

Periodic integration: further results on model selection and forecasting

Philip Hans Franses, Richard Paap*

Received: August 1, 1994; revised version: March 21, 1995

This paper considers model selection and forecasting issues in two closely related models for nonstationary periodic autoregressive time series [PAR]. Periodically integrated seasonal time series [PIAR] need a periodic differencing filter to remove the stochastic trend. On the other hand, when the nonperiodic first order differencing filter can be applied, one can have a periodic model with a nonseasonal unit root [PARI]. In this paper, we discuss and evaluate two testing strategies to select between these two models. Furthermore, we compare the relative forecasting performance of each model using Monte Carlo simulations and some U.K. macroeconomic seasonal time series. One result is that forecasting with PARI models while the data generating process is a PIAR process seems to be worse than *vice versa*.

1 Introduction

Periodic autoregressive [PAR] time series models have proved to be useful in describing seasonally observed times series in such areas as water resources, *cf.* Vecchia & Ballerini (1991) and McLeod (1993) and economics, *cf.* Osborn & Smith (1989) and Franses (1994). The key feature of PAR models is that the autoregressive parameters take different values in different seasons. In order to use conventional identification techniques for univariate seasonal time series, it is necessary to remove one or more stochastic trends from the time series before any analysis. Typically, for quarterly observed times series one uses the first order $(1 - B)$ filter or the fourth order $(1 - B^4)$ filter for this purpose, where the backward shift operator B is defined by $B^k y_t = y_{t-k}$, *cf.* Box & Jenkins (1970). For periodic times series processes one may opt for a third

*This research was sponsored by the Economic Research Foundation, which is a part of the Netherlands Organization for Scientific Research (N.W.O.). The first author thanks the financial support from the Royal Netherlands Academy of Arts and Sciences. Part of the material in this paper has been presented at seminars at Tulane University and Brown University. Several discussions with Eric Ghysels and Hahn Lee proved to be very helpful. Comments from an anonymous referee are gratefully acknowledged.

possibility, *i.e.* that the differencing filter varies with the seasons, like, *e.g.*, $(1 - \alpha_s B)$ under the nonlinear restriction $\prod_{s=1}^4 \alpha_s = 1$, *cf.* Osborn (1988). In the latter case a seasonal time series is called a periodically integrated autoregression, denoted by PIAR, see Boswijk & Franses (1994) for a formal definition of the concept of periodic integration [PI]. Recent studies as Osborn (1988) and Franses & Paap (1994) suggest that PIAR models can be useful for modelling quarterly observed times series. Notice that periodic integration implies that in the long run shocks in different seasons have different impact, since the α_s values are not all equal to each other. In contrast, when a PAR time series requires the $(1 - B)$ filter to remove the stochastic trend, such shocks have the same impact across seasons. In this paper we denote the latter model by PARI. In practice, it is found that for the PIAR model the α_s values in the periodic differencing filter take values that are typically close to unity. Additionally, PIAR models are somewhat more complicated since the periodic differencing filter has to be determined using a nonlinear estimation method. Therefore, we focus in this paper on the selection between PIAR and PARI models.

We compare two selection strategies for univariate quarterly observed times series, which differ with respect to the sequence of tests. The first strategy, which is applied in Franses & Paap (1994) starts with a test for periodic integration. Next, one proceeds with a test whether the periodic differencing filter reduces to the $(1 - B)$ filter. The second strategy, advocated by Ghysels & Hall (1993), tests whether the $(1 - B)$ filter is applicable in a periodic model straightaway. The theoretical implications of both testing strategies are discussed in detail and their power and size properties are investigated using Monte Carlo experiments. We further study the effect on the forecasting performance in case the wrong differencing filter is applied using Monte Carlo experiments. Finally, both the test strategies and the forecast comparison are considered for a set of real life macroeconomic U.K. seasonal time series.

The outline of the paper is as follows. In section 2 we discuss some properties of PAR models. Although we confine ourselves to quarterly time series in this paper, we wish to stress that all tests and methods can also be applied to *e.g.* monthly time series. In section 3 we analyse the two testing strategies to distinguish between PIAR and PARI and apply them to seven U.K. macroeconomic time series. Section 4 deals with a comparison in forecasting performance using Monte Carlo simulations and the same seven time series. The final section provides some concluding remarks.

2 Preliminaries

Periodic autoregressive models are characterised by periodically varying parameters in the autoregressive part of the model. In other words, the dynamic structure varies

with the seasons. A periodic autoregressive model of order p , $\text{PAR}(p)$, for a quarterly observed time series, y_t , $t = 1, \dots, n = 4N$ can be written as

$$y_t = \sum_{s=1}^4 [\mu_s D_{st} + \tau_s D_{st} t] + \sum_{i=1}^p \sum_{s=1}^4 \phi_{is} D_{st} y_{t-i} + \epsilon_t, \quad (1)$$

where ϵ_t is standard white noise, D_{st} represent seasonal dummies, t a time trend and s equals 4 if $(t \bmod 4) = 0$ and $(t \bmod 4)$ otherwise. The ϕ_{is} are seasonally varying autoregressive parameters. An extension to periodically varying variances is straightforward, *i.e.* by including ϵ_{st} instead of ϵ_t .

The simplest model, which is useful to illustrate the properties of PIAR and PARI models, is the $\text{PAR}(2)$ model given by

$$y_t = \phi_{1s} y_{t-1} + \phi_{2s} y_{t-2} + \epsilon_t. \quad (2)$$

where s equals 4 if $(t \bmod 4) = 0$ and $(t \bmod 4)$ otherwise, which we will use in the remaining of the paper. To analyze stochastic trend properties of y_t it is useful to transform (2) into a model with constant parameters. Let $Y_T = (Y_{1T}, Y_{2T}, Y_{3T}, Y_{4T})'$ consist of the quarterly y_t observations, stacked in an annual vector, *i.e.* Y_{sT} is the observation in season s in year T . Then, (2) can be written in a Vector Autoregressive (VAR) representation

$$A_0 Y_T = A_1 Y_{T-1} + \epsilon_T \quad (3)$$

with

$$A_0 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ -\phi_{12} & 1 & 0 & 0 \\ -\phi_{23} & -\phi_{13} & 1 & 0 \\ 0 & -\phi_{24} & -\phi_{14} & 1 \end{pmatrix} \quad \text{and} \quad A_1 = \begin{pmatrix} 0 & 0 & \phi_{21} & \phi_{11} \\ 0 & 0 & 0 & \phi_{22} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (4)$$

where the index T runs from 1 to $N = n/4$ and ϵ_T is a (4×1) vector containing the stacked ϵ_t , see Gladyshev (1961).

