J. R. Statist. Soc. A (1980),
143, *Part* 3, *pp.* 321–337

Seasonal Adjustment by Signal Extraction

By J. P. Burman

Bank of England

Summary

If an ARIMA model has been fitted to a time series, the model spectrum can be partitioned into trend, seasonal and irregular components. The corresponding linear filters are used for signal extraction to provide a theoretically based method of seasonal adjustment. The flexibility, stability and residual seasonality obtained by this method and two others are compared empirically.

Keywords: SEASONAL ADJUSTMENT; ARIMA MODEL; SPECTRUM; SIGNAL EXTRACTION; PARTIAL FRACTIONS; HARMONIC FUNCTION; AUTOCOVARIANCE GENERATING FUNCTION; LINEAR FILTER; FORECAST; BACKCAST; TREND; IRREGULAR; TRANSIENT; EXTREME VALUES; ANNUAL REVISIONS; RESIDUAL SEASONALITY

1. Introduction

Seasonal adjustment as a large-scale practical technique was introduced by the US Bureau of the Census in 1957. Their method developed over the next 8 years until the publication of the X.11 version in 1965. Various other methods appeared in the next 10 years—see Bongard (1960), Burman (1965), Mesnage (1968), Nullau *et al.* (1969), Stephenson and Farr (1972), den Haan (1974) and Durbin and Murphy (1975). Stephenson and Farr is a regression method, Durbin and Murphy partly of this type; but all the others are moving average methods, in which an attempt is made to decompose the series into trend, seasonal and irregular components. This is done by a sequence of linear filters.

Why are there so many competing methods? It is because, although the decomposition is intuitively appealing, none of the methods can be shown to have optimal properties—except, perhaps, for the trend removal filter used in Burman (1965); otherwise they are all *ad hoc* techniques. Indeed, these properties depend on the purpose for which seasonal adjustments have been performed. The most common is to provide an estimate of the current trend so that judgemental short-term forecasts can be made. Alternatively, it may be applied to a large number of series which enter an economic model, as it has been found impracticable to use unadjusted data with seasonal dummies in all but the smallest models: this is often called the historical mode of seasonal adjustment.

It seems that, in the time domain, the decomposition is incapable of precise definition. However, in the frequency domain, the trend and seasonal components of a series can be more clearly defined. The seasonal component comprises the peaks in its spectrum at the basic seasonal frequency and the multiples of this, and the trend is represented by a broad peak at low frequencies. Seasonal adjustment is the operation of removing the seasonal peaks while leaving the rest of the spectrum undisturbed (though this is not entirely feasible—see below). Non-parametric estimation of spectra goes back over 20 years—for example, Blackman and Tukey (1958), but this does not provide an easy way of designing a linear filter to suit each series.

The next development was the theory of optimal filters for estimating an unobservable component of a time series (signal extraction). This was very lucidly explained in the book on prediction by Whittle (1963), which is now unfortunately out of print. Another link in the chain was the development of methods of estimating ARIMA models for time series by Box and Jenkins (1970). This provided, *inter alia*, a way of parametrizing the spectrum of a series, so that signal extraction filters could be derived from it. Box and Jenkins' prime purpose was to fit the models for forecasting, so it is intuitively attractive that they should be used for seasonal adjustment in its forecasting mode. The recent paper by Plosser (1979) shows that X.11 is not particularly helpful for this purpose. However, Statistics Canada have developed a composite method, X.11-ARIMA (Dagum, 1975, 1978), using ARIMA models to forecast and backcast one year, before applying X.11; this seems to perform better than X.11 for all except very noisy series.

The first application of signal extraction to seasonal adjustment was in an unpublished doctoral thesis by Cleveland (1972). An ARIMA model, for which the well-known X.11 method is optimal, was given in Cleveland and Tiao (1976); and another example using the standard $(0, 1, 1)\,(0, 1, 1)_{12}$ ARIMA seasonal model appeared in Box, Hillmer and Tiao (1979). The present paper shows how a general ARIMA model (with one commensense restriction on the parameters) can be used to generate an infinite linear filter which extracts the seasonal component from a series and its forecast and backcast values.

2. PARTITIONING THE SPECTRUM OF ARIMA SERIES

Suppose an infinite time series is believed to consist of two or more independent unobservable components, whose generating processes are known. Then there exist optimal linear filters to separate them, and signal extraction is the estimation and use of these filters.

Let us take an ARIMA seasonal model:

$$z_t = f(B)\,a_t = \frac{\theta(B)}{\phi(B)\,\Phi(B^s)}\,a_t, \tag{1}$$

where B is the lag operator, s the periodicity and a_t white noise; the numerator need not be separable into seasonal and non-seasonal operators. The denominator can be rearranged into components having no common factor: $\psi_m(B)$ for the trend component and $\psi_s(B)$ for the seasonal. Obviously $\phi(B)$, which contains the differencing factors $(1-B)^d$ and the non-seasonal auto-regressive part, will belong to $\psi_m(B)$, but so will part of $\Phi(B^s)$. For a monthly series, the seasonal differencing factorizes as

$$(1-B^{12})^D = (1-B)^D(1+B+B^2, \ldots, +B^{11})^D.$$

The first factor belongs to $\psi_m(B)$, having a root at unity; and the second to $\psi_s(B)$, its roots being the 12th roots of unity. These form conjugate complex pairs which generate peaks in the spectrum of the series at $\pi/6$, $2\pi/6$, ... , $5\pi/6$, plus a real root (-1) producing a peak at π.

The other roots of $\Phi(z) = 0$ are outside the unit circle. If $\Phi(z^{-1}) = 0$ has a real positive root $\Phi_1 < 1$, the factorizing is

$$1-\Phi_1\, B^{12} = (1-\mu B)(1+\mu B+\mu^2\, B^2, \ldots, +\mu^{11}\, B^{11}),$$

where μ is the real positive 12th root of Φ_1.

The first factor contributes to the peak in the spectrum at the origin and belongs to $\psi_m(B)$; the second contains 11 roots contributing to the seasonal peaks, as before, and belongs to $\psi_s(B)$. But, if Φ_1 is complex, say, for example, that

$$(\Phi_1)^{1/12} = \mu\, e^{i\alpha}$$

is the root nearest to unity. Since normally the coefficients in $\Phi(z^{-1}) = 0$ are real, there will be a conjugate $\bar{\Phi}_1$, which has a 12th root $\mu\, e^{-i\alpha}$. Taking the two sets of 12th roots together, they form 12 pairs displaced by an angle α each side of the 12th roots of unity. The result is a series of pairs of spectral peaks on each side of the seasonal frequencies, like the Zeeman effect in physical spectra. If Φ_1 is negative, the extreme case, the peaks occur midway between the seasonal frequencies (which correspond to the odd order harmonics of a 2-year cycle). Thus the distinction between trend and seasonal can be made only for those auto-regressive seasonal models in which $\Phi(z^{-1}) = 0$ has a real positive roots.

The spectrum of z_t is derived from the transfer function between z_t and a_t:

$$
\begin{aligned}
g_z(\omega) &= f(e^{i\omega}) f(e^{-i\omega}) \sigma_a^2 \\
&= \frac{\theta(e^{i\omega})\,\theta(e^{-i\omega})}{\psi_m(e^{i\omega})\,\psi_m(e^{-i\omega})\,\psi_s(e^{i\omega})\,\psi_s(e^{-i\omega})}\sigma_a^2.
\end{aligned} \tag{2}
$$

The conventional division of a time series into trend, seasonal and irregular components may now be made more precise. The trend and seasonal cover the permanent characteristics of the series, responsible for the spectral peaks at the origin and at seasonal frequencies respectively. The irregular component covers the transient characteristics, i.e. it should be white noise or a low order MA process.

