Performance of Seasonal Adjustment Procedures: Simulation and Empirical Results*

Dennis Fok, Philip Hans Franses, Richard Paap

Econometric Institute

Erasmus University Rotterdam

ECONOMETRIC INSTITUTE REPORT EI 2005-30

Abstract

In this chapter we use a simulation experiment to examine whether the seasonal adjustment methods Census X12-ARIMA and TRAMO/SEATS effectively remove seasonality properties from time series data, while preserving other features like the stochastic trend. As data generating processes we use a variety of processes that are actually found in practice. These processes include constant seasonality, changing seasonal patterns due to seasonal unit roots and processes with periodically varying parameters. To check for seasonality, we consider tests for seasonal unit roots, for deterministic seasonality, for seasonality in the variance, and for periodicity in the parameters. Our simulation results show that both adjustment methods are able to remove stochastic seasonal patterns from the data with the exception of changing seasonal patterns due to periodicity in the parameters. On average, the two methods perform equally well.

*This chapter is prepared for volume 1 of the Palgrave Handbook of Econometrics edited by Terence C. Mills and Kerry Patterson. Address for correspondance: Dennis Fok, Econometric Institute (H11-02), Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, The Netherlands, `e-mail`: dfok@few.eur.nl.

1 Introduction

Many quarterly observed macroeconomic time series, such as Gross Domestic Product, Private Consumption, and Industrial production often display (i) an upward trend, (ii) substantial intra-year seasonal variation, (iii) several aberrant observations and (iv) non-linearity. Macroeconomists and policymakers tend to be interested mainly in the trend and in variable-specific business-cycle variation.

Some macroeconomists tend to feel that seasonal variation is likely to blur the view on the trend and the business cycle in macroeconomic time series and therefore they want this variation to be removed from the data before any business cycle analysis. Indeed, a first glance at almost any graph of a quarterly macroeconomic time series immediately indicates that seasonal variation can be quite dominant. Rough calculations, based on regressing the growth rates of such variables on quarterly seasonal dummies, show that almost 80 to 90 per cent of the variation may be attributable to seasonality, see, for example, Miron (1996). Whether this is the best way of summarizing the data is not beyond discussion, see Hylleberg (1994) among others, but it does indicate that business-cycle variation is not immediately and visually obvious in the presence of such substantial seasonality.

There are two main criticisms on (the use of) seasonally adjusted data. The first states that seasonal variation can be important to study in its own right, and it might, for example, be informative concerning which variables lead others into or out of a recession, see Miron (1996), Hylleberg (1994), Ghysels (1994), Franses and Paap (1999) and Matas-Mir and Osborn (2003), among many others. Of course, the analysis of unadjusted data is more involved, as one needs to include specific parameters and variables in the model to capture seasonality. However, recent advances in the area of modelling seasonality show that this analysis need not be that difficult, see Ghysels and Osborn (2001) and Franses and Paap (2004).

The second criticism is that the process of seasonal adjustment may change (dynamic) correlations between macroeconomic variables. Long-run relationships and short-run dynamics in multivariate models tend to differ across models calibrated with unadjusted and with adjusted data. Only in case the seasonal adjustment filter is linear and common to all variables, there is no conflicting inference, see Sims (1974) and Wallis (1974) for early references, and Ghysels and Perron (1993) and Ericsson *et al.* (1994) for more recent evidence. Ghysels *et al.* (1996) however challenge the linearity of the Census X-11 filter.

In this chapter we abstain from a discussion on whether one should seasonally adjust data or not. As starting point we will assume that one is simply interested in seasonally adjusted data and that one needs an automatic adjustment method to remove seasonality from many time series. For this purpose, there are two popular methods for seasonal adjustment. The first method is the Census X12-ARIMA method, see Findley *et al.* (1998). This method is data-based and consists of several steps including outlier correction, trading-day correction and various sequences of moving average filters. The second method TRAMO/SEATS is more model-based, see Gómez and Maravall (1997). There, a reasonably adequate univariate time series model for the data is specified, and the seasonal adjustment filter is derived from the model properties.

To judge the quality of both adjustment methods, in this chapter we consider a simulation experiment. Instead of comparing the adjusted series with the raw series, our main focus is to analyze whether seasonal adjustment methods are able to remove the seasonal patterns in time series in an adequate way while leaving the possible stochastic trend properties of the series untouched. To stay close to reality, we use data generating processes which are likely to be found in practice. These processes display either constant seasonality or changing seasonal patterns due to seasonal unit roots and processes with periodically varying parameters. Plausible parameter values are obtained by estimating the corresponding time series models for fourteen US industrial production series. To search for seasonal patterns before and after correction, we consider tests for seasonal unit roots, for deterministic seasonality, for seasonality in the variance, and for periodicity in the parameters.

The outline of the remainder of this chapter is as follows. In Section 2 we briefly discuss the two seasonal adjustment procedures we apply in this chapter. In Section 3 we discuss several diagnostic and specification tests that we use to evaluate the quality of both seasonal adjustment filters. Section 4 discusses the data generating processes for our simulation experiment. The outcomes of our simulation study are given in Section 5. We conclude in Section 6.

2 Seasonal adjustment procedures

In this section we briefly discuss the two seasonal adjustment methods under scrutiny. We have no intention to be complete and we strongly suggest readers to consult other studies, like Hylleberg (1986), Findley *et al.* (1998), Maravall (1985, 1995) and Harvey

(1989) for more details.

The main assumption of seasonal adjustment is that a seasonally observed time series y_t, $t = 1, \ldots, T$, can be decomposed into two unobserved components, that is,

$$y_t = y_t^{ns} + y_t^s \tag{1}$$

(or $y_t = y_t^{ns} y_t^s$ in case of multiplicative seasonality) with y_t^{ns} denoting the nonseasonal component containing the trend, cycle and all kinds of other features, and y_t^s denoting the seasonal component.

When seasonality is purely deterministic, y_t^s is assumed to be a function of sine and cosine functions. When seasonality is not constant over time, one can consider certain moving average filters to characterize changing seasonality. Preferably, these filters are linear, symmetric and centered around the current observation, see Grether and Nerlove (1970). Denoting the backward shift operator as L, defined by $L^k y_t = y_{t-k}$, $k = 0, \pm 1, \pm 2, \ldots$, such a linear moving average filter is given by

$$C_m(L) = c_0 + \sum_{i=1}^{m} c_i (L^i + L^{-i}), \tag{2}$$

where $c_0, c_1, \ldots, c_m$ are the weights. A simple example is the $C_1(L)$ filter with $c_0 = 1/2$ and $c_1 = -1/4$, which equals $-1/4(L^2 - 2L + 1)L^{-1}$, where it is used that $LL^{-1} = 1$. This filter assumes two unit roots at the nonseasonal frequency because $(L^2 - 2L + 1) = (1 - L)^2$. Hence, it removes the stochastic trend (in fact, it removes two such trends). Generally, when one aims to remove stochastic trends, it holds that $C_m(1) = c_0 + 2\sum_{i=1}^{m} c_i = 0$. Notice that the commonly applied differencing filter $(1 - L)$ is not a symmetric filter.

