SABL: A RESISTANT SEASONAL ADJUSTMENT PROCEDURE WITH GRAPHICAL METHODS FOR INTERPRETATION AND DIAGNOSIS

William S. Cleveland, Douglas M. Dunn, and Irma J. Terpenning
Bell Laboratories

ABSTRACT

SABL (Seasonal Adjustment-Bell Laboratories) is a new seasonal adjustment procedure for monthly time series with a yearly periodicity. From a class of transformations of the data, SABL selects the one that minimizes the lack of stability of the seasonal component. It then decomposes the transformed time series into additive trend, seasonal, and irregular components using resistant-smoothing techniques. The smoothers consist of an initial resistant smooth using moving medians, followed by one or more linear smooths. Graphical displays are provided as diagnostic tools in assessing the nature and adequacy of the decomposition into trend, seasonal, and irregular components.

INTRODUCTION

A Brief Overview of SABL

This paper has been organized to accommodate various types of readers. Those interested only in a nontechnical summary of the goals and general methods of SABL and its relation to other approaches to seasonal adjustment need read only the first two subsections of the introduction. Those readers who, in addition, are interested in understanding the nature of the techniques used, but who do not want a complete description of details, can also read the third subsection of the introduction. Finally, those interested in the details of SABL can find them in the remaining sections of this paper.

Let us suppose that $Z(t)$, for $t=1, \ldots, N$, is a monthly time series which exhibits yearly seasonal fluctuations and that the goal is to seasonally adjust the series. This means computing values which portray, as well as possible, what the series would have been had the seasonal component not been present; that is, with the seasonal part of the series removed. SABL (Seasonal Adjustment-Bell Laboratories) consists of numerical procedures for carrying out the seasonal adjustment and graphical displays for interpreting the adjustment and for diagnosing problems when they exist.

$Z(t)$ is assumed to be the sum of three additive components

$$Z(t)=\hat{T}(t)+\hat{S}(t)+\hat{I}(t)$$

where $\hat{T}(t)$ represents the long-term trend in the series, $\hat{S}(t)$ represents the more-or-less periodic seasonal part, and $\hat{I}(t)$ represents the irregular or noisy part. SABL computes component estimates that will be denoted $T(t)$, $S(t)$, and $I(t)$ and that satisfy

$$Z(t)=T(t)+S(t)+I(t) \qquad (1)$$

The seasonally adjusted series is then

$$Z(t)-S(t)=T(t)+I(t)$$

The data and the three components of (1), computed by SABL, are shown in figure 1 for the logarithm of U.S. money supply (M1). The vertical bars to the right of the three component plots all represent the same number of units and, thus, provide a comparison of the different scalings on the three plots. Similarly, in figure 2, the data and three components are shown for the logarithm of U.S. manufacturing shipments.

The decomposition is achieved through a series of filtering operations. Deviant observations, even if relatively small in number, can substantially distort the results of filtering operations unless precautions are taken. For a specific decomposition task, SABL tailors recently developed filters that are resistant to deviant observations. The overall approach is to first use resistant filters, such as moving medians, that eliminate the distorting effects of outliers and then to use more standard linear filters on these resistantly smoothed results to achieve the desired degree of smoothness.

When variation that properly belongs to one component, say $S(t)$, is incorporated in the variation of another component, say $I(t)$, we shall say that leakage has occurred. One goal of the seasonal adjustment process is to provide an $S(t)$ for which there is as little leakage as possible.

For seasonal time series, the amplitudes of the seasonal

Figure 1. **LOG MONEY SUPPLY**

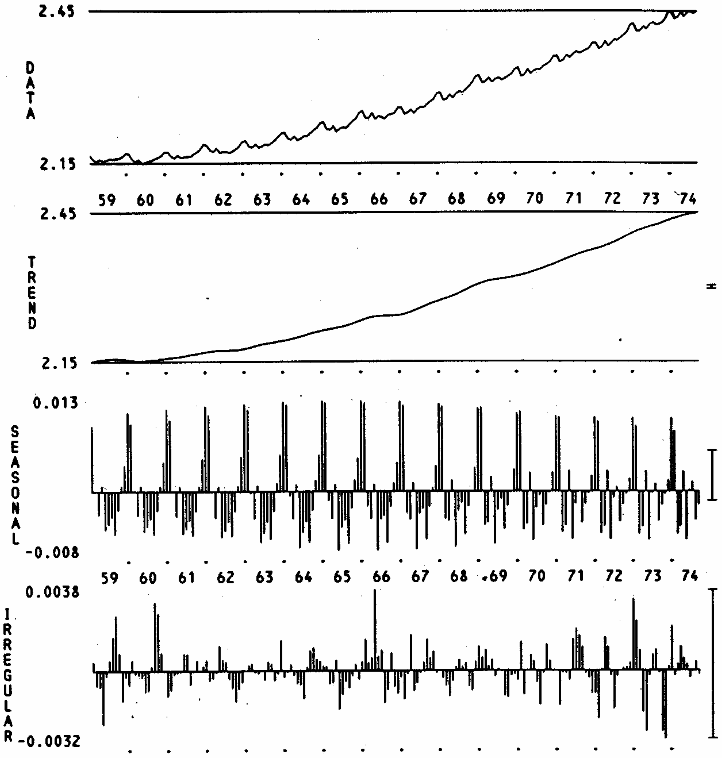

Figure 2. **LOG MANUFACTURING SHIPMENTS**

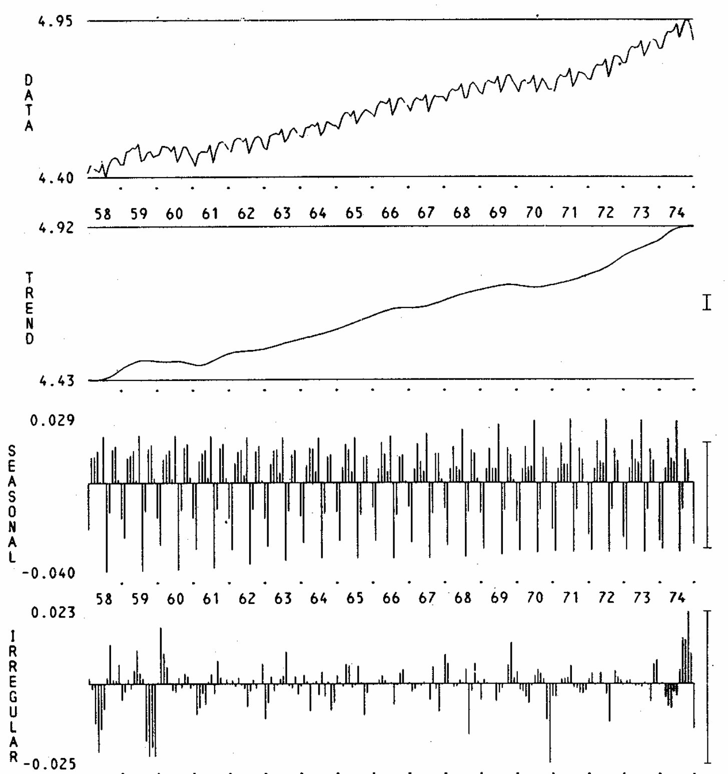

oscillations frequently increase with increasing levels of the trend. This is illustrated in figure 3, which shows the components estimated by SABL for U.S. manufacturing shipments. Such instability is undesirable, since filtering operations perform best when the seasonal is stable and since, more fundamentally, the decomposition in (1) assumes no interaction among the three components. This kind of seasonal instability may be eliminated if the decomposition procedure is applied, not to the raw data, but rather to the data transformed by a power transformation with parameter p, which has the form

$$\begin{array}{ll} (\text{raw data})^p & \text{for } p > 0 \\ \log(\text{raw data}) & \text{for } p = 0 \\ -(\text{raw data})^p & \text{for } p < 0 \end{array}$$

The initial part of the SABL process is an automated procedure to select the power tranformation that minimizes the dependence of the seasonal amplitudes on the trend. For manufacturing shipments the transformation procedure selected the logarithm that, as can be seen by comparing figures 2 and 3, successfully removed much of the lack of stability of the seasonal amplitudes.

In the seasonal adjustment process, it is important to have tools for assessing the adequacy of the adjustment. It is our contention that one of the most effective approaches to assessment is to study graphical displays of the output, $T(t)$, $S(t)$, and $I(t)$. These graphical displays serve a dual purpose. The first is to provide powerful diagnostic aids to assess the adequacy of a three component decomposition. The second is, given that the decomposition is satisfactory, to assess the behavior of each of the three components in determining the makeup of the series. It is important to realize that these graphical methods can be utilized for any three component decomposition, whether based on filtering procedures, such as SABL, or based on a model fit to the series.

It should be emphasized that SABL is still very much in the development stage. While performance for the current version is quite good, changes will undoubtedly be made as a result of continued experimentation with the methodology.

A Comparison of SABL and Other Methods

The single most informative statement we can make to categorize the SABL decomposition procedure among others already in existence is that the philosophy of its overall approach is exactly the same as that used in the X–11 procedure of the Census Bureau [23]. In both cases, filtering operations are utilized to achieve a decomposition in which the tailoring of the decomposition to a particular series is done only through the amount of smoothing of the individual filters. This is in contrast to methods in which a specific model is fit to an individual series.

The differences between SABL and X–11 are in the techniques used to carry out and analyze the decomposition. For example, SABL utilizes resistant filters that, from the start, are not unduly affected by a small number of deviant observations. In contrast, X–11 iterates between nonresistant linear filters and downweighting observations, identified as deviant by examining the irregular. In such an approach, an initial nonresistant linear fit can cause the effect of an outlier to be smeared across several residuals so that good observations appear as outliers [10], or the initial results can be so badly distorted that bad observations do not appear as outliers.

A second major difference between X–11 and SABL lies in the nature of the three component decomposition. X–11 offers two possibilities—an additive procedure and a multiplicative one. SABL offers these two possibilities as a subset of a broader range of possible power transformations of the data.

Finally, SABL provides newly developed graphical methodology, which better enables the user to assess the performance and behavior of the decomposition. The use of any automated statistical procedure without a thorough diagnosis of the results is dangerous. It is only through a critical and penetrating look at the results that one can be assured that they are sensible and that they achieve the desired goals. A detailed comparison of the results of SABL and X–11 decompositions is in progress. An initial comparison, using some of the graphical methods of SABL, is presented for a well-behaved series in the final section of this paper.

The alternative to a general filtering routine, such as SABL, is to fit specific models to individual series. For example, regression models [7; 8; 9; 14; 17; 18; 20] and time series models [13; 19] have been used to model empirically and to seasonally adjust series. While modeling can be successful and provide optimal seasonal adjustment in specific cases [12; 14], any single class of models is unlikely to have the flexibility of a filtering approach. Thus, in a domain where it is necessary to have a decomposition procedure that is easy to apply and that must perform well for many different kinds of series, a filtering approach is likely to be the more desirable. This undoubtedly accounts for the success of X–11 in the past.

However, we do feel that SABL can be useful even in those cases where a modeling approach is the one ultimately utilized. Particularly when the modeling is purely empirical, the output of SABL can be utilized in choosing the form of the model to be fit and to assist in diagnosing possible inadequacies in the fitted model.

A Detailed Summary of SABL

Power transforming the data—It is often the case that the seasonal amplitude is a function of the general level of the series. Since the level is described by the trend, the first step of SABL is to select the power transformation for which the trend and seasonal amplitudes of the transformed series are least related. For example, figures 2 and 3 show that the seasonal amplitudes of the logarithm of manufacturing shipments are more stable than those of manufacturing shipments.

A model which reasonably describes a dependence of

Figure 3. **MANUFACTURING SHIPMENTS**

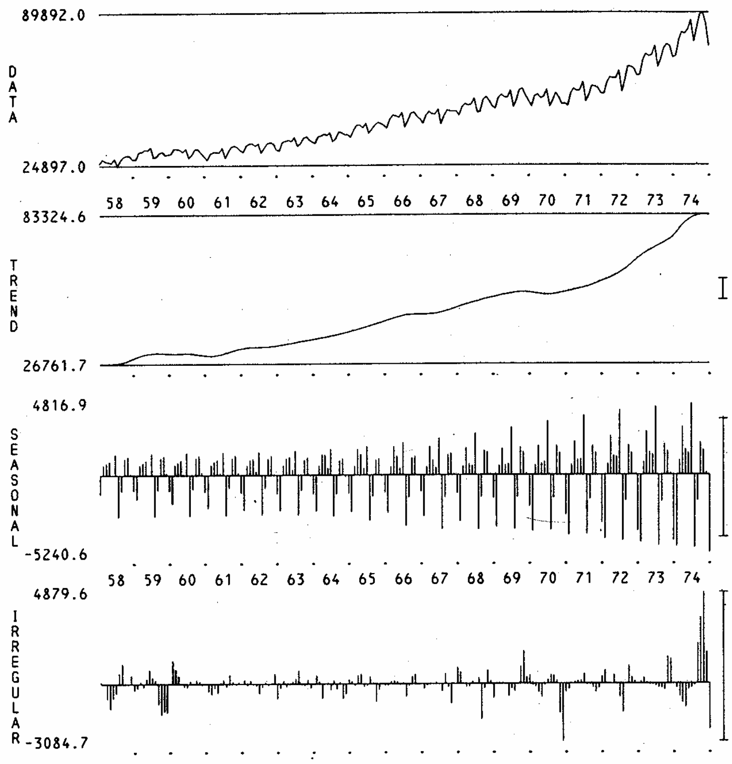

Figure 4. **RESULTS OF SMOOTHING COSINE AND LINEAR FOR FIVE VALUES OF THE SLOPE OF THE TREND**

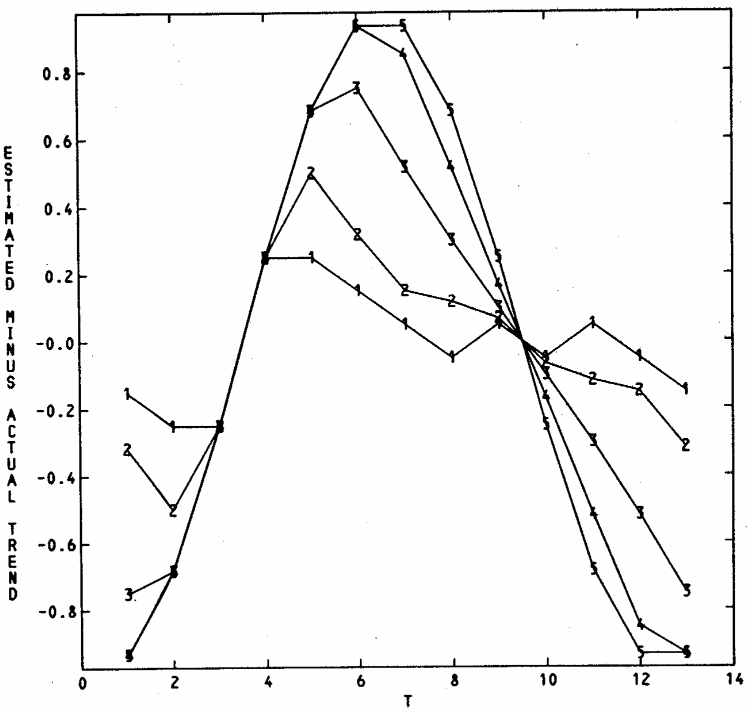

seasonal amplitudes on the trend is the nonadditive model

$$Z(t)=\tilde{T}(t)+\tilde{S}(t)+c(\tilde{S}(t)-\tilde{S}(.))\,(\tilde{T}(t)-\tilde{T}(.))+\tilde{I}(t) \quad (2)$$

where $\tilde{S}(.)$ and $\tilde{T}(.)$ are the means of $\tilde{S}$ and $\tilde{T}$, respectively. We want $\tilde{S}(t)$ to be a component for which the amplitude does not depend on $\tilde{T}(t)$. To ensure such stability, we shall assume $\tilde{S}(t)$ is perfectly stable, so that $\tilde{S}(t) = \tilde{S}(t+12)$.

Here, as elsewhere, we use $Z(t)$ to denote the original series or a power transformation of the original series. For each value of p in the set –1.5(.5)1.5, model (2) is fit to the power transformed data, $Z(t)$. The fit is done by filtering $Z(t)$ to get a trend estimate $T(t)$, computing an estimate, $S(t)$, of $\tilde{S}(t)$ from $Z(t)–T(t)$, and then regressing $Z(t)–T(t)–S(t)$ on $(S(t)–S(.))\ (T(t)–T(.))$ using resistant techniques. The value of p that is selected is the one for which the r^2 for this regression is minimized, i.e., for which the contribution of the product term in the above model is the smallest.

The decomposition into trend, seasonal, and irregular— SABL employs resistant filering operations to carry out the three-component decomposition. Following Tukey [22], these filters are a combination of moving medians that prevent distortion from outliers and the more standard weighted moving averages, which achieve sufficient smoothness. A small fraction of outliers can substantially distort a moving average, whereas the moving median of, e.g., 11, will be resistant, provided no more than 5 outliers occur in any 11 consecutive values of the series.

The details of the SABL decomposition procedure were developed partially as a result of empirical studies on actual and artificial series. In addition, the following overall goals guided the selection of the methods for estimating the trend, seasonal, and irregular:

1. $T(t)$ should reflect the low frequency or trendlike behavior of $Z(t)$.
2. $S(t)$ should reflect the relatively stable behavior that repeats every 12 months.
3. $I(t)$ should contain all behavior that is not seasonal or trendlike.
4. There must be as little leakage as possible from one component into another.
5. Extreme or unusual values of $Z(t)$ should not distort $T(t)$ or $S(t)$. This unusual behavior of $Z(t)$ should be reflected only in $I(t)$.
6. Estimates of trend and seasonal should be as responsive as possible to change in their structure.

The decomposition procedure is iterative in the sense that at each step a new estimate of the trend or seasonal is made, based on the preceding iteration's estimate of the other component. For example, if an estimate, $S(t)$, of the seasonal has just been computed, then an estimate, $T(t)$, of the trend is computed by filtering $Z(t)–S(t)$ with a resistant filter that passes only low frequencies. If an estimate, $T(t)$, of the trend has just been computed, then the estimation of the seasonal is begun by filtering, separately, each of the 12 monthly series of $Z(t)–T(t)$ (i.e., all of the Januarys are filtered as a separate series, then all of the Februarys, etc.). To decrease the danger of trend leaking into this seasonal estimate, a low-pass filter is applied to the estimate, and the result subtracted to form a new seasonal estimate. The only other adjustment in the estimation of the seasonal follows the first seasonal smooth which, for reasons to be described later, has 11 values missing at the end and at the beginning. The missing values are estimated by predicting the seasonal forward and backward.

The stages of the decomposition procedure and their relationship to the transformation and graphical stages of SABL are shown in table 1. The remainder of this section describes these decomposition stages in more detail.

The process begins by computing a trend estimate from $Z(t)$. The design of the trend filter must, in this case, reflect the fact that the seasonal is present in $Z(t)$ and, thus, not allow seasonal to leak into the trend estimate. The procedure is to—

1. Resistantly fit and subtract a line from $Z(t)$ so that an application of moving medians will be more effective.
2. Take moving medians of length 12 of the result of (1).
3. Filter the result of (2) by a weighted moving average that eliminates the variation of seasonal frequencies and passes quadratic polynomials.
4. Add the line fit in (1) to the result of (3). This is the trend estimate, $T(t)$.

