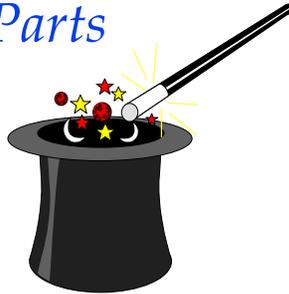


SPC & Selectable Calendars Work Magic For GE Aircraft Engines' Service Parts



by
Jeff Beck
Forecasting/Inventory
Control Analyst
at
**General Electric
Aircraft Engines**
Cincinnati, Ohio

Reprinted from APICS 38th International Conference Proceedings

Introduction

In two years' time, GE Aircraft Engines overhauled its service parts operation. The result? Customer service is about the same or up slightly (depending on how you measure it), while inventory is down 25%. To do this took more people, right? Wrong! Planning personnel was reduced 30%. A better result with less effort! What's the secret? We used statistical process control (SPC) techniques and selectable forecast calendars to renovate the service parts business.

A Pareto analysis convinced us it makes little sense to treat all our 8,000 service parts alike. In order to concentrate our attention on the top 5% of the parts (which comprise 80% of the business), we needed a more powerful way to manage the other 95% of the parts. We implemented a mix of standard and specialized SPC techniques to manage that 95%.

We also describe the applicability of selectable forecast calendars for tough forecasting problems; which provide more accurate forecasts while requiring just a fraction of the effort of monthly forecasting. In our case this technique alone reduced our forecasting workload by 43%.

Acknowledgments

I would like to thank John R. DiPaola, Manager of Inventory Management, and Ron F. Baker Director of Spares Operation, without whom this project could not have succeeded. John led the drive to find a solution for our inventory and forecasting problems, and Ron provided the necessary resources to insure its success.

Background

We provide service parts for GE Aircraft Engines' commercial jet engines. Originally we had 12 inventory planners, each handling a variety of engine parts, by engine section. Each had parts

whose annual usage (at cost) ranged from a dollar to hundreds of thousands of dollars. Each planner also had to deal with parts which had lumpy and intermittent demand, as well as parts with relatively smooth demand. There was little consistency in the practices used by the inventory planners.

We did have a system which computed a forecast, but it was essentially limited to one number: the average monthly demand. There was no provision for demand which was increasing or decreasing, nor was

Summary & Conclusions

❖ Two Principle Methods To Our Success:

- ❖ Selectable Forecast Calendars.
 - ◆ Reduced Forecast Error, Therefore Reducing Safety Stock.
 - ◆ Reduced Our Work load By 43%.
- ❖ Statistical Process Control Tools.
 - ◆ Identifies The Parts That Offer The Greatest Potential Benefit From Taking Action.

❖ In Three Years Time:

- ❖ We Reduced Inventory 25%.
- ❖ Inventory Policy Has Dropped 30%.
- ❖ Customer Service Level Has Improved.

there any way to plan for seasonal variations in demand. The forecast also failed to recognize that with GE's 5-4-4 calendar, there was no such thing as an "average month." We needed to know if our average monthly demand was for an average four-week month or for an average five-week month! With aircraft flying every day, that makes a big difference.

The system was also weak at tracking forecast error. It calculated an error statistic, but that was a mean absolute deviation (MAD), which is a poor substitute for a standard deviation. Furthermore, the MAD was calculated only for the computer-generated forecast, not for forecasts which were sometimes overridden with a technical forecast. This meant that the effort of gathering and inputting accurate intelligence went unrewarded by corresponding safety stock inventory reductions. And if the intelligence made things worse, the increased error was not covered by the safety stock, thereby hurting service. To top it off, the 12 inventory planners followed no consistent practices for setting inventory and service.

The widely-publicized airline industry difficulties forced us to find a way to reduce our costs. Our inventory planners were cut from 12 to 8. At the same time we were under pressure to reduce our service parts inventory. Clearly we had to do something different to reduce our inventory and costs, yet at the same time continue to provide competitive service to our customers. There just wasn't any way we could do this given our inventory planning process at the time.

Organizing For Change

We reorganized our commercial spares inventory organization into two areas. Five planners were assigned the top 5% of the parts which accounted for roughly 80% of our business. We created a "Pull Production" process where the planners essentially micro-manage these parts. Since this is the great majority of our business, it is worth the time spent on these parts.

