AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

February 1984

Preliminary and confidential

CALCULATING THE VARIANCE OF SEASONALLY

ADJUSTED SERIES

Peter Burridge and Kenneth F. Wallis

NUMBER 251

Department of Economics
University of Warwick
Coventry CV4 7AL, England

This paper considers the use of the Kalman filter to perform the seasonal adjustment and to calculate the variance of the signal extraction error in model-based seasonal adjustment procedures. The steady-state filter covariance is seen to provide a convenient basis for obtaining the variances not only of the current adjustment but also of subsequent revisions. The method is applied to the unobserved-components model we have recently proposed as a justification of the X-11 method, and to a real economic time series.

Key words: time series, seasonal adjustment, signal extraction, Kalman filter, unobserved-components model, standard error.

The authors are grateful to William Bell for supplying data. This research was supported by a grant from the Economic and Social Research Council.

1. INTRODUCTION

Public discussion of seasonally adjusted time series usually concentrates on the current data, that is, the seasonally adjusted value of the current month's unemployment or money supply, for example. That such figures will be revised in subsequent months, as new data offer a clearer picture of short-run movements in the series, is increasingly recognised. One might then wish to attach a standard error to the preliminary data as an indication of the likely magnitude of subsequent revisions. Moreover, if the final adjusted value that eventually emerges is regarded as only an estimate of a "true" deseasonalized series, then again, an indication of the likely error is called for. It is the purpose of this paper to propose a framework for the calculation of such quantities, in the context of an approach to seasonal adjustment that is gaining increasing support. While research that would have as an "important byproduct ... estimates of the random variability of seasonally adjusted series" was recommended over twenty years ago by the President's Committee to Appraise Employment and Unemployment Statistics (1962, p.19), the problem still appears to be open.

Traditional seasonal adjustment procedures, of which the best known is the Census Bureau's X-11 method (Shiskin et al., 1967), are often criticised for their ad hoc nature, despite their apparent success in satisfying the demands of users of statistics. More recently, "model-based" methods have been developed, resting on the application of signal extraction theory to stochastic models for the unobserved trend-cycle, seasonal and irregular components of the observed series (Box, Hillmer and Tiao, 1978; Burman, 1980; Hillmer and Tiao, 1982). The present paper utilises a Kalman filter formulation of these methods, which is equivalent to the classical

Weiner-Kolmogorov theory for stationary series, but offers a more appropriate extension to the nonstationary case. We show that in general the Kalman filter offers a convenient approach to the calculation of the variance of the seasonal estimate.

Many of the features of the traditional methods are essentially linear filtering operations, and Wallis (1982) presents the linear filters implicit in the X-11 procedures. Since the current seasonally adjusted value and the adjusted value for the current period that one will calculate in twelve months' time are each linear filters of the original data, so is their difference, the revision, hence given the covariance structure of the original data it is a simple matter to calculate the covariance structure of these revisions. In practice this calculation is not performed, since it rests on an assumption of stationarity, while most economic time series are integrated or difference-stationary series. We argue that for such series the Kalman filter provides a more appropriate foundation, with its explicit treatment of initial conditions and abandonment of the unrealistic notion of an infinite sample.

In Section 2 we describe the model-based or signal extraction approach to seasonal adjustment, and present the Kalman filter set-up for both current adjustment and subsequent revision. This is seen to yield naturally the standard error of the adjusted series. We also present two models that are later used for illustrative purposes, one from our earlier work (Burridge and Wallis, 1983b) and the other from Bell and Hillmer (1983, §7.4). In Section 3 we discuss various questions that arise in practical implementation of the Kalman filter and present numerical results for the two models. For each model the variance of the current adjusted figure is approximately twice that of the eventual final adjusted figure. Section 4 contains concluding comments.

2. SEASONAL ADJUSTMENT AND KALMAN FILTERING

2.1 Seasonal adjustment as signal extraction

We assume that the observed time series, Y_t , is the sum of several unobserved components, not all of which may be of interest. Usually it is convenient to work with three, namely the trend-cycle, seasonal and irregular components, denoted C_t , S_t and I_t respectively, thus

$$(2.1) \quad Y_t = C_t + S_t + I_t.$$

The seasonal adjustment problem is to obtain an estimate, $\hat{S}_t$, of the seasonal component and subtract it from the original series, yielding the seasonally adjusted series. For an adjustment performed using information available at time $t+k$, we write the adjusted value as

$$(2.2) \quad y_{t,t+k}^a = Y_t - \hat{S}_{t,t+k}.$$

The current or preliminary seasonally adjusted figure is $y_{t,t}^a$, and as time goes by and more observations become available this is modified by a sequence of revisions

$$r_t^{(k,k+1)} = y_{t,t+k+1}^a - y_{t,t+k}^a = \hat{S}_{t,t+k} - \hat{S}_{t,t+k+1}, \quad k=0,1,\dots$$