Let

$$z_t = m_t + s_t + r_t,$$

where

$$m_t = \text{trend} = f_m(B)\,b_t; \quad s_t = \text{seasonal} = f_s(B)\,c_t; \; r_t = \text{irregular} = f_r(B)\,d_t$$

and b_t, c_t and d_t are independent white noises. So far the components have not been precisely defined. To do so, consider their spectra. The spectrum of the trend is

$$g_m(\omega) = f_m(e^{i\omega}) f_m(e^{-i\omega}) \sigma_b^2 \tag{3}$$

and $g_s(\omega)$ and $g_r(\omega)$ are similarly defined. Then the independence of b_t, c_t and d_t means that

$$g_z(\omega) = g_m(\omega) + g_s(\omega) + g_r(\omega). \tag{4}$$

In ARIMA models, σ_a^2 is usually defined so that the coefficient of B^0 in $f(B)$ is 1. For the component functions, it is more convenient to define $\sigma_b^2 = \sigma_c^2 = \sigma_d^2 = \sigma_a^2$, which determines the coefficient of B^0 in each case. Since (2) and (3) are symmetric in ω, the model and its components must be rational functions $h_z(x)$, $h_m(x)$, etc. of $x = \cos\omega$. Thus (4) becomes

$$h_z(x) = h_m(x) + h_s(x) + h_r(x). \tag{5}$$

In Burman (1976), the author suggested that one should proceed as follows:
Let

$$U(x) = \theta(e^{i\omega})\,\theta(e^{-i\omega}); \quad V_m(x) = \psi_m(e^{i\omega})\,\psi_m(e^{-i\omega}); \quad V_s(x) = \psi_s(e^{i\omega})\,\psi_s(e^{-i\omega}).$$

Then the polynomial quotient of $h_z(x)$ can be identified with the transient component r_t, and the remainder partitioned into partial fractions identified with m_t and s_t:

$$h_z(x) = \frac{U(x)}{V_m(x)\,V_s(x)} = Q(x) + \frac{R(x)}{V_m(x)\,V_s(x)},$$

where $Q(x)$ is the quotient, assuming $h_z(x)$ is either top heavy or balanced in degree, and $R(x)$ the remainder. Since $V_m(x)$ and $V_s(x)$ have no common factors, we can find functions $\lambda(x)$, $\mu(x)$ by the usual method of calculating an H.C.F. such that

$$\lambda(x)\,V_m(x) + \mu(x)\,V_s(x) \equiv 1,$$

so

$$\frac{\mu(x)}{V_m(x)} + \frac{\lambda(x)}{V_s(x)} \equiv \frac{1}{V_m(x)\,V_s(x)}.$$

Multiply through by $R(x)$ and find quotients and remainders $R_m(x)$ and $R_s(x)$ of the left-hand side. The quotients must cancel, since this is an identity. Thus:

$$h_z(x) \equiv Q(x) + \frac{R_m(x)}{V_m(x)} + \frac{R_s(x)}{V_s(x)} \tag{6}$$

$$\equiv h_r(x)+h_m(x)+h_s(x) \quad \text{(say)}.$$

The first term is MA (q^*-p^*), where q^* is the degree of the numerator of the model and p^* is the degree of the denominator. The second and third terms are of the right character for the spectra of trend and seasonal components, but they may not be positive for all values of ω.

3. MINIMUM SIGNAL EXTRACTION

However, this partition is not unique because constants can be added to the second and third components without altering the character of the spectra. Box *et al.* pointed out that only the minimum amount of variance should be removed from the series in seasonal adjustment. So, if $g_s(\omega)$ has a minimum ε_s, we replace it by the non-negative $g_s^*(\omega) = g_s(\omega)-\varepsilon_s$, and add the same amount to the transient. Similarly, to obtain the smoothest trend, we replace $g_m(\omega)$ by $g_m^*(\omega) = g_m(\omega)-\varepsilon_m$; and $g_r(\omega)$ is replaced by $g_r^*(\omega) = g_r(\omega)+\varepsilon_s+\varepsilon_m$. For bottom-heavy models it turns out that ε_s can be slightly negative, but ε_m is a much larger positive, so the partition still produces valid spectra.

Both Cleveland and Box *et al.* assume that the irregular component is white noise, but for top-heavy models $(q^*>p^*)$ our partition will give a moving average irregular. In fact Cleveland discusses a top-heavy ARIMA model for which the X.11 seasonal adjustment method is optimal:

$$z_t = \frac{1+0{\cdot}26B+0{\cdot}3B^2-0{\cdot}32B^3}{(1-B)^2}b_t+\frac{1+0{\cdot}26B^{12}}{1+B\ldots+B^{11}}c_t+d_t.$$

It would be possible (though difficult) to rearrange our partition of a model with $q^*-p^* = 1$, so as to absorb the first degree term in $h_r(x)$ into the other two components. But the restriction on the irregular component to be white noise seems to be unnecessary. On the other hand, it should be only a low order moving average process: some of the models fitted to monthly series in Dagum (1978) would give an irregular component of order 10 or 11, which could have peaks or troughs in its spectrum at seasonal frequencies. However, later work at Statistics Canada (Lothian and Morry, 1979) indicates that the seasonal operator $(0, 1, 1)_s$ is adequate for all these series and so the irregular is of low order.

Whittle (1963) showed that the best (minimum mean square error) linear estimator of a component m_t given z_t is

$$\hat{m}_t = \frac{f_m(B)f_m(F)}{f(B)f(F)}z_t \quad (\text{since } \sigma_b^2 = \sigma_a^2)$$

Let $g_s^*(\omega)\equiv h_s^*(x)\equiv H_s(B,F)$ where $x = \frac{1}{2}(e^{i\omega}+e^{-i\omega})$ is replaced by $\frac{1}{2}(B+F)$. Then the minimum signal extraction filter for the seasonal component is

$$\begin{aligned}\frac{h_s^*(x)}{h_z(x)} &= \frac{H_s(B,F)}{\psi_s(B)\psi_s(F)}\frac{\psi_s(B)\psi_s(F)\psi_m(B)\psi_m(F)}{\theta(B)\theta(F)}\\ &= \frac{H_s(B,F)\psi_m(B)\psi_m(F)}{\theta(B)\theta(F)} = \frac{C_s(B,F)}{\theta(B)\theta(F)} \quad \text{(say)}, \end{aligned} \tag{7}$$

where $C_s(B,F)$ is a symmetric polynomial in B and F of degree p^*. Similarly, the minimum trend removal filter is

$$\frac{h_m^*(x)}{h_z(x)} = \frac{C_m(B,F)}{\theta(B)\theta(F)} \quad \text{(say)} \tag{8}$$

where $C_m(B,F)$ is also symmetric and of degree p^*. Whittle's original formulation of signal extraction applies to a doubly infinite series and filter. Cleveland (1972) proved that the expected values of the signal series can be obtained by extending the orginal series with forecasts and backcasts. A referee has pointed out that this is a case of the general formula:

$E(X \mid Z) = E[E(X \mid Z, Y) \mid Z]$ where Z is a finite series, Y the unobserved future and past values of the series and X the signal component. In practice, the parameters of seasonal models can be close to the boundary of invertibility, which causes very slow convergence of the filter, so that more than 1000 forecasts and backcasts could be needed for reasonable accuracy.

