

Professional Inventory Management

Geoff Relph, MIOM, Inventory Matters Ltd, Witek Brzeski, IBM UK Ltd, and Gail Bradbear, Ford Motor Company

INTRODUCTION

It can be difficult to determine when stock is at the optimum level as this is influenced by numerous factors including batching policies, the work centre capability, safety policies and production methods. The result is that despite the use of sophisticated MRP / ERP systems, manufacturers typically keep overall inventory levels too high without really achieving high parts availability. In many cases the response in production to poor availability is to over produce with intent to safeguard against future shortages. This paper is the second in a series of three that looks at the alternative options available to optimise inventory held at an aggregate and part level by considering the steps necessary to move towards inventory optimisation. The intent of these papers is to provide a guide to enable managers within manufacturing industries to advance their organisation to the next level.

In our first paper *“The First Steps to Inventory Management”* we looked at the basics beginning with the concept of control and a simple ABC classification system. This paper builds on the concepts developed in *“The First Steps to Inventory Management”* and assumes that inventory is under control and the basics have been established by implementing a 3-class (ABC) system. Here we develop the options for managers to improve inventory performance by introducing the concept of overage, extending inventory analysis to 6 or more classes and considering the use and calculation of safety stock. The third and last paper *“Advanced Inventory*

Management” looks at K-Curve theory and the impact of system parameters on stock levels. Each paper is structured along similar lines and contains background theory, the requirements for operating at the level under discussion, a worked example using data from the manufacturing industry and a concluding summary.

GOOD AND BAD INVENTORY

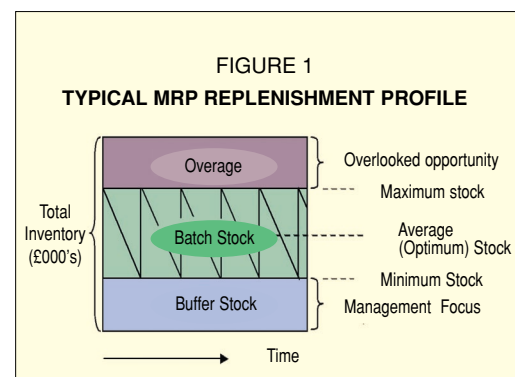
In the first paper we introduced a simple view of inventory consisting of batch and buffer stock. Here we add a further concept, overage, as shown in Figure 1.

Total inventory can thus be viewed as composed of three elements; batch, safety and overage, each of which is required or occurs for a different reason and necessitates different techniques to manage. At a high level inventory can also be viewed as good or bad for the business. Good inventory is that which is required to run the business, while bad inventory is the difference between the actual and the maximum inventory.

Too little inventory, also known as an underage, is any stock deficit below the minimum planned level and can impact manufacturing efficiencies and customer service levels. We discuss later how to manage the risk of underage through a planned use of safety stock. **Overage** is bad inventory in the form of excess stock above the maximum planned level and is a reality of business operations. Overage ties up capital and cash and arises from imperfections in

MRP systems and management processes, combined with the need to be reactive to all changes. Many factors can contribute to overage; concerns with stock-outs, over-optimistic forecasts, lost material, time-fences, damping rules and manual intervention. For example, overage can occur when supply arrives earlier than planned or when demand is less than planned and supplies cannot be delayed. Murphy [1] also argues that there is a relationship between inaccurate data records and excess stock.

Research [2] has shown that a significant proportion of the inventory held by businesses is in the form of overage, which commonly ranges from 10% to 90%, with an average around 40% above the required stock level. Few managers outside of Finance worry about overage because it does not stop production or impact negatively on service levels, and furthermore requires resource to manage. In most businesses overage is thus an overlooked opportunity. Managing overage in a formal manner is addressed in more detail in the third paper. At this stage it is sufficient to recognise that overage



exists, how it can arise and that it presents an opportunity to decrease inventory at little risk to operations or service levels.

