

Advanced inventory planning and forecasting solutions: A case study of the UKTLCS Chinook maintenance programme

Matthew Downing¹, Max Chipulu², Udechukwu Ojiako³ and Konstantinos Kaparis⁴

^{1, 2 and 3} School of Management, University of Southampton

⁴ School of Mathematics, University of Lancaster

This paper which advances on earlier published work by the authors, evaluates the forecasting performance of the inventory planning and forecasting support system used by Boeing's *United Kingdom Through Life Customer Support* (UKTLCS) to manage inventory for the maintenance of the Royal Air force (RAF) fleet of Chinook helicopters. Focussing mainly on a sample of the spare part components considered most influential on maintenance costs, limitations in the current forecasting system are identified. Approaches for improving forecasting performance are explored using Visual Basic for Applications (VBA) based tools: It is found that generic proprietary inventory and forecasting systems can be enhanced by using VBA tools that allow detailed examination of each component's demand time series. For propriety reasons, all data have been sanitised.

Keywords: Forecasting

submitted to Production Planning and Control

1. Introduction

The Chinook helicopter is one of the longest serving military aircraft operated by the Royal Air Force (RAF). A versatile, fast transport helicopter, the Chinook can lift over 10 tonnes of cargo. Boeing, through its *United Kingdom Through Life Customer Support* (UKTLCS) programme, provides maintenance and technical support for the RAF fleet of Chinook helicopters. Part of UKTLCS's maintenance challenge is to ensure that correct parts are delivered on time at the correct location. To ensure the success of this process, UKTLCS employs advanced inventory planning and forecasting systems.

With over 13,000 spare part components in the system for the Chinooks, the complexity and the scale of the supply chain management operation for UKTLCS is vast. For Boeing, the importance of forecasting performance cannot be overstated: If demand for components is systematically under-estimated, then the efficiency of maintenance operations may drop, which may impair Boeing's relationship with the RAF even if contractual obligations are not quite violated. Boeing spends hundreds of millions of dollars per annum maintaining the RAF fleet of Chinooks. As such, if demand for components is systematically over-estimated, large amounts of capital are unnecessarily held up in stock.

At the time of this research project, UKTLCS exclusively used Service Planning and Optimization (SPO) to forecast demand for components. SPO- a proprietary intelligent system developed by MCA solutions- employs a variety of forecasting techniques including specific adapted modelling processes and variants of well-known time series techniques. Although SPO had proved capable of delivering forecast projections, this research arose as a result of senior management at Boeing being interested in exploring changes to SPO that could lead to enhancing forecasting performance and, ultimately, better control of inventory costs.

2. The Study

Substantial research has already established the need for accuracy in military inventory management (Michaels 1999, Tysseland 2009, Johnsen *et al.* 2009) and the role forecasting accuracy plays in this process (Ho and Ireland 1993,

Venkataraman and Nathan 1999, Clark 2005, Graman and Sanders 2009). Based on this study which advances on earlier published work by the authors (see Downing *et al.* 2011), an examination was conducted on real-life data within the supply chain environment of UKTLCS's Chinook maintenance programme. Concentrating on a selected number Chinook components, this paper aims to evaluate and recommend new forecasting techniques specific to UKTLCS's Chinook maintenance programme. Specifically, this paper seeks to advance an earlier outline study by the authors by seeking answers to three research questions not addressed in earlier studies (Downing *et al.* 2011) by the authors:

- What assessment approaches exist to support the evaluation of inventory planning and forecasting support systems?
- How effective is the SPO™ forecasting support system in supporting the Chinook spare components maintenance programme?
- Can a more effective forecasting system be developed to support the UK Chinook spare components maintenance programme?

The remainder of this paper is divided into four sections. Following this introduction, the research methodology is presented. This followed by an exploration of patterns in the selected component demand time series'. In section 5, the forecasting performance of the SPO system is evaluated prior to exploring enhancements to this system in section 6. In the final section of the paper, the research project conclusions are outlined.

3. Research Methodology

3.1. Data Collection

Data for this study were obtained through Participant Observation and Action Research (Argyris and Schon 1991). This approach was possible as access was provided to the principal author by UKTLCS. The researchers were also given access to spread sheets containing component demand and cost data over 29 calendar

months. Much of the initial processing of the data involved merging the component demand and cost datasets. Once merged, the data were discovered to have many months of negative values. Such 'negative' values were, as was explained by the UKTLCS management, situations where components had been moved from one location to another in the event of a stock-out at the end location. Since such negative values did not represent actual demand at the initial location, in months where a location had negative demand recorded, we re-set the demand value to zero. The final data required for the study were earlier forecasts of component demand. However, as these data were not stored by UKTLCS, each of the forecasts had to be extracted individually for each of the selected components and analysed through a Visual Basic for Applications (VBA) based modelling tool. Although slightly tedious, the information and insight from these data was critical to the development of an understanding of previous forecasting performance.

3.2. Selection of Spare Part Components for Evaluation

A decision was made at the outset that the study would focus on the components with the most influence on the maintenance expenditure of UKTLCS. This decision was taken for two reasons: First it allowed us to fulfil Boeing's required key outcome from the project, namely to explore ways to optimise stock holding costs by improving forecasting performance. Second, it provided a critical focus for our research, potentially leading to greater, deeper insight.

Based on advice from specialists working on the UKTLCS maintenance programme (Karr 1958, Denicoff *et al.* 1960), a total of 92 components were selected from the initial list of 13,000 components. The 92 components can be divided into two classes: 'cost drivers' and 'consumables'. A cost driver can be defined as a repairable component with a high annual expenditure. For the purpose of this paper, annual expenditure per cost driver was calculated as the total yearly forecast * unit cost. In total over 40 types of cost drivers were identified, equating to a little over \$200 million of forecasted investment. A consumable can be defined as a component that, once stripped from the aircraft, cannot be repaired and as a result is scrapped. Controlling spending on consumables is just important as it is for the cost drivers

because once consumed, the amount spent on a consumable must be written off. For this research, the annual spend per consumable was calculated as the unit cost*total yearly replenishment rate. The top 15 highest consumables amounted to a forecast investment of over \$10 million. Though only a very small proportion of the total number of components, the contribution to total maintenance costs of the 92 selected components is disproportionately significant, a not unusual inequality in cost when considering the distribution of costs for aircraft components (e.g. Kennedy *et al*, 2002)

4. Evaluation of Component Demand Time Series'

To explore how, for each of the 92 components, demand varies over time, we plotted graphs of each individual time series and conducted univariate analysis (Huberty and Morris 1989, Viertl 2006). The graphs revealed a broad range of patterns. Figure 1 shows some examples of the demand patterns. Given that components perform different, specialised functions in the helicopter, we should expect the observed deviations in demand patterns. Occasionally, spikes in demand are evident. It is our understanding that such spikes may be related to how the helicopters are employed.

Even at 92, the number of items to be forecast is large. In these situations, the literature suggests that the items should be categorised by demand based on factors considered important in forecasting performance. The key factors include the frequency of (non-zero) demand; and the level and variation of the demand value (Williams 1984; Syntetos *et al* 2005; Boylan *et al* 2008). With this in mind, the next step in the data evaluation was to explore if commonalities exist in the patterns of demand by examining the univariate statistics across all 92 components. The statistical summaries of interest are presented in table 1. Though the average demand across the components was low (≈ 3), the dispersion in mean demand was large, ranging from very low (nearly zero) to about 35.