The seasonal smooth is now applied to $Z(t)–T(t)$. Each monthly sequence of $Z(t)–T(t)$ is separately smoothed by a succession of median and moving average filters, each of which has length 4, 3, or 2. As examples, one of the resistant filters takes repeated moving medians of length 3 until there is no change, while one of the linear filters is a moving average with weights 1/4, 1/2, 1/4. This same seasonal smooth is also applied at each iteration. Let $S(t)$ denote the output of the seasonal smooth.

Trend can leak into the seasonal estimate unless special precautions are taken. To prevent this, a low pass filter may be applied to $S(t)$ and the result subtracted from $S(t)$ to form the seasonal estimate. Since $S(t)$ has already incorporated resistant filters, this trend smooth of $S(t)$ need not do so. The procedure is to—

1. Add six values to the beginning of $S(t)$ and six values to the end of $S(t)$ by predicting $S(t)$ forward and backward.
2. Filter the result of (1) using a low pass filter which completely removes the seasonal frequencies.
3. Subtract the result of (2) from $S(t)$ to get the detrended estimate of the seasonal.

Once an estimate, $S(t)$, of the seasonal has been computed, the trend can be estimated by applying a low pass filter to $Z(t)–S(t)$. Unlike the first trend smooth, this filter does not need to be designed to completely remove

Table 1. **SUMMARY OF SABL**

Input	Operation	Output
Raw data	Transform the series	$(\text{raw data})^p$ for $p > 0$ log (raw data) for $p = 0$ $-(\text{raw data})^p$ for $p < 0$
Z = transformed data	First trend smooth	T1 = first trend smooth of Z
Z-T1 = raw seasonal	Seasonal smooth	S1 = seasonal smooth of Z-T1
S1	Seasonal trend smooth and removal	S2 = S1 minus a trend smooth of S1
S2	Seasonal component prediction	S3 = S2 together with 11 values added to the beginning and end
Z-S3 = raw trend	Trend smooth	T2 = trend smooth of Z-S3
Z-T2 = raw seasonal	Seasonal smooth	S4 = seasonal smooth of Z-T2
S4	Seasonal trend smooth and removal	S5 = S4 minus a trend smooth of S4
Z-S5 = raw trend	Trend smooth	T = trend smooth of Z-S5
Z-T = raw seasonal	Seasonal smooth	S6 = seasonal smooth of Z-T
S6	Seasonal trend smooth and removal	S = S6 minus trend smooth of S6
Z, T, S	Compute irregular	I = Z-T-S
Z, T, S, I	Plot data and components	Graphical displays for interpretation and diagnosis

the variation at seasonal frequencies. (It is wise, even here, however, to have variations at seasonal frequencies substantially reduced to guard against leakage of seasonal into $Z(t)–S(t)$.) This trend smooth, like the first, must guard against the distorting effect of outliers. The procedure is to—

1. Fit and subtract a line from $Z(t)–S(t)$, so that an application of moving medians will be more effective.
2. Compute a moving median of length 12 of the result of (1).
3. Smooth the result of (2) by a succession of short length median and moving average filters.
4. Smooth the result of (3) by a weighted moving average that passes quadratic polynomials.
5. Apply steps (2), (3), and (4) to the result of (1) minus the result of (4), and add this new smooth to the result of (4).
6. Add the line fit in (1) to the result of (5).

Step (5), which is called twicing [22], is intended to add back effects that are missed by the initial smooth.

Graphical displays—The graphical methods presented in this paper can be used to assess both the nature and the adequacy of a three-component decomposition. Many of the plots are oriented toward checking whether there is leakage of one component into another. While the plots are used to display the results of the SABL decomposition, they can be used for any decomposition procedure.

Boxplots [22] are used in several of the graphical displays of SABL. Figure 16 is an example. The boxplot is a graphical summarization of the distribution of a set of values. The line overprinted with an "x" in the middle of the box shows the center of the distribution, the upper and lower edges of the box portray quantities, that are much like upper and lower quartiles and the ends of the dotted lines portray tails of the distribution. Values that lie beyond the ends of the dotted lines are shown individually.

Table 2 summarizes the diagnostic and interpretive displays of SABL and identifies a figure exemplifying each. In this section, we attempt to give the general flavor of each display. Details are given in the section on graphical methods for assessing the decomposition.

Figures 1–3. The data and each of the three components from SABL are plotted against time to enable appreciation of the overall change in the components through time. However, much detail which is of interest cannot be seen in these plots.

Figures 9–10. The variability of a given component through time is measured by a moving maximum of length 12 minus a moving minimum of length 12. For the trend and irregular components, this measure is smoothed by a low pass filter. The three measures are plotted against time to show how much of the variation in the original series is attributable to each component and how these relative influences change through time.

Figure 11. For each component, the unsmoothed values of the variability measure described for figure 10 are summarized by boxplots. In addition, the midmean [1] of each monthly series of the seasonal component minus the minimum of these 12 midmeans is portrayed. (The midmean of a set of values is defined to be the average of the order statistics between the upper and lower quartiles.) This display shows the overall influence of each of the three components in determining the variation of the series.

Figure 12. The final seasonal smooth in the SABL decomposition filters the raw seasonal, which is equal to the final seasonal plus the irregular. To check the performance of the final smooth, the raw values are plotted together with the smoothed values for each month.

Figure 13. The trend plot in figure 1 does not allow the detail of local variability to be appreciated, because the long-term change in the trend is large, compared with this variability. To enable appreciation of detail, the first, second, and third differences of the trend are plotted against time.

Figure 14. To check the performance of the transformation in stabilizing the seasonal oscillations, a moving maximum of length 12, a moving minimum of length 12, and a low-pass smooth of the absolute values of the seasonal are plotted. To check for trend in the seasonal, a low-pass smoothing of the seasonal is also plotted.

Figure 15. Each of the 12 monthly series of the seasonal is plotted on the same graph. This allows appreciation of the shape and magnitude of the overall seasonal cycle and also displays the magnitude of changes within each monthly series, compared with the overall seasonal amplitude.

Figure 16. Seasonal residuals are defined to be the seasonal component with the midmean of each monthly series subtracted from the series. Boxplots are used to summarize the distribution of each of the 12 monthly series of the seasonal residuals, thus, portraying the relative variability in the 12 series.

Figure 17. The individual yearly cycles are stacked on top of one another to allow identification of individual months.

Figure 18. To allow better appreciation of the change in the yearly cycle through time, the seasonal residuals defined for figure 16 are plotted as in figure 17.

Figure 19. A two-way display is used to portray the seasonal, with the columns as months, the rows as years, and the absolute value of the seasonal proportional to circle area. Negative values are indicated by a slanted line inside the circle.

Figure 20. To allow better appreciation of the change in the yearly cycle through time, the seasonal residuals defined for figure 16 are plotted as in figure 19.

Figure 21. A measure of the autocorrelation in the irregular and the spectrum of the irregular are plotted. These can point out leakage of seasonal into irregular when it occurs.

Figure 22. A variation of the boxplot is used to summarize the distribution of each of the 12 monthly

Table 2. SUMMARY OF DIAGNOSTIC AND INTERPRETIVE PLOTS

Simultaneous Plotting of Trend, Seasonal, and Irregular	
Figure 1.	Data and three components are plotted against time.
Figure 10.	Variability of each of the three components is plotted against time.
Figure 11.	An overall summary of the variability of each of the three components.
Figure 12.	For each month, the seasonal and the seasonal plus irregular are plotted.
Trend	
Figure 13.	First, second, and third differences of trend are plotted against time.
Seasonal	
Figure 14.	A summary of the change of the seasonal amplitudes through time is plotted together with a seasonal trend smooth.
Figure 15.	Each of the 12 monthly series of the seasonal is plotted on the same scale.
Figure 16.	The monthly midmean is subtracted from each monthly series and the distribution of each monthly series of residuals is summarized by boxplots.
Figure 17.	The yearly cycles of the seasonal are stacked and plotted.
Figure 18.	The monthly midmean is subtracted from each monthly series and the residuals plotted as in figure 17.
Figure 19.	The seasonal is portrayed in a two-way display, with the columns as months and the rows as years.
Figure 20.	The monthly midmean is subtracted from each monthly series and the residuals plotted as in figure 19.
Irregular	
Figure 21.	A measure of autocorrelation of the irregular and an estimate of the spectrum of the irregular are plotted.
Figure 22.	The distribution of each monthly series of the irregular is summarized by a variation of boxplots.
Figure 24.	Smoothed moving quantiles of the irregular are plotted against time.

series of the irregular. This allows appreciation of the dependence of the distribution of the irregular on the month of the year and can aid in an appreciation of the magnitude of seasonal leakage into the irregular when it occurs.

Figure 24. The irregular, a moving upper quartile, and a moving lower quartile of the irregular are smoothed and plotted against time. This allows appreciation of changes in the distribution of the irregular through time, and also allows detection of leakage of the trend into the irregular when it occurs.

THE SABL DECOMPOSITION INTO TREND, SEASONAL, AND IRREGULAR

Table 3 gives a summary of the decomposition procedures used by SABL. In this section, each of these procedures will be described in detail.

First-Trend Smooth (FTS)

The first-trend smooth (FTS) is applied directly to $Z(t)$, which is the data or a power transformation of the data. Since $Z(t)=T(t)+S(t)+I(t)$, FTS must be designed to eliminate $S(t)$ and $I(t)$ in order to produce a trend estimate $T(t)$. FTS must also prevent outliers from distorting the estimated trend. If we attempt to guard against outliers and to remove the seasonal by taking moving medians of length 12 of $Z(t)$, two difficulties arise. The first is that the moving median of length 12 will not in general completely remove the seasonal in the presence of a trend component, even if the seasonal component is perfectly stable. To see this point clearly, suppose $Z(t)$ is made up of a cosine seasonal component plus a linear trend.

$$Z(t)=\cos(\pi t/6)+\alpha t,\ t=0,\ 1,\ 2,\ \ldots$$

The greater the slope, α, the less the seasonal is removed by smoothing with medians of length 12. This is demonstrated in figure 4, which shows the moving median centered at $k+0.5$ minus $\alpha(k+0.5)$ for $k=1, \ldots, 13$, and for $\alpha=0.1(0.1)0.5$. The amplitudes of the curves increase with increasing α. For $\alpha>\pi/6$, $Z(t)$ is monotone, the moving median becomes a moving average of length 2, and the seasonal amplitude is reduced by only 7 percent.

The second difficulty, closely related to the first, is that the outlier-removing capacity of the moving median is reduced for a series with a strong trend component. For example, in a monotonic segment of a series, the moving median cannot decrease the influence of an observation that is unusual, because its increase over the previous observation is relatively small.

Thus, the smoothing capacity of moving medians is greatly reduced if the change in the trend component is large, compared to the variation about the trend. One solution is to fit a smooth function, such as a polynomial, to $Z(t)$ as an initial estimate of the trend, apply smoothing procedures to the residuals, and then take the final estimate of trend to be the smoothed residuals plus the fitted line. Since much of the trend is not present in the residuals, the severity of the two problems, described previously, is not as great.

In carrying out the above fit, we use a linear polynomial, calculated by a resistant procedure called TILT, which is similar to a suggestion in [4]. TILT first divides the independent variable, in this case time, into nine groups, each consisting of approximately $N/9$ consecutive values of the variable. The midmean of the values of the dependent variable, $Z(t)$, in each group, and the midmean of the values of the independent variable are computed. A line is then fit to the resulting nine points using least squares.

The residuals from the line fit by TILT are smoothed first by a moving median of length 12. The results of this resistant smooth will be denoted a_j, $j=6.5,\ 7.5,\ \ldots,\ N-5.5$. Since the smoothing length is 12, a_j contains 11 fewer values than $Z(t)$ and is centered between the original integer time points. a_j is then further smoothed by a weighted moving average

$$\sum_{r=1}^{q} w_r a_{j+r->1}$$

whose weights, w_r, will now be described.

One standard procedure for linear smoothing is to fit polynomials, say quadratics, to successive blocks of points [16]. One property of such a smooth is that quadratic polynomials are preserved by the smooth. We shall retain this property but, in addition, require that a strictly periodic component of period 12 be completely removed. An easy derivation shows that the weights, w_r, can be computed by minimizing

$$\sum_{r=1}^{q} w_r^2$$

subject to the constraint

$$\sum_{r=1}^{q} r^2 w_r = \left(\frac{q-1}{2}\right)^2$$

which ensures that quadratic polynomials are preserved, and also subject to the constraints

$$\sum_j w_{r+12j} = \frac{1}{12}$$

for $r=0, \ldots, 11$, which ensure that seasonals of period 12 are completely removed. The values of w_r can be found using Lagrange multipliers. For the examples in this paper, the value of q was taken to be 24. Since the initial smooth, by moving medians, was also even in length, the linear smooth of a_j, which we denote as b_j, $j=18, \ldots, N-17$, is centered at integer values of time with b_j being centered at time j.

b_j has 23 fewer values than the a_j and 34 fewer values

Table 3. SUMMARY OF THE DECOMPOSITION PROCEDURES OF SABL

First-Trend Smooth (FTS)
- Subtract line fit by TILT
- 12-point moving median
- 24-point weighted moving average with 12-point moving average spliced to ends (QLSS)
- Add line fit by TILT

Seasonal Smooth (SS)
- 4 (3RSR) 2, twice
- Smooth each end value
- Moving average with weights 1/4, 1/2, and 1/4

Seasonal Trend Smooth and Removal (STSR)
- Use SCP to add six values to each end
- QLSS
- 2-point moving average
- Subtract the smooth from the seasonal

Seasonal Component Prediction (SCP)
- Median of three predictions

Trend Smooth (TS)
- Subtract line fit by TILT
- 12-point moving median extended to ends
- 2-point moving average
- 4 (3RSR) 2, twice
- 15-point quadratic smooth
- Twice the result
- Add line fit by TILT

than $Z(t)$ and extends from time points 18 to $N-17$. It is desirable, however, to have the smooth extend as far toward the beginning and end of the series as possible. The shortest length linear smoother (and, therefore, the one that loses the fewest number of values at the beginning and the end) that eliminates a seasonal component of period 12 is an equal-weight moving average of length 12. Let $c_{12}, \ldots, c_{29}$ be the first 18 values of such a smooth applied to a_j. Then $d_{12}, \ldots d_{29}$ is the result of splicing $c_{12}, \ldots, c_{29}$ into $b_{18}, \ldots, b_{29}$ using cosine weights. That is,

$$d_j = c_j \text{ for } 12 \le j \le 17$$

$$d_j = (1-\theta_j)b_j + \theta_j c_j \text{ for } 18 \le j \le 29$$

where $\theta_j = \cos((j-17.5)\pi/12)$. A completely analogous splice is performed at the other end, which results in $d_{N-28}, \ldots, d_{N-11}$. This completes the description of FTS, which has an output of

$$d_{12}, \ldots, d_{29}, b_{30}, \ldots b_{N-29}, d_{N-28}, \ldots, d_{N-11}$$

Since we shall utilize separately the linear portion of FTS (the weighted moving average of length 24 with the moving average of 12 spliced to the ends), we shall refer to it as QLSS (quadratic-linear seasonal stop).

Seasonal Smooth (SS)

The seasonal smoother (SS) is used after the first-trend smooth (FTS) and after each application of the trend smooth (TS). The input to SS is the raw seasonal, calculated as the difference between the data and the current estimate of trend. This raw seasonal is the sum of the smooth seasonal and the irregular components. Thus, the goal of SS is to smooth the raw seasonal in such a way as to eliminate the irregular component. In a large proportion of the series that are seasonally adjusted, the change in the seasonal component from 1 month to the next is not smooth, but each of the 12 monthly series changes slowly from year to year. Hence, a short-length smoother is used to smooth, separately, each of the 12 monthly series of the raw seasonal. This smoother is based on one suggested in [3] that, in turn, is based on ideas in [22].

Let $a_1, \ldots, a_k$ be values of a monthly series of the seasonal. The smoothing procedure consists of the following two steps:

1. Smooth $a_2, \ldots, a_{k-1}$. Denote the smooth by $b_2, \ldots b_{k-1}$.
2. Smooth $a_1, b_2, \ldots b_{k-1}, a_k$.

The reason for treating the end points, a_1 and a_k, in a special way arises from the nature of the trend smooth (TS) used to compute the raw seasonals. The beginning and end values of $T(t)$, e.g., the first and last 12 values, are not as precisely determined as the interior values of $T(t)$, since the effective length of the TS smoother on the interior points is wider than 25 points. The most extreme cases are the end points 1 and N, which can incorporate information only from the right or from the left, whereas the interior points can incorporate information from both sides. We shall first describe the smoothers used in steps (1) and (2) and then describe the end value rules.

The smoother used in step (1) is denoted 4(3RSR)2 twice and consists of the application of a sequence of smoothers of lengths 4, 3, or 2. (The suggestion has been made in [24] that even length-moving medians are preferable to odd.) At each stage, an end value rule is used to extend the smoothed series to the ends. We now describe the procedure as applied to a sequence $c_1, \ldots c_m$ at positions 1 to m. The first smooth of 4(3RSR)2 is a moving median of four, which is denoted in the smoother notation as 4. This yields smoothed values at points 2.5, 3.5, ..., $m-1.5$. End value rule G_1 (to be described later) is used to form smooths at time points 1.5 and $m-0.5$. Then a smoother consisting of medians of three is applied repeatedly until the values do not change. This is denoted as 3R. At each application of the moving median, the smooth at the end point is taken to be the end point itself. After 3R has been applied, end-value rule T_1 is used to smooth the ends. 3R tends to produce local maximums and minimums of length 2 (e.g., 1, 2, 2, 0 or 5, 3, 3, 6). To smooth these maximums and minimums, the following rule is used: Treat the left value of the double maximum or minimum as a right-hand end point and use end-value rule T_1 to smooth it; then, treat the right value as a left-hand end point and use end-value rule T_1 to smooth it. For example, if the sequence is 1, 2, 2, 3, 4, 4, 3, 2, 1, then the left value of 4 is smoothed by applying end value rule T_1 to 2, 3, 4 and the right value of 4 is smoothed by applying end value rule T_1 to 4, 3, 2. After this smoothing of maximums and minimums, called splitting [22], 3R is again used to smooth the sequence. The splitting, followed by 3R, is applied repeatedly until there is no change in the sequence. This is denoted SR.

We have, thus far, applied 4(3RSR) to produce smoothed values at points 1.5, 2.5, ..., $m-0.5$. The result is now further smoothed by a moving average of length 2, denoted in the smoother notation as 2, which produces smoothed values at points 2, 3, ..., $m-1$. End-value rule G_2 is used to provide smooths at points 1 and m. The result is the output of the smoother 4(3RSR)2. Now, twicing [22] is used to recover some of the effects that might have been missed by 4(3RSR)2. Recall that the sequence being smoothed is $c_1, \ldots, c_m$. Let $d_1, \ldots, d_m$ be the result of 4(3RSR)2. 4(3RSR)2, twice is equal to d_j plus 4(3RSR)2 applied to $c_j - d_j$.

The smoother used in step (2), which is applied to $a_1, b_2, \ldots, b_{k-1}, a_k$, first smoothes the ends using end-value rule T_1, then applies a moving average with weights 1/4 1/2, 1/4 (denoted by H) to yield smoothed values at points 2, 3, ..., $m-1$. End-value rule T_2 is then used to extend the smooth to the ends.

