenance spend is 12% of operating costs. Since inventory levels are critical in maintaining service, and inventory is seen as an asset class (not an operating cost) there is upward pressure on inventory levels, even though the high level of investment is seen as a problem.

Delta Airlines holds $1bn in inventory in support of 550 aircraft and 60% of maintenance is unplanned. Delta has implemented a supply chain solution to help allocate its $600M in rotables to the appropriate locations in its network to support unscheduled maintenance (Aviation Week 2003). This is really a simple problem to address, since demand is proportional to the portion of landings occurring in a given location, and does not focus on the decision about how many spares to hold to meet a service level standard. Similarly, America West focuses on supply chain above inventory level optimisation. American Airlines developed a decision support system to allocate parts to different locations (Tedone 1989), enabling service level calculations for respective bases. Once again, the focus is on optimising parts allocation across a network, which is a simpler problem than fleet-wide inventory optimisation.

It is interesting to note that Southwest Airlines, the model low-cost carrier has $360M in inventory at various USA locations, in support of 368 B737 aircraft. This is $1M per aircraft compared to Delta's $2M, however Southwest operates a single aircraft type, which is a narrow-body aircraft, while a large portion of Delta's fleet is wide-body (where costs are typically double). This shows consistency across companies, with little apparent benefit in inventory levels resulting from the single aircraft strategy. Southwest has implemented a new application for planning inventory levels, coupled with a new supply chain solution: they have reduced inventory levels somewhat but see the greater benefit as improvements in service level (Aviation Week 2003).

Recent experience and opinion regarding supply chain integration and outsourcing are mixed (Aviation Week 2006), with the growth of the integrated asset management (IAM) market. While this strategy has worked well in industries like manufacturing and distribution, airlines are slow to outsource ownership and management of spare parts. At the same time, the major OEMs are heavily promoting their IAM programmes, such as Boeing's Integrated Materials Management and Lufthansa Technik's Total Component Support. Clearly there are benefits to be had from pooling demand and entrusting supply to those with better systems, but the trade-off may be a higher operating cost in exchange for asset reduction: after all, the major OEMs are chasing this intermediate market since it provides ongoing profit potential. In this article, one consultant claims that the world rotable investment of an estimated $45bn only needs to be $14bn to meet service level commitments. For consumables and expendables, Boeing claims its IAM program can save an airline with 100 to 200 aircraft $10M a year for 10 years, but this is simply by eliminating retail margins on these parts and does not incorporate rotable optimisation. The pricing model for rotables is typically a by-the-hour arrangement: for each hour flown, the airline pays an agreed fee for complete cover of parts to a specified service level. Apart from passing the financial risk (and upside) to the vendor, a weakness with this model is that the airline is still responsible for downtime costs, so it might handle inventory decisions differently if it is trying to minimise consequential costs, whereas a service provider simply aims to reach a service level and is not penalised for the defined level of failures.

Another emerging practice is engaging repair vendors in IAM programmes: AAR, a major repair agency, offers by-the-hour support arrangements, which they claim can cut repair costs by 15% for an airline.

A more efficient solution than having third-party service providers (even if they are OEMs) managing spares is for airlines to cooperate on demand pooling. One shortcoming is maintaining configuration status, so that parts are really interchangeable. This can be overcome for minimum equipment list (MEL) parts, but airlines will always maintain their own items, such as seats and in-flight entertainment equipment (referred to as buyer-furnished equipment, BFE). There is a conflict in the sense that OEMs want to increase their role (for instance by hosting online exchanges) as they can profit hugely from the aftermarket.

In "Links In The Chain: Changing Dynamics In the MRO Supply Chain" (Aviation Week 2007) it is stated that of an estimated global rotable inventory of $48bn, airline ownership of these parts has fallen from 75% to 61% over an unspecified period. This is a result of moving inventory functions to third parties in a leaner supply chain, learning from automotive and retail industries. Thus in the space of a few years, it appears that control of the supply chain for aircraft spares is moving away from the OEMs in favour of the growing numbers of logistics and systems firms, like SAP, entering the market (Aviation Week 2007).

The relationships between operating efficiency, cost and quality are measured (Lapre 2004). Quality is measured as customer complaints, but it is deduced that high levels of customer complaints, while not an ideal measure, correspond with poor quality practice throughout the operation. This poor practice includes all areas of cost, and the study shows a clear correlation between low quality and poor financial performance. Thus the companies who showed poor quality metrics managed their costs badly and also had lower aircraft utilisation. Indeed, in the 1990s, the airlines with the lowest quality went out of business or went into protective bankruptcy. Of those that survived, their operating costs went up in the majority of cases. The companies at the other end of the spectrum, those with favourable quality metrics, all survived the same period and all showed reduced operating costs. The study also observed that, in undertaking new initiatives to improve operations practice, quality gains were achieved before cost gains. It can be inferred that systematic efforts to manage spares inventory better will therefore lead to improved inventory performance, better operational support and consequent cost reduction.

System vendors are traditionally weak in providing ERP systems to cater for spares, seen as slow-moving service parts (Bailey 2007), which account for 5 to 10% of most companies' investment in different industries. It is more important, however, to review policy critically than to look for external solutions to poor practice. An electrical utility cut 40% of spares inventory over a 4-year period, with a saving of $100M by reviewing practices and improving forecasting, with 70% of the benefit coming from the elimination of excess inventory.

The concept of green supply chain management (Srivastava 2008) applies to the recovery, reuse and recycling of end-of-life products, whose design does not lend itself to refurbishment. Discussion centres mainly on consumer electronics and cars. In this context, the life cycle of an aircraft compares well since aircraft typically have a long life and heavy utilisation, thanks to their design for maintenance. The environmental imperative may be a further pressure on airlines to reduce inventory holdings through demand pooling and better forecasting.

Buxey (2006) argues that current inventory theory, based mainly on economic order quantity (EOQ) calculations is limited in its view of well-defined demand and takes no account of supplier conditions. In the rotable problem, the supplier state should be the focus as demand inputs can be well defined. Buxey further observes that recent supply chain management thinking, and the attendant ICT systems, looks at operational arrangements at the expense of mathematical modelling and optimisation. The ABC system of stock classification, as used in the retail industry, prescribes high levels of safety stocks for the C group (lowest cost) with tight control of the A group. There is no evidence in the literature of this practice having been applied to rotable management, so it is an obvious candidate for improved practice.

Kranenburg (2007) shows that for capital goods in general, it is possible to cut spare parts costs by two thirds by having commonality of equivalent spares in support of different items of equipment. This shows both the effect of scale in reducing cost per machine and the benefit of aircraft fleet rationalisation.

There are several well-known general inventory ordering policies, which are governed by reorder point s, order size q, maximum stock level S, and review period r. The most popular policies (Templemeier 2007) are:

(s, q) – when stock falls to reorder point s, order quantity q

(r, S) – check stock every r days and order the difference between order-up-to level S and current stock

(s, S) – when stock falls to level s, order up to level S

There are variations on these policies, such as (s-1, s), or one-for-one and many differing scenarios, such as multi-echelon systems (with stock held in different locations that supply each other).

It can be inferred that rotable inventory will have a holding quantity value, S but q will be zero in the steady state, i.e., no new parts are ordered in the steady state since repaired items replenish stock. Only when s approaches 0 is some avoiding action taken, since backorders are not acceptable. How this is addressed is a matter of operational policy and varies from one operator to another, but typically airlines will expedite an item under repair, cannibalise a spare or lease a spare from a third party.

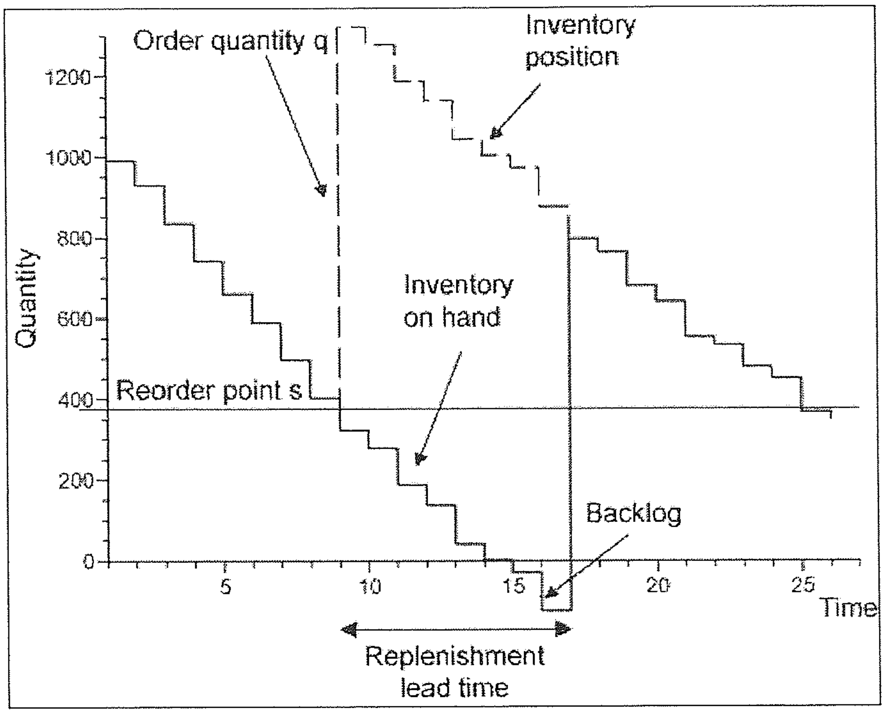

Figure 3.1: (s, q) or trigger point inventory policy (Templemeier 2007)

Figures 3.1, 3.2 and 3.3 show the 3 quoted inventory policies in graphical form and with large quantities. In the case of rotables, quantities are small and order quantities are individual items returning from repair.

Referring to the (s, q) model (Figure 3.1), the normal practice with rotables would be for an order (a repair order) to be triggered when a part is drawn, so s = 1 less than the steady-state inventory level, and q = 1.

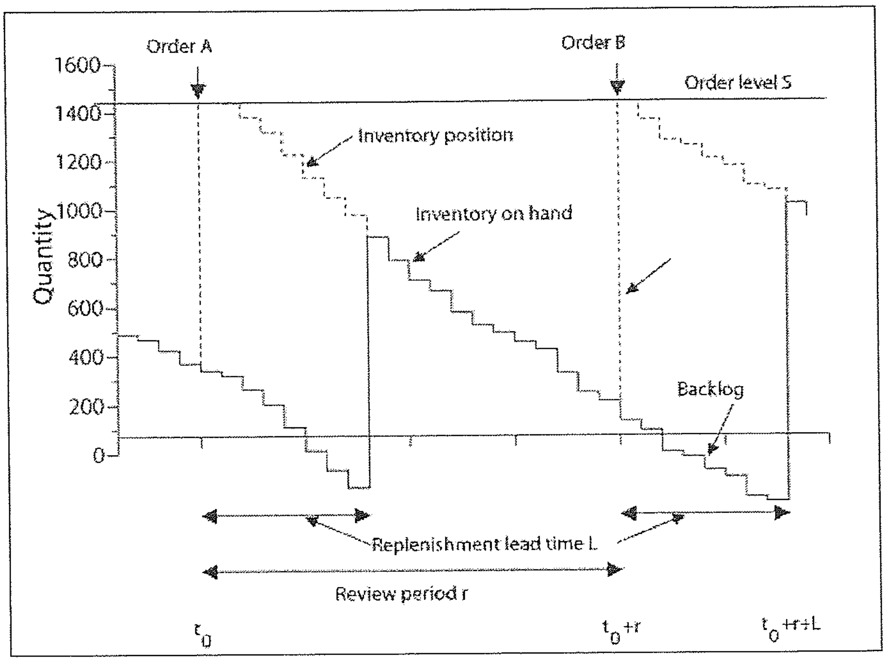

Figure 3.2: (r, S) or periodic review inventory policy (Templemeier 2007)

Considering the periodic review option (r, S), where inventory is replenished to level S every r days, for the rotable case r is variable and the time of a demand event and S is the planned inventory level.

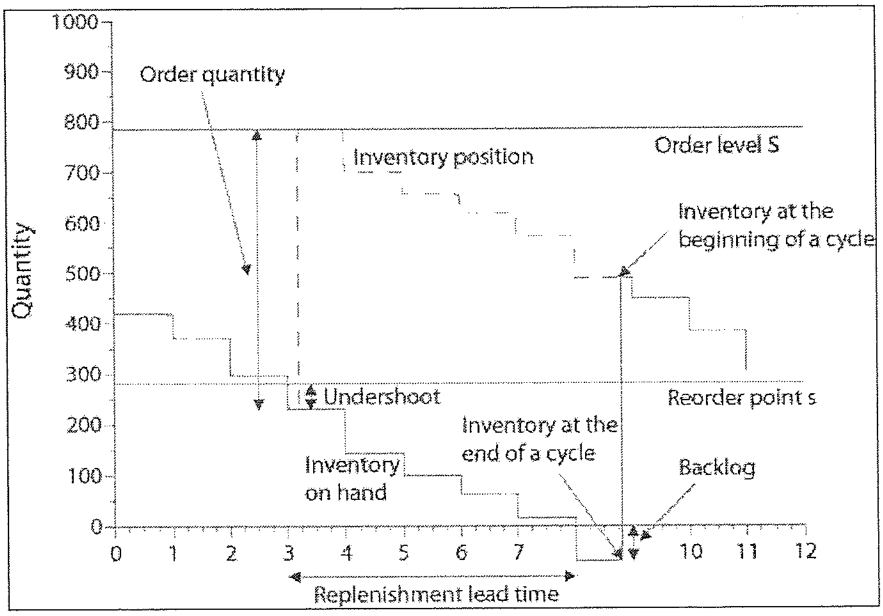

Figure 3.3: (s, S) or or order-up-to inventory policy (Templemeier 2007)

Considering the rotable problem applied to the (s, S) policy, s = 1 less than the planned level S, so it can be considered a simple case of this policy.

From a supply chain perspective, there are models that consider supply chains for product recovery and recycling (Fleischmann 2000, Pati 2008), but these are generally open chains, i.e., recovered products go to a new customer, unlike the closed loop for rotables – see figure 8 below. Thus, while there are network models for recovery, repair and recycling, the closed-loop quality of the rotable problem, with no changes to inventory levels in the medium term, appears to be unusual from in the supply chain field.

The repairable inventory problem, when viewed as a closed loop with steady numbers, can be likened to a Kanban production system, where a fixed number of work cards circulates through production process steps in order to minimise work in progress (Krupp 2002). This is a "pull" system, where stock is drawn in response to demand, as opposed to a "push" system, where stock is added based on a forecast (Singh 1989). There are different versions of Kanban, such as dual Kanban (with different card types for production and stock release (Yang 2000) and more complex versions incorporating advance demand information (Liberopoulos 2005). While these could apply to the single-item rotable inventory-planning problem, it is not apparent how they could model the system-level optimisation developed in this study.

3.3 Models for optimising repairable inventory

This section aims to present the main modelling approaches applied to repairable inventory, with a view to distilling elements of practice that could be applied to new models for rotable inventory optimisation. Inventory planning driven by service level constraints is sparsely covered in the literature, when compared with models based on order timing and volume (Bashyam 1998). Also, inventory planning for service parts, whether disposable or rotable, is not well supported by existing theory and systems (Fortuin 1999) but can be improved by pooling demand.

The Poisson distribution is accepted throughout literature and practice as the best fit to the rotable problem (Kennedy 2002). The Poisson distribution counts the number of discrete occurrences of an event in a given time period, where the mean expected number of events is known (Levine 2008). The Poisson distribution is commonly used for queuing problems and is thus suitable for modelling requests for spare parts. Poisson is available as a function in Microsoft Excel and is compared later with an approximation using a normal distribution. The formula is used in this study in its cumulative form, so that the probability associated with a quantity of 2 includes probabilities for quantities 0 and 1 also. The cumulative Poisson formula is:

$$\mathbf{E(x)} = \sum_{k=0}^{x} \frac{e^{-\lambda}\lambda^{k}}{k!}$$

Where E(x) is the expected probability of a value x, e is the base of the natural logarithm (2.718), k is the discrete integer variable ranging from 0 to x and λ is the mean expected value.

For example, if the mean demand for a part over a year (adjusted for repair time) is 2, then the expected probability of 3 events can be calculated by summing the probabilities for 0, 1, 2 and 3 events:

$$\textbf{E(3)} = \frac{e^{-2}\lambda^0}{0!} + \frac{e^{-2}\lambda^1}{1!} + \frac{e^{-2}\lambda^2}{2!} + \frac{e^{-2}\lambda^3}{3!}$$

$$= \frac{0.135\times1}{1} + \frac{0.135\times2}{1} + \frac{0.135\times4}{2} + \frac{0.135\times8}{6} = \mathbf{0.855}$$

What this means is that, with an average rate of demand of 2 items per year, then if 3 parts are held there is a service level of 85%, or there is an 85% chance that all requests will be met in that period.

While the expected value of a zero holding may appear illogical, what it implies is that if the mean number of events is 2, then there is a chance that there will be exactly 0 events – using a Poisson distribution with a mean of 2, that chance is 13.5%, the first term in the equation above. This means that holding 0 spares would give an implied service level of 13.5%. Clearly is there are any events at all, then they will not be satisfied, but the function says that on average there will be 13.5% since there may be no event. Holding 1 spare gives a service level of 27%, but this should include the probability associated with a holding of 0, so the cumulative chance that there will be 1 event is 13.5 + 27 = 40.5%.

Rotables are treated later in this review as a single-level group of inventory, where items are not subsidiary to others and do not affect each other's performance. This is the view taken of an inventory pool of line-replaceable units used to support aircraft in operation at the airline stage in the supply chain – the base station, supported by the MRO in figure 4 below. The distinction between line support and MRO activity is not well made in the literature, with many models focusing on heavy maintenance. Rotables feature in maintenance but are treated differently: the obvious example being jet engine components, which can be repaired and fitted to engines other than the parent. Further, engines consist of modules, which are theoretically interchangeable among engines, but typically involve heavy maintenance, so the benefit of the modular design is limited. The use of rotables in MRO is deemed too complex to enable full data capture and optimisation (El Hayek 2005) and is addressed by simulation, where modules deteriorate in accordance with the number of heavy maintenance events experienced.

A (T, S) inventory model with compound Poisson demand is described as equivalent to the (r, S) model (Fung 2001), where every period T, an order is placed to replenish stock to level S. Depending on lead time, a service level is predicted for various items. This is compared with an (S, c, s) model, where orders can be triggered at level c and must be triggered at level s. It shows that the service level is slightly lower for (r, S) than for (S, c, s), suggesting that it is better to drive inventory replenishments from inventory level triggers than by planning period.

A further model incorporates sharing among multiple inventory holders with the same type of inventory and the same type of customers. The (S, K) model has a rationing threshold, where $0 \leq K \leq S$, such that if its inventory exceeds K then it will entertain requests to provide inventory to other inventory holders (Zhao 2005). While the model is used for cars held by dealers, it shows that service is improved with less inventory thanks to sharing, but the communications systems need to be in place for the network to function well, and the correct incentives need to be present since there may be competitive and selfish motivations at work. It is important to consider the timing of shared inventory requests being filled: assuming zero transhipment time (from one inventory holder to another) will lead to non-optimal inventory stocking decisions and even where there is a delay, the benefits of having access to pooled inventory are substantial (Wong 2005). The present work does not considering inventory sharing, except in the scenario where utilisation is increased by pooling demand, but multiple locations are not envisaged, since the core operational problem is generally considered at a single location.

Guide (1997) reviews models for repairables, grouping them into single-echelon and multi-echelon and deterministic and stochastic variants. Most single-echelon models are of the (S-1, S) type, meaning that when stock falls by one, a replacement is ordered. Unfortunately, while the analysis is useful, it tends to revert to the production-oriented view from which it aims to differ, i.e., *the focus is on order quantity for items that are not repaired, rather than optimum holding quantities*. The paper recognises in its conclusions that there is a need to be driven by service level imperatives to manage repairables in an appropriate manner, rather than trying to apply order quantity models.

Another perspective is a model using a discrete-time replacement ordering system (Giri 2005), where demand is forecast based on the planned failure time of an item. This could be applied to aircraft spares but the only effect is to relax component repair times when the part is not expected to be needed: as well as being risky, there is no benefit in applying this model to rotables, as it is designed around consumed spares, where demand is forecast, rather than stochastic.

Failure rates may vary for repairable parts over time: they can become less reliable the more they are maintained (El Hayek 2005) or more reliable as new items are improved thanks to operational experience (Jung 1993). On balance, and for simplicity, it will be assumed here that reliability does not change over time, suggesting the use of stationary Poisson processes for predicting demand.

It is important to recognise the stochastic nature of events in planning inventory levels (Keizers 2003), both in the demand function and the repair cycle (Kim 2007). While the occurrence of failures is well represented by the Poisson process, the repair cycle is considered deterministic and must be carefully monitored given the large impact of long lead times on stock levels.

It is important to consider spares requirements as interlinked from a demand perspective: if a system contains many spares then its performance will be governed by the aggregate availability of spares, not the sum of each individual item's availability. System-level demand aggregation can lead to inventory planning with 15% less cost for a given service level than a line-by-line inventory plan for the same service level (Thonemann 2002); this is considered again later in the case of aircraft rotables.

Marginal analysis is reviewed in a complex generalised scenario with hierarchical systems (Zorn 1999) and is presented as a model with diminishing marginal returns, i.e., as the level of an item of inventory is raised, the accruing marginal benefit (divided by cost) falls. It is important that the marginal return decreases in a concave manner, i.e., a given change in marginal return for a change in inventory level should not be greater than an earlier change for a preceding increment in inventory. This is the downfall of the Marginal Analysis approach, which will be seen later in the Results chapter: put simply, Marginal Analysis may choose quantities in the wrong order. Thus if choosing a count of 4 for an item appears more beneficial than a count of 1, 2 or 3, then the final solution set may be logically inconsistent. This problem is explained further in the Analysis chapter.

3.4 Managing aircraft rotables

This section reviews recent publication describing the aircraft rotable management problem, without addressing the detail of mathematical models. It is intended to derive the prevailing industry trends and driving forces at play in order to inform the design of optimisation models.

The world MRO market, i.e., annual expenditure on aircraft maintenance, is estimated at $41bn in 2007, with 4% growth predicted on an ongoing basis. About $8bn is component maintenance, the largest market being the USA, with Europe second ahead of Asia (Overhaul and Maintenance 2007). As the market grows, however, the world fleet is also growing and utilisation is growing. With falling labour costs and improving practices (such as supply chain management and full-service models) unit costs are falling, i.e., the maintenance cost per aircraft seat-mile is falling.

Airbus advises three principles for managing the cost of rotables: commonality, reliability and punctuality (Airbus 2001). The first of these, commonality states that there are clear benefits in growing a fleet with the same aircraft type as that already operated, with a saved investment of around $1M per aircraft added when 10 are already supported. The second point, reliability, refers to both the OEMs continued aim for better component reliability and aircraft design leading to greater system redundancy for improved operational reliability. Finally, punctuality: the importance of managing the repair cycle is stressed, since the time taken to return a part to serviceable stock has a direct bearing on the number of spares needed.

Jet Blue, the American low-fares carrier, operates a single-type fleet of Airbus A320 aircraft. Their Initial Provisioning recommendation from Airbus was for $12M to support 10 aircraft (Aviation Week 2003 c); Airbus acknowledges that customers ignore their recommendations, and Jet Blue purchased just $4M initially. The following year, Airbus recommended another $12M in spares, the airline purchased around $5M as they anticipated growing failure rates. In addition to stock levels, Jet Blue emphasised extra training for their technicians: the A320 is an all-digital aircraft with many spurious electrical and electronic faults, many of which can be cleared with the appropriate knowledge and experience.

Some of the larger MROs have been looking at component management for some time, as a viable business in its own right (Aircraft Technology Engineering and Maintenance 2001). Companies like Lufthansa Technik (LHT) and SR Technics, which originated as the maintenance divisions of Lufthansa and Swiss Air respectively, view the provision of spares to other operators as more lucrative than maintenance services. Indeed, LHT stocks parts for Boeing B737 aircraft, although their parent airline does not operate that model. The business case for stocking parts is considered for market potential, repair capability, technology capability and economies of scale. For the Boeing B737, LHT pools stock with SAS Component. At the time of writing, all of these component management providers employed the Poisson method for demand forecasting (explained in the next section).

British Airways' rotable holding was valued at £500M in 2003 (Jackson 2003); a review of practice took place and new projects were introduced to improve inventory performance. BA continues to use the Poisson model and identified £21M of inventory to be disposed of without adversely affecting service levels. An issue with de-provisioning is the low market value of parts, which are very sensitive to the age of the fleet – this may reduce the incentive to realise inventory reductions.

Revisiting the inventory management situation some years later, there is increased pressure on companies to cut inventory and outsource component support (Aircraft Technology Engineering and Maintenance 2007) and a greater presence by the OEMs looking to share in the ongoing business of product support. However, while there is some use of online systems for component management, there is no reported improvement in planning techniques, and the problem of format standards is still a barrier to e-commerce.

While airlines are aware of the need to cut inventory, the global value of spares is estimated to have risen from $50bn in the late 1990s to $80bn by 2004 (Airline Fleet and Asset Management 2004), with airlines preferring to over-stock than risk Aircraft On Ground (AOG) events due to shortages. There is a lack of systems integration and managing inventory tends to be a reactive, rather than planned, process, which will lead to overstocking. There is disagreement over how inventory performance should be measured, with three proposed measures (Airline Fleet and Asset Management 2004), listed below.

1. Stock turns – how often a part is used. This is irrelevant to rotables, as discussed later. High turns mean high failure rates, which are never desirable, while low turns are hard to interpret for slow-moving items.
2. Material cost per seat mile: clearly, this should be minimised, but this is just a metric, with little relation to inventory levels.
3. Service Level: this is the only measure that should be considered important, as it indicates the performance of the inventory in terms of its usefulness, and is used as the driver for the models developed in this study. A secondary measure (inventory efficiency) is derived later, which is a gauge of difference between the inventory holding providing the optimum performance, and the actual holding.

The world's largest aircraft leasing company, General Electric Capital Aviation Services bought LogisTechs, who had developed inventory-planning software using Marginal Analysis (Flight International 2004). GECAS then launched its Aircraft Component Management Division to provide component management to customer airlines on a power-by-hour basis. This is good for airlines wanting to outsource functions, but given the pricing, it provides no incentive for airlines to take charge of their rotables – for a reasonably-sized airline, this can be expected to prove more expensive in the long term, especially since GECAS (like Boeing and others) see this as a lucrative activity. No airlines were thought to lease Line Replaceable Units before 2004 (LRUs are the rotables that form the focus of this study), but many airlines, including large carriers and especially the low-cost carriers, are entering into long-term support contracts with OEMs (Flight International 2005).

While GECAS operates the Marginal Analysis model for commercial use, a version of the model is offered in the defence sector by LMI consulting as the "Aircraft Sustainability Model, a systems approach to spares management" (LMI Consulting 2006).

The major ERP software provider, SAP, has identified the need for specialist solutions for aircraft MRO with its Aerospace and Defence version of its R/3 product (SAP 2007 a). They recognise the move by the aircraft OEMs into the aircraft after-market in search of new business opportunities. While rotables are catered for with tracking functions, there is no facility for the optimal planning of rotable inventory (SAP 2007 b).

Mabini (2002) developed a model for multi-indentured (hierarchical) modular assemblies with multiple echelons (locations) – this is a development of the problem addressed by the METRIC model (described in the next section): it is designed around the repair and maintenance of a complex system of modules, where components may be replaced, so it is a more general and complex problem than the rotable planning problem.

Weckman (2001) models the reliability of aircraft engines, characterised as complex systems of components, using a Weibull statistical process rather than the normally accepted stationary homogenous Poisson process. This allows for customised distributions to be developed, which change over time. While there have been studies to show that engine release times diminish after successive overhauls, it is assumed here that the reliability of rotables follows a stationary (no change over time) homogenous Poisson process.

3.5 Optimisation models for aircraft rotables

This section looks at published model specifications and experience with specific reference to the aircraft rotable inventory problem. There are two main approaches: planning parts at the individual level, and planning for systems of parts where demand can be combined and considered together.

Demand for spares may arise in several ways (Ghobbar 2003 a): due to hard time constraints (for example, a landing gear assembly must be changed after 500 flights), on condition (an item is inspected against a defined standard, e.g., tyre tread depth) and condition monitoring (real time diagnosis of performance, e.g., brake pad wear). Rotables can generally be considered as arising for maintenance on condition, meaning that their performance is observed to be deficient upon inspection, or often in operation. However, it is best to consider rotables as arising for removal through condition monitoring, since their failure will usually be observed during operation, so the removal does not typically result from a planned inspection. Ghobbar's analysis gives a detailed statistical analysis of parts demand for maintenance items but there is no optimisation involved.

Most companies owning rotable inventory follow practices recommended by the aircraft OEMs, chiefly Airbus and Boeing; these policies are typically limited to individual line item treatments. Airbus advise that 99% of their listed spare parts are rotables and recommends a simple calculation to derive the mean number of expected demand events of a given part in a year (Airbus 1997):

$$\mathbf{E = FH \times n \times N \times (1 / MTBUR \times 365) \times TAT}$$

In the above equation, E = expected demand, FH = average flight hours per aircraft, n = number of units on an aircraft, N = number of aircraft in the fleet, MTBUR = mean time between unscheduled removals and TAT = repair turn around time. Note that E is proportional to repair time, TAT: E is the number of failures multiplied by the proportion of the time that a failed part is unavailable due to being in the repair process. Thus, for example, if the number of failures is 10 and the proportion of time that a spare is in the repair cycle is one-tenth of a year (36 days), then the demand for the inventory calculation is 1.

Using the above calculation for initial provisioning (spares purchase when a fleet is first commissioned), the mean inventory demand figure is applied to a Poisson distribution to give an inventory count that meets a stated probability of demand being satisfied. Thus, if there is a need to meet 90% of requests for the above item, which has a mean expected demand of 1 (when TAT is taken into account), then the number of parts giving a probability of over 90% with a mean of 1 is 2. If 2 parts are held in stock there is a 92% chance that all requests are met, if the average demand over a long time is 1.

Several measures are proposed to reduce the cost of initial provisioning (Airbus 1998): price reduction (not a long-term measure), improved reliability, reduced shop processing time and cost-optimised planning. The last item involves some phasing of the provision of expensive spares, so that not all spares are bought at the outset but are introduced as the fleet accumulates flying time. This assumes that failures follow the Poisson distribution and needs to be managed carefully. The cost-optimised planning principle also introduces reduced service level targets for non-essential items (Airbus 2001). Thus, while the required SL for items of essentiality code 1 ("no-go") is maintained at 94 – 96%, values for items with dependent essentiality ("go-if", i.e., systems with redundancy) are assigned 85 – 92% and non-essential ("go") items are required to be provided 70 – 80% of the time. Thus average values of 95, 89 and 75% are prescribed, as compared with 95, 93 and 90% in common use.

Of an airline's spares inventory, rotables or line replaceable units account for 25% by quantity and 90% by value (Airbus 2001).

The line item-level calculations are developed further in the Models chapter later as the Poisson model (Model 1).

While the above treats items at the individual level, there is also scope for considering the combined demand for parts. Haas (1997) discusses a model for a multi-indentured assembly, an engine, which comprises three levels of indenture. Several stages of supply, or echelons, are also modelled. The model consists of an aggregation of Poisson forecasts for individual part demand to meet an overall service level target. However, there is no account of cost in the model, other than aggregation: since an engine is effectively a hierarchical grouping of modules, there is not much scope for cost optimisation.

Adams (2004) compares item-level and system-level approaches to aircraft rotable optimisation, concluding that item-level forecasting is the least risky but will over-provide spares. Meanwhile Marginal Analysis combines parts, takes account of costs and can be modelled for multi-echelon scenarios. However, while Marginal Analysis gives good results, they are not optimal. A genetic algorithm approach is also tested – this is a complex approach, which may improve on Marginal Analysis but is not necessarily optimal.

The Marginal Analysis method was first developed by Sherbrooke (1968) in a military setting, in the Multi-Echelon Technique for Recoverable Item Control (METRIC) model. It is interesting to note that the US Air Force investment in recoverable items (rotables) is reported at $5bn in 1967. The METRIC model addresses overall optimality of spares stocks and the balanced distribution of spares in a network with two echelons, or levels of supply: bases and depots. Bases are the locations from which aircraft operate (as in the supply chain schematic Figure 2.1) while depots are central inventory locations, usually with comprehensive repair capabilities. The model represents failed requests for parts as back-orders, so that a failed request survives until it is filled. The model is of type (s-1, s), i.e., a replacement is ordered when a part is taken from stock.

The approach adopted in this study differs from the METRIC model because:

(a) demand is presented here in mean terms, not recurring, so that behaviour over a planning period is represented by a Service Level;

(b) back-orders are not tolerated in commercial aviation – some action must be taken to satisfy the demand, usually borrowing a part or expediting delivery in the supply chain;

(c) Sherbrooke acknowledges that the marginal contribution of increasing part numbers should be concave – marginal contribution should reduce as the number of parts increases – but this is not the case and so the Marginal Analysis approach is flawed.

The METRIC model is improved on with MOD-METRIC and VARI-METRIC versions (Sherbrooke 1986), with better forecasting of expected back-order rates. The "best" results, closest to optimal, are derived by simulation and the new versions of the METRIC model are shown to be closer to those that are presumed optimal.

General Electric Rotable Services claim massive savings in inventory through the use of Marginal Analysis – see Figure 3.4. Where current practice achieves 77% service level with $50M in inventory, the same performance is claimed with $28M in inventory, a 44% reduction, through the use of Logistechs k2s (knowledge to spares) solution, in which GE holds a stake. Figure 3.4 also shows that, for the existing $50M in inventory, the operator could increase service level from 77 to 90% through optimal planning. The system provides demand forecasting, optimisation and simulation to customer including Air Canada and America West (Logistechs 2002).

Aston University

Illustration removed for copyright restrictions

Figure 3.4: cost / service level gain using Marginal Analysis (GE Engine Services 2002)

Studies have been performed on demand prediction for aircraft spares (Ghobbar 2004) and the best inventory policies for determining optimal stock levels for spares (Friend 2001, Ghobbar 2003 b) but they generally treat individual line items so are not explored further here, since it is intended to consider only pooled inventory here, i.e., the combined cost and performance of many parts together.

It is claimed that cost reductions of 30% can be achieved by pooling spares among airlines (Kilpi 2004) but there is a trade-off in short-term service level, since it will be necessary to incorporate a delay of typically 12 hours to allow for provision of pooled spares. This work once again considers individual line items.

A genetic algorithm model is used to find optimal inventory levels for multiple locations (Lee 2007) but is again confined to a single line item.

A simulation model is used to determine a re-stocking priority order for multiple bases holding rotables (Lye 2007), but this does not treat groups of parts as a system, focusing on the airline network and its distribution of demand.

In the above instances, there are models of varying sophistication, dealing with the problem of multiple locations. Marginal Analysis seems to be the sole model for grouping parts together with the aim of achieving a desired service level while arranging the inventory selection so as to minimise cost.

Computer World (2005) gives a detailed account of Southwest Airline's supply chain optimisation project, which uses a range of mathematical programming solutions to plan inventory levels for its fleet of 385 Boeing 737 aircraft, with an average utilisation time of 12 hours per day. Mathematical programming is applied to expensive, slow-moving critical parts, but the details of the model are not disclosed. An optimisation-based heuristic, Constrained Marginal Analysis, is used for faster-moving parts: this recognises the problem that Marginal Analysis has with parts with infrequent demand, namely marginal contribution that does not diminish continually. Through its supply chain optimisation project, which included reducing inventory in the supply chain as well as in stores, Southwest cut its 2003 budget for rotable purchase from $26M to $14M, identified $25M of excess inventory and avoided repair costs of $2M, while increasing service level from 92 to 95%.

A model has been proposed using a small example for illustration (MacDonnell 2007), which involves using a linear programming model to pick an optimum selection set of inventory levels for a connected set of parts. The resulting solution should meet a target service level at the least overall cost. It is predicted that this can be achieved on a larger scale by increasing stocks of cheap parts while reducing stocks of expensive parts. The consequential equivalent service levels for the individual inventory items will therefore deviate from the target service level but the global service level is maintained. An issue that remains to be addressed is the different essentiality levels of parts in the same inventory, as they have different service level requirements and should contribute differently to a connected solution. A version of this model has been tested with a large

MRO (maintenance, repair and overhaul service provider) SR Technics (Armac 2007), showing potential for a 25% reduction in capital investment in inventory based on a four-month trial reviewing new purchase requests. As part of the study, the company showed that a 2-day reduction in component repair cycle time would enable a reduction of $7.5M in their UK inventory.

3.6 The gap in the literature

With reference to the categories of literature discussed above, the following summary comments can be made:

3.1 The general background shows the scale and broad application of the problem area;

3.2 The descriptive material on managing repairable inventory sets a context in the field of operations management for the type of problem to be studied here;

3.3 The quantitative modelling of repairable inventory gives some useful general information, mainly for demand forecasting, but does not address the rotable problem directly;

3.4 The descriptive information relating to the rotable problem adds importance to the problem and provides some useful contextual information;

3.5 The quantitative modelling of the rotable problem defines the state of the art.

From the perspective of the rotable problem, the following conclusions are drawn:

1. Most inventory management practice is based on traditional models to determine order size, this is reflected in MRP system design;
2. There is limited recognition of the rotable problem, where there is no net change in inventory over the planning period – most work looks at order size, rather than the steady state;
3. Inventory is usually considered at the level of a single line item, with no relation between different line items with respect to demand satisfaction or cost;
4. Marginal Analysis is the best-developed model for cost-oriented inventory planning with multiple items, but:
    - It bases demand on expected back-orders

- It is not a true optimisation
- It requires marginal contribution to be concave, or constantly diminishing, which is not borne out by reliability distributions (such as Poisson), so is prone to error

The work that follows departs from these issues as follows, with reference to the numbered list above:

1. The problem is analysed in its own context without reference to production-oriented inventory models, so the inventory problem is seen as a joint set of stochastic events;
2. The steady-state problem is considered, where it is assumed that there are no net changes in inventory levels over the planning period;
3. The combined effect of rotable inventory is considered, where reliability is seen as the service level for a system of rotable spares and the relative costs of parts are a factor in determining stock levels;
4. A true optimisation is sought, where there is a single "best" solution (satisfactory performance at minimal cost) to the inventory problem as modelled:
    - Service level drives inventory decisions for the planning period (back orders are not considered);
    - The solution should be fully optimal through the use of mathematically optimal techniques;
    - The solution should be able to cater for Poisson demand distribution, where marginal contributions are not constantly diminishing.

3.7 Cases to be tested

Based on the current typical operating conditions in which rotable inventory is used, a range of cases is proposed for use in testing several optimisation models. The cases below encompass known current practice and a series of improvements, which are suggested in the literature and put forward by experts in the field, as steps to improve the productivity and economy of spare parts stocks held by operators.

The cases are given names (in italics below) for convenience, as these are referred to repeatedly in the experimental section of this work.

1. Current operational practice. ***Base***
2. Wider range of Service Levels prescribed by Airbus. ***Fewer***
3. Reduced Turn Around Time – a conservative TAT reduction of 5 days is proposed as a test case to reflect current drives toward lean supply chain operation and repair shop processing. ***Faster***
4. Larger fleet size – the effect of doubling the fleet size and consequent inventory utilisation is proposed. This is in line with current moves for rotable demand pooling among customer airlines served by consolidated component service providers like AAR and Lufthansa Technik. This case also illustrates the consequences for airlines employing fleet standardisation policies, where the number of aircraft types is reduced to give economies in flight crew rostering and maintenance. ***Bigger***
5. The final case to be tested reflects the above changes combined, namely the wider range of Service Levels, faster repair times and larger fleets. This case will test the likely set of conditions to be encountered by an airline adopting all of the recommendations for current best practice. ***Best***

	Cases – Chapter 3				
	Case 1 *base*	*Case 2* *fewer*	*Case 3* *faster*	*Case 4* *bigger*	*Case 5* *best*
Models – Chapters 4, 5	**25 test runs**				

Table 3.2: the set of cases for testing

Table 3.2 repeats the structure introduced earlier: the cases have been identified in this chapter and the models are introduced in Chapter 4 and developed in Chapter 5.

This chapter has looked at literature in a progressively narrow field approaching the identified research problem. While there are areas of inventory theory that border the rotable management problem, this chapter has demonstrated that there is potential to more directly address the chosen subject area in detail in order to develop new research.

The next chapter, Methodology, proposes optimisation approaches to test current practice and introduce improvements. The models proposed for testing – existing models based on the literature and new models proposed in the Methodology chapter - are then elaborated in the following chapter, Models.

Chapter 4: Methodology

The purpose of this chapter is to describe an area of interest, identify a problem for analysis within the area and develop a strategy for solving the problem. The philosophical approach to the work is also expounded before the detailed consideration of the problem, since the choice of methodological approach should have a bearing on the selection of problem area and the strategy used to process the problem.

The problem area being observed is the aircraft maintenance industry. A set of criteria is presented for the selection of a well-defined problem to be assessed in detail for the purpose of applying theory to effect a measurable improvement. An operational problem is presented and a suite of solutions proposed to perform optimisation on a sample set of data to assess a new solution, which employs linear programming to select a plan from among a large number of alternatives.

In choosing a well-defined quantitative problem it is hoped that a clear benefit can be illustrated by employing Operations Research techniques in a new way to address a substantial commercial decision problem. Namely, the problem highlighted is that of determining adequate inventory levels of rotable spare parts, that is, items that are removed from service upon failure, replaced with an equivalent spare, repaired and returned to stock to be used again. This problem is not well covered by conventional inventory theory, which is based around inventory leaving stock and being replaced on a planned basis.

Having outlined the problem and the data that describe it, several alternative solution methods are put forward. By combining these solutions with a range of operational cases for sensitivity analysis, a plan is established for conducting a sequence of solution runs using software tools developed for the job. The purpose of performing multiple solution runs with a range of cases is to evaluate the relative merits of the solution techniques in the first instance. However, the range of cases also provides an opportunity to consider the operational implications inherent in the different operating conditions likely to be envisaged by the user.

Next, variables are identified to represent solution output for the purpose of comparing solutions against current practice and each other, and to illustrate the variations within the sensitivity analysis.

4.1 Research philosophy and choice of methodology

4.1.1 Discipline – the field of research

Several areas of academic research and activity are applicable to this research:

1. Management Information Systems (MIS) - using Information and Communications Technologies to capture, store and process business information to assist in the running of organisations;
2. Operations Research (OR)- quantitative modelling and optimisation methods are applied to a numerical problem;
3. Management Science - this area can be viewed as more general than Operations Research, since it is not limited to purely quantitative problems and considers decision making and decision support in an operational context;
4. Operations Management (OM), in particular Inventory Management - the need to hold inventory in the problem context presented, and the applicable theory for best practice in inventory planning.

This work originates from a systems perspective, applying MIS thinking to predict potential for better practice in an operational setting. The focus then narrows, so that a practical problem is addressed in a level of detail where tangible commercial value can be presented as a result of a new solution to a significant problem in the field. The context of the problem study and identification is an industry sector and its generic business processes. The core work moves from a descriptive background to a quantitative problem, which has the advantage of deriving new solutions whose performance can be demonstrated in a convincing manner.

The interaction of disciplines is illustrated in Figure 4.1.

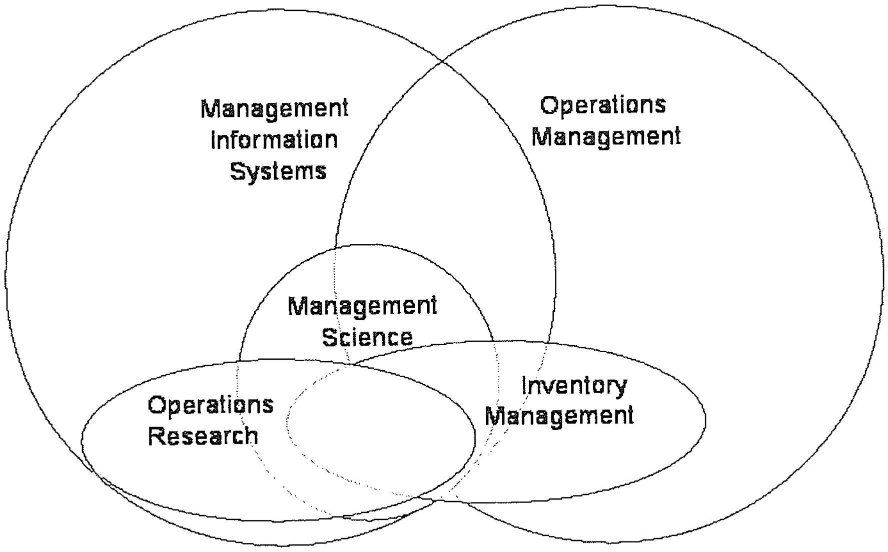

Figure 4.1: disciplines involved in the research

4.1.2 Research philosophy in the MIS, OM and OR disciplines

Positivism is a scientific approach to research, where the researcher observes practice and relates it to theory, in order to add evidence, refute and advance the theory. It is understood that the researcher is gathering data and does not distort the situation being researched (Bryman 2001).

The vast majority of IS research follows a positivist paradigm, although this changed from a reported 97% in the 1980s to 85% in the 1990s (Mingers 2001). Of IS research publications involving case study work, 87% are positivist (Dubé 2003). The vast majority of empirical research in the field of OM is positivistic (Bertrand 2002).

Interpretivism deals with broader situations, typically sociological phenomena, where the role of the researcher may influence the situation and the behaviour of subjects. Goles and Hirschheim (2000) explain interpretivism as follows: 'The interpretivist paradigm seeks explanation within the realm of individual consciousness and subjectivity, and within the frame of reference of the perspective: "social roles and institutions exist as an expression of the meanings which men attach to their world"'.

The body of work performed in this study is a mathematical modelling exercise, where well-defined, numerical operational data is processed in a variety of ways to observe the implications of using different solution approaches and to make recommendations for best practice in the planning of spare part inventory levels. The problem, data and solutions are fully determined and are not in any way shaped by the user or their attitude to the context in which the data is used. These constructs have the same meanings for all users that come into contact with them. Thus the work can be presented as scientific, not subjective, so an interpretivist paradigm is not appropriate.

The work that follows can therefore be said to follow a positivist paradigm, since it deals with objective data and information, which can be readily compared with theory and the resulting theoretical outcomes can be clearly stated in terms of the existing theories. Further, this work looks at a generic quantitative problem, where the nature of the organisation and its members do not have a bearing on the nature of the problem (although they may well influence how it is dealt with).

Finally, since this work aims to contribute to theory in the areas outlined, it seems appropriate to follow the majority in these disciplines where methodology is concerned, to facilitate the successful communication of the work to be performed.

Bell (2002) defines Strategic OR as "OR work that leads to a sustainable competitive advantage". It is hoped that the present work can be demonstrated to be significant in a business context. The work sets out to prove that a large commercial benefit is available through the use of OR in a manner that takes into account the cost–benefit trade-off of such techniques. A common problem with OR practitioners is that the effort to develop solutions may not be sanctioned because a clear business case cannot be proven.

Most research in OM has been quantitative empirical work (Bertrand 2002), which tends to be descriptive rather than normative, i.e., theory is developed to solve some practical problem rather than simply to advance theory in abstraction. Much OR work, especially in the USA, has been more normative, developing theory from theory with risk of becoming less useful in an applied problem area.

This work aims to both have a strategic imperative and follow an empirical, descriptive path where the potential to improve a problem in practice can be clearly shown.

In the field of OM, this work also follows an action research mission (Coughlan 2002), where there is an aim to learn from practice: here, a problem is taken from practice and theory developed around it to improve the practical situation.

Finally, the Strategic OR approach is consistent with the Business Process Re-engineering method used to set the context for this work: BPR seeks to introduce improvements in practice through the study, revision and streamlining of generic operational tasks (Kettinger 1997).

While the experimental work to be specified and carried out in this work is largely quantitative in nature (inventory modelling), there are qualitative aspects to the work, governing how problems are selected for study and how cases are chosen upon which to test new theories and solutions. Eisenhardt (1989) offers a framework for developing theory in concert with case studies, so that a problem observed in the field becomes the motivation for developing new theory. Theory building from cases has become more popular (Eisenhardt 2007) as it aims to benefit from rich qualitative data to help determine the importance of theory development in an applied context. Theory building from case studies recognises that problems are identified in the field and then tied back to the existing literature in order to plan a contribution to theory (Edmonson 2007).

There is justification for mixing quantitative and qualitative methodologies in studying a problem, where better results are obtained by triangulation (Jick 1979) – approaching a problem from alternative perspectives.

The use of empirical data in theory building has grown in recent years; further, the focus on manufacturing industries has lessened, so that there is now a better balance with service-oriented operations (Gupta 2006).

4.2 *Intended contributions to theory*

4.2.1 Operations Management and Inventory Management

In the operations management field, the provision of service spares is important in terms of reliability and cost. Where these spares are valuable and can be repaired economically, the inventory planning problem changes from the norm. This in turn has implications for the supply chain, including reverse logistics.

In the area of inventory management, mainstream practice and theory employ a "consume and replace" model for inventory planning, which often focuses on Economic Order Quantity, minimal holding stocks, stock turns and Just-In-Time supply. There is some treatment of the problem of planning spare part requirements, however this is often still a one-way flow, since many spares are considered consumable. The distinguishing factor in this work is that inventory can cycle or rotate (giving the term "rotable"). The fact that a part can come back into stock calls for important decisions and measurements that are not otherwise called for: the correct amount of inventory to hold and a measure of the efficient performance of a group of parts.

The question of inventory efficiency is interesting, since holding of spares does not correlate with production-oriented inventory models. For example, if a stock of ten spare radios is held in inventory to support a certain aircraft fleet, is this asset group behaving efficiently? The optimal holding number will be determined by the methods considered later. Meanwhile, if there are ten spares at the start of the year, there will be ten spares at the end of the year (plus any items in the repair cycle), so there is no measured consumption. The number of stock turns or items drawn during the planned period can be measured, but what is ideal? An efficiency measure is proposed later: this is simply the difference between actual levels and the optimal solution derived by the techniques here.

Comparing the operational need for spares inventory cover with production scenarios, it is apparent that spares are not productive: in an ideal situation, there would be no spares inventory. This would call for either perfect component reliability or on-the-spot repair schemes for all items (with zero repair times), neither of which is realistic. Thus the inventory items that are not in service should be viewed as a resource to be minimised under all circumstances. This different perspective on inventory presents an interesting problem for research.

Beyond the methods proposed here, a further efficiency improvement would be to manage inventory in a dynamic and responsive manner: given the stochastic nature of demand, it would be possible to hold low stock levels early in the life of a fleet, increasing stocks as demand becomes more likely. Also, it would be possible to hold low stock levels and take action when defined safety stock levels are reached, by sourcing additional stock externally.

It is well understood that, given that demand for spares is stochastic, pooling demand among groups of users will proportionately reduce the need to hold inventory. In other words, twenty aircraft will not need twice as much rotable stock as ten aircraft operating in the same conditions. It is intended to analyse this effect to give a clear measure of this scale effect in order to enumerate the potential gains from pooling demand. Considering operations practice in the industry, it should be worthwhile for maintenance providers to offer rotable inventory cover at a reduced cost to customers adding to their demand. Customers, meanwhile, should be able to outsource rotable support more cheaply than they could provide their own stock. This business model can be fully numerically evaluated using the solutions in this work, facilitating commercial outsourcing decisions.

The intended contributions of this work to the studies of Operations Management and Inventory Management are:

1. Determination of suitable stock levels for rotable spares
2. Measurement of the efficiency of rotable spare pools enabled by the calculation of optimal levels

3. Demonstration of the mass effect of rotable inventory pooling and evaluation of the cost / benefit of this practice

4.2.2 Operations Research

Operations Research offers theoretical optimisation solutions to well-defined, well-bounded problems with stable operating conditions and usually deterministic inputs. The present work uses as its input data a group of stochastic events, converting these into expected demand rates. Thus a goal programming solution is applied to a large-scale probabilistic event space by adapting the problem. Further, a range of heuristics is compared with the theoretical optimisation approach to see how well they approach ideal solutions. Also, these heuristics can be considered useful in similar problem situations where it may not be possible to formulate an optimisation solution.

In summary, the intended contributions of this work to the field of Operations Research are:

1. Linear Programming solution to the rotable inventory planning problem
2. Heuristics for rotable inventory planning
3. Linear Programming formulation incorporating stochastic demand events

4.3 Factors for problem selection in this study

Within the broad industry described in the Industry Background chapter, there are evidently many opportunities for the improved use of ICT solutions. This work is intended to perform an in-depth, rigorous analysis of a selected problem area within the aircraft maintenance industry.

The aircraft maintenance industry is chosen for research in the first instance in light of:

1. observed low adoption of new ICT solutions;
2. the size and high value of the industry and the implied gains offered by better ICT adoption;
3. the author's background experience;
4. availability of data provided by a field research project.

Suitable problems are sought with a combination of the following fields of knowledge:

1. the Aircraft Maintenance Supply Chain Reference Model resulting from an industry-wide business process mapping project, described in the Industry Background chapter;
2. observations and suggestions put forward by industry experts in discussions and arising from the process mapping project;
3. appropriate quantitative techniques - such as discrete event simulation and optimisation using linear programming - and Information Systems applications (such as decision support systems);
4. current practice in the selected problem area – inventory management, reviewed in the Rotable Inventory Management literature review chapter.

Given the observed low state of evolution of enterprise systems and decision support systems in the aircraft maintenance sector in particular, the aims for a focussed study can be stated as follows:

High value: there must be a large potential financial impact offered by studying a problem, in order to justify user involvement. Aircraft maintenance constitutes a large portion of operating costs (Airbus 1998) and more importantly, is considered to be an area where much can be done to control cost by devising new operational models. This is in contrast to other areas of cost, such as fuel, crew costs and airport charges, where there is little scope for the airline to introduce cost-cutting improvements.

Address a *known industry problem*, preferably one suggested by the industry as a current source of "pain". The area studied here is the planning of inventory levels for the more expensive classes of reusable spares inventory, which are not well catered for by ERP systems, as they are based on production-oriented (consume and replace) inventory models. Several participants have identified this as a high-cost area where there is poor planning and analysis.

It should be a *large-scale* problem: the problem must be logically non-trivial to solve. Some OR work looking at Supply Chain decisions – such as where to place inventory among a small number of bases – may be theoretically interesting but easily solved by a simple analysis

It should be a *generic* problem, represented within the generic business process framework for the industry and encountered by all equivalent firm types within the industry. Thus there should be no local or special aspect to the problem area addressed, nor should it be a temporary or transient problem. This implies that the problem will be well known and understood across the industry.