The regular saw-tooth replenishment pattern shown in Figure 1 assumes that demand and order lead-time are predictable and stable allowing repeat orders to be placed and delivered at the correct time. In manufacturing operations this is rarely the case and the real situation appears closer to that shown in Figure 2 with both demand and lead-time varying. Overage at point 'A' may arise from the reasons outlined above while underage at point 'B' may result from a higher than planned demand, a longer than anticipated lead-time or a previous delivery shortage. The combination of variability in demand and lead-time is known as the lead-time usage distribution, the assumption being that demand and lead-time probability are normally distributed.

Buffer stock is used to reduce the risk of a stock-out by compensating for the uncertainties in supply and demand. How to account for these uncertainties and determine the correct level of buffer stock is considered later.

We defined **batch stock** as that ordered on a regular basis to meet demand and allow operations to cope with not making all products simultaneously and looked at how a simple ABC classification could impact on the level of batch stock required. We next continue to refine the management of batch stock through extending stock classification.

EXTENDING STOCK CLASSIFICATION FOR GREATER CONTROL

In paper one we looked at a simple 3-class system for classifying stock based on annual usage value (AUV). There are additional classification factors, related principally to seasonality and demand [3], which can be considered to provide greater control. Seasonality may be due to a number of factors; holiday periods over summer, Easter or Christmas, celebrations such as Valentine's day or Mother's day, Bank Holidays etc.

each of which will require a higher or lower stock level on some goods. The pattern of demand can also be useful when classifying stock to distinguish superseded and obsolete parts from new products, slow moving parts such as spares from normal demand parts and also for highlighting demand trends for normal parts eg. is there a positive, negative or erratic demand trend? It could be argued that to optimise inventory each stock item should be uniquely classed to take into account the above factors, however this would be neither practical nor cost effective.

Increasing the number of classes based on AUV is one of the simplest means to effectively reduce overall inventory and is becoming common practice in more advanced organisations. Increasing the number of classes focuses attention on the most effective re-order period. In a 3-class system the choice seems almost to be determined naturally; a week for class A, a month for class B and a year for class C items. With a larger number of classes this is not the case. Common sense suggests using calendar multiples of a working week or month although for high value items, deliveries several times a week may be appropriate. There are indications [4] that geometric progressions of the order cycle give better results than linear progressions and that optimisation is approached using an 8-class system.

The example below shows how a reduction in batch inventory can be achieved by moving from 3 to 6 classes.

A worked example showing the impact of moving from 3 to 6 classes

Using the spreadsheet for 3-classes developed in paper one the extension to 6-classes is relatively easy. Note though that additional boundaries and order frequencies need to be assigned:

The steps required are:

1. Increase the number of classes to 6 (A1 to C2). Decide break points in percentage terms of cumulative AUV for each class (the example below uses A1=75%, A2=80%, B1=85%, B2=90%, C1=95%, C2=100%. Note that this keeps the upper boundary for A2, B2 & C2 classes at the 3-class starting point previously used) and assign an order frequency to each class (the example below uses A1=5 working days, A2=10, B1=20, B2=40, C1=120, C2=250, based on a 5 day working week and a 250 day working year). As before, calculate:

- a. Average inventory per class and total average inventory. Class inventory is derived from $0.5 \times (\text{sum of AUV for class/Frequency of orders})$
- b. Number of parts per class
- c. Orders per annum per class and total orders.

2. Determine the overall system batch inventory by summing the average inventory per class.

3. Vary the parameters for class-boundaries and/or order frequency per class to arrive at alternative scenarios for batch inventory and total orders.

Figure 3 shows one possible scenario for a 6-class analysis using the above parameters and in Figure 4 the results of this and another for 6-class scenario is compared against the result from our 3-class system starting position (A = 5 days, 80% boundary, B = 20, 90%, C = 250, 100%).

