

A Bayesian Analysis of Periodic Integration

PHILIP HANS FRANSES¹ AND GARY KOOP^{2*}

¹*Erasmus University Rotterdam, The Netherlands*

²*University of Edinburgh, UK*

ABSTRACT

Recent empirical research into the seasonal and trend properties of macro-economic time series using periodic models has resulted in strong evidence in favour of periodic integration (PI). PI implies that the differencing filter necessary to remove a stochastic trend varies across seasons and, hence, that seasonal fluctuations are related to the stochastic trend. Previous studies finding evidence of PI have used classical econometric techniques. In this paper, we investigate the possible sensitivity of this empirical result by using Bayesian techniques. An application of posterior odds analysis and highest posterior density interval tests to several quarterly UK macro-economic series suggests strong evidence for PI, even when we allow for structural breaks in the deterministic seasonals. A predictive exercise indicates that PI usually outperforms other competing models in terms of out-of-sample forecasting. © 1997 John Wiley & Sons, Ltd.

J. forecast., Vol. 16, 509–532 (1997)

No. of Figures: 2. No. of Tables: 6. No. of References: 31.

INTRODUCTION AND MOTIVATION

Periodic autoregressive (PAR) models have proved to be useful in modelling the seasonal and stochastic trend behaviour of time series (see e.g. Birchenhall *et al.*, 1989; Osborn, 1988; Franses, 1994; Franses, 1996b, *inter alia*). PAR models extend conventional autoregressive (AR) models by allowing the AR parameters to take different values across seasons. This extension yields the related concept of periodic integration (PI). Intuitively, in order to remove the stochastic trend in PI models, a differencing filter that varies with the seasons must be used. In contrast, the periodic unit root hypothesis has a first differencing filter applied to all seasons. Boswijk and Franses (1996) proposes a formal test for periodic integration of order one. An application of this test to a

* Correspondence to: Gary Koop, Department of Economics, University of Edinburgh, William Robertson Building, 50 George Square, Edinburgh EH8 9JY, UK.

Contract grant sponsors: The Royal Netherlands Academy of Arts and Sciences, the Social Sciences and Humanities Research Council of Canada.

variety of UK macroeconomic series yields widespread evidence of PI (Franses and Paap, 1994). In fact, other plausible specifications tend to receive very little support relative to PI. The purpose of this paper is twofold: (1) to develop Bayesian methods for investigating the periodic integration hypothesis; and (2) to apply these methods to a wide variety of UK series to investigate whether the rather strong empirical evidence of Franses and Paap (1994) is robust to the adoption of a different methodological perspective (and also to the inclusion of structural breaks).

Before introducing the PAR model, it is useful to motivate interest in such models. The last few years have witnessed a growing interest in analysing quarterly observed macroeconomic quantities in raw format (i.e. not seasonally adjusted). A key motivation for this practice is that economists have come to recognize that seasonal variation seems to be worth while to study in its own right. Specifically, it may contain relevant information that helps us to explain the behaviour of economic agents (see e.g. Ghysels, 1994; Miron, 1996, *inter alia*). A casual look at (the graphs or the estimated autocorrelation functions of) many quarterly observed macroeconomic time series readily reveals that seasonal variation is often quite large. In fact, it seems that for detrended time series, seasonal variation amounts to the dominant source of variation. The results of regressions of growth rates of such variables on four seasonal dummy variables in e.g. Barsky and Miron (1989) and Miron (1996), suggest that over 80% of the variation in a detrended series usually comes from seasonality. This large amount has motivated economic theorists to revise economic models like Real Business Cycle models, as, for example, is done in Braun and Evans (1995) and Chatterjee and Ravikumar (1992). In these theoretical models, seasonality is usually represented as a set of deterministic components which take on different values in different quarters.

In addition to the observation that seasonal variation is substantial in many macroeconomic aggregates, several empirical studies have shown that seasonality is not constant over time (see e.g. Hylleberg, 1994). That is, the amount of seasonal variation may well be still over 80%, but the relative importance of the observations in each quarter may have changed over time. For example, it may be that third quarter observations were highest in the 1960s, while in the 1980s it is the fourth quarter that peaks. Obviously, such changing seasonality cannot be captured by time-series models with constant parameters (e.g. for the model with seasonal dummies). In fact, below we will find that constant parameter models will indeed be outperformed by models with more flexible structures. A class of models that is often used to describe changing seasonality are the seasonal unit root models (see, e.g. Hylleberg *et al.*, 1990; Harvey and Scott, 1994). Such seasonal unit roots correspond with stochastic trends at the seasonal frequencies. Strictly speaking, such unit roots allow for wildly fluctuating seasonality. Given the nature of a seasonal unit root test, it may well be that any changing seasonality is not fluctuating enough in order for one to find seasonal unit roots. Hence, even when seasonality is changing, one may not find that it can be characterized by seasonal unit roots (see, e.g. the findings in Beaulieu and Miron, 1993). However, when one reverses the null hypothesis (i.e. the unit root is no longer the null, but rather the alternative) as is done in e.g. Canova and Hansen (1995), much more evidence of changing seasonality can be found.

A third (and more recently established) feature of many macroeconomic variables is that seasonal variation and nonseasonal variation (in the form of trends and cycles) are often not independent (see Canova and Ghysels, 1994; Franses, 1995, 1996a, *inter alia*). For example, the first two studies show that seasonal variation differs across the business cycle. In Franses (1996a) it is shown through an analysis of consumer confidence indices (which are based on presumably detrended situations) that consumers tend to face difficulties in disentangling trends from

seasonals. Notice that, when seasonal variation and non-seasonal variation cannot be separated, the key assumption underlying seasonal adjustment methods is violated. Indeed, in cases where separation is not possible, one may find traces of seasonality in the presumably adjusted time series. Ooms and Franses (1996) show, for two unemployment series, that there are strong correlations between their estimated seasonal and non-seasonal components. These correlations are in fact so strong that one can make quite precise forecasts of the seasonal component in the next quarter using the 'adjusted' series in the current quarter. Additionally, these correlations imply that the adjusted series can be too optimistic in some seasons and too pessimistic in others. This suggests that, even when one claims only to be interested in forecasting seasonality-free data, one should take into account that the links between seasonal and non-seasonal variation may severely bias forecasts and conclusions.

In Franses (1996b) it is argued that these three features of raw quarterly observed macroeconomic time series (i.e. substantial though changing seasonal variation with possible links with cycles and trends) can be captured by so-called periodic integration models in the univariate case and periodic cointegration models in the multivariate case. In brief, periodic integration is defined by the requirement of a seasonally varying differencing filter to remove the stochastic trend. Osborn (1988) shows that simple periodic integration models may correspond with the permanent income–life cycle hypothesis for consumption once one allows for seasonally varying utility. For the multivariate case, Birchenhall *et al.* (1988) and Franses (1996b) show that periodic cointegration corresponds with the common linear quadratic adjustment costs model with the modification that targets and adjustment costs are allowed to vary seasonally (see also Todd, 1990; Hansen and Sargent, 1995). In particular, modifications to the model in Todd (1990) (to allow for trend) can lead to theoretical models within which non-seasonal shocks may change the behaviour of agents in some seasons though not in others. Simulation results in Franses (1996b) and Ooms and Franses (1996) among others show that, when one neglects periodic parameter variation in the periodically integrated autoregressions (PIAR), one may or may not find seasonal unit roots and may face under- or overadjustment of the 'adjusted' time series. Furthermore, the recent results in Ghysels, Hall, and Lee (1996) and Boswijk, Franses, and Haldrup (1996) show that once one allows for periodicity any evidence for seasonal unit roots tends to disappear. To summarize, the empirical results in Franses and Paap (1994) and Franses (1996b) suggest rather strong empirical evidence of periodic integration in many univariate macroeconomic time series.

