Paper 326-2011

Multifrequency Forecasting with SAS® High-Performance Forecasting Software

Michele A. Trovero, Ed Blair, Michael J. Leonard
SAS Institute Inc., Cary, NC

ABSTRACT

Forecasters often deal with data accumulated at different time intervals (for example, monthly data and daily data). A common practice is to generate the forecasts at the two time intervals independently so as to choose the best model for each series. That practice can result in forecasts that do not agree.

This paper shows how the SAS® High-Performance Forecasting HPFTEMPRECON procedure uses the lower-frequency forecast as a benchmark to adjust the higher-frequency forecast to take the best advantage of both forecasts.

INTRODUCTION

Forecasters often need to produce forecasts for a certain time series at more than one frequency. For example, a company that provides warranty repairs for appliances might want to forecast the number of daily calls for staffing and operational planning, such as ordering supplies. The company might also want to forecast service calls at a monthly frequency to plan long-term expansion and to plan for financial concerns such as the purchase of more vehicles or the hiring of new staff.

This paper deals with the problem of forecasting one time series at different frequencies, with a focus on stock variables. For a stock variable, the low-frequency series is the temporal aggregation of the high-frequency series. The term *accumulation* indicates temporal aggregation, and thus distinguishes it from other forms of aggregation, such as the aggregation of series within a subclass that can take place in a hierarchical forecasting context.

The problem of forecasting at multiple frequencies is easily solved in an ideal world where data are plentiful, series are well behaved (meaning they have mostly nonzero values and and are easily transformed to a covariance stationary series), and the correct model is chosen for each series. In this Lake Wobegon of Statistics[1] the accumulation of the high-frequency forecasts is at least as efficient as the forecasts generated by modeling the low-frequency series, in the sense that the mean squared error of the h-step-ahead prediction of the former is less than or equal to the mean squared error of the h-step-ahead prediction of the latter.

A formal outline of this argument for seasonal ARIMA processes can be found in Wei (1990), Chapter 16. The idea is simple: a forecast (prediction) is the linear projection onto the Hilbert space generated by the observed series. The space spanned by the low-frequency data is a subset of the space spanned by the high-frequency data. Therefore, the accumulation of the projection on the finer space generated by the high-frequency data is at least as "close" to the actual future value as the projection on the coarser space spanned by the low-frequency data. Another way to express the same concept that is simpler and does not require any mathematical jargon is that the accumulation process is a form of compression that involves loss of information. The original high-frequency data cannot be regenerated using only the accumulated data. Therefore, forecasts generated with the restricted information contained in the accumulated data cannot be better than forecasts generated with full information of the non-accumulated data.

Reality, however, rarely comes in textbook format (or radio-show format). Consider the following real-life examples (the name of the companies are retained for confidentiality reasons):

Example 1. The spare-parts branch of a large company operates nationwide and manages more than 40,000 spare parts. Three-months-ahead daily forecasts are needed for each ZIP code for replenishing the repair trucks and for making staffing decisions. Very few parts are needed with regularity. Approximately only 10% of the parts show a somewhat regular demand for each ZIP code. For the remaining parts, the daily demand is almost always zero. Long-term monthly forecasts are needed for part production, hiring purposes, and long-term investments.

Example 2. A large retail store chain collects POS (point-of-sale) data in each store. Hourly forecasts are needed in the medium term for staffing purposes. The hourly data are kept for three months, after which they are discarded due to the cost of storing such a large amount of data. Only data accumulated at daily intervals are kept. Long-term monthly forecasts are needed for expansion and financial planning.

In both examples, forecasts are needed at different frequencies for different purposes. However, there are good reasons to believe that the accumulation of the high-frequency forecasts will not lead to good forecasts for the low-frequency data.

[1] Lake Wobegon is a fictional town in the U.S. state of Minnesota, said to have been the boyhood home of radio-show host Garrison Keillor. It is characterized by the fact that all the women are strong, all the men are good looking, and all the children are above average.

In the first example, most series show intermittent behavior. *Intermittent series* consist mostly of a single value, usually zero. Models for intermittent data, such as the popular Croston (1972) model, cannot capture important features such as trend, seasonality, and dependency on events or other external variables. Additionally, multiple seasonal components might be present in the high-frequency data, whether they are intermittent or not. Modeling and estimating multiple seasonal components simultaneously can be complex and computationally intensive.

In the second example, the duration of the hourly (high-frequency) data is not sufficient to produce monthly (low-frequency) forecasts of any value. Indeed, you can reasonably argue that the information contained in the longer history of the daily data can be used with benefit to forecast the hourly data. For example, when making staffing decisions about the very important winter holiday season, the retailer should use the information contained in the daily data, which covers the previous holiday seasons, and not rely solely on the hourly data forecasts which are based only on the previous three months.