Following the same line of thought, to remove changing seasonality in quarterly data, one may opt for a filter like

$$4C_3(L) = (1 + L + L^2 + L^3)(1 + L^{-1} + L^{-2} + L^{-3}) \tag{3}$$

This $C_3(L)$ filter has two times three seasonal unit roots, that is, two times -1 and two times $\pm$i, see Hylleberg *et al.* (1990). Writing (3) as (2), we have that $4c_0 = 4$, $4c_1 = 3$, $4c_2 = 2$ and $4c_3 = 1$. Generally, for filters that remove changing seasonality, it holds that $c_0 + 2\sum_{i=1}^{m} c_i = 1$ (which also holds for (3) after scaling). More details of the use of linear moving average filters are given in Maravall (1995) and in Grether and Nerlove (1970), where it is shown that filters like (3) have certain optimal properties.

2.1 Census X12-ARIMA

The X12-ARIMA method is one of the most popular seasonal adjustment procedures around. The key references for this approach are Shiskin and Eisenpress (1957) and Shiskin *et al.* (1967). A recent extensive documentation of this method appears in Findley *et al.* (1998). Apart from the treatment of holiday, trading-day and calendar effects, the additive version of the X12-ARIMA method concerns two main actions. The first is the sequential application of a set of linear moving average filters as in (2) to characterize the trend and seasonal fluctuations. The filters have to be selected by the practitioner, that is, one has to select the value of m, where often m equals 5, 7 or 9 for quarterly data. The second and very important action is the removal of outlying observations in several rounds of moving average filtering, and the replacement of these observations by data points that are somehow weighted. Again, this involves decisions that should be made by the practitioner and that will vary across the time series at hand. The outlier weighting part makes the overall procedure an intrinsically nonlinear method in the sense that the weights will depend on the choice of moving average filters. Indeed, Ghysels *et al.* (1996) show that after seasonally adjustment nonlinear features may appear in linear time series.

Neglecting the outlier removal part of the official Census method, it is possible to give a linear symmetric moving average approximation to an often applied sequence of moving average filters in the Census X-11 program. For quarterly time series, the weights in this $C_{28}(L)$ filter are given in Laroque (1977). An approximate version of the $C_{28}(L)$ filter is given in Ghysels and Perron (1993), and a detailed version in Franses (1996, Table 4.1). In Laroque (1977, Table 3) it is shown that the linear $C_{28}(L)$ filter approximately contains the component

$$(1 + L + L^2 + L^3)^2 = (1 + L)^2(1 - \mathrm{i}L)^2(1 + \mathrm{i}L)^2, \tag{4}$$

see also Bell and Kramer (1996). Hence, the resulting seasonal adjustment filter from the Census program approximately encompasses the $C_3(L)$ filter in (3).

In order to seasonally adjust observations at time t with the $C_{28}(L)$ filter, one needs the observations over the sample $y_{t-28}, \ldots, y_{t+28}$. Since such observations are not available at the beginning and at the end of a sample, one needs to obtain backcasts and forecasts of y_t. One approach is now to estimate seasonal ARIMA models for y_t, $t = 1, 2, \ldots, T$, and to generate $\hat{y}_{-27}, \ldots, \hat{y}_0, \hat{y}_{T+1}, \ldots, \hat{y}_{T+28}$, see Dagum (1980) for details. The ARIMA estimation routine is known as regARIMA. This routine also allows for additional regressors to capture, for example, calendar effects and allows for outlier correction.

In our simulation study below, we use the X12-ARIMA procedure with all the default settings. In accordance with the data generating processes we consider, we impose an additive seasonal pattern (so, no natural logs are taken). As regressor variables we only use an intercept. The ARIMA model selection is done using the automatic procedure. We let X12 select the best ARIMA specification out of a (default) set of options.

2.2 TRAMO/SEATS

In response to the possible ambiguities involved in the application and evaluation of the Census X-11 procedure, Hillmer and Tiao (1982) propose the so-called ARIMA-model-based approach to seasonal adjustment, see also Burman (1980) and Gómez and Maravall (1994). A lucid exposition of the model-based method is given in Maravall and Pierce (1987). The most popular seasonal adjustment method in this area is the TRAMO/SEATS (Time Series Regression with ARIMA Noise, Missing Observations, and Outliers/Signal Extraction in ARIMA Time Series) method of Gómez and Maravall (1997). The adjustment method consists of two steps. In the first step (TRAMO) a time series model is estimated. The second part (SEATS) deals with the extraction of the seasonal pattern from the selected ARIMA model.

In a very simple version, it is assumed that a time series can be decomposed as (1). The seasonal component in (1) is described by seasonal ARIMA model as proposed by Box and Jenkins (1970) and Box *et al.* (1994), for example,

$$(1 + L + L^2 + L^3)y_t^s = \psi(L)\eta_t \tag{5}$$

and the nonseasonal part by a nonseasonal ARIMA model like,

$$(1 - L)^d y_t^{ns} = \theta(L)\xi_t, \tag{6}$$

where $\psi(L)$ and $\theta(L)$ are polynomials in L. The two components are imposed to be orthogonal. This routine also allows for outlier correction and for additional regressors to capture, for example, calendar effects. After model selection by TRAMO, in the SEATS part the Wiener-Kolmogorov filter is used to extract the seasonal component from the series.

In our simulation study below, we apply the TRAMO and SEATS procedures with the default settings. To fit the data, we do not use the standard data transformation to logs. With the default settings the TRAMO/SEATS procedure uses an Airline model (see Section 4.4 below) to estimate the seasonal component of a series.

3 Diagnostic tests

There are several criteria that can be used to evaluate the quality of seasonally adjusted data obtained from the above procedures. An extensive discussion of several such criteria is given in Hylleberg (1986, Chapter 3) and Bell and Hillmer (1984). In this chapter we judge the quality of a seasonal adjustment procedure by applying a number of diagnostic and specification tests concerning the presence of seasonal patterns before and after correction. Each test focusses on a property that should (or should not) be present in seasonally adjusted data. We consider tests for the presence of seasonal unit roots, the presence of changing seasonal means, the presence of deterministic seasonality, the presence of correlation at the seasonal lag, the presence of periodicity in the autoregressive parameters and the presence of seasonality in the variance of the series. In this section we consider tests for quarterly data but the tests can easily be extended to monthly data.