The present paper uses a Bayesian approach which offers many different ways to compare models or to test hypotheses. The most common test procedure is to use posterior odds. Posterior odds are essentially weighted likelihood ratios where the weights are given by the prior distribution. As a prior becomes flat, more and more weight is placed in regions of the parameter space where the likelihood is negligible. Hence when working with exact restrictions (such as that implied by the periodic integration hypothesis, as will be shown in the next section), the prior distribution must be informative and proper (i.e. integrate to one). Improper priors typically lead to posterior odds that always favour the restricted model (see Leamer, 1978, pages 110–14 for details). To avoid this problem, we use informative natural conjugate priors for our posterior odds analysis. Given the controversy surrounding the use of informative priors we carry out a sensitivity analysis with respect to the prior hyperparameters. Furthermore, we consider other methods for shedding light on the periodic integration hypothesis which do not require proper priors. In particular, we calculate highest posterior density intervals for the unrestricted periodic model and see whether they contain the PI and periodic unit root models.

Given the criticism often levelled (by both Bayesians and non-Bayesians) at standard hypothesis testing in econometrics, we feel that it is important to supplement conventional test procedures with a predictive analysis. For this reason we compare the predictive power of our models relative to some withheld data.

The paper is organized as follows. After a brief introduction to periodic models is given in the next section and the third section describes the Bayesian methods used in this paper. The fourth section applies these methods to several different quarterly UK series. The final section presents conclusions.

PERIODIC AUTOREGRESSIVE MODELS

The periodic autoregressive model of order p (PAR(p)) for a quarterly univariate time series, y_t ($t = 1, \dots, n$), is given by

$$y_t = \mu_s + \phi_{1s}y_{t-1} + \dots + \phi_{ps}y_{t-p} + \varepsilon_t \quad (1)$$

The ε_t 's are assumed to be Normal white noise.¹ Note that p is the maximum length of the AR orders in each season (i.e. $p = \max_s \{p_s\}$). For detailed accounts of properties of, and estimation methods for, PAR models, the reader is referred to Pagano (1978) and Tiao and Grupe (1980) *inter alia*. Boswijk and Franses (1996) give a detailed discussion of stochastic trends and model selection.

The PAR(p) process can also be written in differenced form:

$$y_t = \mu_s + \phi_s y_{t-1} + \psi_{1s}(y_{t-1} - \phi_{s-1}y_{t-2}) + \dots + \psi_{p-1,s}(y_{t-(p-1)} - \phi_{s-(p-1)}y_{t-p}) + \varepsilon_t \quad (2)$$

where $\phi_{s-4k} = \phi_s$ for $k = 0, 1, 2, \dots$ and $s = 1, 2, 3, 4$. Note that the ψ_{js} 's and ϕ_s 's are non-linear functions of the ϕ_{is} 's in equation (1). The series y_t has a stochastic trend when the non-linear restriction $\phi_1\phi_2\phi_3\phi_4 = 1$ holds. When the ϕ_s 's are not all equal to one, y_t is called periodically integrated. In fact, Boswijk and Franses (1996) shows that, given $\phi_1\phi_2\phi_3\phi_4 = 1$, the multivariate model for the 4×1 vector series $Y_T = (Y_{1T}, Y_{2T}, Y_{3T}, Y_{4T})'$, where Y_{sT} is the observation for season s in year T , has a single unit root. If this restriction holds, we call equation (2) a PIAR(p) process. In the case where $\phi_s = 1$ for all s , the periodic integration model reduces to:

$$y_t - y_{t-1} = \mu_s + \psi_{1s}(y_{t-1} - y_{t-2}) + \dots + \psi_{p-1,s}(y_{t-(p-1)} - y_{t-p}) + \varepsilon_t \quad (3)$$

which we refer to as the periodic unit root model. The main difference between the PI process and the periodic unit root process given in equation (3) is that, for the PI model, the impact of the stochastic trend varies with the season since the appropriate differencing filter, $(1 - \phi_s B)$, depends on s .² In other words, the 'seasonal pattern' can change because of the accumulation of shocks, and shocks in different seasons have different impacts on the future pattern of the time series.

Boswijk and Franses (1996) propose a sequential model selection strategy to investigate trend and seasonal properties of univariate time series using models like equation (1). The first step

¹ One may want to extend equation (1) by allowing for seasonal heteroscedasticity. For simplicity, this extension is excluded here, but could be incorporated in a relatively straightforward manner.

² B is the familiar backshift or lag operator.

involves using a Likelihood Ratio-based test statistic for the non-linear restriction $\phi_1\phi_2\phi_3\phi_4 = 1$ (which follows a Dickey–Fuller type distribution under the null hypothesis). If the null hypothesis of $\phi_1\phi_2\phi_3\phi_4 = 1$ cannot be rejected, the second step is a test for the restrictions $\phi_s = 1$ for $s = 1, \dots, 4$. The relevant test statistic follows a $\chi^2(3)$ distribution.

Franses and Paap (1994) apply this model selection strategy to a large set of UK macro-economic time series, and find widespread evidence in favour of $\phi_1\phi_2\phi_3\phi_4 = 1$. Moreover, even though ϕ_s typically lies near 1 (i.e. values of ϕ_s in the interval [0.90, 1.10] are usual) the restriction $\phi_1 = \dots = \phi_4 = 1$ is invariably massively rejected (typical p -values for this test do not exceed 0.0001).

In the present paper, we employ Bayesian techniques to investigate whether the rather strong results in Franses and Paap (1994) are robust to the assumed statistical methodology and structural breaks. Although many hypotheses involving PAR models can be analysed,³ we focus our attention on:

- H₁: ϕ_1, ϕ_2, ϕ_3 and ϕ_4 unrestricted
- H₂: $\phi_1\phi_2\phi_3\phi_4 = 1$
- H₃: $\phi_1 = \phi_2 = \phi_3 = \phi_4 = 1$

Note that H₃ is nested within H₂. Given that Franses and Paap (1994), using standard order selection criteria, found that $p = 1, 2,$ or 3 for the series considered, we use $p = 1$ and 4 . In practice, our results for these two choices are generally similar.

