Economic Forecasting
Lecture 6: Seasonality

Richard G. Pierse

1 Introduction

Many economic variables, when observed at quarterly or monthly frequencies, exhibit a repeating seasonal pattern. Some firms may face demand that is seasonal (ice cream, suntan lotion, fireworks) or supply that is seasonal (agricultural produce). It is well known that quarterly consumer expenditure peaks in the fourth quarter of each year because of Christmas spending. Employment and unemployment data exhibit seasonal fluctuations because school and university leavers tend to enter the job market at the end of the academic year in September whereas firms tend to recruit at different times in the year.

Seasonality can be thought of as a type of cyclical pattern where the cycles have particular seasonal frequencies corresponding to periods of a fixed number of months or a quarters. However, where economics has developed theories to explain cyclical behaviour, such as real business cycle theory, seasonality has largely been neglected in the economics literature. (One notable exception is the work of Denise Osborn, see *inter alia* Osborn and Smith (1989)). In fact the usual approach to seasonality is to remove the seasonal pattern by seasonal adjustment. Many data series released by statistic agencies have already been seasonally adjusted, usually using the procedure known as *X-12*. For example, the data series on Canadian employment that we used in the previous lecture was seasonally adjusted data, with the seasonal pattern removed. While this may make sense when we are primarily interested in trends (is unemployment going up or going down), seasonal patterns may change over time and may be related across variables. Also, the seasonal fluctuations may be a large part of the variation of the whole variable as with the US liquor sales series illustrated in Figure 1 so that, to remove it would be to throw away most of what is going on. For these reasons, econometricians and forecasters often choose to work with seasonally *unadjusted* data and explicitly to model the seasonal pattern.

As with the models of trends and cycles we have explored in earlier lectures, seasonal models are either deterministic or stochastic. The simplest deterministic

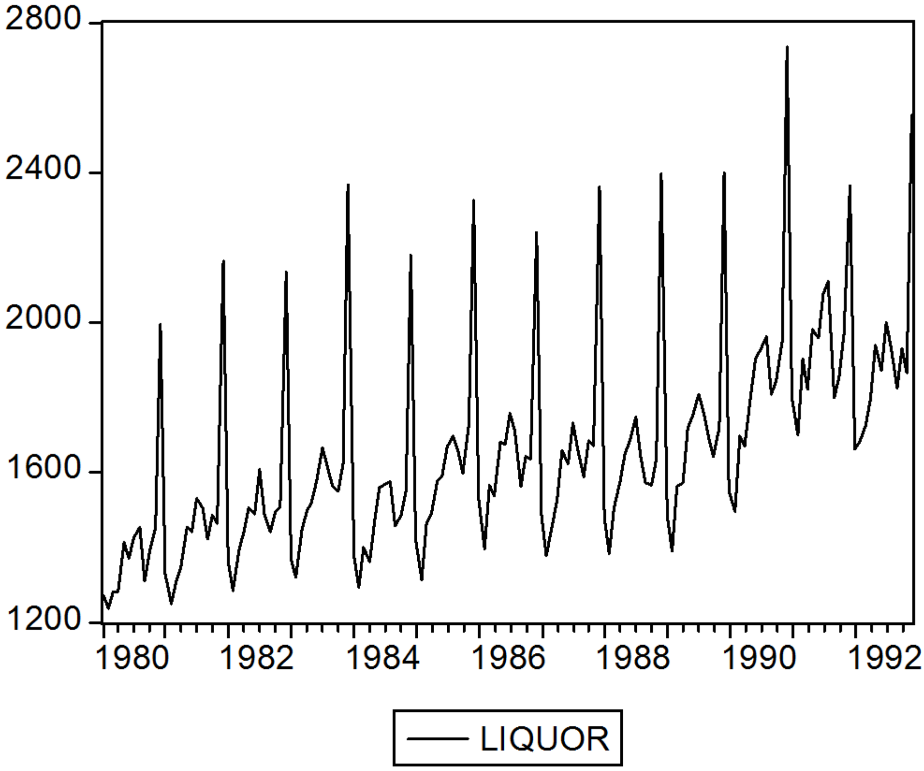

Figure 1: US Liquor Sales: 1980:01–1992:12

model is the seasonal dummy model. The most popular stochastic model is the seasonal *ARIMA* model which extends the Box and Jenkins model to seasonal time series. One interesting issue in seasonal models is the possibility of seasonal unit roots and tests for these have been devised in the literature. Finally, we look at seasonal adjustment procedures and in particular the popular *X-12* seasonal adjustment package.

