WORKING PAPERS AND STUDIES

ISSN 1725-4825

2006 EDITION

Seasonality and unobserved components models: an overview

Conference on seasonality, seasonal adjustment and their implications for short-term analysis and forecasting

Luxembourg, 10-12 May 2006

EUROPEAN COMMISSION

eurostat

THEME
General
and regional
statistics

Europe Direct is a service to help you find answers to your questions about the European Union

Freephone number (*):

00 800 6 7 8 9 10 11

(*) Certain mobile telephone operators do not allow access to 00-800 numbers or these calls may be billed.

A great deal of additional information on the European Union is available on the Internet.
It can be accessed through the Europa server (http://europa.eu).

Luxembourg: Office for Official Publications of the European Communities, 2006

ISBN 92-79-03418-9
ISSN 1725-4825
Catalogue number: KS-DT-06-019-EN-N

Seasonality and unobserved components models: an overview

Andrew Harvey
Faculty of Economics, Cambridge University

December 4, 2006

Abstract

This paper presents an overview of the way in which unobserved components models deal with seasonality and seasonal adjustment. Topics covered include the basic structural model, calendar effects, tests, daily and weekly observations, time-varying splines, robust seasonal adjustment and seasonal specific (periodic) models.

KEYWORDS: Basic structural model, calendar effects, state space form, stationarity tests, time-varying splines, weekly observations.

JEL CLASSIFICATION: C22

1 Introduction

Seasonal adjustment is an exercise in signal extraction. Hence an unobserved components model is a natural starting point. The fundamental reason for building a time series model is that it provides a way of weighting the data that is determined by the properties of the time series.

The basic structural model (BSM) is

$$y_t = \mu_t + \gamma_t + \varepsilon_t, \qquad t = 1, ..., T \tag{1}$$

where μ_t is a stochastic trend, γ_t is a stochastic seasonal component; see Akaike (1980) and Harvey (1989). The irregular component, ε_t, is assumed to be random, and the disturbances in all three components are taken to be mutually and serially uncorrelated. The model may be extended by including explanatory variables other stochastic components, such as cycles. Models of this kind can be used for forecasting and nowcasting as well as signal extraction; see Harvey (2006). The menu-driven STAMP package of Koopman et al (2006) enables many of the methods described here to be implemented.

2 Dummy variable and trigonometric seasonal models

2.1 Deterministic seasonality

A basic requirement of a seasonal component is that when the seasonal effects are fixed, they should sum to zero over a year. Thus if $\gamma_j, j = 1, .., s$, is the seasonal effect in season j,

$$\sum_{j=1}^{s} \gamma_j = 0 \tag{2}$$

The restriction is easily imposed by defining the seasonal effect at time t as

$$\gamma_t = \sum_{j=1}^{s-1} \gamma_j z_{jt}, \qquad t = 1, ..., T$$

where for $t = i, i+s, i+2s, ...,$ and $i = 1, ..., s-1$, the variable z_{jt} is one for $j = i$ and zero for $j \neq i$, while for $t = s, 2s, 3s.., z_{jt} = -1$ for $j = 1, ..., s-1$. In other words γ_t is equal to γ_j in season j and $-\sum_{j=1}^{s-1} \gamma_j$ in season s. The restriction also means that the seasonal effects over the past s periods sum to zero, that is

$$\sum_{j=0}^{s-1} \gamma_{t-j} = 0, \qquad t = 1, ..., T \tag{3}$$

Rather than using a set of dummy variables, a fixed seasonal pattern may be captured by a set of trigonometric terms at the seasonal frequencies, $\lambda_j = 2\pi j/s, \quad j = 1, ..., [s/2]$. The seasonal effect at time t is then

$$\gamma_t = \sum_{j=1}^{[s/2]} \left(\alpha_j \cos \lambda_j t + \beta_j \sin \lambda_j t \right), \quad t = 1, ..., T \tag{4}$$

When s is even, the sine term for $j = s/2$ disappears and so the number of trigonometric parameters, the α_j's and β_j's, is always $s-1$, the same as the number of coefficients in the seasonal dummy formulation. The first frequency, $\lambda_1 = 2\pi/s$, corresponds to a period of twelve months and is known as the *fundamental frequency* while the remaining frequencies are *harmonics*. By using standard trigonometric identies, it is straightforward to show that the seasonal effects over a year sum to zero, as in (3). Provided that the full set of trigonometric terms is included, (4) is equivalent to the dummy variable specification and the estimated seasonal patterns will be identical.

2.2 Stochastic dummies

By introducing a disturbance term into the right-hand side of (3), the seasonal effects can be allowed to change over time. Thus

$$\sum_{j=1}^{s-1} \gamma_{t-j} = \omega_t \quad \text{or} \quad \gamma_t = -\sum_{j=1}^{s-1} \gamma_{t-j} + \omega_t \tag{5}$$

where ω_t is white noise with mean zero and variance σ_ω^2. The bigger the value of σ_ω^2 relative to the variances of other disturbances in the model, the more rapidly the seasonal pattern changes over time and the more rapidly are past observations discounted in constructing a seasonal pattern for the forecast function. The forecasts satisfy the recursion

$$\tilde{\gamma}_{T+l|T} = -\sum_{j=1}^{s-1} \tilde{\gamma}_{T+l-j|T}, \quad l = 1, 2, ... \tag{6}$$

where the starting values are given by the (smoothed) estimates of the seasonal effects, $\gamma_T, ..\gamma_{T-s+2}$, at time T. Thus the seasonal pattern projected into the future is fixed and the seasonal effects sum to zero over any period of one year.

An alternative way of allowing the seasonal dummy variables to change over time is to suppose that the effect of each season evolves as a random walk. This model was introduced by Harrison and Stevens (1976, pp. 217-18). Let γ_{jt} denote the effect of season j at time t. Then

$$\gamma_{jt} = \gamma_{j,t-1} + \omega_{jt}, \quad j = 1, ..., s, \quad t = 1, ..., T \tag{7}$$

where ω_{jt} is a white-noise disturbance term with mean zero and variance σ_ω^2. Although all s seasonal components are continually evolving, only one affects the observation at any particular point in time, that is $\gamma_t = \gamma_{jt}$ for $t = 1, ..., T$, when season j is prevailing at time t. The requirement that the seasonal effects in the forecast function sum to zero over s consecutive time periods is enforced by the restriction that, at any particular point in time, the seasonal components, and hence the disturbances, sum to zero, that is

$$\sum_{j=1}^{s} \gamma_{jt} = 0 = \sum_{j=1}^{s} \omega_{jt}, \quad t = 1, ..., T \tag{8}$$

This restriction is implemented by the correlation structure in

$$Var\left(\boldsymbol{\omega}_t\right) = \sigma_\omega^2 \left(\mathbf{I} - s^{-1}\mathbf{ii}'\right), \quad t = 1, ..., T \tag{9}$$

where $\boldsymbol{\omega}_t = (\omega_{1t}, ..., \omega_{st})'$, coupled with an initial condition requiring that the seasonals sum to zero at $t = 0$. It can be seen from (9) that $Var\left(\mathbf{i}'\boldsymbol{\omega}_t\right) = 0$, so $\mathbf{i}'\boldsymbol{\omega}_t = 0$.

The relationship between the two forms of dummy variable seasonality is examined in Proietti (2000). In practice, it is usually preferable to work with the *balanced dummy variable seasonal model* of (7) though the simplicity of the *single shock* model of (5) can be useful for pedagogic purposes.

2.3 Trigonometric seasonality

A trigonometric seasonal pattern may be allowed to evolve over time by writing the component at each frequency as a recursion and adding disturbances. Thus

$$\gamma_t = \sum_{j=1}^{[s/2]} \gamma_{jt}, \quad t = 1, ..., T \tag{10}$$

and

$$\left.\begin{aligned} \gamma_{jt} &= \gamma_{j,t-1} \cos \lambda_j + \gamma^*_{j,t-1} \sin \lambda_j + \omega_{jt} \\ \gamma^*_{jt} &= -\gamma_{j,t-1} \sin \lambda_j + \gamma^*_{j,t-1} \cos \lambda_j + \omega^*_{jt} \end{aligned}\right\}, \qquad j = 1, ..., [(s-1)/2] \tag{11}$$

where ω_{jt} and ω^*_{jt} are zero mean white-noise processes which are uncorrelated with each other and have a common variance σ_j^2 for $j = 1, ..., [(s-1)/2]$. The component γ^*_{jt} appears as a matter of construction, and its interpretation is not particularly important. When s is even,

$$\gamma_{s/2,t} = \gamma_{s/2,t-1} \cos \lambda_{s/2} + \omega_{s/2,t} = (-1)^t \gamma_{s/2,t-1} + \omega_{s/2,t}. \tag{12}$$

The estimators of the γ_{jT}'s and $\gamma^*_{jT}\prime s$ provide starting values for a projection of the latest seasonal pattern into the future. Hannan, Terrell and Tuckwell (1970) constructed a seasonal model by letting α_j and β_j in (4) evolve as random walks, but it can be shown that this specification is equivalent to the one described above.

Assigning different variances to each harmonic allows them to evolve at varying rates. However, from a practical point of view it is usually desirable to let these variances be the same. Thus

$$Var\left(\omega_{jt}\right) = Var\left(\omega_{jt}^{*}\right) = \sigma_j^2 = \overline{\sigma}_\omega^2, \quad j = 1, ..., \left[\left(s-1\right)/2\right] \tag{13}$$

though for s even,

$$Var\left(\omega_{s/2,t}\right) = \overline{\sigma}_\omega^2/2, \quad t = 1, ..., T \tag{14}$$

As a rule, very little is lost in terms of goodness of fit by imposing this restriction. The model is identical to the balanced dummy variable seasonal model with $\sigma_\omega^2 = 2\overline{\sigma}_\omega^2/s$ for s even and $\sigma_\omega^2 = 2\overline{\sigma}_\omega^2/\left(s-1\right)$ for s odd; see Proietti (2000).

Note that the reduced form of the balanced dummy and trigonometric seasonal models is such that $\sum_{j=0}^{s-1}\gamma_{t-j} = S(L)\gamma_t$ is an $MA(s-2)$ process.

2.4 Daily effects

If some of the seasonal effects are assumed to be the same, the number of dummy variables can be reduced. Since this assumption is particularly relevant for modelling daily effects, it will be described in this context. However, the notation is completely general.

