Random-Coefficient Periodic Autoregression

Philip Hans Franses*
Econometric Institute
Erasmus University Rotterdam

and

Richard Paap
Econometric Institute
Erasmus University Rotterdam

ECONOMETRIC INSTITUTE REPORT EI 2005-34

Abstract

We propose a new periodic autoregressive model for seasonally observed time series, where the number of seasons can potentially be very large. The main novelty is that we collect the periodic parameters in a second-level stochastic model. This leads to a random-coefficient periodic autoregression with a substantial reduction in the number of parameters to be estimated. We discuss representation, estimation, and inference. An illustration for monthly growth rates of US industrial production shows the merits of the new model specification.

Key words: periodic autoregression, random coefficient model
Jel codes: C22, C51.

*We thank Dennis Fok for his comments. Address for correspondence: Philip Hans Franses, Econometric Institute (H11-2), Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, The Netherlands, e-mail: `franses@few.eur.nl`

1 Introduction

Periodic autoregressive time series models [PAR] have become frequently used models to describe and forecast seasonal time series in economics, see, for example, Osborn and Smith (1989), Osborn (1991), Franses and Paap (1994), Franses (1994), Boswijk *et al.* (1997), and Herwartz (1999). These models seem to be particularly considered for time series with low seasonal frequency, like quarterly data within a year or daily data within a week, see Ghysels and Osborn (2001, Chapter 6) and Franses and Paap (2004) for recent surveys of the current state of the art.

Although periodic models may have benefits in terms of fit, forecasting and interpretation, they also have drawbacks. The major problem is that the number of parameters quickly grows with the number of seasons, hence making these models less attractive for application to weekly data and sometimes even monthly data. This is due to the fact that an unrestricted periodic autoregression of order p for seasonal data with frequency S can require pS parameters. Next to potential estimation problems due to a lack of degrees of freedom, the interpretation of such an amount of parameters is also not easy.

In this paper we offer a solution to these problems by proposing a new periodic autoregression, which can easily be used for high frequency seasonal data while preserving interpretability of the parameters. The new model builds on the idea in Jones and Brelsford (1967), where the authors propose to restrict the periodic parameters by imposing a Fourier Series approximation, see also Bloomfield *et al.* (1994) for an application. Indeed, as the periodic parameters themselves show a recurrent pattern, they can be summarized by sums of sine and cosine functions. Jones and Brelsford (1967) also propose to consider a restrictive set of these functions in order to gain degrees of freedom. The latter proposal is taken up in our model as well, though with one major modification. As a smaller set of functions can only amount to an approximation, we introduce an additional error term in our model. Hence, we have a first-level model which contains the time series as the variable to be explained by its own past, and we have a second-level stochastic model for the periodic parameters. We will call the joint model a random-coefficient periodic autoregressive [RCPAR] model.

The outline of our paper is as follows. In Section 2, we describe the representation of the RCPAR model and discuss parameter estimation. In Section 3, we summarize the results of a limited simulation experiment to examine the small sample properties of the parameter estimators. In Section 4, we illustrate our RCPAR model to growth rates of monthly US industrial production. We find convincing evidence that the RCPAR model improves on the deterministic approximation of Jones and Brelsford (1967), as the error term in the second-level model clearly has a non-zero variance. Furthermore, the RCPAR outperforms the deterministic specification and is similar to but much more parsimonious than a conventional PAR model. In Section 5 we conclude with a concise review of further research topics.

2 The model

In this section we discuss the random-coefficient periodic autoregressive model. We consider model representation, parameter estimation, and forecasting.

2.1 Random-coefficient PAR model

Let y_t for $t = 0, \ldots, n = SN$ be a seasonal time series and let S denote the number of seasons with a period of length N. Typically, N amounts to years or weeks and S to months, quarters or weeks. To describe this seasonal time series, one may consider a periodic autoregression of order 1, that is,

$$y_t = \sum_{s=1}^{S} (\mu_s D_{s,t} + \phi_s D_{s,t} y_{t-1}) + \varepsilon_t, \tag{1}$$

where $D_{s,t} = 1$ if t corresponds to season s and 0 otherwise, and where $\varepsilon_t \sim \text{NID}(0, \sigma_\varepsilon^2)$. To ensure stationarity of the time series we impose that $\prod_{s=1}^{S} \phi_s < 1$, see, for example, Franses and Paap (2004, Section 3.2).

If S is large, the number of autoregressive parameters becomes large too. To limit the number of parameters, Jones and Brelsford (1967), amongst other suggestions, propose to describe the ϕ_s parameters by the deterministic function

$$\phi_s = \alpha_0 + \alpha_1 \cos\left(\frac{2\pi s}{S} - \alpha_2 \pi\right), \tag{2}$$

where α_0, α_1, α_2 are unknown parameters. Such a function reduces the amount of parameters from S to 3. Note that for parameter identification we restrict $\alpha_2 \in [0, 1)$ as $\cos(x + k\pi) = (-1)^k \cos(x)$ for $k \in \mathbb{Z}$ and $x \in \mathbb{R}$.

For some economic time series, particularly when S is large, the deterministic specification (2) may be too restrictive. Therefore, we propose to extend (2) with a random error term resulting in

$$\phi_s = \alpha_0 + \alpha_1 \cos\left(\frac{2\pi s}{S} - \alpha_2\pi\right) + u_s, \tag{3}$$

where $u_s \sim \text{NID}(0, \sigma_u^2)$ and $\text{E}[u_s\varepsilon_t] = 0$ for all s, t. This random term distinguishes (3) from the deterministic specification (2). Adding the second-level error term, leads to a random-coefficient specification, see, for example, Swamy (1970) and more recently Maddala *et al.* (1997) and Hsiao (2003, Chapter 6). The difference with a standard random-coefficient approach is that we shrink the periodic autoregressive parameters to the deterministic function (2) instead of a simple mean. We call the PAR(1) model in combination with (3) a random-coefficient periodic autoregression of order 1 RCPAR(1).

