APICS

30th Annual International Conference Proceedings

October 19-23, 1987
St. Louis, Missouri

APICS
Conference
Proceedings

The Library
of American
Production

I15 A SIMPLE TECHNIQUE FOR REDUCING FORECAST ERRORS

John A. Estep, CFPIM
E/Step Software, Inc.

This paper presents a simple but effective method for forecasting lumpy items (i.e., items difficult to forecast accurately). The technique is based on the observation that aggregated data can frequently be forecasted with greater accuracy than individual data. Under this technique, aggregation is across time, rather than across items.

Most organizations attempt to forecast all items on the same calendar, usually on a monthly basis. For some items, monthly usage is so low (e.g., containing frequent zeros) that the monthly forecast is classified as lumpy. Sometimes, however, an item which is lumpy when forecasted monthly has much greater accuracy when forecasted quarterly. Although this gives the user less visibility for this item's forecast (four forecasts per year rather than twelve), a quarterly forecast may be more than adequate in some situations. Similarly, it may be necessary (because of forecast error) and sufficient (for planning purposes) to forecast some items only twice or even only once per year. As an added benefit, less time is spent reviewing forecasts that do not require it.

This paper discusses the circumstances in which one should consider using such a technique. Using data from a service parts manufacturer, it provides examples showing the effects of varying the forecast calendar on the resulting forecast error, as well as on the amount of safety stock required to maintain a particular level of customer service. In the sample studied, half of the items benefitted from less-frequent forecast revision, with an average forecast error improvement of 30% for those items.

1. Forecast Revision Frequency

Finished goods forecasting systems (the statistical part, that is) generally follow the approach of fitting a model to the demand history and projecting that model forward. Then, as time passes, the forecasts are compared to actual demands and the errors are used to revise the model and subsequent forecasts. To accomplish the modelling and revision tasks it is necessary to divide the demand history into time buckets. The resulting model projects future demand using the same time buckets.

In all cases, the frequency with which the forecasts are revised equals the size of the forecast bucket. That is, forecasts in monthly buckets are revised monthly, quarterly forecasts are revised quarterly, and so on.

The time buckets used by individual companies are dictated by a combination of market requirements and available systems. While some fast-moving consumer goods businesses track demand and revise forecasts weekly or even daily, most companies use monthly calendars. "Months" may be calendar months, 4-4-5 type months, etc. For discussion purposes let's assume that "normal" forecast revision is monthly--although it could be more frequent.

With a normal demand-gathering cycle which is monthly, it is usually difficult and often impossible to entertain the thought of revising forecasts more frequently than monthly. But it is certainly possible, and often advisable, to revise the forecasts less frequently. For example, you could accumulate the months of data into quarters, forecast them, and then revise the forecasts each quarter.

2. When to Consider Less Frequent Forecast Revisions

There are two primary reasons for considering less-frequent forecast revisions. The first occurs when sparse or wildly fluctuating demand prevents you from achieving an accurate forecast on a monthly basis. While it is possible to have demand at a high level with wild fluctuations, it is more likely to find such problems with parts which have a low level of demand (for example, when monthly demand is frequently or most often zero).

The service parts environment provides many examples of SKUs ("SKU" is a Stock Keeping Unit--a part at a location) for which achieving an accurate monthly forecast is impossible. If you are not in the service parts business, imagine having to supply an accurate forecast by hour for each of your items--the level of difficulty is about the same! To cite another example, demand for a part at one location may be sparse, even though the demand for the same part at other locations may be greater and more predictable. In this case you may want to forecast the part at the one location less often than at the others.

The second primary reason for considering less-frequent forecast revisions arises when there is simply not enough time to review the forecasts (or at least the exceptions) monthly. If it takes three months to review one month's worth of forecast exceptions, it is unrealistic to insist on monthly forecasts. Of course every finished goods inventory includes a variety of parts, and in most cases it is possible to separate them into groups, each with a appropriate revision frequency. By doing this, you can apply the time you save on the less-frequently review SKUs to the more-frequently reviewed SKUs--which presumably merit greater attention.

If you use ABC logic to categorize your inventory, you are likely to find good candidates for less-frequent forecast revision among the lower levels (e.g., the C's and D's. Keep in mind, however, that the sparseness of the unit demand--not the dollar demand--is the key (unless your units are dollars).

