Estimated Variance of Seasonally Adjusted Series

William P. Cleveland *

February 2002

Keywords: seasonal adjustment, signal extraction

Abstract

For model-based seasonal adjustment, there are explicit formulas for obtaining the variance of the seasonal factors or the seasonally adjusted series. For series adjusted with X-11 or X-12, variance estimates are generally based on a linear approximation of the seasonal adjustment procedure. The work of Pfeffermann (1992) extends earlier work by Wolter and Monseur. This study uses simulated series and comparisons of alternative seasonal adjustment results for a few economic series to assess the accuracy of variance estimates. Pfeffermann's method gives good results when the true seasonal is centered and follows a fairly smooth evolution from year to year. Comparisons with formula-based computations and estimates from the TRAMO-SEATS programs by Maravall and Gomez show the latter can give good variance results for series adjusted with X-11 even if the seasonal factors themselves differ from X-11 factors.

*Industrial Output Section, Federal Reserve Board, Washington, DC 20551. e-mail wcleveland@frb.gov.
The opinions expressed herein are the author's and do not necessarily represent the Board of Governors of the Federal Reserve System or its staff.

1 Introduction

The question of variances for seasonally adjusted series has been addressed in a number of contexts. In the case of model-based seasonal adjustment, there are explicit formulas which can be applied. For series adjusted with X-11 or X-12, the discussion usually takes place in the context of linear approximations to the X-11 estimator. The work of Pfeffermann (1992) extends earlier variance estimates advanced by Wolter and Monseur. An assumption that X-11 produces an unbiased estimate of the true seasonal is required by Pfeffermann, and the method should be used with settings specified for autocorrelations identified in the sampling error of the original series. This study uses simulated series and comparisons of alternative seasonal adjustment results for a few economic series to assess the accuracy of variance estimates. Pfefferman's method gives good results only when the true seasonal is centered and follows a fairly smooth evolution from year to year. Estimates from his method are compared with model-based estimates computed directly with formulas and estimates obtained from the TRAMO-SEATS programs by Maravall and Gomez.

2 Models for Simulated Series

We start with a standard, additive, three component model for a seasonal series,

$$y_t = p_t + s_t + e_t \tag{1}$$

where y_t is the observed series indexed by month t, p_t is the trend or trend-cycle component, s_t is the seasonal component, and e_t is a noise component. The noise is assumed to be uncorrelated unless otherwise indicated. Each of these components has an ARIMA representation.

The trend component will be specified using the lag operator B as

$$(1 - B)(1 - .8B)p_t = a_t \tag{2}$$

which gives the highly autocorrelated trend usually associated with economic time series. The shocks a_t are iid.

Several models will produce highly seasonal series. Three will be used to illustrate which are most appropriate for studying variances of seasonal estimators. Since equation 1 is additive, all seasonal simulations are initialized using a seasonal pattern with mean zero. Consider

$$(1 - B^{12})s_t = b_t \ ,$$

where b_t is white noise. Simulated series from this model have a seasonal with expanding amplitude over time. The model

$$(1 - .95B^{12})s_t = b_t \tag{3}$$

is used to generate the first set of series. It also gives strongly seasonal series, but they approach a nearly constant amplitude. Simulations using this model can have a seasonal mean which differs significantly from zero. Thus, the seasonal generated may differ appreciably from estimation results from X-11.

A closely related model is

$$f_{11}(\phi_s B)s_t = b_t, \tag{4}$$

where $f_{11}(B) = 1 + B + B^2 .. + B^{11}$ or $(1 - B)f_{11}(B) = (1 - B^{12})$. This model does not have the near unit root at frequency zero, so the average of the seasonal pattern over a year remains near zero. In simulations, a value $\phi_s = .995735$ was used, so that $\phi_s^{12} = .95$.

Finally, the seasonal model

$$(1 - \phi_{s2}B^{12})f_{11}(\phi_s B)s_t = b_t \tag{5}$$

is like equation 4, but has smoother, more autocorrelated transitions in the seasonal from year to year. A value of .6666 was used for ϕ_{s2}. The measure sacf12 will indicate the autocorrelations at lag 12 of $(s_t - s_{t-12})$.

The noise model in all cases will be simply

$$e_t = e_t \tag{6}$$

until we have need for more structure. Simulated series were generated with FAME software using 3, 4, and 5 as the seasonal parts of time series models.

