Forecasting Seasonal Time Series*

Philip Hans Franses

Econometric Institute, Erasmus University Rotterdam

September 6, 2004

Abstract

This chapter deals with seasonal time series in economics and it reviews models that can be used to forecast out-of-sample data. Some of the key properties of seasonal time series are reviewed, and various empirical examples are given for illustration. The potential limitations to seasonal adjustment are reviewed. The chapter further addresses a few basic models like the deterministic seasonality model and the airline model, and it shows what features of the data these models assume to have. Then, the chapter continues with more advanced models, like those concerning seasonal and periodic unit roots. Finally, there is a discussion of some recent advances, which mainly concern models which allow for links between seasonal variation and heteroskedasticity and non-linearity.

Key words and phrases: Forecasting, seasonality, unit roots, periodic models, cointegration

*Work on this chapter started when the author was enjoying the hospitality of the Institute of Mathematical Statistics, National University of Singapore (April 2004). Thanks are due to Dick van Dijk and Marnik Dekimpe for their help with the data. All computations have been made using Eviews (version 4.1). Details of estimation results not reported in this chapter can be obtained from the author. A full list of relevant references is also available. The address for correspondence is Econometric Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, email: franses@few.eur.nl

Contents

1 Introduction **1**

2 Seasonal time series **2**

3 Basic Models **15**

4 Advanced Models **24**

5 Recent advances **33**

6 Conclusion **36**

1 Introduction

This chapter deals with seasonal time series in economics and it reviews models that can be used to forecast out-of-sample data. A seasonal time series is assumed to be a series measured at a frequency t where this series shows certain recurrent patterns within a frequency T, with $t = ST$. For example, quarterly data (t) can show different means and variances within a year ($4T$). Similar phenomena can appear for hourly data within days, daily data within weeks, monthly data within years, and so on. Examples of data with pronounced recurrent patterns are quarterly nondurable consumption, monthly industrial production, daily retail sales (within a week), hourly referrals after broadcasting a TV commercial (within a day), and stock price changes measured per minute within a day.

A trend is important to extrapolate accurately when forecasting longer horizons ahead. Seasonality is important to properly take care off when forecasting the next S or kS out-of-sample data, with k not very large, hence the medium term. This chapter reviews models that can be usefully implemented for that purpose. The technical detail is kept at a moderate level, and extensive reference will be made to the studies which contain all the details. Important books in this context are Hylleberg (1992), Ghysels and Osborn (2001), and Franses and Paap (2004). The recent survey of Brendstrup et al. (2004) is excellent, also as it contains a useful and rather exhaustive list of references.

The outline of this chapter is a follows. In Section 2, I discuss some of the key properties of seasonal time series. I use a few empirical examples for illustration. Next, I discuss the potential limitations to seasonal adjustment. In Section 3, I review basic models like the deterministic seasonality model and the airline model, and show features of the data these models assume they have. In Section 4, I continue with more advanced models, like those concerning seasonal and periodic unit roots. Section 5 deals with some recent advances, which mainly concern models which allow for links between seasonal variation and heteroskedasticity and non-linearity. Section 6 concludes with a summary of important future research areas.

2 Seasonal time series

This section deals with various features of seasonally observed time series that one might want to capture in an econometric time series model. Next, the discussion focuses on what it is that one intends to forecast. Finally, I address the issue why seasonal adjustment often is problematic.

How do seasonal time series look like?

In this chapter I use a few series for illustration. The typical tools to see how seasonal variation in a series might look like, and hence which time series models might be considered to start with, are (i) graphs (over time, or per season), (ii) autocorrelations (usually after somehow removing the trend, where it is not uncommon to use the first differencing filter $\Delta_1 y_t = y_t - y_{t-1}$), (iii) the R^2 of a regression of $\Delta_1 y_t$ on S seasonal dummies or on S (or less) sines and cosines, (iv) a regression of squared residuals from a time series model for $\Delta_1 y_t$ on an intercept and $S-1$ seasonal dummies (to check for seasonal variation in error variance), and finally (v) autocorrelations per season (to see if there is periodicity). With periodicity one typically means that correlations within or across variables can change with the season, see Franses and Paap (2004) for an up to date survey. It should noted that these are all just first-stage tools, to see in which direction one could proceed. They should not be interpreted as final models or methods, as they usually do not fully capture all relevant aspects of the time series.

The first set of series concerns private consumption and GDP for Japan, quarterly observed, for the period 1980.1-2001.2 (Data source is www. economagic.com). The graphs of these two series appear in Figure 1. The graphs display an upward moving trend for both series, pronounced intra-year variation, and it seems that this variation is common across the two series.

Figure 2 zooms in on the Japanese consumption series (now in natural logarithms, hence the notation LC) by plotting the quarterly observations against the year (Q_1, Q_2, Q_3 and Q_4). This way one can get a better picture of whether seasonal patterns change over time, as then these lines would intersect. These graphs were introduced

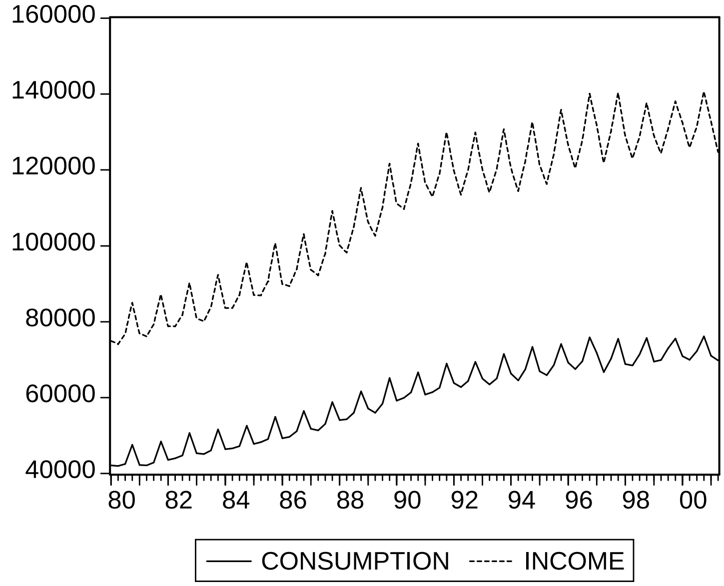

Figure 1: Quarterly consumption and income per quarter in Japan

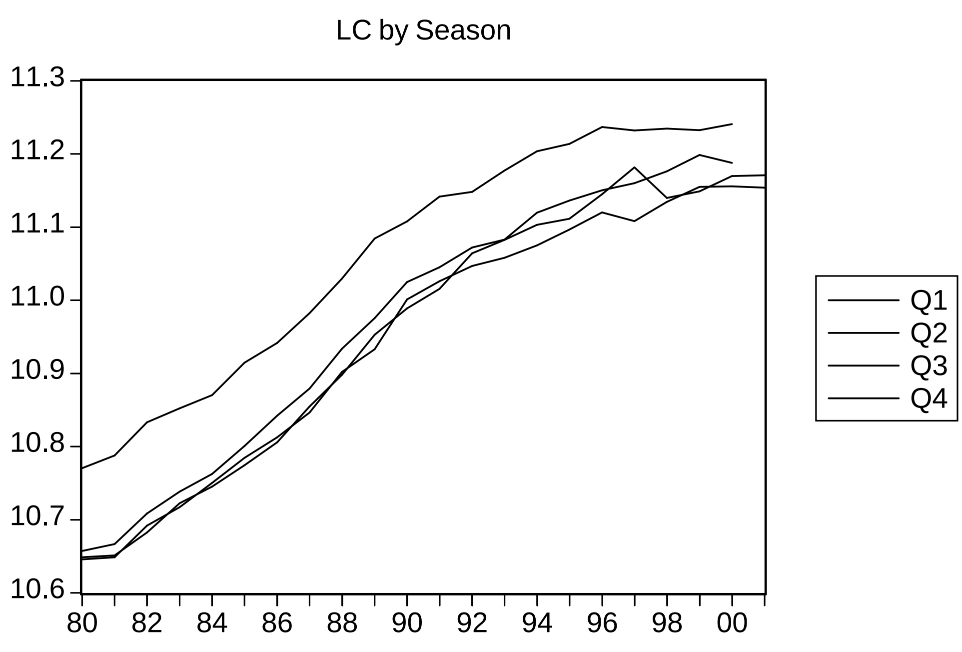

Figure 2: Annual consumption in Japan, observed per quarter

in Franses (1991, 1994) and now appear in Eviews (version 4.1) as "split seasonals". For the Japanese consumption series one can observe that there is a slight change in seasonality towards the end of the sample, but mostly the seasonal pattern seems rather stable over time.

For these two series, after taking natural logs, the R^2 of the "seasonal dummy regression" for $\Delta_1 y_t$, that is,

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

is 0.927 for log(consumption) and for log(income) it is 0.943. The $D_{s,t}$ variables obtain a value 1 in seasons s and a 0 elsewhere. Note that it is unlikely that $\hat{\varepsilon}_t$ matches with a white noise time series, but then still, the values of these R^2 measures are high. Franses, Hylleberg and Lee (1995) show that the size of this R^2 can be misinterpreted in case of neglected unit roots, but for the moment this regression is informative.

A suitable first-attempt model for both Japanese log(consumption) and log(income) is

$$\Delta_1 y_t = \sum_{s=1}^{4} \delta_s D_{s,t} + \rho_1 \Delta_1 y_{t-1} + \varepsilon_t + \theta_4 \varepsilon_{t-4}. \tag{2}$$

where ρ_1 is estimated to be -0.559 (0.094) and -0.525 (0.098), respectively (with standard errors in parentheses), and where θ_4 is estimated to be 0.441 (0.100) and 0.593 (0.906), respectively. Relative to (1), the R^2 of these models have increased to 0.963 and 0.975, respectively, suggesting that constant deterministic seasonality seems to account for the majority of trend-free variation in the data.

The next series is quarterly observed M1 for Australia for 1975.2 to and including 2004.1 (Data source: www.economagic.com), see Figure 3. Again, one can observe a marked upward trend, and there are also signs of seasonality, but this time the type of seasonality is a bit unclear. This might be caused by the dominance of the trend, and hence one might want to have a look at the time series without a trend. There are many ways to de-trend a time series, but for the current purpose it is again convenient to take the natural logarithm and then first differences, approximately amounting to quarterly growth rates. The graph of quarterly growth in M1 appears

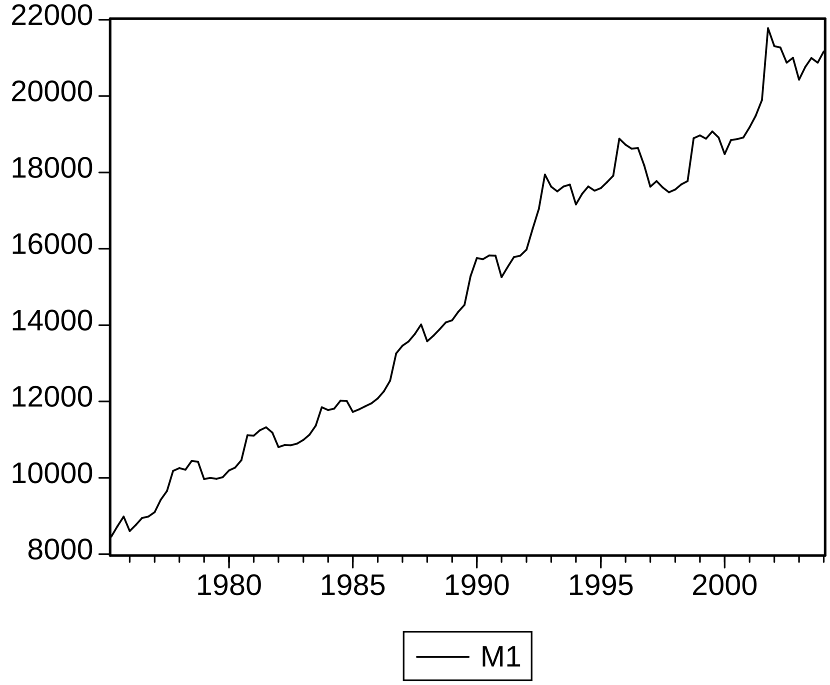

Figure 3: Quarterly M1 in Australia

in Figure 4.