The vector process in (3) is stationary if the roots of the characteristic equation

$$|A_0 z - A_1| = z^2(z^2 - \phi_1^* z - \phi_2^*) = 0, \quad (5)$$

are inside the unit circle, see Lütkepohl (1991), where ϕ_1^* and ϕ_2^* are nonlinear functions of the ϕ_{is} parameters in (2). The process y_t in (2) is then said to be periodically stationary. When one of the solutions to (5) is equal to one, which implies

$$\phi_1^* + \phi_2^* = 1, \quad (6)$$

and all other solutions are inside the unit circle, there are three cointegration relations

between the elements of Y_T in the VAR representation (3). It is easily understood that these cointegration relations can be written as $Y_{1T} - \gamma_1 Y_{4T}$, $Y_{2T} - \gamma_2 Y_{4T}$ and $Y_{3T} - \gamma_3 Y_{4T}$, and hence that $Y_{1T} - \alpha_1 Y_{4T}$, $Y_{2T} - \alpha_2 Y_{4T}$, $Y_{3T} - \alpha_3 Y_{4T}$ and implicitly $Y_{4T} - \alpha_4 Y_{4T}$ with $\alpha_4 = 1/(\alpha_1 \alpha_2 \alpha_3)$ are stationary variables, where the α_i parameters are functions of γ_i . Hence, when Y_T has one unit root, the appropriate differencing filter for y_t equals $(1 - \alpha_s B)$ under the restriction that $\alpha_1 \alpha_2 \alpha_3 \alpha_4 = 1$. In that case (2) can be written as

$$(y_t - \alpha_s y_{t-1}) = \beta_s (y_{t-1} - \alpha_{s-1} y_{t-2}) + \epsilon_t, \quad (7)$$

with $\alpha_1 \alpha_2 \alpha_3 \alpha_4 = 1$ and $\alpha_0 = \alpha_4$ if $s = 1$. The α_s and β_s are functions of the ϕ_{is} parameters. Given (7) and $\alpha_1 \alpha_2 \alpha_3 \alpha_4 = 1$, the characteristic equation (5) becomes

$$\begin{aligned} |A_0 z - A_1| &= z^2 (z - \alpha_1 \alpha_2 \alpha_3 \alpha_4) (z - \beta_1 \beta_2 \beta_3 \beta_4) = 0 \\ &= z^2 (z - 1) (z - \beta_1 \beta_2 \beta_3 \beta_4) = 0 \end{aligned} \quad (8)$$

It is clear from (8) that Y_T has a single unit root, which concern the values of α_s .

Two interesting situations can occur in the case of a single unit root in Y_T . Firstly, when $\alpha_s = 1 \forall s$, the differencing filter $(1 - \alpha_s B)$ reduces to $(1 - B)$. Notice that this corresponds to the parameter restriction $\phi_{1s} + \phi_{2s} = 1$ in (2). In that case model (7) reduces to a PAR(1) model for the first order differenced series $\Delta y_t = y_t - y_{t-1}$,

$$\Delta y_t = \beta_s \Delta y_{t-1} + \epsilon_t, \quad (9)$$

which we will denote as a periodic autoregression for an integrated time series [PARI]. Note that a PARI model is a special case of a PIAR model. The characteristic equation is equal to (8) with $\alpha_s = 1 \forall s$. On the other hand, when $\alpha_s = -1 \forall s$, the series y_t contains the seasonal unit root -1 at the bi-annual frequency, see Hylleberg *et al.* (1990). In this case the $(1 + B)$ filter is needed to remove the stochastic trend and this corresponds to the parameter restriction $\phi_{2s} - \phi_{1s} = 1$ in (2). Notice that $\alpha_s = -1 \forall s$ results again in the characteristic equation (8). In this paper we will not consider such a seasonal unit root. Secondly, when at least two α_s differ from 1 and, the periodic differencing filter $(1 - \alpha_s B)$ with $\alpha_1 \alpha_2 \alpha_3 \alpha_4 = 1$ is required to obtain periodic stationarity, we call y_t periodically integrated [PI], and the PAR model in (7) is then called a periodically integrated autoregressive model of order 2 [PIAR(2)]¹. Notice that a PARI model is a special case of a PIAR model. The key difference between a PIAR and a PARI model is that periodic integration implies that shocks in different

¹Notice that the conjecture in Ghysels & Hall (1993, footnote 3) that in the model $y_t = \delta_s y_{t-1} + \sum_{j=1}^{p-1} \theta_{js} \Delta y_{t-j} + \epsilon_t$ the restriction $\prod_{s=1}^4 \delta_s = 1$ implies periodic integration is not correct, unless p equals 1.

seasons have a different impact on the long run pattern of the time series. In fact, the stochastic trend may change the seasonal pattern of the series permanently, where the seasonal pattern is defined by the sequence of Y_{sT} ranging from highest to lowest within year T . In contrast to PIAR models, in PARI models shocks have the same impact on each Y_{sT} in the long run.

Finally, we discuss one more aspect of a PIAR model. Since a PIAR assumes the adequacy of a periodic differencing filter, one may want to consider fourth order differences in order to get rid of the periodic variation in removing the stochastic trend. For example the model in (7) with the restriction $\alpha_1\alpha_2\alpha_3\alpha_4 = 1$ can be rewritten in the following form

$$\Delta_4 y_t = \beta_s \Delta_4 y_{t-1} + \eta_t, \quad (10)$$

where $\Delta_4 y_t = y_t - y_{t-4}$ and η_t is a periodic moving average [MA] process defined by $\eta_t = \epsilon_t + \alpha_s \epsilon_{t-1} + \alpha_s \alpha_{s-1} \epsilon_{t-2} + \alpha_s \alpha_{s-1} \alpha_{s-2} \epsilon_{t-3}$ with $\alpha_{-2} = \alpha_2$, $\alpha_{-1} = \alpha_3$ and $\alpha_0 = \alpha_4$. If one neglects the periodic MA process, the representation in (10) suggests that a PIAR process is observationally equivalent to a process that is so-called seasonally integrated, see Hylleberg *et al.* (1990). Notice that this equivalence would amount to having 1 unit root in the Y_T process in the PIAR case and 4(!) such unit roots in the Δ_4 case. However, when one writes the η_t process in a stacked form, it can be shown that (10) becomes

$$\beta(B)(Y_T - Y_{T-1}) = \Theta_0 \epsilon_T + \Theta_1 \epsilon_{T-1}, \quad (11)$$

where $\beta(B)$ is a (4×4) matrix with polynomials in B , and $|\Theta_0 z + \Theta_1| = 0$ has three solutions on the unit circle. In other words, an application of the Δ_4 to a PIAR process implies overdifferencing since three redundant unit roots appear in the moving average part of the model for $\Delta_4 y_t$. Additionally, assuming all $\alpha_s = 1$ in a PIAR process, which amounts to the PARI process, it is even more obvious that PIAR, PARI and seasonally integrated processes are not observationally equivalent.

