Seasonality, Forecast Extensions and Business Cycle Uncertainty

Proietti, Tommaso

SEFEMEQ, University of Rome "Tor Vergata"

21 February 2010

Seasonality, Forecast Extensions and Business Cycle Uncertainty

Tommaso Proietti*
Università di Roma "Tor Vergata"

Abstract

Seasonality is one of the most important features of economic time series. The possibility to abstract from seasonality for the assessment of economic conditions is a widely debated issue. In this paper we propose a strategy for assessing the role of seasonal adjustment on business cycle measurement. In particular, we provide a method for quantifying the contribution to the unreliability of the estimated cycles extracted by popular filters, such as Baxter and King and Hodrick-Prescott. The main conclusion is that the contribution is larger around the turning points of the series and at the extremes of the sample period; moreover, it much more sizeable for highpass filters, like the Hodrick-Prescott filter, which retain to a great extent the high frequency fluctuations in a time series, the latter being the ones that are more affected by seasonal adjustment. If a band-pass component is considered, the effect has reduced size. Finally, we discuss the role of forecast extensions and the prediction of the cycle. For the time series of industrial production considered in the illustration, it is not possible to provide a reliable estimate of the cycle at the end of the sample.

Keywords: Linear filters. Unobserved Components. Seasonal Adjustment. Reliability.

1 Introduction

Seasonality is a prominent characteristic of economic time series that are observed at the monthly or quarterly frequency, such as production, sales and employment. Its adjustment serves a variety of useful purposes and satisfies well established information requirements from the users. Indeed, most of the literature on the business cycle (BC) relies on seasonally adjusted data. The latter simplify both the specification and the estimation of the model for the analyst, who can concentrate directly on the business cycle features of interest. Moreover, seasonality is often considered as a nuisance feature and as the most predictable component of a time series.

*Dipartimento S.E.F. e ME.Q., Via Columbia 2, 00133 Roma, Italy.
E-mail: tommaso.proietti@uniroma2.it.

It is nowadays standard practice among statistical agencies to carry out seasonal adjustment (SA) either using X-12-ARIMA (see Findley et al., 1998) and Tramo-Seats (see Maravall and Gomez, 1996, and Caporello and Maravall, 2004). Another popular method is the structural approach popularized by Harvey (1989) and West and Harrison (1997), which is also a model based method, but with the important difference that the decomposition of the time series into unobserved components is specified and estimated at the outset, rather than derived ex-post from an estimated reduced form.

As business cycle measurement and analysis are most often operated on seasonally adjusted series, a fundamental issue is whether seasonal adjustment affects the main stylized facts concerning the business cycle. This paper addresses the issue of its impact on the reliability of cycle estimates obtained by ad-hoc filters.

The paper is organized as follows: section 2 reviews the cycle filters considered in this paper, namely the Baxter and King (1999) filter and the Hodrick-Prescott (1997) filter, which are perhaps the most popular filters, and defines the cycle estimator based on the seasonally adjusted series. Section 3 derives the conditional mean square error of the cycle estimators and a relative measure of its reliability. The evaluation of these quantities using conditional simulation, when a parametric seasonal adjustment method is adopted, is the topic of section 4. Section 5 provides some empirical illustrations. Forecast extensions are discussed in section 6 and in section 7 we draw our conclusions.

2 Cycle filters and Seasonality

The business cycle is often measured by applying an ad hoc filter to seasonally adjusted time series. The aim of the paper is to assess the contribution of seasonal adjustment to the uncertainty of the measurement. A well-known filter has been popularized by Baxter and King (1999): this is a band-pass filter that aims at selecting the fluctuations with a specified range of periodicities, namely those ranging from one and a half to eight years. Thus, if s is the number of observations in a year, the fluctuations with periodicity between $1.5s$ and $8s$ are included. Given the two business cycle frequencies, $\omega_{c1} = 2\pi/(8s)$ and $\omega_{c2} = 2\pi/(1.5s)$, the BK filter cycle filter is

$$\mathrm{w}_{bp}(L) = \frac{\omega_{c2} - \omega_{c1}}{\pi} + \sum_{j=1}^{3s} \frac{\sin(\omega_{c2}j) - \sin(\omega_{c1}j)}{\pi j}(L^j + L^{-j}). \quad (1)$$

up to a proportionality factor, which is $[\mathrm{w}_{bp}(1)]^{-1}$. In the above expression L denotes the lag operator, such that $L^j y_t = y_{t-j}$.

Another important filter in macroeconomics is the Hodrick and Prescott (HP) filter. The HP trend minimises the penalised least squares criterion:

$$\begin{aligned} PLS &= \textstyle\sum_{t=1}^{n}(y_t - \mu_t)^2 + \lambda \sum_{t=3}^{n}(\Delta^2 \mu_t)^2 \\ &= (y-\mu)'(y-\mu) + \lambda \mu' D^{2'} D^2 \mu, \end{aligned}$$

where $\Delta = I - L$ is the difference operator, y and μ are the $n \times 1$ vector with elements $\{y_t, t = 1, \ldots, n\}$ and $\{\mu_t, t = 1, \ldots, n\}$, respectively, and D^2 is the $n-2 \times n$ matrix corresponding to the 2nd differences filter, with $d_{ii} = 1$, $d_{i,i-1} = -2$, $d_{i,i-2} = 1$ and zero otherwise. Differentiating with respect to μ, the first order conditions yield: $\tilde{\mu} = (I_n + \lambda D^{2'} D^2)^{-1} y$. The *smoothness* or *roughness penalty* parameter, λ, governs the trade-off between fidelity and smoothness. HP purposively select the value $\lambda = 1600$ for quarterly time series. Ravn and Uhlig (2002) discuss the choice of λ for any frequency s of observations.