Figure 4 clearly shows there is seasonality in M1 growth[1]. A regression of the growth rates (differences in logs) for this variable on four seasonal dummies gives an R^2 of 0.203. Fitting autoregressive models of order 8, 7, 6 and so on, while checking for the absence of residual correlation, reveals that an AR(5) model fits the data for $\Delta_1 y_t$ reasonably well, although the residuals are not entirely "white". The R^2 increases to 0.519, with strong significant parameters for lags 3, 4, and 5, and hence, seasonality for this series cannot fully be captured by deterministic seasonality.

Next, a regression of the squares of the residuals from the AR(5) model on an intercept and 3 seasonal dummies gives an $F-$value of 5.167, with a $p-$value of 0.002. Hence, this series seems to display seasonal variation in the variance. One cause for this finding is that a better model for this series could be a periodic time series model, which implies seasonal heteroskedasticity if a non-periodic model is fitted,

[1] Interestingly, this seasonality seems to concern intervals of 2 years instead of the usual 1 year, but for the moment this is not pursued any further.

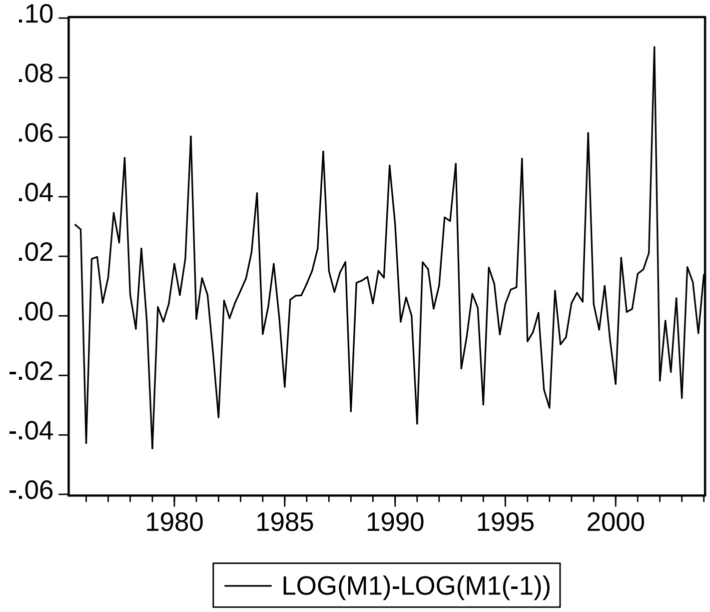

Figure 4: Quarterly growth in M1 in Australia

see Franses and Paap (2004). When I fit an AR(2) model for each of the seasons, that is, I regress $\Delta_1 y_t$ on $\Delta_1 y_{t-1}$ and $\Delta_1 y_{t-2}$ but allow for different parameters for the seasons, then the estimation results for quarters 1 to 4 are (0.929, 0.235), (0.226, 0.769), (0.070, -0.478), and (0.533, -0.203). This suggests that different models for different seasons might be useful, hence a periodic model.

The next series to consider is the index of monthly total industrial production for the USA, covering 1919.01-2004.02 (Data source: www.economagic.com), and its graph is given in Figure 5. Again a trend is clearly visible, and also at times one can observe dips, which are typically assumed to correspond with recessions. There is ample literature on the supposed non-linearity of this time series, see for example Franses and van Dijk (2004) and the references therein, but this is neglected here for the moment, see Section 5. Additionally, as Figure 6 indicates, there seems to be a change in the variance in this series.

A regression of Δ_1 of log(industrial production) on $S = 12$ seasonal dummies gives an R^2 of 0.374. Adding lags at 1, 12 and 13 to this auxiliary regression model

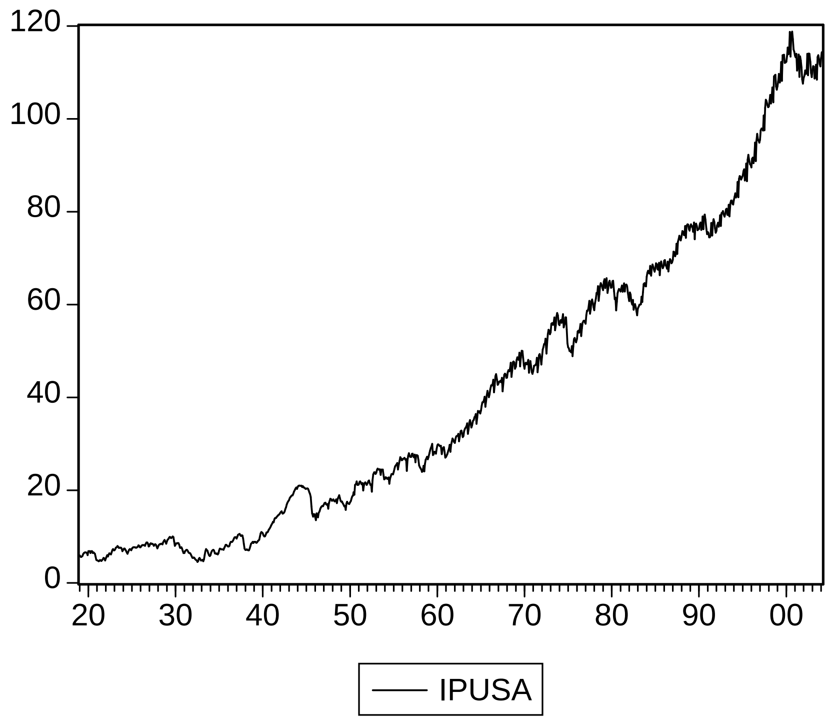

Figure 5: Monthly index of total industrial production in the USA

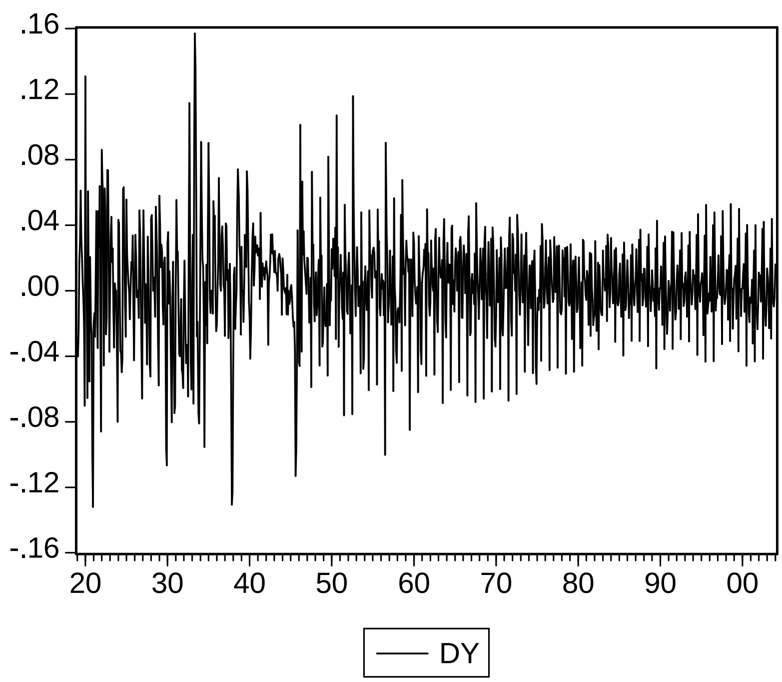

Figure 6: Monthly growth in industrial production in the USA

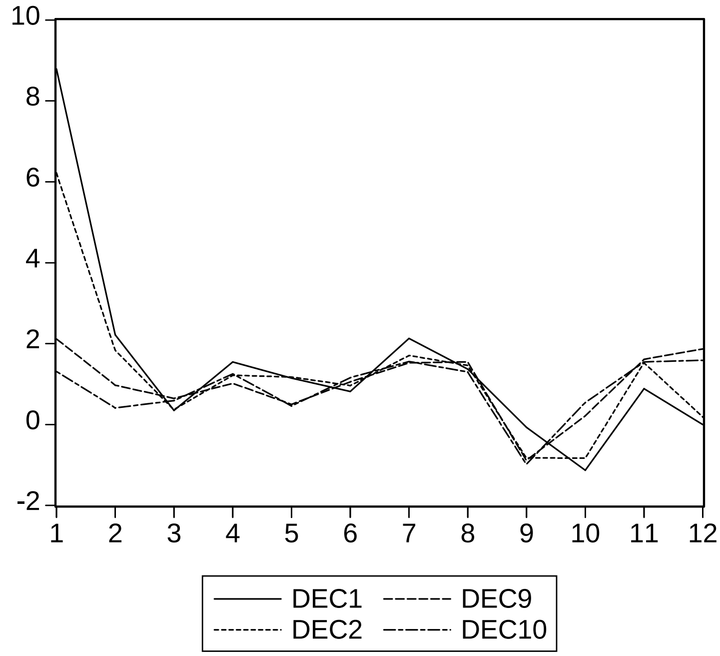

Figure 7: deciles

improves the fit to 0.524 (for 1008 data points), with parameters 0.398, 0.280 and -0.290. There is no obvious residual autocorrelation. The sum of these autoregressive parameters is 0.388. This implies that certainly for forecasts beyond the 12 month horizon, constant seasonality dominates. Testing for seasonal heteroskedasticity in the way outlined above for Australian M1 does not suggest any such variation.

The next series are the monthly returns of 10 decile portfolios (ranked according to market capitalization), for the New York Stock Exchange, ranging from 1926.08 to 2002.12[2]. These returns might be best described by the simple regression model

$$y_t = \sum_{s=1}^{12} \delta_s D_{s,t} + \varepsilon_t \tag{3}$$

$$\varepsilon_t = \rho_1 \varepsilon_{t-1} + u_t. \tag{4}$$

One might expect a "January effect", in particular for the smaller stocks, see Haugen and Lakonishok (1987).

