

A periodic long-memory model for quarterly UK inflation

Philip Hans Franses*, Marius Ooms

Economic Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, Netherlands

Abstract

We consider an extension of the fractionally integrated ARIMA(0, d , 0) model for quarterly UK inflation, where we allow the fractional integration parameter d to vary with the season s . This periodic ARFIMA(0, d , 0) model does not only provide an informative in-sample description, it may also be useful for out-of-sample forecasting. The main result is that the integration parameter in the first two quarters is significantly larger than that in the last two quarters.

Keywords: Fractional integration; Seasonal time series; Periodic models

1. Introduction

This paper analyzes the long-memory properties in the conditional mean of the quarterly inflation rate in the United Kingdom. A particularly useful class of models to describe such properties is the autoregressive fractionally integrated moving average model (ARFIMA), introduced by Granger and Joyeux (1980) and Hosking (1981). Its key feature is the (fractional) differencing filter $(1 - B)^d$, where B is the familiar backward shift operator and d can take values other than 0 or 1. Additional references focusing on estimation methods and applications include Geweke and Porter-Hudak (1983) and Cheung (1993). Baillie (1996), Robinson (1994) and Beran (1992a) present useful surveys of theoretical and practical matters, and Hassler and Wolters (1995) document that several inflation

time-series data, including UK inflation, may have long-memory properties that can be described using the ARFIMA model class.

An additional characteristic of the quarterly UK inflation rate, for example, is its marked seasonal pattern. It is current practice in the literature on long-memory models either to remove seasonal fluctuations by including seasonal dummy variables in the ARFIMA models, or to analyze seasonal ARFIMA models as in Porter-Hudak (1990), where the differencing filter is $(1 - B^S)^d$ with S equal to, for example, 4 or 12. These two approaches assume that the dynamic behavior of the data is equal across seasons, i.e. the covariance function is constant over time (see also Ray, 1991, 1993a). There are, however, some indications for the quarterly UK inflation rate that its dynamics vary across the four seasons, i.e. that a so-called periodic model yields a useful description. In fact, Franses and Paap (1994) successfully fit a periodic autoregression (PAR) to the UK inflation data.

* Corresponding author. Tel.: +3110 4081273; fax: +3110 4527746.

In this paper, we investigate whether we can merge these two characteristics, i.e. the presence of long memory and of periodic dynamics, into a novel model. This model amounts to an extension of the standard ARFIMA(0, d , 0) model, in the sense that it allows for periodic variation in d . Hence, we analyze the usefulness of a so-called periodic ARFIMA(0, d_s , 0) (PARFIMA) model for quarterly UK inflation, where d_s indicates that the value of d can vary with season $s = 1, 2, 3, 4$.

The outline of our paper is as follows. In Section 2, we start with a non-periodic analysis of UK inflation by estimating an AR(4) model and an ARFIMA(0, d , 0) model. An application of a test for periodic autocorrelations in the estimated residuals of the AR(4), and of a test for white noise residuals of the ARFIMA model, reveal that the errors of the two models are not white noise. In Section 3, we therefore consider periodic AR and ARFIMA models, and we evaluate these in one-step-ahead and multi-step-ahead forecasting experiments in Section 4. A comparison of in-sample residual variances indicates that the periodic ARFIMA is to be preferred. For out-of-sample forecasting, we find that the PARFIMA model does not appear to be the overall winner when compared with rival forecasting schemes. In Section 5, we conclude with some remarks.

2. Non-periodic models

The data we analyze for parameter estimation concern the quarterly observed inflation rate for the United Kingdom over the period 1955.2–1988.4. We also collect the data for 1989.1–1994.4 and use these for out-of-sample forecasting evaluation. We construct the inflation rate i_t as $i_t = \log(\text{RPI}_t) - \log(\text{RPI}_{t-1})$, $t = 1, 2, \dots, n$, where RPI_t is the Retail Price Index, all items (where the index is constructed by setting the average observation in 1985 equal to 100). These data can be found in various issues of *Economic Trends*. The i_t data are depicted in Fig. 1.