2.2 Parameter estimation

Parameter estimation of the PAR model with the deterministic specification (2) can be done using concentrated nonlinear least squares [NLS], see Jones and Brelsford (1967). Given the value of $\alpha = (\alpha_0, \alpha_1, \alpha_2)$, we can easily compute the optimal values of the remaining parameters using ordinary least squares [OLS]. Hence, one only has to perform a nonlinear maximization with respect to the three α parameters.

In case of our random-coefficient specification (3) the RCPAR(1) model can be written as

$$y_t = \sum_{s=1}^{S} (\mu_s D_{s,t} + (\alpha_0 + \alpha_1 \cos\left(\frac{2\pi s}{S} - \alpha_2\pi\right)) D_{s,t} y_{t-1}) + v_t, \tag{4}$$

where the error term $v_t = \varepsilon_t + D_{s,t} u_s y_{t-1}$ is heteroskedastic. It is easy to derive that the vector of disturbances $v = (v_1, \ldots, v_n)'$ is normal distributed with mean 0 and covariance matrix

$$\Sigma_v = \sigma_\varepsilon^2 I_n + \Sigma_u, \tag{5}$$

where the (i, j)th element of the $n \times n$ matrix Σ_u equals $\sigma_u^2 D_{s,i} y_{i-1} D_{s,j} y_{j-1}$. To estimate the model parameters we may opt for a feasible generalized NLS [FGNLS] estimator following the lines of Swamy (1970) or a maximum likelihood [ML] estimator.

FGNLS estimator

To construct the GNLS estimator we need consistent estimators for σ_ε^2 and σ_u^2. A consistent estimator for σ_ε^2 is given by $\hat{\sigma}_\varepsilon^2 = (n-2S)^{-1} \sum_{t=1}^{n} \hat{e}_t^2$, where $\hat{e}_t$ are the OLS residuals from

$$y_t = \sum_{s=1}^{S} (\mu_s D_{s,t} + \phi_s D_{s,t} y_{t-1}) + e_t \tag{6}$$

for $t = 1, \ldots, n$. An estimator for σ_u^2 can be obtained by estimating the parameters of

$$\hat{\phi}_s = \alpha_0 + \alpha_1 \cos\left(\frac{2\pi s}{S} - \alpha_2 \pi\right) + \omega_s \tag{7}$$

for $s = 1, \ldots, S$, using NLS, where $\hat{\phi}_s$ is the OLS estimate obtained from (6). If we denote the NLS residuals by $\hat{\omega}_s$ the estimator for σ_u^2 is given by $\hat{\sigma}_u^2 = 1/S \sum_{s=1}^{S} \hat{\omega}_s^2$. Note this estimator is consistent for both S and n going to infinity. In practice S may be relatively small and we may divide by $S-3$ instead of S, see Davidson and MacKinnon (1993, Section 3.2).

The FGNLS estimator follows from minimizing $v' \hat{\Sigma}_v^{-1} v$ with respect to α and $\mu = (\mu_1, \ldots, \mu_S)$, where $\hat{\Sigma}_v$ is given by (5) evaluated in $\hat{\sigma}_\varepsilon^2$ and $\hat{\sigma}_u^2$. Note that given the values of α_0, α_1, α_2, the optimal values of the remaining parameters can be obtained using a FGLS estimator and hence we only have to minimize with respect to the three α parameters. Standard errors of the parameters can be estimated using $\hat{G}' \hat{\Sigma}_v^{-1} \hat{G}$, where $\hat{G}$ is the n-dimensional vector of first-order derivatives of the nonlinear regression mean of (4) with respect to α and μ evaluated in the FGNLS estimates, see Davidson and MacKinnon (1993, Section 9.6). Note that $\mathrm{E}[u_s y_{t-1}] \neq 0$ and hence the FGNLS estimator may be biased. However it is easy to show that the correlation between y_{t-1} and u_s is proportional to the product of at least $S-1$ ϕ_s parameters. For stationary periodic time series, practical values of ϕ_s are about 0.5, and hence this product is approximately zero if S is 12 or higher and so the bias will be small. Simulation results in Section 3 show that the estimator performs very well for relatively small values of S and small values of the product of the individual ϕ_s parameters.

ML estimator

It is also possible to estimate the model parameters using ML. There are two ways to derive the likelihood function. Using the results of the FGNLS approach, it is easy to show that the log of the joint density of $Y = (y_0, \ldots, y_n)$ is given by

$$\ln f(Y; \mu, \alpha, \sigma_\varepsilon, \sigma_u) = -\frac{n}{2} \ln 2\pi - \frac{1}{2} \ln |\Sigma_v| - \frac{1}{2} v \Sigma_v^{-1} v, \tag{8}$$

where Σ_v is given by (5) and v is the vector of residuals of (4). The ML estimator can be obtained by maximizing this log likelihood function with respect to the model parameters. Given the values of α, σ_u and σ_ε the optimal value of the μ parameters is given by a GLS estimator as discussed before. Hence, one only needs a nonlinear optimization with respect to 5 parameters. The disadvantage of the approach is that one has to deal with an $n \times n$ covariance matrix. Although there is an analytical expression for the inverse of Σ_v^{-1} due to its specific structure, the size of the covariance may become very large for large values S and n.

Another way to compute the likelihood function is to integrate with respect to the error terms u_s. The advantage of this approach is that one does not have to deal with the potentially large covariance matrix Σ_v. Consider the density of Y conditional on $\phi = (\phi_1, \ldots, \phi_S)$, that is,

$$f(Y|\phi; \mu, \sigma_\varepsilon) = \prod_{s=1}^{S} \prod_{T=1}^{N} \frac{1}{\sigma_\varepsilon} \phi\left(\frac{y_{s+S(T-1)} - \mu_s - \phi_s y_{s+S(T-1)-1}}{\sigma_\varepsilon} \right), \tag{9}$$

where $\phi(\cdot)$ denotes the pdf of a standard normal distribution. Note that we can consider each of the S seasons separately. The unconditional density of Y is given by

$$f(Y; \mu, \alpha) = \int_{\mathbb{R}^S} f(Y|\phi; \mu, \sigma_\varepsilon) \prod_{s=1}^{S} \frac{1}{\sigma_u} \phi\left(\frac{\phi_s - \alpha_0 - \alpha_1 \cos\left(\frac{2\pi s}{S} - \alpha_2 \pi\right)}{\sigma_u} \right) d\phi_1 \ldots d\phi_S. \tag{10}$$

