

Seasonal Adjustment of Weekly Time Series with Application to Unemployment Insurance Claims and Steel Production

William P. Cleveland¹ and Stuart Scott²

Seasonal adjustment of weekly data poses special problems because the data are not exactly periodic. The workhorse programs X-12 ARIMA, TRAMO/SEATS, and STAMP, are not suitable. Harvey, Koopman, and Riani (1997) introduced a structural model in which the seasonal component is modeled nonparametrically via periodic splines. Pierce, Grupe, and Cleveland (1984) captured a deterministic seasonal component using regression on trigonometric series within an ARIMA framework. The method advanced here uses the same trigonometric components, but adopts a locally weighted regression to capture changing seasonality. The method is illustrated with unemployment insurance claims data published by the U.S. Bureau of Labor Statistics and steel production data. It is being used successfully for these series and for weekly money supply series at the Federal Reserve.

Key words: Unobserved components; weekly data.

1. Introduction

Weekly time series provide timely updates for government and private observers of the economy while they wait for the release of monthly or quarterly economic indicators. The weekly unemployment insurance (UI) claims series produced by the U.S. Department of Labor and money supply series from the Federal Reserve are prominent examples.

These series are compiled for weeks ending on a given day of the week, Saturday in the examples used here. The position of Saturdays within a year varies from year to year, and they may occur 52 or 53 times. This violates the basic periodic time series structure assumed by X-12-ARIMA, TRAMO/SEATS, and STAMP. The SABL method of Cleveland, Dunn, and Terpenning (1978) transforms weekly data to create a series of period 52 and applies robust versions of the seasonal and trend smoothers of X-11. The Kalman filter methods of Gersch and Kitagawa (1983) also assume a fixed number of periods, but could be extended to add multiple regressions and their corresponding hyperparameters.

¹ Industrial Output Section, Federal Reserve Board, Washington DC, 20551, U.S.A. Email: wccleveland@pobox.com

² U.S. Bureau of Labor Statistics, Washington, DC, 20212, U.S.A. Email: scott_s@bls.gov

Acknowledgments: The opinions expressed herein are the authors' and do not necessarily represent those of the Board of Governors of the Federal Reserve System or the U.S. Bureau of Labor Statistics or their staffs. Helpful information on unemployment insurance claims was provided by Tom Stengle, U.S. Department of Labor.

A locally weighted least squares procedure is suggested here, which can be used with a weekly design matrix having 52 or 53 observations in a year. Weighting schemes to achieve localized parameter estimates are determined by signal extraction formulas in conjunction with appropriate objective functions. The method is currently being applied at the U.S. Bureau of Labor Statistics, the Federal Reserve, and the Bank of Canada.

Computational efficiency is important for local weighting procedures, as there is potential for doing a complete regression for each data point. A ten-year design matrix for weekly data may be 522 by 64 (ten years of weeks by 60 trig functions plus slope and outliers), requiring a substantial amount of computing for each regression. Cleveland, Devlin, and Grosse (1988) solve this problem in a more general context. They establish a procedure for selecting points at which to do the regressions and interpolating in between. The method suggested here takes advantage of the periodicity of the seasonal design matrix X . The value of $X'X/n$ is the same for any number of complete years of monthly data containing n months, and is very nearly so for five or more complete years of weekly data. This implies that only one computation of the inverse of $X'X$ is needed to estimate the dynamic regression. The computation can begin at any point within a year, so little flexibility is lost in using complete years of data.

Closely related work appears in Harvey and Koopman (1993), Harvey, Koopman, and Riani (1997), and Poirier (1973). They use state-space models in conjunction with periodic spline smoothing to achieve a similar result. The spacing of the spline knots and the values of the variance ratios have significant implications for the results. Holiday effects may evolve with their model, but the overall procedure will be judged by some as more complex.

2. The Method

We analyze data for weeks ending on Saturday, but a similar line of argument follows for weeks ending on other days. The peculiar feature of weekly data is that Week 1 of a year can be associated with a range of dates from January 1 to January 7, the date of the first Saturday in our case. Further, if Week 1 falls on January 1, there will be 53 weeks in the year rather than 52. This is illustrated in the seasonal factor plot of Figure 1. The solid line is the daily interpolation of weekly factors for part of a year estimated with a model that keeps the factors for a given day of the year the same each year. The boxes show factors for the year 2002. The circles show factors for the year 1999, which would appear quite different without the interpolated values.

