



SAS Publishing



SAS[®] High-Performance Forecasting 9.1

User's Guide

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2004. *SAS® High-Performance Forecasting 9.1 User's Guide*. Cary, NC: SAS Institute Inc.

SAS® High-Performance Forecasting 9.1 User's Guide

Copyright © 2004, SAS Institute Inc., Cary, NC, USA

ISBN 1-59047-235-7

All rights reserved. Produced in the United States of America. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

U.S. Government Restricted Rights Notice: Use, duplication, or disclosure of this software and related documentation by the U.S. government is subject to the Agreement with SAS Institute and the restrictions set forth in FAR 52.227-19, Commercial Computer Software-Restricted Rights (June 1987).

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

1st printing, January 2004

SAS Publishing provides a complete selection of books and electronic products to help customers use SAS software to its fullest potential. For more information about our e-books, e-learning products, CDs, and hard-copy books, visit the SAS Publishing Web site at support.sas.com/pubs or call 1-800-727-3228.

SAS® and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.

Contents

What's New in SAS 9.1 SAS High-Performance Forecasting	1
Chapter 1. The HPF Procedure	3
Chapter 2. Forecasting Process Details	53
Subject Index	81
Syntax Index	85

What's New in SAS 9.1 SAS High-Performance Forecasting

Overview

New features have been added to the following SAS High-Performance Forecasting components:

- PROC HPF

Details

HPF Procedure

The [HPF](#) procedure has new options related to forecast model selection:

- [HOLDOUTPCT=](#) option
- [SEASONTEST=](#) option
- [SELECT=](#) option

The [HPF](#) procedure has new options related to forecast model initialization:

- [NBACKCAST=](#) option

The [HPF](#) procedure has new statement related to intermittent demand models:

- [IDM](#) statement

The [HPF](#) procedure has new options related to out-of-sample forecast performance statistics:

- [BACK=](#) option
- [PRINT=PERFORMANCE](#) option
- [PRINT=PERFORMANCESUMMARY](#) option
- [PRINT=PERFORMANCEOVERALL](#) option

The [HPF](#) procedure has new options related to the processing of the data:

2 ♦ *What's New in SAS 9.1 SAS High-Performance Forecasting*

- `NOTSORTED` option
- `REPLACEBACK` option
- `SORTNAMES` option

The `HPF` procedure has new options related to printed and graphical output:

- `PLOT=` option
- `PRINT=STATES` option

The `HPF` procedure has a new option related to message log:

- `MAXERROR=` option

Chapter 1

The HPF Procedure

Chapter Contents

OVERVIEW	5
GETTING STARTED	6
SYNTAX	9
Functional Summary	9
PROC HPF Statement	11
BY Statement	14
FORECAST Statement	14
ID Statement	19
IDM Statement	22
Smoothing Model Specification Options for IDM Statement	24
Smoothing Model Parameter Specification Options	27
Smoothing Model Forecast Bounds Options	27
DETAILS	28
Accumulation	28
Missing Value Interpretation	30
Diagnostic Tests	30
Model Selection	30
Transformations	31
Parameter Estimation	31
Missing Value Modeling Issues	31
Forecasting	31
Inverse Transformations	31
Statistics of Fit	32
Forecast Summation	32
Comparison to the Time Series Forecasting System	32
Data Set Output	32
OUT= Data Set	32
OUTEST= Data Set	33
OUTFOR= Data Set	33
OUTSTAT= Data Set	34
OUTSUM= Data Set	35
OUTSEASON= Data Set	36
OUTTREND= Data Set	37
Printed Output	37

ODS Table Names	39
ODS Graphics (Experimental)	39
EXAMPLES	42
Example 1.1. Automatic Forecasting of Time Series Data	42
Example 1.2. Automatic Forecasting of Transactional Data	44
Example 1.3. Specifying the Forecasting Model	45
Example 1.4. Extending the Independent Variables for Multivariate Forecasts	46
Example 1.5. Forecasting Intermittent Time Series Data	47
Example 1.6. Illustration of ODS Graphics (Experimental)	49
REFERENCES	52

Chapter 1

The HPF Procedure

Overview

The HPF (High-Performance Forecasting) procedure provides a quick and automatic way to generate forecasts for many time series or transactional data in one step. The procedure can forecast millions of series at a time, with the series organized into separate variables or across BY groups.

- For typical time series, you can use the following smoothing models:
 - Simple
 - Double
 - Linear
 - Damped Trend
 - Seasonal
 - Winters Method (additive and multiplicative)
- Additionally, transformed versions of these models are provided:
 - Log
 - Square Root
 - Logistic
 - Box-Cox
- For intermittent time series (series where a large number of values are zero-valued), you can use an intermittent demand model such as Croston's Method and Average Demand Model.

Experimental graphics are now available with the HPF procedure. For more information, see the [“ODS Graphics”](#) section on page 39.

All parameters associated with the forecast model are optimized based on the data. Optionally, the HPF procedure can select the appropriate smoothing model for you using holdout sample analysis based on one of several model selection criteria.

The HPF procedure writes the time series extrapolated by the forecasts, the series summary statistics, the forecasts and confidence limits, the parameter estimates, and the fit statistics to output data sets. The HPF procedure optionally produces printed output for these results utilizing the Output Delivery System (ODS).

The HPF procedure can forecast both time series data, whose observations are equally spaced by a specific time interval (e.g., monthly, weekly), or transactional data, whose observations are not spaced with respect to any particular time interval. Internet, inventory, sales, and similar data are typical examples of transactional data. For transactional data, the data is accumulated based on a specified time interval to form a

time series. The HPF procedure can also perform trend and seasonal analysis on transactional data.

Additionally, the Time Series Forecasting System of SAS/ETS software can be used to interactively develop forecasting models, estimate the model parameters, evaluate the models' ability to forecast, and display these results graphically. Refer to Chapter 33, "Overview of the Time Series Forecasting System," (*SAS/ETS User's Guide*) for details.

Also, the EXPAND procedure can be used for the frequency conversion and transformations of time series. Refer to Chapter 16, "The EXPAND Procedure," (*SAS/ETS User's Guide*) for details.

Getting Started

The HPF procedure is simple to use for someone who is new to forecasting, and yet at the same time it is powerful for the experienced professional forecaster who needs to generate a large number of forecasts automatically. It can provide results in output data sets or in other output formats using the Output Delivery System (ODS). The following examples are more fully illustrated in the "Examples" section on page 42.

Given an input data set that contains numerous time series variables recorded at a specific frequency, the HPF procedure can automatically forecast the series as follows:

```
PROC HPF DATA=<input-data-set> OUT=<output-data-set>;
  ID <time-ID-variable> INTERVAL=<frequency>;
  FORECAST <time-series-variables>;
RUN;
```

For example, suppose that the input data set **SALES** contains numerous sales data recorded monthly, the variable that represents time is **DATE**, and the forecasts are to be recorded in the output data set **NEXTYEAR**. The HPF procedure could be used as follows:

```
proc hpf data=sales out=nextyear;
  id date interval=month;
  forecast _ALL_;
run;
```

The above statements automatically select the best fitting model, generate forecasts for every numeric variable in the input data set (**SALES**) for the next twelve months, and store these forecasts in the output data set (**NEXTYEAR**). Other output data sets can be specified to store the parameter estimates, forecasts, statistics of fit, and summary data.

If you want to print the forecasts using the Output Delivery System (ODS), then you need to add **PRINT=FORECASTS**:

```
proc hpf data=sales out=nextyear print=forecasts;
  id date interval=month;
  forecast _ALL_;
run;
```

Other results can be specified to output the parameter estimates, forecasts, statistics of fit, and summary data using ODS.

The HPF procedure can forecast both time series data, whose observations are equally spaced by a specific time interval (e.g., monthly, weekly), or transactional data, whose observations are not spaced with respect to any particular time interval.

Given an input data set containing transactional variables not recorded at any specific frequency, the HPF procedure accumulates the data to a specific time interval and forecasts the accumulated series as follows:

```
PROC HPF DATA=<input-data-set> OUT=<output-data-set>;
  ID <time-ID-variable> INTERVAL=<frequency>
  ACCUMULATE=<accumulation>;
  FORECAST <time-series-variables>;
RUN;
```

For example, suppose that the input data set **WEBSITES** contains three variables (**BOATS**, **CARS**, **PLANES**), that are Internet data recorded on no particular time interval, and the variable that represents time is **TIME**, which records the time of the Web hit. The forecasts for the total daily values are to be recorded in the output data set **NEXTWEEK**. The HPF procedure could be used as follows:

```
proc hpf data=websites out=nextweek lead=7;
  id time interval=dtday accumulate=total;
  forecast boats cars planes;
run;
```

The above statements accumulate the data into a daily time series and automatically generate forecasts for the **BOATS**, **CARS**, and **PLANES** variables in the input data set (**WEBSITES**) for the next seven days and store the forecasts in the output data set (**NEXTWEEK**).

The HPF procedure can specify a particular forecast model or select from several candidate models based on a selection criterion. The HPF procedure also supports transformed models and holdout sample analysis.

Using the previous **WEBSITES** example, suppose that you want to forecast the **BOATS** variable using the best seasonal forecasting model that minimizes the mean absolute percent error (MAPE), the **CARS** variable using the best nonseasonal forecasting model that minimizes the mean square error (MSE) using holdout sample analysis on the last five days, and the **PLANES** variable using the Log Winters Method (additive). The HPF procedure could be used as follows:

```
proc hpf data=websites out=nextweek lead=7;
  id time interval=dtday accumulate=total;
  forecast boats / model=bests select=mape;
  forecast cars / model=bestn select=mse holdout=5;
  forecast planes / model=addwinters transform=log;
run;
```

The above statements demonstrate how each variable in the input data set can be modeled differently and how several candidate models can be specified and selected based on holdout sample analysis or the entire range of data.

The HPF procedure is also useful in extending independent variables in (auto) regression models where future values of the independent variable are needed to predict the dependent variable.

Using the **WEBSITES** example, suppose that you want to forecast the **ENGINES** variable using the **BOATS**, **CARS**, and **PLANES** variable as regressor variables. Since future values of the **BOATS**, **CARS**, and **PLANES** are needed, the HPF procedure can be used to extend these variables in the future:

```
proc hpf data=websites out=nextweek lead=7;
  id time interval=dtday accumulate=total;
  forecast engines / model=none;
  forecast boats / model=bests select=mape;
  forecast cars / model=bestn select=mse holdout=5;
  forecast planes / model=addwinters transform=log;
run;

proc autoreg data= nextweek;
  model engines = boats cars planes;
  output out=enginehits p=predicted;
run;
```

The above HPF procedure statements generate forecasts for **BOATS**, **CARS**, and **PLANES** in the input data set (**WEBSITES**) for the next seven days and extend the variable **ENGINES** with missing values. The output data set (**NEXTWEEK**) of the PROC HPF statement is used as an input data set for the PROC AUTOREG statement. The output data set of PROC AUTOREG contains the forecast of the variable **ENGINE** based on the regression model with the variables **BOATS**, **CARS**, and **PLANES** as regressors. See the AUTOREG procedure for details on autoregression.

The HPF procedure can also forecast intermittent time series (series where a large number of values are zero-valued). Typical time series forecasting techniques are less effective in forecasting intermittent time series.

For example, suppose that the input data set **INVENTORY** contains three variables (**TIRES**, **HUBCAPS**, **LUGBOLTS**) that are demand data recorded on no particular time interval, the variable that represents time is **DATE**, and the forecasts for the total weekly values are to be recorded in the output data set **NEXTMONTH**. The models requested are intermittent demand models, which can be specified as **MODEL=IDM**

option. Two intermittent demand models are compared, Croston Model and Average Demand Model. The HPF procedure could be used as follows:

```
proc hpf data=inventory out=nextmonth lead=4 print=forecasts;
  id date interval=week accumulate=total;
  forecast tires hubcaps lugbolts / model=idm;
run;
```

In the above example, the total demand for inventory items is accumulated on a weekly basis and forecasts are generated that recommend future stocking levels.

Syntax

The following statements are used with the HPF procedure.

```
PROC HPF options;
BY variables;
FORECAST variable-list / options;
ID variable INTERVAL= interval options;
IDM options;
```

Functional Summary

The statements and options controlling the HPF procedure are summarized in the following table.

Description	Statement	Option
Statements		
specify BY-group processing	BY	
specify variables to forecast	FORECAST	
specify the time ID variable	ID	
specify intermittent demand model	IDM	
Data Set Options		
specify the input data set	PROC HPF	DATA=
specify to output forecasts only	PROC HPF	NOOUTALL
specify the output data set	PROC HPF	OUT=
specify parameter output data set	PROC HPF	OUTEST=
specify forecast output data set	PROC HPF	OUTFOR=
specify seasonal statistics output data set	PROC HPF	OUTSEASON=
specify statistics output data set	PROC HPF	OUTSTAT=
specify summary output data set	PROC HPF	OUTSUM=
specify trend statistics output data set	PROC HPF	OUTTREND=
replace actual values held back	FORECAST	REPLACEBACK

Description	Statement	Option
replace missing values	FORECAST	REPLACEMISSING
use forecast value to append	FORECAST	USE=
Accumulation and Seasonality Options		
specify accumulation frequency	ID	INTERVAL=
specify length of seasonal cycle	PROC HPF	SEASONALITY=
specify interval alignment	ID	ALIGN=
specify time ID variable values are not sorted	ID	NOTSORTED
specify starting time ID value	ID	START=
specify ending time ID value	ID	END=
specify accumulation statistic	ID, FORECAST	ACCUMULATE=
specify missing value interpretation	ID, FORECAST	SETMISSING=
specify zero value interpretation	ID, FORECAST	ZEROMISS=
Forecasting Horizon, Holdout, Holdback Options		
specify data to hold back	PROC HPF	BACK=
specify forecast holdout sample size	FORECAST	HOLDOUT=
specify forecast holdout sample percent	FORECAST	HOLDOUTPCT=
specify forecast horizon or lead	PROC HPF	LEAD=
specify horizon to start summation	PROC HPF	STARTSUM=
Forecasting Model and Selection Options		
specify confidence limit width	FORECAST	ALPHA=
specify intermittency	FORECAST	INTERMITTENT=
specify forecast model	FORECAST	MODEL=
specify median forecasts	FORECAST	MEDIAN
specify backcast initialization	FORECAST	NBACKCAST=
specify seasonality test	FORECAST	SEASONTEST=
specify model selection criterion	FORECAST	SELECT=
specify model transformation	FORECAST	TRANSFORM=
Intermittent Demand Model Options		
specify model for average demand	IDM	AVERAGE=
specify the base value	IDM	BASE=
specify model for demand intervals	IDM	INTERVAL=
specify model for demand sizes	IDM	SIZE=
Printing and Plotting Control Options		
specify graphical output	PROC HPF	PLOT=
specify printed output	PROC HPF	PRINT=
specify detailed printed output	PROC HPF	PRINTDETAILS

Description	Statement	Option
Miscellaneous Options		
specify that analysis variables are processed in sorted order	PROC HPF	SORTNAMES
limits error and warning messages	PROC HPF	MAXERROR=

PROC HPF Statement

PROC HPF *options*;

The following options can be used in the PROC HPF statement.

BACK= *n*

specifies the number of observations before the end of the data that the multistep forecasts are to begin. The default is BACK=0.

DATA= *SAS-data-set*

names the SAS data set containing the input data for the procedure to forecast. If the DATA= option is not specified, the most recently created SAS data set is used.

LEAD= *n*

specifies the number of periods ahead to forecast (forecast lead or horizon). The default is LEAD=12.

The LEAD= value is relative to the last observation in the input data set and not to the last nonmissing observation of a particular series. Thus, if a series has missing values at the end, the actual number of forecasts computed for that series will be greater than the LEAD= value.

MAXERROR= *number*

limits the number of warning and error messages produced during the execution of the procedure to the specified value. The default is MAXERRORS=50. This option is particularly useful in BY-group processing where it can be used to suppress the recurring messages.

NOOUTALL

specifies that only forecasts are written to the OUT= and OUTFOR= data sets. The NOOUTALL option includes only the final forecast observations in the output data sets, not the one-step forecasts for the data before the forecast period.

The OUT= and OUTFOR= data set will only contain the forecast results starting at the next period following the last observation to the forecast horizon specified by the LEAD= option.

OUT= *SAS-data-set*

names the output data set to contain the forecasts of the variables specified in the subsequent FORECAST statements. If an ID variable is specified, it will also be included in the OUT= data set. The values are accumulated based on the ACCUMULATE=

option and forecasts are appended to these values based on the FORECAST statement USE= option. The OUT= data set is particularly useful in extending the independent variables when forecasting dependent series associated with (auto) regression models. If the OUT= option is not specified, a default output data set DATAn is created. If you do not want the OUT= data set created, then use OUT=_NULL_.

OUTEST= SAS-data-set

names the output data set to contain the model parameter estimates and the associated test statistics and probability values. The OUTEST= data set is particularly useful for evaluating the significance of the model parameters and understanding the model dynamics.

OUTFOR= SAS-data-set

names the output data set to contain the forecast time series components (actual, predicted, lower confidence limit, upper confidence limit, prediction error, prediction standard error). The OUTFOR= data set is particularly useful for displaying the forecasts in tabular or graphical form.

OUTSEASON= SAS-data-set

names the output data set to contain the seasonal statistics. The statistics are computed for each season as specified by the ID statement INTERVAL= option or the SEASONALITY= option. The OUTSEASON= data set is particularly useful when analyzing transactional data for seasonal variations.

OUTSTAT= SAS-data-set

names the output data set to contain the statistics of fit (or goodness-of-fit statistics). The OUTSTAT= data set is particularly useful for evaluating how well the model fits the series. The statistics of fit are based on the entire range of the time series regardless of whether the HOLDOUT= option is specified.

