

Estimation of the business cycle: A modified Hodrick-Prescott filter

Regina Kaiser¹, Agustín Maravall²

¹ Departamento de Estadística y Econometría, Universidad Carlos III de Madrid, calle Madrid 126, E-28903 Getafe (Madrid), Spain (e-mail: kaiser@est-econ.uc3m.es)

² Servicio de Estudios - Banco de España, Alcalá 50, E-28014 Madrid, Spain (e-mail: maravall@bde.es)

Abstract. Hodrick-Prescott (HP) filtering of (most often, seasonally adjusted) quarterly series is analysed. Some of the criticism to the filter are addressed. It is seen that, while filtering strongly affects autocorrelations, it has little effect on crosscorrelations. It is argued that the criticism that HP filtering induces a spurious cycle in the series is unwarranted. The filter, however, presents two serious drawbacks: First, poor performance at the end periods, due to the size of the revisions in preliminary estimators, and, second, the amount of noise in the cyclical signal, which seriously disturbs its interpretation. We show how the addition of two model-based features (in particular, applying the filter to the series extended with proper ARIMA forecasts and backcasts, and using as input to the filter the trend-cycle component instead of the seasonally adjusted series) can considerably improve the filter performance. Throughout the discussion, we use a computationally and analytically convenient alternative derivation of the HP filter, and illustrate the results with an example consisting of 4 Spanish economic indicators.

JEL Classification: C13, C22, E32

Key words: Time series, optimal filtering, trends, seasonality, economic cycles

1 Introduction

There are two different uses of trends in applied work. First, in short-term monitoring and seasonal adjustment, trends are equal to $p_t = x_t - (s_t + u_t)$, where x_t is the observed series, s_t is the seasonal component, and u_t is the irregular component, that typically captures white (or close to it) noise behavior. Examples

The first author acknowledges support by the Spanish grant PB95-0299 of CICYT

of these trends are the ones produced by the Henderson filters in X11 or X12, or the ones obtained in the model-based decomposition of a series, as in programs STAMP or SEATS (see Findley et al 1998; Koopman et al 1996; Gómez and Maravall 1996). Since they only differ from the seasonally adjusted (SA) series by a highly erratic component, often they will contain variation of the series within the range of cyclical frequencies. As a consequence, these trends will only be of interest as a short-term signal to monitor, for example, period-to-period growth. An example is provided by the continuous lines in Fig. 1.1a and 1.1b: The gain of the filter extends over a wide range of cyclical frequencies, and the trend is seen to contain short-term cyclical oscillations. Throughout the paper, these short-term trend will be referred to as trend-cycles, and denoted p_t ; on occasion, they will also be called "noise free" SA series.

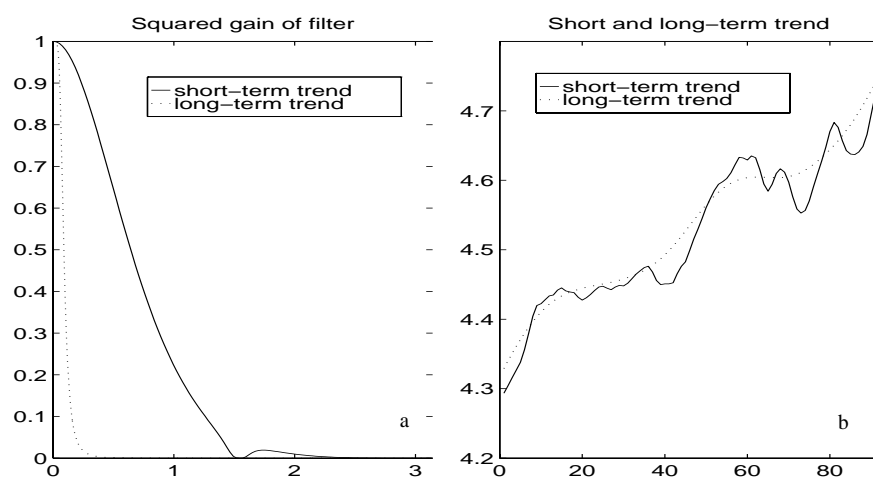


Fig. 1a,b. Short-term versus long-term trends

The second use of trends is in business cycle analysis, where short-term trends cannot be used because they are contaminated with cyclical variation; longer-term trends are needed. Despite its importance, decades of attention have shown that formal modeling of economic cycles is a frustrating issue. Therefore, applied research and work at policy making institutions has relied heavily on ad-hoc filters, the most popular of which is the Hodrick-Prescott (HP) one (see Prescott 1986). Thus a standard procedure to estimate economic cycles is to apply the HP filter to X11-SA series. The dotted lines in Fig. 1.1a and b represent the HP long-term trend gain and estimator. Long-term trends will be called simply trends, and represented by m_t .

The use of the HP filter for business-cycle estimation has been the subject of academic discussion. Criticisms are found in, for example, Canova (1998), Cogley and Nason (1995), Harvey and Jaeger (1993), King and Rebelo (1993), and Maravall (1995). Notwithstanding the criticisms, its widespread use in practice may evidence (besides its simplicity) the empirical fact that, as a first (or

rough) approximation, analysts find the results useful. The decision of which is the cutting point between a trend and a cycle is, ultimately, arbitrary, and to some extent depends on the purpose of the analysis. For example, from a month to month horizon, a periodic 10-year component may well be considered trend; if business cycle is the objective, it should be considered cycle.

Be that as it may, the HP filter presents some serious limitations. First, it is generally accepted that economic cycles have non-linear features (see, for example, Hamilton 1989); in this paper we do not deal with non-linear improvements. We address, first, the well-known criticism of spurious results due to the ad-hoc character of the filter, and the (often ignored yet important) limitation implied by revisions, which produce imprecision in the cycle estimator for recent periods. Then, we show how the integration of some relatively simple ARIMA-model-based (AMB) techniques can produce important improvements in the performance of the cyclical signal.

2 The Hodrick-Prescott filter: Wiener-Kolmogorov derivation

We start by providing an alternative representation of the HP filter that provides an efficient and simple computational algorithm and turns out to be useful for analytical discussion. Let x_t ($t = 1, \dots, T$) denote an observed series. The HP filter decomposes x_t into a smooth trend (m_t) and a residual (c_t), where the trend is meant to capture the long-term growth of the series, and the residual (equal to the deviation from that growth) represents the cyclical component. For the moment we shall assume that the series contains no seasonality.

The HP filter is a low-pass filter and can be seen as a Whittaker-Henderson type A filter and as a member of the Butterworth family of filters (see Gómez 1998). The filter was derived as the solution of a problem that balances a trade-off between fit and smoothness in the following way. In the decomposition

$$x_t = m_t + c_t, \tag{2.1}$$

the HP filter provides the estimator of c_t and m_t such that the expression

$$\sum_{t=1}^T c_t^2 + \lambda \sum_{t=3}^T (\nabla^2 m_t)^2 \tag{2.2}$$

is minimized ($\nabla = 1 - B$ is the difference operator, B is the backward operator $B^j z_t = z_{t-j}$, and F denotes the forward operator, $F^j z_t = z_{t+j}$) The first summation in (2.2) penalizes bad fitting, while the second one penalizes lack of smoothness. The parameter λ regulates the trade-off: when $\lambda = 0$, $\hat{m}_t = x_t$, when $\lambda \rightarrow \infty$, \hat{m}_t becomes a deterministic linear trend. The solution to the problem of minimizing (2.2) subject to the restriction (2.1) is given by (see Danthine and Girardin 1989)

$$\hat{\mathbf{m}} = \mathbf{A}^{-1} \mathbf{x}, \quad \mathbf{A} = \mathbf{I} + \lambda \mathbf{K}' \mathbf{K}, \tag{2.3}$$

where $\hat{\mathbf{m}}$ and \mathbf{x} are the vectors $(\hat{m}_1, \dots, \hat{m}_T)'$ and $(x_1, \dots, x_T)'$ respectively, and \mathbf{K} is an $(n-2) \times n$ matrix with its elements given by $K_{ij} = 1$ if $i=j$ or $i = j+2$, $K_{ij} = -2$ if $i = j+1$, $K_{ij} = 0$ otherwise.

Clearly, the estimator of the trend for a given period depends on the length of the series. Consider the trend for period T , the last observed period. Application of (2.3) yields an estimator to be denoted $\hat{m}_{T|T}$, where the first subindex refers to the period under estimation, and the second to the last observed period. This estimator will be called the concurrent estimator. When one more quarter is observed and \mathbf{x} becomes $(x_1, \dots, x_{T+1})'$, application of (2.3) yields a new estimation of m_T , namely $\hat{m}_{T|T+1}$. As more quarters are added, the estimator is revised. It can be seen that, for large enough k , $\hat{m}_{T|T+k}$ converges to a final or historical estimator, to be denoted \hat{m}_T . Therefore, for a long enough series, the final estimator may be assumed for the central periods, while estimators for the last years will be preliminary. This two sided interpretation of the HP filter seems unavoidable. Because additional correlated new information cannot deteriorate a projection, $\hat{m}_{T|T+1}$ should improve upon $\hat{m}_{T|T}$. Moreover, actual behavior of the US Business Cycle Dating Committee (or similar institutions) reveals in fact a two-sided filter, which starts with a preliminary estimator, and reaches the final decision with a lag of perhaps two years.

As shown in King and Rebelo (1993), the HP filter can be given a model-based interpretation. Let c_t in (2.1) be white noise with variance V_c and m_t follow the model

$$\nabla^2 m_t = a_{mt}, \quad (2.4)$$

where a_{mt} is a white noise variable (with variance V_m) uncorrelated to c_t . Throughout the paper, the expression "white noise" will denote a zero-mean normally identically independently distributed variable. Let $\lambda = V_c/V_m$ so that, without loss of generality, we can set $V_c = \lambda$, $V_m = 1$. The minimum mean squared error (MMSE) estimator of m_t can be obtained in a straightforward manner via the Kalman filter (see Harvey and Jaeger 1993). The interpretation of λ , the HP filter parameter, varies according to the rationalization of the filter. It regulates the trade-off between fitness and smoothness when the function (2.2) is minimized, it is equal to the ratio of the cycle and trend innovations in the model-based approach, and, when expressed as a Butterworth type filter, it is equal to, $\lambda = [4\sin^2(\omega_0/2)]^{-2}$, where ω_0 is the frequency for which 50% of the filter gain has been completed (see Gómez and Maravall 1998).