[2]Data source is http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/

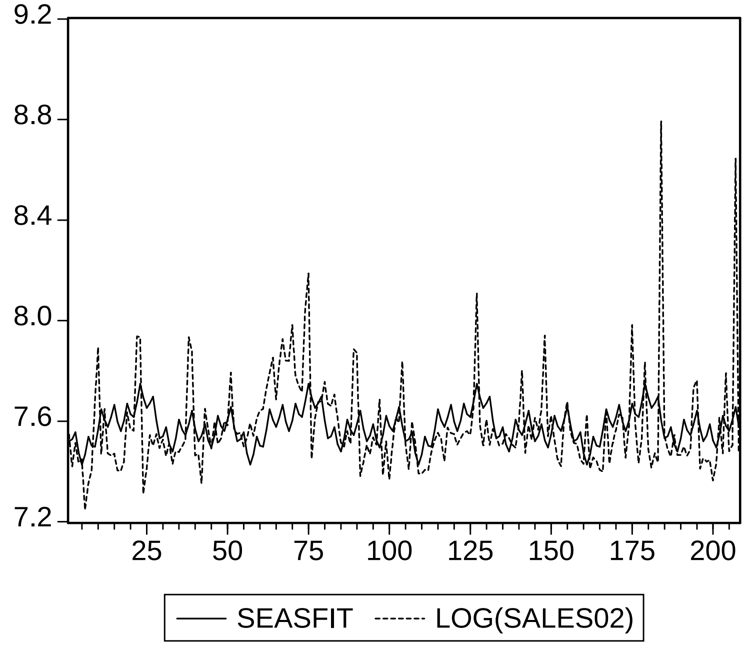

Figure 8: Deterministic seasonal cycles in coffee sales

Comparing the estimated parameters for the decile models and their associated standard errors suggests that only a few parameters are significant. Figure 7 depicts the estimates of $\hat{\delta}_s$ for the first two and for the last two deciles, where the first decile usually gets a significant positive parameter. Also the R^2 values are higher for lower deciles. Clearly, the graphs in Figure 7 suggest the presence of a January effect for the smaller stocks.

The parameters for the seasonal dummies in the regression model in (4) are often not significant. In practice, this is a phenomenon that is quite common for disaggregated data. For example, for weekly sales of instant decaf coffee (observed in the stores of a retail chain in the Netherlands, for 1994 week 29 to 1998 week 28), one might consider

$$y_t = \sum_{s=1}^{52} \delta_s D_{s,t} + \varepsilon_t, \tag{5}$$

but this involves a large amount of parameters, and most likely, many of these will not be statistically relevant. One can then reduce this number by deleting certain $D_{s,t}$ variables, but this might complicate the interpretation. In that case, a more sensible model is

$$y_t = \mu + \sum_{k=1}^{26} [\alpha_k \cos(\frac{2\pi kt}{52}) + \beta_k \sin(\frac{2\pi kt}{52})] + \varepsilon_t, \quad (6)$$

where $t = 1, 2, ...$ A cycle within 52 weeks (an annual cycle) corresponds with $k = 1$, and a cycle within 4 weeks corresponds with $k = 13$. Other interpretable cycles would correspond with 2 weeks, 13 weeks and 26 weeks ($k = 26$, 4 and 2, respectively). Note that $\sin(\frac{2\pi kt}{52})$ is equal to 0 for $k = 26$, hence the intercept μ in (6). One may now decide to include only those cycles that make sense from an interpretation point of view.

Figure 8 shows the fit of the model in (6), where ε_t is assumed an AR(1) process, and where only cycles within 2, 4, 13, 26 and 52 weeks are considered. Hence, there are only 9 variables to characterize seasonality. The R^2 measure is 0.139, suggesting that there is moderate deterministic seasonality in this weekly series.

What do we want to forecast?

For seasonal data with seasonal frequency S, one usually considers forecasting 1-step ahead, $\frac{S}{2}$-steps ahead or S-steps ahead. One may also want to forecast the sum of 1 to S steps ahead, that is, say, a year.

There is no general rule that says that forecasting data at the T frequency is better done using data for that particular frequency than by using data at a higher frequency, and then sum the forecasts. For example, when the purpose is to forecast two years ahead, and one has monthly data, one can choose to use a model for the annual data or for the monthly data. There are less annual data than monthly data, so one has less information to specify a model. On the other hand, monthly data show seasonality that one has to capture and such data also might have more outlying observations which may affect model construction.

Also, if the aim is to forecast a full year ahead, one might perhaps, at least in principle, consider modelling seasonally adjusted data. Of course, these adjusted data should not show any seasonality, and the adjustment method should not have introduced data features that were not there to begin with. As I will discuss next, there are however some problematic aspects of seasonally adjusted data.

Why is seasonal adjustment often problematic?

It is common practice to seasonally adjust quarterly or monthly observed macroeconomic time series, like GDP and unemployment. A key motivation is that practitioners seem to want to compare the current observation with that in the previous month or quarter, without considering seasonality. As many series display seasonal fluctuations which are not constant over time, at least not for the typical time span considered in practice, there is a debate in the statistics and econometrics literature about which method is most useful for seasonal adjustment. Roughly speaking, there are two important methods. The first is the Census X-11 method, initiated by Shiskin and Eisenpress (1957), and the second one uses model-based methods, see for example Maravall (1995). Interestingly, it seems that with the new Census X-12 method, the two approaches have come closer together, see Findley et al. (1998). In Franses (2001) I address the question why one would want to seasonally adjust in the first place, and what follows in this subsection draws upon that discussion.

Except for macroeconomics, there is no economic discipline in which the data are seasonally adjusted prior to analysis. It is hard to imagine, for example, that there would be a stock market index, with returns corrected for day-of-the-week effects. Also, seasonality in sales or market shares is of particular interest to a manager, and seasonal adjustment of marketing data would simply result in an uninteresting time series.

Generally, the interest in analyzing macroeconomic data concerns the trend and the business cycle. In case the data have stochastic trends, one usually resorts to well-known techniques for common trends analysis and cointegration, see for example Engle and Granger (1991). To understand business cycle fluctuations, for example in the sense of examining which variables seem to be able to predict recessions, one can use nonlinear models like the (smooth transition) threshold model and the Markov-switching model, see Granger and Teräsvirta (1993) and Franses and van Dijk (2000) for surveys.

Consider a seasonally observed time series y_t, where t runs from 1 to n. In practice one might be interested in the seasonally adjusted observation at time n or

$n-1$. The main purpose of seasonal adjustment is to separate the observed data into two components, a nonseasonal component and a seasonal component. These components are not observed, and have to be estimated from the data. It is assumed that

$$y_t = \hat{y}_t^{NS} + \hat{y}_t^{S}, \tag{7}$$

where $\hat{y}_t^{NS}$ is the estimated nonseasonal component, and $\hat{y}_t^{S}$ is the estimated seasonal component. This decomposition assumes an additive relation. When this is not the case, one can transform y_t until it holds for the transformed data. For example, if the seasonal fluctuations seem multiplicative with the trend, one typically considers the natural logarithmic transformation.

As said, there are two commonly used approaches to estimate the components in (7). The first is coined Census X-12. This approach applies a sequence of two-sided moving average filters like

$$w_0 + \sum_{i=1}^{m} w_i(L^i + L^{-i}), \tag{8}$$

where L is the familiar backward shift operator, and where the value of m and the weights w_i for $i = 0, 1, \ldots, m$ are set by the practitioner. It additionally contains a range of outlier removal methods, and corrections for trading-day and holiday effects. An important consequence of two-sided filters is that to adjust observation y_n, one needs the observations at time $n+1$, $n+2$ to $n+m$. As these are not yet observed at n, one has to rely on forecasted values, which are then treated as genuine observations. Of course, this automatically implies that seasonally adjusted data should be revised after a while, especially if the newly observed realizations differ from those forecasts. Interesting surveys of this method are given in Bell and Hillmer (1984), Hylleberg (1986), and more recently in Findley et al. (1998).

The second approach involves model-based methods. These assume that the seasonal component can be described by a model like for example

$$(1 + L + L^2 + L^3)y_t^S = \varepsilon_t. \tag{9}$$

With an estimate of the variance of ε_t, and with suitable starting-values, one can estimate the seasonal component using Kalman-filtering techniques, see Harvey (1989).

Given $\hat{y}_t^S$, one can simply use (7) to get the estimated adjusted series.

A few remarks can be made. The first amounts to recognizing that seasonally adjusted data are *estimated data*. In practice this might be forgotten, which is mainly due to the fact that those who provide the seasonally adjusted data tend not to provide the associated standard errors. This is misleading. Indeed, a correct statement would read "this month's unemployment rate is 7.8, and after seasonal adjustment it is 7.5 plus or minus 0.3". The Census X-12 method cannot generate standard errors, but for the model-based methods it is not difficult to do so. Koopman and Franses (2003) propose a method which also allows for business cylce-dependent confidence intervals around seasonally adjusted data.

Obviously, when $\hat{y}_t^{NS}$ is saved and $\hat{y}_t^S$ is thrown away, one cannot reconstruct the original series y_t. Moreover, if the original series y_t can be described by an econometric time series model with innovations ε_t, it is unclear to what extent these innovations are assigned to either $\hat{y}_t^{NS}$, $\hat{y}_t^S$ or to both. Hence, when one constructs an econometric time series model for the adjusted series $\hat{y}_t^{NS}$, the estimated innovations in this model are not the "true" innovations. This feature makes impulse-response analysis less interesting.

The key assumption is the relation in (7). For some economic time series this relation does not hold. For example, if the data can best be described by a so-called periodic time series model, where the parameters vary with the seasons, see Section 4 below, one cannot separate out a seasonal component and reliably focus on the estimated nonseasonal component. There are a few theoretical results about what exactly happens if one adjusts a periodic series, and some simulation and empirical results are available, see Franses (1996), Ooms and Franses (1997) and Del Barrio Castro and Osborn (2004). Generally, seasonally adjusted periodic data still display seasonality.

Given the aim of seasonal adjustment, that is, to create time series which are more easy to analyze for trends and business cycles, it is preferable that seasonally adjusted data (1) show no signs of seasonality, (2) do not have trend properties that differ from those of the original data, and (3) that they do not have other non-linear properties than the original data. Unfortunately, it turns out that most publicly available adjusted data do not have all of these properties. Indeed, it frequently occurs that $\hat{y}_t^{NS}$ can be modeled using a seasonal ARMA model, with highly significant parameters at seasonal lags in both the AR and MA parts of the model. The intuition for this empirical finding may be that two-sided filters as in (8) can be shown to assume quite a number of so-called seasonal unit roots, see Section 3 below. Empirical tests for seasonal unit roots in the original series however usually suggest a smaller number of such roots, and by assuming too many such roots, seasonal adjustment introduces seasonality in the MA part of the model. Furthermore, and as mentioned before, if the data correspond with a periodic time series process, one can still fit a periodic time series model to the adjusted data. The intuition here is that linear moving average filters treat all observations as equal.