In practice, the forecasts for the two or more frequencies of interest are often derived independently from each other by selecting at each frequency a model that provides the best results according to criteria, such as minimizing the MAPE (mean absolute percentage error). However, when the forecasts are derived independently, the accumulation of the high-frequency forecasts is generally different from the forecasts generated by the model for the low-frequency data. Additionally, as in Example 2, you might want to use the low-frequency forecasts to improve the high-frequency forecasts.

This paper shows a method for revising the low-frequency forecasts such that their accumulation at the low frequency is equal to the forecasts generated by the model selected for the low-frequency data. The first section details the method. The second section introduces the HPFTEMPRECON procedure in SAS® High Performance Forecasting and shows how it can reconcile monthly forecasts to daily forecasts for the Box and Jenkins' airline data. The third section presents the results of applying the method to a data set that consists of several time series that exhibit intermittent behavior. Finally, the last section draws the conclusions.

METHOD

The combination of a series of high-frequency data with a series of more reliable but less frequent data is seen often in business statistics. For example, surveys are conducted at quarterly intervals on a subsample of the population of interest to determine the interannual variations, while comprehensive surveys on the whole population are conducted only on a yearly basis. The process of adjusting the more frequent data to match the less frequent but more reliable data is known in the literature as *benchmarking*. Denton (1971) provided the first general framework for benchmarking based on minimizing a quadratic function. A recent and comprehensive review on the topic can be found in Dagum and Cholette (2006).

The lower-frequency forecasts are also referred to as the *benchmark forecasts*. The higher-frequency forecasts are also referred to as the *indicator forecasts*. Benchmarking procedures can be applied more generally to any two series that are measured at different time intervals. Therefore, this paper more generally refers to the *benchmark series* and *indicator series* to indicate the forecasts involved in the benchmarking.

Denote the indicator series by x_t with $t = 1, \ldots < T$, where t is associated with a date. Denote the benchmark series by a_m, $m = 1 \ldots, M$. The benchmarks have a starting date $t_{1,m}$ and ending date $t_{2,m}$ such that $1 \le t_{1,m} \le t_{2,m} \le T$. You want to find an optimal benchmarked series θ_t, $t = 1, \ldots, T$ such that the accumulation of benchmarked series at the frequency of the lower-frequency forecasts is equal to the benchmark series. That is,

$$\sum_{t=t_{1,m}}^{t_{2,m}} \theta_t = a_m, \qquad m = 1, \ldots, M$$

The bias is defined as the expected discrepancy between the benchmark and the indicator series. You can decide whether to adjust the original indicator series to account for the bias. Denote the bias-adjusted indicator series by s_t. When no adjustment for bias is performed, $s_t = x_t$.

The *additive bias* correction is given by:

$$b = \frac{\sum_{m=1}^{M} a_m - \sum_{m=1}^{M} \sum_{t=t_{1,m}}^{t_{2,m}} x_t}{\sum_{m=1}^{M} \sum_{t=t_{1,m}}^{t_{2,m}} 1}$$

In this case, the bias-adjusted indicator is $s_t = b + x_t$.

The *multiplicative bias* correction is given by:

$$b = \frac{\sum_{m=1}^{M} a_m}{\sum_{m=1}^{M} \sum_{t=t_{1,m}}^{t_{2,m}} x_t}$$

In this case, the bias adjusted-series is $s_t = b \cdot x_t$. Note that the multiplicative bias is not defined when the denominator is zero.

Let $\boldsymbol{s} := [s_1, s_2, \ldots, s_T]'$ be the vector of the bias-corrected indicator series, and let $\boldsymbol{\theta} := [\theta_1, \theta_2, \ldots, \theta_T]'$ be the vector of its reconciled values. Let D be the $T \times T$ diagonal matrix whose main-diagonal elements are $d_{t,t} = |s_t|^{\lambda}$, $t = 1, \ldots, T$. Indicate by V the tridiagonal symmetric matrix whose main-diagonal elements are $v_{1,1} = v_{T,T} = 1$ and $v_{t,t} = 1 + \rho^2$, $t = 2, \ldots, T-1$, and whose sub- and super-diagonal elements are $v_{t,t+1} = v_{t+1,t} = -\rho$, $t = 1, \ldots, T-1$. Define $Q := D^{+} V D^{+}$ and $\boldsymbol{c} := -Q\boldsymbol{s}$, where D^{+} indicates the Moore-Penrose pseudo-inverse of D.

The benchmarked (reconciled) series is given by the values θ_t, $t = 1, \ldots, T$ that minimize the quadratic function

$$f(\boldsymbol{\theta}; \lambda, \rho) = \frac{1}{2}\boldsymbol{\theta}' Q \boldsymbol{\theta} + \boldsymbol{c}'\boldsymbol{\theta} \tag{1}$$

under the constraints

$$\sum_{t=t_{1,m}}^{t_{2,m}} \theta_t = a_m, \qquad m = 1, \ldots, M \tag{2}$$

where $0 \le \rho \le 1$ and $\lambda \in \Re$ are parameters that you select.

When $\boldsymbol{s}$ does not contain zeros, the target function described by equation (1) is equivalent to the one proposed by Quenneville et al. (2006).