Alternatively, the same MMSE estimator can be obtained with the so-called Wiener Kolmogorov (WK) filter. In terms of the observations, the previous model can be rewritten as the IMA (2,2) model

$$\nabla^2 x_t = (1 + \theta_1 B + \theta_2 B^2)b_t = \theta_{HP}(B)b_t, \quad (2.5)$$

where b_t are the innovations in the x_t series. The variance of b_t , V_b , and the θ_1, θ_2 -parameters are found by factorizing the spectrum from the identity

$$(1 + \theta_1 B + \theta_2 B^2)b_t = a_{mt} + \nabla^2 c_t. \quad (2.6)$$

As an example, for quarterly series the standard value of λ is 1600, in which case

$$\theta_{HP}(B) = 1 - 1.77709B + 0.79944B^2; \quad V_b = 2001.4. \quad (2.7)$$

For an infinite realization of the series, the MMSE estimator of m_t is given by (see, for example, Maravall 1995)

$$\hat{m}_t = k_{m(HP)} \frac{1}{\theta_{HP}(B)\theta_{HP}(F)} x_t = \nu_{HP}^m(B, F)x_t, \quad (2.8)$$

where $k_{m(HP)} = V_m/V_b$. The filter $\nu_{HP}^m(B, F)$ is symmetric and, since (2.6) implies that $\theta_{HP}(B)$ is invertible, also convergent. Following Cleveland and Tiao (1976), for a finite series, expression (2.8) can still be applied, with x_t replaced by the series extended with forecast and backcasts. A simple and efficient algorithm to apply the filter, based on that in Burman (1980), is given in the Appendix. For the estimator of the cycle,

$$\hat{c}_t = [1 - \nu_{HP}^m(B, F)]x_t = \nu_{HP}^c(B, F)x_t, \quad (2.9)$$

where $\nu_{HP}^c(B, F)$ is also a two-sided centered, symmetric, and convergent linear filter, which can be rewritten as,

$$\hat{c}_t = \nu_{HP}^c(B, F)x_t = \left[k_{c(HP)} \frac{\nabla^2 \bar{\nabla}^2}{\theta_{HP}(B)\theta_{HP}(F)} \right] x_t, \quad (2.10)$$

where $k_{c(HP)} = V_c/V_b$, and a bar over an operator denotes the same operator with B replaced by F . When properly applied, the Danthine and Girardin, the Kalman filter, and the WK solutions are numerically identical (see Gómez, 1999). The last two are considerably more efficient than the first, and can be applied to series of any length. The WK filter turns out to be convenient for analytical discussion.

For seasonal series, since the seasonal variation should not contaminate the cycle, the HP filter is typically applied to X11 SA quarterly series. Throughout the paper, "X11" will denote the default linear filter for an additive decomposition, as in Ghysels and Perron (1993). To adjust a series, the filter X11 will always be applied (in the X11ARIMA spirit) to the series extended at both extremes with ARIMA forecast and backcasts. We shall center attention, first, on historical (or final) estimation. If $\nu_{X11}(B, F)$ denotes the X11-SA filter, and $\nu_{HP}^c(B, F)$ the HP filter (2.9), let $\nu_{HPX}^c(B, F)$ denote the convolution of the two. Because both, the X11 and the HP filters, are symmetric, centered, and convergent, so will their convolution. For seasonal series, the estimator of the cycle (2.9) should thus be replaced by

$$\hat{c}_t = \nu_{HPX}^c(B, F)x_t. \quad (2.11)$$

Throughout the paper we assume quarterly series and denote by S the annual aggregation operator, $S = 1 + B + B^2 + B^3$. Further, in all decompositions of a series into unobserved stochastic components, the components will be assumed orthogonal, and innovations in their models normally distributed.

3 Revisions

3.1 Preliminary estimation of end points and revisions

If $\hat{c}_{T|T}$ denotes the estimator of the cycle for the last observed period (i.e., the concurrent estimator,) as new periods are observed the estimator will be revised to $\hat{c}_{T|T+1}, \hat{c}_{T|T+2}, \dots$ until it converges to the final estimator \hat{c}_T . The difference between the final estimator and the concurrent one measures the revision the latter will undergo, and can be interpreted as a measurement error in the concurrent (more generally, preliminary) estimator. Although the poor behavior of the HP filter for recent periods has often been pointed out (see Baxter and King 1995), the revisions implied by HP filtering have not been analyzed. Two main features of the revision are of interest: a) the magnitude, and b) the duration of the revision process (i.e., the value of k for which $\hat{c}_{T|T+k}$ has, in practice, converged). To look at these features we use the WK version of the filter.

Assume the observed series follows the general ARIMA model

$$\phi(B)x_t = \theta(B)a_t, \quad (3.1)$$

where, without loss of generality, we assume $\phi(B)$ contains the factor ∇^d , $0 \leq d \leq 4$. Because the numerator of ν_{HP}^c in (2.10) cancels the unit roots in x_t , the estimator of the cycle can be expressed as

$$\hat{c}_t = \xi(B, F)a_t, \quad (3.2)$$

where the weights of the polynomial $\xi(B, F)$ can be obtained through the identity $\xi(B, F)\phi(B) = \nu_{HP}^c(B, F)\theta(B)$. Expression (3.2) can be rewritten as

$$\hat{c}_t = \xi^-(B)a_t + \xi^+(F)a_{t+1}, \quad (3.3)$$

where $\xi^-(B) = \sum_{j \geq 0} \xi_{-j} B^j$, and $\xi^+(F) = \sum_{j \geq 0} \xi_j F^j$ are convergent polynomials. The first one contains the effect of the innovations up to and including period t , and the second one includes the effect of innovations posterior to period t . Because $E_t(a_{t-j}) = a_{t-j}$ when $j \geq 0$, and $E_t(a_{t-j}) = 0$ when $j < 0$, the concurrent estimator equal to the expectation at time t of the estimator (3.3), is given by the first term in the right hand side of the equation. The revision in the concurrent estimator will thus be given by

$$r_{t|t} = \hat{c}_{t|t} - \hat{c}_t = \xi^+(F)a_{t+1} = \sum_{j=1}^k \xi_j a_{t+j}, \quad (3.4)$$

where the last equality uses a finite approximation based on the convergence of $\xi^+(F)$. From (3.4), it is straightforward to compute the variance and autocorrelations of the revision process. (We have focussed on the concurrent estimator; the analysis is trivially extended to any preliminary estimator $\hat{c}_{t|T}$.)

Although the filter $\nu_{HP}^c(B, F)$ is fixed, the coefficients of the forward filter $\xi^+(F)$ depend on the ARIMA model for the observed series. Without loss of generality, we set $\text{Var}(a_t) = 1$, so that the variance of the revision

$$\text{Var}(r_{t|t}) = \sum_{j=1}^k (\xi_j)^2 \tag{3.5}$$

is then expressed as a fraction of the variance of the series innovation V_a . To give some examples of the magnitude of the revision error, for a random walk series $\sigma(r_{t|t}) = 0.91\sigma_a$, and for the IMA(2,2) model (2.5) with the parameters equal to (2.7), the model for which the HP filter is optimal, $\sigma(r_{t|t}) = 0.34\sigma_a$. Defining convergence as having removed 95% of the revision variance, in both examples it takes 9 quarters for the estimator to converge. Of more empirical relevance is the case of the HP filter applied to X11-SA series (i.e., the filter $\nu_{XHP}(B, F)$). It is well known that X11, another two-sided filter, also produces revisions. Therefore, the revisions associated with the filter ν_{XHP} will reflect the combined effect of the two filters. For illustration, we select the so-called ‘‘Airline model’’, discussed in Box and Jenkins (1970), given by the expression

$$\nabla \nabla_4 x_t = (1 + \theta_1 B)(1 + \theta_4 B^4) a_t. \tag{3.6}$$

The model fits well many series with trend and seasonality (see the results for 14000 series from 17 countries in Fischer and Planas 1998), and has become a standard example. For the most relevant range for the parameters θ_1 and θ_4 , Table 1 presents the fraction $\sigma(\text{revision})/\sigma(a_t)$ and the number of periods (τ) needed for convergence. The standard deviation of the revision represents between 0.4 and 1.5 of $\sigma(a_t)$, and convergence takes, roughly, between 2 and 5 years. Given that θ_1 close to -1 implies very stable trends, while θ_4 close to -1 implies very stable seasonals, what Table 2 shows is that series with highly moving trends and seasonals will be subject to bigger, longer lasting, revisions. It is worth pointing out that, for the range of values most often found in practice which is the bottom right corner, the revision period lasts between 9 and 15 quarters. The two examples clearly indicate two features. First, that the revision error is quantitatively important, of a magnitude often comparable to that of the 1-period-ahead forecast; and, second, that the revision period lasts more than 2 years.