This unconditional distribution can be split up in S parts which equal

$$\int_{-\infty}^{\infty} \frac{1}{\sigma_u} \phi\left(\frac{\phi_s - \bar{\phi}_s}{\sigma_u} \right) \prod_{T=1}^{N} \frac{1}{\sigma_\varepsilon} \phi\left(\frac{y_{s+S(T-1)} - \mu_s - \phi_s y_{s+S(T-1)-1}}{\sigma_\varepsilon} \right) d\phi_s, \tag{11}$$

with $\bar{\phi}_s = \alpha_0 + \alpha_1 \cos(\frac{2\pi s}{S} - \alpha_2 \pi)$. If we define

$$\tilde{\phi}_s = (\sigma_u^{-2} + \sum_{T=1}^{N} (y_{s+S(T-1)-1}/\sigma_\varepsilon)^2)^{-1} (\bar{\phi}_s/\sigma_u^2 + y_{s+S(T-1)-1}(y_{s+S(T-1)} - \mu_s)/\sigma_\varepsilon^2), \tag{12}$$

we can rewrite (11) as

$$c^{-1} \int_{-\infty}^{\infty} \exp\Big(-\frac{1}{2}\frac{(\bar{\phi}_s - \tilde{\phi}_s)^2}{\sigma_u^2} - \frac{1}{2}\sum_{T=1}^{N} \frac{(y_{s+S(T-1)} - \mu_s - \tilde{\phi}_s y_{s+S(T-1)-1})^2}{\sigma_\varepsilon^2} - \frac{1}{2}(\phi_s - \tilde{\phi}_s)^2/\sigma_{\tilde{\phi}_s}^2\Big) d\phi_s, \tag{13}$$

where $\sigma_{\tilde{\phi}_s}^2 = (\sigma_u^{-2} + \sum_{T=1}^{N}(y_{s+S(T-1)-1}/\sigma_\varepsilon)^2)^{-1}$ and $c = (\sqrt{2\pi})^N \sigma_\varepsilon^N \sigma_u$. Evaluating the integral results in

$$\frac{\sigma_{\tilde{\phi}_s}}{\sigma_u(\sqrt{2\pi}\sigma_\varepsilon)^N} \exp\Big(-\frac{1}{2}\frac{(\bar{\phi}_s - \tilde{\phi}_s)^2}{\sigma_u^2} - \frac{1}{2}\sum_{T=1}^{N} \frac{(y_{s+S(T-1)} - \mu_s - \tilde{\phi}_s y_{s+S(T-1)-1})^2}{\sigma_\varepsilon^2}\Big), \tag{14}$$

and hence the likelihood function can be written as

$$f(Y; \mu, \alpha, \sigma_\varepsilon, \sigma_u) = \prod_{s=1}^{S} \sqrt{2\pi}\sigma_{\tilde{\phi}_s} \frac{1}{\sigma_u}\phi\Big(\frac{\bar{\phi}_s - \tilde{\phi}_s}{\sigma_u}\Big) \prod_{T=1}^{N} \frac{1}{\sigma_\varepsilon}\phi\Big(\frac{y_{s+S(T-1)} - \mu_s - \tilde{\phi}_s y_{s+S(T-1)-1}}{\sigma_\varepsilon}\Big). \tag{15}$$

The maximum likelihood estimator is obtained by maximizing the log-likelihood function (15) with respect to μ, α, σ_ε and σ_u. Note that it is now not possible anymore to construct a concentrated ML estimator and hence we have to maximize with respect to all parameters. Standard errors of the parameters can be obtained from the second-order derivatives of the log-likelihood function.

To describe the periodic autoregressive structure, our RCPAR model requires four parameters, that is, α_0, α_1, α_2 and σ_u. Hence, there should be enough degrees of freedom to apply this model to describe monthly time series, where we should mention that we then only have twelve observations to estimate σ_u. To investigate the small sample properties of the FGNLS and the ML estimator of the RCPAR model, a simulation experiment shall be beneficial. Note that for many practical purposes, like out-of-sample forecasting, we are more interested in consistent estimates of the individual ϕ_s parameters rather than the value of σ_u itself.

2.3 Forecasting

To use the RCPAR model (1) and (3) for forecasting we need the values of the ϕ_s parameters. These parameters are stochastic and unobserved. To estimate the parameters for each of the seasons we use the conditional expected value of ϕ_s given the data Y, which is given by

$$\mathrm{E}[\phi_s|Y;\mu,\alpha,\sigma_\varepsilon,\sigma_u] = \int_{-\infty}^{\infty} \phi_s \frac{f(Y|\phi;\mu,\sigma_\varepsilon)\sigma_u^{-1}\phi((\phi_s-\bar{\phi}_s)/\sigma_u)}{f(Y;\mu,\alpha,\sigma_\varepsilon,\sigma_u)} d\phi_s = \tilde{\phi}_s \qquad (16)$$

for $s = 1,\ldots,S$. If we evaluate these $\tilde{\phi}_s$ as defined in (12) in the parameters estimates, we obtain an estimate for the seasonal ϕ_s parameters. The conditional variance of the ϕ_s is given by $\mathrm{V}[\phi_s|Y;\mu,\alpha,\sigma_\varepsilon,\sigma_u] = \sigma^2_{\tilde{\phi}_s}$ for $s = 1,\ldots,S$, where $\sigma^2_{\tilde{\phi}_s}$ is defined below (13).

Given the estimates of ϕ_s, one can generate forecasts of the periodic autoregressive model in a straightforward manner, see, for example, Franses and Paap (2004, Sections 3.4 & 4.4).

