COMMENTS ON "SEASONAL ADJUSTMENT WHEN BOTH DETERMINISTIC AND STOCHASTIC SEASONALITY ARE PRESENT" BY DAVID A. PIERCE

John P. Burman
Bank of England

This is a very illuminating paper, giving full practical details of seasonal signal extraction, using a particular ARIMA model.

The model $(1, 0, 1)_{12}$ used for the seasonal operator is different from the example in (1). But, both come within the framework of the partial fraction technique, described in the appendix of my discussion on Kuiper. An improvement, suggested by Pierce, is the extraction of any deterministic seasonal component (seasonal mean correction) from the differenced series.

He removes the trend by differencing and a nonseasonal autoregressive filter, instead of including this in a single seasonal ARIMA model; this makes for computational simplicity in finding the minimum of the spectrum of the seasonal component and also the weights of the signal extraction filter. But, if there is any interaction between the seasonal and nonseasonal parts of the model, this may not be the optimal procedure.

Another difference from (1) is that Pierce includes the positive real root of $(1-\Phi B^{12})$ in the seasonal model, although it is usually close to 1 and, thus, generates a spike in the spectrum at zero frequency. For example, for U.S. unemployment, Pierce finds $\Phi=0.547$, and its 12th root is 0.95.

COMMENTS ON "ANALYSIS AND MODELING OF SEASONAL TIME SERIES" BY GEORGE E. P. BOX, STEVEN HILLMER, AND GEORGE C. TIAO

John P. Burman
Bank of England

I found this a most stimulating and original approach to the problem of optimal seasonal adjustment. It combines ARIMA stochastic modeling with signal extraction and illustrates the method by an example. The method is presented in a more compact way, which shows how it can be generalized to include other Box-Jenkins models.

My discussion of Kuiper's paper sets out the partial fraction method of decomposing a seasonal model. Thus, Box, Hillmer, and Tiao's model can be written as follows:

$$\frac{\theta(B)}{\varphi(B)}=\frac{\eta(B)}{d(B)}+\frac{\Psi(B)}{U(B)}$$

where $\varphi(B)=(1-B)(1-B^{12})$

$$\eta(B)=\eta_o-\eta_1B-\eta_2B^2$$

$$d(B)=(1-B)^2$$

$$\Psi(B)=\Psi_o-\Psi_1B---\;-\Psi_{10}B^{10}$$

$$U(B)=1+B+B^2---\;+B^{11}$$

The spectrum of the model $=\sigma_a^2\,\dfrac{\theta(B)\theta(F)}{\varphi(B)\varphi(F)}$

where

$$B=e^{i\omega},\; F=B^{-1} \tag{1}$$

The symmetry makes the operator in (1) a function of terms like $B^n+F^n=2\cos n\omega$, and the latter is expressible as a polynomial of degree n in $x=\cos\omega$. The denominator is a polynomial in x of the same degree as $\varphi(B)$ and can be decomposed into factors of the form $(1-\beta_ix)$, which is equivalent (apart from a numerical factor) to $(1-\alpha_iB)(1-\alpha_iF)$, where α_i is one of the roots of the equation

$$\beta_i=2\alpha_i/(1+\alpha_i^2)$$

Thus, the partial fraction of the function of x leads to a decomposition of the form

$$\frac{\theta(B)\theta(F)}{\varphi(B)\varphi(F)}=d_0+\sum\frac{d_i}{(1-\alpha_iB)(1-\alpha_iF)}$$

with higher powers appearing when there are multiple roots.

For the model used here,

$$\frac{\theta(B)\theta(F)}{\varphi(B)\varphi(F)}=d_0+\frac{d_1}{(1-B)(1-F)}+\frac{d_2}{(1-B)^2(1-F)^2}+\frac{\Psi(B,F)}{U(B)U(F)}$$

$$=\frac{\eta(B,F)}{(1-B)^2(1-F)^2}+\frac{\Psi(B,F)}{U(B)U(F)}$$

where $\Psi(B,F)$ is a symmetric function of degree 10 and $\eta(B,F)$ is a symmetric function of degree 2 in B and F. If d_t and b_t are independent white noises, with $\sigma_d^2=\sigma_b^2=\sigma_a^2$, an equation equivalent to Box et al.'s (47) is produced. The numerator $\Psi(B,F)$ can, in principle, be factorised into $\Psi(B)\Psi(F)$, via a polynomial in cos ω; but, some roots of this may be inside the unit circle, in which case, $\Psi(B,F)$ could be negative and the solution unacceptable. However, with the (0, 1, 1) $(0,1,1)_{12}$ model, this does not occur in practice.

