SEASONALITY, COINTEGRATION, AND THE FORECASTING OF ENERGY DEMAND

Michael P. Clements and Reinhard Madlener

No.484

WARWICK ECONOMIC RESEARCH PAPERS

THE UNIVERSITY OF WARWICK

DEPARTMENT OF ECONOMICS

SEASONALITY, COINTEGRATION, AND THE FORECASTING OF ENERGY DEMAND

Michael P. Clements and Reinhard Madlener
University of Warwick

No.484

September 1997

Seasonality, Cointegration, and the Forecasting of Energy Demand

MICHAEL P. CLEMENTS[1]
University of Warwick, Department of Economics
Coventry CV4 7AL, U.K.
Tel: (++44-1203) 522 990
Fax: (++44-1203) 523 032

REINHARD MADLENER[2]
University of Warwick, ESRC Macroeconomic Modelling Bureau
Coventry CV4 7AL, U.K.
Tel: (++44-1203) 523 277
Fax: (++44-1203) 523 032

September 1997

[1]The author gratefully acknowledges financial research support from the ESRC Macroeconomic Research Consortium (ESRC Grant No. L116251015).

[2]The author is grateful for research funding by the European Commission via the Human Capital and Mobility (HCM) Programme (Contract No. 940613a). The empirical results and graphs reported in this paper were produced using GiveWin and PcGive and PcFiml for Windows: see Doornik and Hendry (1996,1997), and TSP, v4.3A, TSP International, Palo Alto, California.

Abstract

Much of the short–run movement in energy demand in the UK appears to be seasonal, and the contribution of long–run factors to short–run forecasts is slight. Nevertheless, using a variety of techniques, including a recently developed test that is applicable irrespective of the orders of integration of the data, we obtain a long–run income elasticity of demand of about one third, and we are unable to reject a zero price elasticity.

Periodic models that allow for a seasonally varying response of demand to its principle determinants are shown to provide superior short–run forecasts to well–known seasonal time series models, and to univariate periodic models, even *ex ante*, when the determinants themselves have to be forecast. However, the relatively short data sample and small number of forecasts suggests caution in generalising these results.

Keywords: energy demand, seasonality, cointegration, forecasting;

JEL Classification numbers: C52, Q41;

1 Introduction

Over the past decade a number of empirical papers in energy economics have applied cointegration analysis to the modelling of energy demand, thereby accounting for the potential non–stationarity of the data and simultaneously avoiding the loss of valuable long–run information which would result from taking first differences (Table 1). The majority of the studies have employed the two–step procedure of Engle and Granger (1987), where an estimate of the cointegrating relationship is obtained by an ordinary least squares (OLS) regression of the contemporaneous values of the variables. Some have used the potentially more efficient systems approach of Johansen (1988, 1991), which amounts to reduced–rank regression by maximum likelihood (ML) estimation. That the estimates from the 'static regression' of Engle and Granger (EG) may be seriously biased in small samples is well known (see, for example, Hargreaves, 1994, p.88). Equally, the estimates from the Johansen procedure may be difficult to interpret when the cointegrating rank is estimated at greater than one, apart from the greater informational requirements of estimating a vector–autoregression (VAR). For these reasons, we also consider a third approach which asks a slightly different question: is there a long–run relationship between the variables of interest, recognising that this does not pre–suppose that all the variables are integrated of order one, $I(1)$. This is the approach due to Pesaran, Shin, and Smith (1996). It has several advantages for the modelling of energy demand. Firstly, 'temperature' or 'heating degree days' is often found to be an important explanatory variable, so an approach that is agnostic about the orders of integration of the explanatory variables should prove useful. Secondly, the approach requires that there is one 'fundamental' long–run relationship between the variables, which enters the equation for the variable of interest (here, energy demand), but not the equations for the other variables (temperature, income, etc.), in which case we can estimate and test for the long–run relationship using OLS in a single–equation analysis. In the case of modelling energy demand, this assumption appears reasonable. There are few grounds for believing that the main explanatory variables adjust to disequilibrium in the long–run relationship governing the demand for energy.