In the signal extraction literature the three components in (2.1) are treated as uncorrelated random processes, and the signal extraction problem is to estimate S_t , say, from observations on Y . The linear least squares (l.l.s.) solution to this problem is to construct a linear filter

$$(2.3) \quad \hat{S}_{t,t+k} = f_k(L) Y_t \equiv \sum_i f_{k,i} Y_{t-i}$$

so as to minimize the mean square error $E(S_t - \hat{S}_{t,t+k})^2$. The expression $f_k(L)$ is a polynomial in the lag operator L , including negative powers up to L^{-k} . The classical theory assumes that the autocovariances (or equivalently the spectra) of the components and hence of the observed variable are known. Both in theoretical work and in practical implementation this requirement has been met by postulating linear models for the components and expressing the autocovariances as functions of those models' parameters. In the present approach we work directly with the linear models for the components, which are assumed known and of the form

$$(2.4) \quad \phi_s(L)S_t = \theta_s(L)w_{1t}, \quad \phi_c(L)C_t = \theta_c(L)w_{2t}, \quad I_t = v_t$$

where w_{1t} , w_{2t} and v_t are uncorrelated normally distributed white noise series. The l.i.s. results we employ cover the case of correlated noises with finite second moments, but we retain the more restrictive assumptions for ease of presentation.

2.2 The state-space form and Kalman filter

We analyse the model (2.4) in state-space form, writing the state transition equation and the measurement equation as

$$(2.5a) \quad x_{t+1} = Fx_t + Gw_{t+1}$$

$$(2.5b) \quad y_t = H^T x_t + v_t.$$

In general x_t and y_t denote the state vector and the output vector respectively,

and w_t and v_t are independent serially uncorrelated normal random variables with zero means and covariance matrices Q and R . In the present case y_t is a scalar, and denoting the degrees of the lag polynomials in (2.4) by m , n , p and q respectively we obtain the form (2.5) through the following definitions and equivalences:

$$\begin{aligned}
 x_t &= (x_{1t}^T, x_{2t}^T)^T \\
 x_{1t} &= (S_t, S_{t-1}, \dots, S_{t-m+1}, w_{1,t}, w_{1,t-1}, \dots, w_{1,t-n+1})^T \\
 x_{2t} &= (C_t, C_{t-1}, \dots, C_{t-p+1}, w_{2,t}, w_{2,t-1}, \dots, w_{2,t-q+1})^T \\
 F &= \text{block diagonal } [F_1, F_2] \\
 F_1 &= \begin{pmatrix} \phi_{s,1} & \phi_{s,2} & \dots & \phi_{s,m-1} & \phi_{s,m} & -\theta_{s,1} & -\theta_{s,2} & \dots & -\theta_{s,n} \\ 1 & 0 & \dots & 0 & 0 & & & & \\ 0 & 1 & \dots & 0 & 0 & & & 0 & \\ \vdots & & & & \vdots & & & & \\ 0 & 0 & \dots & 1 & 0 & & & & \\ \hline & & & & & 0 & 0 & \dots & 0 \\ & & & & & 1 & 0 & \dots & 0 \\ & & & & 0 & 0 & 1 & \dots & 0 \\ & & & & & \vdots & & & \vdots \\ & & & & & 0 & 0 & \dots & 1 & 0 \end{pmatrix}
 \end{aligned}
 \tag{2.6}$$

F_2 is similarly defined by matching coefficients in the model for C_t to elements of x_{2t} ,

$$\begin{aligned}
 G^T &= \begin{pmatrix} 1 & 0 & \dots & 0 & 1 & 0 & \dots & 0 & 0 & 0 & \dots & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 & 0 & 0 & \dots & 0 & 1 & 0 & \dots & 0 & 1 & 0 & \dots & 0 \end{pmatrix} \\
 H^T &= \begin{pmatrix} 1 & 0 & \dots & & 0 & 1 & 0 & \dots & & & & 0 \end{pmatrix} \\
 w_t &= (w_{1t}, w_{2t})^T \quad Q = \text{diag} [\sigma_{w_1}^2, \sigma_{w_2}^2] \\
 R &= \sigma_v^2 \text{ (scalar)}.
 \end{aligned}$$

The specification is completed by the assumption that in advance of any observations, the initial state vector x_0 is known to be normally distributed with mean $\bar{x}_0$ and covariance $0 \leq P_{0,-1} < \infty$. Usually $\bar{x}_0$ and $P_{0,-1}$ will have to be chosen by the researcher, and appropriate choices are discussed in section 3.

Denoting by Ω_{t+k} the information set comprising the initial conditions together with observations $y_0, y_1, \dots, y_{t+k}$, under our assumptions the l.l.s. estimate of x_t (of which S_t is the first element) is given by the conditional expectation

$$(2.7) \quad \hat{x}_{t,t+k} = E(x_t \mid \Omega_{t+k});$$

this estimate, and its covariance

$$(2.8) \quad P_{t,t+k} = E(x_t - \hat{x}_{t,t+k})(x_t - \hat{x}_{t,t+k})^T,$$

may be obtained recursively from the Kalman filter equations set out below.