The minimum of the seasonal spectrum is given by

$$\epsilon^*=\min_{|B|=1}\left[\frac{\Psi(B,F)}{U(B)U(F)}\right]$$

$$=\min_{|B|=1}\left[\frac{\theta(B)\theta(F)}{\varphi(B)\varphi(F)}-\frac{\eta(B,F)}{d(B)d(F)}\right] \tag{2}$$

With an obvious modification of the paper's notation,

$$\sigma_b^2.\frac{\Psi.(B,F)}{U(B)U(F)}=\sigma_a^2\left\{\frac{\Psi(B,F)}{U(B)U(F)}-\epsilon^*\right\}$$

Hence

$$\sigma_d^2.\frac{\eta.(B,F)}{d(B)d(F)}=\sigma_a^2\left\{\frac{\eta(B,F)}{d(B)d(F)}+\epsilon^*\right\}$$

is the spectrum of the optimal seasonally adjusted series. These equations together are equivalent to (49) and (50). The optimum solution has the troughs of the spectrum of the seasonal component as deep as possible—thus, minimising the loss of power of the spectrum of the adjusted series at interseasonal frequencies.

The extraction filter for seasonal adjustment is given by

$$h(B)=\frac{\sigma^2_{d*}}{\sigma^2_a}\frac{\eta_*(B,F)}{d(B)d(F)}\cdot\frac{\varphi(B)\varphi(F)}{\theta(B)\theta(F)}$$

$$=\frac{\sigma^2_{d*}}{\sigma^2_a}\frac{\eta_*(B,F)U(B)U(F)}{\theta(B)\theta(F)}$$

$$=\frac{\{\eta(B,F)+\epsilon^* d(B)d(F)\}U(B)U(F)}{\theta(B)\theta(F)}$$

Note that factorisation of $\eta_*(B, F)$ into $\eta_*(B)\eta_*(F)$ is required for the elegant method of expansion of the filter, described in the appendix to Box, Hillmer, and Tiao's paper. However, it might be simpler to modify the expansion of $\frac{\varphi(B,F)}{\eta(B)\eta(F)}$ (in the Box et al. appendix notation). (A-1) becomes $\varphi(B, F)=\eta(B)C(B, F)$, and coefficients of $C(B, F)$ can be derived recursively, starting with the leading term F^r (r being the degree of the polynomials). (A-4) becomes

$$C(B, F)=\eta(F)X(B, F)$$

and the first $(r+1)$ coefficients of $X(B, F)$ can be obtained from a set of $(r+1)$ linear equations like (A-5).

Numerical estimation of (2) shows minima, as expected, at $\omega=0$ and close to $(j+{}^1\!/\!_2)\pi/6$ $[j=1, 2---5]$.

The results are in tables 1 and 2. The lowest minimum is always small and positive; for low values of θ, it is at the right-hand end, but increasing with θ; and for high values of θ it is at zero, and decreasing with increasing θ. The results for the model ($\theta=0$, $\Theta=0.75$) agree closely with those given in the paper—though, of course, ϵ^* is different, because we have started with a different acceptable model.

Table 1. **VALUE OF LOWEST MINIMUM ϵ^***

	Θ = 0.5	0.75	0.9
θ = 0	0.01	0.003	0.0004
0.3	.02	.006	.001
0.6	.036	.009	.002
0.9	.02	.005	.0008

Table 2. **POSITION OF LOWEST MINIMUM**

(ω as a multiple of $\pi/6$)

0	5.50	5.50	5.50
0.3	5.50	5.50	5.50
0.6	0	0	0
0.9	0	0	0

COMMENTS ON "A SURVEY AND COMPARATIVE ANALYSIS OF VARIOUS METHODS OF SEASONAL ADJUSTMENT" BY JOHN KUIPER

Dennis Farley and Stephen Zeller
Federal Reserve System

INTRODUCTION

One method of seasonal adjustment, examined by Kuiper in his paper, is X-11 ARIMA, now being employed by Statistics Canada on labor force data.[1] This intuitively appealing method, described in the following, has not been examined at the Board of Governors until now. Our comment evaluates the performance of this method of seasonal adjustment for one time series of particular interest—monthly observations on the narrowly defined money stock, M1.

The difficulty with X-11, or with any other method that employs symmetric moving averages, is that symmetry cannot be preserved at the end points of the sample of data. For example, X-11 estimates trends with a 12-month moving average and also smooths the seasonal component across years with a 3×5-moving average. Thus, symmetry is lost for data within 3 years of the end of the sample. Instead, asymmetric filters are applied, resulting in phase shifts in the adjusted data. What the X-11 ARIMA approach does is to provide X-11 with an augmented sample of data so that all, or most, of the actual data are smoothed with symmetric averages. ARIMA models are, of course, employed in generating the extra observations. The choice of an ARIMA model to generate forecasts is merely one of convenience. Structural models that incorporate seasonality could also be used, although the distinction between these two approaches is somewhat artificial. Under most conditions, there exists a correspondence between the structural and time series representations of an endogenous variable.[2]

We have selected M1 as an example of a series having current seasonally adjusted levels that receive a great deal of attention from the public and press. In addition, these data are used as an input to policy decisions by the Federal Open Market Committee (FOMC). Furthermore, it is well known that these data, as first published by the Board of Governors, Federal Reserve System, often revise substantially as the seasonal factors are reevaluated in light of additional data.[3] Seasonal factors reestimated with several additional years of data lead to a much smoother series than that derived from the first published factors.[4] Accepting these later estimates as correct implies that the current seasonally adjusted data are not providing policymakers with good information about short-run movements in the money stock. In this comment, the major emphasis is on seasonal factors in the current year and those projected for the following year.

