CALCULATING CRITICAL STOCKHOLDING FOR SPARES

JOHN WOODHOUSE

Managing Director,
The Woodhouse Partnership Limited

Introduction

There is a great deal of energy going into supply chain alliances, e-purchasing opportunities and new 'methodologies' (such as JIT, MRP, EAP, EAM etc). Most of these initiatives are focussed on *efficiency* improvement in response to continued pressure to reduce working capital, maintain production continuity and increase output or service levels. What is often missed, however, are the fundamental decisions about what is worth doing or holding in the first place – the right compromise between costs (or capital tied up) and risk exposures (what would happen in event of stockout or supply chain interruption). Getting these decisions right is key – and most of us are getting the balance wrong.

This paper is all about getting back to some basics – i.e. what comprises a good decision and how we can evaluate and compare options on a value-for-money basis. It covers the conclusions and deliverables from a European 'best practice' collaboration project over the last 5 years. This project ("MACRO"[1]) was convened to tackle the core problems of cost/risk/performance trade-off: how do we judge what to spend, when and how, if the available hard data is inadequate. The total project was much wider than just the spares, materials and purchasing areas – it also covered the capital investment, life cycle costing, project evaluation, maintenance, inspection and shutdown optimisation areas.

As in each of the other MACRO areas, the field of spares, materials and purchasing decisions revealed some radical opportunities for improvement. Just developing a structured way of *ensuring that the right questions are asked* made a big difference. In fact the best practice solutions went much further on both the 'soft issues' (the human factors involved) and the hard number-crunching that is needed to determine risks or probability impact and the optimal overall strategies.

One of the very first issues to be faced, however, is the fundamental difference in treatment between slow-moving or 'strategic' inventory ('insurance spares') and the 'normal' consumables and materials flow. Although the final MACRO output provides seamless management of these two families, the underlying concepts and evaluation methods must be quite different.

---

[1] Project EU1488 see www.aptools.co.uk

Basically the **slow-moving items** are characterised by having:

- High purchase cost
- Low likelihood of usage (demand less than 1-5 occasions per year)
- Unpredictable (i.e. random) demand
- Long replacement leadtimes
- Often high operation impact/cost if not available when needed
- A case-by-case re-order policy (when/if each one is used)

Consumables, on the other hand, tend to exhibit quite different features: they are

- Usually low-cost items (with discounts for bulk purchase)
- Needed more often (measured in units per week or month)
- Subject to *patterns* of demand (i.e. not random)
- Available in short leadtimes (hours/days/weeks)
- Less critical, with unavailability impact representing 'hassle' or expediting costs rather than major operational downtime.
- Re-ordered in batches (based on a min/max, EOQ or similar policy)

The first group (slow-moving items) generates decision headaches due to a combination of big numbers (costs, leadtimes, potential operational impact) and great uncertainties (likelihood of needing the item and the size of the potential operational impact). The second group provides further headaches through data overload, lack of consistency/rationalising and through sheer volume of line items. To make quantum improvements, therefore, we must consider the existing headaches, and the emerging opportunities for preventing or mitigating them.

Any simple stock report will show that inventory fits a Pareto distribution: c.80% of stock value is tied up in about 20% of the items. This closely follows the spread of slow- and fast-moving items: the fewer, high-cost but slow-moving items incur disproportionate ownership costs compared to the 'small and many'. For the fewer, high cost cases (which are also the slow-moving, high criticality, data uncertainty items), we can afford to do more extensive, structured "what if?" investigation of the right decisions – evaluating the levels of risk, options for supplier-held spares, pooling and interchangeability, and explore the sensitivities to data assumptions etc. For the bulk consumables we need to go further still, and manage/navigate the data volume, taking account of *patterns* and exploring alternative purchasing points, quantities and management strategies.

These two families are now explored in greater detail individually:

Evaluation of Slow-Moving Spares

The MACRO project has found that approaches used to make slow-moving spares decisions fall into four categories. From an evolutionary perspective, these are:

Evaluation approaches:

1. **Manufacturers' recommendations** (sometimes 'filtered' by engineering judgement);
2. Economic Order Quantities (**EOQ**'s) and '**rules-of-thumb**' approaches (*e.g.* $\sqrt{N} + 1$, where N is number if installed units!)
3. **Target availability calculations**: probabilistic, often simulation-based and geared to achieve a pre-determined service level, such as 90, 95 or 99%
4. **Optimum Total Cost & Risk** calculations: also probabilistic, but putting a price on the impact of unavailability - so the right (optimal) service level can be identified.