However, this difficulty can be completely avoided by a most ingenious suggestion made to the author by Dr G. Tunnicliffe Wilson. First, the two-sided filters (7) and (8) are each partitioned into two one-sided filters:

$$\frac{C(B,F)}{\theta(B)\,\theta(F)} \equiv \frac{G(B)}{\theta(B)} + \frac{G(F)}{\theta(F)}, \tag{9}$$

where $G(\ )$ is a polynomial of degree $r = \max(p^*, q^*)$. This identity gives rise to $(r+1)$ equations which determine the $(r+1)$ coefficients of $G(\ )$. The details are in the Appendix.

Secondly, we apply the forward and backward filters to z_t, which is now assumed to be extended by forecasts and backcasts.

Let

$$x_{1t} = \frac{G(F)}{\theta(F)} z_t = f_1(F)\, z_t \quad \text{(say)},$$

where f_1 is an infinite series. Forecast z_t:

$$\Phi^*(B)\, z_t = \theta(B)\, a_t \quad (t = N+1, \ldots, N+q^*+r),$$

where $\Phi^*(B) = \phi(B)\,\Phi(B^s)$ and N is the number of observations. Construct an intermediate series:

$$w_t = G(F)\, z_t \quad (1 \leqslant t \leqslant N+q^*)$$

Now

$$\begin{aligned} \Phi^*(B)\, x_{1t} &= \Phi^*(B) f_1(F)\, z_t = f_1(F)\, \Phi^*(B)\, z_t \\ &= 0 \text{ for } t \geqslant N+q^*+1. \end{aligned}$$

Thus we have (p^*+q^*) equations to find x_{1t} $(t = N+q^*-p^*+1, \ldots N+2q^*)$:

$$\left.\begin{aligned} \theta(F)\, x_{1t} &= w_t (t = N+q^*-p^*+1, \ldots, N+q^*), \\ \Phi^*(B)\, x_{1t} &= 0 (t = N+q^*+1, \ldots, N+2q^*). \end{aligned}\right\} \tag{10}$$

The remaining x_{1t} can be found recursively from the relation in the first part of (10), working backwards to $t = 1$. The mirror image of these steps is applied to the backcast of z_t to give x_{2t}. Finally, the filtered component is the sum of x_{1t} and x_{2t}. The whole process is applied with $G_m(\ )$ for the trend and $G_s(\ )$ for the seasonal component. Details of the method and the matrix of equations (10) are in the Appendix.

The application of the seasonal filter to the original series, together with its forecast and backcast values, constitutes the Minimum Seasonal Extraction method. It will be called hereafter MSX (as MSE is already in use).

It is important at this stage to examine the reason for the paradox cited by Grether and Nerlove (1970)—that a seasonally adjusted series always has dips in its spectrum at seasonal frequencies. Tukey has pointed out that this is analogous to the fact that the fitted residuals in a linear regression are not an independent white noise series. Let $g_m(\omega) + g_r(\omega) = g_y(\omega)$, the spectrum of the adjusted series. The symmetric filter for extracting the seasonal component has a transfer function which is just the square of the filter:

$$\left\{\frac{g_s(\omega)}{g_z(\omega)}\right\}^2.$$

So the *estimated* spectrum of this component is

$$\hat{g}_s(\omega) = \left\{\frac{g_s(\omega)}{g_z(\omega)}\right\}^2 g_z(\omega) = \frac{\{g_s(\omega)\}^2}{\{g_s(\omega)+g_y(\omega)\}^2} g_z(\omega).$$

The estimated spectrum of the adjusted series is

$$\hat{g}_y(\omega) = \frac{\{g_y(\omega)\}^2}{\{g_s(\omega)+g_y(\omega)\}^2} g_z(\omega)$$

so $\hat{g}_s(\omega)+\hat{g}_y(\omega) < g_z(\omega)$.

There is a "deficiency" in both spectra where $g_s(\omega)\, g_y(\omega) \neq 0$, that is, near the seasonal peaks. Not only are there seasonal dips in the estimated spectrum of the adjusted series, but the peaks in the estimated spectrum of the seasonal component are lower than they should theoretically be. It is suggested that this deficiency should be called "the silent spectrum", since it does not relate to either time series component, but only to the cross-spectrum.

4. The Program

A Fortran program has been written for MSX; this is in two parts. Part 1 estimates an ARIMA model by maximum likelihood, that is, without backcasting, using the very efficient method described by Osborn (1977). (The method is only full ML for a pure IMA model.) The parameter values and a sufficient number of forecasts and backcasts are passed over to Part 2. This partitions the model spectrum, generates the trend and seasonal filters, and applies them to the series together with its forecasts and backcasts. The forward estimates of the parameters are used for the backcasts, since the backcast ML estimates have been found to be almost identical to the forecast estimates.

The program can handle any ARIMA model up to the third order for the non-seasonal parameters, second order for the MA seasonal and first order for the AR seasonal. It would be easy to remove these restrictions on Q and P, but so far no need has been found. The effective number of observations of changes in the seasonal pattern is small, so that number of seasonal parameters that can be identified must also be small.

Many of the operations to obtain the partition of the spectrum $g_z(\omega)$ have been described in terms of polynomial functions of $x = \cos\omega$. But, in writing the computer program, it was realized that a more convenient representation is in terms of harmonic functions, i.e. linear in $\cos\omega$, $\cos 2\omega$, $\cos 3\omega$, etc. There is a (1, 1) correspondence between the two representations, and multiplication and division of functions is only slightly harder in the harmonic forms.

The seasonal filter has a transfer function with a local minimum at zero and others close to $2\pi(j-\frac{1}{2})/s$ $(j = 2, \ldots, s)$. For the model $(0,1,1)(0,1,1)_s$ the minimum minimorum is at either 0 or $\pi(s-1)/s$ (see Burman, 1976); but this is not true generally. The trend filter usually has a simple minimum at π, but a complex model has been constructed which has a second minimum near $\pi/2$.

All widely used methods of seasonal adjustment have a procedure for modifying extreme values. For MSX, we make preliminary estimates $\hat{m}_t$ and $\hat{s}_t$ and set $\hat{r}_t = z_t - \hat{m}_t - \hat{s}_t$. If the RMS of $\hat{r}_t$ is $\hat{\sigma}_r$, multiples α, β of the latter are chosen to determine a modified series:

$$z'_t = z_t - \mathrm{MOD}_t.$$

A term is classed as a partial extreme if $|\hat{r}_t| > \alpha\hat{\sigma}_r$ and a full extreme if $|\hat{r}_t| > \beta\hat{\sigma}_r$.

For an *isolated* extreme the program takes

$$\mathrm{MOD}_t = \lambda_t \hat{r}_t / w_0, \tag{11}$$

where

$$\lambda_t = (|\hat{r}_t| - \alpha\hat{\sigma}_r)/\{(\beta-\alpha)\,\hat{\sigma}_r\} \quad \text{for } \alpha\hat{\sigma}_r < |\hat{r}_t| < \beta\hat{\sigma}_r$$
$$= 1 \quad \text{for } |\hat{r}_t| > \beta\hat{\sigma}_r$$

and

$$w_0 = 1 - 2g_{m0} - 2g_{s0}$$

The scaling factor w_0 is needed because a fraction $2g_{m0}$ of z_t enters the trend estimate, and a fraction $2g_{s0}$ enters the seasonal estimate; where g_{m0} is the coefficient of B^0 in $G_m(B)/\theta(B)$ and g_{s0} similarly in $G_s(B)/\theta(B)$—see equation (9). Thus when r_t is extreme, $\hat{r}_t$ understates the extent of this, unless corrected by w_0.