The end value rules will be described for left-hand end points. Completely analogous rules hold for the right-hand end points. We begin with rules T_1 and T_2. Let the first three values of the sequence be u_1, u_2, u_3. A prediction, u_0, of the sequence one step to the left is achieved by linearly extrapolating u_2 and u_3 to position 0. Thus, $u_0=3u_2-2u_3$. For end value rule T_1, the smooth at position 1 is the median of u_0, u_1, and u_2. For end-value rule T_2, the smooth at position 1 is $\left(\frac{1}{4}u_0+\frac{1}{2}u_1+\frac{1}{4}u_2\right)$.

End-value rule G_1 extends, to position 1.5, the result of applying a moving median of length 4 to $c_1, \ldots, c_m$. Let $e_{2.5}$ and $e_{3.5}$ be the first two values of the smooth. A value, $e_{0.5}$, for position 0.5 is achieved by linearly extrapolating $e_{2.5}$ and $e_{3.5}$ to position 0.5. Thus, $e_{0.5}=3e_{2.5}-2e_{3.5}$. The smooth at position 1.5 is then the median of $e_{0.5}$, c_1, c_2, and $e_{2.5}$.

End-value rule G_2 extends, to position 1, the result of applying a moving average of length 2 to the output of 4(3RSR). Let this output be denoted $f_{1.5}, f_{2.5}, \ldots,$ and let g_2 and g_3 be the results of the moving average of 2 at positions 2 and 3. A value, $g^*_{0.5}$, at position 0.5 is achieved by linearly extrapolating g_2 and g_3 to position 0.5. Thus $g^*_{0.5}=2.5g_2-1.5g_3$. The smooth at position 1 is then $g_1=(g^*_{0.5}+f_{1.5})/2$.

Seasonal Trend Smooth and Removal (STSR)

The iterative process of successively computing trend, then seasonal, is designed to keep either seasonal from entering trend or trend from entering seasonal. Nonetheless, examples show that trend can leak into seasonal, using these techniques as well as others. STSR combats this by applying a trend smoother to $S(t)$, the result of SS, and then subtracting the estimated trend from $S(t)$ to form a new estimate of the seasonal.

This trend-estimation problem is similar to that encountered in applying a trend smoother to $Z(t)$ in that $Z(t)$ and $S(t)$ both contain seasonal variation. However, the estimation problems differ in that, since $S(t)$ is the result of a resistant smoother, nonresistant-linear smoothers may be applied directly to $S(t)$.

The STSR smooth of $S(t)$, $t=1, \ldots, N$ begins by adding six values to the beginning and six values to the end of $S(t)$, using the seasonal-component prediction procedure (SCP) described in the next section. This allows the trend smooth of $S(t)$, which loses six values at each end, to return N values. The extended $S(t)$ is then smoothed by QLSS, the linear smoother of FTS. (See the subsection on first-trend smooth.) A moving average of length 2 is then applied to the result in order to center the output at integer values of time. This smooth is then subtracted from $S(t)$ to form a new estimate of the seasonal.

Seasonal Component Prediction (SCP)

There are two cases where SABL needs to extend (i.e., predict) the seasonal estimate, $S(t)$, forward and backward:

1. The first seasonal estimate that results from applying FTS, then SS, and then STSR has 11 values missing at each end. These are supplied by SCP before applying TS the first time.
2. STSR requires that $S(t)$ be extended by six values at each end.

SCP extrapolates each monthly series, separately, to provide the extended values. For example, suppose the seasonal has been estimated from January 1961 to December 1969 and six additional values are needed at each end: A July 1960 prediction is achieved by extrapolating the July seasonal series one step backward. Since no more than 11 values are needed in either case, each monthly series does not need to be extrapolated by more than one step.

The extrapolation rule for predicting backward will be described. The forward rule is completely analogous. Let a_1, a_2, and a_3 be the first three values of a monthly series at positions 1, 2, and 3. The goal is to provide a value, a_0, for position 0. Each of the following three predictors could serve this purpose:

1. A horizontal extrapolation of a_1, giving the prediction a_1.
2. A linear extrapolation of a_1 and a_2, giving the prediction $2a_1-a_2$.
3. A linear extrapolation using the least squares line fit to a_1, a_2, and a_3, giving the prediction $(4a_1+a_2-2a_3)/3$.

The median of these three predictions is the estimate of a_0.

Trend Smooth (TS)

The trend-smoother (TS) estimates the trend by smoothing the raw trend, $Z(t)-S(t)$, where $S(t)$ is an estimate of the seasonal resulting from SS followed by STSR. Like the first-trend smooth (FTS) TS must be designed to guard against outliers. Unlike FTS, TS does not need to be designed to remove all variation at seasonal frequencies, since most or all of the seasonal will have been removed by the subtraction of $S(t)$. TS does need to be designed to go to the ends of the series in order to have a trend estimate at all time points, $t=1, \ldots, N$.

Like FTS, TS begins by subtracting a line fit by TILT and then smoothing the residuals with a moving median of length 12. Unlike FTS, the moving median is extended to the ends. Let $a_1, \ldots, a_N$ be the values of the residuals. The leftmost value of the moving median of 12 is the median of $a_1, \ldots, a_{12}$, which is centered at position 6.5. The extended-median smooth, at position $k+.5$ for $k=1, \ldots, 5$, is the median of $a_1, \ldots, a_{2k}$. The extension for the other end is completely analogous.

Next, a moving average of 2 is taken to yield smooths $b_2, \ldots, b_{N-1}$. A smooth, b_1, at position 1 is achieved using end-value rule G_2. (See the subsection on seasonal smooth.) This smooth consists of the average of the extended median smooth at position 1.5 and the linear extrapolation to position 0.5 of the moving-average smooth at positions 2 and 3. A smooth at position N is achieved in an analogous manner. These results are now further smoothed by applying the smoother 4(3RSR)2, twice to $b_1, \ldots, b_N$. (See the subsection on seasonal smooth.)

The concept of a trend assumes underlying long-term smooth behavior. Based on this idea, the output of 4(3RSR)2, twice, which will be denoted as $c_1, \ldots, c_N$, is further smoothed by fitting quadratic polynomials to blocks of 15 values. This smoother is similar to the first linear smoother in FTS but does not employ the constraint that completely eliminates variation at seasonal frequencies. Since both of the filters previously used in TS are resistant to outlying observations, a moving average can be safely employed on their output. Let $d_8, \ldots, d_{N-7}$ be the values resulting from the quadratic-polynomial smooth. Thus, d_j consists of the fitted value of the quadratic fit to the block of 15 values of $c_1, \ldots, c_N$, centered at position j. The following procedure is used to extend this smooth to the ends. Let $e_1, \ldots, e_{15}$ be the fitted values of a quadratic-polynomial fit to $c_1, \ldots, c_{15}$. Let $f_1, \ldots, f_{15}$ denote the result of splicing $e_1, \ldots, e_{15}$ into $d_8, \ldots, d_{15}$, using cosine weights. That is,

$$f_j = e_j \text{ for } 1 \leq j \leq 7$$

$$f_j = (1-\theta_j)d_j + \theta_j e_j \text{ for } 8 \leq j \leq 15$$

where $\theta_j = 0.5\cos((j-7.5)\pi/8)+0.5$. A completely analogous splice is performed at the other end, resulting in $f_{N-14}, \ldots, f_N$. The reason for the splice is to allow a smooth transition from e_j to d_j in going from the ends into the interior of the smooth. The final linear smooth is $f_1, \ldots, f_{15}, d_{16}, \ldots, d_{N-15}, f_{N-14}, \ldots, f_N$.

We started with the residuals $a_1, \ldots, a_N$ from the fitted line and successively applied moving medians, 4(3RSR)2, twice, and a linear smooth, which resulted in the smooth $f_1, \ldots, f_N$. The residuals $a_j - f_j$ are now smoothed using this same sequence of three smoothers, and the smoothed residuals are then added to f_j. This is another instance of twicing, which has the intent of recovering from the residuals some effects which might be absent from f_j. These smoothed values are now added to the line fit at the initial stage of TS to produce the output of TS.

Variations in the Amount of Smoothing

The amount of smoothing done at each stage of the SABL decomposition has been chosen on the basis of decomposition results for several time series, including those described in this paper. However, in other cases, more or less smoothing, by any one of the filters, may be desirable. This can be achieved by—

1. Increasing or decreasing the coefficient lengths of the smoothers, which increases or decreases, respectively, the amount of smoothing.
2. Repeatedly applying a smoother, which increases the amount of smoothing.
3. Repeatedly doing twicing, which decreases the amount of smoothing.

The portable software being developed to carry out the SABL operations will provide options for the amount of smoothing.

SELECTION OF A POWER TRANSFORMATION TO REMOVE THE DEPENDENCE OF SEASONAL ON TREND

The Nonadditive Model

When the amplitudes of the seasonal oscillations are a function of trend, the use of a power transformation, described in the first and third subsections of the introduction, has the potential to remove or reduce the lack of stability. In the model (2), the seasonal is made up of a stable seasonal $\tilde{S}(t)$ and an unstable seasonal $c(\tilde{S}(t)-\tilde{S}(.))(\tilde{T}(t)-\tilde{T}(.))$, which has oscillations with amplitudes depending on the level of the trend. It should be emphasized that, while we are prepared to accept this model for the purpose of finding a power transformation, it would typically be too crude for the purpose of performing seasonal adjustment. For seasonal adjustment, SABL allows for more complex changes in the seasonal pattern, described in the section on the SABL decomposition.

Durbin and Murphy [9] have also used a combined additive and multiplicative model. However, rather than attempting to remove the multiplicative part by a transformation, they fit the model and use it as a basis of their seasonal adjustment procedure. Removing multiplicative terms in models by a power transformation has been utilized in other areas of statistics. (See [2; 6]). In these cases, however, the models are completely specified, including assumptions about error terms. In our case, the trend and seasonal are estimated by resistant smoothing operations, rather than by specific parametric forms, in order to have a large amount of flexibility. The probabilistic structure of the error terms is not specified, since it is unlikely that the usual assumptions, such as independence, would hold.

We might have stayed closer to a standard statistical model by treating $Z(t)$ as a month-by-year two-way table and finding the power transformation that minimizes Tukey's one degree of freedom term for nonadditivity [21]. This, however, would have treated the trend as a step function that is constant for 12-month intervals rather than as a continuous process.

Estimation of Trend and Stable Seasonal for the Nonadditive Model

The estimate, $T(t)$, of the trend, $\tilde{T}(t)$, in model (2) is taken to be the result of applying the first-trend smooth, FTS, to $Z(t)$, a power transformation of the data. The stable seasonal, $\tilde{S}(t)$, in the model is estimated by $S(t)$ in the following manner. Let $S(k)$, for $k=1, \ldots, 12$, be the midmean of the k-th monthly series of $Z(t)-T(t)$. That is, $S(k)$ is the midmean of the values $Z(k+12j)-T(k+12j)$ for $j=0, \ldots, n_k-1$, where n_k is the number of times the k-th month occurs. Since the series length is not necessarily a multiple of 12, the n_k are not necessarily equal. The midmean is used, since it provides a resistant location estimate for each of the 12 monthly series. Now $S(t)$, for $t=13, \ldots, N$, is just taken to be a periodic extension of $S(k)$ for $k=1, \ldots, 12$. That is, $S(k+12j)=S(k)$ for $j=1, \ldots, n_k$.

Fitting the Nonadditive Model

The parameter, p, of the power transformation is chosen by doing a linear regression of

$$Y(t)=Z(t)-S(t)-T(t) \text{ on } X(t)=(S(t)-S(.))\,(T(t)-T(.))$$

for a set of values of p and selecting that value of p which minimizes r^2. Since outliers may distort the results of the regression and since errors can have unsymmetric distributions, we employ a resistant regression procedure that applies different weighting schemes to the positive and negative residuals.

The first step in fitting

$$Y(t)=\beta_0+\beta_1 X(t)+\epsilon(t) \tag{3}$$

is to estimate β_0 and β_1 by $\hat{\beta}_0$ and $\hat{\beta}_1$, using TILT (the robust regression algorithm described in the subsection on first-trend smooth). The residuals are

$$e(t)=Y(t)-\hat{\beta}_0-\hat{\beta}_1 X(t)$$

Let s be the median of the positive $e(t)$. Then, each positive $e(t)$ is adjusted by taking

$$e(t)=e(t) \text{ if } e(t) \leq 3s$$

$$e(t)=3s \text{ if } e(t) > 3s$$

The negative residuals are adjusted in a similar manner, using the median of the negative $e(t)$. Adjusted $Y(t)$ are then defined by

$$Y(t)=e(t)+\hat{\beta}_0+\hat{\beta}_1 X(t)$$

Thus, this procedure amounts to a one-step, two-sided huberizing of the residuals. An ordinary least squares regression of $Y(t)$ on $X(t)$ provides the desired r^2 value.

In figure 5 the values of r^2 are plotted for a grid of powers, $-1.5(0.5)1.5$, for each of three series: U.S. money supply (M1), U.S. manufacturing shipments, and AT&T toll revenues. The transformation that minimizes r^2 in each case is—

money supply	–logarithm ($p = 0$)
manufacturing shipments	–logarithm ($p = 0$)
toll revenues	–square root($p = 1/2$)

If the $\epsilon(t)$ in model (3) were independent normal variables with zero mean and constant variance, then the F-distribution could be used to test the significance of the regressions. It is, however, highly unlikely that these assumptions are valid for the $\epsilon(t)$. The operations used to define $Y(t)$ and $X(t)$ will tend to produce $\epsilon(t)$ that are correlated. Furthermore, the $\epsilon(t)$ will often have a skewed distribution. However statistical significance is only one criterion for judging the usefulness of a transformation. We are more concerned with the practical magnitude of seasonal instability than with whether or not the instability is statistically significant. For example, a small value of r^2 for $p \neq 1$ may be highly statistically significant when, for practical purposes, the stability of the seasonal for $p \neq 1$ may be so near to that for untransformed data that the transformation is unwarranted. For this reason, graphical displays, described in the following section, also play a role in the choice of a transformation.

GRAPHICAL METHODS FOR ASSESSING THE DECOMPOSITION INTO TREND, SEASONAL, AND IRREGULAR

Description of Boxplots

In the subsequent sections, it will be useful to have a graphical summary of the sample, or empirical, distribution of a set of numbers. The quantities that will be used to summarize the distribution are the midmean, the semi-midmeans, and the adjacent values. The midmean of a set of values is defined to be the average of the order statistics between the upper and lower quartiles. The upper (lower) semi-midmean is defined to be the midmean of the observations greater (less) than the median. We shall use the midmean and semi-midmeans in place of the more usual median and quartiles. Neither set of statistics is distorted by a small fraction of outliers, but the midmean and semi-midmeans have somewhat greater stability, since they typically involve averaging more order statistics. In addition, they perform better on data with many repeated values (e.g., 2, 2, 3, 3, 3, 5, 5, 6, 6, 6).

Let one step equal 1.5 times the upper semi-midmean minus the lower semi-midmean. The upper (lower) adjacent value is the largest (smallest) observation that is within one step of the upper (lower) semi-midmean. All observations that are either greater than the upper adjacent value or less than the lower adjacent value are referred to as outside.

The distribution is summarized graphically by a boxplot [22], several of which are shown in figure 16. Each boxplot identifies the upper (lower) adjacent value as the upper (lower) dotted horizontal line; the upper (lower) semi-midmean as the upper (lower) edge of the box; and the midmean as the horizontal line inside the box with an "x". Outside values are shown individually. In figure 16, each is identified by year.

Simultaneous Plotting of Data and Components

Plotting the three components against time—connected and vertical line plots—In plotting a time series, $W(t)$, against time, we are plotting a sequence against equally spaced values of the abscissa. In such a situation, two plotting techniques are considered; we may want to use either technique, or both, depending on the needs. The first is to connect successive points of $(t, W(t))$ by straight lines, as in the top plot of figure 6; we shall term this the "connected plot." The second is to portray $(t, W(t))$ by a vertical line, where the other endpoint (t, u) is appropriately chosen, as in the bottom plot of figure 6, where $u=0$. We shall term this the "vertical line plot."

As previously mentioned [5], the vertical-line plot enables one to perceive each individual value, since the line is not a fat character such as "*" and, thus, allows high resolution along the horizontal axis. Figure 6 shows 204 points in each plot. In contrast, the connected plot may not allow sufficient appreciation of individual values. As seen in figure 6, which shows the irregular component, $I(t)$, for log manufacturing shipments, the connected plot does not tell us whether a spike consists of one or many values. The vertical line plot seems more satisfactory for this noisy series, where appreciating patterns of successive values are relatively unimportant.

However, the vertical line plot sacrifices some of the visual perception of connectivity of successive values, and, in some cases, we lose an appreciation of a pattern, defined by a group of values. Such patterns are often appreciated better in the connected plot. Figure 7 illustrates this for log money supply. The pattern is less striking in the vertical line plot, where the highest values of each year protrude and dominate the plot at the expense of the seasonal pattern and the lowest values of each year.

In figure 8, the seasonal component for log manufacturing shipments is displayed by both connected and vertical line plots. In displaying a seasonal component, there may be some advantage to using both techniques. It seems important to display individual values since the seasonal component is not necessarily smooth from t to $(t+1)$ (although it is usually smooth from t to $(t+12)$). However, the vertical line plot underemphasizes values close to zero.

We prefer to use the connected plot to display trend, as in figures 1, 2, and 3, since we want to emphasize the pattern of successive points and since the smoothness in the sequence reduces the need for perceiving the individual values.

In summary, vertical line plots are useful when—

1. We want to appreciate individual values.
2. The sequence being plotted is noisy and regular patterns are not strong or do not need to be appreciated.

Connected plots are useful when—

1. We want to appreciate patterns, defined by groups of successive points.
2. Perceiving individual points is not important.
3. The sequence tends to be smooth.

The change in the variability through time—The variation in each of the three components may be used as a measure of their relative importance in determining the behavior of the series. The X–11 [23] procedure measures the variation of a component as the sum of the absolute values of its first differences. This, however, seems inappropriate for characterizing the variation in the seasonal component, where the amplitudes would seem to be the most important indicator of variation. We, instead, consider how much of the change in the series, over the course of a year, is contributed by each component. Thus, it seems natural to compare the largest value minus the smallest value occurring, during the year, for each component, and this leads to the use of a moving range of length 12 as the measure of variability for each component.

To display the change in the variability through time, the three moving ranges for trend, seasonal, and irregular can be simultaneously plotted. The trend and irregular-moving ranges are smoothed by an altered version of TS (see the subsection on trend smooth), in which the moving median length is changed to 24 and the quadratic smooth length is changed to 21.

Figure 9 shows moving ranges for the log of money supply. The trend increases in influence through time; the seasonal and irregular components are both rather stable with the level of the seasonal well above that for the irregular. Figure 10 shows moving ranges for log manufacturing shipments. Here, the seasonal component contributes the most to the variation of the series.

Comparing overall variability of the three components—The distribution of the values of each of the three moving ranges is summarized by boxplots in figure 11 for log manufacturing shipments. The distribution of the seasonal component is detailed further by plotting the midmean of each of its 12 monthly series minus the minimum midmean. The predominant effect of the seasonal component is seen in figure 10, but the tightness of the distribution of the seasonal moving ranges in comparison to those of the trend and irregular is also clearly seen in figure 11.