The same data are plotted in skip-sampled




Fig. 1. Quarterly inflation in the United Kingdom, 1955.2–1988.4.

format in Fig. 2. This means that the annual time series per season are displayed.

In Fig. 2 we depict the time series $I_{s,T}$, where $s = 1, 2, 3, 4$ and T is the time index for years. In other words, $I_{s,T}$ corresponds to the observation on inflation in year T in season s . It is clear from Fig. 2 that there is a marked seasonal pattern in UK inflation since $I_{2,T}$ is usually the highest, while $I_{3,T}$ is usually the lowest.

In this section, we construct non-periodic




Fig. 2. Inflation in the United Kingdom, 1955–1988, observed per quarter.

models for i_t , where we assume that the seasonal fluctuations can be described using four seasonal dummy variables.

Fig. 3 displays the autocorrelation function (ACF) and the partial ACF (PACF) of i_t , after regressing i_t on four seasonal dummy variables. It is clear that the autocorrelations do not die out quickly. This may indicate long-memory behavior. If we transform i_t by taking first-order differences $\Delta_1 i_t = i_t - i_{t-1}$, the ACF in Fig. 4 indicates that $\Delta_1 i_t$ may be overdifferenced since the sum of the autocorrelations seems close to -0.5 .

These observations suggest the possible usefulness of ARFIMA-type models for i_t .

A useful initial model for i_t would seem to be an autoregressive model of order 4 (AR(4)):

$$(1 - \alpha_1 B - \alpha_2 B^2 - \alpha_3 B^3 - \alpha_4 B^4) i_t = \mu_1' D_t + \mu_2' D_t^* + \varepsilon_t, \quad (1)$$

since all partial autocorrelations after the fourth lag are small, see Fig. 3. The significance of the partial autocorrelation at lag 15 may be owing to one or more outlying observations. Some preliminary estimation results indicate that these outlying observations are 1975.1, 1975.2 and 1979.3. Since we want our models to have




Fig. 3. (Partial) autocorrelation function of quarterly UK inflation (seasonally demeaned).




Fig. 4. (Partial) autocorrelation function of first differences of quarterly UK inflation (seasonally demeaned).

approximately white noise residuals in order to allow for interpretation of the estimated standard errors, we include three dummy variables for these outlying data points. The D_t in (1) denotes a (4×1) vector including a constant and three seasonal dummies, the D_t^* denotes a (3×1) vector including three dummy variables to capture the outlying observations, μ_1 and μ_2 are (4×1) and (3×1) vectors of parameters, respectively, and ε_t denotes a standard homoskedastic white noise process. Below, we will compare the estimation results with and without the inclusion of these D_t^* dummies, where the first is labeled 'outlier correction' and the second case is labelled 'no outlier correction'. In our forecasting exercise, we will also compare the various models for both cases.

The ordinary least squares (OLS) estimation results for the AR(4) model for i_t , as well as the value of the Akaike Information Criterion (AIC), are given in the first panel of Table 1.

When we include higher-order lags in (1), we obtain increasingly less relevant parameters, both in size and in significance. It should be mentioned that when D_t^* is deleted from (1), we find some evidence of ARCH.

The augmented Dickey–Fuller test statistic for the hypothesis of an integer unit root for i_t (that

