BETTER FORECASTING WITH SAS/ETS® SOFTWARE

Bruce R Benseman, DSIR Applied Mathematics

This paper describes our experiences using the SAS/ETS software when forecasting for a large New Zealand primary producer. This company has an annual turnover of $1300 million. And it faces large seasonal fluctuations in the availability and quality of its raw material. The company saves thousands of dollars per day when we forecast these fluctuations better. We will comment on the advantages and problems of different forecasting models and SAS/ETS software.

Problem Background

This paper describes our experiences forecasting for a large New Zealand dairy company. In the 1988/89 season, the company had a record turnover of $1300 million, despite a 7% fall in supplied milkfat. This turnover included $900 million from exported dairy manufacturing. In this season, the company manufactured about 320,000 tonnes of milk powder, cheese, casein, and butter. And each year, it handles over 33% of New Zealand's wholemilk production.

The company's raw material, wholemilk, has a strong seasonal pattern. This means their manufacturing plant cannot be used efficiently all year. Figure 1 shows how the wholemilk supply starts from a winter trough in June and July. It climbs sharply to a peak (called a flush) in late October-early November, then drops gradually over the summer and autumn. The seasonal shape of the wholemilk curve is caused by the cows' lactation cycle. Farmers plan calving to coincide with the spring grass growth.

However, each year, weather can cause significant fluctuations about the wholemilk pattern. These variations about the average pattern can last through to May. A wet summer makes the cows produce more milk. And wholemilk between March and May can vary up to 50% from one year to the next. This represents over $1 million extra for each day in April.

A good forecasting system aims to give company executives at least one month's warning of such a late flush—so they can use it wisely. The company recognises short-term wholemilk forecasting is an important problem. A bad forecast can kill the best of plans.

Seasonal Variation

Like wholemilk volume, wholemilk composition has a seasonal pattern. Wholemilk is about 87% water. The useful components are fat, protein, and lactose. The rest is a small quantity of minerals, ash, and salts. Wholemilk composition directly affects how many tonnes of dairy products can be made from 100 kilolitres of wholemilk. Milk powder yields depend on the fat and solids-non-fat levels in wholemilk. Cheese and casein products depend on the percentage of protein and fat in the wholemilk.

Wholemilk composition depends on the time in the cow's lactation cycle. In general fat levels in wholemilk grow steadily from 4.5% in September to 5.6% in April (see Figure 2). Protein levels dip to a trough of 3.3% in the November flush, but rise to 3.9% by April. These patterns also depend on the geographical region the wholemilk was collected in. For example, central Waikato collection regions like Te Rapa, Waharoa, and Waitoa have higher fat and protein levels than southern regions like Reporoa. Like wholemilk supply, we have to forecast fat, protein, and lactose eight times. Once for the whole company—the rest for the company's seven collection regions.

Unfortunately, wholemilk composition also changes from one season to the next. It depends on climate, farm management, and animal condition. For example, protein levels may vary by 7% from one April to the next.

The company needs to forecast the fat, protein, and lactose levels so it can plan its production and financial operations better. For every 1% increase in fat or protein levels above average, the company can produce hundreds of extra tonnes. So marketing staff need to know when they must sell extra product, and the accountants need to know when their cashflows will be higher.

Forecasting Composition

The company tests the fat, protein, and lactose levels on a daily basis. However, for planning purposes, we use average monthly compositions.

We generally use the simple exponential smoothing technique for forecasting—because it is fast. However, it cannot handle trends, nor the strong seasonal patterns shown in Figures 1 and 2. Double exponential smoothing can anticipate trends, but it looses short-term responsiveness.

So before doing any exponential smoothing, we always isolate the trend and seasonal components from the data series. The trend in our monthly wholemilk composition series is negligible—for example milkfat levels do not change from year to year.

We form a new series of deseasonalised "differences" by subtracting the monthly average (mean) from each actual observation. We calculate the mean using the last 3 years' data.

We give these differences to the FORECAST procedure [2]. This uses the exponential smoothing method (EXPO) to calculate a new series of forecast differences. We tried various smoothing constants and found a weight of a 0.5 gave the smallest forecasting errors. This is close to the auto-correlations we found in the monthly fat, protein, and lactose difference series.