Would seasonal adjustment leave the trend property of the original data intact? Unfortunately not, as many studies indicate. The general finding is that the persistence of shocks is higher, which in formal test settings usually corresponds with more evidence in favor of a unit root. In a multivariate framework this amounts to finding less evidence in favor of cointegration, that is, of the presence of stable long-run relationships, and thus more evidence of random walk type trends. The possible intuition of this result is that two-sided filters make the effects of innovations to appear in $2m + 1$ adjusted observations, thereby spuriously creating a higher degree of persistence of shocks. Hence, seasonal adjustment incurs less evidence of long-run stability.

Non-linear data do not become linear after seasonal adjustment, but there is some evidence that otherwise linear data can display non-linearity after seasonal adjustment, see Ghysels, Granger and Siklos (1996). Additionally, non-linear models for the original data seem to differ from similar models for the adjusted data. The structure of the non-linear model does not necessarily change, it merely concerns the parameters in these models. Hence, one tends to find other dates for recessions for adjusted data than for unadjusted data. A general finding is that the recessions for adjusted data last longer. The intuition for this result is that expansion data are used to adjust recession data and the other way round. Hence, regime switches get smoothed away or become less pronounced.

In sum, seasonally adjusted data may still display some seasonality, can have different trend properties than the original data have, and also can have different non-linear properties. It is my opinion that this suggests that these data may not be useful for their very purpose.

3 Basic Models

This section deals with a few basic models that are often used in practice. They also often serve as a benchmark, in case one decides to construct more complicated models. These models are the constant deterministic seasonality model, the seasonal random walk, the so-called airline model and the basic structural time series model.

The deterministic seasonality model

This first model is useful in case the seasonal pattern is constant over time. This constancy can be associated with various aspects. First, for some of the data we tend to analyze in practice, the weather conditions do not change, that is, there is a intra-year climatological cycle involving precipitation and hours of sunshine that is rather constant over the years. For example, the harvesting season is reasonably fixed, it is known when lakes and harbors are ice-free, and our mental status also seems to experience some fixed seasonality. In fact, consumer survey data (concerning consumer confidence) show seasonality, where such confidence is higher in January and lower in October, as compared with other months. Some would say that such seasonality in mood has an impact on stock market fluctuations, and indeed, major stock market crashes tend to occur more often in the month of October. Other regular phenomena concern calender-based festivals and holidays. Finally, institutional factors as tax years, end-of-years bonuses, and school holidays, can make some economic phenomena to obey a regular seasonal cycle.

A general model for constant seasonality in case there are S seasons is

$$y_t = \mu + \sum_{k=1}^{\frac{S}{2}} [\alpha_k \cos(\frac{2\pi kt}{S}) + \beta_k \sin(\frac{2\pi kt}{S})] + u_t, \tag{10}$$

where $t = 1, 2, ...$ and u_t is some ARMA type process. This expression makes explicit that constant deterministic seasonality can also be viewed as a sum of cycles, defined by sines and cosines. For example, for $S = 4$, one has

$$y_t = \mu + \alpha_1 \cos(\frac{1}{2}\pi t) + \beta_1 \sin(\frac{1}{2}\pi t) + \alpha_2 \cos(\pi t) + u_t, \quad (11)$$

where $\cos(\frac{1}{2}\pi t)$ equals (0,-1,0,1,0,-1,...) and $\sin(\frac{1}{2}\pi t)$ is (1,0,-1,0,1,0,...) and $\cos(\pi t)$ is (-1,1,-1,1,...). The μ is included as $\sin(\pi t)$ is zero everywhere.

The expression in terms of sines and cosines is also relevant as it matches more naturally with the discussion below on filters. For example, if one considers $y_t + y_{t-2}$ for quarterly data, which can be written as $(1 + L^2)y_t$, then $(1 + L^2)\cos(\frac{1}{2}\pi t) = 0$ and $(1 + L^2)\sin(\frac{1}{2}\pi t) = 0$. This means that this filter effectively cancels part of the deterministic seasonal variation. Additionally, it holds that $(1 + L)\cos(\pi t) = 0$, and of course, $(1 - L)\mu = 0$. This shows that deterministic seasonality can be removed by applying the transformation $(1 - L)(1 + L)(1 + L^2) = 1 - L^4$. In words it means that comparing the current quarter with the same quarter last year effectively removes the influence of deterministic seasonality, if there would be any, or that of a trends, again if there would be any. I will return to this transformation later on.

Naturally, there is a one-to-one link between the model in (10) and the model which has the familiar S seasonal dummy variables. For $S = 4$, one has

$$y_t = \sum_{s=1}^{4} \delta_s D_{s,t} + u_t, \quad (12)$$

and it holds that $\mu = \sum_{s=1}^{4} \delta_s$, and that $\alpha_1 = \delta_4 - \delta_2$, $\beta_1 = \delta_1 - \delta_3$ and $\alpha_2 = \delta_4 - \delta_3 + \delta_2 - \delta_1$. There is no particular reason to favor one of the two models, except for the case where S is large, as I mentioned before. For example, when S is 52, a model like in (12) contains many parameters, of which many might turn out to be insignificant in practice. Additionally, the interpretation of these parameters is also not easy. In contrast, for the model in (10) one can choose to assume that some α and β parameters are equal to zero, simply as they are associated with deterministic cycles which are not of interest for the analysis at hand.

The constant seasonality model is applied widely in marketing and tourism. In finance, one might expect seasonality not to be too constant over time, basically as that would imply that traders could make use of it. Further, many macroeconomic data seem to display seasonality that changes over time, as is illustrated by for example Canova and Ghysels (1994) and Canova and Hansen (1995). Seasonal patterns can change due to changing consumption patterns. For example, one nowadays needs to pay for next year's holiday well in advance. Also, one can eat ice in the winter, and nowadays have all kinds of vegetables in any season. It might also be that institutions change. The tax year may shift, the end-of-year bonus might be divided over three periods, and the timing of children's holidays can change. It may also be that behavior changes. For example, one can imagine different responses to exogenous shocks in different seasons. Also, it may be that certain shocks occur more often in some seasons. All these reasons suggest that seasonal patterns can change over time. In the rest of this section, I will discuss three basic models that can describe time series with changing seasonal patterns.

Seasonal random walk

A simple model that allows the seasonal pattern to change over time is the seasonal random walk, given by

$$y_t = y_{t-S} + \varepsilon_t. \tag{13}$$

It might not immediately be clear from this expression that seasonality changes, and therefore it is useful to consider the S annual time series $Y_{s,T}$. The seasonal random walk implies that for these annual series it holds that

$$Y_{s,T} = Y_{s,T-1} + \varepsilon_{s,T}. \tag{14}$$

Hence, each seasonal series follows a random walk, and due to the innovations, the annual series may switch position, such that "summer becomes winter".

From graphs it is not easy to discern whether a series is a seasonal random walk or not. The observable pattern depends on the starting values of the time series, relative to the variance of the error term, see the graphs in Figure 9. When the starting values are very close to each other, seasonal patterns seem to change quite rapidly (the series y) and when the starting values are far apart, the graph of the x series suggests that seasonality is close to constant, at least at first sight.

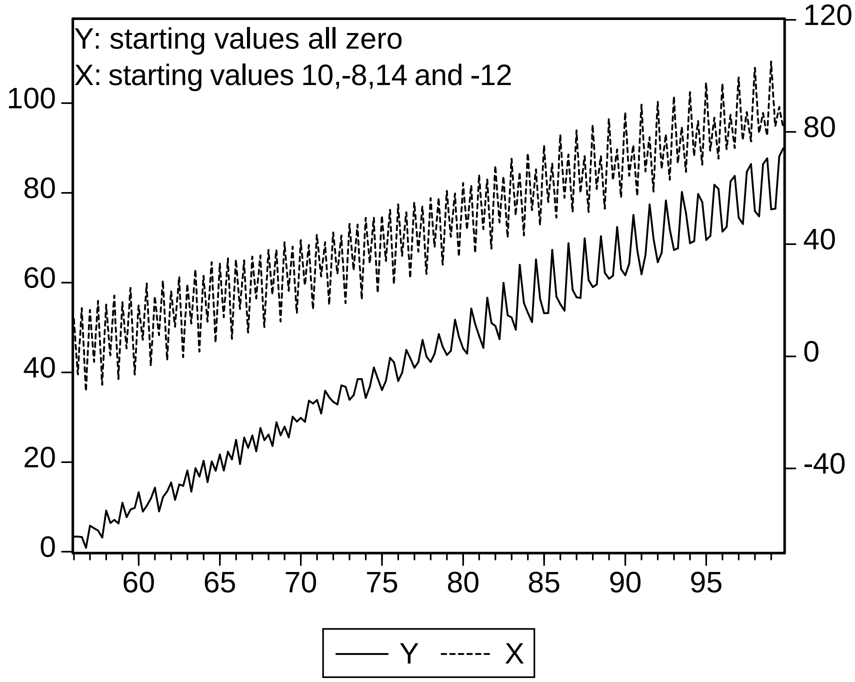

Figure 9: The two quarterly variables are seasonal random walks, with different starting values.

This demonstrates that simply looking at graphs might not be reliable. Here, a look at the autocorrelations of the series $(1 - L^S)y_t$ could be helpful. In the case of a seasonal random walk, the estimated autocorrelation function should look like a white noise series, while such a function for a $(1 - L^S)$ transformed deterministic seasonality series, would result in an error process like $u_t - u_{t-4}$, with a theoretical fourth order autocorrelation of value -0.5, see Franses (1998).

Remarkably, even though the realizations of a seasonal random walk can show substantial within-sample variation, the out-of-sample forecasts are deterministic.

Indeed, at time n, these forecasts are

$$\begin{aligned}
\hat{y}_{n+1} &= y_{n+1-S} \\
\hat{y}_{n+2} &= y_{n+2-S} \\
&\vdots \\
\hat{y}_{n+S} &= y_n \\
\hat{y}_{n+S+1} &= \hat{y}_{n+1} \\
\hat{y}_{n+S+2} &= \hat{y}_{n+2}.
\end{aligned}$$

Another way of allowing for seasonal random walk type changing pattern, is to introduce changing parameters. For example, a subtle form of changing seasonality is described by a time-varying seasonal dummy parameter model. For example, for $S = 4$ this model could look like

$$y_t = \sum_{s=1}^{4} \delta_{s,t} D_{s,t} + u_t, \tag{15}$$

where

$$\delta_{s,t} = \delta_{s,t-S} + \varepsilon_{s,t}. \tag{16}$$

When the variance of $\varepsilon_{s,t} = 0$, the constant parameter model appears. The amount of variation depends on the variance of $\varepsilon_{s,t}$. A illustration is given in Figure 10. Such a model is used as a test vehicle in Canova and Hansen (1995) to diagnose if there is changing seasonality in time series.