Two issues are considered when benchmarking. The first one is to preserve the movement in the high-frequency series as much as possible (movement preservation). The second is to account for the timeliness of the benchmarks, in the sense that the benchmark for the last period might not be available if the indicator series extends beyond the last benchmark value. Bias correction is a way to improve the timeliness of the benchmark in that it attempts to reduce the expected discrepancies between the benchmark and the indicator function. The parameter ρ is a smoothing parameter that controls the movement preservation. The closer ρ is to one, the more the original series movement is preserved. The parameter λ usually takes values 0, 0.5, or 1. For $\lambda = 0$, you have an additive benchmarking model. For $\lambda = 0.5$ and $\rho = 0$, you have a prorating benchmarking model.

In the traditional application of benchmarking, the goal is to regain the additivity of some seasonal adjusted series with respect to the benchmark. In the context of this paper, the goal is to find the optimal forecasts for the high-frequency series that respect the *accumulation constraint* (2). Therefore, it is suggested that you select the bias correction and values of the parameters ρ and λ in such a way as to optimize the selection criteria that was originally used to select the models for the high-frequency data. For example, if the model for the high-frequency data was selected by minimizing MAPE, likewise the parameters ρ, λ, and the bias correction should be chosen to minimize MAPE for the benchmarked forecasts.

When $0 \le \rho < 1$, the constrained minimization problem described by equations (1) and (2) can be derived from the constrained regression problem

$$\begin{aligned} s_t &= \theta_t + c_t e_t \qquad t = 1, 2, \ldots, T \\ e_t &= \rho e_{t-1} + \varepsilon_t \qquad t = 1, 2, \ldots, T \\ \sum_{t=t_{1,m}}^{t_{2,m}} \theta_t &= a_m, \qquad\qquad m = 1, \ldots, M \end{aligned}$$

where ε_t is a white-noise process with variance σ_e^2, and c_t are weights proportional to $|s_t|^{\lambda}$. Therefore, when $\lambda = 0$, the minimization problem is equivalent to a constrained regression problem where the error between the bias-adjusted indicator and the benchmarked series follows an AR(1) process with an autoregressive parameter proportional to ρ.

Let $\mathbf{a} = (a_1, a_2, \ldots, a_M)$. The constraint equation (2) can be rewritten as

$$J\boldsymbol{\theta} = \mathbf{a} \tag{3}$$

where J is a matrix of zeros and ones such that $J\boldsymbol{\theta}$ is the accumulation of the benchmarked series at the frequency of the benchmark.

The solution of the minimization problem then becomes

$$\hat{\theta} = s + C\Sigma_e CJ'(JC\Sigma_e CJ')^{-1}(\mathbf{a} - Js) \tag{4}$$

where C is a diagonal matrix whose main-diagonal elements are c_t, and Σ_e is the covariance matrix of e_t.

When benchmarking can be interpreted as a regression problem, it is also possible to derive the covariance of the reconciled forecasts. See Quenneville et al. (2006) for the details.

A further interpretation of this method is as a way to combine the forecasts at the two frequencies to produce forecasts for the higher frequency. The weights for the combination are derived using the solution of the minimization problem. The lower-frequency forecasts are assigned unit weights since they provide the right-hand side of the constraint equations.

THE HPFTEMPRECON PROCEDURE

Using the method outlined in the preceding section, the HPFTEMPRECON procedure reconciles high-frequency forecasts to low-frequency forecasts in such a way that the accumulation of the reconciled high-frequency forecasts is equal to the low-frequency forecasts. PROC HPFTEMPRECON reconciles forecasts for the same item at two different time frequencies whose intervals are nested in one another. In other words, it reconciles a two-level hierarchy of forecasts in the time dimension. For example, it reconciles monthly forecasts for the Box and Jenkins airline passenger data (in the Sashelp.Air data set) to the quarterly forecasts for the same series. For this reason, the HPFTEMPRECON procedure not only requires two input data sets for the predictions, but also it requires that the two frequencies of the forecasts be specified in two separate statements: the ID statement for the high-frequency data, and the BENCHID statement for the low-frequency data.

SAS High Performance Forecasting procedures are used to generate the forecasts at monthly and quarterly frequencies. These forecasts become the inputs to PROC HPFTEMPRECON . A full discussion about the SAS High Performance Forecasting system is outside the scope of this paper. Details can be found in *SAS High-Performance Forecasting: User's Guide*.