Table 1
Key estimation results of various univariate time-series models

Model ^a		Parameter estimates ^b		
		Outlier correction	No outlier correction	
AR(4)	$\hat{\alpha}_1$:	0.382 (0.074)	0.479 (0.088)	
	$\hat{\alpha}_2$:	0.126 (0.080)	0.143 (0.096)	
	$\hat{\alpha}_3$:	0.113 (0.081)	0.051 (0.096)	
	$\hat{\alpha}_4$:	0.109 (0.072)	0.166 (0.087)	
	k :	11	8	
	AIC:	-1239.78	-1192.14	
ARFIMA(0, d , 0)	\hat{d} :	0.483 (0.066)	0.499 (0.066)	
	k :	8	5	
	AIC:	-1240.82	-1194.73	
PAR(1)	$\hat{\phi}_1$:	1.004 (0.164)	1.134 (0.169)	
	$\hat{\phi}_2$:	0.889 (0.118)	1.051 (0.117)	
	$\hat{\phi}_3$:	0.420 (0.081)	0.466 (0.093)	
	$\hat{\phi}_4$:	0.451 (0.103)	0.451 (0.119)	
	k :	11	8	
	AIC:	-1242.20	-1205.09	
PAR(4)	quarter 1:	ϕ_{11} :	0.824 (0.202)	0.971 (0.217)
		ϕ_{12} :	0.165 (0.145)	0.107 (0.167)
		ϕ_{13} :	-0.058 (0.147)	-0.047 (0.174)
		ϕ_{14} :	0.195 (0.196)	0.217 (0.232)
	quarter 2:	ϕ_{21} :	0.543 (0.168)	0.723 (0.189)
		ϕ_{22} :	0.207 (0.229)	0.337 (0.268)
		ϕ_{23} :	0.325 (0.139)	0.195 (0.158)
		ϕ_{24} :	0.010 (0.099)	0.035 (0.117)
	quarter 3:	ϕ_{31} :	0.224 (0.153)	0.184 (0.181)
		ϕ_{32} :	0.373 (0.196)	0.641 (0.226)
		ϕ_{33} :	-0.094 (0.237)	-0.378 (0.275)
		ϕ_{34} :	0.010 (0.130)	0.046 (0.154)
	quarter 4:	ϕ_{41} :	0.365 (0.131)	0.365 (0.155)
		ϕ_{42} :	0.052 (0.150)	0.052 (0.178)
		ϕ_{43} :	-0.166 (0.208)	-0.166 (0.246)
		ϕ_{44} :	0.479 (0.230)	0.479 (0.272)
		k :	23	20
		AIC:	-1247.74	-1205.47
	PARFIMA(0, d_s , 0)	\hat{d}_1	0.709 (0.158)	0.767 (0.154)
		\hat{d}_2	0.676 (0.124)	0.867 (0.145)
		\hat{d}_3	0.271 (0.070)	0.296 (0.068)
		\hat{d}_4	0.433 (0.189)	0.246 (0.087)
		k :	11	8
		AIC:	-1245.78	-1210.03
		$\hat{d}_1 = \hat{d}_2$	0.688 (0.100)	0.821 (0.105)
		$\hat{d}_3 = \hat{d}_4$	0.305 (0.064)	0.279 (0.056)

Table 1 (Contd.)

Model ^a	Parameter estimates ^b	
	Outlier correction	No outlier correction
k :	9	6
AIC:	–1249.11	–1213.83

^a The AR(4) model is given in Eq. (1), ARFIMA(0, d , 0) in (2) when we allow for seasonal means, PAR(1) in (5) and the PARFIMA(0, d_s , 0) model in (6). The ϕ_{si} for the PAR(4) model refer to season s and lag i . The estimation sample is 1955.1 to 1988.4, while some of the first few observations can serve as starting values. All models are checked for the absence of residual autocorrelation and ARCH, and for the presence of normality. Details of the test results as well as the estimates for the various constants and dummy variables can be obtained from the authors. The AIC values are calculated for the last 131 of the 136 observations (since this is the common estimation sample) as $AIC = 131 \log(RSS/131) + 2k$, where k is the number of parameters and RSS is the residual sum of squares. The model with minimum AIC is to be preferred.

^b Standard errors are given in parentheses.

is, for $d = 1$ versus $d = 0$ in $(1 - B)^d$, where three lags of Δi_t have been included in the auxiliary regression, obtains a value of -4.442 , where the 5% critical value is -2.89 . Hence, the test result implies the rejection of the unit root null hypothesis in (1), even though the autocorrelation function in Fig. 3 indicates non-stationary behavior in i_t . This rejection of a unit root corresponds to the results in Fig. 4 where Δi_t seems to be overdifferenced.