3 Model Selection

Now we turn to the issue of model selection in PAR models. This involves a determination of the order of the PAR and tests for the most appropriate differencing filter.

Determination of the Order p

A first model selection step concerns the decision on the order p of the periodic autoregression. This step can be based upon the Schwarz and Akaike and other Information criteria, on F -tests for parameter restrictions on the parameters of the highest lag,

and on LM tests for (periodic) autocorrelation in the residuals. An LM type test for first order periodic autocorrelation in a PAR(p) model boils down to an F -test for the significance of the four lagged $D_{st}\hat{\epsilon}_t$ in the following auxiliary regression

$$\hat{\epsilon}_t = \sum_{s=1}^4 [\delta_s D_{st} + \theta_s D_{st}t + \sum_{i=1}^p \psi_{is} D_{st}y_{t-i} + \rho_s D_{st}\hat{\epsilon}_{t-1}] + \xi_t. \quad (12)$$

After the appropriate order p has been determined, one can test whether the periodically varying coefficients differ significantly from each other, using a standard F -test. In Boswijk and Franses (1994) it is shown that such an F -test asymptotically follows a standard F distribution. If the null hypothesis of no periodicity cannot be rejected, one can proceed with testing for (non)seasonal unit roots in a nonperiodic model along the lines of Hylleberg *et al.* (1990). Otherwise one proceeds with testing for unit roots in periodic models. In case one expects there to be only a single stochastic trend, as we assume in the present paper, there are two possible strategies to follow. The first strategy is that one starts with a test for the restriction $\alpha_1\alpha_2\alpha_3\alpha_4 = 1$ and then check whether the periodic differencing filter can be simplified to the $(1 - B)$ filter. The second strategy tests the adequacy of the $(1 - B)$ filter straightaway.

Two Separate Tests

Franses & Paap (1994) opt for the approach to test first whether $\alpha_1\alpha_2\alpha_3\alpha_4 = 1$ in a general PAR(p) model written in a format like (7), and then test whether all $\alpha_s = 1$. For simplicity we explain this two-step strategy in the PAR(2) model (7). Extensions to higher order PAR models are straightforward and do not change asymptotic distributions of the tests. We begin testing the following parameter restriction in (7),

$$H_0 : \pi = \prod_{s=1}^4 \alpha_s = 1, \quad (13)$$

see (8), against the alternative $\pi < 1$. This H_0 can be tested using the studentized version of a Likelihood Ratio statistic

$$LR_\tau = \text{sign}(\hat{\pi} - 1) \sqrt{n \log \left(\frac{SSR_0}{SSR_a} \right)}, \quad (14)$$

where SSR_a corresponds to the sum of squared residuals in the unrestricted linear PAR model (2) and SSR_0 to the sum of squared residuals of the nonlinear model like (7) with $\alpha_1\alpha_2\alpha_3\alpha_4 = 1$ obtained after nonlinear least squares. Under the null hypothesis of a single unit root, the LR_τ follows a standard Dickey-Fuller distribution denoted by τ , see Boswijk & Franses (1994). In case one includes four seasonal dummies and/or

four seasonal trends² in the test equation, one should use the distributions for τ_μ and τ_τ tabulated in Fuller (1976, table 8.5.1).

If one rejects the hypothesis in (13), the series y_t is periodically stationary. If the null hypothesis cannot be rejected, *i.e.* $(1 - \alpha_s B)$ is the appropriate differencing filter for y_t , one can proceed with testing for the adequacy of the $(1 - B)$ filter in the series y_t . Given $\prod_{s=1}^4 \alpha_s = 1$, a test for $\alpha_s = 1 \forall s$ amounts to three restrictions in the PIAR model. These restrictions can easily be tested using a standard F -test which is asymptotically F distributed as proved in Boswijk & Franses (1994). If $\alpha_s = 1 \forall s$ cannot be rejected, we have a PARI model as in (9), otherwise we have PIAR model. Additionally a test for $\alpha_s = 1$ for some though not all s , *i.e.* the $(1 - B)$ filter is applicable in a few seasons, can also be performed using F -tests.

Table 1. Simulated fractiles of t -tests for the significance of lagged periodically differenced time series in a PAR(2) process based on 5000 replications. The effective sample size is 120 observations.

DGP ¹				U/R ¹	fractiles						
α_1	α_2	α_3	α_4		5%	10%	20%	50%	80%	90%	95%
1.053	0.888	1.071	0.999	U	-1.65	-1.21	-0.78	0.07	0.90	1.38	1.74
				R	-1.64	-1.26	-0.79	0.05	0.91	1.33	1.70
0.957	1.022	1.032	0.991	U	-1.61	-1.20	-0.76	0.06	0.92	1.39	1.77
				R	-1.61	-1.23	-0.78	0.01	0.92	1.37	1.73
0.744	1.021	1.371	0.960	U	-1.64	-1.26	-0.78	0.06	0.89	1.34	1.71
				R	-1.63	-1.24	-0.79	0.05	0.91	1.34	1.74

¹The DGP is $y_t = \alpha_s y_{t-1} + \epsilon_t$ with $\alpha_4 = 1/(\alpha_1 \alpha_2 \alpha_3)$ and $\epsilon_t \sim N(0, 1)$. The table shows the empirical fractiles of t -tests for the significance of the β_s parameters in the model $(y_t - \alpha_s y_{t-1}) = \beta_s (y_{t-1} - \alpha_{s-1} y_{t-2}) + \epsilon_t$ with the restriction $\alpha_4 = 1/(\alpha_1 \alpha_2 \alpha_3)$ imposed (R) and without this restriction (U).

A final step in this model selection strategy can involve a test for the significance of parameters like the β_s in the PAR model (7). Given that $(1 - \alpha_s B)y_t$ is a (periodically) stationary process, one may expect that t -tests for the significance of the β_s asymptotically follow a standard normal distribution under the null hypothesis. However, since the α_s have to be estimated, one may expect slightly biased distributions in small samples. Table 1 shows simulated fractiles of a t -test for the significance of one of the second order terms, while the DGP is a PIAR(1) for a sample size of 120 observations. These t -tests are performed in a PAR(2) model (7) with and without the parameter restriction of periodic integration, $\prod_{s=1}^4 \alpha_s = 1$. On the left hand side of the empirical

² It is easy to see that one always should include four seasonal dummies and four seasonal trends in the regression since the model $y_t - \mu - \tau t = \phi_s (y_{t-1} - \mu - \tau(t-1)) + \epsilon_t$ implies $y_t = \mu_s + \tau_s t + \phi_s y_{t-1} + \epsilon_t$ where $\mu_s = \mu - \phi_s \mu + \phi_s \tau$ and $\tau_s = (1 - \phi_s) \tau$.

distribution the fractiles closely match their asymptotic values, while on the right hand side the fractiles are only slightly larger.