Table 1. Revisions implied by the HP-X11 filter

	$\theta_4 = 0$		$\theta_4 = -.2$		$\theta_4 = -0.4$		$\theta_4 = -0.6$		$\theta_4 = -0.8$	
	σ_r/σ_a	τ	σ_r/σ_a	τ	σ_r/σ_a	τ	σ_r/σ_a	τ	σ_r/σ_a	τ
$\theta_1 = 0.4$	1.53	19	1.44	18	1.36	17	1.28	9	1.21	9
$\theta_1 = 0.2$	1.34	19	1.26	18	1.18	17	1.12	9	1.06	9
$\theta_1 = 0$	1.15	19	1.08	18	1.02	16	0.96	9	0.90	9
$\theta_1 = -0.2$	0.97	19	0.91	18	0.85	15	0.80	9	0.76	9
$\theta_1 = -0.4$	0.79	18	0.74	17	0.70	14	0.65	9	0.61	9
$\theta_1 = -0.6$	0.64	15	0.60	14	0.55	9	0.51	9	0.47	9
$\theta_1 = -0.8$	0.52	9	0.48	9	0.44	9	0.40	9	0.36	9




Fig. 2. Short-term economic indicator: original series

3.2 An example

An application, that will also be used in later sections, will complete the discussion. We consider four quarterly Spanish economic indicators that could be related to the business cycle. The series are the industrial production index (IPI), cement consumption (CC), car registration (CR) and airline passengers (AP), for the 26-year period 1972/1 - 1997/4, and contain 104 observations. (For the IPI series, the first 12 observations were missing and the period was completed using backcasts). The series were log transformed (following proper comparison of the BIC criteria), and the application will be discussed for the additive decomposition of the logs. So as to facilitate comparisons, we standardize the 4 logged series to have zero mean and unit variance. The 4 series are represented in Fig. 2; their trend and seasonal features are clearly discernible. ARIMA modeling of the 4 series produced similar results: the models were of the type (3.6) and a summary of results is given in Table 2; none of the series appeared to be in need of outlier adjustment. (Estimation was made with the program TRAMO run in an automatic mode, see Gómez and Maravall 1996). Using the ARIMA models to extend the series, the X11 filter was applied to obtain the SA series. Then, the HP ($\lambda = 1600$) filter was applied. The 4 cycles obtained are displayed in Fig. 3. For the series CC and CR the short-term contribution of the cyclical variation is relatively more important than for the series IPI and, in particular, AP.

The cycles of Fig. 3 are a combination of concurrent, preliminary, and final estimators. Since, on occasion, the filter is treated as a one-sided filter (see Prescott 1986), an interesting comparison is the following. Using the first and

Table 2. Summary of ARIMA estimation results

	Parameter Estimates		Residual	BL test	Normality
	θ_1	θ_4	Variance V_a	$Q(< \chi^2_{14})$	$N(< \chi^2_2)$
CC	-0.405	-0.957	0.175	18.4	0.32
IPI	-0.299	-0.721	0.054	23.3	0.14
CR	-0.387	-0.760	0.156	18.7	0.79
AP	-0.392	-0.762	0.017	21.1	2.76

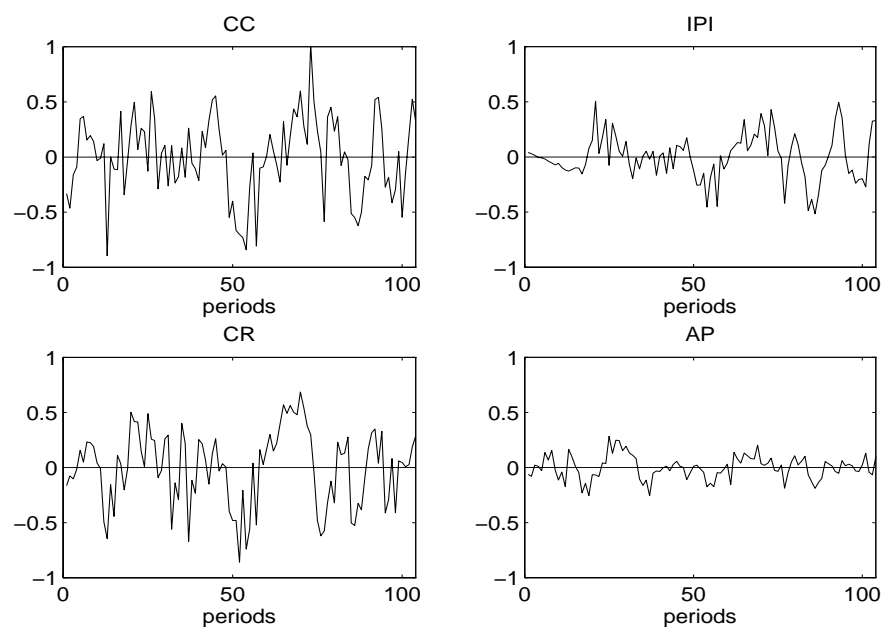


Fig. 3. X11-HP cycles

last 22 periods for safe convergence of the X11 and the HP filters, we obtained the sequence of concurrent and final estimators of the trend and cycle for the 60 central periods of the 4 series. Then, we evaluated the standard loss function of the HP filter, given by (2.2), for the concurrent and final estimators of the trend and cycle; the results are given in Table 3.

Table 3. HP loss-function for concurrent and final estimator

CC		IPI		CR		AP	
Concurrent	Final	Concurrent	Final	Concurrent	Final	Concurrent	Final
624.7	13.3	172.8	2.9	513.4	11.9	43.2	1.0

The improvement achieved by using final estimators instead of concurrent ones is indeed large. Figure 4 compares the series of concurrent and final estimators, for the cycle. A clear phase effect in the concurrent estimator can be observed for the 4 series, a well-known feature of one-sided filters.

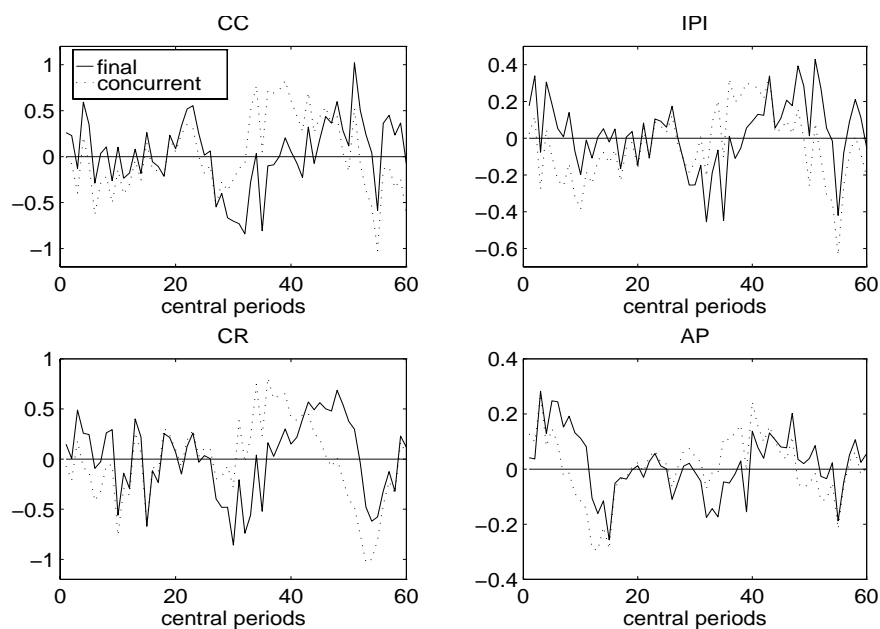


Fig. 4. Concurrent versus final cycle estimator

A point of applied relevance is to assess the imprecision of the estimator of the cycle for recent periods, as measured by the standard error of the revision. In so far as the revision represents a measurement error, its variance can be used to build confidence intervals around the cycle estimator. Figure 5 displays the 95% confidence interval for the 4 series. Direct inspection shows that, although the estimator converges in 2 (at most 3) years, for recent periods it is unreliable. This large increase in the measurement error of the most recent signals implies that forecasts would be of little use.

4 Spurious results

While the problem of revisions has been often overlooked, the danger of obtaining spurious results induced by HP filtering has been frequently mentioned. The squared gain of $\nu_{HPX}^c(B, F)$ is shown in Fig. 6. It displays zeros for the zero and seasonal frequencies. The application of this fixed structure brings the possibility of spurious results. On the one hand, it will affect the autocorrelation structure of the series and spurious correlations between series may be obtained (in the line of Granger and Newbold 1974). On the other hand, the first peak may induce a spurious periodic cycle.

4.1 Spurious crosscorrelation

We performed a simulation in MATLAB, whereby 10.000 independent random samples of 600 observations each were drawn from a $N(0, 1)$ distribution. Each




Fig. 5. 95% confidence intervals for cycle (based on revisions)




Fig. 6. Squared gain: Convolution of HP and X11 filters

white-noise series was filtered through the X11 and HP filters and the last 100 values were selected. Next, 10.000 lag-zero crosscorrelation between two series were sampled (in what follows, all crosscorrelations are lag-zero ones). The same exercise was performed to generate independent random walks series. Table 4 presents the first four moments of the distribution of $\hat{\rho}_0$, the crosscorrelation estimator between the original series and between the cycles.

Table 4. Crosscorrelation between filtered series

a. White noise	Mean	Std.deviation	Skewness	Kurtosis
Original	-0.001	0.11	-0.03	2.9
Cycle	-0.001	0.11	-0.05	2.9
b. Random walk	Mean	Std. deviation	Skewness	Kurtosis
Original	-0.000	0.10	-0.04	2.9
Cycle	0.000	0.19	-0.01	2.8

Clearly, for the white-noise case no spurious crosscorrelations has been induced. For the random walk model the zero-mean normality assumption can still be accepted but the spread of the distribution of $\hat{\rho}_0$ for the cycle becomes wider. A spurious, though moderate, crosscorrelation effect can thus be detected. A similar simulation was performed for the more complex airline model (3.6), with the parameter values set at $\theta_1 = -0.4$ and $\theta_4 = -0.6$. Figure 7 plots the densities of the crosscorrelation estimator for the stationary transformation of the original and SA series and of the X11-SA and HP detrended series. The filter X11 is seen to have virtually no effect while, as before, the HP filter induces a small increase in the spread of the distribution. Still, the HP-X11 filter seems to induce a small amount of spurious crosscorrelations and hence the detection of relatively large crosscorrelation between cycles obtained with it is unlikely to be spurious. (Although the filter will have distorting effects when the series are correlated; see Cogley and Nason 1995.)

4.2 Spurious autocorrelation; calibration

Assume that a theoretical economic model implies that an economic variable follows the 4-year cycle AR(2) process:

$$(1 - 1.293B + 0.490B^2)c_t = a_{ct}, \quad \text{Var}(a_{ct}) = 1, \quad (4.1)$$

with a_{ct} a white-noise innovation, and that simulations of the model yield in fact an ACF for the variable equal to the theoretical ACF of (4.1), shown in the second column of Table 5. The basic idea behind calibration is to validate the economic model by comparing the previous ACF with the one implied by the observed economic variable. To compute the latter, the non-stationary trend and seasonal component need to be removed. (Besides, seasonality and often the trend are typically excluded from the theoretical economic model.)