Airline model

It can be felt that the seasonal random walk model allows for too much variation in the seasonal pattern. Indeed, allowing each season to be a random walk might introduce too much variation. One way to accommodate this is to introduce a correction, for example by having an error term at the lag that corresponds with the filter with a parameter that comes close to unity. An example is

$$y_t = y_{t-S} + \varepsilon_t + \theta_S \varepsilon_{t-S}, \tag{17}$$

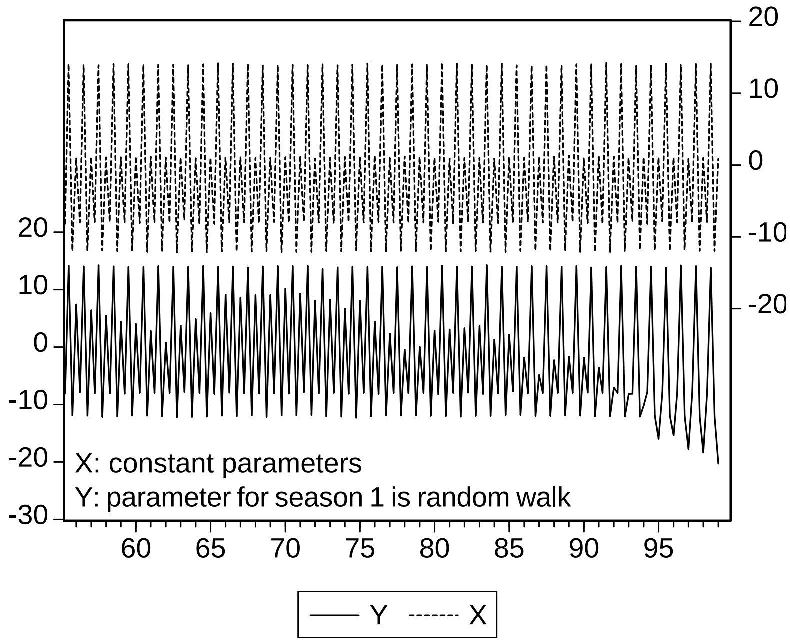

Figure 10: Time series with constant seasonality (x) and with one seasonal dummy parameter as a seasonal random walk (y)

where θ_S can approximate -1. Bell (1987) demonstrates that when $\theta_S = -1$, the model reduces to

$$y_t = \sum_{s=1}^{S} \delta_s D_{s,t} + \varepsilon_t. \tag{18}$$

Clearly, this also gives an opportunity to test if there is constant or changing seasonality.

An often applied model, popularized by Box and Jenkins (1970) and named after its application to monthly airline passenger data, is the airline model. It builds on the above model by considering

$$(1-L)(1-L^S)y_t = (1+\theta_1 L)(1+\theta_S L^S)\varepsilon_t, \tag{19}$$

where it should be noted that

$$(1-L)(1-L^S)y_t = y_t - y_{t-1} - y_{t-S} + y_{t-S-1}. \tag{20}$$

This model is assumed to effectively handle a trend in the data using the filter $(1-L)$ and any changing seasonality using $(1-L^S)$. Strictly speaking, the airline model

Dependent Variable: LC-LC(-1)-LC(-4)+LC(-5)
Method: Least Squares
Date: 04/07/04 Time: 09:48
Sample(adjusted): 1981:2 2001:2
Included observations: 81 after adjusting endpoints
Convergence achieved after 10 iterations
Backcast: 1980:1 1981:1

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	-0.000329	0.000264	-1.246930	0.2162
MA(1)	-0.583998	0.088580	-6.592921	0.0000
SMA(4)	-0.591919	0.090172	-6.564320	0.0000

R-squared	0.444787	Mean dependent var	-5.01E-05
Adjusted R-squared	0.430550	S.D. dependent var	0.016590
S.E. of regression	0.012519	Akaike info criterion	-5.886745
Sum squared resid	0.012225	Schwarz criterion	-5.798061
Log likelihood	241.4132	F-statistic	31.24325
Durbin-Watson stat	2.073346	Prob(F-statistic)	0.000000

Inverted MA Roots	.88	.58	.00+.88i	-.00 -.88i
	-.88			

Figure 11: Airline model estimation results: quarterly log(consumption) in Japan

assumes $S+1$ unit roots. This is due to the fact that the characteristic equation of the AR part, which is,

$$(1-z)(1-z^S)=0, \tag{21}$$

has $S+1$ solutions on the unit circle. For example, if $S=4$ the solutions are (1, 1, -1, i, -i). This implies a substantial amount of random walk like behavior, even though it is corrected to some extent by the $(1+\theta_1 L)(1+\theta_S L^S)\varepsilon_t$ part of the model. In terms of forecasting, it assumes very wide confidence intervals around the point forecasts. On the other hand, the advantages of this model are that it contains only two parameters and that it can describe a wide range of variables, which can be observed from the Eviews output in Figures 11, 12 en 13. The estimated residuals of these models do not obviously indicate mis-specification. On the other hand, it is clear that the roots of the MA polynomial (indicated at the bottom panel of these graphs) are close to the unit circle.

Dependent Variable: LM-LM(-1)-LM(-4)+LM(-5)
Method: Least Squares
Date: 04/07/04 Time: 09:50
Sample(adjusted): 1976:3 2004:1
Included observations: 111 after adjusting endpoints
Convergence achieved after 16 iterations
Backcast: 1975:2 1976:2

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	-0.000192	0.000320	-0.599251	0.5503
MA(1)	0.353851	0.088144	4.014489	0.0001
SMA(4)	-0.951464	0.016463	-57.79489	0.0000

R-squared	0.647536	Mean dependent var	-3.88E-06
Adjusted R-squared	0.641009	S.D. dependent var	0.032112
S.E. of regression	0.019240	Akaike info criterion	-5.036961
Sum squared resid	0.039981	Schwarz criterion	-4.963730
Log likelihood	282.5513	F-statistic	99.20713
Durbin-Watson stat	2.121428	Prob(F-statistic)	0.000000

Inverted MA Roots	.99	.00+.99i	-.00 -.99i	-.35
	-.99			

Figure 12: Airline model estimation results: quarterly log(M1) in Australia

Dependent Variable: LI-LI(-1)-LI(-12)+LI(-13)
Method: Least Squares
Date: 04/07/04 Time: 09:52
Sample(adjusted): 1920:02 2004:02
Included observations: 1009 after adjusting endpoints
Convergence achieved after 11 iterations
Backcast: 1919:01 1920:01

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	2.38E-05	0.000196	0.121473	0.9033
MA(1)	0.378388	0.029006	13.04509	0.0000
SMA(12)	-0.805799	0.016884	-47.72483	0.0000

R-squared	0.522387	Mean dependent var	-0.000108
Adjusted R-squared	0.521437	S.D. dependent var	0.031842
S.E. of regression	0.022028	Akaike info criterion	-4.790041
Sum squared resid	0.488142	Schwarz criterion	-4.775422
Log likelihood	2419.576	F-statistic	550.1535
Durbin-Watson stat	1.839223	Prob(F-statistic)	0.000000

Inverted MA Roots	.98	.85+.49i	.85 -.49i	.49+.85i
	.49 -.85i	.00 -.98i	-.00+.98i	-.38
	-.49 -.85i	-.49+.85i	-.85 -.49i	-.85+.49i
	-.98			

Figure 13: Airline model estimation results: monthly log(industrial production) in the USA

Basic structural model

Finally, a model that takes a position in between seasonal adjustment and the airline model is the Structural Time Series Model, see Harvey (1989). The basic idea is that a time series can be decomposed in various components, which reflect seasonality, trend, cycles and so on. This representation facilitates the explicit consideration of a trend component or a seasonal component, which, if one intends to do so, can be subtracted from the data to get a trend-free or seasonality-free series. Often, a Structural Time Series Model can be written as a seasonal ARIMA type model, and hence, its descriptive quality is close to that of a seasonal ARIMA model.

To illustrate, an example of a structural time series model for a quarterly time series is

$$y_t = \mu_t + s_t + w_t, \quad w_t \sim \mathrm{N}(0, \sigma_w^2) \tag{22}$$

$$(1-L)^2 \mu_t = u_t, \quad u_t \sim \mathrm{N}(0, \sigma_u^2) \tag{23}$$

$$(1 + L + L^2 + L^3) s_t = v_t, \quad v_t \sim \mathrm{N}(0, \sigma_v^2) \tag{24}$$

where the error processes w_t, u_t and v_t are mutually independent, and where the errors are normally and independently distributed. This model contains three unknown parameters, that is, the variances, and of course, also the variables μ_t, s_t and the error term w_t are unobserved. The interest is in estimating the trend and the seasonal component, which are associated with these features due to the lag polynomials $(1-L)^2$ and $(1 + L + L^2 + L^3)$, respectively. For parameter estimation one relies on Kalman filter techniques.

Combining the three equations gives that this y_t can also be described by

$$(1-L)(1-L^4) y_t = \zeta_t \tag{25}$$

where ζ_t is a moving average process of order 5. Notice that this description comes close to that of the airline model above. This can be substantiated by deriving the autocovariances γ_k, $k = 0, 1, 2, \ldots$, of ζ_t, which are

$$\gamma_0 = 4\sigma_u^2 + 6\sigma_v^2 + 4\sigma_w^2 \quad (26)$$

$$\gamma_1 = 3\sigma_u^2 - 4\sigma_v^2 - 2\sigma_w^2 \quad (27)$$

$$\gamma_2 = 2\sigma_u^2 + \sigma_v^2 \quad (28)$$

$$\gamma_3 = \sigma_u^2 + \sigma_w^2 \quad (29)$$

$$\gamma_4 = -2\sigma_w^2 \quad (30)$$

$$\gamma_5 = \sigma_w^2 \quad (31)$$

$$\gamma_j = 0 \quad \text{for } j = 6, 7, \ldots. \quad (32)$$

The only formal differences between this model and the airline model is that the latter implies a zero-valued third order autocovariance, and that $\gamma_3 = \gamma_5$.

Conclusion

There are various ways to describe a time series (and use that description for forecasting) with constant or changing seasonal variation. In the next section, more models will be proposed for describing changing seasonality.

In practice, of course, one needs to make a choice. To make such a choice, one usually zooms in on the key differences between the various models, and these mainly concern the number of unit roots assumed in the autoregressive or moving average polynomials. When these roots are associated with seasonal fluctuations, like for example $(1 + L)$ and $(1 + L^2)$, these roots are called seasonal unit roots. The next section will say more about this selection of models.

To conclude, an important message of this section is that model choice can not just be guided by an informal look at graphs or at autocorrelation functions. Various models can generate data that look very similar, and hence more formal tests are needed.

4 Advanced Models

The previous section reviewed various basic models without checking whether these filters match with the properties of the data. These filters assume a certain amount of unit roots, and it seems sensible to test whether these roots are present or not. In this section I discuss models that allow for a more sophisticated description of seasonal patterns, while allowing for the possible presence of zero frequency trends. Next, I will discuss models that allow the trend and season variation to be intertwined.

Seasonal unit roots

A time series variable has a non-seasonal unit root if the autoregressive polynomial (of the model that best describes this variable), contains the component $1-L$, and the moving-average part does not. For example, the model $y_t = y_{t-1} + \varepsilon_t$ has a first-order autoregressive polynomial $1-L$, as it can be written as $(1-L)y_t = \varepsilon_t$. Hence, data that can be described by the random walk model are said to have a unit root. The same holds of course for the model $y_t = \mu + y_{t-1} + \varepsilon_t$, which is a random walk with drift. Solving this last model to the first observation, that is,

$$y_t = y_0 + \mu t + \varepsilon_t + \varepsilon_{t-1} + \cdots + \varepsilon_1 \tag{33}$$

shows that such data also have a deterministic trend. Due to the summation of the error terms, it is possible that data diverge from the overall trend μt for a long time, and hence one could conclude from a graph that there are all kinds of trends with directions that vary from time to time. Therefore, such data are sometimes said to have a stochastic trend.