Let w be the number of different types of day in a week and let k_j be the number of days of the j-th type for $j = 1, ..., w$. Thus, for example, if all weekdays are alike but both Saturdays and Sundays are different, $w = 3$, $k_1 = 5$ and $k_2 = k_3 = 1$. The effect associated with the j-th type of day is θ_{jt}, where

$$\theta_{jt} = \theta_{j,t-1} + \chi_{jt}, \quad j = 1, ..., w \tag{15}$$

the disturbance term χ_{jt} having zero mean and variance

$$Var\left(\chi_{jt}\right) = \sigma_\chi^2\left(1 - k_j^2/k\right), \quad j = 1, ..., w \tag{16}$$

where

$$k = \sum_{j=1}^{w} k_j^2$$

The covariances between the disturbances are

$$E\left(\chi_{jt}\chi_{ht}\right) = -\sigma_\chi^2 k_j k_h / k, \quad j \neq h, \quad j,h = 1,...,w \tag{17}$$

If $\mathbf{k}$ and $\boldsymbol{\chi}_t$ are $w \times 1$ vectors with j-th elements k_j and χ_{jt} respectively, the covariance matrix can be written as

$$Var\left(\boldsymbol{\chi}_t\right) = \sigma_\chi^2 \left(\mathbf{I} - k^{-1}\mathbf{k}\mathbf{k}'\right), \qquad t = 1,...,T. \tag{18}$$

If all the seasons are assumed to be different, the balanced seasonal model of (7) is obtained; in that case $w = s$ and $k_j = 1$ for $j = 1,...,s$. As in that model, the specification of the covariance matrix of the disturbances ensures that the daily effects sum to zero over a week since $Var\left(\mathbf{k}'\boldsymbol{\chi}_t\right) = 0$ and so $\mathbf{k}'\boldsymbol{\chi}_t = 0$.

Although the daily effects are changing every day, only one of them affects the observations at a particular point in time. Thus if the day at time t is of the $j-th$ type, the daily effect is

$$\theta_t = \theta_{jt}, \quad t = 1,...,T.$$

2.5 State space form

Putting the daily model in state space form is easy. The transition equation can be written in matrix terms as

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \boldsymbol{\chi}_t \tag{19}$$

where $\boldsymbol{\theta}_t$ is a $w \times 1$ vector containing the daily effects and $\boldsymbol{\chi}_t$ is the corresponding vector of disturbances with covariance matrix (18). The only complication is that the $\mathbf{z}_t$ vector in the measurement equation changes over time so that each day it picks out the appropriate daily effect from $\boldsymbol{\theta}_t$.

One of the $\theta_{jt}'s$ may be dropped from the model as it can be reconstructed from the requirement that, at any particular point in time, the sum of the daily effects over a week should be zero. This enables the number of elements in the state vector to be reduced by one. If, without loss of generality, the w-th type of day is dropped then

$$\theta_{wt} = -k_w^{-1} \sum_{j=1}^{w-1} k_j \theta_{jt} \tag{20}$$

For the w-th type of day the j-th element in $\mathbf{z}_t$ is equal to $-k_j/k_w$ for $j = 1, ..., w-1$. For all other types of day there is unity in the j-th position and zeroes elsewhere. However, it is not necessary to drop an element from the state vector in order to enforce the zero sum restriction since (18) combined with an initial diffuse covariance matrix of the form

$$E\left(\boldsymbol{\theta}_0\boldsymbol{\theta}_0'\right) = \sigma_\kappa^2\left(\mathbf{I} - k^{-1}\mathbf{k}\mathbf{k}'\right),$$

with $\sigma_\kappa^2 \to \infty$, will ensure that it holds.

The balanced dummy seasonal model is put into SSF in the same way, with the simplification that $\mathbf{k} = \mathbf{i}$. The trigonometric model has a state vector which, for s even, is given by $(\gamma_{1t}, \gamma_{1t}^*, ..., \gamma_{(s/2)t})'$. The transition matrix has a block diagonal structure while the covariance matrix of the disturbances is diagonal.

3 Basic structural model

3.1 Stationary form

A stochastic trend component is made stationary by applying the first difference operator twice, that is $\Delta^2\mu_t = \Delta\eta_t + \zeta_{t-1}$. The seasonal component is made stationary by the seasonal summation operator. Thus the stationary form of the BSM, (1), is obtained by multiplying through by the first and seasonal difference operators and taking note of the identity $\Delta_s = \Delta.S\left(L\right)$. With the single shock dummy variable seasonal component $S\left(L\right)\gamma_t = \omega_t$ and so

$$\Delta\Delta_s y_t = \Delta_s\eta_t + S\left(L\right)\zeta_{t-1} + \Delta^2\omega_t + \Delta\Delta_s\varepsilon_t \tag{21}$$

The right hand side is an $MA(s+1)$ process since it is the sum of four MA processes, the maximum order of which is $s+1$. The BSM with a trigonometric seasonal component is also such that $\Delta\Delta_s y_t$ is $MA(s+1)$; for quarterly data

$$\begin{aligned}\Delta\Delta_4 y_t &= \Delta_4\eta_t + S(L)\zeta_{t-1} + \Delta^2\left(1+L\right)\left(\omega_{1t} + \omega_{1t}^*\right) \\ &\quad + \Delta^2\left(1+L^2\right)\omega_{2t} + \Delta\Delta_4\varepsilon_t\end{aligned} \tag{22}$$

When the seasonal component depends on only one variance parameter, the unrestricted reduced form contains more parameters than the structural form, namely $s+2$ as opposed to four.

When the slope is deterministic, equal to β in all time periods, the observations are rendered stationary by the seasonal difference operator. For a quarterly trigonometric seasonal model

$$\Delta_4 y_t = \beta + S(L)\eta_t + \Delta\left(1+L\right)\left(\omega_{1t}+\omega_{1t}^*\right) + \Delta\left(1+L^2\right)\omega_{2t} + \Delta_4\varepsilon_t. \quad (23)$$

The autocovariance function of the quarterly BSM with a trigonometric seasonal component can be derived from the stationary form, (22), and is

$$\left.\begin{array}{ll}
\gamma(0) = 2\sigma_\eta^2 + 4\sigma_\zeta^2 + 8\sigma_1^2 + 14\sigma_2^2 + 4\sigma_\varepsilon^2 & \\
\gamma(1) = 3\sigma_\zeta^2 - 2\sigma_1^2 - 12\sigma_2^2 - 2\sigma_\varepsilon^2 & \\
\gamma(2) = 2\sigma_\zeta^2 - 4\sigma_1^2 + 8\sigma_2^2 & \\
\gamma(3) = \sigma_\zeta^2 + 2\sigma_1^2 - 4\sigma_2^2 + \sigma_\varepsilon^2 & \tau = 3, ..., s-2 \\
\gamma(4) = -\sigma_\eta^2 + \sigma_2^2 - 2\sigma_\varepsilon^2 & \\
\gamma(5) = \sigma_\varepsilon^2 & \\
\gamma(\tau) = 0, & \tau \geqslant 6
\end{array}\right\}$$

Imposing the constraint that $\sigma_2^2 = \sigma_1^2/2$ leads to some simplification and establishes that $\gamma(2)$ and $\gamma(3)$ are both non-negative. When the slope is deterministic, the ACF may be similarly obtained from (23).

The corresponding ACFs with single shock seasonal dummies are more easily obtained and are left as an exercise for the reader.

3.2 Seasonal ARIMA models

For modelling seasonal data, Box and Jenkins (1976, ch. 9) proposed a class of multiplicative seasonal ARIMA models. The most important model within this class has subsequently become known as the 'airline model' since it was originally fitted to a monthly series on UK airline passenger totals. The model is written as

$$\Delta\Delta_s y_t = \left(1+\theta L\right)\left(1+\Theta L^s\right)\xi_t \quad (24)$$

where $\Delta_s = 1 - L^s$ is the seasonal difference operator and θ and Θ are MA parameters which, if the model is to be invertible, must have modulus less than one. Box and Jenkins (1976, pp. 305-6) gave a rationale for the airline model in terms of EWMAs at monthly and yearly intervals.

Maravall (1985), compares the autocorrelation functions of $\Delta\Delta_s y_t$ for the BSM and airline model for some typical values of the parameters and finds them to be quite similar, particularly when the seasonal MA parameter, Θ, is close to minus one. In fact in the limiting case when Θ is equal to minus one, the airline model is equivalent to a BSM in which σ_ζ^2 and σ_ω^2 are both zero. It is straightforward to see that this is true from the stationary form of the BSM given in (21): removing the terms in ω_t and ζ_t gives

$$\Delta\Delta_s y_t = \Delta_s \eta_t + \Delta\Delta_s \varepsilon_t = (1 - L^s)(\eta_t + \varepsilon_t - \varepsilon_{t-1})$$

and the last term in parentheses is an MA(1). The airline model thus provides a good approximation to the reduced form when the slope and seasonal are close to being deterministic. If this is not the case the implicit link between the variability of the slope and that of the seasonal component may be limiting.

The plausibility of other multiplicative seasonal ARIMA models can, to a certain extent, be judged according to whether they allow a canonical decomposition into trend and seasonal components; see Hillmer and Tiao (1982). Although a number of models fall into this category the case for using them is unconvincing. It is hardly surprising that most procedures for ARIMA model-based seasonal adjustment are based on the airline model.

Pure AR models can be very poor at dealing with seasonality since seasonal patterns typically change rather slowly and this may necessitate the use of long seasonal lags. A slowly changing seasonal pattern shows up in the airline model when Θ is close to minus one. Note, though, that it is possible to combine an autoregression with a stochastic seasonal component as in Harvey and Scott (1994).

Consumption A model for aggregate consumption provides a nice illustration of the way in which a simple parsimonious STM that satisfies economic considerations can be constructed. Using UK data from 1957q3 to 1992q2, Harvey and Scott (1994) show that a special case of the BSM consisting of a random walk plus drift, β, and a stochastic seasonal not only fits the data but yields a seasonal martingale difference that does little violence to the forward-looking theory of consumption. The unsatisfactory nature of an autoregression is illustrated in the paper by Osborn and Smith (1989) where sixteen lags are required to model seasonal differences. As regards ARIMA models, Osborn and Smith (1989) select a special case of the airline model in which $\theta = 0$. This contrasts with the reduced form for the structural model which has $\Delta_s c_t$ following an $MA(s-1)$ process (with non-zero mean). The seasonal ARIMA model approximates the sample ACF but does not yield forecasts satisfying a seasonal martingale, that is $E[\Delta_s c_{t+s}] = s\beta$.

4 Trading day and calendar effects

It is not unusual for the level of a monthly time series to be influenced by calendar effects. Such effects arise because of changes in the level of activity resulting from variations in the composition of the calendar between years. The two main sources of calendar effects are trading day variation and moving festivals. They may both be introduced into a time series model and estimated along with the other components in the model. Thus, for example, the BSM is extended so as to become

$$y_t = \mu_t + \gamma_t + \tau_t + \varphi_t + \varepsilon_t \tag{25}$$

where τ_t is the trading day variation component and φ_t is the moving festival component.