2.1. An ARFIMA model

Given the values of the parameters in (1), and given the shape of the ACF in Fig. 3, we investigate whether a simple ARFIMA model can be estimated for the inflation data. The ARFIMA(0, d , 0) process:

$$(1 - B)^d i_t = \varepsilon_t, \quad (2)$$

assumes the validity of the differencing operator $(1 - B)^d$, which is defined by the expansion:

$$\begin{aligned} (1 - B)^d &= 1 - dB - (1/2)d(1 - d)B^2 \\ &\quad - (1/6)d(1 - d)(2 - d)B^3 \dots \\ &\quad - (1/j!)d(1 - d)(2 - d) \\ &\quad ((j - 1) - d)B^j - \dots \end{aligned} \quad (3)$$

(see Hosking, 1981). When $0 < d < 0.5$, the process is said to be long-memory stationary, and when $d \geq 0.5$, the process is said to be non-stationary.

To estimate d , we need an estimation method

that does not require a priori knowledge of stationarity, and that can also be modified to allow for seasonal variation in d as in Section 3. We apply Beran's (1995) approximate maximum likelihood (ML) estimator for invertible, and possibly non-stationary, ARFIMA models. This estimator is based on the least-squares criterion, i.e. on minimizing $\sum_{t=2}^n e_t^2(\eta)$, where

$$e_t(\eta) = \sum_{j=0}^{t-1} \phi_j(\eta)(y_{t-j} - \bar{y}_t), \quad (4)$$

with respect to the parameter vector η , given the observations y_1, \dots, y_n with corresponding sample mean \bar{y}_t , where $e_1 = y_1 - \bar{y}_t$ clearly does not depend on η (see Beran, 1995). The $\phi_j(\eta)$ are the coefficients of the AR(∞) representation of the model, see (3). This estimator is asymptotically efficient if the innovations ε_t are Gaussian and it is root- n consistent under milder conditions. In the case of an ARFIMA(0, d , 0) model, the parameter vector η consists of d only.

We subtract the seasonal means $\bar{y}_{t,s}$ from the data. Note that this corresponds to including D_t in model (1). Next, we minimize the residual sum of squares of the AR(∞) representation. We use the concentrated version of this approximate log likelihood function for y_1 to y_n , i.e. $\sum_{t=2}^n e_t^2(\eta)/n$, to compute the standard errors. These are based on the numerical Hessian. Furthermore, we robustify the estimates of the seasonal means by replacing outliers $I_{s,T}$ by $(I_{s,T-1} + I_{s,T+1})/2$. We robustify the estimate of d by deleting the residuals for the outlying observations from the criterion function.

The application of this method to the UK inflation data (where we set $i_1 = 0$) yields the estimation results as presented in the second panel of Table 1. To compare this model with the periodic ARFIMA model below, we again calculate the AIC. It is clear that this ARFIMA model outperforms the AR(4) model on in-sample fit for both outlier correction and no outlier correction.

For the case without outlier correction, the exact Gaussian ML estimation method of Sowell (1992) on the seasonally demeaned data gives an estimate of 0.453 for d with a standard error of 0.044. The approximate frequency domain Gaussian ML method of Boes et al. (1989) results in an estimate of 0.510 with standard error 0.068. Comparing this with the 0.499 (0.066) in Table 1, we observe that for all three estimation methods it is clear that the hypothesis $d = 0.5$ cannot be rejected at any reasonable significance level. This result appears robust to the three outlying observations.

2.2. Periodicity in residual autocorrelations

Before we start estimating a periodic extension of the ARFIMA model in (2), we investigate the residual processes for the presence of systematic, perhaps seasonal, patterns. In fact, Franses and Paap (1994) find that the UK inflation displays periodic autoregressive behavior. Additionally, in Fig. 5 we give the periodic ACF; that is, an

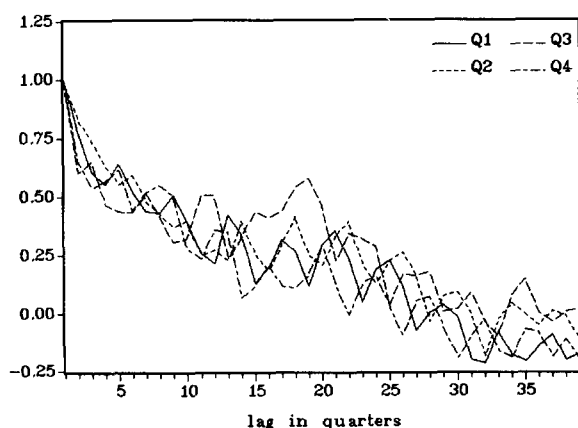


Fig. 5. Periodic autocorrelation function for quarters 1 to 4. (The lags are indicated on the horizontal axis.)

