Journal of Forecasting, Vol. 12, 601–613 (1993)

Periodically Integrated Subset Autoregressions for Dutch Industrial Production and Money Stock

PHILIP HANS FRANSES
Erasmus University Rotterdam, The Netherlands

ABSTRACT

The univariate quarterly Dutch series of industrial production and money stock are both modelled with a periodically integrated subset autoregression (PISA). This model for a non-stationary series allows the lag orders, the values of the parameters and the cyclical patterns to vary over the seasons. The PISA models are found by applying a general-to-simple specification strategy, which deals with non-stationarity and periodicity simultaneously. It is found that the two series show a common asymmetric cyclical behaviour. This paper further proposes a test for periodicity in the errors, with which it is argued that a non-periodic model for the industrial production and money stock is misspecified and that seasonal adjustment does not remove periodicity in the autocorrelation function.

KEY WORDS Non-stationary seasonal time series Periodicity
Seasonal adjustment

INTRODUCTION

The industrial production index and the money stock are important macroeconomic variables since they can reflect business cycle behaviour and changing directions of an underlying trend. Forecasts of these variables are used by many decision makers. Given that decisions sometimes concern time intervals shorter than one year, forecasts for the monthly or quarterly observed industrial production index or money stock can be useful.

For several macroeconomic variables it applies that seasonality, trends, and cycles can be modelled separately, and hence that one may rely on forecasts for seasonally adjusted series. There are, however, variables whose time series show dynamic and cyclical patterns that vary over the seasons. This may be caused by economic agents who have preferences, technologies, constraints, and expectations which are not constant over the seasons. Recently, the construction of economic models in which seasonally varying parameters are allowed has gained some attention (see, for example, Braun and Evans, 1991; Osborn, 1988; Hansen and Sargent, 1993). For an empirical study in which it is found that many consumer confidence indices or expectations in European countries show seasonality, see Franses (1992a).

0277–6693/93/070601–13$11.50

Received August 1992
Revised March 1993

In periodic dynamic patterns it is difficult to separate trend and cycles from seasonality. Alternatively, linear seasonal adjustment filters are not likely to remove the intrinsic seasonality, and hence the 'adjusted' series still shows seasonality. As part of an analysis of a univariate time series, which can include decisions on whether to seasonally adjust or not, it may therefore be sensible to investigate seasonality and trend aspects in an unadjusted series. The present paper illustrates a strategy that can be useful for this investigation. The main feature of this strategy is that it deals with non-stationarity and seasonality simultaneously.

The illustrations of the approach are given by the quarterly Dutch series for money stock and industrial production. It is demonstrated that both series are variables that can be described by four different equations for the seasons. In particular, the two adequate models turn out to be periodically integrated subset autoregressions (PISA). These PISA models should be considered merely as adequate statistical representations of the underlying data-generating process rather than as behavioural equations for economic agents. The models, which will be estimated in the fourth section, allow the dynamics as well as the parameters to vary over the seasons. Before their estimation, the results of standard preliminary univariate time-series analysis are discussed in the next section. An application of these conventional tools shows that autoregressive models for the fourth-order differenced series seem to be adequate. However, the residuals show patterns that can be explained by the presence of dynamic periodicity, which is detected by a formal test for periodicity proposed in the third section. Alternatively, the non-periodic models are misspecified and periodic models would have been more appropriate. In periodicity in a time series, in the sense that the dynamics are different throughout the year, one can easily recognize that an application of linear seasonal adjustment filters does not affect this type of periodicity. This is because these filters that the observations within one year in the same way. An empirical illustration of this phenomenon is given in the third section by showing that the residuals of autoregressions for the seasonally adjusted variables still display periodicity.

Given the outcomes of tests for periodicity in the residuals, one may now want to specify a periodic model for the fourth-order differenced series. A drawback of this procedure is that the choice for the fourth-order filter is likely to be influenced by the initially neglected periodicity (see, for example, Franses, 1992b). Therefore, a more appropriate strategy would be to specify and estimate a general periodic model, and, by means of a sequence of hypothesis tests which involve tests for periodicity as well as for non-stationarity, to select a simplified model. In the fourth section this strategy is pursued along the lines of Boswijk and Franses (1992), and a final adequate model appears to be a PISA. Given that four different models are necessary to describe the variable, it seems worth studying the effects on forecasts, on cyclical behaviour, and on impulse response patterns. The fifth section reports that shocks in the second quarter have a larger effect for both series than those in other quarters. This means that a PISA model can describe asymmetric cyclical patterns and that, in the case of industrial production and money stock, these patterns are similar. One-step-ahead forecasts from the PISA and the rival non-periodic model indicate that an improvement in forecasting can be observed. This confirms the theoretical and empirical findings in Osborn (1991) and Osborn and Smith (1989). The final section presents conclusions.

PRELIMINARY ANALYSIS

A graph of the log of the industrial production index in the Netherlands for the period 1960.1–1989.4 is displayed in Figure 1. A similar graph for the log of the money (M1) stock for 1959.1–1988.4 is depicted in Figure 2. From these figures it can be observed that both series show non-stable seasonal patterns as well as trending behaviour. The first three years will be used as starting values to ensure that all forthcoming models are estimated with the same number of observations, and the last six years are used to evaluate the forecasting performance of some of these models. The quarterly series will be denoted by y_t and m_t, where t runs from