To accurately model these spares decisions, MACRO has identified the factors that must be taken into account in the evaluation. These factors are listed in Table 1, together with indicators to show if are specifically addressed by each of the four approaches mentioned above (Y= included, ?=sometimes given subjective consideration):

Factor — **Evaluation Approach:**	**1**	**2**	**3**	**4**
Demand rate	?	?	Y	Y
Number of operational items in use (supported by the spare)	Y	Y	Y	Y
Number of operational items needed (i.e. installed spare capacity)	?		Y	Y
Operating usage (e.g. hours per year)			Y	Y
Impact of unavailability (downtime impact, lost oppty costs etc.)				Y
Supplier's replacement leadtime	Y	?	Y	Y
Emergency replacement leadtime (downtime costs being incurred)				Y
Chance of refurbishing a failed unit			Y	Y
Refurbishment leadtime			Y	Y
Purchase costs	Y	?	Y	Y
Holding (storage & maintenance) costs			Y	Y
Constrained useful life (projects, shelf life, technology change)				Y

Table 1. Factors addressed by the four approaches to spares evaluation.

From this, the failings of each approach are evident and it is clear that the only approach that embraces all the factors that must be taken into account is approach number **4 - Optimum Total Cost & Risk** calculations. This is the approach that was therefore adopted in the MACRO project (and built into the deliverable tools 'APT-SPARES' and 'APT-STOCK').

The 'Optimum' Spares Holding

The *optimum* spares holding is that level of spares which incurs minimum total impact to the business as a whole *i.e.* the minimum combination of holding costs (*i.e.* storage costs, in-storage maintenance costs etc.) and the risk exposures to the consequences of spares unavailability when needed (usually equipment or production downtime costs).

Slow-moving spares are generally responsible for a large part of the inventory value of any organisation. They form an important part of the inventory management responsibility. The benefits of holding spares, however, lie in operational areas (production or maintenance) - outside the scope or view of warehouse or inventory management functions. To relate the benefits to the ownership costs involved, therefore, it is necessary to consider the impact of alternative holding strategies on the total business profitability, rather than just from a stores (as fewer spares as possible) or operations (as many spares as possible) point of view. This *total business impact* is a combination of costs and risked costs. When this combination is at a minimum, the optimum spares holding has been identified.

While costs such as the purchase cost, storage costs and in-storage maintenance costs can be determined fairly easily (they are actual costs incurred), downtime impact and potential write-off costs (if the spare is not used) are, by comparison, more difficult. Such probabilistic or *risked costs* represent average costs incurred over a long period *i.e.* a combination of no costs for the majority of time, and an occasional large cost event. APT-SPARES is a software tool developed by the MACRO collaborating companies to calculate these risk exposures for each level of spares holding and the optimal combined strategy.

APT-SPARES

APT-SPARES calculates which spares to hold, in what numbers, on either an item-by-item or a batch-review basis. By finding the right combination of ownership costs and risk exposure, this 'what-if?' method can evaluate:

- Opportunities to reduce the spares holding or, if the risk exposure (to the penalties of unavailability) is unacceptable, the necessary additional spares;
- The impact of under- or overstocking.
- The sensitivity of the recommendation to assumptions and base data.
- Supplier options (*i.e.* Supplier A versus Supplier B, offering different prices/leadtimes).
- Spares 'pooling' options (sharing access to common spares, or dedicating)
- Whole unit, modules or component stock options

APT-SPARES was specifically developed to take account of the inherent uncertainties in data (particularly evident in failure rates and downtime consequences). It reveals how much the current uncertainty is costing and, thereby, how much it is worth spending to collect better data. Subsequently, the tool was developed to provide automatic sensitivity testing and other "what if?" features.

Modelling Assumptions

The assumptions employed by APT-SPARES can be conceptualised as a *re-supply chain* that details various routes for the use and replacement of spares. Not just one re-supply chain (which is typical of older, simplistic models using 'queuing theory'), but several, representing normal replacement, emergency/expediting options, workshop repair etc. These routes naturally offer different timescales, costs and risks (*i.e*. the possibilities of repair or of unfulfilled demand for further spares). For a given a level of random demand (the usual case if the spare is a strategic, 'insurance' item), there will be a probability of various spares being at various stages in each of these re-supply chains. APT-SPARES models the movement of spares between these stages by using the relevant components of queuing theory and some advanced probability calculations (to handle the interactions between re-supply chains and avoid double-counting costs or risks). The results are expressed in simple financial terms - the combination of risk exposures and costs for each possible strategy. The analysis can be performed in a manual "what if?" mode, or through a batch review of multiple stock lines, offering global "what if?" capability. The APT-STOCK version even offers automatic sensitivity analysis and the identification of assumptions needed to raise or lower the inventory levels.