3.1 HEGY test

The unit roots in seasonal data, which can be associated with changing seasonality, are the so-called seasonal unit roots, see Hylleberg *et al.* (1990). For quarterly data, these roots are -1, i, and $-$i. For example, data generated from the model $y_t = -y_{t-1} + \varepsilon_t$ would display seasonality. Similar observations hold for the model $y_t = -y_{t-2} + \varepsilon_t$, which can be written as $(1 + L^2)y_t = \varepsilon_t$, where the autoregressive polynomial $1 + L^2$ corresponds to the seasonal unit roots i and $-$i, as these two values solve the equation $1 + z^2 = 0$. Hence, when a model for y_t contains an autoregressive polynomial with roots -1 and/or i, $-$i, the data are said to have seasonal unit roots.

To test for the presence of seasonal unit roots, we consider the approach of Hylleberg *et al.* (1990), henceforth abbreviated by HEGY. The HEGY method amounts to a regression of $\Delta_4 y_t = y_t - y_{t-4}$ on deterministic terms like seasonal dummies and a trend and on $x_{1t} = (1 + L + L^2 + L^3)y_{t-1}$, $x_{2t} = (-1 + L - L^2 + L^3)y_{t-1}$, $x_{3t} = -(1 + L^2)y_{t-1}$, $x_{4t} = -(1 + L^2)y_{t-2}$, and on lags of $\Delta_4 y_t$, where $\Delta_i y_t = y_t - y_{t-i}$. The test regression reads

$$\Delta_4 y_t = \sum_{s=1}^{4} \beta_s D_{s,t} + \gamma t + \pi_1 x_{1t} + \pi_2 x_{2t} + \pi_3 x_{3t} + \pi_4 x_{4t} + \sum_{i=1}^{p} \phi_i \Delta_4 y_{t-i} + \varepsilon_t, \qquad (7)$$

where $D_{s,t} = 1$ if t corresponds to season s and 0 otherwise. The t-test for the significance of the parameter for x_{1t} (π_1) is denoted by t_1, the t-test for π_2 by t_2, and the joint significance test for π_3 and π_4 is denoted by F_{34}. If the π parameters are equal to 0, this corresponds to the presence of the associated root(s), which are 1, −1, and the pair i, −i, respectively. Critical values of these test statistics are given in Hylleberg *et al.* (1990, Table 1).

We argue that in a properly seasonally adjusted times series seasonal unit roots −1, i and −i should not be present. Ideally, the finding of the unit root 1 should not be altered if it is present, as this root is associated with the stochastic trend in the series. The value of p can be determined using an information criterion such as the Bayesian Information Criterion [BIC].

3.2 Canova-Hansen test

The test developed by Canova and Hansen (1995) takes as the null hypothesis that the seasonal pattern is deterministic. To explain the test, consider the process

$$y_t = \sum_{s=1}^{4} \delta_{st} D_{s,t} + \varepsilon_t \tag{8}$$

with

$$\begin{aligned} \delta_{1t} &= \mu_t + \alpha_{1t} - \alpha_{3t} & \qquad \delta_{2t} &= \mu_t - \alpha_{2t} + \alpha_{3t} \\ \delta_{3t} &= \mu_t - \alpha_{1t} - \alpha_{3t} & \qquad \delta_{4t} &= \mu_t + \alpha_{2t} + \alpha_{3t}, \end{aligned} \tag{9}$$

where the stochastic trend is defined as

$$\mu_t = \mu + \mu_{t-1} + \xi_t \tag{10}$$

with $\xi_t \sim N(0, \sigma_\xi^2)$ and the stochastic seasonal terms are given by

$$\alpha_{jt} = \beta_j + \alpha_{j,t-1} + \eta_{jt} \tag{11}$$

with $\eta_{jt} \sim N(0, \sigma_j^2)$ for $j = 1, \ldots, 3$. The process has a stochastic seasonal pattern if one or more $\sigma_j^2 > 0$. If $\sigma_j^2 = 0$ for all j, we have deterministic seasonality. The Canova-Hansen test corresponds to jointly testing for $\sigma_1^2 = \sigma_2^2 = \sigma_3^2 = 0$. The asymptotical critical values are given in Canova and Hansen (1995). For the quarterly case and a significance level of 5%, the critical value is 1.010. We will denote this test by CH in the remainder of this chapter.

The CH test also allows for testing for stationarity of the process itself, that is, testing for $\sigma_\xi^2 = 0$. However, this is not considered here as we only focus on the seasonal properties of the data. In fact, given our data generating process we apply in our simulation experiment the CH test to the first difference of the series to circumvent possible size distortions in the test for the seasonal part, see, for example, Taylor (2003) and Busetti and Taylor (2003).

The null hypothesis in the CH test is rejected in case seasonality of a series is not constant. After seasonal adjustment the CH test therefore should not reject the null hypothesis. Note that having no seasonal pattern at all also implies constant seasonality.

3.3 Test for equal seasonal dummies

A basic test for the presence of seasonality in a time series is to regress the time series on four seasonal dummies. If there is no seasonality in the series, the four coefficients associated with these dummies should be equal. This property can easily be tested with a standard F-test. The test regression equals

$$\Delta_1 y_t = \sum_{s=1}^{4} \beta_s D_{s,t} + \varepsilon_t, \tag{12}$$

where $\Delta_1 y_t = y_t - y_{t-1}$ and $D_{s,t} = 1$ if t corresponds to season s and 0 otherwise. If seasonal adjustment is properly done, and hence there is no seasonality, the F-test for $\beta_1 = \beta_2 = \beta_3 = \beta_4$ should not reject the null hypothesis.

3.4 Test for correlation at the seasonal lag

Seasonal time series typically display autocorrelation at seasonal lags. To test for significant autocorrelation at the seasonal lag we consider the following regression model

$$\Delta_1 y_t = \mu + \phi_1 \Delta_1 y_{t-1} + \phi_2 \Delta_1 y_{t-2} + \phi_3 \Delta_1 y_{t-3} + \phi_4 \Delta_1 y_{t-4} + \varepsilon_t \tag{13}$$

and we test for $\phi_4 = 0$ using a t-test. Insignificant values of the t-test mark the absence of correlation at the seasonal lag. One has to be a little cautious with this approach. Autocorrelation at the seasonal lag does not have to imply seasonality as the true lag-order of the series may be 4 or higher. Note that we do not include seasonal dummies in the test regression as we want to focus on testing for correlation at the seasonal lag. The previous test in (12) should already indicate the presence of unequal seasonal means.