Calendar effects should be modelled so as not to affect the level of the trend. Thus when the forecast function is constructed, they should cancel out under temporal aggregation in the same way as the seasonal component. Furthermore, because they represent what are basically artificial movements in the series, there is a clear case for removing them as part of the process of seasonal adjustment.

4.1 Trading day variation

Trading day variation occurs when the activity of an industry or business varies with the day of the week. Thus for a flow variable, or a time-averaged stock, the observation recorded for a particular month will depend on which days of the week occur five times. Accounting and reporting practices can also create trading day effects in a time series. For example, businesses that perform their bookkeeping on Fridays tend to report higher sales in months with five Fridays than in months with four. Time series other than economic time series may also exhibit analogous effects to trading day variation. For example, road accidents tend to be higher on Fridays and Saturdays.

The trading day component is

$$\tau_t = \sum_{j=1}^{7} \theta_{jt} n_{jt}, \qquad t = 1, ..., T, \tag{26}$$

where $n_{jt}, j = 1, ..., 7$ is the number of times day j occurs in month t and θ_{jt} is an unknown parameter associated with it. The constraint that

$$\sum_{j=1}^{7} \theta_{jt} = 0 \tag{27}$$

ensures the trading day effects are not confounded with the trend. If the $\theta'_{jt}s$ are deterministic, the sum of the trading day effects for each month over a period equal to a whole number of weeks is zero since

$$\underset{t}{\Sigma}\tau_t = \underset{t}{\Sigma}\underset{j}{\Sigma}\theta_j n_{jt} = \sum_j \theta_j \sum_t n_{jt}$$

and $\sum_t n_{jt}$ is the same for all j. Over a year the sum of trading day effects will be almost, but not exactly, equal to zero, as 52 weeks is equal to 364 rather than 365 or 366 days.

When the trading day effects are stochastic, the θ_{jt} evolve as random walks as in the balanced dummy variable seasonal model and the constraint in (27) is imposed by specifying the covariance matrix as in (9). One element may be dropped from the state vector. Indeed with deterministic trading day effects this is the normal way to proceed. If the seventh day is dropped, then

$$\tau_t = \sum_{j=1}^{6} \theta_{jt} \left(n_{jt} - n_{7t} \right) \tag{28}$$

The trading day model may be derived from the daily effects model of (15). If different days give rise to the same effect, a more parsimonious trading day model is obtained, namely

$$\tau_t = \sum_{j=1}^{w} \theta_{jt} n_{jt}, \qquad t = 1, ..., T, \tag{29}$$

where there are w different types of day and the covariance matrix of the disturbances driving the $\theta'_{jt}s$ is as in (18). As before one parameter may be removed. Summing the daily effects over one month, and noting the constraint on θ_{wt} in (20) yields

$$\tau_t = \sum_{j=1}^{w-1} \theta_{jt} n_{jt} + n_{wt} \left(\frac{-1}{k_w} \sum_{j=1}^{w-1} k_j \theta_{jt} \right) = \sum_{j=1}^{w-1} \theta_{jt} \left(n_{jt} - \frac{k_j n_{wt}}{k_w} \right)$$

Thus, for example, if all weekdays are the same, and Saturdays and Sundays are the same, τ_t contains a single variable, that is

$$\tau_t = \theta_{1t} \{ n_{1t} - n_{2t} \left(5/2 \right) \}, \tag{30}$$

where n_{1t} is the number of weekdays in the month and n_{2t} is the number of Saturdays and Sundays. The deterministic trading day model takes $n_{jt} - (k_j/k_w) n_{wt}, j = 1, ..., w-1$, as explanatory variables.

The deterministic version of (28) is used by Kitagawa and Gersch (1984) and Bell and Hillmer (1983), although in the second of these references the remaining part of the model is of the ARIMA form. Bell and Hillmer (1983) also include the variable

$$n_t = \sum_{j=1}^{7} n_{jt} \tag{31}$$

This variable, which is the total number of days in month t, is able to account for effects due to leap year Februaries. However, if there are no leap year effects, or if the February figure is adjusted prior to any model building, n_t becomes superfluous. An implementation of a stochastic trading day effects model can be found in Dagum, Quenneville and Sutradhar (1992).

The trading day model is not satisfactory in all respects. For example, if a particular activity is only carried out on weekdays, it would seem sensible to divide the monthly total by the number of working days before fitting a model. (It may also be multiplied by the average number of working days per month, although this is not important in the present context.) The result will not be the same as using the trading day models given above. Suppose that we amend (30) somewhat, so that it becomes $\tau_t = \theta n_{1t}$. If the structural model is in levels, there is no way that defining τ_t in this way can be equivalent to dividing y_t by n_1. This is the case even if y_t is in logarithms, because there is no value of θ that makes $\exp(\theta n_{1t})$ equal to unity for all n_{1l}.

Another problem is the effect of public holidays on modelling trading day variation. A public holiday could be treated as an eighth day of the week and given its own effect, or it could be treated as a Sunday. There may, however, be other effects arising from public holidays. For example, if a holiday falls on a Friday, the usual Friday effect may be transferred to a Thursday.

4.2 Moving festivals

The month in which certain holidays and religious festivals fall can vary from year to year. A prime example is Easter, which can fall anywhere from 22 March to 25 April. In connection with retail sales, Bell and Hillmer (1983) suggest modelling Easter as

$$\varphi_t = \alpha h_t \tag{32}$$

where h_t is the proportion of the time period H days before Easter that falls in month t. This model can be defined for any positive H and if $H \leqslant 22$ the only months for which h_t will ever be non-zero are March and April. A similar model could be used for road and air traffic except that in this case the time period up to and including Easter Monday might be the relevant one. The value of H would probably be four or five.

As it stands, (32) does not have the property that the φ_t's sum to zero over a year. Fortunately this is easily remedied. In addition, the form of the moving festival component may be generalised. Suppose that φ_t is modelling any moving festival effect, not necessarily Easter, and that h_t is now a weight given to month t. Let the sum of the h_t's over any one year be unity. The pattern of the h_t's depends on the location of the moving festival in question and its postulated effect on the surrounding days. Thus, for example, the weight pattern for a series on road accidents might be derived by assigning initial weights of $\frac{1}{3}$, $\frac{1}{6}$, $\frac{1}{6}$ and $\frac{1}{3}$ to the days from Good Friday to Easter Monday. If Easter Monday were 1 April in a particular year, this would imply that h_t for March would be $\frac{2}{3}$, while for April it would be $\frac{1}{3}$. A moving festival may now be formulated as

$$\varphi_t = \alpha(h_t - 1/s), \tag{33}$$

where s is twelve unless the timing interval is lunar months, in which case it is thirteen.

Further generalisation is possible. The parameter α may be allowed to change over time by modelling it as a random walk. The weight function, h_t, then appears in the corresponding position in the $\mathbf{z}_t$ vector in the measurement equation.

A final issue, which is relevant to both trading day variation and moving festivals, concerns model selection. If there is reason to suspect calendar effects are present, they should be included in the model at the outset. If a model is fitted without calendar effects, then significant Box-Ljung or other serial correlation statistics may be an indication that seasonal effects are present. Specific tests for calendar effects are discussed in sub-section 5.4 below.

As with outliers, calendar effects can distort the correlogram of the original (differenced) series. This is illustrated by Hillmer (1982, p. 388) using a series on the monthly outward station movements (disconnections) of the Wisconsin telephone company from January 1951 to October 1966. An earlier analysis of ths series, using standard ARIMA methodology, had been carried out by Thompson and Tiao (1971). They had obtained the model

$$\left(1 - 0.49L^3\right)\left(1 - 1.005L^{12}\right) y_t = \left(1 - 0.23L^9 - 0.33L^{12} - 0.17L^{13}\right) \xi_t$$

After allowing for trading day variation, Hillmer selected the airline model.

5 Tests

5.1 Seasonal stationarity tests

The basic form of the LBI test against nonstationary stochastic seasonality is obtained for the model

$$y_t = \mu + \gamma_t + \varepsilon_t, \qquad t = 1, ..., T \tag{34}$$

where μ is a constant. The test against the presence of a stochastic trigonometric component at any one of the seasonal frequencies, λ_j, apart from the one at π, is based on the statistic

$$\omega_j = 2T^{-2}\widehat{\sigma}^{-2}\sum_{t=1}^{T}\left[\left(\sum_{i=1}^{t} e_i \cos\lambda_j i\right)^2 + \left(\sum_{i=1}^{t} e_i \sin\lambda_j i\right)^2\right], \qquad j = 1, ..., [(s-1)/2], \tag{35}$$

where $\widehat{\sigma}^2$ is the sample variance of the OLS residuals, $e_t, t = 1, .., T$, from a regression on the seasonal sines and cosines, $\mathbf{z}_t$, and a constant. Following Canova and Hansen (1995), it can be shown that, under the null hypothesis, the asymptotic distribution of this statistic is generalized Cramér-von Mises with two degrees of freedom. If s is even, the statistic at frequency π is

$$\omega_{s/2} = T^{-2}\widehat{\sigma}^{-2}\sum_{t=1}^{T}\left(\sum_{i=1}^{t} e_i(-1)^i\right)^2,$$

and this has an asymptotic distribution which is Cramér-von Mises with one degree of freedom. A joint test against the presence of stochastic trigonometric components at all seasonal frequencies is obtained by summing the individual test statistics, that is

$$\omega = \sum_{j=1}^{[s/2]}\omega_j \tag{36}$$

This statistic has an asymptotic distribution which is generalized Cramér-von Mises with $s-1$ degrees of freedom, denoted $CvM(s-1)$. Canova and Hansen (1995) point out that the same test is obtained if the stochastic seasonal component is of the balanced dummy variable form.

Canova and Hansen show how the above tests can be generalized to handle serial correlation and heteroscedasticity by making a nonparametric correction. However, if the process generating the non-seasonal part of the model is taken as given, the LBI test against stochastic seasonality is constructed from a set of 'smoothing errors'. As shown in Busetti and Harvey (2003, Appendix B) the smoothing errors are, in general, serially correlated but the form of this serial correlation may be deduced from the specification of the model, thereby allowing the construction of a statistic that has a Cramér-von Mises distribution, asymptotically, under the null hypothesis. An alternative possibility is to use the T standardized one-step ahead prediction errors, the innovations, calculated by treating nonstationary and deterministic components as having fixed initial conditions. No correction is then needed; the statistic is of the form (35) and has the same asymptotic distribution. Calculating innovations under the assumption that the initial conditions are fixed requires that the initial conditions be estimated, but a backward smoothing recursions can be avoided simply by reversing the order of the observations and calculating a set of innovations starting from the filtered estimator of the state at the end of the sample. Actually, the forward and backward innovations are not the same and in neither case do the sums, weighted by $\cos\lambda_j t$ and $\sin\lambda_j t$, equal zero, so statistics formed from forward and backward sums are different. Fortunately the asymptotic properties are unaffected. Smoothing errors do not suffer from these ambiguities.