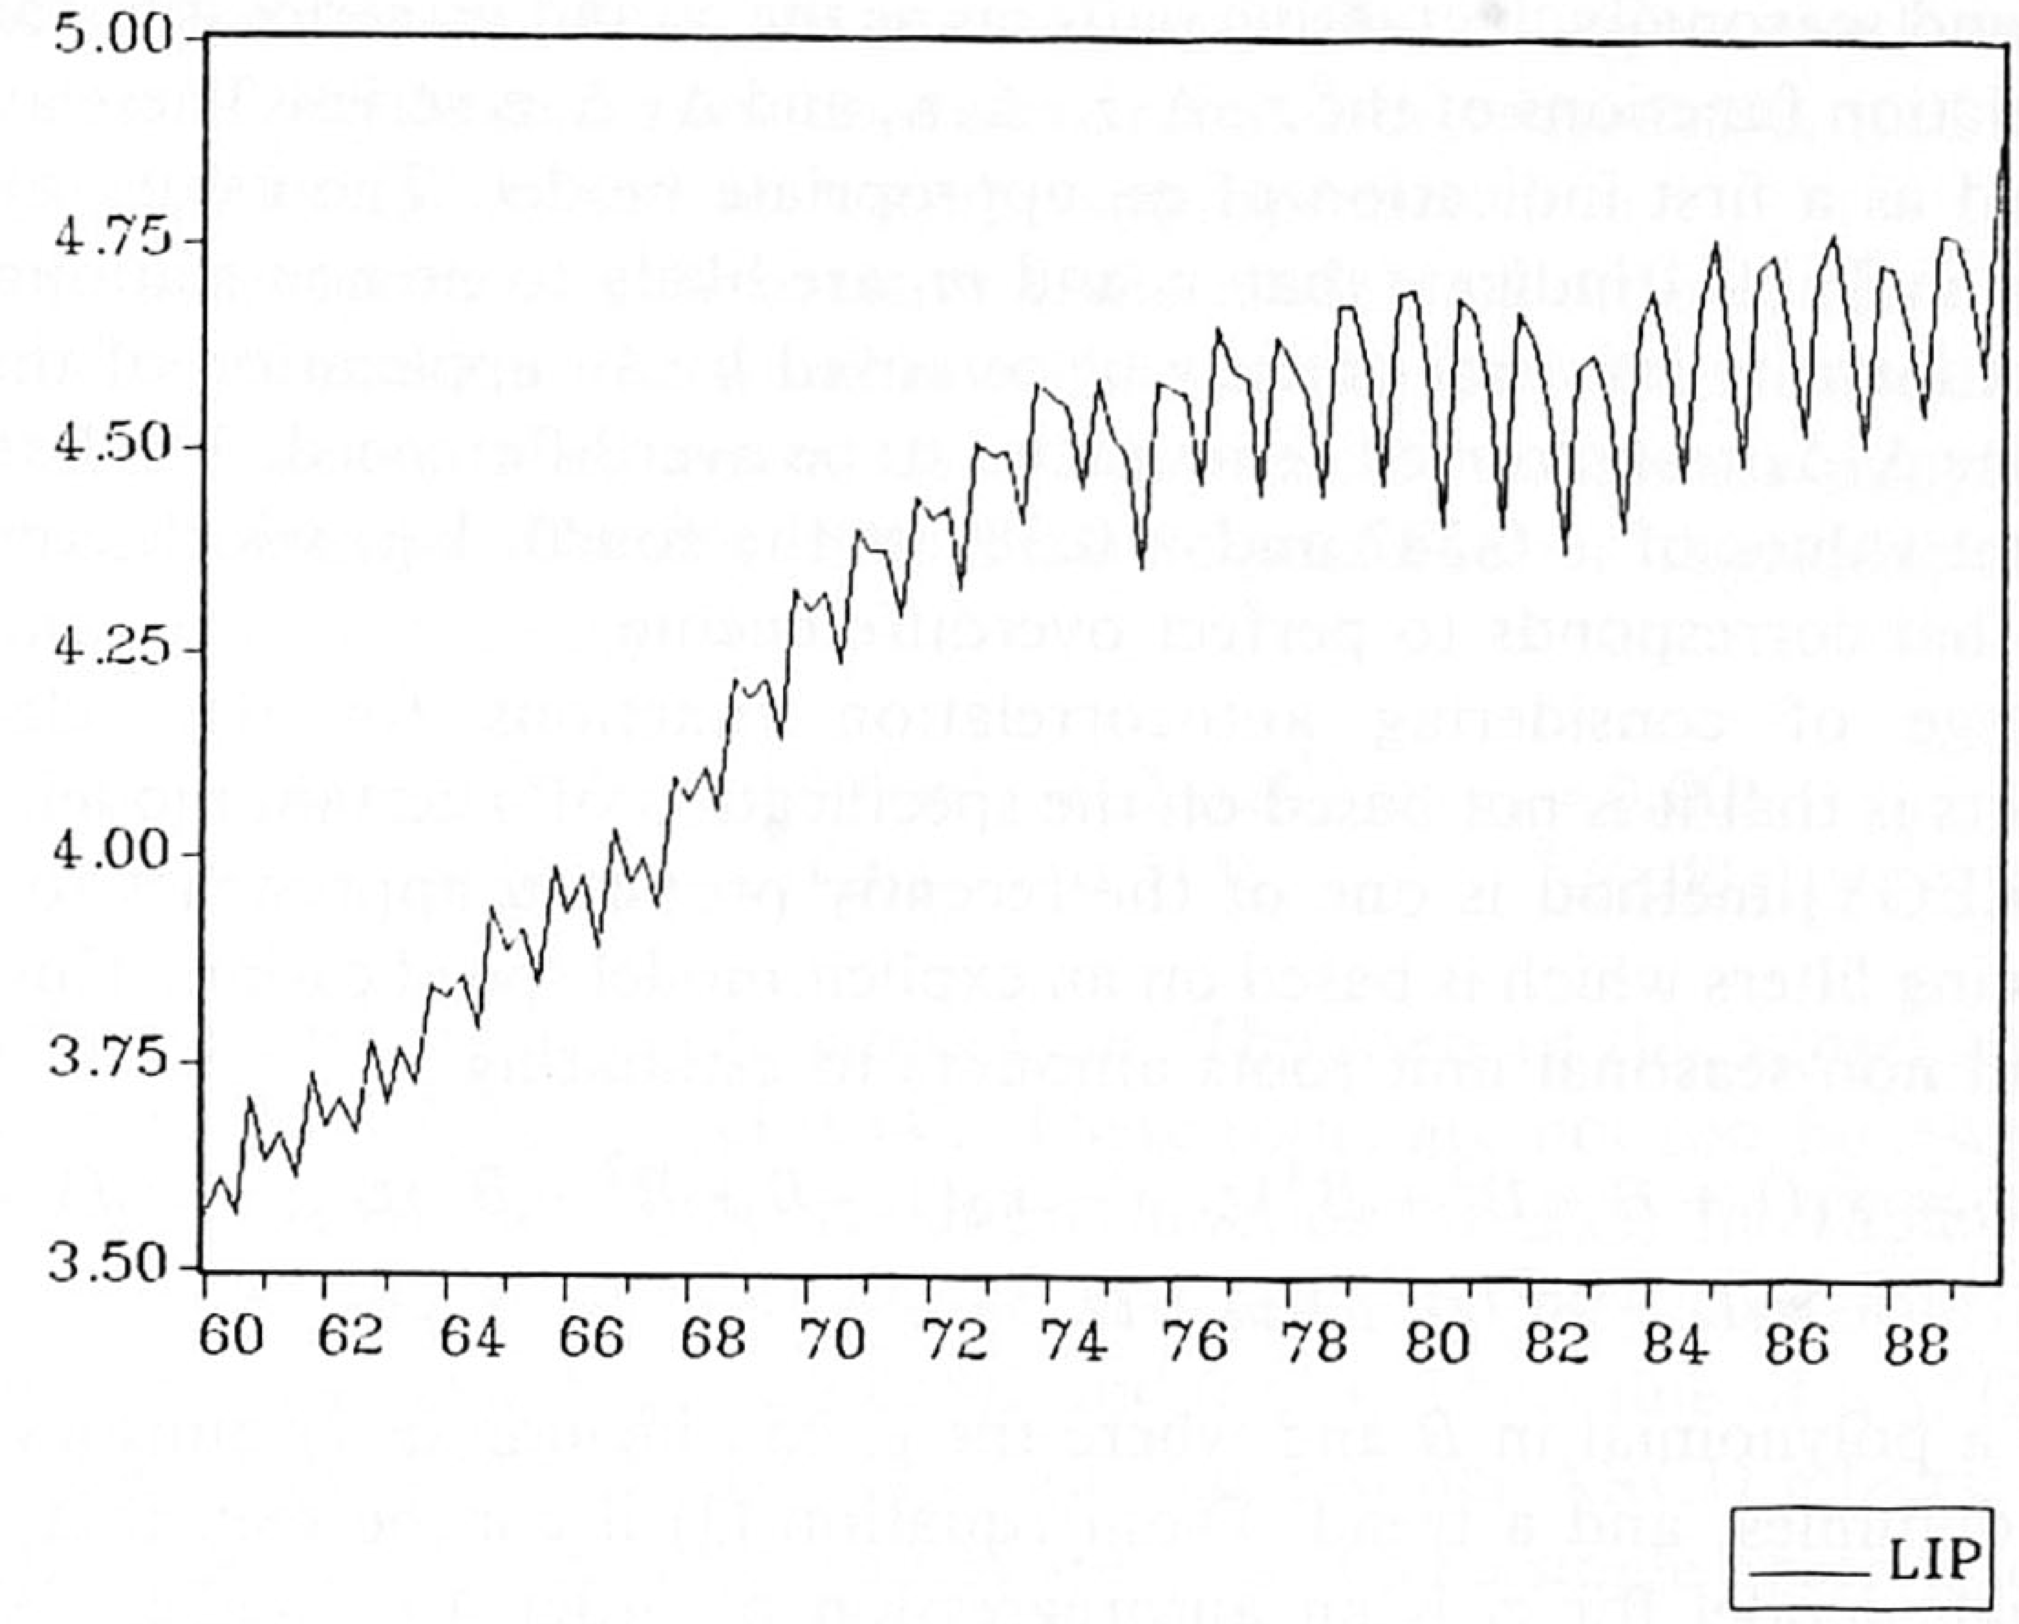

Figure 1. The log of the industrial production index in the Netherlands, 1960.1 to 1989.4