Extremes quite often occur in *pairs* of opposite sign, especially in series of flows. If these represent displacement effects, it seems natural to adjust them together. Thus, if extremes occur at $(t-1)$ and t, a natural choice would be:

$$\left.\begin{aligned} \text{MOD}'_{t-1} &= -\tfrac{1}{2}(\hat{r}_t - \hat{r}_{t-1}),\\ \text{MOD}'_t &= \tfrac{1}{2}(\hat{r}_t - \hat{r}_{t-1}). \end{aligned}\right\} \tag{12}$$

Continuity between isolated and paired extremes (of opposite sign) can be obtained by only applying (12), when $\lambda_t = \lambda_{t-1} = 1$. For intermediate cases one could take an average of (11) and (12):

$$\text{MOD}^*_{t-1} = (1-\lambda_t)\,\text{MOD}_{t-1} - \lambda_{t-1}\,\lambda_t\,\tfrac{1}{2}(\hat{r}_t - \hat{r}_{t-1}),$$
$$\text{MOD}^*_t = (1-\lambda_{t-1})\,\text{MOD}_t + \lambda_{t-1}\,\lambda_t\,\tfrac{1}{2}(\hat{r}_t - \hat{r}_{t-1}).$$

Let

$$\text{CG (centre of gravity)} = \tfrac{1}{2}(\hat{r}_{t-1} + \hat{r}_t).$$

Then

$$\text{MOD}^*_{t-1} = (1-\lambda_t)\,\text{MOD}_{t-1} + \lambda_{t-1}\,\lambda_t(\hat{r}_{t-1} - \text{CG}),$$
$$\text{MOD}^*_t = (1-\lambda_{t-1})\,\text{MOD}_t + \lambda_{t-1}\,\lambda_t(\hat{r}_t - \text{CG}).$$

An alternative, which generalizes more easily to cover triplets, has been embodied in the program:

$$\left.\begin{aligned} \text{CG} &= (\lambda_{t-1}\,r_{t-1} + \lambda_t\,r_t)/(\lambda_{t-1} + \lambda_t),\\ \text{MOD}^*_{t-1} &= (1-\lambda_t)\,\text{MOD}_{t-1} + \lambda_{t-1}(\hat{r}_{t-1} - \text{CG}),\\ \text{MOD}^*_t &= (1-\lambda_{t-1})\,\text{MOD}_t + \lambda_t(\hat{r}_t - \text{CG}). \end{aligned}\right\} \tag{13}$$

There is no obvious reason for the occurrence of pairs of alternating extremes in a series of levels, but they appear nevertheless. More likely *a priori* would be a triplet of alternating extremes, in which the side terms are "shadows" caused by a genuine extreme in the middle. Since such triplets have not yet been met with in practice, no provision has been made for them in the program. Estimates for pairs of the same sign influence each other and are therefore obtained by solving two simultaneous equations, which are an extension of (11). At present α and β are taken as 2·0 and 2·5.

After modification of extremes, the same ARIMA model is refitted to z'_t and revised trend and seasonal components estimated by MSX. The latter are subtracted from the *original* series to give revised residuals, following the convention that extremes should not be modified in a seasonally adjusted series.

The treatment of bias in multiplicative models also needs a mention. For an additive model, the seasonal filter (7) contains a factor $\psi_m(B)$, which renders the series stationary, so the expected values of the z_t and the seasonally adjusted series y_t are equal, apart from seasonal means. Hence the annual arithmetic means (AAM) of the two series have equal expectations. The same argument, applied to the residual filter generating r_t, shows that the AAM of y_t and the trend m_t have equal expectations. These statements, applied to the logarithms of the series in the multiplicative case, show that the annual geometric means (AGM) of the untransformed series have equal expectations. An AM is greater than a GM by an amount which increases with the variance of the terms contributing to the mean. The variance of z_t is greatest and of m_t least, so

$$\mathrm{AAM}(z_t) > \mathrm{AAM}(y_t) > \mathrm{AAM}(m_t).$$

The implied "bias" in y_t relative to z_t can be handled in various ways. In the Bank of England program the AAMs of the seasonal component are scaled separately in calendar years. To avoid the possibility of jumps between calendar years, especially in the trend, it seems better to apply a single bias correction over the whole series. MSX calculates the overall mean of the seasonal factors and the factors representing the irregular component: call these $\bar{b}_1$ and $\bar{b}_2$. Then y_t is scaled up by b_1 and m_t by $\bar{b}_1 \bar{b}_2$. In practice, $\bar{b}_1$ rarely exceeds unity by more than 1 per cent and $\bar{b}_2$ is much smaller.

5. SOME EXAMPLES

Mr Kenny, of the CSO, kindly supplied the author with seven monthly seasonal series:

1. Average earnings (1963–76).
2. Retail sales (1961–76).
3. Commercial vehicles production (1958–76).
4. Unemployment in GB (excluding school leavers under 18) (1958–76).
5. Domestic furniture deliveries (1963–76).
6. Passenger cars production (1958–76).
7. Engineering orders on hand (1958–76).

The lengths of the series vary from 14 to 19 years. Model identification was in two stages: first a standard model $(0,1,1)(0,1,1)_{12}$ was fitted to the full-length series and the first 24 autocorrelations of the residuals examined, using the Ljung and Box (1978) Q-test. For series with substantial growth a logarithmic transformation was first applied. Fits at the 5 per cent significance level were obtained for Series 1, 2 and 6. For Series 4 and 7 the autocorrelations suggested that an AR factor or extra non-seasonal differencing was needed, so the model was extended to $(1,1,2)(0,1,1)_{12}$. The estimation program automatically changes $(1-\phi B)$ into $(1-B)$ if ϕ exceeds 0·96 (about 1 standard error from the stationarity boundary) and re-estimates; it also removes the highest order θ or ϕ, if its coefficient is insignificant (i.e. less than its standard error). For Series 5 a top-heavy model $(0,1,2)(0,1,1)_{12}$ was fitted. At this stage all models had Q-values below the 1 per cent level except that for Series 3 which was a little above, but the first 6 and the 12th autocorrelations of the residuals of this series were all small and no further model extension was indicated.

The series were then progressively truncated and the same models fitted. The models are shown in Table 1: the suffix 12 has been omitted to save space.

Generally the more complex models were still needed for the shorter series, but, for some of 7 and 8 years' length, Θ exceeded 0·96: this causes cancellation of the seasonal factors in the model and the program instead removes a fixed deterministic component (following Pierce, 1976). For the resulting series the $(1,0,0)_s$ operator may be included in the model, but in the few cases tested so far Φ was either negative or insignificantly positive. The program therefore fits the non-seasonal part of the original model, as this is needed by MSX to determine the extreme values.

Attempts to fit more complex seasonal models like $(1,0,1)_{12}$ and $(0,2,2)_{12}$ were unsuccessful and led to ill-determined parameters and a worse fit than $(0,1,1)_{12}$. Our inability to identify $(1,0,1)_{12}$ models contrasts with Pierce's success in doing so. A possible reason for this is that in his method the non-seasonal and seasonal parts of the model are fitted in succession instead of simultaneously.