2 Deterministic Seasonality

The simplest deterministic seasonal model is the seasonal dummy model

$$
\begin{aligned}
S_t &= \delta_1 d_{1t} + \delta_2 d_{2t} + \cdots + \delta_s d_{st} \qquad (1)\\
&= \sum_{i=1}^{s} \delta_i d_{it}
\end{aligned}
$$

where s is the number of seasons and d_{it} is a dummy variable taking the value 1 in the ith season of the year and 0 in all other seasons, $i = 1, \ldots, s$. The parameters δ_1 up to δ_s are coefficients on the seasonal dummy variables. For the quarterly case, $s = 4$, the dummies are

$$
\begin{aligned}
d_1 &= \{1,0,0,0,1,0,0,0,1,0,0,0,\ldots\}\\
d_2 &= \{0,1,0,0,0,1,0,0,0,1,0,0,\ldots\}\\
d_3 &= \{0,0,1,0,0,0,1,0,0,0,1,0,\ldots\}\\
d_4 &= \{0,0,0,1,0,0,0,1,0,0,0,1,\ldots\}\,.
\end{aligned}
$$

The dummies are like intercepts that are different for each season. If there were no seasonal pattern, all the coefficients δ_i would be equal and the model would reduce to

$$S_t = \mu$$

where

$$\mu = \delta_1 = \delta_2 = \cdots = \delta_s.$$

Note that, in general,

$$d_1 + d_2 + d_3 + d_4 = \{1,1,1,1,\ldots\}$$

which is just a standard intercept so that we can always rewrite the seasonal dummy model (1) by dropping any one of the dummies and replacing it with an intercept as in

$$S_t = \mu + \sum_{i=1}^{s-1} \gamma_i d_{it}. \qquad (2)$$

In this formulation, μ is now the intercept for the omitted season s and the coefficients γ_i give the seasonal increase or decrease relative to the omitted season with $\gamma_i = \delta_i - \delta_s$. Note that, when an intercept is included as in (2), one of the seasonal dummies must be dropped because otherwise the intercept and the set of dummies will be perfectly collinear.

The seasonal dummy model (1) can also be represented equivalently in terms of sines and cosines since

$$\sum_{i=1}^{s} \delta_i d_{it} = \mu + \sum_{i=1}^{s/2} \left[\alpha_i \cos(\frac{2\pi i t}{s}) + \beta_i \sin(\frac{2\pi i t}{s})\right] \qquad (3)$$

where the s coefficients α_i and β_i are related to the coefficients δ_i. For the quarterly case, $s = 4$, we have the correspondences

$$\begin{aligned}
\mu &= (\delta_1 + \delta_2 + \delta_3 + \delta_4)/4 \\
\alpha_1 &= (\delta_4 - \delta_2)/2 \\
\alpha_2 &= (\delta_2 + \delta_4 - \delta_1 - \delta_3)/2 \\
\beta_1 &= (\delta_1 - \delta_3)/2
\end{aligned}$$

and $\beta_2 = 0$. Note that the final sine term in (3) is always zero because $\sin(\pi t) = 0$ for all t. The sine and cosine waves in (3) correspond to cycles with periods s, $s/2$, ..., 2. In the quarterly case the cycles are of period 4 and 2, in other words an annual cycle and a half-yearly cycle.