Raw and smooth seasonals—The smoothing of the seasonal is a critical factor in the seasonal adjustment process, because the smoothed seasonal values are subtracted from the data to form the seasonally adjusted series. A seasonal smoother should be flexible enough to follow changes in the seasonal pattern while still being resistant to deviant observations. The seasonal smooth in the SABL decomposition filters the raw seasonal, which is equal to the seasonal plus the irregular. To check the performance of the final smooth,

Figure 5. **R-SQUARED VALUES FOR VARIOUS POWERS**

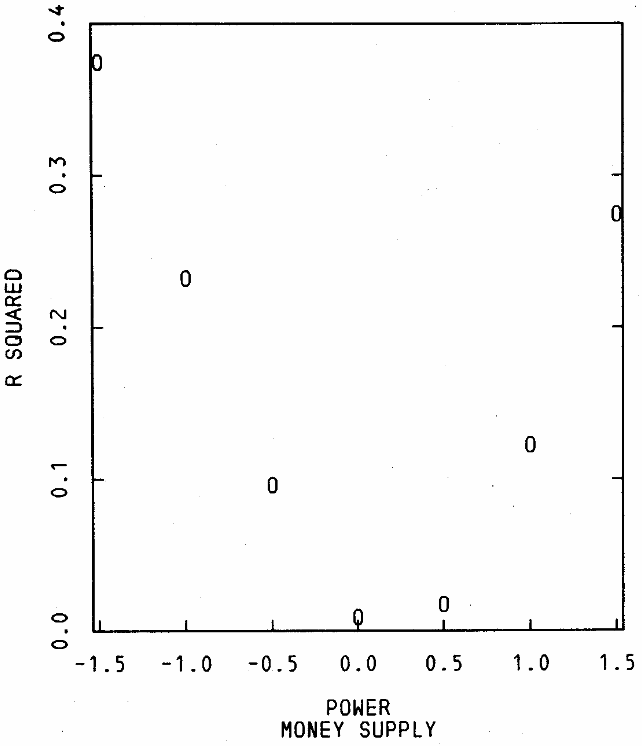

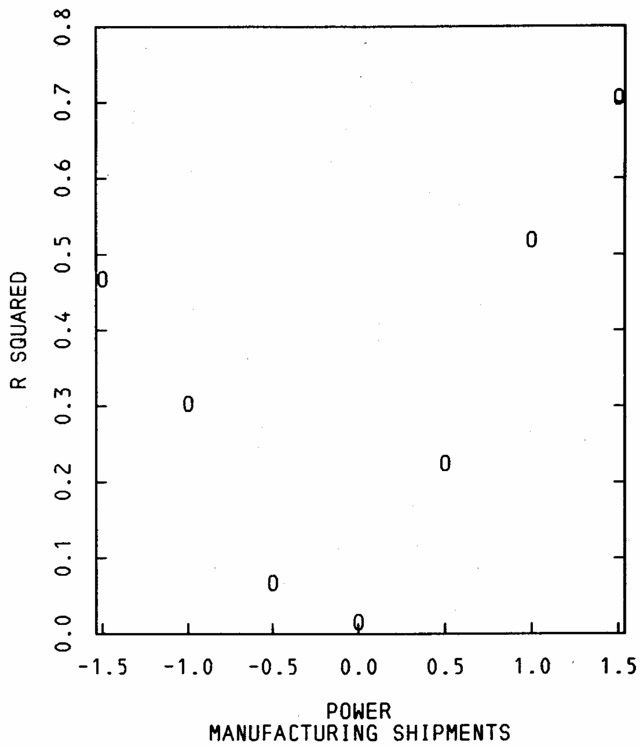

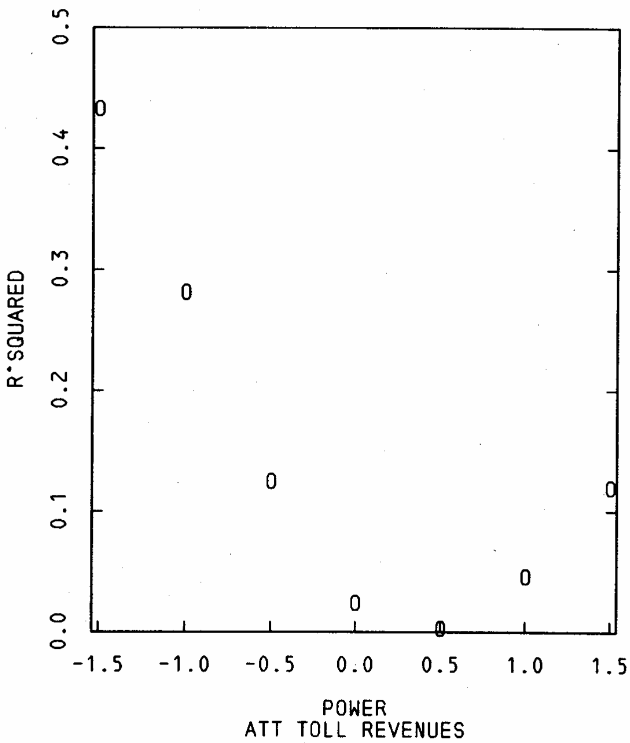

Figure 6. **LOG MANUFACTURING SHIPMENTS, BY IRREGULAR**

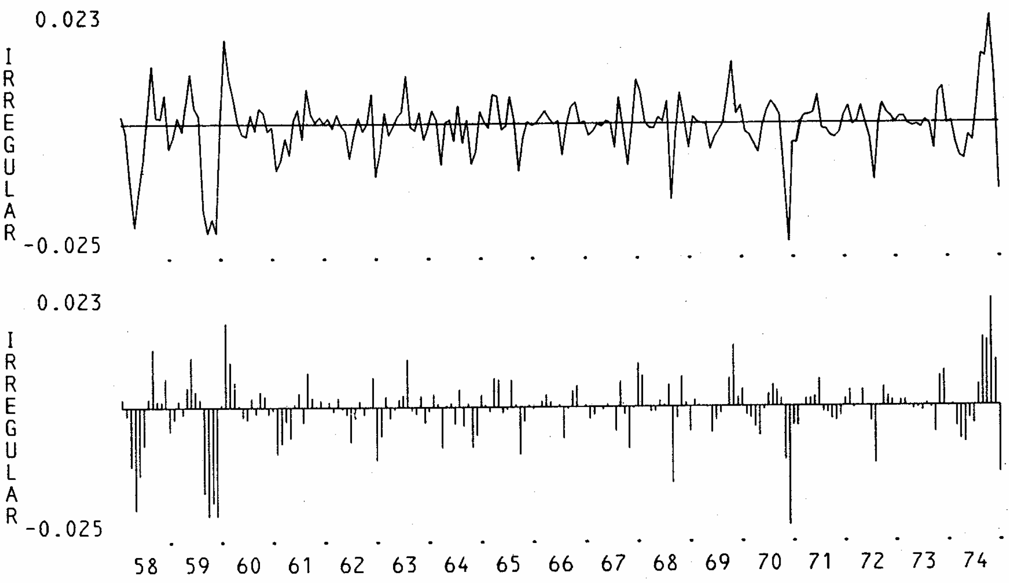

Figure 7. **LOG MONEY SUPPLY, BY DATA**

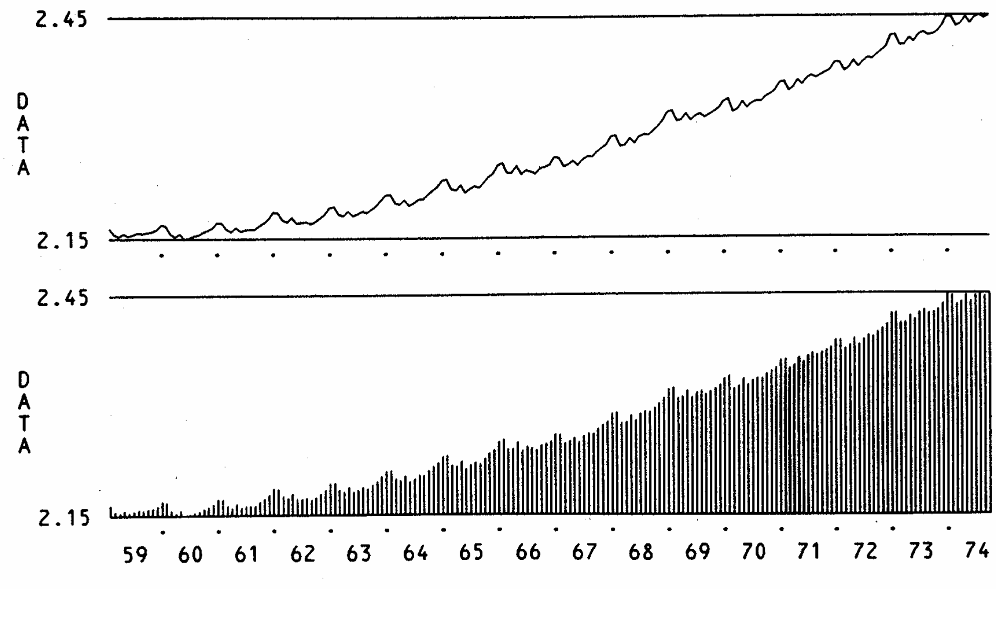

Figure 8. LOG MANUFACTURING SHIPMENTS, BY SEASONAL

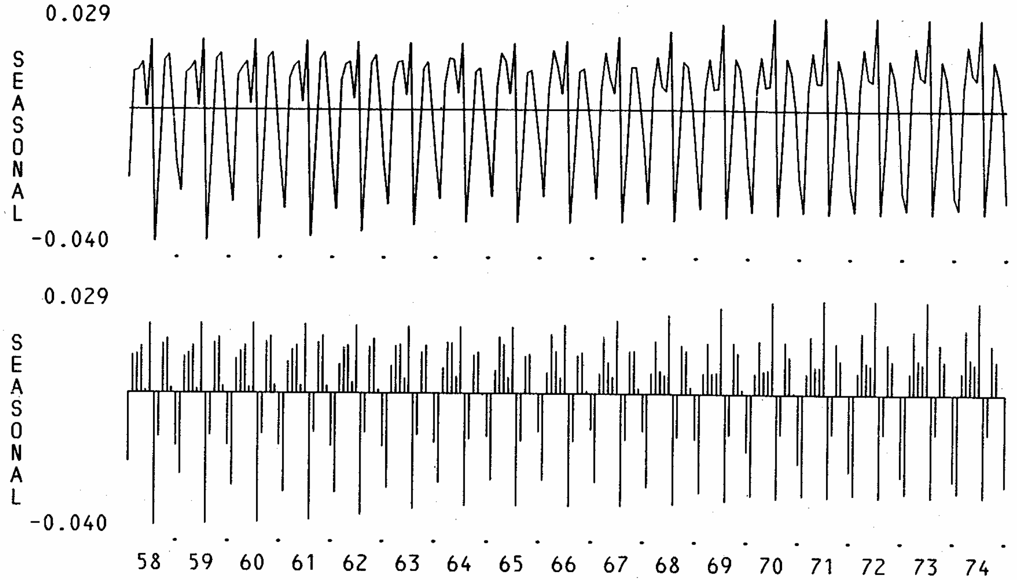

Figure 9. COMPONENTS OF LOG MONEY SUPPLY

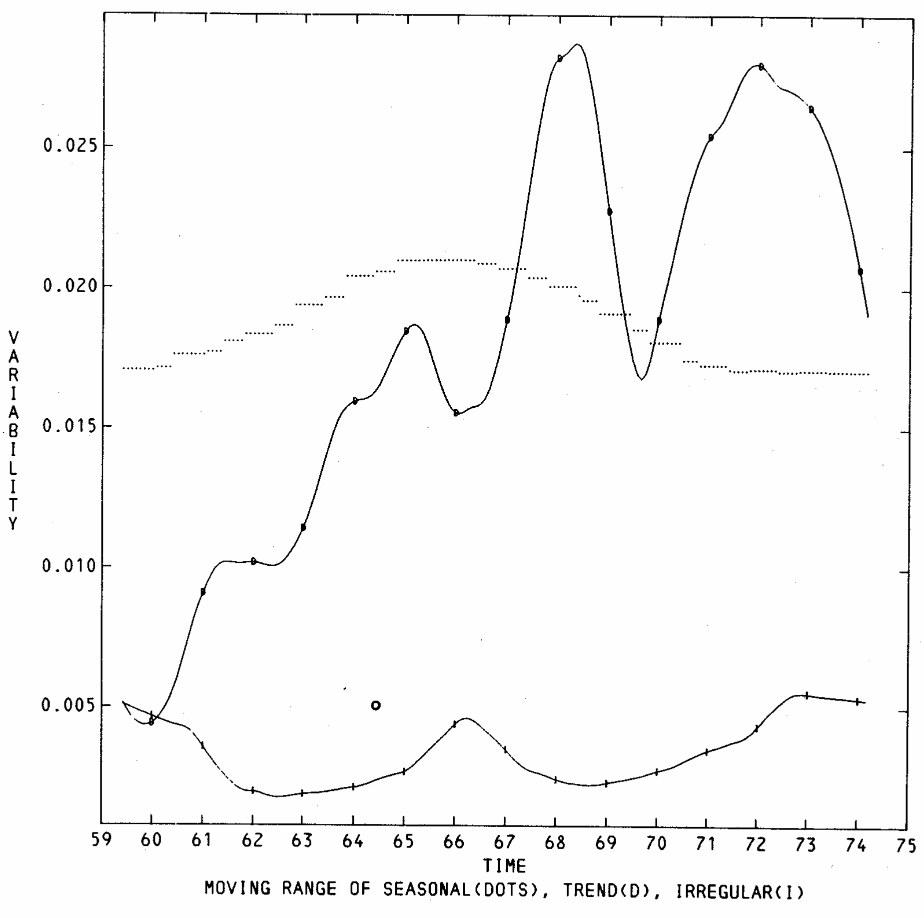

Figure 10. **COMPONENTS OF LOG MANUFACTURING SHIPMENTS, BY VARIABILITY**

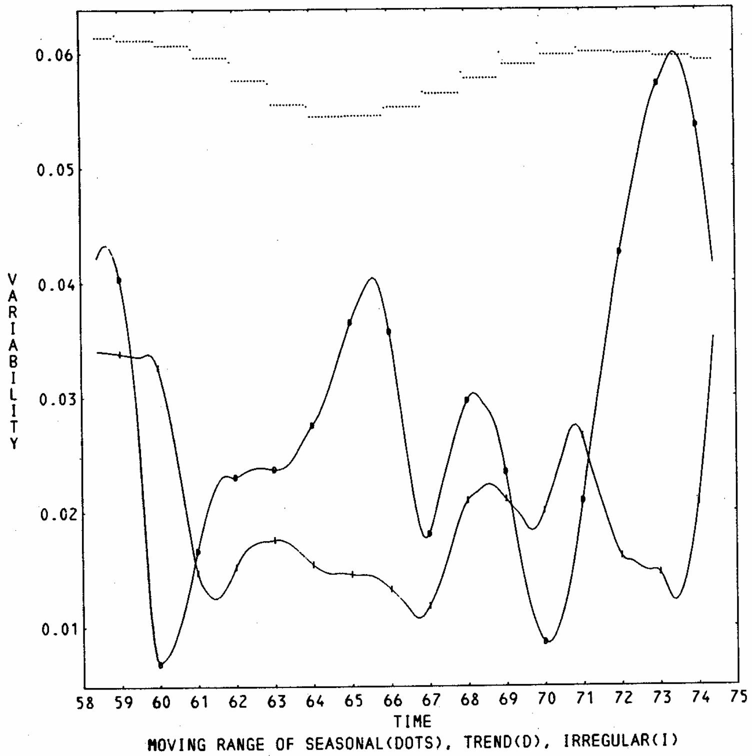

Figure 11. **COMPONENTS OF LOG MANUFACTURING SHIPMENTS, BY OVERALL VARIABILITY**

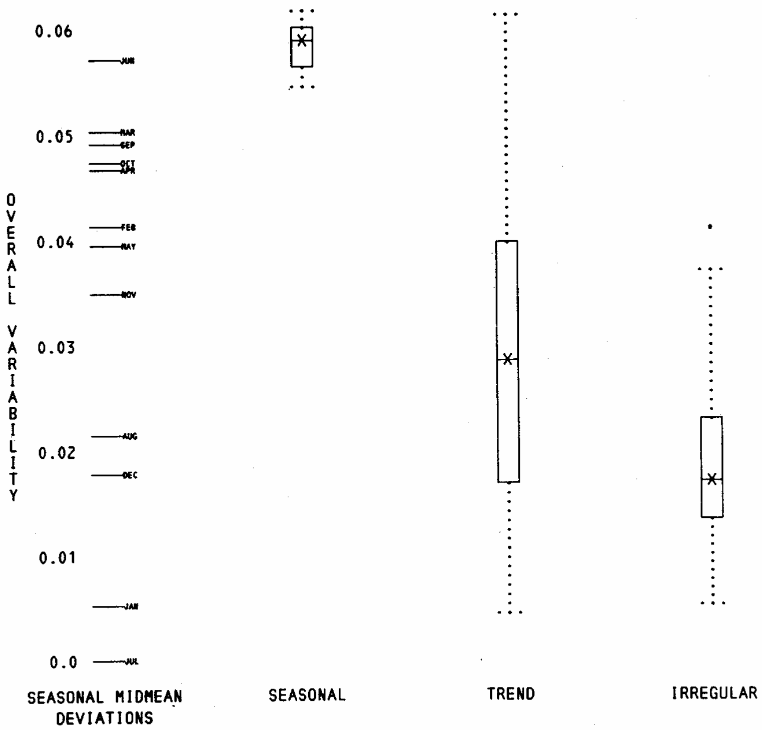

the raw values are plotted together with the smoothed values for each monthly series. For example, figure 12 shows the raw seasonal and the smoothed seasonal from SABL for November for log manufacturing shipments. Because the seasonal smoother in SABL has been designed to be resistant to outliers, the smoothed values are not affected by the outliers.

Plots of Trend: First, Second, and Third Differences of Trend Against Time

Figures 1 and 2 indicate the overall movement of the trend, but local variability in the trend cannot be readily appreciated, because the long-term change in the trend is large compared to its variability. The first, second, and third differences of the trend permit better understanding of local properties. These differences are shown in figure 13 for the logarithm of manufacturing shipments. A lack of smoothness at about 8 positions from each end of the series is apparent in figure 13, suggesting that the trend-estimation procedure (TS) may not be sufficiently smooth in the vicinity of the splice.

Plots of Seasonal

Seasonal amplitude summary and seasonal trend smooth—The plots of $S(t)$ against time (figs. 1 and 2) and the plot of the moving range of $S(t)$ against time (fig. 10) enable us to gain some appreciation of how the seasonal oscillations change through time. However, since the instabilities that are removable by transformation generally change in a smooth way and since it is amplitudes upon which we wish to focus, a useful plot for assessing the need or effect of transformation is to present a smoothed summary of the amplitudes. This is done in figure 14 for two different versions of the seasonal component for log manufacturing shipments; the seasonal used in the bottom plot has the seasonal trend smooth removed by STSR, while that in the top plot does not. Each display shows a 12-point moving maximum of $S(t)$, a 12-point moving minimum, and a QLSS smooth (see the subsection on first-trend smooth), followed by a moving average of length 2, of the absolute values of $S(t)$ plotted as the solid curve without vertical marks. The month in which the maximum or minimum occurs is identified whenever the month changes; in figure 14, the maximum and minimum always occur in June and July, respectively.