3.5 Test for periodicity in AR parameters

Another property which may indicate the presence of seasonality in time series concerns different autoregressive parameters across the seasons, see Franses and Paap (2004). To investigate this periodicity we consider the PAR(p) model

$$\Delta_1 y_t = \mu + \sum_{i=1}^{p} \phi_{is} D_{s,t} \Delta_1 y_{t-i} + \varepsilon_t. \tag{14}$$

Absence of periodicity corresponds with the restriction $\phi_{i1} = \phi_{i2} = \phi_{i3} = \phi_{i4}$ for $i = 1, \ldots, p$. This can be tested with a standard F-test. If the F-statistic is not significant, there is no statistical evidence for periodicity in the autoregressive parameters. The value of p can be determined using an information criterion such as the BIC. Again, the test regression does not contain seasonal dummies as we focus on periodicity in the autoregressive structure. Given the linear and not seasonal-specific structure of the seasonal adjustment procedures we expect that both procedures are not fully capable of removing periodicity from the parameters.

3.6 Test for seasonality in the variance

The previous tests mainly consider the presence of seasonality in the mean of the series. To test for the presence of seasonality in the variance of the series we consider the estimated residuals $\hat{\varepsilon}_t$ of an AR(p) model for $\Delta_1 y_t$

$$\Delta_1 y_t = \mu + \sum_{i=1}^{p} \phi_i \Delta_1 y_{t-i} + \varepsilon_t. \tag{15}$$

The LM-test for seasonality in the residuals amounts to testing for $\beta_1 = \beta_2 = \beta_3 = \beta_4$ in the auxiliary regression

$$\hat{\varepsilon}_t^2 = \sum_{s=1}^{4} \beta_s D_{s,t} + \sum_{i=1}^{p} \rho_i \Delta_1 y_{t-i} + \eta_t \tag{16}$$

using a standard F-test, where $\hat{\varepsilon}_t$ denotes the estimated residuals of (15), see Franses and Paap (2004, p. 40). A significant value of of the F-statistic indicates the presence of seasonality in the variance. The value of p can again be determined using BIC. In the ideal case, seasonal adjustment methods should remove any seasonality in the variance.

The abovementioned diagnostic tests will now be used to analyze the quality of the two seasonal adjustment methods in a simulation experiment. In the next section we discuss the data generating processes we will use in this experiment.

Table 1: US Industrial production series

Series	industry code	sample
Total index	1	1919.1–2000.4
Final products	2	1939.1–2000.4
Total products	30	1939.1–2000.4
Consumer goods	1000	1939.1–2000.4
Automotive products	1001	1947.1–2000.4
Auto parts & allied goods	1002	1947.1–2000.4
Other durable goods	1006	1947.1–2000.4
Clothing	1012	1947.1–2000.4
Chemical products	1016	1954.1–2000.4
Paper products	1017	1954.1–2000.4
Energy products	1018	1954.1–2000.4
Fuels	1019	1954.1–2000.4
Durable consumer goods	1020	1947.1–2000.4
Foods & tobacco	1022	1947.1–2000.4

4 Data generating processes

To analyze whether seasonal adjustment methods are capable of removing seasonal properties from seasonal time series, we perform a simulation experiment. In this section we discuss the five data generating processes we consider. The DGPs are chosen such that they mimic series which are frequently encountered in reality. Plausible values of parameters are obtained by applying the model corresponding to each DGP to the logarithm of fourteen quarterly observed US Industrial production series for different industry codes. These are given in Table 1. The series can be downloaded from `http://www.economagic.com`. A thorough analysis of the seasonal properties of these series can be found in Franses and Paap (2004). All artificial series are generated with standard normal innovations, that is, $\varepsilon_t \sim N(0, 1)$.

4.1 DGP1: Constant annual growth

The first data generating process assumes a constant unconditional yearly growth rate for each quarter. It is an autoregressive [AR] process of order 5 for the annual growth rate, that is,

$$(\Delta_4 y_t - \mu) = \sum_{i=1}^{5} \phi_i (\Delta_4 y_{t-i} - \mu) + \sigma \varepsilon_t. \tag{17}$$

Table 2 displays the parameter settings we use for this DGP, which are based on the true parameter estimates of the fourteen US industrial production series.

This model assumes the presence of three seasonal unit roots, that is, -1 and $\pm$i and hence it allows for a changing seasonal pattern. We expect that both X12-ARIMA and TRAMO/SEATS are capable of removing the changing seasonal pattern from these series.

4.2 DGP2: Deterministic seasonality

The second process we consider is a seasonal autoregressive moving average [ARMA] process for the first difference of the series with different but constant unconditional growth rates per quarter. The exact specification is

$$(1 - \phi_4 L^4)(\Delta_1 y_t - \mu - \delta_1 D_{1,t} - \delta_2 D_{2,t} - \delta_3 D_{3,t}) = (1 + \psi_1 L + \psi_4 L^4)\sigma \varepsilon_t. \tag{18}$$

Table 3 displays the values of the parameters which are used to generate the data. The values corresponds to the parameter estimates of (18) for the fourteen US industrial production series.

This particular specification allows for a nonzero expected growth over an entire year. Nonzero values of the δ_s parameters imply different growth rates in each quarter. The seasonal pattern in these series is however constant over time. Also for this DGP, we expect all seasonal adjustment procedures to perform well although the methods impose seasonal unit roots which should appear as moving average seasonal unit roots in the adjusted series.

4.3 DGP3: Stochastic seasonality

For some economic series the seasonal pattern changes over time. The third DGP in our simulation experiment mimics this feature through stochastic seasonality. We consider a structural time series process with a random walk with drift and trigonometric seasonality,

Table 2: Parameter values for DGP1

parameter	industry code													
	1	2	30	1000	1001	1002	1006	1012	1016	1017	1018	1019	1020	1022
100μ	4.06	3.51	3.45	3.31	3.57	3.95	4.23	0.63	4.94	2.49	3.04	1.69	3.98	2.36
ϕ_1	1.22	1.43	1.46	1.08	0.75	0.93	1.26	1.14	0.78	0.98	0.45	0.55	1.04	0.63
ϕ_2	-0.47	-0.64	-0.73	-0.32	-0.04	-0.02	-0.57	-0.39	-0.08	-0.21	0.03	0.21	-0.29	0.08
ϕ_3	0.38	0.19	0.27	0.18	0.04	-0.05	0.26	0.11	0.15	0.09	0.07	-0.01	0.11	0.12
ϕ_4	-0.77	-0.38	-0.42	-0.59	-0.50	-0.48	-0.57	-0.51	-0.42	-0.36	-0.18	-0.31	-0.49	-0.40
ϕ_5	0.44	0.25	0.26	0.32	0.28	0.32	0.29	0.39	0.38	0.23	0.25	0.24	0.24	0.17
$\ln(\sigma)$	-3.05	-3.57	-3.66	-3.63	-2.23	-3.04	-3.19	-3.58	-3.65	-3.87	-3.29	-3.60	-2.84	-4.22

The DGP is given in (17).