As noted previously, the periodic integration model allows for the stochastic trend to have a varying impact on the series (i.e. an accumulation of shocks can change the seasonal pattern). Changing seasonal patterns can also be generated by a model which has a structural break in the seasonal dummies. Hence, for series which appear to be periodically integrated, we also consider specifications with such a break. That is, we analyse H₁, H₂, and H₃ in light of the model:

$$y_t = \mu_s + \delta_s D_{t > \tau} + \phi_s y_{t-1} + \psi_{1s}(y_{t-1} - \phi_{s-1} y_{t-2}) + \dots + \psi_{p-1,s}(y_{t-(p-1)} - \phi_{s-(p-1)} y_{t-p}) + \varepsilon_t \tag{4}$$

where $D_{t > \tau} = 1$ for $t > \tau$ and $= 0$ otherwise. Given that we do not know the exact breakpoint, τ , we perform our analysis over all possible breakpoints. Within this context, the Bayesian method facilitates empirical inference.

BAYESIAN METHODS

In this section we propose Bayesian techniques to analyse the hypotheses H₁, H₂, H₃. The periodic integration restriction is non-linear. In order to reduce the computational burden we sometimes use an approximation to ensure linearity. A full Bayesian analysis which explicitly incorporated the non-linearity would be possible, but would require computationally intensive techniques such as Gibbs sampling. Since we wish to consider the properties of many series and to carry out a careful prior sensitivity analysis for each series, it is necessary to avoid such computationally burdensome techniques.

³ Given the overwhelming evidence against standard linear AR models with seasonal dummies, we do not consider such models.

The periodic integration restriction, $\phi_1\phi_2\phi_3\phi_4 = 1$, can be linearized by taking logs and noting that $\log(1+x)$ is approximately x for small x . In practice, the ϕ_i 's are near one so that the approximation is a good one. Under this approximation, the periodic integration hypothesis involves the linear restriction $\phi_1 + \phi_2 + \phi_3 + \phi_4 = 4$. For the PIAR(1) model this approximation is adequate to ensure linearity.

The PAR(p) model is non-linear in the parameterization we have used for $p > 1$ (see equation (2)). This parameterization is convenient since it ensures that the periodic unit root restriction remains $\phi_1\phi_2\phi_3\phi_4 = 1$ for all p . To eliminate this non-linearity we replace the ψ_{js} 's, which are nuisance parameters, by maximum likelihood estimates. This amounts to working with the conditional density $p(\phi, \sigma^2 | \text{Data}, \psi = \psi^*)$ instead of $\int p(\phi, \sigma^2, \psi | \text{Data}) d\psi$, where $\phi = (\phi_1, \phi_2, \phi_3, \phi_4)'$. In practice, the former can be expected to be a good approximation to the latter. Asymptotically, these two densities will be identical.

Posterior odds analysis

If we proceed conditionally on p initial observations, then the linear time-series models can be treated as standard linear regression models. Assuming a natural conjugate prior,⁴ posterior odds can be calculated analytically using the formula, for instance, in Judge *et al.* (1985), page 129. Throughout this paper, we adopt natural conjugate or flat priors. As noted, for posterior odds informative priors are required, so the flat prior is not suitable.

Let $y = (y_1, \dots, y_T)'$ and X_j be the appropriate $T \times K_j$ matrix containing all the explanatory variables under H_j ($j = 1, 2, 3$). Then the three models can be written as $y = X_j\beta_j + \varepsilon_j$ where ε_j is $N(0, \sigma_j^2 I_T)$. The natural conjugate prior is the Normal-gamma, i.e. we assume $p(\beta_j, \sigma_j^{-2}) = f_N(\beta_j | \beta_{0j}, \sigma_j^2 \Omega_j) f_G(\sigma_j^{-2} | s_{0j}^{-2}, \nu_{0j})$, where $f_N(\cdot | a, B)$ is the multivariate Normal distribution with mean a and covariance matrix B and $f_G(\cdot | c, d)$ denotes the Gamma distribution with mean c and variance $2c^2/d$. The prior hyperparameters β_{0j} , Ω_j , s_{0j}^{-2} and ν_{0j} must be elicited. Throughout this paper we choose hyperparameters centred over the periodic unit root model (e.g. we set $\beta_{01} = (0, 0, 0, 0, 1, 1, 1, 1)'$). In doing so, we stack the odds, so to speak, in favour of the periodic unit root model H_3 .

To select the other prior hyperparameters we begin by eliciting a prior under H_1 . Since we have little prior information about σ_1^{-2} , prior degrees of freedom are set to 3 (which is the lowest possible value consistent with the existence of the marginal prior mean and covariance of β_1), and $s_{01}^{-2} = 10,000$. In other words, the prior mean of the error precision is 10,000 and its standard deviation is around 5000. Loosely speaking, we are saying we expect the standard deviation of the error to be around 0.01, but are very uncertain about the precise value. The matrix of prior hyperparameters, Ω_{01} , can be elicited by noting that the marginal prior covariance matrix for β_1 is $\nu_{01} s_{01}^2 \Omega_{01} / (\nu_{01} - 2)$. To aid in our prior sensitivity analysis we set

$$\Omega_{01} = \begin{pmatrix} \delta_1 I & 0 \\ 0 & \delta_2 I \end{pmatrix}$$

where δ_1 and δ_2 are scalars. That is, we assume the regression parameters are, *a priori*, independent of each other but allow for the seasonal dummy coefficients to have a different prior standard

⁴ Phillips (1991) suggests using Jeffreys' prior for time-series models. However, this prior is improper and has several undesirable properties (see the discussion following Phillips, 1991). This prior and other candidate priors all have the drawback that analytical results for posterior odds are not available and hence they require the use of computationally intensive methods such as Importance Sampling. We do not make use of these priors here.

deviation than the periodic autoregressive parameters. We focus on prior sensitivity with respect to δ_2 . In particular, we let δ_2 take values between 33.333 and 3333.333, values which imply prior standard deviations between 0.1 and 1. We set $\delta_1 = 3333.333$ (i.e. prior standard deviation equal to one) for the seasonal dummy coefficients to reflect the relative lack of prior knowledge of these parameters, which are common to all models. There are, of course, other directions in which we could carry out a prior sensitivity analysis, but the spread of the prior is a key aspect we wish to focus on.

In theory, we can use any prior we want for H_2 and H_3 . However, it makes sense to select priors which are consistent with that elicited under H_1 . In particular, we use priors of the form $p(\beta_2, \sigma_2^{-2}) = p(\beta_1, \sigma_1^{-2} | \phi_1 + \phi_2 + \phi_3 + \phi_4 = 4)$ and $p(\beta_3, \sigma_3^{-2}) = p(\beta_1, \sigma_1^{-2} | \phi_1 = \phi_2 = \phi_3 = \phi_4 = 1)$. Since the restrictions are both linear, priors for all models are Normal-Gamma (see Raiffa and Schlaifer, 1961 or Poirier, 1985 for details).

Highest posterior density intervals

Since the use of informative priors is somewhat controversial, we also construct Highest Posterior Density Intervals (HPDIs) using the unrestricted model, H_1 , and the standard flat prior: $p(\beta_1, \sigma_1) \propto \sigma_1^{-1}$. The marginal posterior for β_1 takes the form of a multivariate Student-*t* distribution, so HPDIs can easily be calculated. We can then check whether the restriction implied by H_2 lies in the relevant HPDIs. H_3 is a test of a joint hypothesis, so HPDIs will be multi-dimensional. To surmount this problem, note that, since v_1 is typically quite large, the posterior for β_1 is approximately Normal, as is $R_1\beta_1$, where R_1 is a $4 \times K_1$ matrix and the unit root restriction is written as $R_1\beta_1 = (1 \ 1 \ 1 \ 1)'$. Taking the inner product of standardized Normals, we obtain a $\chi^2(4)$ distribution which can be used for the joint test that $\phi_1 = \phi_2 = \phi_3 = \phi_4 = 1$.