First, the HPFESMSPEC procedure generates an exponential smoothing model specification which is then selected by the HPFSELECT procedure:

```
proc hpfesmspec
      rep=work.repo
      specname=myesm;
   esm;
run;

proc hpfselect
      rep=work.repo
      name=myselect;
   spec myesm;
run;
```

Then, forecasts are generated with PROC HPFENGINE at the monthly and the quarterly frequencies using the selected model specification:

```
proc hpfengine
      data=Sashelp.Air
      rep=work.repo
      globalselection=myselect
      out=OutMon
      outfor=OutForMon
      outmodelinfo=OutMod;
   id date interval=month;
   forecast air;
run;

proc hpfengine
      data=Sashelp.Air
      rep=work.repo
      globalselection=myselect
      out=OutQtr
      outfor=OutForQtr
      outmodelinfo=OutModQrt;
   id date interval=qtr accumulate=total;
   forecast air;
run;
```

Note that the variable air appears in the FORECAST statement of both PROC HPFENGINE instances. The INTERVAL= option in the ID statements are different. In the first instance, the time ID interval is month; in the second instance, it is quarter. The monthly forecasts are stored in the PREDICT variable of the OutForMon data set, and the quarterly forecasts are stored in the PREDICT variable of the OutForQtr data set.

Finally, the monthly forecasts are reconciled to the quarterly forecasts using PROC HPFTEMPRECON:

```
proc hpftemprecon
      data=OutForMon
      benchdata=OutForQtr
      outfor=BenFor
      outstat=BenStat
      exp=0.5
      smooth=0.5;
   id date interval=month;
   benchid date interval=qtr;
run;
```

First, notice that the data set of the monthly forecasts is the argument of the DATA= option in the HPFTEMPRECON statement, and the quarterly forecasts data set is the argument of the BENCHDATA= option.

Second, notice that there are two statements to specify the frequency of the data, one for each input data set that contains the predictions. The ID statement is associated with the DATA= data set and specifies the variable that contains the time index of the indicator predictions and its relative frequency (interval). The BENCHID statement is associated with the BENCHDATA= data set and specifies the variable that contains the time index of the benchmark predictions and its relative frequency. Remember that the interval of the ID variable needs to be fully nested in the interval of the BENCHID variable. For example, months are fully nested in quarters. On the contrary, weeks are not fully nested in months, since a week can span two months. Therefore, the frequency of the indicator series cannot be weekly when the benchmark series has a monthly frequency.

The λ and ρ parameters of equation (1) are set by the EXP= and SMOOTH= options, respectively, in the HPFTEMPRECON statement.

Figure 1 displays the first 20 rows of the output data set BenFor, which contains the reconciled forecasts.

Figure 1 Reconciled Forecasts, EXP=0.5, SMOOTH=0.5

BenFor

Obs	_NAME_	DATE	ACTUAL	PREDICT	LOWER	UPPER	ERROR	STD	_RECFLAGS_
1	AIR	JAN1949	112	109.938	88.984	130.892	2.0622	10.6910	00000000
2	AIR	FEB1949	118	118.226	97.272	139.180	-0.2259	10.6910	00000000
3	AIR	MAR1949	132	133.826	112.872	154.780	-1.8256	10.6910	00000000
4	AIR	APR1949	129	131.466	110.512	152.420	-2.4659	10.6910	00000000
5	AIR	MAY1949	121	125.627	104.673	146.581	-4.6269	10.6910	00000000
6	AIR	JUN1949	135	142.449	121.495	163.403	-7.4494	10.6910	00000000
7	AIR	JUL1949	148	158.523	137.569	179.477	-10.5228	10.6910	00000000
8	AIR	AUG1949	148	158.197	137.243	179.152	-10.1975	10.6910	00000000
9	AIR	SEP1949	136	144.722	123.768	165.676	-8.7219	10.6910	00000000
10	AIR	OCT1949	119	124.420	103.466	145.374	-5.4202	10.6910	00000000
11	AIR	NOV1949	104	109.181	88.227	130.135	-5.1808	10.6910	00000000
12	AIR	DEC1949	118	127.755	106.801	148.709	-9.7547	10.6910	00000000
13	AIR	JAN1950	115	133.440	112.486	154.394	-18.4401	10.6910	00000000
14	AIR	FEB1950	126	138.205	117.251	159.159	-12.2048	10.6910	00000000
15	AIR	MAR1950	141	153.241	132.287	174.195	-12.2414	10.6910	00000000
16	AIR	APR1950	135	147.096	126.142	168.050	-12.0956	10.6910	00000000
17	AIR	MAY1950	125	136.073	115.119	157.027	-11.0727	10.6910	00000000
18	AIR	JUN1950	149	149.300	128.346	170.254	-0.2997	10.6910	00000000
19	AIR	JUL1950	170	163.708	142.754	184.662	6.2918	10.6910	00000000
20	AIR	AUG1950	170	166.357	145.402	187.311	3.6435	10.6910	00000000

Figure 2 displays the BenStat data set, which contains the statistics of fit for the reconciled forecasts.