ACF for each of the quarters, which is calculated using the method proposed in McLeod (1994).

It is clear from this graph that the first-order autocorrelation is about 0.8 for quarters 1 and 2, and about 0.6 for quarters 3 and 4. Furthermore, the ACF values die out only very slowly in all quarters. In sum, the models in (1) and (2) may be misspecified with respect to capturing seasonal fluctuations. A further simple diagnostic is to perform an LM-type regression by adding $D_{s,t}\hat{\varepsilon}_{t-1}$ to (1), with $\hat{\varepsilon}_0 = 0$, where $D_{s,t}$ are four seasonal dummies $s = 1, 2, 3, 4$, and to evaluate the joint significance of $D_{s,t}\hat{\varepsilon}_{t-1}$ (see Franses, 1993). For model (1) this yields an F -statistic value of 3.755, which is significant at the 5% level. It seems likely that model (2) also does not pass this test since it is closely related to model (1). This seems to be confirmed by Beran's (1992b) test for white noise residuals in (2) (when estimated with the Boes et al. (1989) method), with obtains a p -value of 0.014.

We interpret these results as tentative evidence that the time series itself and the residuals from non-periodic models (1) and (2) display periodic autocorrelations. This seems to call for a modification of (2) that allows for periodic variation in the parameter d . In the next section, we estimate such a periodic ARFIMA model for the UK inflation data.

3. Periodic models

We start this section with a simple first-order periodic autoregression (PAR(1)):

$$i_t = \phi_s i_{t-1} + \mu_1' D_t + \mu_2' D_t^* + \varepsilon_t, \quad (5)$$

where the AR parameters take a value of ϕ_1 in season 1; that is, when $t \bmod 4 = 1$ for $t = 1, 2, \dots, n$, and values of ϕ_2 , ϕ_3 and ϕ_4 when $t \bmod 4$ equals 2, 3 and 0, respectively. The OLS estimation results of (5) appear in the third panel of Table 1. As could be expected, the ϕ_s parameters differ across the seasons.

It should be mentioned here that residual autocorrelation tests indicate that there may be some residual autocorrelation that can be re-

moved by enlarging model (5). In fact, a PAR(4) model turns out to have better empirical properties. An $F_{12,111}$ -test to compare the PAR(4) with the PAR(1) obtains a value of 3.526, which is significant at the 1% level, and which indicates that (5) is too rough an approximation. We report the parameter estimates for the PAR(4) model in Table 1. Its AIC values show that it clearly outperforms all previous models within the estimation sample.

Obviously, the first-order autoregressive parameters ϕ_s in (5) are not equal, and hence, although (5) can be improved by including additional lags, the estimation results for this PAR(1) model indicate that i_t has periodic dynamics. This confirms the findings for the test for periodicity in the residuals in the previous section. To investigate whether the periodicity in ϕ_s in (5) is caused by neglected seasonal heteroscedasticity, we regress the $\hat{\varepsilon}_t^2$ from the PAR(1) and PAR(4) models on a constant and three seasonal dummies. The F -versions of the TR^2 test statistics obtain values of 0.780 and 0.833, indicating that the null hypothesis of seasonal homoscedasticity cannot be rejected. Hence, there seems no need for our periodic models to resort to weighted least squares to allow for seasonality in the variances.

3.1. A periodic ARFIMA(0, d_s , 0) model

Given the observed long-memory characteristics in the previous section and the periodic patterns in (5), it seems worthwhile to consider the periodic ARFIMA(0, d_s , 0) model,

$$(1 - B)^{d_s} i_t = \varepsilon_t, \quad (6)$$

i.e. a model similar to (2) where now both the mean of i_t and the value of d can vary with the seasons.