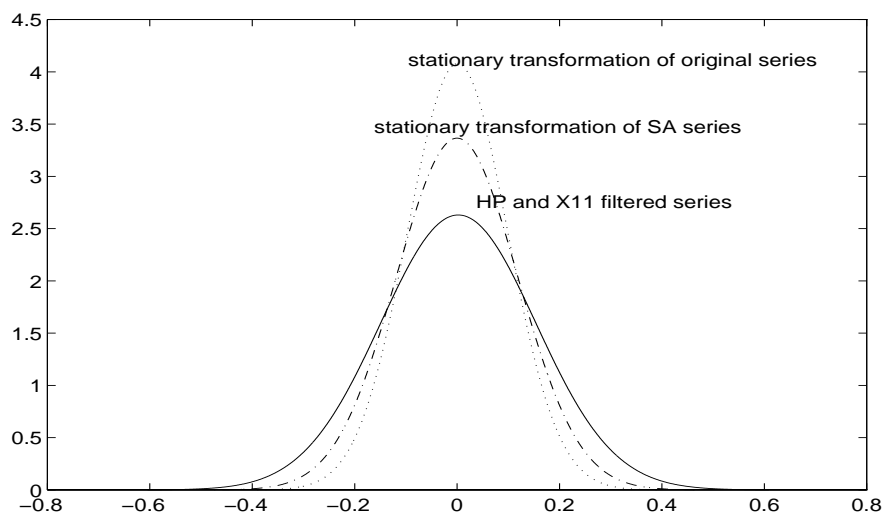


Fig. 7. Density for correlation coefficient: Airline model

Table 5. Theoretical ACF of the component model and of its estimators

Lag-k ACF	True component	X11-HP filtered component			MMSE estimator
		$V_p = 0.1, V_s = 0.1$	$V_p = 0.1, V_s = 1$	$V_p = 1, V_s = 1$	$V_p = 0.1, V_s = 1$
k=1	0.87	0.71	0.19	0.37	0.83
k=2	0.63	0.44	0.22	0.30	0.43
k=3	0.39	0.10	-0.06	0.00	-0.02
k=4	0.20	-0.05	0.22	0.18	-0.35
k=5	0.06	-0.25	-0.23	-0.15	-0.45
k=6	-0.01	-0.30	-0.19	-0.16	-0.43
k=7	-0.05	-0.34	-0.27	-0.26	-0.31
k=8	-0.06	-0.27	-0.01	-0.07	-0.20
k=9	-0.05	-0.25	-0.18	-0.20	-0.10
k=10	-0.04	-0.19	-0.12	-0.16	-0.04
k=11	-0.02	-0.16	-0.17	-0.20	-0.00
k=12	-0.01	-0.09	0.06	-0.03	-0.02
k=13	-0.00	-0.08	-0.07	-0.12	-0.03
k=14	-0.00	-0.05	-0.03	-0.08	-0.03
k=15	-0.03	-0.04	-0.11	-0.13	-0.02
k=16	-0.00	-0.01	-0.13	0.05	-0.01

Let the observed series be generated by the cycle given by (4.1), contaminated by a random walk trend (p_t) and a seasonal component (s_t) as in the Basic Structural Model of Harvey and Todd (1983). Thus the observed series x_t is given by $c_t + p_t + s_t$, where c_t is generated by (4.1), and $\nabla p_t = a_{pt}, Ss_t = a_{st}$, with a_{ct}, a_{pt} and a_{st} mutually orthogonal innovations, with variances V_c, V_p and V_s . Seasonally adjusting with X11 and detrending with the HP filter the observed series, the estimator of the cycle is obtained. Its variance and ACF (the "observed moments"

in the calibration comparison) are straightforward to derive analytically; they are given in the third, fourth and fifth column of Table 5, for the three cases $V_p = V_s = 0.1$; $V_p = 0.1, V_s = 1$; and $V_p = V_s = 1$. Comparing these columns with the second, although the theoretical model is perfectly correct, the second moments obtained from the observed series would seem to indicate the contrary. The distortion that seasonal adjustment and detrending induce also occurs when the components are estimated as MMSE estimators in a model based approach; the ACF of the cycle obtained in this case is given in the 6th row of Table 5. Still, the distortion induced by MMSE estimation is considerably smaller than that induced by HP-X11 filtering. Calibration of models using filtered series seems, thus, an unreliable procedure. If the theoretical economic model is correct, then calibration should not look for similarities between the ACF of the theoretical model and of the empirical series. It should compare instead the empirical moments with the theoretical ones that include the effect of filtering the data. Performing this comparison, however, requires incorporating into the model trend and seasonality.

4.3 Spurious periodic cycle

4.3.1 Random-walk input. We consider the simplest case of a model with trend, namely the random-walk model, $\nabla x_t = a_t$; for which, the cycle is estimated through (2.12). Letting ω denote the frequency in radians, if $\tilde{\nu}_{HPX}^c(\omega)$ is the Fourier transform of $\nu_{HPX}^c(B, F)$, the spectrum of the estimator of the cycle is given by,

$$\hat{g}_{HPX}^c(\omega) = [\tilde{\nu}_{HPX}^c(\omega)]^2 g_x(\omega), \quad (4.2)$$

where $g_x(\omega)$ is the pseudospectrum of x_t (see Harvey 1989); thereafter the term spectrum will also be used to refer to a pseudospectrum. (To simplify notation, all spectra will be implicitly expressed in units of $1/2\pi$.) Figure 8 plots the spectra of the series (dotted line) and of the cycle (continuous line). The latter displays a peak for a frequency in the cyclical range associated with a period of 31-32 quarters, or, approximately, 8 years. Yet, by its own definition, does it make sense to see a random walk as generated by a trend and a 8-year cycle? Is it not rather a case of “overreading” the data? The answer to this question is not quite so obvious, as we proceed to discuss.

4.3.2 Spectral characteristics of the cycle. The cycle obtained with X11-HP filtering displays a stochastic structure which depends on the ARIMA model followed by the observed series, and on the λ -parameter of the HP filter. To look, first at the effect of the model, we set $\lambda = 1600$. Figure 9 compares the cycles obtained when the series follows the IMA(1,1) model $\nabla x_t = (1 + \theta B)a_t$, with $V_a = 1$, for a range of values for θ . In all cases, the period associated with the spectral peak of the cycle is approximately constant, and very close to 8 years. The amplitude of the cycle varies, adapting to the width of the spectral peak for $\omega = 0$ in the series model, which is determined by the parameter θ .

The relative constancy of the period with respect to the model parameter is also shown in Table 6 for a MA(1) and an IMA(2,1) models. What the table seems

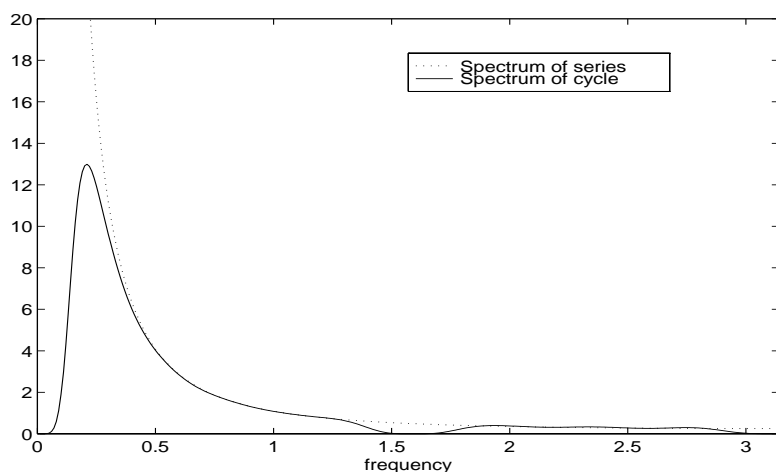


Fig. 8. Spectrum of cycle component in a random walk




Fig. 9. Spectrum of cycle in IMA(1,1) as a function of theta

to indicate is that, for a fixed value of λ , the period of the cycle is determined fundamentally by the order of integration of the series, rather than by the model parameters. As the order of integration increases, so does the period of the cycle.

Table 6. Period of cycle (in years)

Theta	0	-0.3	-0.6
MA(1)	2	3	3.2
IMA(1,1)	7.9	7.9	7.9
IMA(2,1)	10.5	10.5	10.5

When the HP filter is applied to an X11 SA series, a similar effect is seen to occur. For the Airline model (3.6) we computed the period associated with the spectral peak of the cycle for the range $-0.9 < \theta_1 < 0.5$ and $-0.9 < \theta_4 < 0$, and in all cases the period was approximately equal to 10 years. For fixed parameter λ , three conclusions emerge: (1) Given the type of ARIMA model for the series, the associated cyclical period becomes roughly fixed. (2) The period seems to be mostly determined by the order of integration at the zero frequency; the stationary part of the model has little influence. (3) For most actual time series containing a trend ($d=1$ or 2), the standard value of $\lambda = 1600$ implies a period between 8 and 10 years.

Fixing now the model to that of a random walk, the dependence of the cycle period on λ is shown in Fig. 10. The line represents the value of λ associated with the period of the spectral peak of the cycle when only the HP filter is applied. It can be seen that the convolution of X11 has little effect on the period of the cycle spectral maximum (in fact the two figures would be indistinguishable). The relationship between this period and λ is highly nonlinear. When λ is small (and the period of the cycle relatively short) small changes in λ have a very strong impact on the period; for long cycles, very large values of λ need to be used.




Fig. 10. Period of cycle as a function of lambda

The effect of λ is illustrated in Fig. 11, which compares the spectra of the cycles obtained with $\lambda = 1600$ and $\lambda = 25000$ for the same random walk series (the periods associated with the spectral peaks are about 8 and 15 years, respectively). The figure shows that the longer period implies a stochastic cycle more concentrated around its peak (i.e., a more stable cycle). The estimators of the trend and cycle for the two λ values are compared in Fig. 12 and Fig. 13, respectively. The difference between the two trends is seen to consist of a cycle




Fig. 11. Spectrum of a cycle in a random walk




Fig. 12. Estimated trends

with a relatively long period. Comparison of the cycle shows that its short-term profile remains basically unchanged, the main effect being a "pulling away" from the zero line, which allows for longer cycles.