It is well known that, assuming the availability of a doubly infinite sample, $y_{t+j}, j = -\infty, \ldots, \infty$, the above HP filter is equivalent to the Wiener-Kolmogorov optimal signal extraction filter for the trend component μ_t of the following local linear trend model (King and Rebelo, 1993).

$$\begin{array}{lcll} y_t & = & \mu_t + \epsilon_t, & \qquad\qquad t = 1, 2, \ldots, n, \\ \Delta^2 \mu_t & = & \zeta_t, & \zeta_t \sim \mathrm{NID}(0, \sigma_\zeta^2), \\ \epsilon_t & \sim & \mathrm{NID}(0, \lambda\sigma_\zeta^2), & \mathrm{E}(\zeta_t, \epsilon_{t-j}) = 0, \forall j, \end{array} \tag{2}$$

The HP cycle is obtained as $y - \tilde{\mu}$ and the corresponding filter weights are obtained from the rows of the matrix $\lambda D^{2'}(I_n + \lambda D^2 D^{2'})^{-1} D^2$. Given the availability of a doubly infinite sample (see Whittle, 1983):

$$\tilde{\epsilon}_t = w_{HP}(L) y_t, \quad w_{HP}(L) = \frac{\lambda |1 - L|^4}{1 + \lambda |1 - L|^4}, \tag{3}$$

and we have posited $|1 - L|^2 = (1 - L)(1 - L^{-1})$. Kaiser and Maravall (2005) refer to this result as the naïve model based interpretation of the HP filter. This is so since (2) is not usually a plausible model of economic fluctuations (the cycle being a white noise process).

In the sequel we shall generically denote by $\mathrm{w}_c(L)$ the cycle extraction filter under investigation. We mention in passing that the filter may not be time invariant, as in the HP case, although we will continue to use the same notation. The important point is that we are not able to observe y_t, but rather a contaminated version of it,

$$z_t = y_t + \gamma_t,$$

where γ_t is the seasonal component. Let us also assume that the components are orthogonal, so that the pseudo-autocovariance generating function (ACGF) of z_t decomposes as $g_z(L) = g_y(L) + g_\gamma(L)$, where $g_y(L)$ and $g_\gamma(L)$ are ACGFs of the nonseasonal and the seasonal components, respectively, which we assume to have a known parametric form. The minimum mean square linear estimator of the nonseasonal component based on a doubly infinite sample is

$$\tilde{y}_t = \frac{g_y(L)}{g_z(L)} z_t.$$

We mention in passing that with little effort we could cover the case in which there are interactions between the seasonal and the nonseasonal component, for which $g_z(L) = g_y(L) + g_\gamma(L) + g_{y\gamma}(L) + g_{\gamma y}(L)$, where $g_{y\gamma}(L) = g_{\gamma y}(L^{-1})$ is the cross-covariance generating function of (y_t, γ_t).

Business cycle analysis is customarily carried out by applying the filter $\mathrm{w}_c(L)$ to the seasonally adjusted series, $\tilde{y}_t = z_t - \mathrm{E}(\gamma_t | z)$, rather than y_t, which is unobserved. Let

$$c_t = \mathrm{w}_c(L) y_t$$

denote the true cycle, which arises from applying the cycle filter to y_t. The estimator $\mathrm{w}_c(L)z_t$ would have very poor properties due to leakage from the seasonal frequencies; in particular, the fundamental frequency (corresponding to a period of one year) lies very close to the business cycle frequency range that characterizes the BK filter (1.5 years to 8 years), and the HP cycle is a high-pass filter that will retain to great extent the spectral power at the seasonal frequencies.

We thus focus on the estimator of the cycle

$$\tilde{c}_t = \mathrm{E}(c_t|z) = \mathrm{w}_c(L)\tilde{y}_t, \tag{4}$$

where $z = (\ldots, z_{t-1}, z_t, z_{t+1}, \ldots)$, and similarly $y = (\ldots, y_{t-1}, y_t, y_{t+1}, \ldots)$, $\tilde{y} = (\ldots, \tilde{y}_{t-1}, \tilde{y}_t, \tilde{y}_{t+1}, \ldots)$. This is not without consequences for business cycle analysis. In particular, the seasonal component is estimated with nonzero estimation error variance, and this constitutes an additional source of variation for the above estimator.

3 Seasonal Adjustment and the Reliability of the Cycle Estimates

It is important to illustrate how seasonal adjustment affects the reliability of the cycle measurement. Orphanides and van Norden (2003) have stressed the relevance of the uncertainty assessment for the estimation of the output gap, which is often measured by applying an ad hoc filter $\mathrm{w}_c(L)$, such as BK or HP to seasonally adjusted data. Hereto we are not aware of studies that aim at assessing the role of seasonal adjustment on business cycle measurement. The subsequent analysis aims at bridging the gap.

We shall be exclusively concerned with the variability due to estimation of the seasonally adjusted series due to smoothing. We will leave aside the additional source arising from the specification and the estimation of the seasonal model.

Let $e_t = c_t - \tilde{c}_t$ denote the cycle estimation error. Then, $e_t = w_c(L)(y_t - \tilde{y}_t)$ and its ACGF is

$$g_e(L) = |\mathrm{w}_c(L)|^2 \frac{g_y(L)g_\gamma(L)}{g_z(L)},$$

as $g_y(L)g_\gamma(L)/g_z(L)$ is the ACGF of $y_t - \tilde{y}_t$. Here we have written $|\mathrm{w}_c(L)|^2 = \mathrm{w}_c(L)\mathrm{w}_c(L^{-1})$.