The analysis begins with a regression model for a series y , which is the observed series after suitable transformation and detrending. Series used in this article were logged and differenced. Thus, we model y as consisting of a seasonal component and error:

$$y = X\beta + e \quad (1)$$

For the seasonal component of year s , we employ trigonometric variables with fundamental frequency $1/365$ (or $1/366$):

$$X_s(t, 2j - 1) = \sin(2\pi ij/365), \quad X_s(t, 2j) = \cos(2\pi ij/365) \quad (2)$$

where $i = i(s, t)$ is the day of the year s on which week t ends and $j = 1, 2, \dots, p/2$. We choose p sufficiently large to capture the dynamics of the seasonal pattern. The index t runs




Fig. 1. Shifting Day-of- Year of Weekly Seasonal Factors

from 1 to n_s , which is 52 or 53. We stack the yearly matrices into an overall design matrix in X in levels. Now let X represent the first difference of this matrix. Assuming X is defined for K complete years, it has dimension $n \times p$, where $n = \sum_{s=1}^K n_s$. To achieve a weighted regression, we employ an $n \times n$ diagonal weight matrix W and apply the standard solutions:

$$\hat{\beta} = [X'WX]^{-1}X'Wy \tag{3}$$

$$\hat{y} = X[X'WX]^{-1}X'Wy \tag{4}$$

The regression parameter estimates $\hat{\beta}$ minimize $(y - X\beta)'W(y - X\beta)$. The term of (4) requiring an inverse corresponding to a three-year series may be expanded as

$$[X'_1 \quad X'_2 \quad X'_3] \begin{bmatrix} w_1 I_{n_1} & & \\ & w_2 I_{n_2} & \\ & & w_3 I_{n_3} \end{bmatrix} \begin{bmatrix} X_1 \\ X_2 \\ X_3 \end{bmatrix} = \sum_s w_s X'_s X_s \tag{5}$$

For $\sum w_s = 1$ and identical X_s matrices (e.g., for monthly data), $X'WX$ simplifies to $X'_s X_s$. Even though the X_s are different for weekly series, the $X'_s X_s$ are identical for years with 52 weeks. Years with 53 weeks are infrequent enough so we can evaluate Equation (5) simply as $X'X$. Also, we have

$$X'Wy = w_1 X'_1 y_1 + w_2 X'_2 y_2 + w_3 X'_3 y_3 \tag{6}$$

Looking back at (3), we see $\hat{\beta}$ is a weighted sum of regression coefficients for individual years. Our estimated seasonal component for year s becomes

$$\hat{y}_s = X_s [X'X]^{-1} \sum w_i X'_i y_i \tag{7}$$

Use of identical weights w_s would correspond to the results in Pierce, Grupe, and Cleveland (1984). To allow for moving seasonality, we can apply the above method once

for each year, choosing a weight matrix W geared to that particular year and using the results only for that year's seasonal component. Let I_p represent a $p \times p$ identity matrix, and let w_{ij} be the weight applied to year j to estimate year i factors. With three years, we may write

$$\begin{bmatrix} \hat{y}_1 \\ \hat{y}_2 \\ \hat{y}_3 \end{bmatrix} = \begin{bmatrix} X_1[X'X]^{-1} & & \\ & X_2[X'X]^{-1} & \\ & & X_3[X'X]^{-1} \end{bmatrix} \begin{bmatrix} w_{11}I_p & w_{12}I_p & w_{13}I_p \\ w_{21}I_p & w_{22}I_p & w_{23}I_p \\ w_{31}I_p & w_{32}I_p & w_{33}I_p \end{bmatrix} \begin{bmatrix} X'_1y_1 \\ X'_2y_2 \\ X'_3y_3 \end{bmatrix} \tag{8}$$