Table 3: Parameter values for DGP2

	industry code													
parameter	1	2	30	1000	1001	1002	1006	1012	1016	1017	1018	1019	1020	1022
100μ	0.78	0.80	0.37	-0.96	14.27	0.01	5.43	-3.87	-6.57	-3.51	2.38	-0.20	8.10	-5.01
$100\delta_1$	-0.41	-0.16	-0.36	1.54	-10.10	0.94	-7.56	0.99	3.39	2.35	12.00	-3.80	-7.24	2.09
$100\delta_2$	0.97	0.10	1.31	1.56	-10.55	1.67	-4.37	9.57	12.58	5.18	-24.91	3.24	-5.23	9.39
$100\delta_3$	-0.12	0.90	1.48	3.34	-32.89	1.82	-5.81	3.36	14.31	8.59	5.39	2.51	-16.10	10.95
ϕ_4	-0.53	0.03	0.06	0.92	0.81	0.75	0.16	0.95	0.86	0.89	0.88	0.89	0.86	0.85
ψ_1	0.40	0.45	0.47	-0.03	-0.26	0.02	0.58	0.16	-0.28	0.01	-0.75	-0.27	-0.08	-0.36
ψ_4	0.58	0.21	0.23	-0.83	-0.71	-0.79	0.08	-0.83	-0.46	-0.36	-0.14	-0.72	-0.80	-0.47
$\ln(\sigma)$	-1.39	-1.58	-1.61	-1.59	-1.01	-1.37	-1.37	-1.59	-1.59	-1.67	-1.44	-1.60	-1.25	-1.84

The DGP is given in (18).

that is,

$$
\begin{aligned}
y_t &= \mu_t + \delta_{2t} + \delta_{3t} + \sigma\varepsilon_t \\
\mu_t &= \mu + \mu_{t-1} + \sigma_\mu \eta_t \\
\delta_{1t} &= \delta_{2,t-1} + \sigma_1 \xi_{1t} \\
\delta_{2t} &= \delta_{1,t-1} + \sigma_1 \xi_{2t} \\
\delta_{3t} &= -\delta_{3,t-1} + \sigma_3 \xi_{3t},
\end{aligned} \tag{19}
$$

where $\eta_t, \xi_{1t}, \xi_{2t}, \xi_{3t} \sim NID(0,1)$, see, for example, Harvey (1989, p. 41) for a discussion. This DGP is close to the process in DGP1, although now seasonality does not change as quickly. Table 4 displays the parameter values used to generate the series based on the fourteen industrial production series.

DGP3 does not assume seasonal unit roots in the series, but it does assume random walk like patterns in the parameters. When the variances of the error terms are large, it is quite likely that the data from this process can be approximated by a model with seasonal unit roots. When the variances are zero, this process collapses to DGP2. When the variances are very small, the data from this process can display slowly changing seasonal patterns.

4.4 DGP4: Airline model

The fourth data generating process in our simulation experiment is exactly the model underlying the TRAMO/SEATS method, that is, the airline model. This process is specified as

$$\Delta_1 \Delta_4 y_t = (1 + \psi_1 L)(1 + \psi_4 L^4)\sigma\varepsilon_t. \tag{20}$$

DGP4 assumes 3 three seasonal unit roots. Bell (1987) shows that when the MA(4) parameter gets closer to -1, the model generates data that are close to those of DGP2. In principle, the airline model can describe data that show varying patterns of changing seasonality over time.

It is to be expected that TRAMO/SEATS will yield the best seasonally adjusted series for this DGP. The parameter values based on parameter estimates for the fourteen industrial production series are given in Table 5.

Table 4: Parameter values for DGP3

	industry code													
parameter	1	2	30	1000	1001	1002	1006	1012	1016	1017	1018	1019	1020	1022
$\ln(\sigma)$	-20.00	-20.00	-20.00	-20.00	-7.33	-20.00	-20.00	-20.00	-20.00	-20.00	-8.57	-8.85	-20.00	-10.32
100μ	0.92	1.04	1.01	0.88	0.86	0.87	1.03	0.17	1.33	0.65	0.84	0.48	0.98	0.56
$\ln(\sigma_\mu)$	-1.98	-7.08	-7.22	-7.46	-5.10	-6.30	-6.32	-7.41	-7.86	-8.38	-8.05	-8.27	-5.92	-9.43
$\ln(\sigma_1)$	-4.71	-15.60	-15.19	-13.41	-9.28	-12.47	-12.73	-12.12	-10.78	-11.26	-9.54	-12.94	-11.84	-12.06
$\ln(\sigma_2)$	-4.71	-15.60	-15.19	-14.46	-10.59	-14.38	-12.73	-13.77	-13.16	-13.33	-12.07	-13.22	-12.68	-15.72

The DGP is given in (18). In some cases σ converges to a very small value in which case we put σ equal to $\exp(-20)$ and for some series we need to impose that $\sigma_1 = \sigma_2$.

Table 5: Parameter values for DGP4

	industry code													
parameter	1	2	30	1000	1001	1002	1006	1012	1016	1017	1018	1019	1020	1022
ψ_1	0.36	0.45	0.47	0.18	-0.15	0.09	0.50	0.29	-0.33	0.10	-0.80	-0.38	0.17	-0.41
ψ_4	-0.98	-0.84	-0.83	-0.82	-0.73	-0.93	-0.82	-0.69	-0.52	-0.41	-0.42	-0.78	-0.82	-0.61
$\ln(\sigma)$	-3.14	-3.63	-3.702	-3.65	-2.25	-3.08	-3.19	-3.60	-3.63	-3.84	-3.34	-3.68	-2.86	-4.21

The DGP is given in (20).

4.5 DGP5: Periodic autoregressive process

The final DGP we consider is a periodic autoregressive process of order 2, that is,

$$y_t = \sum_{s=1}^{4}(\delta_s D_{s,t} + \tau_s D_{s,t} T_t + \sum_{i=1}^{2} \phi_{is} D_{s,t} y_{t-i}) + \sigma\varepsilon_t, \tag{21}$$

where $T_t = [(t-1)/4] + 1$ where $[\cdot]$ is the integer function, see Franses and Paap (2004) for a survey on periodic models. The values of the autoregressive parameters are different across the seasons. In fact, test results in Franses and Paap (2004, Table 3.2) show that this feature cannot be rejected for any of the fourteen industrial production series. The values of the parameters are displayed in Table 6 and are based on parameter estimates of a periodic autoregression of order 2 for the fourteen series.