We present ‘Bayesian *p*-values’ based on these HPDIs. For instance, for PI we calculate $p(|\phi_1 + \phi_2 + \phi_3 + \phi_4 - M| < 4 - M | \text{Data})$, where M is the posterior mean of $\phi_1 + \phi_2 + \phi_3 + \phi_4$.⁵

Predictive distributions

In this paper, we consider the predictive performance of our three models by withholding ten years of data.⁶ That is, we work with $p(y_{T+n}|y, H_1)$, $p(y_{T+n}|y, H_2)$, and $p(y_{T+n}|y, H_3)$. Note, however, that while analytical expressions for these densities do not exist it is easy to calculate their properties using simple Monte Carlo integration. These predictive densities may be of interest in and of themselves (e.g. whether they vary significantly over the seasons). As to more formal investigation of the periodic integration hypothesis, we calculate the following predictive mean-square error criteria:

$$PMSE(n) = \sum_{i=1}^n E(y_{T+i} - y_{T+i}^*)^2$$

where y_{T+i}^* is the actual realization of the random variable y_{T+i} .

⁵ This can be thought of as a two-sided test. For the joint test of H_3 we present one-sided Bayesian *p*-values.

⁶ Note that our posterior odds and HPDI analysis use the complete data set.

EMPIRICAL RESULTS

In this section we apply our Bayesian methods to nine major UK macroeconomic time series. All data are in logs and observed from 1955:1 to 1988:4 (except for the stock price series which starts in 1963:1: details are given in Table I). These data have been analysed using the periodic model in Franses and Paap (1994) and Franses (1996b). We begin by presenting a few basic properties of the periodic model, following which we present results from HPDI tests, posterior odds, and the prediction exercise.

Table I. Variable notation and definitions (all data are quarterly from 1955:1 to 1988:4, except stock prices which run from 1963:1 to 1988:4)

-
1. *GDP* = GDP at 1985 prices
 2. *TOTCONS* = total consumption at 1985 prices
 3. *CONSDUR* = consumption of durables at 1985 prices
 4. *NONDUR* = consumption = of non-durables at 1985 prices
 5. *TOTINV* = total investment at 1985 prices
 6. *EXPORTS* = exports at 1985 prices
 7. *IMPORTS* = imports at 1985 prices
 8. *PRICES* = index of retail prices (1985 = 100)
 9. *STOCKP* = FT index stock prices (1 July 1985 = 100)
-

Properties of periodic model


Equation (2) presents the unrestricted PAR model. Under a flat prior and the assumptions given above, the posterior distribution of ϕ is multivariate- t . Since it is simple to take random draws from this distribution, it is easy to simulate the properties of any function of ϕ . Accordingly, we use Monte Carlo integration⁷ to plot the distribution of $\phi_1\phi_2\phi_3\phi_4$. If the value 1 lies in a region that receives appreciable posterior probability, then this is informal evidence in favour of PI. Figures 1(a)–1(i) plot these posteriors for the PAR(4) model for the nine series. Results for PAR(1) are qualitatively similar and are omitted for the sake of brevity. The graphs also include the marginal posteriors for the individual ϕ_i 's. Note that this is the only part of the empirical work where we do not rely on the $\phi_1 + \phi_2 + \phi_3 + \phi_4 = 4$ approximation for the PI model.

An examination of the figures reveals widespread evidence of PI. All the series yield posteriors for $\phi_1\phi_2\phi_3\phi_4$ which include one in regions of non-negligible probability. Furthermore, with the exception of prices and stock prices (and possibly imports), there is little evidence for the periodic unit root model, H_3 . That is, at least one of the ϕ_i 's tends to have a distribution which is located far from one. A common pattern in several of the series (especially the major macroeconomic quantity aggregates—GDP, components of consumption and investment) is that dynamics in the second quarter seem stationary (ϕ_2 around 0.92) while the third quarter exhibits explosive behaviour (ϕ_3 around 1.05).

Overall, we draw the following informal conclusions from these figures:

- (1) The variables which measure quantities (GDP, consumption, consumption of durables, consumption of non-durables, investment, exports, and imports) all support H_2 , while the variables which measure prices support H_3 .
- (2) With the exception of prices and stock prices, all the series seem to have dynamics which differ across the seasons.


⁷ For all the Monte Carlo integrations carried out in this paper we used 10,000 replications.



(a)

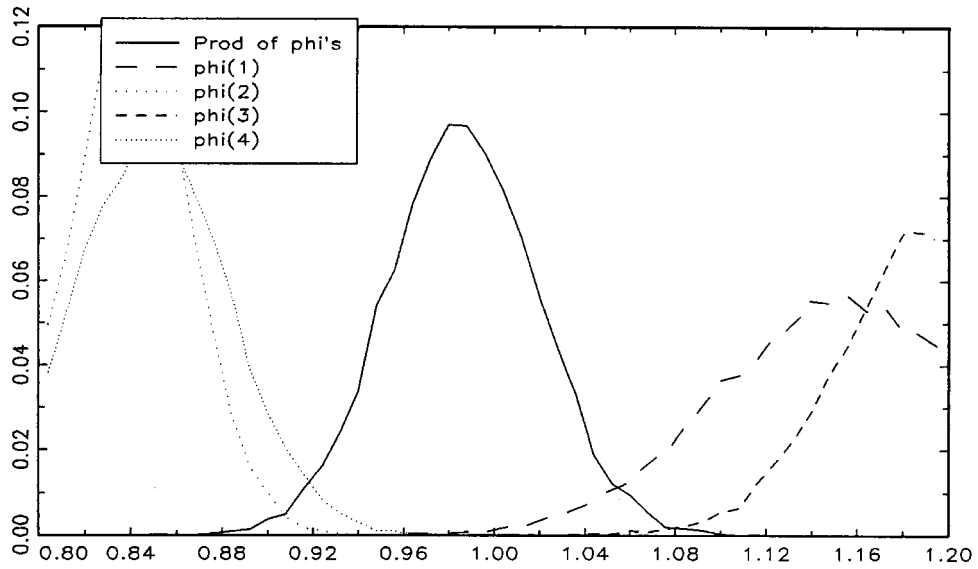
Figure 1. Posterior pdfs based on PAR(4)

Figure 1(a) GDP




(b)

Figure 1(b) total consumption




(c)

Figure 1(c) consumption of durables




(d)

Figure 1(d) consumption of non-durables




(e)

Figure 1(e) total investment




(f)