The same base data has been used in all cases. The starting

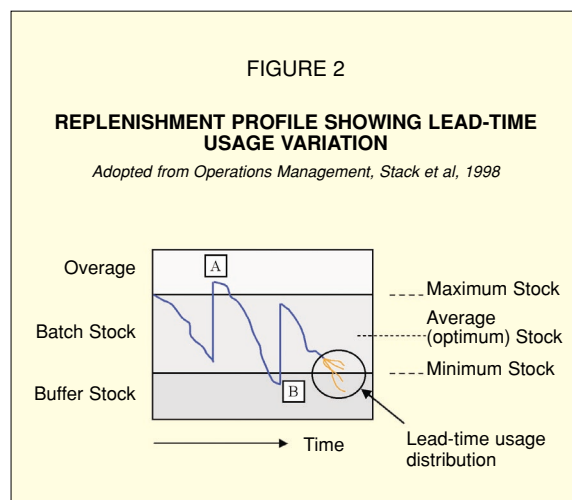


FIGURE 3
EXAMPLE SPREADSHEET ANALYSIS FOR BATCH STOCK IN A 6-CLASS SYSTEM (6-CLASS STARTING POSITION IN FIGURE 4).

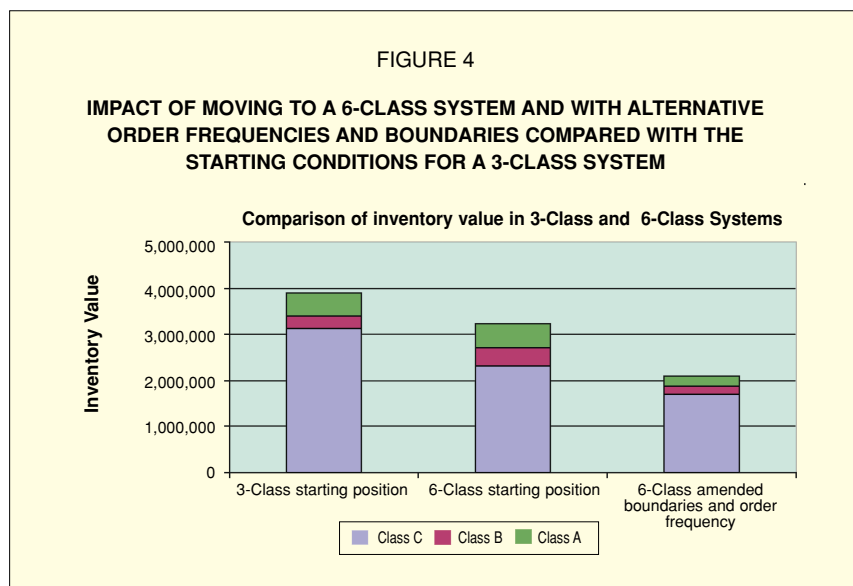
Class	Boundary	% Parts
A1	75%	9.5%
A2	80%	12.1%
B1	85%	15.7%
B2	90%	21.3%
C1	95%	31.6%
C2	100%	100.0%

Working days per annum 250

Class	AUV	Order days	Frequency p.a.
A1	47,020,605	5	50.0
A2	3,142,363	10	25.0
B1	3,139,840	20	12.5
B2	3,134,646	40	6.3
C1	3,140,090	120	2.1
C2	3,135,845	250	1.0

Class	Av. Inventory	# Parts	# Orders
A1	470,206	616	30,800
A2	62,847	163	4,075
B1	125,594	237	2,963
B2	250,772	362	2,263
C1	753,622	664	1,383
C2	1,567,922	4,415	4,415
Total	3,230,963	6,457	45,898

position for 6-classes represents the data shown in Figure 3. As can be seen the total batch inventory $I_{A1} \rightarrow I_{C2}$ has decreased, in this case by €0.66m (17.0%) to €3.23m. At the same time there has been a decrease in administrative burden through a reduction in the number of orders that need to be managed to 45,898 a decrease of 10.9%. The largest decrease in inventory is for C class parts and this is explained by the movement of 664 C parts to an order frequency of twice a year from once a year. This results in half the batch inventory being required for these parts. Some of this decrease is negated by an increase in inventory for 163 'A' class parts and 362 'B' class parts that have moved to an order cycle of 10 days and 40 days respectively (was 5 and 20 days). The decrease in the number of orders comes from decreasing the order frequency on these same A & B parts. Note that the total number of parts per class in the 6-class system has not changed from that in the same class of the 3-class system because the upper boundary limits are set at the same points. A1 + A2 parts in our 6-class



system number 779 as do class 'A' parts in the 3-class system.