As well as comparing alternative methods of estimating the long–run relationship for energy demand, we also consider the usefulness of annual as opposed to quarterly data for

Table 1: Energy demand studies using cointegration techniques (selection)

study	sector/area covered	technique used	data used	estimated ECM adj. coeff./remarks
Engle et al. (1989)	regional commercial (electricity only)	EG 2-step	US monthly data (1975M1–85M5) [125 obs.]	$-1.13/-1.01,^{\circ}$ (highly) significant
Hunt and Manning (1989)	aggregate	EG 2-step	UK annual data (1967–86) [20 obs.]	−0.67, significant
Hunt and Lynk (1992)	industrial	EG 2-step	UK annual data (1948–88) [41 obs.]	−0.80, (highly) significant
Bentzen and Engsted (1993)	aggregate	EG 2-step	Danish annual data (1948–90) [43 obs.]	$-0.24/-0.11,^{*}$ significant (for full sample only)
Engsted and Bentzen (1993)	aggregate	EG 2-step/Johansen/ LQAC model	Danish annual data (1900–91) [92 obs.]	$-0.27/-0.18,^{\dagger}$ (highly) significant
Vaage (1993)	residential (electricity only)	Johansen	Norwegian annual data (1960–89) [30 obs.]	−0.40
Bentzen (1994)	transport (petrol only)	EG 2-step	Danish annual data (1948–91) [44 obs.]	$-0.67/-0.58,^{\ddagger}$ signif.
Fouquet (1995)	residential	EG 2-step	UK quarterly data (1974Q1–94Q1) [84 obs.]	$-0.37/-0.82/-0.71/-0.94,^{\diamond}$ (highly) significant; (mean air temp. included)
Hunt and Witt (1995)	aggregate	Johansen	UK annual data (1967–94) [28 obs.]	−0.65, (highly) significant (average January air temp. included in order to find cointegration)
Madlener (1996b)	residential	Johansen	Austrian annual data (1970–93) [24 obs.]	−0.78; (heating degree days included)

(NOTES: 'EG 2-step' refers to the cointegration analysis procedure suggested by Engle and Granger (1987), 'Johansen' to the multivariate ML approach to cointegration introduced by Johansen (1988, 1991). LQAC stands for 'linear quadratic adjustment cost model'. $^{\circ}$ The first figure is for the error–correction model using monthly long–run forecasts, the second for the one using annual long–run forecasts. * The former value refers to the full sample outcome, the latter to a 1974–90 subsample estimation. † The former value refers to the full sample outcome, the latter to a 1948–91 subsample. ‡ The former value refers to the full sample outcome, the latter to a 1974–91 subsample. $^{\diamond}$ The various values reported are the estimated error–correction adjustment coefficients for changes in coal/petrol/gas/electricity consumption per household, respectively.)

empirical modelling. Due to the rather limited availability of historic time series data for such modelling exercises, the number of annual observations available is often quite small (see Table 1). Nevertheless, the use of higher frequency data tends to be the exception rather than the rule. In particular, among the studies reported, only Engle, Granger, and Hallman (1989) and Fouquet (1995) have used sub–annual data, and only the former has investigated some of the implications of strong seasonal patterns for forecasting (the latter was concerned with the impact of tax increases on energy demand).

By way of contrast, in the more mainstream econometrics literature there has been a great deal of interest in the analysis and modelling of seasonality (see Hylleberg, 1992; Franses, 1996, for a discussion of the concepts and the literature involved), though little of this has as yet filtered down to the energy literature (Madlener, 1996a, provides a recent survey of the methodologies used in the energy demand literature, focusing on studies of residential energy demand). In order to redress this imbalance, in this paper we estimate and generate short–term forecasts from a number of econometric models of quarterly energy demand, which treat seasonality in different ways, and compare the results with forecasts from traditional univariate time–series models.