Figure 2 Reconciled Forecasts Statistics of Fit

BenStat

NAME	_REGION_	DFE	NMISSA	NOBS	N	NPARMS	NMISSP	TSS	SST	SSE
AIR	FORECAST	144	12	156	156	0	0	13371737	2058044.16	22594.43

MSE	RMSE	UMSE	URMSE	MAPE	MAE	RSQUARE	ADJRSQ	AADJRSQ	RWRSQ
156.906	12.5262	156.906	12.5262	3.89198	9.87504	0.98902	0.98910	0.98902	0.86132

AIC	AICC	SBC	APC	MAXERR	MINERR	MAXPE	MINPE	ME
728.013	728.013	728.013	156.906	33.2377	-35.2824	10.5057	-18.2718	0.035721

MPE	MDAPE	GMAPE	MINPPE	MAXPPE	MPPE	MAPPE	MDAPPE	GMAPPE
-0.73963	3.18661	2.56165	-15.4490	11.7390	-0.50276	3.81678	3.21198	2.54583

MINSPE	MAXSPE	MSPE	SMAPE	MDASPE	GMASPE	MINRE	MAXRE	MRE
-16.7422	11.0881	-0.61906	3.84958	3.18773	2.55297	-24.6884	11.9863	-0.16515

MRAE	MDRAE	GMRAE	MASE	MINAPES	MAXAPES	MAPES	MDAPES	GMAPES
0.98866	0.39664	0.35882	0.38186	0.11913	29.4103	8.23151	7.37656	5.45001

You can vary the reconciled forecasts by selecting the values of the SMOOTH= and EXP= options. Figure 3, Figure 4, and Figure 5 show the original forecasts versus the reconciled forecasts for different combinations of the two parameters.