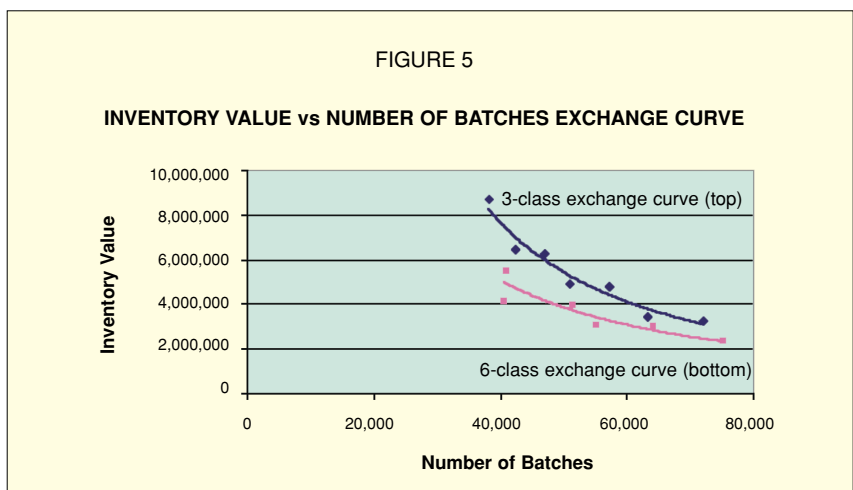
The use of 6-classes provides a greater scope to fine-tune boundary and order cycle parameters and the impact of varying both can be dramatic. The right hand bar chart in Figure 4 uses A1 = 2 working days, 80% boundary, A2 = 5, 83%, B1 = 10, 90%, B2 = 20, 93%, C1 = 120, 95%, C2 = 250, 100%. Here a reduction in inventory of €1.14m (35.3%) is achieved but with a substantial increase in the number of orders required, up by over 73,000 (143%).

Moving from a 3 to a 6-class system can reduce inventory without the need to increase orders. Greater inventory decreases can be made by adjusting the class order frequencies and class boundaries but at the expense of an increased level of orders. This trade off between

a decrease in inventory and an increase in the number of orders was highlighted in our analysis of the 3-class system and is equally in evidence here. Figure 5 shows how this can be graphically represented by an inventory value vs. number of batches exchange curve. Note that with lower inventory and higher batches the curves for 3-classes and 6-classes converge.

Assigning an inventory carrying cost and identifying the order processing capacity can help the business set relevant targets.

Inventory can then be managed through either setting a target inventory level in financial terms, or setting the level of resources for order processing and then deriving a maximum order capacity. However, as the simple scenarios explored here show, the sensitivity of different settings for both the boundary and



order frequency would make it difficult to arrive at the right settings to achieve these targets. Here our initial guides have been common practice and common sense and thereafter trial and error. Whilst inventory improvements can be gained, substantial effort is required to use this methodology in a controlled manner. For a system with 8 or more classes this becomes even more complex.

This trade off between inventory and order level will be explored and the difficulties in establishing the optimum boundary and order periods addressed in the third paper. Meantime we look at the means of establishing the level of our second inventory element - safety stock.

BUFFER STOCK METHODOLOGY

We previously discussed that **buffer stock** is required to compensate for uncertainties in supply and demand although as Burgess [5] points out it is often used for other reasons such as helping to smooth the manufacturing plan, moderating the effect of quality problems or as an order trigger mechanism. The uncertainties may arise from variations in time or quantity, for example orders may be delivered later than required and/or contain fewer items than ordered, while demand may be larger and/or sooner than anticipated. Though the time and quantity dimensions apply to both supply and demand, experience shows that lead-time variations in supply and quantity variations in demand are the main sources of uncertainty. For re-order level systems a further uncertainty centres on the order trigger quantity, which rarely coincides exactly at the order point but typically somewhere below, a feature known as overshoot.

If the function of buffer stock is to mitigate these uncertainties, the question remains to what extent should they be mitigated, or to put it another way how much buffer stock should be carried? To ensure that demand, however variable, can always be met could lead to excess stock in the system and high inventory

costs. Lower levels of buffer stocks would keep inventory costs down but may leave too high a level of demand unfulfilled leading to costly production rescheduling and dissatisfied or lost customers. Due to this relationship with customer satisfaction, both internal and external, buffer stock is often a focus of management attention and as such deserves to be set by more than the 'gut feel' that is often used.