The remaining three planners were given the task of managing the other 95% of the parts in some highly automated fashion. We looked at various forecasting packages, and even experimented with some on the PC. While this activity was useful to us from a forecasting education point of view, we found that most packages suffered from (1) inability to handle large volumes of parts, (2) inability to automate the forecasting process so as to require minimal people time, (3) lack of integration between the forecasting and inventory management processes, (4) lack of a good

solution for dealing with our lumpy demand items (the majority of our parts), or (5) inflexibility.

We installed *The Finished Goods Series* software from E/Step Software Inc. of Tieton, Washington. FGS is a flexible PC-based integrated demand forecasting and inventory management package which, even though it runs on the PC, handles large volumes of data and is designed to exchange information with our mainframe scheduling system. While it is set up to handle large numbers of parts routinely (i.e., hands off), it has SPC tools for identifying and reviewing those exception items which require something other than routine handling.

Selectable Forecast Calendars

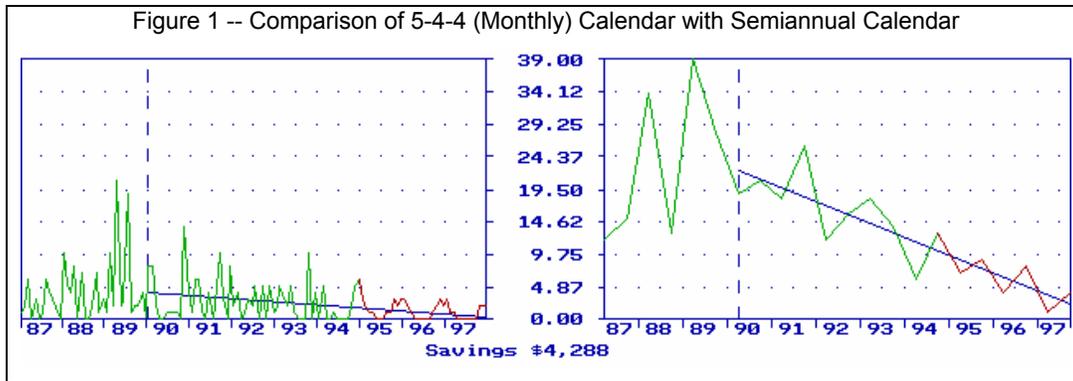
The most glaring characteristic of the parts we trying to forecast was the demand patterns were not smooth. We needed to find a tool that could handle lumpy and intermittent demand patterns. After months of trying different alternatives (i.e. Poisson distribution, Winters exponential smoothing etc.) we found the best way to handle parts with erratic demands was to use less frequent forecast calendars. This tool enabled us to smooth out the large spikes and valleys in our demand which has resulted in not only better forecasts but also reduced error therefore reducing the amount of safety stock required.

Another important benefit of forecasting less frequently is the reduced work effort. By forecasting something quarterly, or semi-annually your work load is reduced 67%, or 83% respectively. For more information on the justification and appropriateness of this technique see Reference 1.

Steps to Implementation

We randomly selected 500 parts which we used to serve as a pilot operation. Our objective was to determine how to use selectable forecast calendars to solve our forecasting problems. From the results of our analysis we developed several "rules of thumb" which we could follow when implementing the rest of the parts. Why not just load everything at once? We could have; but as I mentioned earlier, we had a lot of learning to do, and the most effective way to learn about forecasting and statistical inventory management is to experiment. The more experiments we ran, the more we learned. It was much better to experiment on a few hundred parts, getting the answers quickly, than to experiment on 10,000 parts and have to wait for our answers. Once we had some answers we could scale our results up to the entire group of parts.

Since we have three planners, we divided our parts into separate databases, by inventory class (A,B,C)



rather than by engine section. We also created a fourth database to monitor our “Pull Production” items (the top moving parts), which I’ll refer to as our “A+” database. In this way each inventory planner can work independently, but we have the ability to combine the separate databases by creating a summary part to get reports of grand totals. As you would expect, the A+ database has just 6% of the parts, the A database has 14% of the parts, B has 33%, and C has 47%—close to the classic ABC definition (with A+ and A combined).