A Joint Test

An alternative to the previous, what one can call a LR_T two-step method, is to start testing whether a $(1 - B)$ filter is appropriate straightaway. One approach may be to test *e.g.* in model (2) whether $\phi_{1s} + \phi_{2s} = 1$, using a standard F -test. This test is however not asymptotically F distributed, because the number of unit roots under the null hypothesis is different from the number of unit roots under the alternative hypothesis. Instead, one may use the tests proposed by Ghysels & Hall (1993). Rewrite the $PAR(p)$ model in (1) in the following form

$$\Delta y_t = \delta_s y_{t-1} + \mu_s + \tau_s t + \sum_{j=1}^{p-1} \theta_{sj} \Delta y_{t-j} + \epsilon_t, \quad (15)$$

where δ_s , μ_s , τ_s and θ_{sj} for $j = 1, \dots, p-1$ vary with the season and $\epsilon_t \sim N(0, \sigma^2)$. A first order differencing filter $(1 - B)$ for y_t corresponds to

$$H_0 : \delta_s = 0 \text{ for all } s. \quad (16)$$

This hypothesis is tested against the alternative

$$H_a : \delta_s \neq 0 \text{ for at least some } s, \quad (17)$$

which implies that y_t can be either periodically integrated or periodically stationary or even explosive. This can be viewed as a drawback since under H_a we still do not know whether y_t has a stochastic trend or not. The hypothesis in (16) can be tested using a Wald test. This Wald test can be expressed as

$$W_{i4} = \frac{N}{n - k} \sum_{s=1}^4 t_{\delta_s}^2, \quad (18)$$

where t_{δ_s} represents the t -value for the test $\delta_s = 0$, k is the number of regressors in (15) and $i = 1, 2, 3$ in case (15) contains, 1: no seasonal dummies and no trends, 2: only four seasonal dummies and 3: four seasonal dummies and four seasonal trends. For a similar reason as in footnote 2, we have to include four seasonal dummies and four seasonal trends. Boswijk & Franses (1994) prove that the asymptotic distribution of the Wald test in (18) is the sum of a $\chi^2(3)$ and the square of the Dickey-Fuller distribution, indicating that W_{i4} amounts to a joint test for a single unit root in Y_T and for $\alpha_s = 1 \forall s$. We simulated some critical values of the tests, see table 2. These critical values will be used in the empirical part of this paper.

Table 2. Critical values of the W_{i4} tests based on 10000 Monte Carlo replications^{1,2}.

sample size	W_{14}			W_{24}			W_{34}		
	20%	10%	5%	20%	10%	5%	20%	10%	5%
40	1.74	2.35	2.90	2.63	3.37	4.17	3.61	4.63	5.62
80	1.62	2.15	2.62	2.36	3.01	3.59	3.12	3.84	4.51
120	1.62	2.13	2.60	2.35	2.95	3.52	3.03	3.71	4.33
160	1.59	2.06	2.51	2.30	2.86	3.40	2.97	3.67	4.27

¹The DGP is $y_t = y_{t-1} + \epsilon_t$ with $\epsilon_t \sim N(0, 1)$ and the model to be estimated is $\Delta y_t = \delta_s y_t + \eta_t$, including seasonal dummies and trends if necessary.

² W_{14} corresponds to the test statistic in (18) when the test equation contains no seasonal dummies and trends, W_{24} with seasonal dummies and W_{34} with seasonal dummies and trends.

In order to test the validity of the $(1 - B)$ filter in a $PAR(p)$ model, another possibility is to modify the Dickey-Fuller t -test for periodic autoregressions by considering

$$\Delta y_t = \delta y_{t-1} + \mu_s + \tau t + \sum_{j=1}^{p-1} \theta_{sj} \Delta y_{t-j} + \epsilon_t \quad (19)$$

and a t -test for the significance of δ , τ_{PADF} , *i.e.* a periodic Augmented Dickey-Fuller test [PADF]. Under the null hypothesis, this t -test follows a standard Dickey-Fuller distribution, tabulated in Fuller (1976, table 8.5.1), see Ghysels & Hall (1993). Note that they only allow for a nonseasonal trend term as in (19). If one cannot reject the hypothesis using the τ_{PADF} test the series contains a unit root at the zero frequency. In case one rejects the null hypothesis, there are again several possibilities. The series can be periodically integrated, stationary or even explosive. Note that not every PIAR model of order p can be written in the form (19) with $\delta \neq 0$ due to the restriction that δ is not periodically varying, so not all PIAR processes are captured in the alternative hypothesis. Hence, the τ_{PADF} is in a sense more restrictive than the W_{i4} test.

Some Simulation Results

To compare the two-step and joint model selection strategies we set up a Monte Carlo study. The performance of the test statistics is investigated using six data generating processes [DGPs]. Table 3 shows the DGPs we consider in our simulations. Table 4 shows the outcome of the Monte Carlo simulations. The first block displays the relative number of cases a decision is made if the test equations do not contain any deterministic components, the second if the test equations contain seasonal dummies and the last block if the test equations contain seasonal dummies and seasonal trends except for (19), where we include a nonseasonal trend. We first concentrate on the

Table 3. The parameters of the DGPs, used in the Monte Carlo studies¹.

DGP		α_1	α_2	α_3	α_4	β_1	β_2	β_3	β_4
I	PIAR	1.053	0.888	1.071	0.999	-0.253	-0.352	-0.081	0.331
II	PIAR	0.957	1.022	1.032	0.991	0.009	-0.649	-0.398	-0.646
III	PIAR	0.744	1.021	1.371	0.960	-0.302	-0.539	-0.118	-0.365
IV	PARI	1	1	1	1	-0.315	-0.657	-0.211	0.127
V	PARI	1	1	1	1	-0.037	-0.554	-0.283	-0.663
VI	PARI	1	1	1	1	-0.354	-0.334	-0.027	-0.436

¹Each DGP is $(y_t - \alpha_s y_{t-1}) = \beta_s (y_{t-1} - \alpha_s - 1 y_{t-2}) + \epsilon_t$, with $\alpha_4 = 1/(\alpha_1 \alpha_2 \alpha_3)$ and $\epsilon_t \sim N(0, 1)$. The chosen values of the autoregressive parameters are based upon parameter estimates of a second order PIAR and a first order PARI model for three U.K. macroeconomic time series, Total Investment, Exports and Trade Balance, see Osborn (1990) for a complete description of the data and Franses & Paap (1994) for additional details.