Figure 1 – Sample Component Demand (Units)

Table 1: Descriptive Statistics					
	N	Minimum	Maximum	Mean	Std. Deviation
Mean	92	.0	34.8	3.371	6.4407
Range	92	1	157	14.25	25.584

Median	92	0	28	2.33	5.221
Standard Deviation	92	.18	30.1	3.6	5.9
Variance	92	0.034	915.7	47.5	159.3
% Zeros	92	0%	97%	50%	28.5%
Skewness	92	0.2	5.4	1.987	1.2
Coefficient of Variation	92	0.6	5.4	1.7	1.06
Square Coefficient of Variation	92	.4	29.0	4.1	6.0
Mean Inter-Demand Interval	92	1.0	29.0	4.0	5.8

All but three of the 92 components had experienced periods of no demand as shown in Figure 3. In fact, on average, 50% of the time, there was no demand for components. Despite the afore-mentioned negative demand values, demand cannot be less than zero. Given this clamping effect on the left-tail of the demand distribution, and given that zero demand was frequently observed (figure 3), it is not surprising that all the distributions of demand were asymmetrical, exhibiting right skew (all values skewness were positive). In figure 4, we have plotted skewness against the percentage of zeros (i.e. periods of no demand) observed for the 92 components and labelled each data point with the mean demand value of the component. Roughly, the components cluster around a South-West to North-East diagonal: as the percentage of zeros increased, so did the skewness; and both were concurrent with decreasing mean demand. The majority of the distributions of demand values were leptokurtic, indicating fat tails. This suggests that the high variation in demand might be due to frequently occurring extreme values in both tails: on the one hand, frequently there was no demand; on the other hand, when demand did occur the value was, frequently, significantly dispersed from zero.

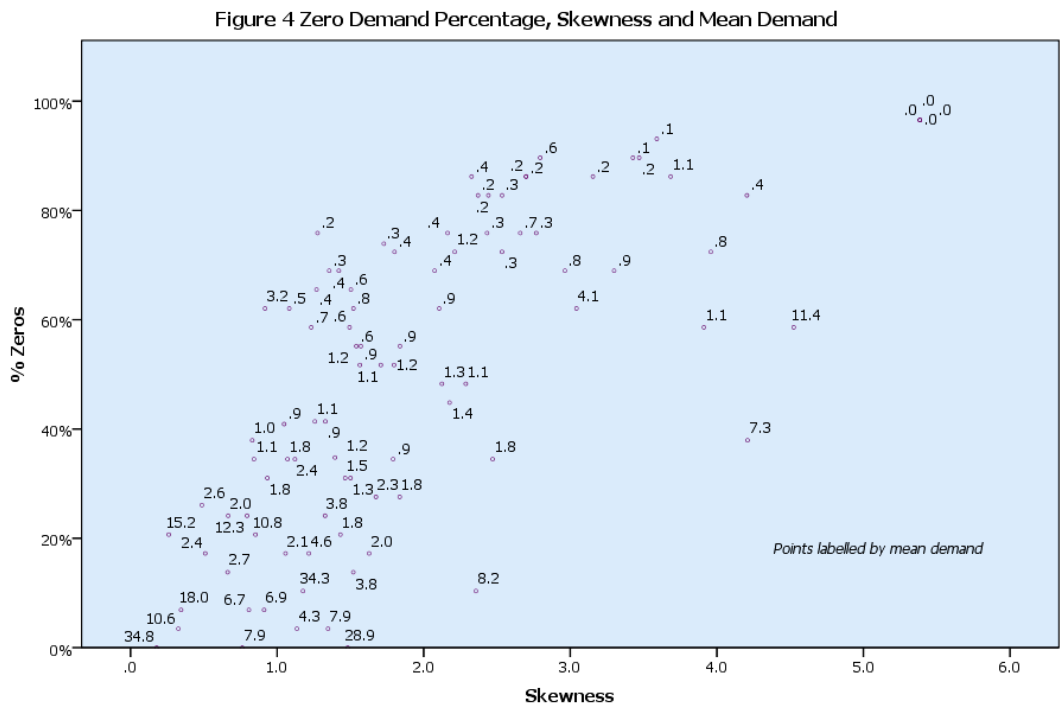
Figure 3 Frequency Distribution of Chinook Demand Components



In table 1, it can be seen that Demand was often infrequent (as indicated by the mean inter-demand interval) and that when demand did occur, values varied widely (as indicated by the dispersion statistics). Thus, vis-à-vis the lead time, demand patterns could, generally, be considered non-normal. Principally, given non-normal demand, the forecasting literature (Syntetos and Boylan 2005; Boylan *et al* 2008.) suggests that the co-variation of the mean inter-demand interval and the square of the coefficient of variation (CV^2) be used to categorise demand for multiple time-series' and to employ the forecasting techniques that have been shown to perform well for the category of demand so identified. It is possible however that the value of the coefficient of variation (CV) of a time series is affected by an underlying trend or seasonality pattern so that a randomised pattern of high and low demand can have a CV value of similar magnitude to another series that exhibits seasonal or trend characteristics. Wallström (2009) suggests that, as a non-scientific rule of thumb, if the discrepancy between the CV and then Mean Absolute Change Scaled ($MACs^5$) is


⁵ There is a confusion in Wallström (2009) in that he uses the same acronym, i.e. $MACs$ to define both Mean Average Change scaled and Mean Absolute Change Rescaled. The $MACs$ calculation we use is the Mean Absolute change Rescaled, Equation 2.29 on page 27 in Wallström (2009)

larger than 10-15%, then it can be concluded that there is some kind of non-random sequence in the time series or that the time series is dominated by a small number of modal values. In figure 4, we plot, for each component, the mean inter-demand interval versus the CV^2 . Each data point in figure 5 is labelled by the discrepancy (as a percentage) between the CV value and $MACs$ values. Following Syntetos et al (2005), we have also categorised each time series as 'Erratic', 'Lumpy', 'Smooth' or 'Intermittent'.




Based on Syntetos et al (2005), therefore, the majority of the times series' could be classified as 'lumpy'. A fair number of them were 'erratic', a few were 'smooth' and none were 'intermittent'. We can observe, in figure 5, a general pattern in the values of the discrepancy between the CV and $MACs$ values: it decreased with higher values of both the CV and Mean Inter-demand interval. However, the majority of the discrepancies were at least 10% and so, under Wallström's (2009) rule of thumb, it can be concluded that generally the observed CV values were due to some kind of pattern (such as the dominance of the time series by a few frequently occurring values) in the data as opposed to random variation.

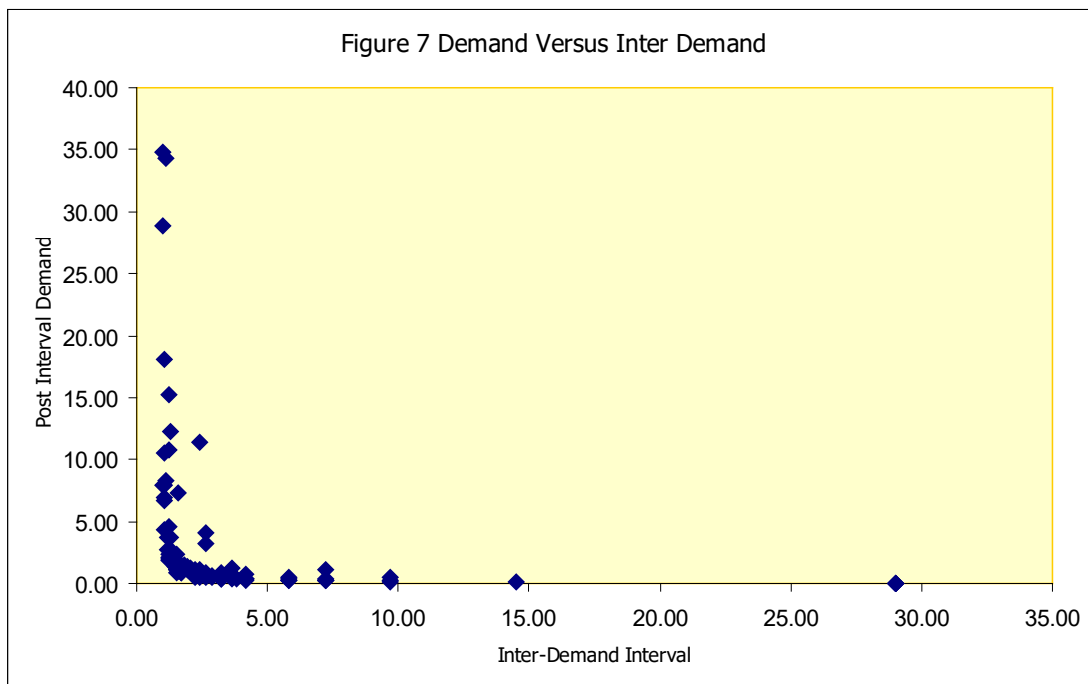
Figure 5 Categorisation of Chinook Component Demand After Syntetos et al (2005)