Accordingly, to answer the question of how much buffer stock should be carried many companies now set target levels for customer service or overall inventory costs as measures to manage buffer stocks, and a number of prevalent models to achieve these targets have emerged. Silver [6] who defines buffer stock as "the average level of the net stock just before a replenishment arrives" suggests four such approaches for managing the balance between carrying inventory and running out of stock depending on the prevailing circumstances for the company and its products:

- A simple approach using a common safety factor or time supply
- Minimising costs by comparing different options to meet customer demand
- Setting safety stocks based on customer service levels
- Setting safety stocks based on aggregate considerations.

Indeed there are a wide variety of methods and formulae to choose from in setting safety stock, but they all have one common aim, which is to compensate for the fluctuations in demand and supply. This section looks at some of the methods used and considerations made for deriving safety stocks with a particular focus on achieving desired customer service levels. Whilst there are numerous mathematically oriented texts on this subject the intent of this article is to present the current practices used from a pragmatic perspective.

SIMPLE BUFFER STOCK

The simplest and most common methods of calculating buffer stock take no account of fluctuations in demand or lead-time. Typically buffer

stock is set either as a number of **safety days** ie. the number of additional working days for which no stock-out occurs based on the forecast annual or monthly demand, or as a **percentage of lead-time demand**, ie. the quantity of demand during the replenishment lead-time multiplied by a percentage factor – often set at 50%.

These methods linearly increase stock levels although a measure of refinement can be introduced by defining different safety day or percentage of lead-time levels per class or even per stock item. The main disadvantages are that they take no account of demand trends or changes in lead-time and are not easily related to satisfying customer service levels. Both these methods typically maintain too much or too little safety stock. In instances where demand and lead-time are stable, safety stock may be unnecessarily held and for products with very short lead-times or high demand variations stock-outs can occur. Figure 6 shows one scenario based on safety days for buffer stock and compares this with results using stochastic safety stock as described below. The methodology used in either case is given later.

STOCHASTIC BUFFER STOCK

Stochastic methodologies are based on statistical data and take account of demand and lead-time fluctuations and as such require historical data on demand patterns, lead-times or variations in quantity delivered. The base assumption is that a normal distribution operates for each attribute although this may not always be the case, eg. spares typically follow a Poisson distribution. Perhaps the most important aspect of using stochastics is that buffer stock levels can be set to achieve desired service levels. Service levels can be viewed from either a vendor or customer perspective and as different definitions exist it is worth clarifying these terms.

One commonly used vendor service level is the **cycle-service level** defined as the percentage of cycles where there will be no shortages [7] or equivalently the

probability of stocking out in any given replenishment cycle [8].

The most widely accepted stochastic method related to vendor service levels addresses the variability in demand (the forecast error) at a constant lead-time (assumed to be equivalent to the forecast period) for a specified service level to calculate safety stock [8], [7].

$$\text{Safety Stock} = Z\sigma\sqrt{L}$$

Z = Standard deviations from mean (also referred to as the safety factor)
 σ = Standard deviation of demand per period
 L = Order or replenishment cycle (lead-time)

The expected unfulfilled orders per cycle ie. the number of unsatisfied customer demands measured in units per replenishment cycle, is given by Unfulfilled Demand = σE_k where E_k is the partial expectation of the distribution [9].

A specific service level can be allocated to each inventory class although often the formula is simplified further by using an overall service level of 97.7% at which value $Z=2$, thus:

$$\text{Safety Stock} = 2\sigma\sqrt{L}$$

The assumption of a constant lead-time removes one variable from the calculation (but see below) though it should be noted that as buffer stock is proportional to the square root of the lead-time the longer the lead-time the greater the requirement for safety stock. To support a lead-time of 5 days requires 4.5σ , for 10 days 6.3σ , for 20 days 8.9σ etc. Lead-time is thus one of the most important factors that drive safety stock and one approach that can be taken to reduce stocks is to reduce the order cycle or the lead-time from the supplier. In considering lead-time to encompass all activities from order initiation to goods receipt, the possibilities to reduce the lead-time can range from capital investment through to sharing and speeding up information flow and improving operating processes. One example of this may be in vendor-certificated deliveries, which do not require incoming inspection at receipt.