Determining the Initial Forecast Calendar

We could have started forecasting everything monthly, but our pilot study taught us that selectable forecast calendars were required to get the best results. We could also have tried every calendar for every part and selected the calendar which resulted in the lowest forecast error. In [1], however, Estep recommends taking other factors into consideration. These include such things as replenishment frequency, how the part is used, etc. In our business, we found we could reliably relate it to the level of demand. Parts with lower demand were also the ones with worse errors and were replenished less often. For these parts, it was perfectly reasonable to forecast them less than monthly. By “level” we mean the rate of demand per month at time=now.

The process we settled on was first to run everything through an automatic model fitting process on our monthly (5-4-4) calendar. If the level that resulted was less than 0.3 (or 4 per year), we put the part on the semiannual calendar and refit. If the level was 0.3 or more but less than 5 (60 per year), we put it on the quarterly calendar. For a level between 5 and 10 (120 per year) we used the bimonthly calendar (6 forecast periods per year). Anything with a level above 10 was kept on the monthly calendar.

We did have to decide how to handle negative levels (Also known as all time supplies) You can get a negative level if the demand history shows demand which has declined for so long that it has gone essentially to zero. From a mathematical point of view, numbers don’t end at zero, but proceed right on to the negative numbers. You can’t use a model with a negative level, however, because its implication is that after demand declines to zero for a part, the customers are going to start shipping parts back!

Going to a less frequent forecasting calendar often causes the negative level (on a monthly calendar) to become zero or positive, giving us a useable forecast. So we tried a quarterly calendar if the level on a monthly calendar was negative. If the level was still negative on quarters, we tried semiannual. Any part which still had a negative level was put on an exception list for manual review. There are other reasons, covered below, for putting parts on exception lists, but the vast majority of the parts were OK on autopilot. This means that these parts did not require any manual review, saving us many of hours of time.

Reexamining the Forecast Calendar Decision

When reviewing exceptions, we gather any available marketing intelligence and use a Simulation facility to evaluate alternative decisions. One of the alternatives we often evaluate is changing the forecast calendar. We tell the system to try all the calendars and rank them by forecast error. Figure 1 is an example of one such comparison. For this item, changing from the 5-4-4 calendar to a semiannual calendar reduced the safety stock by 5 pieces, saving \$4,300.

Calendar	Exceptions	# History		Error	Relative Error
		Periods	Next 12 Months		
Semiannual	none	10	16	3.8	51%
Annual	none	5	16	4.1	55%
Quarter	none	20	17	6.6	88%
Bimonth	Error>Level	30	18	6.7	90%
5-4-4	Error>Level	60	18	7.5	100%

Figure 2—Calendar Comparison for One Part

The evaluation table in Figure 2 compares the forecasts as well as the errors.

The same part on these five different calendars have very similar forecasts, ranging from 16 to 18 pieces per year, but very different errors, ranging from 3.8 to 7.5. The error number is the standard deviation of forecast errors adjusted for the lead time. For any desired service level, the safety stocks are a constant multiple of the error. Thus the semiannual calendar, with a 51% relative error, also needs only 51% of the safety stock compared to the monthly calendar--a 49% inventory savings! The models on bimonthly and monthly calendars are marked "high error" which means the unadjusted standard deviation is greater than the level.

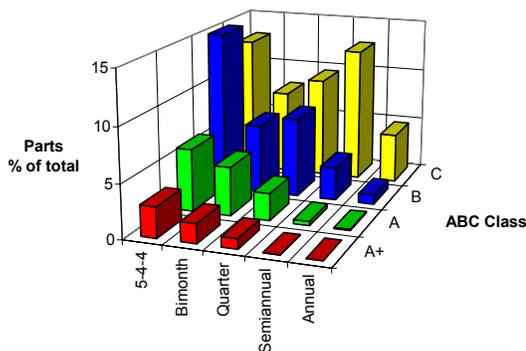
In simulation mode, we see the impact on inventory of each of our forecasting decisions--before we commit to them. Seeing the effect on inventory is the best way to evaluate decisions. The usual alternative is an error percentage, which is little help since a 1% error for an item might be thousands of dollars, while a 200% error on another could be \$1.