The unit roots in seasonal data, which can be associated with changing seasonality, are seasonal unit roots, see Hylleberg et al. (1990) [HEGY]. For quarterly data, these roots are -1, i, and -i. For example, data generated from the model $y_t = -y_{t-1} + \varepsilon_t$ would display seasonality, but if one were to make graphs with the split seasonals, then one could observe that the quarterly data within a year shift places quite frequently. Similar observations hold for the model $y_t = -y_{t-2} + \varepsilon_t$, which can be written as $(1+L^2)y_t = \varepsilon_t$, where the autoregressive polynomial $1+L^2$ corresponds to the seasonal unit roots i and -i, as these two values solve the equation $1+z^2 = 0$.

Testing for seasonal unit roots

In contrast to simply imposing (seasonal) unit roots, one can also test whether they are present or not. The most commonly used method for this purpose is the HEGY method. For quarterly data it amounts to a regression of $\Delta_4 y_t$ on deterministic terms like an intercept, seasonal dummies, a trend and seasonal trends and on $(1 + L + L^2 + L^3)y_{t-1}$, $(-1 + L - L^2 + L^3)y_{t-1}$, $-(1 + L^2)y_{t-1}$, $-(1 + L^2)y_{t-2}$, and on lags of $\Delta_4 y_t$. A $t-$test is used to examine the significance of the parameter for $(1+L+L^2+L^3)y_{t-1}$, and similarly, a $t-$test for $(-1+L-L^2+L^3)y_{t-1}$ and a joint $F-$test for $-(1+L^2)y_{t-1}$ and $-(1+L^2)y_{t-2}$. An insignificant test value indicates the presence of the associated root(s), which are 1, −1, and the pair i, −i, respectively. Asymptotic theory for the tests is developed in Hylleberg et al. (1990), and useful extensions are put forward in Smith and Taylor (1998).

When including deterministic terms, it is important to recall the discussion in Section 2, concerning the seasonal dummies. Indeed, when the seasonal dummies are included unrestrictedly, it is possible that the time series (under the null hypothesis of seasonal unit roots) can display seasonally varying deterministic trends. Hence, when checking for example whether the $(1+L)$ filter can be imposed, one also needs to impose that the α_2 parameter for $\cos(\pi t)$ in (11) equals zero. The preferable way to include deterministics therefore is to include the alternating dummy variables $D_{1,t} - D_{2,t} + D_{3,t} - D_{4,t}$, $D_{1,t} - D_{3,t}$, and $D_{2,t} - D_{4,t}$. And, for example, under the null hypothesis that there is a unit root −1, the parameter for the first alternating dummy should also be zero. These joint tests extend the work of Dickey and Fuller (1981), and are discussed in Smith and Taylor (1999). When models are created for panels of time series or for multivariate series, as I will discuss below, these restrictions on the deterministics (based on the sine-cosine notation) are important too.

Kawasaki and Franses (2003) propose to detect seasonal unit roots within the context of a structural time series model. They rely on model selection criteria. Using Monte Carlo simulations, they show that the method works well. They illustrate their approach for several quarterly macroeconomic time series variables.

Seasonal cointegration

In case two or more seasonally observed time series have seasonal unit roots, one may be interested in testing for common seasonal unit roots, that is, in testing for seasonal cointegration. If these series have such roots in common, they will have common changing seasonal patterns.

Engle et al. (1993) [EGHL] propose a two-step method to see if there is seasonal cointegration. When two series $y_{1,t}$ and $y_{2,t}$ have a common non-seasonal unit root, then the series u_t defined by

$$u_t = (1 + L + L^2 + L^3)y_{1,t} - \alpha_1(1 + L + L^2 + L^3)y_{2,t} \tag{34}$$

does not need the $(1 - L)$ filter to become stationary. Seasonal cointegration at the annual frequency π, corresponding to unit root -1, implies that

$$v_t = (1 - L + L^2 - L^3)y_{1,t} - \alpha_2(1 - L + L^2 - L^3)y_{2,t} \tag{35}$$

does not need the $(1 + L)$ differencing filter. And, seasonal cointegration at the annual frequency $\pi/2$, corresponding to the unit roots $\pm$i, means that

$$w_t = (1 - L^2)y_{1,t} - \alpha_3(1 - L^2)y_{2,t} - \alpha_4(1 - L^2)y_{1,t-1} - \alpha_5(1 - L^2)y_{2,t-1} \tag{36}$$

does not have the unit roots $\pm$i. In case all three u_t, v_t and w_t do not have the relevant unit roots, the first equation of a simple version of a seasonal cointegration model is

$$\Delta_4 y_{1,t} = \gamma_1 u_{t-1} + \gamma_2 v_{t-1} + \gamma_3 w_{t-2} + \gamma_4 w_{t-3} + \varepsilon_{1,t}, \tag{37}$$

where γ_1 to γ_4 are error correction parameters. The test method proposed in EGHL is a two-step method, similar to the Engle-Granger (1987) approach to non-seasonal time series.

Seasonal cointegration in a multivariate time series Y_t can also be analyzed using an extension of the Johansen approach, see Johansen and Schaumburg (1999), Franses and Kunst (1999a). It amounts to testing the ranks of matrices that correspond to variables which are transformed using the filters to remove the roots 1, -1 or $\pm$ i. More precise, consider the $(m \times 1)$ vector process Y_t, and assume that it can be described by the VAR(p) process

$$Y_t = \Theta D_t + \Phi_1 Y_{t-1} + \cdots + \Phi_p Y_{t-p} + e_t, \tag{38}$$

where D_t is the (4×1) vector process $D_t = (D_{1,t}, D_{2,t}, D_{3,t}, D_{4,t})'$ containing the seasonal dummies, and where Θ is an $(m \times 4)$ parameter matrix. Similar to the Johansen (1995) approach and conditional on the assumption that $p > 4$, the model can be rewritten as

$$\begin{aligned} \Delta_4 Y_t &= \Theta D_t + \Pi_1 Y_{1,t-1} && (39) \\ &+ \Pi_2 Y_{2,t-1} + \Pi_3 Y_{3,t-2} + \Pi_4 Y_{3,t-1} && (40) \\ &+ \Gamma_1 \Delta_4 Y_{t-1} + \cdots + \Gamma_{p-4} \Delta_4 Y_{t-(p-4)} + e_t, \end{aligned}$$

where

$$\begin{aligned} Y_{1,t} &= (1 + L + L^2 + L^3) Y_t \\ Y_{2,t} &= (1 - L + L^2 - L^3) Y_t \\ Y_{3,t} &= (1 - L^2) Y_t. \end{aligned}$$

This is a multivariate extension of the univariate HEGY model. The ranks of the matrices Π_1, Π_2, Π_3 and Π_4 determine the number of cointegration relations at each of the frequencies. Again, it is important to properly account for the deterministics, in order not to have seasonally diverging trends, see Franses and Kunst (1999a) for a solution.

Periodic models

An alternative class of models is the periodic autoregression. Consider a univariate time series y_t, which is observed quarterly for N years. It is assumed that $n = 4N$. A periodic autoregressive model of order p [PAR(p)] can be written as

$$y_t = \mu_s + \phi_{1s} y_{t-1} + \cdots + \phi_{ps} y_{t-p} + \varepsilon_t, \tag{41}$$

or

$$\phi_{p,s}(L) y_t = \mu_s + \varepsilon_t, \tag{42}$$

with

$$\phi_{p,s}(L) = 1 - \phi_{1s}L - \cdots - \phi_{ps}L^p, \tag{43}$$

where μ_s is a seasonally-varying intercept term. The $\phi_{1s}, \ldots, \phi_{ps}$ are autoregressive parameters up to order p_s which may vary with the season s, where $s = 1, 2, 3, 4$. For ε_t it can be assumed it is a standard white noise process with constant variance σ^2, but that may be relaxed by allowing ε_t to have seasonal variance σ_s^2. As some ϕ_{is}, $i = 1, 2, \ldots, p$, can take zero values, the order p is the maximum of all p_s.

Multivariate representation

In general, the PAR(p) process can be rewritten as an AR(P) model for the (4×1) vector process $Y_T = (Y_{1,T}, Y_{2,T}, Y_{3,T}, Y_{4,T})'$, $T = 1, 2, \ldots, N$, where $Y_{s,T}$ is the observation of y_t in season s of year T. The model is then

$$\Phi_0 Y_T = \mu + \Phi_1 Y_{T-1} + \cdots + \Phi_P Y_{T-P} + \varepsilon_T, \tag{44}$$

or

$$\Phi(L) Y_T = \mu + \varepsilon_T, \tag{45}$$

with

$$\Phi(L) = \Phi_0 - \Phi_1 L - \cdots - \Phi_P L^P, \tag{46}$$

$\mu = (\mu_1, \mu_2, \mu_3, \mu_4)'$, $\varepsilon_T = (\varepsilon_{1,T}, \varepsilon_{2,T}, \varepsilon_{3,T}, \varepsilon_{4,T})'$, and $\varepsilon_{s,T}$ is the observation on the error process ε_t in season s of year T. The lag operator L applies to data at frequencies t and to T, that is, $Ly_t = y_{t-1}$ and $LY_T = Y_{T-1}$. The $\Phi_0, \Phi_1, \ldots, \Phi_P$ are 4×4 parameter matrices with elements

$$\Phi_0[i,j] = \begin{cases} 1 & i = j, \\ 0 & j > i, \\ -\phi_{i-j,i} & i < j, \end{cases} \tag{47}$$

$$\Phi_k[i,j] = \phi_{i+4k-j,i}, \tag{48}$$

for $i = 1, 2, 3, 4$, $j = 1, 2, 3, 4$, and $k = 1, 2, \ldots, P$. For P it holds that $P = 1 + [(p-1)/4]$, where $[\,\cdot\,]$ is the integer function. Hence, when p is less than or equal to 4, the value of P is 1.

As Φ_0 is a lower triangular matrix, model (44) is a recursive model. This means that $Y_{4,T}$ depends on $Y_{3,T}$, $Y_{2,T}$, and $Y_{1,T}$, and on all variables in earlier years.