Any problem to be solved should be *computationally intensive*: there should be a large number of permutations and decision variables, which could not easily be processed by hand or by a simple calculation. For example, determining an Economic Order Quantity for a consumable material (e.g., engine oil) could easily be estimated from a small amount of operation data

There should be an opportunity to *apply mathematical optimisation techniques*: it is desirable to provide an efficient, theoretically sound approach to solving a problem in a non-obvious way rather than a "brute force" (or full enumeration of all possible variants) approach. In other words, in order to advance the application of the theory, it should not be feasible for a problem to be solved in as efficient a manner without knowledge of the OR theory to be applied. In this project a linear programming solution is used to find an optimal solution from a vast number of permutations, which could not reasonably be determined by brute force and could perhaps only otherwise be explored by simulation, which would not offer an optimal solution.

There should be an opportunity to provide *continuous improvement* through operational decision support on a repeated basis - not a start-up situation but a regular recurring issue. For a solution to be of value it should be applicable to operational situations that change over time, e.g., demand for spares varying with airline fleet usage, changes in fleet size and changes in the reliability, availability and cost of spares.

Access to data: finally, the ability to perform analysis depends on the opportunities to obtain data from the field. This data needs to be complete, usable and representative of realistic operating conditions.

4.4 Problem selection and statement

A large MRO organisation, FLS Aerospace, proposed the problem of rotable inventory management for study as it was launching a special project on the topic and expected that there could be mathematical methods available in which they did not have expertise. On examining the problem, the researcher saw potential for a radically new approach to optimisation using a linear programming model containing operational demand and reliability data.

The selected problem area concentrates on one class of inventory – rotables - there are four recognised inventory classes used in aircraft maintenance:

1. consumables – materials routinely used and replaced in operations, for example oils and hydraulic fluids, filters, brake pads and tyres;

2. expendables – low-value materials usually used only once in the maintenance process, for example fixings (rivets, screws);

3. repairables – materials and components that are not routinely maintained but are sufficiently valuable to warrant repair when damaged, for example fuselage panels, structural components, seats;

4. rotables – substantial functional assemblies that can normally be changed without severe disruption to aircraft operations. Examples include engines (changed overnight), pumps, radios, hydraulic, pneumatic and mechanical actuators, sensors and cabin equipment like fire extinguishers and galley equipment. Note that some literature (usually not specific to aircraft) considers repairables to be the same as rotables – in the case of aircraft maintenance, they are distinct in that repairables may be repaired occasionally whereas rotables are serialised, tracked assemblies with an expected life equal to the life of the fleet. Rotables are treated as an asset class and can be traded.

Line Replaceable Units account for 90% of all spares costs, and rotables (broadly equivalent to LRUs) constitute 9% of an airline's operating expenses (engines 4%, airframe 5%). Although rotables are not consumed, this large expense is attributed to depreciation of the initial purchase and associated financing costs (Airbus 1998).

The rotable planning problem is unusual in that the items of inventory may be put into service, removed, repaired and sent back to stock many times during the life of the fleet. Indeed, it is possible for rotable items to survive longer than any given aircraft. Rotables will be used interchangeably on different aircraft (and to some extent, on different aircraft types). Given that a rotable item is serialised and traced, it is possible for a rotable to have all its components replaced over its lifetime so that the only surviving part is its data plate.

Some rotables are subject to scheduled maintenance, with set intervals for removal from service. However, the majority of rotables are maintained "on condition", i.e., when they fail to meet performance parameters. Airbus (1997) gives the figure for Line Replaceable Units being maintained on condition as high as 99% of all LRUs. No more recent version of this statistic is available, but it can be considered to be representative of continuing practice. The trend is for all rotables, even engines, to be maintained on condition, in which case demand for spares is stochastic rather than predictable.

The selection of rotable inventory meets the problem selection criteria identified in section 4.3 above and as summarised in Table 4.1 below.

Criterion	*Fit*
High value	World rotable stocks are worth $bn
Known industry problem	Airlines and MROs are conscious of rotables being an area where large amounts are tied up and could be reduced MROs are increasingly offering managed component support packages as a value-added service to leverage pooled rotable stocks
Large-scale	Typical aircraft types require several thousand rotable items An airline or MRO will typically support 3 or 4 aircraft types
Generic	Rotable removal and maintenance are standard procedures in the industry, as prescribed largely by the OEMs
Computationally intensive	For several thousand rotable items per type, with spares stock levels up to 50 or more, there is a very large number of permutations of stock needed to meet target SL
Apply mathematical optimisation techniques	Non-obvious goal programming methods can offer fully-determined solutions to a large-scale problem
Continuous improvement	Fleet make-up and utilisation change, also reliability data should be updated, leading to frequent running of a solution
Access to data	Rotable inventory, aircraft operations and reliability information are all highly commercially sensitive and difficult to acquire for research purposes. The research participant, FLS Aerospace, provided operational data with a view to entering a joint venture to develop a solution. No other organisation has been willing to provide equivalent data.

Table 4.1 – problem fit to selection criteria

4.4.1 The chosen problem: Rotable Inventory Optimisation

A complex mathematical model is developed in this study to better predict the required holding of a range of rotable parts to support an aircraft fleet to a required service level. In the case of stock being used as spares, the holding may be the same at the start and end of a planning period, but the efficiency of this holding needs to be considered.

While airlines and their maintenance providers will make statistical calculations for individual holdings, it is possible to reduce the amount of stock needed while maintaining fleet service levels by considering the problem at a system level, rather than at the inventory line item level. This is best described by considering the objective function for spares provisioning. Rather than asking "how many of part X do I need to give a service level of 95% for part X?", the question should be: "how many of each part do I need so that for all requests for spares, a service level of 95% is achieved?". In other words, rather than looking at the reliability of one item in isolation, it makes sense to consider, of all the requests made for spares, what is the performance of the spares holding in total? Even though service level calculations are traditionally performed for each item alone, there are two ways in which the performance of holdings is relative among a pool of parts: (a) the relative frequency of demand and (b) the cost of parts relative to each other.

The mathematical model at the core of this study has been built and implemented with a major MRO (maintenance, repair and overhaul) company with rotable holdings of USD400M. Tests on stock sets of USD10M and USD50M have shown available savings of between 20 and 40% by using a linear programming solution. Some of the saving is explained by surplus inventory accumulated over time in reaction to short-term demand, but much of the benefit is due to the model's approach in assessing joint performance of the pool of parts.

Problem statement:

- determine optimal inventory levels of aircraft line-replaceable spare parts
- target is to meet a service level
- treat all parts together as a system
- minimise total cost for the whole inventory system

4.5 Possible solution methods

The chosen problem, rotable inventory planning, can be addressed in a range of ways, including known approaches used in practice and several proposed innovations. The choice of models and their detailed specifications are laid out in the models chapter.

The use of historical data is a manual method of planning, which is often the approach used in practice. It is discussed here for explanation, but is not included in the list of candidate solution method since it does not involve any type of modelling or optimisation.

Historical data

This item is not included in the numbered list of possible solution methods below, as it will not be considered further beyond this description.

Where a fleet and its spares stock are well established, continue with current stocks if SL is being attained, increase stocks if attained SL is below SL or reduce stocks if target SL is exceeded. Obvious weaknesses with this approach include:

(i) Determining Initial Provisioning stock levels depends entirely on partial vendor advice and uses the Poisson method tested later. There is no allowance for the deterioration in performance of a fleet over time (then again, none of the present methods takes this into account). If the size of a fleet is increased, the OEM will recommend expanding the spares pool proportionately – clearly the incremental spares requirement for a fleet of increasing size diminishes. For instance, if a fleet grows from 10 to 20 aircraft, or flight hours double for the same fleet, the spares requirement is not doubled.

(ii) There is no basis for knowing how good current stock decisions are, other than performance; thus there is little confidence in the planning process, especially when considering new maintenance contracts or fleet changes.

(iii) Evidence suggests that inventory levels creep up over time with this approach: idle stock will not be seen as over-performing. Also, personnel are not judged by having too much stock, rather their performance is seen as inadequate when there is a failure to meet the target SL.

(iv) This approach fails to highlight the consequences of poor performance of the repair cycle – in fact, the levels of stock required are directly proportional to Turn Around Time.

The main advantage of basing inventory levels on historical data is that, as a fleet matures, an airline or MRO builds its own reliability data based on actual failure rates. In the absence of field data, theoretical values are provided by the OEMs, whose engineers determine a "safe" service life or MTBR based on design and test information, which can be expected to be conservative.

An example of historical data being better than theoretical values far is a case from the researcher's own experience. In the early 1990s, Aer Lingus was one of the first customers for the CFM56 engines installed on Boeing B737-300 and -400 aircraft. This engine was one of the first new generation engines with Fully Automated Digital Engine Control, an electronic engine management system that replaced extremely complex and sensitive mechanical controls in earlier engines. The main function of these engine control units (which are themselves rotable assets) is to interpret a set of inputs (ambient air pressure and temperature, several engine state variables and pilot throttle setting) and meter fuel flow to the engine – this is a task far better managed by a computer than by a mechanical system. Also, the engine operational procedure called for "on-condition" maintenance with a very long potential release, or service life to its first major service event. Further, many design and material characteristics, such as a high bypass ratio (large main fan diameter) and an excessive power rating (the engine was designed to power larger aircraft so was operated with low levels of stress on the B737) meant that the performance deterioration of the engine was much slower than expected. While a comparable older-generation engine on a B737-200

(Pratt and Whitney JT8D) could be expected to run for around 6,000 hours before removal, the new engine type regularly doubled this value. This was a surprise to the operator, the maintenance provider and the OEM since there was little historical data available. However, this pattern became consistent and could be attributed to the engine's de-rating (running at a low power setting), pilot practice and in particular the environmental conditions in which the engine operated. Engines generate more thrust (and the aircraft more lift) with cold, damp air so have to work less hard than in hot, dry conditions, where the air is lighter. Also cold, damp air contains less abrasive material, which causes engine wear.

The outcome of this experience was that, contrary to recommended OEM reliability data, the operator used their own very different historical data to give far better forecasting performance. As Aer Lingus set new records for engine durability, the OEM revised its recommended removal data. Aer Lingus was able to reduce its holding of spare engines (which cost around USD5M each), leasing or selling spares to generate extra income.

In reality, while rotable removal rates are forecast (predicted), the actual removal time is stochastic as the decision to remove is based on performance rather than schedule. In the case of an engine, it is the Exhaust Gas Temperature at take-off thrust that determines the removal: when this value becomes close to the limit, the engine is removed. Historical data helps to determine the likely window for this event, but the actual removal time may be in a window of as much as three months, which is roughly the Turn Around Time of an engine. Thus strategic maintenance decisions are crucial in minimising inventory levels: engines will often be removed early in order to stagger removals and control spares utilisation. Given that an airline has a pool of perhaps several hundred engines, each of high value, it is justified to have a team of engineers managing their operation and maintenance, so the removal process is a mixture of predictive, stochastic and strategic decisions. This intervention approach is less viable for increasing numbers of lower-value items, so any model considered here envisages decision support resulting from data analysis, and not manual intervention.

In the case of planned maintenance for complex assemblies with indenture (contained rotables), such as engines, it would be appropriate to consider a simulation model to capture all of the complexity of the decision factors involved. However, since this study focuses on rotables removed on condition (stochastically) in line operations, the analysis is limited to the models outlined below.

4.5.1 Stationary Poisson Process

As observed from the literature, current practice entails a calculation for each part number (line number), whereby a number of parts is calculated to meet a service level target using a Poisson probability distribution. Demand is factored for turn-around time: a part is unavailable while in repair. This approach is deficient in several respects. It is discussed and worked through fully in the models chapter.

4.5.2 Marginal Analysis

The approach with marginal analysis can be stated as: "for each extra dollar spent on spares, choose the part that gives the greatest incremental contribution to service level per dollar".

This method is fully worked through in the models chapter.

4.5.3 Cost-wise skewed holding

This is a proposed heuristic, which simply groups inventory into several bands arranged by costs and applies lower service levels to the more expensive parts and higher service levels to the cheaper parts, with the aim of reaching an overall service level target. This is a trial-and-error approach.

This method is fully enumerated in the models chapter.

4.5.4 Linear Programming

It is proposed to state the rotable problem as a set of relations, which is amenable to full optimisation using linear programming. This is intended to choose the best solution (lowest cost) from all possible feasible solutions (combinations of inventory that meet a performance target).

4.5.5 Linear Programming - split

A second version of the linear programming problem is put forward, to deal with the issue of essentiality in a different way. The problem is split into three groups of parts (as there are three levels of essentiality) and each group is optimised separately.

The Linear Programming formulations and solutions are also fully implemented in the next chapter.

4.5.6 Rationale for model selection

There are two models known to be used in practice: Poisson and Marginal Analysis. These two methods are modelled and compared. The cost-wise skewed holding model is proposed as a heuristic – it is hoped that it will give results comparable to the marginal analysis approach, with far less effort. Finally, linear programming is offered as the theoretically optimum approach to selecting the best from among many combinations.

A further issue is that of different essentialities (level of importance of different spares). Where parts are treated as a system, a scheme of weighting will be used to represent different essentialities. Thus models will use a combined formulation to include all parts in a solution. Depending on the performance of the different models, a further step is possible whereby three solutions can be constructed to separately process each of three levels of essentiality. It is not clear whether this is warranted for all models: if the linear programming approach is far superior in outcome to the others, then the split problem will only be performed for that method.

In summary, there are two known techniques and three new methods are envisaged. Thus five models will be formulated and each tested on the same set of data obtained from operational information. The data will be tested on each model for a range of cases in order to fully test the behaviour of each model and observe the effect of changes in circumstances.

4.6 Experimental design

A sample data set has been obtained, containing operational information for 300 rotable inventory items held by FLS Aerospace to support the Aer Lingus fleet of 22 Boeing B737 aircraft. This sample is about one-tenth of the full data set. Due to the commercially sensitive nature of the data, it was not possible to use the full inventory database for this study, The sample data set is sufficiently large to permit extensive modelling and testing in an experimental environment.

Five models will be built, shown in Table 4.2 below. Each model will be run for five cases, which are derived from the literature and operational data. Thus it will be possible to compare the performance of the models across a range of cases to evaluate the models against each other and observe the effect of changing conditions on the models.

	Case	*1*	*2*	*3*	*4*	*5*
Model		*base*	*fewer*	*faster*	*bigger*	*best*
P	*Poisson*	P1	P2	P3	P4	P5
M	*Marginal Analysis*	M1	M2	M3	M4	M5
C	*Cost-wise skewed*	C1	C2	C3	C4	C5
L	*LP*	L1	L2	L3	L4	L5
L3	*LP3*	L3-1	L3-2	L3-3	L3-4	L3-5

Table 4.2: model run sequence

The five models to be tested are listed as the rows in Table 4.2: the first two models representing current practice as observed in the field and the literature. The last three models are new: the first, Cost-wise skewed holding is a simple heuristic, while the last two are two versions of a linear programming solution intended to achieve an optimal solution. All of the models are specified in detail in the next chapter.

The five cases to be applied to each model are represented by the columns in Table 4.2 and reflect realistic changes in operational circumstances, which, as well as giving insight into those cases, give multiple points at which to compare the models. These cases are summarised as follows:

Case 1 – *base* – the standard operating conditions in current practice;

Case 2 - *fewer* – lower SL values for items with essentiality below level 1;

Case 3 – *faster* – reduce repair times by 5 days in line with suggested improvements;

Case 4 – *bigger* – double the demand for spares to reflect a larger fleet, greater utilisation, or a combination of both;

Case 5 – *best* – combine the changes in Cases 2, 3 and 4, i.e., assume a greater range of SLs, faster repairs and higher utilisation.

The same data set is used as the base for all cases, so that the differences between cases can be easily assessed. The cases are the same as presented to each model, so that it is possible to directly compare the efficacy of each solution directly for each case. Finally, by plotting data for all cases and all models, it is hoped that clear trends will emerge across the cases to show whether certain models perform consistently better than others.

The contents of the data set and its preparation for formulation and solution are fully described in the models chapter.

The models are implemented using available software tools, chiefly Microsoft Excel® and a linear programming solver LP_Solve. While spreadsheet-based models are limited in functionality, they serve well as a test environment for this type of problem (Pasin 2005).

4.7 Results required and criteria for analysis

4.7.1 Expected outcomes from the different approaches

The objective of the study is to find the best method for optimising levels of inventory where the parts are subject to stochastic demand and removed parts are repaired and returned to stock. This arrangement is referred to as rotable inventory.

Several methods have been observed or proposed for solving this problem. All use the same data set and other operational parameters.

The expected outcomes for each method are summarised in Table 4.3 below.

Method	*Expected outcome*
Stationary Poisson Process	Not optimal, doesn't consider relative cost of parts
Marginal Analysis	Good results giving near-optimal cost
Cost-wise skewed holding	Approximation of Marginal Analysis, required trial and error but efficient solution
Linear Programming	Should give theoretically optimum results by finding the optimal permutation from among all possible stock level combinations
LP – split	May give better results but may suffer from less economy of scale

Table 4.3: summary of expected outcomes for the range of solution methods

The data set contains 300 line items taken from a larger inventory database and combined with operational data to give utilisation rates. The characteristics and properties of the data set are discussed in the next chapter.

4.7.2 Alternative cases to be tested

Sensitivity analysis will be performed on the data to give several perspectives on the solutions. The cases are derived from different operating circumstances suggested by the literature to reflect changes in practice and as outlined in the literature chapter.

There are thus five cases to be tested for each model:

1. Base case – *base*
2. Lower service level values for lower-essentiality parts as advised by Airbus - *fewer*
3. Reduced repair times – *faster*
4. Increased utilisation (doubled) to reflect pooling of demand among users or a rationalised fleet strategy by a growing airline – *bigger*
5. A combination of lower service levels, quicker repairs and greater utilisation - *best*

These cases reflect the changing conditions likely to be considered by the rotable planner, with varying target service levels, repair times and utilisation being the variables in the decision-making process.

Common practice in the industry has been to operate with a typical set of SL values of 95, 93 and 90% for no-go, go-if and go items. These figures are provided by FLS Aerospace and are quoted as normally recommended by Boeing (whose fleet relates to the data used here). This set of values is labelled SL_{FLS}:

SL_{FLS} = {0.95, 0.93. 0.90} for essentiality codes 1, 2 and 3

Airbus has recommended different SL value ranges as appropriate to their newer products and revised policy. For each of the three essentiality codes, ranges of 94 – 96, 85 – 92 and 70 – 80% are advised for no-go, go-if and go respectively. Rather than test these full ranges, the mid-point of each is taken for evaluation here, so the set of values is labelled SL_{Airbus}:

SL_{Airbus} = {0.95, 0.89. 0.75} for essentiality codes 1, 2 and 3

If demand can be managed so that target SL can be reduced by a small amount, there may be an appreciable cost saving. This demand management can be achieved in practice by reducing inventory and having an arrangement for emergency support for stock-outs. This is easily forecast by having a safety stock trigger. For example, if the lower SL results in reducing a spares level from 5 to 4, and if a safety stock trigger is set at 2, action can be taken to source spares externally if the level falls to 2. With a safety stock of 2, it is likely that outside spares cover can be arranged in time for 2 further failures to cause a stock-out. In reality, this can be a very cost-effective way of maintaining low stock levels and adopting a strategy where action is taken in response to stock changes in real time. It is therefore instructive to look at changes in the range of SL values to observe their aggregate effect.

Turn Around Time has a direct bearing on stock levels, which is often poorly understood. After all, if TAT could be reduced to zero, there would be no need for spares at all, as failed items would be repaired and replaced without delay. Some operators put great effort into managing rotables carefully but are at the mercy of repair service providers when TAT targets are not achieved. Further, it is important to highlight the entire repair cycle, which includes routing a removed part for repair and receiving it back into stock, as time lost in these internal processes contributes to the cost if inventory.

It is generally accepted that increased utilisation will give better use of spares inventory: for each aircraft added to a fleet, the incremental cost in spares is expected to fall. However, the effect can be predicted to be less in the case of a complete optimisation, such as a linear programming solution, since it should prescribe lower stocks more efficiently for a small fleet.

4.7.3 Verification of models

Quality assurance measures are applied to each of the models to ensure their accuracy. There are three general types of test performed, which are expanded in the Results chapter:

1. Manual checking of sample values – isolated results are taken and traced back through the solution, with calculations checked by hand;

2. Processing of output values in the simple Poisson model – the results from Models 2 to 5 can be put back into Model 1 and overall totals (cost, SL, fills) compared for consistency;
3. Sensitivity analysis – all of the models are run for ranges of SL target values to check that overall results are adjacent and consistent.

4.8 Point of departure

The work that follows develops new models for rotable inventory optimisation and also models reflecting current practice, for comparison. The models and their use differ from current practice as observed and in the literature in the ways described below.

1. Mainstream inventory theory expects ordered deliveries to replace consumed stock; rotable stocks do not change in the medium term, which is the planning horizon. Time scales can be considered in three categories:

    - short term, where there are demand events and replenishments arise from returning repairs;
    - medium term, where a steady state (no changes in stock levels) is observed;
    - and long term, where there are changes to the operating conditions, i.e., varying utilisation due to changing flight schedules, changes in the size and composition of the fleet and changes to some items of inventory and their demand rates, due to modifications, upgrades and new reliability data.

2. Some of the literature provides for back ordering of spares shortages; it is proposed here to model inventory being driven by service level targets. Back ordering – maintaining unfilled demand events until stock is available – is not envisaged in the models to be tested since:

    - the perpetuation of failed demand events is not easily modelled and does not fit logically with the service level model, which does not envisage failed requests repeating (persisting) but deals with mean demand over a planning period;
    - observed practice does not tolerate backorders, which could result in aircraft failing to operate, or Aircraft On Ground

events, with a consequent cost of $50,000 per hour – operators will borrow, lease or expedite a replacement part.

3. Current practice does not always take account of cost, each part being treated independently.
4. There is a method that treats parts together as a system, Marginal Analysis: it is proposed that a better optimisation is possible.

It is proposed to test known models against a new model (linear programming) to see whether predicted improvements can be achieved. For the sake of comparison, a common data set is used for all instances of all models.

In summary, this chapter has identified the discipline areas within which the experimental work is situated and justifies the choice of problem to be studied. Having defined the problem, a range of possible solutions is outlined. These solutions are to be tested on a range of test cases to allow comparison of the results of the solutions with the aim of ranking the solution methods.

Claims are made to the originality of the work by the nature of the solution approaches to be developed and tested.

The specification and implementation of these models is addressed in the next chapter.

Chapter 5: Model formulation and implementation

This chapter is the first part of the results of this work, as new models are developed for evaluation and comparison. Two models are based on known current practice, while three further models are developed to test new ideas for improved solutions. The set of solution values obtained from the models is then presented in the next chapter (Results) and these are compared in detail in the following chapter (Analysis).

	Models	Cases – Chapter 3
Current practice	**1 Poisson** **2 Marginal Analysis**	25 test runs
New models	**3 Cost-wise skewed holding** **4 LP** **5 LP - Split**	

Table 5.1: the set of models to be developed

Table 5.1 shows the 5 models to be specified in this chapter; the models are then tested for each of the cases proposed in Chapter 3 and the results presented in Chapter 6.

This chapter presents details of the data to be used in the analysis of the rotable inventory-modelling problem. The data set is then prepared to make it suitable for formulation into the five solution types proposed.

For all of the models tested, there is any underlying stochastic process, which is taken from current industry practice. The process is as follows: a part fails in service and a replacement is requested. The inventory planning decision is to give a set probability (e.g., 90%) that a demand event occurring at random will be satisfied with an available item. A factor in deciding how many parts to hold is the time taken to maintain and return a removed part to stock. A probability distribution is applied to all items in inventory to predict this demand over a planning period, which in this case and in usual practice is a year. This type of stochastic process is considered suitable for the Poisson distribution, which is recommended for a small number of occurrences of events in a period. Industry practice calls for a Gaussian (normal) distribution to be used instead of Poisson where the number of events exceeds 30 (as advised by Airbus). The effect of using the Normal instead of the Poisson distribution is assessed in the next chapter. Contrary to standard practice, an analysis of these distributions leads to a recommendation to stay with the Poisson distribution for all data points. The rationale for this recommendation is that the difference between the two implies that the Gaussian process will under-estimate events and lead to under-provisioning, which would cause actual service levels achieved to be lower than forecast. Given that finding, and for clarity, the Poisson distribution is used throughout this chapter in the model descriptions.

Five mathematical modelling solutions are proposed for comparison in recommending rotable inventory holding quantities. These are referred to as the Poisson, Marginal Analysis, Cost-Wise Skewed Holding and two Linear Programming models (one using the whole data set and the other with the data set split into classes corresponding to parts' levels of essentiality). The first model represents general known industry practice as recommend by aircraft and component OEMs and practiced by many airlines and maintenance, repair and overhaul operators (MROs). The second solution is a specialist industry solution (whose commercial implementation is now owned by General Electric). The third approach is a heuristic proposed as a new decision support solution in this work, as it should allow a high degree of manipulation for scenario analysis. Finally, the Linear Programming formulation was developed at an early stage in the work described here and as a result has been implemented as a commercial application. The Linear Programming solution is the only theoretically optimising solution – it is designed to choose the optimal configuration among all potential solution values. Two versions of the Linear Programming model are developed in order to cater for different classes of essentiality code, as explained below in the next section. The five solutions are described in detail with schematic diagrams and detailed process specifications. The technical details of the implementation of each solution are fully expounded and discussed.

In order to compare the different solutions, a set of solution parameters is defined, i.e., the controlling variables used in concert with the input data in the formulation and processing of each model. The most important parameters are target Service Level, SL - the required performance of the inventory system - and Turn Around Time, TAT: the time for a removed part to be fixed and replaced in stock, which has a proportionate effect on performance and required inventory levels.

In order to assess how each solution handles perturbation of the above solution parameters and give insight into the likely practical value of each model, a scheme is proposed for sensitivity analysis, with five different scenarios that reflect realistic conditions based on observation of practice.

Given five solution methods and five cases of operating conditions, a plan is developed with twenty-five distinct solution configurations to be run, and a set of output variables recorded.

5.1 Data requirement

To address the problem of determining suitable levels of rotable spares, a well-defined set of data is needed to describe the rotable inventory and its associated operational requirements.

There may be several thousand line items for rotables supporting an aircraft type. The present work will use a selection of several hundred items with the parameters below.

Data is needed at three levels: the part level, the fleet level and the global level.

The main part characteristics are described in general here, and in more detail in the subsequent section, for the actual data set to be tested.

Part data

Part number (includes interchangeable P/Ns) – unique identifier of a rotable assembly. Items should be serialised (have traceable serial numbers) although tracking by serial number is not performed at this level of analysis. It is important to maintain part number status and interchangeable values, since failure to do so may result in duplicate holdings. For example, different aircraft types may share the same engine type, so in some cases engine ancillaries may be compatible between aircraft types.

Item cost (Manufacturer's Current List Price, Gross Book Value, market value or replacement cost). Investment in rotable support typically costs several million dollars per aircraft (although this falls logarithmically with increasing fleet size). It is important to have a realistic value assigned to each item, and this policy must be determined at a company level. Current List Price may apply to modern fleet, but with previous generations of aircraft (say more than 5 years old) a more accurate value is the replacement cost, since a serviceable rotable item will probably be available on the open market at around 50% of CLP. The Gross Book Value may be used for insurance and financial accounting purposes, consisting of purchase price less depreciation, however depreciation is hard to assess for rotables and it makes better sense to value parts at their open market selling price. Airlines may choose to write down the value of rotable stock at the greatest allowed rate in order to write off depreciation against profits; on the other hand, some firms will want to maximise their balance sheet, so may seek minimal depreciation. In either case, it is important to apply realistic values and depreciation to rotables in order to accurately reflect the cost of holding depreciating assets. Where obsolescence is anticipated, either as increased depreciation due to aircraft market value reductions, or where expenditure is needed to upgrade and maintain parts, this should be added prudently to the cost of inventory.

Current holding quantity – the total number of spares of a given part number. This does not include items in service, but includes unserviceable parts in the repair cycle. This number may have originated from the OEM's recommendation for provisioning on commissioning of the fleet and will often have increased due to extra purchases. It is worth noting that the driving metric for rotable performance tends to be service level: failure to keep an aircraft in operation due to a shortage is seen as very costly in operational terms. Thus it is easy for inventory levels to creep up over time as there is less short-term focus on the cost of inventory than there is on operational performance. However, the true cost of inventory (capital investment, depreciation and holding costs) and the theoretical requirement for spares should also be examined.

Essentiality code (1 = no-go, 2 = go-if, 3 = go). Code 1 parts are those without which an aircraft will not be released for operation, such as landing gear. If a code 2 item is unavailable, an aircraft may be cleared to operate based on dependencies and conditions. For example, if a radio is broken, the aircraft may be cleared for operation if 2 other functioning radios are available. If an auxiliary power unit (required to start engines and power systems while engines are shut down) fails, the aircraft may proceed if it can keep an engine running during a stop, or if ground (external) power will be available at a stop. Code 3 refers to items that are not essential to the aircraft's operation, even though they may be normally available. For example, aircraft often operate with faulty in-flight entertainment systems, which, while frustrating for passengers, may be preferable to cancelling a flight.

Mean Time Between Removals – this may be recorded as the manufacturer's prediction of reliability, but should be updated by an analysis of actual performance. The manufacturer's advised MTBR data, based on design and testing, will be used in the Initial Provisioning calculation, where the airline operator and OEM determine suitable spares levels and negotiate the sale of a package of spares. This initial figure will continue to be used in the absence of field data. Note that the key figure is Mean Time Between Removals, which may also be recorded as Mean Time Between Unscheduled Removals, i.e., a stochastic event. From an engineering perspective, Mean Time To Failure is the accepted measure of reliability; however, in the case of aircraft parts, there are frequent instances of parts being removed without having failed, due to mis-diagnosis or the removal required for access to another item. It would possible to refine MTBR data by considering phasing in of new fleet, changing properties of ageing fleet and the tracking of serialised parts, but data is not typically captured in sufficient detail.

Quantity Per Aircraft – the normal configuration complement of an item on an aircraft, for example, in a twin-engined aircraft, most major mechanical items (hydraulic pumps, generators, air conditioning units) will have a QPA of 2.

Repair *Turn Around Time* – the total time from removal of a part from service to its replacement into inventory, where it is once again available for service. This time should include time to diagnose, book repair, ship, repair, receive, inspect and route to stock.

Fleet data

Number of aircraft supported – a homogenous fleet sample is required, i.e., aircraft of the same type, using the same rotables. Rotable groups for different aircraft types are simply treated as distinct optimisation problems.

Hours and cycles operated by each aircraft – in order to calculate the total demand for a given part, it is necessary to combine number of aircraft, quantity per aircraft and hours / cycles flown. This is then applied with the MTBR of the part to a statistical distribution to predict the number of parts needed to achieve a required probability (the service level) of a part being available.

Global data

Target service levels for each essentiality code: it makes sense to apply different service levels to the different essentiality codes, since their impact on operations varies. It is best to apply this at global level, since these variables relate to intended operational performance by the airline and are not a function of part or aircraft type. Typical target service levels could be 95%, 93% and 90% for essentiality codes 1, 2 and 3 respectively. A second set of values (95%, 89% and 75%) is used in some cases.

Assume number of stations = 1. The current model is based on inventory held in one location, with requests arising in the same location. It is a reasonably logically straightforward extension to envisage a model where inventory is distributed around a network in proportion to the number of flights arriving in each location, the cost of holding inventory in each location, and the cost of moving inventory between locations. Future work will incorporate these variables and simulate the problem to take account of the operational effect of time delays in moving inventory around the network, either as a re-distribution of demand or in response to immediate requirements.

Current method for calculating rotable inventory holdings: (i) manufacturer's recommendation, (ii) historic, (iii) stochastic.

(i) Manufacturer's recommendation simply means that, at the time of fleet acquisition, and periodically reviewed, the OEM will recommend a spares package to give a target service level. A simple consideration of the time-based reliability of equipment suggests that it should not be necessary to provide a full set of spares for a fleet at the beginning of its service life: if a part has an expected failure time of 5,000 hours, it should not be necessary to provide 5 spares on the first day or service. There is some conflict of interest here, in that the OEM wants to sell parts. There will usually be a protracted negotiation between airline (and often its MRO) and OEM to agree provisioning for new aircraft – this is referred to as the Initial Provisioning Conference for a new fleet.

(ii) Historically-based provisioning simply looks at requests for parts in the past planning period (typically a year) and, adjusting for changes in operating conditions, updates the requirement – past performance being a strong indicator of future needs, all things being equal. This does not take into account the ageing of a fleet but is simple and likely to be quite dependable. Excessive inventory levels are often overlooked, however, and there is little understanding of which parts are providing value for money, or being used efficiently.

(iii) For stochastic planning, where probability distributions are used, what are the distributions and on what basis are they selected? In practice, most firms use Poisson (discrete normal) distribution curves to predict the failure time where the holding is small in quantity. For larger quantities, a Gaussian (normal) distribution may be used, which is considered more accurate. Boeing and Airbus prescribe different rules of thumb, one advising changing from Poisson to Gauss when the holding quantity exceeds 20, the other when the number exceeds 30. The ideal distribution is one derived wholly from actual historical performance data, either at a fleet level or compiled globally by manufacturers, however this information is not currently available. In practice, given the small numbers of events over a typical planning period of a year, the Poisson distribution is adequate, particularly for fleet-wide decisions where rates of events are relative across a large range of part numbers. Poisson and Gauss distributions are compared in detail in the Results chapter of this study.

5.2 Data preparation

A Maintenance, Repair and Overhaul firm has provided a set of operational data for evaluation. This data set is a sample of 300 parts selected at random from an inventory database of some 3,000 items held by FLS Aerospace in Dublin for the Aer Lingus Boeing B737 Classic (-300, -400 and -500) fleet of around 22 aircraft in 2003. Aer Lingus has since disposed of this fleet. It was confirmed by the rotable inventory manager providing the data that there is no bias in the choice of the sample of data with respect to cost or demand rates.

The size of the data set is considered sufficient to create a demanding optimisation problem where the effects of any ill-fitting items will be absorbed by a sufficiently large sample.

The input data variables are described below.

1. *Sequence number* – an index from 1 to 300 is used for processing purposes.

2. *Part number* – industry standard part numbers are assigned by OEMs. While part numbers are sometimes treated as equivalent to or interchangeable with another item, for the purpose of this exercise part numbers are treated as distinct. Thus demand for a part must be met by that part number and not another. Obviously, part numbers have associated descriptions, but these are left out of this exercise to save space. Given that there is an index created for processing, it is not necessary to use the part number in processing, it is just a label.

3. *Part cost* – values in the data set presented range from USD700 to USD255,000. There are two sources of part cost data:

   (a) *Gross Book Value* – this may be the new value depreciated to reflect the age or remaining life of the part, or the current market value of the part. In the case of rotables, the GBV often reflects the age and currency of the fleet – for instance, early generation Boeing B737 parts are worth less than parts for New Generation models (-700 and -800), which are aircraft with lower noise pollution and lower fuel consumption. Depending on the type of rotable and its reliability information, rotables will sometimes be depreciated by the number of overhauls performed. In the case of complex and high-value items such as engines, the GBV will be a detailed calculation involving time remaining until the next overhaul and the anticipated cost of that overhaul, which takes account of the replacement of high-cost items like rotor disks, which may cost several hundred thousand dollars each. Given their very high value and complex maintenance decisions, engines are managed manually and are not listed in general rotable inventories.

(b) Manufacturer's Current List Price (MLP or CLP): the current or last quoted price for a newly manufactured item. This becomes less realistic for older fleets but is used in case there is no current GBV in the inventory system.

The data set used here contains over 80% GBV values with the remainder consisting of MLP values from the data set provided, and in a small number of cases, current market values obtained from part trader web sites.

The input values used for solutions in this study use GBV where available, with MLP used otherwise. It would be possible to use MLP scaled by the estimated consumed lifetime of the part (component hours / expected total lifetime) but based MLP values are used here as rotable item life figures are not available for separate items of stock.

4. *Essentiality code* – parts are assigned to three levels of importance (1, 2 or 3) with corresponding Service Levels targets, which are 95%, 93% and 90% in the present data set. An alternative set of SL values is used in two of the test cases, following recommendations for change in practice by Airbus. This set of values (95, 89 and 75%) is still a set of three and matches the usage of essentiality code. These

variables may be assigned different values according to operational policy, but will be given these values for comparative assessment in this study, unless otherwise indicated. Most of the solution methods in this study treat the group of inventory items together while seeking an overall target SL. These models need to be adapted to either run separate routines on inventory groups with different essentiality codes, or the figures need to be weighted to compensate for different target SL values.

5. *Service Level* – assigned by the essentiality code above, one of three set values for the problem case.
6. *Owned stock* – the actual number of a given item currently held by the inventory owner, to be compared with solution values.
7. *MTBRIP – Mean Time Between Repairs, Initial Provisioning.* Manufacturer's reliability figure used in the absence of historical data. These figures should be revised over time as fleet experience provides new data. MTBRIP should only be used before historical reliability data becomes available, i.e., when there have been few or no events.
8. *MTBR - Mean Time Between Removals*, the expected time to a demand event, based on actual historic reliability data. This figure is continuously revised and updated by engineering and rotable management personnel based on recent experience. There may be fleet-wide phenomena affecting MTBR. For instance campaign changes, or new repairs designed to improve the release life and performance of an item, will have a large impact on MTBR but can only be applied to parts that have been upgraded. Another factor is No Fault Found removals, i.e., items may be removed from service although they have not failed, due to misdiagnosis or the need to remove a part to enable the replacement of another part. In such cases, it may be possible to return a part to service or return it to serviceable stock following inspection and certification. With many operational issues affecting actual removal rates, it is therefore

important to constantly check and update reliability values to keep inventory levels at appropriate levels. MTBRIP is used where MTBR value is missing (meaning that the part has no failure history in the current set of operational data). MTBR is derived from TCH / REMS so is a calculated figure and is thus not strictly an input. However, it is shown here since a value needs to be chosen from MTBR or MTBRIP to have a complete set of data.

9. *TAT – Turn-Around Time,* the number of days elapsed from an item being removed from service to it being available in stock as a spare. This is the sum of the time taken to route the removed part for inspection, dispatch, repair, return and receipt into stock. Many outside vendors will have contracted guaranteed repair times and will often be obliged to provide either a replacement part or emergency part availability where the agreed repair time is exceeded. Rotable managers usually have a small set of TAT values. In the present data set the following values are used: 20 days where there is in-house repair capability, 28 for fast vendors or simpler parts and 38 for more complex rotables. An alternative case looks at a faster TAT, where these values are reduced by 5 days each.

10. *TCH – Total Component Hours,* the number of parts in service x the number of hours flown by the parent fleet over a set planning period, usually a year. In operation, the inventory system may be updated to give 'trailing twelve month' values, i.e., updated monthly to give the past year's values.

11. *REMS* – number of removals = TCH / MTBR. The REMS value is input by operations staff and reflects the number of requests for a part in the planning period. The number of fills (satisfied requests) divided by REMS gives the SL achieved – this is not included in the records provided. REMS is used to calculate MTBR above; where there are 0 REMS, then MTBRIP is used instead of MTBR. Since the MTBR figure needs to be selected thus, MTBR is used to calculate REMS in the preparation of the solution, so the REMS value is not needed.

A sample of input data is shown in Table 5.2 below.

Seq	Part number	Descr	GBV €	MLP $	Ess	SL	Stk	MTBRIP	MTBR	TAT	TCH	REMS
1	071-01478-0001	CONTROL PANEL, ATC	0	15,000	2	93	2	0	0	38	6,820	0
2	071-01503-2601	ANTENNA, TCAS	0	5,000	2	93	8	0	3,751	38	37,510	10
3	10-61312-9	JACKSCREW, FLAP TE	2,244	3,202	1	95	37	100,000	3,589	28	272,800	76
4	10-671980-1	EXCITER, IGNITION	1,429	2,700	1	95	24	12,500	21,864	28	371,960	17
5	10470-6	PUMP, STANDBY	1,340	54,975	2	93	14	16,700	19,486	28	136,400	7

Table 5.2: sample of first five records of input inventory data

Some observations and inferences may be made from a review of the data:

1. GBV – convert to $. Item 5 is shown at 4% of OEM cost - the reason for this is unknown. If the value is written down to reflect a short remaining life, then replacement may be far more expensive than Gross Book Value. If this is the open market value it will be the replacement cost. However, it may be prudent to revise the data such that the value used is GBV or 50% of MLP, whichever is higher. Where no GBV value is recorded, the full MLP value is used in the absence of any other input. If a part in the data has a value of 0 for both GBV and MLP, efforts will be made to obtain a market value from available sources, namely parts trading web sites. If no cost data is available, the part will be omitted from the model as it cannot be processed meaningfully.

2. Ess – three items are essentiality code 2, "go if", meaning that if another item of the same type is functioning then it is safe to fly without changing the part. These items will have backup systems in operation. However, many items, such as the 'trailing edge flap

jackscrew' seen in item 3, must be functioning for the aircraft to operate and are therefore assigned Ess value 1.

3. SL – this is a global input constant selected according to Ess above.

4. Stk – levels of stock currently held: this data is not used in the calculation of recommended stock holding, but is compared with the output value to give a measure of "efficiency" or how well current stock levels meet the target SL requirements.

5. MTBRIP – manufacturer's reliability data derived from testing and engineering calculations, used only where no actual removals have occurred.

6. MTBR – TCH / REMS, provided here to be selected over MTBRIP. MTBR is based on actual removal and flight time information so can be considered reliable based on operational history. It is important to compile actual MTBR data; where no removal history exists, it is prudent to consult other operators or solicit field data from the OEM as there may be significant divergence from theoretical MTBRIP values. Item 3 achieves only 3.5% of the expected life forecast by the OEM. This may reflect maintenance practice so may vary from one operator's fleet to another, but it can be taken as a solid indicator of future removal rates. Note that item 4 has about twice the reliability suggested by the OEM – in the case of this particular part, ignition exciters are paired and it is common practice to use only one of the pair on each flight. This shows that the OEM's engineering test data is very accurate, but operational practice changes the rate of failure.

7. TCH – total component hours vary widely for several reasons: the number fitted per aircraft, the number of aircraft in the fleet with this exact part number and the number of flight hours by the aircraft models containing this particular part number. There may be several part numbers available with the same part type, where several generations of part are in service. The part number may be changed by an upgrade to the part, so a part number may appear to have low TCH if there has been a recent upgrade.

8. TAT – turn around time is given one of three values in the present data set: 20, 28 or 38 days. This can be associated with two factors: in-house repair capability and the complexity of the typical repair workscope. Item 2, the Traffic Collision Avoidance System antenna (an aircraft-to-aircraft proximity warning device) is a specialist electronic item requiring customised testing and calibration and is given a long TAT, reflecting a long repair, outsourcing or both. Items 3, 4 and 5 have routine repair schemes, which are carried out in house and given 28 day TAT. They are not assigned the shortest TAT as they are somewhat complex to maintain. Simpler items like certain sensors and power supplies are more likely to have a shorter TAT of 20 days.

9. REMS – the number of removals varies widely. Item 1 is a cockpit instrument which is electronic and is not assigned an MTBR value, nor is there any history of failure, nor does the manufacturer provide a failure rate. Item 3, on the other hand, is a motorised moving part, averaging 1.5 removals per week.

10. For items where there is no failure history or expected time to failure (MTBR), there is no benefit in having the item in the model. Given that the information was extracted from a production database and to preserve the information, it is proposed to simply directly set the recommended inventory holding quantity to zero.

Some additional part information fields were provided but the information is omitted from the data set used for solution here, as it is not needed. These fields are listed below.

1. Part description – an abbreviated label, e.g., "PRECOOLER, BLEED AIR", used to describe a part and especially for indexed part searches. While this information is illustrative, it gives a limited description and is not required for processing.

1. SPEC – refers to Air Transport Association reference for parts classification. Category 2 refers to rotables, which applies to all of the parts in the database extracts used in this study. The same

originating production databases contain other groupings, such as chapter/category 1 for consumables.

2. Base stock – minimum stock level to be held at the main operating base. This figure is shown for each item of inventory but is not used in the calculations here, since it is assumed that the SL requirement will establish a suitable stock level. The base stock figure is largely notional, since the stock level advised by current practice is the figure used for operations; if demand events reduce the stock to the base level, there is little that can or should be done in the short term given the stochastic nature of demand. From the data provided, it appears that most base stock levels are 50% of the actual holding or a minimum of 2, which appears arbitrary. Where the base stock may feature in operations is in the use of multiple bases, such that the main base should hold a set level of inventory. However, the present exercise does not take account of multiple operating bases; clearly, aircraft land in different locations with rotable requirements, but the redistribution of inventory around a network is another problem, which should be treated subsequently to the holding quantity calculation. The optimal distribution of inventory around a network can be envisaged as a simple apportioning in proportion to the number of aircraft operations into each location. In practice, rotables are often delivered from the main base on the next flight on busy routes, and in many cases mutual loan arrangements are agreed between airlines with bases in complementary locations.

3. Loan out – some items will be loaned to other operators, usually overhaul customers as part of a commercial arrangement. These items may be available for recall if they are allocated as spares at outlying line stations. However, it is more usual that these items have been supplied in urgent (AOG = Aircraft On Ground) circumstances and are not available to meet demand in the short term. From a theoretical perspective, these parts should probably be replaced. However from a practical point of view the shortage is unlikely to cause a short-term stock-out. In operation, items subject to loan tend

to be expensive and component management staff will manage demand manually, so that there will be a reaction to falling stock levels to head off any imminent stock shortage. Further, the item replaced by the loan item will usually be routed back to the MRO's available stock following repair, so the situation should not persist beyond the part's TAT.

4. Total fitted – the number of a part in operation in the fleet. This number is needed to calculate the total hours operated by a part (TCH) for the purpose of calculating MTBR, given the number of removals.

5. Flight hours per aircraft – the average number of hours flown by the fleet in the planning period just past. This is the total number of hours divided by the fleet size and is subject to some slight error given fleet changes. This data comes from the maintenance engineering database, which gets information as weekly reports from flight operations.

Note that the operational data is reported at the fleet level, so reliability data is based on averages for a part number. In practice, for the more expensive and / or critical rotable items, the part will have a full serialised history maintained by the engineering department. It would therefore be feasible to both consider reliability for an individual part and to phase inventory demand according to the elapsed time for each item. However, this level of detail is not normally processed in inventory planning systems.

The steps below are proposed to prepare the data obtained for formulation into the various models.

1. Part number and description labels removed, these can be added after solution and are only needed for interpretation.

2. GBV € - convert to $ at 0.67 for consistency with MLP values.

3. MLP $ - taken as presented. Where GBV < 0.5 x MLP, the value used as cost for calculation is set at 0.5 x MLP, since some GBV values

appear very low. In the absence of any explanation for the low values, a prudent approach is assumed.

4. Ess – as given, 1, 2 or 3.

5. SL – set at 95, 93 and 90% for Essentiality code (Ess) 1, 2 and 3 respectively, except where indicated otherwise, such as in the cost-wise skewed holding method below. A second set of values is used, SL_{Airbus}, is used as an alternative test case. Also, SL values can be applied for scaling of cost values to weight items according to their essentiality in models where all items are considered in a combined optimisation, namely Marginal Analysis and Linear Programming. This is explained more fully in the section where the Marginal Analysis model is fully developed.

6. Stk – existing stock levels, as given. These values are not used in calculation but used subsequently to compare solution values with current practice. Note that current practice is to use the Poisson method below but stock levels may differ as inventory managers make overriding decisions in individual cases.

7. MTBRIP – used to substitute MTBR when no MTBR data exists

8. MTBR – calculated from TCH and REMS, listed here for selection between MTBRIP and MTBR

9. TAT – as given, used to compute the portion of planning period during which a removed item is in the repair cycle and therefore not available from stock.

10. TCH – as given, the time accumulated by all installed units in a fleet.

11. REMS – actual requests for spares from inventory in the past planning period, usually trailing twelve months where data is available.

The simplified set of clean data for input into the five solution models is therefore of the following form:

1. Index number, 1 to 300. Where a record is removed due to a zero value (see below), the original sequence is preserved to facilitate later comparison

2. Cost – GBV OR MLP OR 0.5 * MLP if greater than GBV OR market value; no zero values
3. SL – 95%, 93% or 90% for Ess code 1, 2 or 3; different values are used for the Airbus SL cases.
4. Actual stock – for analysis of results only, not used in calculations; no zero values
5. MTBR – from MTBR obtained from operations OR MTBRIP if no actual data is available; no zero values
6. TAT – 20, 28 or 38 days as supplied
7. TCH – total component hours in a year as supplied; no zero values

5.3 Combined models – catering for different Service Levels

The data set, in common with industry practice, represents inventory with three classes of essentiality and different corresponding SL values, or required levels of demand satisfaction. The Poisson calculation method - used in current practice and evaluated in the first model tested herein – treats each item individually without reference to the others. In this case, each item can have any SL value as the calculation only concerns that item. This method is deficient in that it ignores cost: it should be possible to maintain fleet-wide SLs for lower total cost by skewing holdings in favour of more heavily providing cheaper parts. Thus the subsequent models all assess the stock as a combined system. In each case there are three SLs corresponding the three Essentiality Codes. This reflects the fact that some parts are more important than others. Therefore, in deciding all inventory levels in a combined model, there is a shortcoming, namely that parts of lower importance will be considered (based on demand and cost) equal to more important parts.

A broad question therefore arises: should parts of different essentiality be treated in the same solution? The argument in favour of all parts being in the same model is that, the larger the inventory set, the more efficient the solution should be, i.e., satisfying say 95% of all demand across all parts is preferable to satisfying 95% of demand in each of three inventory classes. A converse argument applies: if there are different classes of parts according to their essentiality, then their demand and the satisfaction of that demand can be considered independent for each class. In other words, satisfying aggregate demand for "no-go" demand events with "go" or "go-if" parts does not follow logically.

In order to preserve the problem as one full formulation to be solved using different models, it is proposed to apply Service Level scaling to items with lower importance. Thus items with Ess Code = 1 are given their full demand, while items with Ess Code = 2 and 3 have their demand scaled down, or their cost scaled up as appropriate, to make them less attractive to a solution algorithm. Since cost is used in most solutions as a sort field, it may be better to scale demand for different SLs so that the cost data is not distorted. Thus a part with lower essentiality is given an artificially proportionally lower rate of demand to reduce its priority. There are several possibilities for evaluating the best approached to calculating scaling weights; these are presented in the next chapter.

The alternative to SL scaling is to split the data set into three smaller problems, sorted by Ess Code. Rather than doubling the number of problems and dividing each into three, the approach employed here is:

1. **Apply SL scaling to the combined models (all except the first, Poisson);**
2. **Select the solution method that gives the best outcomes;**
3. **Split the data set and develop three models to be solved by the best method, selected above (the three separate models do not employ SL scaling);**
4. **Compare the split problem with the previous best method.**

Testing has shown the Linear Programming approach to give the best results, so a second Linear Programming model set is developed, using a split data set. It is not considered necessary to repeat the evaluation of the three separate data set classes with the other methods since the Linear Programming approach is better by an incontrovertible margin, so there is no case for developing any of the other approaches further.

5.4 Models to be tested – implementation

The problem to be solved is as follows: "given a set of spare inventory items with stochastic demand and other operational data for each item, derive the best combination of inventory to meet a stated level of global demand with minimal cost".

The first obvious step is to identify and replicate known current practice. From interviews and literature, it has been established that Boeing and Airbus, the main commercial aircraft manufacturers, specify a Poisson process to be performed on all items; they do not treat the problem at the fleet, or system, level. This method is listed as the first item in Table 2 below.

Further discussion in the industry refers to a fleet-wide solution, Marginal Analysis, which simply put, says "for every incremental dollar spent on spares inventory, select the part that gives the greatest benefit by satisfying the greatest portion of demand for the least cost". This is the second item in Table 6 and is fully specified later.

It is believed that there are better options than the known solutions above, so two new approaches are developed for testing.

The Cost-Wise Skewed Holding method is a simple partitioning of the problem into a small number of divisions. The parts are grouped by cost and the aim is to over-stock cheap parts and under-stock expensive parts, with the aim of maintaining demand satisfaction levels while reducing cost. This is expected to show a significant gain over current practice, for minimal effort.

Finally, a Linear Programming solution is proposed: the objective can be stated simply as "of all available permutations of stock levels, find the one that meets the stated performance requirement for the least cost". Two versions of this solution are proposed: one combining all parts and the other a set of three problems, each with a different Service Level to reflect the existence of three levels of essentiality.

	Method	*Expected outcome*
1	Stationary Poisson Process	Not optimal, doesn't consider relative cost of parts
2	Marginal Analysis	Good results giving near-optimal cost
3	Cost-Wise Skewed Holding	Approximation of Marginal Analysis, requires trial and error but an efficient solution
4	Linear Programming – combined model	Should give theoretically optimum results by finding the optimal permutation from among all possible stock level combinations
5	Linear Programming – split model	Should give the "purest" solution without SL scaling, may or may not be better than the combined LP model

Table 5.3: summary of solution methods

The methods to be assessed are summarised in Table 5.3 above and are developed in detail below. The models are then implemented and tested on a set of operational data. The software tools used in the implementations are summarised in Table 5.4.

A single data set is used in all tests for the purpose of comparison.

Method	Implemented using
Stationary Poisson Process	Excel – probability functions and look-up tables
Marginal Analysis	Excel – look-up tables, macros, Visual Basic
Cost-Wise Skewed Holding	Excel – look-up tables, macros
Linear Programming (two versions)	Excel for probability functions VBA to output LP formulation LP_Solve Linear Programming solution engine VBA to retrieve LP_Solve solution Excel to analyse solution data

Table 5.4: software tools used for solutions

Each method comprises large sets of calculations, typically a matrix of 300 parts by quantities of 15, but up to quantities of 60. These two-dimensional arrays are well suited to spreadsheet implementations. Multiple variables derived from the arrays yield multiple overlaid calculation sets, which can be represented as three-dimensional arrays. In the interest of expediency, these are implemented as multiple linked spreadsheets. An alternative approach is to develop dynamic arrays generated by a customised software solution. These could be converted to multiple relational database applications. A commercial implementation of the solutions tested here would require dedicated software applications to be developed in order to facilitate automated processing. However, it is quicker to test the models and process the data in a spreadsheet environment for development purposes.

The only solution method that exceeds the capability of a spreadsheet is the Linear Programming approach. A dedicated Linear Programming application is used for this purpose: the problem formulation is generated from a spreadsheet, exported to the solver and the solution output returned to the spreadsheet for interpretation.

Note: the models are specified below; further sample screen captures are shown in Appendix 6. The full set of models, consisting of spreadsheets and LP_Solve model and output files, is available at *http://mmacdonn.ucd.ie/rotable.*

5.5 *Model 1: Poisson – current practice*

The standard practice in the industry today is to perform a line-by-line calculation for each inventory item.