TABLE 1
Models fitted

Series	Final year of estimation period: 7	8	$n-3$	$n-2$	$n-1$	n	Series length = n (years)
1†	(011) (000)	(011) (011)	(011) (011)	(011) (011)	(011) (011)	(011) (011)	14
2†	(011) (011)	(011) (011)	(011) (011)	(011) (011)	(011) (011)	(011) (011)	16
3	(012) (011)‡	(021) (011)‡	(011) (011)	(011) (011)	(011) (011)	(011) (011)	19
4	(112) (000)	(112) (000)	(112) (011)	(112) (011)	(112) (011)	(112) (011)	19
5†	(012) (011)	(012) (011)	(012) (011)	(012) (011)	(012) (011)	(012) (011)	14
5A†§	(011) (011)	(011) (011)	(011) (011)	(011) (011)	(011) (011)	(011) (011)	14
6	(011) (011)	(011) (011)	(011) (011)	(011) (011)	(011) (011)	(011) (011)	19
7	(111) (000)	(111) (000)	(111) (011)	(021) (011)	(021) (011)	(021) (011)	19

† Multiplicative model.
‡ Fixed seasonals on final round.
§ Series 5 adjusted for date of Easter (see below).

TABLE 2a
Preliminary values of Θ

Series	Last year of estimation period: 7	8	$n-3$	$n-2$	$n-1$	n	Series length $n =$
1	Fixed†	0·93	0·79	0·62	0·66	0·71	14
2	0·72	0·68	0·72	0·67	0·67	0·73	16
3	0·96	0·87	0·80	0·82	0·82	0·83	19
4	Fixed†	Fixed†	0·75	0·72	0·72	0·69	19
5A	0·38	0·53	0·48	0·64	0·61	0·69	14
6	0·65	0·69	0·76	0·77	0·81	0·81	19
7	Fixed†	Fixed†	0·75	0·80	0·82	0·83	19

† Reached boundary imposed by program.

TABLE 2b
Final values of Θ

Series	Last year of estimation period: 7	8	$n-3$	$n-2$	$n-1$	n	Series length $n =$
1	Fixed†	0·74	0·68	0·57	0·58	0·65	14
2	0·68	0·62	0·62	0·56	0·67	0·71	16
3	Fixed†	Fixed†	0·74	0·75	0·75	0·77	19
4	Fixed†	Fixed†	0·73	0·68	0·69	0·66	19
5A	0·27	0·50	0·44	0·63	0·57	0·64	14
6	0·62	0·65	0·62	0·65	0·70	0·71	19
7	Fixed†	Fixed†	0·66	0·74	0·78	0·79	19

† Reached boundary imposed by program.

The fitted models were used in MSX, as described above, and the extremes modified. It was noticed that for Series 5 (furniture deliveries) extreme residuals were concentrated in March and April, and that the pattern of these pairs (of opposite sign) was almost perfectly correlated with the position of Easter: normally April is a low month for deliveries, but, when Easter fell in March, that month was low. A new series (5A) was created in which a switch from April to March was made in the 4 years with an early Easter. This switch was estimated from the ratios March/(February, April)$^{\frac{1}{2}}$ over the whole 14 years (a convenient, but not optimal procedure). The parameter θ_2 then became insignificant throughout, and the fit of the model improved considerably, so no further results are given for the original Series 5.

Tables 2a and 2b show the preliminary estimates of Θ and the final ones based on the series modified for extremes.

The final values are always lower than the preliminary ones. This would be an advantage if there really is a moving seasonal pattern, which is easier to detect when some of the noise has been removed. For Series 1 (12–14 years) θ_1 became insignificant on the second round, but dropping it made the fit much worse. It seems that the very slight seasonal pattern of this series makes the model parameters ill-determined. The model with θ_1 was therefore retained. Preliminary and final values are quite close together, except for Series 2 (13–14 years), Series 6 (16–19 years), and—not surprisingly—some of the 7- and 8-year runs. Apart from length of series, the differences are linked with the number and size of extremes—Series 6 has the largest number (14). Even after the Easter adjustment, Series 5A has Θ values which are surprisingly low and variable for less than 12 years' data.

TABLE 3
Absorption of a_t into trend and seasonal components (full length)

	Series						
	1	2	3	4	5A	6	7
σ_r/σ_a	0·35	0·55	0·71	0·13	0·44	0·63	0·23
w_0	0·50	0·68	0·81	0·30	0·59	0·75	0·43

The proportion of the innovation variance absorbed by the irregular component (σ_r/σ_a) normally varies between 45 and 70 per cent for the balanced or top-heavy models—see Table 3, but is much lower for the bottom-heavy models (Series 7). This is because $g_r^*(\omega) = \varepsilon_s + \varepsilon_m$, which is usually fairly small. The proportion is also rather low for Series 1, whose small θ_1 makes it akin to a bottom-heavy model and very low for Series 4, which has a similar tendency, as θ_2 is small. In these cases, the trend filter picks up a large part of any extreme values, so that w_0 in equation (11) is also low.

6. COMPARISON WITH OTHER METHODS

To compare objectively different methods of adjustment is not easy: time series charts, relying on visual judgement, are useless, except for eliminating very inferior methods. Spectra of the adjusted and unadjusted series are a limited help, but tests in the time domain are more sensitive. The most important of these are tests for residual seasonality (has the method removed enough?) and for stability (has it removed too much, i.e. some of the noise?). If the latter is the case, revisions—normally annual—will tend to be larger, and this is something that both producers and users wish to avoid. In this section MSX is compared with the well-known X.11 method and also the Bank of England's official method (called here BE)—see Burman (1965). The same limits for extremes (2·0σ and 2·5σ) were employed throughout.

The stability of any method of seasonal adjustment depends on the degree of smoothing. As explained in Burman (1965), there is a trade-off between less smoothing, less stability, and greater sensitivity in following changes in the pattern, on the one hand, and more smoothing, more stability and less sensitivity on the other hand.

X.11 smooths with a [3] [5] moving average except the last 3 (and first 3) years.

BE smooths each of the harmonic components of the seasonal pattern independently, choosing from a range of filters: fixed, exponential weights and [3] [5]—the most flexible choice.

MSX smooths all components in the same way, the weights being determined by the Θ in the model.

Table 4 gives a measure of the flexibility or movement in the seasonal pattern estimated by the three methods over 8-year series and the full-length series. It was expected that, because of its range of smoothing filters, BE would be less flexible than X.11, but this was only true in 9 out of 14 cases. MSX was less flexible then X.11 in 11 cases. The relatively high flexibility of the MSX seasonals for Series 5A (8 years) reflects the low values of Θ in Table 2b.

TABLE 4
Flexibility
(*Mean absolute year on year changes in seasonal component, expressed as percentage of series mean*†)

	8 years			*Full length*		
Series	*X.11*	*BE*	*MSX*	*X.11*	*BE*	*MSX*
1	0·06	0·05	0·03	0·12	0·11	0·10
2	0·11	0·14	0·12	0·17	0·21	0·11
3	0·49	0·17	0‡	1·08	0·57	0·48
4	0·16	0·11	0‡	0·21	0·16	0·25
5A	0·51	0·61	0·72	0·47	0·56	0·42
6	0·87	0·34	0·86	1·16	0·58	0·79
7	0·08	0·14	0‡	0·12	0·10	0·06

† For multiplicative adjustment, the table shows the mean absolute changes in the seasonal factors (expressed as percentages).
‡ Fixed pattern.

When another year's data are added, the mean absolute revision (m.a.r.) to the seasonals in the last year of the series is likely to be relatively large; for the 2nd last year it should be less, the 3rd last year less still, and so on. Eventually the m.a.r. should settle down at a low level (when estimating filter has become nearly symmetric) or even drop to zero (when the filter is truncated as in X.11). Table 5 shows the m.a.r. for different parts of the series expressed as percentages (m.a.p.r.) of the mean of the shorter series. For multiplicative adjustment, it shows the mean absolute changes in the seasonal factors. The comparisons are made between MSX.1—the method as described so far—and the other two methods (see below for explanation of MSX.2). Column 1 gives the m.a.p.r. for the first 7 years of the 7-year and 8-year series; column 2 contains the same calculation for the 7th year of the two series. Columns 3–6 show the m.a.p.r. for the last 4 years of a series when 1 year is added; but, to reduce sampling fluctuations, they have been averaged over 3 pairwise comparisons (for example, for Series 4, these are 16 *vs* 17 years, 17 *vs* 18 years and 18 *vs* 19 years). Column 7 provides the m.a.p.r. over all years from the comparison $(n-1)$ *vs* n years.