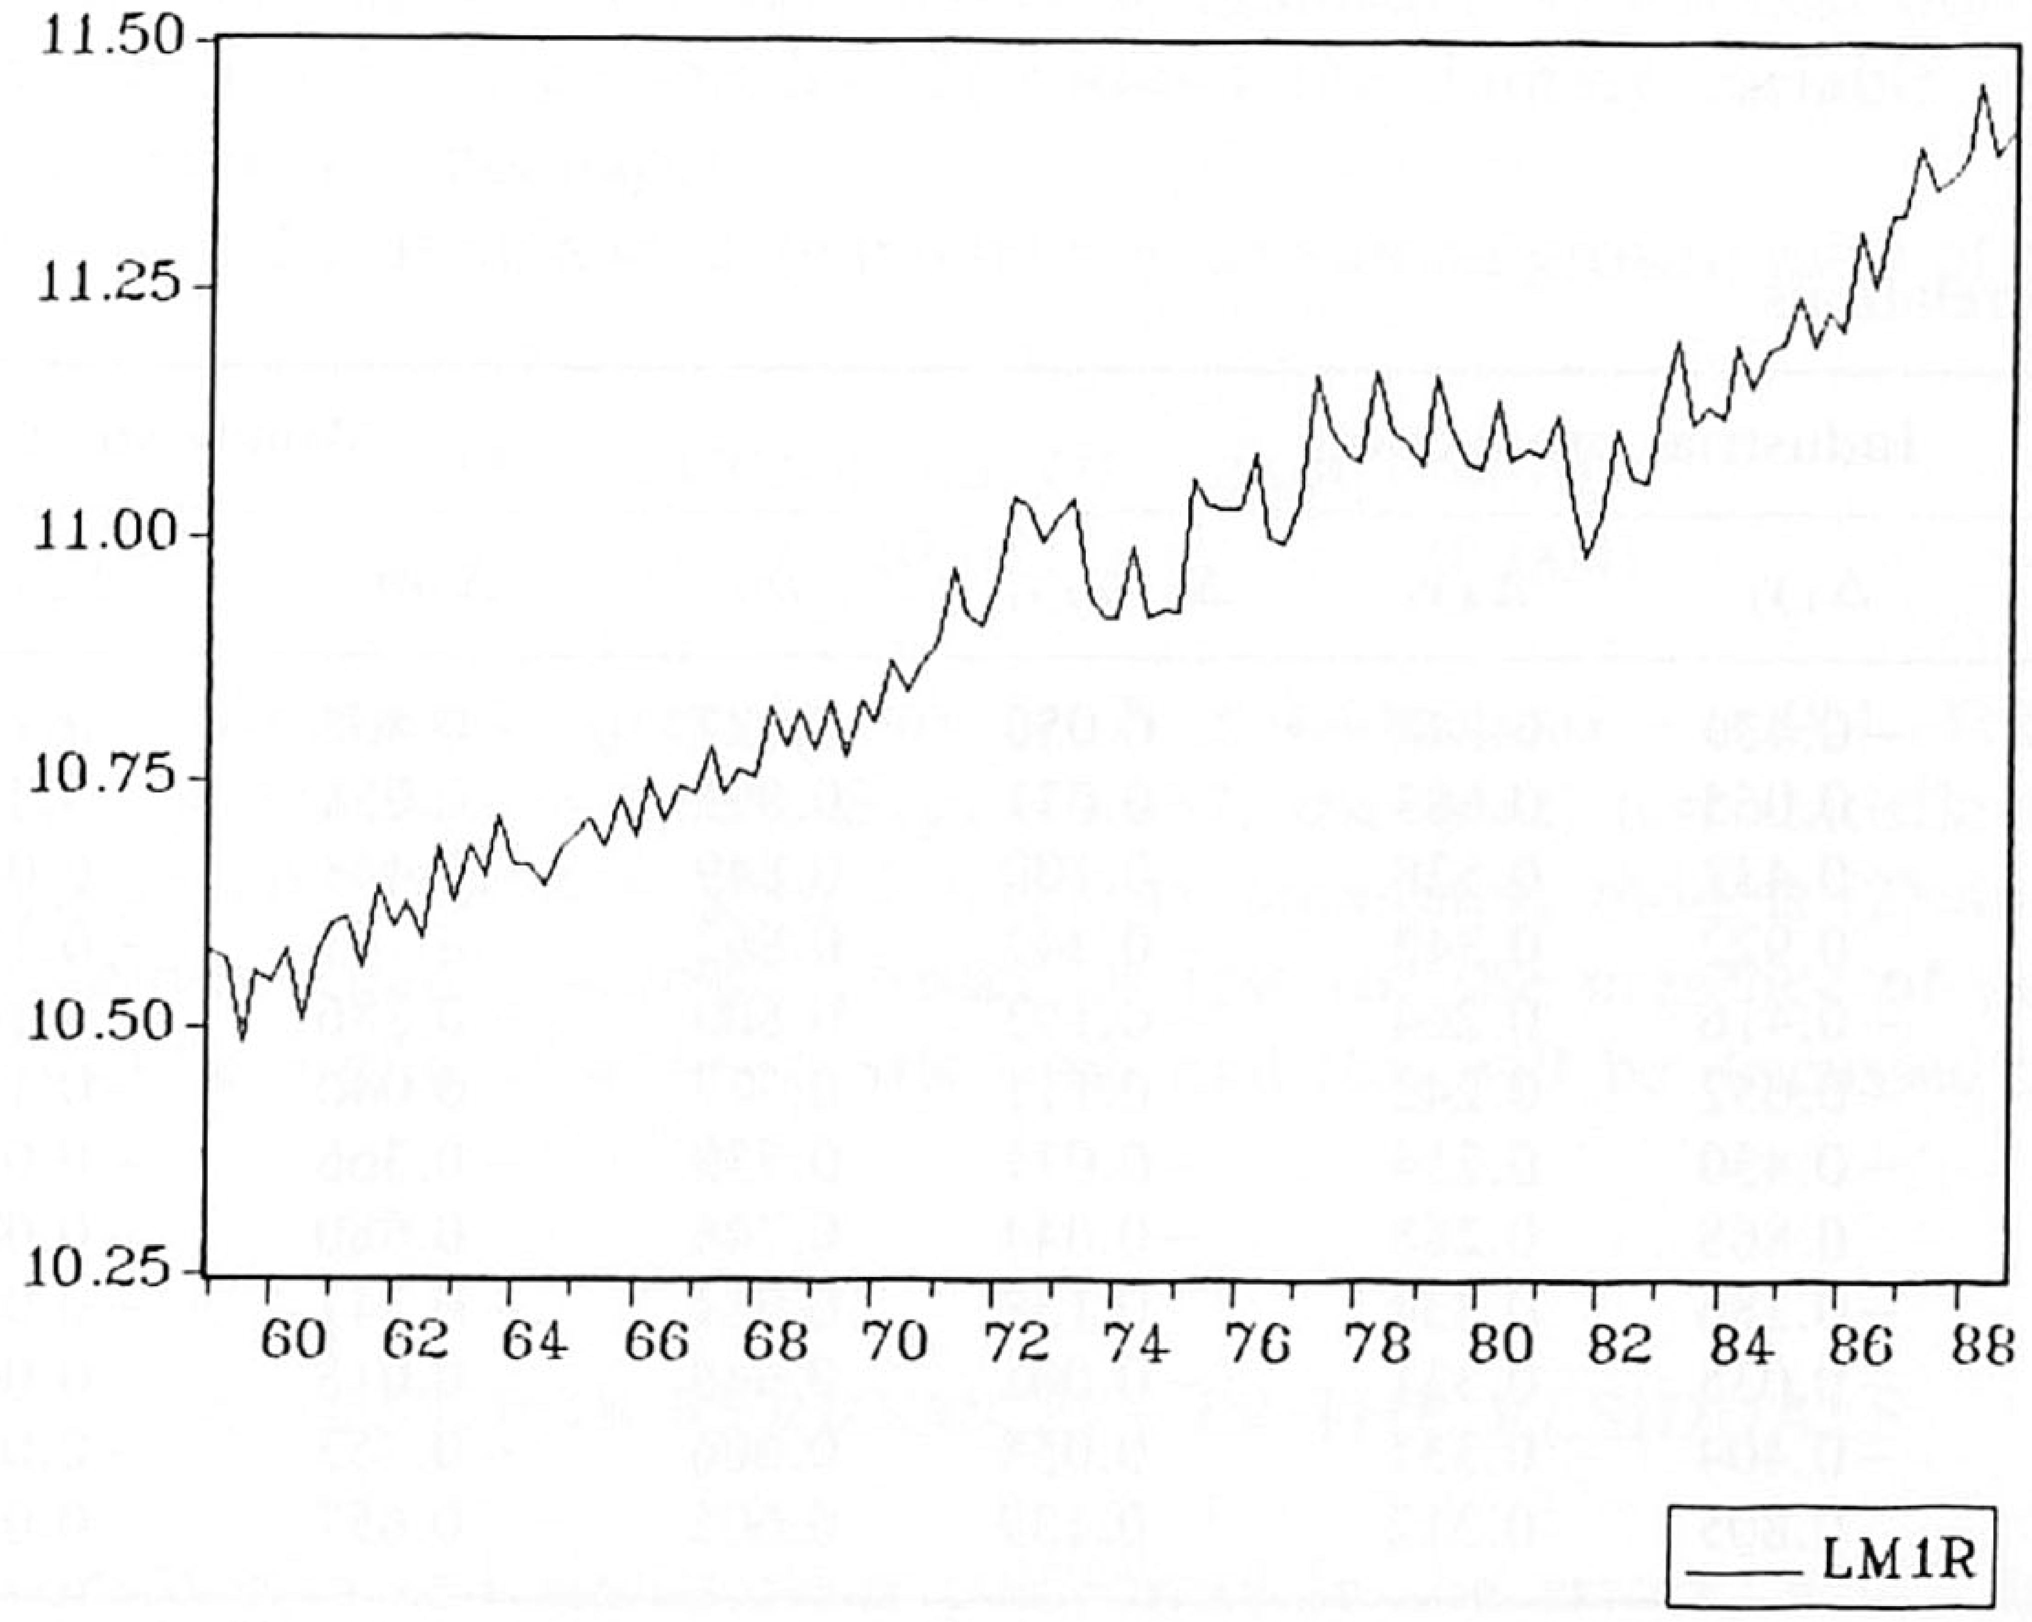

Figure 2. The log of the money stock in the Netherlands, 1959.1 to 1988.4

$1, \ldots, n$, while the annual series containing the observations in the four quarters will be denoted as $Y_{s,T}$ and $M_{s,T}$ where s can be 1, 2, 3, and 4, and where T runs from 1 to $N = n/4$. In the present case, the N and n are equal to 21 and 84, respectively. The vector of $Z_{s,T}$ series will be denoted as $Z_T = (Z_{1,T}, Z_{2,T}, Z_{3,T}, Z_{4,T})'$, where Z_T is Y_T or M_T. The backward shift operator B and the differencing filter Δ are defined by $\Delta_k z_t = (1 - B^k) z_t = z_t - z_{t-k}$, where z_t can be y_t or m_t.