As a consequence, the use of the X11-HP filter (or simply the HP filter) to measure the cycle implies an a-priori choice: The analyst should first decide the length of the period around which he wishes to measure economic activity. Then, given d (the number of unit roots at the zero frequency in the series), he can choose the appropriate value of λ . For example, a business cycle analyst involved in policy making may be interested in using 8 or 10-years cycles; an economic historian looking at several centuries, may be interested in spreading activity over longer periods. Viewed in this way, the HP cycle cannot be seen




Fig. 13. Estimated cycles

as spurious but as an arbitrary yet perhaps sensible way to look at the data. This statement will be made more precise at the end of Section 6.

5 Improving the Hodrick-Prescott filter

In Section 3 we saw that the filter implies large revisions for the last 2 years. The imprecision in the cycle estimator for the last quarters implies, in turn, a poor performance in the detection of turning points. Further, direct inspection of Figure 3 shows another limitation of the filter: the cyclical signal it provides seems rather uninformative. Seasonal variation has been removed, but a large amount of noise remains in the signal. Averaging over the 4 series, the number of times the series crosses the zero line is 31 times, over a period of 26 years! In the next two sections, we show how these two shortcomings can be reduced with relatively simple modifications.

5.1 Reducing revisions

Estimation of the cycle for the end periods of the series by the HP filter implies a truncation of the filter. In terms of the model based interpretation, this truncation is equivalent to the assumption that model (2.5) is always the model that generates forecasts to extend the series at both end points. The assumption will in general be false, and proper optimal forecasts (obtained with the appropriate ARIMA model for the series) can be used instead to improve the filter extension. This idea is the same as the one behind X11 ARIMA (see Dagum 1980) and the HP filter applied to the series extended with ARIMA forecasts will be referred as the Hodrick-Prescott ARIMA (HPA) filter. The poor performance of the HP filter at the end of the series has been often pointed out (see, for example, Apel et al

1996; Baxter and King 1995) and application of the filter to series extended with forecasts is often recommended in practice (see EU Commission 1995).

For any positive integer k , write the final estimator of the cycle as

$$\hat{c}_t = \nu_{HP}^c(B, F)x_t = \sum_{j=0}^{\infty} \nu_{j+k} x_{t+k-j} + \sum_{j=1}^{\infty} \nu_{j+k} x_{t+k+j}, \quad (5.1)$$

and assume a series long enough so as to ignore starting values. Because the preliminary estimator $\hat{c}_{t|t+k}$ is a projection onto a subset of the set onto which \hat{c}_t is projected, it follows that $\hat{c}_{t|t+k} = E_{t+k}(c_t) = E_{t+k}(\hat{c}_t)$, or

$$\hat{c}_{t|t+k} = \sum_{j=0}^{\infty} \nu_{j+k} x_{t+k-j} + \sum_{j=1}^{\infty} \nu_{j+k} E_{t+k}(x_{t+k+j}), \quad (5.2)$$

which expresses the preliminary estimator as a function of the series extended with forecasts. Subtracting (5.2) from (5.1), the revision in $\hat{c}_{t|t+k}$ is equal to

$$r_{t|t+k} = \sum_{j=1}^{\infty} \nu_{j+k} e_{t+k}(j),$$

where $e_{t+k}(j)$ denotes the forecasts error associated with forecasting j periods ahead the variable at time $t+k$. It follows that, reducing these forecasts errors, revisions should decrease (and early detection of turning points should improve).

We performed a simulation exercise. First, we consider the IMA(1,1) model for different values of the θ -parameter. Then, we consider the ARIMA (2,1,1) model, where the AR(2) polynomial is given by $(1 - 0.16B + 0.35B^2)$. This polynomial is the one found in Jenkins(1975) for the mink-muskrat Canadian data, and contains a cycle of period 4.4. The AR(2) structure will therefore produce an increase in the number of turning points. (In order to avoid effects due to the SA filter, no seasonality was entered into the models.) Again, different values of the θ -parameter were considered. A total of 14.000 series of length 100 each were simulated, and for each series the HP filter was compared to the HPA one extended with 16 ARIMA forecasts and backcasts. Table 7 compares the variances of the revision in the concurrent estimator and in the estimator revised after 1, 2, 3 and 4 more years of data are added. It is seen that, in all 70 cases, the HPA filter reduces considerably revisions. This is particularly noticeable for the ARIMA(2,1,1) model, where the use of the standard HP filter may more than triplicate the revision variance.

As for the detection of turning points, we use the following simple criterion (along the lines of method B in Boldin 1994): a turning point is the first of at least two successive periods of negative/positive growth. Table 8 compares the performance of the HP and HPA filters in the first and last 8 observations of the simulated series, both in terms of the mean number of turning points that are dated on the original series and missed by the filtered one, and in terms of the mean number of turning points detected on the filtered series but not present in the original one ("peaks" and "throughs" are considered separately). Of the 56

Table 7. Variance of the revision in estimator. Values are multiplied by 100

Model	Concurrent		1 year rev.		2 year rev.		3 year rev.		4 year rev.	
	HPA	HP	HPA	HP	HPA	HP	HPA	HP	HPA	HP
IMA(1,1)										
$\theta = -0.8$	0.31	0.41	0.08	0.10	0.02	0.03	0.01	0.01	0.00	0.01
$\theta = -0.5$	0.94	1.34	0.24	0.33	0.07	0.11	0.03	0.06	0.02	0.04
$\theta = -0.3$	1.58	2.54	0.39	0.63	0.11	0.19	0.06	0.12	0.04	0.08
$\theta = 0$	2.86	4.84	0.71	1.18	0.21	0.39	0.12	0.24	0.08	0.15
$\theta = 0.3$	4.51	8.29	1.11	2.03	0.32	0.64	0.19	0.38	0.13	0.24
$\theta = 0.5$	5.44	11.02	1.40	2.62	0.44	0.86	0.26	0.50	0.17	0.32
$\theta = 0.8$	7.33	14.89	1.87	3.70	0.60	1.17	0.34	0.70	0.22	0.46
ARIMA(2,1,1)										
$\theta = -0.8$	0.12	0.55	0.05	0.14	0.01	0.03	0.00	0.01	0.00	0.01
$\theta = -0.5$	0.41	1.27	0.15	0.33	0.04	0.09	0.01	0.04	0.00	0.03
$\theta = -0.3$	0.74	2.23	0.25	0.56	0.06	0.15	0.02	0.08	0.01	0.05
$\theta = 0$	1.35	4.26	0.44	1.09	0.12	0.29	0.03	0.15	0.02	0.10
$\theta = 0.3$	2.06	7.00	0.71	1.77	0.18	0.46	0.05	0.25	0.03	0.18
$\theta = 0.5$	2.75	9.68	0.95	2.44	0.23	0.64	0.06	0.33	0.03	0.22
$\theta = 0.8$	3.70	12.95	1.20	3.25	0.31	0.88	0.09	0.47	0.05	0.33

Table 8. Mean number of turning points (First and last 8 observations)

IMA(1,1)	Original		Missed				False Alarms			
	Peaks	Throughs	Peaks		Throughs		Peaks		Throughs	
			HPA	HP	HPA	HP	HPA	HP	HPA	HP
$\theta = -0.8$	1.49	1.50	0.10	0.16	0.10	0.15	0.18	0.19	0.19	0.20
$\theta = -0.5$	1.51	1.52	0.18	0.22	0.19	0.23	0.29	0.31	0.26	0.28
$\theta = -0.3$	1.52	1.52	0.22	0.28	0.24	0.30	0.36	0.37	0.38	0.39
$\theta = 0$	1.62	1.59	0.23	0.32	0.25	0.37	0.49	0.51	0.47	0.49
$\theta = 0.3$	1.69	1.72	0.29	0.40	0.29	0.41	0.49	0.59	0.55	0.63
$\theta = 0.5$	1.77	1.79	0.32	0.46	0.29	0.41	0.54	0.68	0.51	0.63
$\theta = 0.8$	1.83	1.86	0.30	0.43	0.34	0.48	0.52	0.65	0.54	0.69
ARIMA(2,1,1)										
$\theta = -0.8$	1.78	1.79	0.05	0.13	0.05	0.13	0.21	0.18	0.19	0.18
$\theta = -0.5$	1.78	1.77	0.09	0.19	0.11	0.19	0.25	0.24	0.26	0.26
$\theta = -0.3$	1.84	1.77	0.12	0.23	0.10	0.20	0.25	0.31	0.25	0.30
$\theta = 0$	1.87	1.84	0.13	0.25	0.16	0.27	0.27	0.38	0.26	0.36
$\theta = 0.3$	1.93	1.89	0.17	0.27	0.20	0.31	0.25	0.40	0.25	0.41
$\theta = 0.5$	1.94	1.96	0.18	0.31	0.17	0.32	0.24	0.40	0.25	0.42
$\theta = 0.8$	2.02	1.95	0.18	0.32	0.16	0.31	0.22	0.38	0.24	0.41

comparisons, in 53 cases the gain from using the HPA filter is substantial. Table 9 compares the performance of the two filters when all observations in each series are considered, in terms of correctly detected turning points and spurious turning points (indicated by the filtered series and not present in the original one). F_0 is the relative frequency of cases in which the two filters coincide, F_1 denotes the relative frequency of cases in which HPA performs better, while F_{-1} denotes the relative frequency of cases in which HP performs better. The HPA filter performs (in all 56 cases) consistently better.