2.4 RCPAR(p) model

In many applications it is not likely that a first-order PAR model is sufficient to capture dynamics. To allow for higher order autocorrelation we assume that ε_t follows an AR($p-1$) process, that is,

$$\varepsilon_t = \sum_{i=1}^{p-1} \psi_i \varepsilon_{t-i} + \eta_t \qquad (17)$$

with $\eta_t \sim \mathrm{NID}(0,\sigma^2_\eta)$ and $\mathrm{E}[u_s\eta_t] = 0$ for all s,t. The combined RCPAR(p) model can then be written as

$$y_t - \sum_{s=1}^{S} \phi_s D_{s,t} y_{t-1} = \sum_{s=1}^{S} \lambda_s D_{s,t} + \sum_{i=1}^{p-1} \psi_i (y_{t-i} - \sum_{s=1}^{S} \phi_s D_{s,t-i} y_{t-1-i}) + \eta_t \qquad (18)$$

together with (3), where $\lambda_s = \mu_s - \sum_{i=1}^{p-1} \psi_i \mu_{s-i}$ with $\mu_{s-Sk} = \mu_s$ for $k \in \mathbb{Z}$.

Parameter estimation can be done in a similar way as for the RCPAR(1) model. One can either choose for an FGNLS estimator or an ML estimator. Note that the covariance matrix of the disturbances, denoted by Σ_v in the RCPAR(1) representation, will depend on the ψ_i parameters. This is not a problem for the FGNLS estimator as these parameters can be estimated consistently in the first-step regression like (6). In the final FGNLS estimation step, one can still use a concentrated GLS estimator to reduce the dimension of the optimization procedure. Note that if we choose a relatively large value of p, the correlation between y_{t-p} and u_s may become large and hence this may lead to a bias in the FGNLS estimator.

Concentration is not possible anymore in the ML approach. It is therefore computationally more convenient to integrate with respect to the latent u_s variables in constructing the likelihood function. One simply rewrites (18) as

$$(y_t - \sum_{i=1}^{p-1} \psi_i y_{t-i}) = \sum_{s=1}^{S} \Big(\lambda_s D_{s,t} + \phi_s (D_{s,t} y_{t-1} - \sum_{i=1}^{p-1} \psi_i D_{s,t-i} y_{t-1-i})\Big) + \eta_t \tag{19}$$

or

$$\tilde{y}_t = \sum_{s=1}^{S} (\lambda_s D_{s,t} + \phi_s z_t) + \eta_t, \tag{20}$$

where $\tilde{y}_t = (y_t - \sum_{i=1}^{p-1} \psi_i y_{t-i})$ and $z_t = D_{s,t} y_{t-1} - \sum_{i=1}^{p-1} \psi_i D_{s,t-i} y_{t-1-i}$. The analytical integration can be done in a similar way as before by just replacing $y_{s+S(T-1)}$ by $\tilde{y}_{s+S(T-1)}$ and $y_{s+S(T-1)-1}$ by $z_{s+S(T-1)}$ in the relevant equations.

3 Simulations

To investigate the small sample properties of the FGNLS and ML estimator, we consider a simulation experiment. As data generating process [DGP] we consider the RCPAR(1) model as in (1) together with (3). The parameter values are set at $\mu_s = 1$ for $s = 1, \ldots, S$, and $\alpha_0 = \alpha_1 = \alpha_2 = 0.5$, $\sigma_\varepsilon = 1$, $\sigma_u = 0.2$. Without the error term, this DGP implies ϕ_s parameters as displayed in Figure 1 for the case $S = 12$.

We first consider 100 years of monthly observations ($S = 12$) which corresponds to 1200 observations. Table 1 displays the results of the ML estimator. The number of replications is 10 000. It can be seen from this table that the means of the maximum likelihood estimates correspond very well with the true parameters, except for the σ_u parameter which is slightly underestimated. This seems to be due to the small number of seasons. The biases in the estimates of ϕ_s based on (16) are however very small. The maximum bias over the seasons is only 0.005. Furthermore, Table 1 shows that the small sample distribution of the maximum likelihood estimator corresponds reasonably well with the normal distribution, especially for the seasonal intercept parameters μ_s. For the α parameters there is some size distortion. And, as expected, for σ_u, the normal approximation is not too good due to the small sample bias.

Table 2 displays the results of the FGNLS estimator. We can see that the mean of the FGNLS estimator corresponds better to the true values than the ML estimator. The bias in the estimate for σ_u is now much smaller. The variance of the estimator is comparable to the variance of the ML estimator. Furthermore, Table 2 shows that the theoretical size of the small sample distribution of the FGNLS estimator is almost the same as the nominal size with the exception of the α_1 parameter. Hence, we recommend to use the normal approximation in practice.

To see whether matters improve for the ML estimator, we repeat this simulation exercise, where we increase the number of seasons. Table 3 displays the results for the ML estimator for $S = 24$ and $N = 100$. We notice that the small sample bias in the ML estimator for σ_u is substantially smaller than for $S = 12$. The small sample bias in the other parameters is almost 0. The small sample distribution of the ML estimator is closer to normal than for $S = 12$.

In sum, the simulation results suggest that we can reliably draw inference on the parameters in the cosine function, but that we have to be careful with their estimated standard errors. In fact, these errors may be a bit too small for monthly data.

4 Illustration

To illustrate our RCPAR model we consider the monthly growth rates of total industrial production of the United States. The estimation sample is 1920.01–2000.12. Figure 2 displays the times series under scrutiny.

4.1 Model specification

First, we consider a regular PAR(1) model (1) for the series, that is, with unrestricted parameters. The parameter estimates are given in the second column of Table 4. The LM-test for first-order serial correlation in the residuals of a PAR(1) model equals 0.80 with a p-value 0.74. The same test for first-to-fourth order serial correlation equals 2.14 with a

p-value of 0.24. Hence, this model seems to fit the data well. To test for periodicity we test for equal periodic autoregressive parameters ϕ_s in (1) using an F-test. The F-statistic equals 9.04 which is clearly significant at the 5% level and hence there is substantial evidence of periodicity in the autoregressive parameters.

Next we consider the PAR(1) model (1) under the restriction (2), that is, the Jones and Brelsford (1967) type of model. The parameter estimates are given in the fourth column of Table 4. Note that this model only uses three parameters to describe the periodic autoregressive parameters instead of twelve. This restriction has some impact on the estimates of σ_ε and of the seasonal intercepts μ_s, which are clearly different than for the unrestricted model. If we evaluate (2) in $\hat{\alpha}$ we obtain the estimates for ϕ_s, which equal 0.40, 0.32, 0.24, 0.17, 0.14, 0.16, 0.21, 0.29, 0.37, 0.44, 0.47 and 0.45. Figure 3 displays the estimated ϕ_s parameters for the unrestricted and the restricted PAR(1) model. We clearly see a difference between the two specifications, where the restricted model corresponds with the smooth function.