OUTSUM= SAS-data-set

names the output data set to contain the summary statistics and the forecast summation. The summary statistics are based on the accumulated time series when the ACCUMULATE= or SETMISSING= options are specified. The forecast summations are based on the LEAD=, STARTSUM=, and USE= options. The OUTSUM= data set is particularly useful when forecasting large numbers of series and a summary of the results are needed.

OUTTREND= SAS-data-set

names the output data set to contain the trend statistics. The statistics are computed for each time period as specified by the ID statement INTERVAL= option. The OUTTREND= data set is particularly useful when analyzing transactional data for trends.

PRINT= option | (options)

specifies the printed output desired. By default, the HPF procedure produces no printed output. The following printing options are available:

ESTIMATES	prints the results of parameter estimation. (OUTEST= data set)
FORECASTS	prints the forecasts. (OUTFOR= data set)

PERFORMANCE	prints the performance statistics for each forecast.
PERFORMANCESUMMARY	prints the performance summary for each BY group.
PERFORMANCEOVERALL	prints the performance summary for all of the BY groups.
SEASONS	prints the seasonal statistics. (OUTSEASON= data set)
STATISTICS	prints the statistics of fit. (OUTSTAT= data set)
STATES	prints the backcast, initial, and final states.
SUMMARY	prints the summary statistics for the accumulated time series. (OUTSUM= data set)
TRENDS	prints the trend statistics. (OUTTREND= data set)
ALL	Same as PRINT=(ESTIMATES FORECASTS STATISTICS SUMMARY). PRINT=(ALL,TRENDS,SEASONS) prints all of the options listed above.

For example, PRINT=FORECASTS prints the forecasts, PRINT=(ESTIMATES, FORECASTS) prints the parameter estimates and the forecasts, and PRINT=ALL prints all of the above output.

The PRINT= option produces printed output for these results utilizing the Output Delivery System (ODS). The PRINT= option produces results similar to the data sets listed next to the above options in parenthesis.

PRINTDETAILS

specifies that output requested with the PRINT= option be printed in greater detail.

SEASONALITY= *number*

specifies the length of the seasonal cycle. For example, SEASONALITY=3 means that every group of three observations forms a seasonal cycle. The SEASONALITY= option is applicable only for seasonal forecasting models. By default, the length of the seasonal cycle is one (no seasonality) or the length implied by the INTERVAL= option specified in the ID statement. For example, INTERVAL=MONTH implies that the length of the seasonal cycle is twelve.

SORTNAMES

specifies that the variables specified in the FORECAST statements are processed in sorted order.

STARTSUM= *n*

specifies the starting forecast lead (or horizon) for which to begin summation of the forecasts specified by the LEAD= option. The STARTSUM= value must be less than the LEAD= value. The default is STARTSUM=1, that is, the sum from the one-step ahead forecast to the multistep forecast specified by the LEAD= option.

The prediction standard errors of the summation of forecasts takes into account the correlation between the multistep forecasts. The DETAILS section describes the STARTSUM= option in more detail.

BY Statement

BY *variables;*

A BY statement can be used with PROC HPF to obtain separate analyses for groups of observations defined by the BY variables.

FORECAST Statement

FORECAST *variable-list / options;*

The FORECAST statement lists the numeric variables in the DATA= data set whose accumulated values represent time series to be modeled and forecast. The options specify which forecast model is to be used or how the forecast model is selected from several possible candidate models.

A data set variable can be specified in only one FORECAST statement. Any number of FORECAST statements can be used. The following options can be used with the FORECAST statement.

ACCUMULATE= *option*

specifies how the data set observations are accumulated within each time period for the variables listed in the FORECAST statement. If the ACCUMULATE= option is not specified in the FORECAST statement, accumulation is determined by the ACCUMULATE= option of the ID statement. See the ID statement ACCUMULATE= option for more details.

ALPHA= *number*

specifies the significance level to use in computing the confidence limits of the forecast. The ALPHA=value must be between 0 and 1. The default is ALPHA=0.05, which produces 95% confidence intervals.

HOLDOUT= *n*

specifies the size of the holdout sample to be used for model selection. The holdout sample is a subset of actual time series ending at the last nonmissing observation. If the ACCUMULATE= option is specified, the holdout sample is based on the accumulated series. If the holdout sample is not specified, the full range of the actual time series is used for model selection.

For each candidate model specified, the holdout sample is excluded from the initial model fit and forecasts are made within the holdout sample time range. Then, for each candidate model specified, the statistic of fit specified by the SELECT= option is computed using only the observations in the holdout sample. Finally, the candidate model, which performs best in the holdout sample, based on this statistic, is selected to forecast the actual time series.

The HOLDOUT= option is only used to select the best forecasting model from a list of candidate models. After the best model is selected, the full range of the actual time series is used for subsequent model fitting and forecasting. It is possible that one model will outperform another model in the holdout sample but perform less well when the entire range of the actual series is used.

If MODEL=BESTALL and HOLDOUT= options are used together, the last one hundred observations are used to determine whether the series is intermittent. If the series determined not to be intermittent, holdout sample analysis will be used to select the smoothing model.

HOLDOUTPCT= *number*

specifies the size of the holdout sample as a percentage of the length of the time series. If HOLDOUT=5 and HOLDOUTPCT=10, the size of the holdout sample is $\min(5, 0.1T)$ where T is the length of the time series with beginning and ending missing values removed. The default is 100 (100%).

INTERMITTENT= *number*

specifies a number greater than one which is used to determine whether or not a time series is intermittent. If the average demand interval is greater than this number then the series is assumed to be intermittent. This option is used with MODEL=BESTALL option. The default is INTERMITTENT=1.25.

MEDIAN

specifies that the median forecast values are to be estimated. Forecasts can be based on the mean or median. By default the mean value is provided. If no transformation is applied to the actual series using the TRANSFORM= option, the mean and median forecast values are identical.

MODEL= *model-name*

specifies the forecasting model to be used to forecast the actual time series. A single model can be specified or a group of candidate models can be specified. If a group of models is specified, the model used to forecast the accumulated time series is selected based on the SELECT= option and the HOLDOUT= option. The default is MODEL=BEST. The following forecasting models are provided:

NONE	No forecast. The accumulated time series is appended with missing values in the OUT= data set. This option is particularly useful when the results stored in the OUT= data set are subsequently used in (auto) regression analysis where forecasts of the independent variables are needed to forecast the dependent variable.
SIMPLE	Simple (Single) Exponential Smoothing
DOUBLE	Double (Brown) Exponential Smoothing
LINEAR	Linear (Holt) Exponential Smoothing
DAMPTREND	Damped Trend Exponential Smoothing
SEASONAL	Seasonal Exponential Smoothing
WINTERS	Winters Multiplicative Method
ADDWINTERS	Winters Additive Method
BEST	Best Candidate Smoothing Model (SIMPLE, DOUBLE, LINEAR, DAMPTREND, SEASONAL, WINTERS, ADDWINTERS)
BESTN	Best Candidate Nonseasonal Smoothing Model (SIMPLE, DOUBLE, LINEAR, DAMPTREND)

BESTS	Best Candidate Seasonal Smoothing Model (SEASONAL, WINTERS, ADDWINTERS)
IDM CROSTON	Intermittent Demand Model such as Croston's Method or Average Demand Model. An intermittent time series is one whose values are mostly zero.
BESTALL	Best Candidate Model (IDM, BEST)

The BEST, BESTN, and BESTS options specify a group of models by which the HOLDOUT= option and SELECT= option are used to select the model used to forecast the accumulated time series based on holdout sample analysis. Transformed versions of the above smoothing models can be specified using the TRANSFORM= option.

The BESTALL option specifies that if the series is intermittent, an intermittent demand model such as Croston's Method or Average Demand Model (MODEL=IDM) is selected; otherwise, the best smoothing model is selected (MODEL=BEST). Intermittency is determined by the INTERMITTENT= option.

The documentation for [Chapter 2, "Forecasting Process Details,"](#) describes the above smoothing models and intermittent models in greater detail.

NBACKCAST= *n*

specifies the number of observations used to initialize the backcast states. The default is the entire series.

REPLACEBACK

specifies that actual values excluded by the BACK= option are replaced with one-step-ahead forecasts in the OUT= data set.

REPLACEMISSING

specifies that embedded missing actual values are replaced with one-step-ahead forecasts in the OUT= data set.

SEASONTEST= *option*

specifies the options related to the seasonality test. This option is used with MODEL=BEST and MODEL=BESTALL option.

The following options are provided:

SEASONTEST=NONE no test

SEASONTEST=(SIGLEVEL=number) Significance probability value.

Series with strong seasonality have small test probabilities. SEASONTEST=(SIGLEVEL=0) always implies seasonality. SEASONTEST=(SIGLEVEL=1) always implies no seasonality. The default is SEASONTEST=(SIGLEVEL=0.01).

SELECT= *option*

specifies the model selection criterion (statistic of fit) to be used to select from several candidate models. This option would often be used in conjunction with the

HOLDOUT= option. The default is SELECT=RMSE. The following statistics of fit are provided:

SSE	Sum of Square Error
MSE	Mean Square Error
RMSE	Root Mean Square Error
UMSE	Unbiased Mean Square Error
URMSE	Unbiased Root Mean Square Error
MAXPE	Maximum Percent Error
MINPE	Minimum Percent Error
MPE	Mean Percent Error
MAPE	Mean Absolute Percent Error
MDAPE	Median Percent Error
GMAPE	Geometric Mean Percent Error
MINPPE	Minimum Predictive Percent Error
MAXPPE	Maximum Predictive Percent Error
MSPPE	Mean Predictive Percent Error
MAPPE	Symmetric Mean Absolute Predictive Percent Error
MDAPPE	Median Predictive Percent Error
GMAPPE	Geometric Mean Predictive Percent Error
MINSPE	Minimum Symmetric Percent Error
MAXSPE	Maximum Symmetric Percent Error
MSPE	Mean Symmetric Percent Error
SMAPE	Symmetric Mean Absolute Percent Error
MDASPE	Median Symmetric Percent Error
GMASPE	Geometric Mean Symmetric Percent Error
MINRE	Minimum Relative Error
MAXRE	Maximum Relative Error
MRE	Mean Relative Error
MRAE	Mean Relative Absolute Error
MDRAE	Median Relative Absolute Error
GMRAE	Geometric Mean Relative Absolute Error
MAXERR	Maximum Error
MINERR	Minimum Error
ME	Mean Error
MAE	Mean Absolute Error
RSQUARE	R-Square

ADJRSQ	Adjusted R-Square
AADJRSQ	Amemiya's Adjusted R-Square
RWRSQ	Random Walk R-Square
AIC	Akaike Information Criterion
SBC	Schwarz Bayesian Information Criterion
APC	Amemiya's Prediction Criterion

SETMISSING= option | number

specifies how missing values (either actual or accumulated) are assigned in the accumulated time series for variables listed in the FORECAST statement. If the SETMISSING= option is not specified in the FORECAST statement, missing values are set based on the SETMISSING= option of the ID statement. See the ID statement SETMISSING= option for more details.

TRANSFORM= option

specifies the time series transformation to be applied to the actual time series. The following transformations are provided:

NONE	No transformation is applied. This option is the default.
LOG	Logarithmic transformation
SQRT	Square-root transformation
LOGISTIC	Logistic transformation
BOXCOX(<i>n</i>)	Box-Cox transformation with parameter number where number is between -5 and 5
AUTO	Automatically choose between NONE and LOG based on model selection criteria.

When the TRANSFORM= option is specified the time series must be strictly positive. Once the time series is transformed, the model parameters are estimated using the transformed series. The forecasts of the transformed series are then computed, and finally, the transformed series forecasts are inverse transformed. The inverse transform produces either mean or median forecasts depending on whether the MEDIAN option is specified.

The TRANSFORM= option is not applicable when MODEL=IDM is specified.

USE= option

specifies which forecast values are appended to the actual values in the OUT= and OUTSUM= data sets. The following USE= options are provided:

PREDICT	The predicted values are appended to the actual values. This option is the default.
LOWER	The lower confidence limit values are appended to the actual values.

UPPER The upper confidence limit values are appended to the actual values.

Thus, the `USE=` option enables the `OUT=` and `OUTSUM=` data sets to be used for worst/best/average/median case decisions.

ZEROMISS= option

specifies how beginning and/or ending zero values (either actual or accumulated) are interpreted in the accumulated time series for variables listed in the `FORECAST` statement. If the `ZEROMISS=` option is not specified in the `FORECAST` statement, missing values are set based on the `ZEROMISS=` option of the `ID` statement. See the `ID` statement `ZEROMISS=` option for more details.

ID Statement

ID variable *INTERVAL= interval options;*

The `ID` statement names a numeric variable that identifies observations in the input and output data sets. The `ID` variable's values are assumed to be SAS date, time, or datetime values. In addition, the `ID` statement specifies the (desired) frequency associated with the actual time series. The `ID` statement options also specify how the observations are accumulated and how the time `ID` values are aligned to form the actual time series. The information specified affects all variables specified in subsequent `FORECAST` statements. If the `ID` statement is specified, the `INTERVAL=` option must also be specified. If an `ID` statement is not specified, the observation number, with respect to the `BY` group, is used as the time `ID`.

The following options can be used with the `ID` statement.

ACCUMULATE= option

specifies how the data set observations are accumulated within each time period. The frequency (width of each time interval) is specified by the `INTERVAL=` option. The `ID` variable contains the time `ID` values. Each time `ID` variable value corresponds to a specific time period. The accumulated values form the actual time series, which is used in subsequent model fitting and forecasting.

The `ACCUMULATE=` option is particularly useful when there are zero or more than one input observations coinciding with a particular time period (e.g., transactional data). The `EXPAND` procedure offers additional frequency conversions and transformations that can also be useful in creating a time series.

The following options determine how the observations are accumulated within each time period based on the `ID` variable and the frequency specified by the `INTERVAL=` option:

NONE No accumulation occurs; the `ID` variable values must be equally spaced with respect to the frequency. This is the default option.

TOTAL	Observations are accumulated based on the total sum of their values.
AVERAGE AVG	Observations are accumulated based on the average of their values.
MINIMUM MIN	Observations are accumulated based on the minimum of their values.
MEDIAN MED	Observations are accumulated based on the median of their values.
MAXIMUM MAX	Observations are accumulated based on the maximum of their values.
N	Observations are accumulated based on the number of nonmissing observations.
NMISS	Observations are accumulated based on the number of missing observations.
NOBS	Observations are accumulated based on the number of observations.
FIRST	Observations are accumulated based on the first of their values.
LAST	Observations are accumulated based on the last of their values.
STDDEV STD	Observations are accumulated based on the standard deviation of their values.
CSS	Observations are accumulated based on the corrected sum of squares of their values.
USS	Observations are accumulated based on the uncorrected sum of squares of their values.

If the `ACCUMULATE=` option is specified, the `SETMISSING=` option is useful for specifying how accumulated missing values are treated. If missing values should be interpreted as zero, then `SETMISSING=0` should be used. The `DETAILS` section describes accumulation in greater detail.

ALIGN= option

controls the alignment of SAS dates used to identify output observations. The `ALIGN=` option accepts the following values: `BEGINNING` | `BEG` | `B`, `MIDDLE` | `MID` | `M`, and `ENDING` | `END` | `E`. `BEGINNING` is the default.

END= option

specifies a SAS date, datetime, or time value that represents the end of the data. If the last time ID variable value is less than the `END=` value, the series is extended with missing values. If the last time ID variable value is greater than the `END=` value, the series is truncated. For example, `END="&sysdate"D` uses the automatic macro variable `SYSDATE` to extend or truncate the series to the current date. This option and the `START=` option can be used to ensure that data associated with each BY group contains the same number of observations.

INTERVAL= *interval*

specifies the frequency of the input time series. For example, if the input data set consists of quarterly observations, then INTERVAL=QTR should be used. If the SEASONALITY= option is not specified, the length of the seasonal cycle is implied from the INTERVAL= option. For example, INTERVAL=QTR implies a seasonal cycle of length 4. If the ACCUMULATE= option is also specified, the INTERVAL= option determines the time periods for the accumulation of observations.

The basic intervals are YEAR, SEMIYEAR, QTR, MONTH, SEMIMONTH, TENDAY, WEEK, WEEKDAY, DAY, HOUR, MINUTE, SECOND. Refer to SAS/ETS User's Guide chapter on Date Interval, Foremats, and Functions for the intervals that can be specified.

NOTSORTED

specifies that the time ID values are not in sorted order. The HPF procedure will sort the data with respect to the time ID prior to analysis.

SETMISSING= *option* | *number*

specifies how missing values (either actual or accumulated) are assigned in the accumulated time series. If a number is specified, missing values are set to number. If a missing value indicates an unknown value, this option should not be used. If a missing value indicates no value, a SETMISSING=0 should be used. You would typically use SETMISSING=0 for transactional data because no recorded data usually implies no activity. The following options can also be used to determine how missing values are assigned:

MISSING	Missing values are set to missing. This is the default option.
AVERAGE AVG	Missing values are set to the accumulated average value.
MINIMUM MIN	Missing values are set to the accumulated minimum value.
MEDIAN MED	Missing values are set to the accumulated median value.
MAXIMUM MAX	Missing values are set to the accumulated maximum value.
FIRST	Missing values are set to the accumulated first nonmissing value.
LAST	Missing values are set to the accumulated last nonmissing value.
PREVIOUS PREV	Missing values are set to the previous accumulated nonmissing value. Missing values at the beginning of the accumulated series remain missing.
NEXT	Missing values are set to the next accumulated nonmissing value. Missing values at the end of the accumulated series remain missing.

If SETMISSING=MISSING is specified and the MODEL= option specifies a smoothing model, the missing observations are smoothed over. If MODEL=IDM

is specified, missing values are assumed to be periods of no demand, that is, SETMISSING=MISSING is equivalent to SETMISSING=0.

START= option

specifies a SAS date, datetime, or time value that represents the beginning of the data. If the first time ID variable value is greater than the START= value, the series is prepended with missing values. If the first time ID variable value is less than the START= value, the series is truncated. This option and the END= option can be used to ensure that data associated with each by group contains the same number of observations.