There are serious problems in the application of any seasonal adjustment method to the money stock. For example, since M1 is composed of currency and demand deposits, having structural equations that would be specified differently, we probably should use a multivariate approach.[5] Furthermore, for a series that is at least partially controllable, the policymaker's reaction function must be introduced before we can begin to make any meaningful statements about seasonality in these data. Investigation of these issues is clearly beyond the scope of this comment. Instead, we assume that seasonal factors obtained with X-11 from the interior of a data sample are correct seasonal factors.

In the next section monthly ARIMA models for the currency and demand deposit components of M1 are identified. These models are then used to generate forecasts with 1-, 2-, and 3-year horizons. In the third section, seasonal factor estimates based on samples, augmented with forecasts, are compared to those obtained without the forecasts.

THE MODELS

The current Board of Governors' staff procedure is to seasonally adjust the currency and demand deposit components of M1 individually and then to add them together

[1]See [3].
[2]See [8].
[3]The current practice is to reestimate seasonal factors for M1 once a year. Some judgmental review does take place before publication. For a description of this process, see [7].
[4]See [2].
[5]See [4].

Dennis Farley and Stephen Zeller are economists in the Banking Section, Division of Research and Statistics, Board of Governors of the Federal Reserve System. All conclusions and recommendations are the authors' and do not necessarily reflect the opinions held by the Board of Governors.

to get seasonally adjusted M1. For each component, we find that first and seasonal differencing are necessary for stationarity but that a logarithmic transformation is not required. Integrated moving average models of the following general form are then estimated:[6]

$$\nabla_{12}\nabla_1 x_t=\theta(B)a_t \tag{1}$$

where x_t is either the currency or demand deposit component of M1, ∇_1 is the first difference operator ($\nabla_1 x_t=x_t-x_{t-1}$), ∇_{12} is the seasonal difference operator ($\nabla_{12}x_t=x_t-x_{t-12}$), $\theta(B)$ is a polynomial in the backshift operator, $B(B^k a_t=a_{t-k})$, and a_t is white noise. The first estimation period for the currency component of M1 is from July 1953 through June 1965 and for the demand deposit component, from July 1950 through June 1965. These equations are then reestimated, changing the specification slightly, by "rolling up" the sample—adding a year at the end and dropping a year at the beginning. This process is continued nine times until June 1973, so that 3 years of actual data are left outside the last estimation sample.[7] Each equation is then used to generate a 36-month forecast. In general, this forecast reproduces the seasonal pattern quite well, although the level of the forecast after 36 months is often quite different from the actual level.

RESULTS

In assessing the results of the X−11 ARIMA method, we ask the following question: Do the conclusions reached by Statistics Canada with respect to their labor force data, namely increased stability of current and forecasted seasonal factors, also hold for the U.S. money stock? To answer it, we employ a measure of seasonal factor stability, used by Dagum and Kuiper, which will be described. Following Dagum [3] final seasonal factors are those from X−11 when there are available 3 additional years of data. For ordinary X−11, the current-year seasonal factors are just those for the last year in the sample, but we compute current seasonal factors in three additional ways—by augmenting the sample with 1, 2, and 3 years of forecasted data.[8] Recalling that the end point of the first sample of actual data is June 1965 and that this sample is "rolled up" nine times to reach June 1973, there are now 108 (9×12) observations on final and current seasonal factors. Note that there are four sets of current seasonal factors—one from ordinary X−11, one from X−11 ARIMA (1 year), one from X−11 ARIMA (2 years), and one from X−11 ARIMA (3 years)—all referring to the same months and years. The Dagum-Kuiper measure of stability is

$$\frac{1}{12}\sum_{m=1}^{12}\frac{1}{9}\sum_{k=1}^{9}\left|S_{m,k}^{final}-S_{m,k}\right| \tag{2}$$

where S denotes a seasonal factor—either current or forecasted, and the subscripts m and k denote month and year, respectively. The lower this statistic is, the less the current or forecasted seasonal factors are revising.[9]

In addition to current factors, policymakers are interested in forecasted seasonal factors, usually 1 year ahead. In fact, first-published data are seasonally adjusted with forecasted seasonal factors, because X−11 is not rerun until 12 new observations are obtained. These forecasted factors are generated by X−11 as

$$S_{m,k+1}=S_{m,k}+1/2\,(S_{m,k}-S_{m,k-1}),\ m=1,2,\ldots,12 \tag{3}$$

In practice, consecutive differences between seasonal factor estimates for a month are small so that these forecasts are essentially equal to the current factors. One-year-ahead forecasted seasonal factors for the X−11 ARIMA method are simply taken as end-year, next-to-end-year, or third-from-end-year factors in each of the augmented samples.

The results appear in the table. The first row presents the Dagum-Kuiper statistic computed on the X−11 method for the currency and demand deposit components of M1 for both the current and forecasted (1-year-ahead) seasonal factors. The next three rows present these same measures for the X−11 ARIMA method with 1-, 2-, and 3-year forecast horizons. Looking down the columns for current factors, we see that most of the improvement in stability comes from augmenting the sample with just 1 year of data. While the factors for currency are more stable than those for demand deposits, the absolute reduction in the measure of stability is roughly the same for each component.