To make the autoregressive specification more flexible, we consider our RCPAR(1) model (1) with (3). The sixth column of Table 4 displays the ML parameter estimates of this model. The LR-statistic for $\sigma_u = 0$ equals 44.59. As we have a one-sided alternative ($\sigma_u > 0$), the asymptotic distribution of the likelihood ratio statistic is $\frac{1}{2}\chi^2(0) + \frac{1}{2}\chi^2(1)$, see Wolak (1989). Hence, it is clearly significant at the 5% level and our stochastic specification (3) is preferred. Note that the seasonal intercept parameters and the σ_ε are closer to the unrestricted PAR(1) model than those of the non-stochastic specification (2).

The expected values of the ϕ_s parameters (16) are 0.12, 0.29, 0.19, -0.03, 0.42, 0.42, 0.29, -0.14, 0.40, 0.55, 0.59 and 0.74. Note that the values are close to the estimated ϕ_s parameters of the unrestricted PAR specification. Figure 3 also displays the estimates of the RCPAR specification. Evidently, the estimated PAR parameters are close to the estimates of the unrestricted PAR(1) but now estimated using a much smaller number of parameters.

Our RCPAR specification is however not nested in the unrestricted PAR model. Therefore, to compare the fit of both models we use the BIC. As can be seen from the last line of Table 4, our random-coefficient periodic autoregressive model has the smallest value, and hence has the best fit.

The eighth column of Table 4 shows the FGNLS estimates of the RCPAR model The parameter estimates are almost the same as the ML estimates with the exception of the σ_u estimate which is larger. This result corresponds to our simulation results which showed that there is a negative small sample bias in the ML estimator for σ_u.

4.2 Forecasting

Table 5 compares the forecasting performance of the three models, where we use the ML estimates of the RCPAR model[1]. We remove the final H observations from the time series and re-estimate the parameters of the three models for the smaller sample for $H = 12, 24, 36$. Next, we generated H 1-step ahead forecasts and 1 to H-step ahead forecasts for $H = 12, 24, 36$. We compared the forecasts with the out-of-sample realizations using the Root Mean Squared Error [RMSE], Mean Absolute Error [MAE] and the Mean Absolute Percentage Error [MAPE].

The table shows that the forecasting performance of the random-coefficient and the unrestricted PAR model is very similar. This holds for the 1-step ahead predictions as well as the multi-step ahead predictions. The biggest difference can be found for the 12 1-step ahead predictions if we consider the MAPE criterion. This suggest that the parameter reduction from an unrestricted PAR to the random-coefficient PAR specification does not harm forecasting performance. The restricted PAR always performs worse except for the 36 1-step ahead forecasts. Hence, imposing a deterministic structure (2) seems to lead to a decrease in forecasting performance.

5 Conclusion

In this paper we have introduced a new periodic model that should be useful to capture periodic properties of high frequency seasonal data. We illustrated the model for a monthly time series, but in our further work we will explore the relevance of the random-coefficient

[1] There is no substantial difference in forecasting performance if we use the FGNLS estimates.

PAR model for weekly data. This exploration will concern univariate time series data, and also multivariate series. There are various situations where one might think of having a periodic effect of an explanatory variable on the dependent variable.

Another interesting avenue for further research concerns the inclusion of variables other than sine and cosine functions in the second-level of the model. Indeed, an important and meaningful question concerns the nature of the periodicity of the parameters. Perhaps there are economic or institutional reasons why parameters show periodic patterns and the relevance of these reasons can simply be examined using the second-level specification.

Table 1: Small sample properties of the ML estimator and estimated z-ratios. Sample size is 100×12 (1200 data points). Number of replications is 10 000.[c]

θ	true value	$E[\hat{\theta}]$[a]	$V[\hat{\theta}]$[a]	Nominal size z-ratios[b]					
				left tail			right tail		
				0.05	0.10	0.20	0.20	0.10	0.05
α_0	0.50	0.50	0.00	0.13	0.19	0.28	0.25	0.17	0.12
α_1	0.50	0.51	0.01	0.11	0.15	0.24	0.29	0.20	0.14
α_2	0.50	0.50	0.00	0.13	0.18	0.27	0.26	0.17	0.12
μ_1	1.00	1.00	0.03	0.07	0.12	0.22	0.23	0.13	0.08
μ_2	1.00	1.00	0.04	0.07	0.13	0.22	0.23	0.13	0.08
μ_3	1.00	1.00	0.05	0.07	0.12	0.22	0.23	0.13	0.08
μ_4	1.00	1.01	0.07	0.07	0.12	0.22	0.23	0.13	0.08
μ_5	1.00	1.01	0.08	0.07	0.13	0.22	0.23	0.13	0.08
μ_6	1.00	1.00	0.09	0.08	0.13	0.22	0.24	0.14	0.08
μ_7	1.00	1.00	0.08	0.08	0.13	0.22	0.23	0.14	0.09
μ_8	1.00	1.00	0.04	0.08	0.13	0.22	0.23	0.14	0.08
μ_9	1.00	0.99	0.03	0.07	0.13	0.23	0.21	0.12	0.07
μ_{10}	1.00	1.00	0.02	0.07	0.12	0.22	0.22	0.12	0.07
μ_{11}	1.00	1.00	0.02	0.07	0.12	0.22	0.22	0.12	0.07
μ_{12}	1.00	1.00	0.03	0.07	0.13	0.22	0.22	0.12	0.07
σ_ε	1.00	1.00	0.00	0.08	0.14	0.26	0.16	0.07	0.03
σ_u	0.20	0.14	0.00	0.45	0.53	0.64	0.03	0.01	0.00

[a] The mean and variance of the ML estimates.