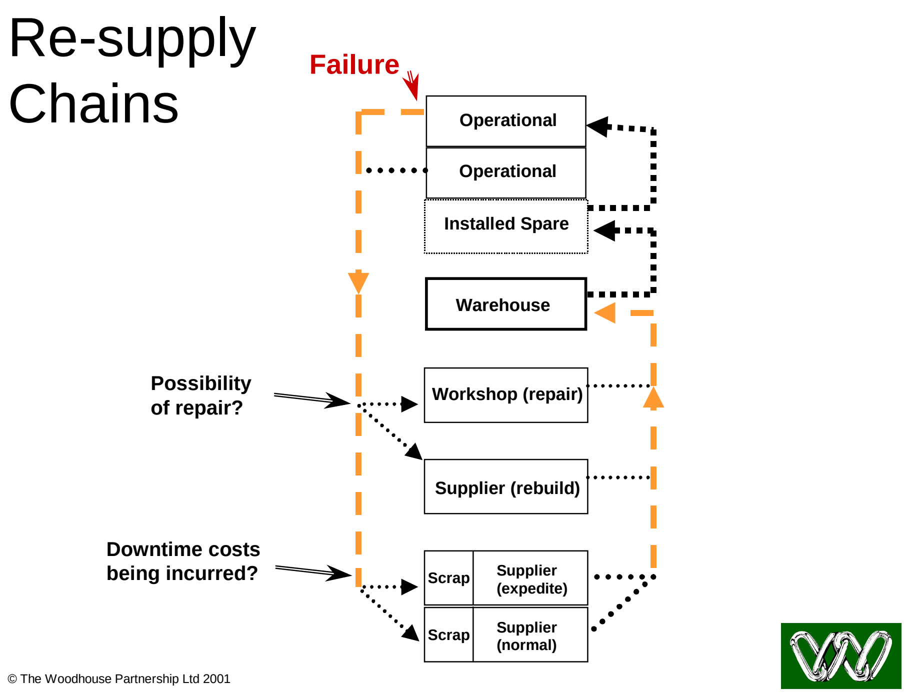

Figure 1.1 Various supply chains of different 'lengths'

Detailed Case Study: Water Pump Motor

APT-SPARES was employed by a large UK company to evaluate whether or not its current policy of holding 1 single spare water pump motor was adequate.

The total purchase price of each pump was £85,000, subject to a replacement leadtime of 6 months. It was anticipated, however, that one could be expedited in 3 months in event of emergency. The pump operates continuously and one spare is held. Failure history for this particular type of pump was very limited. Records for similar pumps were scanty but the engineers estimated an average failure rate of about 1 every 5 years (but stated possible levels between 1-in-3 and 1-in-10 years). Failure of a pump would result in a proportional loss of production operation, with a consequential impact estimated in the range of £3-5,000 per hour. The operators were confident that around 80% of the failed units could be repaired and converted into future replacement spares. The average total repair time in such cases was estimated at 8 weeks. The cost of money tied up in any spares, and the annual costs for storage and maintenance, were provided by the finance and warehouse departments at 12% and 5% of the capital value respectively.

This information represented the 'base' case, and was entered into APT-SPARES as shown in Figure 1.2

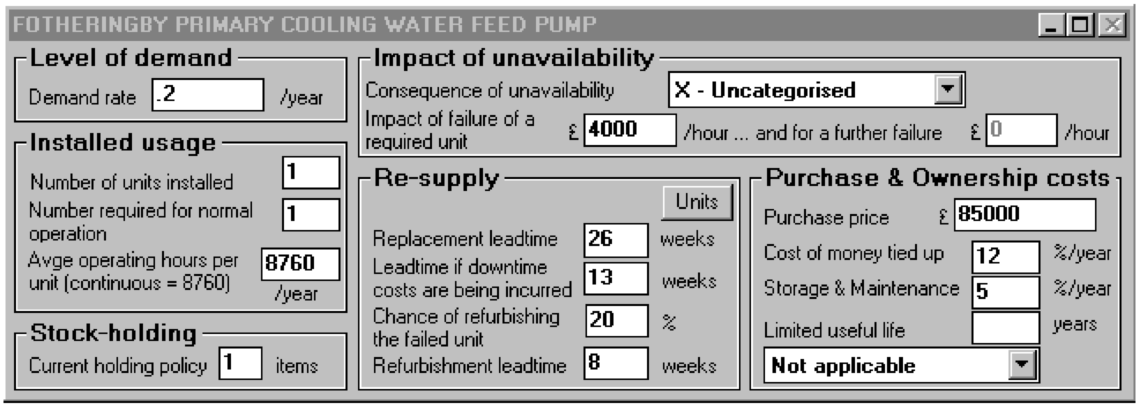

Figure 1.2 APT-SPARES analysis inputs.

The optimum spares policy was calculated and the results displayed in a results table. As shown in Figure 1.3, this displays the annual downtime risks, money tied up, storage and maintenance costs and write-off costs for numerous spares holdings. From this, it can be seen that the optimum spares holding is 2, with a total cost/risk impact of £30,358/year. Compared the current policy of a single spare held (which has a total impact of £96,166/year), this represents a **<u>net</u> improvement of more than £65,000/year**.