For both the smoothing error and innovation forms of the test, nuisance parameters will normally have to be estimated. For stationarity tests, Leybourne and McCabe (1994) argue that this is best done under the alternative using maximum likelihood. Proceeding in this way has the compensating advantage that since there will often be some doubt about a suitable model specification, estimation of the unrestricted model affords the opportunity to check its suitability by the usual diagnostics and goodness of fit tests. Once the nuisance parameters have been estimated, the test statistic is calculated from the innovations obtained with σ_ω^2 set to zero.

The parametric test may be applied in models which include a deterministic trend, a random walk with or without a drift, or a trend with a stochastic slope. In all these cases the asymptotic distribution of the test statistics is unaffected.

We will refer to these tests as *seasonal stationarity* tests. The nonparametric statistics will be denoted $\omega(m)$, where m is the number of lags in the estimator of the spectrum.

5.2 Seasonality test

Seasonal stationarity tests take the null to be deterministic seasonality. Sometimes we may wish to test whether there is any seasonality at all. Busetti and Harvey (2003) suggest two possible tests. Both can be implemented parametrically or nonparametrically.

1) A Wald test of the null hypothesis that the deterministic seasonal coefficients are zero.

2) A seasonal stationarity test without fitting seasonal dummies. Such a test will have also have power against deterministic, as well as stochastic, seasonality. If the test statistic, ω_0, is formed without fitting seasonal dummies, its asymptotic distribution under the null will be a function of Brownian motion rather than of a Brownian bridge, that is $CvM_0(s-1)$. The 5% critical value for three degrees of freedom, as is appropriate for a full test on quarterly data, is 3.46.

Spanish interest rates As an example we consider the logarithm of 3-month money market interest rate in Spain for the period 1977Q1-2001Q4; the source is the Bank of International Settlements (BIS) macroeconomic series database. The series is depicted in the upper panel of figure 1. It is difficult to detect a seasonal pattern from a casual glance at the graph and one would not normally expect one to be present in an interest rate series; however the functioning of the interbank loans market may imply some seasonality.

Fitting the BSM to the series gives a seasonal component as shown in the lower panel of figure 1; the slope variance is estimated to be zero and the estimate of the (fixed) slope is small and insignificant. We have used logarithms of the data only because the diagnostics are better; if the raw series is used, the resulting seasonal pattern is similar.

The chi-square statistic for the seasonals at the end of the series is only 0.09 which is clearly not significant as the 5% critical value for a χ_3^2 is 7.81. However the graph shows a fairly strong seasonal pattern until the mid-eighties. The question is whether the pattern as a whole is in any sense significant.

Setting the seasonal variance to zero and re-estimating the BSM gives a Wald statistic of 4.76, with a p-value of 0.19. This is still not significant. If

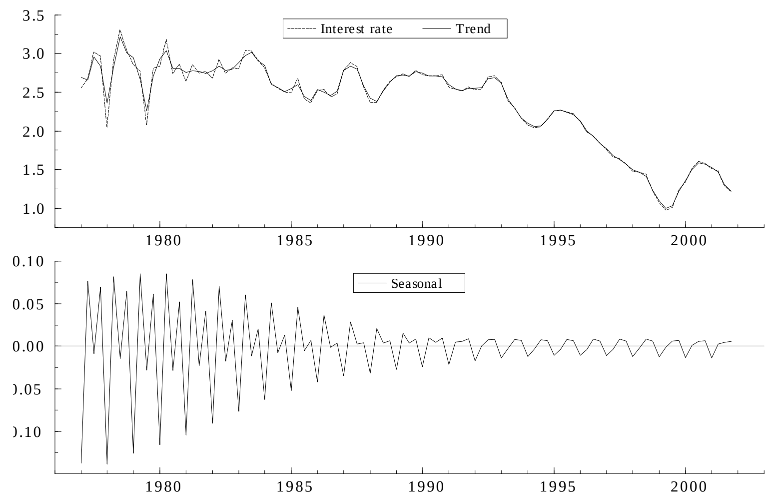

Figure 1: Spanish interest rates

the series is differenced and a nonparametric Wald test is computed using the Newey-West covariance matrix estimator with three lags a similar p-value, 0.17, is obtained. On the other hand, the spectral nonparametric seasonal stationarity test statistic, $\omega_0(m)$, computed using forward summations takes the values 3.83 and 3.01 for $m = 3$ and 6 respectively, rising to 4.64 and 3.89 for $\underline{\underline{\omega}}_0^*(m)$, the preferred form in which the spectrum is estimated after fitting seasonal regressors. As the 5% critical value is 3.46, this test provides a firm rejection of the hypothesis that there is no seasonality in the series.

Finally, for $m = 3$ and 6 the seasonal stationarity test statistic, $\omega(m)$, takes the values 1.17 and 1.02 respectively (against a 5% critical value of 1.00), thereby confirming the presence of stochastic seasonality.

5.3 Seasonal unit root tests

The test of Hylleberg et al. (1990) - HEGY- is testing the null of a nonstationary seasonal against the alternative of a stationary seasonal. Its relationship to the seasonal stationarity test is analogous to that of the relationship between the (augmented) Dickey-Fuller test and KPSS.

The UC seasonal unit root test can be set up by introducing a damping factor into (11) so that each trigonometric term in the seasonal component is modelled by

$$\begin{bmatrix} \gamma_{j,t} \\ \gamma^*_{j,t} \end{bmatrix} = \phi_j \begin{bmatrix} \cos\lambda_j & \sin\lambda_j \\ -\sin\lambda_j & \cos\lambda_j \end{bmatrix} \begin{bmatrix} \gamma_{j,t-1} \\ \gamma^*_{j,t-1} \end{bmatrix} + \begin{bmatrix} \omega_{j,t} \\ \omega^*_{j,t} \end{bmatrix}, \qquad \begin{matrix} j = 1, \ldots, [s/2], \\ t = 1, \ldots, T. \end{matrix} \tag{37}$$

with $\gamma^*_{s/2,t}$ dropping out for s even. The seasonal component, obtained by summing the $\gamma'_{j,t}s$ is then embedded in a general UC model which contains deterministic seasonal trigonometric terms. However, since the forecasts would gradually die down to zero for $\phi_j < 1$, such a seasonal component is not capturing any (non-deterministic) persistent effects of seasonality. In any case the empirical evidence, for example in Canova and Hansen (1995), clearly points to seasonal unit roots as the norm. Nevertheless we may still wish to test the null hypothesis of seasonal unit roots against the alternative of stationary seasonality.

A parametric test of the null hypothesis that the component at a particular frequency is nonstationary against the alternative that it is stationary, that is $H_0 : \phi_j = 1$ against $H_1 : \phi_j < 1$, can be constructed from the null hypothesis innovations as

$$\omega_j = 2T^{-2} \sum_{i=1}^{T} \left[\left(\sum_{t=1}^{i} \widetilde{\nu}_t \cos\lambda_j t \right)^2 + \left(\sum_{t=1}^{i} \widetilde{\nu}_t \sin\lambda_j t \right)^2 \right] < c, \qquad j = 1, ..., [(s-1)/2]. \tag{38}$$

Under the null hypothesis the asymptotic distribution is $CvM_0(2)$ since if the nonstationary seasonal operator, $1 - 2cos\lambda_j L + L^2$, were to be applied it would remove the corresponding deterministic seasonal. For $j = s/2$

$$\omega_{s/2} = T^{-2} \sum_{i=1}^{T} \left(\sum_{t=1}^{i} \widetilde{\nu}_t \cos\pi t \right)^2 = T^{-2} \sum_{i=1}^{T} \left(\sum_{t=1}^{i} (-1)^t \widetilde{\nu}_t \right)^2$$

and this has a $CvM_0(1)$ asymptotic distribution under the null. The full seasonal test statistic is formed by summing the $\omega'_j s$ and its asymptotic distribution under the null is $CvM_0(s-1)$. With seasonal slopes the asymptotic distributions are $CvM_1(.)$; compare Smith and Taylor (1998).

Seasonality tests based on an autoregressive model will tend to perform poorly in situations where an unobserved components model is appropriate. The simulation evidence in Hylleberg (1995) illustrates this point by looking at the results of using the HEGY test for moving average models, which, as Harvey and Scott (1994) note, typically arise as the reduced form of unobserved components models.

A rejection of the null hypothesis in a seasonal unit root test may be an indication of a deterministic seasonal component rather than a stationary seasonal component; see the evidence in Canova and Hansen (1995, p 244). Following the argument in Harvey and Streibel (1998), it can be shown that the appropriate test of the null of deterministic seasonality against the alternative of near-persistent stationary seasonality, that is (37) with the ϕ_j close to one, is, in fact, the seasonal stationarity test. Therefore we may only want to do a test against stationary seasonality if the hypothesis of deterministic seasonality has first been rejected by the seasonal stationarity test.

5.4 Testing for trading day effects

Cleveland and Devlin (1980) showed that peaks at certain frequencies in the estimated spectra of monthly time series indicate the presence of trading day effects. Specifically there is a peak at a frequency of $0.348 \times 2\pi$ radians, with the possibility of subsidiary peaks at $0.432 \times 2\pi$ and $0.304 \times 2\pi$ radians. An option in the output of the X-12-ARIMA program provides a comparison of the estimates of these frequencies with the adjacent frequencies; see Soukup and Findley (2000). However, there is no formal test. Busetti and Harvey (2003) suggest a seasonality test at the relevant frequency or a joint test at all three frequencies. Assuming that no (deterministic) trading day model has been fitted, the asymptotic distribution is CvM_0, as in sub-section 5.2, with the 5% critical value being 2.63 for a test at a single frequency and 5.68 for a test at all three frequencies.

As an example we took the irregular component, obtained from X12-ARIMA, of series s0b56ym, *U.S. Retail Sales of Children's, Family, and Miscellaneous Apparel,* as supplied by the Bureau of the Census. Since the process followed by this irregular component cannot be derived, it was decided to use the nonparametric test. The $\omega(10)$ test statistic for the single main frequency was 7.03. For all three frequencies it was 8.21. Both give a clear rejection of the null hypothesis that there is no trading day effect.