To estimate the values of d_s , we use a periodic version of Beran's (1995) approximate ML estimator, which corresponds to minimizing $\sum_{t=2}^n e_t^2(\eta)$, where now

$$e_t(\eta) = \sum_{j=0}^{t-1} \phi_{j,s}(\eta)(y_{t-j} - \bar{y}_{t,s}), \quad (7)$$

i.e. we allow for periodic AR parameters $\phi_{j,s}$ and we subtract seasonal means $\bar{y}_{t,s}$. Again, e_1 does not depend on η . In our case of UK inflation (where we again set y_1 at 0; that is, we use 136 observations on i_t), this method results in the parameter estimates given in the fifth panel of Table 1. The AIC value in the case of no outlier correction is -1210.03 . The latter value indicates that (6) improves upon the non-periodic AR-FIMA model in (2) with an AIC value of -1194.71 .

The estimation results in (6) suggest that this model can be simplified somewhat further by imposing that $d_1 = d_2$ and $d_3 = d_4$. The relevant estimation results under these restrictions, using the periodic version of the method in Beran (1995), are given in the final panel of Table 1. This model obtains an AIC value of -1213.83 in the no outlier correction case, indicating a further improvement upon (2) and (6). The significant differences between the d_s should not be interpreted as indicating seasonally varying long-memory behavior. In fact, the autocorrelations at high lags do not vary systematically over the seasons, see Fig. 5. The long-memory behavior is not clearly periodic.

4. Forecasting

To investigate whether the PARFIMA model outperforms rival models, we carry out a forecasting exercise for the 24 quarterly observations in 1989 to 1994. As alternative models, we consider the non-periodic ARFIMA model in (2), the unrestricted PAR(4) model, the non-periodic AR(4) model in (1) and the PAR(1) model in (5). The parameters in each model are estimated for the sample until 1988.4, and one-step and k -step ahead (for $k = 1, \dots, 24$) forecasts are generated for 1989.1–1994.4.

Given the true parameter values, there are two ways to generate forecasts from an AR-FIMA model. The first is the naive method, which involves using the $AR(\infty)$ approximation for the series after correction for deterministic terms, where data points before the sample period are set equal to zero (see Peiris and

Perera, 1988; and Ray, 1993b). The second method is to generate the optimal linear forecasts for k periods ahead given n observations. The method is explained in Beran (1994, pp. 164–165). The optimal finite sample method requires an expression for the exact autocorrelations. We use this method only for the non-periodic ARFIMA model. The choice of forecasting method does not affect the general result and we report results for the naive method only. We generate forecasts for the series in deviation from seasonal means, which are estimated over the sample until 1988.4. We add these means to the final forecasts in the last stage.

The forecasting results are summarized in Tables 2 and 3. The differences between the root mean squared errors (RMSEs) of the various models are not large. The PARFIMA model does not appear to outperform the non-periodic models or periodic AR models both for one-step-ahead and multi-step-ahead forecasts. There is no clear winner amongst the five rival models

for the period 1989 to 1994, although the PAR(4) and AR(4) models seem to do rather well. These results seem to add to those in Crato and Ray (1996), where it is found that ARFIMA models have difficulties in improving upon rival ARIMA-type models when forecasting out of sample. Another possibility is that there has been some structural shift in the evaluation period. The relatively large effect of the outlier correction on the multi-period forecast errors favors the latter explanation. The reduction of the effect of extremely high inflation values in the 1970s leads to better forecasts for the low inflation values in the 1990s.

5. Concluding remarks

In this paper, we extended the ARFIMA-(0, d , 0) model by allowing the parameter d to vary with the seasons. An approximate PARFIMA(0, d_s , 0) model was fitted to the

Table 2

Forecast evaluation of PARFIMA and other models for quarterly UK inflation: One-step-ahead forecasts. The models are estimated for the sample until 1988.4, and one-step-ahead forecasts are generated for 1989.1–1994.4. The cells are the Root Mean Squared Errors