Similarly, $Y_{3,T}$ depends on $Y_{2,T}$ and $Y_{1,T}$, and $Y_{2,T}$ on $Y_{1,T}$ and on all observations in past years. As an example, consider the PAR(2) process

$$y_t = \phi_{1s} y_{t-1} + \phi_{2s} y_{t-2} + \varepsilon_t, \tag{49}$$

which can be written as

$$\Phi_0 Y_T = \Phi_1 Y_{T-1} + \varepsilon_T, \tag{50}$$

with

$$\Phi_0 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ -\phi_{12} & 1 & 0 & 0 \\ -\phi_{23} & -\phi_{13} & 1 & 0 \\ 0 & -\phi_{24} & -\phi_{14} & 1 \end{pmatrix} \quad \text{and} \quad \Phi_1 = \begin{pmatrix} 0 & 0 & \phi_{21} & \phi_{11} \\ 0 & 0 & 0 & \phi_{22} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{51}$$

A useful representation is based on the possibility of decomposing a non-periodic AR(p) polynomial as $(1-\alpha_1 L)(1-\alpha_2 L)\cdots(1-\alpha_p L)$, see Boswijk, Franses and Haldrup (1997) where this representation is used to test for (seasonal) unit roots in periodic models. Note that this can only be done when the solutions to the characteristic equation for this AR(p) polynomial are all real-valued. Similar results hold for the multivariate representation of a PAR(p) process, and it can be useful to rewrite (44) as

$$\prod_{i=1}^{P} \Xi_i(L) Y_T = \mu + \varepsilon_T, \tag{52}$$

where the $\Xi_i(L)$ are 4×4 matrices with elements which are polynomials in L.

A simple example is the PAR(2) process

$$\Xi_1(L)\Xi_2(L) Y_T = \varepsilon_T, \tag{53}$$

with

$$\Xi_1(L) = \begin{pmatrix} 1 & 0 & 0 & -\beta_1 L \\ -\beta_2 & 1 & 0 & 0 \\ 0 & -\beta_3 & 1 & 0 \\ 0 & 0 & -\beta_4 & 1 \end{pmatrix},$$

$$\Xi_2(L) = \begin{pmatrix} 1 & 0 & 0 & -\alpha_1 L \\ -\alpha_2 & 1 & 0 & 0 \\ 0 & -\alpha_3 & 1 & 0 \\ 0 & 0 & -\alpha_4 & 1 \end{pmatrix}. \tag{54}$$

This PAR(2) model can be written as

$$(1 - \beta_s L)(1 - \alpha_s L)y_t = \mu_s + \varepsilon_t, . \tag{55}$$

or

$$y_t - \alpha_s y_{t-1} = \mu_s + \beta_s(y_{t-1} - \alpha_{s-1}y_{t-2}) + \varepsilon_t, \tag{56}$$

as, and this is quite important, the lag operator L also operates on α_s, that is, $L\alpha_s = \alpha_{s-1}$ for all $s = 1, 2, 3, 4$ and with $\alpha_0 = \alpha_4$. The characteristic equation is

$$|\Xi_1(z)\Xi_2(z)| = 0, \tag{57}$$

and this is equivalent to

$$(1 - \beta_1\beta_2\beta_3\beta_4 z)(1 - \alpha_1\alpha_2\alpha_3\alpha_4 z) = 0. \tag{58}$$

So, the PAR(2) model has one unit root when either $\beta_1\beta_2\beta_3\beta_4 = 1$ or $\alpha_1\alpha_2\alpha_3\alpha_4 = 1$, and has at most two unit roots when both products equal unity. The case where $\alpha_1\alpha_2\alpha_3\alpha_4 = 1$ while not all α_s are equal to 1 is called periodic integration, see Osborn (1988) and Franses (1996). Tests for periodic integration are developed in Boswijk and Franses (1996) for the case without allowing for seasonal unit roots, and in Boswijk, Franses and Haldrup (1997) for the case where seasonal unit roots can also occur. Obviously, the maximum number of unity solutions to the characteristic equation of a PAR(p) process is equal to p.

The analogy of a univariate PAR process with a multivariate time series process can be used to derive explicit formulae for one- and multi-step ahead forecasting, see Franses (1996). It should be noted that then the one-step ahead forecasts concern one-year ahead forecasts for all four $Y_{s,T}$ series. For example, for the model $Y_T = \Phi_0^{-1}\Phi_1 Y_{T-1} + \omega_T$, where $\omega_T = \Phi_0^{-1}\varepsilon_T$, the forecast for $N + 1$ is $\hat{Y}_{N+1} = \Phi_0^{-1}\Phi_1 Y_N$.

Finally, one may wonder what the consequences are of fitting non-periodic models to periodic data. One consequence is that such a non-periodic model requires many lags, see Franses and Paap (2004) and Del Barrio Castro and Osborn (2004). For example, a PAR(1) model can be written as

$$y_t = \alpha_{s+3}\alpha_{s+2}\alpha_{s+1}\alpha_s y_{t-4} + \varepsilon_t + \alpha_{s+3}\varepsilon_{t-1} + \alpha_{s+3}\alpha_{s+2}\varepsilon_{t-2} + \alpha_{s+3}\alpha_{s+2}\alpha_{s+1}\varepsilon_{t-3}. \tag{59}$$

As $\alpha_{s+3}\alpha_{s+2}\alpha_{s+1}\alpha_s$ is equal for all seasons, the AR parameter at lag 4 in a non-periodic model is truly non-periodic, but of course, the MA part is not. The MA part of this model is of order 3. If one estimates a non-periodic MA model for these data, the MA parameter estimates will attain an average value of the α_{s+3}, $\alpha_{s+3}\alpha_{s+2}$, and $\alpha_{s+3}\alpha_{s+2}\alpha_{s+1}$ across the seasons. In other words, one might end up considering an ARMA(4,3) model for PAR(1) data. And, if one decides not to include an MA part in the model, one usually needs to increase the order of the autoregression to whiten the errors. This suggests that higher-order AR models might fit to low-order periodic data. When $\alpha_{s+3}\alpha_{s+2}\alpha_{s+1}\alpha_s = 1$, one has a high-order AR model for the Δ_4 transformed time series. In sum, there seems to be a trade-off between seasonality in parameters and short lags against no seasonality in parameters and longer lags.

Conclusion

There is a voluminous literature on formally testing for seasonal unit roots in non-periodic data and on testing for unit roots in periodic autoregressions. There are many simulation studies to see which method is best. Also, there are many studies which examine whether imposing seasonal unit roots or not, or assuming unit roots in periodic models or not, lead to better forecasts. This also extends to the case of multivariate series, where these models allow for seasonal cointegration or for periodic cointegration. An example of a periodic cointegration model is

$$\Delta_4 y_t = \gamma_s(y_{t-4} - \beta_s x_{t-4}) + \varepsilon_t, \tag{60}$$

where γ_s and β_s can take seasonally varying values, see Boswijk and Franses (1995).

For example, Löf and Franses (2001) analyze periodic and seasonal cointegration models for bivariate quarterly observed time series in an empirical forecasting study, as well as a VAR model in first differences, with and without cointegration restrictions, and a VAR model in annual differences. The VAR model in first differences without cointegration is best if one-step ahead forecasts are considered. For longer forecast horizons, the VAR model in annual differences is better. When comparing periodic versus seasonal cointegration models, the seasonal cointegration models tend to yield better forecasts. Finally, there is no clear indication that multiple equations methods improve on single equation methods.

To summarize, tests for periodic variation in the parameters and for unit roots allow one to make a choice between the various models for seasonality. There are many tests around, and they are all easy to use. Not unexpectedly, models and methods for data with frequencies higher than 12 can become difficult to use in practice, see Darne (2004) for a discussion of seasonal cointegration in monthly series. Hence, here there is a need for more future research.

5 Recent advances

This section deals with a few recent developments in the area of forecasting seasonal time series. These are (i) seasonality in panels of time series, (ii) periodic models for financial series, and (iii) nonlinear models for seasonal time series.

Seasonality in panels of time series

The search for common seasonal patterns can lead to a dramatic reduction in the number of parameters, see Engle and Hylleberg (1996). One way to look for common patterns across the series $y_{i,t}$, where $i = 1, 2, .., I$, and I can be large, is to see of the series have common dynamics or common trends. Alternatively, one can examine if series have common seasonal deterministics.

As can be understood from the discussion on seasonal unit roots, before one can say something about (common) deterministic seasonality, one first has to decide on the number of seasonal unit roots. The HEGY test regression for seasonal unit roots is

$$\Phi_{p_i}(L)\Delta_4 y_{i,t} = \mu_{i,t} + \rho_{i,1}S(L)y_{i,t-1} + \rho_{i,2}A(L)y_{i,t-1} + \rho_{i,3}\Delta_2 y_{i,t-1} + \rho_{i,4}\Delta_2 y_{i,t-2} + \varepsilon_t, \quad (61)$$

and now it is convenient to take

$$\mu_{i,t} = \mu_i + \alpha_{1,i}\cos(\pi t) + \alpha_{2,i}\cos(\frac{\pi t}{2}) + \alpha_{3,i}\cos(\frac{\pi(t-1)}{2}) + \delta_i t, \quad (62)$$

and where Δ_k is the $k-$th order differencing filter, $S(L)y_{i,t} = (1 + L + L^2 + L^3)y_{i,t}$ and $A(L)y_{i,t} = -(1 - L + L^2 - L^3)y_{i,t}$. The model assumes that each series $y_{i,t}$

can be described by a $(p_i+4)-$th order autoregression. Smith and Taylor (1999) and Franses and Kunst (1999a,b) argue that an appropriate test for a seasonal unit root at the bi-annual frequency is now given by a joint $F-$test for $\rho_{i,2}$ and $\alpha_{1,i}$. An appropriate test for the two seasonal unit roots at the annual frequency is then given by a joint $F-$test for $\rho_{3,i}$, $\rho_{4,i}$, $\alpha_{2,i}$ and $\alpha_{3,i}$. Franses and Kunst (1999b) consider these F-tests in a model where the autoregressive parameters are pooled over the equations, hence a panel HEGY test. The power of this panel test procedure is rather large. Additionally, once one has taken care off seasonal unit roots, these authors examine if two or more series have the same seasonal deterministic fluctuations. This can be done by testing for cross-equation restrictions.

Periodic GARCH

Periodic models might also be useful for financial time series. They can be used not only to describe the so-called day-of-the-week effects, but also to describe the apparent differences in volatility across the days of the week. Bollerslev and Ghysels (1996) propose a periodic generalized autoregressive conditional heteroskedasticity (PGARCH) model. Adding a periodic autoregression for the returns to it, one has a PAR(p)-PGARCH(1,1) model, which for a daily observed financial time series y_t, $t=1,\ldots,n=5N$, can be represented by

$$\begin{aligned} x_t &= y_t - \sum_{s=1}^{5}\left(\mu_s + \sum_{i=1}^{p}\phi_{is}y_{t-i}\right)D_{s,t} \\ &= \sqrt{h_t}\eta_t \end{aligned} \tag{63}$$

with $\eta_t \sim \mathrm{N}(0,1)$ for example, and

$$h_t = \sum_{s=1}^{5}(\omega_s + \psi_s x_{t-1}^2)D_{s,t} + \gamma h_{t-1}, \tag{64}$$

where the x_t denotes the residual of the PAR model for y_t, and where $D_{s,t}$ denotes a seasonal dummy for the day of the week, that is, $s=1,2,3,4,5$.