A usual first step in a conventional approach to modelling a time series that shows similar non-stationary and seasonally fluctuating patterns as the y_t and m_t series is to take a close look at the autocorrelation functions of the z_t, $\Delta_1 z_t$, $\Delta_4 z_t$, and $\Delta_1 \Delta_4 z_t$ series. These autocorrelations can then be used as a first indication of an appropriate model. The values of the estimated autocorrelations in Table I indicate that y_t and m_t are likely to be non-stationary time series, that this non-stationarity may not entirely be removed by an application of the Δ_1 or the Δ_4 filter, but that the $\Delta_1\Delta_4$ transformed series seems to be overdifferenced. The latter is suggested by the significant values of -0.347 and -0.389 at the fourth lag, which come close to the value of -0.5 that corresponds to perfect overdifferencing.

A disadvantage of considering autocorrelation functions for the selection between differencing filters is that it is not based on the specification of a certain model. The Hylleberg *et al.* (1990) [HEGY] method is one of the recently proposed approaches to select between several differencing filters which is based on an explicit model specification. This test procedure for seasonal and non-seasonal unit roots amounts to estimating

$$\phi_h(B)\,\Delta_4 z_t = \pi_1(1 + B + B^2 + B^3) z_{t-1} - \pi_2(1 - B + B^2 - B^3) z_{t-1} - \pi_3(1 - B^2) z_{t-2} - \pi_4(1 - B^2) z_{t-1} + \mu_t + \varepsilon_t \tag{1}$$

where $\phi_h(B)$ is a polynomial in B and where the μ_t can include deterministics as a constant, three seasonal dummies, and a trend. From equation (1) it can be seen that, in general, an assumed adequate model for z_t is an autoregression of order 4 or higher. When all $\hat{\pi}_i$ are insignificant one can apply the Δ_4 filter to obtain stationarity for the z_t series, and when only $\hat{\pi}_1$ is insignificant, one can use the Δ_1 filter for the same reason. For the industrial production

Table I. Autocorrelations

	Industrial production				Money stock			
Lag	y_t	$\Delta_1 y_t$	$\Delta_4 y_t$	$\Delta_1 \Delta_4 y_t$	m_t	$\Delta_1 m_t$	$\Delta_4 m_t$	$\Delta_1 \Delta_4 m_t$
1	0.881	−0.439	0.848	0.050	0.927	−0.405	0.742	0.181
2	0.851	−0.065	0.683	−0.031	0.895	0.058	0.399	0.049
3	0.819	−0.432	0.528	0.109	0.849	−0.463	0.040	−0.284
4	0.863	0.922	0.342	−0.347	0.862	0.743	−0.161	−0.389
5	0.755	−0.416	0.264	−0.193	0.800	−0.386	−0.141	−0.086
6	0.726	−0.052	0.242	−0.111	0.777	0.040	−0.102	−0.040
7	0.692	−0.430	0.254	−0.071	0.739	−0.366	−0.051	0.160
8	0.734	0.868	0.288	−0.044	0.746	0.680	−0.085	−0.072
9	0.628	−0.386	0.334	0.158	0.681	−0.347	−0.083	0.025
10	0.597	−0.035	0.331	−0.001	0.648	0.018	−0.068	−0.060
11	0.559	−0.404	0.331	0.058	0.606	−0.353	−0.026	0.063
12	0.592	0.805	0.312	0.139	0.605	0.657	−0.017	0.040

The estimated standard error for the estimated autocorrelations is 0.109.

series y_t, it emerges that $\phi_h(B)$ is specified as $(1 - \phi_1 B - \phi_4 B^4 - \phi_5 B^5)$, that a constant, seasonal dummies, and trend are included in equation (1), and that the t ratios for the $\hat{\pi}_i$ are -0.855, -1.491, -1.893, and -0.479, respectively. For the money stock series it is found that $\phi_h(B)$ should be specified as $(1 - \phi_1 B - \phi_3 B^3 - \phi_4 B^4)$, that a constant, seasonal dummies, and trend are included, and that the t ratios for the $\hat{\pi}_i$ are -1.311, -3.640, -2.712, and -1.658, respectively. Comparing the obtained t values with the 5% fractiles in Hylleberg *et al.* (1990) indicates that the π_i are all estimated to be insignificantly different from zero, except for the π_2 in the auxiliary regression for m_t. In practice, one would neglect the latter outcome, also since the $(1 - B)(1 + B^2)$ filter for m_t does not seem to be reasonable, and hence one would decide that the Δ_4 filter for y_t and m_t may be appropriate.

The next step is to estimate autoregressive (AR) models of order h for the $\Delta_4 y_t$ and $\Delta_4 m_t$ series. This order h is equal to the number of lags which have been necessary to whiten the residuals of the regression in equation (1). These initial models are simplified by deleting insignificant lagged variables. The final model should pass a set of diagnostic checks. The result obtained for industrial production is

$$\underset{}{(1 - \underset{(0.067)}{0.969}B + \underset{(0.120)}{0.409}B^4 - \underset{(0.113)}{0.281}B^5)\, \Delta_4 y_t = \underset{(0.004)}{0.006}} \tag{2}$$

where the standard errors are given in parentheses. The roots of this subset autoregression are $-0.531 \pm 0.551i$, $0.594 \pm 0.466i$, and 0.842. These roots are not too far away from the unit circle, and this may cause the autocorrelation function of $\Delta_4 y_t$ in Table I not to die out quickly. The values of the F versions of the LM tests for first- and fourth-order residual autocorrelation, i.e. F_{AR1} and F_{AR4}, are 0.001 and 0.613. The value of a $\chi^2(2)$ normality test is 2.782, and those of the LM tests for first- and fourth-order ARCH effects, i.e. F_{ARCH1} and F_{ARCH4}, are 4.054 and 2.353. The latter values are caused by a single observation. Deleting this observation in 1975.4 by including a dummy variable in equation (2) does not significantly affect the parameter estimates and the diagnostic test results for autocorrelation and normality, while the F_{ARCH1} and F_{ARCH4} statistics become insignificant. Given that normality could not be rejected in the first instance, model (2) without the dummy variable is assumed to be adequate in the sequel of this paper.

Similarly, it is found that an adequate model for the annual growth rates of quarterly money stock is

$$(1 - \underset{(0.075)}{0.870}B + \underset{(0.075)}{0.317}B^3)\, \Delta_4 m_t = \underset{(0.004)}{0.010} \tag{3}$$

The roots of this subset autoregression are $0.677 \pm 0.444i$ and -0.484. The application of diagnostic checks yields $F_{\text{AR1}} = 0.666$, $F_{\text{AR4}} = 1.637$, the $\chi^2(2)$ test statistic for normality is 1.035, and $F_{\text{ARCH1}} = 0.103$ and $F_{\text{ARCH4}} = 0.301$. In summary, models (2) and (3) cannot be rejected using conventional diagnostic checks. A test for the presence of periodicity in the residuals can also be useful as a diagnostic tool, and this will be discussed in the following section.

A TEST FOR PERIODICITY IN THE RESIDUALS

This section considers a test statistic for periodicity in the errors, which is related to the autocorrelations of the process and not to the error variances. A test for the latter type of periodicity is given simply by a regression of the squared residuals on a constant and the seasonal dummies D_{qt}, $q = 1, 2, 3$. Similarly, one can construct a test statistic for periodically varying autoregressive conditional heteroscedasticity. An extensive treatment of testing for periodicity in residuals is given in Lütkepohl (1991) and Ghysels and Hall (1992).