[b] The cell denotes the empirical size of the distribution of the z-ratios defined as, $(\hat{\theta} - \theta)/\hat{\sigma}(\theta)$, where $\hat{\sigma}(\theta)$ denotes the estimated standard error of $\hat{\theta}$.

[c] The maximum bias in the estimated ϕ_s parameters based on (16) over the 12 seasons is 0.005.

Table 2: Small sample properties of the FGNLS estimator and estimated z-ratios. Sample size is 100×12 (1200 data points). Number of replications is 10 000.[c]

θ	true value	$\mathrm{E}[\hat{\theta}]$[a]	$\mathrm{V}[\hat{\theta}]$[a]	Nominal size z-ratios[b]					
				left tail			right tail		
				0.05	0.10	0.20	0.20	0.10	0.05
α_0	0.50	0.50	0.00	0.05	0.10	0.20	0.19	0.10	0.05
α_1	0.50	0.51	0.01	0.04	0.08	0.16	0.21	0.11	0.06
α_2	0.50	0.50	0.00	0.06	0.10	0.20	0.19	0.10	0.05
μ_1	1.00	1.00	0.03	0.05	0.10	0.20	0.20	0.10	0.05
μ_2	1.00	1.00	0.04	0.05	0.10	0.20	0.20	0.11	0.06
μ_3	1.00	1.00	0.05	0.05	0.10	0.20	0.21	0.10	0.05
μ_4	1.00	1.00	0.05	0.05	0.10	0.20	0.20	0.10	0.05
μ_5	1.00	1.00	0.06	0.05	0.10	0.20	0.21	0.10	0.05
μ_6	1.00	1.00	0.07	0.05	0.10	0.19	0.20	0.11	0.05
μ_7	1.00	1.00	0.06	0.05	0.09	0.19	0.21	0.10	0.05
μ_8	1.00	1.00	0.04	0.05	0.10	0.20	0.20	0.10	0.05
μ_9	1.00	1.00	0.02	0.05	0.10	0.21	0.20	0.10	0.05
μ_{10}	1.00	1.00	0.02	0.05	0.10	0.20	0.20	0.10	0.05
μ_{11}	1.00	1.00	0.02	0.05	0.10	0.20	0.20	0.10	0.05
μ_{12}	1.00	1.00	0.02	0.06	0.11	0.21	0.20	0.10	0.05
σ_ε	1.00	1.00	0.00						
σ_u	0.20	0.21	0.00						

[a] The mean and variance of the FGNLS estimates.

[b] The cell denotes the empirical size of the distribution of the z-ratios defined as, $(\hat{\theta} - \theta)/\hat{\sigma}(\theta)$, where $\hat{\sigma}(\theta)$ denotes the estimated standard error of $\hat{\theta}$.

[c] The maximum bias in the estimated ϕ_s parameters based on (16) over the 12 seasons is 0.009.

Table 3: Small sample properties of the maximum likelihood estimator and estimated z-ratios. Sample size is 100 × 24 (2400 data points). Number of replications is 10 000.[c]

θ	true value	$E[\hat{\theta}]$[a]	$V[\hat{\theta}]$[a]	Nominal size z-ratios[b]					
				left tail			right tail		
				0.05	0.10	0.20	0.20	0.10	0.05
α_0	0.50	0.50	0.00	0.09	0.14	0.25	0.23	0.13	0.08
α_1	0.50	0.50	0.00	0.08	0.13	0.22	0.26	0.15	0.09
α_2	0.50	0.50	0.00	0.08	0.14	0.24	0.24	0.14	0.08
μ_1	1.00	1.00	0.03	0.06	0.11	0.21	0.21	0.11	0.06
μ_2	1.00	1.00	0.04	0.06	0.11	0.21	0.22	0.12	0.06
μ_3	1.00	1.00	0.04	0.05	0.10	0.20	0.21	0.11	0.06
μ_4	1.00	1.00	0.05	0.05	0.11	0.21	0.21	0.11	0.06
μ_5	1.00	1.00	0.06	0.06	0.11	0.21	0.21	0.11	0.06
μ_6	1.00	1.00	0.07	0.05	0.10	0.20	0.21	0.11	0.06
μ_7	1.00	1.00	0.08	0.05	0.10	0.20	0.21	0.11	0.06
μ_8	1.00	1.00	0.09	0.06	0.11	0.21	0.21	0.11	0.06
μ_9	1.00	1.00	0.09	0.05	0.10	0.20	0.21	0.11	0.06
μ_{10}	1.00	1.01	0.10	0.05	0.10	0.20	0.21	0.11	0.06
μ_{11}	1.00	1.01	0.10	0.05	0.10	0.19	0.21	0.11	0.06
μ_{12}	1.00	1.00	0.09	0.06	0.11	0.20	0.21	0.11	0.06
μ_{13}	1.00	1.00	0.07	0.06	0.11	0.20	0.21	0.11	0.06
μ_{14}	1.00	1.00	0.05	0.06	0.11	0.21	0.21	0.11	0.06
μ_{15}	1.00	1.00	0.03	0.05	0.10	0.21	0.21	0.11	0.06
μ_{16}	1.00	1.00	0.02	0.06	0.11	0.21	0.21	0.11	0.06
μ_{17}	1.00	1.00	0.02	0.06	0.11	0.20	0.21	0.11	0.06
μ_{18}	1.00	1.00	0.02	0.05	0.11	0.21	0.21	0.10	0.05
μ_{19}	1.00	1.00	0.02	0.05	0.10	0.21	0.21	0.11	0.05
μ_{20}	1.00	1.00	0.02	0.05	0.10	0.20	0.20	0.11	0.05
μ_{21}	1.00	1.00	0.02	0.05	0.11	0.21	0.21	0.11	0.06
μ_{22}	1.00	1.00	0.02	0.06	0.11	0.21	0.20	0.10	0.05
μ_{23}	1.00	1.00	0.02	0.06	0.11	0.21	0.21	0.11	0.06
μ_{24}	1.00	1.00	0.03	0.06	0.11	0.21	0.21	0.11	0.06
σ_ε	1.00	1.00	0.00	0.09	0.17	0.30	0.13	0.06	0.03
σ_u	0.20	0.17	0.00	0.34	0.43	0.56	0.04	0.01	0.00

[a] The mean and variance of the ML estimates.