Table 4. Performance of the test strategies, based on 5000 replications. The sample size is 120. The cells report the relative frequencies that a certain decision is made based on the proposed test strategy¹. All tests are evaluated at a 5% significance level.

strategy	decision	DGP					
		I	II	III	IV	V	VI
<i>no constants and no trends</i>							
LR_τ	PI	94.74	95.18	95.18	94.14	94.62	94.68
LR_τ 2-step	PI & no Δ	75.00	26.76	95.14	4.90	4.34	4.94
LR_τ 2-step	PI & Δ	19.74	68.42	0.04	89.24	90.28	89.62
W_{14} test	Δ	24.22	76.80	0.04	94.68	94.50	93.96
PADF	Δ	91.96	95.12	92.76	94.80	94.58	94.32
<i>constants and no trends</i>							
LR_τ	PI	94.94	95.28	95.00	95.40	95.14	95.04
LR_τ 2-step	PI & no Δ	50.14	11.60	94.66	5.02	5.10	4.80
LR_τ 2-step	PI & Δ	44.80	83.68	0.34	90.38	90.04	90.24
W_{24} test	Δ	55.66	91.19	0.32	94.62	94.36	94.96
PADF	Δ	93.78	96.42	95.70	95.50	96.40	95.56
<i>constants and trends</i>							
LR_τ	PI	96.02	96.20	96.74	95.62	96.48	95.80
LR_τ 2-step	PI & no Δ	26.20	7.40	93.52	5.10	5.60	4.80
LR_τ 2-step	PI & Δ	69.42	88.80	3.22	90.52	90.88	91.00
W_{34} test	Δ	81.16	93.60	4.34	96.08	94.60	94.24
PADF	Δ	94.88	93.06	77.56	93.92	97.06	96.30

¹The DGPs are displayed in table 3.

results for the first three DGPs, where the series are periodically integrated. The lag order p in the test regressions is set equal to 2 as in the DGP except for the PADF case in equation (19) which can suffer from serial correlation in the residuals because of the constant δ parameter. In that case we add lagged Δy_t with periodically varying parameters to this equation until there is no significant periodic serial correlation in the residuals using the LM test based on (12). We see that the LR_τ test performs as expected, *i.e.* the relative number of cases that is rejected equals about the significance level of 5%. The PADF test has very low power against periodic integration, since it almost always opts for the $(1 - B)$ filter except for DGP III in case seasonal dummies and a trend are included in the test equation. The test strategies have difficulties with the second DGP, where the α_s are near unity and almost no difficulties with the third DGP, where some of the α_s differ substantially from unity. The LR_τ two-step strategy chooses the inappropriate $(1 - B)$ filter less frequently than the W_{i4} and PADF tests, especially for the first two DGPs. The relative frequency that the test strategies select a $(1 - B)$ filter increases if we include deterministic elements in the test equations.

When the DGPs are PARI processes (the last three DGPs in table 3) the overall results for the empirical size of the LR_τ test do not change. The PADF and the W_{i4} test provide in about 95% of the cases the right decision, which corresponds to the chosen level of significance. Note that now we have set the order p of the regression (19) equal to 2 without a priori testing, since the auxiliary regression corresponds to the DGP. The LR_τ two-step method suffers from the fact that we have to perform two tests and hence only in about 90% of the cases we make the right decision. The results do not change if we include deterministic components in the test equations.

In sum, the empirical performance of the model selection strategies in case the DGP is PARI seems similar. The LR_τ two-step strategy seems to opt for the $(1 - B)$ filter less frequently than the W_{i4} test in case the DGP is PIAR. The PADF test has no power against periodic integration whatsoever. The main theoretical advantage of the LR_τ two-step strategy method is that one obtains information on the stochastic trend properties in case $(1 - B)$ is rejected, while a rejection using W_{i4} tests does not lead to any concrete decision. The practical advantage of the two-step strategy comes from the fact that in case in the first step $\alpha_1\alpha_2\alpha_3\alpha_4 = 1$ cannot be rejected, we can test for the significance of the deterministic components using standard F tests, before continuing with the second step, *i.e.* testing for a $(1 - B)$ filter, which can increase the power of the test strategy. Under the restriction $\alpha_1\alpha_2\alpha_3\alpha_4 = 1$, these F tests are asymptotically F distributed.

Empirical Results

The two test strategies have been applied to seven log transformed U.K. macroeconomic time series. A detailed report of the first stages of the model selection procedure, including determination of model order p and testing for periodicity can be found in Franses & Paap (1994). It turns out that all series can be described by a periodic model. Table 5 displays the outcomes of the LR_τ two-step strategy and the W_{24} , W_{34} and PADF tests. The LR_τ two-step method results suggest that all series are periodically integrated, *i.e.* the $(1 - B)$ filter is not applicable. Results are the same if we add seasonal trends in the test equations, see Franses & Paap (1994). The periodic ADF test indicates that the $(1 - B)$ filter is always useful. This conclusion does not change if we include a deterministic trend in the test equations. Using the W_{24} test we find that $(1 - B)$ is not appropriate for all series. If we however include four seasonal deterministic trends in the test equation, the W_{34} test cannot reject the hypothesis of the presence of a $(1 - B)$ filter except for Imports.

Table 5. The test statistics for the LR_τ two-step strategy and the outcomes of the W_{24} , W_{34} and PADF test for seven log transformed U.K. macroeconomic time series¹.

series ²	order ³	LR_τ	F_Δ^4	lags ³	PADF	lags ³	W_{24}	W_{34}
Imports	1	0.105	4.772**	1	-0.351	-	3.778**	6.126**
Workforce	1	0.584	9.499**	1,4	-0.110	-	10.257**	0.461
Nond. Cons.	1	0.586	31.235**	1,4	-0.358	-	24.857**	1.818
Exports	2	-1.074	6.292**	1	-0.764	1	5.561**	2.570
Total Inv.	2	-1.414	10.155**	1,4	-0.724	1	8.925**	1.113
Private Inv.	1	-0.146	4.467**	1	-0.019	-	3.591**	2.338
Total Cons.	1	1.108	26.768**	1,4	0.831	-	28.910**	1.093

** significant at a 5% level.

¹The samples are 1955.1–1988.4, except for Private Investment, 1962.1–1988.4.

²Each model contains four seasonal dummies and four seasonal trends for W_{34} case. The model for Workforce contains a dummy variable for 1959.2, for Exports in 1967.4 and 1968.1 and for Total Consumption for 1979.3 and 1980.2 to capture outlying observations. The results do not change much when we exclude these dummy variables.