Model	All quarters	Q1	Q2	Q3	Q4
Outlier correction					
AR(4)	0.669	0.887	<u>0.857</u>	<u>0.275</u>	<u>0.439</u>
PAR(1)	0.705	0.607	1.023	0.419	0.630
PAR(4)	<u>0.648</u>	<u>0.590</u>	0.975	0.302	0.537
ARFIMA(0, d , 0)	0.686	0.808	0.880	0.376	0.560
PARFIMA(0, d_s , 0)	0.693	0.693	0.952	0.377	0.627
No outlier correction					
AR(4)	<u>0.646</u>	0.839	<u>0.796</u>	0.374	<u>0.440</u>
PAR(1)	0.745	0.596	1.080	0.542	0.630
PAR(4)	0.675	<u>0.584</u>	1.029	<u>0.363</u>	0.537
ARFIMA(0, d , 0)	0.665	0.820	0.835	0.413	0.481
PARFIMA(0, d_s , 0)	0.688	0.665	0.956	0.441	0.585

Note: All values are multiplied with 100. Underlined are the minimum values of the RMSE over the columns. Note that the restrictions $d_1 = d_2$ and $d_3 = d_4$ are imposed on the parameters of the PARFIMA model.

Table 3

Forecast evaluation of PARFIMA and other models for quarterly UK inflation for forecasts 1 to k steps ahead. The models are estimated for the sample until 1988.4, and one-step-ahead to k -step-ahead forecasts are generated for 1989.1–1994.4. The cells are the Root Mean Squared Errors (averaged over the k forecasts)

Model	All quarters	Q1	Q2	Q3	Q4
Outlier correction					
AR(4)	<u>0.886</u>	<u>1.168</u>	0.962	<u>0.562</u>	<u>0.731</u>
PAR(1)	0.909	1.184	0.940	0.610	0.805
PAR(4)	0.909	1.222	<u>0.914</u>	0.605	0.782
ARFIMA(0, d , 0)	0.980	1.220	1.015	0.693	0.917
PARFIMA(0, d_s , 0)	0.982	1.225	1.008	0.697	0.924
No outlier correction					
AR(4)	1.071	1.407	1.018	0.813	0.951
PAR(1)	<u>0.998</u>	1.307	<u>0.958</u>	<u>0.757</u>	<u>0.912</u>
PAR(4)	1.087	1.451	1.013	0.854	0.940
ARFIMA(0, d , 0)	1.007	<u>1.265</u>	1.018	0.770	0.912
PARFIMA(0, d_s , 0)	1.013	1.275	1.010	0.780	0.922

All values are multiplied with 100. Underlined are the minimum values of the RMSE over the columns. Note that the restrictions $d_1 = d_2$ and $d_3 = d_4$ are imposed on the parameters of the PARFIMA model.

quarterly UK inflation data. To this end, we used a periodic extension of Beran's (1995) estimation approach. We found that for the first half of the year, the d value is significantly larger than 0.5, while for the second part it is significantly smaller. Finally, with respect to forecasting, the differences between the various models are not large, although the PARFIMA model is usually outperformed by rival forecasting models.

An interesting topic of further research is to obtain more insight into the periodic autocovariance function of the PARFIMA model. This would be helpful for deriving standard errors of the estimated seasonal means and for the computation of periodic forecast error variances. It is also important to study the memory characteristics, or stationarity properties, of time series generated by PARFIMA processes. Finally, we may consider the effects of using a non-periodic model in the case of a PARFIMA model as a data-generating process. See Tiao and

Grupe (1980) and Boswijk and Franses (1996) for related results for $I(0)$ and $I(1)$ processes.

Acknowledgements

An earlier version of this paper was presented at the 14th ISF conference in Stockholm (June 1994). Helpful comments from the participants, Christiaan Heij, Bonnie Ray and three anonymous referees are gratefully acknowledged. We thank Falaw Sowell for making available a program for exact ML-estimation. The first author thanks the Royal Netherlands Academy of Arts and Sciences for its financial support.

References

- Baillie, R.T., 1996, Long memory processes and fractional integration in econometrics, *Journal of Econometrics*, 73, 5–60.