The first set give the recursions for $\hat{x}_{t,t-1}$, $\hat{x}_{t,t}$, $P_{t,t-1}$ and $P_{t,t}$, $t=0,1,2,\dots$, the second set give those for $\hat{x}_{j,j+k}$ and $P_{j,j+k}$ for j fixed and $k=0,1,\dots$, and the third set give notational definitions. First,

$$(2.9a) \quad \begin{aligned} \hat{x}_{t,t} &= \hat{x}_{t,t-1} + K_t \tilde{y}_t \\ \hat{x}_{t+1,t} &= F \hat{x}_{t,t} \\ P_{t,t} &= (I - K_t H^T) P_{t,t-1} \\ P_{t+1,t} &= \tilde{F}_t^T P_{t,t-1} \tilde{F}_t + G Q G^T \end{aligned}$$

initialized by $\hat{x}_{0,-1} = \bar{x}_0$, and $P_{0,-1}$. As time passes we obtain a sequence

of estimates, $\hat{S}_{t,t}$, $t=0,1,\dots$, for use in the preliminary adjustment $y_{t,t}^a$, from the first element of $\hat{x}_{t,t}$. To obtain the subsequent revisions, for some chosen time j , we use the next recursions to update $\hat{x}_{j,j+k}$ for $k=0,1,\dots$:

$$\begin{aligned} \hat{x}_{j,j+k} &= \hat{x}_{j,j+k-1} + \tilde{K}_{jk} \tilde{y}_{j+k} \\ (2.9b) \quad \tilde{P}_{j,j+k} &= \tilde{P}_{j,j+k-1} - \tilde{P}_{j,j+k-1}^T H \tilde{\sigma}_{j+k}^{-2} H^T \tilde{P}_{j,j+k-1} \\ \tilde{P}_{j,j+k} &= \tilde{F}_{j+k} \tilde{P}_{j,j+k-1} \end{aligned}$$

initialized at $k=0$ by $\hat{x}_{j,j}$ and $\tilde{P}_{j,j-1} = P_{j,j-1}$ from (2.9a). The intermediate quantities are defined as follows:

$$\begin{aligned} \tilde{y}_t &= y_t - H^T \hat{x}_{t,t-1} \\ K_t &= P_{t,t-1} H \tilde{\sigma}_t^{-2} \\ (2.9c) \quad \tilde{F}_t &= F[I - K_t H^T] \\ \tilde{\sigma}_t^2 &= [H^T P_{t,t-1} H + R] \\ \tilde{K}_{jk} &= \tilde{P}_{j,j+k-1}^T H \tilde{\sigma}_{j+k}^{-2} \end{aligned}$$

We note that $\tilde{\sigma}_t^2$ is the innovation variance of y_t ,

$$(2.10) \quad \tilde{\sigma}_t^2 \equiv E(y_t - H^T \hat{x}_{t,t-1})^2 \equiv E \tilde{y}_t^2,$$

while $\tilde{P}_{j,k}$ is the covariance between the error in the current one-step forecast of the state and that in the estimate of x_j :

$$(2.11) \quad \tilde{P}_{j,k} = E(x_{j+k+1} - \hat{x}_{j+k+1,j+k})(x_j - \hat{x}_{j,j+k})^T.$$

Since all the quantities required to 'smooth' the estimate $\hat{x}_{j,j}$ using (2.9b) are functions of the iterates in (2.9a), it is natural to implement these equations in tandem. That is, we run (2.9a) alone until y_j has been processed, and then augment the recursions with (2.9b). Notice that with the correspondences in (2.6), we can extract 'smoothed' estimates of $S_{t-1}, S_{t-2}, \dots, S_{t-m+1}$ from $\hat{x}_{t,t}$ as we go along, but that further smoothing of $\hat{S}_{t-m+1}$ requires implementation of (2.9b). The recursions have the great merit that the sequences of covariance matrices, which contain the measures of the standard errors of the seasonally adjusted series which we seek, are functions of the initial condition, $P_{0,-1}$, and model parameters alone. In consequence, they may be evaluated off-line, and the nature of their dependence on $P_{0,-1}$ may be examined to guide our choice of starting values.