Over time, a part may be moved from its initial calendar to some other calendar, typically as demand

Calendar	A+	A	B	C	Total
Annual	0	1	3	10	6
Semiannual	1	3	9	26	16
Quarter	17	19	22	20	20
Bimonth	32	34	20	17	21
<u>5-4-4</u>	<u>50</u>	<u>43</u>	<u>46</u>	<u>27</u>	<u>37</u>
Total	100	100	100	100	100

Figure 3--Percentage of Parts per Calendar per Database & Total

becomes more or less sparse. The table in Figure 3 shows the percentage of items currently on each calendar by database. It certainly contradicts the notion that all parts should be treated alike. It shows, as is



to be expected, that as one moves down the inventory classes, the percentage of parts on the less-frequent calendars increases, due chiefly to the increasing sparseness of the demands. The chart below shows the same data in graphical form. I should also point out that the B and C databases contain the vast majority (80%) of the parts. The combination of these two facts means that the workload reduction afforded by the selectable forecast calendars is applied to the majority of the parts. In fact 63% of our parts are currently on less than monthly calendars! The forecasting workload reduction for this distribution of parts and calendars totals 43%. We do about half the work, and still get a better result!

The Forecast Revision Process

Before we can talk about the SPC reports used to identify exceptions, it is important to understand the difference between the initial model generation process described above and the forecast revision process. The former just gets us to a good starting point, while the latter is important for keeping the model up-to-date and warning us about suspected changes.

Our products, once we get them on the right calendar, have demand that for the most part is relatively stable--as is undoubtedly the case for most businesses. Stability means that if a forecast model truly represents the underlying demand for a part, then that model will be effective over some period of time. You would not expect that one model would work one month, and a totally different model is required next month. There are certainly changes, but they are most often incremental changes, not fundamental changes. For example, we would not expect to see a product with a level, trend, and annual seasonality go to just a level and trend next month and to semiannual seasonality the month after that.

The smoothing and error tracking that occurs in the forecast revision process has two purposes. The first is to make those incremental (not fundamental)

changes which keep the model up-to-date with reality. This handles situations such as a trend which is gradually flattening or seasonality which is becoming less conspicuous. The second purpose of the revision process is to identify those parts where the chosen model is suspected of no longer working (i.e., fundamental, not incremental changes). This means that we are alerted when the pattern of demand changes. Knowing that, we can investigate the cause of the change, be it new competition, product changes, or whatever. The fact that a change has occurred

or is suspected of having occurred is the key. This is counter to the belief that one should try every model on every SKU every month.

SPC Tools for Forecast Monitoring

We use SPC tools--virtual control charts--to monitor our forecasts. By "virtual" I mean that we don't always display and look at a control chart. Often the system does it for us, comparing the statistic being monitored with the applicable control limits. Any part which is outside the control limits (potentially a "bad forecast" by some definition) is placed on an exception list for review. We can use the list to call up the exceptions in simulation, or print them on a report, etc. It is important to remember, however, that the SPC tools allow us to ignore the vast majority of our parts which are doing fine on autopilot. It is only a small minority that are identified as exceptions by our control charts.

We use seven principal SPC tools or charts to look for exceptions:

1. Demand Filter Report
2. Tracking Signals
3. Early Warning Report
4. High Total Stock Report
5. Potentially Bad Forecast Report
6. Suspect Forecast Report
7. Average Demand > (or <) Forecast Report

For each of these tools, we create an exception list (in order of descending importance) so that we can call the list up in simulation for review. While most of these produce reports, some only produce a list, because that's all we need. This whole process is automated using macros so that all we have to do is use the lists and reports when reviewing the exceptions. We'll look at each of these control charts in the following sections.

Demand Filter Report

The Demand Filter Report defends against order entry errors corrupting the forecasts and can also help spot trend changes. Each part has a filter sensitivity, expressed in standard deviations, of about 3.5. The smaller the sensitivity value, the more likely the part is to show up on this exception report; i.e., the more scrutiny it receives by the planners. The sensitivity is translated into a minimum and maximum filter value. Actual demand which is less than the minimum or greater than the maximum causes the part to appear

on this report. The parts are sorted in order from largest to smallest error in dollars. That way, if we do not have enough time to review them all, we know we have looked at the most important ones first.

Tracking Signals

The forecast revision process identifies and creates an exception list of parts that are being forecasted using a model that may no longer be appropriate. The method used to detect such bias is the parabolically-masked cumulative sum of errors technique [Reference 2]--which, in spite of its name, is simple to use. Like the filter exceptions, we control the sensitivity by specifying the number of standard deviations required to trip the alarm. The sensitivity can be set by part, to insure that the most important parts get more scrutiny. We currently use a sensitivity of about 3.