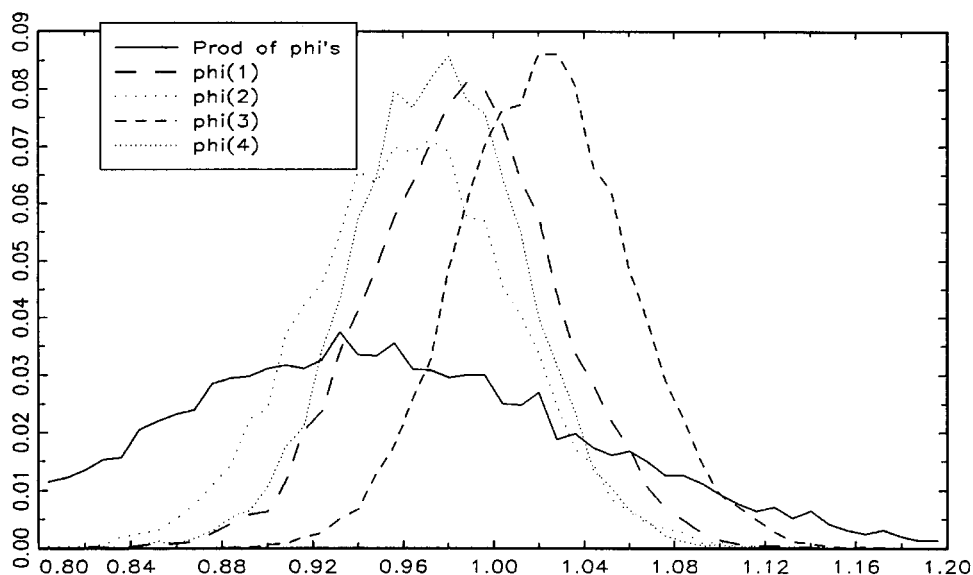
Figure 1(f) exports



(g)
Figure 1(g) imports



(h)
Figure 1(h) prices



(i)

Figure 1(i) stock prices

Highest posterior density interval tests

Table II presents Bayesian p -values based on highest posterior density intervals as described in the previous section. This table strongly reinforces the informal conclusions drawn in that section, which is not surprising since the results are based on the same flat prior. That is, there is strong evidence of PI for all of the quantity series for both the PAR(1) and PAR(4) models. Furthermore, although some individual ϕ_i 's are near one, the hypothesis that $\phi_i = 1$ for all i can be safely rejected for all the quantity series. There is much more evidence of a periodic unit root for the stock price series. The only possible discrepancy between the results in Figures 1(a)–1(i) and those in Table II is for prices. A visual inspection of Figure 1(h) indicates that H_3 has significant support. However, the joint test that $\phi_i = 1$ for all i receives little support. This is undoubtedly largely due to the joint nature of the test, since each of the individual ϕ_i 's seems to be near 1. Nevertheless, we will pay extra attention to this variable in the next empirical investigation.

An informal measure of fit

Table III presents results for an informal measure of fit: how well does a model explain the growth rate of the series being considered? Using the flat prior of the previous section, we calculate the following 'Bayesian R^2 ':

$$R^2 = 1 - \frac{E(\sigma^2 | \text{Data})}{\text{var}(\Delta y_t)}$$

where $\sigma^2 = \text{var}(\varepsilon_t)$. As its name implies, it has an interpretation analogous to the standard R^2 .

Results in Table III are supportive of the general pattern found previously: as compared to the PIAR model, the PAR model adds little explanatory power in most cases. However, in many

Table II. Bayesian p -values based on PAR(1) and PAR(4)

	H_2	H_3	$\phi_1 = 1$	$\phi_2 = 1$	$\phi_3 = 1$	$\phi_4 = 1$
PAR(1)						
<i>GDP</i>	0.71	2.5E-8	0.77	2.6E-7	1.3E-3	0.43
<i>TOTCONS</i>	0.52	2.7E-19	1.4E-3	2.9E-12	5.1E-6	0.75
<i>CONSDUR</i>	0.85	1.6E-7	1.8E-4	1.3E-6	3.5E-5	1.3E-6
<i>NONDUR</i>	0.44	2.8E-19	0.57	4.6E-12	4.8E-4	8.2E-6
<i>TOTINV</i>	0.30	6.7E-11	0.04	1.4E-9	0.02	0.44
<i>EXPORTS</i>	0.61	3.3E-3	1.3E-3	0.10	0.21	0.48
<i>IMPORTS</i>	0.87	7.3E-3	0.05	0.03	0.12	0.12
<i>PRICES</i>	2.0E-3	0.01	0.17	6.6E-3	0.07	0.71
<i>STOCKP</i>	0.94	0.94	0.80	0.88	0.43	0.82
PAR(4)						
<i>GDP</i>	0.94	2.3E-21	0.59	8.2E-6	2.0E-6	0.41
<i>TOTCONS</i>	0.25	7.1E-22	1.8E-3	9.2E-13	4.4E-7	0.74
<i>CONSDUR</i>	0.12	6.1E-37	4.7E-3	1.2E-8	4.9E-6	8.4E-6
<i>NONDUR</i>	0.26	3.5E-20	0.62	1.7E-11	1.4E-4	1.5E-5
<i>TOTINV</i>	0.44	8.5E-11	0.04	5.5E-9	9.0E-3	0.58
<i>EXPORTS</i>	0.64	1.6E-4	4.5E-4	0.29	0.01	0.60
<i>IMPORTS</i>	0.62	5.7E-3	0.02	0.03	0.52	0.11
<i>PRICES</i>	0.34	3.2E-4	0.47	0.21	0.27	0.72
<i>STOCKP</i>	0.73	0.77	0.86	0.60	0.61	0.32

Table III. 'Bayesian R^2 ': $1 - E(\sigma^2|\text{Data})/\text{var}(\Delta y)$

	PAR(4)	PIAR(4)	UNIT(4)	PAR(1)	PIAR(1)	UNIT(1)
<i>GDP</i>	0.867	0.867	0.837	0.814	0.814	0.720
<i>TOTCONS</i>	0.940	0.939	0.917	0.932	0.932	0.879
<i>CONSDUR</i>	0.651	0.644	0.552	0.466	0.466	0.086
<i>NONDUR</i>	0.971	0.971	0.963	0.969	0.969	0.946
<i>TOTINV</i>	0.731	0.729	0.661	0.693	0.691	0.565
<i>EXPORTS</i>	0.520	0.520	0.461	0.354	0.352	0.262
<i>IMPORTS</i>	0.363	0.362	0.311	0.252	0.252	0.190
<i>PRICES</i>	0.670	0.665	0.628	0.203	0.137	0.122
<i>STOCKP</i>	0.306	0.305	0.285	0.112	0.112	0.103

cases the PIAR model adds significant explanatory power relative to the periodic unit root model (labelled UNIT in Table III). In most (but not all cases), adding extra lags provides significant improvements in the Bayesian R^2 .

The last column in Table III (labelled UNIT(1)) is of particular interest. This provides a useful benchmark since the periodic unit root model with $p = 1$ reduces to a model where the seasonal variation in the growth rate of the series is explained solely by seasonal dummies as is the case in e.g. Barsky and Miron (1989). Hence, this column measures the role of purely deterministic seasonality. In some cases (*TOTCONS* and *CONSDUR*) seasonal dummies can explain a great deal, but in most cases (especially *CONSDUR*, *PRICES* and *STOCKP*) the addition of stochastic periodicity greatly improves the fit. For a key series like *GDP*, going from the purely deterministic model to the slightly less restrictive PIAR(1) model causes the Bayesian R^2 to rise from 0.720 to 0.814.