From a customer viewpoint a vendor service level is not a

meaningful measure as the customer is unlikely to know the replenishment cycle and a **customer service level** defined as the proportion of demand met ex-stock during a stated period of time or satisfying product demand immediately out of inventory [8] is a measure more readily understood by customers. Fortunately as Lewis [10] points out when used in a re-order level system the method outlined above invariably results in a customer service level higher than the vendor service level thus addressing the customer viewpoint.

Customer service level as defined above is calculated from:

$$\text{Customer Service Level} = 100 - (100 \div Q) \sum_{D=R+1}^{D_{MAX}} (P_D)(D-R)$$

Q = Order quantity
 D = Average demand during the order cycle
 R = Reorder point
 P_D = Probability of demand of D units during the order cycle

The formula for safety stock shown above can be tailored to specific circumstances and to take account of time differences between order cycle, forecast period and lead-time [11]. A tailored formula may include the following elements:

$$\text{Safety Stock} = Z\sigma\sqrt{L_f} * C_f * D_f * V_f$$

L_f = Factor to account for differences between lead-time and forecast period. Typically L_f = a multiple of the forecast interval
 C_f = Order cycle factor to compensate for over stocking on products with long order cycles.
 D_f = Adjustment for using a forecast mean as opposed to an actual demand mean
 V_f = Adjustment to account for lead-time variations

To achieve optimum safety stock for a given service level and cost will require some experimentation with the best model for the circumstances and constant attention thereafter. Further experimentation may be conducted with different forms of the probability distribution, for example, a Poisson or log-normal distribution may provide a better fit. Finding the best distribution form is particularly relevant if a service level of over 98% is required. Again an approach that continually seeks to determine the best model of reality is recommended.

A worked example showing the impact of simple safety stock and stochastic safety stock

To calculate stochastic safety stocks requires demand data and we have used an alternative data set with this for this section. The spreadsheet is built up as previously and the addition of standard safety stock is then as follows:

For simple safety stock the steps required are:

1. Add a column for safety stock and calculate safety stock per class in terms of the number of days of stock held for buffer. In our example this is accomplished by:

a. Decide the safety days required for each class, here we have used $A1=2$ safety days, $A2=4$, $B1=6$, $B2=8$, $C1=10$, $C2=12$.

b. Calculate the safety stock value per part number and sum to provide a value for safety stock per class and overall.

2. Determine the overall system inventory by summing the average inventory calculated previously with the safety inventory.

The total system inventory I_A , I_B , I_C , for each class is now given by:

$$I_A = 1/2 A U V_A + F_A + \text{Safety Stock}_A$$

For stochastic safety stock the steps required are:

AUV = annual usage value
 F = order or delivery frequency

1. Import the demand data for each part. Ideally at least 30 demand points are required.

2. Add columns for standard deviation and stochastic safety stock and calculate standard deviation using the STDEVP excel function

3. Decide service level requirements per class - here we have used a service level of 97% for all classes. Convert service level to safety factor using the NORMSINV excel function, calculate the stochastic safety per part number from: $\text{Safety Stock} = Z\sigma\sqrt{L_f}$ and sum per class.

In our example we have assumed that $\sqrt{L_f} = 1$ but it would be straight-forward to incorporate the known lead-time into the calculation.

4. Determine the overall system inventory by summing the average inventory calculated previously with the stochastic safety inventory.

Figure 6 shows the parameters used and the results for simple and stochastic safety stocks. It can be seen that the level of safety stocks is a significant factor in determining total inventory. In our example, accounting for safety stocks has increased inventory from €1.41m to €2.61m and €2.82m for simple and stochastic calculations respectively.