The table also illustrates the difficulty of obtaining good forecasts of seasonal factors. For demand deposits, the stability measure jumps by one-third for all seasonal factor forecasts. The situation is slightly worse for currency, where seasonal factor forecasts are half again as unstable as current factors in all cases. As we read down the columns for forecasted factors, there are overall gains in stability, for both demand deposits and currency, of 20–25 percent, but they occur at different forecast horizons. These results suggest that significant improvements may be had for demand deposits by using X−11 ARIMA (3 years), but, for currency, a 1- or 2-year horizon is best.

[6]For a discussion of ARIMA model fitting, see [1, especially chs. 6–9].

[7]Coefficient estimates and summary statistics for these models are available on request.

[8]In all of these adjustments, the total sample size is restricted to 10 years. Options, in effect, are: Standard multiplicative run, with 1.5- to 2.5-sigma range for graduation of extremes, 9-term Henderson average for the trend cycle, 3×3- and 3×5-moving average smoothing of seasonal irregular ratios, and no preliminary trading-day adjustments.

[9]Compared to other criteria, such as root-mean square revision, this statistic does not give as much weight to large revisions. Since policymakers are probably more sensitive to large, rather than small, revisions in published data, we computed root-mean square revisions as well, with no qualitative change in the results of the table.

Table 1. **SEASONAL FACTOR STABILITY: RESULTS OF THE DAGUM-KUIPER STATISTICS FOR CURRENT AND FORCASTED FACTORS, BY PERCENT**

Model	Demand component		Currency component	
	Current	Forecasted	Current	Forecasted
X-11	0.174	0.231	0.091	0.131
X-11 ARIMA (1 year) . . .	.155	.208	.068	.105
X-11 ARIMA (2 years) . .	.149	.202	.076	.104
X-11 ARIMA (3 years) . .	.142	.193	.072	.114

CONCLUSIONS

This exercise has shown that, on the average, increased stability of current and forecasted seasonal factors is to be derived from using X–11 ARIMA, rather than ordinary X–11, to seasonally adjust U.S. money stock data. However, there are a number of points to consider before adopting the X–11 ARIMA procedure. First, the method is not fully automatic—an important consideration for an agency that must seasonally adjust hundreds of series. An ARIMA model for the series must be obtained, usually with a substantial investment of time for specifying, fitting, and testing. Second, the model chosen must provide good forecasts of the series. Forecast accuracy is needed so that X–11 is operating on a series that is consistent in terms of its seasonal pattern. For an analysis ex post facto, there is no problem, since forecasting performance may be checked with actual data; but, for use ex ante, there are no actual data against which to test the forecasts. One must rely on goodness of fit within sample or on a judgmental assessment of the forecasts as reasonable. Third, the gain in seasonal factor stability (i.e., the amount of revision) should be balanced against the cost of achieving it. For instance, the greatest improvement in the table for current factors for demand deposits comes from using X–11 ARIMA with 3 years of forecasts. The difference versus ordinary X–11 is 0.032 percent. This means that, for a level of demand deposits of $230 billion, the numbers adjusted by X–11 ARIMA are, on the average, $74 million closer to the final numbers than are those adjusted by ordinary X–11. In terms of levels, this average improvement is not overwhelming. However, the average is somewhat misleading, since improvements up to 0.50 percent, or $1.2 billion, occur for particular months.

In conclusion, the X–11 ARIMA approach is to be recommended for those series for which reasonable ARIMA models can be built and where the gain in stability justifies the expenditure of resources. (Quite often, such models will already have been estimated for other purposes.) For series that are highly visible economic indicators and where small changes assume political significance, any gain in stability is probably worth the effort needed to achieve it. X–11 ARIMA is also to be recommended to individual researchers who want to seasonally adjust relatively few series, while avoiding some of the asymmetries implicit in X–11.

REFERENCES

1. Box, George E. P., and Jenkins, G. M. *Time Series Analysis: Forecasting and Control.* Rev. ed. San Francisco: Holden-Day, Inc., 1976.

2. Brookings Institution. "Improving Monetary Control," by W. Poole and C. Liberman. In *Brookings Papers on Economic Activity*. No. 2. Washington, D.C.: Brookings Institution, 1972, pp. 293–335.

3. Dagum, Estella Bee. "Seasonal Factor Forecasts from ARIMA Models." *Proceedings of the International Statistical Institute, 40th Session, Warsaw, 1975.* Warsaw: International Statistical Institute, 1976, pp. 206–219.

4. Geweke, John. "The Temporal and Sectoral Aggregation of Seasonally Adjusted Time Series."
   Included in this report.

5. Kuiper, John. "A Survey and Comparative Analysis of Various Methods of Seasonal Adjustment."
   Included in this report.

6. U.S. Department of Commerce, Bureau of the Census. *The X-11 Variant of the Census Method II Seasonal Adjustment Program,* by Julius Shiskin, Allan H. Young, and J. C. Musgrave. Technical Paper 15. Washington, D.C.: Government Printing Office, 1967.