Table IV. Posterior model probabilities for H_2 (H_1 in parentheses) corresponding to different prior standard deviations under H_1 in PAR models of order p

	$p = 1$			$p = 4$		
Prior st. dev.	0.10	0.20	1.00	0.10	0.20	1.00
<i>GDP</i>	0.81 (0.20)	0.89 (0.11)	0.98 (0.02)	1.00 (0.00)	1.00 (0.00)	0.96 (0.00)
<i>TOTCONS</i>	0.83 (0.17)	0.90 (0.10)	0.98 (0.02)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)
<i>CONSDUR</i>	0.65 (0.35)	0.79 (0.21)	0.95 (0.05)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)
<i>NONDUR</i>	0.85 (0.15)	0.92 (0.08)	0.98 (0.02)	1.00 (0.00)	1.00 (0.00)	0.97 (0.00)
<i>TOTINV</i>	0.64 (0.36)	0.78 (0.22)	0.95 (0.05)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)
<i>EXPORTS</i>	0.77 (0.23)	0.83 (0.12)	0.13 (0.00)	1.00 (0.00)	0.98 (0.00)	0.32 (0.00)
<i>IMPORTS</i>	0.79 (0.20)	0.80 (0.10)	0.06 (0.00)	0.83 (0.06)	0.47 (0.02)	0.01 (0.00)
<i>PRICES</i>	0.00 (0.01)	0.00 (0.01)	0.00 (0.00)	0.03 (0.00)	0.00 (0.00)	0.00 (0.00)
<i>STOCKP</i>	0.33 (0.19)	0.08 (0.02)	0.00 (0.00)	0.40 (0.23)	0.12 (0.00)	0.00 (0.00)

Posterior odds

Table IV presents posterior model probabilities based on posterior odds (assuming equal prior odds) for different values of δ_2 , for $p = 1$ and 4.⁸ In particular, we choose δ_2 such that, under H_1 , the prior standard deviation takes on one of three values: 0.1, 0.2 and 1. We believe that a prior standard deviation of 0.2 is most reasonable; nevertheless we consider these other values in order to investigate prior sensitivity of our results. Priors under H_2 and H_3 are induced from the prior under H_1 as described in the previous section.

Results for the first five series can be dealt with quickly: there is overwhelming evidence for PI. Exports and imports require a bit more care since the periodic unit root model (H_3) gets significant posterior probability when the prior standard deviation under H_1 is 1. Our tendency is to dismiss this finding on the grounds that the resulting prior is very flat and, as indicated before, tends to favour more restricted models such as H_3 .

Posterior odds also reinforce our informal intuition that stock prices and prices can be described by a periodic unit root model. In fact, unlike with the HPDI tests, there is overwhelming evidence for this hypothesis.⁹ This finding reinforces our initial mistrust of HPDI tests in high-dimensional problems.

Overall, on the basis of all our in-sample analysis, we would conclude that periodic integration occurs for all of our quantity variables, but not for our price variables. Prices and stock prices exhibit simple periodic unit root behaviour. Note that these findings are similar to those given in Franses and Paap (1994), who use classical methods and find periodic integration for all the series

⁸ We only present $p(H_1|Data)$ and $p(H_2|Data)$. $p(H_3|Data)$ is, of course, one minus the sum of the previous two.

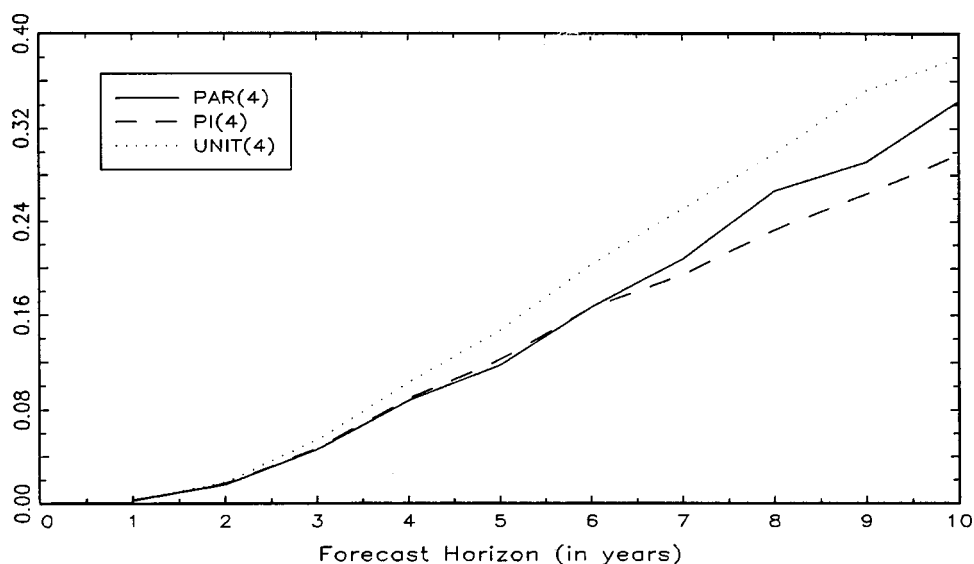
⁹ Although we do not consider this extension, it is possible that this model may be reduced to a non-periodic unit root model.

considered here except the price and stock price series, which they find to exhibit non-periodic unit root behaviour.

Predictive analysis


Figures 2(a)–2(i) contain predictive mean squared errors (PMSEs) for different forecast horizons as described in the previous section for models $p = 4$. Results for models with $p = 1$ tend to be qualitatively similar but provide more evidence in favour of the periodic integration hypothesis. Before discussing these figures in detail, it should be stressed here that predictive analysis is very different from the sort of in-sample measures considered previously. If the structure of the model is changing in a way that is not captured by the model, then a good in-sample fit will not guarantee good out-of-sample forecasts. Furthermore, results in this section use only a subset of the data used previously. (With the exception of stock prices, for the in-sample calculations $n = 136$. Here we use $n = 136 - 40 = 96$ observations and reserve 40 observations of hold-out data to compare with the predictions made using the 96 observations.) Note also, that we are using data through 1978:4 to predict 1979:1 to 1988:4. In other words, we are using pre-Thatcher data to try to predict behaviour during the Thatcher years.