In each plot of figure 14, the curve with the vertical marks is again a QLSS smooth, followed by a moving average of length 2, of the seasonal component that aids in detecting leakage of trend into the seasonal. By comparing this QLSS smooth of the seasonal with the moving maximum and minimum, we can see whether any trend is present that is large, compared to the oscillations of the seasonal. There is a small amount of trend in the seasonal, resulting from SABL without STSR, as shown in the top plot. When STSR is used, the seasonal shows no trend at all.

Monthly seasonal plots—The seasonal amplitude plots, discussed in the previous section, enable us to appreciate the overall kind of instability that occurs when the seasonal amplitudes increase with the level of the series. But, other, more subtle, kinds of nonremovable instabilities may occur and must be recognized for complete appreciation of the seasonal pattern. Such instabilities involve changes in the levels of the cycles, often in different ways, as a result of a change, through time, in the mechanism that causes the seasonal behavior. However, in plots $S(t)$ versus t, it is difficult to perceive the behavior of a particular month, since our eyes cannot readily pick out, e.g., the January values. The behavior of the individual monthly series of the seasonal component can be assessed by displaying each series on the same graph, as in figure 15. Such a monthly seasonal plot allows the magnitudes of changes across the years of each monthly series to be compared with the amplitude of the seasonal oscillations. For each month, the midmean is shown as a horizontal line; each successive yearly value for that month is shown as a vertical line, emanating from the midmean.

Figure 15 shows the seasonal component of log-manufacturing shipments plotted in this manner. Noticeable changes in the yearly patterns are the decreases during each of the months (January, February, April, October, and December) and the increases during March, May, June, and July. The changes within several of these months are not negligible compared with the change from month-to-month, indicating an evolution of the seasonal pattern of manufacturing shipments over years.

Boxplots of seasonal residuals—The monthly seasonal plot reveals sufficiently detailed information about seasonal cycles to allow an understanding of changes in the seasonal component, when they exist. It is useful, however, to summarize the variation in each monthly series. This is done in figure 16 by making a boxplot of each monthly series of the seasonal residuals. The seasonal residuals are defined to be the values of the seasonal component, with the midmean of each monthly series subtracted from the series. The monthly midmeans are adequately displayed in figure 15. Plotting residuals from these midmeans allows more effective comparison of the variability of the monthly series. Since month-of-the-year is a circular variable with the end (December) immediately preceding the beginning (January), the months January, February, and March are plotted twice to permit better appreciation of the shape of the cycle at the beginning and end of the year. Figure 16 clearly shows that the seasonal component is most variable in December and somewhat more variable in April, June, and July than in the other months.

Parallel plots of seasonal and seasonal residuals—It is useful to have a display that presents the yearly cycles intact in such a way that overall behavior and change, through time, can be appreciated. Plots of $S(t)$ versus t

Figure 12. **SEASONAL FOR LOG MANUFACTURING SHIPMENTS**

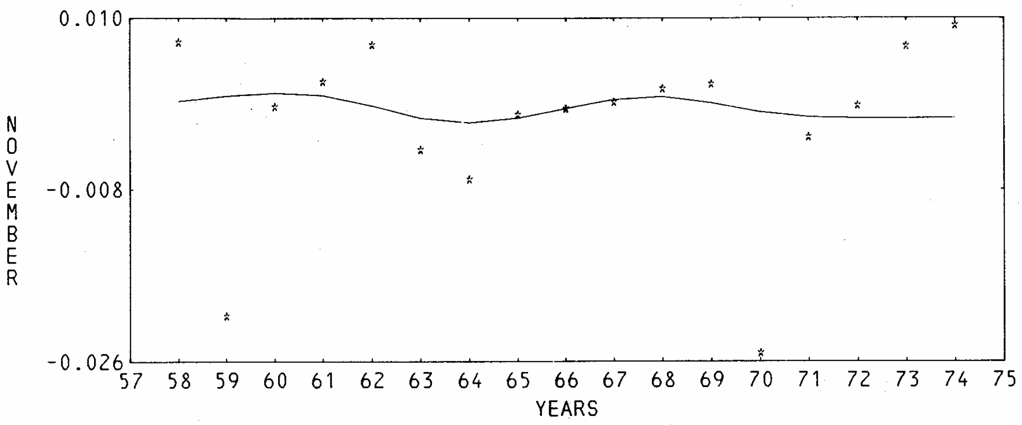

Figure 13. **TREND FOR LOG MANUFACTURING SHIPMENTS**

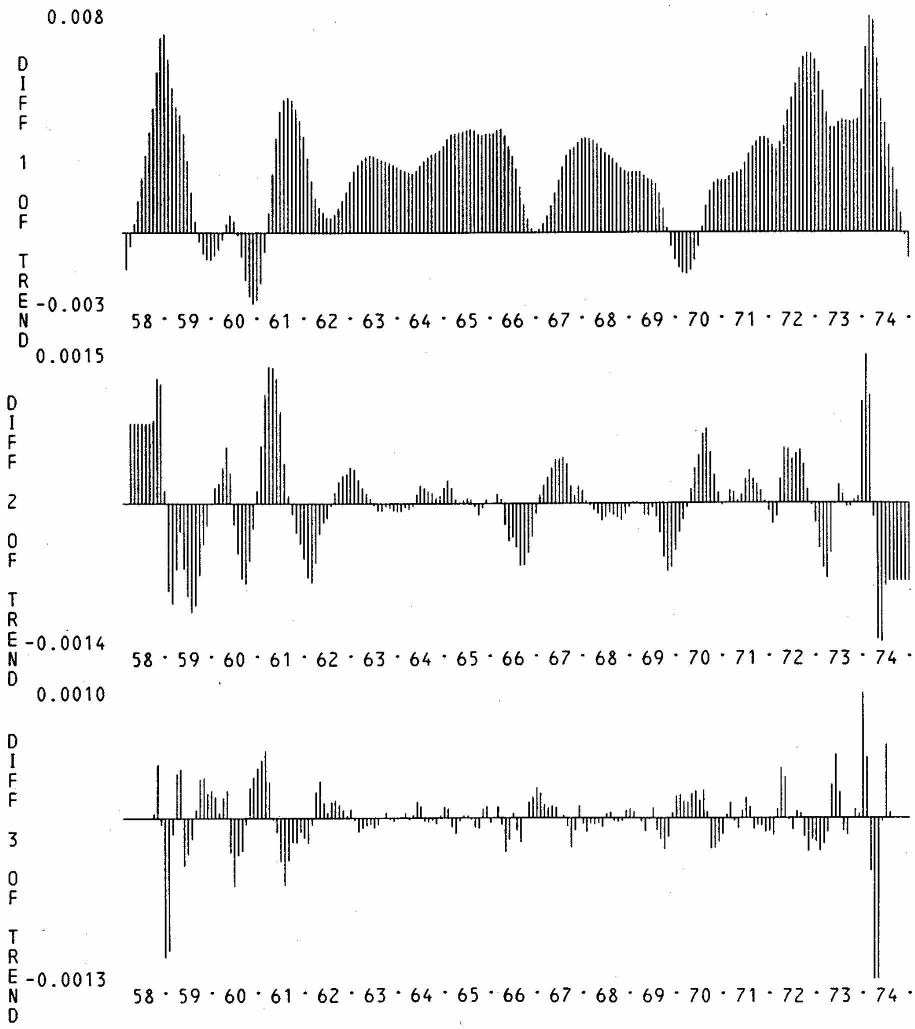

Figure 14. **SEASONAL AMPLITUDE PLOTS FOR LOG MANUFACTURING SHIPMENTS**

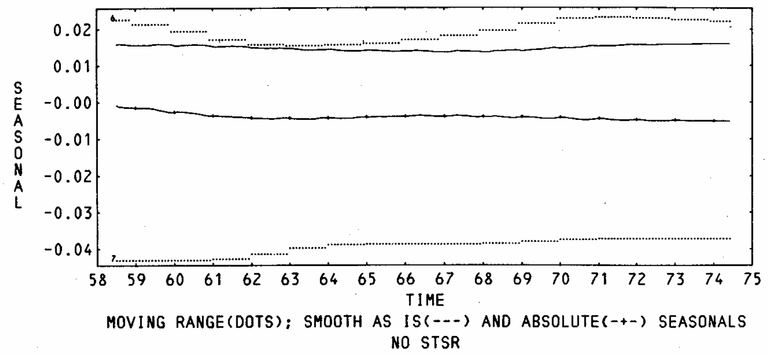

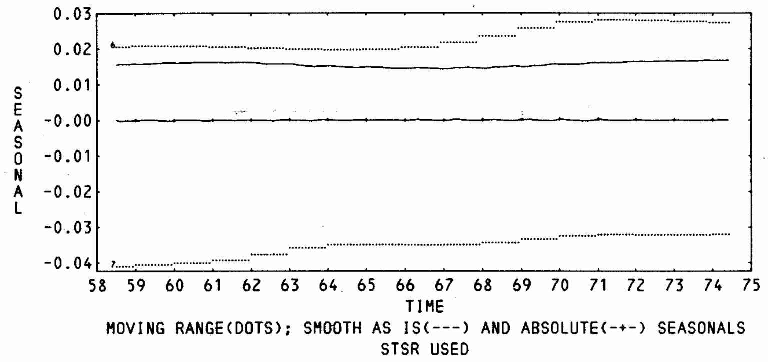

Figure 15. **MONTHLY SEASONAL PLOTS FOR LOG MANUFACTURING SHIPMENTS**

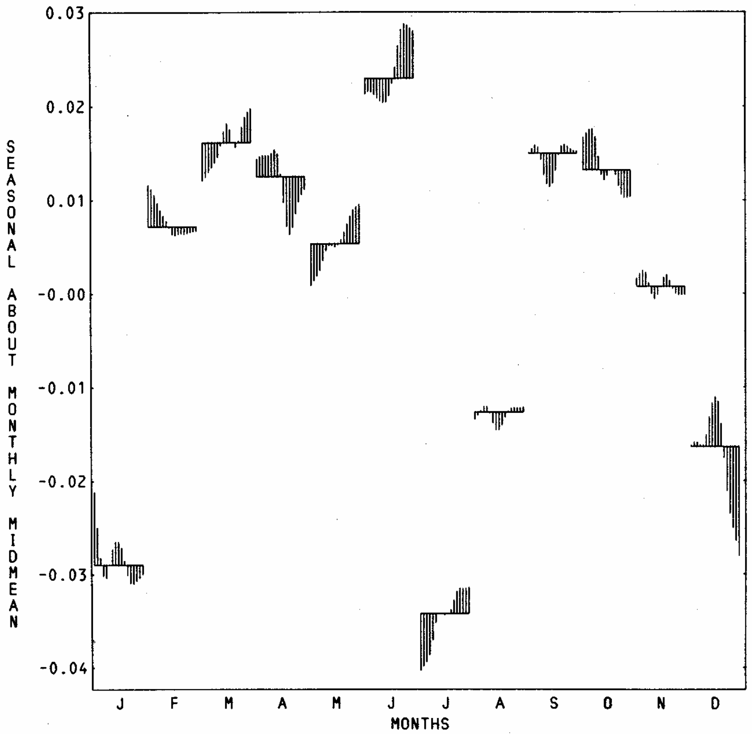

Figure 16. **BOXPLOTS, SEASONAL RESIDUALS OF LOG MANUFACTURING SHIPMENTS**

(Residuals from monthly midmeans)

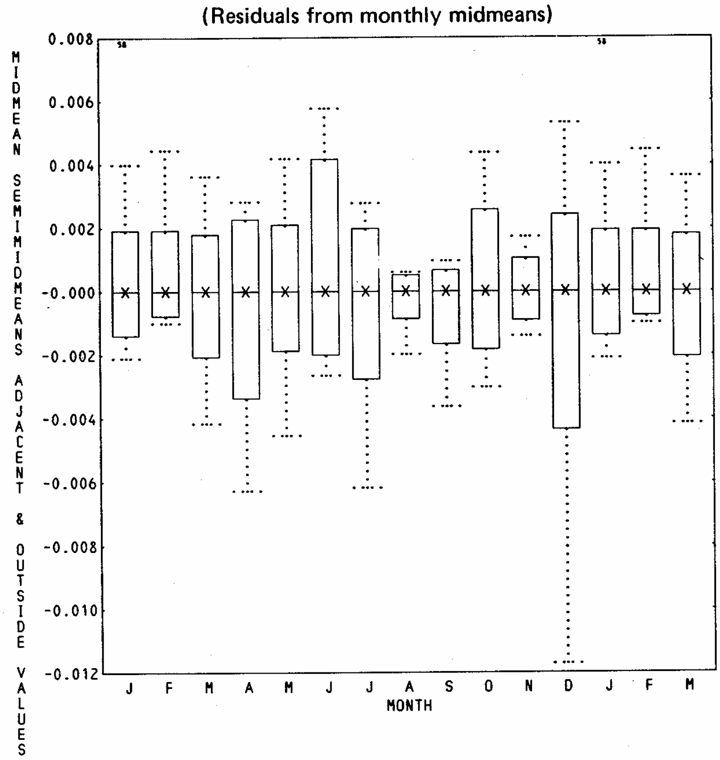

Figure 17. **PARALLEL PLOTS, SEASONAL OF LOG MONEY SUPPLY**

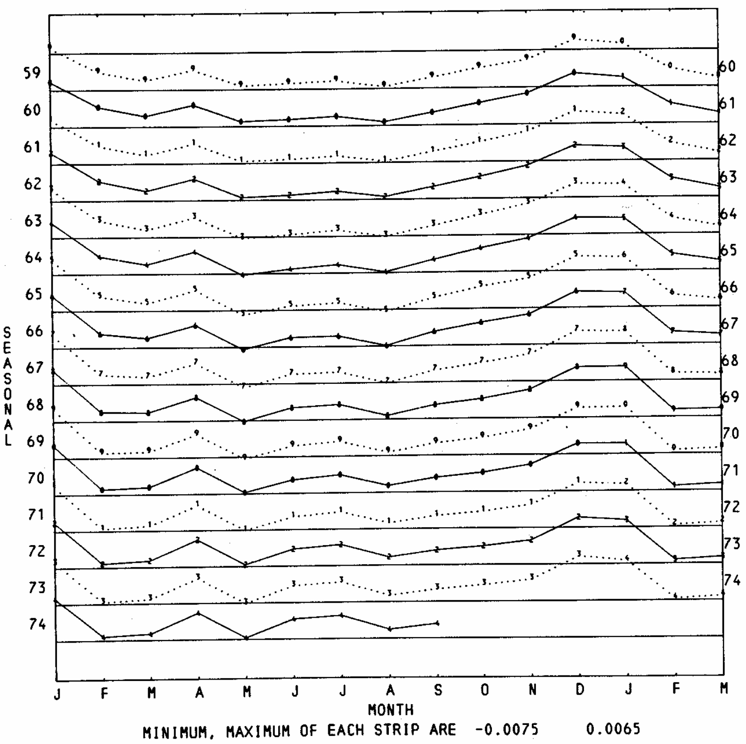

present yearly cycles intact and allow us to compare the general levels of the cycles, but difficulty in identifying individual months interferes with a full appreciation of shape. In the parallel plot of figure 17, we leave yearly cycles intact and also identify individual months. Each row shows 15 months of data, with January, February, and March of the following year included to allow better appreciation of shape at the beginning and end of the row. To assist in identifying the curves, the last digit of each year is plotted.

Tukey [22] has described an important principle in displaying data. To appreciate departures from an overall or gross effect it is useful to plot residuals from the effect. This serves as an enlarger, blowing up the residual behavior and allowing better recognition of its properties. This has been done in figure 18, by making a parallel plot of the seasonal residuals.

In Figure 17, the parallel plot for the seasonal component of log money supply shows emerging peaks in April and July. The residual parallel plot of figure 18 shows this emergence more clearly and also displays an evolution from peak to trough in February.

Circle plots of seasonal and seasonal residuals—It is also desirable to plot the seasonal component in a two-way display so that the yearly cycles are treated in the same manner as the 12 monthly series, as is done when numerical values are shown in a two-way table. An attempt at achieving such symmetry is the circle plot of figure 19. The years are rows and the months are columns. The circle areas are proportional to the absolute values of the seasonal component of log money supply. Negative values are designated by a slanted line inside the circle. Again, each row contains 15 months to allow better appreciation of shape at the beginning and end of the row.

There is less symmetry in the perception of rows and columns of figure 19 than one might have thought, even though the circle centers are equally spaced across rows and across columns. We tend to see the display as forming columns, because we tend to organize the diagram by grouping together similar symbols. This kind of perceptual, visual organization has been intensively studied by Gestalt psychologists (see [11]). The lesson to be learned is that we need to expend some effort to appreciate fully the effects in both rows and columns in circle plots. That is, we can expect, in the sense of Julesz [15], to have to utilize cognitive processes rather than relying solely on visual perception to interpret such displays.

As with the parallel plots, it is useful here also to plot the seasonal residuals to allow appreciation of departures from overall or gross effects. This is done in figure 20. The changes, through time, of each monthly series now become more apparent, but the magnitude of the changes is not easily compared to the amplitude of the seasonal oscillations.

Plots of Irregular

Spectrum and autoassociation of the irregular—The leakage of seasonal into the irregular can, in some cases, be revealed by the plot of irregular against time. But, seasonal oscillations may be obscured by variation at other frequencies. Looking at spectra and autocorrelations can greatly assist in revealing seasonal oscillations in the irregular. As a measure of autocorrelation at a given lag, we shall use the square of a robust estimate [10] of the correlation coefficient times the sign of the estimate and refer to it as the autoassociation. Figure 21 shows the spectrum and some autoassociations for the irregular component of the square root of toll revenues. The autoassociations show a pattern with a period of 3 months. The specturm peaks at one cycle per 3 months. These periodicities suggest that some seasonal has indeed leaked into the irregular.

Monthly distributions of the irregular—The autoassociation and spectrum plots in figure 21 aid in detecting seasonal in irregular but do not give information about the shape of this seasonal contamination. This is done in figure 22, using a variant of the boxplot. The distribution of each monthly series of the irregular is summarized by the same statistics described in the subsection on the description of boxplots, but the plotting technique differs in that each box is collapsed into a solid vertical line, the middle horizontal line is removed, and the successive monthly midmeans are connected by straight lines. This plotting variant is used, because we are less concerned with comparing month-to-month variability and more concerned with detecting a pattern in the monthly midmeans; thus, the pattern in the midmeans is emphasized by connecting them. This may be contrasted with figure 16, where the goal was to compare variability.

The nature of the seasonal leakage into the irregular component for square-root toll revenues is seen in figure 22 by the pattern of peaking nearly every 3 months. This leakage had been detected by the spectrum and autoassociations in figure 21. But, comparison of figure 22 with figure 23, which shows boxplots of the seasonal, reveals that this leakage is small in magnitude, compared to the variation of the seasonal component. Figure 22 also conveys information about the dependence of the spread of the irregular on the month-of-the-year and information about outliers.