The contribution of this paper is at least threefold: (i) we compare alternative ways of estimating the long–run relationship for energy demand, and compare some of these methods on annual and quarterly data; (ii) we formulate short–term forecasting methods on the quarterly data; and (iii) we assess the forecasting performance of the econometric and rival time–series models.

The structure of the paper is as follows: Section 2 investigates the time series properties of the time series. Section 3 reports on the multivariate cointegration analysis and the approach of Pesaran, Shin, and Smith (1996). Section 4 describes our econometric model specifications, and Section 5 the time series models. Section 6 concludes.

2 Analysing the time series properties of the data

Four variables are used in this study (all in logarithms): domestic energy consumption, q, real disposable income, y, real energy price, p, and the temperature variable 'heating degree days', h. The data samples used are from 1975Q4 to 1996Q3 for the quarterly data and from 1976 to 1995 for the annual data (for a detailed data description see Appendix 1).

2.1 Annual data

Although we only have twenty observations we decided to test for the order of integration of the four data series using the (Augmented) Dickey–Fuller tests. Perhaps not unsurprisingly given the small sample period, the tests were in some cases unduly sensitive to the deterministic terms (constant, linear trend) included in the regression. When we ran a regression of the form

$$\Delta x_t = \mu + \gamma t + \phi x_{t-1} + \sum_{j=1}^{n} \beta_j \Delta x_{t-j} + \epsilon_t \tag{1}$$

for q and h (with n set to a low number to preserve degrees of freedom), ϕ was clearly significantly different from zero, using the appropriate critical values. Hence q and h both appear to be stationary around a deterministic trend. This appears somewhat implausible, at least for h, and omitting the γt term from the above regression indicated that both variables are $I(1)$. For both p and y we failed to reject $H_0 : \phi = 0$ in Eq.(1). We also failed to reject $H_0 : \phi = \gamma = 0$, and then testing $H_0 : \phi = 0$ in Eq.(1) without the γt term also failed to reject the null of $I(1)$. Given the small sample size it seemed inadvisable to test the variables for higher orders of integration. Consequently, y and p appear to be $I(1)$, which accords with the general body of evidence, at least for y. That the evidence for q and h is less decisive is consistent with the fact that it can be very difficult to distinguish between a trend–stationary and a unit root process in finite samples (see, e.g., Campbell and Perron, 1991, p.157). After examining the quarterly series, we carry out a cointegration analysis on the annual data to provide a rough check on the analysis on the quarterly data.

2.2 Quarterly data

The analysis of the quarterly data series requires that we find an adequate characterisation of the strong seasonal patterns evident in q and h, and to a lesser extent in y (see Figure 1). Many of the studies on annual data have found unit roots in the series, i.e., the series appear to be $I(1)$. We wish to allow for the possibility that the seasonality in the series we analyse is stochastic and non–stationary, that is, there are unit roots at the seasonal frequencies. A procedure for testing for roots at the seasonal frequencies (and also at the long–run, or zero frequency) has recently been developed by Hylleberg, Engle, Granger, and Yoo (1990) (HEGY), and we apply this to the four time series used in the study. The plots of the series suggest that p is not seasonal, let alone a non–stationary stochastic seasonal process, and moreover, it makes little sense for h to contain roots at any frequency. Nonetheless, for completeness we report results for testing all the series.