3. When Not To Consider Less Frequent Forecast Revisions

There are at least two reasons for not using a less-frequent calendar. The most obvious is when doing so yields a less accurate forecast. This is most likely to occur when the heuristics employed by the modelling logic (which select the models or weight the history) result in very different modelling decisions on one calendar versus another.

Secondly, business requirements may demand a more-frequent calendar. For example, a service part critical to the operation of a larger piece of apparatus may require monthly review, even though the part may have a slightly better forecast error on a quarterly calendar than it does on a montly calendar. There are many other instances where this can occur, driven by such things as competitive pressures, public safety, contractual requirements, etc.

4. Introduction to Examples and Methodology

The examples cited in this paper are from a Fortune 500 company which supplies service parts for heavy industrial machinery. Fifty SKUs were chosen (somewhat randomly) from a population of over 15,000. There was no attempt to select a statistically stratified sample; the intent was simply to pick a few SKUs from different inventory classes to see how they would fare on various forecast calendars. The 50 SKUs comprise 30 different parts; ten of the parts have three SKUs each (two satellite warehouses plus the total demand), while the other 20 parts have one SKU (total demand), each. The parts are generally from the B and C classes with a few A's and D's. Most of the SKUs had six or seven years of demand history and none had less than four.

The SKUs were loaded into the database of a finished goods forecasting and inventory management system. Then, using each of five calendars in succession, the system's automatic modelfitting logic was used to find the best model for each SKU. The resulting forecast error (standard deviation of forecast errors over the lead time) was saved for each of the calendars. These numbers can be directly compared since each SKU's lead time is the same regardless of the forecast revision frequency, and the errors are converted to cover this length of time. When all calendars were processed, the system identified the calendar with the lowest error for each SKU. Then the models for the best calendar in each case were reinstated and several summary reports were run.

5. The Calendars Evaluated

The parts were forecasted using five different calendars: monthly, bimonthly, quarterly, semiannual, and annual. Companies typically consider only monthly and quarterly forecast calendars, but, as this sample shows, the others can be very useful too.

The monthly calendar happens in this case to be a 5-4-4 calendar, but the forecasting logic takes this into consideration so the results have no impact on the comparisons in this paper. The bimonthly calendar has six revisions per year, or one every other month. It is useful in situations where monthly forecasts are too inaccurate, but quarterly forecasts too infrequent, to allow adequate review. The quarterly calendar has one revision every third month, that is, four revisions per year. SKUs on this calendar receive substantially less attention than those on a monthly calendar.

Dropping back even further, there are times when it is appropriate to review forecasts twice per year (semiannual calendar) or only once per year (annual). These are typically for very slow-moving items, but the review time involved is only a small fraction of the time required by the monthly calendar.

6. Choosing The Best Calendars

For half of the 50 SKUs, the best calendar (i.e., the one resulting in the lowest forecast error) was monthly. The next most frequently recommended calendar was annual. This is perhaps not too surprising, given the service parts nature of the business and the resulting sparseness of the demands.

The complete tabulation is shown in Table 1.

Best Calendar	Number of SKUs
Monthly	25
Bimonthly	10
Quarterly	1
Semiannual	1
Annual	13

Table 1

It is clear that for this group of parts at least, there is no middle ground!

In some cases, the differences in forecast errors between calendars are substantial; in others they are very small. Table 2 lists some of the SKUs for which this method favored retaining the monthly calendar.

Part Numb	Loc	Monthly Error	Bimonthly Error	Quarterly Error	Semiannual Error	Annual Error
01	A	3.93	4.72	4.76	5.19	4.10
01	B	4.13	7.82	8.16	9.83	11.35
01	T	9.00	14.01	14.72	16.35	20.54
13	T	13.64	150.24	148.10	141.96	169.39
15	T	132.73	163.97	169.89	1514.11	1518.00
26	T	1673.25	2218.86	2290.17	2447.76	3242.69

Table 2

In these examples, locations A and B are satellite warehouses, while location T is the total demand.

For part 01, while the monthly calendar gave the lowest error for all locations, at location A there is very little difference between the monthly and annual calendars. A decision on the appropriate calendar would have to take into consideration other factors besides forecast error--such as necessary visibility. Locations B and T exhibit a steady decrease in forecast accuracy as the revision frequency decreases. (For other parts, this method assigned different calendars to different locations for the same part.)

Part 13 is interesting simply because of the huge difference in errors between monthly and the other calendars. The closest calendar has an error over ten times larger than that for the monthly calendar. For parts 15 and 26 the difference is also apparent, although less striking. In both cases, the accuracy decreases steadily with the transition from monthly to annual calendars.