The parameters in the dummy variable model may be estimated from the regression

$$y_t = \delta_1 d_{1t} + \delta_2 d_{2t} + \cdots + \delta_s d_{st} + \varepsilon_t$$

or the regression

$$y_t = \mu + \gamma_1 d_{1t} + \gamma_2 d_{2t} + \cdots + \gamma_{s-1} d_{s-1t} + \varepsilon_t. \tag{4}$$

In either case, the regression residuals

$$e_t = y_t - \widehat{y}_t$$

define a *seasonally adjusted* series

$$y_t^{SA} = y_t - \widehat{y}_t.$$

Figure 2 illustrates the US liquor sales series seasonally adjusted using the additive seasonal dummy model (4). Although most of the seasonal variation has been removed, notice that seasonal spikes remain in the early part of the time series. This illustrates a potential drawback of the seasonal dummy model which is that the seasonal effect is assumed to be constant and not to vary with the level of the series. In a trended series like Figure 2, this assumption may be unrealistic.

An alternative model is given by the seasonal trend model

$$y_t = \mu + bt + \sum_{i=1}^{s-1} (\gamma_i + \beta_i t) d_{it} + \varepsilon_t. \tag{5}$$

where both the intercept and the trend are allowed to be different in each season with the intercept in season i given by $\mu + \gamma_i$ and the trend in season i given by $(b + \beta_i)t$. Figure 3 illustrates the liquor series seasonally adjusted using the model (5). Note that the seasonal spikes in the early part of the series have disappeared and the adjusted series looks much more satisfactory.

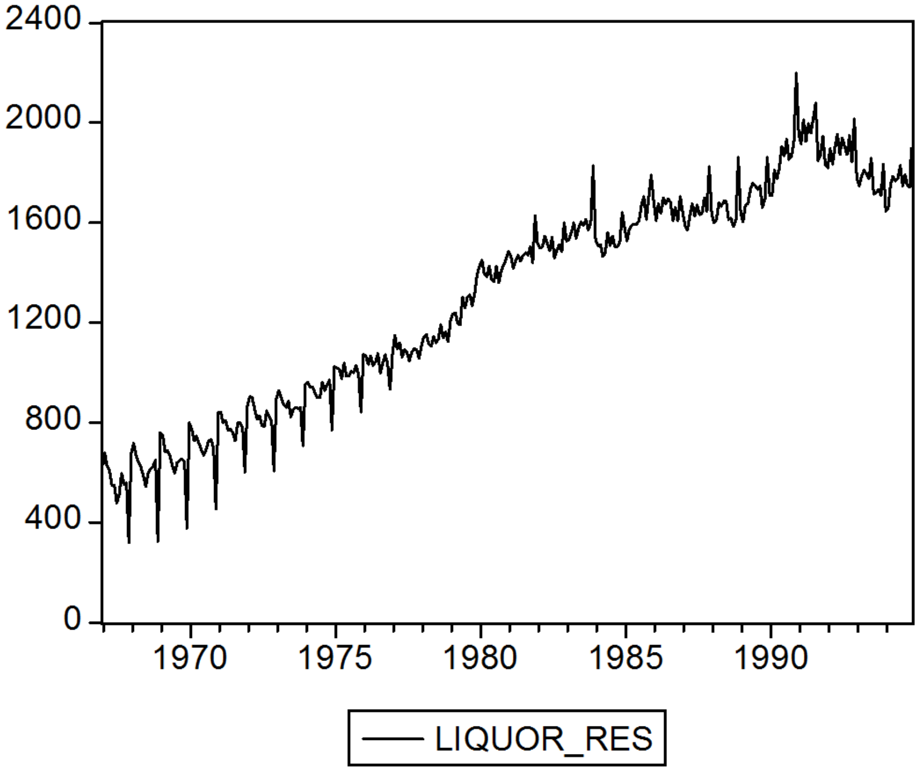

Figure 2: US Liquor sales: seasonally adjusted using additive seasonal dummies

3 Calendar Effects

In addition to a regular pattern, seasonal data can be subject to calendar effects due to *holiday variation* and *trading day variation*. *Holiday variation* is due to the effect of moveable feasts such as Easter or Ramadan, which can occur on different dates in different years. Since these holidays can affect sales, shipments, hours worked etc., it is important to take account of them in a forecasting model of a seasonal series. *Trading day variation* is due to the fact that different months (and even weeks) have different numbers of business days. This can be very important, for example, when building a monthly forecasting model of volume traded on the London Stock Exchange.