ZEROMISS= option

specifies how beginning and/or ending zero values (either actual or accumulated) are interpreted in the accumulated time series. The following options can also be used to determine how beginning and/or ending zero values are assigned:

NONE	Beginning and/or ending zeros unchanged. This is the default.
LEFT	Beginning zeros are set to missing.
RIGHT	Ending zeros are set to missing.
BOTH	Both beginning and ending zeros are set to missing.

If the accumulated series is all missing and/or zero the series is not changed.

IDM Statement

IDM options;

The IDM statement is used to specify an intermittent demand model. An intermittent demand series can be analyzed in two ways: individually modeling both demand interval and size component or jointly modeling these components using the average demand component (demand size divided by demand interval). The IDM statement specifies the two smoothing models to be used to forecast the demand interval component (INTERVAL= option) and the demand size component (SIZE= option), or specifies the single smoothing model to be used to forecast the average demand component (AVERAGE= option). Therefore, two smoothing models (INTERVAL= and SIZE= options) must be specified or one smoothing model (AVERAGE= option) must be specified. Only one statement can be specified.

The following examples illustrate typical uses of the IDM statement:

```
/* default specification */
idm;

/* demand interval model only specification */
idm interval=(transform=log);

/* demand size model only specification */
idm size=(method=linear);
```

```

/* Croston's Method */
idm interval=(method=simple)
    size      =(method=simple);

/* Log Croston's Method */
idm interval=(method=simple transform=log)
    size      =(method=simple transform=log);

/* average demand model specification */
idm average=(method=bestn);

```

The default specification uses both the INTERVAL= option and SIZE= option defaults for the decomposed (Croston's) demand model and the AVERAGE= option defaults for the average demand model.

The following example illustrates how to automatically choose the decomposed demand model using MAPE as the model selection criterion:

```

idm interval=(method=simple transform=auto select=mape)
    size      =(method=simple transform=auto select=mape);
forecast sales / model=idm select=mape;

```

The preceding fits two forecast models (simple and log simple exponential smoothing) to both the demand interval and size components. The forecast model that results in the lowest in-sample MAPE for each component is used to forecast the component.

The following example illustrates how to automatically choose the average demand model using MAPE as the model selection criterion:

```

idm average=(method=simple transform=auto select=mape);
forecast sales / model=idm;

```

The preceding fits two forecast models (simple and log simple exponential smoothing) to the average demand component. The forecast model that results in the lowest in-sample MAPE is used to forecast the component.

Combining the above two examples, the following example illustrates how to automatically choose between the decomposed demand model and the average demand model using MAPE as the model selection criterion:

```

idm interval=(method=simple transform=auto select=mape)
    size      =(method=simple transform=auto select=mape)
    average   =(method=simple transform=auto select=mape);
forecast sales / model=idm select=mape;

```

The preceding automatically selects between the decomposed demand model and the average demand model as described previously. The forecast model that results in the lowest in-sample MAPE is used to forecast the series.

The following options can be specified in the IDM statement:

INTERVAL=(smoothing-model-options)

specifies the smoothing model used to forecast the demand interval component. See smoothing model specification options described below.

SIZE=(smoothing-model-options)

specifies the smoothing model used to forecast the demand size component. See smoothing model specification options described below.

AVERAGE=(smoothing-model-options)

specifies the smoothing model used to forecast the demand average component. See smoothing model specification options described below.

BASE=AUTO | number

specifies the base value of the time series used to determine the demand series components. The demand series components are determined based on the departures from this base value. If a base value is specified, this value is used to determine the demand series components. If BASE=AUTO is specified, the time series properties are used to automatically adjust the time series. For the common definition of Croston's Method use BASE=0 which defines departures based on zero. The default is BASE=0.

Given a time series, y_t , and base value, b , the time series is adjusted by the base value to create the base adjusted time series, $x_t = y_t - b$. Demands are assumed to occur when the base adjusted series is nonzero (or when the time series, y_t , departs from the base value, b).

When BASE=AUTO, the base value is automatically determined by the time series median, minimum, and maximum value and the INTERMITTENT= option value of the FORECAST statement.

Smoothing Model Specification Options for IDM Statement

The smoothing model options describe how to forecast the demand interval, size, and average demand components (INTERVAL= option, SIZE= option, and AVERAGE= option).

If the smoothing model options are not specified, the following are the defaults for the demand interval, size, and average components.

```
interval=(transform=auto method=bestn
          levelrest=(0.0001 0.9999)
          trendrest=(0.0001 0.9999)
          damprest =(0.0001 0.9999) select=rmse bounds=(1,.));

size      =(transform=auto method=bestn
          levelrest=(0.0001 0.9999)
          trendrest=(0.0001 0.9999)
          damprest =(0.0001 0.9999) select=rmse);

average   =(transform=auto method=bestn
          levelrest=(0.0001 0.9999)
```

```
trendrest=(0.0001 0.9999)
damprest =(0.0001 0.9999) select=rmse);
```

The above smoothing model options provide the typical automation in intermittent demand model selection.

The following describes the smoothing model options:

TRANSFORM= option

specifies the time series transformation to be applied to the demand component. The following transformations are provided:

NONE	No transformation is applied.
LOG	Logarithmic transformation
SQRT	Square-root transformation
LOGISTIC	Logistic transformation
BOXCOX(<i>n</i>)	Box-Cox transformation with parameter number where number is between -5 and 5
AUTO	Automatically choose between NONE and LOG based on model selection criteria. This option is the default.

When the TRANSFORM= option is specified, the demand component must be strictly positive. Once the demand component is transformed, the model parameters are estimated using the transformed component. The forecasts of the transformed component are then computed, and finally, the transformed component forecasts are inverse transformed. The inverse transform produces either mean or median forecasts depending on whether the MEDIAN option is specified.

MEDIAN

specifies that the median forecast values are to be estimated. Forecasts can be based on the mean or median. By default the mean value is provided. If no transformation is applied to the actual series using the TRANSFORM= option, the mean and median component forecast values are identical.

METHOD= method-name

specifies the forecasting model to be used to forecast the demand component. A single model can be specified or a group of candidate models can be specified. If a group of models is specified, the model used to forecast the accumulated time series is selected based on the SELECT= option of the IDM statement and the HOLDOUT= option of the FORECAST statement. The default is METHOD=BESTN. The following forecasting models are provided:

SIMPLE	Simple (Single) Exponential Smoothing
DOUBLE	Double (Brown) Exponential Smoothing
LINEAR	Linear (Holt) Exponential Smoothing
DAMPTREND	Damped Trend Exponential Smoothing

BESTN Best Candidate Nonseasonal Smoothing Model (SIMPLE, DOUBLE, LINEAR, DAMPTREND)

NOSTABLE

specifies that the smoothing model parameters are not restricted to the additive invertible region of the parameter space. By default, the smoothing model parameters are restricted to be inside the additive invertible region of the parameter space.

LEVELPARM= *number*

specifies the level weight parameter initial value. See the smoothing model parameter specifications options below.

LEVELREST=(*number,number*)

specifies the level weight parameter restrictions. See the smoothing model parameter specifications options below.

TRENDPARM= *number*

specifies the trend weight parameter initial value. See the smoothing model parameter specifications options below.

TRENDREST=(*number,number*)

specifies the trend weight parameter restrictions. See the smoothing model parameter specifications options below.

DAMPPARM= *number*

specifies the damping weight parameter initial value. See the smoothing model parameter specifications options below.

DAMPREST=(*number,number*)

specifies the damping weight parameter restrictions. See the smoothing model parameter specifications options below.

NOEST

specifies that the smoothing model parameters are fixed values. To use this option, all of the smoothing model parameters must be explicitly specified. By default, the smoothing model parameters are optimized.

BOUNDS=(*number,number*)

Specifies the component forecast bound. See the smoothing model forecast bounds below.

SELECT= *option*

specifies the model selection criterion (statistic of fit) to be used to select from several candidate models. This option would often be used in conjunction with the HOLDOUT= option specified in the FORECAST statement. The default is SELECT=RMSE. The statistics of fit provided are the same as those provided in the FORECAST statement.

Smoothing Model Parameter Specification Options

The parameter options are used to specify smoothing model parameters. If the parameter restrictions are not specified the default is (0.0001 0.9999), which implies that the parameters are restricted between 0.0001 and 0.9999. Parameters and their restrictions are required to be greater than or equal to -1 and less than or equal to 2. Missing values indicate no lower and/or upper restriction. If the parameter initial values are not specified, the optimizer uses a grid search to find an appropriate initial value.

Smoothing Model Forecast Bounds Options

Specifies the demand component forecast bounds. The forecast bounds restrict the component forecasts. The default for demand interval forecasts is BOUNDS=1. The lower bound for the demand interval forecast must be greater than one. The default for demand size forecasts is BOUNDS=(.,.) or no bounds. The demand size forecasts bounds are applied after the forecast is adjusted by the base value.

Details

The HPF procedure can be used to perform trend and seasonal analysis on transactional data. For trend analysis, various sample statistics are computed for each time period defined by the time ID variable and `INTERVAL=` option. For seasonal analysis, various sample statistics are computed for each season defined by the `INTERVAL=` or the `SEASONALITY=` option. For example, suppose the transactional data ranges from June 1990 to January 2000, then the trend statistics are computed for every month: June 1990, July 1990, ..., January 2000. The seasonal statistics are computed for each season: January, February, ..., December.

The HPF procedure can be used to forecast time series data as well as transactional data. If the data is transactional, then the procedure must first accumulate the data into a time series before it can be forecast. The procedure uses the following sequential steps to produce forecasts, with the options that control the step listed to the right:

- | | |
|---------------------------------|---|
| 1. Accumulation | <code>ACCUMULATE=</code> option |
| 2. Missing Value Interpretation | <code>SETMISSING=</code> option |
| 3. Diagnostic Tests | <code>INTERMITTENT=</code> and <code>SEASONTEST=</code> options |
| 4. Model Selection | <code>MODEL=</code> , <code>HOLDOUT=</code> , <code>HOLDOUTPCT=</code> , and <code>SELECT=</code> options |
| 5. Transformations | <code>TRANSFORM=</code> option |
| 6. Parameter Estimation | <code>MODEL=</code> option |
| 7. Forecasting | <code>MODEL=</code> and <code>LEAD=</code> options |
| 8. Inverse Transformation | <code>TRANSFORM=</code> and <code>MEDIAN</code> options |
| 9. Statistics of Fit | <code>SELECT=</code> option |
| 10. Summation of Forecasts | <code>LEAD=</code> and <code>STARTSUM=</code> options |

Each of the above steps is described below.

Accumulation

If the `ACCUMULATE=` option is specified, data set observations are accumulated within each time period. The frequency (width of each time interval) is specified by the `INTERVAL=` option. The ID variable contains the time ID values. Each time ID value corresponds to a specific time period. Accumulation is particularly useful when the input data set contains transactional data, whose observations are not spaced with respect to any particular time interval. The accumulated values form the actual time series, which is used in subsequent analyses.

For example, suppose a data set contains the following observations:

```

19MAR1999    10
19MAR1999    30
11MAY1999    50
12MAY1999    20
23MAY1999    20

```


If the INTERVAL=MONTH is specified, all of the above observations fall within three time periods of March 1999, April 1999, and May 1999. The observations are accumulated within each time period as follows:

If the ACCUMULATE=NONE option is specified, an error is generated because the ID variable values are not equally spaced with respect to the specified frequency (MONTH).

If the ACCUMULATE=TOTAL option is specified:

O1MAR1999	40
O1APR1999	.
O1MAY1999	90

If the ACCUMULATE=AVERAGE option is specified:

O1MAR1999	20
O1APR1999	.
O1MAY1999	30

If the ACCUMULATE=MINIMUM option is specified:

O1MAR1999	10
O1APR1999	.
O1MAY1999	20

If the ACCUMULATE=MEDIAN option is specified:

O1MAR1999	20
O1APR1999	.
O1MAY1999	20

If the ACCUMULATE=MAXIMUM option is specified:

O1MAR1999	30
O1APR1999	.
O1MAY1999	50

If the ACCUMULATE=FIRST option is specified:

O1MAR1999	10
O1APR1999	.
O1MAY1999	50

If the ACCUMULATE=LAST option is specified:

```
O1MAR1999    30
O1APR1999    .
O1MAY1999    20
```

If the ACCUMULATE=STDDEV option is specified:

```
O1MAR1999    14.14
O1APR1999    .
O1MAY1999    17.32
```

As can be seen from the above examples, even though the data set observations contained no missing values, the accumulated time series may have missing values.

Missing Value Interpretation

Sometimes missing values should be interpreted as unknown values. The forecasting models used by the HPF procedure can effectively handle missing values (see the “[Missing Value Modeling Issues](#)” section on page 31). But sometimes missing values are known, such as when missing values are created from accumulation and no observations should be interpreted as no (zero) value. In the former case, the SETMISSING= option can be used to interpret how missing values are treated. The SETMISSING=0 option should be used when missing observations are to be treated as no (zero) values. In other cases, missing values should be interpreted as global values, such as minimum or maximum values of the accumulated series. The accumulated and interpreted time series is used in subsequent analyses.

Diagnostic Tests

The INTERMITTENT= option set the thresholds for categorizing a series as intermittent or non-intermittent. The SEASONTEST= option set the thresholds for categorizing a series as seasonal or non-seasonal.

Model Selection

When more than one candidate model is specified, forecasts for each candidate model are compared using the model selection criterion specified by the SELECT= option. The selection criterion is computed using the multistep forecasts in the holdout sample range if the HOLDOUT= or HOLDOUTPCT= options are specified, or the one-step ahead forecasts for the full range of the time series if the HOLDOUT= and HOLDOUTPCT= options are not specified. The candidate model with the best selection criterion is selected to forecast the time series.

Transformations

If the TRANSFORM= option is specified, the time series is transformed prior to model parameter estimation and forecasting. Only strictly positive series can be transformed. An error is generated when the TRANSFORM= option is used with a nonpositive series.

Parameter Estimation

All parameters associated with the model are optimized based on the data with the default parameter restrictions imposed. If the TRANSFORM= option is specified, the transformed time series data are used to estimate the model parameters.

Missing Value Modeling Issues

The treatment of missing values varies with the forecasting model. For the smoothing models, missing values after the start of the series are replaced with one-step-ahead predicted values, and the predicted values are applied to the smoothing equations. See [Chapter 2, “Forecasting Process Details,”](#) for greater detail on how missing values are treated in the smoothing models. For MODEL=IDM, specified missing values are assumed to be periods of no demand.

The treatment of missing values can also be specified by the user with the SETMISSING= option, which changes the missing values prior to modeling.

Even though all of the observed data are nonmissing, using the ACCUMULATE= option can create missing values in the accumulated series.

Forecasting

Once the model parameters are estimated, one-step ahead forecasts are generated for the full range of the actual (optionally transformed) time series data, and multistep forecasts are generated from the end of the observed time series to the future time period after the last observation specified by the LEAD= option. If there are missing values at the end of the time series, the forecast horizon will be greater than that specified by the LEAD= option.

Inverse Transformations

If the TRANSFORM= option is specified, the forecasts of the transformed time series are inverse transformed. By default, the mean (expected) forecasts are generated. If the MEDIAN option is specified, the median forecasts are generated.

Statistics of Fit

The statistics of fit (or goodness-of-fit statistics) are computed by comparing the actual time series data and the generated forecasts. If the TRANSFORM= option was specified, the statistics of fit are based on the inverse transformed forecasts.

Forecast Summation

The multistep forecasts generated by the above steps are summed from the STARTSUM= number to the LEAD= number. For example, if STARTSUM=4 and LEAD=6, the 4-step through 6-step ahead forecasts are summed. The predictions are simply summed. However, the prediction error variance of this sum is computed taking into account the correlation between the individual predictions. The upper and lower confidence limits for the sum of the predictions is then computed based on the prediction error variance of the sum.

The forecast summation is particularly useful when it is desirable to model in one frequency yet the forecast of interest is another frequency. For example, if a time series has a monthly frequency (INTERVAL=MONTH) and you want a forecast for the third and fourth future months, a forecast summation for the third and fourth month can be obtained by specifying STARTSUM=3 and LEAD=4.

Variance-related computations are only computed when no transformation is specified (TRANSFORM=NONE).

Comparison to the Time Series Forecasting System

With the exception of Model Selection, the techniques used in the HPF procedure are identical to the Time Series Forecasting System of SAS/ETS software. For Model Parameter Estimation, the default parameter restrictions are imposed.

Data Set Output

The HPF procedure can create the OUT=, OUTEST=, OUTFOR=, OUTSTAT=, OUTSUM=, OUTSEASON=, and OUTTREND= data sets. In general, these data sets will contain the variables listed in the BY statement. In general, if a forecasting step related to an output data step fails, the values of this step are not recorded or are set to missing in the related output data set, and appropriate error and/or warning messages are recorded in the log.

OUT= Data Set

The OUT= data set contains the variables specified in the BY, ID, and FORECAST statements. If the ID statement is specified, the ID variable values are aligned and extended based on the ALIGN= and INTERVAL= options. The values of the variables specified in the FORECAST statements are accumulated based on the ACCUMULATE= option and missing values are interpreted based on the SETMISSING= option. If the REPLACEMISSING option is specified, missing values embedded missing values are replaced by the one step-ahead forecasts.

These variable values are then extrapolated based on their forecasts or extended with missing values when the MODEL=NONE option is specified. If USE=LOWER is specified, the variable is extrapolated with the lower confidence limits; if USE=UPPER, the variable is extrapolated using the upper confidence limits; otherwise, the variable values are extrapolated with the predicted values. If the TRANSFORM= option is specified, the predicted values will contain either mean or median forecasts depending on whether or not the MEDIAN option is specified.

If any of the forecasting steps fail for particular variable, the variable values are extended by missing values.