This DGP displays a slowly changing seasonal pattern. As the seasonal adjustment filters do not use periodic filters, we expect that the seasonal adjustment methods are not able to fully remove this seasonal pattern from the series.

5 Simulation results

In this section we discuss the results of our simulation experiment. The setup of our experiment is as follows. For each DGP in Section 4 we simulate 1000 time series with the fourteen different parameter settings and hence we obtain 5 times 14000 seasonal time series with different properties. Each time series contains 60 years of quarterly data. The first ten years are discarded to initialize the data generating process. The analysis below is based on the remaining 50 years. All series are seasonally adjusted using X12-ARIMA and TRAMO/SEATS using default options. We apply the diagnostic tests discussed in Section 3 to the raw series and both seasonally adjusted series.

The results of our simulation experiment are presented in Table 7[1]. The table displays the rejection frequencies of the diagnostic tests for the raw data and the seasonally adjusted data using X12-ARIMA and TRAMO/SEATS for the five DGPs. All tests are performed with a 5% level of significance. For the ease of interpretation of Table 7, Table 8 displays the desired outcomes of the diagnostic tests after seasonal adjustment.

[1] All simulations were done in Ox 3.4 (Doornik, 1999). The actual seasonal adjustment was done through calls to the original procedures of CENSUS X12-ARIMA and TRAMO/SEATS which are shipped with EViews 5.

Table 6: Parameter values for DGP5

	industry code													
parameter	1	2	30	1000	1001	1002	1006	1012	1016	1017	1018	1019	1020	1022
$100\delta_1$	3.16	7.73	6.53	1.04	66.82	15.68	-19.61	17.21	6.71	-20.44	4.15	34.72	15.96	-59.58
$100\delta_2$	13.37	0.08	1.79	-12.56	44.49	28.59	3.34	-71.13	-1.57	-6.04	11.72	1.40	15.17	10.41
$100\delta_3$	15.98	28.45	26.27	25.95	38.65	16.34	35.41	75.20	-6.20	-8.07	-17.74	-9.25	39.78	80.57
$100\delta_4$	0.42	11.90	11.56	18.07	39.76	-12.30	42.41	14.60	22.73	65.06	27.14	27.71	31.69	5.65
$100\tau_1$	0.06	0.14	0.12	0.13	1.19	0.53	0.16	-0.30	0.43	-0.09	0.39	-0.09	0.65	-0.31
$100\tau_2$	0.21	0.00	0.01	-0.19	0.69	0.51	0.30	0.17	-0.23	-0.11	-0.37	0.18	0.44	-0.03
$100\tau_3$	0.35	0.47	0.46	0.34	0.95	0.20	0.57	0.07	0.04	-0.14	-0.03	-0.15	0.79	0.53
$100\tau_4$	0.00	0.08	0.03	0.11	0.54	-0.13	0.46	-0.05	0.30	0.50	0.58	0.08	0.19	0.02
ϕ_{11}	1.35	1.34	1.33	1.40	0.80	1.42	1.66	1.22	1.03	1.24	1.24	0.74	1.23	0.85
ϕ_{12}	1.13	1.25	1.20	1.10	0.78	0.85	1.00	1.17	0.57	0.73	-0.21	0.41	0.92	0.98
ϕ_{13}	1.28	1.39	1.33	1.09	1.10	0.93	1.52	1.23	1.33	2.12	1.27	0.88	1.18	0.79
ϕ_{14}	1.40	1.64	1.71	0.74	0.00	1.22	0.95	1.34	0.02	0.16	0.09	0.63	0.62	0.54
ϕ_{21}	-0.37	-0.37	-0.36	-0.42	-0.12	-0.53	-0.65	-0.22	-0.11	-0.18	-0.28	0.19	-0.37	0.32
ϕ_{22}	-0.19	-0.25	-0.21	-0.04	0.03	0.01	-0.05	-0.03	0.48	0.30	1.23	0.56	-0.02	0.01
ϕ_{23}	-0.37	-0.52	-0.46	-0.19	-0.37	0.01	-0.69	-0.41	-0.30	-1.07	-0.20	0.17	-0.41	-0.04
ϕ_{24}	-0.40	-0.69	-0.75	0.21	0.82	-0.18	-0.10	-0.38	0.88	0.62	0.79	0.30	0.29	0.44
$\ln(\sigma)$	-3.14	-3.71	-3.78	-3.69	-2.32	-3.14	-3.24	-3.71	-3.68	-3.76	-3.45	-3.77	-2.93	-4.10

The DGP is given in (21).

Table 7: Rejection frequencies of diagnostic tests

		HEGY				seasonality	seasonal	periodicity	seasonality
DGP	series	root 1	root −1	roots ±i	CH	in mean	lag	in AR	in variance
	Unadjusted data	5%	9%	11%	89%	100%	100%	13%	6%
DGP1 (17)	X12-ARIMA	6%	100%	100%	0%	0%	63%	8%	4%
	TRAMO/SEATS	4%	100%	100%	0%	0%	97%	5%	6%
	Unadjusted data	6%	88%	95%	14%	84%	79%	32%	8%
DGP2 (18)	X12-ARIMA	6%	100%	100%	0%	0%	37%	7%	5%
	TRAMO/SEATS	6%	100%	100%	0%	0%	15%	5%	5%
	Unadjusted data	7%	78%	74%	60%	97%	90%	30%	8%
DGP3 (19)	X12-ARIMA	6%	100%	100%	0%	0%	57%	7%	5%
	TRAMO/SEATS	6%	100%	100%	0%	0%	36%	4%	6%
	Unadjusted data	6%	10%	11%	89%	100%	100%	31%	9%
DGP4 (20)	X12-ARIMA	6%	100%	100%	0%	0%	100%	20%	9%
	TRAMO/SEATS	12%	100%	100%	8%	0%	78%	9%	10%
	Unadjusted data	36%	78%	86%	54%	90%	81%	72%	13%
DGP5 (21)	X12-ARIMA	38%	100%	100%	0%	0%	55%	65%	23%
	TRAMO/SEATS	39%	100%	100%	0%	0%	28%	61%	27%

The diagnostic tests are discussed in Section 3.