In order to investigate the properties of the conditional variance model, it is useful to define $z_t = x_t^2 - h_t$, and to write it as

$$x_t^2 = \sum_{s=1}^{5}(\omega_s + (\psi_s+\gamma)x_{t-1}^2)D_{s,t} + z_t - \gamma z_{t-1}. \tag{65}$$

This ARMA process for x_t^2 contains time-varying parameters $\psi_s+\gamma$ and hence strictly speaking, it is not a stationary process. To investigate the stationarity properties of x_t^2, (65) can be written in a time-invariant representation. Franses and Paap (2000) successfully fit such a model to the daily S&P 500 index, and even find that

$$\Pi_{s=1}^{5}(\psi_s + \gamma) = 1. \tag{66}$$

In other words, they fit a periodically integrated GARCH model.

Models of seasonality and nonlinearity

It is well known that a change in the deterministic trend properties of a time series y_t is easily mistaken for the presence of a unit root. In a similar vein, if a change in the deterministic seasonal pattern is not detected, one might well end up imposing seasonal unit roots, see Ghysels (1994), Smith and Otero (1997), Franses, Hoek and Paap (1997) and Franses and Vogelsang (1998).

Changes in deterministic seasonal patterns usually are modelled by means of one-time abrupt and discrete changes. However, when seasonal patterns shift due to changes in technology, institutions and tastes, for example, these changes may materialize only gradually. This suggests that a plausible description of time-varying seasonal patterns is

$$\phi(L)\Delta_1 y_t = \sum_{s=1}^{4} \delta_{1,s} D_{s,t}(1 - G(t;\gamma,c)) + \sum_{s=1}^{4} \delta_{2,s} D_{s,t} G(t;\gamma,c) + \varepsilon_t, \tag{67}$$

where $G(t;\gamma,c)$ is the logistic function

$$G(s_t;\gamma,c) = \frac{1}{1+\exp\{-\gamma(s_t - c)\}}, \qquad \gamma > 0. \tag{68}$$

As s_t increases, the logistic function changes monotonically from 0 to 1, with the change being symmetric around the location parameter c, as $G(c-z;\gamma,c) = 1-G(c+z;\gamma,c)$ for all z. The slope parameter γ determines the smoothness of the change. As $\gamma \to \infty$, the logistic function $G(s_t;\gamma,c)$ approaches the indicator function $I[s_t > c]$, whereas if $\gamma \to 0$, $G(s_t;\gamma,c) \to 0.5$ for all values of s_t. Hence, by taking $s_t = t$, the model takes an "intermediate" position in between deterministic seasonality and stochastic trend seasonality.

Nonlinear models with smoothly changing deterministic seasonality are proposed in Franses and van Dijk (2004). These authors examine the forecasting performance of various models for seasonality and nonlinearity for quarterly industrial production series of 18 OECD countries. They find that the accuracy of point forecasts varies widely across series, across forecast horizons and across seasons. However, in general, linear models with fairly simple descriptions of seasonality outperform at short forecast horizons, whereas nonlinear models with more elaborate seasonal components dominate at longer horizons. Simpler models are also preferable for interval and density forecasts at short horizons. Finally, none of the models is found to be the best and hence, forecast combination is worthwhile.

To summarize, recent advances in modeling and forecasting seasonal time series focus at (i) models for panels of time series and at (ii) models which not only capture seasonality, but also conditional volatility and non-linearity, for example. To fully capture all these features is not easy, also as various features may be related. More research is needed in this area too.

6 Conclusion

Forecasting studies show that model specification efforts pay off in terms of performance. Simple models for seasonally differenced data forecast well for one or a few steps ahead. For longer horizons, more involved models are much better. These involved models address seasonality in conjunction with trends, non-linearity and conditional volatility. Much more research is needed to see which models are to be preferred in which situations.

There are at least two well articulated further research issues. The first concerns methods to achieve parsimony. Indeed, seasonal time series models for monthly or weekly data contain a wealth of parameters, and this can reduce efficiency dramatically. The second concerns the analysis of unadjusted data for the situation where people would want to rely on adjusted data, that is, for decisions on turning points. How would one draw inference in case trends, cycles and seasonality are related? Finally, in case one persists in considering seasonally adjusted data, how can we design methods that allow for the best possible interpretation of these data, when the underlying process has all kinds of features?

It is hoped that this chapter has contributed to get an interest in modeling and forecasting seasonal time series data, and that it has stimulated an interest in doing further research in this fascinating area.

References

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

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

Bollerslev, T. and E. Ghysels (1996), Periodic Autoregressive Conditional Heteroscedasticity, *Journal of Business and Economic Statistics*, 14, 139-151.

Boswijk, H. P. and P. H. Franses (1995), Periodic Cointegration – Representation and Inference, *Review of Economics and Statistics*, 77, 436-454.

Boswijk, H. P. and P. H. Franses (1996), Unit Roots in Periodic Autoregressions, *Journal of Time Series Analysis*, 17, 221-245.

Boswijk, H. P., P. H. Franses and N. Haldrup (1997) Multiple Unit Roots in Periodic Autoregression, *Journal of Econometrics*, 80, 167-193.

Brendstrup, B., S. Hylleberg, M.O. Nielsen, L. L. Skippers, and L. Stentoft (2004), Seasonality in Economic Models, *Macroeconomic Dynamics*, 8, 362-394.

Canova, F. and E. Ghysels (1994), Changes in Seasonal Patterns: Are they Cyclical?, *Journal of Economic Dynamics and Control*, 18, 1143-1171.

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, 237-252.

Darne, O. (2004), Seasonal Cointegration for Monthly Data, *Economics Letters*, 82, 349-356.

Del Barrio Castro, T. and D.R. Osborn (2004), The Consequences of Seasonal Adjustment for Periodic Autoregressive Processes, *Reconometrics Journals*, to appear.

Dickey, D.A. and W.A. Fuller (1981), Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root, *Econometrica*, 49, 1057-1072.

Engle, R.F. and C.W.J. Granger (1987), Cointegration and Error Correction: Representation, Estimation, and Testing, *Econometrica*, 55, 251-276.

Engle, R.F. and C.W.J. Granger (1991, eds.), *Long-Run Economic Relationships: Readings in Cointegration*, Oxford: Oxford University Press.

Engle, R. F., C. W. J. Granger, S. Hylleberg and H. S. Lee (1993), Seasonal Cointegration: The Japanese Consumption Function, *Journal of Econometrics*, 55, 275-298.

Engle, R.F. and S. Hylleberg (1996), Common Seasonal Features: Global Unemployment, *Oxford Bulletin of Economics and Statistics*, 58, 615-630.

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

Franses, P. H. (1991), A Multivariate Approach to Modeling Univariate Seasonal Time Series, Econometric Institute Report 9101, Erasmus University Rotterdam.

Franses, P. H. (1994), A Multivariate Approach to Modeling Univariate Seasonal

Time Series, *Journal of Econometrics*, 63, 133-151.

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

Franses, P.H. (1998), *Time Series Models for Business and Economic Forecasting,* Cambridge: Cambridge University Press.

Franses, P.H. (2001), Some Comments on Seasonal Adjustment, *Revista De Economia del Rosario (Bogota, Colombia),* 4, 9-16.

Franses, P. H., H. Hoek and R. Paap (1997), Bayesian Analysis of Seasonal Unit Roots and Seasonal Mean Shifts, *Journal of Econometrics*, 78, 359-380.

Franses, P. H. and S. Hylleberg and H. S. Lee (1995), Spurious Deterministic Seasonality, *Economics Letters*, 48, 249-256.

Franses, P. H. and R.M. Kunst (1999a), On the Role of Seasonal Intercepts in Seasonal Cointegration, *Oxford Bulletin of Economics and Statistics*, 61, 409-433.

Franses, P.H. and R.M. Kunst (1999b), Testing Common Deterministic Seasonality, with an Application to Industrial Production, Econometric Institute Report 9905, Erasmus University Rotterdam.

Franses, P. H. and R. Paap (2000), Modelling Day-of-the-week Seasonality in the S&P 500 Index, *Applied Financial Economics*, 10, 483-488.

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

Farnses, P.H. and D.J.C. van Dijk (2000), *Non-linear Time Series Models in Em-*

pirical Finance, Cambridge: Cambridge University Press.

Franses, P.H. and D.J.C. van Dijk (2004), The Forecasting Performance of Various Models for Seasonality and Nonlinearity for Quarterly Industrial Production, *International Journal of Forecasting*, to appear.

Franses, P. H. and T. J. Vogelsang (1998), On Seasonal Cycles, Unit Roots and Mean Shifts, *Review of Economics and Statistics*, 80, 231-240.

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

Ghysels, E., C.W.J. Granger and P.L. Siklos (1996), Is Seasonal Adjustment a Linear or a Nnnlinear Data-Filtering Process? (with Discussion), *Journal of Business and Economic Statistics* 14, 374-397.

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

Granger, C.W.J. and T. Terasvirta (1993), *Modelling Non-Linear Economic Relationships*, Oxford: Oxford University Press.

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

Haugen, R.A. and J. Lakonishok (1987), *The Incredible January Effect: The Stock Market's Unsolved Mystery*, New York: McGraw-Hill.

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

Hylleberg, S. (1992), *Modelling Seasonality*, Oxford: Oxford University Press.

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

Johansen, S. (1995), *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*, Oxford: Oxford University Press.

Johansen, S. and E. Shaumburg, (1999), Likelihood Analysis of Seasonal Cointegration, *Journal of Econometrics*, 88, 301-339.

Kawasaki, Y. and P.H. Franses (2003), Detecting Seasonal Unit Roots in a Structural Time Series Model, *Journal of Applied Statistics*, 30, 373-387.

Koopman, S.J. and P.H. Franses (2003), Constructing Seasonally Adjusted Data with Time-Varying Confidence Intervals, *Oxford Bulletin of Economics and Statistics*, 64, 509-526.

Löf, M. and P. H. Franses (2001), On Forecasting Cointegrated Seasonal Time Series, *International Journal of Forecasting*, 17, 607-621.

Maravall, A. (1995), Unobserved Components in Economic Time Series, in H. Pesaran, P. Schmidt and M. Wickens (eds.), *Handbook of Applied Econometrics* (Volume 1), Oxford: Basil Blackwell.

Ooms, M. and P.H. Franses (1997), On Periodic Correlations between Estimated Seasonal and Nonseasonal Components in German and US Unemployment, *Journal of Business and Economic Statistics* **15**, 470-481.

Osborn, D. R. (1988), Seasonality and Habit Persistence in a Life-Cycle Model of Consumption, *Journal of Applied Econometrics*, 3, 255-266.

Osborn, D. R. (1990), A Survey of Seasonality in UK Macroeconomic Variables, *International Journal of Forecasting*, 6, 327-336.

Osborn, D. R. (1991), The Implications of Periodically Varying Coefficients for Seasonal Time-Series Processes, *Journal of Econometrics*, 48, 373-384.

Osborn, D. R. and P. M. M. Rodrigues (2001), Asymptotic Distributions of Seasonal Unit Root Tests: A Unifying Approach, *Econometric Reviews*, 21, 221-241.

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

Smith, J. and J. Otero (1997), Structural Breaks and Seasonal Integration, *Economics Letters*, 56, 13-19.

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.

Smith, R. J. and A. M. R. Taylor (1999), Likelihood Ratio Tests for Seasonal Unit Roots, *Journal of Time Series Analysis*, 20, 453-476.