Depending on their function in the helicopter, components can be categorised into six types (1 to 6). The RAF fleet of helicopters operate from four locations (A, B, C and D). Both the type of component and location may influence the demand pattern. Therefore we investigated whether visually discernable clusters of components by type/location existed in the univariate statistics. Generally, we were not able to identify clustering by type/location on the univariate statistics. There was one exception: We found that type 2 components generally exhibited higher mean demand levels than other types (figure 6).



Besides exploring the variation of the univariate statistics across the 92 components and whether the type of component and location were influential, we also explored if (i) large inter-demand intervals preceded by unusually large demand values; and (ii) If flying hours and average demand were concomitant. In both cases, we were not able to observe a clear pattern indicative of such relationships (Figures 7 and 8, below.)





Overall, we concluded from the above evaluation of the data that the demand for the selected g2 will be difficult to forecast and may not be accurately forecasted if particular attention is not paid to the non-normality of the patterns, particularly the combined effects of the amount of variation in the values and inter-demand intervals. Below we explore the performance of the SPO forecasts, keeping in mind these patterns in demand.

5. SPO™ Forecasting Performance Evaluation

5.1. Overview of the SPO system

From a supply chain perspective, UKTLCS's goal is to ensure that the right components are delivered to the right locations, in sufficient quantities, on time. Currently Boeing UKTLCS outperforms all contractual service requirements, whilst performing specialist over and above tasks (O & A) on the side. Given the enormous number of components, the diverse demand patterns, the numerous locations and types of components; UKTLC, necessarily, uses an expert system to successfully manage the supply chain operation. The system used is Service Planning & Optimisation (SPO) by MCA Solutions.

The SPO Product Suite is an advanced inventory planning and forecasting solution that allows UKTLCS to monitor, modify and optimise the inventory level of mission-critical materials (*MCA Solutions*). Using a series of integrated graphical user interfaces (GUI's), SPO provides a number of solution modules including strategy, ERP and CRM. Demand for contract (or non-contract) components can be forecasted under the strategy module. To begin forecasting, the user first selects the appropriate component number from the GUI. From here s/he can assess the component's demand history, change the forecasting parameters and see the projected forecasts for the next 12 months. Table 2 below presents a summary of the forecasting techniques available in SPO's strategy module.

Table 2 Forecasting Techniques Available in SPO (Source: MCA Solutions)


Forecasting Technique	Description and Usage at the time of this research
Moving Average Constant	Provides a simple way to smooth historical demand information. The default moving average is 6.
MCA Smoothing	A variant of the exponential smoothing method with a damped trend. It uses the parameter α [$0 \leq \alpha \leq 1$] to smooth historical demand. Currently, $\alpha = 0.1$, the default level.
Trend	A variant of the single exponential smoothing method that is tailored towards intermittent demand time series. It can be used for all types of time series, except where there is trend or seasonality. Smoothing is based on the parameter α . The default level of α is 0.1, but the level has been changed, by management, within the last 12 months to 0.01.
Seasonality	A variant of the double exponential smoothing method which takes into account a trend. The parameter α is used to smooth the time series, with parameters β and ϕ to dampen the effects of the trend on the forecast. The default and the current level of the parameters are 0.1, 0.05, 0.8 for α , β and ϕ , respectively.

Table 2 Forecasting Techniques Available in SPO (Source: MCA Solutions)

	parameter α to smooth and γ to allow for the seasonal variation. The current and default values for α and γ are 0.1 and 0.6, respectively.
Trend	Incorporates both a trend and a seasonal effect in the data using the parameters α , β , ϕ and γ . The default levels of these parameters are 0.1, 0.05, 0.8 and 0.6, respectively.
Seasonality	
MCBF/MTBF	Uses Mean Causal Between Failure (MCBF) divided by Mean Time Between Failure (MTBF) to compute a forecast. MTBF forecast for a component is based on the contract population history and total contract demand in the historical periods used; it places a higher weight on the most recent demand levels. Forecasted MCBF for a component is based on contract causal history and the total contract demand in the periods of history used. The MCBF/MTBF forecast for each location-component combination is used to generate a causal forecast.
MCA	Blends a time series forecast with MCBF forecast such
Composite	that: $\text{CompositeForecast} = \gamma * (\text{TimeSeriesForecast}) + (1 - \gamma) * (\text{MCBF} / \text{MTBFForecast})$ <p>The nature of the equation means that an ever increasing emphasis is created as γ increases. Suggested values for γ are between 0.25-0.5.</p>
Manual	SPO allows the user to input her/his own forecasting technique
Override	projections.
Replacement	Using the replacement rate forecast based on the parent component's usage. If certain conditions are met, then the replacement forecast for the child is generated as a sum of all its parents local usage multiplied by the corresponding replacement rates .

5.2. *Evaluation of the Existing SPO Forecasting Performance*

To explore the possibility of developing a new forecasting strategy, we conducted a diagnosis of UKTLCS's existing forecasting system. Of the SPO forecasting techniques shown in table 2, at the time of this research, the MCA smoothing forecasting method was used for most of the UKTLCS components. Very little is known about the MCA smoothing forecasts as the equations of the technique have never been revealed. However, we can make assumptions about its likely formulation. The user guide of SPO has suggested that MCA Smoothing is a variant of Simple Exponential Smoothing (SES), and from investigation of the nature of the technique, the forecast value does not update in periods of zero demand. Our initial assumption was that MCA Smoothing was most likely to be using Croston's (1972) methodology. To test this assumption we compared, for all the 49 series for which MCA Smoothing was used, the forecasting errors obtained using MCA Smoothing against those obtained using Croston's Method based on the same smoothing alpha value of 0.01. Figure 9 shows a plot of the Mean (signed) Error obtained from each series for the two methods. It can be seen that that forecasting accuracy are not identical. However, the patterns of forecasting accuracy across the series are similar, with both methods tending to over-estimate demand. Thus, while the results failed to confirm that the MCA Smoothing technique is indeed Croston's, they did appear to indicate that the formulation of MCA Smoothing is very similar in manner to Croston's.