The FORECAST procedure forecasts 12 months ahead. When we have the forecast differences, we add them to the averages to get the forecast actuals. Unfortunately, over the winter months, May to August, the

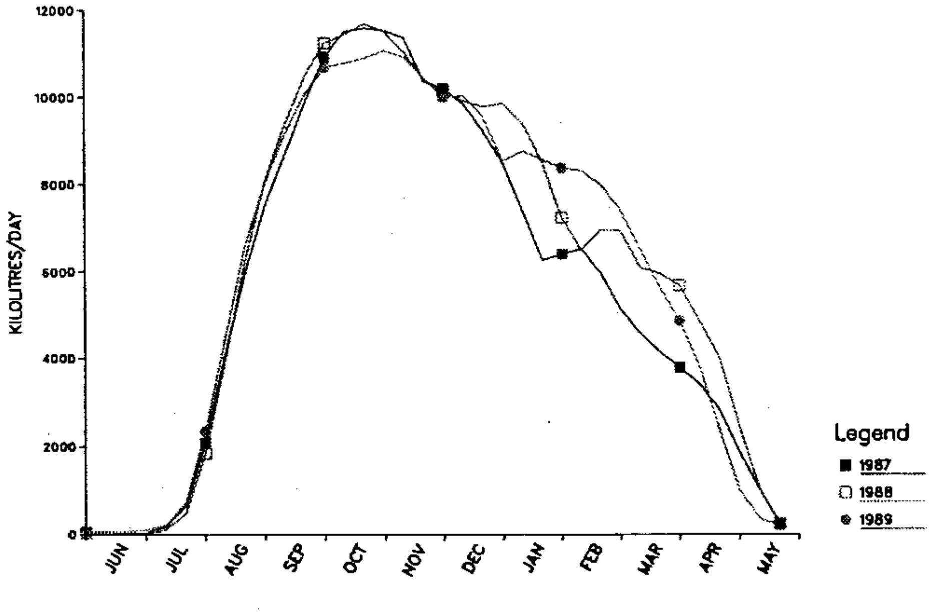

Figure 1: Total Company Wholemilk Volumes

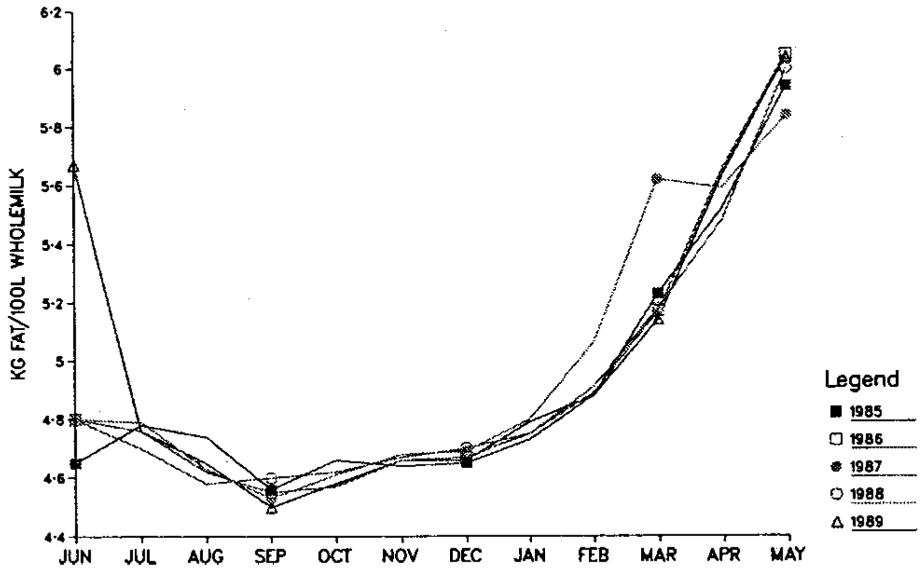

Figure 2: Average Company Milkfat Levels

wholemilk composition data series are very noisy. Exponential smoothing over these months gives unreliable results. So we use just the means for these months as our forecasts.