2.3 The steady state of the Kalman filter

At first sight this recursive signal extraction algorithm may seem inconsistent with the l.l.s. results for stationary processes presented, for example, by Whittle (1963), applied to the seasonal adjustment problem by Grether and Nerlove (1970), and extended to difference-stationary processes by Cleveland and Tiao (1976) and Pierce (1979). In Whittle's treatment the filter (2.3) is obtained by inverting the covariance matrix of the data, $\underline{y} = (y_0, y_1, \dots, y_{t+k})^T$, and estimating S_t as its projection on $\underline{y}$:

$$(2.12) \quad \hat{S}_{t,t+k} = \hat{S}_{t,-1} + \text{cov}(S_t, \underline{y})^T [\text{var}(\underline{y})]^{-1} (\underline{y} - \hat{\underline{y}}_{-1})$$

where $\hat{S}_{t,-1}$ and $\hat{\underline{y}}_{-1}$ are the unconditional expectations of S_t and $\underline{y}$ (usually zero). As the length of the available record increases indefinitely ($t, k \rightarrow \infty$) so the coefficients on successive data values come to depend on

their distance from t but not on t itself, and they are given as functions of the model parameters by the formulae to be found in chapter 6 of Whittle (1963). The fact that these formulae yield well-defined stable filters when stationary models of the form (2.4) are extended to cover difference-stationary models led Whittle and subsequently Pierce (1979) to conclude that they delivered the l.l.s. estimates in these more general circumstances. The same result was obtained by Cleveland and Tiao (1976) by explicitly assuming the existence of a doubly infinite sample $\{y_t; -\infty < t < \infty\}$ and showing that the coefficients in (2.12) approach appropriate limits. We prefer not to invoke the latter assumption, which does not yield the required result for more general non-stationary cases, and which begs the question of whether such a sample is possible. However, as we have shown elsewhere (Burridge and Wallis, 1983a) the results obtained by these authors are substantially correct, and may be rigorously derived using the Kalman filter apparatus.

The key result is that if

(i) the parameterization (2.4) is parsimonious, that is, that each of the pairs of polynomials $\{\phi_s(L), \theta_s(L)\}$ and $\{\phi_c(L), \theta_c(L)\}$ have no common factors, and

(ii) the autoregressive lag polynomials $\phi_s(L)$ and $\phi_c(L)$ have no unstable common factor (that is, a common factor $(1 - \lambda L)$ with $|\lambda| \geq 1$),

then from any finite non-negative definite $P_{0,-1}$ the sequence $\{P_{t,t-1}; t=0,1,\dots\}$ tends to a limit given by the unique positive definite steady-state covariance, P , which satisfies the fourth equation of (2.9a)

$$(2.13) \quad P = \tilde{F} P \tilde{F}^T + G Q G^T.$$

The correspondence between the Kalman filter and the limiting form of (2.12)

is then obtained by taking as initial conditions $\bar{x}_0 = 0$, and $P_{0,-1} = P$. The significance of this result is two-fold: first, it shows that a particular assumption about initial conditions is implicit in the usual formulation, and secondly it indicates the precise sense in which results for stationary processes can be extended to cover the non-stationary case. We note finally that when P satisfies (2.13), all subsidiary quantities in (2.9) are also in steady state (that is, the filter (2.9a) does not depend on t , nor (2.9b) on j).

2.4 Model specification for seasonal adjustment

Our numerical illustrations use two models that are specializations of (2.4). The first, presented in Burridge and Wallis (1983b), is chosen to correspond to a decomposition of Y_t implicit in the Census Bureau's X-11 seasonal adjustment method. That is, the optimal signal extraction filter for this model closely approximates the symmetric or "historical" X-11 linear filter. Writing the model initially in terms of a seasonal component, S_t , and a combined non-seasonal component, $N_t = C_t + I_t$, we have

Model 1

$$(2.14) \quad \begin{aligned} (1 + L + L^2 + \dots + L^{11})S_t &= (1 + 0.71L^{12} + 1.00L^{24})w_{1,t} \\ (1 - L)^2 N_t &= (1 - 1.59L + 0.86L^2)\eta_t \\ \sigma_{w_1}^2 / \sigma_{\eta}^2 &= 0.017. \end{aligned}$$

The second model is that fitted to the U.S. Bureau of Labor Statistics series 'employed non-agricultural males aged 20 and over' by Bell and Hillmer (1983). It is obtained from a seasonal ARIMA model for the observed series using the 'canonical decomposition' method of Hillmer and Tiao (1982).

Again the model is initially written in two-component form:

Model 2

$$(2.15) \quad (1 + L + L^2 + \dots + L^{11})S_t = (1 + 2.093L + 2.722L^2 + 2.977L^3 + 2.869L^4 + 2.581L^5 + 2.169L^6 + 1.670L^7 + 1.206L^8 + .745L^9 + .411L^{10} - .007L^{11})w_{1,t}$$

$$(1 - .260L)(1 - L)^2N_t = (1 - .990L + .000699L^2 - .00001L^3)\eta_t$$

$$\sigma_{w_1}^2 = 82.11, \quad \sigma_{\eta}^2 = 14,412.$$

Bell and Hillmer find this model an interesting one for various purposes because it is different from the models that justify the use of X-11.