³Order denotes the order of the PAR model, which is used to test for periodic integration. Lags denote which lagged Δy_t are used in the test regression while testing for the $(1 - B)$ filter using the PADF and the W_{i4} tests. The number of lags is determined using LM type diagnostics for periodic autocorrelations in the residuals.

⁴ F_Δ denotes the outcomes of an F -test for the validity of a $(1 - B)$ filter in a PIAR model. This F_Δ test asymptotically follows a standard F distribution.

Table 6 shows the parameter estimates of the PIAR models for the seven series together with some diagnostic test statistics. The test for serial correlation in the residuals in the nonlinear PIAR models are computed using the Gauss-Newton re-

Table 6. Parameter estimates of the PIAR models¹ for seven U.K. macroeconomic time series.

par.	Imports	Workf. ²	N. Cons.	Exports ²	Tot. Inv.	Priv. Inv.	T. Cons. ²
μ_1	0.347 (0.138)	0.568 (0.116)	-0.141 (0.082)	0.390 (0.126)	-0.506 (0.192)	0.573 (0.237)	-0.432 (0.098)
μ_2	-0.229 (0.140)	-0.189 (0.123)	0.762 (0.076)	0.086 (0.130)	0.817 (0.182)	0.302 (0.237)	0.912 (0.089)
μ_3	0.267 (0.134)	-0.564 (0.119)	-0.297 (0.083)	-0.384 (0.133)	-0.528 (0.218)	-0.632 (0.249)	-0.573 (0.099)
μ_4	-0.261 (0.139)	0.155 (0.112)	-0.341 (0.081)	-0.086 (0.104)	0.270 (0.210)	-0.252 (0.234)	0.058 (0.093)
α_1	0.965 (0.014)	0.944 (0.011)	1.004 (0.008)	0.957 (0.013)	1.053 (0.022)	0.925 (0.026)	1.033 (0.009)
α_2	1.038 (0.015)	1.019 (0.012)	0.932 (0.007)	1.022 (0.016)	0.888 (0.019)	0.972 (0.264)	0.918 (0.008)
α_3	0.974 (0.014)	1.056 (0.012)	1.030 (0.008)	1.032 (0.015)	1.071 (0.017)	1.073 (0.028)	1.057 (0.009)
β_1				0.009 (0.143)	-0.253 (0.214)		
β_2				-0.649 (0.166)	-0.351 (0.137)		
β_3				-0.398 (0.124)	-0.080 (0.160)		
β_4				-0.646 (0.153)	0.331 (0.237)		
diagnostic test statistics ³							
J-B	1.000 (0.606)	1.353 (0.508)	5.521 (0.063)	2.211 (0.331)	0.599 (0.741)	2.067 (0.356)	4.280 (0.112)
LM ₁	1.966 (0.163)	0.268 (0.605)	0.050 (0.823)	3.200 (0.076)	0.764 (0.384)	3.417 (0.068)	0.324 (0.570)
LM ₄	1.286 (0.279)	2.168 (0.077)	1.400 (0.238)	0.795 (0.531)	2.068 (0.090)	2.075 (0.090)	0.891 (0.472)

¹The models are $y_t = \sum_{s=1}^4 D_{st}[\alpha_s y_{t-1} + \mu_s + \beta_s(y_{t-1} - \alpha_{s-1} y_{t-2})] + \epsilon_t$ with $\alpha_1 \alpha_2 \alpha_3 \alpha_4 = 1$, where D_{st} represent seasonal dummies (standard errors between brackets).

²We include dummy variables for Workforce in 1959.2, for Exports in 1967.4 and 1968.1 and for Total Consumption for 1979.3 and 1980.2 to capture outlying observations.

³J-B is the Jarque-Bera normality test, LM₁ the F -version of a test on first order serial correlation in the residuals and LM₄ the same test for first-to-fourth order serial correlation (p-values between brackets). The LM tests are calculated using (21).

Table 7. Parameter estimates of the PARI models¹ for seven U.K. macroeconomic time series.

par.	Imports	Workf. ²	N. Cons.	Exports ²	Tot. Inv.	Priv. Inv.	T. Cons. ²
μ_1	0.010 (0.007)	0.002 (0.001)	-0.076 (0.017)	-0.027 (0.008)	-0.006 (0.014)	- 0.102 (0.009)	-0.056 (0.014)
μ_2	0.039 (0.007)	0.001 (0.001)	-0.008 (0.016)	0.029 (0.007)	-0.038 (0.009)	0.055 (0.009)	- 0.016 (0.013)
μ_3	0.015 (0.009)	0.003 (0.001)	0.034 (0.008)	0.005 (0.008)	0.012 (0.012)	0.026 (0.009)	0.034 (0.007)
μ_4	-0.013 (0.007)	-0.000 (0.001)	0.019 (0.008)	0.026 (0.006)	0.042 (0.017)	0.067 (0.009)	0.038 (0.011)
β_{11}	-0.221 (0.172)	0.563 (0.212)	0.161 (0.152)	0.037 (0.150)	-0.216 (0.218)		0.226 (0.250)
β_{12}	-0.320 (0.164)	-0.354 (0.158)	-0.233 (0.160)	-0.554 (0.146)	-0.590 (0.130)		- 0.421 (0.154)
β_{13}	-0.009 (0.165)	-0.110 (0.118)	-0.292 (0.113)	-0.283 (0.126)	-0.113 (0.116)		- 0.375 (0.110)
β_{14}	-0.257 (0.172)	-0.022 (0.132)	0.323 (0.191)	-0.663 (0.159)	0.134 (0.206)		0.044 (0.139)
β_{41}		0.334 (0.142)	0.279 (0.160)		0.185 (0.132)		0.436 (0.160)
β_{42}		0.039 (0.125)	0.691 (0.102)		0.528 (0.109)		0.557 (0.102)
β_{43}		0.615 (0.142)	0.119 (0.206)		0.631 (0.217)		0.287 (0.168)
β_{44}		0.193 (0.201)	0.490 (0.158)		0.207 (0.211)		0.128 (0.016)
diagnostic test statistics ³							
J-B	0.061 (0.970)	6.700 (0.035)	1.399 (0.497)	0.300 (0.861)	6.076 (0.048)	1.155 (0.561)	10.630 (0.005)
LM ₁	0.011 (0.916)	0.388 (0.535)	1.373 (0.244)	3.821 (0.053)	0.032 (0.858)	2.953 (0.089)	0.022 (0.883)
LM ₄	0.904 (0.464)	0.325 (0.861)	1.228 (0.303)	2.153 (0.078)	0.473 (0.756)	0.098 (0.258)	1.152 (0.336)