The estimator $\tilde{c}_t$ is (un)conditionally unbiased, since $\mathrm{E}(e_t|z) = 0$, but the variance of $\tilde{c}_t$ is smaller than $\mathrm{Var}(c_t)$. Denoting by $g_y(\omega), \omega \in [0, \pi]$, the spectral generating function (SGF) of y_t, and by $\mathrm{w}_c(\omega) = |\mathrm{w}_c(e^{-\imath\omega})|^2$ the squared gain of the cycle filter, i.e. the squared modulus of the frequency response function of the filter, with $e^{-\imath\omega} = \cos(\omega) - \imath\sin(\omega)$, the unconditional variance of the true cycle is

$$\mathrm{Var}(c_t) = \int_{-\pi}^{\pi} \mathrm{w}_c(\omega)g_y(\omega)d\omega.$$

Also, the SGF of $\tilde{y}_t$ is $g_{\tilde{y}}(\omega) = [g_y(\omega)]^2/g_z(\omega)$, as the estimator of the SA series is $\tilde{y}_t = g_y(\omega)/g_z(\omega)z_t$.

The variance of the cycle estimator is thus:

$$\begin{aligned}\mathrm{Var}(\tilde{c}_t) &= \int_{-\pi}^{\pi} \mathrm{w}_c(\omega) g_{\tilde{y}}(\omega) d\omega \\ &= \int_{-\pi}^{\pi} \mathrm{w}_c(\omega) \frac{g_y(\omega)}{g_z(\omega)} g_y(\omega) d\omega \\ &\leq \mathrm{Var}(c_t)\end{aligned}$$

The last inequality follows from the following basic identity:

$$\begin{aligned}\mathrm{Var}(c_t) &= \mathrm{Var}\left[\mathrm{E}(c_t|z)\right] + \mathrm{E}\left[\mathrm{Var}(c_t|z)\right] \\ &= \mathrm{Var}(\tilde{c}_t) + \mathrm{Var}(e_t|z) \\ \mathrm{Var}(e_t|z) &= \mathrm{Var}(c_t|z) = \int_0^{\pi} \mathrm{w}_c(\omega) \frac{g_y(\omega) g_\gamma(\omega)}{g_z(\omega)} d\omega\end{aligned} \tag{5}$$

where the expectations are taken with respect to the distribution of y given z. A first conclusion is that seasonal adjustment implies an underestimation of the cycle volatility, i.e. the amplitude of the cycle estimate is lower than the true amplitude.

The components $\mathrm{Var}(c_t)$ and $\mathrm{Var}(e_t|z)$ can be evaluated as a by-product of the SA modeling effort, as it will be illustrated shortly. Hence, in the sequel, we will assume that a parametric unobserved components model is available, postulating a decomposition $z_t = y_t + \gamma_t$.

In particular, $\mathrm{Var}(e_t|z)$ quantifies the uncertainty surrounding the measurement of the business cycle that is due to the fact that y_t, and thus c_t, is unobserved, and has to be estimated.

4 Empirical Evaluation of the Reliability of the Cycle Estimates

The estimation error variance could be computed analytically; however, the analytic expressions are valid for the cycle estimates in the middle of a long time series; analogous formulae could be derived using optimal prediction theory (see e.g. Whittle, 1983), but are extremely cumbersome and analytically intractable.

When the cycle filter admit a model-based representation we could assess the reliability of the cycle estimates using the embedding strategy proposed by Kaiser and Maravall (2005). This requires, however that the central symmetric filter, $\mathrm{w}_c(L)$, admits a spectral factorization $\mathrm{w}_c(L) = \varphi(L)\varphi(L^{-1})$, which is not the case of the BK filter (see Proietti, 2009).

When a model-based decomposition of the series is available, we propose to evaluate the reliability of the cycle estimates in finite samples via Monte Carlo methods, by setting up the following simulation scheme.

1. Formulate the model $z_t = y_t + \gamma_t$ and estimate it by maximum likelihood, under the assumption of Gaussianity. Obtain $\tilde{y} = \mathrm{E}(y|z)$ using the Kalman filter and smoothing algorithm.

2. For $i = 1, \dots, M$, run the following simulation smoother (Durbin and Koopman, 2004):

   (a) Draw $y^{(i)}, \gamma^{(i)}$, from $g_y(L)$ and $g_\gamma(L)$, respectively. Obtain $z^{(i)} = y^{(i)} + \gamma^{(i)} \sim g_z(L)$.

(b) Obtain the estimate of the nonseasonal component $\tilde{y}^{(i)}, i = 1, \ldots, M$, (using the Kalman filter and smoother as in step 1).

(c) Compute $\dot{y}^{(i)} = \tilde{y} + (y^{(i)} - \tilde{y}^{(i)}) \sim y|z$.
$\dot{y}^{(i)}$ is a draw from the conditional distribution of y given z. The latter is normal with mean $\tilde{y}$.

(d) Compute

$$\dot{c}_t^{(i)} = w_c(L)\dot{y}_t^{(i)} \sim c_t|z$$

($\dot{c}^{(i)}$ is a draw from the conditional distribution of c_t given z).

3. Estimate $\mathrm{Var}(c_t|z)$ using

$$\widehat{\mathrm{Var}}(c_t|z) = \frac{1}{M}\sum_i \left(\dot{c}_t^{(i)} - \tilde{c}_t\right)^2$$

The estimate of the conditional variance, $\widehat{\mathrm{Var}}(c_t|z)$, provides the assessment of the reliability required.

The actual implementation of this scheme entails a parametric model for the decomposition $z_t = \mu_t + \gamma_t$. This is the case of ARIMA model based seasonal adjustment (see Hillmer and Tiao, 1982) and structural time series models (Harvey, 1989).

5 Empirical illustration

We present an empirical illustration which deals with the extraction of the cycle by the BK filter and the HP filter for three monthly industrial production series referring to Germany, France and Italy. The series are taken from the Eurostat database EuroInd and are available for the sample 1990.1-2009.4 (France and Italy) and 1991.1-2009.4 (Germany). We restrict our analysis to the period up to 2007.12.