Smoothed-moving quantiles—Trend can leak into the irregular if the trend smoother is not sufficiently flexible. Vertical line plots of the irregular, e.g., figures 1 and 2, may reveal such leakage. But, if the trend in the irregular is small compared to the total variability of the irregular, leakage is more easily seen in a smoothed version of the irregular. In figure 24, the middle curve consists of the irregular smoothed two times by an altered version of TS (see the subsection on trend smooth) in which the moving median length is changed to 24, and the quadratic smooth length is changed to 21. The second smoothing of the first smooth is done to further reduce the variability in the final smoothed curve. To aid in comparing the magnitude of the trend present in this smoothed irregular, with the variability of the irregular, and to gain information about

Figure 18. **PARALLEL PLOTS, SEASONAL RESIDUALS OF LOG MONEY SUPPLY**

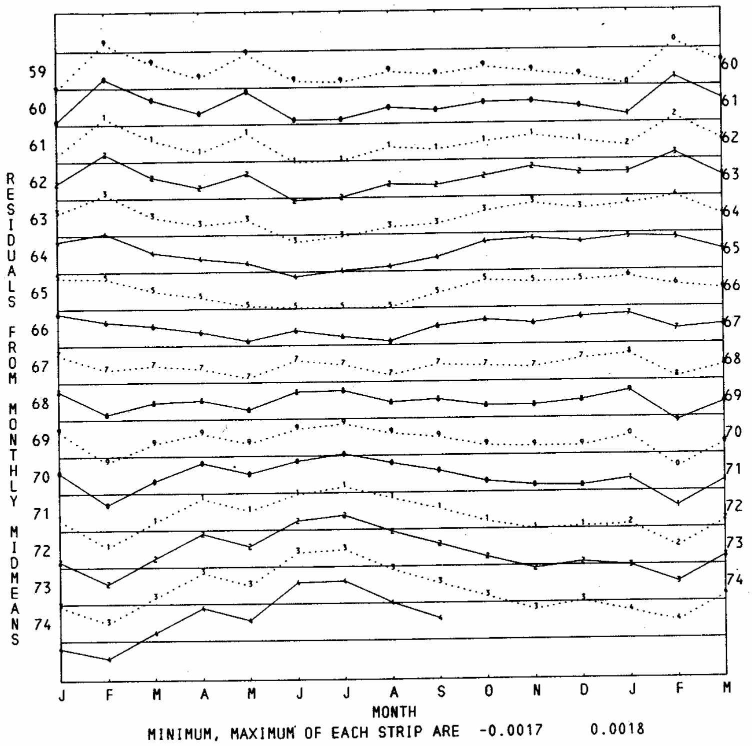

Figure 19. **CIRCLE PLOTS, SEASONAL OF LOG MONEY SUPPLY**

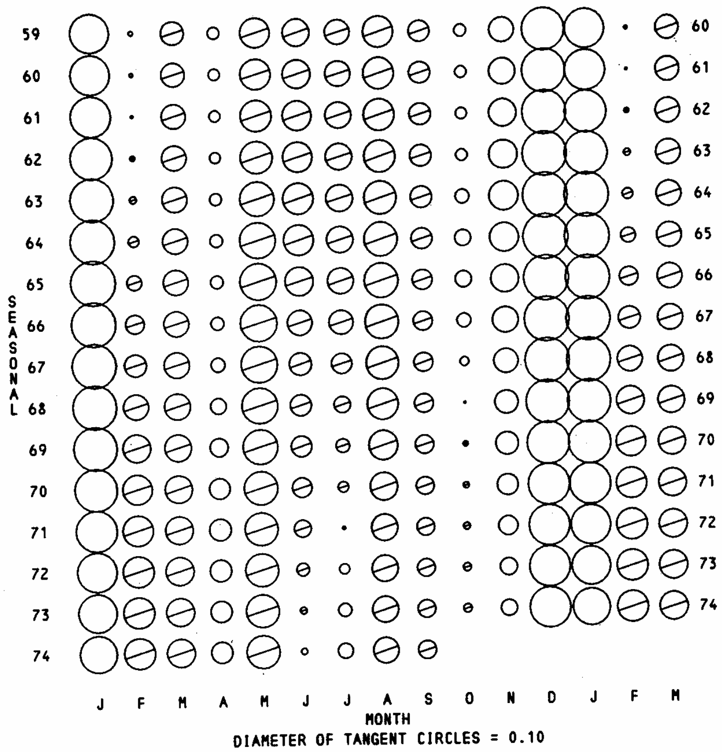

Figure 20. **CIRCLE PLOTS, SEASONAL RESIDUALS OF LOG MONEY SUPPLY**

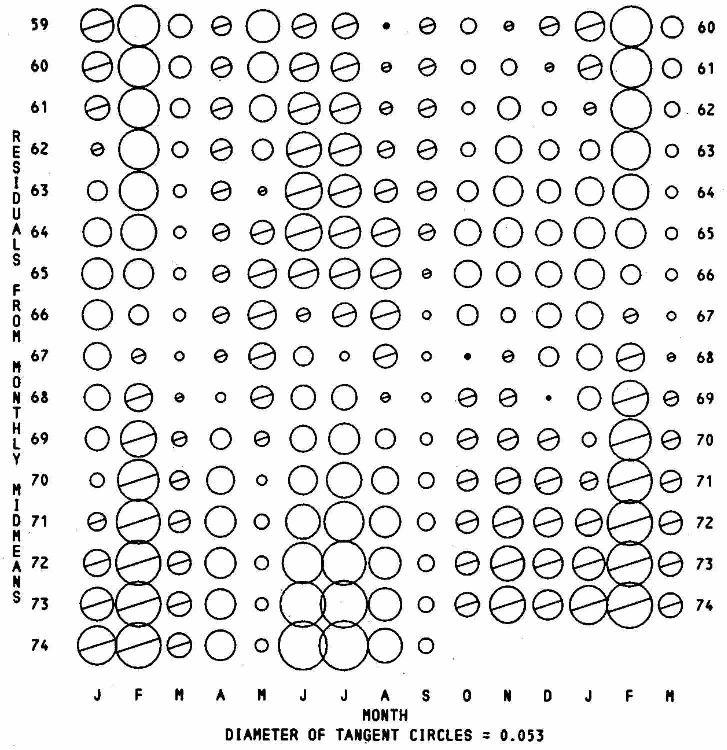

Figure 21. **IRREGULAR OF SQUARE-ROOT TOLL REVENUES, BY CYCLES AND LAGS**

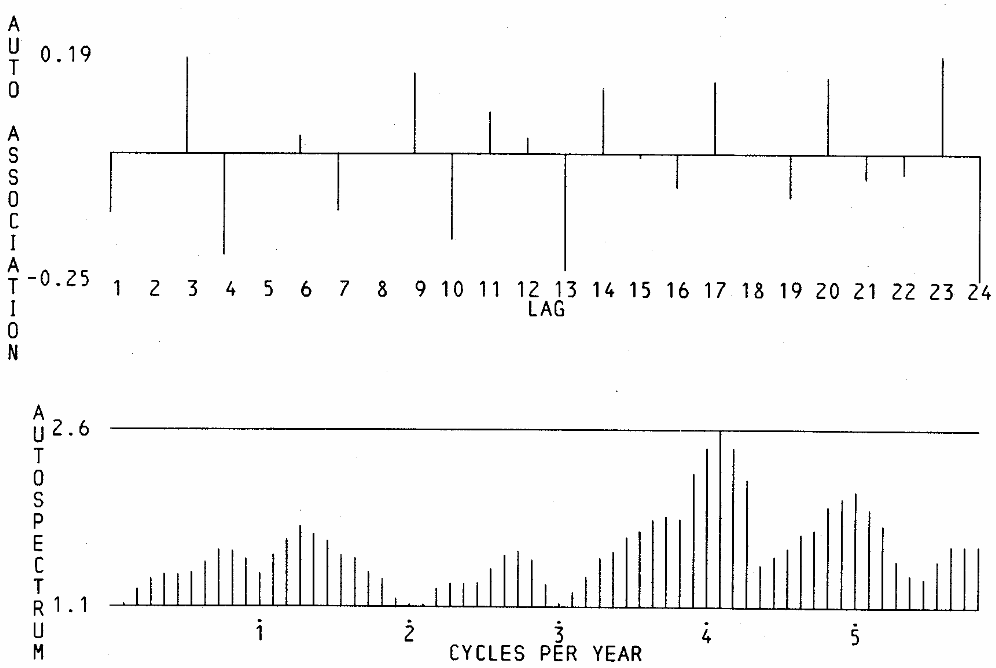

Figure 22. **IRREGULAR OF SQUARE-ROOT TOLL REVENUES, BY MONTH**

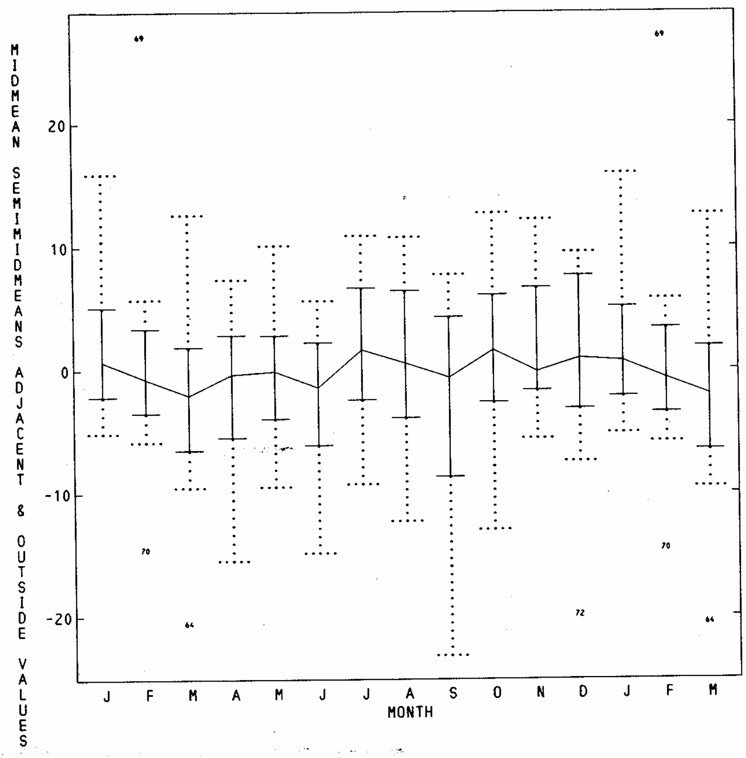

Figure 23. **SEASONAL OF SQUARE-ROOT TOLL REVENUES**

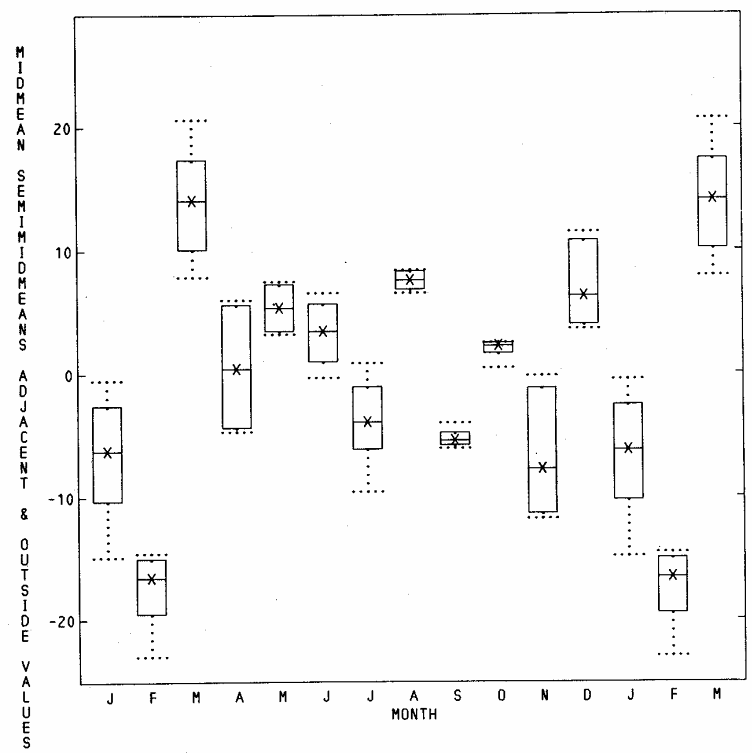

Figure 24. **IRREGULAR OF LOG MONEY SUPPLY**

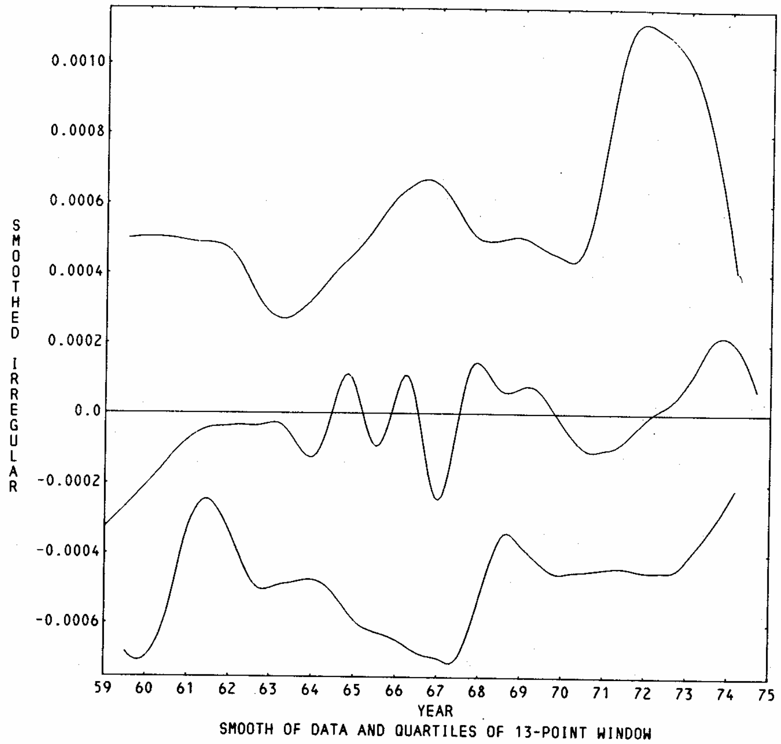

Figure 25. **MANUFACTURING SHIPMENTS USING MULTIPLICATIVE X–11**

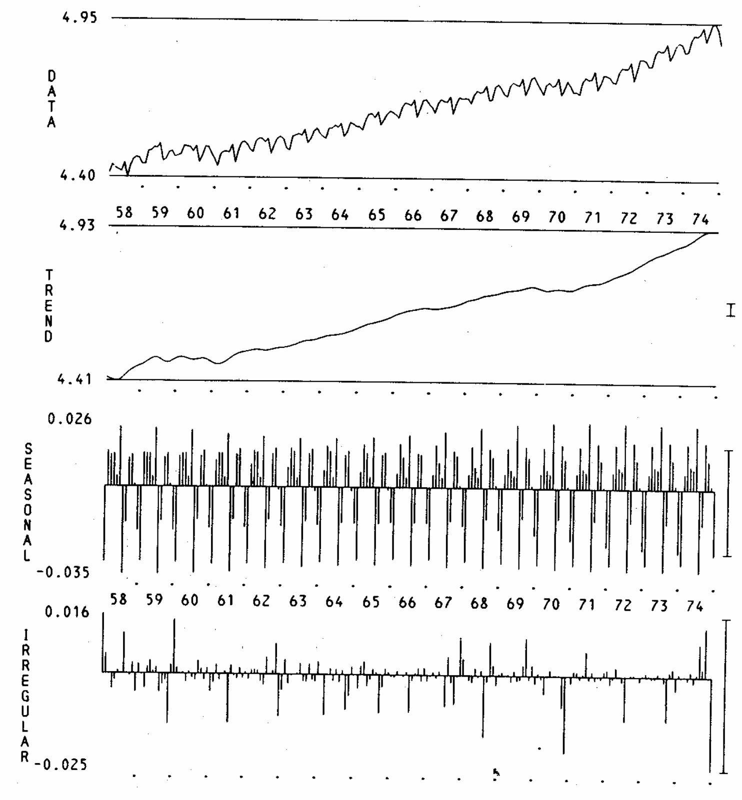

Figure 26. **LOG SEASONAL MANUFACTURING SHIPMENTS USING MULTIPLICATIVE X–11**

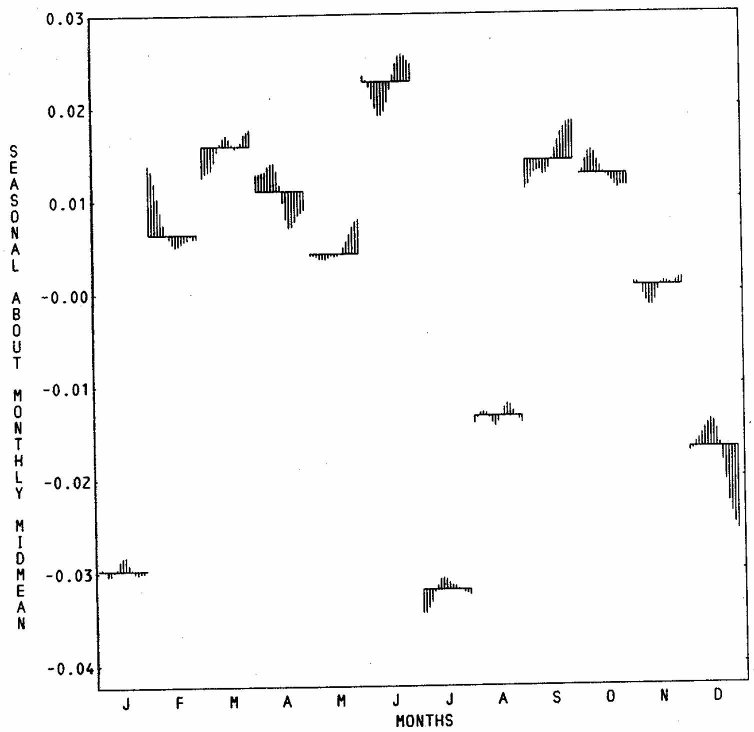

Figure 27. **SABL X–11 SEASONAL DIFFERENCES, LOG MANUFACTURING SHIPMENTS**

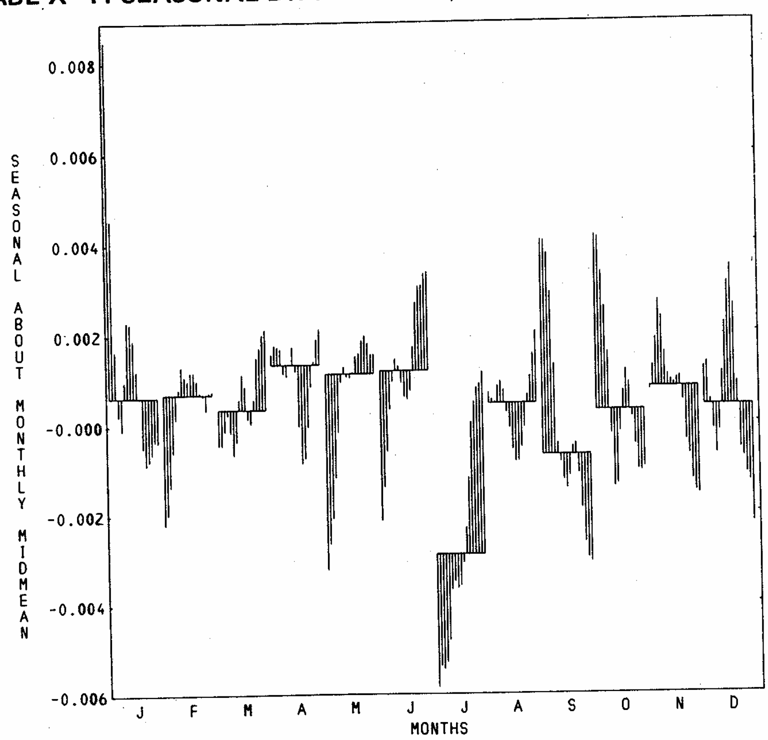

the change in the distribution of the irregular, through time, other smoothed moving quantiles are also plotted. The upper and lower curves of figure 24 are moving upper and lower quartiles, respectively, of length 13 that are then smoothed two times with this same altered version of TS. Figure 24 shows the described moving statistics for the irregular of log money supply. No trend is obvious in the irregular, but its variability seems to be increasing, with time.