Consider the case in which one assumes the appropriateness of a constant-parameter model, while the data-generating process has periodic autoregressive parameters. For example, one fits

$$x_t = \alpha x_{t-1} + \varepsilon_t \tag{4}$$

while

$$x_t = \alpha_s x_{t-1} + v_t \tag{5}$$

is the data-generating process, where α_s denotes a parameter with values that vary with the season s, and where v_t is assumed to be white noise. When $\alpha_1\alpha_2\alpha_3\alpha_4$ is sufficiently smaller than unity, the autocorrelation function of model (5) looks like that of a non-periodic AR(1) model, and hence the case of fitting model (4) to a time series generated by model (5) may not be an unusual occurrence in practice. However, when $\alpha_1\alpha_2\alpha_3\alpha_4$ approaches unity, one may need many more lags to whiten the residuals (see Osborn and Smith, 1989). Theoretically, the residuals ε_t are not white noise, since

$$\varepsilon_t = (\alpha - \alpha_s)x_{t-1} + v_t \tag{6}$$

An LM test applied to model (4) for first-order residual autocorrelation may yield the impression that the model is not misspecified. However, an LM test for first-order periodic autoregressive patterns (PAR1) in the residuals constructed from the auxiliary regression

$$\hat{\varepsilon}_t = \gamma x_{t-1} + \sum_{s=1}^{4} \beta_s D_{st}\hat{\varepsilon}_{t-1} + \xi_t \tag{7}$$

as an $F_{\mathrm{PAR1}}(4, n-4-1)$ test will indicate that a periodic model would have been more appropriate.

An application of such an LM test to the residuals of the model in equation (2) results in an $F_{\mathrm{PAR1}}(4, 75)$ test statistic value of 4.173, which is significant at a 5% level. Comparing this value to that of the F_{AR1} statistic value of 0.001 seems to confirm the impression that a periodic model would have been more appropriate for $\Delta_4 y_t$ or y_t. The $F_{\mathrm{PAR1}}(4, 76)$ test applied to model (3) obtains a value of 3.498, which is also significant at a 5% level. Hence, a similar conclusion can be drawn for the money stock.

These results suggest that it may be appropriate to start an analysis of seasonal time series with estimating periodic models instead of non-periodic models. If the null hypothesis of non-periodicity cannot be rejected one may then proceed along the lines of Hylleberg *et al.* (1990) to decide on the non-stationary aspects of a time series. The HEGY method has been extended to monthly time series in Franses (1991) and Beaulieu and Miron (1993). These authors find that many monthly seasonal time series only need a Δ_1 filter to remove non-stationary patterns. However, in Franses (1992b) it is shown that an application of the HEGY method to periodic autoregressive time series can yield the inappropriate impression that such a Δ_1 filter is adequate.

Before turning to the specification of periodic autoregressive models, it may be interesting to see whether the periodicity is also not removed by the application of a linear seasonal adjustment filter. The application of a linearization of the Census-X11 filter (see Laroque, 1977) uses a 57-period moving average filter. This implies that only seasonally adjusted series are constructed for 1967.1 to 1982.4 for y_t and for 1966.1 to 1981.4 for m_t. These adjusted time series will be denoted as ya_t and ma_t. One could have used forecasts for y_t and m_t, generated from models (2) and (3), to construct longer seasonally adjusted series. This is not, however, done here since the F_{PAR1} tests have indicated that these models may not be appropriate. An adequate model for ya_t is

$$\underset{(0.012)}{ya_t = 0.942\,ya_{t-1}} + \underset{(0.053)}{0.267} \tag{8}$$

for which $F_{\text{AR1}} = 0.222$ and $F_{\text{AR4}} = 0.361$, but $F_{\text{PAR1}}(4, 56) = 2.847$. The latter test statistic is significant at the 5% level. Alternatively, the periodicity in y_t is not removed by a linear filter. For the adjusted money stock series an adequate model appears to be

$$\Delta_1 ma_t = \underset{(0.127)}{0.309}\ \Delta_1 ma_{t-1} \tag{9}$$

for which $F_{\text{AR1}} = 0.019$ and $F_{\text{AR4}} = 0.986$, while $F_{\text{PAR1}}(4, 55)$ obtains a value of 1.848. This F_{PAR1} value is significant only at a 14% significance level.

The test statistic calculated from model (7) can be used to have an indication of the kind of misspecification once one has estimated a non-periodic model. Of course, the auxiliary regression in model (7) can be enlarged by including variables such as $D_{st}x_{t-1}$. This regression comes close to the estimation of a fully specified periodic model. An alternative and preferable strategy may then be to start a univariate analysis by specifying a general periodic model, and to select a simplified model using tests for parameter restrictions (see Franses, 1992c). An easily applicable strategy is proposed in Boswijk and Franses (1992), and in the next section it will be applied to the macroeconomic time series under consideration.

PERIODICALLY INTEGRATED SUBSET AUTOREGRESSION (PISA)

A general expression for a periodic autoregression of order p is

$$y_t = \sum_{i=1}^{p} \sum_{s=1}^{4} \phi_{is} D_{st} y_{t-i} + \sum_{s=1}^{4} D_{st}(\mu_s + \tau_s T_t) + \varepsilon_t \tag{10}$$

where ϕ_{is}, μ_s, and τ_s are periodically varying parameters, of which μ_s and τ_s refer to a seasonally varying mean and trend T_t. The lag order is not necessarily p for all seasons, i.e. the orders may be p_s in season s, and p is then equal to $\max(p_s)$. Further, not all ϕ_{is} or μ_s, τ_s have to be unequal to zero. Finally, note that y_t is untransformed.

The initially specified model for the industrial production y_t appears to be of order 5, while τ_s as well as ϕ_{2s} and ϕ_{3s} can be set equal to zero for all s. This model is found by estimating a model such as (10) with p such that several diagnostic checks do not indicate misspecification, and by deleting insignificant variables. A test for seasonal heteroscedasticity also does not suggest modification of model (10). A first step is now to test for periodicity, or testing the validity of the hypothesis that $\phi_{js} = \phi_j$ for j is 1, 4, and 5. This joint hypothesis can be tested using an F-test which asymptotically follows a standard distribution since the parameter restrictions do not affect the number of unit roots in the system of Y_T (see Franses, 1992b). Here, the $F(9, 68)$ test statistic obtains a value of 4.464, and the null of no periodicity can be rejected. The estimation results for the restricted version of model (10) indicate that some of the parameters ϕ_{4s} and ϕ_{5s} are equal to zero. The model can then be simplified to

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ -\phi_{12} & 1 & 0 & 0 \\ 0 & -\phi_{13} & 1 & 0 \\ 0 & 0 & -\phi_{14} & 1 \end{bmatrix} \begin{bmatrix} Y_{1,T} \\ Y_{2,T} \\ Y_{3,T} \\ Y_{4,T} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & \phi_{11} \\ \phi_{52} & \phi_{42} & 0 & 0 \\ 0 & \phi_{53} & 0 & 0 \\ 0 & 0 & 0 & \phi_{44} \end{bmatrix} \begin{bmatrix} Y_{1,T-1} \\ Y_{2,T-1} \\ Y_{3,T-1} \\ Y_{4,T-1} \end{bmatrix} + \mu + \varepsilon_T \quad (11)$$

and the parameters are estimated as

$$\begin{array}{ll} \text{quarter 1:} & y_t = \underset{(0.015)}{1.043} y_{t-1} - \underset{(0.066)}{0.216} \\ \text{quarter 2:} & y_t = \underset{(0.136)}{0.853} y_{t-1} + \underset{(0.199)}{0.744} y_{t-4} - \underset{(0.147)}{0.646} y_{t-5} + \underset{(0.100)}{0.212} \\ \text{quarter 3:} & y_t = \underset{(0.124)}{1.321} y_{t-1} - \underset{(0.113)}{0.405} y_{t-5} + \underset{(0.085)}{0.241} \\ \text{quarter 4:} & y_t = \underset{(0.096)}{0.573} y_{t-1} + \underset{(0.078)}{0.451} y_{t-4} + \underset{(0.099)}{0.023} \end{array} \quad (12)$$

It can be seen from model (12) that a large proportion of the variation in y_t can be explained by the seasonally varying autoregressive polynomials. Further, it is clear that the models are indeed quite different for the various seasons. This may explain the results obtained when applying the descriptive techniques discussed in Barsky and Miron (1989). A regression of $\Delta_1 y_t$ on four seasonal dummies gives a coefficient of determination with a value of 0.876. However, (unreported) graphs of the recursive estimates of the parameters for these dummies show that they are not constant over time. Also, seasonality seems to change over time for this y_t series, a phenomenon that can be modelled with a model like (12).