In summary, the results of the simulation exercise show that applying the HP filter to the series extended at both ends with appropriate ARIMA forecasts and

Table 9. Relative performance of HP vs HPA: captured and spurious turning points

Model	Capt. peaks			Capt. Throughs			False peaks			False Throughs		
	F_1	F_0	F_{-1}	F_1	F_0	F_{-1}	F_1	F_0	F_{-1}	F_1	F_0	F_{-1}
IMA(1,1)												
$\theta = -0.8$	0.07	0.92	0.01	0.07	0.93	0.00	0.06	0.91	0.03	0.05	0.91	0.04
$\theta = -0.5$	0.07	0.91	0.02	0.07	0.89	0.04	0.09	0.85	0.06	0.10	0.84	0.06
$\theta = -0.3$	0.11	0.84	0.05	0.10	0.86	0.04	0.12	0.81	0.07	0.13	0.78	0.09
$\theta = 0$	0.14	0.80	0.06	0.16	0.80	0.04	0.15	0.74	0.11	0.15	0.74	0.11
$\theta = 0.3$	0.17	0.78	0.05	0.17	0.78	0.05	0.22	0.68	0.09	0.20	0.70	0.10
$\theta = 0.5$	0.18	0.76	0.06	0.19	0.74	0.07	0.23	0.68	0.09	0.20	0.70	0.10
$\theta = 0.8$	0.18	0.76	0.06	0.20	0.75	0.05	0.25	0.64	0.11	0.24	0.65	0.11
ARIMA(2,1,1)												
$\theta = -0.8$	0.11	0.88	0.01	0.11	0.88	0.01	0.07	0.86	0.07	0.06	0.84	0.10
$\theta = -0.5$	0.14	0.84	0.02	0.13	0.84	0.03	0.10	0.80	0.10	0.10	0.81	0.09
$\theta = -0.3$	0.17	0.79	0.04	0.16	0.80	0.04	0.15	0.77	0.08	0.14	0.78	0.08
$\theta = 0$	0.18	0.77	0.04	0.18	0.77	0.05	0.19	0.74	0.07	0.19	0.75	0.06
$\theta = 0.3$	0.18	0.77	0.05	0.20	0.73	0.07	0.28	0.66	0.06	0.25	0.67	0.08
$\theta = 0.5$	0.21	0.73	0.06	0.23	0.71	0.06	0.23	0.70	0.07	0.23	0.70	0.07
$\theta = 0.8$	0.21	0.74	0.05	0.22	0.73	0.05	0.28	0.66	0.06	0.25	0.69	0.06

backcasts is likely to provide a cycle estimator for recent periods that requires smaller revisions and improves detection of turning points.

5.2 Improving the cyclical signal

Concerning erraticity of the cycle estimator, illustrated in Figure 3, one possible improvement could come from using a more appropriate SA procedure. Since the width of the spectral peaks associated with seasonal frequencies vary across series, fixed filters such as X11 may over or underestimate seasonality. Having obtained an ARIMA model for the series, one could use, instead of X11, an ARIMA-model-based (AMB) type of adjustment, following the approach of Burman (1980) and Hillmer and Tiao (1982). We use the program SEATS (Gómez and Maravall, 1996) to seasonally adjust the 4 series of the example in Section 3. Figure 14 compares the cycles obtained by applying the HP filter to the AMB and X11 SA series, and Fig. 15 exhibits the spectra of the two cycles for the 4 series. It is seen that the estimates of the cycle obtained with the two SA series are close: Turning points are basically unchanged and the cyclical signal remains very noisy. (Figure 15 illustrates the overestimation of seasonality implied by the X11 filter for the case of the CC series: The "holes" that X11 induces for the seasonal frequencies are obviously too wide. The AMB method, instead, adjusts the width of the hole to the width of the peak.)

Given that the SA series produces a cyclical signal with too much noise it would seem that this signal could be improved by removing the noise from the SA series, n_t . Thus we decompose n_t as in

$$n_t = p_t + u_t, \tag{5.3}$$

where for the case of the Airline model, the trend-cycle p_t follows an IMA (2,2) model, say,




Fig. 14. HP cycle based on X11 and SEATS SA series

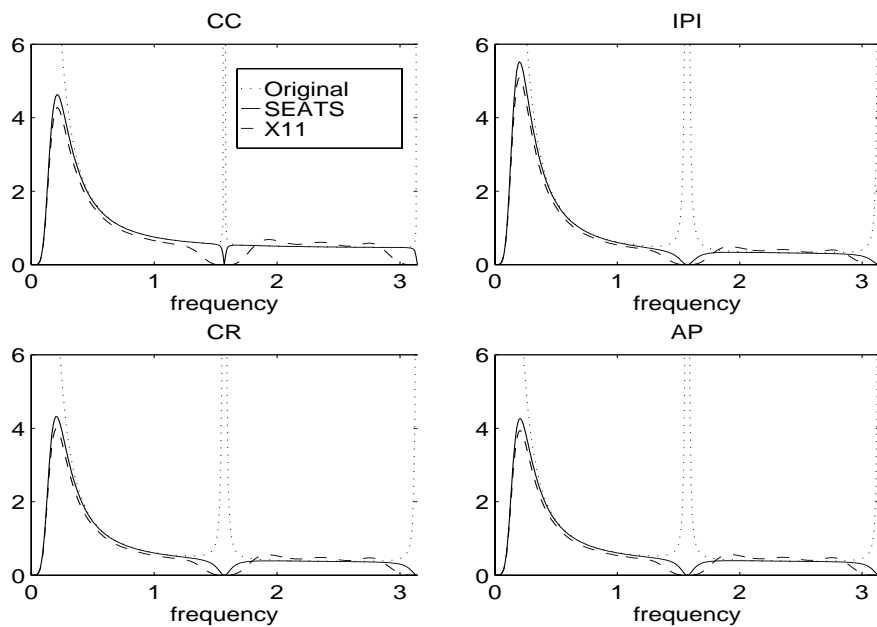


Fig. 15. Spectrum of cycle based on X11 and SEATS SA series

$$\nabla^2 p_t = \theta_p(B)a_{pt}, \quad \text{Var}(a_{pt}) = V_p, \quad (5.4)$$

where a_{pt} is white-noise and the polynomial $\theta_p(B)$ can be factorized as $(1 - \alpha B)(1 + B)$, with the second root reflecting a spectral zero for the frequency π , and α not far from 1. The irregular u_t , and the component innovations a_{st} , and a_{pt} are mutually orthogonal white noises. If $\theta(B) = (1 + \theta_1 B)(1 + \theta_4 B^4)$, the MMSE estimator of p_t is given by,

$$\hat{p}_t = \left[k_p \frac{\theta_p(B)S}{\theta(B)} \frac{\theta_p(F)\bar{S}}{\theta(F)} \right] x_t, \quad (5.5)$$

with $k_p = V_p/V_a$. From (5.3), the trend-cycle component p_t is seen to be the noise-free SA series. Using the HP filter on the trend-cycle estimator \hat{p}_t , the estimated cycles are displayed in Fig. 16 (continuous line). The use of the trend-cycle instead of the SA series drastically improves the cyclical signal, which becomes much cleaner. Figure 17 compares the spectra of the cycles obtained with the two series (p_t and n_t). The difference is due to the fact that the cycle based on p_t has removed variance associated with frequencies of no cyclical interest and, as shown in Fig. 18, the spectrum of the difference is close to that of white noise. So to speak, the band-pass features of the cycle are much better defined. This improvement of the cyclical signal allows for a clearer comparison of cycles among series, as is evidenced by comparing Figs. 19 and 20. (Figure 16 showed that for the series AP the cyclical component has become very small and hence it is not included.) Figure 20 shows that the series CC, IPI, and CR have fairly similar cyclical patterns, moving roughly in phase.

One further advantage of using the more stable signal p_t is that it produces a decrease in the size of the revisions in the cyclical estimate for the last periods, as shown in Fig. 21. Although the full revision process takes close to 10 years, in practice after two years most of the revision has been completed. Finally, Fig. 22 displays the 95% confidence interval for the cycle estimator for the full period, based on the associated revisions when the trend-cycle component is used as input. In our view, if cyclical analysis of the 4 series had to be summarized in one figure, Fig. 22 would be the appropriate choice.

6 Final remarks

We have seen how two serious drawbacks of the standard application of HP-X11 filtering to estimate cycles, namely the poor behavior of the estimator at the end of the series and the excessive noise in the cyclical signal, can be significantly reduced with two modifications that are straightforward to incorporate. The "modified HP filter" consists of applying the HPA filter to the trend-cycle component of the series, which requires extending the trend-cycle with optimal backcasts and forecasts. One simple way to implement the procedure is to run first TRAMO-SEATS in an automatic manner (this also yields forecasts of the trend-cycle); then, to apply the HP filter to the extended trend-cycle series. The




Fig. 16. HP cycle based on SEATS trend and on X11 SA series

two steps can be done as a single one in the following manner. Let the model identified for the series be

$$\nabla\nabla_4x_t = \theta(B)a_t, \tag{6.1}$$

where $\theta(B)$ is invertible. Convoluting the HP filter (2.10) with the filter (5.5) for the trend-cycle component yields the estimator of the cycle as a linear filter of the observations,

$$\hat{c}_t = \left[k \frac{\theta_p(B)\nabla\nabla_4}{\theta_{HP}(B)\theta(B)} \frac{\theta_p(F)\bar{\nabla}\bar{\nabla}_4}{\theta_{HP}(F)\theta(F)} \right] x_t, \tag{6.2}$$

where $k = V_p k_{c(HP)}$. (In the discussion, we use the expression for the final estimator. For preliminary estimators it is assumed that the series has been extended with the appropriate ARIMA forecasts.) Direct inspection shows that the filter in (6.2) is the autocovariance function (ACVF) of the stationary model

$$\theta_{HP}(B)\theta(B)z_t = \theta_p(B)\nabla\nabla_4b_t, \tag{6.3}$$

with $Var(b_t) = k$. It is well-known (see, for example, Maravall 1987) that if an ARIMA model is decomposed into signal plus white noise, the filter that yields the MMSE estimator of the noise is given by the ACVF of the inverse model (multiplied by the variance of the noise). The inverse model is the one that results from interchanging the AR and MA parts, that is, the expression in brackets in (6.2) is the ACVF of the inverse model of