OUTEST= Data Set

The OUTEST= data set contains the variables specified in the BY statement as well as the variables listed below. For variables listed in FORECAST statements where the option MODEL=NONE is specified, no observations are recorded for these variables. For variables listed in FORECAST statements where the option MODEL=NONE is not specified, the following variables contain observations related to the parameter estimation step:

NAME	Variable name
MODEL	Forecasting Model
TRANSFORM	Transformation
PARM	Parameter Name
EST	Parameter Estimate
STDERR	Standard Errors
TVALUE	<i>t</i> -Values
PVALUE	Probability Values

If the parameter estimation step fails for a particular variable, no observations are recorded.

OUTFOR= Data Set

The OUTFOR= data set contains the variables specified in the BY statement as well as the variables listed below. For variables listed in FORECAST statements where the option MODEL=NONE is specified, no observations are recorded for these variables. For variables listed in FORECAST statements where the option MODEL=NONE is not specified, the following variables contain observations related to the forecasting step:

NAME	Variable name
TIMEID	Time ID values
PREDICT	Predicted Values
STD	Prediction Standard Errors

LOWER	Lower Confidence Limits
UPPER	Upper Confidence Limits
ERROR	Prediction Errors

If the forecasting step fails for a particular variable, no observations are recorded. If the TRANSFORM= option is specified, the values in the variables listed above are the inverse transform forecasts. If the MEDIAN option is specified, the median forecasts are stored; otherwise, the mean forecasts are stored.

OUTSTAT= Data Set

The OUTSTAT= data set contains the variables specified in the BY statement as well as the variables listed below. For variables listed in FORECAST statements where the option MODEL=NONE is specified, no observations are recorded for these variables. For variables listed in FORECAST statements where the option MODEL=NONE is not specified, the following variables contain observations related to the statistics of fit step:

NAME	Variable name
SSE	Sum of Square Error
MSE	Mean Square Error
UMSE	Unbiased Mean Square Error
RMSE	Root Mean Square Error
URMSE	Unbiased Root Mean Square Error
MAPE	Mean Absolute Percent Error
MAE	Mean Absolute Error
RSQUARE	R-Square
ADJRSQ	Adjusted R-Square
AADJRSQ	Amemiya's Adjusted R-Square
RWRSQ	Random Walk R-Square
AIC	Akaike Information Criterion
SBC	Schwarz Bayesian Information Criterion
APC	Amemiya's Prediction Criterion
MAXERR	Maximum Error
MINERR	Minimum Error
MINPE	Maximum Percent Error
MAXPE	Minimum Percent Error
ME	Mean Error
MPE	Mean Percent Error

If the statistics of fit step fails for particular variable, no observations are recorded. If the TRANSFORM= option is specified, the values in the variables listed above are computed based on the inverse transform forecasts. If the MEDIAN option is specified, the median forecasts are the basis; otherwise, the mean forecasts are the basis.

OUTSUM= Data Set

The OUTSUM= data set contains the variables specified in the BY statement as well as the variables listed below. The OUTSUM= data set records the summary statistics for each variable specified in a FORECAST statement. For variables listed in FORECAST statements where the option MODEL=NONE is specified, the values related to forecasts are set to missing. For variables listed in FORECAST statements where the option MODEL=NONE is not specified, the forecast values are set based on the USE= option.

Variables related to summary statistics are based on the ACCUMULATE= and SETMISSING= options:

NAME	Variable name
STATUS	Forecasting Status. Nonzero values imply that no forecast was generated for the series.
NOBS	Number of Observations
N	Number of Nonmissing Observations
NMISS	Number of Missing Observations
MIN	Minimum Value
MAX	Maximum Value
MEAN	Mean Value
STDDEV	Standard Deviation

Variables related to forecast summation are based on the LEAD= and STARTSUM= options:

PREDICT	Forecast Summation Predicted Values
STD	Forecast Summation Prediction Standard Errors
LOWER	Forecast Summation Lower Confidence Limits
UPPER	Forecast Summation Upper Confidence Limits

Variance-related computations are only computed when no transformation is specified (TRANSFORM=NONE).

Variables related to multistep forecast based on the LEAD= and USE= options:

_LEAD n _ Multistep Forecast (n ranges from one to LEAD=number). If USE=LOWER, this variable will contain the lower confidence limits; if USE=UPPER, this variable will contain the upper confidence limits; otherwise, this variable will contain the predicted values.

If the forecast step fails for a particular variable, the variables related to forecasting are set to missing. The OUTSUM= data set contains both a summary of the (accumulated) time series and optionally its forecasts for all series.

OUTSEASON= Data Set

The OUTSEASON= data set contains the variables specified in the BY statement as well as the variables listed below. The OUTSEASON= data set records the seasonal statistics for each variable specified in a FORECAST statement.

Variables related to seasonal statistics are based on the INTERVAL= or SEASONALITY= options:

NAME	Variable name
TIMEID	Time ID values
SEASON	Seasonal index
NOBS	Number of Observations
N	Number of Nonmissing Observations
NMISS	Number of Missing Observations
MIN	Minimum Value
MAX	Maximum Value
RANGE	Range Value
SUM	Summation Value
MEAN	Mean Value
STDDEV	Standard Deviation
CSS	Corrected Sum of Squares
USS	Uncorrected Sum of Squares
MEDIAN	Median Value

The above statistics are computed for each season.

OUTTREND= Data Set

The OUTTREND= data set contains the variables specified in the BY statement as well as the variables listed below. The OUTTREND= data set records the trend statistics for each variable specified in a FORECAST statement.

Variables related to trend statistics are based on the INTERVAL= and SEASONALITY= options:

NAME	Variable name
TIMEID	Time ID values
SEASON	Seasonal index
NOBS	Number of Observations
N	Number of Nonmissing Observations
NMISS	Number of Missing Observations
MIN	Minimum Value
MAX	Maximum Value
RANGE	Range Value
SUM	Summation Value
MEAN	Mean Value
STDDEV	Standard Deviation
CSS	Corrected Sum of Squares
USS	Uncorrected Sum of Squares
MEDIAN	Median Value

The above statistics are computed for each time period.

Printed Output

The HPF procedure optionally produces printed output for these results utilizing the Output Delivery System (ODS). By default, the procedure produces no printed output. All output is controlled by the PRINT= and PRINTDETAILS options associated with the PROC HPF statement. In general, if a forecasting step related to printed output fails, the values of this step are not printed and appropriate error and/or warning messages are recorded in the log. The printed output is similar to the output data set and these similarities are described below.

PRINT=SUMMARY

prints the summary statistics and forecast summaries similar to the OUTSUM= data set.

PRINT=ESTIMATES

prints the parameter estimates similar to the OUTEST= data set.

PRINT=FORECASTS

prints the forecasts similar to the OUTFOR= data set. For MODEL=IDM, a table containing demand series is also printed.

If the MODEL=IDM option is specified, the demand series predictions table is also printed. This table is based on the demand index (when demands occurred).

PRINT=PERFORMANCE

prints the performance statistics.

PRINT=PERFORMANCESUMMARY

prints the performance summary for each BY group.

PRINT=PERFORMANCEOVERALL

prints the performance summary for all BY groups.

PRINT=STATES

prints the backcast, initial, and final smoothed states.

PRINT=SEASONS

prints the seasonal statistics similar to the OUTSEASON= data set.

PRINT=STATISTICS

prints the statistics of fit similar to the OUTSTAT= data set.

PRINT=TRENDS

Prints the trend statistics similar to the OUTTREND= data set.

PRINTDETAILS

The PRINTDETAILS option is the opposite of the NOOUTALL option.

Specifically, if PRINT=FORECASTS and the PRINTDETAILS options are specified, the one-step ahead forecasts, throughout the range of the data, are printed as well as the information related to a specific forecasting model such as the smoothing states. If the PRINTDETAILS option is not specified, only the multistep forecasts are printed.

ODS Table Names

The table below relates the PRINT= options to ODS tables:

Table 1.1. ODS Tables Produced in PROC HPF

ODS Table Name	Description	Option
ODS Tables Created by the PRINT=SUMMARY option		
DescStats	Descriptive Statistics	
DemandSummary	Demand Summary	MODEL=IDM option only
ForecastSummary	Forecast Summary	
ForecastSummmation	Forecast Summation	
ODS Tables Created by the PRINT=ESTIMATES option		
ModelSelection	Model Selection	
ParameterEstimates	Parameter Estimates	
ODS Tables Created by the PRINT=FORECASTS option		
Forecasts	Forecast	
Demands	Demands	MODEL=IDM option only
ODS Tables Created by the PRINT=PERFORMANCE option		
Performance	Performance Statistics	
ODS Tables Created by the PRINT=PERFORMANCESUMMARY option		
PerformanceSummary	Performance Summary	
ODS Tables Created by the PRINT=PERFORMANCEOVERALL option		
PerformanceSummary	Performance Overall	
ODS Tables Created by the PRINT=SEASONS option		
SeasonStatistics	Seasonal Statistics	
ODS Tables Created by the PRINT=STATES option		
SmoothedStates	Smoothed States	
DemandStates	Demand States	MODEL=IDM option only
ODS Tables Created by the PRINT=STATISTICS option		
FitStatistics	Statistics of Fit	
ODS Tables Created by the PRINT=TRENDS option		
TrendStatistics	Trend Statistics	

The ODS table ForecastSummary is related to all time series within a BY group. The other tables are related to a single series within a BY group.

ODS Graphics (Experimental)

This section describes the use of ODS for creating graphics with the HPF procedure. These graphics are experimental in this release, meaning that both the graphical results and the syntax for specifying them are subject to change in a future release.

To request these graphs, you must specify the ODS GRAPHICS statement. In addition, you can specify the **PLOT=** option in the HPF statement according to the

following syntax. For more information on the ODS GRAPHICS statement, refer to Chapter 9, “Statistical Graphics Using ODS” (*SAS/ETS User’s Guide*).

PLOT= option | (options)

specifies the graphical output desired. By default, the HPF procedure produces no graphical output. The following printing options are available:

ERRORS	plots prediction error time series graphics.
ACF	plots prediction error autocorrelation function graphics.
PACF	plots prediction error partial autocorrelation function graphics.
IACF	plots prediction error inverse autocorrelation function graphics.
WN	plots white noise graphics.
MODELS	plots model graphics.
FORECASTS	plots forecast graphics.
MODELFORECASTSONLY	plots forecast graphics with confidence limits in the data range.
FORECASTSONLY	plots the forecast in the forecast horizon only.
LEVELS	plots smoothed level component graphics.
SEASONS	plots smoothed seasonal component graphics.
TRENDS	plots smoothed trend (slope) component graphics.
ALL	Same as specifying all of the above PLOT= options.

For example, PLOT=FORECASTS plots the forecasts for each series. The PLOT= option produces printed output for these results utilizing the Output Delivery System (ODS). The PLOT= statement is experimental for this release of SAS.

ODS Graph Names

PROC HPF assigns a name to each graph it creates using ODS. You can use these names to reference the graphs when using ODS. The names are listed in [Table 1.2](#).

To request these graphs, you must specify the ODS GRAPHICS statement. In addition, you can specify the **PLOT=** option in the HPF statement. For more information on the ODS GRAPHICS statement, refer to Chapter 9, “Statistical Graphics Using ODS” (*SAS/ETS User’s Guide*).

Table 1.2. ODS Graphics Produced by PROC HPF

ODS Graph Name	Plot Description	Statement	PLOT= Option
DemandErrorsPlot	Average Demand Errors	PROC HPF	PLOT=ERRORS
DemandForecastsPlot	Average Demand Forecasts	PROC HPF	PLOT=FORECASTS
DemandIntervalHistogram	Demand Interval Histogram	PROC HPF	PLOT=MODELS
DemandIntervalPlot	Demand Interval Forecast Plot	PROC HPF	PLOT=MODELS

Table 1.2. (continued)

ODS Graph Name	Plot Description	Statement	Option
DemandSizeHistogram	Demand Size Histogram	PROC HPF	PLOT=MODELS
DemandSizePlot	Demand Size Forecast Plot	PROC HPF	PLOT=MODELS
ErrorACFNORMPlot	Standardized autocorrelation of Prediction Errors	PROC HPF	PLOT=ACF
ErrorACFPlot	Autocorrelation of Prediction Errors	PROC HPF	PLOT=ACF
ErrorHistogram	Prediction Error Histogram	PROC HPF	PLOT=ERRORS
ErrorIACFNORMPlot	Standardized inverse autocorrelation of Prediction Errors	PROC HPF	PLOT=IACF
ErrorIACFPlot	Inverse autocorrelation of Prediction Errors	PROC HPF	PLOT=IACF
ErrorPACFNORMPlot	Standardized partial autocorrelation of Prediction Errors	PROC HPF	PLOT=PACF
ErrorPACFPlot	Partial autocorrelation of Prediction Errors	PROC HPF	PLOT=PACF
ErrorPlot	Plot of Prediction Errors	PROC HPF	PLOT=ERRORS
ErrorWhiteNoiseLogProbPlot	White noise log probability plot of Prediction Errors	PROC HPF	PLOT=WN
ErrorWhiteNoiseProbPlot	White noise probability plot of Prediction Errors	PROC HPF	PLOT=WN
ForecastsOnlyPlot	Forecasts Only Plot	PROC HPF	PLOT=FORECASTONLY
ForecastsPlot	Forecasts Plot	PROC HPF	PLOT=FORECAST
LevelStatePlot	Smoothed Level State Plot	PROC HPF	PLOT=LEVELS
ModelForecastsPlot	Model and Forecasts Plot	PROC HPF	PLOT=MODELFORECAST
ModelPlot	Model Plot	PROC HPF	PLOT=MODELS
SeasonStatePlot	Smoothed Season State Plot	PROC HPF	PLOT=SEASONS
StockingAveragePlot	Stocking Average Plot	PROC HPF	PLOT=FORECASTS
StockingLevelPlot	Stocking Level Plot	PROC HPF	PLOT=FORECASTS
TrendStatePlot	Smoothed Trend State Plot	PROC HPF	PLOT=TRENDS

Examples

Example 1.1. Automatic Forecasting of Time Series Data

This example illustrates how the HPF procedure can be used for the automatic forecasting of time series data. Retail sales data is used for this illustration.

The following DATA step creates a data set from data recorded monthly at numerous points of sales. The data set, **SALES**, will contain a variable **DATE** that represents time and a variable for each sales item. Each value of the **DATE** variable is recorded in ascending order and the values of each of the other variables represent a single time series:

```
data sales;
    format date date9.;
    input date date9. shoes socks laces dresses coats shirts ties
        belts hats blouses;
    datalines;
    ... data lines omitted ...
;
run;
```

The following HPF procedure statements automatically forecast each of the monthly time series.

```
proc hpf data=sales out=nextyear;
    id date interval=month;
    forecast _ALL_;
run;
```

The above statements automatically select the best fitting model and generate forecasts for every numeric variable in the input data set (**SALES**) for the next twelve months, and stores these forecasts in the output data set (**NEXTYEAR**).

The following GPLOT procedure statements plot the forecasts related to shoe sales:

```
title1 "Shoe Department Sales";
axis2 label=(a=-90 r=90 "items" );
symbol1 v = dot      i = join l = 1;
symbol2 v = star      i = join l = 2;
symbol3 v = circle i = join l = 2;

proc gplot data=nextyear;
    plot shoes * date = 1
        socks * date = 2
        laces * date = 3 / overlay
        haxis= '01JAN1994'd to '01DEC2000'd by year
        href= '01JAN1999'd
        vaxis=axis2;
run;
```

The GPLOT procedure results are shown in [Figure 1.1](#). The historical data is shown left the horizontal reference line and the forecasts for the next twelve monthly periods is shown to the right.




Figure 1.1. Retail Sales Forecast Plots

The following HPF procedure statements are identical to the statements above with the exception that the PRINT=FORECASTS option is specified:

```
proc hpf data=sales out=nextyear print=forecasts;
  id date interval=month;
  forecast _ALL_;
run;
```

In addition to automatically forecasting each of the monthly time series, the above statements print the forecasts using the Output Delivery System (ODS), which is partially shown in [Output 1.1.1](#). This output shows the predictions, prediction standard errors and the upper and lower confidence limits for the next twelve monthly periods.

Output 1.1.1. Forecast Tables

Forecasts for Variable shoes					
Obs	Time	Forecasts	Standard Error	95% Confidence Limits	
62	FEB1999	7548.0041	607.5238	6357.2792	8738.7289
63	MAR1999	7177.1472	699.4400	5806.2701	8548.0244
64	APR1999	5497.5595	780.7609	3967.2964	7027.8227
65	MAY1999	4838.2001	854.5169	3163.3778	6513.0224
66	JUN1999	6728.4521	922.5244	4920.3375	8536.5668
67	JUL1999	6786.1094	985.9738	4853.6362	8718.5826
68	AUG1999	5853.9650	1045.6953	3804.4399	7903.4900
69	SEP1999	7517.0144	1102.2949	5356.5561	9677.4728
70	OCT1999	7100.2489	1156.2315	4834.0769	9366.4210
71	NOV1999	7224.6449	1207.8618	4857.2793	9592.0106
72	DEC1999	6357.1556	1257.4701	3892.5594	8821.7518
73	JAN2000	6492.2657	1305.2871	3933.9500	9050.5815

Example 1.2. Automatic Forecasting of Transactional Data

This example illustrates how the HPF procedure can be used to automatically forecast transactional data. Internet data is used for this illustration.

The following DATA step creates a data set from data recorded at several Internet Web sites. The data set, WEBSITES, will contain a variable TIME that represents time and the variables ENGINE, BOATS, CARS, and PLANES that represent Internet Web site data. Each value of the TIME variable is recorded in ascending order, and the values of each of the other variables represent a transactional data series.

```
data websites;
  format time datetime.;
  input time datetime. boats cars planes engines;
  datalines;
  ... data lines omitted ...
;
run;
```

The following HPF procedure statements automatically forecast each of the transactional data series:

```
proc hpf data=websites out=nextweek lead=7;
  id time interval=dtday accumulate=total;
  forecast boats cars planes;
run;
```

The above statements accumulate the data into a daily time series and automatically generate forecasts for the BOATS, CARS, and PLANES variables in the input data set (WEBSITES) for the next week and stores the forecasts in the output data set (NEXTWEEK).