Given utilisation data as Total Component Hours, TCH:

TCH = number of aircraft x quantity per aircraft x hours flown per aircraft

and reliability data, Mean Time Between Removals, MTBR

then the expected number of removals is:

REMS = TCH / MTBR

For example, 2 radios each on 20 aircraft, each flying 4000 hours gives:

TCH = 2 x 20 x 2000 = 80,000 hours

If the item is expected to fail at 4,000 hours (MTBR) then demand, or the number of removals in a year is:

REMS = 80,000 / 4,000 = 20

Thus it is expected that there will be 20 requests for a replacement radio during a year. However, given that rotables are repaired and returned to stock, the number of spares needed at any given time can be scaled down in proportion to the portion of the year for which an item is unavailable while it is in the repair cycle. In other words, if it takes a month to return an item to stock, then the actual number of spares needed is one-twelfth of the total number of requests for the year. Therefore the number of removals can be scaled down by an Un-Availability Factor (UAF), where TAT is Turn Around Time, the time taken between an item being drawn from stock and the item that it replaced being returned to available stock:

UAF = TAT / 365

For example, if TAT = 38 days, then

UAF = 38 / 365 = 0.104

In this example, the number of spares needed is then

20 x 0.104 = 2.08

Given that the failure is stochastic, it is necessary to apply a Poisson distribution, which is used for a relatively small number of events such as arrivals in a queue, or in this case, part failure around a mean value.

Cumulative expected Poisson values for a mean of 2 are shown in Table 5.5.

X	*Probability*
1	0.38
2	0.65
3	0.84
4	0.94
5	0.98
6	0.99

Table 5.5: Poisson probability values for mean of 2

What this means is that, if the spares requirement is 2, then if 2 items are stocked there is a 65% likelihood of a request being met. If 3 spares are held, the likelihood of a satisfied request, i.e., Service Level, increases to 84%. Therefore in order to exceed a Target Service Level of 95% for this item, it is necessary to hold a stock of 5 parts. Even though the scaled-down demand for spares is just 2, the Poisson distribution allows for the likelihood of 2 failures occurring while spares are unavailable.

The above calculation is performed for each item of inventory in turn, to give a full inventory plan to meet the target SL for each part. This is the method used extensively in the industry and as the Airbus Initial Provisioning calculation (Airbus 1997) for new fleet acquisition.

Drawbacks with this approach are:

(i) Granularity – since each part is dealt with alone and the stock numbers are small, the required stock number may exceed the target

SL by a significant amount. In the example above, the holding quantity of 5 required to exceed the target SL of 95% actually gives a SL of 98%, which is a waste. What is often done in practice is that the application performing the calculation gives a value, x, to exceed the target SL and also gives the SL achieved by x-1. In this case, x-1 = 4 gives a SL of 94%, so the inventory planner would be likely to choose this holding quantity, especially if it is a high-value item. This requires subjective review of the data. In the case that the solution values are followed, the achieved SL may be a point higher than the target SL, e.g., an aggregate level of 96% to attain a target of 95%, which could imply a cost of up to 20% more in inventory holding than the theoretical value to exactly meet the target SL.

(ii) There is no relation between parts – failures and demand for a given item are treated in isolation. This approach does not consider the relative cost of parts: it is desirable to hold larger stocks of cheap parts, so that they contribute more to satisfying demand for spares. An ideal situation would be that, for a 95% SL, or 95% of demand events satisfied, the 5% failed requests would be for the most expensive items – this would be the cheapest way to meet the SL requirement.

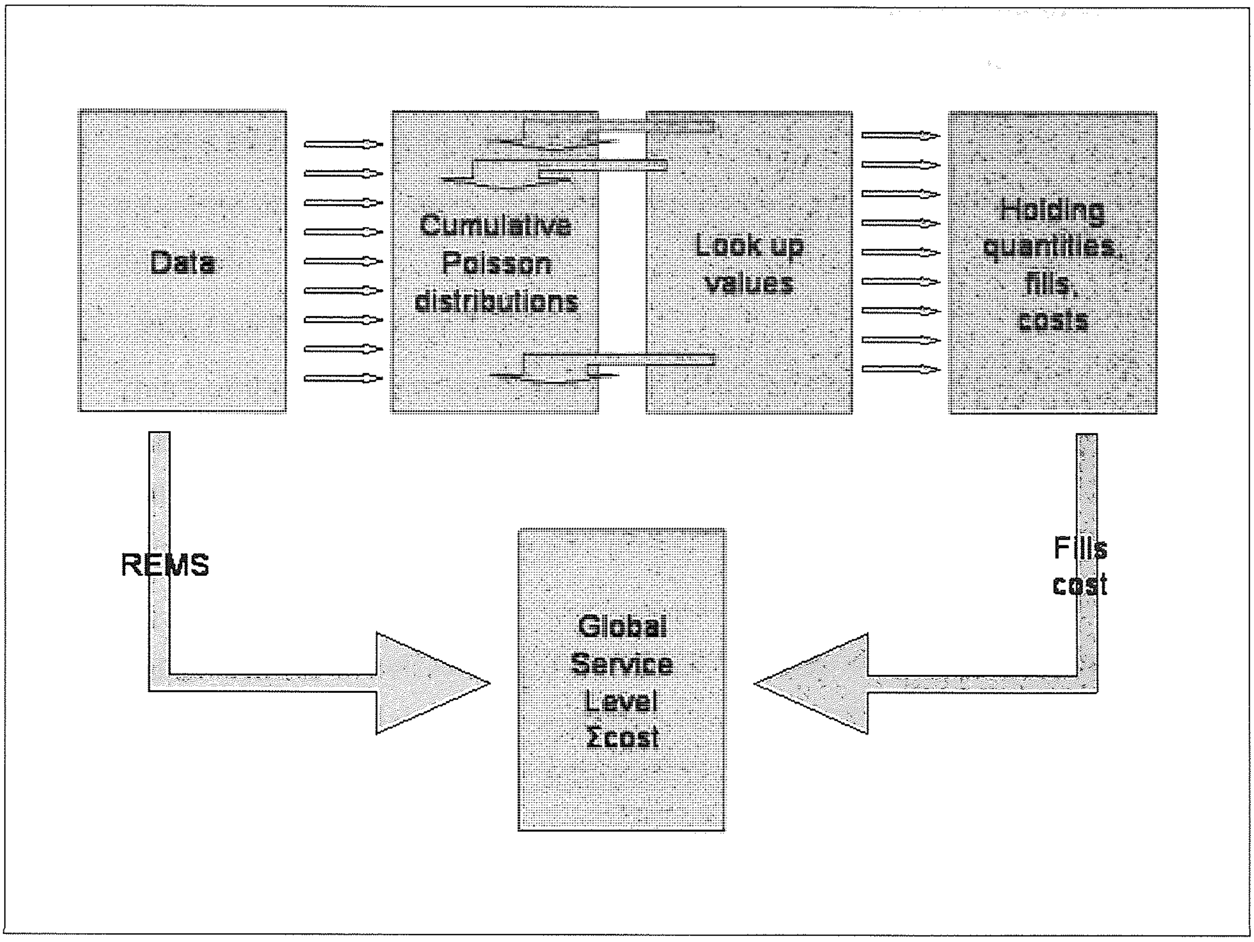

Figure 5.1: Poisson solution implementation schematic

In Figure 5.2, the Data box represents a list of 300 items extracted from inventory and reliability databases with the following attribute set (value range and units in {}):

Data [**Sequence number** {1 to 300}, **Part number** {string}, **Part cost** {integer $}, **Essentiality** {1, 2, 3}, **SLReq** (Service Level Required) {90%, 93%, 95%} or {95%, 89%, 75%}, **Owned Stock** {0 to 99}, **MTBRIP** {hours}, **MTBR** {hours}, **TAT** {days}, **TCH** {hours}]

The following is calculated for each item:

1. Mean Time Between Removals:

    IF MTBR > 0

    THEN MTBR = MTBR

    ELSE MTBR = MTBRIP

2. Removals:

Removals = REMS {integer} = TCH / MTBR

3. Un-Availability Factor, the portion of the planning period (one year) for which a part is not available while it is in the repair cycle:

   UAF = TAT / 365 {0.000 to 1.000}

4. Required inventory level, the mean number of spares required to fulfil demand:

   REQ {0.00 to 99.99} = REMS x UAF

   This number is left as a real number with 2 decimal places since, although it as an inventory holding quantity, it is a mean value and will be used as the input mean value in the Poisson distribution calculation, so rounding would distort the output.

5. A range of cumulative Poisson distribution values based on REQ: the Poisson function returns probabilities for values from 0 to 15 based on REQ as the mean value. 15 is chosen by trial and error as an adequate number to give a sufficient stock holding in 99% of cases to meet the target SL. With the current data set, there are a few parts needing larger holdings to satisfy demand – this small number of cases is processed with greater quantities to find optimal holdings.

   Sample cumulative Poisson values are shown in table 5.6.

		Expected value, x									
		1	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*
mean	*2*	0.41	0.68	0.86	0.95	0.98	1	1	1	1	1
	3	0.20	0.42	0.65	0.82	0.92	0.97	0.99	1	1	1
	4	0.09	0.24	0.43	0.63	0.79	0.89	0.95	0.98	0.99	1
	5	0.04	0.12	0.27	0.44	0.62	0.76	0.87	0.93	0.97	0.99

Table 5.6: sample cumulative Poisson distribution values

What the values in Table 5.6 indicate is that, for a given mean, and for each expected value, the cumulative probability of an event occurring is the corresponding fraction in the table. For example, reading down column 5 means that if we choose to hold 5 parts in inventory (expected value x), then there is a 98% chance of meeting

2 requests, a 92% chance of meeting 3 requests, a 79% chance of meeting 4 requests and a 62% chance of meeting 5 requests. Reading across the table, if the mean expected demand is 4, then in order to meet a 95% Service Level, we need to hold 7 spares. For each mean demand value, a holding is identified (highlighted in table 9) to meet a 95% SL requirement. Thus the stock requirements to satisfy 95% of requests for spares with mean demand 2, 3, 4 or 5 are 4, 6, 7 and 9 parts respectively. The same distribution is used for all part, where the mean is the number of demand events scaled down for the portion of a year that a failed item spends in the repair cycle.

6. A look-up process is performed to obtain the recommended holding quantity HQ from the Poisson table for each item in turn.

7. Another look-up process is performed to extract the corresponding actual SL **act SL** from the table, since the target SL may be exceeded by the HQ selected – this is the granularity effect explained in the Methodology chapter. For example, referring to table 9 above, where the mean demand is 3, the SL offered by 5 parts is 92% and by 6 parts is 97%. Therefore to meet a target SL of 95%, it is necessary to stock 6 parts, with a SL of 97%. This part will therefore contribute above its target requirement and will therefore increase the global SL achieved.

   Note: in current practice, the solution presents the holding quantity below the number meeting target SL, i.e., (HQ – 1), together with SL for (HQ -1). In the above example, the solution used in industry would return the following values:

   HQ = 6

   SL (HQ) = 97%

   SL (HQ – 1) = 92%

   The rotable inventory manager will then make a decision whether to advise holding HQ or (HQ – 1): this is a subjective decision and will be based on the manager's knowledge of the history of reliability and demand for the part, the value of the part and the firm's ability to source replacements or expedite repairs as a response to low stock levels.

8. The number of fills achieved is calculated:

   Fills {0.0 to 99.9} = REMS x act SL

   The number of fills is the number of successful requests for spares. Thus if there are 10 removals in a year and the SL achieved is 97%, then the number of fills is 9. Statistically, the number of fills is 9.7. In this model, the number of fills is left as a real number with one decimal place.

9. Global SL can now be calculated as:

   Global SL {90 to 99%} = Σfills / ΣREMS

   Given the granularity issue above (that the HQ value will often be associated with a SL greater than target SL), it can be predicted that the global SL will be above target SL. Initial results indicate that, even where there is a mixture of 3 SL values (90, 93 and 95%), the overall SL exceeds the highest target SL of 95%. The ideal target SL could be computed as a weighted sum of all parts:

Target global SL	=	**((SL1 x ΣREMS for all parts with SL1)**
	+	**(SL2 x ΣREMS for all parts with SL2)**
	+	**(SL3 x ΣREMS for all parts with SL3))**
	/	**ΣREMS**

10. The cost of each line item, and the total holding cost are calculated:

    Holding Cost {integer $} = HQ x Part cost

    Total Holding Cost = Σ Holding Cost

This solution method is implemented in a straightforward manner as a set of spreadsheets, one for each case tested. The cases are defined as test cases in the next section, and provide sensitivity analysis of the solution to reflect different sets of operation conditions observed in practice.

Since each line item is calculated independently of the others, there is no fleet-level optimisation offered by this solution. Further, since SL calculations are performed individually, there is no need for SL scaling, which reduces the relative importance of parts with essentiality codes 2 and 3, where code 1 represents parts that are essential for aircraft operation.

Case 1: SLFLS

SL 0.96 Cost 15285581

TAT reduction | utilisation

line	PartMaster	Stk	Ess	SL Req	GBV/RR	GBV $	MLP $	MTBR	TCH	TAT	adj TAT	REMS	adj REMS	value $	UAF = TAT / 365	demand	0	1		
2	071-01503-2601	7	2	0.93	0	0	9000	3751	37510	38	38	10	10	2500	0.104	1.04	0	0.72	[illegible]	[illegible]
3	10-61312-9	34	1	0.95	2244	3366	3202	3589	272800	28	28	76	76	3366	0.077	5.83	0	0.02	[illegible]	[illegible]
4	10-617980-1	21	1	0.95	1429	2144	2700	21864	371690	28	28	17	17	2144	0.077	1.3070	0.63	0.86	[illegible]	
5	10470-6	8	2	0.93	1340	2010	54975	19486	136400	28	28	7	7	27488	0.077	0.54	0	0.90	0.98	[illegible]
6	107484-5	1	1	0.95	1507	2261	6050	8974	170500	28	28	19	19	3025	0.077	1.46	0	0.57	0.82	[illegible]
7	107492-2	4	1	0.95	3413	5120	7291	5602	156860	28	28	28	28	5100	0.077	2.15	0	0.37	0.64	[illegible]
8	108032-6	16	2	0.93	277	416	1835	7199	64790	28	28	9	9	918	0.077	0.69	0	0.85	0.97	[illegible]
9	109486-5-1	17	1	0.95	2423	3635	14180	1705	105710	28	28	62	62	7090	0.077	4.76	0	0.05	0.15	[illegible]
10	114029	6	3	0.9	508	762	3000	28741	201190	38	38	7	7	1500	0.104	0.73	0	0.83	0.96	[illegible]
11	1211175-011	2	2	0.93	0	0	9000	40920	81840	38	38	2	2	4500	0.104	0.21	0	0.94	1.00	[illegible]

Figure 5.2: Poisson implementation in Excel

Figure 5.2 shows the top left corner of a Poisson solution (Case 1). Each row represents a line of inventory. For each line there is a required SL, based on essentiality (for the first item, an Essentiality Code of 2 corresponds to SL = 93%). Looking at the first item (line item 1 is discarded as it has removals recorded as 0), there are 10 REMS expected per year. Un-Availability Factor is the portion of the period during which a part is in the repair cycle, in this case 0.104 or 10% of the time. Thus the requirement is calculated as 10 * 0.104 = 1.04 items to be held in stock on average. This mean is applied in the Poisson distribution (with expected values 1, 2, 3, 4,...). Given a target SL of 93%, it can be seen that the first value to exceed the target SL is a quantity of 3, highlighted in Figure 5.3.

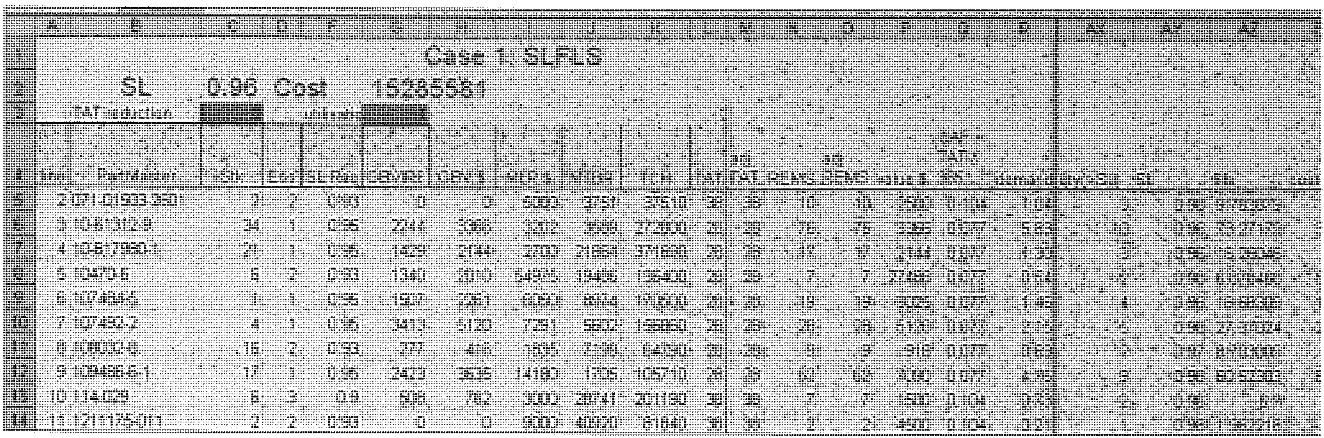

Case 1: SLFLS

SL 0.96 Cost 15285581

Item	PartNumber	+Stk	Ess	SL Req	GBVREI	GBV $	MTBUR $	MTBR	TCH	TAT	adj TAT	REMS	adj DEMS	value $	SAF = TAT/365	demand	Qty>SL	SL	fills
2	071-01933-350	2	2	0.93	0	0	5000	3751	37510	38	38	10	10	2500	0.104	1.04	3	0.98	9.782075
3	10-613129	34	1	0.95	2244	3366	3202	3589	272800	28	28	76	76	3365	0.077	5.83	10	0.96	73.27135
4	10-617950-1	21	1	0.95	1429	2144	2700	21864	371680	28	28	17	17	2144	0.077	1.30	3	0.96	16.28345
5	10470-6	6	2	0.93	1340	2010	64975	19406	135400	28	28	7	7	27486	0.077	0.54	2	0.98	6.870486
6	107484-5	1	1	0.95	1507	2261	6060	8974	170600	28	28	19	19	3025	0.077	1.46	4	0.96	18.68305
7	107492-2	4	1	0.95	3413	5120	7291	5602	156860	28	28	28	28	5100	0.077	2.15	5	0.98	27.37024
8	108000-8	16	2	0.93	277	416	1695	7198	64790	28	28	9	9	918	0.077	0.69	2	0.97	8.703006
9	109486-6-1	17	1	0.95	2423	3635	14180	1705	105710	28	28	62	62	3000	0.077	4.76	9	0.98	60.52303
10	114029	6	3	0.9	508	762	3000	28741	201190	38	38	7	7	1500	0.104	0.73	2	0.96	[illegible]
11	121117S-011	2	2	0.93	0	0	9000	40920	81840	38	38	2	2	4500	0.104	0.21	1	0.98	1.962216

Figure 5.3: Poisson solution output values

Figure 5.3 shows, on the right-hand side, the solution output values derived from the Poisson distribution calculations. The quantity is selected as the number of items for which the required SL is exceeded – in the first line, a quantity of 3 exceeds a SL of 93%. This quantity gives a SL value of 98%. Multiplying the number of removals (demand events), denoted REMS, by the achieved SL gives a fill rate of 9.78 for the first item. Thus, if there are 10 requests for this item over the planning period (a year) then an average of 9.78 will be satisfied. Finally, the extended cost is given for the calculated quantity.

The global SL value (96%) and cost are shown at the top of the spreadsheet, the SL being the total number of fills divided by the total number of removals.

5.6 Model 2: Marginal Analysis

This method treats the inventory pool as a whole and looks for the best economic value in allocating spares. Like the cost-wise skewed holding heuristic that follows, Marginal Analysis aims to over-provide cheap parts, under-provide expensive parts and meet the overall target SL for the combined pool of stock, rather than seeking to meet an objective SL for each respective part. Logistechs Inc developed the Marginal Analysis solution commercially and provided a consultancy service to rotable owners until being bought by General Electric. As a major engine OEM, GE's strategy is increasingly an inclusive service-based model, where an engine packaged is sold with maintenance. Many airlines now choose the Power By the Hour (PBH) option with engine providers, where a set amount (e.g., USD300) is paid for each engine hour flown to cover ownership and maintenance of the engine. In this way the engine provider aims to generate continuous revenue from their assets. GE is now a major user of Marginal Analysis for its own rotable stock. Meanwhile, GE claims to have helped airlines reduce rotable stock levels without harming SL by 20 to 40% using Marginal Analysis (GE 2002).

The Marginal Analysis approach works as follows:

1. Get utilisation (Total Component Hours), cost and reliability (Mean Time Between Removals) data for each rotable part number and calculate respective expected number of removals (REMS)
2. Factor unavailability due to repair time – Un-Availability Factor:

$$UAF = \text{Turn Around Time} / 365$$

3. Compute mean stock level required:

$$x = REMS \times UAF$$

4. Calculate cumulative probability values using a Poisson distribution to give Expected values for each item of stock

Up to this point, the process is the same as the basic Poisson process in 3.3.1 above. At this point, the basic process chooses each stock level to exceed target SL. Where the present method differs is in viewing the performance of the entire stock pool.

5. Calculate the number of fills, or successful demand events, for each quantity of each part, called fill rate.
6. Divide the incremental fill rate by the cost of part i:

$$\Delta fill / cost_i$$

this is the marginal contribution of each quantity of each part, i.e., the contribution to global SL made by each count of each part

7. Having calculated $\Delta fill / cost_i$ for each number of each stock item, sort these values in descending order
8. Allocate items of inventory in descending order of $\Delta fill / cost_i$. Each time an item is allocated, calculate the number of fills (successful requests for stock):

$$fills = \sum fill$$

9. The target SL is attained when:

$$\sum fills = \sum REMS \times SL$$

The disadvantages of Marginal Analysis are:

(i) it is computationally intensive and complex compared to the Poisson model;

(ii) it is not as theoretically sound as the Linear Programming approach below: it gives good results but is not based on a fully optimal solution;

(iii) it doesn't permit mixed SLs: in practice, rotables are held with 3 levels of essentiality (no-go, go-if, go) with typical SLs of 95, 93 and 90%;

(iv) it is possible that the model will recommend a zero holding where an item is expensive and has a low rate of demand: while this may work logically, it may not be acceptable in practice not to hold a spare. The model is modified here to apply a policy of having a minimal holding of one – this adds complexity to the solution algorithm;

(v) the method may pick items in the wrong order of quantity – for example, it may advise holding 3 of a given item without the cumulative quantities (1 and 2) as well. This is discussed further later.

However, Marginal Analysis can expected to give very good results, which will be close to optimal values achieved by the Linear Programming solution.

A further possibility is to run two versions of the Marginal Analysis model:

(i) treat all items together with one target SL, regardless of essentiality code. This risks over-stocking lower essentiality code items but should improve overall fill rates and therefore global SL, however the part costs can be weighted by SL to give a bias to higher-SL items;

(ii) run separate models for each essentiality code: this will give a more precise solution but will be less efficient as each pool is smaller.

It is proposed to assess the first option here for the purpose of comparison. If the results are close to the best solution among the models tested, then a further set of split problems (as (ii) above) can be tested. As a heuristic, the following will be evaluated: scale down the marginal contribution (Δfill / $cost_i$) value for each item by the SL corresponding to its essentiality code.

This method (and the other methods that follow) departs from the conventional (Poisson) approach described above in that it considers the performance of the entire inventory set, rather than treating each item in isolation. The objective then changes from "satisfy x% of requests for part Y" to "satisfy x% of requests for all parts". This presents the opportunity then to bias the inventory pool for cost: if the aim is to meet 950 of 1,000 requests for spares over a planned period, then it is cheaper to have the 50 failed requests to be for the most expensive items. In other words, some portion of requests for inventory will fail, so it is desirable to meet the performance requirement by filling requests for cheaper items.

Having determined that it makes sense to fill requests with the cheaper part, the question then becomes: "for each incremental expenditure on stock, which item gives the greatest likelihood of meeting an inventory request, for the least amount of money?"

From the probability distribution obtained in the Poisson process above, it can be seen that for each increment in the quantity of an item, there is an attendant increase in SL or ΔSL. For example, referring to Table 5 above, a part with a mean demand rate of 3 and a holding of 4 has a SL of 82%. Increasing the holding by 1 to 5 changes the SL to 92%, giving a ΔSL of 10% for this change in holding.

If two parts have the same probability distribution with the values in the examples above, then given the opportunity to increase the holding of one or the other from 4 to 5 will give the same ΔSL. However, if one part has a higher removal rate than the other, it makes sense to pick the part that will be called for the most often, over the part with the lower removal rate. Thus fill rate and change in fill rate are used:

Δfill(j) = ΔSL x REMS = (SL(j) – SL(j-1)) x REMS for part i

or

Δfill(j) = fill(j) – fill(j-1) for part i

Further, cost is a factor: for two parts with the same Δfill, it stands that choosing the part with the lower cost will give better value.

Thus for each Δfill for each inventory item, it is appropriate to divide ΔSL by the cost of the item to see which parts give the best value. The incremental increase in SL divided by cost is called the marginal contribution of a stock item:

Marginal contribution, MC = $\Delta fill_{ij}$ / $cost_i$,

where i is a part in the inventory list

and j is a quantity of part i

The marginal contribution needs to be computed for each quantity of each item. Having calculated MC for each quantity of each part, these values can then be sorted in descending order. The inventory list is then filled in order of decreasing MC. The number of fills, or successful requests for inventory, is computed and summed at each step. When the number of fills divided by the total number of removals, REMS, exceeds the target SL, the overall holding is sufficient.

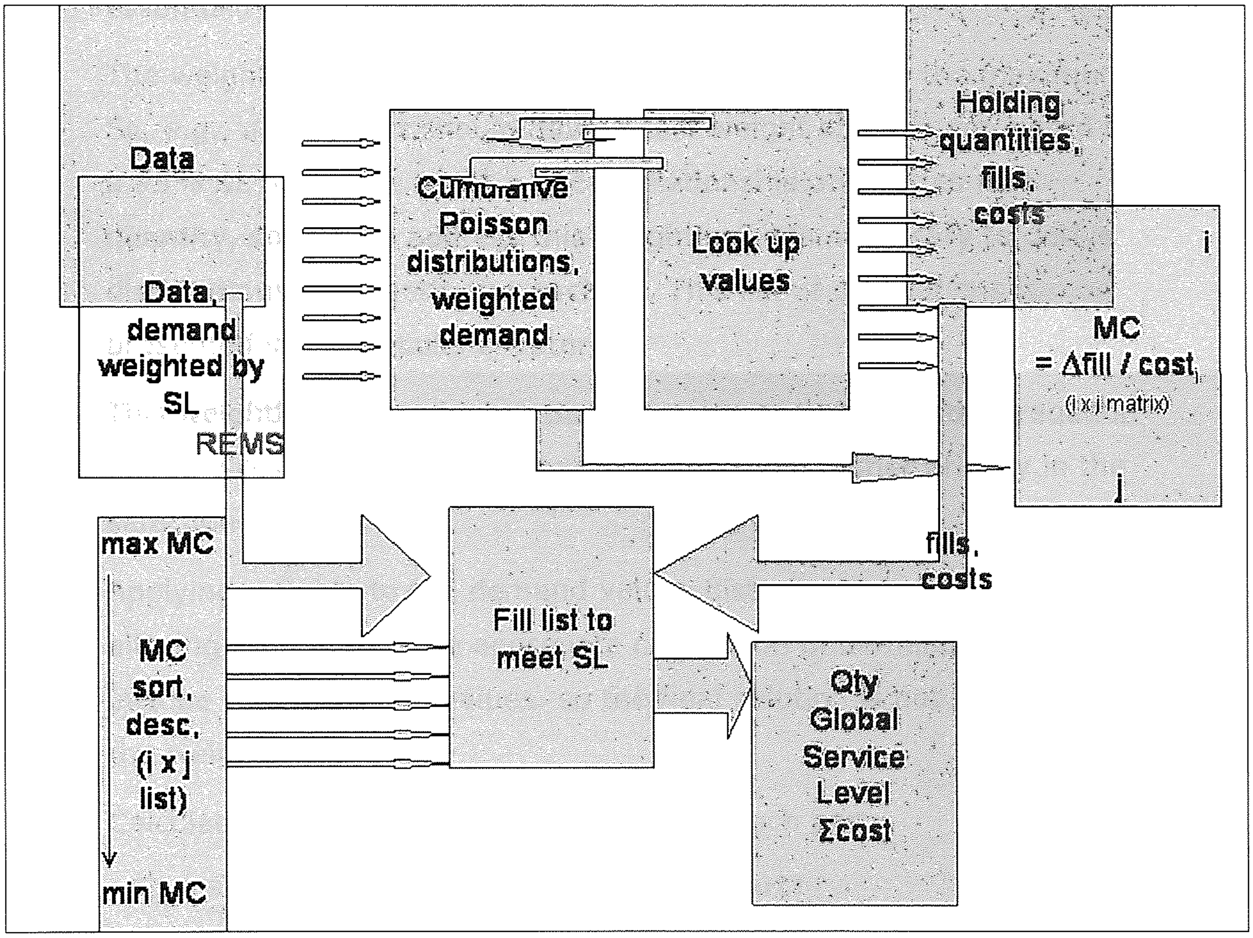

Figure 5.4: Marginal Analysis solution implementation schematic

In Figure 5.4, the Data box is the database extract of 300 items used as inputs to the model; the cost values for each item are weighted by their SL. Given three SL values representing three essentiality levels, it makes sense to divide the cost by SL so that the items with higher SL will be ranked proportionately higher in a solution – in other words, a part with a higher essentiality will be chosen over a part of similar value but lower essentiality. This weighting is to be used for ranking only; care must be taken to use the original cost values in the final total cost calculation.

The Marginal Analysis solution can be specified as follows:

1. Poisson distributions are calculated for each part as in the Poisson method above, and the holding quantities and corresponding SLs extracted.

2. In order to weight SL outcomes by essentiality code, demand rates are weighted according to the mean number of removals and the essentiality code of the item.

   The weighting is a function of the target SL relative to the maximum SL, e.g., a 90% part should have a cost weight above 1 compared to a 95% SL part. This is not a linear relationship and varies with quantity, so to fully address this weighting, it is necessary to derive distributions by quantity for each SL. This set of calculations is presented in the Results chapter.

   This weighting makes higher SL items appear to have more frequent demand relative to lower SL parts, giving them higher priority in the selection process.

   Applying weights to the demand values distorts the solution without altering other data. The aggregate calculation of demand does not include the SL scaling values, so the final solution is not distorted by the weights.

3. Calculate the fill rate for each quantity of each part:

   fill = SL x REMS for each quantity j of each part i

   where

   fill = number of satisfied requests

   SL = SL for QTY j

   REMS = number of removals of part i

4. Calculate the incremental fill rate:

$\Delta fill_j$ = fill(j) – fill (j-1) for each quantity j of each part i

5. Calculate Marginal Contribution, the increase in fill rate scaled for cost:

MC = $\Delta fill_j$ / $wcost_i$ for each quantity j of each part i

For the higher SL items weighted in step 2, the MC will be higher relative to a lower SL item.

There is now an (i x j) matrix of MC values.

6. This matrix is then split into a list of length i x j and sorted in decreasing order of MC. Each MC value represents one part.

7. Parts are picked starting from the top of the list and moving down. For each part picked, the number of fills can be computed as:

Fills = Δfill x $REMS_i$, for each quantity j of each part i

8. A question arises here with respected to global target SL, i.e., the desired performance of the whole inventory pool. Since there is a mixture of SL values with different rates of removal and different values, what is the appropriate target SL? An approximation can be made by multiplying all removals REMS by corresponding SL figures and dividing by the total number of REMS:

target SL = Σ ($REMS_i$ x SL_i) / Σ REMS

9. A total of all fills is kept and the routine is complete when:

SL = Σ(REMS x Δfill) $\geq$ target SL

11. The selection can then be re-sorted to show the sequence of parts and their quantities.

12. The total cost is computed as the sum of individual costs x quantities. It is important to use $cost_i$, not $wcost_i$, which is only used to give a weighting to the MC evaluation so that higher SL parts will be picked proportionately more than lower SL parts.

An example of Marginal Analysis is shown for two parts selected from the data set, shown in Table 5.7. The two parts are chosen with the same SL so that the weighted cost step is not needed. The parts are chosen with significantly different costs and removal rates. For ease of viewing, calculations are shown for part quantities up to a value of 9 in Table 5.8, which is sufficient for the demand rates shown. In the full-scale model a range of 15 is used, with a small number of high-demand parts being processed separately with quantities up to 30 – these are selected at the Poisson demand distribution stage and typically number less than 5 for set of 300 parts.

index	*Part no*	*descr*	*REMS*	*cost*	*TAT*	*Mean demand*
29	172625-7	VALVE, ASSY ANTI-ICE	33	12072	28	2.5
4	10-617980-1	EXCITER, IGNITION	17	1429	28	1.3
			Total 50			

Table 5.7: 2-part sample for Marginal Analysis example

QTY	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*
Part 1 SL	0.28	0.54	0.75	0.89	0.96	0.98	1.00	1.00	1.00
Part 2 SL	0.63	0.86	0.96	0.99	0.99	0.99	0.99	0.99	1.00
					Multiply REMS by SL				
Part 1 fills	9.27	17.68	24.78	29.27	31.54	32.50	32.85	32.96	32.99
Part 2 fills	10.63	14.55	16.26	16.82	16.96	16.99	17.00	17.00	17.00
					Subtract fills (j) from fills (j-1)				
Δfill 1	9.27	8.41	7.10	4.49	2.27	0.96	0.35	0.11	0.03
Δfill 2	10.63	3.92	1.71	0.56	0.15	0.03	0.01	0.00	0.00
					Divide Δfill by cost = Marginal Contribution				
MC1 = Δfill 1 / cost 1	0.000767848	0.000697	0.000588	0.000372	0.000188378	7.948E-05	2.87434E-05	9.1E-06	2.56E-06
MC2 = Δfill 2 / cost 2	0.007439645	0.002746	0.001194	0.000389	0.000101494	2.206E-05	4.10977E-06	6.7E-07	9.71E-08

Table 5.8: Marginal Contribution calculations for two parts

All of the part number and quantity pairs are then listed, together with Marginal Contribution and Δfill values. This list is then sorted in descending order by Marginal Contribution in Table 5.9.

Selection order

Part no	HQ	MC	Δfill	fills	Total SL	Total cost
2	1	0.00743965	10.6	10.63	0.212	1429
2	2	0.00274566	3.9	14.55	0.291	2858
2	3	0.00119355	1.7	16.26	0.325	4287
1	1	0.00076785	9.3	25.53	0.510	16359
1	2	0.00069670	8.4	33.94	0.678	28431
1	3	0.00058790	7.1	41.04	0.820	40503
2	4	0.00038913	0.6	41.59	0.831	41932
1	4	0.00037207	4.5	46.09	0.921	54004
1	5	0.00018838	2.3	48.36	0.967	66076
2	5	0.00010149	0.1	48.50	0.970	67505
1	6	0.00007948	1.0	49.46	0.989	79577
1	7	0.00002874	0.3	49.81	0.996	91649
2	6	0.00002206	0.0	49.84	0.996	93078
1	8	0.00000910	0.1	49.95	0.999	105150
2	7	0.00000411	0.0	49.96	0.999	106579
1	9	0.00000256	0.0	49.99	0.999	118651
2	8	0.00000067	0.0	49.99	0.999	120080
2	9	0.00000010	0.0	49.99	0.999	121509

Table 5.9: part quantities sorted by Marginal Contribution

The result given by this solution is:

QTY(part 1) = 5

QTY(part 2) = 4

Total cost = 66,076

Global SL = 96.7%

The highlighted SL value in the Table 5.9 shows that the mean value of 48.36 fills, with a total demand of 50, gives a SL of 96.72%, which exceeds the target SL of 95%. Note that the previous quantity gives a SL of just 92.1%.

What emerges from this example is that the Marginal Analysis model, by choosing an excess of cheaper parts to increase SL, may then end up choosing a more expensive part to meet the SL, ending up with an excess of parts overall. In the example, the routine chooses more parts than the Poisson process, which, while achieving a higher SL, creates a more expensive solution. The Poisson process would choose 5 of part 1 and 3 of part 2, both with SL 96% and a combined cost of 64,647. The Marginal

Analysis solution chooses that, for an extra cost of 1,429, a SL of 96.7% is achieved. Ideally then, the solution should reduce the quantity of part 2 to 3 having chosen the 5th item of part 1 and revising its solution. This is an inherent weakness in the model, namely that by choosing in order of MC / cost, the solution may overshoot the target SL by over-providing the cheaper parts. It is expected that this effect will diminish with a large number of parts. Also, it is predicted that the model will give a total SL closer to the target than will the Poisson model, which must exceed target SL for every part.

Based on this simple example, the Marginal Analysis approach gives a more costly solution than Poisson, albeit with a higher SL.

It can be seen from the two-part example that this is a complex and intensive processing operation and is difficult to formulate and implement. To implement this model on a large scale (300 parts) using Excel, it is necessary to program Basic procedures to extract and sort the data. For 300 parts with a quantity of up to 15 each, the sort list will contain 4,500 records. It would be possible to reduce the size of the model by removing values where SL is high, however the model may continue to allocate inexpensive parts even if they have a high part-level SL as they continue to contribute to global SL. A better approach is to remove items with a low MC, however a cut-off needs to be established and MC needs to be calculated in the first place. It may be safer and simpler to calculate all values; the sort list will have a long tail, or balance of values not included in the selection.

There is a further risk with the Marginal Analysis solution: the above example shows that, for low numbers, the change in Poisson value decreases as the quantity rises. However, for a higher mean value, the change in probability (and thus the Marginal Contribution) may occur such that the values are not in decreasing order, as shown in Figure 5.5 below.

This method is implemented in a spreadsheet, using routines to generate the {line – quantity} pairs with corresponding MC and Δfill, as shown in Figure 5.5. Each case tested is developed as a new set of spreadsheets.

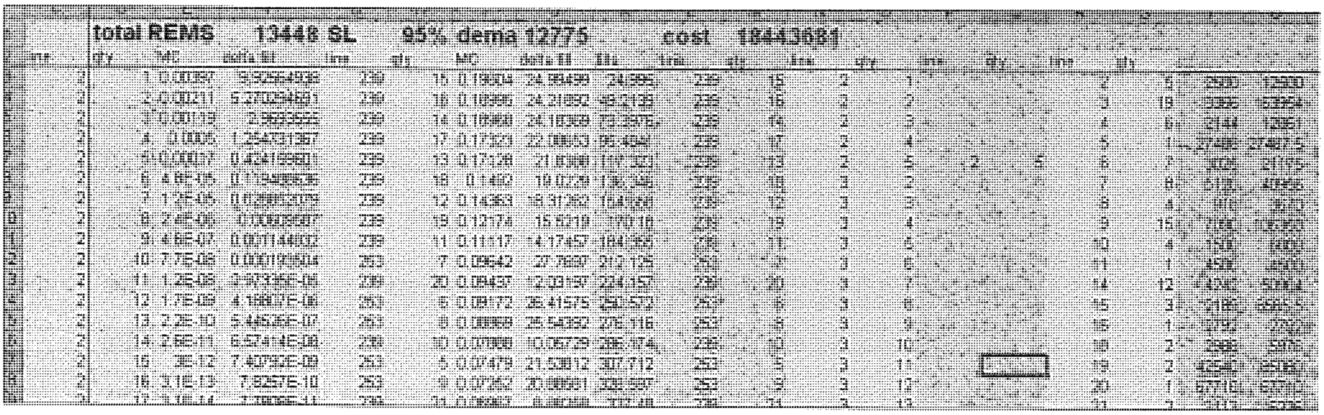

Figure 5.5: Marginal Analysis implementation in Excel

Figure 5.5 shows the sorting processes used to derive a solution. The first two columns on the left show the {line – quantity} pairs, from 1 to j for each item from 1 to i. Thus there are i x j pairs to be calculated and sorted. There are 8,220 pairs in the example shown (Case 5).

Listed with each pair are the respective MC value and Δfill, i.e., the relative value of allocating each part and the number of satisfied demand requests.

The first highlighted area shows the pairs sorted in decreasing order of MC, so that the part and quantity allocations with the greatest benefit are ranked highest.

The "fills" column accumulates Δfill values until the total exceeds the number of demand events to be met in order to meet the target global SL. In this case, there are 13,448 removals, so a 95% SL calls for 12,775 fills.

The second highlighted area shows the pairs sorted by line and quantity, with the adjacent area picking the maximum value of each quantity.

The highlighted area to the right shows each line item once, together with its calculated quantity.

Two disadvantages of this approach are apparent from the example shown:

1. Due to the shape of the Poisson distribution, the MC values do not occur in ascending order for a given item. Thus the method chooses 15 of part number 239 first, then 16, 14, 17 and so on. If the solution set does not include lower values (1, 2, 3,...) then the solution is not

logically consistent. In such a case, choosing the indicated values where smaller quantities are absent will lead to over-provisioning. This can be expected to be worse for high mean values. This is examined further in later chapters.

2. Where an item is expensive and has low demand, it may not be included at all in the solution set since it will have low MC for all values. A policy is assumed whereby all items must have a minimal stock level of 1. This will lead to over-provisioning.

5.7 Model 3: Cost-Wise Skewed Holding

Viewing the pool of spares inventory as a whole, the objective of the stock system can be stated as a Service Level: for all stock requests, at least x% must be satisfied, where x is the target Service Level. A simple solution is to increase the Service Level for the cheapest parts, so that they are more likely to satisfy a request. Thus it is more acceptable (and cheaper) to have a stock demand failure (a "miss") for an expensive item than for a cheap one. This method requires trial and error with SL variations to give a satisfactory solution and is a quicker approximation of the Marginal Analysis approach. To implement this approach, a band of SL values is applied for a range of inventory value bands. For example, if the global target SL is 95%, then the inventory pool may be skewed as shown in Table 5.10 below.

Band	*Value band*	*SL*
1	Lowest 20% by value	98%
2	20 - 40% rank	96%
3	40 - 60% rank	95%
4	60 - 80% rank	90%
5	80 - 100% rank	80%

Table 5.10: sample Target Service Levels for cost-wise skewed holding

By calculating line-by-line stock levels required to meet target SL in each band, the overall SL can then be calculated simply as:

SL = total satisfied requests / total removals for the period

The parts could be divided either by arranging bands in size of equal combined value, or bands with equal numbers of parts. The problem with partitioning by combined value is that the quantities of each part are unknown, so the value of a single part is not meaningful. Thus the partitioning is carried out by allocating an equal number of parts to each band.

The disadvantages of this approach are:

(i) the precision of the approach is limited by the division of the stock into an arbitrarily convenient number of bands;

(ii) dividing the parts into groups of equal numbers means that the final combined value of each band can be expected to be very different, which may give an unequal treatment where cost is the measured outcome

(ii) as a rough partitioning of the problem, this approach can be expected to give results that are some way off the optimum derived by the theoretically complete Linear Programming approach below;

(iii) the range of SL values to be used is arbitrary and needs to be varied by trial and error to give good performance for minimal cost.

The advantage of this approach is its simplicity. As this method is a new solution, there is no experience available to suggest how well it will work.

A refinement of this approach would be to force parts with lower essentiality codes into the lower value bands to give better reliability. This can be approximated by scaling the value of each part by its target SL for the purpose of sorting.

This method is proposed as a convenient heuristic and should be less complex to implement and compute than the Marginal Analysis solution. This approach aims to simply group items by cost into a small number of categories, or bands, to which different target SLs can be applied. The individual holding quantities are taken from the Poisson method above, requiring limited further formulation and calculation. It is proposed to divide the holding into five bands of ascending cost and descending SL. This allows case analysis to be performed by varying each of the target SL levels to observe the overall system-wide effect on performance (SL attained) and total cost. Since this approach divides the inventory group into 5 bands rather than 300 individual items, it can be expected to give a less efficient result than Marginal Analysis above, or Linear Programming below. However, the model allows wide variation of SL for the bands and could offer potential as a useful decision support tool for fleet planning.

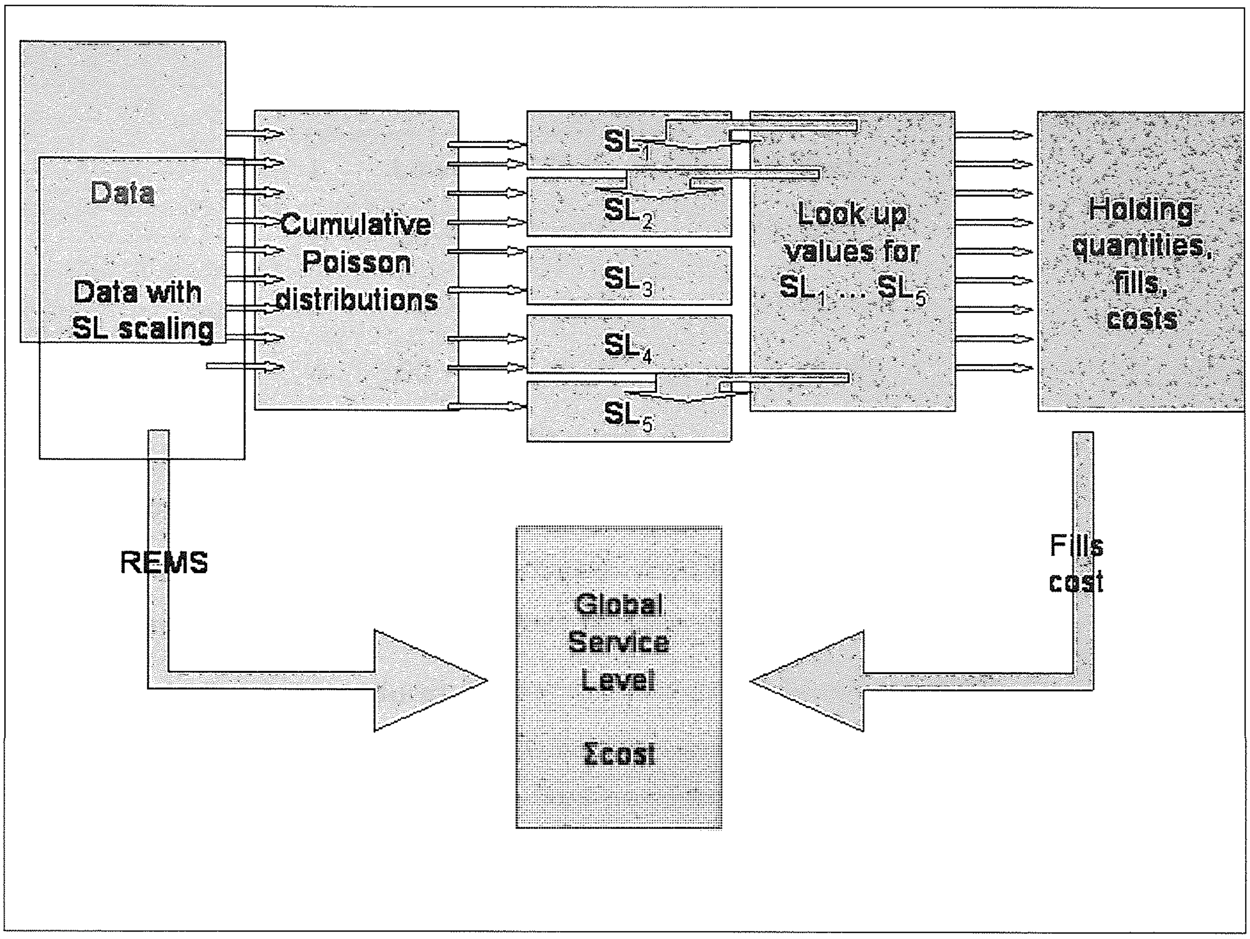

Figure 5.6: Cost-Wise Skewed Holding solution implementation schematic

The cost-wise skewed holding solution is illustrated in Figure 5.6 and specified below.

1. Compute demand event probabilities by Poisson distribution for each quantity j of each part i.
2. Weight demand for each part according to essentiality.
3. Sort the data by part value
4. Split the list into 5 partitions (or another arbitrary small number).

   The list can be divided such that each partition contains one-fifth of the number of parts, or one-fifth of the parts by value. However, since the total value of parts is not known in advance, it is proposed to simply divide the list numerically in the first instance.

5. Apply different target SLs to each partition, for instance the values shown in Table 5.11 below.

Band	*Value band*	*SL*
1	Lowest 20% by value	highest
2	20 – 40% rank	≤ SL(1)
3	40 – 60% rank	≤ SL(2)
4	60 – 80% rank	≤ SL(3)
5	80 – 100% rank	≤ SL(4)

Table 5.11: Target Service Level value ranking for cost-wise skewed holding

Note that the SL values are chosen arbitrarily and can be varied for scenario analysis and the effect on cost and SL observed.

6. Use look-up procedures to find the individual SL values and corresponding quantities QTY for each part i such that the individual SLs exceed the target SL for that partition

7. Multiply SL by REMS for each part i and calculate global SL:

   global SL = Σ (SL(i) x REMS(i)) / Σ REMS

8. Multiply QTY by cost for each part I and calculate total cost

   cost = Σ (QTY (i) x cost(i))

Note that this solution does not seek a global SL, rather it calculates holdings for the range of SLs provided and gives the consequential global SL value. The component SLs are varied until the global SL is observed to meet the target SL required.

Implementation of this solution in Excel is straightforward, although there are now two SLs to be considered for each part: the input target SL and the variable SL used as a solution parameter.

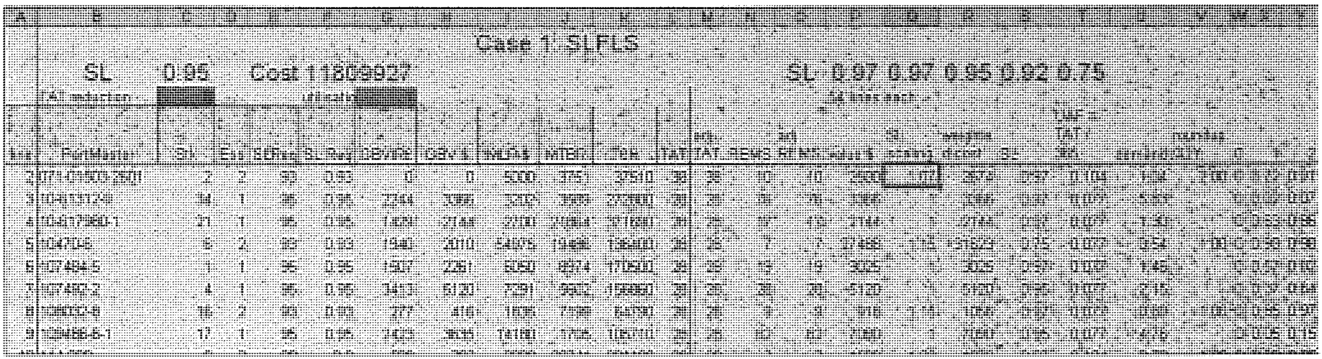

Figure 5.7: Cost-Wise Skewed Holding implementation in Excel

Figure 5.7 shows the solution layout for this method: the SL and cost values at the top left are outputs. The five SL values at the top right are global variables. With the data set divided into 5 categories sorted by cost, each category is linked to a SL value. The 5 values are then altered manually with the aim of meeting target SL at minimum cost. The 5 SL values are shown in order of increasing part value, so that the left-most value (0.97) is attached to the least expensive one-fifth division of the set. The right-most SL value (0.75) is attached to the most expensive division.

SL scaling is applied to reflect lower essentiality of some items. Where a part is of essentiality level 1, the SL scaling factor is 1. For other essentiality levels, a SL scaling value is looked up in an array where values are calculated depending on mean and SL. Thus for the first item, with SL required (column F in Figure 16) of 93% and a rounded-up demand of 2 (column V) a SL scaling value of 1.07 is applied. The cost of the item is then scaled up by this amount to make it appear proportionally more expensive and will thus appear in a different order in the cost-wise sort.

With the SL values set at 5 levels, the resulting quantities are selected as for the Poisson method, with fills and extended costs contributing to the global SL and cost calculations.

Two obvious shortcomings with this method are the choice of number of partitions to be used, and need for trial-and-error selection of SL values. The ideal is a solution that would use a partition for each individual part, and then assess all possible part quantities and resultant SL values. This is the aim of the final approach below.

5.8 Model 4: Linear Programming – combined model

The final solution technique proposed and tested in this study uses Linear Programming to select a solution from among the large number of stock permutations. The solution is optimised for cost and has as a constraint that it must meet the global target SL. The objective of this solution can be stated as: "which is the cheapest combination of inventory holdings to meet the SL criterion?" Given that LP should look at all possible combinations and choose the best, this technique is the only one among the four solutions presented here that can be considered theoretically optimal. Two versions of the LP method are used, to address different essentiality codes in two different ways. The present model includes all parts in a single formulation; the subsequent model separates parts into three formulations according to essentiality code.

This is a solution for full system optimisation, optimised for cost or service level. This approach can be outlined as follows:

1. for each rotable stock item, perform the calculations used in the methods described above
2. if cumulative probability values are calculated up to some arbitrary global number, say 30, then there will be a n x 30 matrix of stock values for n part numbers. The data set used here contains 300 line items, so there are 9,000 stock values.
3. choose the best stock combination to satisfy either:

minimise Σcost subject to ΣSL > target SL

or

maximise ΣSL subject to Σcost < budget

subject to a binary constraint:

for each line item, only one stock number can be selected

As with the other approaches above, the question arises of whether to separate parts by essentiality codes. Simply put, should the demand for essential parts be combined with the demand for non-essential or less-essential items? From this point of view, it makes sense to run separate solutions for each code as they have different SLs. However, when the overall performance of the inventory pool is considered, it makes sense to treat all items together as subject to aggregate demand. Further, the smaller groups of parts (if separated by essentiality) will reduce the effect of the solution, compared with treating the group of parts together to maximise SL while minimising cost. A possible solution for the different SLs attached to the essentiality codes is to weight part cost in inverse proportion to SL. For example, if SL values for essentiality codes 1, 2 and 3 are 95, 93 and 90% respectively, then dividing the part cost by its associated SL will give a higher weight to an item with a lower SL, making it less attractive for optimisation. Clearly, the final total cost calculation needs to use original cost data, not weighted values. SL scaling is evaluated in the next chapter.

The alternative to SL scaling is to split the problem to avoid mixing parts with different essentialities - this is outlined in the next section, the LP - split model.