$$\theta_p(B)\nabla\nabla_4x_t = \theta(B)\theta_{HP}(B)d_t. \tag{6.4}$$

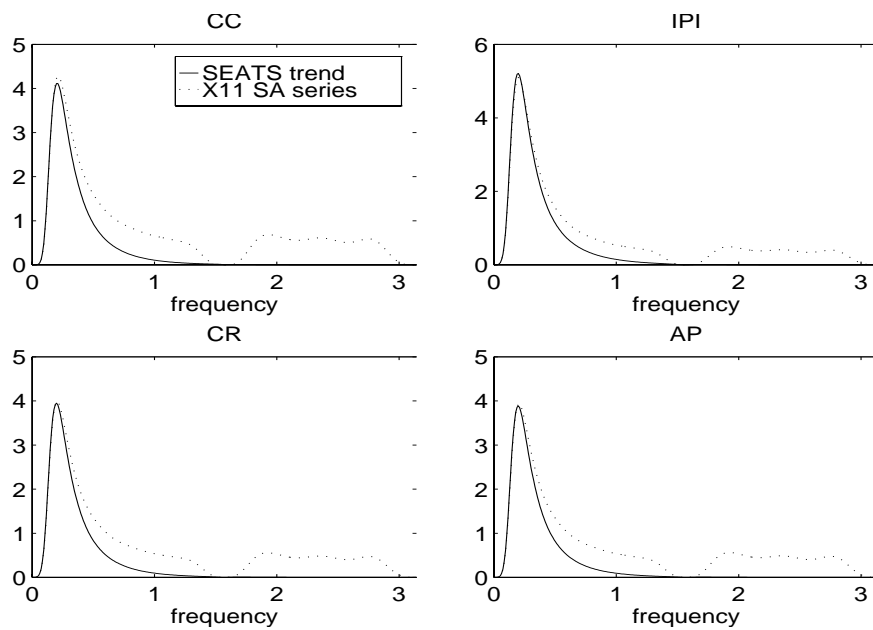


Fig. 17. Spectrum of cycle (SEATS trend and X11 SA series)




Fig. 18. Difference between cycles (SEATS trend and X11 SA series)

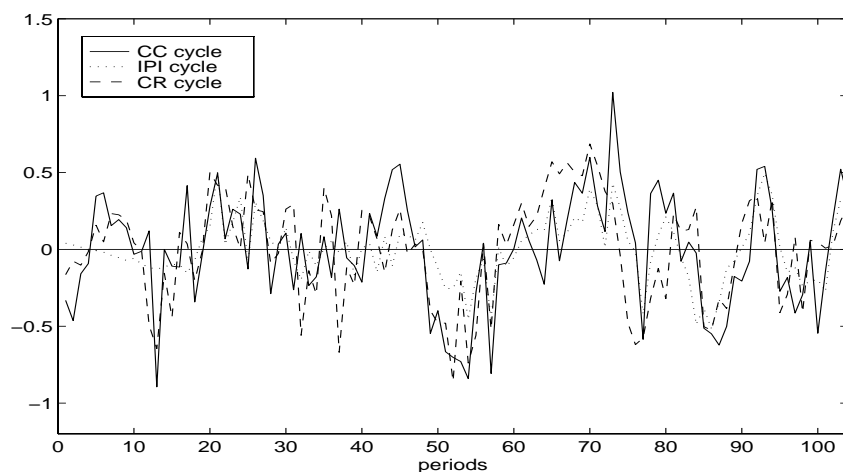


Fig. 19. HP cycles based on X11 SA series




Fig. 20. HP cycles based on SEATS trend

It follows that $\hat{\epsilon}_t$, given by (6.2), is the estimator of the noise in the decomposition of (6.4) into signal plus white-noise when the variance of the latter is k . This model-based interpretation of the modified HP filter provides a convenient algorithm (ACVF of ARIMA models are easy to compute; see the appendix in Box et al. 1978), but does not provide a sensible interpretation since, to start with, the cycle is not white noise. A full model-based interpretation of the complete decomposition into trend, cycle, seasonal, and irregular components, where the last two are the same as in the standard AMB decomposition, the trend and the cycle aggregate into the trend-cycle component of the AMB decomposition, the cycle is the modified HP filter advocated in this paper, and the components aggregate into the model specified for the series, is presented in Kaiser and Maravall (1999). Basically, meaningful model-based interpretations of the filter can

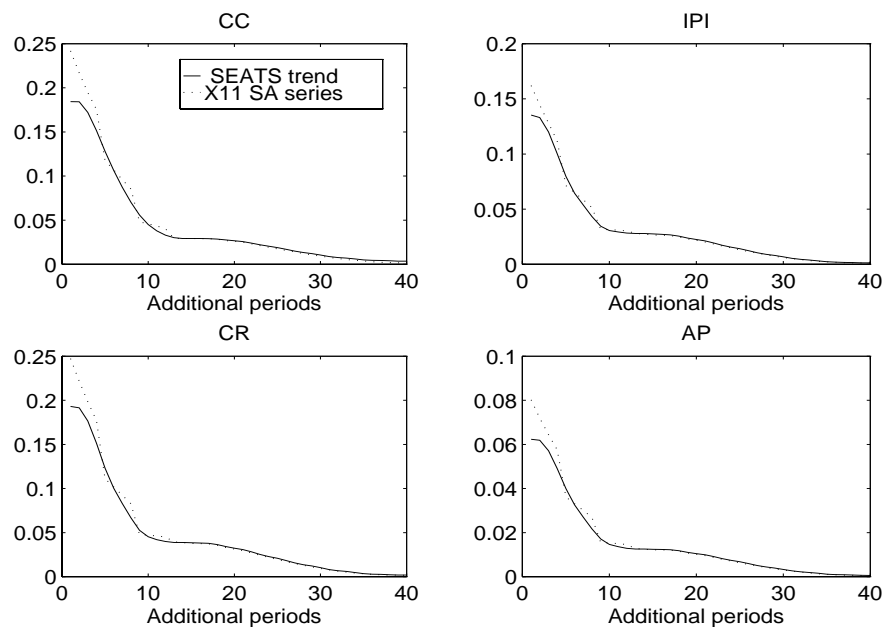


Fig. 21. Standard deviation of revision from concurrent to final estimation




Fig. 22. 95% CI for HP cycle (based on revisions)

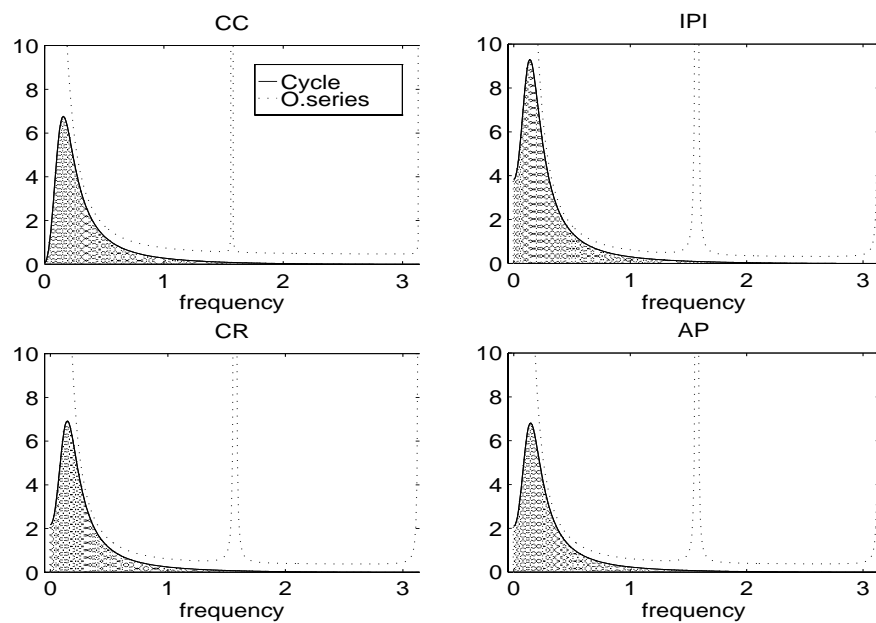


Fig. 23. Spectra of the series and of the cycle

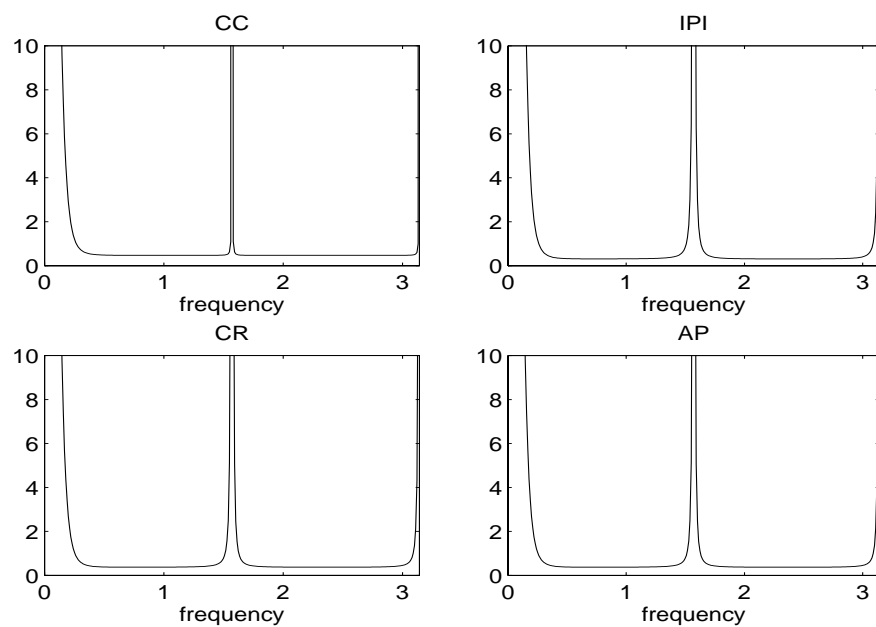


Fig. 24. Spectra of the difference (original series minus cycle)

be found by noticing that (6.2) is also the MMSE estimator of a cycle that follows the model

$$\theta_{HP}(B)c_t = \theta_p(B)a_{ct}, \quad (6.5)$$

when the ARIMA model for the series is (6.1) and $\text{Var}(a_{ct})/\text{Var}(a_t) = k$. Figure 23 displays the spectrum of the series x_t and of the cycle c_t . Figure 24 shows what is left of the series once the cycle is removed. In all 4 cases it consists of a stable trend, seasonal effects, and noise. This brings us back to the spuriousness issue. Assuming the ARIMA model identified for the series is acceptable, so will be the decomposition of Figs. 23 and 24, which aggregate into the series spectrum. One may prefer other components, but there would be nothing spurious about the results.