7. U.S. Federal Reserve System, Board of Governors. "Seasonal Adjustment of M1: Currently Published and Alternative Methods," by Edward R. Fry. In *Staff Economic Study No. 87.* Washington, D.C.: Government Printing Office, 1976.

8. Zellner, Arnold, and Palm, F. "Time Series Analysis and Simultaneous Equation Econometric Models." *Journal of Econometrics* 2 (May 1974): 17–54.

COMMENTS ON "A SURVEY AND COMPARATIVE ANALYSIS OF VARIOUS METHODS OF SEASONAL ADJUSTMENT" BY JOHN KUIPER

Martin M. G. Fase
De Nederlandsche Bank N.V.

On the outset, I would like to emphasize that I found Kuiper's paper quite interesting and useful. This is particularly true, because practitioners need guidelines since so many methods are available. It was for this very reason that, 5 years ago, we, at the special studies section of De Nederlandsche Bank (i.e., the Dutch central bank), did a similar study, as Kuiper has presented now. I believe Kuiper's work is in the same spirit and follows the same methodology as we applied. We compared nine different methods, including the methods Kuiper compared.[1] However, keeping in mind our own results, I cannot believe Kuiper's main finding. This seems to be that it is not possible to discriminate between different methods of seasonal adjustment.

Our analysis, based on five representative Dutch series, employing the same criteria as Kuiper used, did suggest that the X-11 and the Burman methods perform well. This was particularly so, because these two methods produce stable seasonals. Stability, in this context, means that the seasonals did not change very drastically when new data became available. To study this property, it is useful to add, to a particular series, observations over 12 months successively over a reasonable number of years. (We took 5 years.) I think Kuiper did not follow this procedure quite well. Therefore, his remark that significant differences occurred for the recent period seems, to me, unjustified.

Finally, I would like to add that, for an analysis to employ the additive or multiplicative model, the search procedure, referred to by Durbin and Kenny [1], which, incidentally, is quite common in practice, is more appropriate and simpler than Kuiper's strategy on this point.

REFERENCES

1. Durbin, J., and Kenny, P. "Seasonal Adjustment when the Seasonal Component Behaves Purely Multiplicatively nor Purely Additively."
   Included in this report.

2. Fase, Martin M. G., Koning, J., and Volgenant, A. F. "An Experimental Look at Seasonal Adjustment." *De Economist* 121 (1973): 441–480.

COMMENTS ON "ANALYSIS AND MODELING OF SEASONAL TIME SERIES" BY GEORGE E. P. BOX, STEVEN HILLMER, AND GEORGE C. TIAO

Agustin Maravall
Federal Reserve System

IDENTIFICATION

When decomposing an observed time series into unobserved components, it is well known that underidentification may be a problem. Let $[T_t]$, the observed time series, be the outcome of the process

$$\rho_p(B)T_t=\eta_q(B)d_t \tag{1}$$

and let π_t and e_t denote the two unobserved components, so that

$$T_t=\pi_t+e_t \tag{2}$$

where e_t is white noise. Thus, the ARMA process generating π_t is of the form

$$\rho_p(B)\pi_t=\alpha_{q^*}(B)c_t \tag{3}$$

where $q^*\leq \max(p, q)$. Model (3) is not uniquely determined from (1) and (2). In order to achieve identification of (3), Box, Hillmer, and Tiao (BHT) assume $q^*=\max(p, q)$, and select the model for which σ_e^2 is maximum [1]. In this note, we mention some situations where the appropriateness of these assumptions may be questioned and where the BHT procedure may lead to nonparsimonious specifications.

DECOMPOSITION

In an effort to estimate the permanent and transitory component of M1 [2], a similar decomposition was used. The AR polynomial in (1), $\rho_p(B)$, was estimated as $(1-\rho B)\Delta_{13}\Delta_{52}$.[1] In investigating the MA specification, it was found that the MA polynomial $\eta(B)$ in (1) could be explained simply by the noise-induced moving average $\rho(B)e_t$, resulting from substituting the model (3) for π_t into (2). Quite nicely, the sample ACF of $[T_t]$ could be explained by the specification

$$(1-\rho B)\,\xi_t=c_t$$

[1] BHT apply their smoothing procedure to a trend-seasonal-irregular decomposition. But the permanent transitory decomposition does not require any seasonal extraction.

where $\xi_t=\Delta_{13}\Delta_{52}\pi_t$, together with (2). Thus, a decomposition, such as (2), allowed a fairly parsimonious model to explain a fairly complicated autocorrelation structure. Also, as in this example, $p=66$, $q^*=1$, identification of the model did not require any additional assumption concerning the variance of the noise, because the lemma 1 (from [2]) could be applied.

Lemma 1

Let the zeroes of $\alpha(B)$ lie on or outside the unit circle. Then, model (3) is identified iff $p>q^*$. When $p>q^*+1$ the model is overidentified.

Thus, an empirical type of consideration may lead to a different smoothing strategy in which a more parsimonious overidentified model is obtained. Of course, one cannot expect all empirical applications to explain the observed MA polynomial by the noise-induced one. Yet, even when this simplification is not possible, an exactly identified model with $p=q+1$ can always be found. Since ARMA $(p, p-1)$ models can be rationalized as the discrete time representation of continuous processes, lemma 2 can be easily proven.