In order to interpret the results of the unit root tests, a brief outline of the testing procedure suggested by HEGY is given. The procedure is based upon the regression equation

$$\alpha(L)\Delta_4 x_t = \pi_1 z_{1,t-1} + \pi_2 z_{2,t-1} + \pi_3 z_{3,t-2} + \pi_4 z_{3,t-1} + \mu_t + \epsilon_t, \tag{2}$$

where

$$\mu_t = \mu_1 + \mu_2 t + \sum_{i=1}^{3} \mu_{2+i} Q_{i,t},$$

$\alpha(L)$ is a stationary r^{th} order lag polynomial, and

$$\begin{aligned} z_{1t} &= (1 + L + L^2 + L^3)x_t = S(L)x_t \\ z_{2t} &= -(1 - L)(1 + L^2)x_t \\ z_{3t} &= -(1 - L^2)x_t = -(1 - L)(1 + L)x_t. \end{aligned}$$

The tests for unit roots at different frequencies are based on the significance of the estimated $\hat{\pi}_i$ coefficients. For example, if $\hat{\pi}_1$ is significantly different from (typically being less than) zero, then the null of a unit root at the zero frequency is rejected. A rejection of the null that π_2 equals zero indicates the absence of a unit root at the bi–annual frequency. Roots at the annual frequency can be tested by an F-test of the null that $\pi_3 = \pi_4 = 0$. The appropriate critical values are given in Hylleberg et al. (1990, pp.226–227, Tables 1a and 1b).

The choice of n, the order of the polynomial $\alpha(L)$, has to be made with care. Under–fitting (n too small) can result in the unit root test statistics being substantially over–sized, so that we may fail to find roots that are a feature of the data generating process (DGP). Conversely, over–fitting may reduce the power of the test statistics, leading to the discovery of spurious roots. Taylor (1997) considers a number of strategies that have been proposed in the literature for determining the appropriate augmentation, such as general–to–simple and

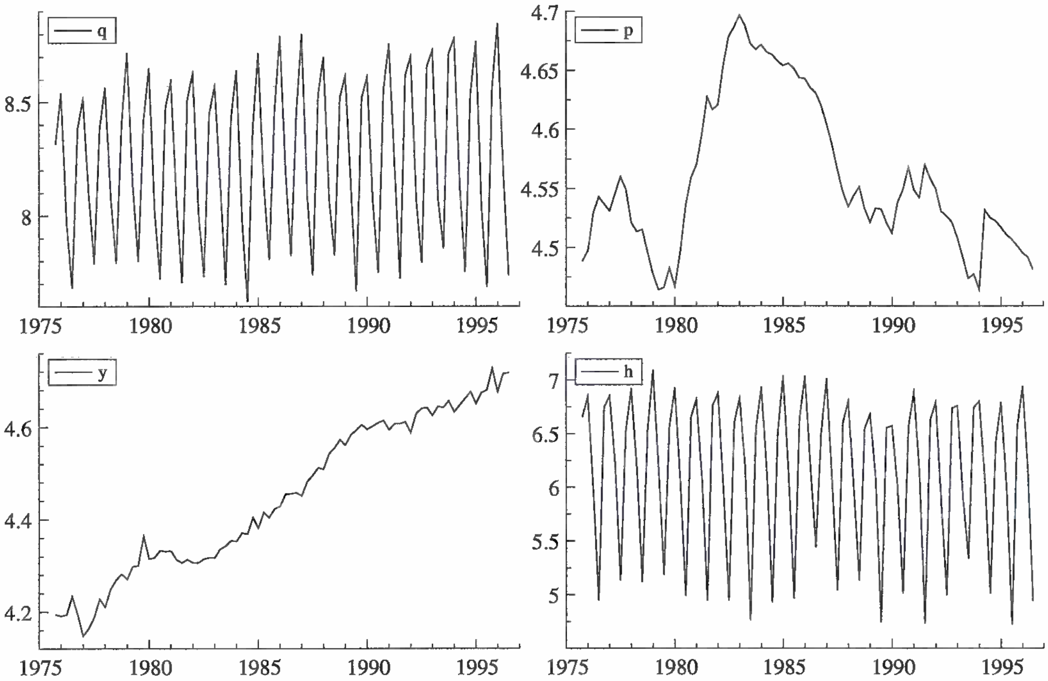

Figure 1: Quarterly data plots (all series in logs)

information criteria. In both cases he shows that there may be a tendency to underfit, and this will be exacerbated when deterministic variables (e.g., constant, trend and seasonals) are included. Our approach is to begin with $n = 5$ and consider the values of the test statistics for all values of n down to zero. To save space, in Table 2 we only report the results for the value of n obtained by sequentially simplifying the model by deleting the longest lag whenever its t-value is less than 2 in absolute terms. But we also signal those instances when inference appears to be sensitive to the lag order.