By creating a VBA based modelling tool that allowed efficient processing of key parameters, we analysed the forecasting performance of the SPO system. There exist a large number of accuracy measures for assessing forecasting techniques in literature. Not all are appropriate for assessing forecasts in intermittent demand situations (Syntetos and Boylan 2005). We based our selection of accuracy measures on two criteria: (i) simple and intuitive so that it can be easily communicated to non-experts in forecasting such as the majority of the decision-taking senior managers at Boeing; and (ii) appropriate for use with intermittent data. To assess the accuracy of forecasts for each individual time-series, we adopted the use of the Mean Square Error (MSE), the calculation of which simultaneously prevents cancellation of over-estimation and under-estimation errors and imposes larger penalties on large errors. MSE, while not the most straightforward in its interpretation, is the most closely related to the variance, which is important for stock control purposes. To compare forecasting performance across groups (by

forecasting method/strategy) of time series', we elected to use the Mean (signed) Error (ME) which has been shown not to be strongly scale dependent (Syntetos and Boylan 2005). Both the MSE and ME satisfy our twin criteria for choice of accuracy measures. The results were as follows:

Table 3 SPO Performance Mean Errors by Forecasting Method

Accuracy Measure	Mean Error (ME)		Mean Square Error (MSE)	
	FH/MCBF	MCA Smoothing	FH/MCBF	MCA Smoothing
Descriptive Statistics				
Number of Time Series ⁶	40	49	40	49
Mean	-1.8	-0.6	101.2	39.1
5% Trimmed Mean	-0.9	-0.5	37.3	10.2
Median	-0.4	-0.3	4.7	1.0
Variance	28.2	6.6	112950.7	30493.9
Std. Deviation	5.3	2.6	336.1	174.6
Minimum	-32.0	-10.3	0.0	0.0
Maximum	0.7	11.1	1794.9	1205.4
Range	32.8	21.4	1794.8	1205.4
Interquartile Range	1.1	0.7	13.8	5.1
Skewness	-5.1	0.4	4.1	6.5
Kurtosis	28.4	12.9	17.8	43.8

Table 3 presents summary statistics for the ME and MSE values of the component time series by SPO forecasting method. By inspection of the ME summaries, it can be seen that, overall, the MCA smoothing method performed better than other forecasting techniques⁷ such as the Mean Causal Between Failure and Mean Time Between Failure (MCBF & MTBF) methods. The performance of the MCA smoothing forecasts was also more consistent, less variable. Both methods tended to over-estimate demand. Indeed, inspection of individual time series' errors shows that often demand was over-estimated by as much as 100%, sometimes hundreds of percentages over. Croston's (1972) has been shown to be biased so that it tends to over-estimate demand (Syntetos and Boylan 2001). Therefore, the pervasiveness of over-estimation in the MCA smoothing based forecasts lends

⁶Items 69, 79 and 83 were excluded as we were not able to establish with exactitude the forecasting method that been used

⁷ Based on MCBF and flying hours

support to our supposition that the MCA smoothing method is based on Croston's original formulation.

Next, to gain deeper insight into when the MCA smoothing method performs very well or very badly, we inspected its performance on individual time series'. We took samples of the worst and best performing items (tables 4 and 5). Noteworthy among the worst performing items is item 88, which showed substantial inaccuracy. This is likely to be due to the lumpy nature of the demand pattern. A series of zero demand months were followed by a month of 157 demanded units, followed again by zero demand. Although not many forecasting techniques will be able to deal with this type of time series, the underlying method of MCA Smoothing performs particularly poorly. This suggests the model is not able to react to changing time series patterns. At the other end of the scale, the items where the MCA smoothing method performed better were either the patterns with one dominating value of a very low value or time series with low intermittence (mean inter-demand interval below 1.5). Generally, the results from the best performing items showed no anomalies at the magnitude exhibited by the poor performing items. Overall, the inspection of individual time series' suggests that the magnitude of the average inter-demand interval and the presence of outliers could both be influential on forecasting performance.

Table 4 A selection of worst performing items

Item	Location	Mean	High	Standard. Deviation	Zero Frequency	Skewness	Average Inter- Demand Interval
88	A	11.41	157	29.70	17	4.52	2.42
29	B	0.72	3	1.03	17	1.23	2.42
45	B	1.07	7	1.60	14	2.29	1.93
57	B	1.79	11	2.30	10	2.47	1.53
63	B	1.14	4	1.13	10	0.84	1.53
13	C	0.45	3	0.83	21	1.80	3.63
54	C	3.76	16	3.99	7	1.33	1.32
62	C	6.93	22	5.22	2	0.91	1.07
66	C	0.45	2	0.69	19	1.27	2.90

Table 5. A selection of best performing items

Item	Location	Mean	High	Standard. Deviation	Zero Frequency	Skewness	Average Inter- Demand Interval
10	B	1.00	3	1.04	11	0.83	1.61
49	B	0.35	2	0.65	17	1.73	3.83
72	B	1.79	6	1.72	9	0.93	1.45
89	B	18.03	40	12.79	2	0.34	1.07
24	C	0.59	3	0.95	19	1.50	2.90
44	C	0.31	3	0.66	22	2.77	4.14
67	C	7.93	27	7.22	1	1.35	1.04

5.3. Impact of MCA smoothing alpha parameter on performance

A decision was made by UKTLCS management to reduce the value of alpha in MCA smoothing from the recommended 0.1 value to 0.01. This, they reasoned, would calm down fluctuations in the forecasted values and in return lower the effect on the spiking target stock level (TSL). A comparison between the forecast values of 0.1 and 0.01 was undertaken to investigate any detriment on performance as a result of the change in parameter value.

Using another VBA based modelling tool, performance of the MCA smoothing method with the two separate alpha values was compared against the performance of a simple forecasting technique (in this case Simple Exponential Smoothing). The choice to adopt this approach was based on earlier work by Markland (1970), which suggests that forecasting the demand for military helicopter spare components is best conducted using exponential smoothing models. The VBA tool would then compare the MSE of each to look at the change in performance depending on the value of alpha. The results (Table 6) provided a strong indication that increasing the alpha value back towards the recommended 0.1 will increase the performance and in turn could generate a reduction in inventory held by UKTLCS. Theoretically, the lower the alpha value the more akin a total average the forecast values become, with older demand values having a larger bearing on the forecasts. The poorer forecasting accuracy (at the very low alpha value of 0.01) suggested, however, that, in fact, a larger smoothing constant was required to give larger importance to recent changes in the demand patterns.

Table 6. Number of items falling into parameter categories

Alpha Parameter Value	MSE(In)	MSE(Out)
0.1	69	56
0.01	23	36

5.4. Impact of flying hours (and other factors) on demand

At location D, as part of the MCBF/MTBF forecasting technique, a causal linear relationship is assumed between flying hours and demand. In theory, increase in flying hours should increase the service requirements of the Chinooks, thus increasing the demand for spare parts. To determine if a significant relationship does exist between flying hours, operational temperature⁸ and demand, we examined the correlation coefficients among the variables and conducted a linear regression analysis where the dependent variable was demand; and flying hours and operational temperature were the predictors. We found that neither flying hours nor operational temperature was significantly related to demand. It is questionable, then, that causal forecasts are applied in these circumstances.

5.5. Cluster analysis of choosing forecasting method based on location

When choosing the appropriate SPO forecasting method, the location of the Chinooks is thought important so that, for example, at location D, MCBF/MTBF techniques are used because they are considered most appropriate for that location. In principle, this rationale is sound. The location determines the usage pattern of the Chinooks, which can be expected to impact on service requirements and demand for components. However, the evidence from the field appears to suggest forecasting performance is poor as a result of this choice, with asset managers frequently reporting 'over-buying'. To explore if demand patterns for the spare components are similar based on location so that it is justifiable to select forecasting techniques based on location, we conducted cluster analysis of components based on demand in the statistical package PASW17 (formerly SPSS) as follows:

⁸ Maintenance experts believe temperature changes affect the service needs of the Chinook

- *Step 1*- We calculated the proximities between all 13, 000 components using 'seuclid'⁹ as the distance metric in the *proximities* syntax. The proximities values were then summed for each component to give a total square *Euclidean* distance- an overall measure of how dissimilar each component is from all the other UKTLCS components.
- *Step 2* –We then ranked the components at each location by (ascending) size of the total square Euclidean distance. From the ranked-component lists, we selected the first 400 at location D and 200 from each of other three locations, giving a total of 1000 components for analysis. 1000 is the maximum number of cases the cluster algorithm in PASW17 can process.
- *Step 3* – Using the proximities matrix for the 1000 selected components as input, we conducted hierarchical cluster analysis, starting with the simplest two cluster solution and successively incrementing the number of clusters to 10.