The choice of the w_{ij} for each year is to be determined. Years close to the year being estimated should get most weight. We use a formula from signal extraction theory (see Cleveland and Tiao (1976) and references). The seasonal factors for a given week of the year (or month for monthly data) are assumed to follow an autocorrelated random walk. The detrended data are this seasonal part plus white noise, which means no autocorrelations in the detrended, seasonally adjusted series of lag one year. Given the model

$$y_t = u_t + e_t \tag{9}$$

$$(1 - B)(1 - \phi B)u_t = a_t \tag{10}$$

with white noise terms e_t and a_t , the weights to estimate u_t given y form the desired W^* matrix. These are obtained from

$$E[u|y] = \left[I + \nu \sum_u^{-1} \right]^{-1} y = W^*y \tag{11}$$

where $\nu = \sigma_e^2 / \sigma_a^2$ and \sum_u is the autocorrelation matrix of u . The values w_{ij}^* are the weights w_{ij} in Equation (8). The weights are more concentrated (distant years have less effect) for smaller values of ϕ and ν , but the pattern is much more sensitive to changes in ν . Two examples of W^* for a series of length 9 years are given in Table 1 and Table 2. The rows are labeled for the year being estimated and contain the weights for that year (bold) and adjacent years. Note that the last rows show symmetric weight patterns. For seasonal series with more noise, it makes sense to use more data to extract an estimate of the seasonal signal. With $\nu = 10$, the first three years provide more than 80 percent of the weight for year 1; with $\nu = 24$, they provide about two-thirds of the weight.

Table 1. Year Weights for Given ϕ and ν

	$\phi = .5$		$\nu = 24$						
Year									
1	0.285	0.226	0.165	0.116	0.078	0.052	0.035	0.024	0.018
2	0.226	0.215	0.174	0.130	0.092	0.064	0.044	0.031	0.024
3	0.165	0.174	0.177	0.148	0.113	0.083	0.060	0.044	0.035
4	0.116	0.130	0.148	0.160	0.138	0.109	0.083	0.064	0.052
5	0.078	0.092	0.113	0.138	0.155	0.138	0.113	0.092	0.078
Phase shift									
	2.001	1.319	0.788	0.366	0.000	-0.366	-0.788	-1.319	-2.001

Table 2. Year Weights for Given ϕ and ν

	$\phi = .5$	$\nu = 10$							
Year									
1	0.384	0.271	0.169	0.096	0.049	0.023	0.008	0.002	-0.002
2	0.271	0.264	0.197	0.127	0.074	0.039	0.019	0.007	0.002
3	0.169	0.197	0.218	0.170	0.114	0.068	0.037	0.019	0.008
4	0.096	0.127	0.170	0.203	0.163	0.111	0.068	0.039	0.023
5	0.049	0.074	0.114	0.163	0.200	0.163	0.114	0.074	0.049
Phase shift									
	1.254	0.706	0.358	0.146	0.000	-0.146	-0.358	-0.706	-1.254

Tables 3 and 4 show the weights for X-11 seasonal adjustment with 3×5 and 3×9 seasonal filters, respectively (cf. Shiskin, Young, and Musgrave 1967). While there is a rough correspondence with the filters from the signal extraction formula, the signal extraction weights exhibit exponential decay and tend to concentrate on the year being estimated more.

The results of these operations might be termed a seasonal kernel regression, with the shape of the kernel and smoothing parameter determined by ν (or ϕ and ν if ϕ is allowed to vary). Use of the signal extraction formula automatically supplies correct kernel shapes at the ends of the series. The problem of choosing ν remains. In kernel or smoothing spline regressions, the smoothing parameter is optimized by minimizing some sort of penalized residual sum-of-squares, or a cross-validation technique, Silverman (1984). As suggested in Härdle et al. (1988), convergence is slow and the surface rather flat. As there is no simple objective criterion for a seasonally adjusted series, we have chosen to create sets of seasonal factors for a grid of ν and p values and compute the smoothness of the resulting seasonally adjusted series using concurrent and projected factors and the size of revisions when new data are added. Thus, part of our evaluation is like the approach in Grillenzoni (1994) of minimizing one-step-ahead projection errors. We would expect the methods set forth in Harvey, Koopman, and Riani (1997) to be capable of similar results, but do not have the software to make a direct comparison possible.