The following GPLOT procedure statements plot the forecasts related to the Internet data:


```

title1 "Website Data";
axis2 label=(a=-90 r=90 "Websites" );
symbol1 v = dot      i = join l = 1;
symbol2 v = star     i = join l = 2;
symbol3 v = circle   i = join l = 2;

proc gplot data=nextweek;
  plot boats * time = 1
        cars * time = 2
        planes * time = 3 / overlay
  haxis= '13MAR2000:00:00:00'dt to
        '18APR2000:00:00:00'dt by dtweek
  href= '11APR2000:00:00:00'dt
  vaxis=axis2;
run;

```

The GPLOT procedure results are shown in [Figure 1.2](#). The historical data is shown to the left of the horizontal reference line and the forecasts for the next twelve monthly periods are shown to the right.




Figure 1.2. Internet Data Forecast Plots

Example 1.3. Specifying the Forecasting Model

In the previous example, the HPF procedure was used to automatically select the appropriate forecasting model using the root mean square error (RMSE) as the default selection criterion. This example illustrates how the HPF procedure can be used to more narrowly specify the possible candidate models. Internet data from the previous example are used for this illustration.

In this example, we will forecast the **BOATS** variable using the best seasonal forecasting model (**BESTS**) that minimizes the mean absolute percent error (**MAPE**), the **CARS** variable using the best nonseasonal forecasting model (**BESTN**) that minimizes the mean square error (**MSE**) using holdout sample analysis, and the **PLANES** variable using Log Winters (additive). The following HPF procedure statements forecast each of the transactional data series based on these requirements:

```
proc hpf data=websites out=nextweek lead=7;
  id time interval=dtday accumulate=total;
  forecast boats / model=bests select=mape;
  forecast cars / model=bestn select=mse holdout=5;
  forecast planes / model=addwinters transform=log;
run;
```

Example 1.4. Extending the Independent Variables for Multivariate Forecasts

In the previous example, the HPF procedure was used to forecast several transactional series variables using univariate models. This example illustrates how the HPF procedure can be used to extend the independent variables associated with a multiple regression forecasting problem. Specifically, PROC HPF is used to extend the independent variables for use in forecasting a regression model.

In this example, we will accumulate and forecast the **BOATS**, **CARS**, and **PLANES** variables as illustrated in the previous example. In addition, we will accumulate the **ENGINES** variable to form a time series that is then extended with missing values within the forecast horizon with the specification of **MODEL=NONE**.

```
proc hpf data=websites out=nextweek lead=7;
  id time interval=dtday accumulate=total;
  forecast engines / model=none;
  forecast boats / model=bests select=mape;
  forecast cars / model=bestn select=mse holdout=5;
  forecast planes / model=winters transform=log;
run;
```

The following AUTOREG procedure statements are used to forecast the **ENGINES** variable by regressing on the independent variables (**BOATS**, **CARS**, and **PLANES**).

```
proc autoreg data= nextweek;
  model engines = boats cars planes / noprint;
  output out=enginehits p=predicted;
run;
```

The output data set (**NEXTWEEK**) of the PROC HPF statement is used as an input data set for the PROC AUTOREG statement. The output data set of PROC AUTOREG contains the forecast of the variable **ENGINES** based on the regression model with the variables **BOATS**, **CARS**, and **PLANES** as regressors. See the AUTOREG procedure for details on autoregression models.

The following GPLOT procedure statements plot the forecasts related to the ENGINES variable:

```
proc gplot data=enginehits;
  plot boats * time = 1
       cars * time = 2
       planes * time = 3
       predicted * time = 4 / overlay
  haxis= '13MAR2000:00:00:00'dt to
         '18APR2000:00:00:00'dt by dtweek
  href= '11APR2000:00:00:00'dt
  vaxis=axis2;
run;
```

The GPLOT procedure results are shown in Figure 1.3. The historical data is shown left the horizontal reference line and the forecasts for the next four weekly periods is shown to the right.

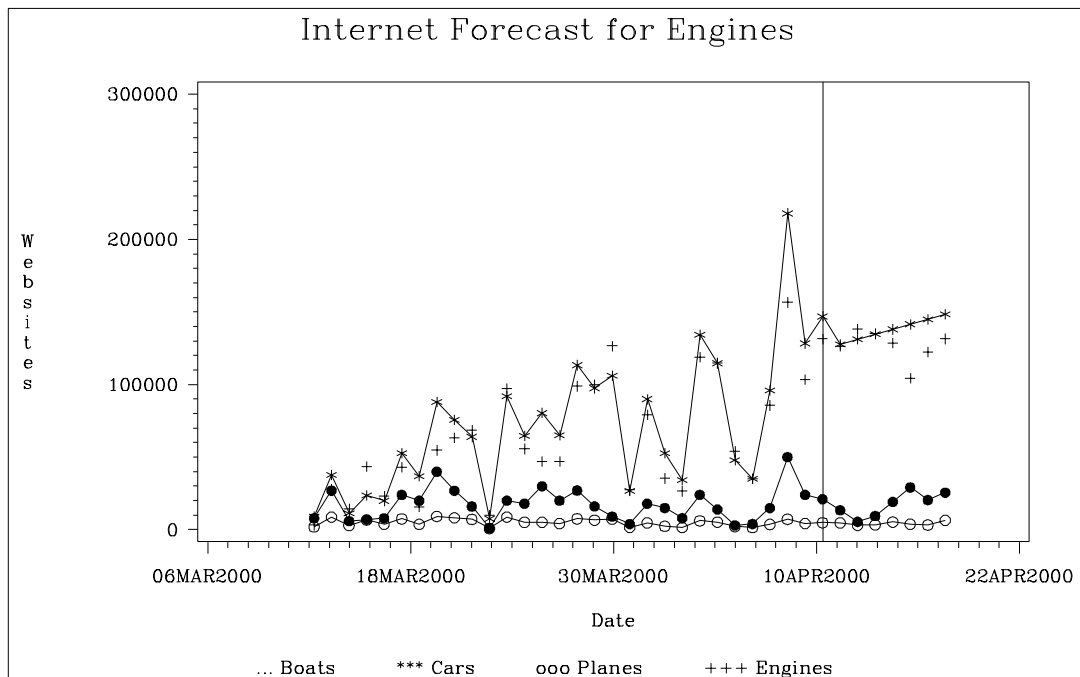


Figure 1.3. Internet Data Forecast Plots

Example 1.5. Forecasting Intermittent Time Series Data

This example illustrates how the HPF procedure can be used to forecast intermittent time series data. Inventory demand is used for this illustration.

The following DATA step creates a data set from inventory data recorded at no particular frequency. The data set, INVENTORY, will contain a variable DATE that represents time and the demand variables (TIRES, HUBCAPS, and LUGBOLTS), which represent inventory items. Each value of the DATE variable is recorded in ascending order, and the values of each of the other variables represent a transactional data series.

```

data inventory;
  format date date9.;
  input date date9. tires hubcaps lugbolts;
  datalines;
  ... data lines omitted ...
  ;
run;

```

The following HPF procedure statements forecast each of the transactional data series using an intermittent demand model:

```

proc hpf data=inventory out=nextmonth lead=4 print=forecasts;
  id date interval=week accumulate=total;
  forecast tires hubcaps lugbolts / model=idm;
run;

```

The above statements accumulate the data into a weekly time series, and generate forecasts for the **TIRES**, **HUBCAPS**, and **LUGBOLTS** variables in the input data set (**INVENTORY**) for the four weekly periods, and store the forecasts in the output data set (**NEXTMONTH**). The **PRINT=FORECAST** option produces the results partially shown in [Output 1.5.1](#). The first table records the demand series and predictions. The second table represents forecasts or recommended stocking levels.

Output 1.5.1. Forecast Tables

Demands for Variable tires							
Index	Demand	Time	Demand Intervals		Demand Size		
	Actual		Predicted	Actual	Predicted		
1	Sun, 31 Aug 1997		14	14.0000	6.00000	4.42151	
2	Sun, 26 Oct 1997		8	14.0000	4.00000	4.74216	
3	Sun, 1 Mar 1998		18	12.8920	2.00000	4.59140	
4	Sun, 26 Apr 1998		8	13.8353	2.00000	4.06498	
5	Sun, 31 May 1998		5	12.7577	2.00000	3.64550	
6	Sun, 27 Sep 1998		17	11.3251	6.00000	3.31123	
7	Sun, 3 Jan 1999		14	12.3731	.	3.85743	
Error = (Interval Actual)*(Predict) - (Size Actual)							
Demands for Variable tires							
Index	Demand	Time	Estimate of Mean Demand per Period				
			Predict	Std	95% Confidence Limits		Error
1	Sun, 31 Aug 1997		0.31582	0.11634	0.08780	0.54384	-1.578494
2	Sun, 26 Oct 1997		0.33873	0.11634	0.11071	0.56674	-1.290193
3	Sun, 1 Mar 1998		0.35614	0.11634	0.12813	0.58416	4.410578
4	Sun, 26 Apr 1998		0.29381	0.11634	0.06579	0.52183	0.350503
5	Sun, 31 May 1998		0.28575	0.11634	0.05773	0.51377	-0.571254
6	Sun, 27 Sep 1998		0.29238	0.11634	0.06436	0.52040	-1.029543
7	Sun, 3 Jan 1999		0.27553	0.11634	0.04751	0.50355	.
Error = (Interval Actual)*(Predict) - (Size Actual)							
Forecasts for Variable tires							
Obs	Time	Forecasts	Standard				
			Error	95% Confidence Limits			
84	Sun, 3 Jan 1999	0.27553	0.11634	0.04751	0.50355		
85	Sun, 10 Jan 1999	0.27553	0.11634	0.04751	0.50355		
86	Sun, 17 Jan 1999	0.27553	0.11634	0.04751	0.50355		
87	Sun, 24 Jan 1999	0.27553	0.11634	0.04751	0.50355		

Example 1.6. Illustration of ODS Graphics (Experimental)

This example illustrates the use of experimental ODS graphics.


The following statements utilize the `SASHELP.AIR` data set to automatically forecast the time series of international airline travel.

The graphical displays are requested by specifying the experimental ODS `GRAPHICS` statement and the experimental `PLOT=` option in the `PROC HPF` statement. In this case, all plots are requested. [Output 1.6.1](#) through [Output 1.6.4](#) show a selection of the plots created.

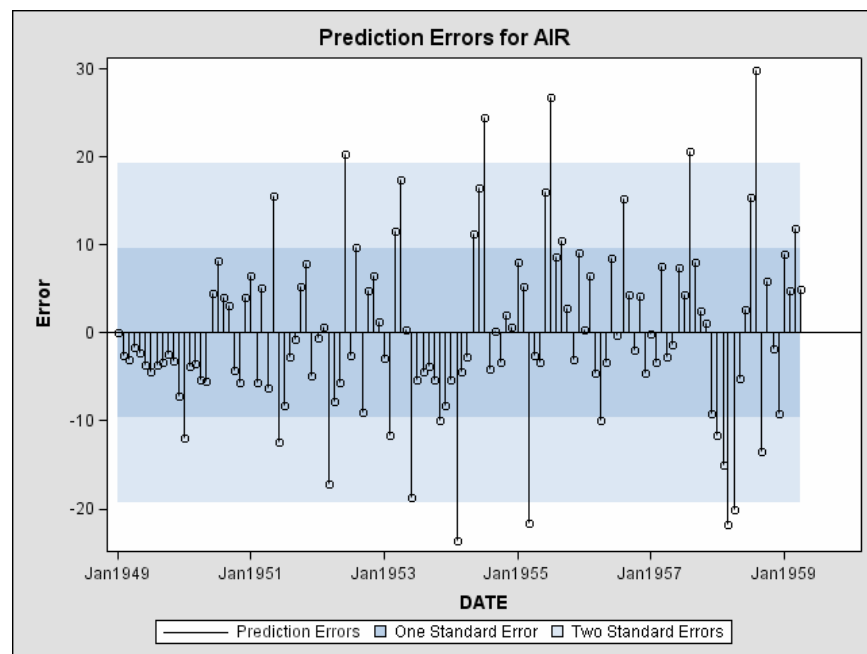
For general information about ODS graphics, refer to Chapter 9, “Statistical Graphics Using ODS” (*SAS/ETS User’s Guide*). For specific information about the graphics available in the HPF procedure, see the “[ODS Graphics](#)” section on page 39.

```
ods html;  
ods graphics on;  
  
proc hpf data=sashelp.air out=_null_ lead=20 back=20 print=all plot=all;  
  id date interval=month;  
  forecast air / model=best transform=auto select=mape;  
run;  
  
ods graphics off;  
ods html close;
```


Output 1.6.1. Smoothed Trend Plot (Experimental)




Output 1.6.2. Prediction Error Plot (Experimental)



Output 1.6.3. Prediction Error Standardized ACF Plot (Experimental)



Output 1.6.4. Forecast Plot (Experimental)

References

Pyle, D. (1999), *Data Preparation for Data Mining*, San Francisco: Morgan Kaufman Publishers, Inc.

Chapter 2

Forecasting Process Details

Chapter Contents

FORECASTING PROCESS SUMMARY	55
SMOOTHING MODELS	55
Smoothing Model Calculations	56
Missing Values	57
Predictions and Prediction Errors	57
Smoothing Weights	58
Equations for the Smoothing Models	59
INTERMITTENT MODELS	68
Intermittent Time Series	68
Intermittent Series Decomposition and Analysis	68
Croston's Method	69
Average Demand Method	70
Time Indexed Holdout Samples versus Demand Indexed Holdout Samples	71
Automatic Intermittent Demand Model Selection	72
SERIES TRANSFORMATIONS	72
SERIES DIAGNOSTIC TESTS	74
STATISTICS OF FIT	75
REFERENCES	77

Chapter 2

Forecasting Process Details

This chapter provides computational details on several aspects of the SAS High Performance Forecasting.

Forecasting Process Summary

This section summarizes the forecasting process.

Parameter Estimation

Computational details for the forecasting models are provided in the following sections "Smoothing Models" and "Intermittent Models." The results of the parameter estimation process are printed in the Parameter Estimates table or stored in the OUTEST= data set.

Model Evaluation

Model evaluation is based on the one-step-ahead prediction errors for observations within the period of evaluation. The one-step-ahead predictions are generated from the model specification and parameter estimates. The predictions are inverse transformed (median or mean) and adjustments are removed. The prediction errors (the difference of the dependent series and the predictions) are used to compute the statistics of fit, which are described in section "Statistics of Fit." The results generated by the evaluation process are printed in the Statistics of Fit table or stored in the OUTSTAT= data set.

Forecasting

The forecasting process is similar to the model evaluation process described in the preceding section, except that k -step-ahead predictions are made from the end of the data through the specified forecast horizon, and prediction standard errors and confidence limits are calculated. The forecasts and confidence limits are printed in the Forecast table or stored in the OUTFOR= data set.

Smoothing Models

This section details the computations performed for the exponential smoothing and Winters method forecasting models.

Smoothing Model Calculations

The descriptions and properties of various smoothing methods can be found in Gardner (1985), Chatfield (1978), and Bowerman and O'Connell (1979). The following section summarizes the smoothing model computations.

Given a time series $\{Y_t : 1 \leq t \leq n\}$, the underlying model assumed by the smoothing models has the following (additive seasonal) form:

$$Y_t = \mu_t + \beta_t t + s_p(t) + \epsilon_t$$

where

μ_t	represents the time-varying mean term.
β_t	represents the time-varying slope.
$s_p(t)$	represents the time-varying seasonal contribution for one of the p seasons
ϵ_t	are disturbances.

For smoothing models without trend terms, $\beta_t = 0$; and for smoothing models without seasonal terms, $s_p(t) = 0$. Each smoothing model is described in the following sections.

At each time t , the smoothing models estimate the time-varying components described above with the *smoothing state*. After initialization, the smoothing state is updated for each observation using the *smoothing equations*. The smoothing state at the last nonmissing observation is used for predictions.

Smoothing State and Smoothing Equations

Depending on the smoothing model, the *smoothing state* at time t will consist of the following:

- L_t is a smoothed level that estimates μ_t .
- T_t is a smoothed trend that estimates β_t .
- $S_{t-j}, j = 0, \dots, p-1$, are seasonal factors that estimate $s_p(t)$.

The smoothing process starts with an initial estimate of the smoothing state, which is subsequently updated for each observation using the *smoothing equations*.

The smoothing equations determine how the smoothing state changes as time progresses. Knowledge of the smoothing state at time $t-1$ and that of the time-series value at time t uniquely determine the smoothing state at time t . The *smoothing weights* determine the contribution of the previous smoothing state to the current smoothing state. The smoothing equations for each smoothing model are listed in the following sections.

Smoothing State Initialization

Given a time series $\{Y_t : 1 \leq t \leq n\}$, the smoothing process first computes the smoothing state for time $t = 1$. However, this computation requires an initial estimate of the smoothing state at time $t = 0$, even though no data exists at or before time $t = 0$.

An appropriate choice for the initial smoothing state is made by backcasting from time $t = n$ to $t = 1$ to obtain a prediction at $t = 0$. The initialization for the backcast is obtained by regression with constant and linear terms and seasonal dummies (additive or multiplicative) as appropriate for the smoothing model. For models with linear or seasonal terms, the estimates obtained by the regression are used for initial smoothed trend and seasonal factors; however, the initial smoothed level for backcasting is always set to the last observation, Y_n .

The smoothing state at time $t = 0$ obtained from the backcast is used to initialize the smoothing process from time $t = 1$ to $t = n$ (refer to Chatfield and Yar 1988).