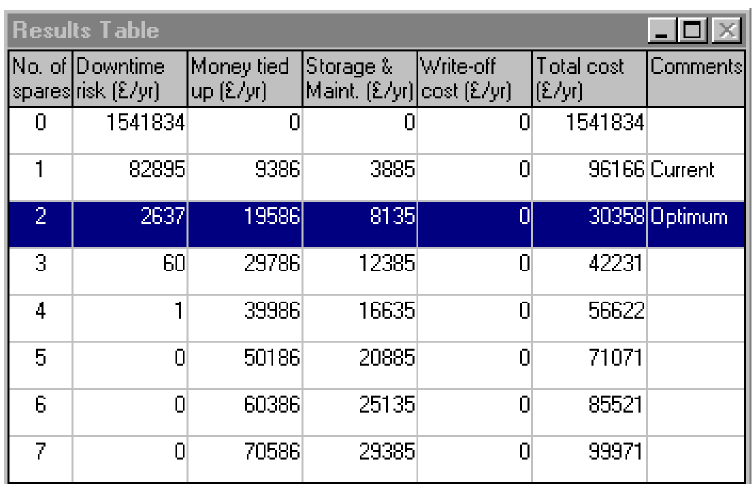

Results Table

No. of spares	Downtime risk (£/yr)	Money tied up (£/yr)	Storage & Maint. (£/yr)	Write-off cost (£/yr)	Total cost (£/yr)	Comments
0	1541834	0	0	0	1541834	
1	82895	9386	3885	0	96166	Current
2	2637	19586	8135	0	30358	Optimum
3	60	29786	12385	0	42231	
4	1	39986	16635	0	56622	
5	0	50186	20885	0	71071	
6	0	60386	25135	0	85521	
7	0	70586	29385	0	99971	

Figure 1.3 APT-SPARES results table.

In addition to a tabular format, the results were displayed as a 3-dimensional graph. This is shown in Figure 1.4 below.

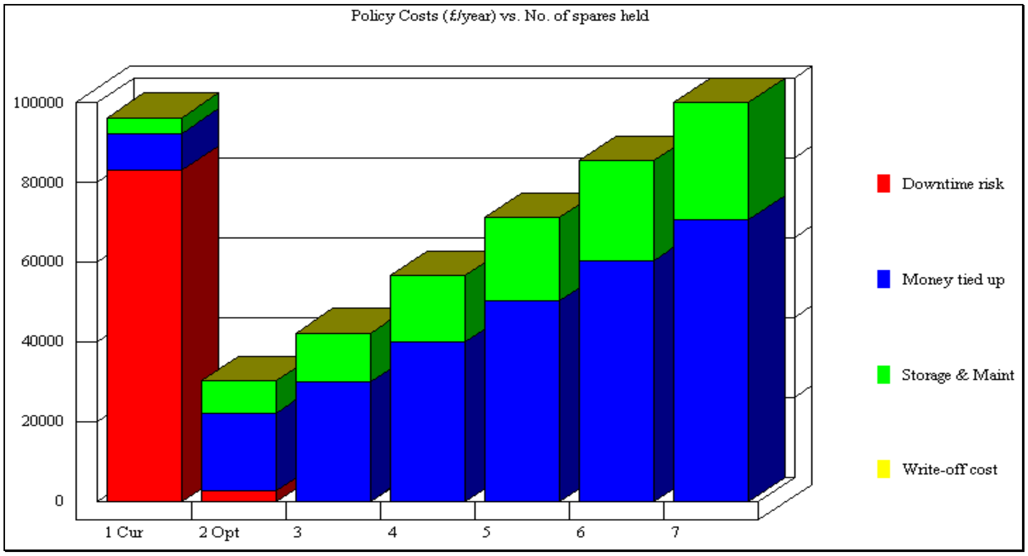

Figure 1.4 APT-SPARES results graph.

Notice the combination of costs and risks incurred by the current policy (one spare). The risk shrinks greatly by the acquisition of a second spare, and the increased ownership costs are small by comparison. In other evaluations, the reverse effect can be revealed (i.e. the current policy incurs too great an ownership cost and a higher level of risk is worth taking). APT-SPARES makes some observations about each analysis and presents them as conclusions as shown in Figure 1.5.

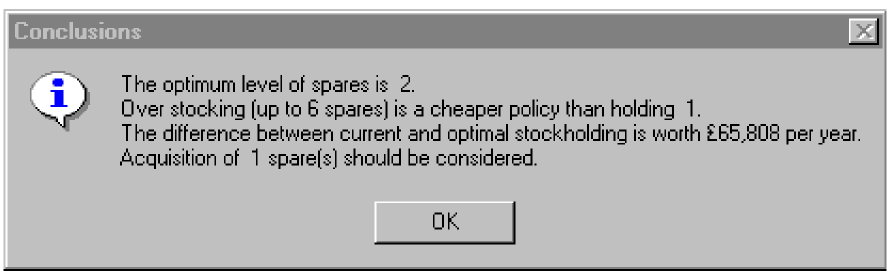

Figure 1.5 Initial analysis conclusions.