6 Seasonal adjustment

Once a STM has been fitted, the seasonally adjusted series is obtained by signal extraction using the Kalman filter and associated smoother (KFS). Efficient algorithms are described in Durbin and Koopman (2001, pp. 70-73). The KFS automatically adjusts the weighting pattern to give optimal (MMSE) estimates at the ends of the series. The KFS is much easier to implement than the Wiener-Kolmogorov (WK) filter and is more general[1] - for example it can be used with models that are not time invariant. Unlike WK, the weights are implicit, but they can be calculated by the algorithm of Koopman and Harvey (2003). Figure 2 shows weights for the model fitted to the logarithms of gas consumption by 'Other final users in the UK' as displayed by STAMP 7. The weights for the seasonally adjusted series are $1 - w_s(L)$ where $w_s(L)$ is the polynomial of weights for extracting the seasonal. The gain shows the effect of the filter on a stationary series. The gain of the seasonal adjustment filter is one minus the gain of the seasonal filter. This is zero at π and $\pi/2$ in order to remove the non-stationary stochastic seasonal component: the pseudo-spectrum of the seasonal at those frequencies is infinity.

Figure 3 shows the weights at the end of the series. In this case there is a phase shift.

7 Breaks in the seasonal pattern

The seasonal pattern sometimes changes as the result of an intervention. Modelling an effect of this kind requires the introduction of $s - 1$ dummy variables into the measurement equation, starting at time τ. These dummies are constrained to sum to zero over s consecutive time periods. alternatively pulse dummies can be added to the part of the state vector associated with the seasonal, that is $\boldsymbol{\gamma}_t$. As an example, figure 4 shows the number of marriages in the UK every quarter. Estimating (1) with a random walk trend

[1] In the second edition of his celebrated text describing the WK filter, Whittle (1984, p xi) writes 'In its preoccupation with the stationary case and generating function methods, the 1963 text essentially missed the fruitful concept of state structure. This ..... has now come to dominate the subject.' Nevertheless WK formulae can give interesting theoretical insights into the weighting structure of smoothed estimators and as such provide the basis for excellent examination questions.

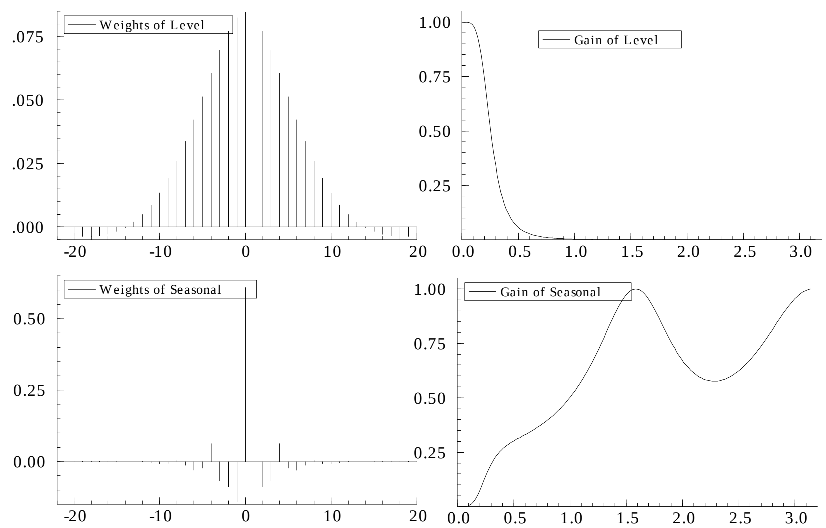

Figure 2: Weights for quarterly gas series in the middle of the sample.

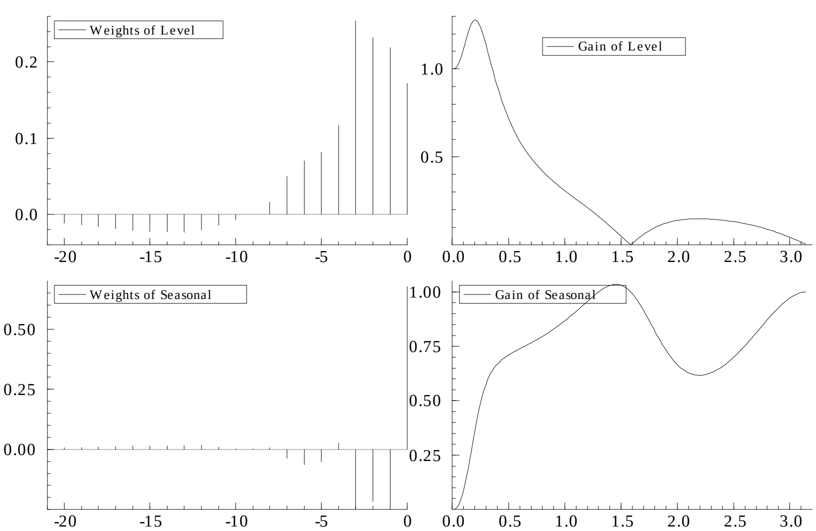

Figure 3: Weights for quarterly gas series at the end of the sample.

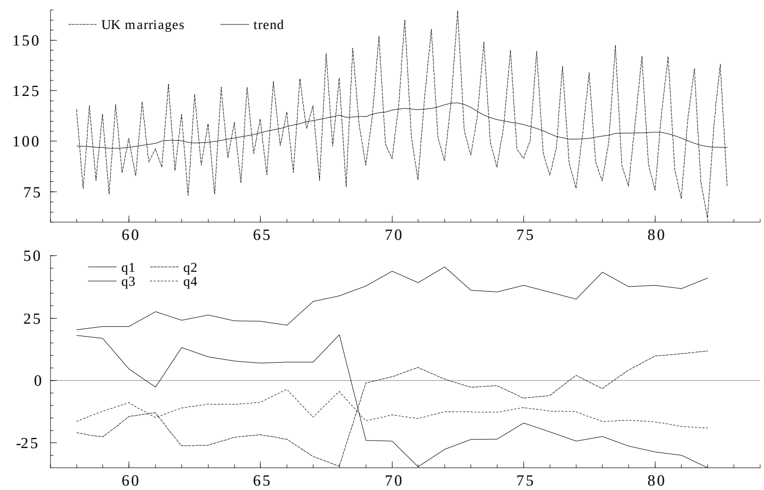

Figure 4: UK marriages and plot of seasonal for each quarter

using the STAMP program gives

$$\widetilde{\sigma}_{\varepsilon} = 0.00 \qquad \widetilde{\sigma}_{\eta} = 1.61 \qquad \widetilde{\sigma}_{\omega} = 2.69$$

with an equation standard error (the standard deviation of the innovations), $\widetilde{\sigma}$, of 7.91. The lower panel in figure 4, the plot of individual seasons, displays a dramatic switch between the first and second quarters starting in 1969. Indeed the parametric seasonal stationarity test statistic, constructed from the Kalman filter innovations, is 6.96 which is a very decisive rejection of the null hypothesis of a constant seasonal pattern. The reason is that there was a change in the tax law. Up to the end of 1968 couples were allowed to claim the married persons tax allowance retrospectively for the entire year in which they married. As the tax year begins in April this arrangement provided an incentive to marry in the first quarter of the calendar year, rather than in the spring. The abolition of this rule led to a marked decrease in the number of weddings in quarter one and a compensating rise in quarter two.

Adding a set of three seasonal break dummy variables[2], starting in the first quarter of 1969, to take account of a complete change in the seasonal

[2] Actually it is only necessary to have a single switch dummy for q1 and q2 to capture the change.

pattern leads to the following estimates of the parameters:

$$\widetilde{\sigma}_{\varepsilon} = 2.42 \qquad \widetilde{\sigma}_{\eta} = 1.59 \qquad \widetilde{\sigma}_{\omega} = 1.36$$

with

$$Q(9,7) = 12.54 \qquad and \qquad \widetilde{\sigma} = 5.66$$

where $Q(P, f)$ is the Box-Ljung statistic based on P residual autocorrelations but with f degrees of freedom. The $t-$statistics for the seasonal break dummies are -8.33, 7.58 and 2.09 respectively. There is a big reduction in the estimate of the seasonal parameter, σ_{ω}, which no longer needs to be such as to allow the stochastic seasonal model to accommodate the change, and the equation standard error, $\widetilde{\sigma}$, has fallen considerably.

When a full set of seasonal break dummy variables is included in the model the distibution of the seasonal stationarity test statistic is $CvM(2s-2)$. The parametric test statistic, calculated from the Kalman filter innovations, is 2.42, giving a strong indication that there is still stochastic seasonality present. This is backed up by the fact that estimating the model with a fixed seasonal gives a significant Box-Ljung statistic of $Q(9,8) = 22.38$ while the fourth order residual autocorrelation, $r(4)$, is 0.33.

If the breakpoint is unknown Busetti and Harvey (2003) show that running the seasonal break tests with an estimated breakpoint leads to an asymptotically valid procedure.

8 Time-varying periodic splines

With quarterly or monthly observations it is rarely necessary or desirable to cut down on the number of terms used to model the seasonal pattern. On the other hand, when s is large, as for example with weekly data, there are both statistical and computational reasons for wanting to constuct a more parsimonious seasonal model.

The dimensions of the dummy variable seasonal model can be reduced by assuming that certain seasons are the same, as in the daily model of sub-section 2.4. However, while this may be suitable for daily observations, it is generally inappropriate for a pattern within the year. Similarly, the trigonometric seasonal model can be straightforwardly cut down by excluding pairs of sines and cosines at certain frequencies, usually the higher ones, but while this may be a sensible option for a slowly changing seasonal effect, such as temperature, it may not be satisfactory for modelling economic variables, where there are often sharp peaks at certain times of the year. What is needed is a model that is parsimonious but flexible enough to capture marked variations in a periodic pattern while retaining a reasonable degree of continuity. This leads to the notion of periodic time-varying splines. These are described, with applications to weekly and intra-daily observations, in Harvey, Koopman and Riani (1996) and Harvey and Koopman (1993).

The first step in obtaining a time-varying periodic spline is to show how spline models in general, and periodic splines in particular, can be set up in a regression. Following Harvey and Koopman (1993), time variation is introduced in the usual way by allowing the parameters to follow stochastic processes.

Piecewise regression using cubic splines Suppose there are n pairs of observations (x_j, y_j), $j = 1, ..., n$, and that we wish to set up a nonlinear regression model of the form

$$y_j = f(x_j) + \varepsilon_j, \qquad j = 1, ..., n, \tag{39}$$

where the ε_j's are mutually uncorrelated disturbances with zero mean and constant variance, σ^2. In a cubic spline regression model, $f(x_j)$ is constructed by putting together polynomials of degree at most three in such a way as to preserve continuity in second derivatives. The h individual cubics are joined at the co-ordinates $\left(x_i^{\dagger}, \gamma_i^{\dagger}\right)$, $i = 0, ..., h$. The set of x values, $x_0^{\dagger} < x_1^{\dagger} < \cdots < x_h^{\dagger}$, is known as a *mesh*; the $h + 1 \geq 3$ individual points are called *knots*. The setup is completed by making assumptions about the spline at its end points.