Figure 3 Original versus Reconciled Forecasts, SMOOTH=0.9, EXP=0

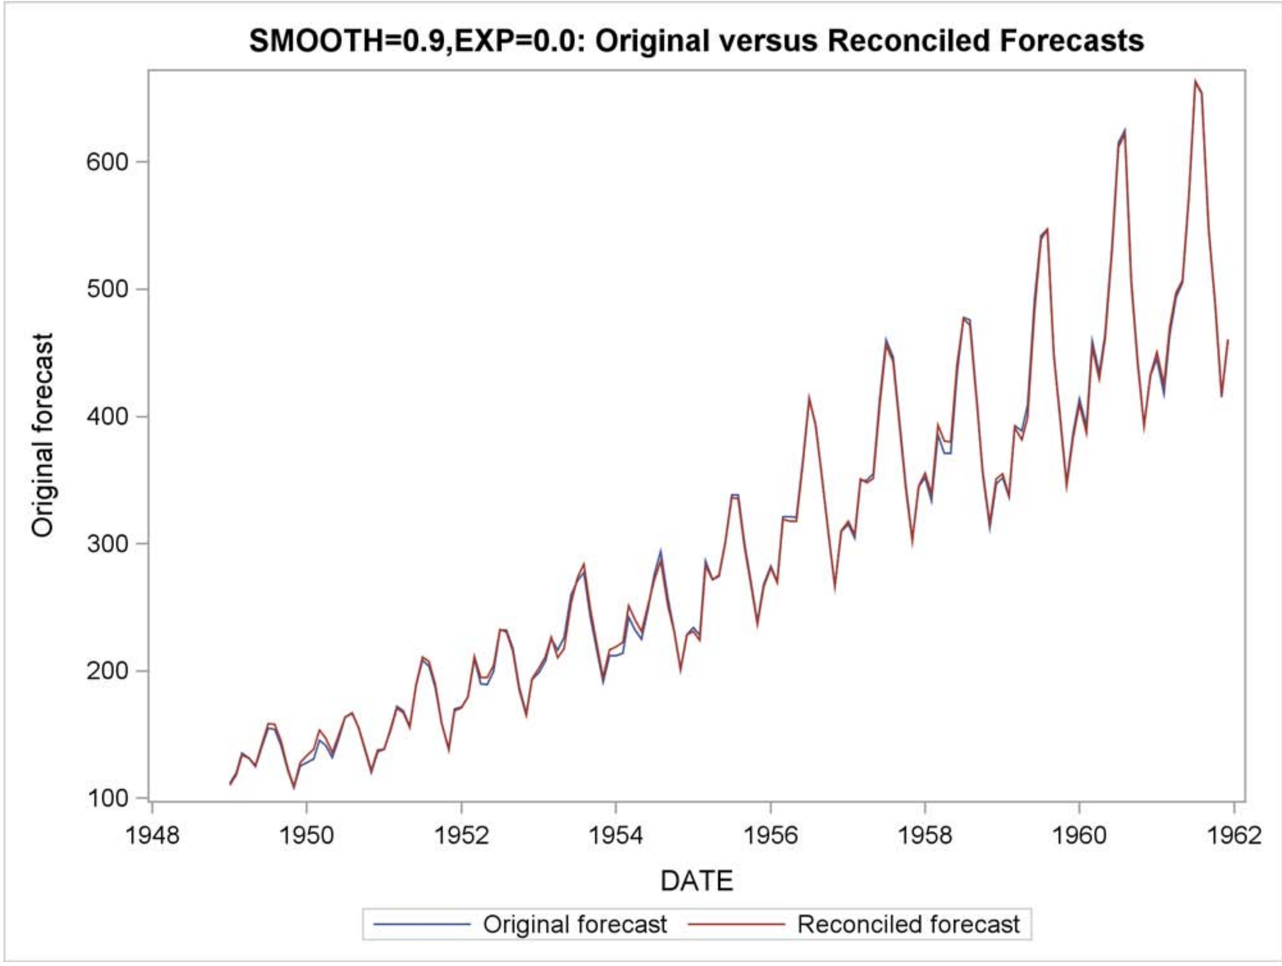

Figure 4 Original versus Reconciled Forecasts, SMOOTH=0.5, EXP=0.5

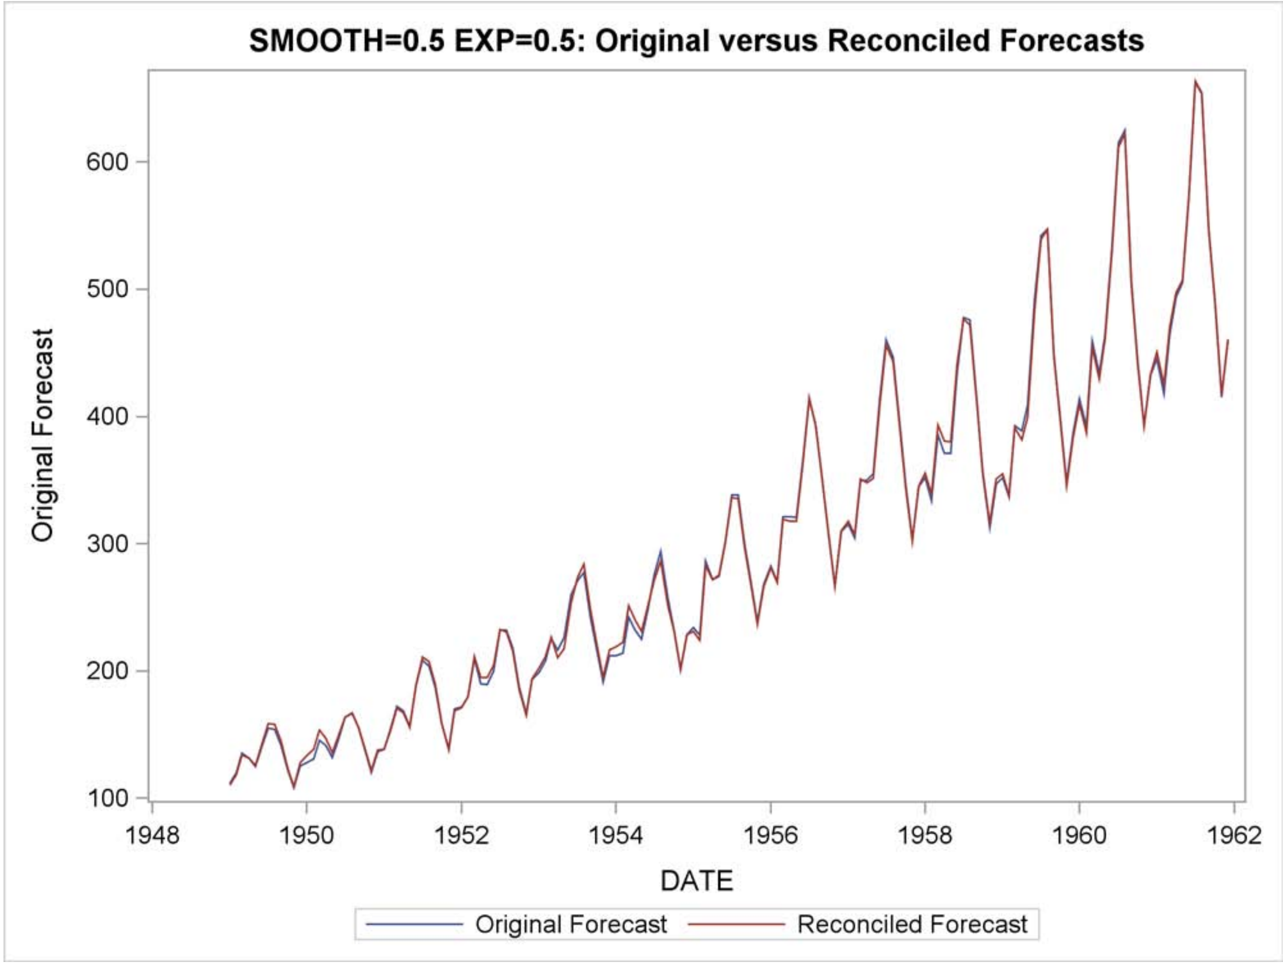

Figure 5 Original versus Reconciled Forecasts, SMOOTH=0.9, EXP=0.99

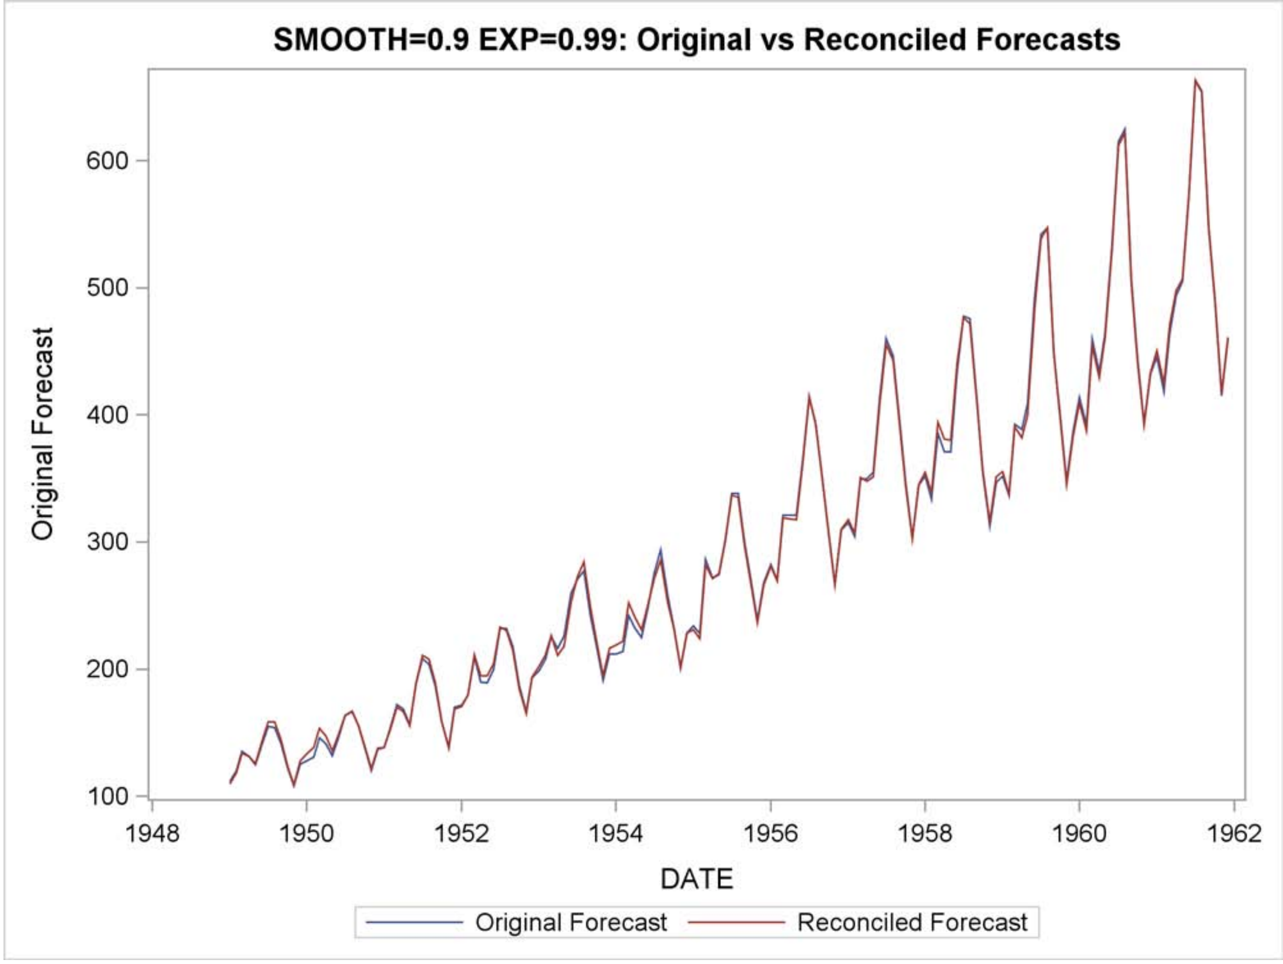

DATA ANALYSIS

This section applies the method discussed in the preceding sections to a data set of real data that consists of several time series, most of which show intermittent behavior. The data represent six years of monthly demand for 753 parts at the British Royal Air Force (RAF), between July 1992 and June 1998, for a total of 72 observations. Demand for spare parts is a typical example in which intermittent demand is usually encountered. And, indeed, a majority of the series in this collection exhibit intermittent behavior.

First, forecasts are generated independently at the monthly and quarterly intervals. Two years of data are used to fit the model. One year is used for out-of-sample model selection. After model selection, the model parameters are estimated again to use the full three years of data. That leaves two years of data for evaluation of the performance of the forecasts. SAS Forecast Server is used to perform model selection. The full details of the model selection procedure it uses can be found in Leonard (2002).