As expected, the m.a.p.r. in columns 3–6 descend from right to left, X.11 usually more steeply than the others, but the revisons for the 4th last year are still larger than the average of all years (column 7). For the 7–8 year runs the m.a.p.r. are larger than those for the longer runs, and again they are larger for the 7th year than for the average of all 7 years (except when MSX produces a fixed pattern).

TABLE 5

Mean absolute percentage revisions to seasonal component on adding one year's data

	Series	*Lengths of estimation periods*: 7 *vs* 8 — *All years*	7 *vs* 8 — *7th*	*Averages of* 3 *pairs: (n – 3) vs (n – 2), (n – 2) vs (n – 1), and (n – 1) vs n* — *4th last*	*3th last*	*2nd last*	*Last*	(n – 1) *vs n* — *All years*	*Full length of series n =*
		Years of comparison							
1	X.11	0·14	0·16	0·07	0·14	0·21	0·30	0·06	14
	BE	0·06	0·07	0·07	0·13	0·20	0·28	0·07	
	MSX.1	0·06	0·10‡	0·07	0·13	0·22	0·32	0·08	
	MSX.2	0·09	0·10‡	0·09	0·15	0·24	0·33	0·08	
2	X.11	0·06	0·19	0·12	0·19	0·31	0·43	0·07	16
	BE	0·09	0·13	0·21	0·31	0·32	0·39	0·14	
	MSX.1	0·09	0·21	0·17	0·22	0·33	0·47	0·07	
	MSX.2	0·12	0·26	0·15	0·21	0·30	0·41	0·09	
3	X.11	1.62	3·57	0·48	0·79	1·23	1·80	0·28	19
	BE	1·30	1·38	0·76	0·87	1·05	1·27	0·44	
	MSX.1	0·69†	0·69†	0·55	0·69	0·89	1·15	0·42	
	MSX.2	1·02	1·13	0·55	0·64	0·75	0·90	0·39	
4	X.11	0·37	0·82	0·20	0·40	0·60	0·80	0·11	19
	BE	0·44	0·43	0·39	0·46	0·59	0·76	0·24	
	MSX.1	0·43†	0·43†	0·27	0·38	0·54	0·78	0·12	
	MSX.2	0·43†	0·43†	0·27	0·37	0·52	0·72	0·12	
5A	X.11	0·17	0·57	0·30	0·53	0·82	1·15	0·20	14
	BE	0·45	0·83	0·51	0·60	0·88	1·05	0·40	
	MSX.1	0·75	1·51	0·38	0·62	0·90	1·26	0·34	
	MSX.2	0·55	1·22	0·40	0·63	0·89	1·22	0·36	
6	X.11	1·30	3·45	0·88	1·60	2·32	2·97	0·36	19
	BE	1·68	2·20	0·99	1·22	1·26	1·53	0·66	
	MSX.1	1.21	2.35	1·07	1·35	1·72	2·25	0·46	
	MSX.2	1·25	2·20	1·04	1·22	1·43	1·66	0·48	
7	X.11	0.11	0·27	0·10	0·16	0·25	0·36	0·04	19
	BE	0·22	0·49	0·13	0·18	0·26	0·32	0·08	
	MSX.1	0·09†	0·09†	0·08	0·10	0·12	0·15	0·04	
	MSX.2	0·09†	0·09†	0·08	0·12	0·11	0·12	0·05	

† Fixed seasonals for 7 and 8 years.
‡ Fixed seasonals for 7 years.

For *Series 1* revisions are small for all three methods, which is not surprising, since the series displays very little seasonality. For *Series 2*, which has stronger seasonality, again all methods are very stable, with X.11 and MSX.1 close together. For *Series 3, 4 and 7*, MSX.1 is the most stable for the longer runs (columns 4–6), though twice X.11 overtakes it in the 4th last year; and MSX.1 is far more stable for the 7–8 year runs, since it selects fixed patterns. *Series 5A's* results are contradictory: X.11 does best by a small margin for the longer runs, but is by far the most stable for 7–8 years. MSX.1's poor showing in the latter case is due to its low values of Θ (see Table 2b), leading to high flexibility (Table 4). Finally, for *Series 6* (longer runs) the order of stability is: BE, MSX.1, X.11, until the 4th last year is reached, when X.11 is best.

Summing up, if stability near the end of a series is more important than the average over the whole series, MSX.1 does very well, in most cases. X.11 is more stable on average over a whole series, but the largest revisions in the table occur with this method (over 3 per cent in column 2, Series 3 and 6). However, MSX.1 has two unsatisfactory features: firstly, it can become very unstable for a short series (Series 5A); and secondly, for bottom-heavy, or nearly bottom-heavy, models, the irregular component is much smaller than for the other methods, because so much goes into the trend. Consequently, it identifies fewer extremes in Series 4 (unemployment) than X.11 or BE in the exceptionally cold winter of 1962–63, and the modifications are much smaller. The second defect does not seem very serious, because MSX.1 is no less stable than the other methods for series 4 and 7.

The instability in Series 5A is undoubtedly partly due to the remaining variability in March and April, but its behaviour does suggest two general lines for further research. Firstly, at what length of series does one start to estimate *moving* seasonality? For BE it is introduced at 7 years; and probably the same should apply for MSX. It would then be natural to introduce a lower bound for Θ in short series, which would be gradually relaxed as the length increased to (say) 10 years; and above that it would be dropped.

Secondly, a more fundamental point, the modification of extremes lowers Θ for every series, thus increasing flexibility and lowering stability. Examination of the preliminary and final seasonal components shows that modification of extremes nearly always produces a more flexible pattern in the months in which they occur. By contrast, with BE the smoothing parameters λ vary both ways between preliminary and final seasonals, and the (simple) averages of these values are close together.

7. EXTREMES, STABILITY AND RESIDUAL SEASONALITY

If there are no prior reasons for expecting an extreme in certain months (e.g. strikes, exceptional weather), is there any evidence that extremes represent departures from normality? For the seven series in our sample, the absolute values of the residuals in MSX exceed $2\hat{\sigma}_r$ for 4–6 per cent of the observations; but they exceed $2{\cdot}5\hat{\sigma}_r$ for an average of $2\frac{1}{2}$ per cent of the observations—twice the expected proportion for a normal distribution. So the modification of extremes can be partly justified on the grounds that (with full replacement above $2{\cdot}5\sigma$) they stand outside the process generating the main series. Against this must be set the size of revisions to estimates of the extremes, which are the immediate cause of the instability of the seasonal component in certain cases.

Is there any trade-off for the lower stability caused by modifying extremes? Stability is only one half of performance; the other half is the absence of residual seasonality. There is no generally accepted test of the latter. Idempotency is one possibility (Fase, Koning and Volgenant, 1973)—running an adjusted series through the same procedure again. Another (used in BE) is to calculate von Neumann ratios for the residuals of each month, and to check whether the number of significantly low ratios in a group of series exceeds the number expected. Yet another test is to see if the spectrum of an adjusted series is smooth and removes no power at inter-seasonal frequencies. Spectra were estimated for the three seasonally adjusted versions of each full-length series, and proved to be very close together—the only interesting feature being the appearance of the predicted Grether–Nerlove "dips" at some of the seasonal frequencies. (Copies of the charts may be obtained from the author.)