For models with seasonal terms, the smoothing state is normalized so that the seasonal factors S_{t-j} for $j = 0, \dots, p - 1$ sum to zero for models that assume additive seasonality and average to one for models (such as Winters method) that assume multiplicative seasonality.

Missing Values

When a missing value is encountered at time t , the smoothed values are updated using the *error-correction form* of the smoothing equations with the one-step-ahead prediction error, e_t , set to zero. The missing value is estimated using the one-step-ahead prediction at time $t - 1$, that is $\hat{Y}_{t-1}(1)$ (refer to Aldrin 1989). The error-correction forms of each of the smoothing models are listed in the following sections.

Predictions and Prediction Errors

Predictions are made based on the last known smoothing state. Predictions made at time t for k steps ahead are denoted $\hat{Y}_t(k)$ and the associated prediction errors are denoted $e_t(k) = Y_{t+k} - \hat{Y}_t(k)$. The *prediction equation* for each smoothing model is listed in the following sections.

The *one-step-ahead predictions* refer to predictions made at time $t - 1$ for one time unit into the future, that is, $\hat{Y}_{t-1}(1)$, and the *one-step-ahead prediction errors* are more simply denoted $e_t = e_{t-1}(1) = Y_t - \hat{Y}_{t-1}(1)$. The one-step-ahead prediction errors are also the model residuals, and the sum of squares of the one-step-ahead prediction errors is the objective function used in smoothing weight optimization.

The *variance of the prediction errors* are used to calculate the confidence limits (refer to Sweet 1985, McKenzie 1986, Yar and Chatfield 1990, and Chatfield and Yar 1991). The equations for the variance of the prediction errors for each smoothing model are listed in the following sections.

Note: $\text{var}(\epsilon_t)$ is estimated by the mean square of the one-step-ahead prediction errors.

Smoothing Weights

Depending on the smoothing model, the smoothing weights consist of the following:

α	is a level smoothing weight.
γ	is a trend smoothing weight.
δ	is a seasonal smoothing weight.
ϕ	is a trend damping weight.

Larger smoothing weights (less damping) permit the more recent data to have a greater influence on the predictions. Smaller smoothing weights (more damping) give less weight to recent data.

Specifying the Smoothing Weights

Typically the smoothing weights are chosen to be from zero to one. (This is intuitive because the weights associated with the past smoothing state and the value of current observation would normally sum to one.) However, each smoothing model (except Winters Method – Multiplicative Version) has an ARIMA equivalent. Weights chosen to be within the ARIMA additive-invertible region will guarantee stable predictions (refer to Archibald 1990 and Gardner 1985). The ARIMA equivalent and the additive-invertible region for each smoothing model are listed in the following sections.

Optimizing the Smoothing Weights

Smoothing weights are determined so as to minimize the sum of squared one-step-ahead prediction errors. The optimization is initialized by choosing from a predetermined grid the initial smoothing weights that result in the smallest sum of squared, one-step-ahead prediction errors. The optimization process is highly dependent on this initialization. It is possible that the optimization process will fail due to the inability to obtain stable initial values for the smoothing weights (refer to Greene 1993 and Judge et al 1980), and it is possible for the optimization to result in a local minima.

The optimization process can result in weights to be chosen outside both the zero-to-one range and the ARIMA additive-invertible region. By restricting weight optimization to additive-invertible region, you can obtain a local minimum with stable predictions. Likewise, weight optimization can be restricted to the zero-to-one range or other ranges.

Standard Errors

The standard errors associated with the smoothing weights are calculated from the Hessian matrix of the sum of squared, one-step-ahead prediction errors with respect to the smoothing weights used in the optimization process.

Weights Near Zero or One

Sometimes the optimization process results in weights near zero or one.

For Simple or Double (Brown) Exponential Smoothing, a level weight near zero implies that simple differencing of the time series may be appropriate.

For Linear (Holt) Exponential Smoothing, a level weight near zero implies that the smoothed trend is constant and that an ARIMA model with deterministic trend may be a more appropriate model.

For Damped-Trend Linear Exponential Smoothing, a damping weight near one implies that Linear (Holt) Exponential Smoothing may be a more appropriate model.

For Winters Method and Seasonal Exponential Smoothing, a seasonal weight near one implies that a nonseasonal model may be more appropriate and a seasonal weight near zero implies that deterministic seasonal factors may be present.

Equations for the Smoothing Models

Simple Exponential Smoothing

The model equation for simple exponential smoothing is

$$Y_t = \mu_t + \epsilon_t$$

The smoothing equation is

$$L_t = \alpha Y_t + (1 - \alpha)L_{t-1}$$

The error-correction form of the smoothing equation is

$$L_t = L_{t-1} + \alpha e_t$$

(Note: For missing values, $e_t = 0$.)

The k -step prediction equation is

$$\hat{Y}_t(k) = L_t$$

The ARIMA model equivalency to simple exponential smoothing is the ARIMA(0,1,1) model

$$(1 - B)Y_t = (1 - \theta B)\epsilon_t$$

$$\theta = 1 - \alpha$$

The moving-average form of the equation is

$$Y_t = \epsilon_t + \sum_{j=1}^{\infty} \alpha \epsilon_{t-j}$$

For simple exponential smoothing, the additive-invertible region is

$$\{0 < \alpha < 2\}$$

The variance of the prediction errors is estimated as

$$\text{var}(e_t(k)) = \text{var}(\epsilon_t) \left[1 + \sum_{j=1}^{k-1} \alpha^2 \right] = \text{var}(\epsilon_t)(1 + (k-1)\alpha^2)$$

Double (Brown) Exponential Smoothing

The model equation for double exponential smoothing is

$$Y_t = \mu_t + \beta_t t + \epsilon_t$$

The smoothing equations are

$$L_t = \alpha Y_t + (1 - \alpha)L_{t-1}$$

$$T_t = \alpha(L_t - L_{t-1}) + (1 - \alpha)T_{t-1}$$

This method may be equivalently described in terms of two successive applications of simple exponential smoothing:

$$S_t^{[1]} = \alpha Y_t + (1 - \alpha)S_{t-1}^{[1]}$$

$$S_t^{[2]} = \alpha S_t^{[1]} + (1 - \alpha)S_{t-1}^{[2]}$$

where $S_t^{[1]}$ are the smoothed values of Y_t , and $S_t^{[2]}$ are the smoothed values of $S_t^{[1]}$. The prediction equation then takes the form:

$$\hat{Y}_t(k) = (2 + \alpha k / (1 - \alpha))S_t^{[1]} - (1 + \alpha k / (1 - \alpha))S_t^{[2]}$$

The error-correction form of the smoothing equations is

$$L_t = L_{t-1} + T_{t-1} + \alpha e_t$$

$$T_t = T_{t-1} + \alpha^2 e_t$$

(Note: For missing values, $e_t = 0$.)

The k -step prediction equation is

$$\hat{Y}_t(k) = L_t + ((k-1) + 1/\alpha)T_t$$

The ARIMA model equivalency to double exponential smoothing is the ARIMA(0,2,2) model

$$(1-B)^2 Y_t = (1-\theta B)^2 \epsilon_t$$

$$\theta = 1 - \alpha$$

The moving-average form of the equation is

$$Y_t = \epsilon_t + \sum_{j=1}^{\infty} (2\alpha + (j-1)\alpha^2) \epsilon_{t-j}$$

For double exponential smoothing, the additive-invertible region is

$$\{0 < \alpha < 2\}$$

The variance of the prediction errors is estimated as

$$\text{var}(e_t(k)) = \text{var}(\epsilon_t) \left[1 + \sum_{j=1}^{k-1} (2\alpha + (j-1)\alpha^2)^2 \right]$$

Linear (Holt) Exponential Smoothing

The model equation for linear exponential smoothing is

$$Y_t = \mu_t + \beta_t t + \epsilon_t$$

The smoothing equations are

$$L_t = \alpha Y_t + (1 - \alpha)(L_{t-1} + T_{t-1})$$

$$T_t = \gamma(L_t - L_{t-1}) + (1 - \gamma)T_{t-1}$$

The error-correction form of the smoothing equations is

$$L_t = L_{t-1} + T_{t-1} + \alpha e_t$$

$$T_t = T_{t-1} + \alpha\gamma e_t$$

(Note: For missing values, $e_t = 0$.)

The k -step prediction equation is

$$\hat{Y}_t(k) = L_t + kT_t$$

The ARIMA model equivalency to linear exponential smoothing is the ARIMA(0,2,2) model

$$(1 - B)^2 Y_t = (1 - \theta_1 B - \theta_2 B^2) \epsilon_t$$

$$\theta_1 = 2 - \alpha - \alpha\gamma$$

$$\theta_2 = \alpha - 1$$

The moving-average form of the equation is

$$Y_t = \epsilon_t + \sum_{j=1}^{\infty} (\alpha + j\alpha\gamma) \epsilon_{t-j}$$

For linear exponential smoothing, the additive-invertible region is

$$\{0 < \alpha < 2\}$$

$$\{0 < \gamma < 4/\alpha - 2\}$$

The variance of the prediction errors is estimated as

$$\text{var}(e_t(k)) = \text{var}(\epsilon_t) \left[1 + \sum_{j=1}^{k-1} (\alpha + j\alpha\gamma)^2 \right]$$

Damped-Trend Linear Exponential Smoothing

The model equation for damped-trend linear exponential smoothing is

$$Y_t = \mu_t + \beta_t t + \epsilon_t$$

The smoothing equations are

$$L_t = \alpha Y_t + (1 - \alpha)(L_{t-1} + \phi T_{t-1})$$

$$T_t = \gamma(L_t - L_{t-1}) + (1 - \gamma)\phi T_{t-1}$$

The error-correction form of the smoothing equations is

$$L_t = L_{t-1} + \phi T_{t-1} + \alpha e_t$$

$$T_t = \phi T_{t-1} + \alpha \gamma e_t$$

(Note: For missing values, $e_t = 0$.)

The k -step prediction equation is

$$\hat{Y}_t(k) = L_t + \sum_{i=1}^k \phi^i T_t$$

The ARIMA model equivalency to damped-trend linear exponential smoothing is the ARIMA(1,1,2) model

$$(1 - \phi B)(1 - B)Y_t = (1 - \theta_1 B - \theta_2 B^2)\epsilon_t$$

$$\theta_1 = 1 + \phi - \alpha - \alpha \gamma \phi$$

$$\theta_2 = (\alpha - 1)\phi$$

The moving-average form of the equation (assuming $|\phi| < 1$) is

$$Y_t = \epsilon_t + \sum_{j=1}^{\infty} (\alpha + \alpha \gamma \phi(\phi^j - 1)/(\phi - 1))\epsilon_{t-j}$$

For damped-trend linear exponential smoothing, the additive-invertible region is

$$\{0 < \alpha < 2\}$$

$$\{0 < \phi\gamma < 4/\alpha - 2\}$$

The variance of the prediction errors is estimated as

$$\text{var}(e_t(k)) = \text{var}(\epsilon_t) \left[1 + \sum_{j=1}^{k-1} (\alpha + \alpha\gamma\phi(\phi^j - 1)/(\phi - 1))^2 \right]$$

Seasonal Exponential Smoothing

The model equation for seasonal exponential smoothing is

$$Y_t = \mu_t + s_p(t) + \epsilon_t$$

The smoothing equations are

$$L_t = \alpha(Y_t - S_{t-p}) + (1 - \alpha)L_{t-1}$$

$$S_t = \delta(Y_t - L_t) + (1 - \delta)S_{t-p}$$

The error-correction form of the smoothing equations is

$$L_t = L_{t-1} + \alpha e_t$$

$$S_t = S_{t-p} + \delta(1 - \alpha)e_t$$

(Note: For missing values, $e_t = 0$.)

The k -step prediction equation is

$$\hat{Y}_t(k) = L_t + S_{t-p+k}$$

The ARIMA model equivalency to seasonal exponential smoothing is the ARIMA(0,1,p+1)(0,1,0)_p model

$$(1 - B)(1 - B^p)Y_t = (1 - \theta_1 B - \theta_2 B^p - \theta_3 B^{p+1})\epsilon_t$$

$$\theta_1 = 1 - \alpha$$

$$\theta_2 = 1 - \delta(1 - \alpha)$$

$$\theta_3 = (1 - \alpha)(\delta - 1)$$

The moving-average form of the equation is

$$Y_t = \epsilon_t + \sum_{j=1}^{\infty} \psi_j \epsilon_{t-j}$$

$$\psi_j = \begin{cases} \alpha & \text{for } j \bmod p \neq 0 \\ \alpha + \delta(1 - \alpha) & \text{for } j \bmod p = 0 \end{cases}$$

For seasonal exponential smoothing, the additive-invertible region is

$$\{\max(-p\alpha, 0) < \delta(1 - \alpha) < (2 - \alpha)\}$$

The variance of the prediction errors is estimated as

$$\text{var}(e_t(k)) = \text{var}(\epsilon_t) \left[1 + \sum_{j=1}^{k-1} \psi_j^2 \right]$$

Winters Method – Additive Version

The model equation for the additive version of Winters method is

$$Y_t = \mu_t + \beta_t t + s_p(t) + \epsilon_t$$

The smoothing equations are

$$L_t = \alpha(Y_t - S_{t-p}) + (1 - \alpha)(L_{t-1} + T_{t-1})$$

$$T_t = \gamma(L_t - L_{t-1}) + (1 - \gamma)T_{t-1}$$

$$S_t = \delta(Y_t - L_t) + (1 - \delta)S_{t-p}$$

The error-correction form of the smoothing equations is

$$L_t = L_{t-1} + T_{t-1} + \alpha e_t$$

$$T_t = T_{t-1} + \alpha \gamma e_t$$

$$S_t = S_{t-p} + \delta(1 - \alpha)e_t$$

(Note: For missing values, $e_t = 0$.)

The k -step prediction equation is

$$\hat{Y}_t(k) = L_t + kT_t + S_{t-p+k}$$

The ARIMA model equivalency to the additive version of Winters method is the $\text{ARIMA}(0,1,p+1)(0,1,0)_p$ model

$$(1-B)(1-B^p)Y_t = \left[1 - \sum_{i=1}^{p+1} \theta_i B^i\right] \epsilon_t$$

$$\theta_j = \begin{cases} 1 - \alpha - \alpha\gamma & j = 1 \\ -\alpha\gamma & 2 \leq j \leq p-1 \\ 1 - \alpha\gamma - \delta(1-\alpha) & j = p \\ (1-\alpha)(\delta-1) & j = p+1 \end{cases}$$

The moving-average form of the equation is

$$Y_t = \epsilon_t + \sum_{j=1}^{\infty} \psi_j \epsilon_{t-j}$$

$$\psi_j = \begin{cases} \alpha + j\alpha\gamma & \text{for } j \bmod p \neq 0 \\ \alpha + j\alpha\gamma + \delta(1-\alpha), & \text{for } j \bmod p = 0 \end{cases}$$

For the additive version of Winters method (see Archibald 1990), the additive-invertible region is

$$\{\max(-p\alpha, 0) < \delta(1-\alpha) < (2-\alpha)\}$$

$$\{0 < \alpha\gamma < 2 - \alpha - \delta(1-\alpha)(1 - \cos(\vartheta))\}$$

where ϑ is the smallest nonnegative solution to the equations listed in Archibald (1990).

The variance of the prediction errors is estimated as

$$\text{var}(e_t(k)) = \text{var}(\epsilon_t) \left[1 + \sum_{j=1}^{k-1} \psi_j^2\right]$$

Winters Method – Multiplicative Version

In order to use the multiplicative version of Winters method, the time series and all predictions must be strictly positive.

The model equation for the multiplicative version of Winters method is

$$Y_t = (\mu_t + \beta_t t) s_p(t) + \epsilon_t$$

The smoothing equations are

$$L_t = \alpha(Y_t/S_{t-p}) + (1 - \alpha)(L_{t-1} + T_{t-1})$$

$$T_t = \gamma(L_t - L_{t-1}) + (1 - \gamma)T_{t-1}$$

$$S_t = \delta(Y_t/L_t) + (1 - \delta)S_{t-p}$$

The error-correction form of the smoothing equations is

$$L_t = L_{t-1} + T_{t-1} + \alpha e_t / S_{t-p}$$

$$T_t = T_{t-1} + \alpha \gamma e_t / S_{t-p}$$

$$S_t = S_{t-p} + \delta(1 - \alpha)e_t / L_t$$

(Note: For missing values, $e_t = 0$.)

The k -step prediction equation is

$$\hat{Y}_t(k) = (L_t + kT_t)S_{t-p+k}$$

The multiplicative version of Winters method does not have an ARIMA equivalent; however, when the seasonal variation is small, the ARIMA additive-invertible region of the additive version of Winters method described in the preceding section can approximate the stability region of the multiplicative version.

The variance of the prediction errors is estimated as

$$\text{var}(e_t(k)) = \text{var}(\epsilon_t) \left[\sum_{i=0}^{\infty} \sum_{j=0}^{p-1} (\psi_{j+ip} S_{t+k} / S_{t+k-j})^2 \right]$$

where ψ_j are as described for the additive version of Winters method, and $\psi_j = 0$ for $j \geq k$.

Intermittent Models

This section details the computations performed for the intermittent forecasting models.

Intermittent Time Series

Intermittent time series have a large number of values that are zero. These types of series commonly occur in Internet, inventory, sales, and other data where the demand for a particular item is intermittent. Typically, when the value of the series associated with a particular time period is nonzero, *demand* occurs; and, when the value is zero (or missing), *no demand* occurs. Since it is entirely possible that the number of time periods for which *no demand* occurs is large, many of the series values will be zero. Typical time series models (for example, smoothing models) are inadequate in the case of intermittent time series because many of the series values are zero. Since these models are based on weighted-summations of past values, they bias forecasts away from zero. Unlike the smoothing models that provide forecasts for future time periods, intermittent forecasting models provide recommended *stocking levels* or *estimated demand per period* that are used to satisfy future demand.