[b] The cell denotes the empirical size of the distribution of the z-ratios defined as, $(\hat{\theta} - \theta)/\hat{\sigma}(\theta)$, where $\hat{\sigma}(\theta)$ denotes the estimated standard error of $\hat{\theta}$.

[c] The maximum bias in the estimated ϕ_s parameters based on (16) over the 24 seasons is 0.008.

Table 4: Parameter estimates and estimated standard errors of the unrestricted, the restricted and the random-coefficient PAR(1) model for the growth rates of US industrial production.

	unrestricted PAR(1)		restricted PAR(1)		random-coefficient PAR(1) ML		random-coefficient PAR(1) FGNLS	
	$\hat{\theta}$	s e	$\hat{\theta}$	s e	$\hat{\theta}$	s e	$\hat{\theta}$	s e[a]
μ_1	1.17	0.33	1.96	0.29	1.26	0.33	1.23	0.33
μ_2	1.93	0.28	1.89	0.27	1.92	0.27	1.92	0.28
μ_3	0.27	0.39	0.13	0.30	0.23	0.36	0.25	0.37
μ_4	−0.35	0.27	−0.52	0.26	−0.39	0.26	−0.38	0.27
μ_5	0.76	0.26	0.62	0.26	0.73	0.26	0.74	0.26
μ_6	1.49	0.26	1.67	0.26	1.52	0.26	1.51	0.26
μ_7	−4.41	0.31	−4.25	0.28	−4.39	0.30	−4.40	0.30
μ_8	2.87	0.42	4.78	0.32	3.12	0.42	3.03	0.42
μ_9	0.62	0.41	0.73	0.32	0.62	0.39	0.62	0.40
μ_{10}	−0.88	0.33	−0.61	0.28	−0.85	0.32	−0.86	0.32
μ_{11}	−2.06	0.26	−2.01	0.26	−2.05	0.25	−2.05	0.26
μ_{12}	−1.04	0.31	−1.67	0.28	−1.13	0.30	−1.10	0.31
ϕ_1	0.08	0.09						
ϕ_2	0.29	0.10						
ϕ_3	0.17	0.13						
ϕ_4	−0.09	0.12						
ϕ_5	0.49	0.12						
ϕ_6	0.47	0.10						
ϕ_7	0.30	0.10						
ϕ_8	−0.20	0.09						
ϕ_9	0.40	0.09						
ϕ_{10}	0.56	0.10						
ϕ_{11}	0.62	0.10						
ϕ_{12}	0.79	0.09						
σ_ε	2.26	0.04	2.36	0.05	2.27	0.05	2.28	
α_0			0.30	0.03	0.32	0.07	0.32	0.09
α_1			−0.16	0.04	−0.16	0.10	−0.16	0.13
α_2			0.86	0.08	0.82	0.20	0.82	0.25
σ_u					0.23	0.06	0.30	
max. lik.	-2167.89		-2209.64		-2187.35			
BIC	4510.1		4529.5		4490.6			

[a] Standard errors for $\hat{\sigma}_\varepsilon$ and $\hat{\sigma}_u$ are not available for the FGNLS estimator.

Table 5: Out-of-sample forecasting performance of the unrestricted, the restricted and the random-coefficient PAR(1) model for the growth rates of US industrial production[a].

	unrestricted PAR(1)	restricted PAR(1)	random-coef. PAR(1)
1-step ahead predictions 2000.01–2000.12			
RMSE	1.09	1.12	1.10
MAE	0.93	0.93	0.94
MAPE	0.60	0.60	0.62
1-step ahead predictions 1999.01–2000.12			
RMSE	1.13	1.14	1.13
MAE	0.95	0.95	0.95
MAPE	0.81	0.81	0.81
1-step ahead predictions 1998.01–2000.12			
RMSE	1.20	1.18	1.20
MAE	0.98	0.98	0.98
MAPE	1.13	1.13	1.12
1–12 step ahead predictions 2000.01–2000.12			
RMSE	0.99	1.01	0.99
MAE	0.84	0.89	0.85
MAPE	0.71	0.76	0.71
1–24 step ahead predictions 1999.01–2000.12			
RMSE	1.01	1.02	1.01
MAE	0.84	0.86	0.84
MAPE	0.72	0.74	0.72
1–36 step ahead predictions 1999.01–2000.12			
RMSE	1.04	1.05	1.04
MAE	0.85	0.86	0.85
MAPE	0.90	0.94	0.91

[a] Out-of-sample forecasts are constructed using parameter estimates obtained from in-sample observations only.