Both types of calendar effects can be modelled by the use of dummy variables, in a similar way to standard seasonal dummies. For example, in a monthly model an Easter dummy can be defined taking the value 1 when the month contains

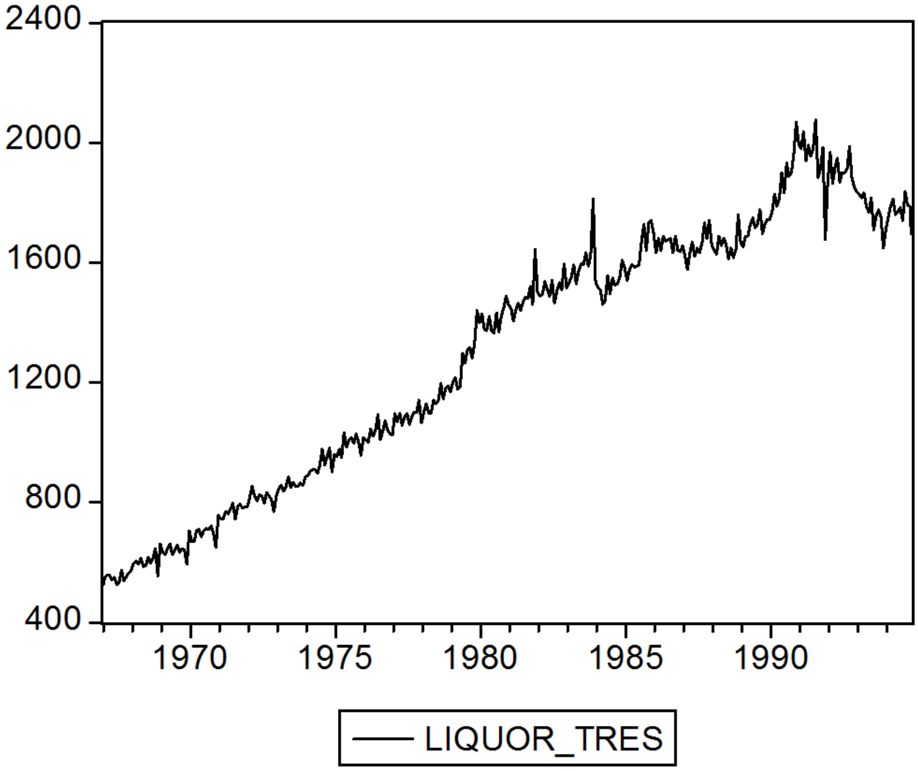

Figure 3: US Liquor sales: seasonally adjusted using seasonal trend model

Easter and 0 otherwise. Each holiday would be modelled in a similar way by a special dummy. Similarly, in a monthly model we can define a dummy equal to the number of trading days in the month to pick up trading day effects. Adding these two sets of dummies to the model (2) leads to the regression model model

$$y_t = \mu + \sum_{i=1}^{s-1} \gamma_i d_{it} + \sum_{i=1}^{n} \delta_i d_{it}^{hol} + \eta d_t^{td} + \varepsilon_t$$

where n is the number of holidays and d_{it}^{hol} is a dummy variable for the i-th holiday with associated coefficient δ_i and d_t^{td} is the trading day dummy with associated coefficient η. The parameters in this model can be estimated by ordinary least squares.

4 Seasonal *ARMA* Models

Seasonal effects can be modelled stochastically through the seasonal *ARMA* model. Consider the first order seasonal autoregressive *SAR(1)* model

$$y_t = \rho y_{t-s} + \varepsilon_t$$

or

$$(1 - \rho L^s) y_t = \varepsilon_t .$$

In this model y_t is related to its value in the same season of the previous year. As usual, for stationarity we require that $|\rho| < 1$ and for the repeating patterns that we observe in seasonal series such as Figure **??** we would normally expect $\rho > 0$. We can also consider the first order seasonal moving average *SMA(1)* model

$$y_t = \varepsilon_t + \psi \varepsilon_{t-s}$$

where y_t is related to the disturbance in the same season of the previous year and where, for identifiability we require that $|\psi| \leq 1$.