The advantage of using stochastic methodology can be clearly seen in the visible relationship to service levels allowing management decisions to be made concerning the right balance of desired levels of service and inventory carrying cost. Improving demand forecasts thus decreasing sigma, the standard deviation of demand per period, leads to a decrease in stochastic safety stocks. Reducing lead-times also impacts demand accuracy, as the period of time the demand forecast is required to cover is reduced improving the quality of the forecast, and again decreasing sigma. Alternatively increasing the quantity per order and reducing the number of orders also reduces the occasion when there is a risk of shortages thus increasing the service level.

Focusing solely on demand variability however should be used with caution, as Benjaafar and Kim [12] show that higher levels of demand variability do not necessarily require higher levels of safety stocks (the fraction of safety stock due to total stock decreases with demand variability). Thus strategies such as demand pooling and advanced order information for coping with lead-time demand variability may be less useful when the demand variability is high. In multi-level systems one pitfall to avoid is the unnecessary creation of safety stock throughout the different levels of materials eg. components and raw materials as well as finished stocks. Techniques using an accumulated demand error through the bill of material are being applied. Another approach to ensure high finished goods availability involves

FIGURE 6

EXAMPLE SPREADSHEET ANALYSIS FOR A 6-CLASS SYSTEM WITH SAFETY STOCK USING SIMPLE AND STOCHASTIC METHODOLOGY

Class	Boundary	% Parts	Customer Service	Safety Factor	Safety days
A1	50.0%	0.90%	97.0%	1.88	2
A2	75.0%	3.42%	97.0%	1.88	5
B1	90.0%	9.81%	97.0%	1.88	10
B2	96.0%	19.43%	97.0%	1.88	15
C1	99.0%	37.41%	97.0%	1.88	20
C2	100.00%	100.0%	97.0%	1.88	25

Class	AUV	Order Days	Frequency p.a.	Stochastic Safety	Simple Safety
A1	35,697,518	2	140.0	454,294	254,982
A2	17,844,666	5	56.0	282,050	318,655
B1	10,762,291	10	28.0	212,080	384,368
B2	4,286,659	20	14.0	129,376	229,642
C1	2,143,471	120	2.3	85,081	153,105
C2	714,820	250	1.1	39,375	63,823
Total	71,449,424			1,202,256	1,404,575

Class	Av. Inventory	# Parts	# Orders	Inventory + Stochastic SS	Inventory + Simple SS
A1	127,491	53	7,420	581,785	382,473
A2	159,327	149	8,344	441,378	477,982
B1	192,184	378	10,584	404,263	576,551
B2	153,095	569	7,966	282,471	382,737
C1	459,315	1,064	2,483	544,396	612,420
C2	319,116	3,702	4,146	358,491	382,939
Total	1,410,528	5,915	40,943	2,612,785	2,815,103

stocking long lead-time components at service levels of 99.9% or higher.

The calculation of safety stock to reflect real world uncertainties is complex because of its dependence on a number of factors with varying probabilities, which is one reason many companies still use the simple methods available. These however do not provide an optimal solution. Though the mathematics on which more sophisticated techniques are based was developed some years ago their application is still relatively recent and requires tailoring to account for the variables relevant to a particular supply chain. Because of the significant impact of safety stock on overall inventory and service levels it is worth the effort to find the right model for your business.

SUMMARY OF THE STEPS TO PROFESSIONAL INVENTORY MANAGEMENT

In paper one the first steps towards basic inventory control and management were discussed:

Step 1. Inventory Control

- Instigate appropriate business processes, minimising manual intervention

- Tie system data on stock location and quantity to physical stock. Insist on accurate data records and check using cycle counts

- Set targets and measure the degree of control through inventory variances (book to actual) and instances of physical inventory shortfalls

- Educate staff.

Step 2. Basic Inventory Management

- Classify inventory using a 3-class ABC system based on Annual Usage Value

- Determine appropriate order parameters for each inventory class

- Analyse the results at aggregate, class and part level and take appropriate action

This article has looked at the next steps that can be taken to increase the level of inventory management.

Step 3. Focused Inventory Analysis

- Understand the breakdown of inventory into three elements; batch stock, buffer stock and overage, and the need to manage each breakdown into inventory classes separately

- Recognise the existence of **overage** and the reasons that lead to it

- Increase the number of classes from 3 to 6 or more

- Take into account relevant seasonality or demand factors

- Set targets for inventory level and/or order capacity and set boundary and order frequency parameters to manage **batch stock**

- Consider the requirements for **buffer stocks** and their relation to customer service levels. Set buffer stock levels by applying simple safety days at class or part level or better still through the use of demand based stochastic techniques.