A formal test for non-stationarity, or in this case of a periodic model for periodic integration, is developed in Boswijk and Franses (1992). This reduces to checking whether for $z = 1$ the following equality holds:

$$\begin{vmatrix} 1 & 0 & 0 & -\phi_{11}z \\ -\phi_{12} - \phi_{52}z & 1 - \phi_{42}z & 0 & 0 \\ 0 & -\phi_{13} - \phi_{53}z & 1 & 0 \\ 0 & 0 & -\phi_{14} & 1 - \phi_{44}z \end{vmatrix} = 0 \quad (13)$$

Some rewriting allows the null hypothesis of periodic integration to be written as a non-linear restriction on the parameters, or

$$\phi_{11} = \frac{(1 - \phi_{42})(1 - \phi_{44})}{\phi_{14}(\phi_{12} + \phi_{52})(\phi_{13} + \phi_{53})} \quad (14)$$

This restriction is imposed on model (12), and the parameters are now estimated with the non-linear least squares routine in MicroTSP (version 7). The one-sided likelihood ratio test statistic LR_τ, which asymptotic distribution is the same as that of the usual Dickey–Fuller τ_μ statistic here, obtains a value of -2.305. The null hypothesis of periodic integration thus cannot be rejected. A full account of this LR_τ statistic is given in Boswijk and Franses (1992). In case one wants to test for specific parameter restrictions while assuming that periodic integration is appropriate, one adds these to model (14). One interesting example is given by checking whether the $(1 - B)$ filter can be applied in each equation of model (12) given the validity of periodic integration. The parameter restrictions for this hypothesis are $\phi_{11} = 1$, $\phi_{12} + \phi_{42} + \phi_{52} = 1$, $\phi_{13} + \phi_{53} = 1$, and $\phi_4 + \phi_{44} = 1$. Since the number of unit roots in the system for Y_T remains one under this hypothesis, the $F(3, 73)$ test value of 8.882 can be compared to the critical values of a standard F-distribution.

In summary, an adequate univariate model for industrial production is a periodically integrated subset autoregression. For each of the models in (12) one can calculate the characteristic roots of the autoregressive polynomial and the corresponding cycle length. Assuming the validity of the restriction in model (14), these results are summarized in Table II. It can be seen that the cyclical movements in y_t are indeed varying over the seasons. The model in (11), i.e. $A_0 Y_T = A_1 Y_{T-1} + \mu + \varepsilon_T$, can be rewritten as a first-order vector autoregressive model, or $Y_T = A_0^{-1} A_1 Y_{T-1} + \mu^* + \varepsilon_T^*$. For the production series the eigenvalues of $A_0^{-1} A_1$ are 0, 0.234 and $0.817 \pm 0.047i$.

Similarly, an initially adequate periodic autoregressive model for m_t appears to be of order 4. An $F(12, 64)$ test for the equality of the parameters over the seasons yields a value of 7.218. Some of the estimated parameters appear to be insignificant, and the model can be simplified to

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ -\phi_{12} & 1 & 0 & 0 \\ -\phi_{23} & -\phi_{13} & 1 & 0 \\ 0 & 0 & -\phi_{14} & 1 \end{bmatrix} \begin{bmatrix} M_{1,T} \\ M_{2,T} \\ M_{3,T} \\ M_{4,T} \end{bmatrix} = \begin{bmatrix} 0 & 0 & \phi_{21} & 0 \\ 0 & \phi_{42} & \phi_{32} & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} M_{1,T-1} \\ M_{2,T-1} \\ M_{3,T-1} \\ M_{4,T-1} \end{bmatrix} + \mu + v_T \quad (15)$$

and the parameters are estimated as

$$\begin{aligned} &\text{quarter 1:} \quad m_t = \underset{(0.024)}{0.893} m_{t-2} + \underset{(0.265)}{1.171} \\ &\text{quarter 2:} \quad m_t = \underset{(0.231)}{1.409} m_{t-1} - \underset{(0.325)}{1.146} m_{t-3} + \underset{(0.227)}{0.828} m_{t-4} - \underset{(0.391)}{0.975} \\ &\text{quarter 3:} \quad m_t = \underset{(0.154)}{1.346} m_{t-1} - \underset{(0.172)}{0.463} m_{t-2} + \underset{(0.338)}{1.198} \\ &\text{quarter 4:} \quad m_t = \underset{(0.026)}{0.815} m_{t-1} + \underset{(0.281)}{2.056} \end{aligned} \quad (16)$$

A regression of $\Delta_1 m_t$ on four seasonal dummies has an R^2 of 0.644. However, the graphs of the recursive estimates for the corresponding parameters indicates that these estimates are not constant. Furthermore, these graphs intersect, which implies that seasonality seems to change

Table II. The roots of the autoregressive polynomials in the periodically integrated subset autoregressions

	Industrial production		Money stock	
Quarter	Roots	Cycle length	Roots	Cycle length
1	1.040		0.925	
2	$0.001 \pm 0.991i$, -0.991, 1.000, 0.875	4.0	$0.645 \pm 0.628i$, -0.936, 1.118	8.1
3	$1.083 \pm 0.151i$, $-0.083 \pm 0.709i$, -0.673	45.4	$0.675 \pm 0.090i$	47.5
4	$0.137 \pm 0.779i$, 1.012, -0.702	4.5	0.815	

significantly throughout the sample. The LR_τ test statistic calculated along the lines of models (13) and (14) for periodic integration obtains a value of -1.737. Given the parameter estimates and lag lengths in model (16), a test for the appropriateness of the $(1 - B)$ filter does not seem to be relevant. Hence, an appropriate model for m_t is also a periodically integrated subset autoregression. The roots of the various polynomials under the periodic integration restriction, and their corresponding cycles, are displayed in the second part of Table II. Again, one can observe asymmetric cyclical patterns. The eigenvalues of the $A_0^{-1}A_1$ matrix, where the parameter estimates from model (16) have been substituted, are 0, 0, -0.367, and 0.933.

CYCLES, SHOCKS AND FORECASTS

Models (12) and (16) can now be used for forecasting and evaluating the effects of exogenous shocks. From Table II it can be seen that the cyclical behaviour of the y_t and m_t series varies over the seasons, although the patterns of the first, second, and fourth quarter are similar. The third quarter, however, shows a much larger cycle, i.e. of about 11 years, than those in the other quarters. One of the implications of these varying cycle lengths is that for these macroeconomic variables the turning points can occur more frequently in some seasons than in others (see Ghysels, 1991, for related evidence for the composite leading indicator index in the United States). Further, an interesting result is that the Dutch industrial production and money stock variables have common asymmetric patterns since the cycle lengths in the third quarter are almost equal.

To be more concrete, given the large cycle in the third quarter, and also given the high value of the parameter for y_{t-1} in that season, it can be seen that an exogenous shock in the second quarter has the largest effect on both the future patterns of industrial production and money