More generally, we can define the seasonal *ARMA(P,Q)* model

$$\rho(L^s) y_t = \psi(L^s) \varepsilon_t \tag{6}$$

where

$$\rho(L^s) = \rho_1 L^s + \rho_2 L^{2s} + \cdots + \rho_P L^{Ps}$$

is a P-th order autoregressive process and

$$\psi(L^s) = \psi_1 L^s + \psi_2 L^{2s} + \cdots + \psi_Q L^{Qs}$$

is a Q-th order moving average process.

The *seasonal ARMA* model (6) can be combined with the standard *ARMA* model of lecture 4 to define the *multiplicative seasonal ARMA(p,q)(P,Q)* model

$$\rho(L^s)\phi(L) y_t = \psi(L^s)\theta(L)\varepsilon_t \tag{7}$$

where

$$\phi(L) = 1 - \phi_1 L - \phi_2 L^2 - \cdots - \phi_p L^p$$

and

$$\theta(L) = 1 + \theta_1 L + \theta_2 L^2 + \cdots + \theta_q L^q .$$

In this model the seasonal *SAR(P)* and *SMA(Q)* polynomials multiply the non-seasonal *AR(p)* and *MA(q)* polynomials.

5 Seasonal Unit Roots

Consider the model

$$y_t = y_{t-s} + \varepsilon_t$$

or

$$(1 - L^s)y_t = \Delta_s y_t = \varepsilon_t$$

where Δ_s is the seasonal difference operator defined by $\Delta_s x_t = x_t - x_{t-s}$. This is a unit root process that is known as the *seasonal random walk*. In fact there are s random walks, one for each season of the year. This can be seen by factorising the process for the quarterly case, $s = 4$ which gives

$$\begin{aligned} \Delta_4 y_t &= (1 - L^4)y_t = (1 - L)(1 + L + L^2 + L^3)y_t \\ &= (1 - L)(1 + L)(1 + iL)(1 - iL)y_t \end{aligned} \tag{8}$$

which has *four* unit roots: one at frequency 0 which is the conventional random walk and three at seasonal frequencies corresponding to cycles of 2 quarters, and 4 quarters (the pair of imaginary roots) respectively.

For time series including seasonal unit roots the seasonal *ARIMA(P,D,Q)* model is

$$\rho(L^s)\Delta_s^D y_t = \psi(L^s)\varepsilon_t \tag{9}$$

and the *joint multiplicative seasonal ARIMA(p,d,q)(P,D,Q)* model is defined by

$$\rho(L^s)\phi(L)\Delta_s^D \Delta^d y_t = \psi(L^s)\theta(L)\varepsilon_t \tag{10}$$

where D is the order of seasonal differencing. Note that this model combines all three structural components of a time series: trend, cycle and seasonal in a single model.

Testing for seasonal unit roots is important and two main tests have been developed: the Dickey-Hasza-Fuller test which is a joint test for all of the unit roots and the Hylleberg-Engle-Granger-Yoo test which tests the roots separately.

5.1 Dickey-Hasza-Fuller (*DHF*) test

Dickey, Hasza and Fuller (1984) derive a test of the hypothesis $\alpha_s = 0$ in the model

$$\Delta_s y_t = \alpha_s y_{t-s} + \varepsilon_t$$

against the alternative that $\alpha_s < 0$. The test statistic is simply the t-value on $\widehat{\alpha}_s$ and critical values for this test are presented in their paper (reprinted in Hylleberg (1992)) for the cases $s = 2$, 4, and 12. As with standard Dickey-Fuller tests, deterministic components (constant and trend) can be added to the specification but do affect the distribution of the statistic. Lagged values of $\Delta_s y_t$ can be added to 'whiten' the errors without affecting the distribution.