- Understand that the models used attempt to replicate reality and should be continually adjusted to improve fit. Understand how demand varies through the year and periodically reset safety levels as necessary.

In the third and final paper “*Advanced Inventory Management*” we introduce the concept of the K-curve to manage batch stock, discuss the impact and control of overage and look at alternative inventory management strategies.

REFERENCES

1. Murphy, N., “The Right Amount of Stock?” *BPICS ‘Control’*, February/March, 1990.
2. Relph, G & Barrar, P., “Overage Inventory – How Does it Occur and Why is it Important?”, (2000), *Int. J. Production Economics* 81-82 (2003), pp 163-171.
3. Murdoch, J., “Classification of Stock Items for Effective Management Control”, *BPICS ‘Control’* October/November 1988.
4. Relph, G, Browne, N, Dupernex, S & Dance, R., “Improved Inventory Performance into the 21st Century”, *BPICS ‘Control’*, June/July, 1994.
5. Burgess, M., “Safety Stocks or How to Manage an Inefficient Business”, *BPICS ‘Control’*, November, 1995.
6. Silver, E A, Pyke D F & Peterson, “*Inventory Management and Production Planning and Scheduling*”, 3rd Edition. John Wiley, New York 1998.
7. Wilkinson, S., “Service Level and Safety Stock Based on Probability”, *BPICS ‘Control’*, April, 1996 (and Lewis, C., “Safety Stocks – Yes Probably”, *BPICS ‘Control’*, September, 1996.)
8. Vollmann T E, Berry WL, Whybark DC, “*Manufacturing Planning and Control Systems*,” 2nd Edition, Richard D Irwin 1984.
9. Brown, R , “*Advanced Service Parts Inventory Control*”, 2nd Edition, Materials Management Systems Inc., Vermont 1982.
10. Lewis, C., “Safety Stocks – Yes Probably”. *BPICS ‘Control’*, September, 1996.
11. Piasecki, D., “Optimizing Safety Stock”., *Inventoryops.com* (http://www.inventoryops.com/safety_stock.htm)
12. Benjaafar, S & Kim J, “*When Does Higher Demand Variability Lead to Lower Safety Stocks?*” University of Minnesota, Minneapolis, USA 2001.

(IOM: formerly known as BPICS)

About the authors

Witek Brzeski is a management consultant in IBM's Business Consulting Services specialising in the design and implementation of business processes. He has worked within manufacturing, purchasing and R&D in supply chain and technical roles and has experience in project management, manufacturing operations, supply chain management, consumer product development and change management. Witek has an MBA from Warwick Business School.

Geoff Relph, MIOM is Director of Inventory Matters Limited specialising in Manufacturing Planning, Control and Inventory Management. His 25 years of 'hands on' experience in Production Control, MRP Operations, and Inventory Management has been gained in large blue chip companies - Rolls-Royce Motors, Crosfield Electronics (De La Rue Group), GEC-Marconi and IBM. Geoff is recognised as an industry leader in MRP operations and inventory management and has been instrumental in the development of innovative inventory management techniques, through both academic research and development

in response to specific client requirements. The academic work has been with Aston Business School, Sunderland University and more recently with Manchester Business School working on his own doctoral research on Inventory Management in Business Systems. Geoff regularly speaks for the IOM, and has published in several Journals and spoke at the 12th Symposium on Inventory Research in Budapest in 2002. Geoff is the Vice Chairman of the Institute of Operations Management.

Gail Bradbear, joined the IOM in 1997 and has been a member of the Editorial committee for over 2 years. She has worked in production and supply chain management for over 7 years first as Production Controller and Materials Manager at Hanovia, Slough and then with a number of market leading manufacturing organisations as a Consultant as part of IBM's Industrial Practice, where her particular interests lay in the field of stock control and inventory management. Gail is currently working for Ford as a Business Analyst on a pan European IT implementation project.