3. Unemployment Insurance Claims

Among closely watched economic series is Initial Claims from the Unemployment Insurance program. Claims data come from administrative records collected from individual government employment offices across the U.S., first assembled at the state level, and then forwarded to the Department of Labor (DOL), Washington, DC. Each

Table 3. Year Weights for 3×5 Seasonal Filter in X-11

Year									
1	0.283	0.283	0.283	0.150	0.000	0.000	0.000	0.000	0.000
2	0.250	0.250	0.250	0.183	0.067	0.000	0.000	0.000	0.000
3	0.150	0.217	0.217	0.217	0.133	0.067	0.000	0.000	0.000
4	0.067	0.133	0.200	0.200	0.200	0.133	0.067	0.000	0.000
5	0.000	0.067	0.133	0.200	0.200	0.200	0.133	0.067	0.000

Table 4. Year Weights for 3×9 Seasonal Filter in X-11

Year									
1	0.246	0.221	0.197	0.173	0.112	0.051	0.000	0.000	– 0.000
2	0.208	0.192	0.176	0.160	0.144	0.092	0.028	0.000	0.000
3	0.173	0.163	0.154	0.143	0.133	0.123	0.079	0.032	0.000
4	0.141	0.137	0.132	0.128	0.123	0.117	0.113	0.075	0.034
5	0.111	0.111	0.111	0.111	0.111	0.111	0.111	0.111	0.111

weekly report reflects activity from Sunday through Saturday. Initial Claims (IC) are reported the following Thursday and Continuing Claims a week later. With less than a week's time lag, Initial Claims are clearly one of the most timely indicators of the state of the economy. Press releases and many media accounts contain both weekly figures and less volatile 4-week moving averages. Especially during times of greater economic uncertainty, the claims data have received more than usual attention.

We estimate the seasonal component using 60 variables, sine and cosine terms for 30 seasonal frequencies, $2\pi i/365$, $i = 1$ to 30. This is enough to capture periodic effects as short as two weeks. With holiday effects and outliers, the total number of regression terms is below 90, not excessive with data spans of length 600 and above. The choice of 30 frequencies was made by examining reductions in residual sums of squares for varying sets of frequencies. Table 5 shows the reductions from adding to the number of frequencies (nf) in groups of 6 from 6 to 30. The residual sums of squares (ss) after fitting a nonseasonal ARIMA model to the seasonally adjusted series are shown in the second row and the mean square differences (msd) from adding frequencies in the next. The F values are shown in the final row.

The number of degrees of freedom associated with a weighted regression is $\text{tr}(X(X'\Omega X)^{-1}X'\Omega)$, where Ω is a diagonal matrix containing the weights. For our application this equals the trace of the annual weight matrix partially illustrated in Tables 1–2. This is in agreement with calculations described in Zhang (2003, Section 3) for varying coefficient models. The residual F tests justify our choice of 30 frequencies. The 95 percent significance value for the F test is below 2, so all the additions are highly

Table 5. Ratios of Seasonal Factors, Post-Thanksgiving to Thanksgiving Weeks

End-date	Week	Year	r_T
11/24	47	2001	1.44
11/25	48	2000	1.45
11/25	47	1995	1.43
11/26	48	1994	1.42
11/27	48	1999	1.41
11/27	48	1993	1.40
11/28	48	1998	1.36
11/29	48	1997	1.31
11/30	48	2002	1.28
11/30	48	1996	1.26

Table 6. Residual SS Reductions from Adding Frequencies for Initial Claims

	6	12	18	24	30
nf	6	12	18	24	30
SS	3,020	1,581	1,204	944	773
msd	na	120	31.4	21.6	14.2
F	na	60	15.7	14.4	10.7

significant. Holiday and outlier variables were included when selecting the number of frequency components to avoid biasing the result toward using higher frequencies.

The next step was to select ν , with the fixed-weight regression regarded as a very large value of ν . Table 6 contrasts use of a fixed-weight regression with $\nu = 16$ in local weighting. Series root mean square differences at lags 1, 4, 13, and 26 were computed over the series spans ending in 1999, 2000, and 2001; these differences are used as indicators of residual seasonality. It is clear that the corresponding values for $\nu = 16$ are all smaller, particularly for monthly and quarterly differences. Within each column the values get larger for longer lags, as longer lags capture more trend. Corresponding results for the year $K + 1$ (when projected factors would apply) show smaller differences between methods, suggesting that concurrent adjustment would be best. As one would expect, factor revisions as new data come in are smaller for the fixed-weight regression model. A value of $\nu = 10$ gave only slightly smoother results in the current period and larger numbers for revisions and year-ahead smoothness statistics.