The cluster analysis results were inconclusive: The primitive, two cluster, solution showed that the majority of components from location D form one cluster. Unfortunately, however, many of the components from location C also appeared in the same cluster as location D. When we increased the number of clusters, we were still not able to see a clear separation of the clusters for components from locations C and D. Overall, therefore, these results suggest that the demand patterns from both locations are too similar (or not dissimilar enough) to be separated, so that rather than location, per se, the commonalities may be due to other factors. This in turn suggests that it may be erroneous to place too much weight on the location criterion when choosing forecasting methods.

5.6. A Monte Carlo simulation approach to parameter estimation

To study how, for each of the 92 time series', forecasting performance may vary with the parameters used in the forecasting method, we simulated the demand for each of the time series based on its historical demand probability distribution. For each component, we simulated the demand for 18 periods, and then calculated the

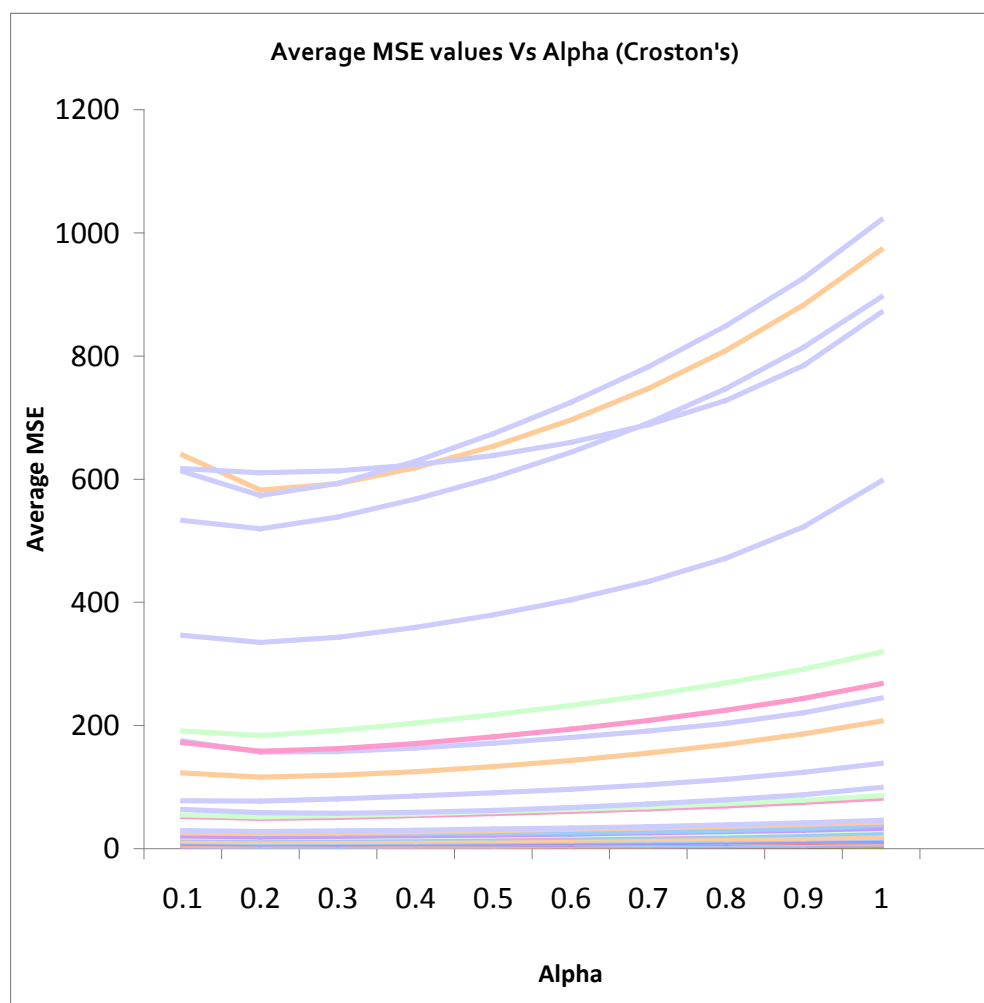
⁹ This calculated the square of the Euclidean distance between any two component demand values as the as the measure of proximity between them. The reason we squared the Euclidean distance was to magnify the dissimilarities between components, giving us the best chance to extract clear clusters

mean square error for forecast of the simulated demand values based on (i) Croston's (Croston 1972)¹⁰ and (ii) Simple Exponential Smoothing (SES) (Gardner 2006, Billah *et al.* 2006). Each simulation was then replicated 10,000 times and the average MSE for both forecasting methods calculated for each component. This process was then repeated for different parameter values in both forecasting methods.

The simulation results were insightful. When the SES forecasting method was used, the variation in the average MSE values suggests that the optimal alpha value is in the range 0.1-0.3, and close to a value of 0.2. Given that this optimal alpha range is fairly narrow, using SES with a selected alpha setting within the optimal range, say $\alpha = 0.2$, could therefore lead to sufficient forecasting performance.

¹⁰ For the purpose of this study the two parameters for demand size and demand interval will be treated as the same.

Figure 10 MSE Curves of Monte Carlo Results for Croston's Method



However, as suggested before, the MCA smoothing method is most likely a variant of Croston's intermittent demand method, making the exploration of the performance of Croston's method more important. Figure 10 shows the average MSE results obtained using Croston's method with different alpha values. The optimal range (0-0.3) of parameter values for Croston's was wider than it was for SES. However, the turning point in the curve is around the same value, i.e. 0.2. Outside of the suggested optimal range, MSE values often increased substantially, showing a need to set alpha values within the optimum range. Exploring the distribution of the univariate statistics of interest with the MSE values based on Croston's revealed that when the skewness was large, the results tended to suggest that the parameter should be set as low as possible, even at a 0-0.1 level.

Overall, the differences in performance between Croston's and SES were marginal. SES performed better with very slow moving items, whilst Croston's performed better with faster moving items with some intermittence.

5.7. Evaluation of SPO Forecasting performance- Summary

In summary the results in this section provided an insight into the performance of the existing forecasting techniques, and suggested ways in which SPO performance could be enhanced. First, the evidence suggests that setting the MCA smoothing parameter at 0.01 lessens forecasting performance. Second, while using a single parameter value may be sufficient to obtain good forecasting performance, forecasting may be improved by dynamically fine tuning the parameters for each component around the indicated optimum. Third and fourth, there appears no justification to base forecasting technique entirely on the location of the components, nor is there significant statistical evidence to support causal forecasting based on flying hours or other presumed important factors such as operational temperature.

6. Development of Enhanced Forecasting Models

Having explored the limitations of existing forecasts, the next stage of the study was to explore possible enhancements to the SPO forecasting. VBA was used to create a fully dynamic model that can compare forecasting techniques. Using the imbedded excel solver, the tool can also attempt to find the optimum parameters for each forecasting method for each time series based on the MSE. By inputting demand data, month and flight hours, the user can then change the settings (apart from the number of months) to suit their requirements. From here the tool does the rest. The time series data are put through each of the forecasting techniques, located in the relevant tabs; and the model fit statistics and all the forecasted projections are displayed back onto the front screen¹¹. The multiple model button is an add-on, for the purpose of creating forecasts on the large scale for a large number of components. Using multiple models however has its drawbacks in that the user may

¹¹ The Holt Winters Multiplicative Method will only work with non-zero data. The system will recognise this fact and produce no forecast statistic results.

not be able to assess the individual time series autocorrelations and graphical patterns very easily. Figure 11 demonstrates the sample VBA coding, showing how the forecast dates¹² can be incremented.