Finally, we store the forecast wholemilk composition —in table form for the company's production planning system [1]. This system finds the best manufacturing programme for their powder, casein, cheese, and butter products. Its reports and maps show the company how to allocate the wholemilk from their collection regions to their various processes, and how to handle the resulting cream and whey. These decisions are based on information on freight costs, manufacturing capacities and costs, product yields, prices, and market constraints.

Forecasting Volume

Forecasting wholemilk volume is similar to composition, but involves some extra calculations. Firstly, we need to convert the data from period totals to period averages. The dairy company plans in ten day periods—three periods per month. So the last period has between 8 and 11 days. Wholemilk data is stored as a period total. For planning purposes, we convert it to a daily average.

As before, we deseasonalise the series by subtracting the period average from each actual observation. For January to May periods, we use a three year mean—to overcome the climatic variation between years. However for June to December periods, we use just the previous last year's level—to model the ever sharpening peak.

Next we get the FORECAST procedure to forecast 36 periods ahead. We found a smoothing constant of 0.8 gave the best forecasting results. (The REG procedure showed us the differences have a 0.8 auto-correlation). After we have forecast the differences, we add them to the average wholemilk volumes. Except however for the winter months, June and July, where we use just the averages—because the series is predominantly noise. We also find the wholemilk flush in November is best forecast with an average—to avoid overshooting the peak.

We actually weight the forecast differences before adding them to the averages. These weights decrease with time—because we do not expect today's volume difference of say 10% to last more than a month or two. So we bring the forecasts back to the average, within a specified time.

And for pragmatic rather than statistical reasons, we also force each summer forecast to lie between 70% and 105% of the previous year's actual. This limits any overreaction the EXPO forecasts makes to swings in actual wholemilk.

At this stage, we let the users adjust the volume forecasts. Often, they want to increase the forecasts for the November flush—to ensure they have enough processing capacity. And they prefer pessimistic forecasts towards the end of the season—to avoid optimistic financial predictions.

Finally we blend the 8 forecasts to improve their robustness. A consensus of forecasts from different approaches (or people) helps avoid extreme predictions. Remember we are forecasting the company total and 7 collection regions. We find the difference between the company forecast and the sum of the 7 regional forecasts. Then we move half this difference from the company forecast, and share it among the 7 regions.

Other Forecasting Methods

It took some time finding the best way to deseasonalise the data. We tried 28 different models. We used relative differences (dividing observations by averages); but this series was too erratic. When we smoothed it with running means, we found it forecast slower and worse than the naive 3 year average forecast—mainly because the large winter differences upset the exponential smoothing forecasts for several months.

When we ignored or dampened the winter noise, our forecasts improved significantly. We also changed from relative to absolute differences (that is, additive rather than multiplicative models). Absolute differences have the advantage of being measured in the same units as the original series—for example, kilograms of fat. Moreover we could stop worrying about divides by zero. We also put upper and lower bounds on the composition data to filter its outliers.

When we found that the basic forecasts did not respond to recent shocks fast enough, we increased the exponential smoothing constant. Currently we use 0.5 as the smoothing constant for all 24 composition monthly series we forecast—and 0.8 for the 8 volume 10-day series. That is, we can use a more responsive weight for a series with more frequent observations.

To compare the different forecasting models, we calculate the mean absolute forecasting error over the last three seasons, but excluding the winter months. We used the MEANS procedure to find the average, standard deviation, minimum and maximum errors for each model.

We compared all forecasting models to the 'naive' forecast—one based on a three year average. Figure 3 shows how EXPO forecasts wholemilk volumes much better than one based on means. We found that the best exponential smoothing models reduced the naive forecast mean absolute errors as follows:

- by 23% for the fat level in wholemilk
- by 42% for the protein level in wholemilk
- by 49% for the wholemilk volume series

STEPAR and ARIMA

We tried the STEPAR method of the FORECAST procedure. STEPAR fits a time trend plus a autoregressive model to the data. That is, the difference at any time is a function of both time and of earlier differences. The trend term captures long term growth. The autoregressive components captures the short term weather fluctuations. STEPAR usually found only two lag-terms significant, and performed no better than the EXPO method. We decided not to use the STEPAR method because it is twice as slow, and we have 24 series to forecast.