Lemma 2

Let the zeroes of $\alpha(B)$ lie on or outside the unit circle. There is one, and only one, model (3) that satisfies (1) and (2), for which there is an underlying continuous stationary stochastic process.

Thus, a model-based consideration (continuity) leads to an alternative smoothing strategy, where the assumption σ_e^2=maximum is substituted by the assumption $q^*=p-1$ (i.e., $\alpha_p=0$).[2]

[2] The CPI example, considered by BHT can certainly be represented by the model $\Delta\pi_t=(1-\alpha B)c_t$, with $\alpha=-1$ and $\sigma_e^2=0.00161$. But, it is also perfectly consistent with $\Delta\pi_t=c_t$ and $\sigma_e^2=0.00159$.

The views expressed are not necessarily related to those of the Federal Reserve System. I would like to express my gratitude to David A. Pierce.

REFERENCES

1. Box, George E. P., Hillmer, Stephen C., and Tiao, George C. "Analysis and Modeling of Seasonal Time Series."
Included in this report.

2. U.S. Federal Reserve System, Board of Governors. *On Modeling Unobserved Components With Time Series,* by Agustin Maravall. Special Studies Paper No. 93. Washington, D.C.: Government Printing Office, 1976.

COMMENTS ON "ANALYSIS AND MODELING OF SEASONAL TIME SERIES" BY GEORGE E. P. BOX, STEPHEN HILLMER, AND GEORGE C. TIAO

Christopher A. Sims
University of Minnesota

As Box, Tiao, and Hillmer point out, models of the form

$$\Psi(B^s)\Theta(B)X=\delta(B^s)\gamma(B)\epsilon \tag{1}$$

with ϵ white noise and Ψ, Θ, δ and γ low-order polynomials in their arguments, are successful in producing forecasting models of a wide range of time series, including seasonal components with period s.

This is all the more remarkable a finding since, when we examine these models in the frequency domain, this class of models implies some restrictions on the properties of X, which rule out a large part of all the stationary processes displaying Granger's "property S".[1]

In particular, the spectral density of X is given by

$$S_x(w)=\left|\frac{\delta(e^{-siw})\gamma(e^{-iw})}{\Psi(e^{-siw})\Theta(e^{-iw})}\right|^2\sigma_\epsilon^2 \tag{2}$$

Setting

$$H(w)=\left|\frac{\delta(e^{-siw})}{\Psi(e^{-siw})}\right|^2$$

$$G(w)=\left|\frac{\gamma(e^{-iw})}{\Theta(e^{-iw})}\right|^2\sigma^2$$

we have

$$S_x(w)=H(w)G(w)$$

Now, H is periodic with period $2\pi/s$, while G, if γ and Θ are of as low an order as usual, is smooth. Thus, $\log S_x$ is the sum of an exactly periodic function and a smooth function. If $\log S_x$ has sharp narrow peaks at the seasonal frequencies and if S_x has the form (2), then the peaks are all of exactly the same height and width. That the peaks be all the same height rules out the possibility that the annual seasonal pattern of the series be consistently smooth or consistently rough. That the peaks be all the same width rules out the possibility that some frequencies in the seasonal pattern might be less stable than others, e.g., that the monthly first difference of the seasonal pattern might show more tendency to change from year to year than 3-month moving averages of the seasonal pattern.

This limitation on spectral densities of the form (2) does not mean that models of the form (1) will yield poor predictions. In fact, the property S can be interpreted in the time domain as asserting simply that there is a slowly changing annual pattern to a component of the series. Hence, any forecasting scheme that estimates the average annual pattern over the past several years and projects that average pattern into the future with little change will produce forecasts which are good in an absolute sense. Differences between forecasts based on the model (1) and those based on the correct model, when the series has the property S but is not well represented in the form (1), will show up mainly in differences in the accuracy with which changes in the seasonal pattern (by construction, small to start with) are forecast. There are some applications where this could be important, especially where we have a structural model of relations between seasonal patterns in different series and wish to extract seasonal components accurately to study their interrelations.

For example,

$$X_t=1.71473\,X_{t-1}+0.9801\,X_{t-2}+\epsilon_t$$

$$=2(0.99)\cos\frac{\pi}{6}X_{t-1}+(0.99)^2X_{t-2}+\epsilon_t$$

This process consists of a strong 12-period seasonal and very little else. However, it has a peak only at $\frac{\pi}{6}$, thus, the seasonal pattern is purely a sinusoid of period 12. Obviously this seasonal is better forecast from its own recent past than from its value 12 months ago.

An interesting open question is whether there is any convenient way to expand the class of models considered by Box, Tiao, and Hillmer to avoid these restrictions.

Obviously, the frequency-domain methods of seasonal decomposition and forecasting used by Geweke[2] are not subject to the limitations discussed here. They may, in turn, of course, be limited by difficulty in handling extremely sharp seasonal peaks properly.

[1] See [1].

[2] See [2].