For this problem, the LP formulation is a special case as it is presented as a binary or {0, 1} Integer Linear Programming problem. Referring to Table 5.6 above (basic Poisson value calculations), with the sample set of four items with mean demand values of 2, 3, 4 and 5, the solution may only select one value from each row. This logical constraint means that, for a vector of variables of quantities for each row, exactly one variable will be assigned a value of 1 and the rest will be assigned a value of 0. This is equivalent to the highlighted cells in Table 5 showing the selected quantity to meet the SL target. However the difference with the LP solution is that it is seeking the target SL at the global level (for all 300 parts together) rather than at the part level. However, the part-level SL values are used as inputs to the formulation.

Like the Marginal Analysis solution, this approach includes all inventory items in a single calculation space. Since these items have different SL targets, it is appropriate to weight the demand of lower-essentiality items to make them commensurately less attractive to the solution. This weighting is performed in the same way as for Marginal Analysis and uses distributions of values to alter different mean values for different SLs.

In order to process the LP solution, it is necessary to first create a large-scale formulation of the form:

Minimise Σ ((cost(i) x QTY(i)) x X_{ij}) ***Minimise cost***

Subject to Σ(REMS(i) x SL(ij) x X_{ij}) ≥ targetSL x ΣREMS ***Reach target SL***

and Σ X_{ij} = 1 for all values of j for each part i

Choose exactly one quantity for each part

and $X_{ij} \in \{0,1\}$ ***Binary variable (can't have fractional quantities)***

For a 300-part data set, i is given values from 1 to 300.

The typical maximum part QTY value is 15, so j has values from 1 to 15, although this may need to be increased if individual part SLs do not comfortably exceed the maximum SL for the problem set. Therefore, as with the Marginal Analysis approach, there will be a small number of parts with higher QTY than 15 required to give individual SL exceeding target SL, so these parts will be given larger quantities, say 30, to satisfy target SL.

For convenience in labelling the X_{ij} variable, j is assigned a letter value. This makes it easier to create a Linear Programming formulation in a simple notation. Values of j are assigned letters in ascending order:

QTY(j){1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15}

→ QTY(j){a, b, c, d, e, f, g, h, i, j, k, l, m, n, o}

The formulation is illustrated with an example as shown below, using data used in the Marginal Analysis example earlier.

Assume there are two parts with quantities from 1 to 5 each:

i ∈ {1, 2} – part numbers are 1 and 2

j ∈ {1, 2, 3, 4, 5} → {a, b, c, d, e} – 5 part quantities labelled a to e

Assume further that costs for parts 1 and 2 are 12072 and 1429 units respectively:

cost(i) = {12072, 1429}

And target SL values are 95% for each part and for the whole solution:

SL(i) = {0.95, 0.95}; target SL = 0.95

Removals for the two parts are:

REMS(i) = {33, 17}

The REMS values give corresponding SL values for each value of i and j (using TAT, UAF and a cumulative Poisson distribution):

SL(1j) = {0.28, 0.54, 0.75, 0.89, 0.96}

SL(2j) = {0.63, 0.86, 0.96, 0.99, 0.99}

Finally, create 0, 1 binary variable X, where X can only have a non-zero value once for each part:

X_{ij} = 1 for j = 1 to 5 and for each value of i

The formulation is:

Minimise

12072 X1a + 24144 X1b + 36216 X1c + 48288 X1d + 60360 X1e + 1429 X2a

+ 2858 X2b + 4287 X2c + 5716 X2d + 7145 X2e

Subject to

0.28 x 33 X1a + 0.54 x 33 X1b + 0.75 x 33 X1c + 0.89 x 33 X1d + 0.96 33 X1e

+ 0.63 x 17 X2a + 0.86 x 17 X2b + 0.96 x 17 X2c + 0.99 x 17 X2d + 0.99 x 17 X2e

≥ 0.95 (33 + 17) = 47.5

and Xij = 0, 1

and X1a + X1b + X1c +X1d + X1e = 1

and X2a + X2b + X2c + X2d + X2e = 1

The problem is shown formulated in a Linear Programming application LPSolve, in Figure 5.8 below.

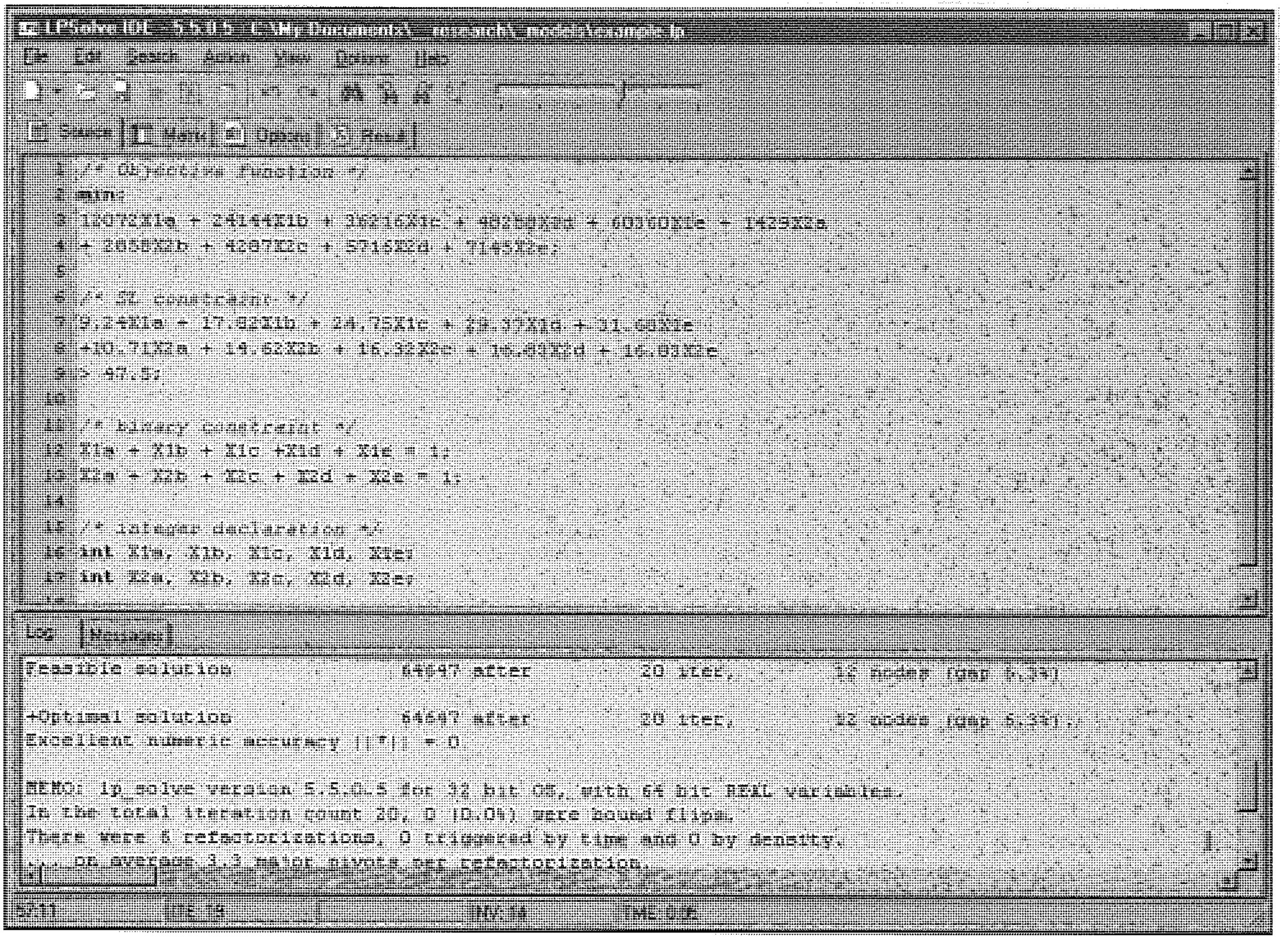

Figure 5.8: LPSolve formulation of sample problem, with i=2 and j=5

The solution set is shown in Figure 5.9 below.

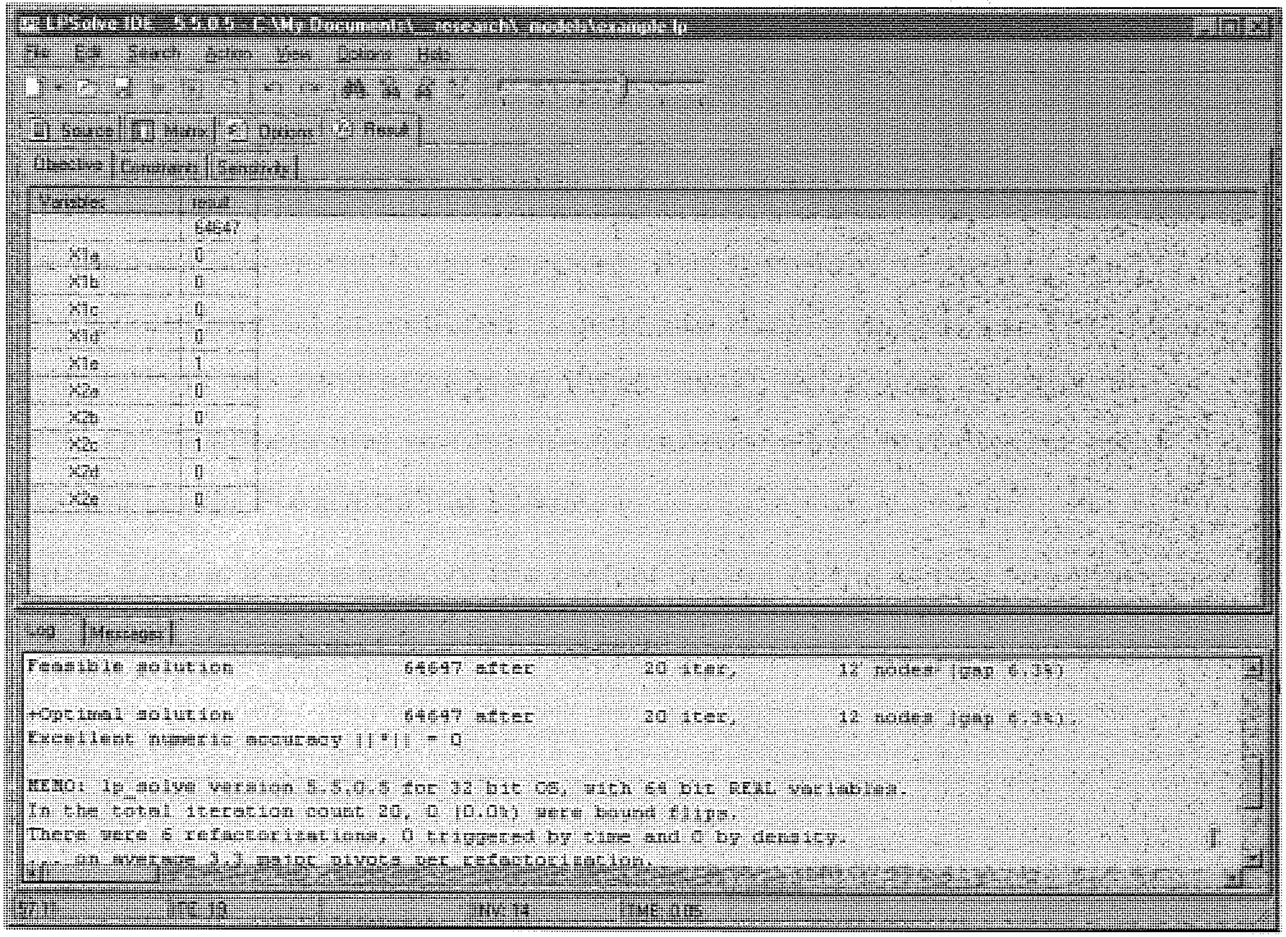

Figure 5.9: LPSolve solution of sample problem, with i=2 and j=5

Thus the solution provided is:

X1e = 1, or X1 has QTY = 5

X2c = 1, or X2 has QTY = 3

SL constraint:

REMS = 0.96 x 33 + 0.96 x 17 = 31.68 + 16.32 = 48

SL achieved is:

SL = REMS / ΣREMS = 48 / (33 + 17) = 0.96

Total cost of the solution (the value of the objective function):

Total cost = 60360 + 4287 = 64647

The above example shows the formulation for i = 2 and j = 5, where there are $j^i = 5^2 = 25$ possible outcomes. The formulation for this small problem is quite long: there are 5 x 2 terms in the objective function, 5 x 2 terms in the Service Level constraint and 5 terms in each of 2 {0, 1} constraints, or a total of 3 x ij terms. Thus for a data set with 300 parts and a quantity range of 15 for each, there will be 3 x 300 x 15 = 13,500 terms to be formulated. Given that there are 5 cases to be tested, there are then 67,500 terms to be compiled. It is therefore necessary to develop automated procedures to generate these expressions. The model formulations are generated using Visual Basic, extracting variables from Excel spreadsheets and writing the formulations into text files for input into the Linear Programming solver.

Figure 5.10 is a schematic of the logical design of the Linear Programming model.

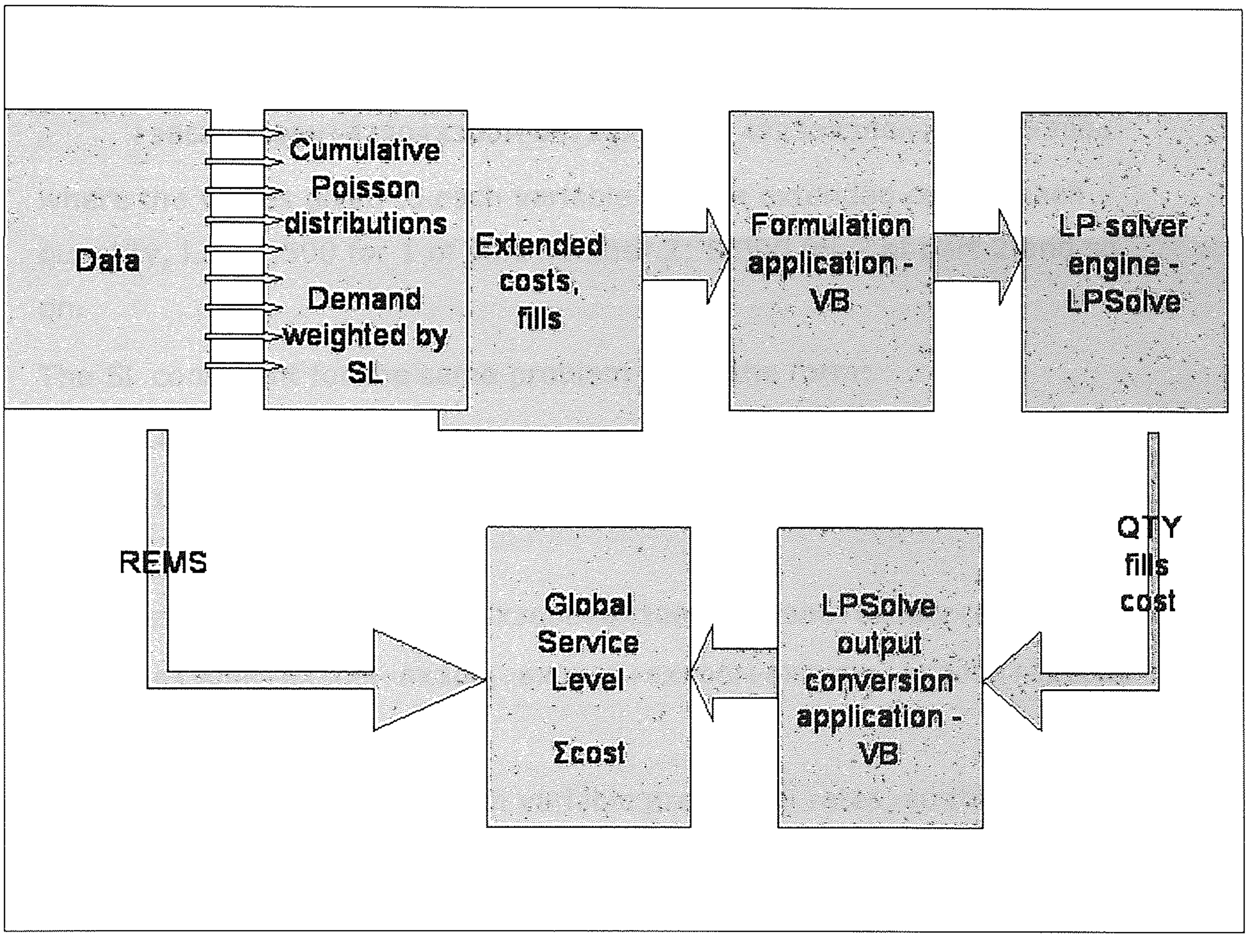

Figure 5.10: Linear Programming solution implementation schematic

The model formulation produces several arrays to give:

objective function – minimise cost;

SL constraint – meet target demand;

binary constraint – choose exactly one quantity value for each line;

integer declaration – the value of 1 in the binary constraint cannot be split.

For the set of parts tested, the first labelled X2, the last labelled X300 and each assigned a range of quantity values from 1 to 15 denoted by letters from "a" to "o", the objective function is of the form:

minimise

2500X2a+5000X2b+7500X2c+10000X2d+12500X2e+15000X2f+17500X2g

+20000X2h+22500X2i+25000X2j+27500X2k+30000X2l+32500X2m

+35000X2n+37500X2o

...

+3517X300a+7034X300b+10551X300c+14068X300d+17585X300e

+21102X300f+24619X300g+28136X300h+31653X300i+35170X300j

+38687X300k+42204X300l+45721X300m+49238X300n+52755X300o

where the values given to each variable are the extended cost of each quantity, i.e., 2,500 for 1 of part number 2; 5,000 for 2 of part 2 and so on.

The SL constraint for the same problem takes the form:

7.46X2a+9.25X2b+9.83X2c+9.97X2d+9.99X2e+10X2f+10X2g+10X2h

+10X2i+10X2j+10X2k+10X2l+10X2m+10X2n+10X2o

...

+5.8X300a+5.98X300b+6X300c+6X300d+6X300e+6X300f+6X300g

+6X300h+6X300i+6X300j+6X300k+6X300l+6X300m+6X300n+6X300o

> 6387

where the values given to each variable are the fill rates, so that quantity 1 of part number 2 will fill 7.46 demand events. The right-hand side (> 6387) is the requirement that the sum of all fills exceed 95% of the total number of removals.

The binary constraints are as follows:

X2a+X2b+X2c+X2d+X2e+X2f+X2g+X2h+X2i+X2j+X2k+X2l+X2m+X2n +X2o=1;

...

X300a+X300b+X300c+X300d+X300e+X300f+X300g+X300h+X300i+X300j+X300k+X300l+X300m+X300n+X300o=1;

namely, each part number can only have one of the choice of quantity values from 1 to 15, represented by "a" to "o".

Finally, so that the solution cannot split the value of 1 in the binary constraint, all decision variables are declared integers:

int X2a,X2b,X2c,X2d,X2e,X2f,X2g,X2h,X2i,X2j,X2k,X2l,X2m,X2n,X2o;

...

int X300a,X300b,X300c,X300d,X300e,X300f,X300g,X300h,X300i,X300j, X300k,X300l,X300m,X300n,X300o;

The objective function and constraints are loaded into a Linear Programming solver application, LPSolve (as shown in Figure 5.8) and the application is run. The solution set is output as the range of all decision variables with a binary value 1 next to the selected variables (Figure 5.9).

The solution set is sorted to remove variables with zero values, then the variable labels are parsed to give the solution quantities, as shown in Figure 11. Since each part number has exactly one binary variable with a value of 1, there will be i of these variables and (i x j) - i zero values, where i is the number of parts and j is the maximum quantity of each part. Where i = 300 and j = 15, there are 300 binary variables with value 1 and 4200 with value 0.

The quantities for each part are derived from the binary variables and are then used to calculate the total number of fills (satisfied demand events), the global SL and total cost, shown in Figure 5.11.

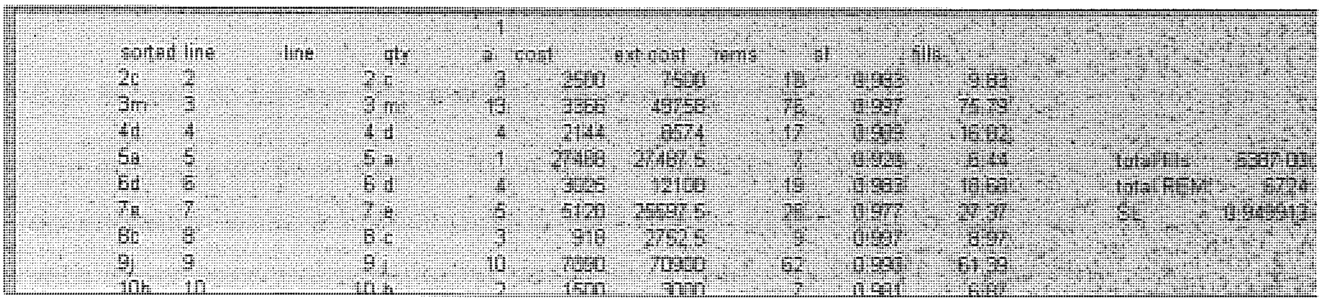

sorted	line	line	qty	1 a	cost	ext cost	rems	sl	fills		
2c	2		2 c	3	2500	7500	10	0.983	9.83		
3m	3		3 m	13	3366	43758	76	0.997	75.79		
4d	4		4 d	4	2144	8574	17	0.993	16.02		
5a	5		5 a	1	27468	27467.5	7	0.920	6.44	total fills	6387.03
6d	6		6 d	4	3025	12100	19	0.983	18.68	total REM	6724
7e	7		7 e	5	5120	25597.5	28	0.977	27.37	SL	0.949913
8c	8		8 c	3	918	2752.5	9	0.997	8.97		
9j	9		9 j	10	7090	70900	62	0.990	61.38		
10h	10		10 h	2	1500	3000	7	0.981	6.87		

Figure 5.11: LP – combined model solution in Excel

5.9 Model 5: Linear Programming – split model

The previous method applies SL scaling to process all parts in the same optimisation while prioritising parts with higher essentiality over those with lower essentiality. Another possibility is to formulate three separate models, each one containing parts of the same essentiality. This is a theoretically better solution in that there is no need to distort data to cater for different essentialities; however it creates three smaller problems (which may therefore benefit less from the scale of the solution) and requires three times the data preparation, processing and analysis.

This solution therefore comprises three separate models, using parts grouped by essentiality code, so there is one formulation each for essentiality codes 1, 2 and 3, each with a set SL and no SL scaling. Each model is input into the Linear Programming solver, solved and parsed separately. The results of the three solution runs can then be merged to give the optimised full inventory set.

The implementation of this method is the same as for the combined model, with the extra steps of splitting the data set into three groups, and without the SL scaling process.

This model is denoted LP3, to reflect that the solution comprises three LP formulations.

Figure 5.12 shows the problem space split into three separate formulations, which are solved individually and the results are then combined.

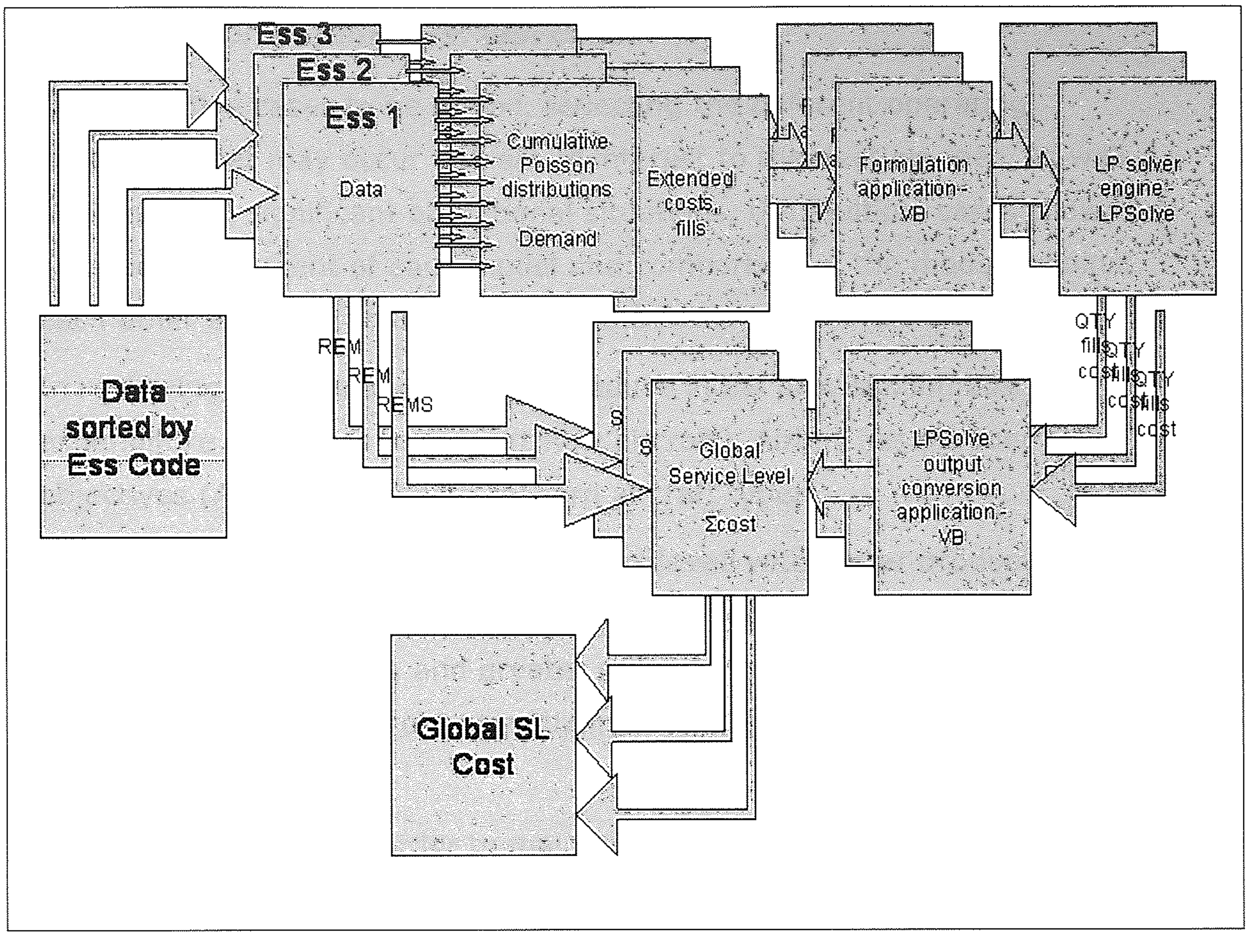

Figure 5.12: LP3 solution implementation schematic

This model produces a full set of outputs, which are processed using their exact target SLs, so this is the most logically correct solution and can be expected to yield both the best and the most reliable results.

This model is outlined in the following steps:

Split the data set into partitions: Ess 1, Ess2 and Ess 3

For each partition:

Minimise Σ ((cost(i) x QTY(i)) x X_{ij})

Subject to Σ(REMS(i) x SL(ij) x X_{ij}) ≥ targetSL x ΣREMS

and Σ X_{ij} = 1 for all values of j for each part i

and $X_{ij} \in \{0,1\}$

Join solutions for each partition for overall solution, cost and SL

5.10 Solution parameters: test cases

There are five cases to be tested, reflecting the different operating conditions that may occur and are of interest to the fleet planner. As well as comparing significant operational situations, the multiple cases allow detailed assessment of each model and comparison of the models against each other and with actual data from the field.

In reviewing practice and aiming to reduce initial provisioning cost (the expenditure on spares to support new aircraft), Airbus has pursued four objectives (Airbus 1998) with modern aircraft types:

(i) Improve reliability – this is being achieved with better product and system design, including greater use of electronic systems in place of mechanical devices, and greater redundancy of components, so that more items move from essentiality level 1 (no-go) to essentiality level 2 (go-if);

(ii) Freeze prices – in order to maintain customer loyalty, Airbus held price increase for several years, partly in response to customer complaints of overpricing;

(iii) Reduce shop processing time – improve handling processes and design better repair schemes;

(iv) Manage Service Levels – broaden the range of values used.

The first two items above are initiatives by the manufacturer and are outside the control of the operator or maintenance provider. The last two items, Shop Processing Time, or Turn Around Time and Service Level, can be addressed in operation.

Case 1: **base** case – normal conditions presented in the data set are applied. $SL_{FLS} = \{0.95, 0.93, 0.90\}$.

Case 2: **fewer**, reduced SL $SL_{Airbus} = \{0.95, 0.89, 0.75\}$. This reflects current thinking in the industry and a move to manage stock more actively. Thus items with essentiality code 1 ("no-go") are maintained at

95% as before, but lower-essentiality items are given substantially lower SL targets. With essentiality code 2 and 3 parts having lower target SLs, there will be a greater incidence of flights operating with failed items. However, there is scope for rotable management staff to react to demand by increasing inventory in the short term, leasing, buying or borrowing spares, so that the achieved SL will be far higher. The reduced SL approach also takes into account increasing redundancy being designed into the aircraft at a system level. For example, where cockpit instruments are normally duplicated, good design practice will increasingly include the same function represented as a secondary function on a related device. The use of integrated electronic systems provides greater coverage of essential functions throughout the aircraft, reducing the number of components that are essential to flight on their own.

Case 3: **faster**, reduced TAT = Case 1 TAT – 5 days. There is a general awareness that TAT is important among MROs, however the full impact of TAT overruns is not always appreciated. Airbus have set as a target reduced Shop Processing Times by designing rotables that are simpler to diagnose and repair, with greater modularity and increasing use of on-board electronics. Further, airlines and MROs can avail of process-improving technologies like RFID tagging and e-commerce to speed up routing decisions and logistics. A reduction of 5 days is chosen as achievable and significant. It can be predicted that inventory requirements should be scaled down in proportion to the lead-time reduction.

Case 4: **bigger,** double utilisation = double the number of removals (demand events, REMS). This is intended to demonstrate the effect of a larger fleet using one inventory pool: since demand is stochastic, it stands to reason that a larger fleet will not need a correspondingly larger spares stock. Since the test data set was provided, Aer Lingus, like many other airlines, has rationalised their fleet, with one narrow-body aircraft type for short haul and one wide-body type for long-haul. One of the factors in fleet rationalisation is improved spares usage. Another trend is for the outsourcing of component support, where an airline pays a service charge for spares cover at a specified SL. The provider of the spares can use the same stock to cover several customers, getting greater utilisation from their inventory.

Case 5: **best,** all combined = best-case case. This case assume lower target SLs, quicker repair cycle and larger fleet, which are the conditions to be aimed for by an efficient airline.

5.11 Solution run plan

The full set of model executions is shown in Table 5.12 below, to be presented with full output values in the Results Chapter.

	1	*2*	*3*	*4*	*5*
Case	*base*	*fewer*	*faster*	*bigger*	*best*
Poisson	P1	P2	P3	P4	P5
Marginal Analysis	M1	M2	M3	M4	M5
Cost-wise skewed	C1	C2	C3	C4	C5
LP	L1	L2	L3	L4	L5
LP3	L3-1	L3-2	L3-3	L3-4	L3-5

Table 5.12: solution run plan

Each case uses exactly the same data set with the same parameters, so the values for each instance in each column of Table 5.12 can be directly compared for the value of the solution output. Each solution uses the same routine to process each case of the data, so reading across each row in table 15 gives a comparison of outcomes for the different cases solved in the same manner. Thus it is only meaningful to compare methods for the same case (columns) and cases by the same method (rows) in the solution set.

5.12 Output variables – test results

The following results are required for each of run in Table 5.12 above. These variables are defined here as they must be designed into each solution.

The output values attained for each of these variables for each solution run are presented in the Results Chapter and interpreted in the Analysis chapter.

Service Level – whether the service level is fixed (as an input) or results from the calculation. For instance, in the case of the Poisson calculation, each part must exceed the target SL. This will result in a total SL > target SL due to granularity. Total SL = $\sum$fills / $\sum$removals.

Total cost – the sum of all quantities x inventory item costs.

Total inventory count – the sum of all inventory quantities $\sum$HQ required to fulfil the SL requirement. For cost-skewed solutions (all except Poisson) a higher count of cheaper parts can be expected.

Average item value – again, where the solution is skewed for cost, it can be expected that this figure will be lower. The best solution will have both the lowest average cost and the lowest total cost.

(Current stock-recommended holding) – the difference between the stock quantities actually held in operation and the stock quantities recommended by each solution.

(Total cost of current stock-total cost of recommended holding) – the difference between the combined value of actual stock derived from operational data and the combined value of the holding recommended by each solution.

Matching Metric 1 – the sum of the absolute value of the quantity of actual stock less the recommended holding quantity for each line item:

Σ|(actual holding - recommended holding quantity)|

This metric shows the match between the actual and the ideal stock levels, regardless of which is the larger number.

The matching metric can be normalised and expressed as a ratio by dividing the above by the total count of the actual holding:

$$\Sigma \left| \text{(actual holding - recommended holding quantity)} \right| \; / \; \Sigma \text{ (actual holding)}$$

The smaller this number, the closer is the actual holding to the prescribed level. This is a relative guide of scale between actual and recommended holdings.

Matching Metric 2 – each solution is compared against the best solution for comparison:

$$\Sigma \left| \text{(holding by method X - recommended holding by best method)} \right| \; / \; \Sigma \text{ (holding by best method)}$$

In the case that the actual holding is far greater than that recommended by any of the solutions, this metric will give a more precise comparison between methods. Also, as the best method has a value of 1, then the percentage value of each metric will give a true indication of the relationship between the solutions by each method.

Matching Metric 3– multiplying the holding difference in Matching Metric 1 above gives a basis for comparing the cost performance of the different solutions:

$$\left(\Sigma \left| \text{(actual holding - recommended holding quantity)} \right| \text{ x cost}\right) \; / \; \Sigma \text{ (actual holding x cost)}$$

Rather than simply comparing the total cost difference between methods, this method incorporates differences at the line item level and aggregates them.

Matching Metric 4 – as Matching Metric 2, comparing each method with the best method tested, but incorporating cost into holding quantity differences:

(Σ|(holding by method X - recommended holding by best method)| x cost)

/ Σ (holding by best method x cost)

This chapter has given detailed descriptions of a range of models to test current practice and new solutions. Each solution is described at a logical level and then specified for implementation as software with which to perform solutions using the test data set and the prescribed range of test cases.

The results of these model solutions are presented in the next chapter, Results.

Chapter 6: Results

Five solutions are built as models and run for five cases each, giving a solution set of twenty-five iterations, with defined output measures to facilitate interpretation of the behaviour and performance of each model in generating a solution. The first two models are based on known practice (Poisson and Marginal Analysis), while the other three are new models.

This chapter is structured around the models: the test cases are used to give a range of realistic perspectives from which to evaluate the models and it is the relative performance of the models that is of primary interest. The key results are shown in Table 6.6.

By running each model with the same data set and the same sensitivity parameters, it is possible to compare the performance of each model relative to the others. Testing different cases allows assessment of the impact of operational decisions. Finally, the results can be compared with actual operational data to see the degree to which the different approaches achieve the planning objective and how efficiently this is achieved in terms of cost.

These output variables, the test results, are compared and analysed in order to see which models give the 'best' results, how close the results are for the different models, and how the results compare with the actual operational state of the data set provided. Of particular interest is the set of theoretically optimal solutions provided by the Linear Programming solution compared with the others – it is expected that the Marginal Analysis approach will provide a close-to-optimal set of recommendations.

Before generating the results from the solutions, it is necessary to process some global data to be used in the solutions.

The first set of calculations is an assessment of the probability distributions to be used in forecasting demand data. There are distributions used as standard practice, which are reviewed and assessed for their suitability to the problem and with respect to the quantities (mean values) used as inputs, since ranges of mean values require changes in distribution according to industry practice.

Second, it is necessary to develop ratios for the scaling of Service Level (SL) values: in order to treat all parts in one solution, items with a lower SL need to be de-emphasised relative to items with higher SL. Given the stochastic nature of demand, it is unlikely that a linear relationship can be applied to scaling items with different SL values: the best approach is to generate a distribution of ratios, or different scale factors for respective demand quantities. Applying these ratios to the respective cost values for each item will weight the less essential items unfavourably so that they will be chosen with less frequency by a solution. Thus demand data and recommended inventory levels are preserved correctly. Taking the output from a solution, it will then be necessary to reverse the SL weighting applied to cost data so that the actual total cost values can be generated, i.e., the distortion introduced by scaling is eliminated.

Finally, the range of cases to be tested is defined. There are five cases (base, lower SL, faster repair, increased utilisation, all combined) based on practice and recommendations for improvement.

Running the five solutions (Poisson, Marginal Analysis, Cost-Wise Skewed Holding and two variants of Linear Programming) for each of the five cases gives a set of twenty-five solution runs and test results. These results can be compared against each other and against actual holding levels for the data set provided, enabling conclusions to be drawn about the relative performance of the models and weaknesses in any of the models. It is then possible to make recommendations with regard to best practice for inventory planning.

The solutions are subjected to sensitivity analysis in order to give insight into varying decision parameters, and also to provide a measure of quality assurance in the functionality of the models and the correct processing of the data. The sensitivity analysis is performed by varying the objective function and observing the consequent change in total cost. For each model, it is confirmed that varying the SL value (for essentiality code 1 parts) between 90 and 98% gives a sharply increasing total cost. The LP solutions are transposed, so that if, for example, a model with target SL of 95% gives a total cost of $10.3M, then formulating the problem with budgets of $10M and $11M as total cost objective functions will return SL values that confirm the first solution by interpolation.

6.1 Evaluation of probability distributions

All items are assumed to follow the same mode of failure and behave in an average way. While there may be differences in the probability distributions between electronic, mechanical, hydraulic, pneumatic, electro-mechanical and engine systems, this is not represented in the data. There is no account taken of component or fleet age.

Current practice in planning rotable inventory employs a cumulative Poisson discrete probability distribution. The Poisson distribution is used to predict the chance of an occurrence in a given time period where the mean approaches zero as the time period becomes shorter (Levine 2008). The time period should be chosen such that mean values are meaningful (the time should not be too short) and values are low, representing infrequent events (the time period should not be too short).

Conventionally, aircraft part failures are measured over a period of a year, with most item failure rates below 50. A histogram of demand rates from the data set is shown in table 1 below. Note that the first column label, 'below' is the upper limit of the range, so 'below 10' counts values between 5 and 10, 5 being the upper limit below the current value. Frequency is then the number of events in the respective band, and portion is the percentage of events in that band. 'Cumulative' is the cumulative frequency up to that band.

Below	Frequency	Portion	cumulative
5	113	38%	38%
10	55	18%	56%
15	19	6%	62%
20	20	7%	69%
30	20	7%	76%
40	18	6%	82%
50	16	5%	87%
60	12	4%	91%
70	4	1%	92%
80	5	2%	94%
90	4	1%	95%
100	4	1%	97%
200	8	3%	99%
300	2	1%	100%
More	0	0%	100%

Table 6.1: histogram of demand events for the sample data set

It can be seen from Table 6.1 that event frequencies are low, with the majority below 10 events per year, and 87% at 50 or fewer events.

Actual inventory holdings required will be far lower than the rate of failures, since items are repaired and returned to stock. Thus providing inventory at the same rate at which failures occur would assume no replacement of repaired items into stock. The stock level needed is in proportion to the time taken to repair and return a part to stock, Turn-Around Time (TAT). Thus if the TAT is one-tenth of a year, then the stock needed to support demand should be one-tenth of the number of events on average. For each item, the demand is scaled down by TAT / 365, referred to as Un-Availability Factor (UAF).

There are differences in opinion about how to deal with increasing quantities, since Poisson is intended for fairly small numbers of independent events – but what is a small number? Airbus recommends Poisson and normal distributions for calculations in its published Initial

Provisioning formulation (Airbus 1997), changing from Poisson to a normal distribution when the mean number of expected events exceeds 30. This crossover value is anecdotally reported as 50 according to Boeing advice. The normal distribution is a continuous distribution, usually referred to in the industry as the Gaussian distribution.

Both Poisson and normal distributions were calculated, graphed and subtracted using a range of mean values in steps of 10 from 10 to 100 and expected values in steps of 1 from 1 to 100. The normal distribution was calculated with the standard deviation set as the square root of the mean: as the mean increases, the Poisson distribution approaches this normal distribution.

Graphs of Poisson and normal mass function and cumulative distributions are shown in Appendix 2. The difference between the distributions is also graphed.

As expected, the greatest difference between Poisson and Gaussian (normal) distributions occurs when the mean value is smallest. However there is a significant error across all tested values, ranging from 8% at a mean value of 10 to 3% at a mean value of 100 as shown in Table 6.2 below. The greatest deviation occurs when the expected value equals the mean in each case. These differences are summarised in table 17 below.

Mean	(Normal) – (Poisson) distribution *Maximum difference*	*Minimum difference*	*Occurs at expected value*
10	0	-0.08	10
20	0	-0.06	20
30	0	-0.05	30
40	0	-0.04	40
50	0	-0.04	50
60	0	-0.03	60
70	0	-0.03	70
80	0	-0.03	80
90	0	-0.03	90
100	0	-0.03	100

Table 6.2: difference between cumulative normal and Poisson distributions for a range of mean values from 10 to 100

The following conclusions can be drawn from this analysis:

1. The normal and Poisson cumulative distribution are significantly different for mean values up to 100 and should therefore not be considered equivalent.
2. The normal and Poisson mass functions (non-cumulative) show less error (less than 1%) and would suffice as equivalents for a different application.
3. The differences between normal and Poisson cumulative values are negative in each case, indicating that the normal probability values are lower than the Poisson values. For example, with a mean of 10, an expected value of 10 has a 58% likelihood with a Poisson distribution but only a 50% chance with normal. What this means for inventory holding is that if the mean number of failures of an item in a year is 10 and there are 10 spares in stock, then the Poisson distribution says that 58% of requests will be satisfied, as opposed to 50% under normal.
4. The Poisson distribution therefore predicts greater success than the normal distribution for a given stock holding.
5. Given that the Poisson distribution is recommended as suitable for a discrete process with small numbers it can be assumed that the Poisson outcomes for small mean values are reliable.
6. For increasing mean values, Poisson continues to forecast greater success for a given stock level. It is therefore considered appropriate to use the Poisson process across the entire range of values required for this study. This makes the Poisson forecast more optimistic, giving more conservative holding quantities. If the Poisson process can be assumed to be safe, then it will provide a more economic solution than the normal distribution.

Ultimately, the true test of which distribution is the most appropriate is a large-scale discrete event simulation to facilitate comparison of the rate and timing of events with actual historical data. If sufficient data could be collated and analysed, it would ultimately be possible to generation a real probability distribution profile based on historical performance. However, it is likely that this level of analysis would show differences in the failure modes of different types of equipment, such as electronic equipment (once it works initially, failure is random and should follow the Poisson distribution) compared with mechanical items (which will wear and deteriorate with use. Therefore the ultimate solution would have bespoke distributions for every different part: it would require massive volumes of failure history data to produce reliable distributions.

It would be possible to validate the use of the Poisson process through discrete event simulation. However, in the absence of detailed operational data, using mean values will give a solution whose output is of no greater quality than the forecasts produced by the predictive methods assessed here.

A further step would be to use a tailored distribution, such as a Weibull distribution (El Hayek 2005; see also http://www.weibull.com) for each part number, where the shape of the distribution curve can be altered to fit the data, but this would require very detailed reliability data over a long, fixed planning period.

6.2 Service Level Scaling

Of the four models to be evaluated, the first (Poisson) and third (Cost-Wise Skewed Holding) calculate demand at the line item level. Each line item will have one of three SL values, determined by its Essentiality Code. Further, different sets of SL values are tested:

$$SL_{FLS} = \{0.95, 0.93, 0.90\}$$

$$SL_{Airbus} = \{0.95, 0.89, 0.75\}$$

Where the stock is treated as a pool in the second (Marginal Analysis) and fourth (Linear Programming) models, it is necessary to consider how different items with different SLs can be equated in the same model. It is not appropriate to treat all parts as if they had the same SL, since the resulting recommendations would treat parts with lower essentiality in the same way as more important parts.

There are several possibilities for representing items with different SL requirements in a model that combines demand for all items – these are considered below.

1. Set a low global SL: give all parts a SL value of say 90%. While this is simple, it will under-provide parts that are essential to operations and over-provide part that would not disrupt operations.

2. Split the stock into SL groups: where there are three SL values, split the stock into separate problem spaces for Essentiality Codes 1, 2 and 3. This provides a simpler solution and should give good accuracy for each respective group. However, the effect of aggregate demand is lessened: a request for inventory can be for a part from any of the three groups. While this approach should give good results, it is better to aim for a combined model. The difference between the split stock and combined stock will be tested for the Linear Programming solution. Note that the Cost-Wise Skewed Holding (model 3) splits the stock into groups, but does so in order to over-stock low-cost items and under-stock more expensive ones. It will still be appropriate to weight items for different essentialities in model 3.

3. Scale data in proportion to SL values: where a model chooses using cost as a factor, scale the cost of the part up for lower-SL items to make them less attractive to the solution. Thus scale factors can be computed as follows:

For Ess 1, weighting = 1

For Ess 2 (SL = 93%), weighting = 95 / 93 = 1.021

For Ess 3 (SL = 90%), weighting = 95 / 90 = 1.055

Scaling the cost of an item will reduce its benefit (per unit cost) in a fleet-wide solution. Thus if a part with Ess 3, or SL = 90%, costs 1,000, then the weighted cost is 1,055. Thus a part of higher SL and the same base cost will be chosen first.

This approach is reasonably simple to apply, but is flawed in that the subject statistical distributions are not linear, i.e., the stock needed to meet a 95% target SL compared to the stock needed to meet a 90% requirement are not necessarily in a proportion of 95:90. In fact, the requirement to meet 95% demand can normally be expected to be more than 1- (95 / 90) = 5.5% greater than that to reach 90%. Therefore, while this approach does address the difference in SLs, it under-represents the difference.

4. Scale data in proportion to the expected values of the Poisson distribution: in order to proportion correctly for the different SLs, the cumulative distribution value for each SL can be used to create ratios between the SLs.

The ratios needed are obtained from the expected values that give the specific SL values sought. Since the distributions are based about a mean, a typical mean value must be chosen. Values of 30, 10 and 5 cover the typical values encountered and are used to estimate weights below.

Taking 30 as a mean value (since it gives a good range of integer expected values), the expected values and probabilities around 90% are as shown in Table 6.3.

X	36	37	38	39	40	41
Poisson	**0.880373**	**0.910987**	**0.935156**	**0.953747**	**0.96769**	**0.977893**

Table 6.3: Poisson distribution values with a mean of 30

For this mean value of 30, it is then possible to interpolate to find the expected values for 90, 93 and 95%:

X(90%) = (0.90 – 0.880373) / (0.910987 – 0.880373) + 36 = 36.641

X(93%) = (0.93 – 0.910987) / (0.935156 – 0.910987) + 37 = 37.787

X(95%) = (0.95 – 0.935156) / (0.953747 - 0.935156) + 38 = 38.798

Ratios can now be computed to give weightings for 90 and 93% parts compared to 95% SL parts:

W(90%) = 38.798 / 36.641 = 1.059

W(93%) = 38.798 / 37.787 = 1.027

These values are very close to the linear ratios derived in 3 above.

Repeating the calculations with a mean of 10 gives the following expected values:

X(90%) = 13.682

X(93%) = 14.388

X(95%) = 14.964

These values give weights of:

W(90%) = 14.964 / 13.682 = 1.094

W(93%) = 14.964 / 14.388 = 1.040

These values are almost double the values based on linear ratios.

Finally, the calculations are carried out for a mean of 5:

X(90%) = 7.511

X(93%) = 7.971

X(95%) = 8.499

Giving weights of:

W(90%) = 8.499 / 7.511 = 1.131

W(93%) = 8.499 / 7.971 = 1.066

As expected, these values are larger again than the linear ratios

In summary, the calculated ratios are shown in Table 6.4 below.

mean	*Linear*	*30*	*10*	*5*
SL = 90%	+5.5%	+5.9%	+9.4%	+13.1%
SL = 93%	+2.1%	+2.7%	+4.0%	+6.6%

Table 6.4: SL weights for a range of means

Given the large discrepancy for different mean values, it appears necessary to calculate ratios for all mean values, so that each part is adjusted for its SL at the correct mean value.

This exercise needs to be repeated for the second case of SL values, SL_{Airbus}, with values 95, 89 and 75%.

5. Generate aggregate scale data for Poisson values using the entire data set:

   As an approximation, the aggregate expected values can be generated from the data set following the Poisson method (method 1) using the Total Inventory Count (TIC) as follows:

   for the data set with mixed SLs:

   set all SLs to Ess1 value = 95% and record the total value of parts required = TIC(95%)

   set all SLs to Ess2 value = 93% (FLS) or 89% (Airbus) and record the total value of parts required = TIC(93%)

   set all SLs to Ess3 value = 90% (FLS) or 75% (Airbus) and record the total value of parts required = TIC(90%)

   Ratios for weighting parts with SL of 93 and 90% can then be calculated as:

   W(90%) = TIC(95%) / TIC(90%)

   W(93%) = TIC(95%) / TIC(93%)

   The following values are obtained from the data set for Case 1:

   TIC(95%) = 1040

   TIC(93%) = 971

   TIC(90%) = 895

Giving weights of:

W(90%) = 1040 / 895 = 1.162

W(93%) = 1040 / 971 = 1.071

This exercise is repeated for the second case of SL values, SL_{Airbus}, with values 95, 89 and 75%:

TIC(89%) = 874

TIC(75%) = 704

W(89%) = 1040 / 874 = 1.190

W(75%) = 1040 / 704 = 1.477

This approach will give proportionate weight within the solution space to each group of SL values; however this is still an approximation and will not represent values as well as would specific calculations for all quantities as in 4 above.

The most accurate method will be 4 above, where differences are calculated for each SL value for each mean demand value in the Poisson distribution. Method 5 would serve as a quick approximation, which could be built into a production application.

A full set of values is calculated by method 4: for each required SL and each mean value, a ratio is derived to represent the difference in expected value between the given SL and SL = 95%. It can be seen that these ratios tend to constants for large numbers. Values for means from 1 to 10 and 50 are shown in table 5 below.

	SL			
mean	*0.75*	*0.89*	*0.9*	*0.93*
1	2.31	1.36	1.32	1.15
2	*1.69*	*1.21*	*1.17*	*1.07*
3	1.57	1.2	1.17	1.08
4	1.47	1.17	1.14	1.05
5	1.44	1.16	1.13	1.07
6	1.39	1.14	1.12	1.05
7	1.36	1.13	1.12	1.05
8	1.34	1.13	1.11	1.05
9	1.31	1.12	1.1	1.05
10	1.29	1.11	1.09	1.04
50	1.13	1.05	1.05	1.02

Table 6.5: SL weights for mean values 1 to 10 and 50

Table 6.5 mostly shows weights for low mean values, where the difference is greatest. The values signify the difference in inventory quantity to satisfy the various SLs: looking at SL = 75% and mean = 3, the value of 1.57 means that for a mean value of 3, it will be necessary to have 1.57 times the amount of stock to satisfy 95% of demand, compared with meeting 75% of demand. The ratios are taken to even out at 50, so weights for mean values above 50 are given the numbers shown for a mean of 50.

The weights are applied to the item cost or demand (depending on the solution) to distort the input data in order of SL. Weighted cost is denoted $wcost_i$. The weights must then be applied in reverse to the solution values to obtain correct cost data.

6.3 Verification of results: quality assurance

It is necessary to apply measures to all of the models tested here to minimise errors in data preparation, calculations and processing of results.

6.3.1 Model 1: Poisson

Since the Poisson event calculations are the basis of all subsequent calculations for all of the models, it is of critical importance that these are correct and accurate. However, in the Poisson model, the calculations are simple since they are performed only at the line level.

Check 1: Given the nature of the calculations (especially granularity, or rounding up), the total solution SL value will always exceed the target SL.

Check 2: for parts with similar demand rates, the holding values should be the same or very close.

Check 3: a full model is implemented using the Normal distribution (with standard deviation set as the square root of the mean) and the difference between models assessed. From table 2 above, using a median value of 20, it can be predicted from the evaluation of distributions that the Normal distribution will prescribe a total inventory count around 6% greater than the Poisson solution.

6.3.2 Model 2: Marginal Analysis

This method is complex and produces a large number of calculations. The algorithm stops selecting parts when the target SL is reached. However, there is a flaw in this approach in that part quantities are not always picked in order of quantity, leading to situations where a quantity of a part is selected without the lower quantities of that part being included in the solution. This is logically inconsistent but cannot be resolved within this technique. The consequence of this flaw, which results from the shape of the distribution, is that in order to include the parts chosen for their high Marginal Contribution, it is necessary to include the lower quantities of the same part, which may not be in the solution set. The result of this action is that the method will over-provide parts in all cases. It is not therefore possible to use the global SL as a check on the number of fills resulting from the solution, as this number will exceed the theoretical number. Thus for example, if the total number of removals is 10,000 and target SL 95%, then the algorithm will continue selecting parts until the cumulative fill value is 9,500. However, due to the order in which part quantities are selected, the resulting solution set may contain 9,800 fills, giving an effective SL of 98%. Thus it is not possible to check that the total number of parts allocated gives exactly the expected number of fills (or satisfied demand events).

This model is inherently difficult to check other than by comparing solutions with the results from the other models, but a limited number of checks is proposed:

Check 1: calculate Δfill by two means as outlined in Chapter 4:

Δfill = fill(j) – fill(j-1),

where fill(j) = SL(j) x REMS(i) for quantity j of part i

Δfill = (SL(j) – SL(j-1)) x REMS(i)

Check 2: solve for a range of SL values, e.g., 90, 91, 92, 93, 94, 95 and check that total values are in line.

6.3.3 Model 3: Cost-Wise Skewed Holding

This method seeks a solution to match on objective, so the set of holding quantities produced should give a number of fills to equal the target number, or:

Σfills = SL x ΣREMS

This is a conclusive test and is sufficient to confirm that the solution set is consistent.

Check 1: take the solution set (holding quantities), apply Poisson values, calculate fills for each part and confirm that the sum matches the Σfills value calculated.

6.3.4 Model 4: Linear Programming – combined

This method also aims to meet an objective, so the solution set should give an exact match to the objective function value.

Check 1: take the solution set (holding quantities), apply Poisson values, calculate fills for each part and confirm that the sum matches the Σfills value calculated.

Check 2: perform sensitivity analysis by varying the objective function (SL) and observe the change in the solution values.

6.3.5 Model 5: Linear Programming – split

Likewise, this model resolves an objective function, so its output can be tested to give the required performance to exactly match the objective. A difference with this model is that the problem space is split into three, with three target SLs, so each is tested against its objective function.

Check 1: for each partition of the problem space (grouped by essentiality code) take the solution set (holding quantities), apply Poisson values, calculate fills for each part and confirm that the sum matches the Σfills value calculated.