We now present graphical evidence of the improvement. Figure 2 illustrates the difference between the locally weighted and fixed regression seasonal factors. At various points in the graph we see more smoothness with local weights (solid) than with fixed weights (dashed). Note in particular that in years 2000, 2001, and 2002 the use of local weighting eliminates a sharp rise and drop around the middle of the year (present with fixed weighting) and makes the December-January period smoother. The spurious mid-year rise and fall helped spur the investigation into adopting local weighting. Part of the change in the summer claims filing pattern can be explained by a change in annual model change-over practices in the automotive industry.

Another view of the changing seasonality is given in Figure 3, which shows the evolution of the first 24 trigonometric coefficients characterizing the seasonal pattern. For each frequency component there is a plot showing changes in its coefficient over time about a horizontal line giving its mean value. For example, the first cosine coefficient is

Table 7. Effect of Local Weighting for Initial Claims on Root Mean Square Differences

	Fixed weights			$\nu = 16$		
	1999	2000	2001	1999	2000	2001
Year	1999	2000	2001	1999	2000	2001
1st diff	14.6	15.1	15.6	12.8	13.0	12.9
4th diff	20.6	20.7	21.4	17.0	17.1	17.5
13th diff	28.8	29.2	30.0	25.0	24.8	26.1
26th diff	35.1	35.1	35.1	32.8	32.8	34.9

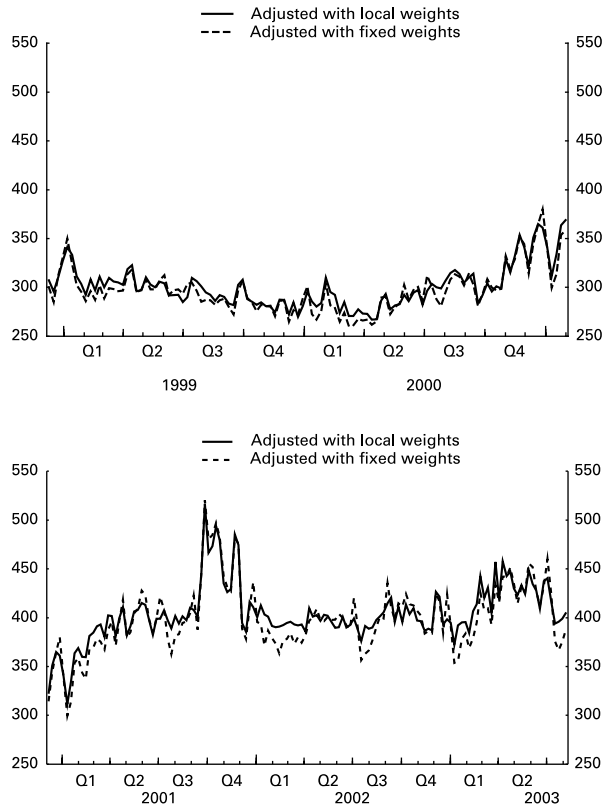


Fig. 2. Seasonally Adjusted Initial Claims

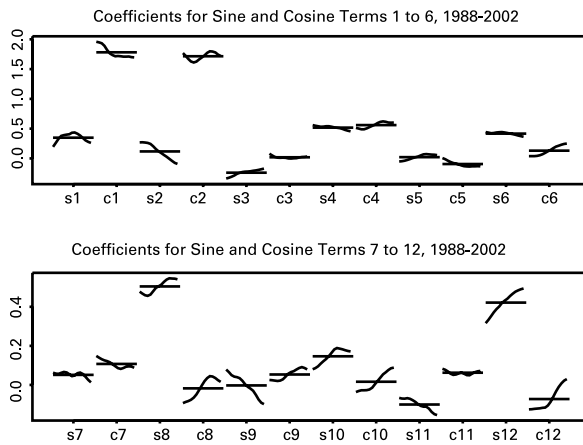


Fig. 3. Evolution of Trigonometric Coefficients for Initial Claims

near 2 initially and then declines to about 1.7. A number of the coefficients show marked changes over time.

Every government holiday is potentially significant for this series because it represents loss of a day for filing claims. All require special treatment, as they affect different weeks in different years. Thanksgiving can occur in either week 47 or 48. New Year's Day is by definition in Week 1. When it falls on Sunday, corresponding to a late Week 1, its effects are split between Weeks 1 and 2; otherwise, its effects come in Week 2. While not a Federal holiday, Easter affects claims activity and moves between Weeks 12–17. Special variables are introduced for occurrence of July 4 on Wednesday and for particular patterns related to Christmas. For examining holiday effects, one may create year-over-year plots of not seasonally adjusted series and seasonally adjusted series not accounting for that holiday.