From (8) it can be seen that the *DHF(4)* test is a *joint test of four unit roots* against an alternative of no unit roots. In particular it tests for a unit root at zero frequency (i.e. the long run) at the same time as testing for seasonal unit roots.

5.2 Hylleberg-Engle-Granger-Yoo (*HEGY*) test

Hylleberg, Engle, Granger and Yoo (1990) develop a framework in which it is possible to *separately* test the four unit roots in (8) in the quarterly case ($s = 4$). This is based on constructing the model:

$$\begin{aligned}\Delta_4 y_t &= \pi_1(1 + L + L^2 + L^3)y_{t-1} \\ &\quad -\pi_2(1 - L + L^2 - L^3)y_{t-1} \\ &\quad -\pi_3(1 - L^2)y_{t-2} - \pi_4(1 - L^2)y_{t-1} + \varepsilon_t\,.\end{aligned}$$

The t-ratio on $\widehat{\pi}_1$ is a test of the null of a unit root at zero frequency and can be shown to follow a Dickey-Fuller distribution. The t-ratio on $\widehat{\pi}_2$ is a test of a unit root at the semi-annual frequency which also has a Dickey-Fuller distribution. The t-value on $\widehat{\pi}_3$ is a test for a unit root at the annual frequency, *conditional on the hypothesis that* $\pi_4 = 0$, and follows a *DHF(2)* distribution. Finally, a joint test of the hypothesis that $\pi_3 = 0$ and $\pi_4 = 0$ can be constructed from the F-statistic for a test of this restriction. The distribution of this last statistic is close to the standard $F_{2,T-k}$ distribution and critical values are tabulated in *HEGY*. As usual, adding lagged values of $\Delta_4 y_t$ to the regression does not change the distributions. However, if deterministic seasonal dummies are included in the regression, then this does affect the distribution of the tests of π_2, π_3, and π_4 leading to fatter tails.

The *HEGY* tests have been extended to the monthly case by Beaulieu and Miron (1993) and Franses (1991). Alternative tests for seasonal unit roots have been developed by Kunst (1997) and Osborn, Chui, Smith and Birchenhall (1988) and are discussed in Ghysels and Osborn (2001).

6 Seasonal Adjustment

Seasonal adjustment is the removal of the seasonal component from a data series. We have seen that in a deterministic seasonal model, the variables can be seasonally adjusted by regressing on the appropriate seasonal dummies and taking the residuals from the regression. In practice, statistical agencies that publish seasonally adjusted data use much more sophisticated methods to model and remove the seasonal patterns in the data. The most widely used seasonal adjustment procedure is the *X-12* algorithm developed by the U.S. Census Bureau. An alternative procedure based on the computer programs *TRAMO* (*Time series Regression with ARIMA noise, Missing observations and Outliers*) and *SEATS* (*Signal Extraction in ARIMA Time Series*) has been developed by the Spanish statistician Agustín Maravall. This is used in some European countries. Both procedures are available in *EViews*.

6.1 The *X-12* Seasonal Adjustment Program

Since the 1960s the U.S. Census Bureau have been developing different versions of a method for the automatic seasonal adjustment of published data series. The latest version of this method is called *X-12* and is widely used by statistical agencies around the world, including the UK. The complete procedure is quite complicated and consists of two stages and several steps with many variants available.

The first stage of the procedure, which is known as *X-12-ARIMA*, is the building of a *multiplicative seasonal ARIMA* model based on (10) but including deterministic components:

$$\rho(L^s)\phi(L)\Delta_s^D\Delta^d(y_t - \sum_{i=1}^{k}\beta_i x_{it}) = \psi(L^s)\theta(L)\varepsilon_t. \tag{11}$$

Here the x_{it} represent a set of deterministic components which generally comprise an intercept and linear trend, additive seasonal dummies plus holiday and trading day dummies. In addition, outlier dummies may be included to take out the effect of particular outlying observations. This model is then used to increase the sample by forecasting and backcasting observations outside the original sample. These observations are then used in the second main stage of the procedure which is the *X-11* program.