The results in Table 2 suggest that q, p and y have a zero–frequency root. Again, the finding of a zero–frequency root for h as well is puzzling, and we explore this further below. None of the series appear to have roots at the seasonal frequencies, except perhaps for q at the bi–annual frequency when we keep in the lag with length three, $n = 3$, although the sequential simplification procedure suggests that a zero augmentation is acceptable. Nonetheless, as noted above, such procedures tend toward an under–fitting of the lagged fourth–difference terms.

Figure 2 depicts UK domestic energy consumption, plotted for each of the quarters separately. It is apparent that the underlying trend in energy consumption has been positive in the first and fourth quarters, and the fourth–quarter quarterly growth rate is the most pronounced.

3 Analysis of the long–run relationship

3.1 The Johansen ML procedure

Because of the two major drawbacks of the two–step EG procedure (i.e. the possibility of seriously biased estimates in small samples and the requirement that all variables are $I(1)$), we refrained from using their approach. However, we tested the quarterly and annual data for cointegration using the maximum likelihood approach of Johansen (1988, 1991). For the quarterly data we began with a third–order vector–autoregression (VAR), and for the annual data a second–order system (as a compromise between whitening the residuals and saving on degrees of freedom). In both instances q, y and p were treated as endogenous, while h was assumed to be exogenous (entered unrestrictedly). Moreover, we included a

Table 2: Testing for roots at the zero and seasonal frequencies

lag augmentation	t_{π_1}	t_{π_2}	t_{π_3}	t_{π_4}	F_{π_3,π_4}	t-values of lagged depend. variables
			q			
lag 3 only	−3.52	−2.82	−4.11*	−1.58	10.48**	1.72
0	−3.03	−3.59*	−4.42**	−0.92	10.33**	
			p			
0	−1.77	−5.13**	−4.58**	−5.57**	42.07**	
			y			
lags 1 and 2	−2.81	−3.82**	−4.30**	−2.59*	15.11**	$-0.40, 2.51$
			h			
0	−3.01	−3.86**	−5.12**	1.05	14.25**	
			$(q - 0.485h - 5.277)$			
0	−3.70*	−4.95**	−4.26**	0.12	9.07**	

(NOTES: * denotes significance at the 5% level, ** at the 1% level. Deterministic terms included: constant, seasonals, non–seasonal trend. To see whether the inference of a zero frequency root for h was due to low power of the test from the inclusion of the (unnecessary) linear trend, the regression was run again omitting the trend term. The inferences were unaltered. Moreover, dropping the first lag of the fourth difference for y did not change the test outcomes. The last entry in the table refers to the regression reported in Section 4, Table 5.)