We assume that z_t follows a basic structural model:

$$\begin{aligned} z_t &= y_t + \gamma_t, \\ y_t &= \mu_t + \epsilon_t, \\ \gamma_t &= s_t + x'_{it}\delta_i, \end{aligned}$$

where the nonseasonal component, y_t is a local linear trend plus irregular,

$$\begin{array}{lllllll} \mu_t & = & \mu_{t-1} & + \ \beta_{t-1} & + & \eta_t, & \eta_t \sim \mathrm{NID}(0, \sigma^2_\eta) \\ \beta_t & = & & \ \ \ \beta_{t-1} & + & \zeta_t, & \zeta_t \sim \mathrm{NID}(0, \sigma^2_\zeta) \end{array} \tag{6}$$

and $\epsilon_t \sim \mathrm{NID}(0, \sigma^2_\epsilon)$. The seasonal component is made up of a purely seasonal trigonometric cycle plus the calendar component, which is obtained from the regression on suitable explanatory variables.

The component s_t arises from the combination of six stochastic cycles defined at the seasonal frequencies $\lambda_j = 2\pi j/12$, $j = 1, \ldots, 6$, λ_1 representing the fundamental frequency (corresponding to a period of 12 monthly observations) and the remaining being the five harmonics (corresponding to periods of 6 months, i.e. two cycles in a year, 4 months, i.e. three cycles in a year, 3 months, i.e. four cycles in a year, 2.4, i.e. five cycles in a year, and 2 months):

$$s_t = \sum_{j=1}^{6} s_{jt}, \qquad \begin{bmatrix} s_{j,t} \\ s^*_{j,t} \end{bmatrix} = \begin{bmatrix} \cos\lambda_j & \sin\lambda_j \\ -\sin\lambda_j & \cos\lambda_j \end{bmatrix} \begin{bmatrix} s_{j,t-1} \\ s^*_{j,t-1} \end{bmatrix} + \begin{bmatrix} \omega_{j,t} \\ \omega^*_{j,t} \end{bmatrix}, j = 1, \ldots, 5, \tag{7}$$

and $s_{6,t} = -s_{6t} + \omega_{6t}$. The disturbances ω_{jt} and ω^*_{jt} are normally and independently distributed with common variance σ^2_ω for $j = 1, \ldots, 5$, whereas $\text{Var}(\omega_{6t}) = 0.5\sigma^2_\omega$ (see Proietti, 2000, for further details). Calendar effects are accounted for by six trading day regressors, measuring the number of days of type j, $j = 1, \ldots, 6$, occurring in month t, in excess of the number of Sundays. One regressor picks up the Easter effect, and an additional one captures the the length of month (LOM) effect. All the disturbances in the model are mutually and serially uncorrelated.

The model is fitted to each of the time series by maximum likelihood. Conditional on the maximum likelihood estimates, we compute the estimate of the cycle (4). To assess its reliability, we perform the simulation scheme proposed in the previous section to get the estimated conditional variance of $\tilde{c}_t$.

Figure 1 displays the 95% interval estimates of the seasonally adjusted series, the HP cycle, the BK cycle, and the HP bandpass cycle. It should be noticed that the SA series displays a lot of variation at the high frequencies. This is responsible for the high variability of the HP cycle estimates (see the top right panel). The latter is obtained by adopting the value $\lambda = 129600$ for the smoothness parameter, which, according to Ravn and Uhlig (2002), is the exact monthly analogue of the traditional HP with $\lambda = 1600$ for quarterly time series. As a matter of fact, the HP cycle filter is a high-pass filter that retains to a great extent the amplitude of the high frequency fluctuations. As a result a lot of high frequency variation leaks from the SA series to the cycle estimates.

The reliability improves a lot when we consider cycle measures that suppress the high frequency variation. The bottom left plot displays the point estimates of the BK cycle, along with their 95% confidence bands. These estimates are little affected by the adjustment, being much more stable. It should be noticed that the BK filter is two-sided and fails to produce the estimates at the beginning and at the end of the sample. These limitations are overcome by the bandpass version of the HP filter, which originates from the difference of two low-pass HP filters with cut-off periods equal to 18 months (1.5 years) and 92 months (8 years), corresponding respectively to $\lambda = 68.7$ and 54535. The bottom right plot confirms that in the middle of the sample seasonal adjustment does not contribute much to the variability of the estimates. However, as it might be anticipated, the variance of the estimates is much higher at the extremes of the sample period.

Figure 2 compares the 95% interval estimates of the HP bandpass cycle for Germany, Italy and France. The German series is available starting from 1991. The plot reveals that the uncertainty is about the same, the series behaving very similarly, and that the business cycles are very synchronized.

Figure 1: France, Index of Industrial Production (Manuf. total, Source Eurostat). Interval estimates of the seasonally adjusted series, the HP cycle with $\lambda = 129600$, the BK cycle, and the HP bandpass cycle.

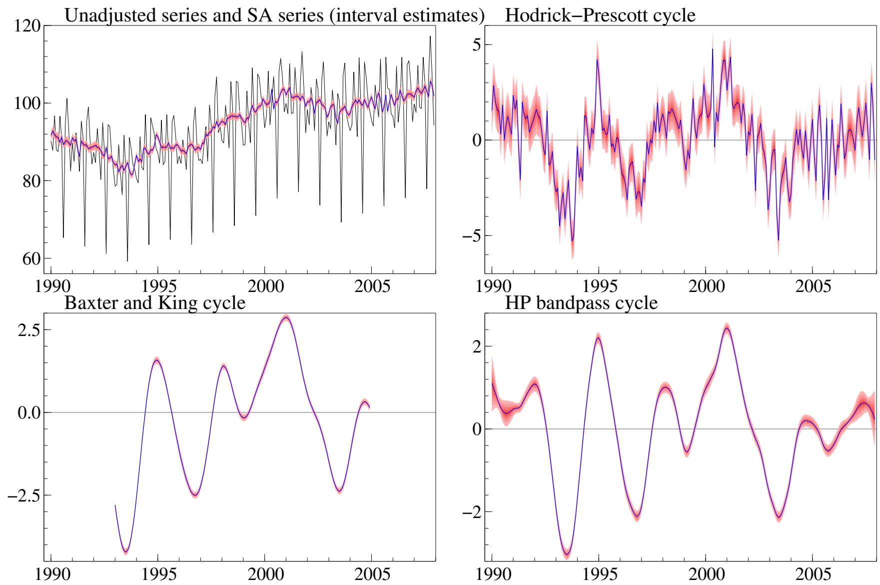

Figure 2: Germany, France and Italy, Index of Industrial Production (Manuf. total, Source Eurostat). Comparison of the estimates of the HP bandpass cycle.