The *X-11* program is the previous incarnation of *X-12*. It has been extensively studied by statisticians and can be approximated by a linear procedure. We will look at the case of monthly data where $s = 12$. The procedure is based on the decomposition of a time series into trend, cycle, seasonal and irregular components of the form

$$y_t = T_t + C_t + S_t + u_t$$

(additive) or

$$y_t = T_t C_t S_t u_t$$

(multiplicative) that we discussed in the first lecture. In the *X-11* procedure, the trend and cycle component are modelled together as the combined trend/cycle component TC_t and (initially) the seasonal and irregular components are combined into a combined component SI_t. In the final stage, the seasonal and irregular components are separated. The procedure consists in a series of moving average filters applied to the series to extract the components. There is no estimation involved. There are three steps to the process which successively refine the estimates of the components.

6.1.1 Step 1

The first step is the estimation of an initial trend and cycle component using the centred moving average filter

$$TC_t^1 = (1/24)(1+L)(L^{-6} + L^{-5} + \cdots + L^4 + L^5)y_t = M(L)y_t.$$

This equation takes a moving average of y_t using 12 observations centred on the current observation. This is then averaged again over two observations. The remaining seasonal plus irregular component is then defined as a residual using either $SI_t^1 = y_t - TC_t^1$, (additive) or $SI_t^1 = y_t/TC_t^1$ (multiplicative). This seasonal plus irregular component is then filtered using another centred moving average as

$$S_t^{F1} = (1/9)(L^{-12} + L^{-11} + \cdots + L^{11} + L^{12})SI_t^1.$$

The initial seasonal component is then defined by either

$$S_t^1 = S_t^{F1} - M(L)S_t^{F1}$$

or

$$S_t^1 = \frac{S_t^{F1}}{M(L)S_t^{F1}}.$$

This ensures that the seasonal component sums to unity over the year. Finally, the first step seasonally adjusted series is defined as either the additive residual

$$y_t^{SA1} = y_t - S_t^1$$

or the multiplicative residual

$$y_t^{SA1} = \frac{y_t}{S_t^1}.$$

6.1.2 Step 2

The second step refines the estimates of the seasonally adjusted series from the first step. Firstly, a new estimate of the trend/cycle component is defined by

$$\begin{aligned} TC_t^2 &= (\gamma_h L^{-h} + \gamma_{h-1}L^{-h+1} + \cdots + \gamma_0 + \cdots + \gamma_{h-1}L^{h-1} + \gamma_h L^h)y_t^{SA1} \\ &= H(L)y_t^{SA1} \end{aligned}$$

where the $\gamma_0, \ldots, \gamma_h$ are a series of fixed symmetric weights. This transformation is called the $2h+1$ term *Henderson filter* and the default value of h is 6. Then the second stage seasonal plus irregular component is defined by either $SI_t^2 = y_t - TC_t^2$, (additive) or $SI_t^2 = y_t/TC_t^2$ (multiplicative). The second stage seasonal plus irregular component is then filtered using the seasonal moving average filter

$$S_t^{F2} = (1/15)(L^{-12} + 1 + L^{12})(L^{-24} + L^{-12} + 1 + L^{12} + L^{24})SI_t^2.$$

and the second stage seasonal factor and adjusted series are defined by either

$$S_t^2 = S_t^{F2} - M(L)S_t^{F2}$$

and

$$y_t^{SA2} = y_t - S_t^2$$

or

$$S_t^2 = \frac{S_t^{F2}}{M(L)S_t^{F2}}$$

and

$$y_t^{SA2} = \frac{y_t}{S_t^2}.$$

6.1.3 Step 3

In the third and final step, the estimates of the second step are used to define a new estimate of the trend/cycle component and define the final decomposition of the series into trend and cycle, seasonal and irregular components. Firstly, the final estimate of the trend plus cycle is derived by applying the Henderson filter again (possibly using a different value of h) to the seasonally adjusted series y_t^{SA2} to give