Given the knots and the associated values of the ordinates, $x_0^{\dagger}, ..., x_h^{\dagger}$, it can be shown that any point on the spline function is a linear combination of the $\gamma_i^{\dagger}$'s. Thus at the observation points, we can write

$$f(x_j) = \mathbf{w}_j' \boldsymbol{\gamma}^{\dagger}, \qquad j = 1, ..., n, \tag{40}$$

where $\mathbf{w}_j$ is an $(h+1) \times 1$ vector that depends on the position of the knots and the distance between them, as well as on the observed value x_j, and $\boldsymbol{\gamma}^{\dagger} = \left(\gamma_0^{\dagger}, \gamma_1^{\dagger}, ..., \gamma_n^{\dagger}\right)'$. If x_j corresponds to a knot, $x_j = x_i^{\dagger}$, then all the elements in $\mathbf{w}_j$ are zero, apart from the ith which is unity, and $f(x_j) = \gamma_i^{\dagger}$.

Substituting (40) in (39) gives the cubic spline regression model

$$y_j = \mathbf{w}_j' \boldsymbol{\gamma}^{\dagger} + \varepsilon_j, \qquad j = 1, ..., n \tag{41}$$

Given the assumptions on $\varepsilon_j, \boldsymbol{\gamma}^\dagger$ is estimated by ordinary least squares.

Periodic splines Now suppose that the explanatory variable is time and that there is a pattern repeated over a stretch of s observations, so that we have s periodic effects $\gamma_j, j = 1, ..., s$. These effects can be modelled by a spline of the form (40), in which $n = s, x_j = j$ and continuity from one period to the next is preserved by the condition that

$$\gamma_0^\dagger = \gamma_h^\dagger \tag{42}$$

together with the conditions that the first and second derivatives at 0 and h are the same. This removes the need for further assumptions about the end conditions. The implications for the $\mathbf{w}_j$ vectors, which are now $h \times 1$, corresponding to $\gamma_1^\dagger, ..., \gamma_h^\dagger$, are easily worked out. Full details can be found in Poirier (1976, pp. 43-47) and Harvey, Koopman and Riani (1997, appendix). The periodic spline is therefore

$$\gamma_j = \mathbf{w}_j' \boldsymbol{\gamma}^\dagger, \qquad j = 1, ..., s, \tag{43}$$

where $\boldsymbol{\gamma}^\dagger$ is $h \times 1$.

As with any seasonal component, the periodic effects should sum to zero over a complete period so as not to be confounded with the trend. Thus

$$\sum_{j=1}^{s} \gamma_j = \sum_{j=1}^{s} \mathbf{w}_j' \boldsymbol{\gamma}^\dagger = \mathbf{w}' \boldsymbol{\gamma}^\dagger = 0, \tag{44}$$

where $\mathbf{w}$ is the $h \times 1$ vector

$$\mathbf{w} = \sum_{j=1}^{s} \mathbf{w}_j. \tag{45}$$

The restriction can be enforced by arbitrarily dropping one of the elements of $\boldsymbol{\gamma}^\dagger$. If $\gamma_h^\dagger$ is dropped, then substituting

$$\gamma_h^\dagger = -\sum_{i=1}^{h-1} (w_i / w_h)\, \gamma_i^\dagger, \tag{46}$$

where w_i is the $i-$th element of $\mathbf{w}$, in (43) gives

$$\gamma_j = \sum_{i=1}^{h-1} (w_{ji} - w_{jh} w_i / w_h)\, \gamma_i^\dagger, \qquad j = 1, ..., s, \tag{47}$$

where w_{ji} is the ith element of the vector $\mathbf{w}_j$.

Time-varying effects When the periodic effects evolve over time, (43) becomes

$$\gamma_t = \gamma_{jt} = \mathbf{w}_j'\boldsymbol{\gamma}_t^{\dagger}, \qquad j = 1, ..., s, \qquad t = 1, ..., T, \tag{2.15}$$

where γ_{jt} is the jth periodic effect at time t and

$$\boldsymbol{\gamma}_t^{\dagger} = \boldsymbol{\gamma}_{t-1}^{\dagger} + \boldsymbol{\omega}_t^{\dagger}, \tag{2.16}$$

where $\boldsymbol{\omega}_t^{\dagger}$ is an $h \times 1$ vector of serially uncorrelated random disturbances each with mean zero. The zero-sum constraint, $\mathbf{w}'\boldsymbol{\gamma}_t^{\dagger} = 0$, is as in (44), and the implied constraint on the disturbances, $\mathbf{w}'\boldsymbol{\omega}_t^{\dagger} = 0$, is reflected in the covariance matrix

$$E\left(\boldsymbol{\omega}_t^{\dagger}\boldsymbol{\omega}_t^{\dagger\prime}\right) = \sigma_\omega^2 \left[\mathbf{I} - \frac{1}{\mathbf{w}'\mathbf{w}}\mathbf{w}\mathbf{w}'\right] \tag{48}$$

because $Var(\mathbf{w}'\boldsymbol{\omega}_t^{\dagger}) = 0$.

The above covariance matrix is of the same form as (18) in the daily model. As in that model, an element can be dropped from $\boldsymbol{\gamma}_t^{\dagger}$. If the $h-th$ is dropped, γ_t can be obtained from the first $h-1$ elements of the vector $\gamma_t^{\dagger}$ and the state space form is set up with a component of the form (47) featuring in the measurement equation. However, as was pointed out in sub-section 2.6, the constraint can be enforced without dropping an element by combining (48) with a corresponding diffuse initial covariance matrix.

9 Robust seasonal adjustment

Simulation techniques of the kind of the kind described in Durbin and Koopman (2001) are relatively easy to use when the measurement and transition equations are linear but the disturbances are non-Gaussian. Allowing the disturbances to have heavy-tailed distributions provides a robust method of dealing with outliers and structural breaks. While outliers and breaks can be dealt with *ex post* by dummy variables, only a robust model offers a viable solution to coping with them in the future.

Allowing ε_t to have a heavy-tailed distribution, such as Student's t, provides a robust method of dealing with outliers. An outlier is defined as an observation that is inconsistent with the model. By employing a heavy-tailed distribution, such observations are consistent with the model whereas with a Gaussian distribution they would not be. Treating an outlier as though it were a missing observation effectively says that it contains no useful information. This is rarely the case except, perhaps, when an observation has been recorded incorrectly.

Gas consumption in the UK Estimating a Gaussian BSM for gas consumption produces a rather unappealing wobble in the seasonal component at the time North Sea gas was introduced in 1970. Durbin and Koopman (2001, p 233-5) allow the irregular to follow a t-distribution and estimate its degrees of freedom to be 13. The robust treatment of the atypical observations in 1970 produces a more satisfactory seasonal pattern around that time.

Another example of the application of robust methods is the seasonal adjustment paper of Bruce and Jurke (1996).

In small samples it may prove difficult to estimate the degrees of freedom of a t-distribution. A reasonable solution then is to impose a value, such as six, that is able to handle outliers. Other heavy tailed distributions may also be used; Durbin and Koopman (2001, p 184) suggest mixtures of normals and the general error distribution. De Rossi and Harvey (2006) effectively assume a double exponential distribution in their algorithm for estimating a time-varying median.

Heavy tailed distributions may also be used to provide models that are robust to breaks in the trend or seasonal.

10 Seasonal specific models

Periodic models were originally introduced to deal with certain problems in environmental science, such as modelling river flows; see Hipel and McLeod (1994, ch. 14). The key feature of such models is that separate stationary AR or ARMA model are constructed for each season. Econometricians have developed periodic models further to allow for nonstationarity within each season and constraints across the parameters in different seasons; see the monograph by Franses and Papp (2004). These approaches are very much within the autoregressive/ARIMA paradigm. The structural framework offers a more general way of capturing periodic features by allowing periodic components to be combined with components common to all seasons. These common components may exhibit *seasonal heteroscedasticity*, that is have different values for the parameters in different seasons. Such models have a clear interpretation and make explicit the distinction between an evolving seasonal pattern of the kind typically used in a structural time series model and genuine periodic effects.

The first sub-section below introduces seasonal heteroscedasticity and we then move on to define periodic models. The relationship between nonstationary periodic models and STMs is then examined and it is concluded that the large number of parameters in periodic models means that there have to be strong reasons for wanting to use them. The third sub-section introduces *partly periodic* models, in which some or all of the components are periodic with respect only to groups of seasons. These may be handled within a state space framework and generalised further to include seasonal heteroscedasticity. Overall we have a class of what might be called *seasonal specific* models.

10.1 Seasonal heteroscedasticity

If the variance hyperparameters in a STM are different in different seasons, the model is said to exhibit *seasonal heteroscedasticity (SH)***.** Such models are not time-invariant but this poses no problem if they are handled using the SSF.

Irregular - The simplest example of SH is when the irregular component has a different variance in each season, that is

$$Var\left(\varepsilon_t^{(j)}\right) = \sigma^2_{\varepsilon,j}, \quad t = 1, ..., T, \qquad j = 1, .., s \tag{49}$$

Proietti (1998) gives an example involving monthly water usage. The trend is modelled as a random walk plus drift while the seasonal ends up being deterministic even though allowance was made for it being stochastic. The seasonal heteroscedasticity is not part of the seasonal component since it reflects the transitory effects which are more volatile at certain times of the year. For example, unusually hot weather in the summer can give rise to much higher water consumption than would normally be expected.

Another example of seasonal heteroscedasticity is the level of activity in the construction industry in countries such as the USA where the occasional severe winter leads to a much higher variance in the winter quarter.

Trend - It is possible that there is more scope for a permanent change in a series at certain times of the year. Effects of this kind could be allowed for by letting the level and slope variances be seasonal specific. They do not belong in the seasonal component since they do not give rise to a seasonal pattern in the forecast function.