With this in mind, consider the results in Figures 2(a)–2(i). The strongest impression is that the unrestricted PAR model does very poorly (except for the investment series), probably due to overparameterization. At short horizons, all the models have similar forecasting performance, but at medium to long horizons differences appear. For most of the series (i.e. GDP, consumption, consumption of non-durables, exports, and imports) the PI model does best or nearly best. However, the periodic unit root model does very well for most of the series. In many cases, it predicts just as well as the PI model, and in some cases (e.g. consumption of non-durables) it does better. One possible reason for this finding is that some sort of structural break occurred, a



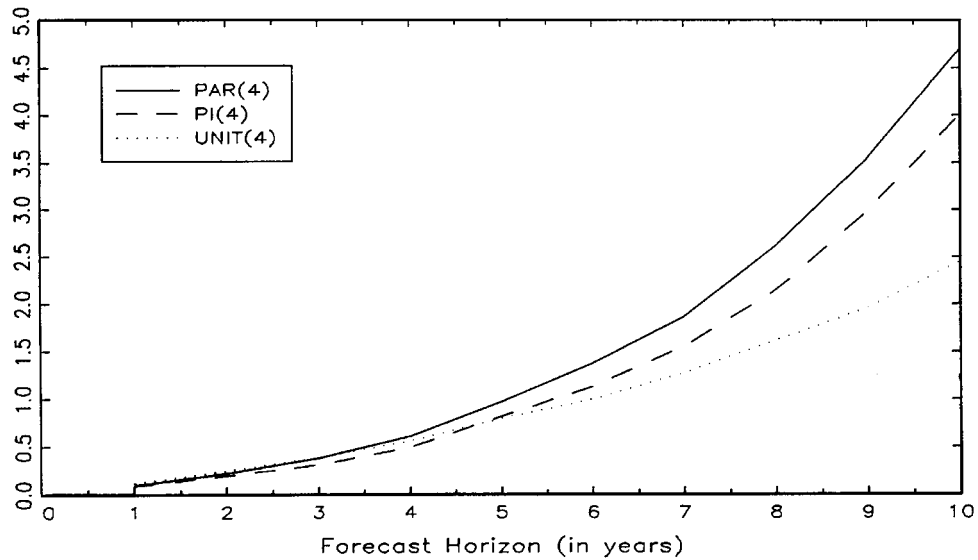
(a)

Figure 2. Predictive MSEs for (a) *GDP*




(b)

Figure 2(b) *TOTCONS*



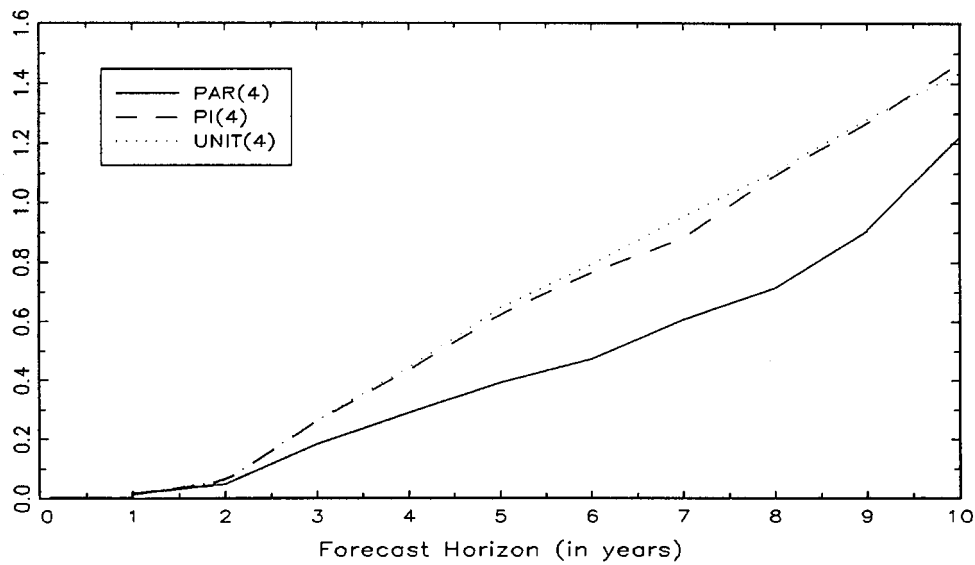
(c)

Figure 2(c) *CONSDUR*




(d)

Figure 2(d) *NONDUR*




(e)

Figure 2(e) *TOTINV*




(f)

Figure 2(f) EXPORTS




(g)

Figure 2(g) IMPORTS



(h)

Figure 2(h) *PRICES*



(i)

Figure 2(i) *STOCKP*

possibility which will be examined shortly. However, a more likely reason is the superior parsimony that the periodic unit root model has relative to the other models.

Structural break model

As discussed earlier, the structural break model given in equation (4) can possibly exhibit similar behaviour to the PI model. We are interested in whether we have mistakenly chosen PI when in reality equation (4) is the correct specification. This suggests the following strategy: choose the breakpoint which yields the least evidence of periodic integration. To keep things simple, we use a flat prior and report Bayesian p -values for PI as well as for the joint test that all seasons have no structural break (i.e. $\delta_s = 0$ in equation (4)). Furthermore, given the results of our predictive analysis, we focus on $p = 4$. Results are reported in Table V. Note that we now use the full sample again. For all series (except the two price series, for which there is little evidence of PI), allowing for a structural break in the seasonal dummies decreases the evidence for PI. However, there still is substantial evidence for PI for many series. Interestingly enough, with some exceptions, in the last column of Table V there is no overwhelming evidence for a structural break.

The strategy used in Table V is liable to the criticism of data mining. That is, we are deliberately looking over dozens of breakpoints for the one which will yield the strongest evidence against PI. An alternative strategy would be to find the breakpoint which maximizes the likelihood function, which is done in Table VI. In this table, there is much evidence for periodic integration—as well as substantial evidence of the presence of a structural break. The conclusion of this table is that a

Table V. Bayesian p -values for structural break which minimizes evidence in favour of periodic integration model

	Breakpoint	PI restriction	Struct. break
<i>GDP</i>	1973 : 4	0.14	0.13
<i>TOTCONS</i>	1973 : 2	0.02	0.03
<i>CONSDUR</i>	1960 : 3	0.11	0.25
<i>NONDUR</i>	1973 : 2	0.01	0.09
<i>TOTINV</i>	1982 : 2	0.05	0.03
<i>EXPORTS</i>	1968 : 1	0.00	0.00
<i>IMPORTS</i>	1974 : 1	0.12	0.06
<i>PRICES</i>	1981 : 3	0.02	0.10
<i>STOCKP</i>	1980 : 3	0.01	0.09

Table VI. Bayesian p -values for structural break which maximizes likelihood

	Breakpoint	PI restriction	Struct. break
<i>GDP</i>	1967 : 4	0.44	0.00
<i>TOTCONS</i>	1977 : 2	0.91	0.00
<i>CONSDUR</i>	1958 : 1	0.31	0.01
<i>NONDUR</i>	1977 : 2	0.81	0.00
<i>TOTINV</i>	1982 : 3	0.06	0.03
<i>EXPORTS</i>	1968 : 2	0.01	0.00
<i>IMPORTS</i>	1967 : 1	0.96	0.03
<i>PRICES</i>	1961 : 2	0.62	0.00
<i>STOCKP</i>	1973 : 2	0.17	0.05

structural break did occur (which could account for some of the predictive results noted before), but that periodic integration still is a plausible hypothesis for most of the series.

The likelihood function tends to be a fairly flat and multimodal function of the breakpoint, so it is hazardous to discuss its exact timing. Nevertheless, note that the timing of the breakpoint seems (with a few exceptions) to be clustered around three points in time: (1) the increases in VAT which occurred in 1968 and 1973; (2) the oil shock of 1973 : 4; and (3) the recession in the early years of the Thatcher era. We take this as mildly supportive evidence that our analysis is reasonable; that is, that sensible breakpoints have been found.