Check 2: combine the holding quantities from the three solutions, apply Poisson values, calculate fills for each part and confirm that the sum matches the global Σfills value calculated.

Check 2: perform sensitivity analysis by varying the objective function (SL) in each solution and observe the change in the individual and aggregate solution values.

6.3.6 Overall checks

It can be predicted that the quality of the solutions is in the order below. Major deviations will suggest inconsistencies.

Check 1: compare total cost values for all solutions in each case, with the following anticipated descending order:

1. Actual holding
2. Poisson
3. Cost-Wise Skewed Holding
4. Marginal Analysis
5. Linear Programming - combined
6. Linear Programming – split

Check 2: graph cost and efficiency metrics for each method (y-axis) against five cases (x-axis) – if graphs cross or deviate, look for inconsistencies.

Check 3: perform random manual checks of calculated values at line item level, taking holding quantities prescribed by the various solutions and calculating the corresponding probability (Poisson value) to check that the appropriate Service Level value is obtained.

Check 4: transpose optimisation problems – for the linear programming models, exchange objective function (minimise cost) with the Service Level constraint by maximising SL for a fixed cost. This is evaluated in 6.6.3 and 6.6.4 below and proves that the models arrive at the same results from independent and distinct starting points.

Check 5: split the data set, perform separate solutions for each partition and combine the results. This is a simple and reliable check for consistency, which works for all techniques.

6.3.7 Validation of results

The detailed outputs of the models were reviewed with the provider of the test data and deemed to be feasible outputs, generally in line with current holdings and consistent with recent demand history.

The logic of the models and the overall results were presented to, and reviewed with several expert groups, including faculty at the MIT Department of Aeronautics and Astronautics and Engineering Systems Division, the Lean Aircraft Initiative at MIT, the Product Support Technical Committee of the American Institute of Aeronautics and Astronautics, HTS

Limited (a Maintenance, Repair and Overhaul provider) and Bombardier Aerospace (an aircraft manufacturer and provider of spares support). All of these participants approved the validity of the approach and the logical approach to the development of the models presented.

The theoretical problem, the solution approach and preliminary findings were also presented and reviewed at several relevant academic conferences. Supporting published working papers are listed in Appendix 4: Publications.

6.4 Solution outputs – test results

These are the measures of the output of each model as defined in the previous chapter. The values and relevance of these variables is presented herein for each of the 25 solution runs laid out in Chapter 4.

Below are descriptions of each of the table headings used to present the main test results in table 6.

Service Level % – whether the service level is fixed (as an input) or results from the calculation. For instance, in the case of the Poisson calculation, each part must exceed the target SL. This will result in a total SL > target SL due to granularity. Total SL = Σfills / Σremovals. The Marginal Analysis method will overshoot the target SL as discussed above. The CWSH and LP methods aim to meet the SL objective exactly.

Total cost $M – the sum of all quantities x inventory item costs for each run.

Actual – cost $M – the difference between the total cost of the actual holding and the total cost of the proposed holding. This is not shown for cases 4 and 5 since they assume a doubling of fleet utilisation so a comparison with actual operational data is not meaningful.

Total inventory count – the sum of all inventory quantities ΣHQ required to fulfil the SL requirement. For cost-skewed solutions (all except Poisson) a higher count of cheaper parts can be expected.

Average item value $ – again, where the solution is skewed for cost, it can be expected that this figure will be lower. The best solution will have the lowest average cost and the lowest total cost.

Count, compared to Poisson – since all of the solutions are so different from the actual case, comparison is made with Poisson, being the current practice for deriving theoretical values. The number of parts recommended by each solution is compared with the Poisson recommendation. The count may be higher for solutions that over-provide cheaper parts to meet SL target. Thus a high count is not a negative if total cost is favourable.

Saving vs Poisson $M – again, taking Poisson as the base case for calculated values, the improvement for each fleet-wide method is measured. This is then the expected benefit from each method as compared against current industry practice.

Saving vs Poisson % - the saving offered by each method expressed as a percentage of the cost of the Poisson solution.

Match with LP3, count

As discussed in Chapter 4, it is desirable to assess methods against each other. Since the actual holding is vastly different from the solution sets obtained by solution, it is more meaningful to compare each method against the best method, having confirmed that the best method offers a valid and correct solution.

The first metric counts the match in holding quantities over all items and expresses this relative to the total holding by the preferred method:

$$\mathbf{1 - (\Sigma|\text{recommended holding quantity} - \text{best holding}| \;/\; \Sigma\,(\text{best holding}))}$$

This metric indicates the proportion of all parts held that are at the "ideal" level compared with the best solution.

Match with LP3, cost

The second matching metric factors cost into the discrepancies between holding quantities from each method compared with the best method.

$$\mathbf{1- (\Sigma|(\text{recommended holding quantity} - \text{best holding})| \times \text{cost}) \;/\; \Sigma\,(\text{best holding} \times \text{cost}))}$$

This metric indicates the proportion of total cost that is spent correctly, or on the "ideal" levels of each part compared with the best solution.

The main solution values are shown in Table 6.6; the significance of the results and the performance of the models is assessed in Chapter 7.

Note that shaded areas in Table 6.6 reflect entries that would not be meaningful: the costs results for Cases 4 and 5 are not compared with actual data (from the test data set) since these cases represent twice the utilisation of the base case. Also, comparisons with Model 1 and Model 5 are shaded where they would be compared with themselves.

	Run	SL %	Total Cost $M	Actual – cost $M	Total inventory count	Average item value $	Count, compared to Poisson	Saving vs Poisson, $M	Saving vs Poisson, %	Match with LP3, count	Match with LP3, cost
Actual		89	32.9		2325	14164	1330	-17.6	-115	34%	-153%
Poisson	P1	96	15.3	17.6	995	15362				76%	33%
	P2	94	14.0	19.0	930	15006				76%	33%
	P3	96	13.2	19.8	865	15220				76%	36%
	P4	93	25.2		1649	15294				76%	32%
	P5	94	21.0		1338	15713				73%	29%
Marginal Analysis	M1	96	11.1	21.8	1050	10597	55	4.2	27	96%	86%
	M2	95	10.6	22.4	1006	10509	76	3.4	24	92%	79%
	M3	96	9.5	23.4	916	10395	51	3.6	28	93%	86%
	M4	99	25.2		2060	12227	411	0.0	0	84%	47%
	M5	97	18.4		1623	11364	285	2.6	12	85%	58%
Cost-wise skewed	C1	95	11.8	21.1	979	12063	-16	3.5	23	84%	65%
	C2	95	11.8	21.0	982	12178	52	2.0	14	83%	58%
	C3	95	10.6	22.3	820	12984	-45	2.5	19	82%	63%
	C4	95	23.0		1634	14083	-15	2.2	9	79%	45%
	C5	95	18.4		1389	13260	51	2.6	12	79%	49%
LP	L1	94	10.3	22.6	1019	10106	24	5.0	33	97%	93%
	L2	93	9.8	23.2	972	10035	42	4.2	30	94%	88%
	L3	94	8.6	24.3	882	9807	17	4.5	34	96%	96%
	L4	95	17.5		1785	9786	136	7.8	31	97%	93%
	L5	95	14.6		1513	9665	175	6.4	30	91%	84%
LP3	L3-1	94	9.7	23.2	1006	9662	11	5.6	36		
	L3-2	91	8.7	24.2	932	9381	2	5.2	37		
	L3-3	94	8.3	24.6	854	9709	-11	4.9	37		
	L3-4	94	16.4		1734	9469	85	8.8	35		
	L3-5	91	12.2		1381	8858	43	8.8	42		

Table 6.6: solution run results

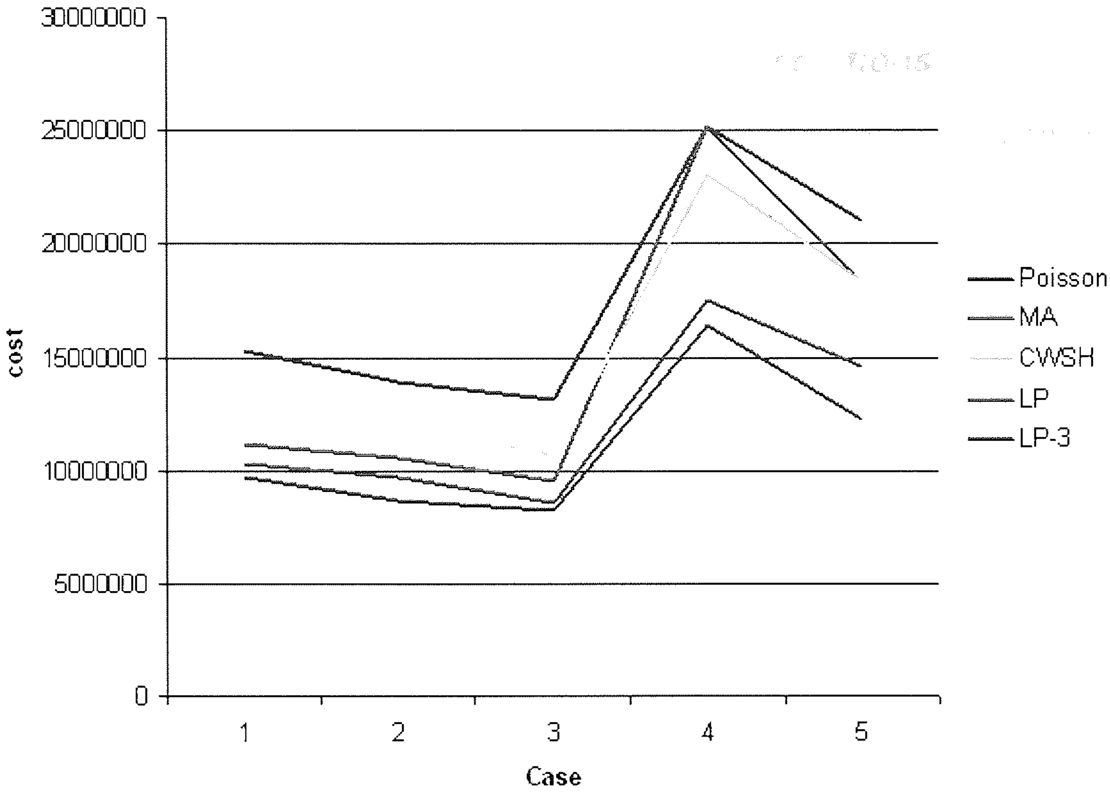

Figure 6.1: total cost ($M) results for each model for 5 cases

Figure 6.1 shows the total cost value obtained from each solution model for each of the 5 cases tested. Note that cases 4 and 5 entail a doubling of fleet utilisation so are not directly compatible with cases 1 to 3. However, the graph shows consistency for the different models, which serves as a quality check.

The values and relationship between solution outputs are discussed in the next chapter.

6.5 *Scale and computational intensity of the solutions*

It is worth considering the size of the solution space for each method since:

1. The data set and formulations used in this study are an order of magnitude smaller than a typical full-sized inventory database, which will contain several thousand parts (compared with 300 here);
2. The computational intensity of a solution will have a bearing on how easily it can be implemented, configured, connected to production systems and run frequently for decision support applications;
3. From a theoretical perspective, in general, a larger solution space can be assumed to evaluate a greater number of permutations so can be expected to be closer to optimal than a small-scale solution.

The four models proposed range widely in complexity. It is interesting to consider the value of each model:

- does a complex solution give a better result than a simple one?
- for a very complex solution, is the solution sufficiently better than a simple solution to justify the effort?
- which solution gives the best result for the data set used?
- which solution gives the best value (for the data set used): quality of result relative to the effort needed to generate the result?

6.5.1 Poisson

The Poisson solution processes one line at a time, so the increase in complexity with increasing size is simply linear. Thus a 300-part data set creates a problem space 10 times larger than a 30-part problem. This method picks one Poisson value per part so is very computationally simple. Thus where there are i parts, the complexity of this method can be given as

Complexity (Poisson) ~ i

6.5.2 Marginal Analysis

The Marginal Analysis solution creates a marginal contribution variable MC for each number j of each part i, so the set of values created is of size i x j. Thus for 300 parts with values from 1 to 15, the solution space is of size 300 x 15 = 4500. The 2-dimensional array of MC values is parsed into a list, together with Part Number identifier, quantity and Δfill rate values. This list is then sorted by descending MC and values picked until the target Service Level is achieved:

$$\Sigma(\Delta\text{fill rate}) \geq SL \times \Sigma REMS$$

Since the solution space is O(i x j), it increases in a linear fashion for a large problem. For example, 3,000 parts with 15 values gives a sort list of 45,000 values, which is a very efficient solution.

Complexity (Marginal Analysis) ~ i x j

6.5.3 Cost-Wise Skewed Holding

This method takes a small number of groups of parts ranked by cost – in the trials herein there are 5 cost groupings. This is a simple sort on the number of parts. Once the parts are grouped, the model performs a simple selection of values as in the Poisson method above, with different Service Level values for each group. The consequential global SL value is calculated and compared with the target SL. Adjustments are then made to the problem parameters until a satisfactory global SL level is attained.

This method is a partitioned version of Poisson, so, overlooking the initial sort, the complexity will be the same as the Poisson method:

Complexity (cost-wise skewed holding) ~ i

6.5.4 Linear Programming

In the LP solution, there is a logical constraint that one non-zero value may be assigned to each part i. There are j options for each part. The solution seeks to assign j for each part i such that the overall target SL is reached with the solution set giving the lowest-cost permutation among all available combinations. Each part i has j possible quantities. For 1 part, there are j possible solutions. For 2 parts there are j solutions for the second part for each of the j solutions for the first part, so there are j^2 solutions. Thus the problem space is multiplied by j for each extra part number:

Complexity (Linear Programming) ~ j^i

So for 300 parts each with quantities from 1 to 15, there are 15^{300} combinations, an unfeasibly large number.

While it might appear useful to limit the solution space by eliminating outlying values, this is not really recommended since the LP solution may pick extreme quantities based on cost. Thus for example, the solution may pick 1 as a quantity (assuming that no zero values are allowed) for a high-cost item even if the individual part SL is low (say 50%). Meanwhile the solution may pick very high quantities for low-cost parts, with implied individual SL values over 99%, since these higher quantities will add to SL performance by satisfying demand events. It may be possible to "trim" the solution set by say picking the half of the part list with lower costs and eliminating the quantities with implied SL below target SL. This may be worthwhile as the problem set becomes larger.

The number of iterations for each LP solution is shown in Table 6.7 below. Note that the larger problems have small numbers of iterations, suggesting there is greater "slack" in the solution space, in other words it is easier to find the optimal solution. While this appears counterintuitive, validation of the results confirms that the solutions give good results in line with expectations.

iterations	*LP*	*LP3*
Case 1	290383	3618
Case 2	232266	3132
Case 3	94790	3432
Case 4	1351	2992
Case 5	2603	3141

Table 6.7: iterations for Linear Programming solution runs

6.6 Solution sensitivity analysis

In addition to testing the models on multiple cases, it is useful to perform sensitivity analysis on the models. This analysis is limited to the two Linear Programming techniques for the following reasons:

1. The Poisson solution takes assigned SL values, which are problem inputs. The sensitivity analysis used here is performed by varying the problem's objective function, SL.

2. The Poisson problem is a comparatively simple method of calculation so can be expected to be stable without testing.

3. The Marginal Analysis model has an inherent logical flaw. Its solution values (as tested with the present data set) are lower than the LP solution values to the extent that the performance of the technique can be considered inferior, although testing with further data is needed in order to compare the methods more rigorously. Further, the Marginal Analysis model is cumbersome to implement so it would take a large effort to perform sensitivity analysis for little benefit.

4. The Cost-Wise Skewed Holding model incorporates sensitivity analysis since the user is required to alter the set of 5 SL inputs in order to get the best solution (reach SL for minimal cost).

6.6.1 LP combined model

This model considers all parts together, aiming to meet a stated SL value for the lowest cost.

The general target SL is 95%: this is high on aggregate since the problem sets use a combination of values, of which 95% is the highest. SL scaling aims to distort the solution towards the higher-essentiality parts, but the overall SL should still be lower. There is no simple means of establishing an appropriate lower value, so the highest value is maintained.

Sensitivity analysis is performed by varying the target SL, and the consequent objective function (total cost) values recorded. This has the effect of evaluating the financial impact of changes in operational criteria, and serves as a rigorous quality check on the integrity of the model.

The values obtained are presented in Table 6.8 below and are plotted graphically in Figure 6.2.

SL	*Case 1*	*Case 2*	*Case 3*	*Case 4*	*Case 5*
99%	16604534	15226447	13164531	25841427	21665764
98%	13587526	12726970	11240256	22661758	19032285
97%	12158655	11448480	10170770	20589652	17255650
96%	11164468	10538056	9327826	18930383	15846071
95%	10298520	9753910	8649440	17467965	14622747
94%	9628634	9101200	8149919	16279881	13619039
93%	9088736	8645625	7774285	15320969	12846059
92%	8673546	8279183	7472836	14498928	12209163
91%	8340153	7975537	7230206	13785024	11628991
90%	8059076	7719282	7031881	13182903	11151060
89%	7831146	7523178	6866818	12685824	10759722
88%	7641177	7349931	6713172	12277815	10434371

Table 6.8: sensitivity analysis for the LP combined model: total cost

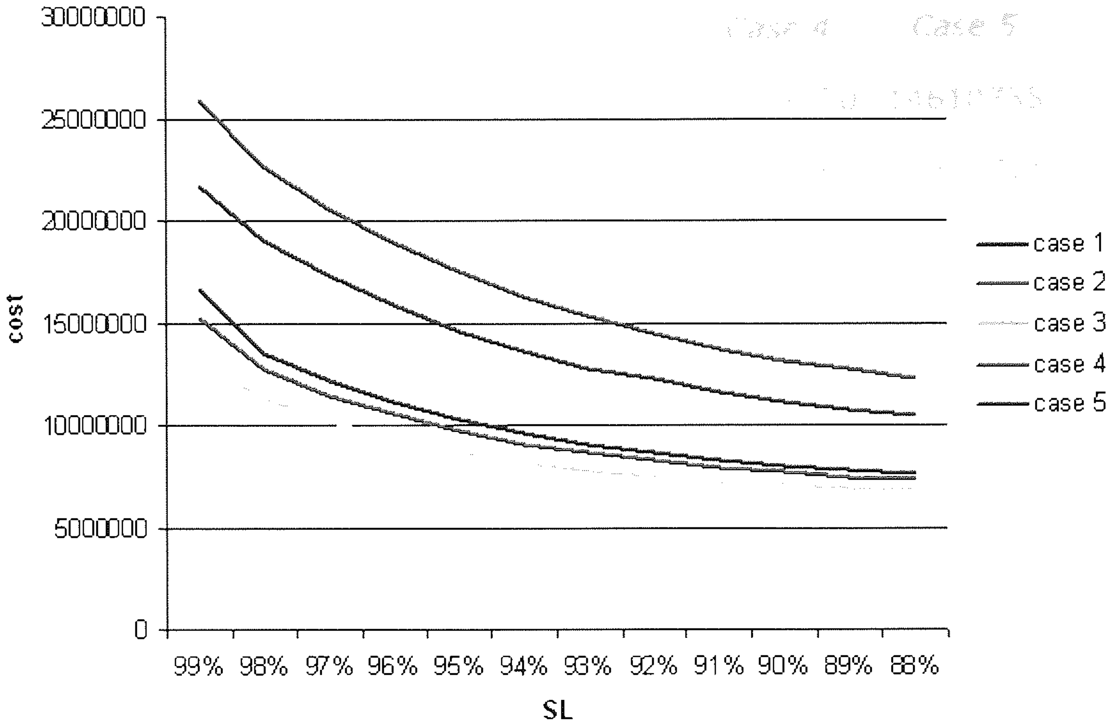

Figure 6.2: sensitivity analysis for the LP combined model

6.6.2 LP split model

This model divides the data set into three by essentiality codes. Each essentiality code has a corresponding SL. There are thus three problem formulations for each case. The results are then combined and an overall SL observed.

Performing sensitivity analysis on this model requires varying each of the three SL values in the separate model formulations. Since this presents a large number of permutations, a simple approach is adopted here: all three values are varied by the same amount together. The results of this testing are shown in Table 6.9 below.

SL	Case 1	Case 2	Case 3	Case 4	Case 5
3%	12198319	10366639	10274518	20561870	14610755
2%	11172505	9675445	9437239	18989905	13707711
1%	10357448	9159622	8777574	17542206	12909800
0%	9720018	8743446	8291202	16419984	12232908
-1%	9198482	8385689	7905505	15485949	11660340
-2%	8797708	8099653	7601996	14670135	11190008
-3%	8468745	7863297	7357385	13964843	10788308

Table 6.9: sensitivity analysis for the LP split model: total cost

Note that in Table 6.9 there are different sets of SL values used: in Cases 1, 3 and 4 the set of base values is SL_{FLS} = {95%, 93%, 90%}, while in Cases 2 and 4 the base values used are SL_{Airbus} = {95%, 89%, 75%}. The set of values is then varied in unison, so SL + 3% in Case 1 = {98%, 96%, 93%}. These extreme values are not meaningful in operational terms but serve to test the stability of the model in checking for expected responses to sensitivity changes.

Figure 6.3 below shows this sensitivity data in graphical form.

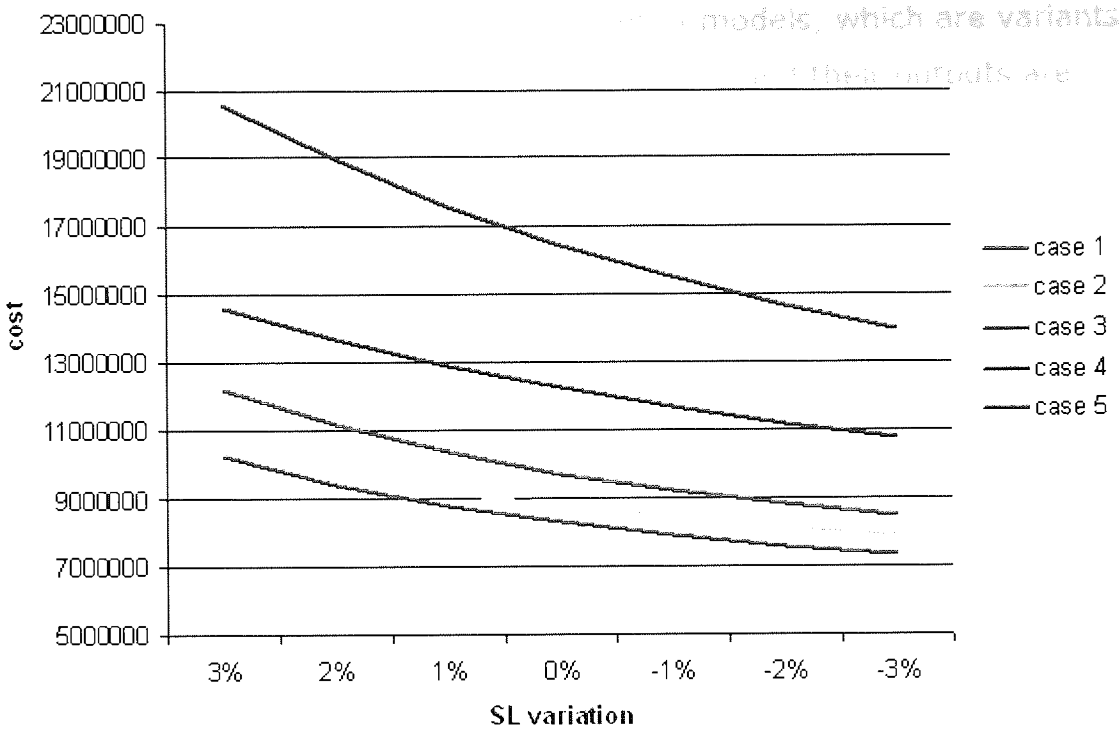

Figure 6.3: sensitivity analysis for the LP split model

In Figure 6.3, the x-axis represents all SL values being reduced in unison reading from left to right, with the resulting total solution costs shown by the graphed curves. The lower 3 lines are models where the utilisation and resulting demand are the same; the top two lines reflect twice the level of utilisation and demand.

The two LP solutions above are driven by the cost minimisation objective, with a stated SL target as a constraint. A futher option, one that would be of interest in practice, is to transpose the problem, i.e., fix the cost and maximise SL. In other words, the problem statement can be changed from:

minimise cost for a stated SL

to

maximise SL for a set budget.

As well as presenting the problem in a useful configuration, rearranging the formulation in this manner provides a further measure of sensitivity analysis.

This alternative perspective calls for a new set of models, which are variants on the LP models discussed so far. The new models and their outputs are discussed below.

6.6.3 LP combined model, cost-oriented

The objective function of the LP model is changed to a maximisation of the number of fills, which will maximise SL, and a constraint introduced:

Σ (cost x selected quantity j) for each part i < budget

where budget is some arbitrary amount, a test value.

Having performed the first set of LP solutions above, the required budget to reach a target SL is known: for instance, case 1 requires a budget of $10.3M to reach a 95% target SL. It is therefore of interest to set budgets in steps above and below this value, to see the consequent SL value.

The SL values obtained for a range of budgets are shown in Table 6.10 and plotted in Figure 6.4 below.

Budget $M	Case 1	Case 2	Case 3	Case 4	Case 5
5	62%	64%	70%	42%	48%
6	76%	78%	83%	55%	62%
7	84%	86%	90%	64%	71%
8	90%	91%	94%	71%	78%
9	93%	94%	96%	77%	83%
10	95%	95%	97%	81%	87%
11	96%	97%	98%	84%	90%
12	97%	97%	98%	87%	92%
13	98%	98%	99%	90%	93%
14	98%	99%	99%	91%	94%
15	99%	99%	99%	93%	95%
16	99%	99%	100%	94%	96%
17	99%	99%	100%	95%	97%
18	99%	99%	100%	95%	97%
19	99%	99%	100%	96%	98%
20	99%	100%	100%	97%	98%
21	99%	100%	100%	97%	99%

Table 6.10: transposed LP combined model: SL values for set budgets

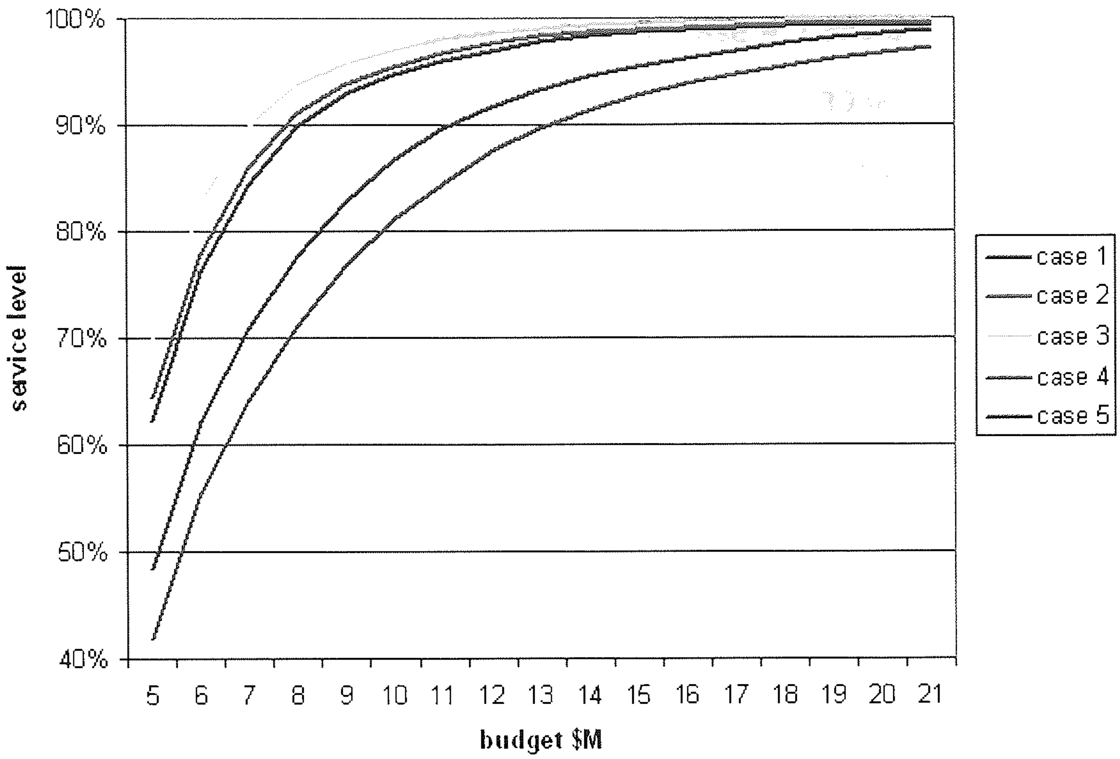

Figure 6.4: transposed LP combined solution values

6.6.4 LP split model, cost-oriented

The LP split model, denoted LP3, is transposed and formulated for a range of budget values. Using the same values as for the LP combined model, with each problem split into three essentiality groups, there are:

3 Ess codes x 5 cases x 17 budget sensitivity values = 255 models.

An assumption is made in relation to the budget allocated to each Ess code: this could be varied, but for consistency, the budget is apportioned in ratios derived from the results achieved for the first-run LP3 solutions (i.e., with cost as the objective function), shown in Table 6.11 below.

	Case 1	Case 2	Case 3	Case 4	Case 5
Ess1	35%	38%	35%	34%	39%
Ess2	63%	59%	62%	63%	58%
Ess3	3%	2%	3%	3%	2%

Table 6.11: budget ratios derived from LP3 solutions

Table 6.11 shows the share of each budget used in each solution, which consists of three LP formulations. Thus, if a budget of $10M is set, then the models for case 1 are allocated $3.5M, $6.3M and $0.3M for Ess code 1, 2 and 3 respectively.

The solution formulations return a number of fills for the set budget. The fills are added for each of the three solution sets and the sum divided by the total number of removals to give the overall SL:

SL = (fills(ess1) + fills(ess2) + fills(ess3))/ REMS

The resulting SL values are shown in Table 6.12 for a range of budgets.

Budget $M	Case 1	Case 2	Case 3	Case 4	Case 5
5	61%	58%	68%	41%	46%
6	75%	74%	81%	55%	60%
7	83%	83%	89%	63%	69%
8	89%	88%	93%	70%	76%
9	92%	92%	95%	75%	81%
10	94%	94%	96%	80%	85%
11	95%	95%	97%	83%	88%
12	97%	96%	98%	86%	90%
13	97%	97%	99%	89%	92%
14	98%	98%	99%	91%	93%
15	98%	98%	99%	92%	95%
16	99%	99%	100%	93%	95%
17	99%	99%	100%	94%	96%
18	99%	99%	100%	95%	97%
19	99%	99%	100%	96%	97%
20	99%	99%	100%	96%	98%
21	100%	99%	100%	97%	98%

Table 6.12: transposed LP split model: SL values for set budgets

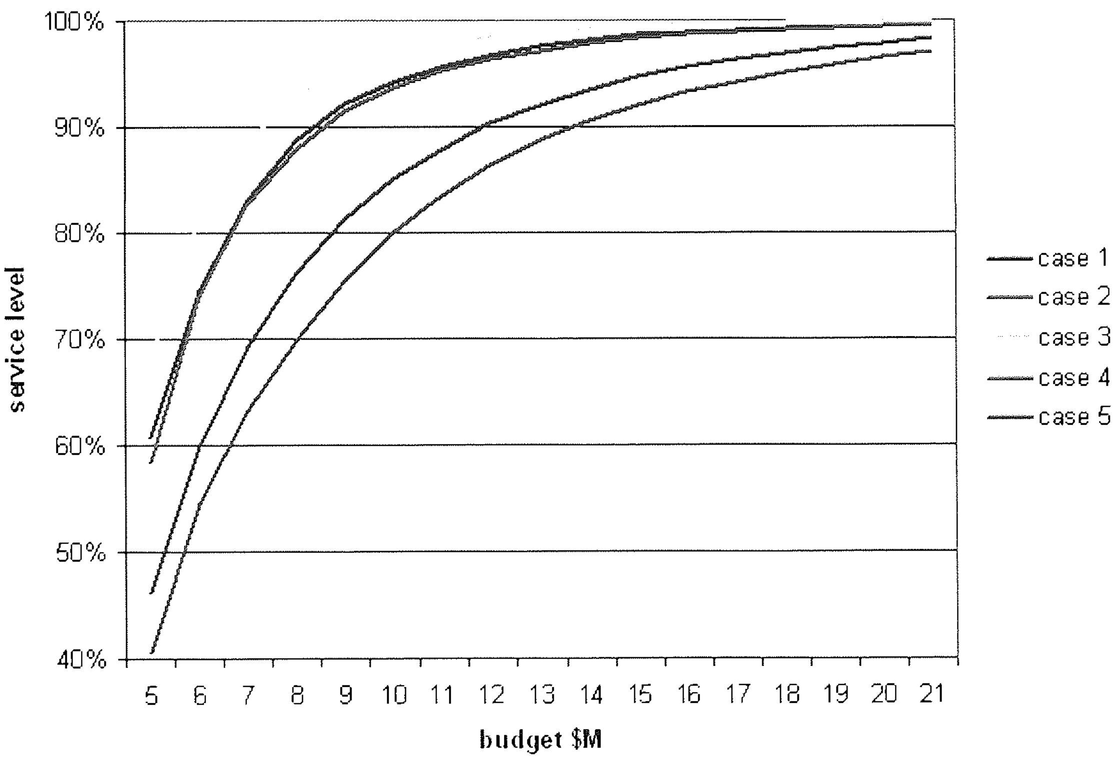

Figure 6.5: transposed LP split solution values

Figures 6.5 and 6.6 allow the user to see what SL is available for a range of budgets: the steeper the curve, the greater is the gain in SL for each incremental unit of budget allocated. This is almost a transposed version of the curves obtained for cost against target SL values earlier, however, these values are the result of a different set of model formulations and solutions. It serves as a quality check to observe that values coincide: in case 1 for the SL-constrained solution, the cost was $10.3M, while the values here show that SL ranges from 95 to 96% for budgets of $10M and $11M. Interpolation of the detailed SL values (94.56 and 95.82%) gives a result of $10.35M for 95.00%.

There is a discrepancy between the combined and split models: for the SL-constrained solutions earlier, LP3 had a clear cost advantage over LP, but it did so with lower global SL since it takes better account of the different essentialities. In the cost-oriented solutions, the combined model reports higher SLs than does the split model, for the same budget. In the combined model, where SL scaling is applied to items with essentiality codes 2 and 3, these parts may still have higher SLs than their target, if their cost is favourable (low) to the solution. In the split model, the optimisation is more accurate since each model has the relevant SL as a constraint. In any case, it is fair to assume that the LP split model will be more accurate, although it is not clear how budget should best be divided among essentiality codes.

This chapter has presented the outputs obtained by processing the same set of five test cases with five solutions. The results show consistency across the cases, suggesting that the model implementations and solution results are valid. The results also show progression through the models: as the models become more sophisticated, the results improve. The impact of these results is considered in Chapter 7 below.

Chapter 7: Analysis of model outputs

This chapter interprets the results put forward in the last chapter, assessing the performance of each modelling approach in its own right and in comparison to the other models. The quality of the solutions is considered and compared against the best solution and the advantages and disadvantages of each approach are reviewed. The range of cases tested is reviewed for each model and over all. The actual operational data set is measured against the range of solution results.

Finally, the potential improvement offered by the models is identified and recommendations are made about the best approach to solving the subject problem based on the experience of testing the range of Models.

7.1 Summary of test cases

Five cases are tested in each solution model. The characteristics of the cases are reiterated here, along with the rational for choosing these cases. Note that cases are also referred to as cases in this analysis.

Base – data as obtained, using current SL values {95, 93, 90%}. This is the case applying to the data set provided and can be compared directly with the actual operational stock holding.

Fewer – same data set, using a wider range of SL values as prescribed by Airbus {95, 89, 75%}. It is expected that by reducing the SL targets for parts whose failure does not prevent an aircraft from flying, it should be feasible to cut stock with little impact on operations.

Faster – same data set, with a 5-day reduction in repair Turn-Around Times. In order to assess the impact on stock levels of faster repair processing, a realistic TAT reduction is introduced. This is a target stated by Airbus in seeking to gain rotable support contracts with their airline customers. It is important to emphasise the effect of repair times on stock levels: in theory, if repair times were zero, then there would be no need for any spare stocks.

Bigger – double the fleet size for the base case. The data tested represents an inventory pool supporting an average of 22 aircraft. There is potential to use the stock more efficiently by increasing the size of the supported fleet. Due to the stochastic nature of demand, it is anticipated that the marginal cost of providing spares support for additional aircraft will be small. This is important to an MRO like FLS (who provided the data) since they will routinely bid for contracts to support compatible fleet. The case presented here involves twice the level of utilization in order to determine the required increase in stock levels and attendant costs to meet target SLs.

Best – double the fleet size, apply the FLS SL values and 5-day TAT reduction. If all of the proposed measures could be achieved together, namely, a revised set of SL values, faster repair handling and supporting a larger fleet, it can be expected that the overall performance of the inventory pool will be more efficient and the unit cost the least.

7.2 *Interpretation of results*

Base data: current inventory value

The existing inventory pool on which comparisons are based has a SL value of 89% at a cost of $33M and an inventory count of 2325. Compared with all solution methods, this represents a lower SL at over twice the cost and twice the number of parts held. This is explained by a long history of uninformed decisions over many years, where repair times were not well managed and no overall control was applied to cost. Since the cost of inventory is an asset cost, not a running cost, it is not usually the focus for operational improvement. Also there is an attitude that it is preferable to spend more money on spares than to ground an aircraft. There may be epidemics of part failure or campaign changes, leading to a rush on certain spares, which are not subsequently depleted. Since there is no imperative to reduce assets, it is rare that an operator will aggressively cut inventory levels by selling off excess.

This study compares a range of techniques for solving the rotable problem in isolation of the historical behaviour of inventory managers: while the actual inventory holdings are shown for comparison, it is the comparison among models that is of interest.

The level of the actual inventory holding is so far out of line with the model solutions that it is not used extensively for comparison below, rather most comparisons are made against the best solution.

7.2.1 Model 1: Poisson

	SL %	*Total cost $M*	*Count*	*Average part value $*
Case 1	96	15.3	995	15362
Case 2	94	14.0	930	15006
Case 3	96	13.2	865	15220
Case 4	93	25.2	1649	15294
Case 5	94	21.0	1338	15713

Table 7.1: Model 1 results

Table 7.1 shows the SL, total cost, total part count and average part cost for each of 5 cases tested with the Poisson Model. The Poisson method gives a large improvement over the actual inventory holding reported in the data set: half the cost and half the number of items, with a far higher SL. This indicates that the actual holding, as well as holding too many parts at too great a cost, is holding the *wrong* parts to give a high SL. This model can therefore be seen to give a sound set of recommendations, with two deficiencies. The first is that cost is ignored by this approach so there is no bias toward stocking relatively more cheap parts than expensive ones. The second shortcoming of this model is granularity, or rounding up of each small quantity. This granularity results in a SL of 96% for the first case, when the SL values in use are 95, 93 and 90%. Likewise in the second case, with SL values of 95, 89 and 75%, the overall SL is still high. This calls for more detailed consideration, so the average SLs for each Ess Code are obtained for cases 1 and 5 (being the most different cases) and shown in Table 7.2 below.

Case	*Ess Code*	*Target SL %*	*SL achieved*
1 - base	1	95	97
	2	93	96
	3	90	94
	Total	95	96
5 - best	1	95	96
	2	89	92
	3	75	82
	Total	95	94

Table 7.2: Service Levels set and attained for Model 1

Table 7.2 shows target and attained SLs for the base case, Case 1, and the largest case with the greatest spread in SL values, Case 5. In both cases the attained SLs differ substantially from the target values, due to granularity in the line-level calculations. The Total values refer to the global SL set and achieved for the combined stock in both cases.

Given the excessive SLs achieved it must be concluded that this technique, in addition to ignoring cost, is less than optimal since it prescribes aggregate inventory levels that are far above the ideal. It would be possible to lower the input global SL by trial and error until the attained SL meets the target, however there is no assurance that this is a good solution, especially since the model does nothing to relate demand events to each other across the inventory set.

7.2.2 Model 2: Marginal Analysis

	SL %	*Total cost $M*	*Count*	*Average part value $*
Case 1	96	11.1	1050	10597
Case 2	95	10.6	1006	10509
Case 3	96	9.5	916	10395
Case 4	99	25.2	2060	12227
Case 5	97	18.4	1623	11364

Table 7.3: Model 2 results

What is immediately apparent from the results in Table 7.3 is that, however good the cost outcome may be, the overall SL obtained is very high. Given that this model takes a combined approach to the data, it is necessary to set the global target SL at the highest value, in this case 95%. Thus items with lower SLs will still be considered for this target. SL scaling mitigates this to some extent and is discussed further below.

Apart from the inefficiency inherent with the combined SL approach, there is still a large overshoot of the SL, giving a solution that appears to be more expensive than it should be. This could possibly be addressed by iteratively reducing the global SL input until the output SL meets the required level. However, that would not then strictly be an optimisation approach and would contain some degree of uncertainty.

The reason for the SL being exceeded in each case is the order in which the model chooses part quantities: as mentioned before, due to the shape of the probability distribution, the algorithm may pick part-quantity pairs ranked by marginal contribution (MC), which do not occur in ascending order of quantity. Thus if the lower quantities have MC values which fall outside the solution set (the MC is less than the lowest MC counted to reach target SL), then the solution set includes broken sequences of quantities – see Table 7.4 below. It is then necessary to stock extra parts to complete the appropriate part number of each quantity. There is no obvious way to address this flaw; a more sophisticated model is called for.

line	*quantity*	*MC*	*Δfill*	*Action*
216	4	0.00013	14.6091	Included in the solution
216	5	0.00012	13.4697	
216	3	0.00011	12.6759	
216	6	9.3E-05	10.3493	
216	2	7.4E-05	8.24893	Discarded by the algorithm
216	7	6.1E-05	6.81579	
216	1	3.9E-05	4.35497	
216	8	3.5E-05	3.92763	
216	9	1.8E-05	2.01183	
216	10	8.4E-06	0.92746	

Table 7.4: Marginal Analysis quantity order selection

The selection shown in Table 7.4 is based on a cut-off MC value of 7.8E-05: part-quantity pairs are selected in descending order of MC until

Σ(Δfill) > REMS x SL

or the required number of fills is counted to meet the SL requirement.

However, as can be seen from Table 7.4, the quantities 1 and 2 of part 216 are not counted in the solution. For the solution to be logical, however, these lower quantities must be added, with the result that there are too many parts picked. In the case of this part, the algorithm picked a quantity that would give a SL of 51% for that part (since it chooses parts considering cost and fill rate), whereas the count in the final solution gives a SL of 82%.

The fact that the part quantities do not give MC in descending order is a function of the probability curve: Figure 7.1 shows the Poisson mass function (non-cumulative) curve for this part, which has a mean value of 4.61.

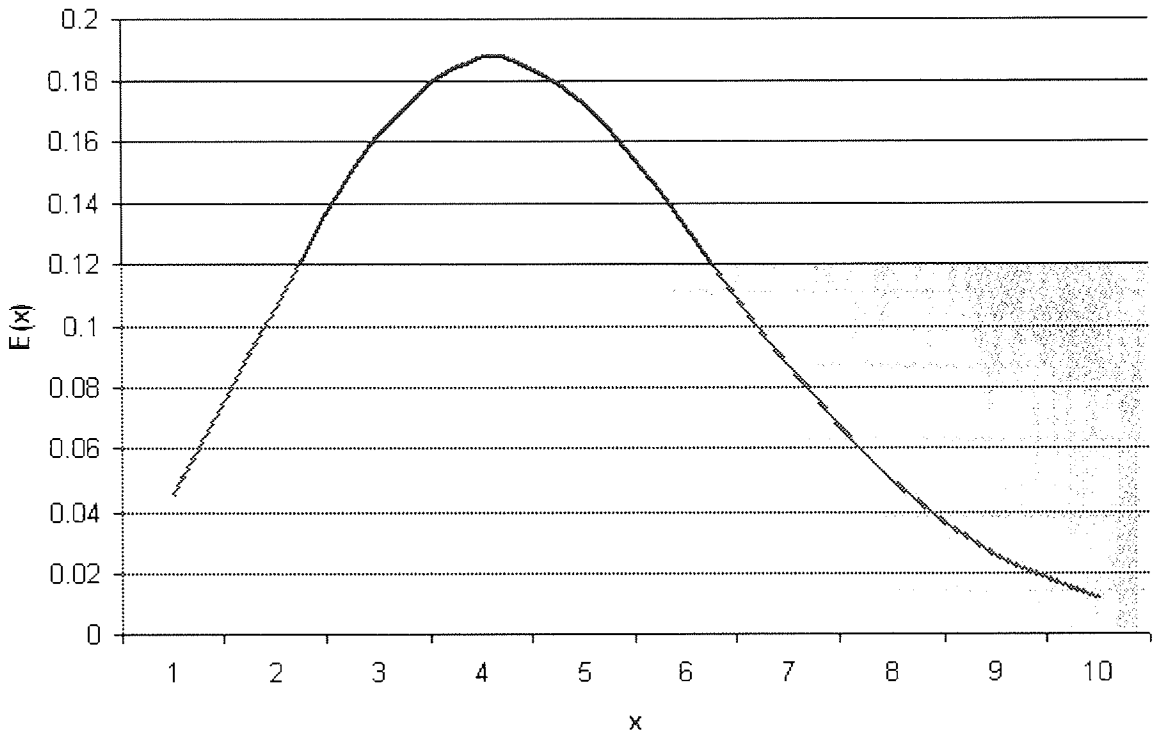

Figure 7.1: Poisson distribution for a mean value of 4.61

The MC cut-off value used in the model can be estimated as corresponding to an expected value around 0.12: quantities with higher values (3, 4, 5 and 6) are high enough to be included, while lower values (shown in the shaded area in Figure 7.1) are discarded.

The logical inconsistency inherent in this approach is exacerbated by size: as part quantities increase, so the fill rates (and therefore MC) for low numbers of a part with a high mean value become lower and less likely to be selected.

It would be possible to formulate a split version of this model, i.e., with three problems corresponding to each level of essentiality. However, given the extent to which the solution overshoots the target SL, it is not considered a strong candidate for the best optimisation method and so does not warrant this further exploration.

As mentioned above, given that this model over-stocks, it could be run with lower input global SL values until the outcome approaches the target SL, however, given the selection process, it is not clear that such a solution would necessarily be close to optimal.

The SL values attained by the marginal analysis method are plotted in order of increasing part cost in Figure 7.2.

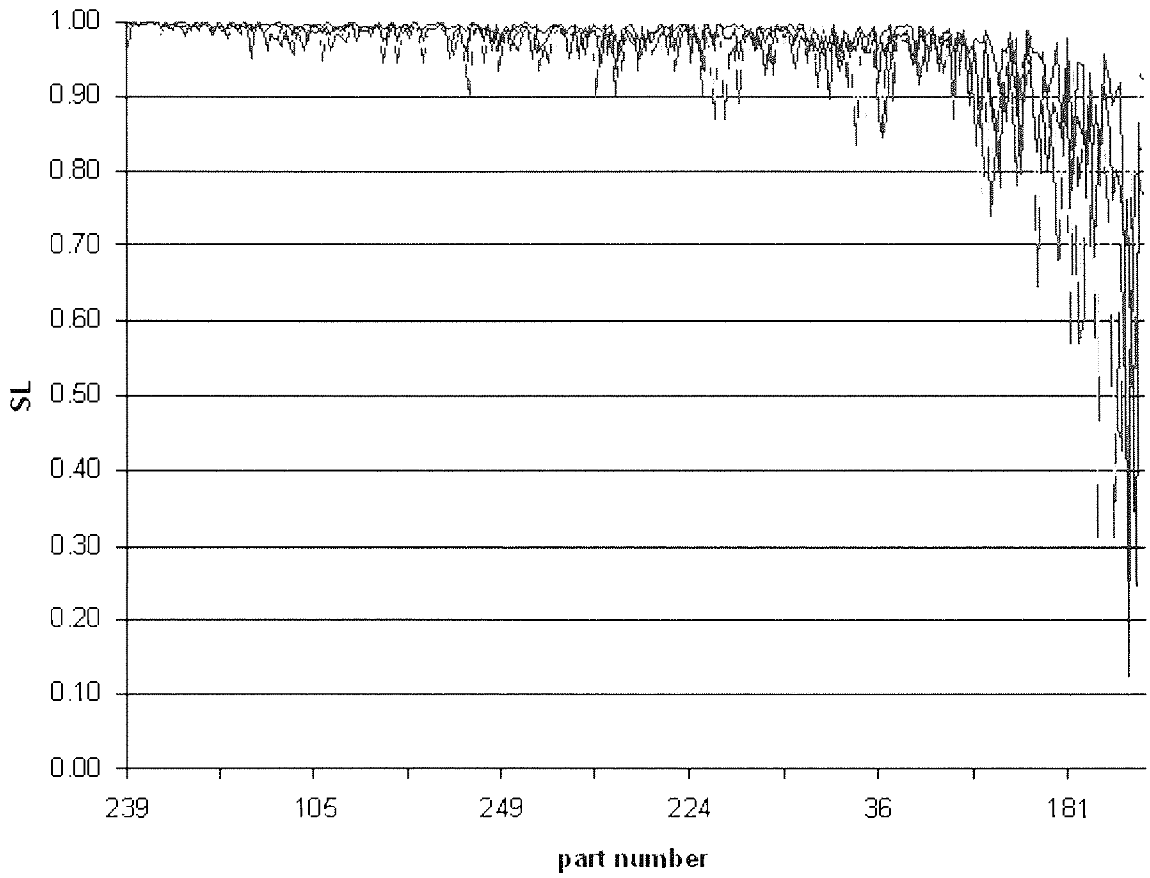

Figure 7.2: SL plotted against part numbers sorted by cost, Model 2

7.2.3 Model 3: Cost-Wise Skewed Holding

	SL %	*Total cost $M*	*Count*	*Average part value $*
Case 1	95	11.8	979	12063
Case 2	95	11.8	982	12178
Case 3	95	10.5	820	12984
Case 4	95	21.3	1634	14083
Case 5	95	17.8	1389	13260

Table 7.5: Model 3 results

This model uses SL values at two levels – in the first instance to apply SL scaling weightings to cost in order to group items into 5 bands ranked by cost. The user then varies a set of 5 variable SL values and a consequent global SL obtained. An advantage with this approach is that the model can be aimed to meet a target SL quite accurately. However, given that the model aims at a single target SL, it will over-provide the items with lower essentiality to meet the target SL. Also, the defined SL value set (FLS and Airbus) is used only to rank items so the effect of their difference is minimal.

Table 7.5 shows the total cost results for each case by this method.

Thus, the difference between Cases 1 and 2, with different SL ranges, is minimal because the internal variable SL values are applied to each cost group, and in both cases the best results are obtained with the same 5 SL values.

The set of SL values used for each cost band is given in Table 7.6 below: in each case, these parameters yield an overall SL of 95%.

SL %	Band 1	Band 2	Band 3	Band 4	Band 5
Case 1	97	97	95	92	75
Case 2	97	97	95	92	75
Case 3	97	97	95	95	75
Case 4	97	97	97	97	75
Case 5	97	97	95	95	75

Table 7.6: SL parameters applied to parts grouped into 5 cost bands

The range of SL values appears in Table 7.6 in ascending order of cost, i.e., for Case 1, 97% is applied to the band of one-fifth of the parts with the lowest weighted cost, while 75% is applied to the band with the highest cost.

While this model is easy to use and allows sensitivity analysis, it is not a pure optimisation, since trial and error evaluation is needed, and there are only 5 partitions. Also, the combined problem set does not allow sufficiently for different SL values. Again, the global SL values have a diminished role in the solution (they are used for weighting and ranking only) since the SL values used in calculations are those varied by the user.

Since the model seeks a goal by trial and error, and the number of partitions is small (and arbitrary) it can be assumed that it is some way from being optimal. However the model is of value for quick calculations and sensitivity analysis of SL performance. It may be more realistic to lower the global SL target, since the value used here is 95%, but this is employed in order to facilitate direct comparison against the Poisson method.

It is worth observing that the outcomes from Cases 1 and 4, which differ by doubling utilisation, indicate a near doubling of the total cost. This is an unexpected result, since there is an expectation that doubling utilisation makes more efficient use of the stock pool and would give economy of scale, requiring an increase in stock less than the increase in utilisation. This may be explained to an extent by the scale effect experienced in Model 2, and is a limitation of the Poisson distribution, which is aimed at small numbers of events over a planning period. As discussed earlier, there are several rules of thumb whereby inventory planners move from Poisson to Gauss (normal) distributions at some agreed point, such as 20 or 50. However, reviewing the data shows that, while some removal rates are high, the demand rates (obtained by scaling down by Un-Availability Factor, or the proportion of time during which an item is in repair) used as inputs to the probability calculations are low, with only 3 values above 20.

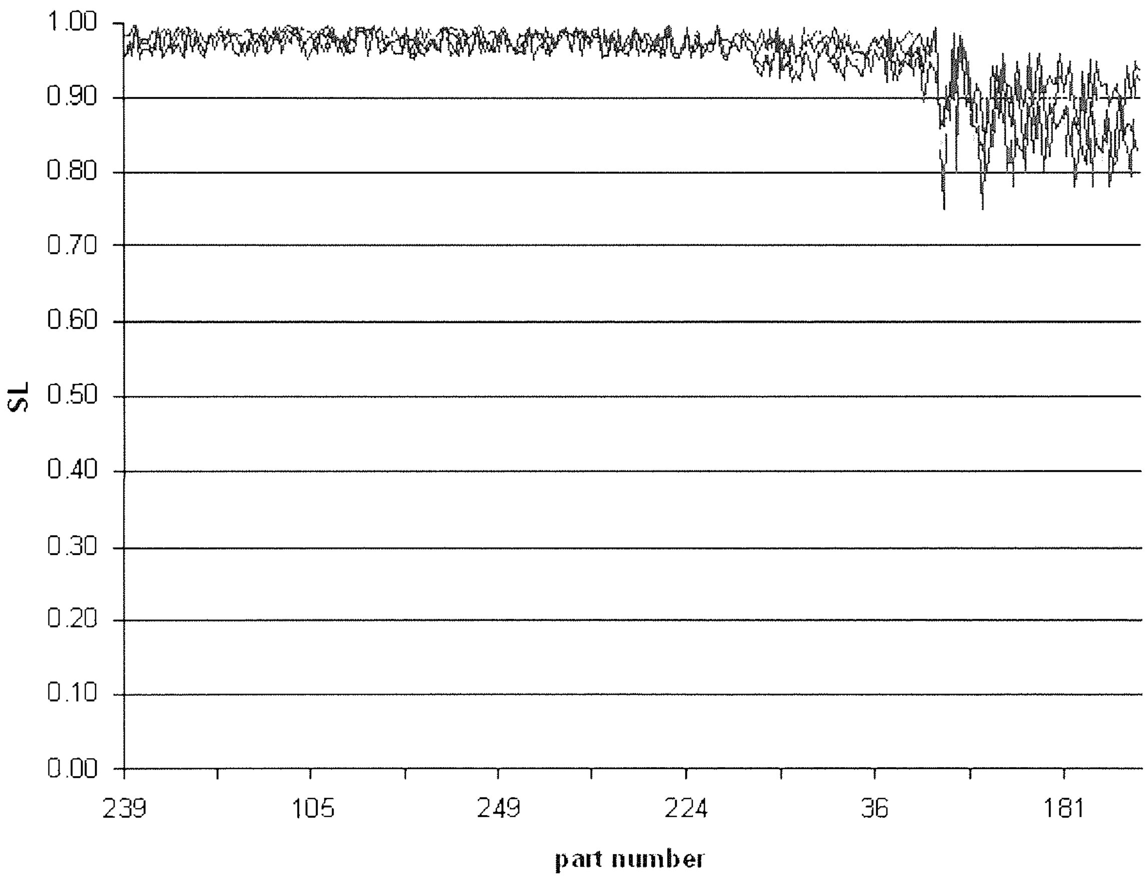

Figure 7.3: SL plotted against part numbers sorted by cost, Model 3

Figure 7.3 shows SL values for parts ranked by cost in ascending order. The one-fifth of the data to the right of the graph represents the highest-cost band.