Seasonal - The seasonal component itself may be subject to seasonal specific effects insofar as the way in which it evolves depends on the time of year. The time series of Italian industrial production provides a good example. The series is typically very low in August due to holidays and it is also very variable. One possibility is to increase the variance of the irregular component in August. However, Proietti (1998) argues that the variability arises because the seasonal effect associated with August changes at a faster rate than the seasonal effect associated with other months. In order to capture this kind of phenomenon, he generalizes the (balanced) dummy variable seasonal model. The covariance matrix of the disturbances becomes

$$Var\left(\boldsymbol{\omega}_t\right) = \mathbf{W} - \frac{1}{\mathbf{i}_s'\mathbf{W}\mathbf{i}_s}\mathbf{W}\mathbf{i}_s\mathbf{i}_s'\mathbf{W} \tag{50}$$

where $\mathbf{W}$ is a diagonal matrix with $\sigma^2_{\omega,j}, j = 1, .., s$, as the $j-$th diagonal term. (Note that at least two variances must be non-zero if $\mathbf{W}$ is to be non-null.) Since $Var\left(\mathbf{i}_s'\boldsymbol{\omega}_t\right) = 0$, the seasonals sum to zero over a year. The above covariance matrix reduces to that in (9) if $\sigma^2_{\omega,j} = \sigma^2_{\omega}$ for all $j = 1, .., s$. In the case of Italian industrial production, Proietti sets all $\sigma^2_{\omega,j}$'s the same except for August, $\sigma^2_{\omega,8}$.

Cycle - cyclical component may exhibit different properties in different seasons. Effects of this kind may be captured in a stochastic cycle, not only by having the disturbance variances differ across seasons but also by letting the frequency,λ_c, and damping factor, ρ, be season specific.

10.2 Periodic models

A classic example of the need for a periodic model arises in hydrology when there are monthly observations on river flow. If melting snow is an important factor in river flow in the spring, the correlation between the flow in successive months may be negative whereas at other times of the year it is positive. These features show up in the periodic autocorrelation functions, as illustrated in Hipel and McLeod (1994, p504-5). The autocorrelation function for period (season) j is

$$\rho_j(\tau) = Cov(y_t^{(j)}, y_{t-\tau}) / \sqrt{Var(y_t^{(j)}) Var(y_{t-\tau})}, \quad j = 1, .., s,$$

where the superscript in $y_t^{(j)}$ indicates that the time t is in period j.

The periodic ARMA model can be written as:

$$y_t = \mu_j + u_t^{(j)}, \quad t = 1, ...T, \quad j = 1, .., s, \tag{51}$$

where the mean in season j, μ_j, is constant over different years and $u_t^{(j)}$ follows a stationary ARMA (p_j, q_j) process, that is

$$u_t^{(j)} = \phi_{j,1} u_{t-1} + .. + \phi_{j,p} u_{t-p} + \xi_t^{(j)} + \theta_{j,1} \xi_{t-1} + ... + \theta_{j,q} \xi_{t-q},$$

with $Var\left(\xi_t^{(j)}\right) = \sigma_j^2$, $j = 1, .., s$.

The above model is denoted as $PARMA\,(p_1, q_1, p_2, q_2, ..., p_s, q_s)$. Most attention has been focussed on the periodic AR model, denoted $PAR\,(p_1, ..., p_s)$. In the simple first-order case

$$y_t^{(j)} = \mu_j + \phi_j \left(y_{t-1} - \mu_{j-1}\right) + \varepsilon_t^{(j)}, \quad Var\left(\varepsilon_t^{(j)}\right) = \sigma_j^2, \quad j = 1, .., s \tag{52}$$

For economic time series, the model needs to be able to capture a long-term trend. This can be done by allowing the $u_t^{(j)}$'s to be nonstationary. Osborn (1988) fits a model of the form (52) to UK consumption. She imposes the constraint

$$\prod_{j=1}^{s} \phi_j = 1, \tag{53}$$

so the model has s seasonal and nonseasonal unit roots with a reduced form such that $\Lambda_s y_t \sim MA\,(s-1)$. However, she reports that the model exhibits some residual serial correlation and so is not entirely satisfactory. Harvey and Scott (1994 p.1331-2) argue that this periodic model is overelaborate and unconvincing. They show that a similar fit to U.K. consumption data can be obtained by a non-periodic model consisting of a random walk and a stochastic seasonal. This structural model also has a reduced form in which $\Lambda_s y_t \sim MA\,(s-1)$ but it is much more parsimonious, with only two parameters. It has a straightforward statistical interpretation and is consistent with the rational expectations-permanent income hypothesis of economic theory.

It is important to note that different acf's for different seasons are not necessarily evidence for a periodic model. They can, for example, be produced by seasonally heteroscedastic models; see Proietti (1998 p9).

10.3 Relationship between periodic models and trend plus seasonal models

Sensible application of non-stationary periodic models requires that their relationship to models with stochastic trends and seasonals be properly understood. The first step is set up a periodic generalisation of the local level model:

$$y_t = \mu_t^{(j)} + \varepsilon_t^{(j)}, \qquad t = 1, ..., T, \tag{54}$$

where $Var(\varepsilon_t^{(j)}) = \sigma^2_{\varepsilon,j}$, as in (49), and

$$\mu_t^{(j)} = \mu_{t-1}^{(j)} + \eta_t^{(j)}, \qquad Var(\eta_t^{(j)}) = \sigma^2_{\eta,j}, \qquad j = 1, ..., s. \tag{55}$$

The model is equivalent to a $PARIMA(0, 1, 1, 0, 1, 1, ..., 0, 1, 1)$ model. If it were generalised so that

$$\mu_t^{(j)} = \phi_j \mu_{t-1} + \eta_t^{(j)}, \qquad Var(\eta_t^{(j)}) = \sigma^2_{\eta,j}, \qquad j = 1, ..., s, \tag{56}$$

and the irregular dropped, then (52) would be a special case. However, a slight difference will be introduced with respect to conventional periodic models in that the levels in (55) are assumed to evolve in all time periods, not just in the one in which they directly affect the observations.

The periodic local level model in (54) contains $2s$ hyperparameters. This may be contrasted with a three parameter structural model consisting of a random walk trend, a seasonal and irregular. The question is whether the former contains the latter as a special case.

A single level can be constructed as the average of the individual levels, that is

$$\mu_t = (1/s)\mathbf{i}'\boldsymbol{\mu}_t, \tag{57}$$

where $\boldsymbol{\mu}_t$ contains the s level components, $\mu_t^{(j)}, j = 1, ..., s$. The seasonals are then defined as deviations from this average, that is

$$\boldsymbol{\gamma}_t = \boldsymbol{\mu}_t - \mathbf{i}\mu_t = \boldsymbol{\mu}_t - (1/s)\mathbf{i}\mathbf{i}'\boldsymbol{\mu}_t, \tag{58}$$

and so

$$\boldsymbol{\gamma}_t = \boldsymbol{\gamma}_{t-1} + \boldsymbol{\omega}_t, \qquad \boldsymbol{\omega}_t = \boldsymbol{\eta}_t - (1/s)\mathbf{i}\mathbf{i}'\boldsymbol{\eta}_t, \tag{59}$$

where $\boldsymbol{\eta}_t = (\eta_t^{(1)}, ..., \eta_t^{(s)})'$. Thus

$$y_t = \mu_t + \gamma_t + \varepsilon_t, \qquad t = 1, ..., T \tag{60}$$

where in a type j period, $\gamma_t = \gamma_{jt}$ and $\varepsilon_t = \varepsilon_t^{(j)}$. If the $\sigma_{\eta,j}^{2\prime}s$ in (54) are all the same, and denoted as σ_ω^2,then $Var(\boldsymbol{\omega}_t)$ is as in (9). The variance of the disturbance driving the level, μ_t, is σ_ω^2/s. This restriction, that is tying the level variance to that of the seasonals, can be removed if by introducing a disturbance in the levels common to all seasons, that is

$$\mu_t^{(j)} = \mu_{t-1}^{(j)} + \overline{\eta}_t + \eta_t^{(j)}, \qquad Var(\eta_t) = \overline{\sigma}_\eta^2, \quad Var(\eta_t^{(j)}) = \sigma_{\eta,j}^2, \qquad j = 1, ..., s, \tag{61}$$

where $\overline{\eta}_t$ and the $\eta_t^{(j)\prime}s$ are mutually independent. The vector of levels then becomes

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t-1} + \mathbf{i}\overline{\eta}_t + \boldsymbol{\eta}_t, \qquad Var(\boldsymbol{\eta}_t) = \mathbf{D} = diag(\sigma_{\eta,1}^2, ..., \sigma_{\eta,s}^2) \tag{62}$$

If the common level and the seasonal vector are defined as in (57) and (59) above, the transition equation for the level becomes

$$\mu_t = \mu_{t-1} + \eta_t,$$

with $\eta_t = \overline{\eta}_t + (1/s)\mathbf{i}'\boldsymbol{\eta}_t$. This disturbance is only independent of $\boldsymbol{\omega}_t$ if $Var(\boldsymbol{\eta}_t)$ is a scalar matrix. Setting $Var(\boldsymbol{\eta}_t)$=$\sigma_\omega^2\mathbf{I}$ gives (60) with the level variance equal to σ_η^2+ σ_ω^2/s. If $Var(\boldsymbol{\eta}_t)$ is diagonal, but not necessarily scalar,

$$Var(\boldsymbol{\omega}_t) = \mathbf{D} + s^{-2}\mathbf{ii}'\mathbf{Dii}' - s^{-1}\mathbf{ii}'\mathbf{D} - s^{-1}\mathbf{Dii}'$$

This is slightly different from the specification in (50), though both have the disturbances summing to zero over a year; in the case of the model presented above this is by construction, though, of course, it is also true that $Var(\mathbf{i}'\boldsymbol{\omega}_t) = 0$.

The general model consisting of (54) and (61) therefore contains both the pure periodic and the structural model as special cases. In principle, an LR test of $Var(\boldsymbol{\eta}_t)$ being scalar can be carried out quite easily, with the test statistic being asymptotically χ_{s-1}^2 under the null. A constant variance for the irregular involves another $s-1$ degrees of freedom. Overall $2s-2$ restrictions are needed to give the time invariant structural model. On the other hand, a purely periodic model is obtained if $\overline{\sigma}_\eta^2 = 0$. A LR test statistic for this hypothesis has an distribution which is an even mixture of chi-squares with zero and one degrees of freedom.

If the model is time invariant, the signal extraction filters are (in a doubly infinite sample) the same in all seasons. When periodic features are present this is no longer the case. This has implications for seasonal adjustment using standard procedures. On the other hand, if a model with periodic features is handled using the SSF there is no problem. The measurement equation for (60) is

$$y_t = \mathbf{z}_t'\boldsymbol{\mu}_t + \varepsilon_t^{(j)}, \qquad t = 1, ..., T,$$

with the $s \times 1$ vector $\mathbf{z}_t$ having zero elements everywhere except in position j. A decomposition into level and seasonals is then made from (57) and (58). If MSEs are required then the model is best set up with μ_t and $\boldsymbol{\gamma}_t$ in the state vector. The correlation between the level and seasonal disturbances is

$$Var(\eta_t\boldsymbol{\omega}_t') = Var((\overline{\eta}_t + (1/s)\mathbf{i}'\boldsymbol{\eta}_t)(\boldsymbol{\eta}_t - (1/s)\mathbf{ii}'\boldsymbol{\eta}_t)') = s^{-1}\mathbf{i}'\mathbf{D} - s^{-2}\mathbf{i}'\mathbf{Dii}'$$

The model can easily be extended to include a slope in the trend; (62) is replaced by

$$\begin{aligned} \boldsymbol{\mu}_t &= \boldsymbol{\mu}_{t-1} + \mathbf{i}\overline{\eta}_t + \mathbf{i}\beta_{t-1} + \boldsymbol{\eta}_t, \\ \beta_t &= \beta_{t-1} + \zeta_t \end{aligned}$$

The slope is the same in all seasons and is incorporated into the equation for μ_t.The model could be extended to allow for seasonal specific slopes if desired; see Proietti (1998).