Early Warning Report

Using selectable calendars gives up some visibility, so this report is our way of getting it back. For example, let's assume we expect to sell 2 of an item in the next year; the part is on an annual calendar; and the safety stock is 1. Each month we post the demand for all items (no matter what calendar they're on), but the annual items get revised only after the end of December. Now what happens if by March the year-to-date demand is 4? Normally we wouldn't see this for another 9 months when we do the next revision of parts on the annual calendar.

Part Number	Forecast Calendar	Std Cost	Period-to-date Demand	Next Period Forecast	Safety Stock	Fcst + Sfty Stk	Excess Demand (units)	Excess Demand (dollars)
CB-667	BIMONTH	1,285.71	870	263	481	744	126	161,793
AB-885	BIMONTH	982.22	170	34	125	159	11	11,094
CB-540	SEMIANNL	1,356.51	102	49	50	99	3	4,621
AA-868	SEMIANNL	13.2	148	9	11	20	128	1,695
CC-913	SEMIANNL	116.78	4	2	1	3	1	107
CC-934	SEMIANNL	38.03	1	0	0	0	1	38

Figure 4 -- Early Warning Report

Since we post the demand every month, it's easy to print a list of all items where the demand exceeds the forecast plus safety stock. This lets us discover the situation immediately and take corrective action now, rather than waiting until the end of the forecast period. We use forecast plus safety stock as the control limit, rather than just the forecast, because we would normally expect the demand to exceed the forecast at least sometime in the period about 50% of the time. To help us prioritize our time, we sort the report in descending order of the dollar amount of the excess

demand. We run this report on the parts which are not on monthly calendar, since the Demand Filter Report accomplishes the same thing for monthly parts. In the example in Figure 4, we are most concerned about the top parts on the report. We'll get to the others only if time permits.

We run a second version of this report that catches parts where demand is too low. That's a bit difficult on parts with very sparse demands, but we have arrived at a method which works well. If we are going to worry about parts where period-to-date demand exceeds forecast plus safety stock, then it (in some sense, at least!) makes sense to be concerned with parts for which demand is less than forecast minus safety stock. We make this test a little tougher to satisfy by limiting it to those whose demand is less than 30% of the forecast minus the safety stock. This helps identify parts with decreasing demand.

High Total Stock Report

This report catches those parts with a planned inventory in excess of a year's demand. We compare the total of safety stock plus working stock. The reason for this is that in our business a great many parts are replenished infrequently. If a part is replenished once every six months, it only has two opportunities per year to stock out. This means that sometimes the working stock alone buys us substantial customer service,

Part Number	12 Month		-----Total Average Stock-----		
	Fcst (units)	Std Cost	Dollars	Months	Days
CB-667	1,420	1,285.71	694,490	24.1	723
CG-040	7,812	329.35	252,409	12.0	365
BA-710	3,241	153.22	106,375	15.2	456
BG-584	2,578	244.36	96,118	16.1	483
BC-177	151,741	3.34	66,424	12.9	387

Figure 5--High Total Stock Report

perhaps enough to reach our service target without the need for any safety stock. Furthermore, there are frequently instances in the service parts environment where the forecast errors are higher than the EOQ. In such cases one can minimize the total stock (working plus safety) by raising the lot size to equal the forecast error. The increase in working stock is more than offset by the decrease in safety stock. [3]

Figure 5 shows an example of our High Total Stock Report. We use a one year control limit because

when you have many items with annual usage of 1 and safety stock of 1, it's easy to have a large amount of slow-moving inventory (in pieces, if not in dollars). Notice that the report is sorted in descending order by planned inventory dollars. We could also have sorted it by days. Each is valuable for identifying parts which merit the planner's attention.

Potentially Bad Forecast Report

This report shows us all parts where the forecast error (standard deviation of forecast errors not adjusted for lead time) is greater than 80% of the next year's forecast. These are truly large errors, when you con-

Part Number	36 Month Avg Hist	60% of Avg Hist	ABS(Avg Hist-	Forecast /Rev.	12 Month Forecast	Stdcost Times Std Dev/LT
			(12 Mo. Fcst /Rev))			
CB-667	19.67	11.80	23.17	42.83	257	35,080.77
AB-885	17.70	10.62	23.47	41.17	247	25,161.36
CB-540	1.88	1.13	1.38	0.50	6	11,334.97
AA-868	1.33	0.80	0.83	0.50	3	5,052.17
CC-913	32.25	19.35	13.17	45.42	545	1,158.87
CC-934	66.17	39.70	50.17	116.33	698	775.48

Figure 6--Potentially Bad Forecast Report

sider that the definition of a lumpy forecast is one where the error is greater than the level! We sort this report in descending dollars of forecast error. See Figure 6.