The composite models for the observed series corresponding to these unobserved-components models can be readily obtained; indeed, Bell and Hillmer's derivation of Model 2 began with the empirical composite model. The models are

$$\text{Model 1} \quad (1 - L)(1 - L^{12})Y_t = \beta(L)\epsilon_t$$

$$\text{Model 2} \quad (1 - .26L)(1 - L)(1 - L^{12})Y_t = (1 - 0.88L^{12})\epsilon_t$$

$$\sigma_{\epsilon}^2 = 16,164.$$

The moving average operator in Model 1, $\beta(L)$, is of degree 26, its complexity being the counterpart of the relative simplicity of (2.14), the formulation in which the model was first derived. By contrast, the second composite model has a simple seasonal ARIMA form with two estimated parameters, and a relatively complex component form (2.15). Its error variance differs slightly from Bell and Hillmer's original estimate. Since the error variances of the

first model can be scaled arbitrarily, the X-11 filters determining only their relative values, for subsequent comparative purposes we assign the same value, $\sigma_{\varepsilon}^2 = 16,164$, for the innovation variance of Y.

The non-seasonal components in both models have AR and MA operators, ϕ_n and θ_n , say, of equal degree, so these components may be further decomposed into the sum of a white noise irregular component, I_t , and a trend-cycle component C_t of the same form. This is accomplished by choosing a variance, R , for the white noise irregular component, and then obtaining $\theta_c(z)$ and $\sigma_{w_2}^2$ by canonical factorization of the covariance generating function of $\phi_c(L)C_t = \phi_n(L)(N_t - I_t)$, namely

$$(2.16) \quad \sigma_{w_2}^2 \theta_c(z) \theta_c(z^{-1}) = \sigma_n^2 \theta_n(z) \theta_n(z^{-1}) - R \phi_n(z) \phi_n(z^{-1}).$$

Setting $R=1.0$ scarcely alters the parameters, and we obtain, for Models 1 and 2 respectively,

$$(2.17a) \quad (1 - L)^2 C_t = (1 - 1.59L + .86L^2) w_{2t}, \quad \sigma_{w_1}^2 = 180.8, \quad \sigma_{w_2}^2 = 10,631$$

$$(2.17b) \quad (1 - .26L)(1 - L)^2 C_t = (1 - .989L + .00686L^2 + .00000804L^3) w_{2t},$$

$$\sigma_{w_1}^2 = 82.11, \quad \sigma_{w_2}^2 = 14,409.$$

From the point of view of seasonal adjustment, which requires estimation of S_t , it is irrelevant whether we work with N_t or $C_t + I_t$. These are observationally equivalent, and the further decomposition is somewhat arbitrary. However experience suggests (Schmidt, 1976) and our own experiments confirm that the recursions (2.9) can become numerically unstable when $R=0$. To avoid this problem we therefore employ the C_t specifications given in (2.17a,b) together with a unit variance white noise I_t , in place of the N_t specifications given in (2.14), (2.15) respectively.

The seasonal component in both models has $\phi_s(L) = 1 + L + \dots + L^{11}$, thus the sum of twelve consecutive terms meanders slowly around zero. While the pseudo-spectrum of Y has spikes at frequencies $k\pi/6$, $k=0,1,\dots,6$, this specification entirely associates the spike at the origin with the non-seasonal component. Then $\phi_s(L)$ and $\phi_c(L)$ have no unstable common factors, so ensuring the convergence of the filter covariance to a steady state.

3. NUMERICAL RESULTS

3.1 Passage to the steady state

The practical significance of the theoretical results cited in section 2.3 about convergence of the covariance sequence $\{P_{t,t-1}\}$ to a steady state P depends on the rate of this convergence in the present models. An associated practical question concerns the choice of initial condition, $P_{0,-1}$. If the initial condition is the null matrix, $P_{0,-1} = 0$, then we also have the result (Caines and Mayne, 1970) that under the present assumptions, the matrices $P_{t+1,t} - P_{t,t-1}$ are non-negative definite, and so diagonal elements of $P_{t,t-1}$, representing the variances of the state estimates, converge monotonically. To illustrate the speed of this convergence, we plot the second diagonal element of $P_{t+1,t}$, that is $\text{var}(\hat{S}_{t,t})$, for Models 1 and 2 in Figures 1 and 2 respectively. As the figures indicate, $\text{var}(\hat{S}_{t,t})$ reaches its steady-state value in Model 1 after five years, while convergence in Model 2 is rather slower, taking twenty years to be within 1% of the steady-state value.