The method finally chosen for testing residual seasonality was as follows: fit a non-seasonal ARIMA model to each seasonally adjusted series, and calculate r_{12}, r_{24} and r_{36} for the residuals; then the analogue of the Ljung–Box test is

$$Q_s = n(n+2) \sum_{j=1}^{3} r_{12j}^2/(n-12j)$$

where n is the number of terms in the differenced series. The probability distribution of Q_s is not known exactly—see Pierce (1976); it is thought to lie between $\chi^2(2)$ and $\chi^2(3)$ for a one-parameter model, most probably close to $\chi^2(3)$. Table 6 shows the values of Q_s for each method, using Series 2–7: first, the preliminary seasonal adjustments applied to the original series; second, the final seasonal adjustments applied to the *modified* series. All but one of the r_{12} are negative, three-quarters of the r_{24} and all the r_{36}—this apparent "over-adjustment" is the time-domain equivalent of the Grether–Nerlove effect. We note first that the large majority of the series have Q_s significant at the nominal 1 per cent level, but since the true distribution is unknown, it is difficult to draw any conclusions from this. Secondly, it is remarkable that, for X.11, Q_s increases from the preliminary to the final round. For BE and MSX, Q_s goes in both directions, but the shift is small with MSX.1 for Series 4 and 7, for which the estimated extremes are small; and for Series 5A the instability noted for this method in Table 5 is accompanied by a worsening of Q_s. Thirdly, the final adjustments for BE have consistently lower residual seasonality than either MSX.1 or X.11.

TABLE 6
Residual seasonality test (Q_s)

		Series					
		2	3	4	5A	6	7
MSX.1	Preliminary	15·74‡	15·43‡	18·62‡	13·02‡	20·06‡	10·88†
	Final	12·78‡	11·13†	18·09‡	18·80‡	21·67‡	11·09†
MSX.2	Final	11.77‡	7·07	15·46‡	16·63‡	10·86†	7·85
BE:	Preliminary	12·09‡	12·73‡	1·81	7·91	11·49†	19·69‡
	Final	5·83	11·10†	6·73	14·44‡	3·31	10·34†
X.11:	Preliminary	12·61‡	20·19‡	7·80	15·96‡	24·91‡	28·43‡
	Final	15·94‡	30·85‡	12·81‡	19·42‡	27·37‡	33·37‡

† Significant at 5 per cent point for $\chi^2(3)$.
‡ Significant at 1 per cent point for $\chi^2(3)$.

Since MSX.1—modifying extremes and re-estimating the model—does not show a uniform reduction in Q_s, a better trade-off between stability and residual seasonality may be possible. MSX.2 is defined to be the same as MSX.1 up to the modification of extremes: but the preliminary parameter values are used again on the second round to obtain forecasts and backcasts. Table 6 shows, surprisingly, that MSX.2 has lower Q_s values than the final seasonals of MSX.1. Referring back to Table 5, MSX.2 has almost the same stability as MSX.1 for *Series 1*, but is slightly better for *Series 2*. For *Series 3* MSX.2 is definitely more stable on the longer runs, but less on the shorter ones, because there is not a fixed seasonal pattern on the first round of estimation. For *Series 4 and 7* there is little difference between the two variants, as the identified extremes are small. The same is true of *Series 5A*, except for the 7- and 8-year runs, where the marked improvement stems from avoidance of low Θ values. Finally, for the longer runs of *Series 6* there is a dramatic improvement: it is now nearly as stable as BE.

We conclude that MSX.2 is to be preferred, and that its stability for recent observations is substantially better than that of X.11, except in the case of *Series 5A*. MSX.2 is, of course, quicker to run than MSX.1: for example, on an IBM 370/158, a 12-year monthly series takes about 8 seconds CPU time (provided reasonable starting values of the model parameters in the estimation can be obtained). This compares with 7 seconds for X.11 and 6 seconds for BE.

8. BOTTOM-HEAVY MODELS

When this paper was virtually complete, it was realized that there was a simple way of moving more of the high frequency spectral power from the trend to the irregular, which could be applied to bottom-heavy models (Series 4 and 7). Using the notation of Section 3, we first observe that $g_m(\omega)$ declines more slowly towards its minimum (say ε_0) at $\omega = \pi$ than it does for balanced models. We therefore assume that the irregular component may follow a first order moving average: $\varepsilon_s + \varepsilon_0 + \varepsilon_1(1+\cos\omega)$; and determine ε_1 by minimizing $\{g_m(\omega) - \varepsilon_0\}/(1+\cos\omega)$. In the two cases examined so far, the minimum of the latter expression is also at $\omega = \pi$, so the final trend spectrum—

$$g_m^*(\omega) = g_m(\omega) - \varepsilon_0 - \varepsilon_1(1+\cos\omega)$$

is still monotone decreasing. The new values of w_0 (see Table 3) for Series 4 and 7 are 0·45 and 0·54, and the new values of σ_r/σ_a are 0·21 and 0·29 respectively. The average magnitudes of the extremes are now only a little less than those of X.11 and BE, though the individual extremes picked out by the three methods are not always the same.

The above procedure could be applied to remove still more from the trend, if a second order moving average were acceptable for the irregular component.

9. CONCLUSIONS

What conclusions can be drawn at this stage about the use of Signal Extraction for seasonal adjustment? Obviously they must be very tentative until a much larger number of series has been tested:

(i) Trend and seasonal signal extraction filters can be derived for all suitable seasonal ARIMA models; and, although these filters are doubly infinite, their effect on the original series can be obtained by simple finite operations. It is desirable to replace or modify extreme values of the series and re-apply the filters (Section 4).

(ii) Fitting ARIMA models to a large number of series is quite practical: first using the standard $(0,1,1)(0,1,1)_s$ model and then extending it to the $(1,1,2)(0,1,1)_s$ family, if the diagnostic checks suggest it. For the latter, the non-seasonal operator is automatically simplified to $(1,1,1)$ or $(0,2,1)$ if appropriate. Initially some skilled resources are needed for model identification, but thereafter the same model can generally be used for a number of years.

(iii) Only one seasonal operator is needed in practice, but there is some variety of non-seasonal operators, leading to balanced, top-heavy or bottom-heavy models. For bottom-heavy models, too much of the current observation is absorbed into the trend, compared with X.11 and the Bank of England method, resulting in implausibly small modifications for extreme values. This "defect" can probably be overcome by assuming a moving average irregular component (Section 8).

(iv) For short series the seasonal model may be degenerate, implying a fixed seasonal pattern, so that the adjustment consists simply of subtraction of seasonal means.

(v) Generally, both signal extraction (MSX) and the Bank of England method are less flexible than X.11 and give considerably smaller revisions to the *last 3 years* of a series, without leaving any more residual seasonality than X.11. MSX often produces larger average revisions than X.11 over the *whole* series (except in the case of the shorter series where a fixed seasonal pattern is obtained). However, this disadvantage could be nullified by a sensible publication policy: for example, in official publications no amendments (solely due to revised seasonal adjustment) would be made to data more than 4 years' old; but complete revisions could be made available to research workers on request.

(vi) MSX.2, which omits re-estimation of the model parameters, seems to be more stable and, at the same time, leaves less residual seasonality than MSX.1 (which includes re-estimation).