Figure 11 Example VBA Coding Dynamic Forecasting Model

```

'Increments Forecast Date into Cells

With Sheets("LES")

    Range("K9").Value = DateAdd("m", 1, .Range("A" & _
    Sheets("Data_Series_Input").Range("Num_Months").Value + 8).Value)

    For count = 1 To Sheets("Data_Series_Input").Range("Forecast_Period").Value

        .Range("L" & count + 8).Value = count

        If count = Sheets("Data_Series_Input").Range("Forecast_Period").Value Then

            Else

                .Range("K" & count + 9).Value = DateAdd("m", 1, .Range("K" & count + 8).Value)

            End If

        Next count

'*****
'Sets the Range of the Holdout and Test Sets

Set Rnge = .Range("H" & Sheets("Data_Series_Input").Range("Num_Months").Value + 9 _
- Sheets("Data_Series_Input").Range("Hold").Value, "H" & _
Sheets("Data_Series_Input").Range("Num_Months").Value + 8)

    Rnge.Name = "HoldOut"

Set Rnge = .Range("Hg", "H" & Sheets("Data_Series_Input").Range("Num_Months").Value +
9 - Sheets("Data_Series_Input").Range("Hold").Value - 1)

    Rnge.Name = "Test"

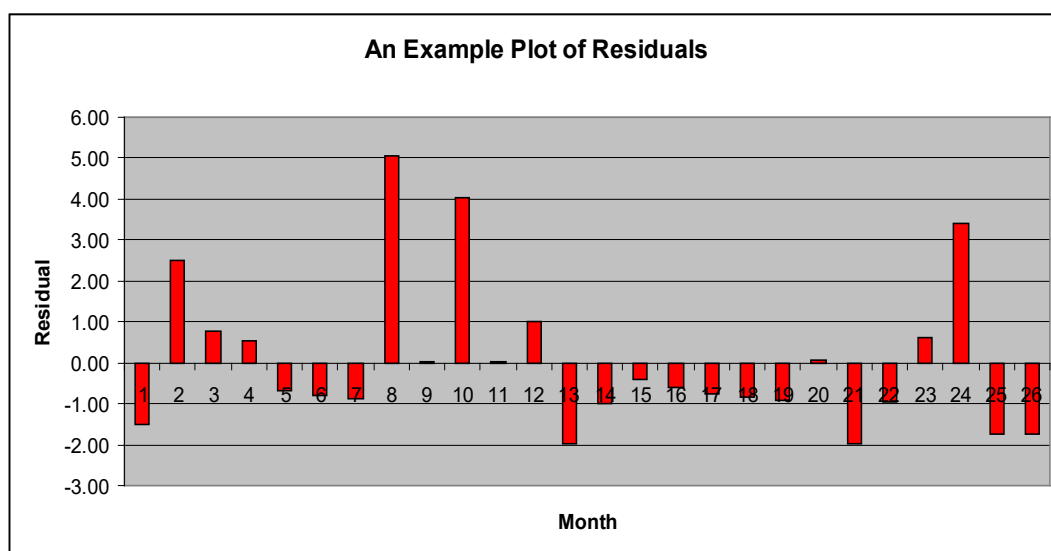
.Range("MSE").Formula = "=AVERAGE(HoldOut)"
.Range("MSE IN").Formula = "=AVERAGE(Test)"

```

In section 5.6, it was suggested that when it comes to selecting the forecasting parameters, two strategies could be adopted: (i) to use a generic fixed (non-optimised) parameter value in all forecasts or (ii) to use optimised parameter

values. Using the VBA tool both strategies were explored. Although it has been claimed that adaptive models can learn and can therefore forecast more accurately, this claim is not supported by empirical evidence (Makridakis *et al.* 1982, Carmo *et al.* 2004). The non-optimised models were run with all parameter values set at a default level of 0.1, normally the recommended level in this type of demand. All models were run with a hold out period of 20 months and nine months of initialisation, with the residuals of each forecast being examined (Figure 12).

Figure 12 An example plot of residual for a selected forecasting technique



6.1. Results: Optimised parameter Models

The optimised parameter estimation models gave an indication of the best fitting model for each time series, based on minimising the MSE (Out). The results were as shown in Table 6, indicating the number of components for which each forecasting method was indicated as the best fitting.

Table 7 Results of the optimised parameter estimations

Forecast Type	Number of Time Series'	Av. Demand	Av. Theil	Av. MAPE	Av. Zero
Naïve	5	7.20	-	3.49%	24.80
SES	4	22.25	0.858	3.63%	15.50
LES	1	12.00	0.791	2.98%	19.00
MA(3)	1	29.00	0.661	3.86%	14.00
MA(6)	2	11.00	0.887	5.94%	20.50
MA(9)	3	154.33	0.849	3.95%	4.33
MA(12)	29	30.24	0.950	7.33%	13.55


Croston's	8	144.88	0.650	3.41%	10.25
SBA Croston's	35	126.23	0.815	4.43%	15.06
Modified Croston's	4	39.50	0.815	3.88%	12.00
HW Additive	0	-	-	-	-
HW Multiplicative	0	-	-	-	-

On the whole, the results tended to go along with expectations. Often, an underlying Croston (1972) method came out on top as the best fitting model. However, it was not Croston's method itself but the bias reduction method of Syntetos and Boylan (2005, 2006), denoted SBA in table 7, which produced the better results. This suggests that the overestimation problems, outlined by Syntetos and Boylan (2005, 2006), were overcome for the components for which SBA is indicated as the best fitting model. Components which suggested the use of a Naive forecast were often those with early demand followed by an extended period of no demand. When using the Croston's (1972) method with this type of data, another update of the demand was not realised over the observed life of the component. This generally meant an overestimation of the overall demand later on. In contrast, using a Naive forecast, an update in the forecasted demand was created straight after the demand realisation of the previous demand magnitude.

Expectations of good results for SES were slightly dampened; however, in general, it did perform well in each of the time series. Its versatility therefore represents its key benefit, but with an abundance of dedicated tools, it should only be considered when the number of series is large and time is short.

The surprising results of the optimised parameter models came from the number of series for which the best fitting model was a 12-month Moving Average of [MA (12)]. However, with the highest average Mean Percentage Error (MAPE), this, in fact, suggests difficulties in trying to fit forecasts to that type of series. Many of the components which suggest the use of MA (12) have very similar time series characteristics. Figure 13 gives an indication of the type of time series.

Figure 13 Typical time series fit with MA(12)



Many of the patterns show higher demand values in earlier months (12-18 old). This suggests that intermittent modelling was not necessarily the best method when demand values have decreased. We observed that outliers (and the length of the series) may have a significant effect on the results produced. To test this proposition four items which produced a *MA(12)* forecasting method were selected and processed in finer detail. Using a judgement approach on the series length and *windsorisation* to detect outliers, the new and old methods were compared¹³. The results (Table 8) give an indication that being dynamic with the forecasts will increase performance and often decrease holding costs. All the *MSE (Out)* values decreased significantly and dedicated intermittent modelling techniques proved very good candidates for forecasting the time series'. As such, the use of an outlier detection method can be a big factor in increasing the accuracy of the results.