In practical situations the initial condition $P_{0,-1} = 0$ might be thought to be unreasonable, since it represents perfect knowledge of the initial state vector. Since we treat this as a random variable, a more realistic assumption might be $P_{0,-1} \gg P$, and in Figures 3 and 4 we provide similar plots for the initial assumption $P_{0,-1} = 10 P$. Although for a general $P_{0,-1}$ we no longer have results about the monotonicity of the sequence, we see that in this case $\text{var}(\hat{S}_{t,t})$ again approaches its steady-state value monotonically. The principal difference between the behaviour of the filter covariances for the two models is the slower convergence in Model 2. This is essentially due to the relatively small innovation variance of its seasonal

component, $\sigma_{w_1}^2$; increasing this by a factor of 3 reduces the time taken for convergence, reported above, from twenty to eleven years. Nevertheless in both models as originally specified, the rate of convergence is such as to ensure that, for practical purposes, discussions of the variance of seasonally adjusted series can proceed in terms of the steady-state covariance matrices. Also the results indicate that a practical way of calculating P is simply to run the recursions until convergence (remembering that this calculation is independent of an observation sequence).

An alternative initial choice arises if it is assumed that the economic process described by the data originated earlier than the starting point of the time series under consideration, a finite record of which is available for analysis, for then our results about the speed of convergence to the steady state might suggest the choice $P_{0,-1} = P$, having in mind the stationarity of the steady-state signal extraction error. The calculation of P can be performed off-line, moreover the remaining filter calculations are simplified if $P_{t,t-1}$ is replaced by P from the beginning. However the use of the Kalman filter to calculate the seasonally adjusted figures raises a further initialization question, namely the choice of an initial state vector, and this is discussed in the next section. We return to the results on the variance of the seasonally adjusted series in section 3.3.

3.2 Current adjustment with the Kalman filter

Use of the recursions (2.9a) for data processing requires the choice of an initial condition $\bar{x}_0$, as well as $P_{0,-1}$. In this section we report our experience in working with the original series to which Model 2 was fitted by Bell and Hillmer, plotted in Figure 5.

First, we note that the non-stationarity of x_t makes inappropriate the choices that might be natural in a stationary environment, namely $\bar{x}_0 = 0$, the unconditional mean, together with $P_{0,-1} = P$, the steady-state covariance, or $P_{0,-1}$ equal to the unconditional covariance of x_t . Also the choice $P_{0,-1} = 0$, used for illustrative purposes in the previous section, is irrelevant for practical purposes, as already noted. In the absence of information about the levels the components had reached when observations began, the difficulty is that a poor guess at x_0 together with too small an initial variance causes the filter to attach too much weight to this misleading information, and a long time could elapse before the data come to dominate the initial choice. In the present case the "stationary" choice $\bar{x}_0 = 0$, $P_{0,-1} = P$ results in an "adjusted" series that is still obviously seasonal at the end of the record (176 observations).

It is more realistic in the present model, in which the trend-cycle component principally accounts for the drift in the series, to assign an initial state vector in which the p components relating to C_t are the correct order of magnitude. Also the uncertainty surrounding the initial choice should be reflected by setting $P_{0,-1} > P$. For illustrative purposes we consider the choice $\hat{x}_{0,-1} = (0, \dots, 0, y_2, y_1, y_0, 0, 0, 0)^T$, ensuring that the level of the trend-cycle is about right but taking no prior view on the shape of the seasonal pattern, and we reflect this uncertainty by choosing $P_{0,-1} = 10 P$, as in a previous illustration. Applying the filter appropriate to Model 2 produces the current-adjusted series, $y_{t,t}^a$, plotted in Figure 6. After an initial "learning" period the filter settles down, and the general appearance of the latter part of the series is acceptable.

To approximate the application of X-11 to this series, we also apply the filter appropriate to Model 1, with similar initial choices

$\hat{x}_{0,-1} = (0, \dots, 0, y_1, y_0, 0, 0)^T$ and $P_{0,-1} = 10 P$. The resulting current-adjusted series is plotted in Figure 7. While at first sight this might also appear acceptable as a seasonally adjusted series, closer scrutiny indicates important differences between the two filters. A subsidiary output from the recursions (2.9a) is an estimated innovation series, $\tilde{y}_t$, and in Figures 8 and 9 we plot the autocorrelation functions of the innovation series corresponding to Figures 6 and 7 respectively, having discarded the first three years of each series. We see that while the first correlogram, corresponding to the Model 2 filter, does not exceed twice the standard error, indicating that the innovation series is approximately white noise, the second indicates substantial residual autocorrelation. In effect, this provides evidence for the unsurprising conclusion that adjustment based on an inappropriate model is less than optimal; more generally we recommend this calculation as a simple check on the performance of the filter.