¹The models are $\Delta y_t = \sum_{s=1}^4 D_{st}[\mu_s + \sum_{i=1}^{p-1} \beta_{is} \Delta y_{t-i}] + \epsilon_t$, where D_{st} represents seasonal dummies.²We include dummy variables for Workforce in 1959.2, for Exports in 1967.4 and 1968.1 and for Total Consumption for 1979.3 and 1980.2 to capture outlying observations.³J-B is the Jarque-Bera normality test, LM₁ the F -version of a test on first order serial correlation in the residuals and LM₄ the same test for first-to-fourth order serial correlation (p-values between brackets).

gression, see Davidson & MacKinnon (1993). Write the PIAR model in the following form

$$y_t = x_t(\Phi) + \epsilon_t, \quad (20)$$

where $x_t(\Phi)$ denotes the explanatory part of the PIAR model and Φ a parameter vector containing say $(\mu_s, \alpha_s, \beta_s, \forall s)$. Testing for serial correlation of order one in the $\hat{\epsilon}_t$ process boils down to testing whether ρ differs significantly from zero using a F -test in the following linear regression

$$\hat{\epsilon}_t = X_t(\hat{\Phi})b + \rho\hat{\epsilon}_{t-1} + \eta_t, \quad (21)$$

where $X_t(\hat{\Phi})$ is a vector containing the first derivatives of the regression function x_t with respect to the elements of Φ evaluated at the ML estimates $\hat{\Phi}$. The diagnostic test results in table 6 indicate that the PIAR models do not seem to be misspecified. Furthermore, note that the α_s are estimated close to unity indeed. Finally, we do not restrict some β_s values to zero for illustrative purposes.

Table 7 shows the PARI models we have estimated for the seven series together with tests for serial correlation and normality of the residuals. For comparability purposes, we only include four seasonal dummies in the PARI models, similar to what we do for the PIAR models. We will use the estimates in table 7 for our forecasting experiment in the next section. Note that we often need the fourth lag of Δy_t to get rid of serial correlation in the residuals. In the following section we investigate whether the imposed restriction $(1 - B)$ as in a PARI model has implications for forecasting using Monte Carlo simulations and forecasts generated by the empirical models constructed in this section.

4 A Forecasting Comparison

In the previous section we considered model selection strategies for PAR models. We now turn to the question whether it matters for forecasting if one selects either one of the models. A Monte Carlo experiment can give some insight in this subject matter. We also investigate how PIAR and PARI models perform in practice when comparing the forecasting performance of the models for the seven U.K. series analysed in the previous section.

Monte Carlo Results

First, we set up a Monte Carlo experiment to analyze the forecasting performance of both models when the DGP is either the PARI or the PIAR process. We use the six

Table 8. Forecasting performance of PIAR and PARI models, based on 5000 replications. The cells report the relative frequencies that a PIAR model forecasts better than a PARI model, based on the MSPE of the one-step ahead forecasting errors. The sample size used for estimation is 120.

DGP ¹		number of forecasts ¹	I	II	III	IV	all quarters
PIAR ²	I	96	53.62	90.76	51.80	61.80	86.96
		48	49.74	82.66	47.44	60.18	82.58
		12	54.66	45.12	65.04	47.40	73.20
	II	96	64.44	50.38	46.90	40.52	51.20
		48	62.62	51.98	46.62	42.68	52.42
		12	56.76	51.52	46.90	45.72	52.74
	III	96	97.80	96.64	90.40	83.24	99.90
		48	92.16	91.56	76.16	74.92	99.50
		12	69.36	70.98	53.18	57.76	94.38
PARI ²	IV	96	42.16	42.96	43.64	42.26	17.76
		48	45.54	45.20	46.48	46.36	25.30
		12	48.38	48.68	48.74	48.06	37.58
	V	96	40.94	42.92	40.78	43.74	17.42
		48	43.82	45.54	44.82	46.18	25.76
		12	48.02	47.68	48.78	49.30	39.22
	VI	96	41.56	43.32	42.00	42.18	17.22
		48	45.20	46.08	45.78	44.54	25.48
		12	48.54	48.90	48.92	48.14	37.20

¹In each replication 216 observations are generated using the DGPs displayed in table 3. The first 120 observations are used to estimate a PIAR and a PARI model, while the last 96 observations are used to evaluate the one-step ahead forecasts based on the estimated models.

²When the DGP is PARI(1), we estimate a PIAR(2) and PARI(1) model, while when the DGP is PIAR(2) we estimate a PIAR(2) and a PARI model of order k , where the order k is the smallest order for which the model does not contain periodic serial correlation in its residuals based on the periodic LM test, see (12).

DGPs of table 3 to generate 216 observations. The first 120 observations are used to estimate a PIAR and a PARI model and the last 96 are used to evaluate forecasts. Table 8 shows the outcomes of the simulations. The cells of the table show the relative frequency that the PIAR model forecasts better than the PARI model. The comparison is based upon the mean squared of one-step ahead forecasting errors [MSPE] in each season separately and over all seasons. Hence, the cells of table 8 show the relative number of times that the MSPE of the PIAR forecasts is smaller than that of the PARI forecasts. So, it is possible that the forecasting performance of the PARI model in each season separately is better than in the four quarters together and *vice versa*.

The first three DGPs correspond with a PIAR process. Especially for DGP I and III, where one of the α_s differs substantially from 1, we observe that it matters if one selects the wrong model. The outcomes of the simulations depend on the chosen parameters. For the third DGP the PIAR model forecasts better in more than 90% of the cases, while for the second DGP, where the α_s values are near 1, there is not so much difference in forecasting performance.

These results are in contrast to those for the last three DGPs, which are PARI processes, *i.e.* the outcomes for each DGP are roughly the same. A PIAR model forecasts better in each quarter in about 45% of the cases. However, the forecasting performance of PIAR models decreases the more forecasts one evaluates. Notice further that when we evaluate the forecasts within each season, we get a different picture than when we consider the forecasting performance in the four seasons together. In that case the PIAR model may be better only in between 17% and 40% of the cases, depending on the number of forecasts.

Empirical Results

Another evaluation of the relative forecasting performance of both models can be obtained using the models for the U.K. series discussed previously in tables 6 and 7.