Reviewing the model and its results, it appears that the grouping of parts into even divisions (with equal numbers of parts) may not be very effective, since the effect of cost is greatest in the highest-cost band. This is reflected in Table 7.6, where SLs are high for the low-cost bands (as expected) but the best results seem to arise with a sharp drop in SL for the top band only. It is therefore of interest to repeat the experiment with a different sort rule. It is not clear that the use of 5 bands is near optimal but this number is maintained for convenience and comparison.

Instead of:

> **5 bands divided into equal number of parts ranked by cost**

it is proposed to repeat the model with:

> **5 bands with equal total cost**

The weakness with this formulation is determining a suitable basis for total cost: using a quantity of 1 for each part is arbitrary and meaningless. It is therefore proposed to use the results from model 1, the Poisson distribution, since this will give appropriate rates of demand to each part. The inventory set could then be divided in one of two ways:

> **allocate one fifth of the total to groups of parts in ascending order of extended value (cost x quantity) as this reflects each part's importance in the solution;**

or

> **sort the parts by cost (for a quantity of one) and allocate them into bands according to their extended cost (cost x quantity): this will rank parts in order of cost but will size the bands by portion of total cost.**

The second rule is chosen since it ranks parts by their cost and the model will change the quantities. Thus a total cost of $15M is divided into 5 bands of $3M. Allocating parts by total cost, with the parts sorted by unit cost, the resulting bands contain 190, 41, 24, 8 and 8 parts respectively in order of ascending part cost.

This improved model allows for greater reduction of the highest-cost parts while maintain overall SL – it is possible to reduce the target SL of the band accounting for the top one-fifth of cost down to 0, with 95% global SL achieved at a lower cost than the first model. Note that the model is set up to always allocated a minimum of one of each part, so setting the top band to 0 will give a quantity of 1 for each part. This operational condition ensures that there will be no short-term stock-outs. The SL achieved for the top band will be greater than 0: it transpires that the SL achieved for the top band, with a minimum quantity of 1, averages 59%.

This model is applied to the data set and the results are shown in Table 7.7 below.

	SL %	*Total cost $M*	*Count*	*Average part value $*
Case 1	95	11.0	1026	10722
Case 2	95	11.0	1026	10723
Case 3	95	9.7	896	10894
Case 4	95	19.5	1749	11174
Case 5	95	16.4	1467	11194

Table 7.7: improved Model 3 results

Table 7.7 shows the overall cost, count and average part cost for the improved version of the Cost-Wise Skewed Holding model. The set SL values used to reach these solutions are shown in Table 33 below.

Table 7.8 shows the SL values giving good results from the improved model: these are trial-and-error values so cannot therefore be proven to be optimal.

SL %	Band 1	Band 2	Band 3	Band 4	Band 5
Case 1	97	95	80	50	0
Case 2	97	95	80	50	0
Case 3	97	95	80	50	0
Case 4	97	97	90	50	0
Case 5	97	95	90	50	0

Table 7.8: SL parameters used in the improved Cost-Wise Skewed Holding Model

Figure 7.4 shows SL plotted for parts arranged in order of increasing value: the SL drops off rapidly for the top cost bands, which appears to give better overall results than the original model with equally-sized groups of parts.

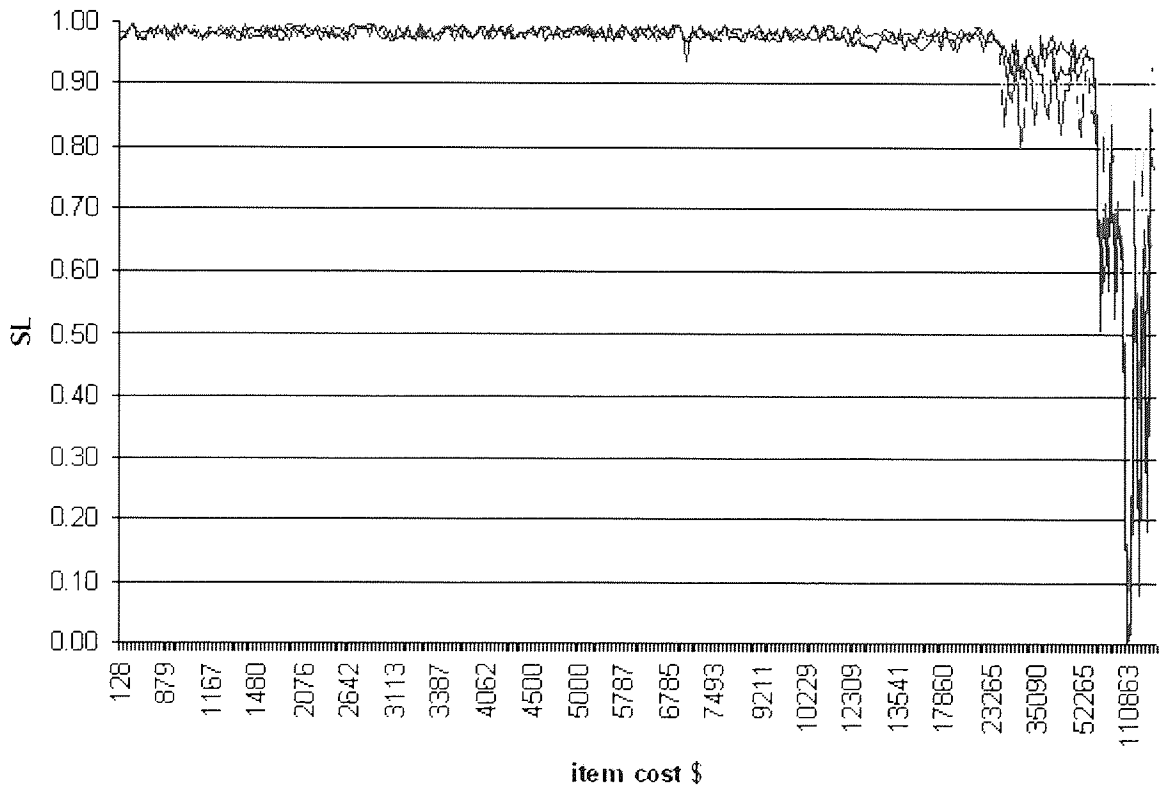

Figure 7.4: SL for parts with increasing value, improved version of Model 3

A limitation with this model is its handling of essentialities: since it groups all parts together, target SLs are applied the same way in each band. While SL scaling is used, its effect is less than the result of processing each group of parts separately for different essentialities. Thus the Cost-Wise Skewed Holding Model with all bands set to 95% SL will give a more costly solution than the Poisson Model, which maintains correct SLs for each essentialiity group.

7.2.4 Model 4: Linear Programming – combined

	SL %	*Total cost $M*	*Count*	*Average part value $*
Case 1	94	10.3	1019	10106
Case 2	93	9.8	972	10035
Case 3	94	8.6	882	9807
Case 4	95	17.5	1785	9786
Case 5	95	14.6	1513	9665

Table 7.8: Model 4 results

Table 7.8 gives the overall results for Model 4, showing decreasing average part value with increasing case size (utilisation).

The LP-combined Model uses SL scaling to group parts with a range of essentialities into a problem that seeks a single global SL as its goal. In each case, the global SL is 95% and the solution achieves a lower value in some cases, due to the effect of SL scaling, which reduces demand for lower-essentiality parts.

A major risk with LP models is their sensitivity and potential instability (depending on the structure of the formulation): any slight error in data entry, calculations or formulation will render the solution either infeasible or wrong, possibly giving a sub-optimal feasible solution. It is therefore especially important to check the consistency and accuracy of LP solutions as comprehensively as possible, since they can choose any feasible permutation of outcomes and thus the constraint formulation needs to be designed with care. The range of values obtained for the 5 cases is consistent with expectations: the total cost values are in line with the changing conditions from one case to the next. As discussed in the previous chapter, a wide-ranging sensitivity analysis was performed on all LP formulations, giving consistent results.

Finally, using the recommended quantities to determine SL, fill rates and total cost in a separate calculation independently validates the LP output.

One concern arising from the results for the LP Model is the fact that SL achieved for Cases 4 and 5 is 95%, while it is lower for the first 3 cases. This can be attributed to the SL scaling values being smaller for larger part quantities, due to the differences between probabilities becoming smaller for the different SLs as quantities increase. Thus the effect of SL scaling is lessened to the extent that SL scaling gives no overall reduction in the overall cost, whatever about the individual part quantities.

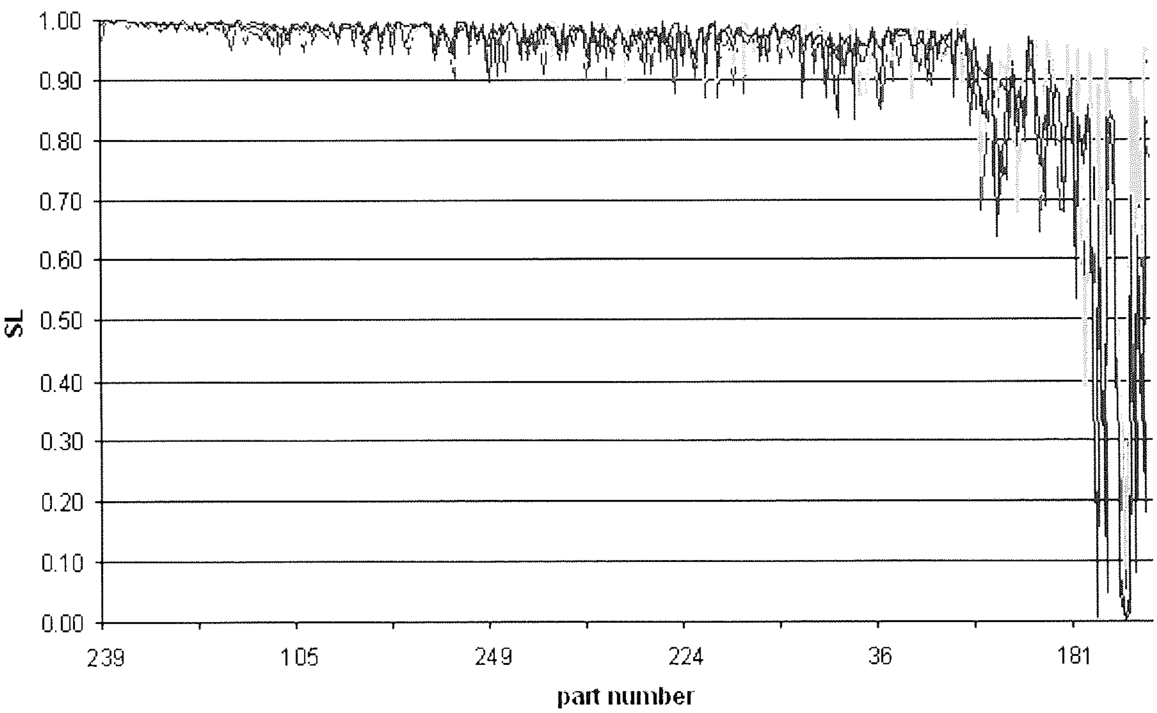

Figure 7.5: SL plotted against part numbers sorted by cost, Model 4

Figure 7.5 shows a sharp drop in SL for the more expensive parts, which confirms the model's objective. Beyond the overall trend, it is not meaningful to distinguish between models since there is granularity for the more expensive parts, i.e., the quantities for small parts are low and rounded up, so there are large variations in derived SLs.

7.2.5 Model 5: Linear Programming – split

	SL %	*Total cost $M*	*Count*	*Average part value $*
Case 1	94	9.7	1006	9662
Case 2	91	8.7	932	9381
Case 3	94	8.3	854	9709
Case 4	94	16.4	1734	9469
Case 5	91	12.2	1381	8858

Table 7.9: Model 5 results

Table 7.9 shows the overall results for model 5: note the falling average part costs are significantly lower than for model 4, indicating a far more efficient solution.

Model 5 is the most theoretical approach, solving a separate LP formulation for each SL value. The SL achieved corresponds with the ranges of SL values assigned, with consequently lower overall values in cases 2 and 5. The benefits of a wider spread of SLs in Case 2 and TAT reduction in Case 3 are evident in significantly lower total cost results. Note in particular the reduced part count in Case 3 resulting from TAT reduction.

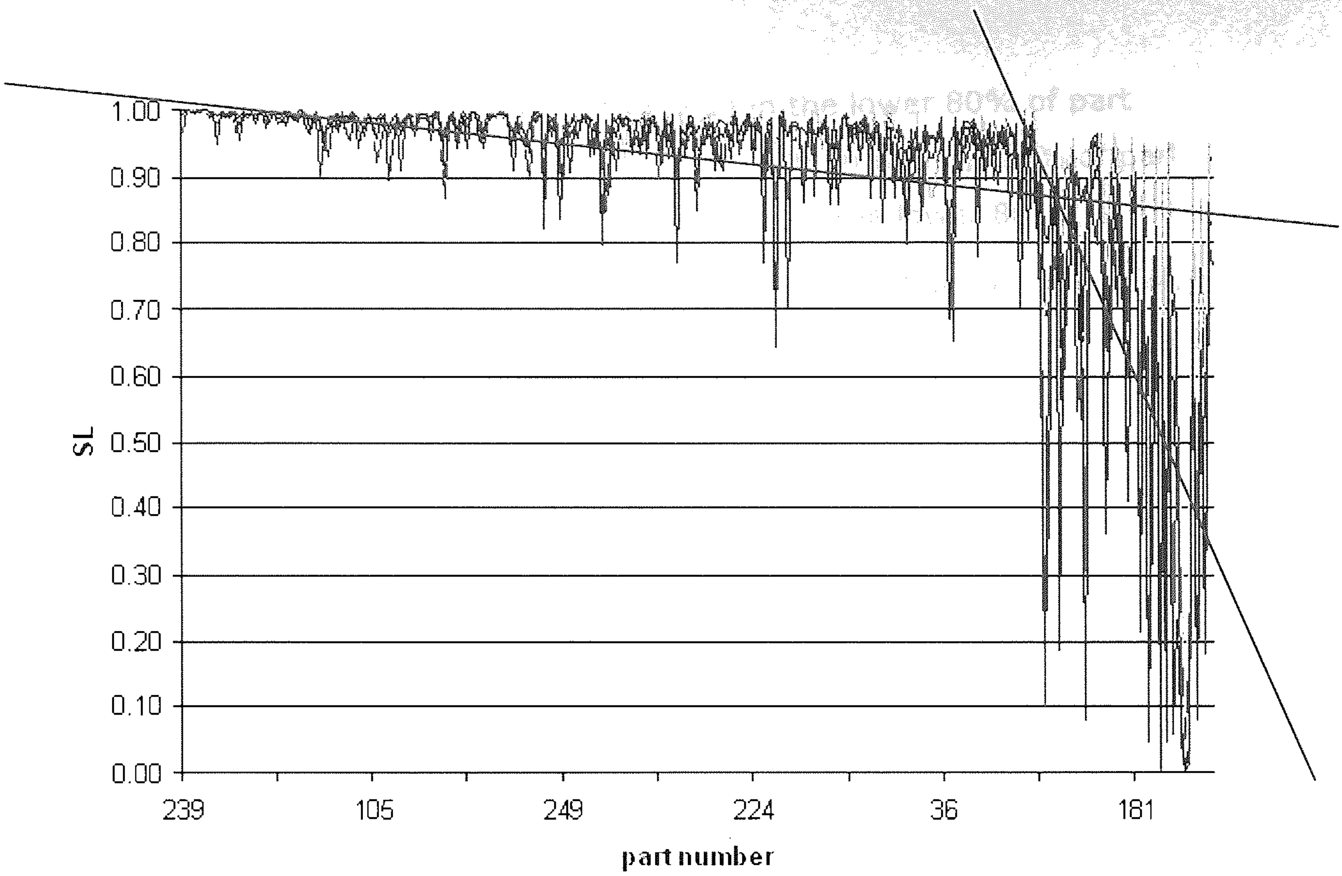

Figure 7.6: SL plotted against part numbers sorted by cost, Model 5

Figure 7.6 shows SL for each part in ascending order of part cost: there is a very pronounced drop in SL as part costs increase. While this effect can be seen across the graph (shown by the pink line in figure 5), it becomes far greater in the top two divisions on the x-axis, or the top one-fifth of the inventory set (shown by the red line).

The slopes of the lines above can be calculated as the change in SL over the lower 80% and top 20% of the inventory set sorted by cost.

Taking moving averages of 10 parts, SL for the parts with lowest cost is 0.9945, ranging up to 0.9482 for the part with cost rank 215 (= the number of parts, 271 * 80% - 1). The slope is then (0.9945 - 0.9482) / 0.8 = 0.058.

The slope for the top 20% of parts is calculated in the same way with a result of (0.9482 - 0.4917) / 0.2 = 2.28.

The sum of all part costs (with a quantity of 1) in the lower 80% of part numbers is 1212898, while the sum of all part costs in the top 20% of part numbers is 2944559, or 2.42 the combined value of the lower 80%.

The extended cost of parts in the lower 80% is calculated as the sum of each part cost multiplied by the recommended holding. This is performed for each of the 5 cases for LP3. The same calculations are performed for the extended costs of parts in the top 20% for each case, and the results compared.

Total $	*Case 1*	*Case 2*	*Case 3*	*Case 4*	*Case 5*
Lower 80%	4906770	4565142	4092968	8600505	6857904
Top 20%	4813250	4178305	4198235	7819481	5375006
Top 20%, ratio	50%	48%	51%	48%	44%

Table 7.10: division of total cost between lower 80% and upper 20% of parts ranked by value

Table 7.10 shows how the total recommended inventory investment is split between the most expensive 20% of the inventory set and the remainder. It is interesting to note that, even with declining SLs, the top 20% still accounts for half of the inventory cost. Clearly, the cases with wider ranges of SL for non-essential parts (the Airbus recommendation), namely cases 2 and 5, show a further reduction in the allocation of more expensive parts.

7.3 Comparison of results from the different Models

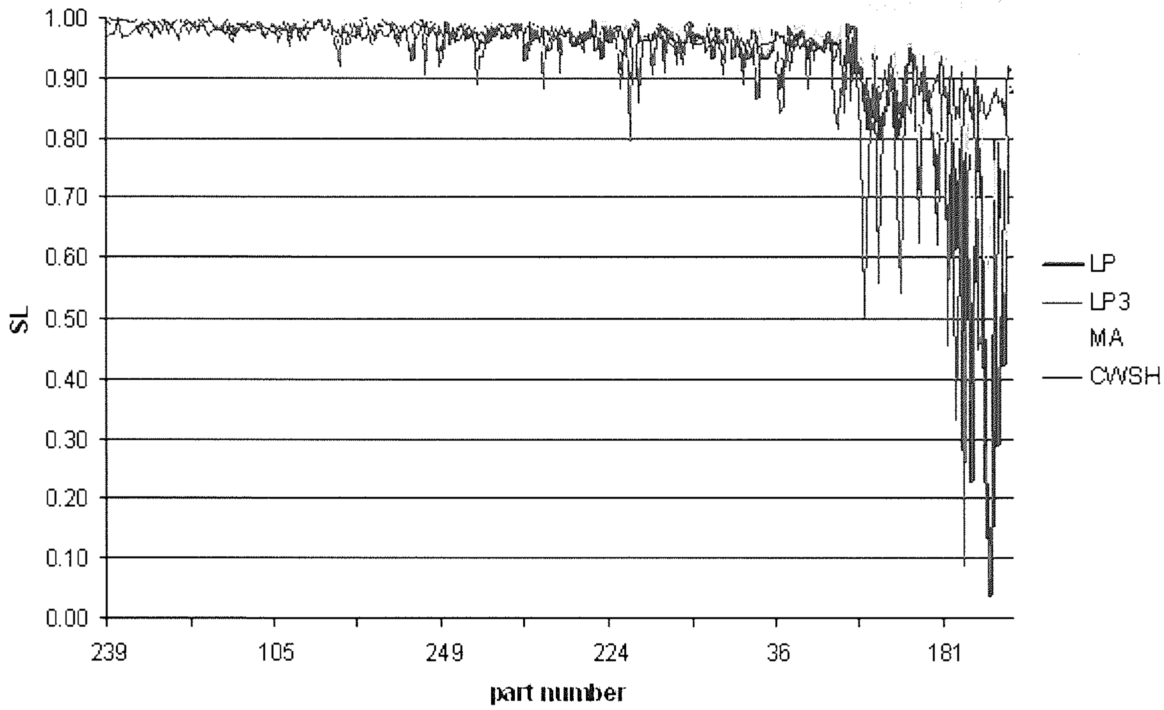

Figure 7.7: SL plotted against part numbers sorted by cost, averages for Models 2,3 4 and 5

In Figure 7.7, each line represents the average SL value for each part from all 5 cases for Models 2, 3, 4 and 5. Model 1 is not plotted since it does not cater for cost, so SL values will be scattered at random about the mean.

Figure 7.7 illustrates graphically the extent to which each model introduces bias to recommended inventory levels in favour of less expensive parts. This was the objective of the process of evaluating models at the outset, namely looking for the inventory holding that would meet SL requirements for the minimum expenditure.

Looking at the curves in Figure 7.7, the most pronounced drop scanning to the right (in the direction of increasing part cost) is clearly for LP3 (Model 5), making it the solution most biased by cost. Next in order is LP (Model 4), the combined LP Model, which is not as biased as Model 5 but better than the others.

Model 2, marginal analysis, shows a strong cost bias, but less than the LP solutions.

Model 3, Cost-Wise Skewed Holding, shows a strong cost bias but less than the LP solutions. The improved version of Model 3 is not depicted in Figure 7.7 for clarity: this improved version gives better results than the first version of the Cost-Wise Skewed Holding Model. The improved version of Model 3 performs with roughly the same results as Model 2, marginal analysis, but does much better with the larger cases (4 and 5). Model 3, as expected, shows an average drop in SL for the right-most one-fifth band of data, without further discrimination within the band. Thus for Model 3, the cheapest part in the top band is treated the same as the most expensive.

Figure 7.8 shows conclusively that Model 5 is clearly better than the other solutions, which supports total cost and average cost data. Overall, given the large advantage that the LP-split Model has over the other models, it makes sense to discount the other models as being sub-optimal. This is theoretically sound since the Model 5 selects the best solution set from each of three well-defined and constrained problem formulations. Given the clear difference, there is limited merit in performing extensive comparisons of the other techniques against each other.

LP3 is a clear winner in terms of overall cost, showing that it produces the most optimum set of results, while LP is an approximation of this set of results. Further, LP3 maintains a lower overall SL than does LP for the same data as its formulation allows different Ess Code models to be fully solved independently.

Apart from the LP solutions, the next best model overall appears to be Model 3, Cost-Wise Skewed Holding in its improved version, where total cost is spread equal over a small number of bands. It is not an optimisation, but based on the present results, appears to perform nearly as well as marginal analysis for base cases, and performs far better for the cases with higher levels of utilisation. While the improved Model 3 is still inferior to the LP solutions, it gives useful results with minimal processing.

As well as looking at total cost, it is interesting to consider the average part value for each solution: if a solution recommends a larger total inventory count than another, but does so at a lower average part value, then the total cost may not be higher. Figure 7.8 (also shown in Chapter 6) shows the total cost of each model formulation for each case: there is a clear separation between the solutions, with LP3 giving consistently superior results.

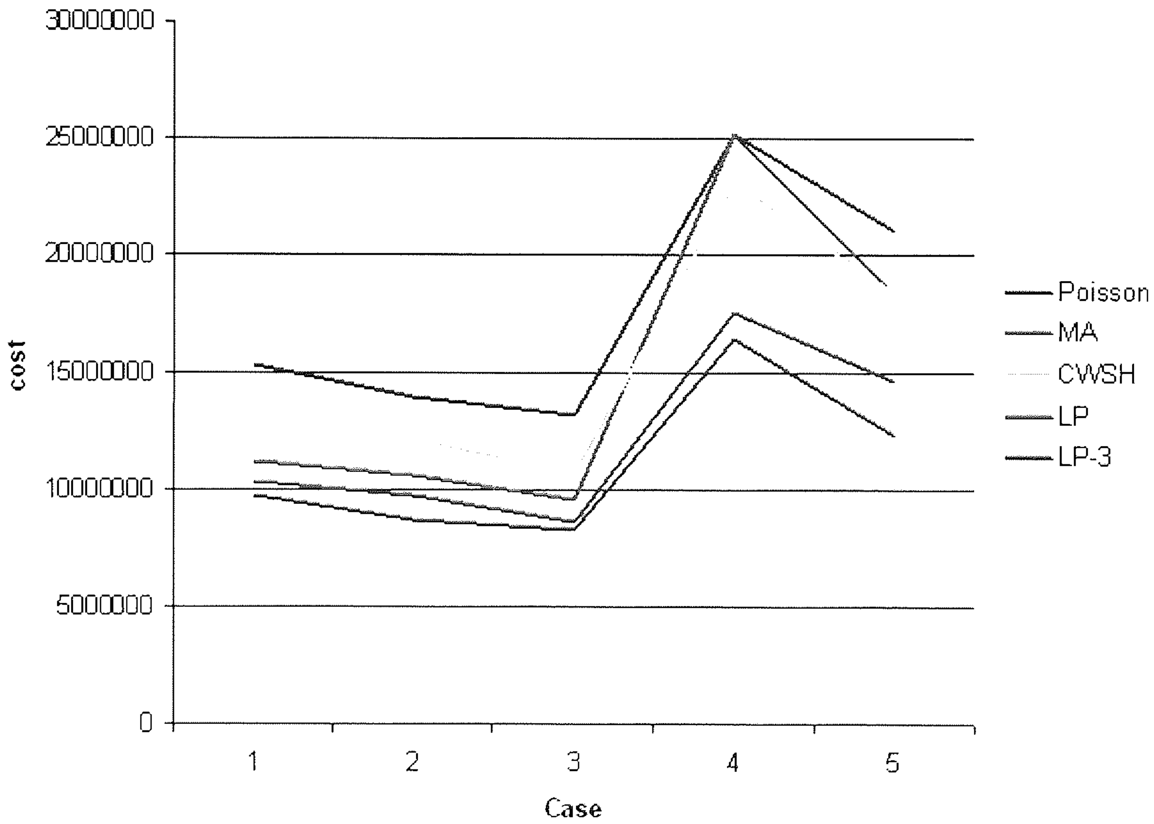

Figure 7.8: total cost ($M) for each Model and each case

The following observations can be made from Figure 7.8:

- Poisson is the worst (most expensive) model in all cases.
- CWSH gives a good improvement over Poisson. However, the margin is less for Case 2 (fewer), since CWSH uses its own SL values for

calculation purposes, only using the input SL values for scaling, so the effect of lower SLs is largely lost. Also the margin between Poisson and CWSH is less for Cases 4 (bigger) and 5 (best), the cases with higher utilisation. It can be inferred that the effect of granularity (rounding up to the nearest integer for every line) is less for a higher global rate of demand.

- MA is significantly better than Poisson and CWSH for Cases 1 to 3, it is the joint worst for Case 4 and improves slightly for Case 5. As discussed earlier, the order in which MA selects parts quantities leads to over-provisioning, and this effect turns out to get worse with increasing volume.
- LP is far better than the preceding three models in all cases, with the gap increasing for the larger problem spaces, supporting the concept that this model chooses the best outcome from a broad range of possibilities, which MA does less well as the number of possibilities increases.
- LP3 is the best model in all cases, generally following the same trend as LP but giving a lower overall cost since it achieves its three separate SL targets to give a lower global SL. The biggest gap between LP and LP3 is for case 5 (best), which has the highest level of utilisation and the wider range of SL values, which it handles more efficiently than the LP Model, with the result that SL for LP3 is 91% for Case 5, while LP has a SL of 95% for the same case.

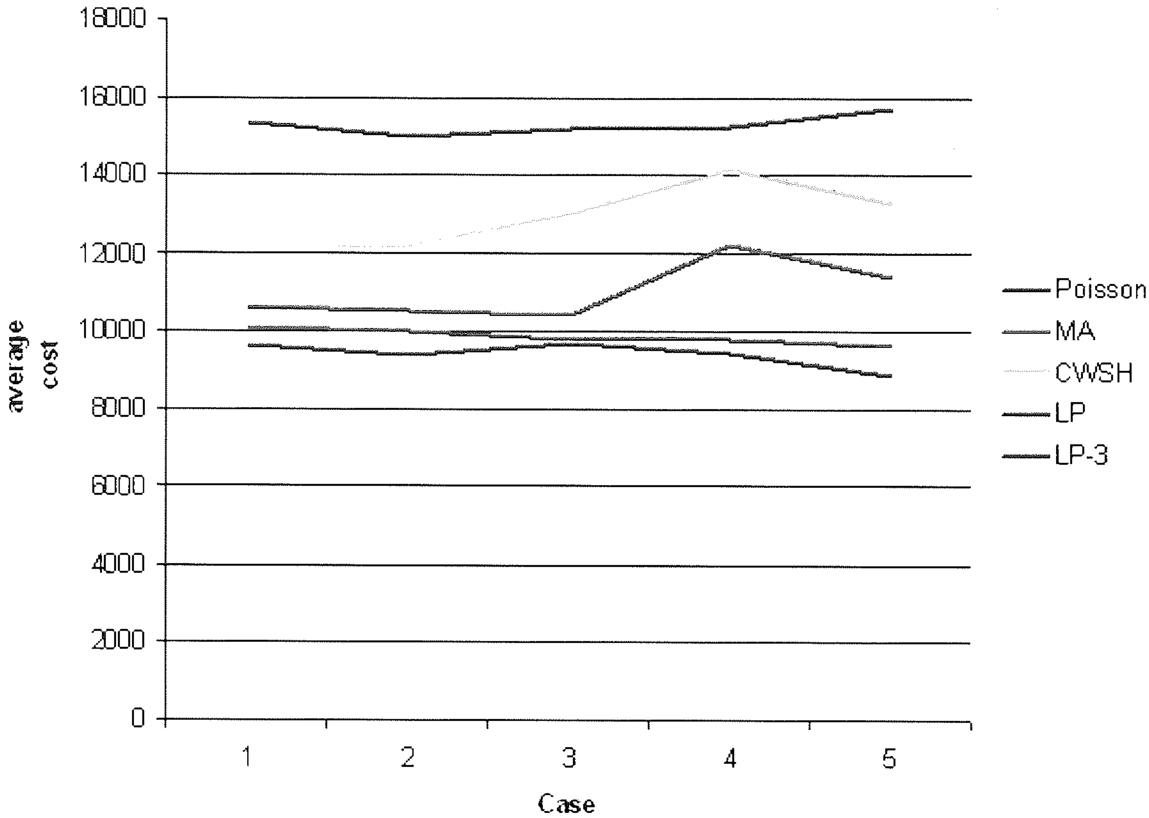

Figure 7.9: average part cost for each Model and each case

Figure 7.9 shows the average part cost by each model for each case. It shows a consistency across all of the cases since no lines cross. The solution with the lowest average part cost can be considered the best, and the graph is consistent with other findings, namely that the result rank LP3 the best with Poisson the worst outcome. This average individual part cost assessment removes the variability of different models attaining different SLs (total cost does not make this distinction).

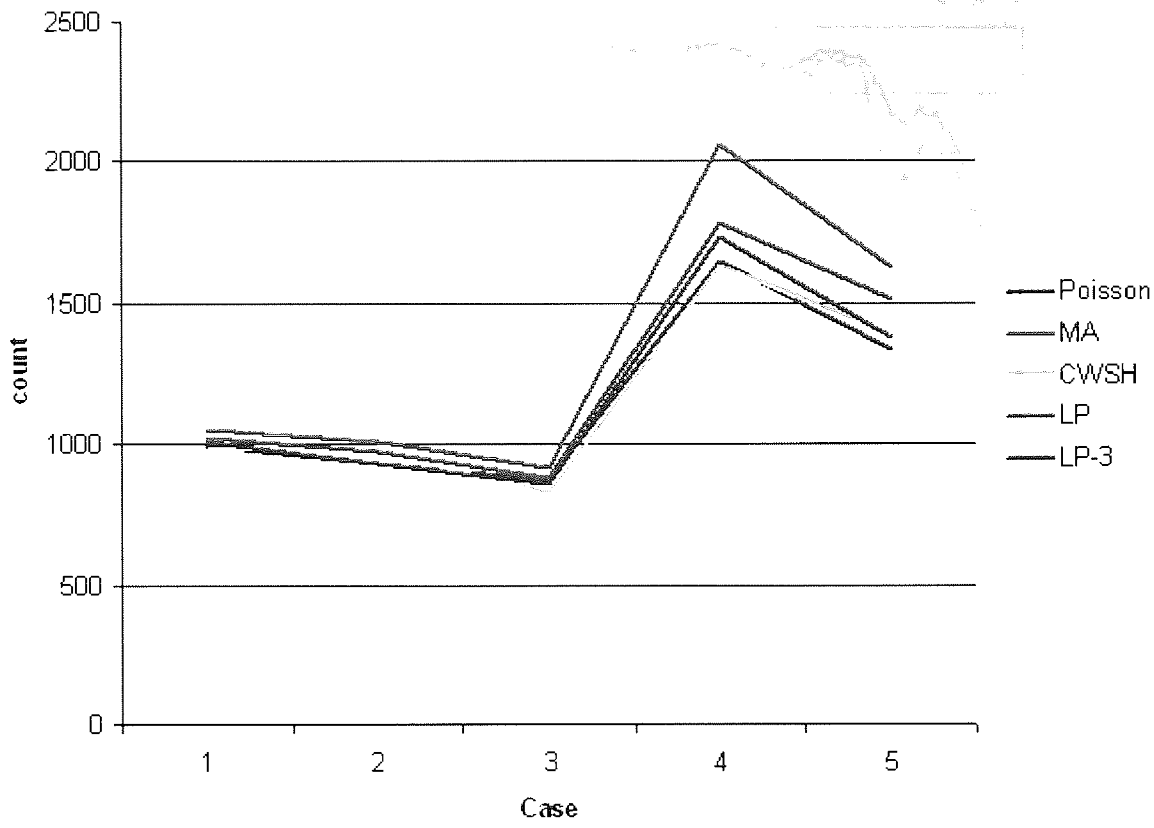

Figure 7.10: total inventory count for each Model and each case

Figure 7.10 shows the total number of parts prescribed by each model for each case: while MA runs higher than the rest in all cases, the other models are close in their number of parts. What is particularly interesting is that the Poisson Model, the least successful, has low numbers of parts in all cases, and the lowest in Case 5 (best). What this emphasises is that choosing the right parts is more important than choosing the fewest.

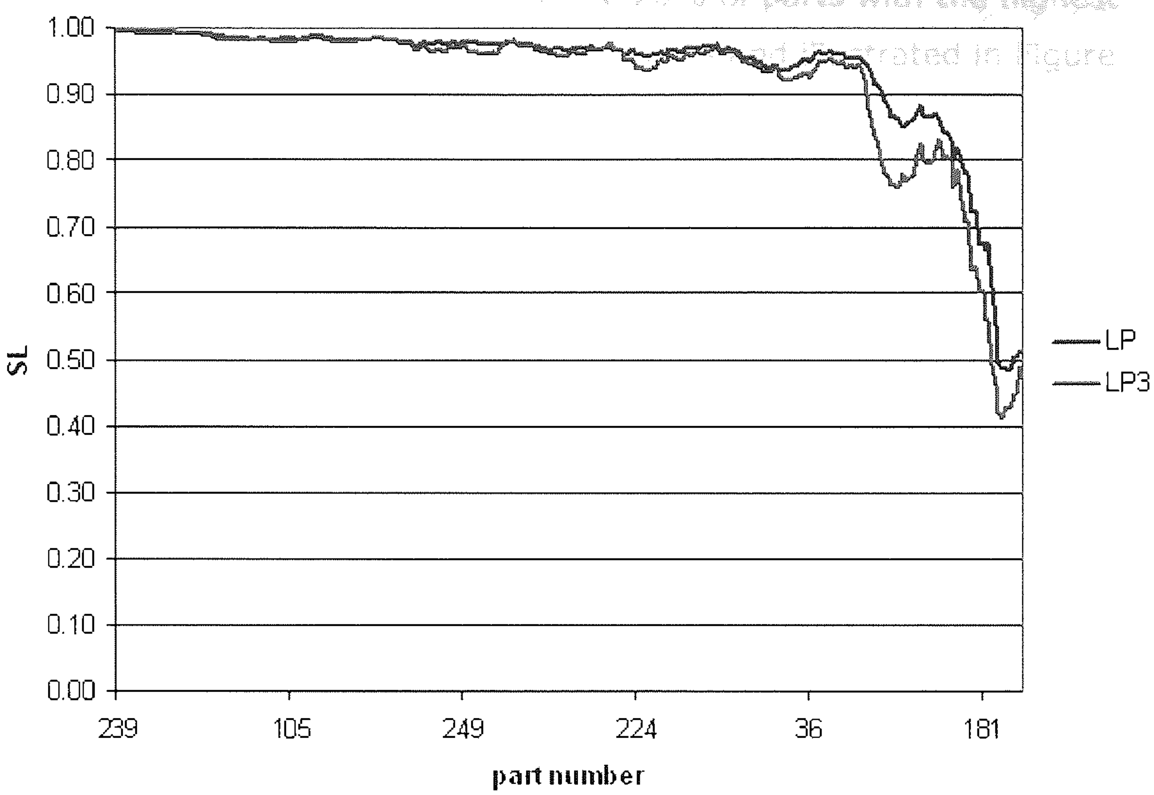

Figure 7.11: 10-part moving average SL against part cost for Models 4 and 5

Figure 7.11 shows 10-part moving average SL values against increasing part cost for the two LP models, clearly illustrating the sharp drop in performance for the most expensive parts. Also, LP3 (Model 5) is seen to consistently out-perform LP (Model 4) in reducing the inclusion of more expensive parts. A question arises from the graph: given the logarithmic trend of the curve, would there be a better solution that would give a linear curve, i.e., a constant declining slope from left to right? This is not evaluated here: it could perhaps be approximated by using the square root of the part cost to reduce the level of priority given to the higher parts. However, since the LP Model performs optimisation on the stated costs it can be assumed that it makes an optimal selection. Also, SL is based on a cumulative probability, so the number of parts will increase exponentially as SL approaches 1, so the behaviour of the problem is non-linear anyway.

In the review of results from Model 5 earlier, the proportion of expenditure allocated to the most expensive parts was considered: the measure derived was the portion of budget allocated to the 20% of parts with the highest value. This evaluation is repeated for all models and illustrated in Figure 7.12.

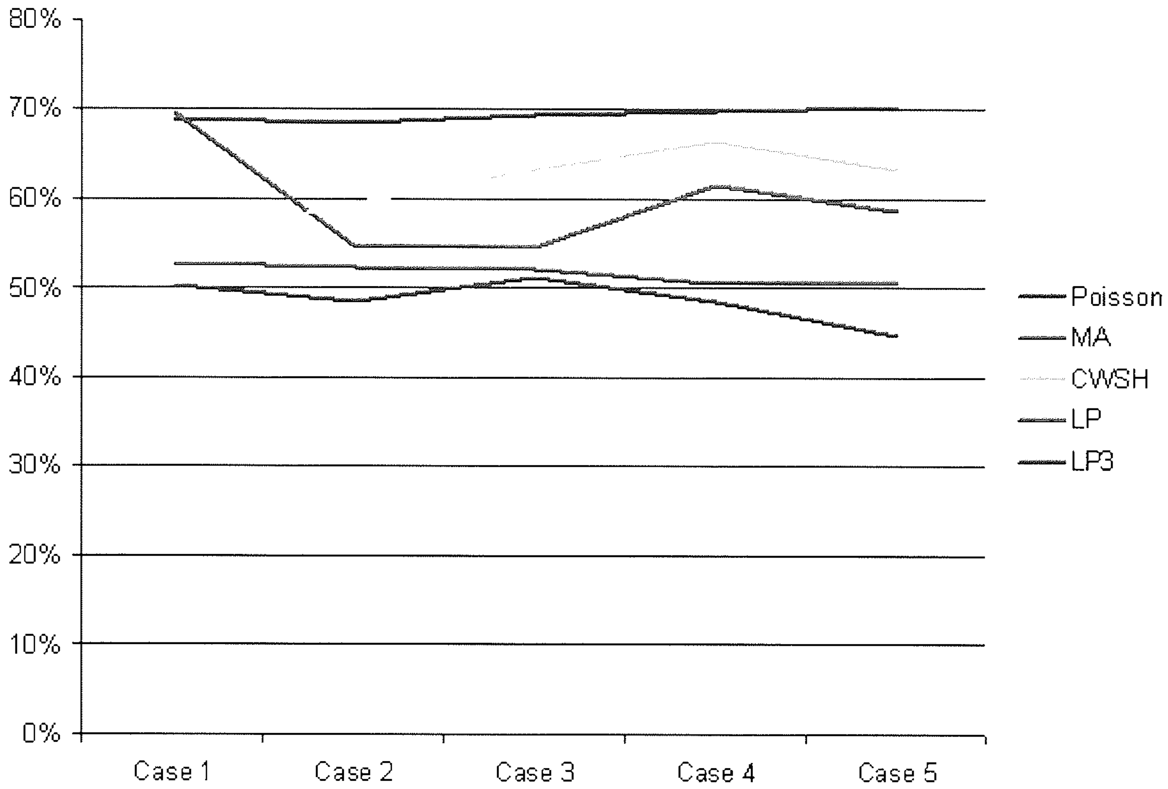

Figure 7.12: proportion of budget allocated to the most expensive 20% of parts

Figure 7.12 shows the proportion of the total cost for each solution allocated to the most expensive one-fifth of the inventory group. The baseline is the Poisson Model, which takes no account of cost and consistently allocates 70% of budget to the top 20%. The other models all improve on Model 1 with the two linear programming models showing the best and most consistent outcomes. What this graph represents is the tendency of each model to avoid the most expensive parts: thus the less spent on expensive parts, the more the model fills demand with cheaper parts. Model 5 (LP3) gives the best performance in this respect and shows an improving trend as the problem size grows (for Cases 4 and 5).

Another measure of performance is the cost slope derived for Model 5 earlier. This is the measure of the rate at which SL falls with increasing part cost. This is not measured for Model 1 since cost is not assessed, so the slope will tend to be 0 (flat). For Model 3, Cost-Wise Skewed Holding, the slope will tend to 0 within each of the 5 SL bands since there is no ranking for cost within each band. Slope values are given in Table 7.11 for the remaining Models, 2, 4 and 5.

	Model 2	Model 4	Model 5
SL1	0.996	0.9955	0.9945
SL2	0.9733	0.9607	0.9482
Slope 1 = SL1 - SL2 / 0.8	0.0284	0.0435	0.0579
SL3	0.9721	0.9617	0.9482
SL4	0.6903	0.5146	0.4917
Slope 2 = SL3 - SL4 / 0.2	1.41	2.24	2.28
Slope 2 / slope 1	49.6	51.5	39.4

Table 7.11: SL / part cost slopes for Models 2, 4 and 5

The slopes for each of these Models are plotted in Figure 7.13.

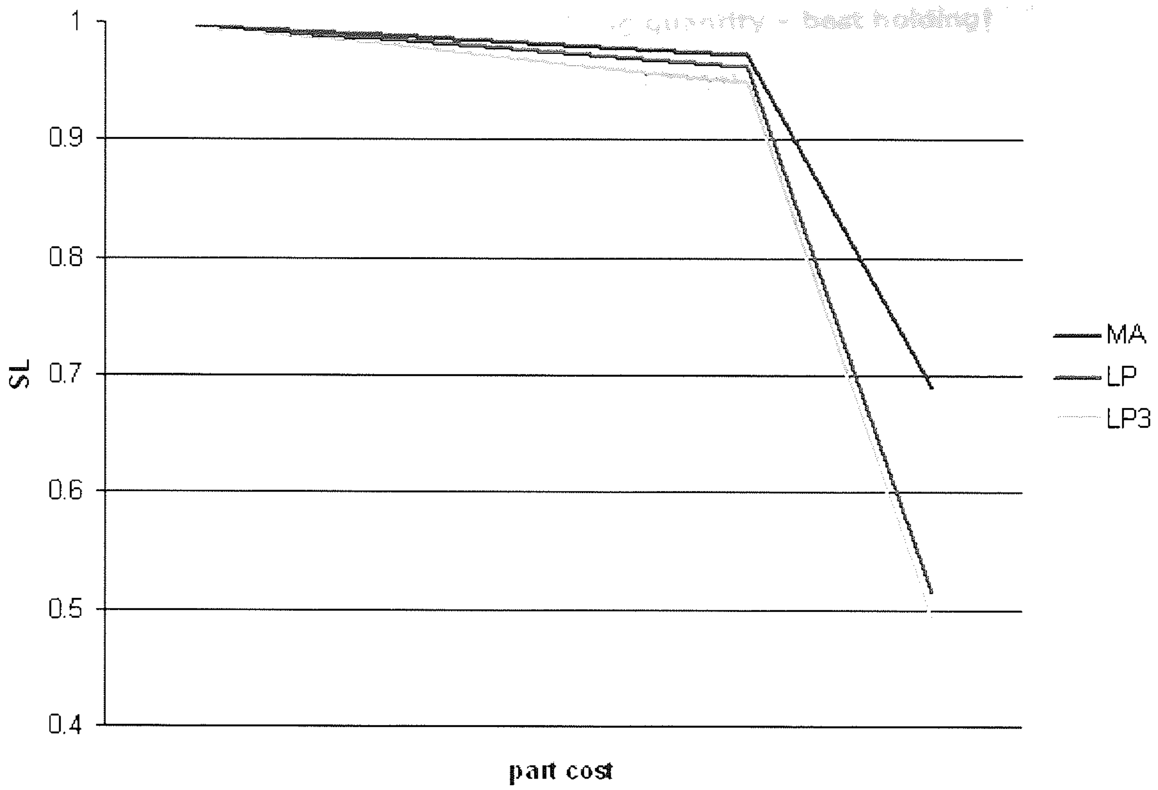

Figure 7.13: slope of declining SL with increasing part cost for Models 2, 4 and 5

Figure 7.13 shows the slope of declining SL as part cost increases for the MA, LP and LP3 Models, with slopes evaluated either a break point at 80% of the inventory set ranked by cost. Thus the slope to the left of the figure shows the drop in SL with increasing cost for the lower 80% of parts ranked by cost, with the right-hand side slope reflecting a far steeper rate of decline for the most expensive 20% of the parts. LP and LP3 have similar slopes, both being substantially steeper than MA, with LP3 giving the greatest decline over all.

Two further measures were presented in Chapter 6: Match with LP3, count and Match with LP3, cost. The first of these is a measure of the degree to which a solution agrees with the best solution in terms of the recommended holding of each line item:

1 – (Σ|recommended holding quantity – best holding|)

/ Σ (best holding))

The results of this analysis are shown in Table 7.12 below.

	Stk	*Poisson*	*MA*	*CWSH*	*LP*
Case 1	34%	76%	96%	84%	97%
Case 2		76%	92%	83%	94%
Case 3		76%	93%	82%	96%
Case 4		76%	84%	79%	97%
Case 5		73%	85%	79%	91%
average	34%	75%	90%	81%	95%

Table 7.12: holding match with LP3

Table 7.12 shows how well each solution aligns with the optimal holding in terms of part quantities (1 = complete match). Stk refers to the actual holding reported in the data set, which is a holding with twice as many parts as the optimal solution and so shows a large discrepancy. The four models compared against LP3 show results that agree with other findings: models 1 to 4 give improving results in that order.

Only Case 1 is compared with the actual holding in Table 7.12, since Case 1 is the actual operating case and it is not meaningful to compare with the others, so they are blank in Table 7.12.

The second matching metric looks at budget allocation, comparing each solution against the optimal. By factoring cost into the previous measure above and below the line, the assessment now gauges how closely the investment recommendation from each solution corresponds to the optimal:

$$1-\left(\Sigma\left|(\text{recommended holding quantity} - \text{best holding})\right| \times \text{cost}\right) / \Sigma\,(\text{best holding} \times \text{cost})$$

The results of this analysis are shown in Table 7.13.

	Stk	Poisson	MA	CWSH	LP
Case 1	-153%	33%	86%	65%	93%
Case 2		33%	79%	58%	88%
Case 3		36%	86%	63%	96%
Case 4		32%	47%	45%	93%
Case 5		29%	58%	49%	84%
average	-153%	33%	71%	56%	91%

Table 7.13: budget match with LP3

It can be seen from Table 7.13 that, when cost is taken into account, the discrepancies from the optimal solution are very large (1 = complete match). Thus Model 4 gives a satisfactory solution, Model 2 is a distant third place and the others should not be considered acceptable in comparison. Note the deterioration in Model 2 (MA) for the larger Cases 4 and 5. What Table 7.13 illustrates is, for each unit of budget allocated, the degree to which the allocation agrees with the best solution.

In Table 7.13, only Case 1 (base) is compared with actual stock levels since the comparison with other cases is not meaningful, so the values showing comparison of actual stock against Cases 2 to 5 are blank.

Transposed LP Models

The LP problems were re-formulated with a set cost, with maximum fill rate, leading to maximum SL as the objective. It is interesting to consider the marginal gain in SL for increasing budget: what is the point at which diminishing returns indicate that the budget should be constrained?

Table 7.14 shows the SL for each incremental $1M allocated to budget, the cost constraint in the transposed LP3 (LP split) Model. data presented in

Table 7.15 shows the change in SL, scaled in proportion to the change in budget.

Budget $M	Case 1	Case 2	Case 3	Case 4	Case 5
5	61%	58%	68%	41%	46%
6	75%	74%	81%	55%	60%
7	83%	83%	89%	63%	69%
8	89%	88%	93%	70%	76%
9	92%	92%	95%	75%	81%
10	94%	94%	96%	80%	85%
11	95%	95%	97%	83%	88%
12	97%	96%	98%	86%	90%
13	97%	97%	99%	89%	92%
14	98%	98%	99%	91%	93%
15	98%	98%	99%	92%	95%
16	99%	99%	100%	93%	95%
17	99%	99%	100%	94%	96%
18	99%	99%	100%	95%	97%
19	99%	99%	100%	96%	97%
20	99%	99%	100%	96%	98%
21	100%	99%	100%	97%	98%

Table 7.14: SL values for increasing budget for 5 cases, transposed LP3 Model

Table 7.14 shows the SL value for each case with budget ranging from $5M to $21M – remember that cases 4 and 5 have twice the fleet size or utilisation of cases 1, 2 and 3, so the comparison is not a direct one. The colour coding in Table 15 represents:

Yellow – gain in SL of 2% per $1M increase in budget

Blue – gain in SL of between 1.5% and 2% per $1M increase in budget

Red – gain in SL of between 1% and 1.5% per $1M increase in budget

Thus if the criterion applied is that each increase of $1M must deliver an increase of 2% in SL, then the appropriate budget for case 1 is $9M, $10M and so on. Whether the corresponding SL values of 92%, 94% and so on are appropriate is a further decision to be made based on judgment of operational needs against financial limitations.

Table 7.14 is not proportional with cost: each \$1M increment is a smaller percentage of the budget. This is addressed using the data presented in Table 7.15.

Budget \$M	Case 1	Case 2	Case 3	Case 4	Case 5
6	83%	93%	78%	84%	81%
7	60%	61%	52%	60%	67%
8	45%	42%	34%	54%	53%
9	31%	33%	19%	50%	47%
10	20%	21%	14%	46%	37%
11	15%	16%	12%	37%	32%
12	12%	12%	9%	34%	30%
13	12%	11%	7%	33%	22%
14	9%	10%	5%	26%	19%
15	7%	8%	4%	22%	18%
16	6%	6%	2%	19%	15%
17	4%	5%	2%	17%	12%
18	3%	3%	1%	15%	11%
19	2%	3%	1%	13%	10%
20	1%	2%	1%	14%	9%
21	2%	2%	0%	12%	9%

Table 7.15: ΔSL / Δbudget values for increasing budget for 5 cases, transposed LP3 Model

Table 7.15 shows the change in SL divided by the change in budget for each increment in budget and for each case. Referring back to Case 1 in table 40, the change in SL for the increase from \$5M to \$6M is:

ΔSL(75-61) = 14%.

The change in budget for the same data point is:

Δbudget 1 - (5/6) = 17%.

Therefore the change in SL proportional to the change in budget is:

ΔSL / Δbudget = 14% / 17% = 83%.

The colour coding in Table 7.15 corresponds to:

Yellow - gain in ΔSL/Δbudget of 20% per \$1M increase in budget

Blue - gain in ΔSL/Δbudget of between 15% and 20% per \$1M increase in budget

Red - gain in ΔSL/Δbudget of between 10% and 15% per \$1M increase in budget

It can be seen that these values decline in a similar fashion to SL in Table 7.14, however they are better scaled. Also, the difference between Case 4 and Case 5 is greater, showing that the marginal benefit of increasing budget in Case 5 falls away quicker than it does for Case 4. This is illustrated in Figure 7.14 below.