3.3 The variance of the adjusted series

The steady-state variance of the current-adjusted values $y_{t,t}^a$, conditional on the data, is given by $\text{var}(\hat{S}_{t,t})$ already discussed in section 3.1. Conditioning on the data restricts attention to the error variance of the seasonal component estimate, which is the appropriate measure of the "accuracy" of the adjusted data, and we emphasise that this quantity should not be interpreted as an unconditional variance of the adjusted series, even for a stationary series. Expressing these values alternatively in terms of σ_ε^2 , the innovation variance of Y , or $\sigma_{w_1}^2$, the innovation variance of the seasonal component, we have

$$\text{Model 1:} \quad \text{var}(y_{t,t}^a) = 2441.7 = .151 \sigma_\varepsilon^2 = 13.51 \sigma_{w_1}^2$$

$$\text{Model 2:} \quad \text{var}(y_{t,t}^a) = 2506.4 = .155 \sigma_\varepsilon^2 = 30.52 \sigma_{w_1}^2$$

The improvement in the seasonal adjustment of a given data point as more data are obtained can be studied by turning to the recursions (2.9b), and considering $\text{var}(\hat{S}_{t,t+j})$, $j=0,1,\dots$. For Model 1 this is plotted in Figure 10. It is seen that the variance of the "historical" or "final" adjusted figure $y_{t,t+h}^a$, say, is approximately half that of the current or preliminary adjusted figure. This reduction is achieved after an additional three years' data have been observed, after which further reductions are negligible. This corresponds to the familiar property of the X-11 procedure, which Model 1 represents in a signal extraction context: although the half-length of the X-11 filters is approximately seven years, the remote weights are very small, and it is not usually necessary to revise seasonally adjusted data that are more than three years old. The variance of the final figure, taking $h=36$, is

$$\text{Model 1} \quad \text{var}(y_{t,t+h}^a) = 1118.0 = .0692 \sigma_{\varepsilon}^2 = 6.183 \sigma_{w_1}^2.$$

The behaviour of the variance as additional observations are obtained is seen to be a small reduction as each of the first two observations arrive, followed by a pause until the next observation on the same month arrives, when there is a striking fall. The further reductions to the historical value are mainly achieved at $j=24$ and $j=36$, again indicating the particular value of additional observations on the month in question. These results allow a user to attach a standard error to an adjusted value $y_{t,t+j}^a$; it is the standard error of $\hat{S}_{t,t+j}$ viewed as an estimate of S_t , and varies with j but not t .

Attention might alternatively be focussed on the revision process, and at some intermediate point a user might wish to know by how much an adjusted figure is likely to be revised. Since in the present case, as in

Pierce (1980), the revision is independent of the error in the final estimate, the variance plotted in Figure 10 can be decomposed and we have

$$\text{var}(r_t^{(j,h)}) = \text{var}(Y_{t,t+h}^a - Y_{t,t+j}^a) = \text{var}(\hat{S}_{t,t+j}) - \text{var}(\hat{S}_{t,t+h}),$$

$$j=0,1,\dots,h.$$

Thus if the intermediate figure $Y_{t,t+j}^a$, $j < h$, is regarded as an estimate or forecast of the final adjusted value $Y_{t,t+h}^a$, then the appropriate standard error is the square root of the difference between the solid and dotted lines in Figure 10.

The corresponding results for Model 2 are plotted in Figure 11.

The historical variance is

$$\text{Model 2: } \text{var}(Y_{t,t+h}^a) = 1242.8 = .0769 \sigma_\varepsilon^2 = 15.136 \sigma_{w_1}^2,$$

which is again approximately one-half of the variance of the preliminary estimate, although a much longer period is required to achieve this than in Model 1. The signal extraction filter corresponding to Model 2 is relatively long, as noted by Bell and Hillmer, and the variance is not within 1% of the final value until eighteen years have elapsed. (This accords with our experience in running the recursions (2.9b) with the actual data: revisions to the early seasonal estimates continue throughout the record.) Nevertheless standard error calculations for this model can proceed exactly as in the preceding case.

4. CONCLUSION

This paper has presented an approach to model-based seasonal adjustment using the Kalman filter. For the computation of seasonally adjusted data, this represents an alternative to existing methods based on classical signal extraction theory that is computationally more convenient and offers a more rigorous development in the nonstationary case. The steady-state Kalman filter coincides with the classical signal extraction filter in the case in which the latter is defined, namely the stationary case.

An important byproduct of the approach, and the main contribution of the paper, is the calculation of the variance of seasonally adjusted data. For two models results are presented which show how the variance of the seasonally adjusted data is reduced from a preliminary value, through subsequent revisions as more data become available, to a final or historical value. The results also allow a user to attach a standard error of revision to a preliminary adjusted value, as a caution against placing too much reliance on the first-announced values for a given month.

The time series models underlying the method have been assumed to be known, and the contribution of the uncertainty surrounding the identification and estimation of a model for the observed series and its subsequent decomposition into models for the components remains to be investigated. In general such a contribution would be expected to be positive, thus the estimates obtained by the present method might be expected to provide a lower bound to the overall uncertainty in the seasonally adjusted data. Our first model is one that provides a signal extraction interpretation of the X-11 filter, and so the results for this model provide measures of the

variability of the X-11 method when it is applied to data for which it is optimal. We caution against interpreting these results as "X-11" results in more general settings. The application of a sub-optimal filter to a given series will produce inappropriately adjusted data, and it is not possible to say whether the calculated variance of the seasonal estimate is too great or too small in general, simply that it is wrong. The situation is analogous to that which arises in least squares regression: under classical assumptions the standard least squares calculations yield an unbiased estimator of the covariance matrix of least squares coefficients; when these assumptions are relaxed the standard calculations yield a biased estimator of the covariances of the least squares coefficients, and the bias cannot be signed without further detailed assumptions. Similarly, a case-by-case analysis is necessary to assess the actual performance of X-11 or our signal extraction approximation when employed in inappropriate situations.