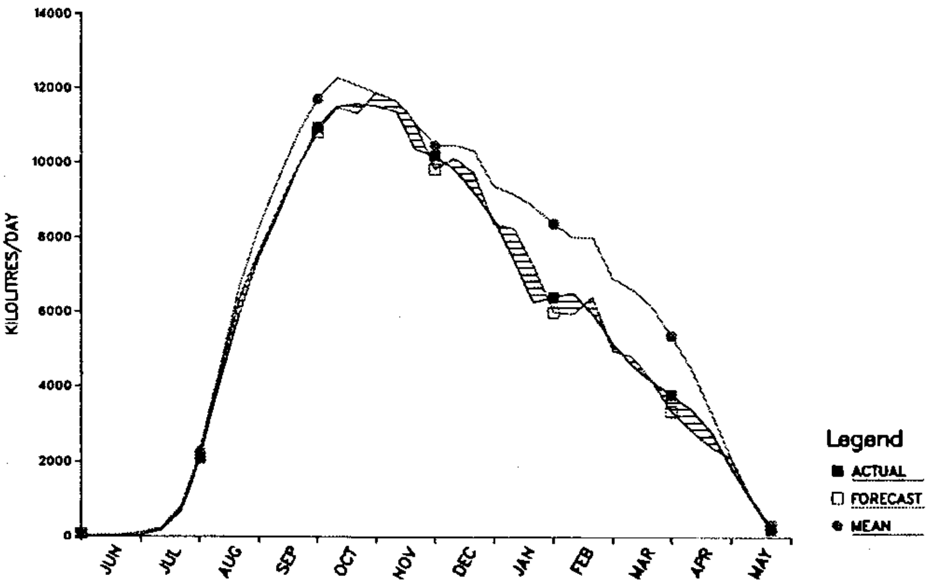

Figure 3: Forecast and Actual Company Wholemilk

We also tried the ARIMA procedure in the SAS/ETS module. ARIMA has the advantage of choosing the best regression model for our exponential smoothing approach. We don't have to run dozens of models looking for the best exponential smoothing weight. We had mixed results. For some series, ARIMA outperformed the EXPO forecasting results. However for series with strong winter noise, ARIMA performed worse. We have to reduce or eliminate this noise, if we want ARIMA to work properly. In other words, ARIMA cannot handle unusually noisy parts of the data series. Finally, there is the CPU cost. The ARIMA procedure is several times slower than the FORECAST one.

Conclusion

This paper has examined several forecasting models and SAS procedures. The user is generally only interested in the results. The models we developed give the company a month's warning about unusual wholemilk volume and composition. If the milk volumes or yields are higher than normal, the company can plan to produce more profitable products. For example, if company planners can win a wholemilk powder contract (instead of making the extra wholemilk into skimpowder stocks) they earn more revenue—about $15,000 per day. Conversely, if summer-autumn wholemilk volumes falls below average, the company can close its factories earlier. This too saves thousands of dollars per day in energy and labour costs.

We had previously programmed the company forecasts in FORTRAN. After we converted the forecasting models to the SAS system, we reached the following conclusions:

- the FORTRAN compiler is a fraction of the cost of the SAS system rental
- the FORTRAN task runs nearly 10 times faster than the SAS software task
- SAS programs are more compact—for example, SAS software gives us reporting and plotting procedures
- the SAS system lets scientists develop mathematical models faster
- the SAS system makes it easier for users to modify the forecasts and to tailor the computer output.

Overall, SAS software makes the whole forecasting system much friendlier to use. We like working with the SAS products.

References

1. Benseman, B.R. "Production Planning in the New Zealand Dairy Industry", *Journal of the Operational Research Society*, Vol. 37, No. 8, 1986, pp. 747-754.
2. *SAS/ETS User's Guide: Econometric and Time Series Library* Version 5 Edition. SAS Institute Inc., Gary, NC, USA.

Contact: If you have any questions or comments on this paper, you can write to me at:

DSIR Applied Mathematics
PO Box 1335
Wellington
New Zealand
or telephone NZ (04) 727 855