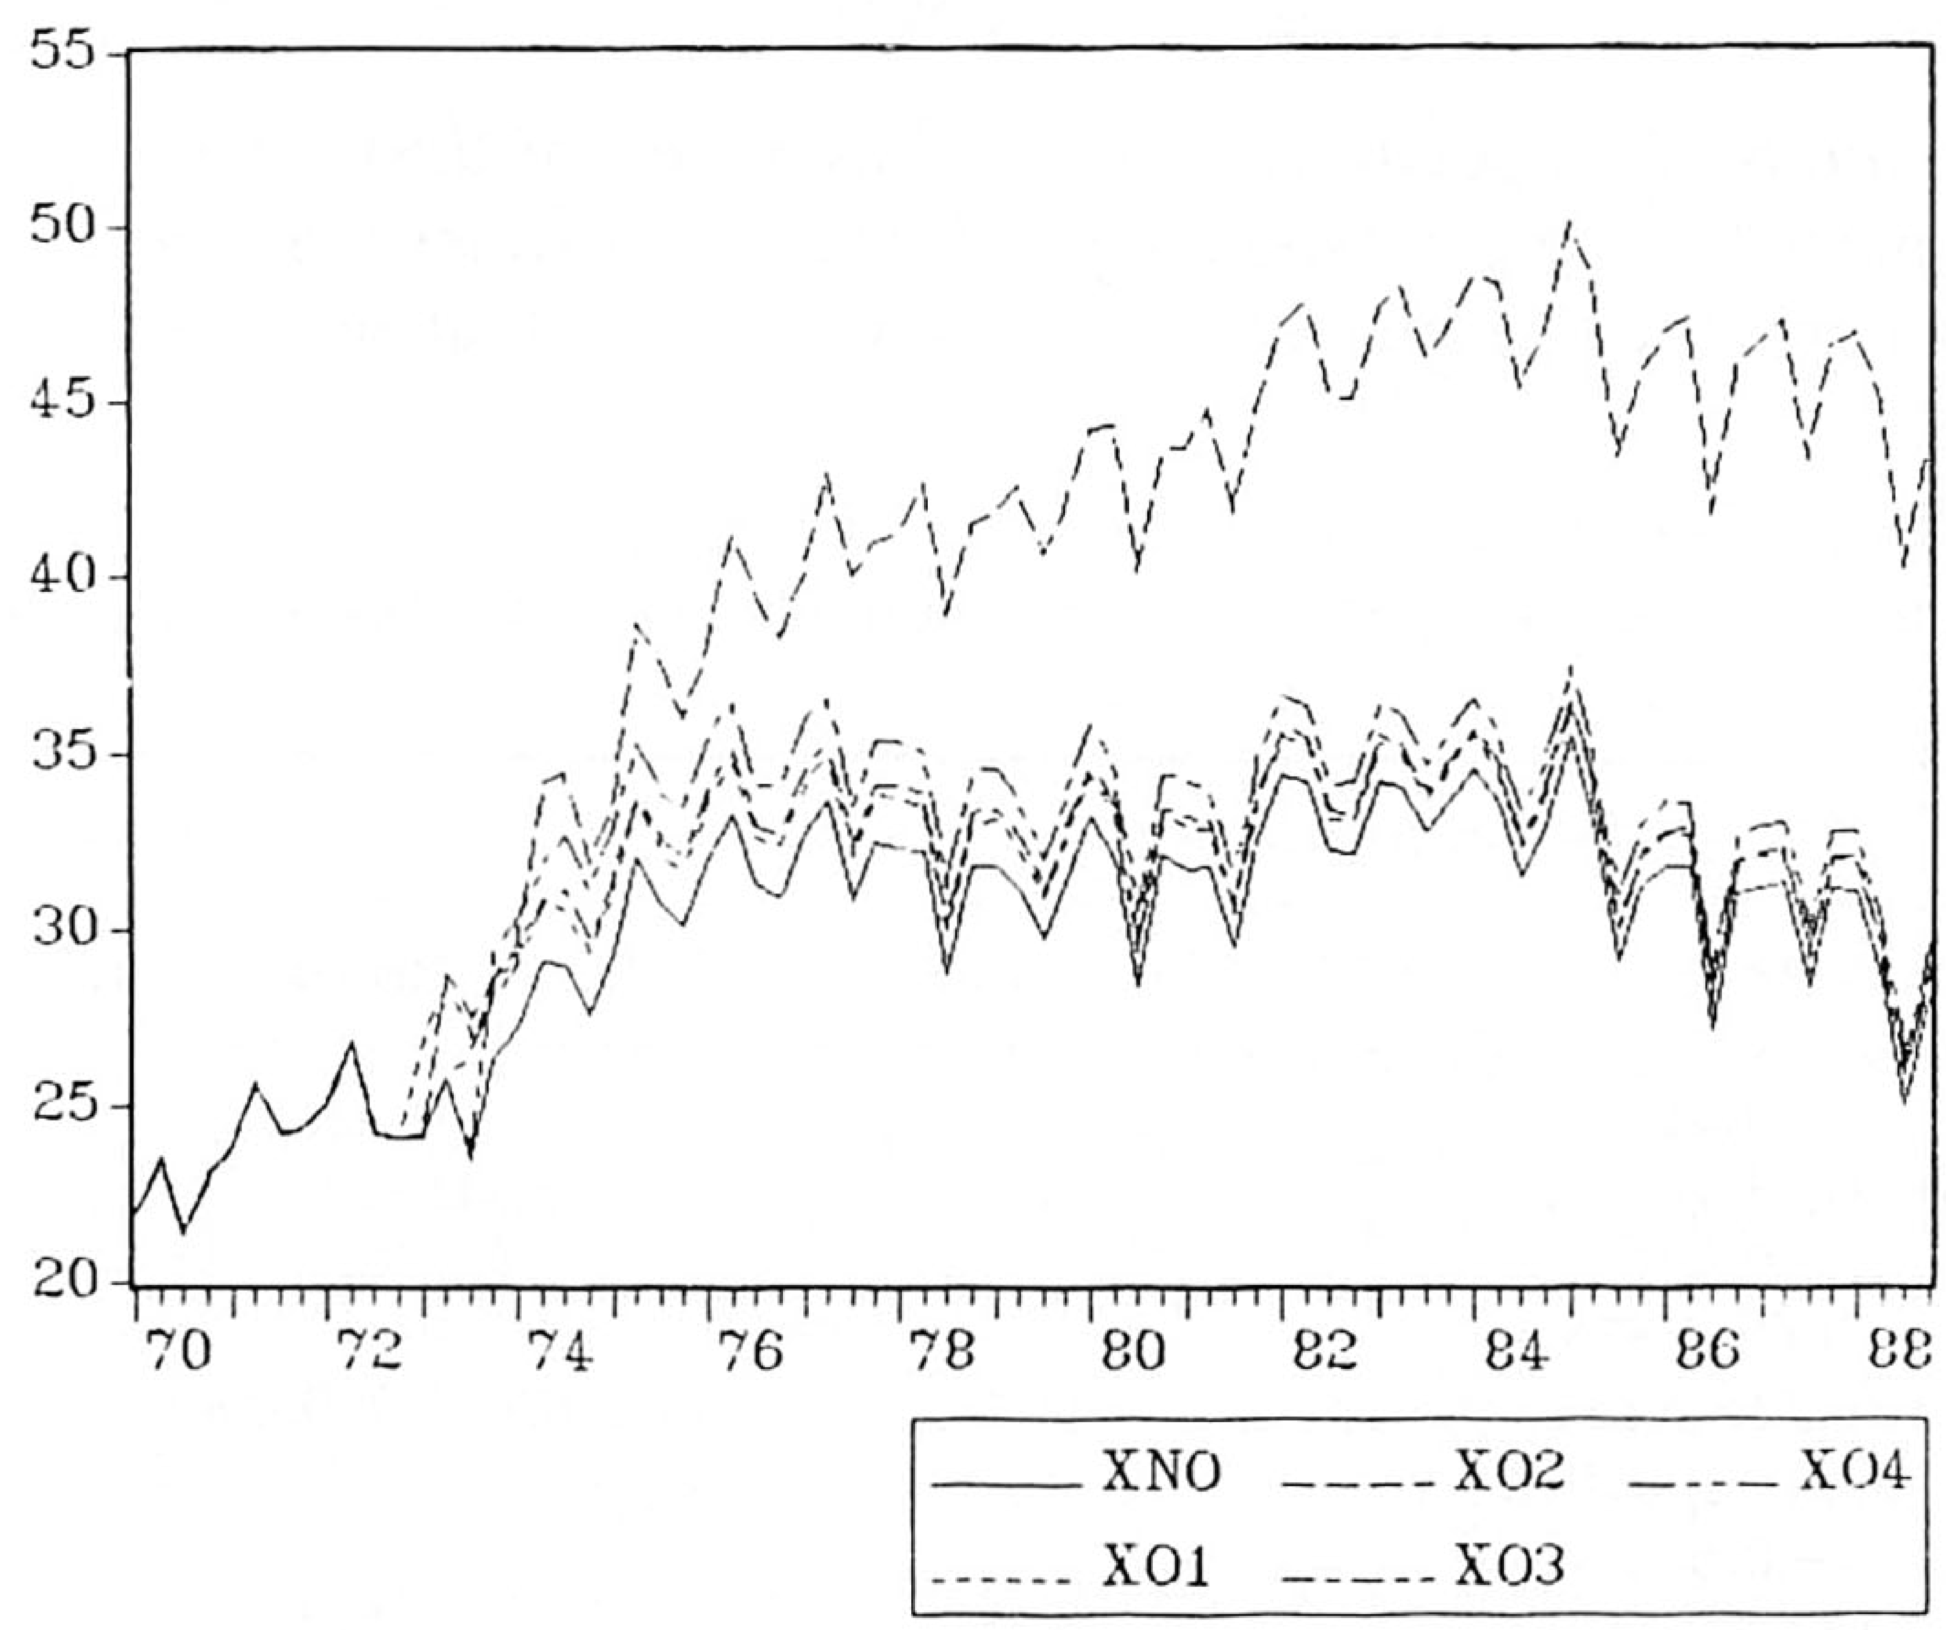

Figure 3. The effect of an innovative outlier in quarters 1, 2, 3, or 4. XNO is the series without an innovative outlier

Table III. An evaluation of the forecasting performance on the basis of 24 one-step-ahead forecasts and the root mean squared error (RMSE)

	Industrial production		Money stock	
RMSE (×100)	Model (2)	Model (12)	Model (3)	Model (16)
Total	4.117	3.856	6.373	6.441
Quarter 1	4.718	5.966	4.005	8.769
Quarter 2	4.738	2.424	4.907	5.431
Quarter 3	3.300	2.347	8.064	4.434
Quarter 4	3.496	3.576	11.772	11.085
SIGN	2.041[a]		0.408	

[a] Significant at the 5% level.
Note: Models (2) and (3) are non-periodic models, and models (12) and (16) are periodically integrated subset autoregressions. SIGN denotes the sign test for the equality of the squared forecasting errors from the two rival models.

stock. This occurrence for a univariate series is easily illustrated by simulating a quarterly series for, say, 1950.1–1999.4, from the model as it is given in equation (12) with the restriction in model (14) imposed, where ε_t is drawn from a standard normal distribution, and by adding an innovation outlier in, say, each of the quarters of 1973. The graphs of the effects of these outliers, with a magnitude of three times the standard deviation, are displayed in Figure 3. It is clear that an exogenous shock in the second quarter has a large impact. A similar result emerges for the model in equation (16).