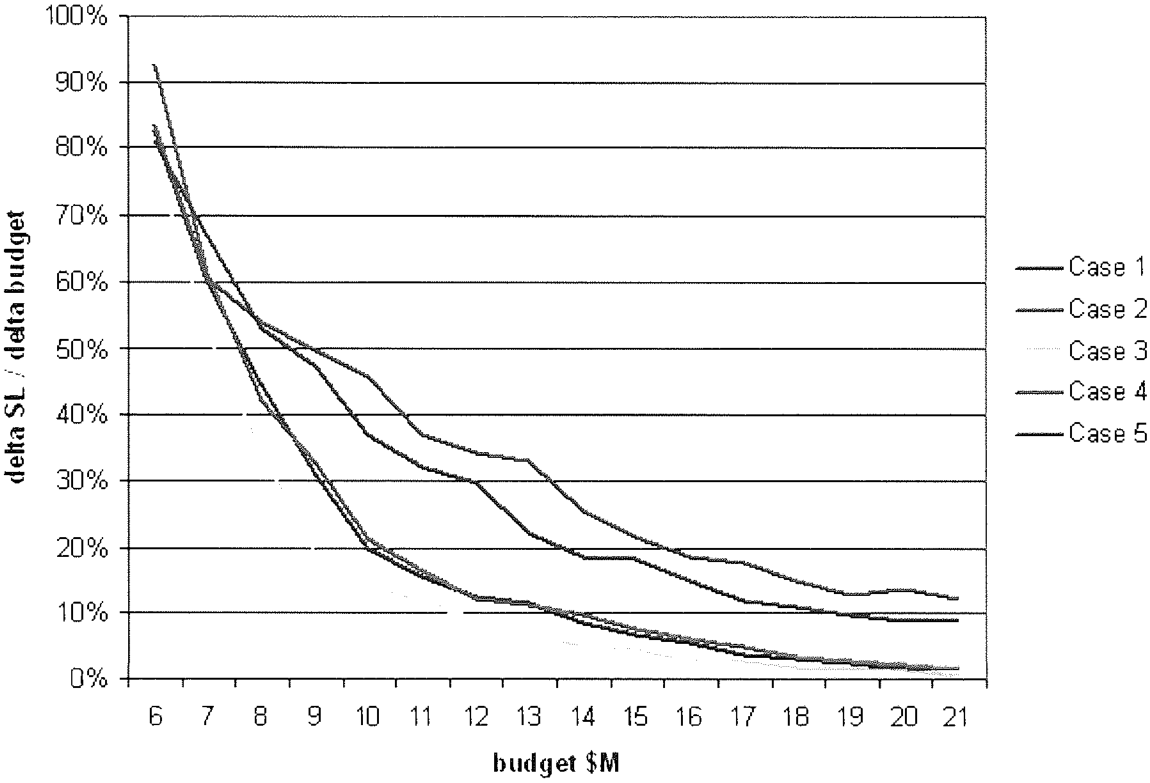

Figure 7.14: ΔSL / Δbudget for increasing budget for 5 cases, transposed LP3 Model

Figure 7.14 shows the diminishing marginal return for each case in the LP3 Model. Cases 1, 2 and 3 have the same utilisation, while Cases 4 and 5 show twice the utilisation of Cases 1, 2 and 3. Cases 1 and 2 show equivalent diminishing returns, while Case 3 falls faster, so can be considered more efficient (this is the case with faster repair times). Case 5 behaves significantly better than Case 4, again reflecting the benefit of shorter turn-around times and lower SLs for non-essential items. Referring to Table 7.15 and Figure 7.14, the colour bands in the table reflect values at 10, 15 and 20% in the figure, so the point at which the marginal return for Case 5 falls below 20% is $13M.

7.4 Review of SL scaling

The aim of SL scaling is to demote the value of parts with essentiality less than 1, i.e., parts with essentiality 2 or 3 ("go-if" or "go"). Weightings were derived with respect to the part mean value of demand and the SL target relative to the SL for Ess Code 1 items. For example, if SL(Ess Code 3) = 75%, SL(Ess Code 1) = 95% and mean demand = 2 then a weighting of 1.69 is applied to the part. This weight is applied either to the demand, so the demand is scaled down, or the cost, so that the cost is scaled up. Thus, depending on the solution method, this part is 1.69 times less attractive to the solution (either in its contribution to meeting demand or its cost) than a part with the same demand and cost but Ess Code 1.

A shortfall of this scaling approach is that the different models optimise in different ways, so it cannot be assumed that the effect of scaling works equally across the models for the purpose of comparison. For instance, Model 3 applies the weighting to cost, then it sorts items into 5 bands with equal numbers of line items, arranged by cost. Thus the weighting may alter the selection of items into bands. In practice this effect was very slight. It is not appropriate to apply weighting to demand in Model 3 since it uses varying SLs to determine global SL.

Model 2 (Marginal Analysis) calculates change in fill rate divided by cost, so it is possible to weight either cost or demand and expect the same outcome.

The first model (Poisson) does not employ scaling since each line item is calculated in isolation, so the three SL values are applied directly according to essentiality.

Model 4 (LP-combined) uses SL scaling on demand values. This model has been run for Case 5 without the scaling process applied, in order to observe the difference. The rationale for choosing Case 5 is that it is the largest problem space with the widest range of SL values, so the effect of scaling should be greatest. The results of this model with and without scaling are shown in Table 7.16 below.

	Global SL	*Total cost*	*Number of parts*
Model 4 with SL scaling	95%	14622747	1513
Model 4 without SL scaling	95%	15092087	1561
Difference		3.2%	3.2%

Table 7.16: Model 4 (LP), Case 5 with and without SL scaling

Model 5 (LP3) does not employ SL scaling since it treats each essentiality level in isolation and so is a theoretically better decision set. With reference to Table 7.9 above, Model 5 obtains a cost of $12.2M and 91% SL, which is an accurate combination of the mixed SL values assigned, and 17% cheaper than the Model 4 solution.

The difference between the Model 4 and Model 5 solutions highlights another inherent weakness in treating the inventory pool as a single set: it is not clear what the target SL should be. If the global SL is set as SL (Ess Code 1) = 95% here, then there is wasteful over-provisioning of the items with lower essentiality.

The motivation for using SL scaling in the first place was to have a combined problem space to afford economy of scale. Also, in one sense it can be argued that there is a relationship in demand between parts of different essentiality, although the counter-argument can also be made. If the requirement of the inventory system is to satisfy an average number of demand events to meet a global SL, then either all parts are treated together, or the problem must be split and the set of solutions reassembled to give a global data set.

Having evaluated both the combined model using SL scaling and the split model, which obviates SL scaling, it is observed that the split model produces far better results and so should be recommended as the better method. The only downside of the split model is the additional processing needed; the lack of economy of scale does not appear to be significant, and the SL scaling method does not give results that are close enough to be acceptable.

In summary, the arguments against using SL scaling are:

1. It is not clear how to derive optimal scaling values;
2. The effect of scaling varies according to the optimisation technique used;
3. The effect of scaling is generally less than expected;
4. Using scaling prevents the setting of accurate target SLs and will therefore over-provide parts with lower SLs than the global SL value.

7.5 Computational intensity

Models 1 (Poisson) and 3 (Cost-wise skewed holding) are computationally straightforward, with a set of probability calculations for each line item.

Model 2 (Marginal Analysis) sorts an i x j list by marginal contribution: in the case of the larger cases resulting in a sort list with 8220 rows.

Models 4 (LP) and 5 (LP3) performed a linear programming solution with a large solution space. The number of iterations taken to solve each case for the two models is shown in Table 7.17 below (shown earlier in Chapter 6).

iterations	*LP*	*LP3*
Case 1	290383	3618
Case 2	232266	3132
Case 3	94790	3432
Case 4	1351	2992
Case 5	2603	3141

Table 7.17: linear program solution iterations for Models 4 and 5

Model 5 (LP3) has a low number of solution iterations in all cases. This model solves three separate formulations, each being smaller than the combined models used in the LP-combined (Model 4) approach. The low number of iterations suggests that the solutions are easily found: increasing target SL from 95% to 99% roughly doubles the number of iterations required. Cases 4 and 5 have far fewer iterations than the smaller formulations, again suggesting that the optimal solution is more easily reached by the solver.

This chapter has assessed and compare the performance of the models implemented to represent current practice and new solutions. Chapter 8, below, considers the implications of these results for theory and practice in the field.

Chapter 8: Discussion

This chapter makes practical recommendations for better inventory management based on the models tested. As well as promoting the best method as observed from analysis of the results, there are some general recommendations, such as assigning cost to repair times and managing expensive inventory interactively.

The findings and recommendations are then compared to the literature to see how well they are supported by current knowledge and whether the new ideas make new contributions to the field.

The limitations of the work are considered, both in light of the internal modelling and analysis and with reference to literature in the field.

Finally, suggestions are put forward for the extension and development of the work performed, identifying potential for greater application in the subject field and also in other fields with similar characteristics.

8.1 Recommendations for improved practice

Based on the input data, its operational context and the findings from the analysis of results, several ideas are put forward that offer potential for improved service and reduced cost in the management of rotable inventory.

$M	Case 1 base	Case 2 fewer	Case 3 faster	Case 4 bigger	Case 5 best
Poisson	15.3	14.0	13.2	25.2	21.0
Marginal Analysis	11.1	10.6	9.5	25.2	18.4
Cost-wise skewed	11.8	11.8	10.6	23.0	18.4
Cost-wise, improved	11.0	11.0	9.7	19.5	16.4
LP	10.3	9.8	8.6	17.5	14.6
LP3	9.7	8.7	8.3	16.4	12.2

Table 8.1: summary of results – total cost

Table 8.1 shows the total cost of the inventory holding recommended by five different models (there are two versions of the Cost-wise skewed model) in each of five different cases. The first two models are in general use Poisson (as recommended by the OEMs and used generally by airlines and MRO service providers) and Marginal Analysis (developed by Logistechs and now owned by General Electric). The other models are developed and tested here, and fall into two types: Cost-wise skewed (a simple partitioned heuristic) and Linear Programming (with two versions).

Table 8.1 shows that, for the same objective, the new models offer large cost savings over known practice. Looking at the results for different cases, it can be seen that there is consistent behaviour by the models, so that in general, the best solution (SL target met at least cost) for one case tends to be the best for all.

Some recommendations can be drawn from these results and observation of the behaviour of the models, as discussed below.

8.1.1 Optimise inventory using the best method

Given the substantial benefits offered by the LP solution, it is well worthwhile preparing data and processing inventory decisions using this technique. The gains over current practice (Poisson and Marginal Analysis) are significant and consistent. The split LP solution is significantly and consistently better than the combined model with service level scaling.

The LP solution presented here has been tested successfully in a commercial environment (Armac 2008) with part lists of several thousand parts, although data describing the performance of the solution has not been made available. SR Technics holds 22,000 part numbers in the UK, although not all are rotables and there are several aircraft types involved, so a typical problem formulation will not have more than several thousand parts (Armac 2007).

8.1.2 Perform frequent reviews of operational data and re-run optimisation

Once the processes are in place to run automated opimisation models, it is worth reviewing changing operational data on a regular schedule and re-evaluating optimisation decisions. A consequence of performing fleet-wide optimisation, and seeking the optimal recommendation across the whole inventory system, is that any small changes to operational inputs may yield a quite different optimal inventory set.

8.1.3 Assign opportunity cost to repair times

The effect of changing TAT has been shown here and can be measured as around 15% of the total inventory value (by the best method) for a 5-day reduction. It is therefore advantageous to pay more for a faster repair (or less for a slower repair): as well as using this effect in negotiating pricing with repair vendors, this cost saving should be fully counted in examining decisions to develop in-house repairs.

8.1.4 Treat groups of parts with different essentiality codes as separate inventory pools

While it is usually accepted that treating all parts together creates a larger inventory system, which should therefore enjoy economy of scale and be more efficient, what the modelling results show here is that it is more cost-effective to split the inventory by essentiality code and optimise each group separately.

8.1.5 Reduce costs further by "cheating" the optimal holding recommendation

Taking say the top 20% by value of the inventory holding, it is worth manually intervening in the supply of stock in reaction to demand events. In other words, if it is possible to quickly source spares externally, it makes sense to under-provide the most expensive parts, only taking action when a demand event occurs. Thus if the optimisation recommends a holding level of 2 for a part with a cost of $200,000 (for example an Airbus flight management computer) then it may be worthwhile to hold only one spare. When a demand event occurs, the rotable manager can then seek a further part externally. The rotable manager is then gambling that (a) the part may not fail and there may not be any demand event in a planning period and (b) that if there is a failure, there will not be a second failure in the interval required to source another part. This practice becomes better still as numbers increase: if the recommended holding is 5, then the inventory manager can hold one spare as long as further spares can be sourced quickly (which is easily established).

Taking Case 5 as an example and using the best model, LP-split, the most expensive 50 parts are selected from the data set. A rule is then applied:

Halve the recommended holding, maintaining a minimum holding of 1

The result of this reduction is a drop of 5% in the number of parts held, a global SL reduction from 90% to 87.6% and a cost reduction of $4M or 32%.

Given that the present data set is a representative sample – about one tenth - of a larger set, the potential of the heuristic can be extended to the larger group of around 3,000 parts. Assuming that the average costs reflect the larger data set, a cost saving of up to $40M can be extrapolated by manually intervening in the planning of the 500 most expensive parts. This potential saving – over and above using the best optimisation model – easily justifies the extra effort required in having rotable planning managers monitor and control these items, buying or leasing emergency spares as required.

8.1.6 Determine diminishing marginal return for inventory investment

Using the cost-oriented LP solution in the previous chapter, the change in SL per change in budget was determined for each case in the best model (LP3). Referring to Figure 7.14 **in** the analysis chapter, this measure, ΔSL / Δbudget, falls as budget increases, reflecting a diminishing marginal return. While there is no clear break in the trend, a rule may be arrived at by observation, which suggests the best point at which to stop allocating budget. Table 8.2 shows the resulting SL and budget for decision points at ΔSL / Δbudget values of 20%, 15% and 10%.

	ΔSL / Δbudget	Case 1	Case 2	Case 3	Case 4	Case 5
budget $M	20%	9	10	8	15	13
	15%	11	11	9	17	15
	10%	13	13	11	21	18
SL %	20%	92	94	93	92	92
	15%	95	95	95	94	95
	10%	97	97	97	97	97

Table 8.2: budget and SL values for thresholds of marginal return

Table 8.2 does not give a single recommendation for the 'best' level of budget but shows the rate of change of marginal return. For instance, looking at Case 5, an increase of $2M gives an increase in SL of 3%, but a further $3M brings an increase in SL of 2%. As illustrated in the Results chapter, the marginal return shown in Table 8.2 can be shown graphically as a slope – as the slope becomes flatter, the marginal return diminishes.

Given that the ΔSL / Δbudget value of 20% gives global SL values between 92 and 94%, this is likely to be acceptable in practice, so this 20% threshold is recommended as sufficient.

The actual holding in the data set achieves 89% SL with a cost of $33M; the above gives SL of 92% for $9M in Case 1 using the LP3 model.

8.2 Comparison of results with literature

With reference to the Literature Review chapter, Ghobbar (2003 a), Jackson (2003), Kilpi (2004) and the industry press confirm the scale and importance of the rotable inventory problem for aircraft operations. Therefore the significance of the problem is confirmed. Adams (2004) and Sherbrooke (1986) recognise the value in viewing the collection of rotables as a system, where the performance of the inventory is measured as the performance of the system, rather than considering parts in isolation, as do many of the more complex planning models (Kim 2007, Lee 2007). These references to the concept of system-level optimisation indicate the potential for improvement to be gained from new solutions.

The literature presents the Poisson model as the industry standard (Airbus 1997, Haas 1997) – this is used as the base case here. Although these sources are a decade old, observation has confirmed that mainstream industry practice has not changed substantially in the interim. Marginal Analysis is presented by Sherbrooke (1986) and Logistechs (2006) as a cost-oriented system-level approach. Note that, while Sherbrooke's references are dated, that author is now affiliated with Logistechs and has concentrated on commercial exploitation of his research in recent years. The Marginal Analysis technique is tested here: as the size of the problem space, and the frequency of demand, grows, the logic fails the Marginal Analysis model. This is acknowledged by Sherbrooke as a limitation, as well as the fact that Marginal Analysis is not a full optimisation. Given the more complex case for which Sherbrooke developed the METRIC models, Marginal Analysis probably works better in that context than in the single-echelon, non-indentured cases in this work, since there may be lower stock levels in the military scenario upon which the work is based.

The key articles in the literature are revisited in turn below, and compared with the work carried out here.

Ghobbar (2003 a) confirms the scale of the problem, and gives useful description of the operational issues relating to the arising of failures and the need to provide rotable support. However, demand for spares is only addressed at the part level.

Jackson (2003) confirms the scale of investment in rotables – GBP500M for British Airways - and the perceived waste, but again only considers the problem at the part level.

Airbus (1997) provides the Initial Provisioning formula, using Poisson distributions. This is employed in all of the models in this study, which extend the work by considering the problem at the system level. There do not appear to be any more recent recommendations by aircraft manufacturers for rotable inventory planning in the public domain. Once again, inventory systems are viewed only at the line item level.

Haas (1997) confirms the Poisson method as the most commonly used, however this work focuses on planned maintenance using the METRIC models, as opposed to rotable spares planning.

Adams (2004) acknowledges the importance of system-level optimisation; system-level solutions include Marginal Analysis, Genetic Algorithms and simulation. A small example (three parts) shows that the Genetic Algorithm approach gives similar outcomes to the Marginal Analysis approach and is hard to formulate. Neither can be considered a true optimisation, i.e., a technique that chooses a solution with some optimal characteristic (cost) from a large space of potential solutions.

Sherbrooke (1986) recognises that Marginal Analysis works best with very low demand rates, which are more usually found in military applications, as opposed to commercial operations. This is described as concavity, meaning that the probability distribution should descend constantly from the lowest value (one), i.e., the curve should not peak at some other value, which would lead to the algorithm choosing some quantity greater than one as being better value than a quantity of one.

Sherbrooke's work is updated by Logistechs (2006), who report large gains in system-wide cost performance for the Marginal Analysis approach. However, high demand rates are seen to produce poorer results.

Some further work has been carried out on more complex supply chain configurations. Kim (2007) looks at a more complex supply chain with depot spares, i.e., distribution centres for more efficient inventory holding in a large network. However, this work looks at spares provision in general and so is not directly focused on the line replacement problem.

Lee (2007) set out an evolutionary algorithm solution to the multi-echelon problem with random data – this goes beyond the line replacement problem and may be of interest for further study. However, this approach is hard to formulate and may be too complex to build for a large system of parts.

Kilpi (2004) promotes the concept of pooling spares among airlines. However some of the benefits are offset by additional cost and delay and the argument may not be as strong as it appears, especially if stronger optimisation methods, such as linear programming models, are used.

In summary, the models tested in this study, and reflecting mainstream practice, show that the Poisson method works consistently but takes no account of cost. The Marginal Analysis approach, in limited use in the industry, is better than Poisson, although the difference between Poisson and Marginal Analysis is less when demand increases (Cases 4 and 5).

The new models evaluated here show that the simple heuristic, cost-wise skewed holding, should be used in the improved form and is roughly as good as Marginal Analysis, with less effort and better logical consistency. The cost-wise skewed holding model is not reported in the literature.

The linear programming models give the best results, with the split model, LP3 (which treats each essentiality group as a separate problem) giving by far the best results and the most logical solution. This serves as a benchmark against which to test other solutions and can be considered optimal. The linear programming models are proposed as a test case in MacDonnell (2007).

The work carried out here builds on the current literature. In particular, description of the linear programming models, and the evaluation of further test results, hold potential for further contributions to the published literature.

.

8.3 Conclusions

There is potential for improvement in current practice for managing aircraft spare parts:

- The Poisson method, as recommended by OEMs is far from optimal;
- The Marginal Analysis, promoted as a new industry solution, is better than Poisson as it takes account of cost, but does not give an optimal solution.

There is room in the literature for new theory for optimising reusable / repairable / rotable inventory:

- Much of the literature struggles to deal with the rotable situation, where there is no net change in inventory levels over the planning period;
- Some models in the literature are concerned primarily with optimal positioning of inventory where there is a choice of locations – this is a simple proportional problem; these models do not advance the manner in which appropriate inventory levels are determined;
- The main observed method that takes account of cost, Marginal Analysis, is developed for military applications, where multi-echelon internal supply chains are the norm and stock shortages are preserved as back orders: these conditions do not fully reflect the commercial environment. Further, the Marginal Analysis model suffers from a logical flaw as quantities increase.

It is possible to develop a pure optimisation model for rotable planning, formulated as a binary linear program.

- Better results are obtained by splitting rotables into groups with matching levels of essentiality; there is no scale benefit in combining these groups.

- The optimisation model offers large improvements in inventory performance relative to cost.
- The optimisation model takes some effort to formulate and process, however there are no additional data requirements and the potential benefit of using the model in operation far outweighs the effort required to implement it.

8.4 Contributions to theory

As proposed in 4.2 in the Methodology chapter, contributions to theory are revisited and described below.

1. The main development in this work is a formulation of the rotable inventory management problem such that a recommendation is made for a set of inventory holdings where a required performance target (service level) is maintained at minimal total cost. This formulation takes the form of a binary integer linear programme and is solved by commonly available LP solving software. This contribution is in the fields of inventory management and operations management.

   In short, the aim of the work can be stated as the determination of suitable levels of rotable inventory.

   Modelling the rotable problem as described also permits review of current practice against a newly modeled holding for the purpose of comparison.

   The modelling solution further enables scenario analysis, for example the expected change in inventory levels that would result from an increase in fleet size, changing aircraft utilisation or pooling of spares among airlines.

Further contributions are claimed in the field of operations research:

2. A large-scale LP formulation representing the rotable inventory planning problem;

3. The use of an LP solution incorporating stochastic inventory demand.

8.5 Limitations of this study

8.5.1 The data set

A single data set is used here for the purpose of analysis. While this data is provided as average, typical and without bias, it cannot be claimed to be representative of all cases without proof. This would require the use of multiple data sets for different operators, different aircraft types and different operating conditions. Nonetheless, it is possible to claim a marked improvement in decision-making using the models developed here and the sample data set provided.

Access to data, due to its commercial sensitivity, is a restriction with this type of work: in general, using commercial data to develop theory is limited by the data provider's need for privacy. The willingness of an industrial case study subject to provide operational data is motivated in this case by both their awareness of the importance of the problem area and the potential for improvement.

The data set used is taken from a database with about ten times as many line items: this makes it easier to develop solutions and test multiple configurations of the data but prevents a complete assessment of the full range of operational data. The data set provided was reviewed and declared to be representative and unbiased by the inventory planning manager who provided the data. The sample was obtained by simply extracting the first 300 items from the larger set.

It is not possible to prove that the data set used did not have peculiar characteristics, which would distort the performance of optimisation models used here. However, the extensive verification and the consistency of the results across multiple solution methods gives confidence in the normal performance of the respective methods when used in combination with the present data set.

Given the limited access to data, and in order to perform an in-depth analysis of the problem and comparison of the solutions, this single data set is used throughout all 25 solutions – five techniques tested on each of five cases. The results are consistent for each technique and for each case.

Extensive sensitivity analysis, in particular for the linear programming solutions, shows consistency in the results, suggesting that there are no significant data anomalies.

8.5.2 Simplified view of the problem context

Rotable planning in practice uses simplified reliability data, as used in this study: all parts are given an average reliability measure (mean time between removals) and the same distribution is used for all parts (Poisson). However, this facilitates large-scale modelling, the benefit of which outweighs the lack of more detailed historical failure information.

The operational situation for the inventory studied here assumes a single-echelon system, where parts are held in one location. This may appear simplified when considering that aircraft operate from many locations, but it reflects the normal operating case for the bulk of inventory. It would be possible to model demand arising in multiple locations, which is a simple problem, but this is dealt with in most cases by shipping spares from one location to another. The multi-echelon model described in much of the literature reflects military operations, where there is a central store (depot) and multiple operating bases – this does not reflect commercial conditions.

The solution developed here is based on reducing stocks of costly items and increasing the service level performance of less expensive spares: this assumes that the impact of a stock-out is the same (within the same essentiality code) for a part regardless of cost, which is a fair assumption. However, the consequential costs of sourcing emergency spares are not captured here, so there may be higher costs resulting from stock-outs of more expensive parts. In summary, it is assumed here that there are no significant operational or cost downsides to skewing the inventory by part cost.

8.5.3 The techniques used

Given the one-tenth sample used as the data set, it is not possible to show here whether there are scale limitations on the optimisation model, i.e., deteriorating performance that would arise from the size of the formulation. However, it is known that larger data sets have been tested successfully in a commercial implementation of this solution.

Five techniques are tested as five models in this work: the first two reflect known practice, while the others are proposed and tested here as likely candidates for providing better solutions. The generation and selection of new solution techniques is necessarily subjective, based on the author's knowledge and discussion of the problem with domain experts. The cost-wise skewed holding is a simple partition of the first model (Poisson), which is the prevailing industry practice. While the cost-wise skewed holding technique is very simple, it offers a clear step forward in terms of solution quality, while being simple enough to be seen as an obvious progression based on the Poisson model.

The linear programming solutions formulated and tested are proposed by the author as simply an optimisation model that would select the best (lowest cost) combination of outcomes from a defined solution space (a feasible set of inventory holdings). The difference between the two LP models reflects the way in which different essentiality codes are represented, and it is necessary to evaluate both types of LP model for comparison.

The LP formulations used consist of binary problem structures, where a part-quantity pair is chosen from a small range (1 to 15 or 1 to 30). A preferable design would be an LP model that would directly calculate the optimum holding for each line item without constraint, but this was not considered to be a feasible problem statement, since the model would have to optimise against many objective functions (one for each line item).

It is not possible to categorically state that the model used is the best available, since its choice is limited by the author's awareness of suitable techniques. It is therefore conceivable that either a brute-force heuristic (evaluate all possible scenarios) or perhaps a search algorithm (such as a genetic algorithm) could find equal or better outcomes. However, given that the LP models used are true optimisations, it can be asserted that they choose the optimal solution within the defined space.

8.6 Further work

This work has been performed on a single data set, which is useful in performing detailed analyses for multiple combinations of models and scenarios. However, it is desirable to apply this work to a new data set to build further evidence of the success of the optimisation model developed here.

A larger data set is needed to perform a full optimisation: the work here is based on a data set of about one tenth of the rotable pool in support of an aircraft type. It has not been possible to obtain a full set of operational data at this time for the purpose of study and publication, due to its commercially sensitive nature. Commercial work has followed this study, with the participation of the inventory owner, but the data remains confidential. It is hoped to extend this study using a full data set for a different aircraft type to the one studied here.

Pooling inventory between operators offers obvious potential to increase utilisation and make better use of a spares inventory. However, the scale benefit is less than commonly thought if the base inventory is well optimised in the first place. In the results obtained here, doubling the fleet size calls for the spares inventory to be increase by two-thirds. Given the downsides of pooling (complex management, quality concerns, repair and usage pricing, logistics), it is not clear that a moderately sized operator will benefit greatly from pooling. This problem merits further study.

More complex reliability data would permit detailed analysis of the data and more accurate modelling. Given enough data, it would be possible to derive tailored distribution for different part types or even individual components using Weibull analysis, whereby a distribution curve is given shape parameters to reflect mean time to failure and changes over time. However, this is not worth pursuing until optimisation of the inventory is performed in the first place.

A simulation study could be performed to forecast a period of operation. However, this would need more detailed reliability data in order to have significance. Simulating with the input data used here would give no further insight than the optimisation solution. The optimal solution attained by the linear programming model, LP3, can be tested by relaxing constraints to check that a better solution is not feasible.

A longitudinal study is of interest. This would consist of enacting inventory reductions, reviewing inventory changes and compiling historical demand and SL data over a period of several years in order to confirm a real cost reduction while maintaining acceptable standards of service.

There are other industries using rotable inventory: it would be of interest to study requirements and identify the potential benefits of rotable inventory optimisation for ships, trains, power generation, mining, manufacturing and medical equipment using expensive repairable spares.

Appendix 1: References

Adams 2004. "Inventory optimisation techniques, system vs item level inventory analysis". C Adams, IEEE Reliability and Maintainability Symposium. , 55-60.

Adesola 2005. "Developing and evaluating a methodology for business process improvement". S Adesola, T Baines, Business Process Management Journal. 11(1), 37-46.

Airbus 1997. "Spare parts = cost benefit management". J Rutledge, Airbus Fast. (21), 25-29.

Airbus 1998. "Common, reliable and punctual: the path to lower spares costs". O Panayiotou, Airbus Fast. (23), 12-19.

Airbus 2001. "Cost reduction for initial spares investment". M Haupt, Airbus Fast. (27), 11-14.

Aircraft Technology Engineering and Maintenance 2001. "Trends in component maintenance". Aircraft Technology Engineering and Maintenance. Oct / Nov, 60-67.

Aircraft Technology Engineering and Maintenance 2007. "When less is more - the ins and outs of inventory management". , Aircraft Technology Engineering and Maintenance. June / July, 48-55.

Airline Fleet and Asset Management 2004. "Too much, too little, too late". , Airline Fleet and Asset Management. Mar / Apr, 44-48.

Aras 2006. "Coordination and priority decisions in hybrid manufacturing / remanufacturing systems". N Aras, V Verter, T Boyaci, Production and Operations Management. 15(4), 528-543.

Armac 2007. "Case study on SR Technics UK". Armac Systems, Armac Systems. , http://www.armacsystems.com, July 2008

Aviation Week 2000. "Online trading exchanges leave many suppliers cold". A Velocci, Aviation Week & Space Technology. 153(20), http://www.aviationweek.com, July 2008.

Aviation Week 2003 a. "Managing Technical Operations The JetBlue Way". B Rosenberg, Aviation Week. Jan 10, http://www.aviationweek.com/aw/generic/story_generic.jsp?channel=om&id=news/om1202JBlu.xml, July 2008.

Aviation Week 2003 b. "Companies Get Creative In Their Inventory Management Solutions". P Brown, Aviation Week. Apr 13, http://www.aviationweek.com, July 2008.

Aviation Week 2003 c. "Internet Trading: Generation 3". B Rosenberg, Aviation Week & Space Technology. 158(9), http://www.aviationweek.com, July 2008.

Aviation Week 2004 a. "For Boeing, Partnerships Are MRO's Future". MRO 04, Aviation Week & Space Technology. 161(10), http://www.aviationweek.com, July 2008.

Aviation Week 2004 b. "IBM Emphasizing Service Lifecycle Management". MRO 04, Aviation Week & Space Technology. 160(16), http://www.aviationweek.com, July 2008.

Aviation Week 2004 c. "SUPPLIERS: GEAR UP!". Anselmo, Joseph C, Aviation Week & Space Technology. 161(19), http://www.aviationweek.com, July 2008.

Aviation Week 2005 a. "Replacing Inventory with Information". J Croft, Aviation Week. Jun 20, http://www.aviationweek.com, July 2008.

Aviation Week 2005 b. "A Fresh InterFACE". Lufthansa, Aviation Week & Space Technology. 163(2), http://www.aviationweek.com, July 2008.

Aviation Week 2006. "Links in the Chain". J Croft, Aviation Week. Dec 06, http://www.aviationweek.com/aw/generic/story_generic.jsp?channel=om&id=news/om308scs.xml, July 2008.

Aviation Week 2007 a. "Links In The Chain: Changing Dynamics In the MRO Supply Chain". B Trebilcock, Aviation Week. Dec 13, http://www.aviationweek.com/aw/generic/story_generic.jsp?channel=mro&id=news/om1207link.xml&headline=Links%20In%20The%20Chain:%20Changing%20Dynamics%20In%20the%20MRO%20Supply%20Chain, July 2008.

Aviation Week 2007 b. "Strengthening Supply Links". , Aviation Week & Space Technology. 166(3), http://www.aviationweek.com, July 2008.

Bailey 2007. "MRO inventory reduction - challenges and management: a case study of the Tennessee Valley Authority". G Bailey, M Helms, Production Planning and Control. 18(3), 261-270.

Bashyam 1998. "Optimization of (s,S) Inventory Systems with Random Lead Times and a Service Level Constraint". S Bashyam, M Fu, Management Science. 44(12), 243-256.

Beamon 1999. "Measuring supply chain performance". BM Beamon, International Journal of Operations and Production Management. 19(3), 275-292.

Bell 2002. "In Search of Strategic Operations Research". P Bell, C Anderson, Interfaces. 32(2), 28-40.

Bertrand 2002. "Operations management research methodologies using quantitative modelling". J Bertrand, J Fransoo, International Journal of Operations and Production Management. 22(2), 241-264.

Brown 1984. "On the reliability of repairable systems". M Brown, Operations Research. 32(3), 607-615.

Burgess 2006. "Supply chain management: a structured literature review and implications for future research". K Burgess, PJ Singh, R Koroglu, International Journal of Operations and Production Management. 26(7), 703-729.

Buxey 2006. "Reconstructing inventory management theory". G Buxey, International Journal of Operations and Production Management. 26(9), 996-1012.

Cassady 1998. "Comprehensive fleet maintenance management". Cassady, C.R.; Murdock, W.P.; Nachlas, J.A.; Pohl, E.A, IEEE International Conference on Systems, Man and Cybernetics. 5, 4665-4669.

Cheung 2005. "An aircraft service scheduling model using genetic algorithms". A Cheung, W Ip, D Lu, C Lai, Journal of Manufacuring Technology Management. 16(1), 109-119.

Cheung 2005 b. "Expert system for aircraft maintenance services industry". A Cheung, W Ip, D Lu, Journal of Quality in Maintenance Engineering. 11(4), 348-358.

Computer World 2005. "Southwest Supply Chain Optimization Project". , Computer World. , http://www.cwhonors.org/Search/his_4a_detail.asp?id=4916, July 2008.

Coughlan 2002. "Action research for operations management". P Coughlan, D Coghlan, International Journal of Operations and Production Management. 22(2), 220-240.

Dubé 2003. "Rigor in information systems positivist case research: current practices, trends and recommendations". L Dubé, G Paré, MIS Quarterly. 27(4), 597-635.

Edmonson 2007. "Methodological fit in management field research". AC Edmonson, SE McManus, Academy of Management Review. 32(4), 1155-1179.

Eisenhardt 1989. "Building theories from case study research". KM Eisenhardt, Academy of Management Review. 14(4), 532-550.

Eisenhardt 2007. "Theory building from cases: opportunities and challenges". KM Eisenhardt, ME Graebner, Academy of Management Journal. 50(1), 25-32.

El Hayek 2005. "Optimizing life cycle cost of complex machinery with rotable modules using simulation". M El Hayek, E van Voorthuysen, D Kelly, Journal of Quality in Maintenance Engineering. 11(4), 333-347.

Fleischmann 2000. "A characterisation of logistics networks for a product recovery". M Fleischmann, H Krikke, R Dekker, S Flapper, Omega. 28, 653-666.

Flight International 2003. "Slow exchange". P Conway, Flight International. Oct, http://www.flightglobal.com/articles/2003/10/01/171801/slow-exchange.html, July 2008.

Flight International 2004. "GECAS launches inventory control unit". , Flight International. 17 Aug, http://www.flightglobal.com/articles/2004/08/17/185967/gecas-launches-inventory-control-unit.html, July 2008.

Flight International 2005. "Taking stock". , Flight International. 11 Jan, http://www.flightglobal.com/articles/2005/01/11/192341/taking-stock.html, July 2008.

Fortuin 1999. "Control of service parts". L Fortuin, H Martin, International Journal of Operations and Production Management. 19(9), 950-971.

Friend 1999. "Extending Visual Basic for Applications to MRP: low-budget spreadsheet alternatives in aircraft maintenance". C Friend, A Ghobbar, Production and Inventory Management Journal. 40(4), 9-20.

Friend 2001. "A predictive cost model in lot-sizing methodology, with specific reference to aircraft parts inventory: an appraisal". C Friend, A Swift, A Ghobbar, Production and Inventory Management Journal. 42(3/4), 24-33.

Fung 2001. "(T, S) policy for coordinated inventory replenishment systems under compound Poisson demands". R Fung, X Ma, H Lau, Production Planning and Control. 12(6), 575-583.

Garg 2006. "Maintenance management: literature review and directions". A Garg, S Deshmukh, Journal of Quality in Maintenance Engineering. 12(3), 205-238.

GE Engine Services 2002. "GE Rotable Solutions, LogisTechs deliver innovative LRU materials solutions". E Krimmons, GE Engine Services Service Solutions. 2(5), http://www.geae.com/services/information/servicesolutions/pdf/v02i05.pdf, July 2008.

Ghobbar 2003 a. "Evaluation of forecasting methods for intermittent parts demand in the field of aviation: a predictive model". A Ghobbar, C Friend, Computers and Operations Research. 30, 2097-2114.

Ghobbar 2003 b. "Comparison of Lot-Sizing Methods in Aircraft Repairable Component Inventory Systems". A Ghobbar, C Friend, Journal of Aircraft. 40(2), 378-383.

Ghobbar 2004 a. "The material requirements planning system for aircraft maintenance and inventory control: a note". A Ghobbar, C Friend, Journal of Air Transport Management. 10(3), 217-221.

Ghobbar 2004 b. "Forecasting Intermittent Demand for Aircraft Spare Parts: A Comparative Evaluation of Methods". A Ghobbar, Journal of Aircraft. 41(3), 665-673.

Giri 2005. "A discrete-time order-replacement model with time discounting and spare part provisioning". B Giri, T Dohi, N Kaio, Journal of Quality in Maintenance Engineering. 11(3), 190-205.

Graves 2003. "Process flexibility in supply chains". S Graves, B Tomlin, Management Science. 49(7), 907-919.

Guide 1997. "Repairable inventory theory: models and applications". V Guide, R Srivastra, European Journal of Operational Research. 102, 1-20.

Gupta 2006. "Empirical research published in Production and Operations Management (1992-2005): trends and future directions". S Gupta, R Verma, L Victorino, Production and Operations Management. 15(3), 432-448.

Haas 1997. "Target setting for the departments in an aircraft repairable item system". H Haas, J Verrijdt, European Journal of Operational Research. 99, 596-602.

Jackson 2003. "Half a billion pounds worth of nuts and bolts: BA Project Name - Review of engineering inventory models and processes". S Jackson, M Samkuban, I Luck, OR Society Handbook: OR45 conference. http://www.orsoc.org.uk/conf/previous/or45/or45_Handbook.doc, July 2008.

Jain 2001. "Analyzing the Supply Chain for a Large Logistics Operation using Simulation". Sanjay Jain, IEEE Proceedings of the Winter Simulation Conference. , 1123-1128.

Jick 1979. "Mixing qualitative and quantitative methods: triangulation in action". T D Jick, Administtrative Science Quarterly. 24, 602-611.

Jung 1993. "Recoverable inventory systems with time-varying demand". W Jung, Production and Inventory Management Journal. 34(1), 77-81.

Keizers 2003. "Diagnosing order planning performance at a navy repair and maintenance organisation, using logistic regression". J Keizers, J Bertrand, J Wessels, Production and Operations Management. 12(4), 445-463.

Kennedy 2002. "An overview of recent literature on spare parts inventories". W Kennedy, J Patterson, J Fredendall, International Journal of Production Economics. 76, 201-225.

Kettinger 1997. "Business Process Change: A Study of Methodologies, Techniques, and Tools". W Kettinger, J Teng, MIS Quarterly. 21(1), 55-80.

Kilpi 2004. "Pooling of spare components between airlines". J Kilpi, A Vepsalainen, Journal of Air Transport Management. 10(2), 137-146.

Kim 2007. "An algorithm for repairable item inventory system with depot spares and general repair time distribution". J Kim, T Kim, S Hur, Applied Mathematical Modelling. 31, 795-804.

Kranenburg 2007. "Effect of commonality on spare parts provisioning costs for capital goods". A Kranenburg, G van Houtum, International Journal of Production Economics. 108, 221-227.

Krupp 2002. "Integrating Kanban and MRP to reduce lead time". JAG Krupp, Production and Inventory Management Journal. 43(3/4), 78-82.

Lapre 2004. "Performance improvement paths in the US airline industry: linking trade-offs to asset frontiers". M Lapre, G Scudder, Production and Operations Management. 13(2), 123-134.

Lee 2007. "Multi-objective simulation-based evolutionary algorithm for an aircraft spare parts allocation problem". L Lee, E Chew, S Teng, Y Chen, European Journal of Operational Research. 189, 476-491.

Levine 2008. "Statistics for managers". D Levine, D Stephan, T Krehbiel, M Berenson, Pearson, New Jersey. 5e, 197-199.

Liberopoulos 2005. "Tradeoffs between base stock levels, numbers of kanbans and planned supply lead times in production / invnetory systems with advance demand information". G Liberopoulos, S Koukoumialos, International Journal of Production Economics. 96, 213-232.

LMI consulting 2006. "The aircraft sustainability model". , LMI consulting. Sep, citrix1.lmi.org/asm/IntroSys.pdf, July 2008.

Logistechs 2006. "Expect results". Logistechs, http://www.rotable.com, 2007.

Lye 2007. "A virtual warehouse simulation tool for aerospace rotables management". K Lye, L Chan, IEEE Aerospace Conference. 1-7.

Mabini 2002. "Controlling multi-indenture repairable inventories of multiple aircraft parts". M Mabini, A Christer, Journal of the Operational Research Society. 53(12), 1297-1307.

MacDonnell 2004. "Managing Information Technology Investment for Aircraft Sustainment". M MacDonnell and BT Clegg. Systems Engineering Symposium, MIT Engineering Systems Division. Cambridge, USA, March 2004. http://esd.mit.edu/symposium/pdfs/papers/macdonnell.pdf, July 2008.

MacDonnell 2007. "Designing a support system for aerospace maintenance supply chains". M MacDonnell, B Clegg, Journal of Manufacturing Technology Management. 18(2), 139-152.

Manos 2006. "Value stream mapping - an introduction". T Manos, Quality Progress. June, 64-69.

McKone 2002. "Guidelines for implementing predictive maintenance". K McKone, E Weiss, Production and Operations Management. 11(2), 109-124.

Michaelides 2006. "Improving transport aircraft utilisation and minimizing downtime through the development of an e-MRO solution". Z Michaelides, S Naseri, J Keith, EurOMA proceedings. 1-10.

Mingers 2001. "Combining IS Research Methods: Towards a Pluralist Methodology". J Mingers, Information Systems Research. 12(3), 240-259.

Nurmilaakso 2002. "XML-based supply chain integration: a case study". J Nurmilaakso, J Kettunen, I Seilonen, Integrated Manufacturing Systems. 13(8), 586-595.

Overhaul and Maintenance 2007. "MRO market is up and down". F Jackman, Overhaul and Maintenance. April, http://www.aviationweek.com/aw/generic/story_generic.jsp?channel=mro&id=news/omcvr407.xml&headline=MRO%20Market%20Is%20Up%20And%20Down, July 2008.

Pasin 2005. "Capacity planning and scheduling in services - a spreadsheet application". F Pasin, H Giroux, Production and Inventory Management Journal. 44(III), 1-12.

Pati 2008. "A goal programming model for paper recycling system". R Pati, P Vrat, P Kumar, Omega. 36, 405-417.

Piplani 2005. "A coordination framework for supply chain inventory alignment". R Piplani, Y Fu, Journal of Manufacuring Technology Management. 16(6), 598-614.

Porter 2008. "The five competitive forces that shape strategy". ME Porter, Harvard Business Review. Jan, 78-93.

Samaranayake 2002. "Development of engineering structures for scheduling and control of aircraft maintenance". P Samaranayake, G Lewis, E Woxvold, D Toncich, International Journal of Operations and Production Management. 22(8), 843-867.

SAP 2007 a. "Assuring the shortest possible aircraft maintenance turnaround time". , SAP Solution Brief. http://www.sap.com, June 2008.

SAP 2007 b. "MRO and M&E service providers". , SAP for Aerospace and Defence. http://www.sap.com, June 2008.

Savaskan 2004. "Closed-Loop Supply Chain Models with Product Remanufacturing". R Savaskan,S Bhattacharya, L Van Wassenhove, Management Science. 50(2), 239-252.

Scarf 2002. "Inventory theory". H Scarf, Operations Research. 50(1), 186-191.

Sherbrooke 1968. "METRIC - a multi-echelon techqnique for recoverable inventory control". C Sherbrooke, Operations Research. 16(1), 122-141.

Sherbrooke 1986. "VARI-METRIC: improved approximations for multi-indenture, multi-echelon availability models". C Sherbrooke, Operations Research. 34(2), 311-319.

Sherwin 2000. "A review of overall models for maintenance management". D Sherwin, Journal of Quality in Maintenance Engineering. 6(3), 138-164.

Singh 1989. "The development of a Kanban system: a case study". N Singh, KH Shek, D Meloche, International Journal of Operations and Production Management. 10(7), 28-36.

Smith 2001. "e-commerce and operations research in airline planning, marketing and distribution". B Smith, D Gunther, B Venkateshwara, R Ratliff, Interfaces. 31(2), 37-55.

Soh 2006. "Electronic marketplaces and price transparency: strategy, information technology and success". C Soh, M Markus, K Goh, MIS Quarterly. 30(3), 705-723.

Spengler 2003. "Closed-Loop Supply Chains - a System Dynamics Approach". T Spengler, M Schroter, Interfaces. 33(6), 7-17.

Srivastava 2008. "Network design for reverse logistics". S Srivastava, Omega. 36, 535-548.

Tedone 1989. "Repairable Part Management". M Tedone, Interfaces. 19(4), 61-68.

Templemeier 2007. "Inventory Management in Supply Networks, Norderstedt (Books on Demand) 2007". H Templemeier, http://www.advanced-planning.eu/advancedplanninge-237.htm.

Teunter 2008. "Dynamic inventory rationing strategies for inventory systems with two demand classes, Poisson demand and backordering". R Teunter, W Haneveld, European Journal of Operational Research. 190, 156-178.

Thonemann 2002. "Easy quantification of improved spare parts inventory policies". U Thonemann, A Brown, W Hausman, Management Science. 48(9), 1213-1225.

Weckman 2001. "Modeling the reliability of repairable systems in the aviation industry". G Weckman, R Shell, J Marvel, Computers and Industrial Engineering. 40, 51-63.

Wong 2005. "Inventory pooling of repairable spare parts with non-zero lateral transshipment time and delayed lateral transshipments". H Wong, D Cattrysse, D van Oudheusden, European Journal of Operational Research. 165, 207-218.

Yang 2000. "Managing and flow line with single-Kanban, dual-Kanban or Conwip". KK Yang, Production and Operations Management. 9(4), 349-366.

Zhao 2005. "Inventory sharing and rationing in decentralized dealer networks". H Zhao, V Deshpande, J Ryan, Management Science. 51(4), 531-547.

Zorn 1999. "Modeling dimininshing marginal returns in a hierarchical inventory system of repairable spare parts". W Zorn, R Deckro and L Lehmkuhl, Annals of Operations Reseach. 91, 319-337.

Appendix 2: Comparison of Poisson and Gauss (normal) distribution functions

The Poisson distribution is acknowledged in the literature as the most appropriate for estimating the failure of items with long intervals between failures. The Poisson distribution is prescribed by the main aircraft OEMs, Airbus and Boeing, as the function to use in estimating demand for the purpose of planning rotable inventory levels. However, both OEMs recommend changing to a Gauss (normal) distribution when the rate of failure is expected to be higher than an arbitrary threshold (20, 30 or 50) in the planning period (generally accepted as one year).

Poisson and normal distributions are evaluated here for a range of means. The functions are shown as mass functions and in cumulative form.

The cumulative Poisson formula is:

$$E(x) = \sum_{k=0}^{x} \frac{e^{-\lambda}\lambda^{k}}{k!}$$

where E(x) is the expected probability of a value x, e is the base of the natural logarithm (2.718), k is the discrete integer variable ranging from 0 to x and λ is the mean expected value.

The Poisson distribution is a discrete distribution, calculated for integer values of x.

The normal distribution is calculated as:

$$E(x) = \frac{1}{\sqrt{2\pi\sigma}} e^{-(1/2)[(x-\mu)/\sigma]^2}$$

where E(x) is the expected probability of a value x, ∏ is the mathematical constant 3.1412, σ is the standard deviation (width of the curve), e is the constant 2.718 as before and μ is the mean.

As an approximation to the Poisson function, the normal distribution uses the square root of the mean as the standard deviation, which is an input shape parameter for the normal distribution. The Poisson distribution is a function of mean only and does not have a variable shape parameter. The normal distribution is a continuous distribution, calculating all cumulative values up to the value of x.

The Poisson mass function is plotted in Figure A2.1 below: note that this is a discrete function, so the graphs show curves joining data points at integer values for illustration.

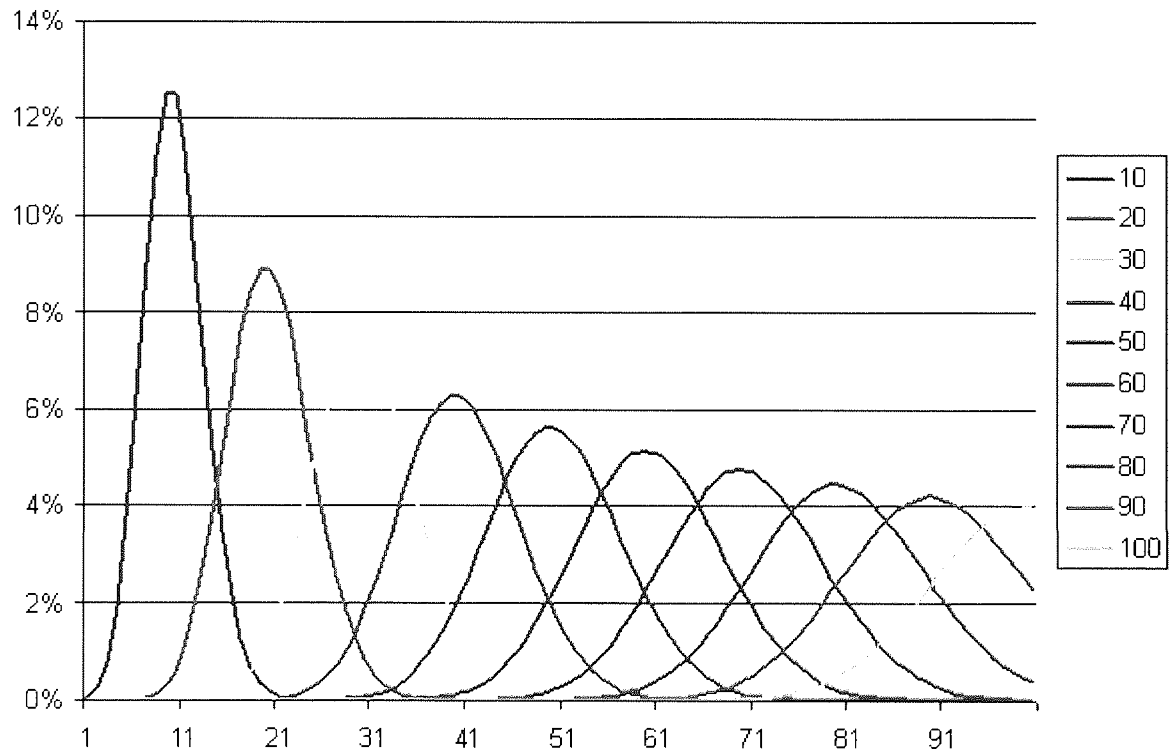

Figure A2.1: Poisson mass function

Figure A2.1 shows expected values on the x-axis, with each curve calculated for the mean value shown.

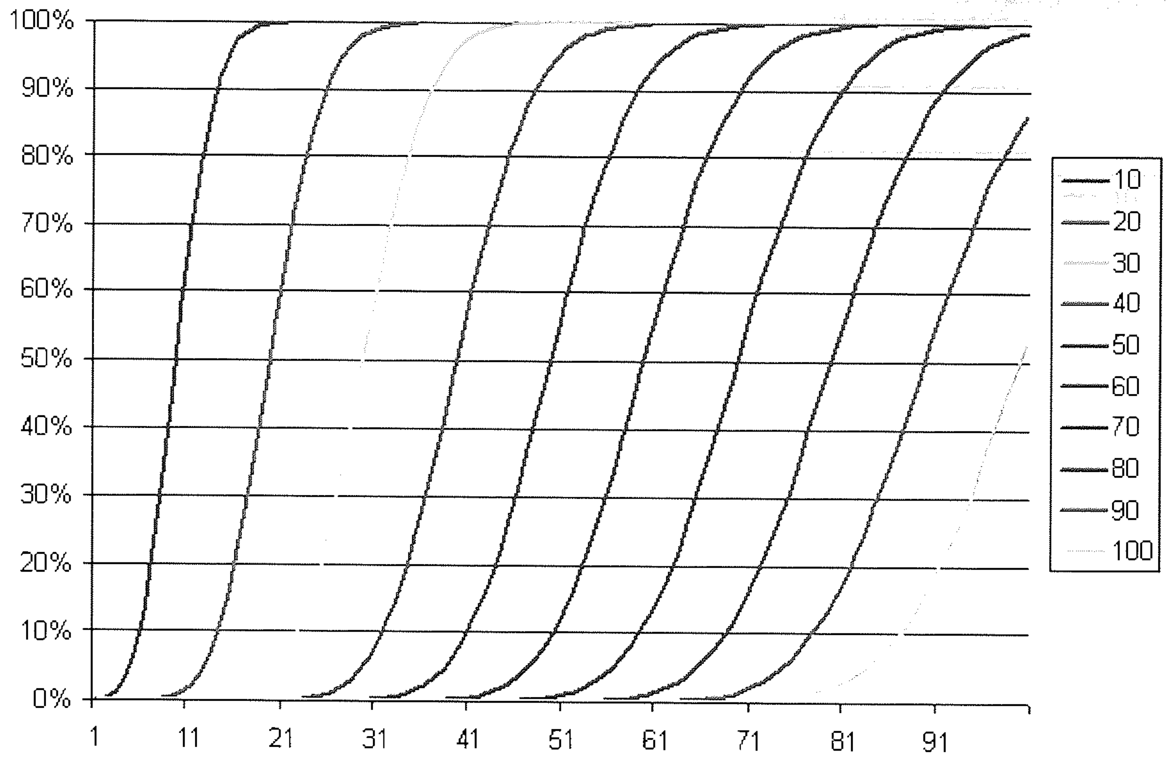

Figure A2.2: cumulative Poisson function

Figure A2.2 is the cumulative Poisson function, showing the sum of expected values up to and including the point on the x-axis, for each of the means shown in the legend.

The normal function, used with $\sigma = \sqrt{\mu}$, looks very similar to the Poisson curves in Figures A2.1 and A2.2.

Figures A2.3 and A2.4 show the difference between normal and Poisson values for the range of expected values and mean values from 10 to 100.