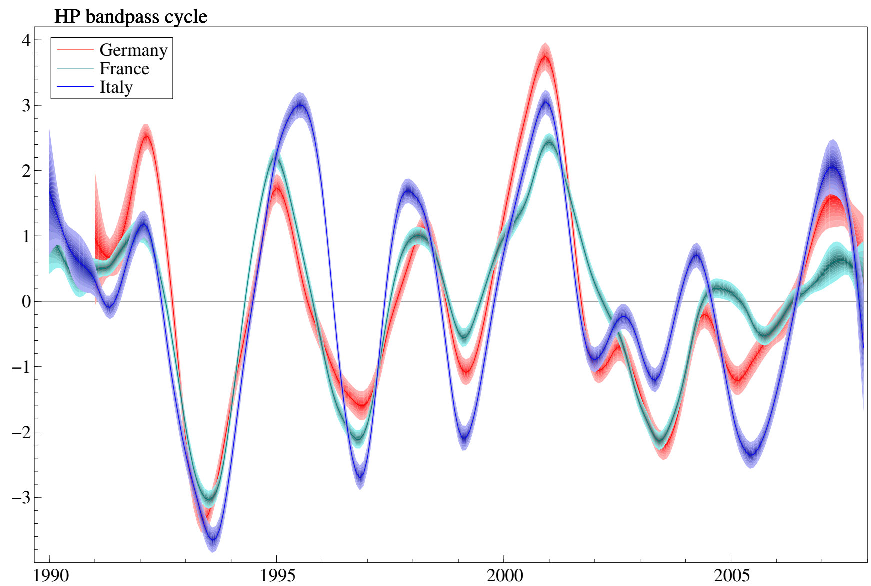

6 Forecast Extensions and Estimation in Real Time

The estimation of the cycle in real time is perhaps the most problematic aspect of business cycle measurement. If, on the one hand, it is relatively straightforward to provide a solution to the problem of band-pass filtering when the availability of a doubly infinite sample is assumed, the adaptation of the optimal filter to a finite realization is indeed an issue (see Percival and Walden, 1993, Pollock, 2003). We have seen, for instance, that the BK filter does not handle estimation in real time (more generally, for the last three years of the sample), but it makes feasible only the historical estimation of the cycle, instead. Christiano and Fitzgerald (2003) provided a solution to this problem, which takes into account the nature of the filtered series and its possible nonstationarity. Their proposal rests on the idea of extending the series by optimal forecasts.

For both the highpass and bandpass HP filter we were able in the previous sections to provide interval estimates also at the extremes of the sample, by means of Monte Carlo simulation. However, the assessment of the reliability at the end of the sample refers to the asymmetric HP cycle filter obtained either from the penalized least squares estimation criterion or, equivalently, from the Kalman filter and smoothing algorithm adapted to the local linear trend model (2).

In particular, the end of sample weights are obtained from the last rows of the matrix $\lambda D^{2'}(I_n + \lambda D^2 D^{2'})^{-1} D^2$. The real time estimates of the cycle at time t, built upon the information available up to and including that time, can be represented as follows,

$$\tilde{c}_{t|t} = \frac{\theta_2}{\theta(L)} \Delta^2 \tilde{y}_t, \tag{8}$$

and computed via the recursive formula

$$\tilde{c}_{t|t} = -\theta_1 \tilde{c}_{t-1|t-1} - \theta_2 \tilde{c}_{t-2|t-2} + \theta_2 \Delta^2 \tilde{y}_t, \tag{9}$$

$t = 3, \dots, n$, with starting values $\tilde{c}_{1|1} = \tilde{c}_{2|2} = 0$. Here $\theta(L) = (1 + \theta_1 L + \theta_2 L^2)$ denotes the moving average polynomial of the ARIMA(0,2,2) reduced form of the local linear trend model (2). The moving average parameter θ_2 is related to λ by

$$\lambda = \frac{\theta_2(1+\theta_2)^2}{(1-\theta_2)^4},$$

whereas θ_1 is obtained from θ_2 as

$$\theta_1 = \frac{-4\theta_2}{1+\theta_2}.$$

A proof of the recursion (9) is provided in the Appendix.

Hence, it must be remarked that the previous assessment of business cycle uncertainty at the end of the sample referred to the direct asymmetric HP filter (8). This filter is one-sided and implies a phase shift, i.e. turning points are displaced along the time axis. Since the model (2), for which this asymmetric filter is optimal, can be considered as a naïve approximation to a sensible model of economic fluctuations, it has been proposed to apply the filter to the series extended by optimal forecasts. The HP filter would then be applied to the seasonally adjusted series extended by its forecasts. The benefits are twofold: a reduction of the effect of displacing turning points, a reduction of the size of the cycle revision as new observations become available.

The issue that we address in this section deals with the reliability of a given symmetric cycle filter, like BK or the central HP filter given in (4), when applied to the seasonally adjusted series extended by forecasts. This analysis is more relevant than that carried out in the previous section for the assessment of the business cycle in real time. For the BK cycle, it enables to obtain interval estimates of the cycle for the last three years of the sample period; as far as HP is concerned, the analysis focuses on the interval estimation of the cycle component as extracted by the fixed filter 4, interpreted as a highpass filter. Hence, it deals with the second-order properties of the projection of $c_t = \mathrm{w}_c(L)y_t$ onto the available finite sample, z.