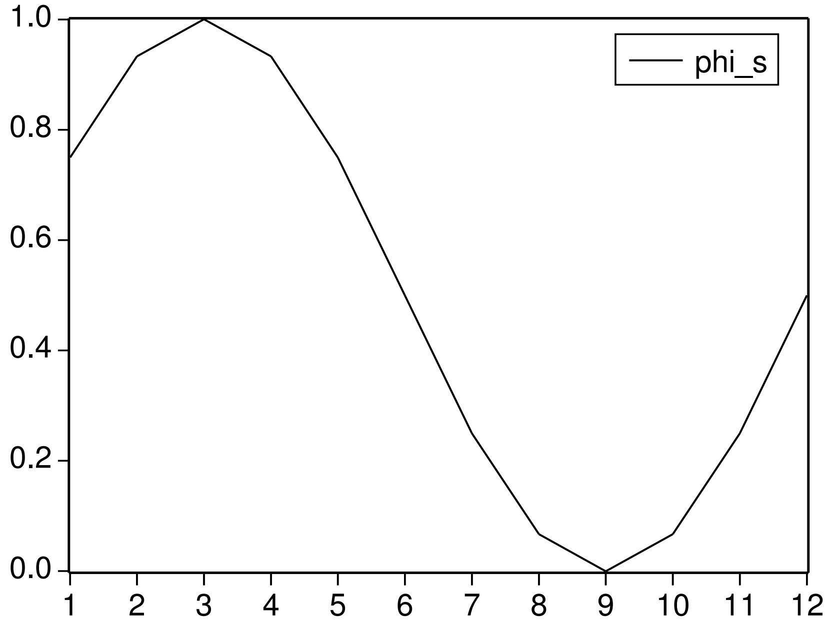

Figure 1: PAR(1) parameters of DGP

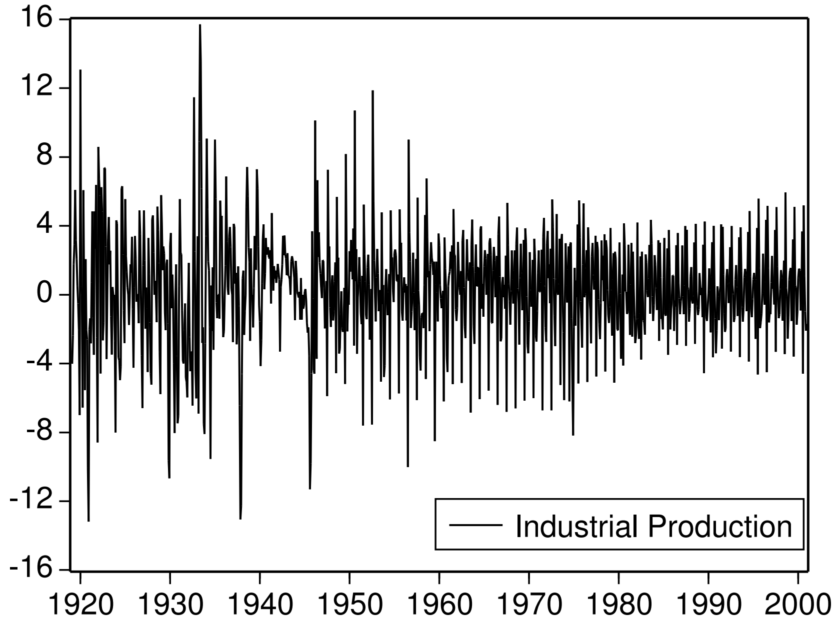

Figure 2: Monthly growth percentages of US industrial production 1920.01–2000.12

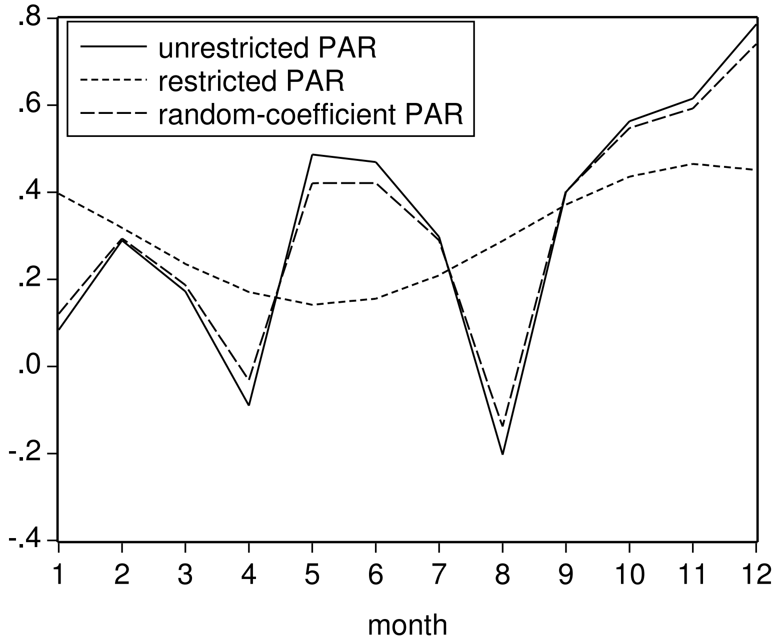

Figure 3: Estimated PAR(1) parameters for the three model specifications

References

Bloomfield, P., H. L. Hurd, and R. B. Lund (1994), Periodic Correlation in Stratospheric Ozone Data, *Journal of Time Series Analysis*, **15**, 127–150.

Boswijk, H. P., P. H. Franses, and N. Haldrup (1997), Multiple Unit Roots in Periodic Autoregression, *Journal of Econometrics*, **80**, 167–193.

Davidson, R. and J. G. MacKinnon (1993), *Estimation and Inference in Econometrics*, Oxford University Press, Oxford.

Franses, P. H. (1994), A Multivariate Approach to Modeling Univariate Seasonal Time Series, *Journal of Econometrics*, **63**, 133–151.

Franses, P. H. and R. Paap (1994), Model Selection in Periodic Autoregressions, *Oxford Bulletin of Economics and Statistics*, **56**, 421–440.

Franses, P. H. and R. Paap (2004), *Periodic Time Series Models*, Oxford University Press, Oxford.

Ghysels, E. and D. R. Osborn (2001), *The Econometric Analysis of Seasonal Time Series*, Cambridge University Press, Cambridge.

Herwartz, H. (1999), Performance of Periodic Time Series Models in Forecasting, *Empirical Economics*, **24**, 471–301.

Hsiao, C. (2003), *Analysis of Panel Data*, Econometric Society Monographs, 2nd edn., Cambridge University Press.

Jones, R. H. and W. M. Brelsford (1967), Time Series with Periodic Structure, *Biometrika*, **54**, 403–407.

Maddala, G. S., R. P. Trost, H. Li, and F. Joutz (1997), Estimation of Short-run and Long-run Elasticities of Engergy Demand from Panel Data using Shrinkage Estimators, *Journal of Business and Economic Statistics*, **15**, 90–100.

Osborn, D. R. (1991), The Implications of Periodically Varying Coefficients for Seasonal Time-Series Processes, *Journal of Econometrics*, **48**, 373–384.

Osborn, D. R. and J. P. Smith (1989), The Performance of Periodic Autoregressive Models in Forecasting Seasonal U.K. Consumption, *Journal of Business and Economic Statistics*, **7**, 117–127.

Swamy, P. A. V. B. (1970), Efficient Inference in a Random Coefficient Regression Model, *Econometrica*, **38**, 311–323.

Wolak, F. A. (1989), Local and Global Testing of Linear and Nonlinear Inequality Constraints in Nonlinear Econometric Models, *Econometric Theory*, **5**, 1–35.