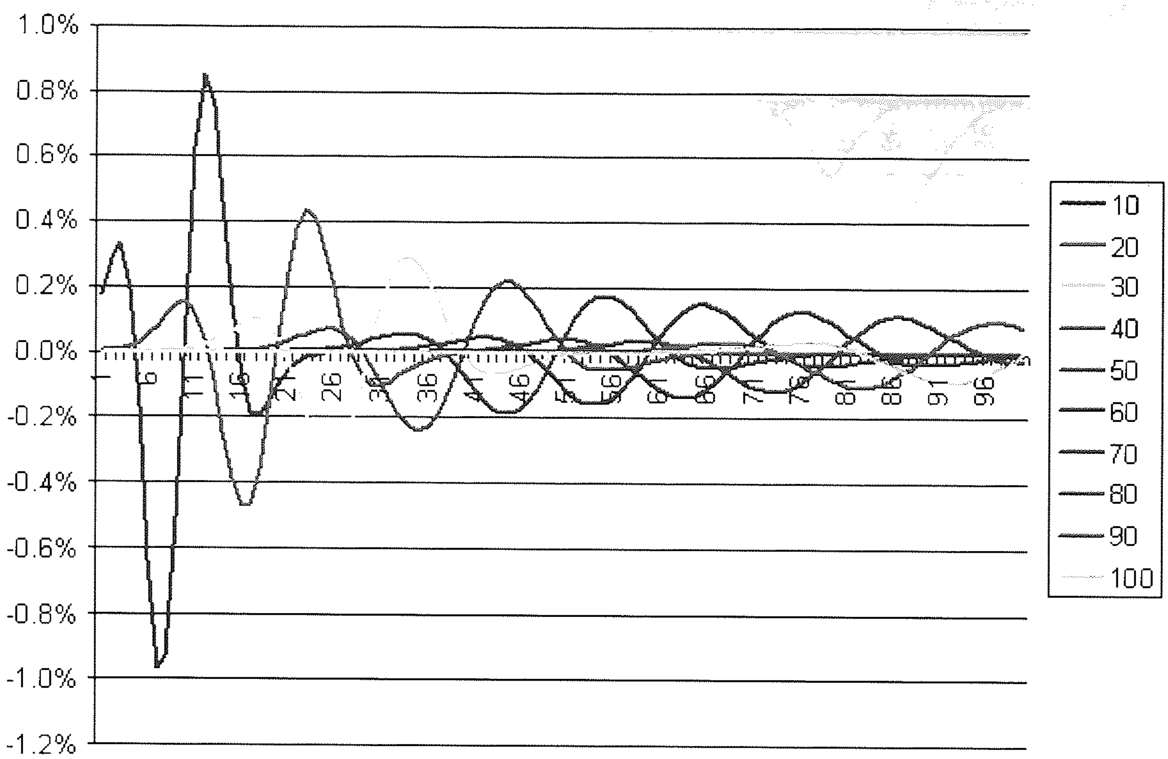

Figure A2.3: (normal-Poisson) mass functions

Figure A2.3 shows the Poisson mass function subtracted from the normal mass function. The biggest departure is for the lowest mean value, 10, the range where Poisson is normally recommended. Thus using the normal distribution as an approximation for low mean values will give results that are significantly different.

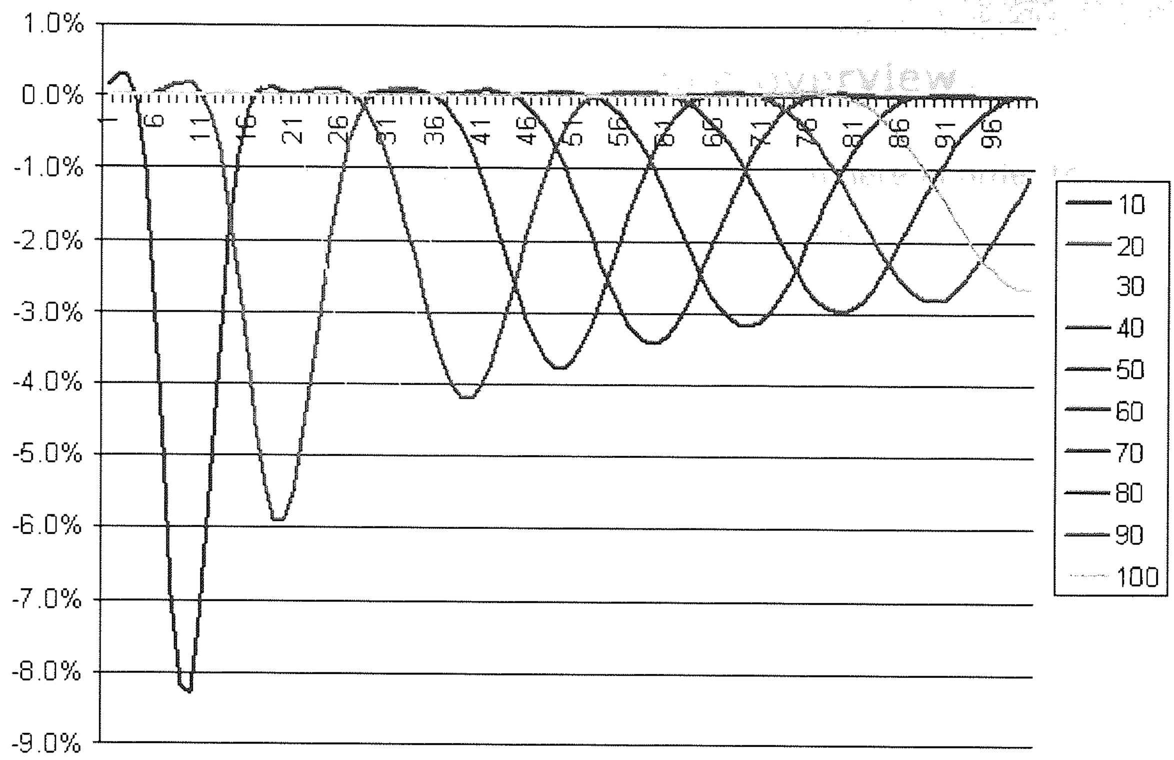

Figure A2.4: cumulative (normal-Poisson) difference

Figure A2.4 shows the more pronounced difference between normal and Poisson when the functions are shown in cumulative form (which is the relevant version for calculating inventory demand). What the graph shows is that, for example if 10 parts are held for a mean rate of demand of 10 items, the normal distribution predicts 8% less chance of success than does the Poisson. The actual values for a mean of 10 and an expected value of 10 are 58% for Poisson and 50% for normal. Thus Poisson predicts greater success, so will show better performance. Since the effect becomes less as the mean value increases, it is decided to use the Poisson distribution for all mean values.

Appendix 3: Aircraft maintenance overview

A description of aircraft maintenance processes is given here in order to show the complex yet generic nature of this problem area.

Airframe and engine maintenance are generally treated quite separately, however in the context of rotable spares they can be considered as part of the same issue.

A3.1 Airframe maintenance

Scheduled airframe maintenance events are classified by the following levels of work, or workscopes:-

A check – monthly, performed overnight at the line station.

B check – approximately every three months, also performed overnight on the line.

C check – approximately every 12 – 18 months, performed by an MRO in a hangar and taking several weeks.

D check – approximately every 5 years, performed by an MRO in a hangar and taking several months, comprising complete stripping, inspection, repair, rebuilding, testing and re-certification of the aircraft. All cabin and cockpit interiors, equipment and fittings are removed and overhauled, as are the landing gear assemblies.

All of these events are scheduled and determined by the number of hours and cycles flown by the aircraft.

Airframe maintenance checks A to D are specified differently for each aircraft type and also vary by airline depending on the equipment fitted to the aircraft and airline operating procedures. However, the generic classifications are useful in planning and sourcing maintenance services. Thus, while the exact cost of a D check is not know in advance, due to the many repair and replacement events that may occur, MROs will negotiate contracts around D checks such that an airline will agree a maximum or average budget with a limit on the number of manhours chargeable. This allows the airline to negotiate competitively among MROs and puts the MRO in the unfavourable position of setting manhour limits on checks for pricing, while typically accruing 20 or 30% extra manhours for which it cannot charge. Thus airframe overhaul is heavily costed around labour and is very competitive, so it is very difficult to operate an airframe MRO profitably.

Airframe maintenance may be performed at a range of locations, although the cost and time involved in moving aircraft for maintenance are significant for smaller jobs. It is usual for an airline to have a full-time engineering representative on site with the MRO during C and D checks and there may be pooling and exchange of parts, which suggest consolidating maintenance tasks in the smallest possible number of locations.

By its nature, airframe overhaul must be carried out at an airport: where the airport is busy, it will usually only be feasible to provide maintenance to aircraft that are normally based there. In London Heathrow Airport, for example, where landing slots are fully utilised, it would not be feasible to fly aircraft in for maintenance. On the other hand, it makes sense to maintain aircraft there if they are usually based there, since to move them to a quieter airport requires more slots to leave and return to Heathrow. Line checks, A and B checks will obviously be performed where an aircraft is parked overnight, but it also makes sense to have heavier checks performed at the home base. A restriction on this approach is the availability of hangar space at the base airport, but this is usually less of a problem than landing rights since there has historically been ample hangar space at main airports, where national flag carriers usually operated their own MRO facilities.

In choosing where to source MRO services, airlines have flexibility, notwithstanding the cost and other issues associated with ferrying aircraft solely for maintenance activities. Large buyers of MRO services may award contracts to different MROs for different aircraft types but will seek economies of scale for the same aircraft type. Heavy checks (C and D) are usually sourced separately from line operations, which may be kept in house by a large operator. Generally, an operator will look for the lowest unit labour cost and will expect other terms (standard numbers of hours for defined workscopes, material mark-ups) to be the same among different supplier. Thus local input costs for the MRO will play a large part in competitiveness. For long-haul operators there may be potential to source overhaul services in destinations away from their base: for example, Virgin Atlantic, who source most of their heavy checks in the UK and Ireland, can avail of lower costs and currency conditions in Asia and the USA. However, this requires a long-term work plan so that aircraft coming out of service can be replaced with aircraft coming into service, reducing the need for ferrying empty aircraft.

Airframe MRO providers bill heavy checks in three ways:

1. Labour – manhours x an agreed hourly rate. There may be agreed standard workscopes so that the labour charge for a given job is effectively fixed. Some MROs now offer fixed prices for the constituent tasks in a larger job, so that prices are fixed but dependent on the work required. Given the constraint on chargeable hours, there is little profit potential in selling labour. Also, it is hard to pass on rises in labour costs in an openly competitive market.

2. Subcontract – some components are routed for outside work, but most large MROs will aim to create as many repair capabilities as possible in house. These activities offer more profit potential than C and D check work. Also, once the capability is in place, repair services can be sold to other users (MROs, airlines and traders). Large MROs may have workshop facilities for radios, avionics (cockpit equipment), sheet metal panels, composite repairs, landing gear, electrical equipment, hydraulics, pneumatics and machine shops for structural components. Further areas where work may be insourced by larger MROs include seat

overhaul and In-Flight Entertainment equipment repair. Some repair facilities can be moved off-airport where the cost of facilities is lower and it may be possible to use staff who are not certified for aircraft work. For example, seats and IFE may be repaired by specialists and inspected by aircraft technical staff prior to reinstallation.

3. Materials – in airframe heavy checks, the bulk of materials charges are small consumable items – the customer will not normally permit the replacement of expensive items without approval.

While heavy checks are the ostensible function of an airframe MRO, these businesses are typically divided into two areas for operational purposes:

1. Aircraft overhaul – performing prescribed checks. The main management tasks are labour scheduling and documentation management.
2. Component management – an MRO's greatest opportunity for profit is in materials management, and particularly component management. This refers to buying, selling, stocking, reworking and supplying items with global demand and a sufficient inherent value. Using in-house repair capabilities, an MRO can purchase used stock and refurbish it as buyers are found. Another service is to monitor, control and maintain customer spares stock. The MRO may enter into an agreement where they undertake to repair an airline's removed components and aim to provide full cover of replacement parts. A further development on this idea is to provide component support where the airline holds no spares of its own: the MRO holds suitable spares and provides a defined level of service to the airline. If an MRO can cover the same parts requirements for several customers with one stock, it can make large profits by managing its stocks well.

When considering aircraft parts, items are classed in the following inventory categories, in increasing order of value:

1. Consumable – material that is used up in operations, such as oils and hydraulic fluids.
2. Expendable – materials used once, such as gaskets, seals, connectors, batteries and fixings. The most numerous part in a metallic airframe is rivets – once used they must be replaced and are unsuitable for repair.
3. Repairable – items that are not regularly replaced or maintained, but are worth repairing. They will normally be reinstated on the parent aircraft. Examples include fuselage sections, seats and cabin panels.
4. Rotable – components designed to be exchanged without disruption to service, or Line Replaceable Units. These are usually changed overnight and do not require the aircraft to be removed from service. Examples include hydraulic pumps, avionics and galley equipment. These items are typically quite expensive and worth repairing. Since they are replaced on the line, removed items are not reinstated by are held in rotable stock to be used for a future requirement. Since items are repaired and re-stocked, the demand for these items, and the appropriate stock levels, are determined by the reliability of the part type and the level of utilisation of the part in the fleet. Thus determining suitable levels of spare parts inventory is a peculiar problem, and the focus of the detailed work in this project.

A3.2 Engine maintenance

Engine maintenance arisings originate in a different manner to airframe checks and give rise to different types of maintenance activity.

The pilot monitors engine performance and engine data is gathered and analysed by powerplant engineers, who may be airline technical staff, MRO staff, or both. This data is collected at the end of each flight and may be conveyed by paperwork or by telemetry over a GSM link.

The pilot observes Engine Pressure Ratio (EPR), the increase in air pressure at the back of the engine over the intake pressure. This must be above a prescribed value for take-off to proceed: the pause in acceleration at the start of a take-off roll is the point at which the pilot is checking for sufficient thrust, shown as an EPR reading above a defined limit, which is set for a given aircraft and operating conditions. In some cases, depending on the aircraft and engine, the pilot is governed by the rotational speed of the fan, while confirming that EPR is acceptable. In order for an engine to deliver the required effort, it is burning fuel and the spools (compressor and turbine assemblies) are rotating. At the point of take-off, fuel burn rate is not considered a sufficiently precise measure of engine operation, so the Exhuast Gas Temperature (EGT), is measured by a set of probes in the exhaust nozzle. A more efficient engine will deliver the required effort (EPR) with a lower exhaust temperature and thus lower fuel burn. As an engine wears and loses compression, it must burn more fuel and rotate faster to generate the required thrust. Thus there are limits for rotational speed, vibration and EGT for safe operation. Excessive EGT will cause metal damage to the engine: the engine parts in the path of the hot turbine gases encounter temperatures far above the melting point of the parent metals. This is possible by covering 'hot section' parts to thermal barrier coatings containing mineral and ceramic compounds. Also, the engine contains inner air channels to route cooling air through airfoils and into the hot gas stream. There is a defined EGT limit, the safe operating temperature of the engine, and the EGT achieved at required EPR is measured, the difference from the limiting value being called EGT margin. When the EGT margin value is low (single figures) or is dropping rapidly over successive flights, the engine will be removed for overhaul.

In summary, an engine can be removed for 'on-condition' causes as follows: EPR not reached (lack of compression and thrust), overspeed (excessive rotational speed needed for take-off thrust), vibration outside limits (caused by airfoil wear or damage or bearing misalignment), low EGT margin or poor EGT trend (margin falling rapidly), excessive fuel consumption. In addition, engine diagnostic systems may signal failure of components or poor performance.

Engine maintenance events fall into two categories, described below.

1. *Repair* – when an engine falls below its performance threshold, or if it malfunctions, it will be removed for maintenance. Failures requiring minor work include external oil leaks (for example leaking hoses or fittings) or fan blades dented beyond accepted limits. It may be possible to return the engine to service by performing a module swap, for example, changing the fan or main gearbox. Failing that, it may be necessary to perform a module disassembly and carry out repairs on major internal assemblies. This will then require rebuilding and testing of the engine. By definition, repair events, termed *unscheduled engine removals*, are unpredictable and will range from an external repair to a module swap to performance restoration to full overhaul, depending on the diagnosis. There are many dependencies in the engine repair process, so that repairing a core element of an engine will trigger mandatory actions on related parts. For example, changing a fan (the largest set of compressor airfoils) will call for a visual inspection of the exposed second stage compressor and an internal inspection (boroscope) of the remaining compressor stages.

2. *Overhaul* –maintenance events are driven by on-condition criteria, which are established by performance monitoring, with the exception of arising time limits on life-limited parts. Where there is a general deterioration in performance, and the engine has had a respectable release time since overhaul, it will usually be necessary to overhaul the engine – it is a combination of good planning and luck to have minimal time remaining on life-limited parts.

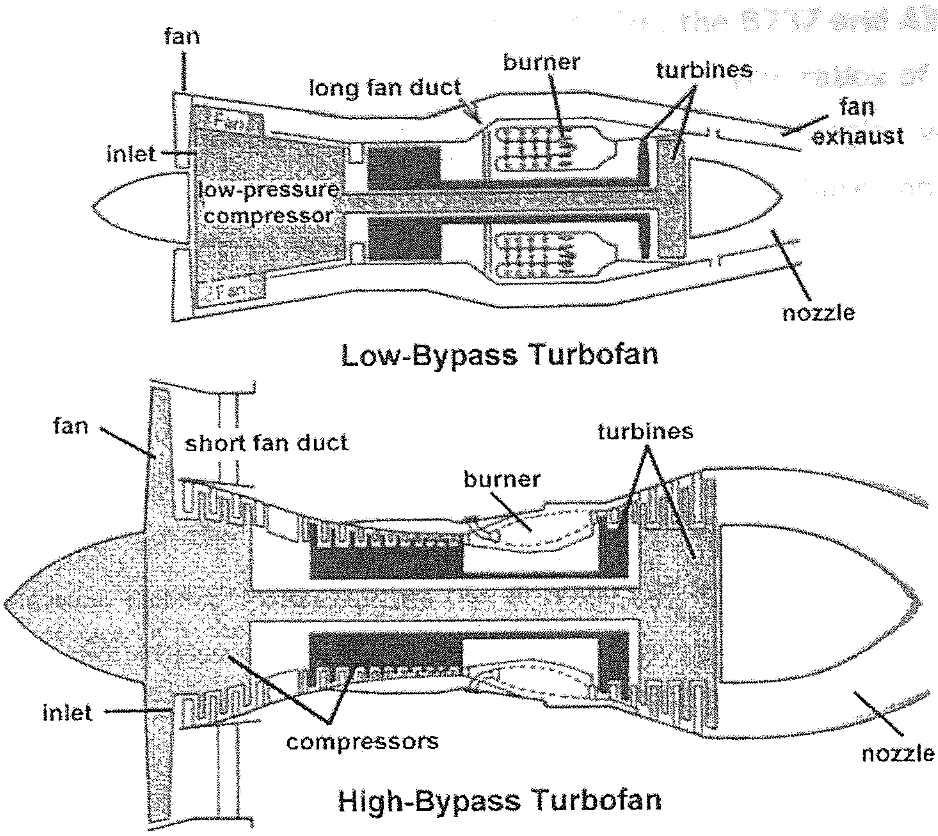

Figure A3.1: early (top) and modern (bottom) jet engine designs

(source: aerospaceweb.org)

There has been a general shift in jet engine design in the past twenty years: Figure A3.1 above shows the older design (top) compared with the newer design (bottom). Most engines manufactured since the early 1990s are referred to as *modern* generation engines and have the following main distinguishing features compared with earlier engines:

1. Platform design – engines are no longer designed for a single aircraft type, but as a platform to serve many aircraft and upon which many variants can be based, with different thrust ratings and applications. For instance, the most common engine in commercial aviation, the CFM56, is found on Boeing 737s in various power settings (-3 with approximately 20,000 pounds thrust), equivalent Airbus A320 family narrow-body twin engine aircraft, and in a higher-thrust version on the four-engine wide-body Airbus A340 (-5C at approximately 30,000

pounds thrust). As well as differing in size, the B737 and A320 are short-haul aircraft, with typical flight-hour to cycle ratios of 1, while the A340 can fly sectors up to 16 hours long. Different engine versions in the same family are not interchangeable but may share components and maintenance processes.

2. FADEC, Full Authority Digital Engine Control refers to the use of solid-state electronic control systems as an integral part of an engine. The main Engine Control Unit performs a calculation, which has as its inputs altitude, air temperature, engine speed, pressure at several points in the engine and the pilot input, or Power Lever Angle. Combining these inputs, the ECU determines the correct fuel flow rate to meet the requested power setting. This electronic system supercedes earlier mechanical fuel controls, which can take one person three months to overhaul and calibrate.

3. High-bypass fan – see figure A3.1 above: modern engines combine the fuel efficiency of propeller engines at low speeds with the high-speed capability of jets by having a fan bypass ratio of around 6, meaning that 6 times as much air from the fan bypasses the engine as goes through the engine core.

4. Noise regulation compliance – high-bypass fans and improved exhaust nozzle design allow modern engines to meet requirements for operation in built-up areas.

5. On-condition maintenance – engine designers aim to minimise the use of life-limited parts and scheduled maintenance tasks so that the engine can operate for as long as it is performing well. CFM56 engines used since the late 1980s on Boeing B737 aircraft have an expected first run life (time to first overhaul) of 16,000 hours, but many reach 25,000 hours (cfm56.com). Older generation engines typically ran for 8,000 to 10,000 between major shop visits.

6. Modular construction – newer engines can be split into modules, which are serialised and tracked separately. It is therefore possible to re-combine modules in a more optimal manner at the repair facility in order to minimise downtime and maximise service life. For example, a serviceable (fit for use) fan and main gearbox from an incoming engine can be fitted (following a minimal workscope) to an engine whose other modules are complete, possible completing the engine earlier. Also, modules with life limited parts can be combined with comparable modules, so that an engine will not be removed from service due to one module hitting a time limit. Further, the workscopes required for different modules may be managed differently – thus the 'hot section' modules (combustor and turbines) will require different maintenance routines to the other parts. By purchasing extra hot section modules it may be possible to reduced engine shop visit times and reduce the number of spare engines needed to cover engines in the maintenance cycle. The cost of spares is large: a typical overhaul takes three months, so if an engine is overhauled on average every three years, an airline will need an extra engine for every twelve engines installed (3 / 36 months = 1 / 12), just to cover overhauls. Many airlines will typically have twice that number of spares to cover failures, engines being removed during a time of above-average utilisation and more engines being removed at the same time for overhaul than there are spares.

Most modern engines have the following principal modules:

1. Fan – a disk and around 40 profiled blades in matched pair, attached to the front of the low-pressure compressor and driven by the low-pressure turbine
2. Fan casing – the containment ring around the fan, which will be the main structural support for the engine
3. Low-pressure Compressor LPC – several stages (around 5) of axial compressor, driven by the low-pressure turbine. Each stage consists of a

toothed disk connected to adjacent disks and populated with between 40 and 80 blades. Between each rotor stage is a set of stationary blades, or guide vanes, which straighten the airflow from each rotor stage, preventing the body of air from spinning and send it in an axial direction.

4. High-pressure Compressor HPC– several stages (around 7) of axial compressor, driven by the high-pressure turbine, with successively smaller cross-sectional areas. Again, there are stator stages between the rotors.

5. Combustor – an annular enclosure, or a set of cylindrical 'cans' with fuel injector nozzles immediately downstream from the last set of stator vanes, or exit vanes, at the back of the HPC. A fine fuel spray is injected into highly compressed air and the resulting expanding gases flow at higher pressure from the back of the combustor.

6. High-pressure Turbine HPT – the combustion gases impinge on the first set of nozzle guide vanes, then onto the first stage of the HPT, which is caused to rotate. There may be a second HPT rotor stage, but many engines have only one stage. The stationary and rotating airfoils in the HPT are coated with ceramic thermal barrier coatings and cooled by internal air flow, which is bled out the back of airfoil to join the main air flow through the engine. The HPT shaft is joined to the HPC and drives the HPC.

7. Low-pressure Turbine LPT – several stages of turbine increase in cross-sectional area as the gas pressure drops. The engine is designed such that the LPT removes as much kinetic energy as possible from the gas flow so that there is very little thrust from the engine exhaust, the main propulsive thrust being provided by the fan, which is driven by the LPT. By designing the engine with little exhaust thrust, it uses energy more efficiently by driving the fan. Also, the engine can be modified for other

uses, such as a turboprop or turboshaft (which can drive a helicopter, ship or generator), where exhaust thrust is not desirable.

8. Low-pressure turbine shaft - the LPT is connected by the LPT shaft to the LPC. The LPT shaft runs concentrically inside the high-pressure spool so that low-pressure and high-pressure spools are mounted independently and are free to run at different speeds.

9. Main gearbox – a tower shaft is driven by a bevel gear on the LPT shaft, at the back of the LPC, and housed in a fan casing strut. The gearbox is mounted on the outside of the fan casing and is used to drive auxiliary components such as a hydraulic pump, engine fuel pump and an electrical generator via a speed regulator called a constant speed drive

Together, the fan, LPC, LPT shaft and LPT form the low-pressure rotor, or spool, which is fixed together and mounted on radial and axial bearings. The HPC and HPT are linked (shown dark in figure A3.1) and span the combustor, fixed in rotation and forming the high-pressure rotor. In operation, the speeds of the low-pressure spool and high-pressure spool are indicated to the pilot as N1 and N2. N1 tells the pilot the speed of the fan, which generates most of the thrust, and is used at the thrust criterion for some engines. N2 tells the pilot what the core of the engine is doing – this should be in step with N1. If N1 is high relative to N2 it means that the engine is not developing thrust. There are bleed valves between high- and low-pressure compressor stages to prevent vacuums and surges caused by the delay in N1 responding to changes in N2. These valves are operated for a few seconds to let air in (on acceleration) or out (on deceleration). Engine speeds, which can be 30,000 rpm for N2 (far lower for N1) are indicated as percentages of the recommended maximum values. Thus for a CFM-powered aircraft the pilot will look for 100% N1, and will check that N2 is in line.

Engines are not considered transport-sensitive for MRO operations: an overhaul may take three months and cost $2M, while transport will cost around $10,000 in the same region and take 2 or 3 days. Engines are often shipped by road but can be flown for maintenance. It is unusual, however, for an engine to be sent outside its own continent for maintenance. Engine overhaul facilities may be located in industrial areas removed from airports where the cost of facilities may be lower and access easier. A large engine overhaul facility will occupy 30,000 to 40,000 m^2 of workshop facilities and will usually have a separate test cell of around 5,000 m^2. The test cell must be away from built-up areas and built to very high standards, as it must be able to contain an uncontrolled engine failure safely. Also, the exhaust from a test cell, which emerges vertically, must be away from air traffic areas as the blast is dangerous to aviation.

Engine maintenance is a specialist activity and usually operated on its own – engine MROs do not typically engage in other businesses or airframe overhaul. They will usually have allied activities to maximise their resources, namely engine trading and leasing and parts trading. The engine trading and leasing activity may be useful in providing spares cover to customers, which may be part of a negotiated contract. Equally, the MRO is in an ideal position to buy and hold unserviceable engines, which it can then overhaul at cost using slack production time, or it can strip old engines for parts, which it may then re-stock or sell on their own. An engine MRO can trade parts very profitably by buying unserviceable parts and holding them until until a buyer is found. By holding buffer stocks, the MRO can delay overhaul activity until the part is needed to replace stock being sold or consumed.

Engine maintenance charges are billed in three categories:

1. Labour – a number of hours at an agreed rate, negotiated as part of a long-term contract with the customer. The MRO doesn't make much profit on labour and is typically restricted by competitive standards to a given number of hours for a given type of work.

2. Subcontract – where components are sent for outside repair, the vendor charges are passed on, with a small mark-up. Depending on the part, there may be an exchange item provided to save time.

3. Materials – most of the profit on engine overhaul is in the sale of materials, both for the OEM and the MRO. The OEM will often agree the sale of engines at cost in order to secure the future demand for spares. The MRO will usually charge a mark-up of around 15% on OEM prices where the MRO supplies the material. Some customers seek to buy material directly but this causes delays and confusion and is discouraged. Where an MRO can make more profit is in providing overhauled replacement parts, for which it will charge a proportionate price.

Appendix 4: Publications

The following are published peer-reviewed works associated with this study.

Journal article

MacDonnell M and BT Clegg. "Designing a support system for aerospace maintenance supply chains", Journal of Manufacturing Technology Management, 18(2), pp 139 – 152, 2007.

Book chapter

MacDonnell M and BT Clegg. "Management of Rotable Aircraft Spares Inventory: Review of Practice and Development of New Solutions" in "Recent developments in supply chain management", R de Koster and W Delfman, pp 149 – 158. Helsinki School of Economics, 2008 Helsinki.

Published conference proceedings

MacDonnell M and BT Clegg. "IT Support for Managing Aircraft Spares in a Closed-loop Supply Chain", in M Khosrow-Pour (ed) "Emerging Trends and Challenges in Information Technology Management", Proceedings of the 2006 Information Resources Management Association Conference, May 2006. Washington DC.

MacDonnell M and BT Clegg. "Management of Rotable Aircraft Spares Inventory: Review of Practice and Development of New Solutions" in E Anderson et al (eds) "POMS College of Product Innovation and Technology Management: First Conference 2006", proceedings, Production and Operations Management Society, May 2006. Boston.

MacDonnell M and BT Clegg. "Management of rotable aircraft spares inventory: Development of a new solution" in DJ Bennett, B Clegg, A Greasley, and P Albores (eds) "Technology and Global Integration", Proceedings of Second European Conference on Management of Technology Aston Business School & IAMOT. Aston Business School, 10- 12 September 2006 pp 470-477.

MacDonnell M and BT Clegg. "Developing an inventory management system for aircraft operations support" in K Soliman and R Connolly (eds) "Information Management in the New Economy", Proceedings of 8th Conference of the International Business Information Management Association, June 2007, Dublin.

Workshop

MacDonnell M and BT Clegg. "Managing Information Technology Investment for Aircraft Sustainment". Systems Engineering Symposium, MIT Engineering Systems Division. Cambridge, USA. http://esd.mit.edu/symposium/pdfs/papers/macdonnell.pdf. March 2004.

Appendix 5: Data sample

Below is a sample of the operational data set used for testing: the data set used consisted of 300 line items. Note that there is no reliability data for line items 1, 12, 13 and 17: in all, 27 items were omitted from the data set for the models tested as there was insufficient data, leaving 273 items to test.

line	PartMaster	DESCR	Stk	Ess	SLReq	GBV $	MLP $	MTBR	TCH	TAT
1	071-01478-0001	CONTROL PANEL, ATC	2	2	93	0	15000	0	6820	38
2	071-01503-2601	ANTENNA, TCAS	2	2	89	0	5000	3751	37510	38
3	10-61312-9	JACKSCREW, FLAP TE	34	1	95	3366	3202	3589	272800	28
4	10-617980-1	EXCITER, IGNITION	21	1	95	2144	2700	21864	371690	28
5	10470-6	PUMP, STANDBY	6	2	89	2010	54975	19486	136400	28
6	107484-5	REGULATOR, HIGH STAGE	1	1	95	2261	6050	8974	170500	28
7	107492-2	REGULATOR, BLEED AIR	4	1	95	5120	7291	5602	156860	28
8	108032-8	REGULATOR, AIR PRESS	16	2	89	416	1835	7199	64790	28
9	109486-6-1	VALVE, BLEED SHUT-OFF	17	1	95	3635	14180	1705	105710	28
10	114-029	INDICATOR, POSITION O/F VLV	6	3	75	762	3000	28741	201190	38
11	1211175-011	ACTUATOR, VARIABLE STATOR	2	2	89	0	9000	40920	81840	38
12	123266-2-1	VALVE, SUPPLY DUCT CHECK 3 1/2	4	2	93	0	3378	0	153450	28
13	123268-1-1	VALVE, SUPPLY DUCT CHECK 3 1/2	5	2	93	0	2744	0	129580	28
14	129666-2	SENSOR, PRECOOLER CONT VALVE	14	2	89	4242	3802	4498	211420	28
15	129694-2	THERMOSTAT, BLEED AIR O/TEMP	5	2	89	2189	3677	27848	167090	28
16	1316200-3	VALVE, CONTROL BRAKE METERING	3	1	95	2660	5584	368280	368280	28
17	152050	COOLER, OIL	7	2	93	0	3861	0	109120	28
18	152LMA18	INDICATOR, EGT	3	2	89	2853	5976	10230	51150	28
19	158300-101	ACTUATOR, AUTO PILOT AILERON C	7	1	95	42540	69522	37023	518320	20
20	162300-103	COMPUTER, ELEV FEEL	3	1	95	67710	92577	10798	64790	28

Table A5.1: data sample

Line = sequence number
Descr = rotable description
Ess = essentiality code {1, 2, 3}
GBV $ = Gross Book Value in USD
MTBR = Mean Time Between Removals
TAT = Turn Around Time {20, 28, 38}

PartMaster = industry standard part number
Stk = number of spares owned
SLReq = SL, according to Ess
MLP $ = Manufacturer's List Price
TCH = Total Component Hours

Appendix 6: Sample model code

Note: the full set of spreadsheet and LP_Solve files used in this work is available at *http://mmacdonn.ucd.ie/rotable.*

Model 1: Poisson

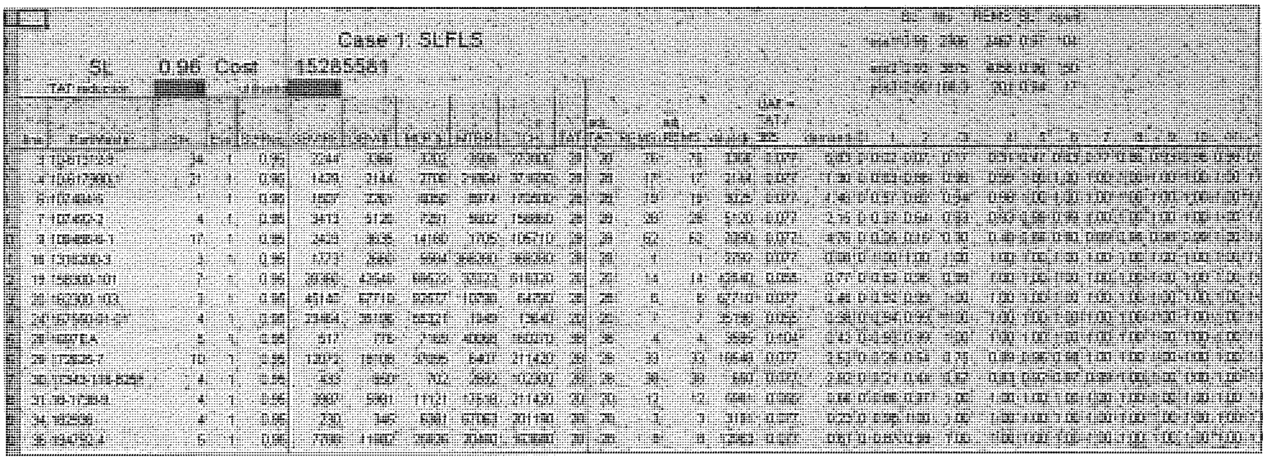

Figure A6.1: Model 1 spreadsheet sample

Figure A6.1 shows expected values for a range of quantities for each line item – the algorithm selects the quantity of each part for which the target SL is exceeded.

Model 2: Marginal Analysis

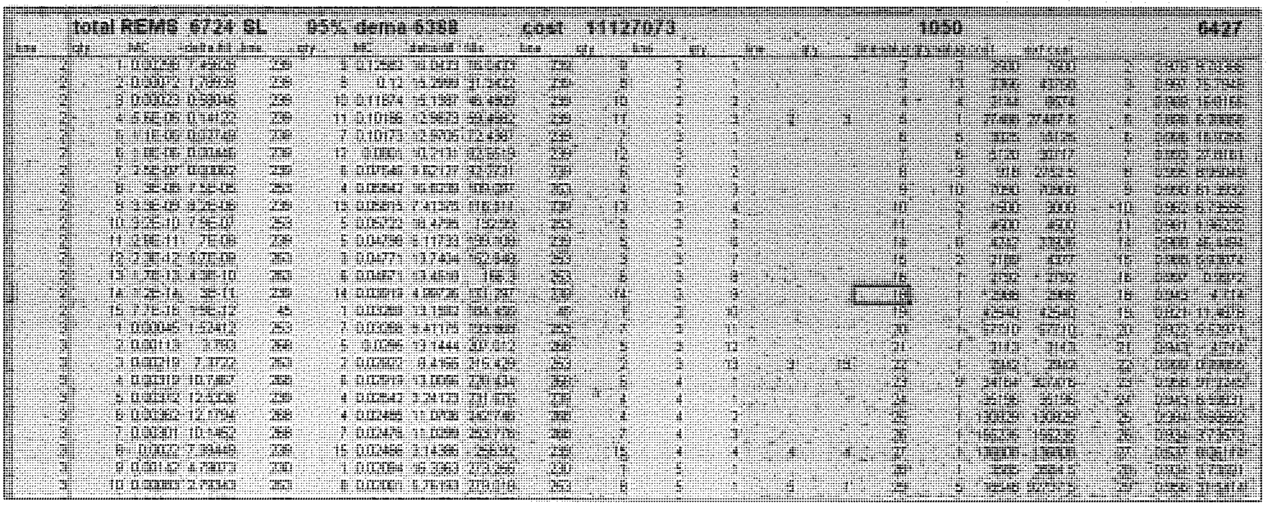

Figure A6.2: Model 2 spreadsheet sample

Figure A6.2 shows marginal contribution (MC) and incremental fill rate for quantities 1 to 15 for line items 2, 3, ... in the first four columns. These values are then sorted by descending (MC) in the next four columns (highlighted). Fills are then summed until target SL is reached (95% of total removals, 6724, is 6388). Line numbers and quantities are re-sorted (highlighted, quantity 1 repeats), then the maximum quantity for each line item is extracted to give the solution. The actual number of fills is 6427, which exceeds 6388 due to the logical flaw of this approach (quantities are not chosen in ascending order).

Model 3: Cost-Wise Skewed Holding

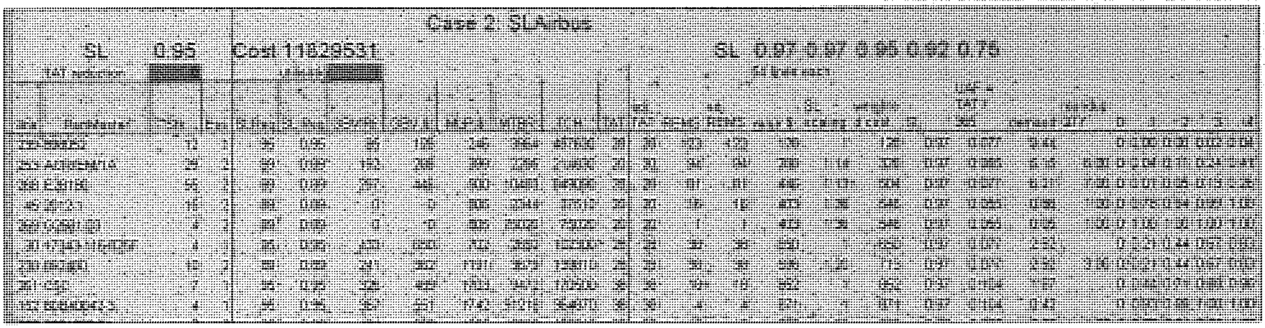

Figure A6.3: Model 3 spreadsheet sample

Figure A6.3 shows the first parts, which are in the lowest-cost band and have a SL target of 0.97. Remaining items, sorted by cost, will have decreasing SLs (0.75 for the highest-cost band) according to their cost band.

Model 4: Linear Programming

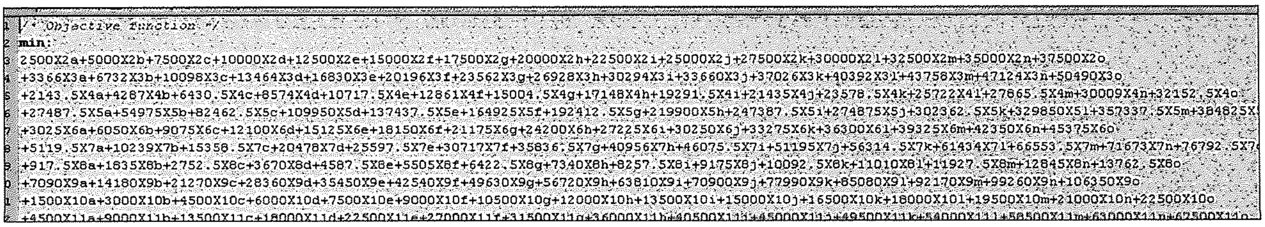

```
/* Objective function */
min:
2500X2a+5000X2b+7500X2c+10000X2d+12500X2e+15000X2f+17500X2g+20000X2h+22500X2i+25000X2j+27500X2k+30000X2l+32500X2m+35000X2n+37500X2o
+3366X3a+6732X3b+10098X3c+13464X3d+16830X3e+20196X3f+23562X3g+26928X3h+30294X3i+33660X3j+37026X3k+40392X3l+43758X3m+47124X3n+50490X3o
+2143.5X4a+4287X4b+6430.5X4c+8574X4d+10717.5X4e+12861X4f+15004.5X4g+17148X4h+19291.5X4i+21435X4j+23578.5X4k+25722X4l+27865.5X4m+30009X4n+32152.5X4o
+27487.5X5a+54975X5b+82462.5X5c+109950X5d+137437.5X5e+164925X5f+192412.5X5g+219900X5h+247387.5X5i+274875X5j+302362.5X5k+329850X5l+357337.5X5m+384825X
+3025X6a+6050X6b+9075X6c+12100X6d+15125X6e+18150X6f+21175X6g+24200X6h+27225X6i+30250X6j+33275X6k+36300X6l+39325X6m+42350X6n+45375X6o
+5119.5X7a+10239X7b+15358.5X7c+20478X7d+25597.5X7e+30717X7f+35836.5X7g+40956X7h+46075.5X7i+51195X7j+56314.5X7k+61434X7l+66553.5X7m+71673X7n+76792.5X7
+917.5X8a+1835X8b+2752.5X8c+3670X8d+4587.5X8e+5505X8f+6422.5X8g+7340X8h+8257.5X8i+9175X8j+10092.5X8k+11010X8l+11927.5X8m+12845X8n+13762.5X8o
+7090X9a+14180X9b+21270X9c+28360X9d+35450X9e+42540X9f+49630X9g+56720X9h+63810X9i+70900X9j+77990X9k+85080X9l+92170X9m+99260X9n+106350X9o
+1500X10a+3000X10b+4500X10c+6000X10d+7500X10e+9000X10f+10500X10g+12000X10h+13500X10i+15000X10j+16500X10k+18000X10l+19500X10m+21000X10n+22500X10o
+4500X11a+9000X11b+13500X11c+18000X11d+22500X11e+27000X11f+31500X11g+36000X11h+40500X11i+45000X11j+49500X11k+54000X11l+58500X11m+63000X11n+67500X11o
```

Figure A6.4: Model 4 objective function sample

Figure A6.4 shows the objective function for Model 4, which is to minimise the sum of extended costs (total cost), given that only one quantity can be chosen for each part number.

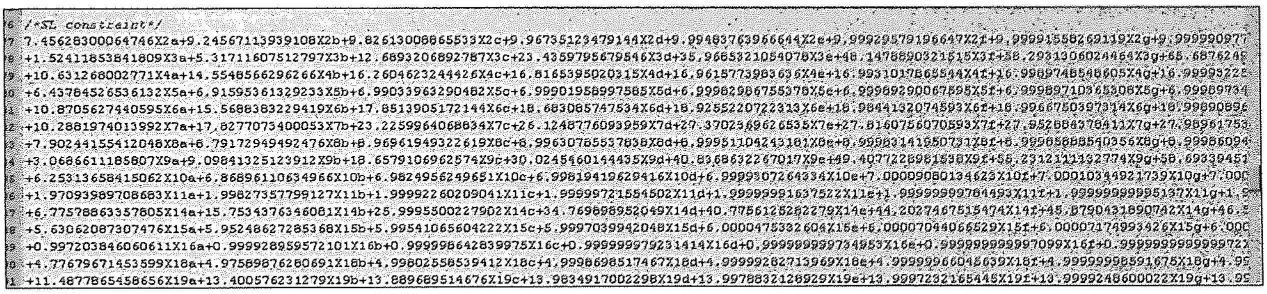

```
/*SL constraint*/
7.45628300064746X2a+9.24567113939108X2b+9.82613008665533X2c+9.96735123479144X2d+9.99483763966644X2e+9.99929579196647X2f+9.99991558269119X2g+9.99999097
+1.52411853841809X3a+5.31711607512797X3b+12.6893206892787X3c+23.4359795679546X3d+35.9665321054078X3e+48.1478890321515X3f+58.2931306024464X3g+65.68762
+10.631268002771X4a+14.5548566296266X4b+16.2604623244426X4c+16.8165395020315X4d+16.9615773983636X4e+16.9931017865544X4f+16.9989748548605X4g+16.999922
+6.43784526536132X5a+6.91595361329233X5b+6.99033963290482X5c+6.99901958997585X5d+6.99982986755378X5e+6.99989290067595X5f+6.99989710365308X5g+6.99989734
+10.8705627440595X6a+15.5688383229419X6b+17.8513905172144X6c+18.6830857475 34X6d+18.9255220722313X6e+18.9844132074593X6f+18.9966750397314X6g+18.99890896
+10.2881974013992X7a+17.6277073400053X7b+23.2259964068634X7c+26.1248776093959X7d+27.3702569626535X7e+27.8160756070593X7f+27.9526843784 11X7g+27.98961753
+7.90244155412048X8a+8.79172949492476X8b+8.96961949322619X8c+8.99630785537838X8d+8.99951104243181X8e+8.99983141950731X8f+8.99985688540356X8g+8.99986094
+3.0686611185807X9a+9.09841325123912X9b+18.6579106962574X9c+30.0245460144435X9d+40.8368632267017X9e+49.4077226981638X9f+55.2312111132774X9g+58.69339451
+6.25313658415062X10a+6.86896110634966X10b+6.9824956249651X10c+6.99819419629416X10d+6.99993072643 34X10e+7.00009080134623X10f+7.00010344921739X10g+7.000
+1.97093989708683X11a+1.99827357799127X11b+1.99992260209041X11c+1.99999721554502X11d+1.99999991637522X11e+1.99999999784493X11f+1.99999999995137X11g+1.9
+6.77578863357805X14a+15.7534376346081X14b+25.9995500227902X14c+34.7698989520 49X14d+40.7756125282279X14e+44.2027467515474X14f+45.8790431890742X14g+46.5
+5.63062087307476X15a+5.95248627285368X15b+5.99541065604222X15c+5.99970399420 48X15d+6.00004753326 04X15e+6.00007044066529X15f+6.00007174993426X15g+6.000
+0.997203846060611X16a+0.999928959572101X16b+0.99999864283997 5X16c+0.99999997923 1414X16d+0.99999999973 4953X16e+0.999999999997099X16f+0.999999999999972X
+4.77679671453599X18a+4.97589876260691X18b+4.99802558539412X18c+4.99986985174 67X18d+4.99999282713969X18e+4.99999966045639X18f+4.99999998591675X18g+4.99
+11.4877865458656X19a+13.400576231279X19b+13.889689514676X19c+13.9834917002298X19d+13.9978832128929X19e+13.9997232165445X19f+13.9999248600022X19g+13.99
```

Figure A6.5: Model 4 SL constraint sample

Figure A6.5 is the SL constraint for Model 4: the sum of all fills (subject to there being only one selected quantity for each part number) exceeds the target SL * total number of removals.

```
/*binary constraint*/
X2a+X2b+X2c+X2d+X2e+X2f+X2g+X2h+X2i+X2j+X2k+X2l+X2m+X2n+X2o=1;
X3a+X3b+X3c+X3d+X3e+X3f+X3g+X3h+X3i+X3j+X3k+X3l+X3m+X3n+X3o=1;
X4a+X4b+X4c+X4d+X4e+X4f+X4g+X4h+X4i+X4j+X4k+X4l+X4m+X4n+X4o=1;
X5a+X5b+X5c+X5d+X5e+X5f+X5g+X5h+X5i+X5j+X5k+X5l+X5m+X5n+X5o=1;
X6a+X6b+X6c+X6d+X6e+X6f+X6g+X6h+X6i+X6j+X6k+X6l+X6m+X6n+X6o=1;
X7a+X7b+X7c+X7d+X7e+X7f+X7g+X7h+X7i+X7j+X7k+X7l+X7m+X7n+X7o=1;
X8a+X8b+X8c+X8d+X8e+X8f+X8g+X8h+X8i+X8j+X8k+X8l+X8m+X8n+X8o=1;
X9a+X9b+X9c+X9d+X9e+X9f+X9g+X9h+X9i+X9j+X9k+X9l+X9m+X9n+X9o=1;
X10a+X10b+X10c+X10d+X10e+X10f+X10g+X10h+X10i+X10j+X10k+X10l+X10m+X10n+X10o=1;
X11a+X11b+X11c+X11d+X11e+X11f+X11g+X11h+X11i+X11j+X11k+X11l+X11m+X11n+X11o=1;
X14a+X14b+X14c+X14d+X14e+X14f+X14g+X14h+X14i+X14j+X14k+X14l+X14m+X14n+X14o=1;
X15a+X15b+X15c+X15d+X15e+X15f+X15g+X15h+X15i+X15j+X15k+X15l+X15m+X15n+X15o=1;
X16a+X16b+X16c+X16d+X16e+X16f+X16g+X16h+X16i+X16j+X16k+X16l+X16m+X16n+X16o=1;
X18a+X18b+X18c+X18d+X18e+X18f+X18g+X18h+X18i+X18j+X18k+X18l+X18m+X18n+X18o=1;
X19a+X19b+X19c+X19d+X19e+X19f+X19g+X19h+X19i+X19j+X19k+X19l+X19m+X19n+X19o=1;
X20a+X20b+X20c+X20d+X20e+X20f+X20g+X20h+X20i+X20j+X20k+X20l+X20m+X20n+X20o=1;
```

Figure A6.6: Model 4 binary constraint sample

Figure A6.6 is the binary constraint for Model 4: all of the variables for different quantities of the same line item must sum to 1.

```
/*integer declaration*/
int X2a,X2b,X2c,X2d,X2e,X2f,X2g,X2h,X2i,X2j,X2k,X2l,X2m,X2n,X2o;
int X3a,X3b,X3c,X3d,X3e,X3f,X3g,X3h,X3i,X3j,X3k,X3l,X3m,X3n,X3o;
int X4a,X4b,X4c,X4d,X4e,X4f,X4g,X4h,X4i,X4j,X4k,X4l,X4m,X4n,X4o;
int X5a,X5b,X5c,X5d,X5e,X5f,X5g,X5h,X5i,X5j,X5k,X5l,X5m,X5n,X5o;
int X6a,X6b,X6c,X6d,X6e,X6f,X6g,X6h,X6i,X6j,X6k,X6l,X6m,X6n,X6o;
int X7a,X7b,X7c,X7d,X7e,X7f,X7g,X7h,X7i,X7j,X7k,X7l,X7m,X7n,X7o;
int X8a,X8b,X8c,X8d,X8e,X8f,X8g,X8h,X8i,X8j,X8k,X8l,X8m,X8n,X8o;
int X9a,X9b,X9c,X9d,X9e,X9f,X9g,X9h,X9i,X9j,X9k,X9l,X9m,X9n,X9o;
int X10a,X10b,X10c,X10d,X10e,X10f,X10g,X10h,X10i,X10j,X10k,X10l,X10m,X10n,X10o;
int X11a,X11b,X11c,X11d,X11e,X11f,X11g,X11h,X11i,X11j,X11k,X11l,X11m,X11n,X11o;
int X14a,X14b,X14c,X14d,X14e,X14f,X14g,X14h,X14i,X14j,X14k,X14l,X14m,X14n,X14o;
int X15a,X15b,X15c,X15d,X15e,X15f,X15g,X15h,X15i,X15j,X15k,X15l,X15m,X15n,X15o;
int X16a,X16b,X16c,X16d,X16e,X16f,X16g,X16h,X16i,X16j,X16k,X16l,X16m,X16n,X16o;
int X18a,X18b,X18c,X18d,X18e,X18f,X18g,X18h,X18i,X18j,X18k,X18l,X18m,X18n,X18o;
int X19a,X19b,X19c,X19d,X19e,X19f,X19g,X19h,X19i,X19j,X19k,X19l,X19m,X19n,X19o;
int X20a,X20b,X20c,X20d,X20e,X20f,X20g,X20h,X20i,X20j,X20k,X20l,X20m,X20n,X20o;
int X21a,X21b,X21c,X21d,X21e,X21f,X21g,X21h,X21i,X21j,X21k,X21l,X21m,X21n,X21o;
int X22a,X22b,X22c,X22d,X22e,X22f,X22g,X22h,X22i,X22j,X22k,X22l,X22m,X22n,X22o;
int X23a,X23b,X23c,X23d,X23e,X23f,X23g,X23h,X23i,X23j,X23k,X23l,X23m,X23n,X23o;
int X24a,X24b,X24c,X24d,X24e,X24f,X24g,X24h,X24i,X24j,X24k,X24l,X24m,X24n,X24o;
int X25a,X25b,X25c,X25d,X25e,X25f,X25g,X25h,X25i,X25j,X25k,X25l,X25m,X25n,X25o;
int X26a,X26b,X26c,X26d,X26e,X26f,X26g,X26h,X26i,X26j,X26k,X26l,X26m,X26n,X26o;
int X27a,X27b,X27c,X27d,X27e,X27f,X27g,X27h,X27i,X27j,X27k,X27l,X27m,X27n,X27o;
int X28a,X28b,X28c,X28d,X28e,X28f,X28g,X28h,X28i,X28j,X28k,X28l,X28m,X28n,X28o;
int X29a,X29b,X29c,X29d,X29e,X29f,X29g,X29h,X29i,X29j,X29k,X29l,X29m,X29n,X29o;
```

Figure A6.7: Model 4 integer declaration sample

Figure A6.7 is the integer declaration for Model 4, which together with the binary constraint ensures that exactly one quantity will be selected for each line item.

```
Variables;MILP Better;MILP Better;MILP Better;MILP Better;MILP Better;MILP Better;MILP B
;314117276027.354;10311087.9999997;10307204.9999997;10306992.9999997;10306969.499
X100o;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X101e;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X102a;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X103a;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X104h;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X105c;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X106b;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;1;1;1
X107b;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X108c;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;0;1;1;1
X109b;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X10b;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;0;1;1;0;0;0;0;0;0;0;0;0;0;0;0;0;0;1;1;1;1;1;1;1;1
X110a;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X111e;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X112e;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X113e;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;0;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X114e;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X115a;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X116b;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X117h;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X118d;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X119e;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
X11a;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1;1
```

Figure A6.8: Model 4 LP solution sample

Figure A6.8 is the solution of Model 4: only variables with values of 1 are shown in this sorted list, so that X100o means that line item 100 should have a quantity of 15. The 0 and 1 values show the solution set iterations, so that X106b was only in the last three solutions before the optimal solution was reached.

Model 5: Linear Programming – split

Model 5 looks the same as Model 4, the difference being that the problem set is split into three groups by essentiality. The separate solution sets are then combined to give the overall solution.