The expectation of the cycle conditional on the observed sample, i.e. the optimal estimator of the cycle given the data is $\tilde{c}_t = \mathrm{w}_c(L)\mathrm{E}(y_t|z)$, and is computed on the seasonally adjusted series augmented by its forecasts and backcasts, obtained from the maintained seasonal model. The reliability of the cycle estimator, as measured by $\mathrm{Var}(c_t|z)$ is evaluated by Monte Carlo simulation using the same algorithm illustrated in section 4. The vector y now includes also the past and future values of the series. Their conditional expectations given the observed series, z, are obtained from a simple modification of the Kalman filter updating equations (see Durbin and Koopman, 2001). The algorithm enables to draw repeated independent samples from $y|z$ and thus from $c|z$. These samples are used to estimate the mean square error of the cycle estimator.

In figure 3 we report the interval estimates of c_t for the HP highpass filter ($\lambda = 129600$), the BK filter and the HP bandpass filter, applied to the German Industrial Production series. In the simulation experiment we have simulated sample from the predictive distribution of the seasonally adjusted series up to ten years ahead. The plot also presents, in a different color, the 95% prediction intervals of the cycle. A few important facts emerge.

- The estimation error variance grows very rapidly at the end of the sample.
- The estimation error variance is very large at the end of the sample, so that the cycle is not significantly different from zero.
- For the HP cycle the estimation error variance at the extremes of the sample is larger than that arising from the assessment of the reliability made in the previous section (e.g. compare with the interval estimates reported in figure 2).
- The predicted values of the cycle converge to zero: for the BK filter they converge to zero oscillating around it with decreasing amplitude (this phenomenon is related to impulse response function of the filter).
- The prediction error variance of the cycle converges to the unconditional variance, $\mathrm{Var}(c_t)$.

Figure 3: Germany, Index of Industrial Production. Interval estimates of the HP cycle with $\lambda = 129600$, the BK cycle, the HP bandpass cycle, using forecast extensions.

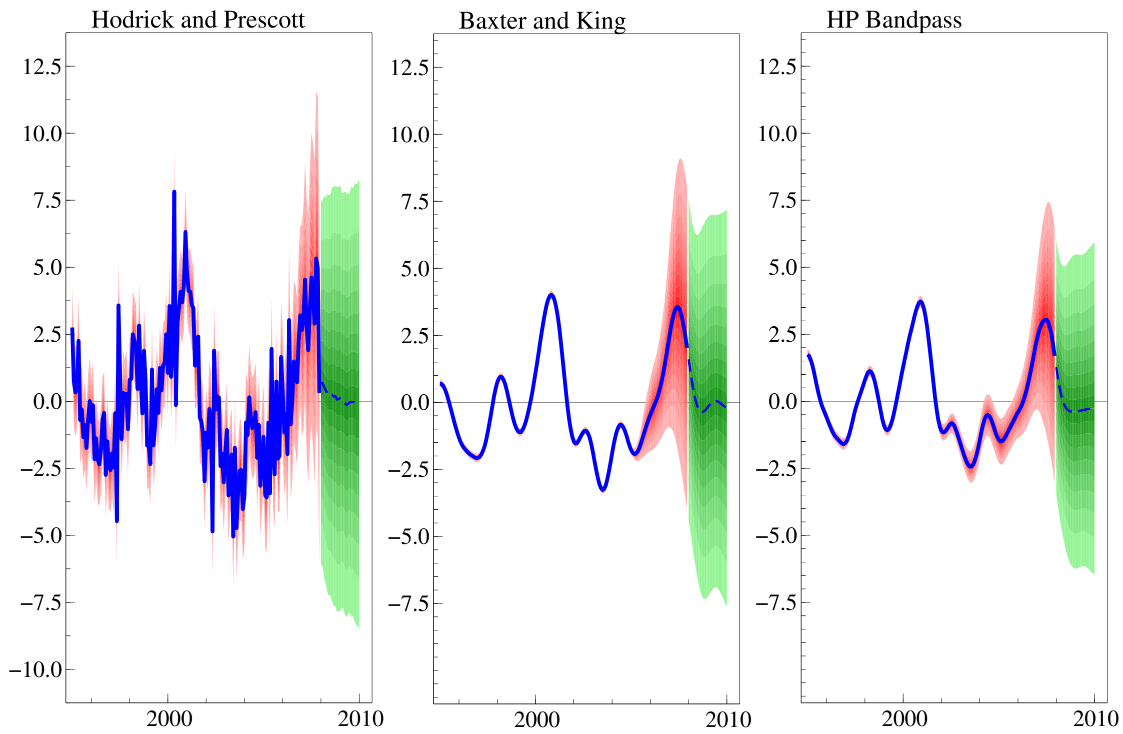

Figure 4: Germany, Index of Industrial Production. Point estimates of the HP bandpass cycle using forecast extensions and the direct asymmetric filter.