Attention has been restricted to the revisions in seasonally adjusted data that arise solely from improvements in the seasonal decomposition as more data become available. In practice many economic series are revised after their first appearance, in an attempt to reduce errors arising from other sources. The integration of these two kinds of revision is a subject for further research.

REFERENCES

- BELL, W. and HILLMER, S. (1983). Issues involved with the seasonal adjustment of economic time series. Unpublished paper, U.S. Bureau of the Census.
- BOX, G.E.P., HILLMER, S.C. and TIAO, G.C. (1978). Analysis and modelling of seasonal time series. In Seasonal Analysis of Economic Time Series (A. Zellner, ed.), pp. 309-344. Washington, DC: Bureau of the Census.
- BURMAN, J.P. (1980). Seasonal adjustment by signal extraction. Journal of the Royal Statistical Society A, 143, 321-337.
- BURRIDGE, P. and WALLIS K.F. (1983a). Signal extraction in nonstationary series. Warwick Economic Research Paper No. 234.
- BURRIDGE, P. and WALLIS K.F. (1983b). Unobserved-components models for seasonal adjustment filters. Warwick Economic Research Paper No. 244.
- CAINES, P.E. and MAYNE, D.Q. (1970). On the discrete time matrix Riccati equation of optimal control. International Journal of Control, 12, 785-794.
- CLEVELAND, W.P. and TIAO, G.C. (1976). Decomposition of seasonal time series: a model for the Census X-11 program. Journal of the American Statistical Association, 71, 581-587.
- GRETHER, D.M. and NERLOVE, M. (1970). Some properties of "optimal" seasonal adjustment. Econometrica, 38, 682-703.
- HILLMER, S.C. and TIAO, G.C. (1982). An ARIMA-model-based approach to seasonal adjustment. Journal of the American Statistical Association, 77, 63-70.
- PIERCE, D.A. (1979). Signal extraction error in nonstationary time series. Annals of Statistics, 7, 1303-1320.
- PIERCE, D.A. (1980). Data revisions with moving average seasonal adjustment procedures. Journal of Econometrics, 14, 95-114.
- PRESIDENT'S COMMITTEE TO APPRAISE EMPLOYMENT AND UNEMPLOYMENT STATISTICS (R.A. Gordon, Chairman) (1962). Measuring Employment and Unemployment. Washington, DC: U.S. Government Printing Office.
- SCHMIDT, G.T. (1976). Linear and nonlinear filtering techniques. In Control and Dynamic Systems, Vol. 12 (C.T. Leondes, ed.), pp. 63-98. New York: Academic Press.
- SHISKIN, J., YOUNG, A.H. and MUSGRAVE, J.C. (1967). The X-11 variant of the Census Method II seasonal adjustment program. Technical Paper No. 15 (revised). Washington, DC: Bureau of the Census.
- WALLIS, K.F. (1982). Seasonal adjustment and revision of current data: linear filters for the X-11 method. Journal of the Royal Statistical Society, A, 145, 74-85.
- WHITTLE, P. (1963). Prediction and Regulation by Linear Least-Square Methods. London: English Universities Press.

Figure 1 $\text{Var}(\hat{S}_{t,t})$, $t=0, \dots, 100$ for Model 1 with $P_{0,-1}=0$.

Figure 2 $\text{Var}(\hat{S}_{t,t})$, $t=0, \dots, 600$ for Model 2 with $P_{0,-1}=0$.

Figure 3 $\text{Var}(\hat{S}_{t,t})$, $t=0, \dots, 175$ for Model 1 with $P_{0,-1} = 10 P$.

Figure 4 $\text{Var}(\hat{S}_{t,t})$, $t=0, \dots, 175$ for Model 2 with $P_{0,-1} = 10 P$.

Figure 5 Non-agricultural employed males aged 20 and over (U.S.) 1965(1) through 1979(8), millions.

Figure 7 $y_t^a, t=0, \dots, 175$ for Model 1 with $P_{0,-1} = 10 P$, millions

Figure 8 Autocorrelation function of $\tilde{y}_t$ for Model 2 with $P_{0,-1} = 10 P$.

Figure 9 Autocorrelation function of $\bar{y}_t$ for Model 1 with $P_{0,-1} = 10 P$.

Figure 10 $\text{Var}(\hat{S})$ vs. n for Model 1

Figure 11 $\text{Var}(\hat{S}_{t,t+j})$, $j=0, \dots, 440$ for Model 2