$$TC_t^3 = H(L)y_t^{SA2}.$$

Then the unexplained irregular component is defined by

$$u_t = y_t^{SA2} - TC_t^3$$

or

$$u_t = y_t^{SA2}/TC_t^3.$$

The final decomposition of the original series into trend/cycle, seasonal and irregular components is then either

$$y_t = TC_t^3 + S_t^2 + u_t$$

or

$$y_t = TC_t^3 S_t^2 u_t.$$

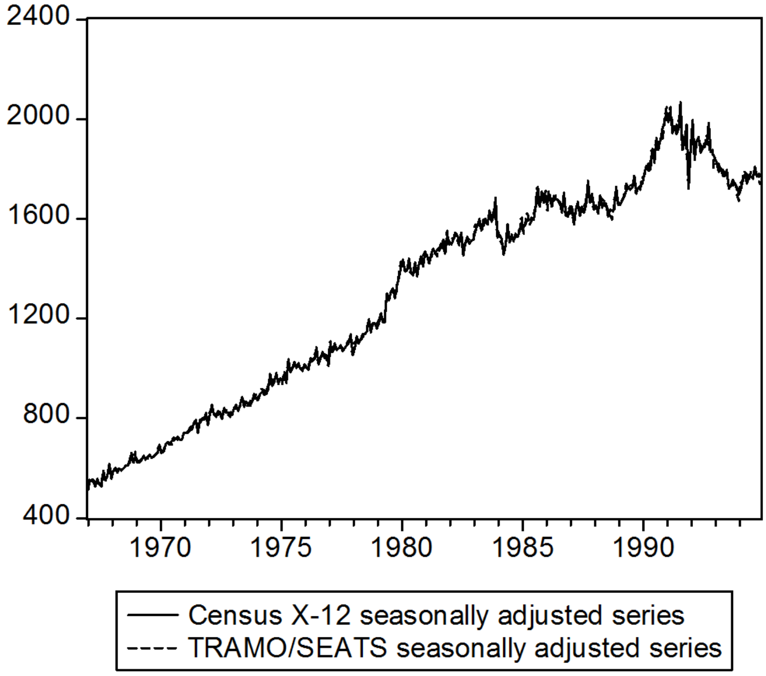

Figure 4: US Liquor sales: seasonally adjusted with *X-12* and *TRAMO/SEATS*

6.2 Example

Figure 4 applies the *X-12* procedure to the US liquor sales series to derive a seasonally adjusted series. The graph also shows for comparison a seasonally adjusted series derived by applying the *TRAMO/SEATS* procedure. In this case it appears that the two procedures give very similar results.

References

[1] Beaulieu, J.J. and J.A. Miron (1993), 'Seasonal unit roots in aggregate U.S. data', *Journal of Econometrics*, 55, 305–328.

[2] Dickey, D.A., D.P. Hasza and W.A. Fuller (1984), 'Testing for unit roots in seasonal time series', *Journal of the American Statistical Association*, 79, 355–67.

[3] Franses, P.H. (1991), 'Seasonality, nonstationarity and the forecasting of monthly time series', *International Journal of Forecasting*, 7, 199–203.

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

[5] Hylleberg, S. (1992), *Modelling Seasonality*, Oxford University Press, Oxford.

[6] Hylleberg, S., R.F. Engle, C.W.J. Granger and B.S. Yoo (1990), 'Seasonal integration and cointegration', *Journal of Econometrics*, 44, 215–38.

[7] Kunst, R.M. (1997), 'Testing for cyclical non-stationarity in autoregressive processes', *Journal of Time Series Analysis*, 18, 123–135.

[8] Osborn, D.R., A.P.L.Chui, J.P.Smith and C.R. Birchenhall (1988), 'Seasonality and the order of integration for consumption', *Oxford Bulletin of Economics and Statistics*, 50, 361–377.

[9] Osborn, D.R. and J.P. Smith (1989), 'The performance of periodic autoregressive models in forecasting seasonal UK consumption', *Journal of Business and Economic Statistics*, 7, 117–127.