3 Series Descriptors

In order to characterize the series realized using the different models and relate them to each other, some basic measures are required. The average values of a seasonal pattern are obtained from a 2 x 12 moving average filter, as in X-11. The mean of the absolute values of the averaged seasonal patterns for a set of series will be called scenter, the degree to which the simulated seasonal patterns are centered on zero. Since X-11 routinely centers its seasonal estimate, one cannot expect small deviations of the seasonal estimates from the true seasonal unless the true seasonal is also centered. The amplitude of the simulated seasonal will be measured by its variance over a number of complete years, svar. In order to have comparable results from the three models, the trend model and the expected value of svar are kept the same for all simulations. The rate at which the seasonal pattern changes is measured by the variance of $s_t - s_{t-12}$, sdvar.

The fundamental measure of the accuracy of the seasonal estimate will be $mean(s_t - \hat{s}_t)^2$, its mean squared error called here smse. The lower limit of this measure is affected by the character of seasonal estimators relative to the true seasonal. It will be smaller when the estimator has a close relationship to the data generation processes for the series components. In practice, the seasonal data generation process can only be known to the extent that it is consistent with the observed series. Both the contributions of the fundamental variance in the seasonal estimator and the underlying misspecification of the seasonal component in the estimator are important aspects of the overall mean squared error. Whether using model-based estimators or X-11, the estimator is some form of moving average of the observed series. The estimates obtained may have smoother, more autocorrelated changes in the seasonal for a given month than were present in the true seasonal. This study suggests smooth changes in the true seasonal must be assumed for Pfeffermann's variance estimates to be accurate. The bias measure

$$sbias = \underset{t}{mean}(abs(\underset{i}{mean}(s_{it} - \hat{s}_{it})))$$

reflects both whether the true seasonal and its estimate have the same average value, and the ability of the estimator to follow the changes in the true seasonal. The index i runs over the number of simulated series. One would expect the mean of $(s_{it} - \hat{s}_{it})$, the seasonal estimate errors for each t, to approach zero for a higher number of simulations if an estimator is capable of replicating the true seasonal.

4 Variances

Model-based analysis begins with the two-component decomposition

$$y_t = n_t + s_t \ , \tag{7}$$

where

$$n_t = p_t + e_t. \tag{8}$$

The variance of the seasonal component estimate of the seasonally adjusted series is given in Cleveland and Tiao (1976) as

$$\left[\Sigma_s^{-1} + \Sigma_n^{-1}\right]^{-1}. \tag{9}$$

The estimator for the seasonal component is

$$E[s|y] = \left[\Sigma_s^{-1} + \Sigma_n^{-1}\right]^{-1} \Sigma_n^{-1} y \tag{10}$$

In this study both the theoretical value from 9 and the smse value obtained by applying 10 to the simulated series are presented, using computations performed in SAS IML. SAS computations were done using matrices of dimension 11 years. The TRAMO-SEATS programs may also be used to give a variance estimate like 9. However, these programs model the observed series directly and then compute implied component models. This process will not generally lead to the same models used in the simulations. It turns out that the variance estimates from SEATS frequently agree quite well with those obtained from 9 and with the smse computed using seasonal factors from SEATS. This suggests that a reasonable variance estimate for a series seasonally adjusted with X-12 may be obtained by using TRAMO-SEATS on the same series, even if the adjustments are not quite the same.

5 Results

All three model specifications used equation 2 for the trend with $\sigma_a^2 = .036$. Models 1 - 3 in Table 1 use the seasonal models in equations 3 to 5, respectively. The noise variances for σ_b^2 were .1, .05, and .0064. Expanding the covariance generating functions for the seasonal models gives a variance of 1 for s_t in each case. For the models 2a and 3a, e_t had variance 0.3, while e_t had variance 0.1 for models 2b and 3b. For model 3c, e_t had an MA(1) model with variance .3 and covariance .15. Sixteen series were simulated for each model condition. New noise terms were used for all three components in each simulation. While sixteen is not enough to achieve asymptotic results, the standard deviations of the measures showed this to be enough to make the desired distinctions. The numbers in parentheses are the standard deviations of the 16 mean square estimates (smse) obtained. Most of the entries in the left column of Table 1 have been defined. The entry "var mod" is the central value from equation 9, while "smse mod" refers to the results of equation 10 compared with the true seasonal. Similarly, "var SEATS" is the average of the variance estimates given by the SEATS program, and "smse SEATS" is the average mean squared error of the seasonal estimates computed by the TRAMO-SEATS programs. Series were simulated for the period 1970 through 1999, and seasonal adjustment runs in X-12 and Tramo-SEATS were from 1977 through 1996. Only the values from 1982 through 1993 were used for mse computations to correspond with the model-based calculations carried out in SAS and eliminate end effects. Of course, variances are larger at the end of a seasonally adjusted series. Computations of variance estimates using Pfeffermann's procedure assumed a correlation at lag 1 for the sampling error. This would generally raise the estimate over using lag 0 and this assumption might be used as a precaution where the correlations are unknown.