Table 8 Comparison of items, Pre & Post Manipulation

	<i>Item 18</i>	<i>Item 54</i>	<i>Item 60</i>	<i>Item 61</i>
Time Series Length	15	19	17	19
Outlier Detection	Yes	No	Yes	No
Original Method	MA(12)	MA(12)	MA(12)	MA(12)
MSE (Out) Original	3.38	13.33	16.44	44.00
Forecast	1.67	3.33	4.08	4.00
New Method	MA(12)	Croston's	SES	SBA
MSE (Out) New	1.58	12.09	9.56	6.54
Forecast	1.58	3.06	3.16	4.02
Difference	-0.09	-0.27	-0.92	+0.02
Percent Difference	-5.4%	-8.0%	-22.5%	+0.5%

¹³ Due to the inevitable decrease in time series size, seasonality methods were not compared as the best fitting modelling methodologies.

It is interesting to note that the parameters were often set, by excel solver, at very conflicting values. In the case of second parameter estimations which fit trends in the data, we mostly saw beta values close to zero. This means that, based on the best MSE (Out), excel solver was trying to suggest there were limited cyclical trends. This suggestion is backed up by the fact that very few models were fit with a beta trend factor.

6.2. Results: Non-optimised parameter Models

Table 9 shows the summary of the results when the models were run with a generic non-optimised parameter value of 0.1. Compared to the results of the optimised parameter models (table 8), it can be seen that a shift from intermittent and exponential smoothing methods has occurred, with many of the time series' fitting a moving average model. However, we also observe that the average MAPE is far higher in the moving average models, suggesting lower accuracy when the demand data are fit with these models.

Table 9. Results of the non-optimised parameter estimations

<i>Forecast Type</i>	<i>Number of Time Series</i>	<i>Av. Demand</i>	<i>Av. Theil</i>	<i>Av. MAPE</i>	<i>Av. Zero</i>
<i>Naïve</i>	0	-	-	-	-
<i>SES</i>	1	12.00	0.831	3.45%	19.00
<i>LES</i>	12	238.50	0.855	3.57%	11.08
<i>MA(3)</i>	2	123.00	0.763	2.91%	6.50
<i>MA(6)</i>	6	12.00	0.861	5.04%	20.17
<i>MA(9)</i>	4	117.50	0.866	4.49%	8.75
<i>MA(12)</i>	42	26.78	1.008	8.84%	14.17
<i>Croston's</i>	0	-	-	-	-
<i>SBA Croston's</i>	0	-	-	-	-
<i>Modified Croston's</i>	3	119.33	0.660	3.40%	8.00
<i>HW Additive</i>	14	113.64	0.751	4.25%	14.07
<i>HW Multiplicative</i>	8	67.88	0.841	4.97%	15.13

6.3. Overview: Optimised Versus Non-Optimised Models and Original SPO Forecasts

In general, the intermittent demand models worked well when parameters are both optimised and non-optimised. Instead of the Croston's method however, it is the revised version of this, the SBA model, which performed best overall. With

non-optimised parameters, for the majority of the time series', the best fitting models were moving average forecasting techniques, but with less accuracy.

Next we compared how well the optimised models perform with respect to the forecasts based on the existing SPO system. Table 10 shows the summary statistics of the ME values across the 89 time series that we were able to compare.

Table 10 A comparison of Enhanced (Optimised parameters) models with Original SPO Forecasts

<i>Location</i>	<i>Summary Statistic</i>	<i>Enhanced (optimised parameters)</i>	<i>Existing SPO Forecasts</i>
All locations (N = 89*)	Mean	-0.25	-1.12
	Std. Deviation	1.03	4.06
A (N = 8)	Mean	0.44	0.09
	Std. Deviation	1.89	5.76
B (N = 2)	Mean	-1.00	-0.65
	Std. Deviation	1.41	0.50
C (N = 39)	Mean	-0.39	-0.70
	Std. Deviation	1.00	1.45
D (N = 40)	Mean	-0.21	-1.80
	Std. Deviation	0.76	5.31
* As before three time series' were not included due to inability to establish the exact SPO method used.			

It can be seen in table 10 that, overall, the optimised parameter models performed better than the original SPO models and also exhibited less variance in accuracy across all time series'. Series by series, the optimised parameter models were better 72% of the time based on ME values and 65% better based on MSE values. It is noteworthy that the biggest improvement in forecasting performance occurred at location D, the location where previously forecasts had been chosen based on location. As the components in the analysis were not randomly selected, we did not test whether the perceived superiority of the optimised parameter models over SPO models is statistically significant. Therefore, these results cannot be generalised to other situations.

Overall, therefore, this comparative study suggests that minimising with respect to MSE can increase the 'fit' towards the demand pattern and it was recommended

that the method of optimisation be used to determine the best model in the analysis tool provided, which can be used as enhancement to the SPO system.

7. Conclusions

The aim of this project was to explore and recommend any enhancements to the SPO system that Boeing UKTLCS utilises to forecast demand for components used in the maintenance of Chinook helicopters on behalf of the RAF. Focussing mainly on the 92 components that Boeing consider most influential on maintenance cost, we conducted research in three parts.

First, we explored the time series' of the 92 components, individually and simultaneously. We uncovered diverse demand patterns. The diversity was evident in how intermittent demand is, in the number demanded when demand does occur and in the distribution of non-demand and demand periods. It is clear, as a result, that for Boeing, forecasting demand for components is an onerous task as the likelihood of misclassification when identifying classes of demand patterns is considerable.

Second, we examined how well the extant SPO system performs. We examined, in particular, the relative merits of the criteria used to select forecasting methods, the forecasting methods chosen and how, when forecasting methods are chosen, the forecasting parameters are selected. Using cluster analysis, we looked at the use of location as a selection criterion. We were unable to achieve clear clusters between one location where location is used to select forecasting methods and another location where forecasting techniques are chosen independently of the location. While we note that this does not preclude the existence of considerable differences in other inventory characteristics between the two locations, the lack of cluster separation along with reports of 'over-buying' from asset managers about the former location leads to the suggestion that the location criterion ought to weigh less in selecting forecasting methods. Another criterion used to select forecasting methods is the assumption of causality of component demand by the number of flying hours of the Chinooks and other operational environment variables such as

temperature. We were not able to find a strong or even weak relationship between flying hours and demand. Nor did we find correlation between operational temperature and demand. Therefore we did not find any evidence to vindicate the decision to use causal forecasting for some components. Comparing the methods used to forecasting demand, we found that, overall, the MCA smoothing method which is used for most of the SPO forecasting was more accurate than other methods used. When we turned our attention to the selection of parameter values for forecasting, we were able to demonstrate that a decision taken to calm down forecast values by reducing the alpha value could have been inadvertently reducing forecasting accuracy.

Beyond revealing the listed deficiencies, the evaluation of SPO also suggested ways in which SPO enhancements could be achieved in the third and final part of the project. Notably, Monte Carlo simulation indicated that the trough of the MSE curves using Croston's and SES techniques, for all 92 series, occurred for a reasonably invariant range of forecasting parameter values. As such, it may be sufficient to pick a single fixed value within this optimal range of parameter values, as was the practice at the time of this research. Concurrently, the simulation results indicated that while the optimal range of parameter values was similar for all 92 series, the actual optimal point did vary across time series. Therefore a second way to choose parameter values could be by working out optimal parameter values for each individual time series, i.e. dynamic forecasting. We created a VBA tool to investigate not only the relative merits of each of these approaches but also allowed us to indicate which, among all (both for non-intermittent and intermittent data situations) commonly used forecasting techniques was the best fit for each time series. Overall, dynamic forecasting was superior. When we inspected the best fitting method for each time series, the results were at once expected and surprising. For most time series, the best fitting method was either Croston's itself or a variant of it, in particular SBA. This is as was expected given that intermittence was observed in the data. But in a surprising number of time series', the less sophisticated SMA was a better fit. We investigated this in finer detail and came to the conclusion that SMA was a better fit when a time series suffered from skewness.