REFERENCES

1. Granger, Clive W. J. "Seasonality: Causation, Interpretation, and Implications."
   Included in this report.

2. Geweke, John. "The Temporal and Sectoral Aggregation of Seasonally Adjusted Time Series."
   Included in this report.

COMMENTS ON "ANALYSIS AND MODELING OF SEASONAL TIME SERIES" BY GEORGE E.P. BOX, STEVEN HILLMER, AND GEORGE C. TIAO

Kenneth F. Wallis
University of Warwick

Two of the conference papers contain an empirical example based on the same data set (one made available by the organizers), and it is of interest to compare the results of Box, Hillmer, and Tiao (BHT) with my own results for the manufacturers' shipments, inventories, and orders series, seasonally unadjusted. An initial handicap is that BHT work with the logarithms of the variables, whereas I do not, being influenced by the consideration that a full model describing the determination of these and other relevant variables would contain the identities given in the first paragraph of the section on manufacturers' shipments, inventories, and orders in my paper, "Seasonal Adjustment and Multiple Time Series Analysis," and, hence, would be easier to handle by linear methods if cast in terms of the untransformed variables. Despite this apparent lack of comparability, some similarities between the two sets of results do emerge.

First, in the univariate analyses, similarities in the models for the shipments and new orders series can be observed. Both BHT and I choose $d=D=1$, and once a strong seasonal autocorrelation has been accommodated, relatively little remains to be modeled: The comparison is closest if the MA representations of seasonality, given in footnote 1 of my paper, are considered. For the inventories series the choice of model seems less clearcut, and there is less agreement: BHT, again, choose $d=D=1$ and estimate a nonseasonal AR component of $(1-0.85L)$, whereas I am ambivalent between $d=1$ and $d=2$, with $D=0$ in each case, but, nevertheless, the results with $d=2$ yield a seasonal AR operator that contains the factor $(1-0.78L^{12})$.

In their multivariate analyses, BHT estimate what econometricians term "reduced forms," but the univariate models can still be regarded as (final equation) solutions of this system, and appropriate comparisons can be made. For this purpose, I consider the second, restricted set of estimates. (It should be noted that, in the first set of estimates, the matrix $\hat{\Phi}$ has a root outside the unit circle, suggesting that these estimates were not obtained by an exact maximum likelihood procedure.) The estimated autoregressive matrix is triangular, and the equation for I_t can be immediately separated and written as

$$(1-0.90L)\Delta\Delta_{12}\log I_t=(1-0.40L)(1-0.75L^{12})a_{3t}$$

which is close to BHT's univariate estimate. For the new orders series, the solution is

$$\begin{aligned}(1-0.97L)(1-0.90L)\Delta_{12}\log NO_t \\ =(1-0.36L)(1-0.90L)(1-0.74L^{12})a_{2t} \\ -0.345L^2(1-0.75L^{12})a_{3t}\end{aligned}$$

and, if the contribtuion of a_{3t} to the right-hand side is ignored (its variance is but $1^1/_2$ percent of that of a_{2t}), then cancellation and slight simplification yields

$$\Delta\Delta_{12}\log NO_t=(1-0.36L)(1-0.74L^{12})a_{2t}$$

which, again, corresponds well with their univariate results. In the residual covariance matrix, the most striking feature is the strong correlation between the shipments and new orders errors—our two estimates virtually agree on a coefficient of 0.7, though it must be admitted that my estimate is based on a slightly different multiple time-series representation. The small correlation of the inventories residuals with the shipments and orders residuals is noted by BHT and is also present in my own results: This, together with the decomposability of their reduced form model leads BHT to suggest that the inventories series behaves independently of the other two series. However, an economist's structural model would no doubt include the direct relationship between shipments and inventories, referred to previously, and, while it is true that the failure of my final form representation to pass the test of a common autoregression could be due to decomposability, inspection of the coefficients does not support the view that the inventories series is causally prior.

COMMENTS ON THE NBER-CENSUS CONFERENCE ON SEASONAL ANALYSIS OF ECONOMIC TIME SERIES

Allan H. Young
Bureau of Economic Analysis

I would like to comment on several points concerning the X–11 method of seasonal adjustment.

The statement concerning trading-day variation in "An Overview of the Objectives and Framework of Seasonal Adjustment," by Shirley Kallek, views the subject too narrowly. The statement implies that reasonable empirical daily weights which express the percent of the week's activity that occurs on each day of the week are preferred to fitted weights, computed by the trading-day adjustment option in X–11. The research that provided the basis for the trading-day adjustment option, however, indicated that fitted weights perform better within and beyond the sample period than do reasonable empirical weights. It should be noted that, among other things, the research utilized daily retail sales in obtaining empirical weights—a better basis for such weights than would usually be available. The reason for the superior performance of fitted weights is that a substantial proportion of economic activity occurs on the basis of monthly plans and schedules that are drawn up with little or no attention to the number of trading days in calendar months and/or is recorded and reported on a basis that takes little account of the number of trading days in calendar months. Therefore, allowance must be made for the relation of each day of the week to the monthly volume of activity, rather than for the relation of the daily activity to the weekly volume of activity. The relation of each day of the week to the monthly volume can vary with the calendar composition of the month, while the actual daily rates that are proportions of the weekly volume remain fixed. Thus, Saturday and/or Sunday can be assigned substantial weight, even when simple observation indicates the activity is shut down. (See [3] for a further discussion of this point.)