Intermittent Series Decomposition and Analysis

An intermittent time series (demand series) can be decomposed into two components: a demand interval series and a demand size series. Both of these component series are indexed based on when a *demand* occurred (demand index) rather than each time period (time index). The demand interval series is constructed based on the number of time periods between *demands*. The demand size series is constructed based on the size (or value) of the *demands* excluding zero (or base) demand values. Using these two component series, the average demand series is computed by dividing the size component values by the interval component values.

When a *demand* occurs is subjective; but typically depends on a *base* value. Typically, the base value is zero (default), but it can be any constant value and can be automatically determined based on the characteristics of the demand series.

Given a time series y_t , for $t = 1$ to T , where t is the time index, suppose that there are N nonzero demands occurring at times $t = t_i$, where $t_{i-1} < t_i$, for $i = 1$ to N . The time series is dissected into the demand interval series and the demand size series as follows:

$$\begin{array}{lll} \text{(Demand Interval Series)} & q_i = t_i - t_{i-1} & \text{for } i = 2 \text{ to } N \\ \text{(Demand Size Series)} & d_i = y_{t_i} - \text{base} & \text{for } i = 1 \text{ to } N \\ \text{(Average Demand Series)} & a_i = d_i / q_i & \text{for } i = 2 \text{ to } N \end{array}$$

For the beginning of the demand series, q_1 is assigned to t_1 , which assumes that a demand just occurred prior to the first recorded observation. For the end of the demand series, q_{N+1} is assigned to $(T + 1 - t_N)$, which assumes that demand will occur just after the last recorded observation. The next future demand size, d_{N+1} , is always set to missing.

After decomposition, descriptive (summary) statistics can be computed to gain a greater understanding of the demand series including those statistics based on the season index.

For statistical analysis and model fitting, q_i and a_i for $i = 2$ to N and d_i for $i = 1$ to N are used. For forecasting, q_i for $i = 1$ to $N + 1$, a_i for $i = 1$ to N , d_i for $i = 1$ to N are used.

Croston's Method

Croston's Method models and forecasts each component independently, then combines the two forecasts. The following provides a description of how Croston's Method is used in SAS High Performance Forecasting. More detailed information on this method can be found in Croston (1972) and Willemain, Smart, and Shocker (1994). The following description of Croston's Method is based on the perspective of a person familiar with typical time series modeling techniques such as smoothing models.

By treating each component of the demand series as a time series based on the demand index, optimal smoothing parameters can be estimated and predictions for each component can be computed using nonseasonal exponential smoothing methods (simple, double, linear, damped-trend) as well as their transformed versions (log, square-root, logistic, Box-Cox).

For example, the following simple smoothing equations are used to generate predictions for demand size and interval components:

$$\begin{aligned} \text{(Smoothed demand interval series)} \quad L_i^q &= L_{i-1}^q + \alpha_q(q_{i-1} - L_{i-1}^q) \\ \text{(Smoothed demand size series)} \quad L_i^d &= L_{i-1}^d + \alpha_d(d_{i-1} - L_{i-1}^d) \end{aligned}$$

The demand interval parameter, α_q , and demand size parameter, α_d , and the starting, intermediate, and final smoothing level states, L_i^q and L_i^d , are estimated from the data using simple exponential smoothing parameter estimation. For the starting state at $i = 1$, $L_1^q = \max(q_1, L_0^q)$ where L_0^q is the final backcast level state. For $i > 1$, the one-step-ahead prediction for demand interval q_i is $\hat{q}_i = L_{i-1}^q$. For $i > 0$, the one-step-ahead prediction for demand size d_i is $\hat{d}_i = L_{i-1}^d$.

Other (transformed) nonseasonal exponential smoothing methods can be used in a similar fashion. For linear smoothing, $L_1^q = \max(q_1 - T_0^q, L_0^q)$ where T_0^q is the final backcast trend state. For damp-trend smoothing, $L_1^q = \max(q_1 - \phi_q T_0^q, L_0^q)$ where ϕ_q is the damping parameter and T_0^q is the final backcast trend state. For double smoothing, $L_1^q = \max(q_1 - T_0^q/\alpha_q, L_0^q)$ where α_q is the weight parameter and T_0^q is the final backcast trend state.

Using these predictions based on the demand index, predictions of the average demand per period can be estimated. Predicted demand per period is also known as "stocking level," assuming that disturbances affecting q_i are independent of d_i . (This assumption is quite significant.)

$$\begin{aligned}
\text{(Estimated demand per period)} \quad y_i^* &= \hat{d}_i / \hat{q}_i \\
E[y_i^*] &= E[d_i] / E[q_i] = E[z_{i-1}] / E[p_{i-1}] = \mu_d / \mu_q \\
\text{(Variance)} \quad Var(y_i^*) &= (\hat{d}_i / \hat{q}_i)^2 (Var(d_i) / \hat{d}_i^2 + Var(q_i) / \hat{q}_i^2)
\end{aligned}$$

where μ_d , \bar{d} , and s_d are the mean, sample average, and standard deviation of the non-zero demands, and μ_q , \bar{q} , and s_q are the mean, sample average, and standard deviation of the number of time periods between demands.

For the beginning of the series, the denominator of y_1^* is assigned q_i or the starting smoothing state p_0 , whichever is greater. For the end of the series, the denominator of y_{N+1}^* is assigned $q_{N+1} = (T + 1 - t_N)$ or the final smoothing state p_N , whichever is greater.

Once the average demand per period has been estimated, a stocking level can be recommended:

$$\begin{aligned}
\text{(Recommended stocking level)} \quad \hat{y}_t &= y_i^* && \text{when } t_i \leq t < t_{i+1} \\
\text{(Variance)} \quad Var(\hat{y}_t) &= Var(y_i^*) && \text{when } t_i \leq t < t_{i+1}
\end{aligned}$$

Since the predicted demand per period will be different than typical time series forecasts, the usual way of computing statistics of fit should also be different. The statistics of fit are based on the difference between the recommended stocking levels between demands and the demands:

$$\begin{aligned}
\text{(Accumulated recommended stocks)} \quad s_t &= \sum_{i=0}^t (\hat{y}_t - y_i^*) \\
\text{(Estimate - Demand)} \quad e_{t_i} &= \hat{d}_i q_i - d_i && \text{when time } t_{i+1} \text{ has demand}
\end{aligned}$$

Croston's Method produces the same forecasts as simple exponential smoothing when demand occurs in every time period, $q_i = 1$ for all i , but different (lower) prediction error variances. Croston's Method is recommended for intermittent time series only.

Average Demand Method

Similar to Croston's Method, the Average Demand Method is used to forecast intermittent time series; however, the Average Demand Method forecasts the average demand series directly, whereas Croston's Method forecasts average demand series indirectly using the inverse decomposition of the demand interval and size series forecasts.

By treating the average demand series as a time series based on the demand index, optimal smoothing parameters can be estimated and predictions for average demand can be computed using nonseasonal exponential smoothing methods (simple, double, linear, damped-trend) as well as their transformed versions (log, square-root, logistic, Box-Cox).

For example, the following simple smoothing equations are used to generate predictions for the average demand series:

$$(\text{Smoothed Average Demand Series}) \quad L_i^a = L_{i-1}^a + \alpha_a(a_{i-1} - L_{i-1}^a)$$

The average demand level parameter, α_a , and the starting, intermediate, and final smoothing level states, L_i^a , are estimated from the data using simple exponential smoothing parameter estimation. For the starting state at $i = 1$, $L_1^a = \max(a_1, L_0^a)$ where L_0^a is the final backcast level state. For $i > 1$, the one-step-ahead prediction for a_i is $\hat{a}_i = L_{i-1}^a$.

Other (transformed) nonseasonal exponential smoothing methods can be used in a similar fashion. For linear smoothing, $L_1^a = \max(a_1 - T_0^a, L_0^a)$ where T_0^a is the final backcast trend state. For damp-trend smoothing, $L_1^a = \max(a_1 - \phi_a T_0^a, L_0^a)$ where ϕ_a is the damping parameter and T_0^a is the final backcast trend state. For double smoothing, $L_1^a = \max(a_1 - T_0^a/\alpha_a, L_0^a)$ where α_a is the weight parameter and T_0^a is the final backcast trend state.

Using these predictions based on the demand index, predictions of the average demand per period are provided directly, unlike Croston's Method where the average demand is predicted using a ratio of predictions of the demand interval and size components.

$$\begin{aligned} (\text{Estimated demand per period}) \quad & y_i^* = \hat{a}_i + base \\ & E[y_i^*] = E[a_i] \\ (\text{Variance}) \quad & \text{see the exponential smoothing models} \end{aligned}$$

For the beginning of the series, \hat{a}_1 is derived from the starting level smoothing state and starting trend smoothing state (if applicable).

Once the average demand per period has been estimated, a stocking level can be recommended similar to Croston's Method.

The Average Demand Method produces the same forecasts as exponential smoothing when demand occurs in every time period, $q_i = 1$ for all i , but different (lower) prediction error variances. The Average Demand Method is recommended for intermittent time series only.

Time Indexed Holdout Samples versus Demand Indexed Holdout Samples

Holdout samples are typically specified based on the time index. For example, "hold-out the last six months data." For a demand series, the demand indexed holdout refers to the "demands that have occurred in the last six months." If there are four demands in the last six months, the demand indexed holdout is four for a time indexed holdout of six. If there are no demands in the time indexed holdout, the demand indexed holdout is zero and in-sample analysis is used. For intermittent demand model selection, demand indexed-based holdouts are used for model selection.

Automatic Intermittent Demand Model Selection

The exponential smoothing method to be used to forecast the intermittent demand series can be specified, or it can be selected automatically using a model selection criterion and either in-sample or holdout sample analysis. The exponential smoothing method for each demand series component (interval, size, and average) can be automatically selected as well as the choice between Croston's Method and the Average Demand Method.

For Croston's Method, the exponential smoothing methods used to forecast the demand interval and size components are automatically selected independently. For the Average Demand Method, the exponential smoothing methods used to forecast the average demand component are automatically selected, again independently. Based on the model selection criterion, the selection is based on how well the method fits (in sample) or predicts (holdout sample) the demand series component by treating the demand index as a time index. The following equations describe the component prediction errors associated with each of the demand series components that are used in component model selection:

$$\begin{aligned} \text{(Demand Interval Series)} \quad e_i^q &= q_i - \hat{q}_i & \text{for } i = 2 \text{ to } N \\ \text{(Demand Size Series)} \quad e_i^d &= d_i - \hat{d}_i & \text{for } i = 1 \text{ to } N \\ \text{(Average Demand Series)} \quad e_i^a &= a_i - \hat{a}_i & \text{for } i = 2 \text{ to } N \end{aligned}$$

Once the exponential smoothing methods are selected for each demand series component, the predictions for either Croston's Method, (\hat{d}_i/\hat{q}_i) , the Average Demand Method, \hat{a}_i , or both are computed based on the selected method for each component.

When choosing between Croston's Method and the Average Demand Method, the model is selected by considering how well the model predicts average demand with respect to the time. The following equations describe the average prediction errors associated with the predicted average demand that are used in model selection:

$$\begin{aligned} \text{(Croston's Method)} \quad e_i^c &= (d_i/q_i) - (\hat{d}_i/\hat{q}_i) & \text{for } i = 2 \text{ to } N \\ \text{(Average Demand Method)} \quad e_i^a &= a_i - \hat{a}_i & \text{for } i = 2 \text{ to } N \end{aligned}$$

Series Transformations

For forecasting models, transforming the time series may aid in improving forecasting accuracy.

There are four transformations available, for strictly positive series only. Let $y_t > 0$ be the original time series, and let w_t be the transformed series. The transformations are defined as follows:

Log is the logarithmic transformation

$$w_t = \ln(y_t)$$

Logistic is the logistic transformation

$$w_t = \ln(cy_t/(1 - cy_t))$$

where the scaling factor c is

$$c = (1 - e^{-6})10^{-\text{ceil}(\log_{10}(\max(y_t)))}$$

and $\text{ceil}(x)$ is the smallest integer greater than or equal to x .

Square Root is the square root transformation

$$w_t = \sqrt{y_t}$$

Box Cox is the Box-Cox transformation

$$w_t = \begin{cases} \frac{y_t^\lambda - 1}{\lambda}, & \lambda \neq 0 \\ \ln(y_t), & \lambda = 0 \end{cases}$$

Parameter estimation is performed using the transformed series. The transformed model predictions and confidence limits are then obtained from the transformed time-series and these parameter estimates.

The transformed model predictions \hat{w}_t are used to obtain either the minimum mean absolute error (MMAE) or minimum mean squared error (MMSE) predictions \hat{y}_t , depending on the setting of the forecast options. The model is then evaluated based on the residuals of the original time series and these predictions. The transformed model confidence limits are inverse-transformed to obtain the forecast confidence limits.

Predictions for Transformed Models

Since the transformations described in the previous section are monotonic, applying the inverse-transformation to the transformed model predictions results in the *median* of the conditional probability density function at each point in time. This is the minimum mean absolute error (MMAE) prediction.

If $w_t = F(y_t)$ is the transform with inverse-transform $y_t = F^{-1}(w_t)$, then

$$\text{median}(\hat{y}_t) = F^{-1}(E[w_t]) = F^{-1}(\hat{w}_t)$$

The minimum mean squared error (MMSE) predictions are the *mean* of the conditional probability density function at each point in time. Assuming that the prediction errors are normally distributed with variance σ_t^2 , the MMSE predictions for each of the transformations are as follows:

Log is the conditional expectation of inverse-logarithmic transformation.

$$\hat{y}_t = E[e^{w_t}] = \exp(\hat{w}_t + \sigma_t^2/2)$$

Logistic is the conditional expectation of inverse-logistic transformation.

$$\hat{y}_t = E \left[\frac{1}{c(1 + \exp(-w_t))} \right]$$

where the scaling factor $c = (1 - 10^{-6})10^{-\text{ceil}(\log_{10}(\max(y_t)))}$.

Square Root is the conditional expectation of the inverse-square root transformation.

$$\hat{y}_t = E[w_t^2] = \hat{w}_t^2 + \sigma_t^2$$

Box Cox is the conditional expectation of the inverse Box-Cox transformation.

$$\hat{y}_t = \begin{cases} E[(\lambda w_t + 1)^{1/\lambda}], & \lambda \neq 0 \\ E[e^{w_t}] = \exp(\hat{w}_t + \frac{1}{2}\sigma_t^2), & \lambda = 0 \end{cases}$$

The expectations of the inverse logistic and Box-Cox ($\lambda \neq 0$) transformations do not generally have explicit solutions and are computed using numerical integration.

Series Diagnostic Tests

This section describes the diagnostic tests that are used to determine the kinds of forecasting models appropriate for a series.

The series diagnostics are a set of heuristics that provide recommendations on whether or not the forecasting model should contain a log transform, trend terms, and seasonal terms or whether or not the time series is intermittent. These recommendations are used by the automatic model selection process to restrict the model search to a subset of the model selection list.

The tests that are used by the series diagnostics will not always produce the correct classification of the series. They are intended to accelerate the process of searching for a good forecasting model for the series, but you should not rely on them if finding the very best model is important to you.

The series diagnostics tests are intended as a heuristic tool only, and no statistical validity is claimed for them. These tests may be modified and enhanced in future releases of the SAS High Performance Forecasting. The testing strategy is as follows:

1. **Intermittent test.** Compute the average time interval between demands. If the time average time interval is greater than a preset limit, an intermittent forecasting model is used.
2. **Seasonality test.** The resultant series is tested for seasonality. A seasonal dummy model with AR(1) errors is fit and the joint significance of the seasonal dummy estimates is tested. If the seasonal dummies are significant, the AIC statistic for this model is compared to the AIC for an AR(1) model without seasonal dummies. If the seasonal model has a lower AIC, a seasonal forecasting model is used.

Statistics of Fit

This section explains the goodness-of-fit statistics reported to measure how well different models fit the data. The statistics of fit for the various forecasting models can be printed or stored in a data set.

The various statistics of fit reported are as follows. In these formula, n is the number of nonmissing observations and k is the number of fitted parameters in the model.

Number of Nonmissing Observations.

The number of nonmissing observations used to fit the model.

Number of Observations.

The total number of observations used to fit the model, including both missing and nonmissing observations.

Number of Missing Actuals.

The number of missing actual values.

Number of Missing Predicted Values.

The number of missing predicted values.

Number of Model Parameters.

The number of parameters fit to the data. For combined forecast, this is the number of forecast components.

Total Sum of Squares (Uncorrected).

The total sum of squares for the series, SST, uncorrected for the mean: $\sum_{t=1}^n y_t^2$.

Total Sum of Squares (Corrected).

The total sum of squares for the series, SST, corrected for the mean: $\sum_{t=1}^n (y_t - \bar{y})^2$, where \bar{y} is the series mean.

Sum of Square Errors.

The sum of the squared prediction errors, SSE. $SSE = \sum_{t=1}^n (y_t - \hat{y}_t)^2$, where \hat{y} is the one-step predicted value.

Mean Square Error.

The mean squared prediction error, MSE, calculated from the one-step-ahead forecasts. $MSE = \frac{1}{n} SSE$. This formula enables you to evaluate small holdout samples.

Root Mean Square Error.

The root mean square error (RMSE), \sqrt{MSE} .

Mean Absolute Percent Error.

The mean absolute percent prediction error (MAPE), $\frac{100}{n} \sum_{t=1}^n |(y_t - \hat{y}_t)/y_t|$. The summation ignores observations where $y_t = 0$.

Mean Absolute Error.

The mean absolute prediction error, $\frac{1}{n} \sum_{t=1}^n |y_t - \hat{y}_t|$.

R-Square.

The R^2 statistic, $R^2 = 1 - SSE/SST$. If the model fits the series badly, the model error sum of squares, SSE , may be larger than SST and the R^2 statistic will be negative.

Adjusted R-Square.

The adjusted R^2 statistic, $1 - (\frac{n-1}{n-k})(1 - R^2)$.

Amemiya's Adjusted R-Square.