7. SKUs With Non-Monthly Recommended Calendars

Table 3 lists the comparison across calendars for the 25 SKUs for which this method recommended non-monthly calendars. For those SKUs for which the best calendar was not monthly, the average percentage reduction in the forecast error (when compared to using the monthly calendar) was 30%. This is indeed a significant reduction.

Part Numb	Loc	Monthly Error	Bimonthly Error	Quarterly Error	Semiannual Error	Annual Error	Best Frequency	% Error Reduction
02	A	5.32	2.86	3.88	6.34	7.94	6	46.23
03	A	96.30	86.42	139.91	166.96	197.90	6	10.26
04	B	1.72	1.74	1.46	1.39	1.27	1	26.55
05	A	8.06	7.56	8.18	7.92	7.84	6	6.10
05	B	21.20	16.85	29.29	32.74	28.15	6	20.51
06	A	2.74	2.42	2.59	3.10	2.20	1	19.60
06	B	0.61	0.49	0.48	0.47	0.42	1	32.26
06	T	5.88	4.97	5.33	6.37	5.08	6	15.48
07	B	6.24	6.81	6.64	7.21	5.35	1	14.30
07	T	13.57	15.72	13.53	12.20	7.65	1	43.62
08	B	47.59	1.28	1.14	0.67	0.62	1	98.70
08	T	106.48	26.45	39.66	36.92	58.52	6	75.16
09	A	1.06	1.29	1.47	1.67	0.70	1	33.54
09	B	1.57	0.89	1.70	1.54	2.07	6	43.35
09	T	1.94	1.72	2.14	2.51	2.52	6	11.55
10	A	214.71	192.47	238.63	226.74	166.41	1	22.50
12	T	64.36	64.23	69.37	72.26	74.80	6	0.19
14	T	17.30	16.38	18.22	18.29	21.23	6	5.30
18	T	36.05	34.04	27.33	103.78	121.56	4	24.19
22	T	140.61	124.50	120.08	120.77	79.52	1	43.45
24	T	22.87	23.44	30.55	22.33	26.06	2	2.36
27	T	19.66	20.21	18.15	23.24	17.40	1	11.51
28	T	103.49	174.75	154.86	155.40	50.24	1	51.45
29	T	257.30	243.72	219.50	170.19	135.57	1	47.31
30	T	95.00	131.17	125.51	114.86	47.71	1	49.78
Average Error Reduction (%) For 25 SKUs								30.21

Table 3

A closer analysis of the results in Table 3 provides additional insights. For example, location B for part 08 scored a 98% reduction in forecast error, yet there is almost no difference between the semiannual and annual errors. Similarly, for parts 12 and 14 the bimonthly calendar is best, but there is very little difference between the monthly and bimonthly calendars. And for part 6 at locations A and B there is, with only minor exceptions, very little difference in the errors for any of the calendars. One concludes that the final calendar determination must take other factors into consideration (such as competitive pressures, service requirements, etc.). Yet for several SKUs (for example, parts 28, 29, and 30) there is no question that the chosen calendars are substantially superior to the others.

8. Effect On Inventory Value

Safety Stock

Although the average reduction in forecast errors was 30% for the non-monthly SKUs, the effect on safety stock may be different because of the differing costs and service levels. For these 25 SKUs the effect was a 45% reduction in the total dollar value of the safety stock (in standard cost dollars) required to maintain the same level of customer service as compared to the safety stock for a monthly forecast. This is certainly a significant amount. Of course, depending upon the nature of your own business, you may see a safety stock value savings which is either greater or less than your average forecast error reduction.

Working Stock

A decision to reduce the forecast revision frequency does not directly cause an increase in working stock. Working stock is controlled by the replenishment frequency. If you are setting replenishment frequency using the EOQ (Economic Order Quantity) method, you are looking at the total annual usage--which in most cases is not affected (or only slightly affected) by the change in forecast frequency.

You can run into a problem, however, when, for example, you have a part which is replenished monthly but forecasted annually. This is not too likely since the same low demand which drives it toward an annual forecast calendar, also pushes the EOQ toward less-frequent replenishments. But if it does occur, you could have a part with only a month or two's stock on hand for which the only forecast is an annual value. It is just possible (but perhaps not probable) that you could receive demands for the entire year's amount, leaving you out-of-stock and severely backordered (or with lost sales).