We assume three different forecasting horizons. The first part of the sample is used to estimate the model, while the second part is used to evaluate the forecasts. From table 9 we see that for the separate quarters in 20 of the 84 cases the MSPE of the PIAR models is significantly smaller than the MSPE of the PARI models, using nonparametric sign test at a 10% level. In 13 cases it is the other way around. In 51 cases there is no significant difference in forecasting performance, although the PIAR model forecasts better in 27 of these 51 cases. Especially, in the third quarter the PIAR model produces better forecasts than the PARI model. When we look at the forecasting performance of the four quarters together, we see that in 5 cases the PIAR models forecast significantly better and in 2 cases the PARI model. The balance

Table 9. Forecasting comparison of PIAR versus PARI models. The cells report whether the MSPE of one-step ahead forecasts from the PIAR model is smaller (–) or larger (+) than those from the PARI model. Significant differences are checked using a nonparametric sign-test.

series	number of forecasts	quarters				total
		I	II	III	IV	
Imports	72	–	+	–	–	–
	48	–	–	+*	–	–
	24	–	–*	+*	–**	+*
Workforce	72	+	+	–**	–	–**
	48	+	+	–**	–	+
	24	+*	+	–*	+	+
Nondurable Consumption	72	–	–*	+	+**	–
	48	–	+	–	–	–
	24	+*	+	–	–**	–
Exports	72	–	+	–*	+**	+
	48	–	+*	–*	+	+
	24	–	+*	–	+*	+
Total Investment	72	+	+*	–	–	+
	48	+	+	+	–*	+
	24	+	+	+*	–*	+
Private Investment	72	+	+	+	+*	+*
	48	–*	+	–*	–**	–**
	24	–*	+	–**	–**	–**
Total Consumption	72	–	–	–	+*	+
	48	–*	–	–	–	–*
	24	–**	+	–*	–	–**

* significant at 10% and ** significant at 5% using a two sided nonparametric sign-test.

¹The models are estimated until 70.4, 76.4 and 82.4, while forecasts are generated for 1971.1–1988.4 (72), 1977.1–1988.4 (48) and 1983.1–1988.4 (24), respectively.

between + and - signs is roughly equal. The number of cases that the PIAR model forecasts (significantly) better is larger when the forecasting horizon is smaller. The forecasting performance of the PARI model is slightly better when the oil crisis in 1973 is included in the forecasting evaluation period *i.e.* for 72 out of sample forecasts, while the PIAR model forecasts significantly better when the crisis period is in the estimation sample.

To summarize, on the basis of Monte Carlo experiments we conjecture that in general, forecasting with PARI models, when the DGP is PIAR, is worse than *vice versa*, especially when at least one of the α , parameter differs substantially from unity. The empirical results in table 9 are however not as convincing as the Monte Carlo results.

5 Conclusion

In this paper we considered model selection and forecasting issues of two restricted periodic autoregressive models. The first is a periodic model in first differences, the second a periodically integrated model, which contains a stochastic trend with an impact that varies with the seasons. The impact of a shock in the latter model depends on the quarter in which the shock occurs and the impact is different for each season. In the former model shocks have the same impact on each quarter of the series in the long run. Therefore, the choice between the two models becomes relevant if one wants to make forecasts based on an estimated model. Since a periodic model in first differences is a special case of a periodically integrated model, a possible selection strategy is to test for periodic integration first and then to test whether the periodic differencing filter can be simplified to a first difference filter. In case one is only interested in the question whether the series contains a unit root at the zero frequency, one can start directly with testing for such a unit root. A drawback of this approach is that, in case of rejection one obtains no insight into the possible presence of stochastic trends. Of course, such a result may complicate the construction of multivariate models since this typically involves aspects as cointegration and common stochastic trends.

The two model selection strategies and the forecasting performance of both models have been compared using both Monte Carlo simulations and empirical time series. The main recommendation is that we strongly suggest not to use the periodic ADF test due to its low power against periodic integration. The LR_τ two-step strategy seems to opt less frequently for the $(1 - B)$ filter in case the DGP is PIAR than the W_{i4} test. The theoretical disadvantage of the W_{i4} test is the lack of information on the (stochastic) trend in case the presence of the $(1 - B)$ filter is rejected. Besides, the LR_τ two-step strategy gives opportunities for testing for the significance of deterministic components

after the hypothesis of a periodic stochastic trend cannot be rejected, which can give some increase in power. For the forecasting performance of the models it seems that forecasting with a periodic model in first differences when a periodically integrated model is appropriate is worse than *vice versa*. This is especially the case when (one of) the parameters in the periodic differencing filter differs substantially from unity.

References

- Boswijk, H.P. & Franses, P.H. (1994) Unit Roots in Periodic Autoregressions, Tinbergen Institute discussion paper TI 94-4, Rotterdam.
- Box, G.E.P. & Jenkins, G.M. (1970) *Time Series Analysis, Forecasting and Control*, San Francisco: Holden-Day.
- Davidson R. & MacKinnon, J.G. (1993) *Estimation and Inference in Econometrics*, New York: Oxford University Press.
- Franses, P.H. (1994) A Multivariate Approach to Modelling Univariate Seasonal Time Series, *Journal of Econometrics*, 63, 133-152.
- Franses, P.H. & Paap, R. (1994) Model Selection in Periodic Autoregressions, *Oxford Bulletin of Economics and Statistics*, 56, 421-439.
- Fuller, W.A. (1976) *Introduction to Statistical Time Series*, New York: Wiley.
- Ghysels, E. & Hall, A. (1993) On Periodic Time Series and Testing the Unit Root Hypothesis, Discussion paper 2693, C.R.D.E., University of Montreal.
- Gladyshev, E.G. (1961) Periodically Correlated Random Sequences, *Soviet Mathematics*, 2, 385-388.
- Hylleberg, S., Engle, R.F., Granger, C.W.J. & Yoo, B. (1990) Seasonal Integration and Cointegration, *Journal of Econometrics*, 44, 215-238.
- Lütkepohl, H. (1991) *Introduction to Multiple Time Series Analysis*, Berlin: Springer-Verlag.
- McLeod, A.I. (1993) Parsimony, Model Adequacy and Periodic Correlation in Time Series Forecasting, *International Statistical Review*, 61, 387-393.
- Osborn, D.R. (1988) Seasonality and Habit Persistence in a Life-Cycle Model of Consumption, *Journal of Applied Econometrics*, 3, 255-266.
- Osborn, D.R. (1990) A Survey of Seasonality in UK Macroeconomic Variables, *International Journal of Forecasting*, 6, 327-336.
- Osborn, D.R. & Smith, J.P. (1989) The Performance of Periodic Autoregressive Models in Forecasting Seasonal U.K. Consumption, *Journal of Business and Economic Statistics*, 7, 117-127.
- Vecchia, R.L. & Ballerini, R. (1991) Testing for Periodic Autocorrelations in Seasonal Time Series Data, *Biometrika*, 78, 53-63.

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

Richard Paap
Tinbergen Institute
Erasmus University Rotterdam
Oostmaaslaan 950-952
NL-3063 DM Rotterdam
The Netherlands