Generally, such time series' were more difficult to forecast accurately, although in some of them accuracy could be increased by using outlier detection.

Overall, based on averaged ME values across the time series', it was shown that forecasts based on the VBA tool that aids the selection of forecasting method and finds the optimal parameter values for the selected method would lead to more accurate forecasts than those generated by the existing the SPO system. When 'what if analysis' (not presented in this paper) incorporating the VBA tool results in the MCA system was conducted, it was shown that the overall savings in inventory costs of adopting the tool would run into millions of dollars. Boeing adopted the VBA tool as an enhancement to the SPO system.

In terms of originality, this study is an enhancement to earlier outline studies undertaken by the authors (see Downing *et al.* 2011). In the first place, in this study, an evaluation of the component demand time series' is undertaken. Secondly, the authors assessed the accuracy of forecasts for each individual time-series adopted in addition to conducting an inspection of the performance of individual time series'. The authors also as an original contribution, conducted a robust evaluation of the performance of the SPO forecasting system by rigorously interrogating data to examine (i) the performance of the SPO forecasting system beyond a proposition stage (ii) the impact of MCA smoothing alpha parameter on performance, (iii) how flying hour's impacts on demand, and (iv) how forecasting methods could be chosen based on location.

This project, thus, demonstrates the value of examining the demand pattern of each component in detail in a case such as Boeing's UKTLCS maintenance where the contribution to the overall maintenance cost of components is so unequal that a very small proportion of components contribute the most. We envisage two ways in which this research can be built upon in future: one extension would be to investigate how based on the forecasting methods found to be best fitting for the 92 selected components, the VBA tool can be used to set forecasting methods for the remainder of the 12, 000+ components in the Boeing portfolio. The second

would be to investigate how the results reported here can be integrated into a maintenance model, combining the analysis of demand for components with maintenance decisions both strategic and operational about buying/leasing, reliability, usage, repair/replacement, etc., of components.

References

- Argyris, C. and Schon, D., 1991. Participatory action research and action science compared: a commentary. *In: W. Whyte, ed. Participatory Action Research*. Pub. Sage
- Boylan, J., Syntetos, A. and Karakostas, G, 2008. Classification for forecasting and stock control: a case study. *Journal of the Operational Research Society*, 59, 473-481.
- Carmo, J., António, A. and Rodrigues, J., 2004. Adaptive forecasting of irregular demand processes. *Engineering Applications of Artificial Intelligence*, 17 (2), 137-143.
- Clark, A., 2005. Rolling horizon heuristics for production planning and set-up scheduling with backlogs and error-prone demand forecasts. *Production Planning and Control*, 16 (1), 81-97.
- Croston, J., 1972. Forecasting and Stock Control for Intermittent Demands. *Operations Research Society*, 23, 289-304.
- De Gooijer, J. and Hyndman, R., 2006. 25 years of time series forecasting. *International Journal of Forecasting*, 22 (3), 443-473.
- Denicoff, M., Fennell, J. and Solomon, H., 1960. Summary of a method for determining the military worth of spare parts. *Naval Research Logistics Quarterly*, 7 (3), 221-234.
- Downing, M., Chipulu, M., Ojiako, GU. and Kaparis, K., 2011. Supply chain forecasting for the UK Chinook fleet. *OR Insight*, 24, 110-130.
- Graman, G. and Sanders, N., 2009. Modelling the tradeoff between postponement capacity and forecast accuracy. *Production Planning and Control*, 20 (3), 206-215.

- Ho, C. and Ireland, T., 1993. A diagnostic analysis of the impact of forecast errors on production planning via MRP system nervousness. *Production Planning and Control*, 4 (4), 311-322.
- Huberty, C. and Morris, J., 1989. Multivariate Analysis Versus Multiple Univariate Analyses. *Psychological Bulletin*, 105 (2) (March), 302-308.
- Johnsen, T., Howard, M. and Miemczyk, J., 2009. UK defence change and the impact on supply relationships. *Supply Chain Management: An International Journal*, 14 (4), 270-279.
- Karr, H., 1958. A method of estimating spare-part essentiality. *Naval Research Logistics Quarterly*, 5 (1), 29-42.
- Kennedy, W. J., Patterson, J. W. and Fredendall, L. D., 2002. An overview of recent literature on spare parts inventories. *International Journal of Production Economics*, 76(20), 201-215.
- Makridakis, S., Wheelwright, S. C. and Hyndman, R. J., 1998. *Forecasting: Methods and Applications*. John Wiley & Sons, Inc.
- Markland, R., 1970. A comparative study of demand forecasting techniques for military helicopter spare parts. *Naval Research Logistics Quarterly*, 17 (1), 103-119.
- Michaels, L., 1999. The making of a lean aerospace supply chain. *Supply Chain Management: An International Journal*, 4 (3), 135-144.
- Syntetos, A. A. and Boylan, J. E., 2001. On the bias of intermittent demand estimates, *International Journal of Production Economics*, 71, 457-466
- Syntetos, A. A. and Boylan, J. E., 2005. The Accuracy of Intermittent Demand Estimates. *International Journal of Forecasting*, 21, 303-314,
- Syntetos, A. and Boylan, J., 2006. Comments on the Attribution of and Intermittent Demand Estimator. *International Journal of Forecasting*, 22, p.202.
- Syntetos, A. A., Babai, M. Z., Dallery, Y. and Teunter, R., 2009. Periodic Control of Intermittent Demand Items: Theory and Empirical Analysis. *Journal of the Operational Research Society*, 60, 611-618.
- Syntetos, A.A., Boylan, J. E. and Disney, S.M., 2009a Forecasting for Inventory Planning: a 50-review. *Journal of the Operational Research Society*, 60, 149-160.

- Syntetos, A., Nikolopoulos, K., Boylan, J., Fildes, R. and Goodwin, P., 2009b. The effects of integrating management judgement into intermittent demand forecasts. *International Journal of Production Economics*, 118 (1), 72-81.
- Theil, H., 1966. *Applied Economic Forecasting*. Amsterdam-Holland Publishing Company.
- TysseLand, B., 2009. Spare parts optimization process and results: OPUS10 cases in the Norwegian Defence. *International Journal of Physical Distribution & Logistics Management*, 39 (1), 8-27.
- Venkataraman, R. and Nathan, J., 1999. Effect of forecast errors on rolling horizon master production schedule cost performance for various replanning intervals. *Production Planning and Control*, 10 (7), 682-689.
- Viertl, R., 2006. Univariate statistical analysis with fuzzy data. *Computational Statistics & Data Analysis*, 51 (1), 33-147.
- Wahab, D., 2008. *Neural Network Models Using SAS Enterprise Miner*, 22nd Annual SAS Malaysia Forum, Kuala Lumpur Convention Centre: University Malaya.
- Wallstrom, P., 2009. *Evaluation of Forecasting Techniques and Forecast Errors - With Focus on Intermittent Demand*, Licentiate Thesis. Lulea University of Technology.

Appendix A Abbreviations

ARIMA	Autoregressive Integrated Moving Average
COV	Covariance
HW	Holt Winter's
HYD/DYN	Hydraulic Dynamic
ID	Inter Demand Interval
LES	Linear Exponential Smoothing
MA	Moving Average
MK ₂	Mark 2
MODCR	Modified Croston's Method
MPE	Mean Percentage Error
MSE	Mean Squared Error
NF ₁	Naive Forecast
OBS	Observations
PB	Percentage Best
SES	Simple Exponential Smoothing

SYNBOS

Syntetos & Boylan Approximation Method

TSL

Target Stock Level

VBA

Visual Basic for Applications