Using uncertain data – sensitivity analysis

After determining such a 'base case', APT-SPARES was employed in a 'what-if?' role to test the sensitivity of the outcome to uncertainties in data. Although the base case utilised a failure rate of 1 in every 5 years, it was considered that the actual value could lie anywhere between 1-in-10 to 1-in-3 years. In addition, uncertainties were also expressed about the value of production downtime that could be affected. This value could lie anywhere between £3,000 and £5,000 per hour.

From this, 'best' (*i.e.* fewer spares needed due to a lower demand rate and lower downtime costs) and 'worse' (*i.e.* more spares needed due to a higher demand rate and higher downtime costs) cases were developed. These are detailed in Table 2.

Case	Inputs		Results	
	Demand Rate	**Downtime Value**	**Optimum Stock**	**Total Cost**
Best	0.1/yr	£3,000/hr	**2 spares**	£28,569/yr
Likely	0.2/yr	£4,000/hr	**2 spares**	£30,358/yr
Worst	0.33/yr	£5,000/hr	**2 spares**	£40,960/yr

Table 2. Best, base and worst cases.

This reveals that the outcome was *not sensitive* to the uncertainty in demand rate and downtime value, and that despite the range estimates, more accurate data is not required.

The scenario was taken one stage further to establish the demand rate for which the current policy (just one spare) would be justified. By a sequence of automatic "what if?" calculations, this was found to be 0.08 (*i.e.* 1 failure in every 13 years) – the level of risk low enough for a single spare to be adequate. This level of reliability was dismissed as overly optimistic. Meanwhile, a second spare pump was ordered, with a most likely level of net savings worth over **£65,000/year.**

Consumables policies

The day-to-day administration of warehouse inventory, supplier chain issues and demand backlog is dominated by the 'small and frequent' variety of materials. This is why purchasing and inventory management 'systems' have gained such high profile – they get the paperwork under control. But they also frequently represent a case of 'the tail wagging the dog': they make the administration more efficient, but do not question what stock or inventory should be held in the first place. They very rarely even record the consequences of NOT having items in stock when they are needed. A simple counting of 'stockouts' does very little to help determine what money is worth spending to avoid them!

To set the scene for a more rigorous approach to inventory decisions, we need to organise some of the options and potential influences first. They include:

Purchasing policies: the 'context'

- Are orders placed on a min/max basis, on a regular cycle, with variable or pre-determined order quantities and with/without allowing pre-assignment ('backordering')?

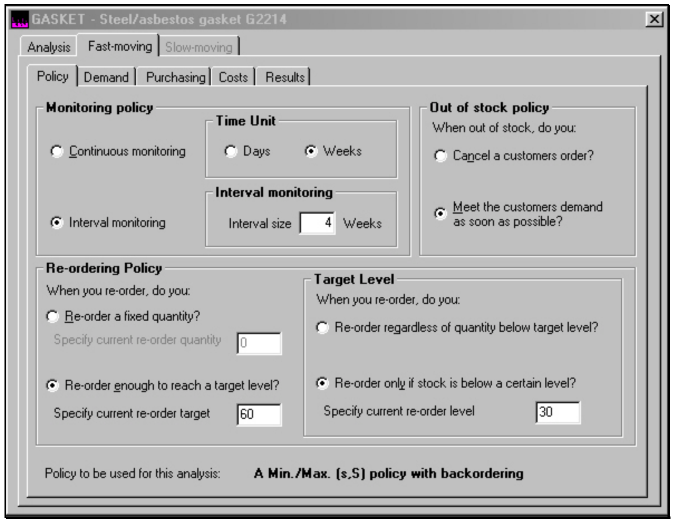

Figure 2.1 Some of the inventory/purchasing policy options

Demand patterns and importance: the 'outbound'

- Is the usage rate fixed, random, 'clustered', seasonal, or any other sort of pattern?
- What are the consequences of unavailability, definitions of 'criticality' etc?
- What alternatives are there for repair, modification or interchangeability if no stock is available?

Supply options and variability: the 'inbound'

- Purchase quantity constraints, bulk discounts, e-purchasing opportunities, supplier-held stock, sharing or 'pooling' options, leadtimes (fixed, variable, patterns) etc.

Inventory management costs: the 'buffer'

- Purchase order costs, expediting, warehousing costs, in-storage maintenance, cost of capital tied up

These issues are not new and should be familiar to most parts of the supply chain. However they are not generally being taken into proper account – we tend to quantify the easiest to quantify, and try to avoid the nastier questions (such as "what would be the $$ impact of a stockout?" or "what is the risk associated with this leadtime not being fulfilled?"). Putting numbers to all of these factors, *including the uncertainties*, can sort out which items have what effect upon the optimal solution, and what data is needed for collection in the future.