In the service parts business, even in this unlikely scenario, all is not lost. By inquiring into the source and purpose of the demand, you may discover that it is an order to replenish the customer's stock (rather than needed now to fix some machinery). In this case you want to ship just the stock required now and deliver the rest later, leaving some stock still on the shelf to cover your other customers.

But the best solution is to compare your forecast and replenishment frequencies. When there is a wide disparity, either revert to a more frequent forecast, or increase your working stock. You can calculate which of these is going to cost more for inventory, given some target customer service level. For most parts, however, you should find that those on the annual or semiannual calendars are replenished from between once every six months to once every two years.

9. Effect on Work Load

Reviewing forecasts or forecast exceptions less frequently obviously takes less time than reviewing them every month. The time savings range from 50% on bimonthly items to over 90% for annual items. The time saved on the less-frequent SKUs can be applied to the more-frequent ones, which usually require more attention and which tend to pay greater dividends for such increased attention.

In some environments, however, you may find a large number of SKUs on the less-frequent calendars. This happens quite often in the service parts environment. It could be that your review of the SKUs on annual calendars takes 3 months. This is unacceptable since some parts actually go 15 months without review, plus you have other (non-annual) parts to review as well.

A forecast analyst or product planner's time is a finite resource just like a lathe or milling machine. You must plan this resource. In this example, you could plan to review one-twelfth of the SKUs on annual calendars every month. This spreads the work evenly throughout the year and avoids periodic crunches. You can follow the same approach for the other calendars: review one-sixth of the semiannual SKUs every month, etc.

10. Master Production Scheduling Considerations

When transferring the non-monthly forecast to your scheduling system, you may need to re-examine its assumptions. For example, you could have five parts with forecasts for the period ending 12/31/87--but each for a different calendar: one for the month of December, one for November and December, one for the 4th quarter, one for the last half of 1987, and one for all of 1987! You may find your MPS system assuming that each is the forecast just for December. In general you can solve this problem by using period-beginning dates, rather than period-ending dates (or use the period-ending date for the prior period).

You should also examine how you consume the forecast. An annual forecast must be maintained for an entire year with actual demand posted against it every month. You do not want the system dropping the annual forecast after the end of January.

11. Other System Considerations

The technique of less-frequent forecast revisions (where appropriate) is easy to implement with manual systems. You need only to adjust your procedures, and possibly forms, to put the changes into effect.

For automated systems there are a few things to keep in mind. First is the need to accumulate the "normal" monthly demand over two to 12 months. Depending on your systems, this may be tackled in your order entry (demand gathering) system, your forecasting system, or in the interface between the two.

If your forecasting system does not allow SKUs to use different calendars within a single database, you will have to put the SKUs on each calendar into a separate database. This handles the problem but makes it harder to get a summarized forecast report across all SKUs. If your system does not allow for summarizing across databases, you will need to prepare separate summaries for each calendar and consolidate them manually or by using a spreadsheet program.

12. Summary

The use of less-frequent finished goods forecast revision calendars, where appropriate, is a good way to achieve simultaneously:

1. Improved Forecast Accuracy
2. Reduced Work Load

The improved accuracy translates directly into lower inventory to achieve the same level of customer service. The technique is particularly applicable to the service parts environment. In the sample studied, half of the items benefitted from less-frequent forecast revision, with an average forecast error improvement of 30% for those items. The reduction in safety stock value associated with these SKUs was 45%. However, as discussed, not all SKUs can benefit from this method. Service requirements, competitive pressures, public safety, and other concerns may dictate retention of more frequent forecast revisions, and an initial analysis of results must be done before any change is mandated across groups of parts.

John A. Estep is president of E/Step Software Inc., of Hopkinton, Massachusetts, a firm specializing in consulting and the development of software for manufacturers and distributors. With over 17 years' experience in consulting and systems development, Mr. Estep's expertise is concentrated in the areas of database design, demand forecasting, inventory management, distribution, master scheduling, and quality control. In addition, Mr. Estep has taught courses for APICS certification examinations. Before launching his own business in 1983, Mr. Estep spent 6 years with STSC Inc., as Product Manager for Materials Management Systems and as Materials Management Consultant. Prior to joining STSC, Mr. Estep worked in consulting, systems development, and operations research for several firms including Talon Inc., and Bolt, Beranek & Newman.

An APICS Certified Fellow, Mr. Estep holds an M.S. degree in Electrical and Computer Engineering from the University of Massachusetts, and a B.S. in Mathematics from Allegheny College. He is a member of APICS and the Boston Computer Society.