Table 8: Null hypotheses of tests and favorable outcomes after seasonal adjustment

Test	H_0	Favorable rejection frequency
HEGY (zero frequency, 1)	presence of unit root	unchanged
HEGY (seasonal frequency, -1, $\pm$i)	presence of seasonal unit root	high
CH	stationary seasonal process	low
Seasonality in mean	equal seasonal dummies	low
Seasonal lag	absence of correlation	low
Periodicity in AR parameters	absence of periodicity	low
Seasonality in variance	absence of seasonality	low

DPG1

The first panel of Table 7 displays the simulation results of DGP1. Columns 4–6 show the results of the HEGY tests. The rejection frequency of the test for the root at the zero frequency is 5% for the unadjusted series as expected. For the seasonally adjusted series they are about 5% and hence both seasonal adjustment methods do not seem to affect the unit root in the series. The rejection frequencies for the roots -1 and $\pm$i are about 10%. The slight size distortion is due to the fact that we select the lag-order of the test regression using BIC to mimic reality. If we fix the lag-order at the true value the size is 5%. Both seasonal adjustment methods remove the seasonal unit roots from the series leading to 100% rejection frequencies. The CH test rejects constant seasonality in 89% of the cases for the unadjusted series. After seasonal adjustment constant seasonality (if present) cannot be rejected and hence this suggests that both seasonal adjustment methods remove the seasonal unit roots in an adequate way. The seventh column shows that the presence of different seasonal means is rejected after applying both seasonal adjustment methods. X12-ARIMA seems to remove fourth order correlation from the series in a better way than TRAMO/SEATS where in 97% of the cases a zero coefficient for seasonal lag is rejected. The absence of periodicity in the AR parameters and the absence of seasonality in the variance is rejected in about 5% of the cases for the adjusted and unadjusted data. Note that there is a slight size distortion in the test for periodicity for the raw series which is again due to the fact that we select the lag-order of the test regression using BIC.

DPG2

The second panel of Table 7 displays the results for DGP2. Again we reject the presence of the root 1 in about 5% of the cases for the unadjusted and adjusted data. The presence of seasonal unit roots -1 and ±i is rejected in more than 88% of the cases for the unadjusted series and always rejected after correction. The CH tests for constant seasonality is rejected in 14% of the cases for the unadjusted series. The small size distortion is due to the fact that we have a large MA component which is not completely captured by the nonparametric estimate of the serial correlation in the series. After seasonal correction the rejection frequency is zero. The pattern of the outcomes of the remaining tests corresponds to the results for DGP1. However, the rejection frequency for a zero parameter at the seasonal lag is now higher for the X12-ARIMA than for the TRAMO/SEATS corrected series. Hence, TRAMO/SEATS performs slightly better.

DPG3

The HEGY procedure rejects the presence of seasonal unit roots in about 75% of the cases as can be seen from the third panel of Table 7. This rejection frequency is 100% for the adjusted series. The presence of the nonseasonal unit root is rejected in about 6% of the cases for both the unadjusted and adjusted series. Constant seasonality is rejected in 60% for the raw series and never rejected for the adjusted series. The outcomes of the remaining tests correspond to the results for DGP2. Hence, the performance of both seasonal adjustment methods is about the same.

DPG4

TRAMO/SEATS uses the airline model to remove seasonality from a series. Hence, we expect that this correction should perform best for this DGP, see fourth panel of Table 7. Remarkably, the presence of a unit root at the zero frequency is rejected in 12% of the cases after applying TRAMO/SEATS, while X12-ARIMA reports a rejection frequency of about 5%. Seasonal unit roots are removed properly as the rejection frequencies after seasonal adjustment are 100%. Note that we have a little size distortion for the seasonal unit roots tests for the raw series which is again due to the fact that we select the lag order of the test regression using BIC. The CH test rejects constant seasonality in 8% of the cases after applying TRAMO/SEATS, while X12-ARIMA never rejects constant seasonality. The parameter belonging to the seasonal lag remains significant after seasonal correction with X12-ARIMA. For TRAMO/SEATS, however, we reject in 78% of the cases. Hence, TRAMO/SEATS seems to perform a little bit better. The outcomes of the remaining test are as expected.

DPG5

The final panel of Table 7 displays the results for DGP5. The presence of a unit root is rejected in about 38% of the cases for the adjusted and unadjusted series. This is due to the fact that many of the parameter settings correspond to processes which are close to unit root type behavior. Seasonal unit roots are rejected in about 80% of the cases for the raw series and in 100% of the cases for the adjusted data. The CH test reports constant seasonality after seasonal correction. A clear difference with the previous DGPs is that the test for equal autoregressive parameters is rejected in more than 60% of the cases. This holds for both the unadjusted and the adjusted series. Hence, both adjustment filters do not remove this type of seasonality from the series. After seasonal adjustment there also seems to be more seasonality in the variance. This is not a surprise as periodic time series with constant variance of the error term, may have different variances across the season, see, for example, Franses and Paap (2004, p. 31–33). Finally, although the DGP5 is a second order autoregressive model, we reject in about 80% of the cases a zero parameter at the seasonal lag. After seasonal adjustment this percentage is reduced for both seasonal adjustment methods but TRAMO/SEATS performs better.

In sum, we conclude that both seasonal adjustment methods remove stochastic seasonal patterns due to seasonal unit roots or stochastic trigonometric seasonality in an adequate way. Rejection frequencies of seasonal unit roots are 100% after applying the seasonal adjustment filters. The CH test for constant seasonality is never rejected after applying the seasonal adjustment filter except for DGP4 where we reject constant seasonality in 8% of the cases after applying TRAMO/SEATS. Both adjustment methods do not seem to affect the presence of a unit root although for DGP4 there is a slight increase in the rejection frequency after applying TRAMO/SEATS. Different means across the seasons are fully captured by both methods. We detect significant correlation at the seasonal lag after applying TRAMO/SEATS in fewer cases than after applying X12-ARIMA. Applying TRAMO/SEATS also leads to less periodicity in the autoregressive parameters but the differences with X12-ARIMA are relatively small.

6 Concluding Remarks

In this chapter we have demonstrated that, when averaged over many realistic DGPs and large samples, the CENSUS X12-ARIMA and TRAMO/SEATS methods seem to perform about equally well. We acknowledge the possibility that for specific series the adjusted series may well be different across methods, but on average our simulations do not indicate a preference for either one of the two methods.

Hence, in the end, our results suggest that a preference for one of the methods merely amounts to a matter of taste. We must say though that an advantage of the TRAMO/SEATS method is that it easily allows for the construction of confidence bounds around seasonally adjusted data, see Koopman and Franses (2002). This feature seems to do justice to the fact that, after all, seasonally adjusted data are estimates which are based on real data.