(vii) For shorter series—under 7 years—probably no attempt should be made to find moving seasonality. For intermediate lengths (7–10 years) some lower bound could be placed on the values of Θ, to improve stability, and this restriction would be gradually relaxed as the series lengthened.

(viii) Signal extraction is now ready for large-scale trials by other statisticians. Progress from ad hoc to more theoretically optimal methods of seasonal adjustment could be rapid in the next few years.

REFERENCES

BLACKMAN, R. B. and TUKEY, J. W. (1958). *The Measurement of Power Spectra.* New York: Dover Publications.

BONGARD, J. (1960). Some remarks on moving averages. In *Seasonal Adjustment on Electronic Computers*, pp. 361–390. Paris: OECD.

BOX, G. E. P., HILLMER, S. C. and TIAO, G. C. (1979). Analysis and modelling of seasonal time series. NBER–CENSUS *Conference on Seasonal Analysis of Economic Time Series*, Washington, D.C., September 1976.

BOX, G. E. P. and JENKINS, G. M. (1970). *Time Series Forecasting and Control.* San Francisco: Holden-Day.

BOX, G. E. P. and TIAO, G. C. (1975). Intervention analysis with applications to economic and environmental problems. *J. Amer. Statist. Ass.*, **70**, 70–79.

BURMAN, J. P. (1965). Moving seasonal adjustment of economic time series. *J. R. Statist. Soc.* A, **128**, 534–558.

—— (1979). Discussion on paper by Box, Hillmer and Tiao. NBER–CENSUS *Conference on Seasonal Analysis of Economic Time Series*, Washington, D.C., September 1976.

CLEVELAND, W. P. (1972). Analysis and forecasting of seasonal time series. University of Wisconsin (unpublished Ph.D. Thesis).

CLEVELAND, W. P. and TIAO, G. C. (1976). Decomposition of seasonal time series: a model for the Census X-11 program. *J. Amer. Statist. Ass.*, **71**, 581–587.

DAGUM, E. B. (1975). Seasonal factor forecasts from ARIMA models. *International Statistical Institute. Proceedings of 40th Session*, Vol. 3, Warsaw, pp. 206–219.

—— (1978). Modelling, forecasting and seasonally-adjusting economic time series with the X-11-ARIMA method. *The Statistician*, **27**, 203–216.

DURBIN, J. and MURPHY, M. J. (1975). Seasonal adjustment based on a mixed additive–multiplicative model. *J. R. Statist. Soc.* A, **138**, 385–410.

FASE, M. M. G., KONING, J. and VOLGENANT, A. F. (1973). An experimental look at seasonal adjustment. *De Economist*, **121**, 441–480.

GEWEKE, J. (1979). The temporal and sectoral aggregation of seasonally-adjusted time series. NBER–CENSUS *Conference on Seasonal Analysis of Economic Time Series*, Washington, D.C., September 1976.

GRETHER, D. and NERLOVE, M. (1970). Some properties of "optimal" seasonal adjustment. *Econometrica*, **38**, 682–703.

HAAN, R. J. A. DEN (1974). *A Mechanised Method of Seasonal Adjustment.* The Hague: Central Planning Bureau.

LJUNG, G. M. and BOX, G. E. P. (1978). On a measure of lack of fit in time series models. *Biometrika*, **65**, 297–303.

LOTHIAN, J. and MORRY, M. (1978). Selection of models for the automated X-11-ARIMA seasonal adjustment program. *Statistics Canada Research Papers*, No. 78–10–003.

MESNAGE, M. (1968). Élimination des variations saisonnières: la nouvelle méthode de l'OSCE. *Études et enquêtes statistiques*, No. 1, 7–78.

NULLAU, B., HEILER, S., WASCH, P., MEISNER, B. and FILIP, D. (1969). The Berlin method, a contribution to time series analysis. *Contributions to structural research*, Vol. 7. Berlin; German Institute for Economic Research.

OSBORN, D. R. (1977). Exact and approximate maximum likelihood estimates for vector moving average processes. *J. R. Statist. Soc.* B, **39**, 114–118.

PIERCE, D. A. (1979). Seasonal adjustment when both deterministic and stochastic seasonality are present. NBER–CENSUS *Conference on Seasonal Analysis of Economic Time Series*, Washington D.C., September 1976.

PLOSSER, C. I. (1979). Short-term forecasting and seasonal adjustment. *J. Amer. Statist. Ass.*, **74**, 15–24.

STEPHENSON, J. A. and FARR, H. T. (1972). Seasonal adjustment of economic data by application of the general linear statistical model. *J. Amer. Statist. Ass.*, **67**, 37–45.

WHITTLE, P. (1963). *Prediction and Regulation by Linear Least-square Methods.* Princeton: D. Van Nostrand.

APPENDIX: TUNNICLIFFE WILSON ALGORITHM

Let $G(B) = g_0 + g_1 B \ldots + g_r B^r$ and $C(B,F) = c_0 + c_1(B+F) \ldots + c_r(B^r + F^r)$. From (9):

$$\theta(F)\, G(B) + \theta(B)\, G(F) \equiv C(B,F).$$

Equating coefficients of B^r, B^{r-1}, etc.,

$$\begin{aligned}
&\theta_0 g_r + \theta_r g_0 &&= c_r,\\
&\theta_1 g_r + \theta_0 g_{r-1} + \theta_r g_1 + \theta_{r-1} g_0 &&= c_{r-1}\\
&\ldots\\
&2(\theta_r g_r + \theta_{r-1} g_{r-1} \ldots + \theta_0 g_0) &&= c_0.
\end{aligned}$$

So if A is defined as

$$A = \begin{bmatrix} \theta_0 & 0 & \cdots & \cdots & 0 \\ \theta_1 & \theta_0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \theta_r & \theta_{r-1} & \cdots & \cdots & \theta_0 \end{bmatrix} + \begin{bmatrix} 0 & \cdots & \cdots & 0 & \theta_r \\ 0 & \cdots & 0 & \theta_r & \theta_{r-1} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \theta_r & \theta_{r-1} & \cdots & \cdots & \theta_0 \end{bmatrix}.$$

The equations become

$$A(g_r\, g_{r-1}, \ldots, g_0)' = (c_r\, c_{r-1}, \ldots, c_0)'.$$

(The ordering of the matrix columns puts ones in most elements of the leading diagnonal for easier inversion.)

For the second stage, let $\Phi^*(B) = \phi_0 + \phi_1 B \ldots + \phi_p B^p$ (with p written for p^*). Equations (10) are (with q written for q^*):

$$\begin{bmatrix} \theta_0 & \theta_1 & \cdots & \theta_q & 0 & \cdots & \cdots & 0 \\ 0 & \theta_0 & \theta_1 & \cdots & \theta_q & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & 0 & \theta_0 & \theta_1 & \cdots & \cdots & \theta_q \\ \phi_p & \cdots & \cdots & \phi_1 & \phi_0 & 0 & \cdots & 0 \\ 0 & \phi_p & \cdots & \cdots & \phi_1 & \phi_0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & 0 & \phi_p & \cdots & \cdots & \phi_1 & \phi_0 \end{bmatrix} \begin{bmatrix} x_{N+q-p+1} \\ \vdots \\ \vdots \\ x_{N+q} \\ \vdots \\ \vdots \\ \vdots \\ x_{N+2q} \end{bmatrix} = \begin{bmatrix} w_{N+q-p+1} \\ \vdots \\ \vdots \\ w_{N+q} \\ 0 \\ \vdots \\ \vdots \\ 0 \end{bmatrix}.$$

The matrix is of order $(p^* + q^*)$ and has ones in the leading diagonal.