Overall, there seems to be significant evidence of a structural break in many of the series. Although this weakens only slightly the evidence of periodic integration, the structural break could account for the poor forecasting performance of the PI model in some cases.

CONCLUSIONS

In this paper we describe Bayesian methods for providing evidence for the presence of periodic integration. Informal posterior analysis, HPDI tests, posterior odds, and prediction are all used to shed light on this important hypothesis. A careful empirical analysis of the periodic integration hypothesis in UK macroeconomic time series is carried out. Our in-sample findings support the findings of Franses and Paap (1994) that many UK series seem to be periodically integrated. This finding is of empirical relevance to UK macroeconomists. In PI models the seasonal and non-seasonal components of the data are not independent of each other so that standard seasonal adjustment procedures will produce inadequate seasonally corrected time series. Furthermore, the finding of periodic integration means that traditional unit root models will be unsuited to capturing the stochastic trend properties of the series.

ACKNOWLEDGEMENTS

We would like to thank an anonymous referee for many detailed and helpful comments. This paper was initiated while the first author visited the Department of Economics of the University of Toronto in 1994, and a first version was completed when the second author visited the Tinbergen Institute in Rotterdam in 1995. We thank both institutions for their hospitality. In addition, both authors would like to thank the Royal Netherlands Academy of Arts and Sciences and the Social Sciences and Humanities Research Council of Canada, respectively, for financial support.

REFERENCES

- Barsky, R. and Miron, J., 'The seasonal cycle and the business cycle', *Journal of Political Economy*, **97** (1989), 503–34.
- Beaulieu, J. and Miron, J., 'Seasonal unit roots in aggregate US data', *Journal of Econometrics*, **55** (1993), 305–28.
- Birchenhall, C., Bladen-Hovell, R., Chui, A., Osborn, D., and Smith, J. P., 'A seasonal model of consumption', *Economic Journal*, **99** (1989), 837–43.
- Boswijk, H. P. and Franses, P. H., 'Unit roots in periodic autoregressions', *Journal of Time Series Analysis*, **17** (1996), 221–45.

- Boswijk, H. P., Franses, P. H., and Haldrup, N., 'Multiple unit roots in periodic autoregression', *Journal of Econometrics*, forthcoming.
- Braun, R. and Evans, C., 'Seasonality and equilibrium business cycle theories', *Journal of Economic Dynamics and Control*, **19** (1995), 503–31.
- Canova, F. and Ghysels, E., 'Changes in seasonal patterns: are they cyclical?' *Journal of Economic Dynamics and Control*, **18** (1994), 1143–71.
- Canova, F. and Hansen, B., 'Are seasonal patterns constant over time? A test for seasonal stability', *Journal of Business and Economic Statistics*, **13** (1995), 237–52.
- Chatterjee, S. and Ravikumar, B., 'A neoclassical model of seasonal fluctuations', *Journal of Monetary Economics*, **29** (1992), 59–86.
- Franses, P. H., 'A multivariate approach to modeling univariate seasonal time series', *Journal of Econometrics*, **63** (1994), 251–76.
- Franses, P. H., 'Quarterly USA unemployment: cycles, seasons and asymmetries', *Empirical Economics*, **20** (1995), 717–25.
- Franses, P. H., 'Recent advances in modelling seasonality', *Journal of Economic Surveys*, **10** (1996a), 299–345.
- Franses, P. H., *Periodicity and Stochastic Trends in Economic Time Series*, Oxford: Oxford University Press, 1996b.
- Franses, P. H. and Paap, R., 'Model selection in periodic autoregressions', *Oxford Bulletin of Economics and Statistics*, **56** (1994), 421–40.
- Ghysels, E., 'On the economics and econometrics of seasonality', in Sims, C. A. (ed.), *Advances in Econometrics, Sixth World Congress of the Econometric Society*, Cambridge: Cambridge University Press, 1994.
- Ghysels, E., Hall, A., and Lee, H., 'On periodic structures and testing for seasonal unit roots', *Journal of the American Statistical Association*, forthcoming, 1996.
- Hansen, L. and Sargent, T., *Recursive Linear Models of Dynamic Economies*, Stanford, CA: Hoover Institution, 1995.
- Harvey, A. and Scott, A., 'Seasonality in dynamic regression models', *Economic Journal*, **104** (1994), 1324–45.
- Hylleberg, S., 'Modelling seasonal variation', in Hargreaves, C. P. (ed.), *Nonstationary Time Series Analysis and Cointegration*, Oxford: Oxford University Press, 1994.
- Hylleberg, S., Engle, R., Granger, C., and Yoo, B. S., 'Seasonal integration and cointegration', *Journal of Econometrics*, **44** (1990), 215–38.
- Leamer, E., *Specification Searches*, New York: Wiley, 1978.
- Judge, G., Griffiths, W., Hill, R., Lutkepohl, H. and Lee, T., *The Theory and Practice of Econometrics*, Second Edition, New York: John Wiley and Sons, 1985.
- Miron, J., *The Economics of Seasonal Cycles*, Cambridge, MA: MIT Press, 1996.
- Ooms, M. and Franses, P. H., 'On periodic correlations between estimated seasonal and nonseasonal components in German and U.S. unemployment', *Journal of Business and Economic Statistics*, forthcoming, 1996.
- Osborn, D., 'Seasonality and habit persistence in a life-cycle model of consumption', *Journal of Applied Econometrics*, **3** (1988), 255–66.
- Pagano, M., 'On periodic and multiple autoregressions', *Annals of Statistics*, **6** (1978), 1310–17.
- Phillips, P. C. B., 'To criticise the critics: An objective Bayesian analysis of stochastic trends', *Journal of Applied Econometrics*, **6** (1991), 333–64.
- Poirier, D., 'Bayesian hypothesis testing in linear models with continuously induced conjugate priors across hypotheses', in Bernardo, J. M., DeGroot, M. H., Lindley, D. V. and Smith, A. F. M. (eds), *Bayesian Statistics 2*, Amsterdam: Elsevier Science/North-Holland, 1985.
- Raiffa, H. and Schlaifer, R., *Applied Statistical Decision Theory*, Cambridge, MA: MIT Press, 1961.
- Tiao, G. and Grupe, M., 'Hidden periodic autoregressive-moving average models in time series data', *Biometrika*, **67** (1980), 365–73.
- Todd, R., 'Periodic linear-quadratic methods for modeling seasonality', *Journal of Economic Dynamics and Control*, **14** (1990), 763–95.

Authors' biographies:

Philip Hans Franses is currently at the Econometric Institute, Erasmus University, Rotterdam. He is a time series econometrician. Most of his recent publications have applied time series techniques to issues relating to seasonality.

Gary Koop is currently at the Department of Economics, University of Edinburgh. He is a Bayesian econometrician. His recent publications are in the areas of time series, qualitative choice and stochastic frontier models.

Authors' addresses:

Philip Hans Franses, Econometric Institute, Erasmus University, Rotterdam, The Netherlands.

Gary Koop, Department of Economics, University of Edinburgh, William Robertson Building, 50 George Square, Edinburgh, EH8 9JY, UK.