A GRAPHICAL COMPARISON OF SABL AND X–11

It is natural to ask how any new seasonal adjustment procedure, e.g., SABL, compares with the existing standard, X–11. Unfortunantely, no clear-cut and generally accepted guidelines exist as a basis for such a comparison. Some properties of SABL and X–11 were contrasted in the subsection on a comparison of SABL and other methods, but a careful comparison of two seasonal adjustment methods requires detailed examination of many more facets of the adjustments. A detailed comparison of SABL and X–11 is in progress and will be the subject of a future paper. However, in order to give users a basic feel for the similarities and differences between the two methods, we present a preliminary graphical comparison of the components of manufacturing shipments from SABL and X–11.

In making this first comparison, manufacturing shipments is the data series selected for two reasons. First, it has a rather regular seasonal pattern. Second, the power transformation (log) selected for that series by SABL leads to a natural and simple comparison with X–11, using the X–11 multiplicative model and the logs of the X–11 components in the comparison. Figures 2 and 25 show the data and three components from SABL and X–11, respectively, for log manufacturing shipments.

SABL is designed with the same philosophy as X–11. Hence, one might expect (and, indeed, hope) the seasonally adjusted series $T(t)+I(t)=Z(t)-S(t)$ from the two decompositions would be similar for well-behaved data with a regular seasonal pattern. The trend and irregular, however, can differ considerably, depending on how much of the middle- and high-frequency variation is allowed into the trend. The sequences of nonlinear and linear filters, used by SABL, tend to result in a smoother trend than that formed by X–11, and an irregular which has occasional sequences of values with the same sign. Comparison of figures 2 and 25 shows that the trend from SABL is smoother than from X–11, while the SABL irregular shows more structure.

Figures 15 and 26 show the seasonal components of log-manufacturing shipments from SABL and X–11, respectively, using the monthly seasonal plot. A finer comparison is offered in figure 27, where the differences (SABL minus X–11) of the seasonal components are plotted. Note that the range of this difference plot is smaller by a factor of 5 than for the seasonal plot. The 12 monthly midmeans of the differences are near zero, with the exception of July for which the SABL estimate is lower than the X–11 estimate in the earlier years. The differences for September are large, compared to those for, e.g., March and April. In addition, the first two January values are much larger for SABL than for X–11.

In summary, this very modest comparison for a relatively stable series indicates good general agreement between SABL and X–11 in the level of the seasonal components. Differences in the shapes of evolving seasonal patterns are evident and need to be examined in greater detail to determine their causes.

REFERENCES

1. Andrews, D. F., Bickel, P. J., Hampel, F. R., Huber, P. J., Rogers, W. H. and Tukey, J. W. *Robust Estimates of Location: Survey and Advances*. Princeton, N.J.: Princeton University Press, 1972.
2. Anscombe, F. J. and Tukey, J. W. "The Examination and Analysis of Residuals." *Technometrics* 5 (1963): 141–160.
3. Bell Laboratories. "A Monte Carlo Investigation of the Mean Square Error, Bias, and Correlations of Some Nonlinear Smoothers." Memorandum by W. S. Cleveland, A. M. Gross, and R. Guarino, 1976.
4. ———. "A Robust Procedure for Simple Linear Regression." Memorandum by W. A. Larsen and I. J. Terpenning, 1972.
5. ———. "Three Specialized Plotting Programs." Memorandum by S. M. Bruntz, W. S. Cleveland, B. Kleiner, and I. J. Terpenning, 1974.
6. Box, G. E. P. and Cox, D. R., "An Analysis of Transformations." *Journal of the Royal Statistical Society*, ser. B, 26 (1964): 211–243.
7. Cleveland, W. P. and Tiao, G. C. *Decomposition of Seasonal Time Series: A Model for the Census X–11 Program*. Technical Report No. 332 Madison, Wis.: University of Wisconsin, Department of Statistics, 1974.
8. Cowden, D. J. "Moving Seasonal Indexes." *Journal of the American Statistical Association* 37 (1942): 523–524.
9. Durbin, J. and Murphy, M. J. "Seasonal Adjustment Based on a Mixed Additive-Multiplicative Model." *Journal of the Royal Statistical Society*, ser. A, 138 (1975): 385–410.
10. Gnanadesikan, R. and Kettenring, J. R. "Robust Estimates, Residuals, and Outlier Detection with Multiresponse Data." *Biometrics* 28 (1972): 81–124.
11. Gregory, R. L. *Intelligent Eye*. New York: McGraw-Hill, 1970.
12. Grether, D. M. and Nerlove, M. "Some Properties of 'Optimal' Seasonal Adjustment." *Econometrica* 38 (1970): 682–703.
13. Hannan, E. J. "The Estimation of a Changing Seasonal Pattern." *Journal of the American Statistical Association* 59 (1964): 1063–1077.
14. Jorgenson, Dale W. "Minimum Variance, Linear, Unbiased Seasonal Adjustment of Economic Time Series." *Journal of the American Statistical Association* 59 (1964): 681–724.
15. Julesz, Bela. "Experiments in the Visual Perception of Texture." *Scientific American* 23 (1975): 34–43.
16. Kendall, M. G., and Stuart, A. *The Advanced Theory of Statistics*. Vol. 3. London: C. Griffin & Co., 1966.
17. Lovell, Michael C. "Seasonal Adjustment of Economic Time Series and Multiple Regression Analysis." *Journal of the American Statistical Association* 58 (1963): 993–1010.
18. Mendershausen, Horst. "Eliminating Changing Seasonals by Multiple Regression Analysis." *Review of Economics and Statistics* 21 (1939): 171–177.
19. Parzen, Emanuel. *An Approach to Time Series Modeling and Forecasting Illustrated by Hourly Electricity Demands*. Technical Report No. 37. Buffalo, N.Y.: Statistical Science Division, State University of New York at Buffalo, 1976.
20. Stephenson, A. J., and Farr, H. T. "Seasonal Adjustment of Economic Data by Application of the General Linear Statistical Model." *Journal of the American Statistical Association* 67 (1972): 37–45.
21. Tukey, J. W. "One Degree of Freedom For Non-additivity." *Biometrics* 5 (1949): 232–242.
22. ———. *Exploratory Data Analysis*. Reading, Mass.: Addison-Wesley Publishing Company, 1977.
23. U.S. Department of Commerce, Bureau of the Census. *The X–11 Variant of the Census Method II Seasonal Adjustment Program*. Technical Paper 15 (rev. ed.). Washington, D.C.: Government Printing Office, 1967.
24. Velleman, Paul F. *Robust Non-linear Data Smoothing*. Technical Report No. 89. Princeton, N.J.: Department of Statistics, Princeton University, 1975.

COMMENTS ON "SABL: A RESISTANT SEASONAL ADJUSTMENT PROCEDURE WITH GRAPHICAL METHODS FOR INTERPRETATION AND DIAGNOSIS" BY WILLIAM S. CLEVELAND, DOUGLAS M. DUNN, AND IRMA J. TERPENNING

Peter Bloomfield
Princton University

The significant properties of SABL, the seasonal adjustment procedure developed by the authors, are described in the title of the paper. The first property is robustness, namely the property of being insensitive to small departures from the assumptions.

ROBUSTNESS AND RESISTANCE

In the context of seasonal adjustment, it seems wisest to make few assumptions anyway, and it would be more precise to describe the procedure as resistant. Resistance is the property of being insensitive to gross perturbations of a small part of the data. This is essentially a numerical property, related to continuity. It is more widely useful than robustness, which is a statistical property and can only properly be defined in well-specified contexts.

There are two main ways of constructing resistant procedures. In the first, one starts with a conventional, perhaps nonresistant, procedure that is used to obtain a trial fit to the data. Data points that deviate widely from this trial fit are inspected to determine whether they are valid. The trial fit is then modified by downweighting any suspect data points. In some versions, the modified fit is used as it stands; in others, it is regarded as a new trial fit, and the process is repeated until no further substantial change occurs. The identification and downweighting of stray data points may of course be done algorithmically, according to mathematical criteria. In the context of estimating the location of a single sample, the M– and m– estimates described and studied by Andrew et al [2] are of this form, and Beaton and Tukey [3] give such an algorithm for polynomial regression.

This approach has a serious, though not fatal, weakness, that is pointed out by the authors. It is that there are unusual, but not unheard-of situations, where one deviant data point in a particularly sensitive position may, if the trial fit is not resistant, distort the trial fit to the extent that the deviant point is not identified. Since the easiest way to implement this approach is often to use a least-squares method to obtain the trial fit and since least squares methods are notably nonresistant, this is a major drawback. The use of the median, rather than the mean, as the trial value in M– and m– estimates is one way to avoid this problem. (See Andrews et al [2]). Finding appropriate resistant trial fits is often difficult with more highly structured data. An example is given by Andrews [1] for the multiple regression problem.

The second approach to the construction of resistant algorithms is the one adopted in SABL. The first step of each phase is designed to produce a first fit (or smooth) that is resistant to deviant values. Subsequent steps work to refine this fit (in this case, by further smoothing operations) but do not use the original data values. The later steps may, thus, be sophisticated nonresistant algorithms (such as the cubic-pass-seasonal-stop filter used in SABL) without damaging the overall resistance of the procedure. In contrast with the first approach, in SABL, the iterations are designed to improve the separation of the fit into trend plus seasonal, rather than to improve the identification of deviant data points.

GRAPHICAL DISPLAYS

The second notable feature of SABL is the battery of graphical displays available to its users. Some of these are conventional, such as the graphs of the trend and seasonal and irregular components produced by SABL. The two-way array of circles used as an alternative display of the seasonal component is a more continuous variant of the coded plots used by Tukey [6] and others. As the authors point out, its usefulness is less than might have been expected because of the tendency of the eye to follow columns rather than rows. This defect might have been avoided by separating the rows a little.

The graphs of moving statistics (moving range, maximum, etc.) are valuable ways of describing the behavior of the various components as a function of time. I must, however, put in a plea for graphs that are less sensitive to the fluctuations in extreme values. The graph of moving average absolute seasonal components for log manufacturing shipments (fig. 16 in SABL) shows far less variability than the moving range (fig. 13). The moving range appears to be just the June-July difference, which is not necessarily very representative of the magnitude of the seasonal variation for a whole year.

The boxplots of the seasonal components are also useful in giving a quick visual summary of the overall shape and the consistency of the seasonal variation.

However, they need to be viewed together with the related plot of the irregular; for instance, a comparison of figures 25 and 24 of SABL suggests that the months in which the irregular is most variable are those where the seasonal component is least variable. Whatever the implications of behavior, such as this, it would be easier to detect if the graphs could be combined in some way.

AIMS OF SEASONAL ADJUSTMENT

Durbin and Murphy [5] identify two distinct aims in seasonal adjustments.

> Seasonal adjustment has two main aspects, namely the historical adjustment of past data using all the data available, and the current adjustment of each new observation. From a practical point of view the latter aspect is more important since it is the currently adjusted values which are most relevant for policy purposes.

SABL is concerned mainly with historical adjustment. It does produce a seasonal component and a seasonally adjusted component at the ends of the data, but they are found according to some apparently rather arbitrary rules. The various options that are available to the user are also selected on the basis of the performance of SABL in historical adjustment in the case of the intial power transformation of the data that is entirely appropriate. Various forms of smoothing are also available, and the adequacy of any choice must be judged from diagnostic graphical displays that are necessarily dominated by the body of the data. However, each choice of smoothing carries with it a rule for obtaining end values, and it is not clear that a choice that appears to be satisfactory in the body of the data would necessarily be associated with a good end-value rule. It would be desirable to see some assessment of the degree of matching of end-value rules. Some comments made by the authors, for example, with respect to the sine function splicing of the end values of the trend component, suggest that there may be some problems.

OTHER PROBLEMS AND OTHER METHODS

The authors describe a number of checks for leakage, which is the appearance of an effect that should be confined to one component in another. I would like to suggest one more form of leakage that deserves some attention, and that is leakage of the irregular component into the seasonal component. The authors state that the "seasonal smoother should be flexible enough to follow changes in the seasonal pattern while still being resistant to deviant observations." In addition, the seasonal smoother should be no more flexible than is necessary to follow changes in the seasonal pattern, which is a different requirement from being resistant to deviant observations.

The danger of such leakage is not the gross bias that may occur from other forms of leakage but that the computed seasonal component may be more variable than is necessary. This is a particular problem in current adjustment, since currently adjusted values are updated as new data become available. If there is leakage of the irregular component into the seasonal component, these modifications will be larger than necessary, which reduces the value of the current seasonally adjusted data.

The appropriate level of firmness of the seasonal smoother may be found from the spectrum of the detrended data (seasonal plus irregular). If the seasonal pattern were, in fact, constant, the peaks in the spectrum at the fundamental seasonal frequency and its harmonics would be as narrow as possible (being images of the spectral window if a conventional quadratic spectrum estimate is used; see [4, ch. 7]). In this case, the smoother may be as firm as possible and should just replace each monthly value by a (resistant) estimate of location of the values for that month. If the seasonal pattern were to vary, these peaks in the spectrum would be wider than the minimum, and the additional width would indicate the bands of frequencies surrounding the seasonal frequencies that contain the seasonal effect. In this case, the seasonal smoother should be chosen to have a transfer function that passes power at those frequencies but annihilates or satisfactorily attenuates all others. If it passes power at more frequencies than is necessary, it will allow more of the irregular to leak into the seasonal component than is necessary.

Diagnosis of this problem, after the fact, is not easy. The spectrum of the irregular component usually shows troughs at the seasonal frequencies. (See, for instance, fig. 23.) This merely indicates the widths of the bands from which power has been removed (and passed into the seasonal component) but not whether there was seasonal power across those bands in the first place. The spectrum of the seasonal component is not normally computed and would presumably just show power across the same bands. Only in extreme cases would the true seasonal peaks be narrow enough to be distinguished.

MORE COMPLEX SEASONAL BEHAVIOR

My final comment relates to the problem of seasonal adjustment, in general, and not specifically to SABL. Figure 1 shows a series with a strong seasonal component. It is not of economic origin, being actually the height of the ozone column at Arosa, Switzerland. There is essentially no trend in the data, at least in comparison with the magnitude of the seasonal variation. Figure 2 shows boxplots (familiar to readers of SABL but constructed here from medians and quartiles) for the monthly series. The graphs show a very clear seasonal effect, not only in the location of the values for each month but also in the scale. Furthermore, the correlation between values in adjacent months also varies seasonally.

Such complex seasonal dependence should presumably be considered in making seasonal adjustments. For in-

Figure 1. **HEIGHT OF OZONE COLUMN AT AROSA, SWITZERLAND**

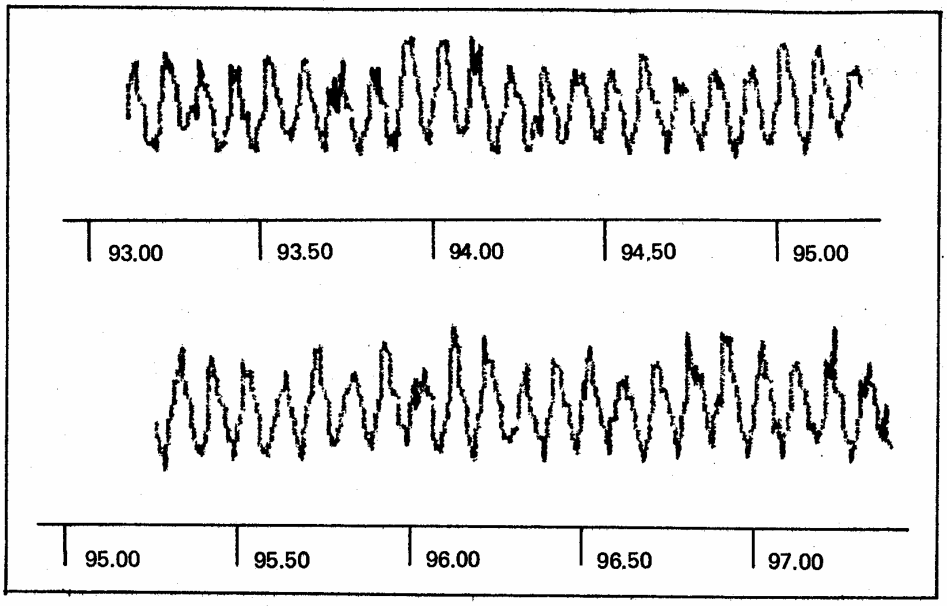

Figure 2. **BOXPLOT OF MONTHLY SERIES FROM FIGURE 1**

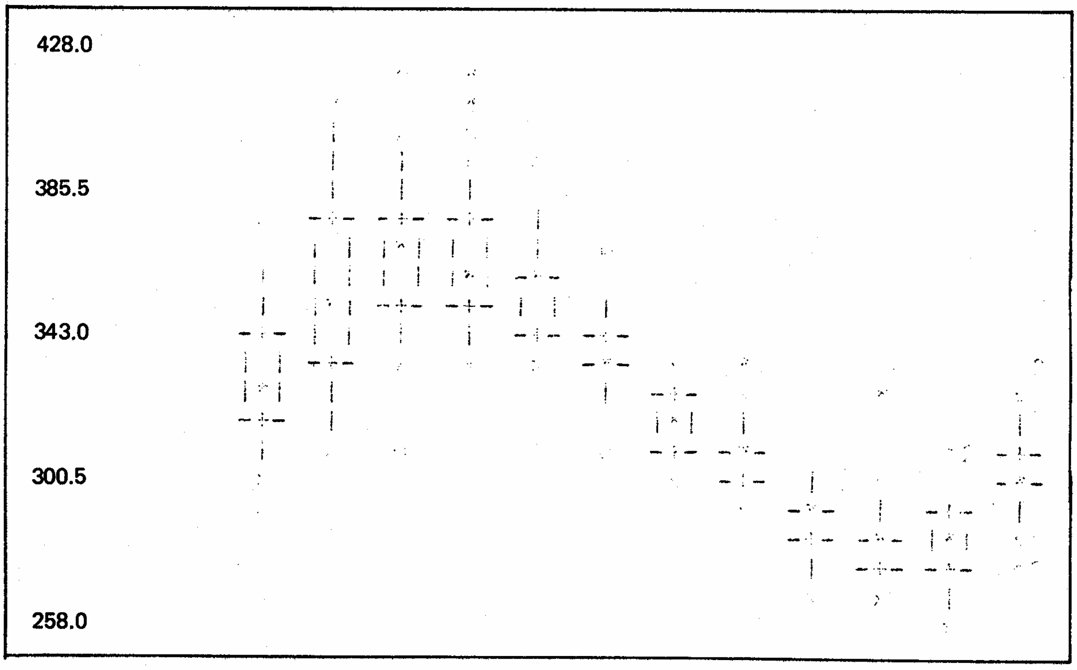

stance, the low variability of the summer months means that a trend smoother should place more weight on those than on the winter months. The seasonal variation in the correlation structure has more obscure implications.