Amemiya's adjusted R^2 , $1 - (\frac{n+k}{n-k})(1 - R^2)$.

Random Walk R-Square.

The random walk R^2 statistic (Harvey's R^2 statistic using the random walk model for comparison), $1 - (\frac{n-1}{n})SSE/RWSSE$, where $RWSSE = \sum_{t=2}^n (y_t - y_{t-1} - \mu)^2$, and $\mu = \frac{1}{n-1} \sum_{t=2}^n (y_t - y_{t-1})$.

Akaike's Information Criterion.

Akaike's information criterion (AIC), $n \ln(SSE/n) + 2k$.

Schwarz Bayesian Information Criterion.

Schwarz Bayesian information criterion (SBC or BIC), $n \ln(SSE/n) + k \ln(n)$.

Amemiya's Prediction Criterion.

Amemiya's prediction criterion, $\frac{1}{n} SST(\frac{n+k}{n-k})(1 - R^2) = (\frac{n+k}{n-k})\frac{1}{n} SSE$.

Maximum Error.

The largest prediction error.

Minimum Error.

The smallest prediction error.

Maximum Percent Error.

The largest percent prediction error, $100 \max((y_t - \hat{y}_t)/y_t)$. The summation ignores observations where $y_t = 0$.

Minimum Percent Error.

The smallest percent prediction error, $100 \min((y_t - \hat{y}_t)/y_t)$. The summation ignores observations where $y_t = 0$.

Mean Error.

The mean prediction error, $\frac{1}{n} \sum_{t=1}^n (y_t - \hat{y}_t)$.

Mean Percent Error.

The mean percent prediction error, $\frac{100}{n} \sum_{t=1}^n \frac{(y_t - \hat{y}_t)}{y_t}$. The summation ignores observations where $y_t = 0$.

References

- Aldrin, M. and Damsleth, E. (1989), "Forecasting Non-seasonal Time Series with Missing Observations," *Journal of Forecasting*, 8, 97-116.
- Anderson, T.W. (1971), *The Statistical Analysis of Time Series*, New York: John Wiley & Sons, Inc.
- Archibald, B.C. (1990), "Parameter Space of the Holt-Winters' Model," *International Journal of Forecasting*, 6, 199-209.
- Bartolomei, S.M. and Sweet, A.L. (1989), "A Note on the Comparison of Exponential Smoothing Methods for Forecasting Seasonal Series," *International Journal of Forecasting*, 5, 111-116.
- Bowerman, B.L. and O'Connell, R.T. (1979), *Time Series and Forecasting: An Applied Approach*, North Scituate, Massachusetts: Duxbury Press.
- Box, G.E.P. and Cox D.R. (1964), "An Analysis of Transformations," *Journal of Royal Statistical Society B*, No. 26, 211-243.
- Box, G.E.P. and Jenkins, G.M. (1976), *Time Series Analysis: Forecasting and Control*, Revised Edition, San Francisco: Holden-Day.
- Brocklebank, J.C. and Dickey, D.A. (1986), *SAS System for Forecasting Time Series, 1986 Edition*, Cary, North Carolina: SAS Institute Inc.
- Brown, R.G. (1962), *Smoothing, Forecasting and Prediction of Discrete Time Series*, New York: Prentice-Hall.
- Brown, R.G. and Meyer, R.F. (1961), "The Fundamental Theorem of Exponential Smoothing," *Operations Research*, 9, 673-685.
- Chatfield, C. (1978), "The Holt-Winters Forecasting Procedure," *Applied Statistics*, 27, 264-279.
- Chatfield, C. and Yar, M. (1988), "Holt-Winters Forecasting: Some Practical Issues," *The Statistician*, 37, 129-140.
- Chatfield, C. and Yar, M. (1991), "Prediction intervals for multiplicative Holt-Winters," *International Journal of Forecasting*, 7, 31-37.
- Cogger, K.O. (1974), "The Optimality of General-Order Exponential Smoothing," *Operations Research*, 22, 858.
- Cox, D. R. (1961), "Prediction by Exponentially Weighted Moving Averages and Related Methods," *Journal of the Royal Statistical Society, Series B*, 23, 414-422.
- Croston, J.D. (1972), "Forecasting and Stock Control for Intermittent Demands," *Operations Research Quarterly*, 23, No. 3.
- Davidson, J. (1981), "Problems with the Estimation of Moving Average Models," *Journal of Econometrics*, 16, 295.

- Fair, R.C. (1986), "Evaluating the Predictive Accuracy of Models," in *Handbook of Econometrics*, Vol. 3., Griliches, Z. and Intriligator, M.D., eds., New York: North Holland.
- Fildes, R. (1979), "Quantitative Forecasting – the State of the Art: Extrapolative Models," *Journal of Operational Research Society*, 30, 691-710.
- Fuller, W.A. (1976), *Introduction to Statistical Time Series*, New York: John Wiley & Sons, Inc.
- Gardner, E.S., Jr. (1984), "The Strange Case of the Lagging Forecasts," *Interfaces*, 14, 47-50.
- Gardner, E.S., Jr. (1985), "Exponential Smoothing: the State of the Art," *Journal of Forecasting*, 4, 1-38.
- Granger, C.W.J. and Newbold, P. (1977), *Forecasting Economic Time Series*, New York: Academic Press, Inc.
- Greene, W.H. (1993), *Econometric Analysis*, Second Edition, New York: Macmillan Publishing Company.
- Hamilton, J. D. (1994), *Time Series Analysis*, Princeton University Press: Princeton.
- Harvey, A.C. (1981), *Time Series Models*, New York: John Wiley & Sons, Inc.
- Harvey, A.C. (1984), "A Unified View of Statistical Forecasting Procedures," *Journal of Forecasting*, 3, 245-275.
- Hopewood, W.S., McKeown, J.C. and Newbold, P. (1984), "Time Series Forecasting Models Involving Power Transformations," *Journal of Forecasting*, Vol 3, No. 1, 57-61.
- Ledolter, J. and Abraham, B. (1984), "Some Comments on the Initialization of Exponential Smoothing," *Journal of Forecasting*, 3, 79-84.
- Ljung, G.M. and Box, G.E.P. (1978), "On a Measure of Lack of Fit in Time Series Models," *Biometrika*, 65, 297-303.
- Makridakis, S., Wheelwright, S.C., and McGee, V.E. (1983), *Forecasting: Methods and Applications*, Second Edition, New York: John Wiley & Sons.
- McKenzie, Ed (1984), "General Exponential Smoothing and the Equivalent ARMA Process," *Journal of Forecasting*, 3, 333-344.
- McKenzie, Ed (1986), "Error Analysis for Winters' Additive Seasonal Forecasting System," *International Journal of Forecasting*, 2, 373-382.
- Montgomery, D.C. and Johnson, L.A. (1976), *Forecasting and Time Series Analysis*, New York: McGraw-Hill Book Co.
- Morf, M., Sidhu, G.S., and Kailath, T. (1974), "Some New Algorithms for Recursive Estimation on Constant Linear Discrete Time Systems," *I.E.E.E. Transactions on Automatic Control*, AC-19, 315-323.
- Nelson, C.R. (1973), *Applied Time Series for Managerial Forecasting*, San Francisco: Holden-Day.

- Newbold, P. (1981), "Some Recent Developments in Time Series Analysis," *International Statistical Review*, 49, 53-66.
- Pankratz, A. and Dudley, U. (1987), "Forecast of Power-Transformed Series," *Journal of Forecasting*, Vol 6, No. 4, 239-248.
- Priestly, M.B. (1981), *Spectral Analysis and Time Series, Volume 1: Univariate Series*, New York: Academic Press, Inc.
- Roberts, S.A. (1982), "A General Class of Holt-Winters Type Forecasting Models," *Management Science*, 28, 808-820.
- Sweet, A.L. (1985), "Computing the Variance of the Forecast Error for the Holt-Winters Seasonal Models," *Journal of Forecasting*, 4, 235-243.
- Winters, P.R. (1960), "Forecasting Sales by Exponentially Weighted Moving Averages," *Management Science*, 6, 324-342.
- Yar, M. and Chatfield, C. (1990), "Prediction Intervals for the Holt-Winters Forecasting Procedure," *International Journal of Forecasting*, 6, 127-137.
- Willemain, T.R, Smart, C.N, Shocker, J.H (1994) "Forecasting Intermittent Demand In Manufacturing: Comparative Evaluation of Croston's Method," *International Journal of Forecasting*, 10, 529-538.

Subject Index

A

- additive-invertible region
 - smoothing weights, [58](#)
- adjusted R-square
 - statistics of fit, [76](#)
- AIC,
 - See Akaike's information criterion
- Akaike's information criterion
 - AIC, [76](#)
 - statistics of fit, [76](#)
- Amemiya's prediction criterion
 - statistics of fit, [76](#)
- Amemiya's R-square
 - statistics of fit, [76](#)

B

- BIC,
 - See Schwarz Bayesian information criterion
- boundaries
 - smoothing weights, [58](#)
- Box Cox
 - transformations, [72](#)
- Box Cox transformation,
 - See transformations
- Brown smoothing model,
 - See double exponential smoothing
- BY groups
 - HPF procedure, [14](#)

C

- calculations
 - smoothing models, [56](#)
- corrected sum of squares
 - statistics of fit, [75](#)
- Croston's,
 - See intermittent models

D

- damped-trend exponential smoothing, [63](#)
 - smoothing models, [63](#)
- diagnostic tests, [74](#)
- double exponential smoothing, [60](#)
 - Brown smoothing model, [60](#)
 - smoothing models, [60](#)

E

- error sum of squares
 - statistics of fit, [75](#)

- exponential smoothing,
 - See smoothing models

F

- forecasting, [55](#)
- forecasting models
 - intermittent models, [68](#)
 - smoothing models, [55](#)

G

- goodness-of-fit statistics,
 - See statistics of fit

H

- Holt smoothing model,
 - See linear exponential smoothing
- Holt-Winters Method,
 - See Winters Method
- HPF procedure
 - BY groups, [14](#)
 - ODS graph names, [40](#)

I

- initializations
 - smoothing models, [57](#)
- intermittent models
 - Croston's Method, [68](#)
 - forecasting models, [68](#)

L

- linear exponential smoothing, [61](#)
 - Holt smoothing model, [61](#)
 - smoothing models, [61](#)
- log
 - transformations, [72](#)
- log test, [74](#)
- log transformation,
 - See transformations
- logistic
 - transformations, [72](#)

M

- mean absolute error
 - statistics of fit, [76](#)
- mean absolute percent error
 - statistics of fit, [75](#)
- mean percent error
 - statistics of fit, [76](#)

mean prediction error
 statistics of fit, 76
 mean square error
 statistics of fit, 75
 missing values
 smoothing models, 57
 MMAE, 73
 MMSE, 73
 model evaluation, 55

N

nonmissing observations
 statistics of fit, 75
 number of observations
 statistics of fit, 75

O

ODS graph names
 HPF procedure, 40
 optimizations
 smoothing weights, 58

P

parameter estimation, 55
 predictions
 smoothing models, 57

R

R-square statistic
 statistics of fit, 76
 random walk R-square
 statistics of fit, 76
 root mean square error
 statistics of fit, 75

S

SBC,
 See Schwarz Bayesian information criterion
 Schwarz Bayesian information criterion
 BIC, 76
 SBC, 76
 statistics of fit, 76
 seasonal exponential smoothing, 64
 smoothing models, 64
 seasonality test, 74
 series diagnostics, 74
 simple exponential smoothing, 59
 smoothing models, 59
 smoothing equations, 56
 smoothing models, 56
 smoothing models
 calculations, 56
 damped-trend exponential smoothing, 63
 double exponential smoothing, 60
 exponential smoothing, 55
 forecasting models, 55
 initializations, 57
 linear exponential smoothing, 61
 missing values, 57

 predictions, 57
 seasonal exponential smoothing, 64
 simple exponential smoothing, 59
 smoothing equations, 56
 smoothing state, 56
 smoothing weights, 58
 standard errors, 58
 underlying model, 56
 Winters Method, 65, 67
 smoothing state, 56
 smoothing models, 56
 smoothing weights, 58
 additive-invertible region, 58
 boundaries, 58
 optimizations, 58
 smoothing models, 58
 specifications, 58
 weights, 58
 specifications
 smoothing weights, 58
 square root
 transformations, 72
 square root transformation,
 See transformations
 standard errors
 smoothing models, 58
 statistics of fit, 75
 adjusted R-square, 76
 Akaike's information criterion, 76
 Amemiya's prediction criterion, 76
 Amemiya's R-square, 76
 corrected sum of squares, 75
 error sum of squares, 75
 goodness-of-fit statistics, 75
 mean absolute error, 76
 mean absolute percent error, 75
 mean percent error, 76
 mean prediction error, 76
 mean square error, 75
 nonmissing observations, 75
 number of observations, 75
 R-square statistic, 76
 random walk R-square, 76
 root mean square error, 75
 Schwarz Bayesian information criterion, 76
 uncorrected sum of squares, 75

T

transformations
 Box Cox, 72
 Box Cox transformation, 72
 log, 72
 log transformation, 72
 logistic, 72
 square root, 72
 square root transformation, 72

U

uncorrected sum of squares

- statistics of fit, [75](#)
- underlying model
 - smoothing models, [56](#)

W

- weights,
 - See smoothing weights
- Winters Method, [65](#), [67](#)
 - Holt-Winters Method, [65](#)
 - smoothing models, [65](#), [67](#)

Syntax Index

A

ACCUMULATE= option
 FORECAST statement (HPF), [14](#)
 ID statement (HPF), [19](#)
ALIGN= option
 ID statement (HPF), [20](#)
ALPHA= option
 FORECAST statement (HPF), [14](#)
AVERAGE= option
 IDM statement (HPF), [24](#)

B

BACK= option
 PROC HPF statement, [11](#)
BASE= option
 IDM statement (HPF), [24](#)
BOUNDS= option
 IDM statement (HPF), [26](#)
BY statement
 HPF procedure, [14](#)

D

DAMPPARM= option
 IDM statement (HPF), [26](#)
DAMPREST= option
 IDM statement (HPF), [26](#)
DATA= option
 PROC HPF statement, [11](#)

E

END= option
 ID statement (HPF), [20](#)

F

FORECAST statement
 HPF procedure, [14](#)

H

HOLDOUT= option
 FORECAST statement (HPF), [14](#)
HOLDOUTPCT= option
 FORECAST statement (HPF), [15](#)
HPF, [5](#)
HPF procedure, [9](#)
 syntax, [9](#)

I

ID statement
 HPF procedure, [19](#)
IDM statement
 HPF procedure, [22](#)
INTERMITTENT= option
 FORECAST statement (HPF), [15](#)
INTERVAL= option
 ID statement (HPF), [21](#)
 IDM statement (HPF), [24](#)

L

LEAD= option
 PROC HPF statement, [11](#)
LEVELPARM= option
 IDM statement (HPF), [26](#)
LEVELREST= option
 IDM statement (HPF), [26](#)

M

MAXERROR= option
 PROC HPF statement, [11](#)
MEDIAN option
 FORECAST statement (HPF), [15](#)
 IDM statement (HPF), [25](#)
METHOD= option
 IDM statement (HPF), [25](#)
MODEL= option
 FORECAST statement (HPF), [15](#)

N

NBACKCAST= option
 FORECAST statement (HPF), [16](#)
NOEST option
 IDM statement (HPF), [26](#)
NOOUTALL option
 PROC HPF statement, [11](#)
NOSTABLE option
 IDM statement (HPF), [26](#)
NOTSORTED option
 ID statement (HPF), [21](#)

O

OUT= option
 PROC HPF statement, [11](#)
OUTEST= option
 PROC HPF statement, [12](#)
OUTFOR= option

PROC HPF statement, [12](#)
 OUTSEASON= option
 PROC HPF statement, [12](#)
 OUTSTAT= option
 PROC HPF statement, [12](#)
 OUTSUM= option
 PROC HPF statement, [12](#)
 OUTTREND= option
 PROC HPF statement, [12](#)

P

PLOT= option
 PROC HPF statement, [40](#)
 PRINT= option
 PROC HPF statement, [12](#)
 PRINTDETAILS option
 PROC HPF statement, [13](#)
 PROC HPF statement, [11](#)

R

REPLACEBACK option
 FORECAST statement (HPF), [16](#)
 REPLACEMISSING option
 FORECAST statement (HPF), [16](#)

S

SEASONALITY= option
 PROC HPF statement, [13](#)
 SEASONTEST= option
 FORECAST statement (HPF), [16](#)
 SELECT= option
 FORECAST statement (HPF), [16](#)
 IDM statement (HPF), [26](#)
 SETMISSING= option
 FORECAST statement (HPF), [18](#)
 ID statement (HPF), [21](#)
 SIZE= option
 IDM statement (HPF), [24](#)
 SORTNAMES option
 PROC HPF statement, [13](#)
 START= option
 ID statement (HPF), [22](#)
 STARTSUM= option
 PROC HPF statement, [13](#)

T

TRANSFORM= option
 FORECAST statement (HPF), [18](#)
 IDM statement (HPF), [25](#)
 TRENDPARM= option
 IDM statement (HPF), [26](#)
 TRENDREST= option
 IDM statement (HPF), [26](#)

U

USE= option
 FORECAST statement (HPF), [18](#)

Z

ZEROMISS= option
 FORECAST statement (HPF), [19](#)
 ZEROMISSING= option
 ID statement (PROC HPF), [22](#)

Your Turn

If you have comments or suggestions about *SAS® High-Performance Forecasting 9.1 User's Guide*, please send them to us on a photocopy of this page or send us electronic mail.

For comments about this book, please return the photocopy to

SAS Publishing
SAS Campus Drive
Cary, NC 27513
E-mail: **yourturn@sas.com**

For suggestions about the software, please return the photocopy to

SAS Institute Inc.
Technical Support Division
SAS Campus Drive
Cary, NC 27513
E-mail: **suggest@sas.com**