APT-STOCK

The number-crunching, in these cases, is beyond the abilities of a single set of probability maths. – the potential permutations of patterns, uncertainties, weighted cost influences etc requires a simulation approach. This involves describing the various factors and their variabilities (or uncertainties), then simulating the interactions of demand, resupply, leadtimes, costs etc. tens of thousands of times to observe which combination of assumptions is the best (results in minimum total business impact). APT-STOCK was commissioned by the MACRO project for this purpose, and it uses state-of-the-art simulation technology for all the possible inventory and purchasing policies. In a single analysis, it will explore batches of about 50 combinations of, for example, reorder point and reorder quantity (or min/max or reorder cycle/quantity etc), *each combination* explored for the equivalent of 200 years of operation, and all 50 possible strategies analysed in less than 2 seconds! This sort of performance allows us to do wide-ranging "what if?" studies, and find out what data is critical and what makes no difference to our strategies.

The data volume problem

The technology has even been designed to interface to external data sources (such as stock records) and run in batch mode – in one customised case, the method reviews 28,000 stock lines automatically every month, and displays suggestions for adjustment (for manager review/authorisation). Of course this level of embedded optimisation needs some procedural customisation each time (the 'rules' used to fill the inevitable gaps in local data, how to handle apparent 'spikes' of demand or cost, the review/approval process and the IT interfacing). However, with case studies such as 60% inventory savings under our belts, this sort of set-up investment pales into insignificance!

Case study: Gaskets – optimal inventory & e-procurement

This is a typical case of consumables used in various locations for various purposes. A general consumption pattern has emerged, with an average of 6 units per week, but this varies quite widely with equipment shutdowns, major overhauls and other 'spikes'. A 'per unit' cost has been negotiated with the supplier on the basis of this consumption (but alternative bulk discounts need to be explored). The supplier is not totally reliable – his normal 4 week delivery sometimes extends to 6, or even occasionally 12 weeks if there are production or import problems. Stockout consequences, however, are quite low – we can always make up a suitable temporary gasket in the workshops (at a premium price, but it avoids operational impact).

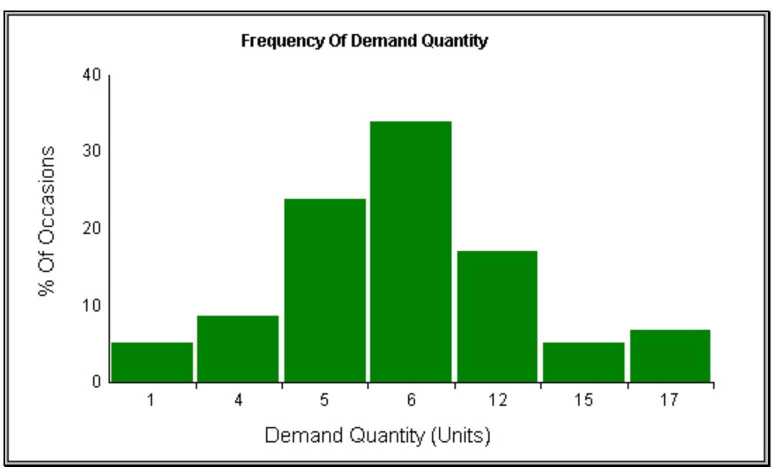

Figure 2.2 Distribution of demand (units per week) for the gaskets

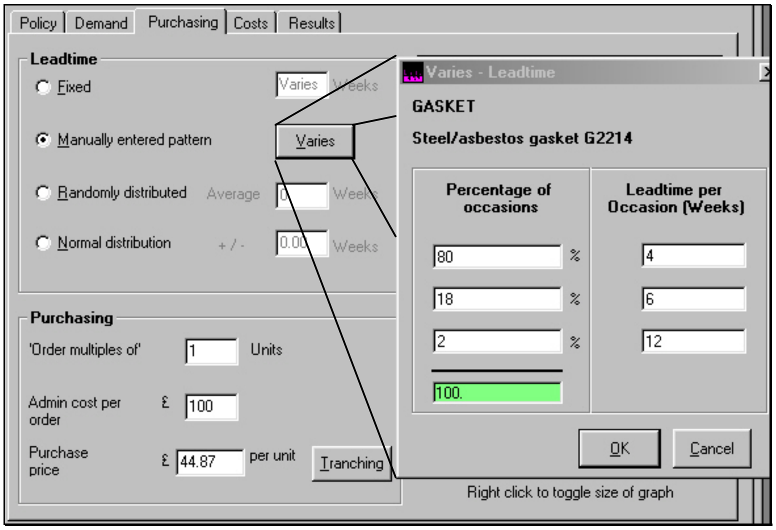

Figure 2.3 Description of uncertain leadtimes, purchase costs etc.

This information is entered into the model, with appropriate corporate rates for cost of capital tied up, the cost of raising purchase orders etc. Then the simulations are performed – exploring different combinations of *when* the orders are placed and for *how many*. The optimum combination of costs and stock-out risks is identified and highlighted in blue: this represents the *optimum service level* – the right amount of risk to take in view of the costs involved.