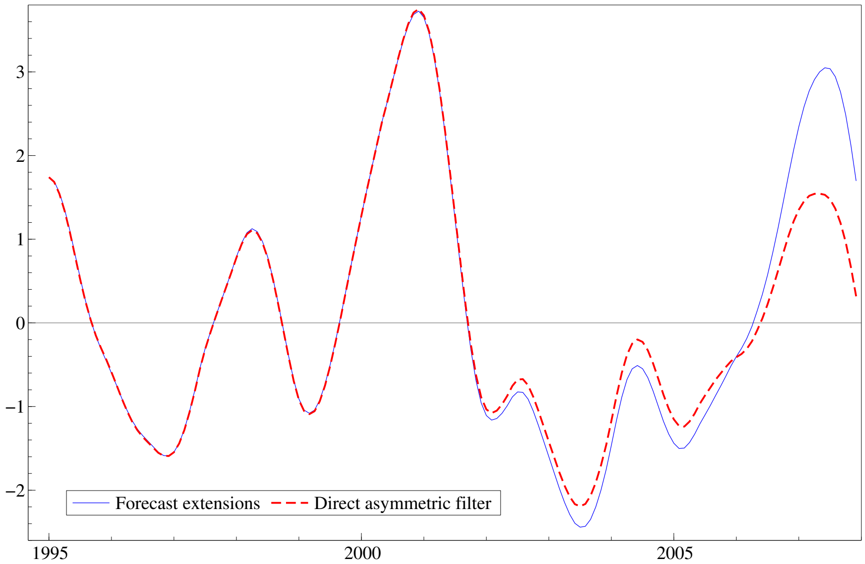

The third result is interesting: extending the series by forecasts is useful for reducing the bias in the assessment of the cycle, which here translates into the phase shift and the underestimation of the amplitude that is produced by the one sided filter (8); however, this is at odds with the increase in the estimation error variance.

Figure 4 compares the point estimates of c_t using forecast extensions with those obtained in the previous section by applying the direct HP asymmetric filter. Both the amplitude and phase effects (especially with respect to the last turning point) can be appreciated from the graph.

7 Conclusions

Seasonality and the adjustment thereof contributes to the uncertainty of the business cycle in several ways. The seasonally adjusted series is estimated and its estimation error has to be taken into account. Moreover, seasonal adjustment is a source of data revision: as a new observation is added, the estimates are revised to take into account the new information.

This paper has provided a way to assess the contribution of seasonal adjustment as an often neglected source of unreliability of the cycle estimates. The main conclusion is that this contribution is sizable for highpass filters, like the Hodrick-Prescott filter which are strongly affected by the high frequency fluctuations in a time series. These are the one that are more affected by seasonal adjustment. If a band-pass component is considered, the effect is much less sizable, although it is stronger at the extremes of the sample period.

In general, the contribution of seasonal adjustment to the uncertainty of the cycle estimates will depend on the extent to which the seasonal component γ_t contaminates the observations z_t. A more variable seasonal component will imply a larger estimation error for the seasonally adjusted series, which in turn will yield a less reliable cycle estimate.

The extraction of the business cycle from the seasonally adjusted series neglects an important source of unreliability, i.e. the fact that the series is estimated, rather than observed. The correct assessment of the uncertainty can be made, within a model based framework, by working with the original time series, and formulating a data coherent decomposition into a seasonal component and a non-seasonal one. The latter component may as well be fully agnostic about the presence of the cycle. In fact, the perspective considered here is that we do not model the cycle, which is extracted by a fixed filter.

We have also considered the strategy of extending the seasonally adjusted series by its optimal forecasts. We conclude that this offers the correct perspective, non only for reducing the revisions of the cycle estimates and ameliorating the displacement of the turning points of the cycle, but also for assessing the reliability of the cycle estimates at the end of the sample and for forecasting the cycle. Unfortunately, it turns out that for the Industrial Production series considered in our illustrations the estimation error variance at the end of the sample is so large, that we cannot reject the null that the cycle equals zero.

Acknowledgements

I would like to thank Dominique Ladiray, Gianluigi Mazzi and James Mitchell for their comments on a preliminary version of the paper.

Appendix: Derivation of the real time HP filter

Let us focus on the naïve model based interpretation of the HP filter, i.e. as the optimal signal extraction filter for the local linear trend model (2); see King and Rebelo (1993). The model has an ARIMA(0,2,2) reduced form

$$\Delta^2 y_t = \xi_t + \theta_1 \xi_{t-1} + \theta_2 \xi_{t-2}, \xi_t \sim \text{NID}(0, \sigma^2),$$

where NID denotes Gaussian white noise and $\Delta = 1 - L$ is the difference operator.

Let $\gamma(k)$ denote the autocovariance of $\Delta^2 y_t$ at lag k. Equating the lag 2 autocovariance of the structural and reduced form yields: $\sigma^2_\epsilon = \lambda \sigma^2_\zeta = \theta_2 \sigma^2$, which is true for $\theta_2 \geq 0$. Also, $\gamma(1) + 4\gamma(2) = 0$, which implies $\theta_1 = -4\theta_2/(1 + \theta_2)$. Moreover, $\sigma^2_\zeta = [\theta(1)]^2 \sigma^2$, which in turn gives the noise to signal ratio $\lambda = \sigma^2_\epsilon / \sigma^2_\zeta = \theta_2 (1 + \theta_2)^2 / (1 - \theta_2)^4$.

Under the assumption that (2) is the true model, The joint distribution of ϵ_t and y_t, given $Y_{t-1} = \{y_s, s \leq t-1\}$ is

$$\begin{array}{c} y_t \\ \epsilon_t \end{array} \Bigg| Y_{t-1} \sim \mathrm{N}\left[\begin{pmatrix} 2y_{t-1} - y_{t-2} + \theta_1\xi_{t-1} + \theta_2\xi_{t-2} \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2 & \sigma^2_\epsilon \\ \sigma^2_\epsilon & \sigma^2_\epsilon \end{pmatrix}, \right]$$

where $\mathrm{Cov}(y_t, \epsilon_t|Y_{t-1}) = \sigma^2_\epsilon$ can be seen on writing $y_t - \mathrm{E}(y_t|Y_{t-1}) = \mu_t - \mathrm{E}(\mu_t|Y_{t-1}) + \epsilon_t$. Denoting $\tilde{c}_{t|t} = \mathrm{E}(\epsilon_t|y_t, Y_{t-1}) = \mathrm{E}(\epsilon_t|Y_t)$, using a standard lemma on conditioning for Gaussian random variables we obtain:

$$\tilde{c}_{t|t} = \frac{\sigma^2_\epsilon}{\sigma^2}(y_t - 2y_{t-1} + y_{t-2} - \theta_1\xi_{t-1} - \theta_2\xi_{t-2})$$

so that, replacing $\sigma^2_\epsilon = \theta_2\sigma^2$, and rearranging,

$$\begin{aligned} \tilde{c}_{t|t} &= \theta_2\left[\Delta^2 y_t - (\theta_1 L + \theta_2 L^2)\xi_t\right] \\ &= \theta_2[\theta(L)]^{-1}\Delta^2 y_t \end{aligned}$$

which proves that the real time estimates of the cycle are obtained from the second order recursion:

$$\tilde{c}_{t|t} = -\theta_1\tilde{c}_{t-1|t-1} - \theta_2\tilde{c}_{t-2|t-2} + \theta_2\Delta^2 y_t,$$

which for a finite sample can be started off at $t = 3$, with starting values $\tilde{c}_{1|1} = \tilde{c}_{2|2} = 0$.