Table 1: Simulation Results

	Model 1	Model 2a	Model 3a	Model2b	Model3b	Model 3c
scenter	.201	.011	.007	.011	.007	.008
sbias	.084	.061	.039	.055	.030	.040
svar	.899	1.007	.953	1.007	.953	.948
sdvar	.107	.109	.023	.109	.023	.023
sacf12	-.016	-.018	.545	-.018	.545	.545
var Pfef	.064 (.011)	.063 (.009)	.053 (.008)	.031 (.003)	.021 (.003)	.058 (.010)
smse X-12	.166 (.070)	.100 (.018)	.060 (.019)	.072 (.011)	.032 (.009)	.065 (.025)
var mod	.176	.081	.049	.048	.027	.055
smse mod	.158 (.067)	.083 (.018)	.053 (.019)	.048 (.010)	.029 (.011)	.059 (.023)
var SEATS	.100 (.012)	.110 (.021)	.088 (.016)	.082 (.023)	.073 (.022)	.084 (.016)
smse SEATS	.171 (.071)	.098 (.021)	.057 (.019)	.073 (.022)	.045 (.018)	.090 (.054)

Notes:
$scenter = \underset{t}{mean}(\underset{i}{mean}(abs(\underset{t}{mave}12(s_{it})))),\ sbias = \underset{t}{mean}(abs(\underset{i}{mean}(s_{it} - \hat{s}_{it})))$
$svar = \underset{i}{mean}(var(s_{it})),\ sdvar = \underset{i}{mean}(var(s_{it} - s_{i,t-12}))$
$sacf12 = \underset{i}{mean}(cor(s_{it} - s_{i,t-12}, s_{i,t-12} - s_{i,t-24}))$
$smse(\hat{s}_{it}) = \underset{i}{mean}(mse(\hat{s}_{it} - s_{it}))$

As expected, the measure scenter is relatively large for Model 1, where the true seasonal is not centered. The large smse values for the three estimates of the seasonal pattern also reflect the not-centered true seasonal. Of course the variance estimated by SEATS under Model 1 is a small value like that of Pfeffermann, reflecting the centered seasonal assumption by both. Note that the "var mod" calculation from equation 9 is close to the actual smse. More generally, the "smse mod" values from equation 10 using the optimal filters are consistently smallest and agree well with the theoretical variance from the SAS calculations, as they should.

The most important finding in this table is the relatively small variance estimate by Pfeffermann's method in Model 2. Only with the autocorrelated seasonal changes of Model 3 do his variance estimates have good correspondence to the smse of the X-12 method. Whether true seasonals in actual economic series are more like those of model 2 or model 3 is a matter of opinion. The estimated seasonal variance from Pfeffermann's method for Model 3c with an MA1 error is quite good, responding well to the correlated error. The variance estimates produced by SEATS are good for Model 2, but high for models 3b and 3c. The true variance of the seasonal estimate is likely to lie between

Pffermann's estimate and that of SEATS. Model 3b is close to models which have been advanced as having estimators close to the default linear X-11 estimator. The smse from X-12 is smallest for this model and not far from the Pfeffermann variance estimate.

This result may not be strong enough for a statistical agency using X-12 to publish either the Pfeffermann or the SEATS variance estimate as the variance of the seasonally adjusted series, but it gives the data analyst an idea of how sensitive decisions should be to a particular seasonal adjustment.

6 Analysis of Selected Series

To get a feeling for the implications of estimated seasonal factor variances for seasonally adjusted economic series, five aggregate series were analyzed, with the results in Table 2. As measured by the variation in the seasonal compared to the variation in the X-12 irregular, the seasonal patterns range from strong to fairly mild. In general, the agencies publishing these series do not adjust them as aggregates but as a sum of adjusted components, so the analyses here do not correspond to the seasonal adjustment procedures used. However, the seasonally adjusted series obtained here using Tramo-SEATS and X-12 are as close to the published ones as they are to each other, giving some assurance that the variance estimates are reasonable. In some situations the real focus of attention is on monthly growth rates or ratios. Variances of log differences depend on the correlation of adjacent log level estimates. If this correlation is .5, then the level and growth rate variances are the same. Results contained in the output of the SEATS program suggest correlations at least this high. Entries in the table are 100 times the standard deviations of log measures. They can be interpreted as one standard deviation of percent error in level or (assuming .5 correlation of adjacent errors) growth rate.