Figure 6 Model Selection and Evaluation

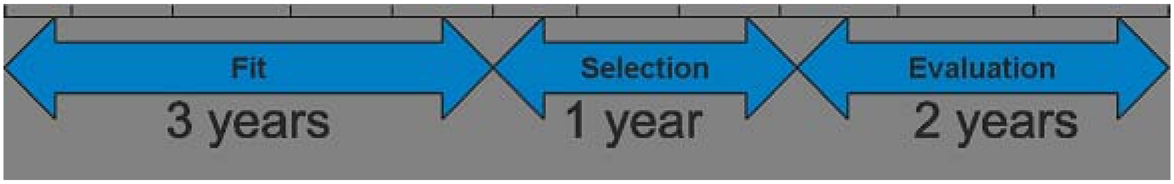

For each series at each frequency, the best model is selected as the one that minimizes the out-of-sample root mean squared error (RMSE) among a large class of models that include ARIMA (autoregressive integrated moving average), ESM (exponential smoothing models), UCM (unobserved component models), and models for intermittent series.

RMSE is chosen as selection criterion because it can be computed unequivocally regardless of the value of the series. The most common selection criterion in the forecasting practice, the mean absolute percentage error (MAPE), is not meaningful with intermittent series.

Figure 7 and Figure 8 display the model family selected for the monthly and the quarterly data, respectively. You can see that for approximately 50% of the monthly series, a model for intermittent data is selected. This proportion is dramatically reduced for the quarterly data.

Figure 7 Model Family Distribution for Monthly Data

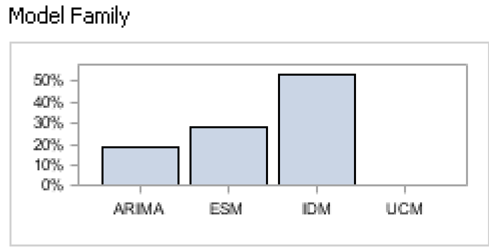

Figure 8 Model Family Distribution for Quarterly Data

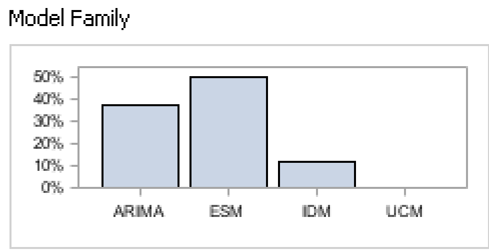

The monthly forecasts are reconciled to the quarterly forecasts for a grid of values of ρ and λ, with $\rho = 0, 0.1, \ldots, 0.9, 1$ and $\lambda = 0, 0.5, 1$. For each series the set of values of ρ and λ is selected as those that minimize the out-of-sample RMSE in the selection interval. Finally, the RMSE of the reconciled forecasts is compared to the RMSE of the original model forecasts in the two-year evaluation period.

The RMSE of the reconciled monthly forecasts for the selected values of ρ and λ is improved for 562 of the 753 series when compared to the original model RMSE. The average improvement for these 562 series is 52%.

CONCLUSION

This paper presents a method for reconciling higher-frequency forecasts to lower-frequency forecasts for a time series accumulated in a hierarchy of time intervals. The method is a based on the minimization of a quadratic loss function subject to the constraint that the reconciled lower-frequency forecasts accumulate to the higher-frequency intervals. Under certain circumstances, the problem can also be interpreted as a regression problem. This method is implemented in the SAS HPFTEMPRECON procedure.

The target function depends on two parameters whose selection can depend on the same criteria that are used to select the models for the forecasts at the two frequencies.

The application of this method can lead to more accurate forecasts when the data at higher frequency are mostly intermittent and therefore are not suitable for models that include features such as input variables, events, and seasonal components.

REFERENCES

Box, G. E. P. and Jenkins, G. M. (1976), *Time Series Analysis: Forecasting and Control*, Revised Edition, San Francisco: Holden-Day.

Croston, J. D. (1972), "Forecasting and Stock Control for Intermittent Demands," *Operations Research Quarterly*, 23, No. 3.

Dagum, E. B. and Cholette, P. A. (2006), *Benchmarking, Temporal Distribution, and Reconciliation Methods for Time Series*, volume 186 of *Lecture Notes in Statistics*, Springer Verlag.

Denton, F. (1971), "Adjustment of Monthly or Quarterly Series to Annual Totals: An Approach Based on Quadratic Minimization," *Journal of the American Statistical Association*, 82, 99–102.

Leonard, M. J. (2002), "Large Scale Automatic Forecasting: Millions of Forecasts," in *International Symposium of Forecasting*.

Quenneville, B., Fortier, S., Chen, Z.-G., and Latendresse, E. (2006), "Recent Developments in Benchmarking to Annual Totals in X-12-ARIMA and at Statistics Canada," in *Proceedings of the Eurostat Conference on Seasonality, Seasonal Adjustment, and Their Implications for Short-Term Analysis and Forecasting*, Luxembourg.

Wei, W. W. S. (1990), *Time Series Analysis: Univariate and Multivariate Methods*, Redwood: Addison-Wesley.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author:

Michele A. Trovero, Ph.D.
SAS Institute Inc.
1430 Route 206 Suite 200
Bedminster, NJ, 09721
Tel.: (908) 470-0142 x52715
Fax: (919) 444-7777
E-mail: Michele.Trovero@sas.com