Interestingly, the smoothed estimates can be computed by the same filter applied to the sequence $\tilde{c}_{t|t}$ in reverse order, i.e. the recursive formula

$$\tilde{c}_t = -\theta_1\tilde{c}_{t+1} - \theta_2\tilde{c}_{t+2} + \Delta^2\tilde{c}_{t|t}.$$

In finite samples the backwards recursion is run for $t = n, \ldots, 1$, with starting values $\tilde{c}_{n+1} = \tilde{c}_{n+2} = 0$. This remarkable result arises from the direct application of the Wiener-Kolmogorov filter for estimating ϵ_t given the availability of a doubly infinite sample.

References

Baxter, M., and King, R.G. (1999). Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series. *The Review of Economics and Statistics*, 81, 575–593.

Bell W. R. and Martin D.E.K. (2004). Computation of asymmetric signal extraction filters and mean squared error for ARIMA component models. *Journal of Time Series Analysis*, 25, 603–623.

Bell W. R. and Hillmer S.C. (1984). Issues involved with the seasonal adjustment of economic time series, *Journal of Business and Economic Statistics*, 2, 291320.

Caporello, G. and Maravall, A. (2004). Program TSW. Revised Manual Version May 2004. Documentos Occasionales N8 0408, Madrid: Banco de Espana.

Cleveland, W.P. and Tiao, G. C. (1976) Decomposition of Seasonal Time Series A Model for the Census X-11 Program, Journal of the American Statistical Association, 71, 581-587

Cleveland, W.S. (1983), Seasonal and Calendar Adjustment, D. R. Brillinger and P. R. Krishnaiah, eds., Handbook of Statistics, Vol. 3, Elsevier Science Publishers B.V. 0983) 39-72

Christiano L.J., Fitzgerald T.J., (2003), The band pass filter, *International Economic Review*, 44, 435-465.

Doornik, J.A. (2001). *Ox 3.0 - An Object-Oriented Matrix Programming Language*, Timberlake Consultants Ltd: London.

Durbin J., and Koopman, S.J. (2001). *Time Series Analysis by State Space Methods*, Oxford University Press: New York.

Durbin, J., and S.J. Koopman (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika*, 89, 603-615.

Findley, D. F. (2005), Some Recent Developments and Directions in Seasonal Adjustment, Journal of Official Statistics, Vol. 21, No. 2

Findley, D. F., Monsell, B. C., Bell, W. R., Otto, M. C. and Chen, B. (1998), 'New capabilities and methods of the X-12-ARIMA seasonal adjustment program', Journal of Business and Economic Statistics 16, 127-177.

Gomez, V., Maravall, A. (1996), Programs TRAMO and SEATS, Banco de Espana, Servicios de Estudios, Documento de trabajo n. 9628.

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

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

Hodrick R.J., and Prescott, E.C. (1997). Postwar U.S. Business Cycles: an Empirical Investigation. *Journal of Money, Credit and Banking*, 29, 1-16.

Kaiser, R. Maravall A., 2001. *Measuring Business Cycles in Economic Time Series*, Lecture Notes in Statistics, 154, Springer-Verlag, New York.

Kaiser, R. Maravall A., (2005). Combining filter design with model-based filtering (with an application to business-cycle estimation). *International Journal of Forecasting*, 21, 691–710.

King, R. G., and Rebelo, S. T. (1993). Low frequency filtering and real business cycles. *Journal of Economic Dynamics and Control*, 17, 207-233.

Koopman S.J., Shepard, N., and Doornik, J.A. (1999), Statistical algorithms for models in state space using SsfPack 2.2, Econometrics Journal, 2, 113- 166.

Maravall, A. (2005). Brief description of the programs. Madrid. Banco de Espana. [pdf] Nerlove, M., Grether, D.M. and Carvalho, J.L. (1979), Analysis of Economic Time Series: A Synthesis, New York: Academic Press.

Orphanides A, van Norden S (2002) The Unreliability of Output Gap Estimates in Real Time. *The Review of Economics and Statistics*, 84, 569-583

Percival D., Walden A. (1993). *Spectral Analysis for Physical Applications.* Cambridge University Press.

Pollock, D. S. G. (2000). Trend estimation and de-trending via rational square wave filters. *Journal of Econometrics*, 99, 317 334.

Pollock, D. S. G. (2003). Improved frequency selective filters. *Computational Statistics and Data Analysis*, 42, 279 297.

Proietti, T. (2000), Comparing seasonal components for structural time series models, *International Journal of Forecasting*, 16, 247-260.

Proietti T. (2009), On the Model Based Interpretation of Filters and the Reliability of Trend-Cycle Estimates, Econometric Reviews, vol. 28, 186208.

Ravn, M.O. and Uhlig, H. (2002). On adjusting the HodrickPrescott filter for the frequency of observations. *The Review of Economics and Statistics*, 84, 371376.

West, M., Harrison, J. (1997), *Bayesian Forecasting and Dynamic Models*, 2nd edition, New York, Springer-Verlag.

Whittle P. (1983) *Prediction and Regulation by Linear Least Squares Methods*, Second edition. Basil Blackwell, Oxford.