Thanksgiving in particular illustrates the need to consider holidays along with the day of the year. In addition to the Thursday holiday, Friday is likely to be a lighter than average day for initial claims. In compensation, the following week will have increased activity. Table 7 contains ratios of seasonal factors $r_T = sf_{T+1}/sf_T$, where T denotes the week containing Thanksgiving and sf_T is the seasonal factor for that week. These are arranged by end-date of Thanksgiving week. The ratios decrease monotonically as the end-date advances (with 2001 being a slight exception). The ratios are all well above 1, but range from 1.26 to 1.45, a substantial difference.

Outlier and intervention specification was an iterative process involving (1) automatic outlier detection with X-12-ARIMA, using residuals from the fixed regression method, and (2) outlier significance testing with the fixed regression program. Economic effects of the September 11, 2001 attacks varied in timing and magnitude, depending on the type of business. After testing several alternatives, eight weeks between September and November 2001 were classified as additive outliers (AO's). All had t -statistics between 2.9 and 5.0. An important criterion for their selection was not biasing the seasonal pattern existing prior to 9/11/01. The dashed lines in Figure 4 plot seasonal factors for the last half of 2000 with no outlier treatment for the post-9/11 period and with treatment using the 8 AO's. Seasonal factors based on an input span through 8/27/01 are shown as a solid line. Especially in late September and October we see better agreement between the solid and dashed lines using the outlier treatment.

4. Steel Production

Raw steel production data from the American Iron and Steel Institute are received weekly at the Federal Reserve. Figure 5 shows raw steel production since 1995. The 1995 and 1996 patterns reflect earlier patterns when production tended to be high toward the end of the first quarter and then decline through the year's end. In later years the third quarter is relatively higher and production falls off more sharply at the end of the year. Periods of shutdowns at the end of years 2000 and 2001 are evident. The seasonal pattern shift can be seen in the chart of seasonal factors, Figure 6. This shows the daily pattern of the trigonometric expansion using the locally estimated values for the years shown.

Table 8 gives some statistics on residuals from seasonally adjusted series, after modeling remaining autocorrelations, for various choices of numbers of trigonometric




Fig. 4. Effect of Post-9/11/01 Specification on Weekly Seasonal Factors for Initial Claims

components keeping ϕ fixed at 0.5 and $\nu = 20$. The root mean square differences of the seasonally adjusted series for corresponding numbers of frequencies are given in Table 9. The top half of Table 9 gives values for the entire series through the ending year T . The lower half of the table gives values for year $T + 1$, corresponding to projected




Fig. 5. Weekly Production of Raw Steel

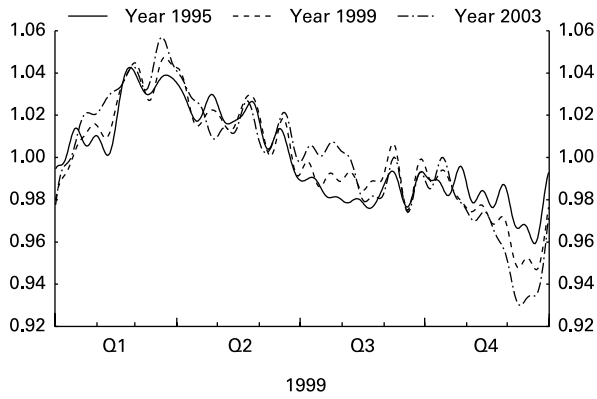


Fig. 6. Seasonal Factors for Raw Steel

Table 8. Residual SS Reductions from Adding Frequencies for Steel

nf	6	12	18	24	30
SS	6,973	6,114	5,585	5,049	4,587
msd	na	47.7	29.3	29.7	25.6
F	na	6.8	4.2	4.9	4.5

factors. We chose $nf = 24$ as getting sufficiently smooth series without too much evidence of overfitting suggested by increasing values in the lower half of the table. Most of the seasonality at monthly or lower frequencies is removed with only 24 trigonometric components.