In his discussion of the "Overview," Lawrence Klein was concerned about the introduction of the Slutsky-Yule effect in the seasonally adjusted series because of iteration in X–11. Iteration is used in X–11 to improve the identification of extreme values and the measurement of the trend cycle. There are three rounds of two iterations, each, as follows:

1. A, B
2. A, B
3. A, B

Iteration A is based on the 12-month moving average; iteration B is based on the Henderson curve. The purpose of rounds 1 and 2 is to find extreme values. The only effect carried over from round 2 to round 3 is the modification of extremes in the original series. Durbin [2] showed that iteration B has no effect on a stable seasonal factor, obtained as an average over all years, except for an end-point correction. Young [4] showed that iteration B has only a small effect on the moving seasonal factor in X–11. Thus, it does not seem that the Slutsky-Yule effect plays much role in the seasonal factors, although I am not aware that anyone has fully examined possible effects arising from the treatment of extreme values.

It is important that the record be straight concerning a point in the "Overview," which was commented on by Klein. The "Overview" suggests that statistical agencies, such as the Census Bureau, apply the seasonal adjustment procedure in such a way that the impact of strikes on the data is removed. That is not the case. The statement in the "Overview" pertains to the X–11 strike adjustment option that is available for improving the estimate of the trend cycle in strike periods. With either that option or the standard option, the impact of the strike remains in the seasonally adjusted data.

There are several statements by the authors and discussants that X–11 has been shown to be a fairly good approximation to an ARIMA model for some economic series, but not for others. These statements are based on work most recently reported in an article by Cleveland and Tiao [1]. In that article, X–11 is described as performing creditably on an airline passenger series and is compared with an ARIMA model that was fit to the series. However, X–11 performed poorly on a telephone disconnection series for which a different ARIMA model is appropriate. The poor performance is based on a misspecification of X–11 for the telephone data. The telephone disconnection series contains substantial trading-day variation. This is revealed by the autocorrelation structure of the irregulars, as shown in chart F of the article by Cleveland and Tiao. The autocorrelation structure corresponds closely with that of trading-day factors, based on equal activity for Monday–Friday, with near zero activity on Saturday and Sunday. If the trading-day option had been used, the autocorrelation structure of the irregulars would not have shown this pattern, and X–11 probably would have performed about as well on the telephone data as on the airline data.

Julius Shiskin indicated, in the keynote address, that several variants of method II have served as the official method of seasonal adjustment. It may be helpful to describe how these variants differed with respect to the calculation of the current and year-ahead seasonal factors. The seasonal factors in the X–3 variant reduced the size of revisions in many series by roughly one-third, compared to the original specification of method II. This reduction resulted from a change in the procedure for extending the SI ratios for a given month to the 3 years beyond the end of the series in order that the 3×5-moving average could be computed to the end of the series. The original procedure used an average of the SI ratios for the last 2 years as an estimate of the ratios for the next 3 years. This was replaced in X–3 with an average of the SI ratios for the last 4 years. The X–3 technique was retained in X–9 and is the standard option in X–11. The X–10 variant, which was developed in cooperation with the OECD, was never the official program in the United States. It fit a different curve to each month, ranging from a very short moving average to a stable seasonal, depending on a signal-to-noise ratio. This option is available in X–11 and, if used intelligently, can lead to reduced revisions for some series. To the best of my knowledge, the only major use of this option by the U.S. statistical agencies is for the import and export series adjusted by the Census Bureau. For these series, a 3×9-moving average is used for each month in place of the 3×5-moving average.

The X–11 ARIMA variant, developed by Statistics Canada, is another approach to tailoring the procedure for obtaining current seasonal adjustment factors to the series. The reduction in revisions, reported at the conference, of about one-fifth for X–11 ARIMA is important news. One hopes that the statistical agencies will follow up and test the method on U.S. series. As with X–10, the ARIMA variant requires considerable skill and judgment. One possibility, which should not be overlooked, is that tests with X–11 ARIMA may indicate that much of the possible improvement could be captured with a very limited number of suboptimal modifications of the X–11 weights. If so, this would facilitate the uniform application of the method on a large scale.

REFERENCES

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

2. Durbin, John. "Trend Elimination for the Purpose of Estimating Seasonal and Periodic Components of Time Series." In *Proceedings of the Symposium on Time Series Analysis*. Edited by M. Rosenblatt. New York: John Wiley and Sons, Inc., 1963, pp. 3–13.

3. U.S. Department of Commerce, Bureau of the Census. *Estimating Trading-Day Variation in Monthly Economic Time Series*, by Allan H. Young. Technical Paper 12. Washington, D.C.: Government Printing Office, 1965.

4. Young, Allan H. "Linear Approximations to the Census and BLS Seasonal Adjustment Methods." *Journal of the American Statistical Association* 63 (June 1968): 445–471.