References

Bell, W. R. (1987), A Note on Overdifferencing and the Equivalence of Seasonal Time Series Models With Monthly Means and Models With $(0, 1, 1)_{12}$ Seasonal Parts When $\theta = 1$, *Journal of Business & Economic Statistics*, **5**, 383–387.

Bell, W. R. and S. C. Hillmer (1984), Issues Involved with the Seasonal Adjustment of Economic Time Series (with discussion), *Journal of Business & Economic Statistics*, **2**, 291–320.

Bell, W. R. and M. Kramer (1996), Towards Variances for X-11 Seasonal Adjustments, Tech. Rep. 96/07, Statistical Research Division, Bureau of the Census.

Box, G. E. P. and G. M. Jenkins (1970), *Time Series Analysis; Forecasting and Control*, Holden-Day, San Francisco.

Box, G. E. P., G. M. Jenkins, and G. C. Reinsel (1994), *Time Series Analysis, Forecasting and Control*, 3 edn., Prentice Hall, Englewood Cliffs.

Burman, J. P. (1980), Seasonal Adjustment by Signal Extraction, *Journal of the Royal Statistical Society A*, **143**, 321–337.

Busetti, F. and A. M. R. Taylor (2003), Testing Against Stochastic Trend and Seasonality in the Presence of Unattended Breaks and Unit Roots, *Journal of Econometrics*, **117**, 21–53.

Canova, F. and B. E. Hansen (1995), Are Seasonal Patterns Constant over Time? A Test for Seasonal Stability, *Journal of Business & Economic Statistics*, **13**, 237–252.

Dagum, E. B. (1980), The X-11-ARIMA Seasonal Adjustment Method, Tech. Rep. 12-564E, Statistics Canada, Ottawa.

Doornik, J. A. (1999), *Object-Oriented Matrix Programming Using Ox*, 3rd edn., London: Timberlake Consultants Press and Oxford: www.nuff.ox.ac.uk/Users/Doornik.

Ericsson, N. R., D. F. Hendry, and H. A. Tran (1994), Cointegration, Seasonality, Encompassing, and the Demand for Money in the United Kingdom, in C. P. Hargreaves (ed.), *Nonstationary Time Series Analysis and Cointegration*, Oxford University Press, Oxford.

Findley, D. F., B. C. Monsell, W. R. Bell, M. C. Otto, and B.-C. Chen (1998), New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program, *Journal of Business & Economic Statistics*, **16**, 127–177.

Franses, P. H. (1996), *Periodicity and Stochastic Trends in Economic Time Series*, Oxford University Press, Oxford.

Franses, P. H. and R. Paap (1999), Does Seasonality Influence the Dating of Business Cycle Turning Points?, *Journal of Macroeconomics*, **21**, 79–92.

Franses, P. H. and R. Paap (2004), *Periodic Time Series Models*, Oxford University Press, Oxford.

Ghysels, E. (1994), On the Periodic Structure of the Business Cycle, *Journal of Business & Economic Statistics*, **12**, 289–298.

Ghysels, E., C. W. J. Granger, and P. L. Siklos (1996), Is Seasonal Adjustment a Linear or Nonlinear Data Filtering Process?, *Journal of Business & Economic Statistics*, **14**, 374–386.

Ghysels, E. and D. R. Osborn (2001), *The Econometric Analysis of Seasonal Time Series*, Cambridge University Press, Cambridge.

Ghysels, E. and P. Perron (1993), The Effect of Seasonal Adjustment Filters on Tests for a Unit Root, *Journal of Econometrics*, **55**, 57–98.

Gómez, V. and A. Maravall (1994), Program SEATS Signal Extraction in ARIMA Time Series: Instructions for the User, Tech. Rep. ECO 94/28, European University Institute, Florence, Italy.

Gómez, V. and A. Maravall (1997), Programs TRAMO and SEATS, Instructions for the User, Beta Version: November 1997.

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

Harvey, A. C. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press, Cambridge.

Hillmer, S. C. and G. C. Tiao (1982), An ARIMA-Model-Based Approach to Seasonal Adjustment, *Journal of the American Statistical Association*, **77**, 63–70.

Hylleberg, S. (1986), *Seasonality in Regression*, Academic Press, Orlando.

Hylleberg, S. (1994), Modelling Seasonal Variation, in C. Hargreaves (ed.), *Nonstationary Time Series Analysis and Cointegration*, Oxford University Press, Oxford.

Hylleberg, S., R. F. Engle, C. W. J. Granger, and B. S. Yoo (1990), Seasonal Integration and Cointegration, *Journal of Econometrics*, **44**, 215–238.

Koopman, S. J. and P. H. Franses (2002), Constructing Seasonally Adjusted Data with Time-Varying Confidence Intervals, *Oxford Bulletin of Economics and Statistics*, **64**, 509–526.

Laroque, G. (1977), Analyse d'une Méthode de Désaissonnalisation: Le Programme X-11 du Bureau of Census, Version Trimestrielle, *Annales de l'INSEE*, **28**, 105–127.

Maravall, A. (1985), On Structural Time Series Models and the Characterization of Components, *Journal of Business & Economic Statistics*, **3**, 350–355.

Maravall, A. (1995), Unobserved Components in Economic Time Series, in H. Pesaran, P. Schmidt, and M. Wickens (eds.), *The Handbook of Applied Econometrics*, vol. 1, Basil Blackwell, Oxford, pp. 12–72.

Maravall, A. and D. A. Pierce (1987), A Prototypical Seasonal Adjustment Model, *Journal of Time Series Analysis*, **8**, 177–193.

Matas-Mir, A. and D. R. Osborn (2003), Seasonal Adjustment and the Detection of Business Cycle Phases, The School of Economics Discussion Paper Series 0304, School of Economics, The University of Manchester.

Miron, J. A. (1996), *The Economics of Seasonal Cycles*, MIT Press, Cambridge, Masachusetts.

Shiskin, J. and H. Eisenpress (1957), Seasonal Adjustment by Electronic Computer Methods, *Journal of the American Statistical Association*, **52**, 415–449.

Shiskin, J., A. H. Young, and J. Musgrave (1967), The X-11 Variant of the Census Method II Seasonal Adjustment Program, Technical report, Bureau of the Census, US Department of Commerce, Washington DC, USA.

Sims, C. A. (1974), Seasonality in Regression, *Journal of the American Statistical Association*, **69**, 618–627.

Taylor, A. M. R. (2003), Robust Stationarity Tests in Seasonal Time Series Processes, *Journal of Business & Economic Statistics*, **21**, 156–163.

Wallis, K. F. (1974), Seasonal Adjustment and Relations between Variables, *Journal of the American Statistical Association*, **69**, 18–31.