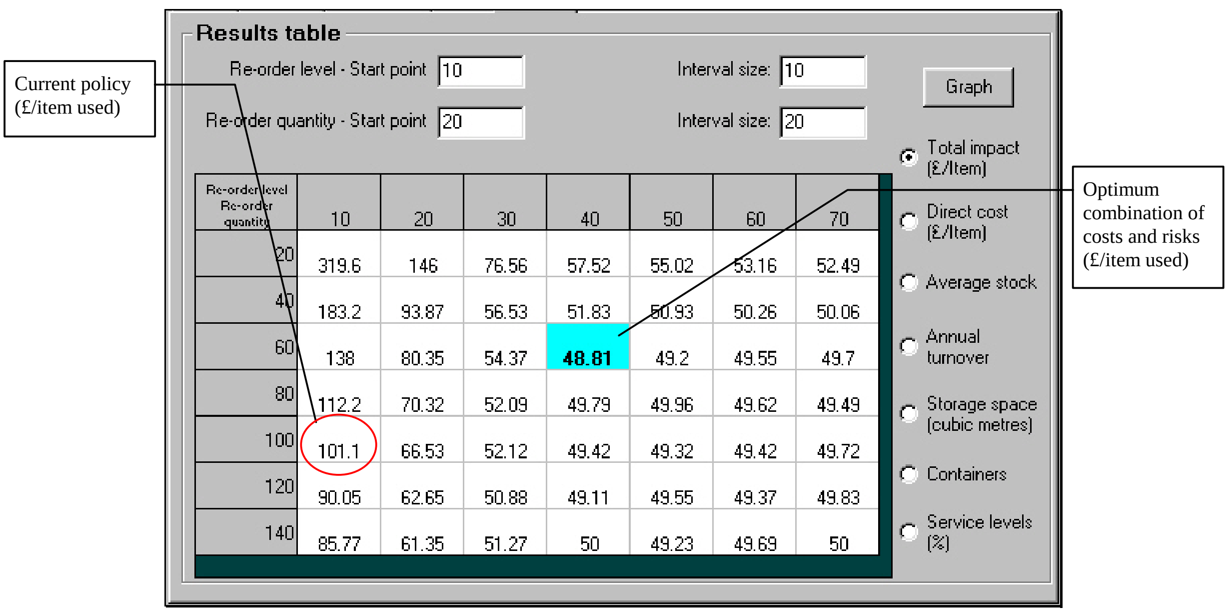

Figure 2.4 Results table (total impact view)

The initial results show that the optimal policy is to place orders when stock drops to 40 units, and reorder in quantities of 60. This contrasts with the current policy of a reorder level of 10 and order quantity of 100 (i.e. rarer, bigger orders). The difference between these policies is (£101.10 – £48.81) x 6 units/week = **£16,300/year**. And that is just one small gasket decision!

Note also that, in this case, overstocking/ordering is cheaper than understocking/ordering (the cost per unit rises more slowly toward the bottom right of the table, that towards the risky/expensive left and top).

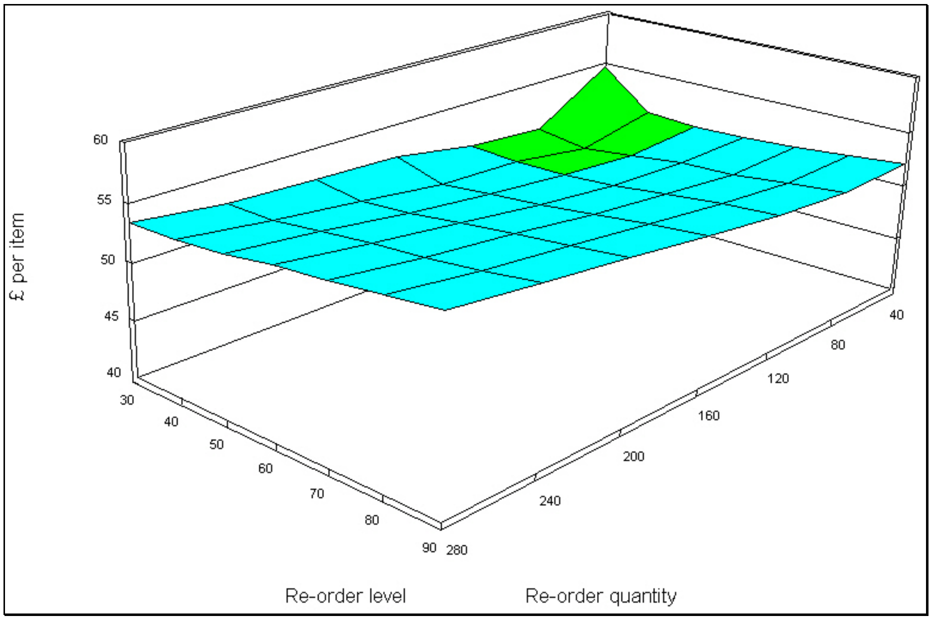

Figure 2.5 Results graph, showing best value-for-money area (darker/green shading).

The key to this optimisation, of course, is the inclusion of risk into the equation. The above results take account of both the purchase/holding costs, *and* the likelihood and consequences of having insufficient stock to meet demand. This is revealed if we examine the Service Levels associated with different policies: the optimum result has a service level of 99.7% whereas the current policy meets only 86.55% of demands from stock.