- Beran, J., 1992a, Statistical methods for data with long-range dependence, *Statistical Science*, 7, 404–416.
- Beran, J., 1992b, A goodness-of-fit test for time series with long range dependence, *Journal of the Royal Statistical Society B*, 54, 749–760.
- Beran, J., 1994, *Statistics for Long Memory Processes* (Chapman and Hall, London).
- Beran, J., 1995, Maximum likelihood estimation of the differencing parameter for invertible short and long memory autoregressive integrated moving average models, *Journal of the Royal Statistical Society B*, 57, 654–672.
- Boes, D.C., R.A. Davis and S.N. Gupta, 1989, Parameter estimation in low order fractionally differenced ARMA processes, *Stochastic Hydrology and Hydraulics*, 3, 97–110.
- Boswijk, H.P. and P.H. Franses, 1966, Unit roots in periodic autoregressions, *Journal of Time Series Analysis*, 17, 221–245.
- Cheung, Y-W., 1993, Long memory in foreign exchange rates, *Journal of Business and Economic Statistics*, 11, 93–101.
- Crato, N. and B.K. Ray, 1996, Model selection and forecasting for long-range dependent processes, *Journal of Forecasting*, 15, 107–125.
- Franses, P.H., 1993, Periodically integrated subset autoregressions for Dutch industrial production and money stock, *Journal of Forecasting*, 12, 601–613.
- Franses, P.H. and R. Paap, 1994, Model selection in periodic autoregressions, *Oxford Bulletin of Economics and Statistics*, 56, 421–439.
- Geweke, J. and S. Porter-Hudak, 1983, The estimation and application of long-memory time series models, *Journal of Time Series Analysis*, 4, 221–238.
- Granger, C.W.J. and R. Joyeux, 1980, An introduction to long-memory time series models and fractional differencing, *Journal of Time Series Analysis*, 1, 15–39.
- Hassler, U. and J. Wolters, 1995, Long memory in inflation rates: International evidence, *Journal of Business and Economic Statistics*, 13, 37–46.
- Hosking, J.R.M., 1981, Fractional differencing, *Biometrika*, 68, 165–176.
- McLeod, A.I., 1994, Diagnostic checking of periodic autoregression models with application, *Journal of Time Series Analysis*, 15, 221–233.
- Peiris, M.S. and B.J.C. Perera, 1988, On prediction with fractionally differenced ARIMA models, *Journal of Time Series Analysis*, 9, 215–220.
- Porter-Hudak, S., 1990, An application of the seasonally fractionally differenced model to the monetary aggregates, *Journal of the American Statistical Association*, 85, 338–344.
- Ray, B.K., 1991, Fractionally differenced ARMA processes: Seasonality and forecasting issues, Ph.D. thesis, Columbia University, New York City.
- Ray, B.K., 1993a, Long-range forecasting of IBM product revenues using a fractionally differenced ARMA model, *International Journal of Forecasting*, 9, 255–269.
- Ray, B.K., 1993b, Modeling long-memory processes for optimal long-range prediction, *Journal of Time Series Analysis*, 14, 511–526.
- Robinson, P.M., 1994, Time series with strong dependence, in: C. Sims, ed., *Advances in Econometrics, Sixth World Congress*, Volume 1 (Cambridge University Press, Cambridge, UK).
- Sowell, F., 1992, Maximum likelihood estimation of stationary univariate fractionally integrated time series models, *Journal of Econometrics*, 53, 165–188.
- Tiao, G.C. and M.R. Grupe, 1980, Hidden periodic autoregressive moving average models in time series data, *Biometrika*, 67, 219–228.

Biographies: Philip Hans FRANCES is Associate Professor in Econometrics at the Econometric Institute, and Director of Research at the Rotterdam Institute for Business Economic Studies, both of the Erasmus University Rotterdam. His research interests include modelling and forecasting seasonal time series. On this topic, he has published in various journals and in *Periodicity and Stochastic Trends in Economic Time Series* (Oxford University Press, Oxford).

Marius OOMS is Assistant Professor in Econometrics at the Econometric Institute of the Erasmus University Rotterdam. He graduated in econometrics in 1985 at the University of Amsterdam. He received his Ph.D. in 1993 on a thesis entitled 'Empirical Vector Autoregressive Modelling'. His current research interests include empirical modelling of seasonality, and long memory and testing for parameter stability in time-series models.