10.4 Partly periodic models

Having different trends in all seasons, as in (54), may be a little extreme and it certainly requires a large number of parameters to be estimated. This suggests a class of models which are partly periodic. There are two ways in which this can arise. Firstly only a subset of components may be periodic. Secondly, periodicity may apply to groups of seasons rather than to all seasons. The two ideas can be combined in that only particular components are periodic with respect to groups of seasons. The SSF allows such models to be handled because the measurement equation picks out the relevant components at any particular point in time.

When all the components are periodic we will refer to the model as being *purely partly periodic*. We begin by considering such models. Suppose there are just two groups for which separate models are to be constructed, with s_1 seasons in the first group and s_2 in the second and $s_1 + s_2 = s$. Then

$$y_t = \mu_t^{(k)} + \gamma_t^{(k)} + \varepsilon_t^{(k)}, \qquad k = 1, 2, \tag{63}$$

where the seasonal, $\gamma_t^{(k)}$, is modelled by a set of $s_k, k = 1,2$ time-varying dummies which embody the zero sum restriction over the group. If a group has only one season then the seasonal component is not needed. The trends can be assumed to have slopes for more generality so that

$$\begin{aligned}\mu_t^{(k)} &= \mu_{t-1}^{(k)} + \beta_{t-1}^{(k)} + \eta_t^{(k)}, \qquad Var(\eta_t^{(k)}) = \sigma^2_{\eta,k}, \qquad k = 1,2. \\ \beta_t^{(k)} &= \beta_{t-1}^{(k)} + \zeta_t^{(k)}, \qquad Var(\zeta_t^{(k)}) = \sigma^2_{\zeta,k}.\end{aligned}$$

The overall trend is given by

$$\mu_t = (s_1/s)\mu_t^{(1)} + (s_2/s)\mu_t^{(2)}.$$

This may be useful in the context of smoothing and forecasting. A seasonal component for the model as a whole can be defined as

$$\gamma_t = \mu_t^{(k)} - \mu_t + \gamma_t^{(k)}, \qquad k = 1,2.$$

In a pure partly periodic model, the disturbances in the two groups are mutually independent. However, some correlation could be introduced between the disturbances in the trends. For example suppose ρ is the correlation between η_{1t} and η_{2t} and between ζ_{1t} and ζ_{2t}. If $\rho = 1$ together with $\sigma^2_{1\eta} = \sigma^2_{2\eta}$ and $\sigma^2_{1\zeta} = \sigma^2_{2\zeta}$, the model reduces to the seasonal heteroscedastic formulation of (49), while if $\sigma^2_{1\omega} = \sigma^2_{2\omega}$ as well, the BSM is obtained. (There is a constant difference between the two trends, which can be transferred to the seasonals).

Now suppose the group periodic effects only apply to certain components. At the simplest level this might mean letting the variance of the irregular in (63) be the same in all seasons. This restriction is easily enforced by the joint treatment of the two groups within the SSF; in fact there is one less parameter to estimate. If a common cycle were added the model would become

$$y_t = \mu_t^{(k)} + \gamma_t^{(k)} + \psi_t + \varepsilon_t, \qquad k = 1,2, \tag{64}$$

Further generalisation might involve the introduction of seasonal heteroscedasticity. This does not need to be connected in any way with the periodic groupings.

A test of seasonal stability within a group, that is $\sigma^2_{k\omega} = 0$, can be carried out using a seasonal specific version of the seasonal stationarity test. To test that the seasonals in the first group are fixed, let $\mathbf{A}$ be an $s \times (s_1 - 1)$ full rank matrix with each of the first $(s_1 - 1)$ columns containing a one and a minus one while the remaining elements are all zero. Then $\mathbf{A}'\mathbf{y}_t$ is stationary under the null hypothesis that $\sigma^2_{1\omega} = 0$ and the asymptotic distribution of the test statistic is $CvM(s_1 - 1)$.

Italian industrial production Since August behaves so differently to the other months it is worth letting it have its own trend. Thus the model is as in (63) but with slopes in the trend components and no seasonal component, as such, for August. Other seasons may be deterministic, while August is not. Such a feature is not possible in the seasonally heteroscedastic model of (50). Note that the August trend is allowed to be correlated with the (common) trend in the other seasons.

REFERENCES

Akaike, H. (1980). Seasonal adjustment by a Bayesian modeling. *Journal of Time Series Analysis* 1: 1-13.

Bell, W.R., and S.C. Hillmer (1983). Modeling time series with calendar variation. *Journal of the American Statistical Association* 78: 526-34.

(1984). Issues involved with the seasonal adjustment of economic time series [with discussion]. *Journal of Business and Economic Statistics* 2: 292-349.

Box, G.E.P., and G.M. Jenkins (1976). *Time Series Analysis: Forecasting and Control*, revised edn. San Francisco: Holden-Day.

Bruce, A. G., and S. R. Jurke (1996). Non-Gaussian seasonal adjustment: X-12-ARIMA versus robust structural models. *Journal of Forecasting* 15: 305-28.

Busetti, F.and A.C. Harvey (2003), Seasonality tests. *Journal of Business and Economic Statistics, 21, 420-36.*

Canova, F., and B.E. Hansen (1995), Are seasonal patterns constant over time? A test for seasonal stability. *Journal of Business and Economic Statistics* 13**:** 237-52.

Cleveland, W.S., and S.J. Devlin (1980). Calendar effects in monthly time series: detection by spectrum analysis and graphical methods. *Journal of the American Statistical Association* 75: 487-96.

Dagum, E.B., Quenneville, B. and Sutradhar, B. (1992) Trading-day multiple regression models with random parameters. *International Statistical Review,* 60, 57-73.

De Rossi, G. and Harvey, A C (2006). Time-varying quantiles. Cambridge, Faculty of Economics, CWPE 0649.

Durbin, J., and S.J. Koopman (2001). *Time series analysis by state space methods.* Oxford University Press, Oxford.

Franses, P. H. and R. Papp (2004). *Periodic time series models.* Oxford University Press, Oxford.

Hannan, E.J., R.D. Terrell and N. Tuckwell (1970). The seasonal adjustment of economic time series. *International Economic Review* 11: 24-52.

Harrison, P.J., and C.F. Stevens (1976). Bayesian forecasting. *Journal of the Royal Statistical Society, Series B* 38: 205-47.

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

Harvey A.C. (2006). Forecasting with Unobserved Components Time Series Models. *Handbook of Economic Forecasting*, edited by G Elliot, C Granger and A Timmermann. North Holland, 327-412.

Harvey, A. C., and S. J. Koopman (1993). Forecasting hourly electricity demand using time-varying splines. *Journal of American Statistical Association* 88: 1228-36.

Harvey, A. C., S. J. Koopman, and M. Riani (1997). The modeling and seasonal adjustment of weekly observations. *Journal of Business and Economic Statistics* 15: 354-68.

Harvey, A. C., and A. Scott (1994). Seasonality in dynamic regression models. *Economic Journal* 104: 1324-45.

Harvey, A. C., and M. Streibel (1998). Tests for deterministic versus indeterministic cycles. *Journal of Time Series Analysis* 19: 505-29.

Hillmer, S.C. (1982). Forecasting time series with trading day variation. *Journal of Forecasting* 1: 385-95.

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.

Hipel, R. W., and McLeod, A. I. (1994). *Time series modelling of water resources and environmental systems.* Developments in Water Science, 45, Elsevier, Amsterdam.

Hylleberg, S. (1995). Tests for seasonal unit roots: general to specific or specific to general? *Journal of Econometrics*, 69, 5-25.

Hylleberg, S., Engle, R., Granger, C.W.J and B.S. Yoo (1990). Seasonal integration and co-integration. *Journal of Econometrics* 44, 215-38.

Kitagawa, G., and W Gersch (1984). A smoothness priors-state space modeling of time series with trend and seasonality. *Journal of the American Statistical Association* 79: 378-89.

Koopman, S.J. and Harvey, A.C. (2003) Computing observation weights for signal extraction and filtering. *Journal of Economic Dynamics and Control*, 27, 1317-33.

Koopman, S.J., A.C. Harvey, J.A. Doornik and N. Shephard (2006). *STAMP 7.0 Structural Time Series Analyser, Modeller and Predictor*, London: Timberlake Consultants Ltd.

Leybourne, S.J. and B.P.M. McCabe (1994), A consistent test for a unit root, *Journal of Business and Economic Statistics* 12, 157-166.

Maravall A. (1985). On structural time series models and the characterization of components. *Journal of Business and Economic Statistics* 3: 350-5.

Osborn, D. R. (1988). Seasonality and habit persistence in a life cycle model of consumption. *Journal of Applied Econometrics,* 3: 255-66.

Osborn, D. R and J. R. Smith (1989). The performance of periodic autoregressive models in forecasting U.K consumption. *Journal of Business and Economic Statistics*, 7, 117-27.

Poirier, D.J. (1976) The Econometrics of Structural Change: With Special Emphasis on Spline Functions Amsterdam: North-Holland.

Proietti, T. (1998). Seasonal heteroscedasticity and trends, *Journal of Forecasting* 17: 1-17.

Proietti, T. (2000). Comparing seasonal components for structural time series models, *International Journal of Forecasting,* 16: 247-60.

Smith, R.J. and A.M.R. Taylor (1998), Additional critical values and asymptotic representations for seasonal unit root tests, *Journal of Econometrics* 85, 269-288.

Soukup, R.J. and D.F. Findley (2000), On the Spectrum Diagnostics used by X-12- ARIMA to Indicate the Presence of Trading Day Effects After Modeling or Adjustment, @ 1999 Proceedings of the Business and Economic Statistics Section, American Statistical Association, 144-149.

Thompson, H.E., and G.C. Tiao (1971). Analysis of telephone data: a case of forecasting seasonal time series. The Bell System Journal of Economics and *Management Science* 2: 515-41.

Whittle, P. (1983) *Prediction and Regulation*, 2nd ed. Oxford: Blackwell.