Re-order level / Re-order quantity	10	20	30	40	50	60	70
20	36.32	66.32	89.99	97.76	99.36	99.6	99.85
40	65.97	83.18	95.42	99.41	99.76	99.87	99.99
60	77.68	88.88	97.13	**99.7**	99.86	99.99	100
80	83.23	91.63	97.79	99.77	99.81	100	100
100	86.55	93.24	98.11	99.58	100	99.96	99.97
120	89.04	94.56	98.4	99.77	99.97	99.99	99.92
140	90.29	94.95	98.66	99.76	99.95	99.98	99.98

- Direct cost (£/Item)
- Average stock
- Annual turnover
- Storage space (cubic metres)
- Containers
- Service levels (%)

Figure 2.6 Service Levels achieved by different purchasing policies

From this point, there are many avenues to explore – the impact on average stock levels, the turnover, storage space required etc (see options on the right of these tables). However the results also give us the scope for analysing alternative scenarios, such as "what if we developed an e-purchasing agreement for this item?". In such a case, the unit purchase cost should go down, so would the overhead cost of each purchase order. The leadtimes might be also affected (either way – it depends on the degree of shared access to the supplies and the logistics of the supplier) and the supplier may offer different bulk quantity discounts. If, for example, the average purchase price dropped from £44 to £38 per unit, and the overhead reduced from £100 to £40 per purchase order, then the inventory strategy is re-optimised to reveal that the order quantity would rise to 100 but the reorder point should remain at 40. The net further savings, compared to the earlier optimum result, amount to (£48.81 - £42.43) x 6 units/week = £1,990/year.

Re-order level / Re-order quantity	10	20	30	40	50	60	70
20	176	89.21	55.96	46.45	45.33	45.03	45.67
40	108.8	63.96	46.52	43.19	43.1	43.31	43.63
60	84.2	54.21	43.91	43.22	42.57	43.2	43.44
80	74.95	51.66	43.74	42.64	42.66	43.09	43.61
100	66.48	49.59	44.39	**42.43**	43.08	43.35	43.81
120	63.59	50.44	44.74	42.93	43.23	43.85	44.19
140	61.19	47.89	43.63	43.5	44.16	44.11	44.56

Total impact (£/Item)
Direct cost (£/Item)
Average stock
Annual turnover
Storage space (cubic metres)
Containers
Service levels (%)

Figure 2.7 Re-optimised results for e-purchasing option

Conclusions

The scientific treatment of inventory and purchasing decisions holds great scope for cost/performance improvement. The MACRO project conclusively showed that data availability is not the key issue – it is *what we do with the data* that matters. Modern, high-speed computer tools can enable us to explore, in "what if?" mode, all the strategies, risks, uncertainties and critical influences on inventory decisions. The feedback from those companies who have 'rolled out' such techniques is staggering. As a bonus, the methods have provided a forum for buyers, suppliers, holders and users to negotiate their inevitably conflicting interests. When we continue to measure and reward inventory managers by 'how little they hold' (i.e. reductions in inventory value), yet simultaneously pressure the maintenance department to fix operational plant without delays, we need a means of finding the right compromise. The MACRO project has successfully generated these procedures, the underlying maths, and some quite handy "what if?" evaluators to put numbers to it all!

John Woodhouse
john.woodhouse@twpl.co.uk
February 2003

APPENDIX

MACRO Project summary

THE MACRO PROJECT

Cost/Risk Evaluation of Asset Management decisions

The European multi-industry collaboration project to collate best practices and develop risk-based methods for the management of physical assets.

EEC Eureka Project EU1488

What is the MACRO project?

The MACRO project is a 3-year, £1.3 Million programme of investment in the tools and methods for making better decisions. Leading European companies are collaborating to define the correct Asset Management strategy.

The results combine the best of structured thinking and range estimating methods with reliability and risk optimisation techniques into a completely new way of handling decisions when data is poor or uncertain.

Target Areas

Project Evaluation, Asset Replacement & Life Cycle Costing;

Cost/benefit evaluation and prioritising of modifications, project screening, capex/opex trade-off, repair versus replacement, life extension & refurbishment options.

Operating, Maintenance & Inspection Strategies;

Shutdown programming, inspection, test and maintenance intervals, optimisation of efficiency, reliability, performance and equipment lifespan, condition monitoring strategy and methods.

Resourcing & Materials Strategies;

Spares and materials stock levels, supply chain decisions, min/max policies, logistics and warehousing strategies.

Who is involved?
Anglesey Aluminium
Asset Performance Tools
Brown & Root Energy Services*
Det Norske Veritas
ICI Eutech
Institute of Asset Management
Intevep
National Power
Railtrack
Sector
Shell Norway
The National Grid Company*
The Woodhouse Partnership*
UK Government (DTI)*
Yorkshire Electricity*
* Core Sponsors