Table 9. Smoothness Measures with Increasing Numbers of Trig Components for Steel

	Year = T				
nf	6	12	18	24	30
1st diff RMS	52	50	47	44	42
4th diff RMS	74	69	68	66	66
13th diff RMS	108	107	105	105	104
26th diff RMS	156	155	154	154	153
	Year = T + 1				
nf	6	12	18	24	30
1st diff RMS	63	65	67	69	73
4th diff RMS	110	109	110	111	112
13th diff RMS	142	142	142	143	143
26th diff RMS	164	162	160	160	159

5. References

- Bell, W.R. and Hillmer, S.C. (1983). Modeling Time Series with Calendar Variation. *Journal of the American Statistical Association*, 78, 526–534.
- Box, G.E.P. and Jenkins, G.M. (1970). *Time Series Analysis: Forecasting and Control*. Holden Day: San Francisco.
- Burman, J.P. (1980). Seasonal Adjustment by Signal Extraction. *Journal of the Royal Statistical Society, Series A*, 143, 321–337.
- Cleveland, W.P. and Grupe, M.R. (1983). Modeling Time Series When Calendar Effects Are Present. *Applied Time Series Analysis of Economic Data*, A. Zellner (ed.). U.S. Department of Commerce, Washington, DC, 57–67.
- 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, 581–587.
- Cleveland, W.S. (1979). Robust Locally Weighted Regression and Smoothing Scatterplots. *Journal of the American Statistical Association*, 74, 829–836.
- Cleveland, W.S., Devlin, S.J., and Grosse, E. (1988). Regression by Local Fitting: Methods, Properties, and Computational Algorithms. *Journal of Econometrics*, 37, 87–114.
- Cleveland, W.S., Dunn, D.M., and Terpenning, I.J. (1978). SABL: A Resistant Seasonal Adjustment Procedure with Graphical Methods for Interpretation and Diagnosis. With Comments in *Seasonal Analysis of Economic Time Series*, Arnold Zellner (ed.). U.S. Department of Commerce, Washington, DC, 201–241.
- Dagum, E.B. (1980). *The X-11-ARIMA Seasonal Adjustment Method*. Statistics Canada, Ottawa.
- Gersch, W. and Kitagawa, G. (1983). The Prediction of Time Series with Trends and Seasonalities. *Journal of Business and Economic Statistics*, 1, 253–264.
- Grillenzoni, C. (1994). Optimal Recursive Estimation of Dynamic Models. *Journal of the American Statistical Association*, 89, 777–787.
- Härdle, W., Hall, P., and Marron, J.S. (1988). How Far Are Automatically Chosen Regression Smoothing Parameters from Their Optimum? *Journal of the American Statistical Association*, 83, 86–95.
- Härdle, W. (1990). *Applied Nonparametric Regression*. Cambridge, UK: Cambridge University Press.
- Harvey, A. and Koopman, S.J. (1993). Forecasting Hourly Electricity Demand Using Time-Varying Splines. *Journal of the American Statistical Association*, 88, 1228–1236.
- Harvey, A., Koopman, S.J., and Riana, M. (1997). The Modeling and Seasonal Adjustment of Weekly Observations. *Journal of Business and Economic Statistics*, 15, 354–368.
- 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.
- Liu, L.M. (1980). Analysis of Time Series with Calendar Effects. *Management Science*, 26, 106–122.
- Pierce, D.A., Grupe, M.R., and Cleveland, W.P. (1984). Seasonal Adjustment of the Weekly Monetary Aggregates: A Model-Based Approach. *Journal of Business and Economic Statistics*, 2, 260–270.

- Poirier, D.J. (1973). Piecewise Regression Using Cubic Splines. *Journal of the American Statistical Association*, 68, 515–524.
- Shiskin, J., Young, A.H., and Musgrave, J.C. (1967). The X-11 Variant of the Census Method II Seasonal Adjustment Program. Technical Paper 15, U.S. Bureau of the Census, Washington, DC.
- Silverman, B.W. (1984). A Fast and Efficient Cross-Validation Method for Smoothing Parameter Choice in Spline Regression. *Journal of the American Statistical Association*, 79, 584–589.
- Zhang, C. (2003). Calibrating the Degrees of Freedom for Automatic Data Smoothing and Effective Curve Checking. *Journal of the American Statistical Association*, 98, 609–628.

Received April 2005

Revised February 2007