Is it enough, however, to correct for the seasonal dependence of the mean? To be fully deseasonalized, should the seasonal dependence of the scale also be removed and by what form of adjustment? Or, should we merely be careful to quote seasonally dependent scale estimates along with seasonally adjusted values? In such a series, the whole issue of seasonal adjustment seems to be a more difficult problem.

REFERENCES

1. Andrews, David F. "A Robust Method for Multiple Linear Regression." *Technometrics* 16 (1974): 523–531.

2. ———, Bickel, P. J., Hampel, F. R., Huber, P. J., Rogers, W. H., and Tukey, John W. *Robust Estimates of Location*. Princeton, N.J.: Princeton University Press, 1972.

3. Beaton, Albert E., and Tukey, John W. "The Fitting of Power Series, Meaning Polynomials, Illustrated on Band-Spectroscopic Data." *Technometrics* 16 (1974): 147–185.

4. Bloomfield, Peter. *Fourier Analysis of Time Series: An Introduction*. New York: Wiley, 1976.

5. Durbin, James, and Murphy, M. J. "Seasonal Adjustment Based on a Mixed Additive-Multiplicative Model." *Journal of the Royal Statistical Society*, ser. A, 138 (1975): 385–410.

6. Tukey, John W. *An Introduction to Exploratory Data Analysis*. Reading, Mass: Addison-Wesley, 1977.

COMMENTS ON "A RESISTANT SEASONAL ADJUSTMENT PROCEDURE WITH GRAPHICAL METHODS FOR INTERPRETATION AND DIAGNOSIS" BY WILLIAM S. CLEVELAND, DOUGLAS M. DUNN, AND IRMA J. TERPENNING

William P. Cleveland
Duke University

ABSTRACT

The graphical methods presented in the paper discussed are found to be of value for assessing the significance of seasonal and trend patterns and for judging the quality of a seasonal decomposition. The new seasonal adjustment procedure, SABL, is evaluated for the results it produces and its simplicity of application. Initial reaction is positive on both counts, but more comparisons with other methods are requested in order that adequate judgements can be made.

INTRODUCTION

The paper by Cleveland, Dunn, and Terpenning, as I see it, makes valuable contributions in three areas. One area is methods of display of trends and seasonal components, however arrived at, in such a way as to facilitate comprehension of individual component patterns and their relative importance. Second is the provision of tools for determining whether a decomposition was well done in the sense of the usual objectives described in the following discussion. Third, a new method of achieving decomposition of seasonal time series is given. I shall discuss these three areas in the order given, since I feel that the first two are a necessary context for discussing the method. Although most of the graphs facilitate both pattern comprehension and evaluation of the quality of the decomposition, I feel that the emphasis varies from one graph to another, and the conceptual viewpoints are sufficiently distinct so that it is useful to discuss some graphs from both points of view. Figure numbers given refer to figures in the paper being discussed.

COMPREHENSION OF THE COMPONENTS

The first plots one thinks of in this area are the connected plots and vertical line plots of the overall series, the trend components, and the seasonal component, as in figures 1–3. While I agree, in general, with their comments on the circumstances under which each type of plot is useful, a vertical line plot of the trend might reveal something about turning point locations. Vertical line plots of first differences of the trend, as in figure 15, are certainly helpful in this regard. They give particular emphasis to periods of boom and bust. Higher order trend differences speak more to smoothness and presence or absence of certain frequency components and, thus, seem more relevant to evaluation of the decomposition.

Though seasonal amplitude plots, figure 16, are discussed in terms of evaluating the success of an amplitude-stabilizing transformation, they would provide an equally successful characterization of the untransformed series. I was particularly impressed by the effectiveness of the transposed seasonal vertical line plots in figure 17 in revealing the character of the changes in the seasonal pattern and their relationship to the overall seasonal pattern. The 2-year parallel plots of the seasonal component in figures 18 and 19 are good displays of seasonal pattern for any year but not as good on its evolution as they stand. I thought they were improved by omitting every other line and drawing each curve centered on its mean value, a feat easily accomplished with some tracing paper. Since the seasonal is smoothed to produce only minor changes from year-to-year, the character of the change is emphasized more by the omissions. The mean lines give a better reference to the eye to see the pattern. Figure 20 seems redundant and less effective than the transposed plots. The circle plots are an interesting idea, but my eyes read area and fail to note the line reading negative. Hence, I get an absolute value plot from them. I would like to see all the types of plots for one series. Assessing the contribution of the various displays is difficult when different series are used to illustrate different points.

The box plots of figures 21, 24, and 25 address a topic of interest to Tiao and myself, the variability of different months. The relative variability of the seasonal and irregular components in the different months is immediately evident here. Our work indicates that differences in variability from month-to-month might well call for different models for different months.

Finally, the relative importance of seasonal trend and irregular components is well described by the plots of variability comparisons using moving ranges. The box plots of the variation in the trend seasonal and irregular components are also valuable for this purpose.

VALIDITY OF THE DECOMPOSITION

Most lists of objectives for seasonal decomposition procedures take me through a two-level process of conceptualization. First, there is the ideal seasonal series and, secondly, there are the permitted departures from it to allow for reality. This ideal series displays underlying concepts of trend and seasonality. In the time domain, the series has a linear or low-order polynomial trend. For monthly series, the seasonal component has the same relative values of the different months each year. In the frequency domain, the trend power is largely in the low frequencies and tapers off smoothly in the seasonal range. The seasonal power is concentrated at the seasonal frequencies (1/12, 2/12, ..., 6/12). The noise, or irregular component, is uncorrelated, or white noise.

It is in this context that statements of objectives, such as—

1. The trend component should appear trendlike
2. The seasonal component should be seasonal
3. The irregular component should look irregular

make some sense. A second, but closely related, class of objectives are that—

1. The trend should not contain elements of the seasonal or irregular (it should be smooth and without periodic components in the seasonal range).
2. The seasonal should not contain elements of the trend or irregular (it should be smooth from year-to-year and have zero average).
3. The irregular should not contain elements of the trend or seasonal.

Even at this ideal level, there is some potential conflict between time domain and frequency domain conceptualizations. A trigonometric series of finite length does not consist precisely of spikes at the given frequencies, and one does not need to restrict the representation of trend to polynomials over the span of the series to achieve a smooth spectrum in which low frequencies predominate.

This leads, naturally, to my second level of conceptualization. Trends generally refuse to follow a low-order polynomial for any great length of time. One only needs to observe the series discussed in the paper under review. Thus, the spectral definition of trend is emphasized. In the time domain, smoothness is emphasized. Whether or not continuous first derivatives are desirable depends on the usefulness of turning point analysis. A substantial amount of power in the seasonal domain is permitted. It has also been found that seasonal patterns do not repeat exactly from year-to-year. Here is additional reason for the seasonal spectrum to spread out from the strictly seasonal frequencies. In the time domain, a phase shift among the frequency components is permitted from year-to-year or, equivalently, the relative monthly values change from year-to-year.

A third class of objectives is related to these discrepancies between the ideal and real series. These can be summarized as follows:

1. The trend estimate should follow the real trend.
2. The seasonal should follow the real seasonal.

A large part of discussions about the adequacy of seasonal adjustments can be viewed in terms of how much deviation from the ideal components is appropriate or how the real differs from the ideal. The remainder of these discussions is about two goals not yet mentioned and not raised to any degree in this paper. One is that seasonal adjustment should not affect predictions. Trend predictions from the trend component or adjusted series should not differ from the trend component of overall predictions. This problem is closely related to the handling of end effects. In general, parametric approaches have an advantage here. The other objective is that relationships between series should not be distorted by seasonal adjustments.

Considerable attention has been given to frequency domain criteria for the results of decomposition procedures. (See [5; 15; 17; 19].) Plots of the power spectrum of the trend estimate and of the seasonal estimate are often made. Cross-power spectra of the original series and the seasonally adjusted series or of the original series and its components yield coherence and phase plots. These are used to indicate how much of the power of the original series was transferred to each component at various frequencies and how much phase or lead-lag distortion may be present in the components. Though helpful, these tools have not proved entirely satisfactory. One reason is the uncertainty about objectives mentioned previously. How smooth should a trend be? How fast can a seasonal pattern change? Each of these has implications in the frequency domain. The difficulty of assessing the impact of deviations from ideal patterns in the frequency domain on the actual estimates has led to more than one call for time domain strategies. Further, the results in [5] reconfirm the inconsistency in frequency and time domain criteria. In an artificial series, where optimal estimates in the minimum mean squared error sense were obtained for the components, dips in the trend spectrum at the seasonal frequencies were obtained that are not consistent with the idea of having the smooth trend spectrum in this range. Thus, the optimal trend spectrum is not uniquely defined.

The graphical methods presented in the paper under discussion are a welcome addition to our techniques for resolving decomposition issues. Although a few comparisons of SABL with X–11 are made here, these methods have not yet been fully tried on several competing versions of estimates or placed along side spectral comparisons.

Hence, their real power to reveal the information desired cannot yet be fully appreciated. The discussion which follows is largely an underlining of the strengths of particular graphs.

Trend Assessment

The smoothness of a trend can be seen directly in a connected plot of the trend, but differences between closely related estimates could be seen more readily in the amplitude of line plots of their second differences. These, of course, translate directly into rates of slope change. Third differences are harder to interpret and show some indication of generating confusion rather than yielding simplicity in figure 15. A power spectrum of the second difference would reveal just as much about presence of seasonal frequencies as a power spectrum of the third difference. The autocorrelation function may also provide information on the presence of seasonal frequencies. Coherence plots will reflect the smoothness of the estimate and degree to which the low-frequency behavior in a series is captured in the trend estimate, but I would agree with the authors that time domain representation of the trend and its differences is more useful. Phase shift is more complicated. Although zero-phase shift, at all frequencies, seems desirable, and it is best seen in a phase shift plot, it is difficult to tell what departures from this mean. Plots of the first difference of the original and adjusted series might be helpful. When comparing two adjustment procedures, the first difference plots of both trend estimates would provide an interesting comparison in this regard. Stephensen and Farr [19] showed, in a simulated series, more phase shift between the original series and true trend component than between the original series and the estimated trend component. One wonders how the first different plots of the true trend and estimated trend components would have compared. Another approach would be to generate two correlated series and see if the decomposition recovers the correct relationship, as in [21].

One thing to keep in mind when using additive stochastic models, such as the ones appearing in [5; 19], is that simulating a seasonal with a model in shift operators of degree 12 and fitting a series with such a model are not precisely the same. Such a model is consistent with a seasonal component having the property that moving averages of width 12 give a value of zero. However, simulated seasonal component series may deviate from this considerably and thus be defined as having trend components by a subsequent smoothing process. Perhaps these stochastic component models do not adequately reflect our ideas about seasonality. When one is simply forecasting the sum of the seasonal and trend of the series, as in Box and Jenkins [1] or Thompson and Tiao [20], such decomposition issues become irrelevant, and the overall models are perfectly adequate.

Seasonal Assessment

The seasonal estimate is evaluated in two kinds of ways. First, what does its subtraction from the total series leave behind? Does the seasonally adjusted series still contain seasonal effects? This kind of evaluation goes hand-in-hand with trend studies. Does the spectrum of the second difference of the adjusted series reveal seasonal effects? Second, the seasonal estimate itself may be examined. The smoothness of year-to-year changes in the seasonal pattern may be seen in the transposed plots and 2-year plots mentioned under interpretation of the seasonal pattern. Trend leakage is usually studied by 12-term plus 2-term moving averages, often referred to as centered 12-term moving averages. The authors here have adopted a 24-point weighted average with a 12-point uniform average spliced at the ends. To this is added a 12-point moving maximum and corresponding moving minimum. Thus, the importance of some trend leakage, relative to the overall seasonal amplitude, is presented. No criteria for how much leakage is too much are suggested. Perhaps the ratio of these values to the trend values would be helpful if the adjusted series is the objective of the analysis.

Irregular Assessment

Computing the autocorrelations or power spectrum of the irregular to detect seasonal components is fairly standard procedure. Figure 24 shows midmeans and semi-midmeans revealing the strength of the seasonal pattern and the variability attributable to each month. The moving quartiles and averages of the irregular to detect trend are also useful. The need for a special filter for this purpose is not clear. The trend filter contained in SABL should serve equally well. The box plots of the irregular for the different months offer valuable insight concerning whether the decomposition gives a uniformly accurate summary of the behavior of the series for each month.

One important means of assessing the relationship between series is to relate the residuals of a model for one to the values or residuals of another. In this regard, the phase properties of residual estimates become important. Though Rosenblatt [17] generated some phase-shift plots for irregulars, little attention has been given to this topic, since no one knows just what they should look like. The simulation approach of Grether and Nerlove [5] would allow coherence and phase plots of estimated versus true irregular components. Vertical line plots of the estimated and true irregulars could provide additional information. As mentioned earlier, pairs of simulated series could also be analyzed.

Comparisons of Adjustment Methods

Any method of evaluating a single seasonal adjustment procedure can, of course, be used to compare procedures. However, when comparing adjustment methods, one has the opportunity to generate differences between corresponding component estimates. If vertical line plots of trend or seasonal component differences were scaled to represent percent of the trend range, some perspective on the importance of the differences would then be achieved. An average of the two trend curves might be used to get the denominator. The differences might also be divided by the difference between some of appropriate upper and lower percentile points of the irregulars. This would indicate if the variability of the estimates between methods was in line with the estimate variability suggested by the irregulars of a single decomposition.

THE SABL PROCEDURE

How should SABL be judged? The principal objectives of a seasonal adjustment procedure are to produce satisfactory estimates of the seasonal and trend components. Properties of these have been discussed. The plots presented in the paper for SABL decomposition of several series look quite good. The authors emphasized robustness against outliers, and the irregular plots indicate that some large values and short groups of such values were ignored. The slight peculiarities at the end of one series, seen in figure 15, are the only hint of need for further adjustments.

Though the evidence presented here indicates this is a good procedure, much more remains to be done. The authors tell us that more comparisons with X–11 results are planned. There are indications that the performance of adjustment techniques based on successive applications of given filters may vary according to the overall structure of the observed series in Cleveland and Tiao [3]. Hence, it would be well to select several series with rather different autocorrelation structures for comparisons. Parametric models could also be used on these series. Comparisons with parametric methods near the ends of the series would be particularly interesting, since the prediction formulas implied by the models are used to a greater extent there. I would also like to see how the end estimates would change if the X–11 or SABL programs were applied to a series that had first been forecasted for a year by the methods of Box and Jenkins. The filters used in SABL are such that study by linear approximation, as has been done with X–11, does not appear reasonable.

My impression is that this program would be slightly harder to use than X–11 but not too much. It is more complicated, but the kinds of decisions required are similar. Is this smooth enough or should another iteration be used? Is a trend filter of width 15 or 21 appropriate? This kind of decision is required at more stages of SABL. I could imagine an interactive version of SABL with the results of each stage displayed on a video terminal. One could say "go on," "one more time," or, perhaps, "use the wider filter next time." Decomposition by parametric models requires a model fitting process that is often less straightforward than either X–11 or SABL.

The use of additive stochastic models presumes not only a model-fitting process but resolution of severe identification problems. The discussion in Cleveland [2] suggests that, by maximizing certain error variance ratios to give smooth estimates, the range of models is reduced considerably. Perhaps, any of the permissible models would do well, since they will differ generally only in moving average parameters. Defining the appropriate space within which the values would be selected is not easy. However, a nonlinear least squares routine could be used to get a model. The autocorrelation function of a component model with specific parameter values would be generated. This would be compared with the autocorrelation function with the model identified for the overall series and the parameter values of the component model adjusted to arrive at a match of the autocorrelation functions.

In any case, an evaluation problem remains. There are no theoretical results that say a specific set of filters or a specific class of models is good for all series and for all purposes. Thus, any series should be subjected to multiple techniques of analysis. Perhaps, the tools are now available to make intelligent judgements concerning the best decomposition for a given purpose.

REFERENCES

1. Box, George E. P., and Jenkins, G. M. *Time Series Analysis: Forecasting and Control.* San Francisco: Holden Day, 1970.

2. Cleveland, William P. "Analysis and Forecasting of Seasonal Time Series." Ph.D. dissertation, University of Wisconsin, 1972.

3. ———. and Tiao, George C. "Decomposition of Seasonal Time Series: A Model for the Census X–11 Program." *Journal of the American Statistical Association* 71 (1976): 581–587.

4. Durbin, James, and Murphy, M. J. "Seasonal Adjustment Based on a Mixed Additive-Multiplicative Model." *Journal of the Royal Statistical Society*, ser. A, 138 (1975): 385–410.

5. Grether, D. M., and Nerlove, M. "Some Properties of 'Optimal' Seasonal Adjustment." *Econometrica* 38 (1970): 682–703.

6. Hannan, E. J. "The Estimation of a Changing Seasonal Pattern." *Journal of the American Statistical Association* 59 (1964): 1063–1077.

7. ———. "The Estimation of Seasonal Variation." *Australian Journal of Statistics* 2 (1960): 1–15.

8. ———. "The Estimation of Seasonal Variation in Economic Time Series." *Journal of the American Statistical Association.* 58 (1963): 31–44.

9. Hinch, M. J., and Talwar, P. P. "A Simple Method for Robust Regression." *Journal of the American Statistical Association* 70 (1975): 113–119.

10. Jorgenson, D. W. "Minimum Variance Linear Unbiased Seasonal Adjustment of Economic Time Series." *Journal of the American Statistical Association* 59 (1964): 681–724.

11. ———. "Seasonal Adjustment of Data for Econometric Analysis." *Journal of the American Statistical Association* 62 (1967): 147–150.

12. Mehra, R. K. "Identification of Stochastic Linear Dynamic Systems Using Kalman Filter Representation." *American Institute of Aeronautics and Astonautics Journal* 9 (1971): 28–31.

13. ———. "On the Identification of Variances and Kalman Filtering." *IEEE Transactions on Automatic Control* 15 (1970): 175–184.

14. Nerlove, M. "A Comparison of Modified 'Hannan' and the BLS Seasonal Adjustment Filters." *Journal of the American Statistical Association* 60 (1965): 442–491.

15. ———. "Spectral Analysis of Seasonal Adjustment Procedures." *Econometrica* 32 (1964): 241–286.

16. Pagan, A. "A Note on Extraction of Components from Time Series." *Econometrica* 43 (1975): 163–167.

17. Rosenblatt, Harry M. "Spectral Evaluation of BLS and Census Revised Seasonal Adjustment Procedure." *Journal of the American Statistical Association* 63 (1968): 472–501.

18. Sims, Christopher A. "Seasonality in Regression." *Journal of the American Statistical Association* 69 (1974): 618–626.

19. Stephenson, J. A., and Farr, H. T. "Seasonal Adjustment of Economic Data by the Application of the General Linear Statistical Model." *Journal of the American Statistical Association* 67 (1972): 37–45.

20. Thompson, H. E., and Tiao, George C. "Analysis of Telephone Data: A Case Study of Forecasting Seasonal Time Series." *The Bell Journal of Economics and Management Science* 2 (1971): 515–541.

21. Wallis, K. F. "Seasonal Adjustment and Relationships Between Variables." *Journal of the American Statistical Association* 69 (1974): 18–31.