Appendix

Wiener-Kolmogorov version of the Hodrick-Prescott filter

We present an algorithm to compute the HP trend with the Wiener Kolmogorov filter applied to the finite series $[x_1, \dots, x_T]$ using an approach similar to the one in Burman (1980). The algorithm is explained in Kaiser and Maravall (1999).

I. Prior computations

a) Given λ (the HP-filter parameter), from (2.6) the system of equations

$$\begin{aligned} (1 + \theta_1^2 + \theta_2^2)V_b &= 1 + 6\lambda \\ \theta_1(1 + \theta_2)V_b &= -4\lambda \\ \theta_2V_b &= \lambda \end{aligned}$$

yields the parameters θ_1, θ_2 and V_b . [They can also be easily obtained from the spectral factorization of (2.6).] For $\lambda = 1600$, the solution is given by (2.7). Removing the subscript "HP", the WK filter to estimate m_t , given by (2.8), can be expressed as

$$\nu(B, F) = \frac{k_m}{\theta(B)\theta(F)} = k_m \left[\frac{G(B)}{\theta(B)} + \frac{G(F)}{\theta(F)} \right], \quad (A.1)$$

where $G(B) = g_0 + g_1B + g_2B^2$. Removing denominators in the above identity and equating the coefficients of the terms in B^0, B^1 and B^2 , yields a system of equations that can be solved for g_0, g_1 and g_2 . If

$$A = \begin{bmatrix} 1 & 0 & 0 \\ \theta_1 & 1 & 0 \\ \theta_2 & \theta_1 & 1 \end{bmatrix} + \begin{bmatrix} 0 & 0 & \theta_2 \\ 0 & \theta_2 & \theta_1 \\ \theta_2 & \theta_1 & 1 \end{bmatrix},$$

the solution is given by

$$[g_2 \ g_1 \ g_0]' = A^{-1}[0 \ 0 \ 1]'. \quad (\text{A.2})$$

(For $\lambda = 1600$, one obtains $g_0 = -44.954$, $g_1 = 11.141$, and $g_2 = 56.235$.) Compute the matrix

$$H = \begin{bmatrix} 1 & -2 & 1 & 0 \\ 0 & 1 & -2 & 1 \\ 1 & \theta_1 & \theta_2 & 0 \\ 0 & 1 & \theta_1 & \theta_2 \end{bmatrix}^{-1}.$$

Thus the expressions $\theta(B)$, k_m , $G(B)$ and H are obtained simply from the value of λ . They can be stored for further use, since they will be the same for all series.

b) Using an ARIMA model for the series x_t , extend the series with 4 backcasts and 4 forecasts. The extended series is given by

$$[x_{-3}, \dots, x_0, x_1, \dots, x_T, x_{T+1}, \dots, x_{T+4}].$$

When the ARIMA model is correctly specified, the HPA filter is obtained; when model (2.5) and (2.7) are used, the filter becomes the standard HP filter.

II. The Algorithm

Step I. For $t = 1, \dots, T + 2$, compute

$$y_t = g_0 x_t + g_1 x_{t+1} + g_2 x_{t+2},$$

$$[x_{T+1}^F, \dots, x_{T+4}^F]' = H[0, 0, y_{T+1}, y_{T+2}]',$$

and, for $t = T, \dots, 1$, obtain recursively

$$x_t^F = -\theta_1 x_{t+1}^F - \theta_2 x_{t+2}^F + y_t.$$

Step II. For $t = -1, 0, 1, \dots, T + 4$ compute

$$z_t = g_0 x_t + g_1 x_{t-1} + g_2 x_{t-2},$$

$$[x_0^B, x_{-1}^B, x_{-2}^B, x_{-3}^B]' = H[0, 0, z_0, z_{-1}]',$$

and, for $t = 1, \dots, T + 4$, obtain recursively

$$x_t^B = -\theta_1 x_{t-1}^B - \theta_2 x_{t-2}^B + z_t.$$

Step III. For $t = 1, \dots, T + 4$, obtain

$$\hat{m}_{t|T} = k_m [x_t^F + x_t^B].$$

This yields the MMSE estimator of the trend for the sample period $t = 1, \dots, T$, and the forecasting period $t = T + 1, \dots, T + 4$, equal to $E(m_t | x_1, \dots, x_T)$. The algorithm is fast and reliable, even for a series with (say) a million observations. It is remarkable that 4 forecasts and backcasts are enough to reproduce the full effect of the infinite filter.

References

1. Apel, M., Hansen, J., Lindberg, H. (1996) Potential Output and Output Gap. *Quarterly Review of the Bank of Sweden* 3: 24–35
2. Baxter, M., King, R.G. (1995) Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series. NBER Working Paper No.5022
3. Boldin, M.D. (1994) Dating Turning Points in the Business Cycle. *Journal of Business* 67: 97–131
4. Box, G.E.P., Jenkins, G.M. (1970) *Time Series Analysis: Forecasting and Control*. Holden-Day, San Francisco
5. Box, G.E.P., Hillmer, S.C., Tiao, G.C. (1978) Analysis and Modeling of Seasonal Time Series. In: Zellner, A. (ed.) *Seasonal Analysis of Economic Time Series*. U.S. Department of Commerce, Bureau of the Census, Washington, DC, pp. 309–334
6. Burman, J.P. (1980) Seasonal Adjustment by Signal Extraction. *Journal of the Royal Statistical Society A* 143: 321–337
7. Canova, F. (1998) Detrending and Business Cycle Facts. *Journal of Monetary Economics* 41: 475–512
8. Cleveland, W.P., Tiao, G.C. (1976), Decomposition of Seasonal Time Series: a Model for the X-11 Program. *Journal of the American Statistical Association* 71: 581–587
9. Cogley, T., Nason, J.M. (1995) Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research. *Journal of Economic Dynamics and Control* 19: 253–278
10. Dagum, E.B. (1980) The X11 ARIMA Seasonal Adjustment Method. Statistics Canada Catalogue 12-564E
11. Danthine, J.P., Girardin, M. (1989) Business Cycles in Switzerland: a Comparative Study. *Economic Review* 33: 31–50
12. EU Commission (1995) The Commission Services Method for Cyclical Adjustment of Government Budget Balances. *DG II / 401 / 95-EN*
13. Findley, D.F., Monsell, B.C., Bell, W.R., Otto, M.C. Chen, B.C. (1998) New Capabilities and Methods of the X12 ARIMA Seasonal Adjustment Program (with discussion). *Journal of Business and Economic Statistics* 16: 127–177
14. Fischer, B., Planas, C. (1998) Large Scale Fitting of ARIMA Models and Stylized Facts of Economic Time Series. Eurostat Working Paper n 9/1998/A/8
15. Granger, C.W.J., Newbold, P. (1974) Spurious Regressions in Econometrics. *Journal of Econometrics* 2: 111–120
16. Gómez, V. (1998) Butterworth Filters: a New Perspective. Working Paper D-98005, DGAPP, Ministerio de Economía y Hacienda, Madrid
17. Gómez, V. (1999) Three Equivalent Methods for Filtering Nonstationary Time Series. *Journal of Business and Economic Statistics* 17: 109–116
18. Gómez, V. and Maravall, A. (1998) Seasonal Adjustment and Signal Extraction in Economic Time Series. Working Paper 9809, Servicio de Estudios, Banco de España
19. Gómez, V., Maravall, A. (1996) Programs TRAMO and SEATS; Instructions for the User. Working paper 9628, Servicio de Estudios, Banco de España
20. Ghysels, E., Perron, P. (1993) The Effect of Seasonal Adjustment Filters on Tests for a Unit Root. *Journal of Econometrics* 55: 57–98
21. Hamilton, J.D. (1989) A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica* 57: 357–84
22. Harvey, A.C. (1989) *Forecasting Structural Time Series and the Kalman Filter*. Cambridge University Press, Cambridge
23. Harvey, A.C., Todd, P.H.J. (1983) Forecasting Economic Time Series with Structural and Box-Jenkins Models; A Case Study. *Journal of Business and Economic Statistics* 1: 299–306
24. Harvey, A.C., Jaeger, A. (1993) Detrending, Stylized Facts and the Business Cycle. *Journal of Applied Econometrics* 8, 231–247
25. Hillmer, S.C., Tiao, G.C. (1983) An ARIMA-Model Based Approach to Seasonal Adjustment. *Journal of the American Statistical Association* 77: 63–70
26. Jenkins, G.M. (1975) The Interaction between the Muskrat and the Mink Cycles in North Canada. Proceedings of the 8th International Biometric Conference, Editura Academiei Republicii So-

- cialiste Romania, pp. 55–71
27. Kaiser, R., Maravall, A. (1999) Short-Term and Long-Term Trends, Seasonal Adjustment, and the Business Cycle. Working Paper 99-10(2). Statistics and Econometrics Series, Universidad Carlos III de Madrid
 28. King, R.G., Rebelo, S.T. (1993) Low Frequency Filtering and Real Business Cycles. *Journal of Economics Dynamics and Control* 17: 207–233
 29. Koopman, S.J., Harvey, A.C., Doornik, J.A., Shephard, N. (1996) *Stamp: Structural Time Series Analyser, Modeller and Predictor* Chapman and Hall, London
 30. Maravall, A. (1995) Unobserved Components in Economic Time Series. In: Pesaran, H., Wickens, M. (eds.) *The Handbook of Applied Econometrics*, vol. 1. Basil Blackwell, Oxford
 31. Maravall, A. (1987) On Minimum Mean Squared Error Estimation of the Noise in Unobserved Component Models, *Journal of Business and Economic Statistics* 5: 115–120
 32. Prescott, E. (1986) Theory ahead of Business Cycle Measurement. *Carnegie-Rochester Conference Series on Public Policy* 25: 11–66