Estimates of the standard deviation of the seasonal factors from Pfeffermann's method and from SEATS are given in the third and fourth rows of the table. They track each other pretty well. For the CPI and M2 series, these deviations are relatively large compared with the standard deviation of the seasonal pattern. These series also have the most noise in relation to the seasonal pattern. As an additional benchmark, the root mean square of the differences (RMSD) between the two estimated seasonal patterns for each series are presented in row 5. These values compare fairly closely with the standard deviations of the factor estimates, though somewhat smaller for CPI and M2 where the seasonals have less variation. To check out the impact on growth rates implied by these numbers, standard deviations of month-to-month ratios of the seasonally adjusted series were computed. If these are large compared to the standard deviations of the seasonal factor estimates, then errors in the seasonal factors would not distort growth rates much. The ratios of the values in the last row to those in the third row are about 3.5, except a 4.5 for CPI. Using a model where the variance of $\hat{s}$ is a component of the variance of the estimated seasonally adjusted ratios, a ratio of 3.5 implies the correct sign for the true growth rate about 95 percent of the time.

Table 2: Estimated Variances for Selected Series

	IP	Retail	CPI	Labor	M2
St. Dev $\hat{s}$	1.50	7.91	0.15	1.07	0.29
St. Dev $\hat{I}$	.37	1.25	.10	.14	.17
St. Dev. Pfeff	.22	.64	.07	.08	.10
St. Dev. SEATS	.27	.55	.08	.08	.13
RMSD	.34	.84	.04	.07	.05
St. Dev. SAR	.74	2.12	.32	.27	.36

Notes to the table:

IP = Industrial Production index, Federal Reserve Board
Retail = Retail Sales, Commerce Department
CPI = Consumer Price Index (all urban, all items), Labor Department
Labor = Civilian Labor Force, Labor Department
M2 = M2 Index of Money Supply, Federal Reserve Board

St. Dev $\hat{s}$ = standard deviation of the estimated seasonal pattern in logs (amplitude of the seasonal pattern)

St. Dev $\hat{I}$ = standard deviation of the X-12 irregular in logs

St. Dev. Pfeff = seasonal factor standard deviation estimates from Pefferman's method, percent

St. Dev. SEATS = seasonal factor standard deviation estimates from the TRAMO-SEATS program, percent

RMSD = root mean square of the difference between the Tramo-SEATS and X-12 estimates of the seasonal factors

St. Dev. SAR = 100 times the standard deviation of the seasonally adjusted month-to-month ratios.

REFERENCES

1. Bell, W.R., and Hillmer, S.C. (1983), "Modeling time series with calendar variation," Journal of the American Statistical Association, 78, 526-534.

2. Bell, W.R., and Hillmer, S.C. (1984), "Isses involved with seasonal adjustment of economic time series,", Journal of Business and Economic Statistics, 2, 291-349.

3. Box, G.E.P., and Jenkins, G.M. (1970), Time Series Analysis: Forecasting and Control, Holden Day: San Francisco.

4. Burman, J.P. (1980). "Seasonal adjustment by signal extraction, "Journal of the Royal Statistical Society, Series A, 143, 321-337.

5. Burridge, P. and Wallis, K. F. (1985). "Calcualting the variance of seasonally adjusted series," Journal of the American Statistical Association, 80, 541-52.

6. Cleveland, W.P. and Grupe, M.R. (1983). "Modeling time series when calendar effects are present," Applied Times Series Analysis of Economic Data, ed. A. Zellner, U.S. Department of Commerce, Washington, D.C., 57-67.

7. Cleveland, W.P. and Tiao, G.C. (1976). "Decomposition of seasonal time series: a model for the X-11 program," Journal of the American Statistical Association, 71, 581-87.

8. Dagum, Estela Bee (1980). The X-11-ARIMA Seasonal Adjustment Method, Statistics Canada, Ottawa.

9. Gomez, V. and Maravall A. (1996). Programs Tramo and Seats, Banco de Espana, Documento de Trabaho 9628.

10. Hillmer, S.C. and Tiao, G.C. (1982). "An ARIMA model-based approach to seasonal adjustment," Journal of the American Statistical Association, 77, 63-70.

11. Pfeffermann, D. (1994). "A general method for estimating the variance of X-11 seasonally adjusted estimators," Journal of Time Series Analysis, 15, 85-116.

12. Shiskin, J., Young, Alan H., and Musgrave, J.C. (1967). "The X-11 variant of the census method II seasonal adjustment program," Technical Paper 15, Bureau of the Census, Washington, D.C.

13. Wolter, K. M. and Monsour, N. J. (1981)."On the problem of variance estimation for a deseasonalized series," in Current Topics in Survey Sampling: An International Symposium (eds. D. Krewski, R. Platek and J. N. K. Rao), New York: Academic Press, 367-403.