Finally, one-step-ahead forecasts from the non-periodic models in equations (2) and (3) and from the PISAs in equations (12) and (16) are generated for 24 quarters. The results for the overall, as well as the per quarter, root mean squared error are displayed in Table III. The outcomes for the mean absolute prediction error are broadly similar and hence not reported. Also, the results for the non-parametric SIGN test are reported. This test is used to decide whether the squared errors from the rival models are significantly different. The (unreported) results for the Wilcoxon rank-sum test yield similar outcomes. From Table III it can be concluded that the periodic model for industrial production yields improved forecasts in two of the four seasons, and that the forecasts from the models for money stock are not significantly different. The results confirm the theoretical results in Osborn (1991) that a periodic model should outperform, or at least perform equal to, a non-periodic model when the first is appropriate.

CONCLUSIONS

One of the conclusions of this paper is that the quarterly Dutch industrial production index and money stock can be described by periodically integrated subset autoregressions. This means that each quarter is described by a different model with different lag structures and that cyclical patterns vary with the seasons. Interestingly, it turns out that both series have common asymmetric cyclical movements, i.e. for both series a shock in the second quarter has the largest effect and this effect is similar. Out-of-sample forecasts indicate that a periodic model can outperform a non-periodic model. The latter model, however, is rejected for both series since it is misspecified. Using a test for periodicity in the errors, it is shown that there is such periodicity in the non-periodic models as well as in the models for the seasonally adjusted series. This emphasizes the conjecture that linear seasonal adjustment filters may not remove all periodicity in the autocorrelation function.

The question is how one should proceed in practice when one wants to check for periodicity in economic time series. The two options are, first, to estimate a non-periodic model and to check for periodicity in the residuals, and second, to estimate a general periodic model and to use tests for parameter restrictions related to periodicity to obtain a simplified model, as in Franses (1992c). In the present paper it is argued that it seems most appropriate to work along the lines of Boswijk and Franses (1992), since this approach deals with periodicity and non-stationarity simultaneously. The major argument is that initially neglected periodicity can blur correct inference with respect to non-stationarity. In the present paper, for example, it is shown that periodicity causes one erroneously to use fourth-order differencing while no filter at all is appropriate.

As already noted, the estimated PISA models may not be very useful as representations of behaviour of economic agents. They should be regarded as descriptive statistical models which can be used as a starting point for the construction of multivariate dynamic econometric models. In fact, the revealed non-stationarity and periodicity may well be caused by periodic equilibrium relations between sets of variables and by periodically varying adjustment to disequilibrium errors. This multivariate extension of periodic integration is called periodic cointegration. Recently, an empirical modelling strategy for this type of cointegration has been proposed and applied in Franses and Boswijk (1992).

Finally, there remains the issue of seasonal adjustment in relation to periodic time series. The results in this paper suggest that linear seasonal adjustment filters applied to periodic time series may not yield series that are free from seasonal fluctuations. Hence, one conclusion is that seasonal adjustment for a periodically integrated series may not be possible without a specified model which relates trends, cycles, and seasons. Given the non-trivial relations between these components, this is not likely to be a straightforward exercise. Future research can be directed towards finding accurate and useful correction methods. However, the finding of periodic integration does not automatically preclude the use of periodic models when one wants, for one reason or another, to seasonally adjust a time series along standard lines. It may well be that forecasts from PISA type models can be used for adjusting current observations, and that this yields smaller data revision errors. Further theoretical developments and Monte Carlo simulations may shed some light on this issue. These investigations should then also consider the case when one wants to forecast a seasonally adjusted time series. A natural question is whether it is more appropriate to adjust a raw series and then to forecast using a periodic model, or to forecast the raw series using a PISA model and then to adjust the forecasts. Given that the outcomes in the present paper suggest that adjusting prior to modelling removes some but not all periodicity, it seems that the second strategy may be favourable. This conjecture can be verified via a thorough investigation of the theoretical effects of linear seasonal adjustment filters on periodically integrated time series.

NOTE

Discussions with Peter Boswijk, Eric Ghysels, Marius Ooms, and Timo Teräsvirta have proved to be very helpful, as well as the comments from Niels Haldrup and Svend Hylleberg at an Institute of Economics seminar in Aarhus. In particular, the extensive suggestions from an anonymous referee improved the paper in several respects. The financial support from the

Royal Netherlands Academy of Arts and Sciences is also gratefully acknowledged. The data used in this paper, as well as the unpublished papers in the References, are available from the author upon request.

REFERENCES

Barsky, R. B. and Miron, J. A., 'The seasonal cycle and the business cycle', *Journal of Political Economy*, **97** (1989), 503–34.

Beaulieu, J. J. and Miron, J. A., 'Seasonal unit roots in aggregate U.S. data', *Journal of Econometrics*, **55** (1993), 305–28.

Boswijk, H. P. and Franses, P. H., 'Testing for periodic integration', Econometric Institute Report 9216, Erasmus University Rotterdam, 1992.

Braun, R. A. and Evans, C. L., 'Seasonality and equilibrium business cycle theories', Institute for Empirical Macroeconomics Discussion Paper 45, Federal Reserve Bank of Minneapolis, 1991.

Franses, P. H., 'Seasonality, non-stationarity and the forecasting of monthly time series', *International Journal of Forecasting*, **7** (1991), 199–208.

Franses, P. H., 'Seasonality in consumer confidence in some European countries', Econometric Institute Report 9261, Erasmus University Rotterdam, 1992a.

Franses, P. H., 'A multivariate approach to modeling univariate seasonal time series', *Journal of Econometrics*, forthcoming, 1992b.

Franses, P. H., 'Testing for seasonality', *Economics Letters*, **38** (1992c), 259–62.

Franses, P. H. and Boswijk, H. P., 'Testing for periodic cointegration', Econometric Institute Report 9201, Erasmus University Rotterdam, 1992.

Ghysels, E., 'Are business cycle turning points uniformly distributed throughout the year?' Discussion paper 3981, CRDE, Montral, 1991.

Ghysels, E. and Hall, A., 'Testing periodicity in some linear macroeconomic models', Unpublished manuscript, CRDE, Montreal, 1992.

Hansen, L. P. and Sargent, T. J., 'Seasonality and approximation errors in rational expectations models', *Journal of Econometrics*, **55** (1993), 21–55.

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

Laroque, G., 'Analyse d'une méthode de désaisonnalisation: le programme X11 du US Bureau of Census, version trimestrielle', *Annales de l'INSEE*, **28** (1977), 105–27.

Lütkepohl, H., *Introduction to Multiple Time Series Analysis*, Berlin: Springer-Verlag, 1991.

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

Osborn, D. R., 'The implications of periodically varying coefficients for seasonal time-series processes', *Journal of Econometrics*, **28** (1991), 373–84.

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

Author's biography:

Philip Hans Franses is a research fellow of the Royal Netherlands Acadamy of Arts and Sciences, and he is affiliated to the Econometric Institute of the Erasmus University Rotterdam. He obtained an MSc degree in Economics from the University of Groningen and a PhD in Econometrics from the Erasmus University Rotterdam. His main research interests include time series, seasonality, and model selection. On these topics he has published in several journals, including *Economics Letters*, *Econometric Reviews*, *Journal of Marketing Research*, *International Journal of Forecasting*, and *Oxford Bulletin of Economics and Statistics*.

Author's address:

Philip Hans Franses, Econometric Institute, Erasmus University Rotterdam, PO Box 1738, NL-3000 DR Rotterdam, The Netherlands.