Suspect Forecast Report

The Suspect Forecast Report catches items which are potentially not so bad as the previous report, but still may be worth reviewing. The criteria used for selecting parts on this report is all parts where the forecast exceeds 160% of the average demand in the past 24 months. Parts are sorted in descending order by dollar value of the forecast error.

Average Demand > (or <) Forecast Report

This report is similar to the Early Warning Report in concept, but differs in execution. Here we compare the average demand for the past 12 months to the next period's forecast plus safety stock. This is actually two reports and two exception lists. The first looks for parts where the average demand is greater than the control limit. The second covers parts where the forecast is less than 60% of the average demand. As with the Early Warning Report, we only do this for the parts which are not on months.

Reviewing Exceptions

Earlier we discussed changing the forecast calendar in response to an exception. This is just one of many actions we could take. The idea is to discover why the exception occurred and then fix it. Often it is

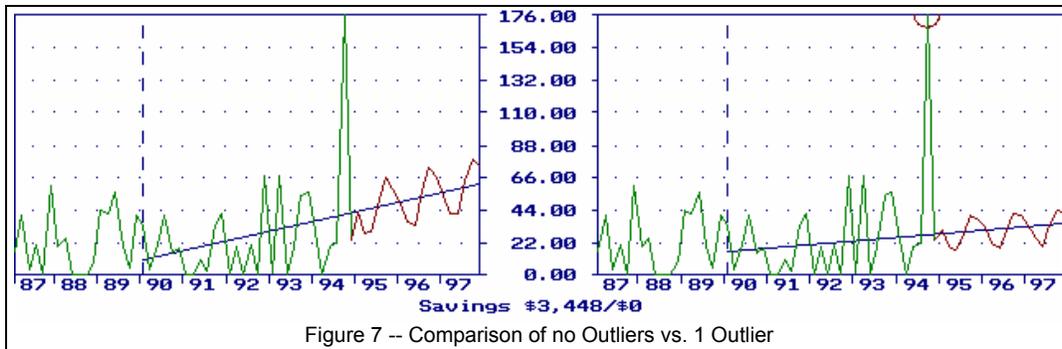


Figure 7 -- Comparison of no Outliers vs. 1 Outlier

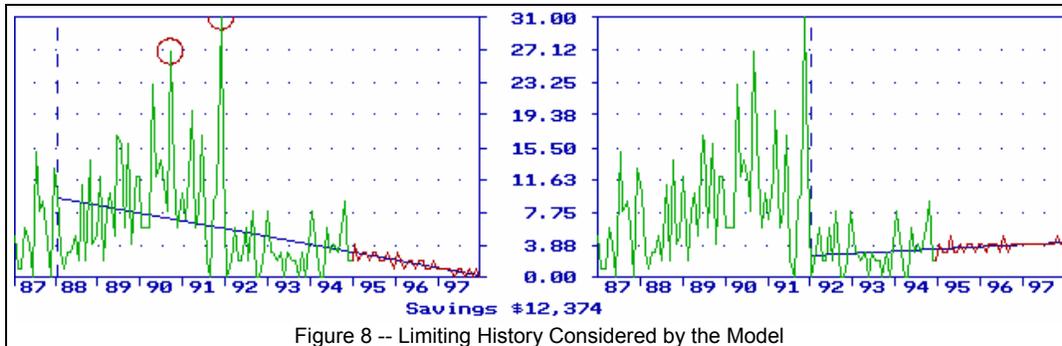


Figure 8 -- Limiting History Considered by the Model

necessary to obtain some outside (i.e., from someone not in the forecasting department) marketing intelligence. When we do this, we use a Marketing Intelligence Evaluation Report to tell us whether the override to the forecast helped or hurt. Frequently the cause of the exception is an incipient change in the pattern of demand. When we discover the cause for the change, the action we take is often that of discarding outliers or setting a demand history limit.

Discarding Outliers

Outliers are periods of demand history which are so far from normal (either high or low) that they can be ignored as irrelevant. This could be because of a one-time retrofit program, for example. The system contains an outlier sensitivity for each part which allows us to control how likely it is for a period of demand to be ignored as an outlier. The sensitivity is calibrated in standard deviations, with a typical value of about 4. If we want to exclude a particular point or points we just tighten the value and refit in simulation. Figure 7 shows an example of tightening the sensitivity to exclude 2 periods. The resulting forecast is much flatter than before. The error in the forecast is also less, giving us an inventory savings of \$4,400. Of course one should not discard data which is, in fact, representative. We can tell when we're going too far because while a sensitivity of 3 or 4 standard deviations is reasonable, when we have to get down below 2 standard deviations to discard a value, we are likely making a mistake.

Demand History Limits/Pattern Changes

When the pattern of demand changes we need to react to correct the model. Looking at the demand history for the part in Figure 8, the demand was first climbing for 4 years; then it declined for 1-2 years; and leveled off for the last 2-3 years. Using demand history limits we fine tune the precise date when the change became apparent. This allows the system to use the maximum history possible to get a better model, but without using history which is inappropriate to the present.

In Figure 8 the system identified a pattern change but was not aggressive enough in discarding past history. The graph clearly shows the leveling off of the demand, so we moved the date of change to January of 92. The resulting forecast better fits the relevant data, and the forecast error reduced by 48%. This 48% reduction also applies to the safety stock, saving us the tidy sum of \$12,400. Not bad for five minutes' work!

Summary and Conclusions

In three years of use, we reduced our inventory by 25%. Our inventory policy has gone down 30%, but with as many slow-moving parts as we have, it takes a while for actual inventory to fall to the new policy level. At the same time we have held customer service level constant or raised it slightly, depending on which service measurement is used. We have achieved all this even though we were limited to a smaller staff!

The two principal methods which have led to this achievement are our use of selectable forecast calendars and the SPC tools to focus our attention where it does the most good. The former alone has reduced our workload by 43% while contributing to a greatly reduced forecasting error. The latter has also contributed to reducing our forecast error, as we only review and correct the forecasts which our SPC tools identify as the most error-prone and where there is the greatest potential benefit from taking action.

These techniques are applicable to any service parts environment, not just commercial jet engine parts. They are also applicable to fast moving businesses, except that instead of forecasting less often, the need exists to forecast more often. So instead of using calendars ranging from months to years, one might need to forecast using months, weeks, days, or anything in between. The point is to find the calendar which works best for each product, rather than blindly following the paradigm of forecasting all parts on the same calendar. The same lesson applies to the use of SPC tools to identify and prioritize the Parts for scrutiny. The idea is to use these tools to enable the planner to treat Parts differently--each according to its need and the value generated by acting on that need.

References

1. Estep, J.A., "A Simple Technique for Reducing Forecast Errors," *APICS 30th Annual International Conference Proceedings*, APICS, 1987, pp. 291-293.
2. Greene (Ed.), *Production and Inventory Control Handbook* (2nd ed.), APICS, 1987, p. 29.15.
3. Brown, R.G., *Advanced Service Parts Inventory Control*, APICS, 1982, p. 305.

About the Author

Jeff Beck is a team leader in Forecasting/Inventory Planning for commercial spare parts at GE Aircraft Engines. He is responsible for approximately 8,000 part numbers for both inventory level and customer service. Beck has been working for the past five years as a forecasting/inventory planner, and has been team leader for the past two years. Beck was hired into GE on the Logistic Development training program. While on this program he had assignments in Sales Forecasting, Spares Configuration, and Inventory Planning. He has an MBA with a concentration in Information Systems Management from Xavier University in Cincinnati, and a BS in Marketing from The Ohio State University.

The above is a reprint of an article submitted by Jeff Beck, GE Aircraft Engines, which appeared in the APICS 38th International Conference (October 22-27, 1995, Orlando, Florida) Proceedings. Several charts and the accompanying text, which were deleted from the published article for production reasons, have been restored in this reprint.

For additional information contact:

E/Step Software Inc.
12015 Summitview Road
Yakima, WA 98908

Telephone: (509) 853-5000
Email: John.Estep@EstepSoftware.com
Web site: www.EstepSoftware.com