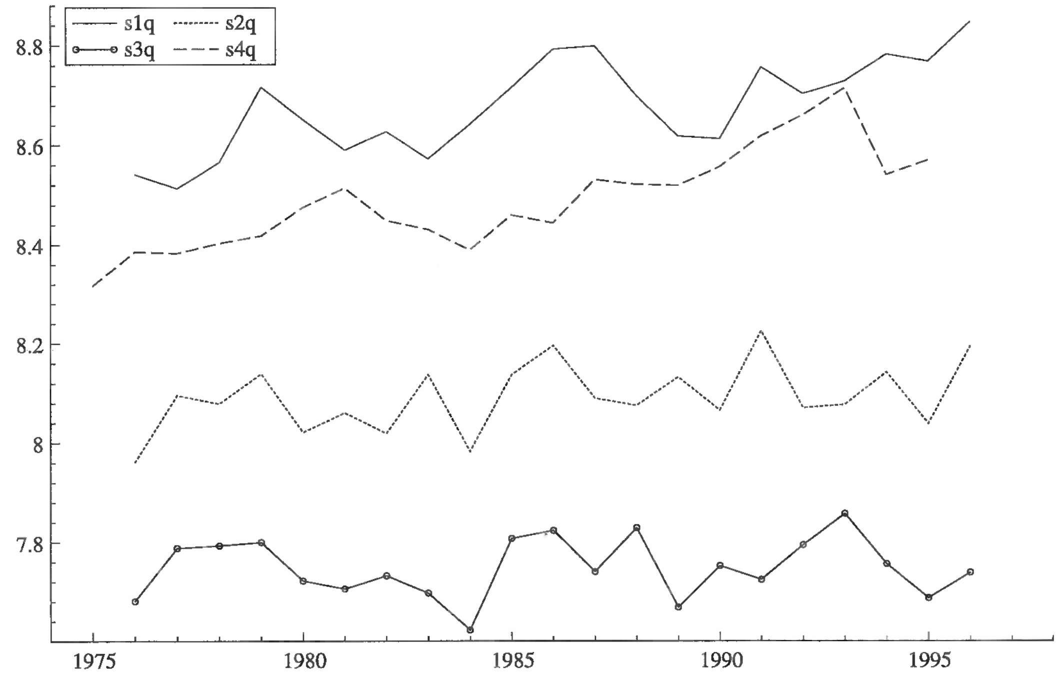

Figure 2: UK domestic energy consumption, plotted for each quarter for 1975–96 (in logs)

deterministic trend, restricted to enter the long run only. As it turned out, the order of both models could be reduced by one, and the deterministic trend was insignificant in both cases.

The results for the quarterly data depicted in Table 3 indicate a single cointegrating vector on the basis of the maximum eigenvalue and trace test statistics, regardless of whether Reimers (1992) small–sample correction is employed. Both the estimated long–run price and income elasticities have the expected signs, but the former is particularly small in magnitude (-0.01 and 0.33 respectively).

The results of the cointegration analysis for the annual data are reported in Table 10 in the Appendix. Apart from some evidence that there might be more than one cointegrating vector, the results broadly corroborate the findings from the analysis of the quarterly data: the estimated long–run elasticities for price and income are -0.06 and 0.49 respectively.

3.2 The Pesaran–Smith–Shin approach

A testing procedure for the existence of a long–run relationship that can be applied irrespective of whether the underlying variables are $I(1)$ or $I(0)$ has recently been introduced by Pesaran, Shin, and Smith (1996). This procedure avoids the problems inherent in pre–testing for unit roots prior to testing for cointegration (see, e.g., Cavanaugh, Elliott, and Stock, 1995), and is particularly attractive in the present context. It is a bounds test procedure based on the usual F- or Wald–test of the lagged levels terms in an 'error-correction formulation' (c.f., Kremers, Ericsson, and Dolado, 1992), and allows a conclusive decision to be drawn, without needing to know the order of integration or cointegration rank of the variables, whenever the test statistics falls outside the critical value bounds.

Because of the power and size problems of unit root tests alluded to above, the approach may be informative when, as in our case, there are doubts as to whether the variables are $I(0)$ or $I(1)$.[1]

The 'unrestricted' error–correction model (ECM) considered is

$$\Delta y_t = a_0 + a_1 t + \phi y_{t-1} + \boldsymbol{\delta}'\mathbf{x}_{t-1} + \sum_{i=1}^{n_1-1} \psi_i \Delta y_{t-i} + \sum_{i=0}^{n_2-1} \boldsymbol{\varphi}_i \Delta \mathbf{x}_{t-i} + \xi_t, \tag{3}$$

Table 3: Cointegration statistics (quarterly data)

μ_i		l_i				rank r
		903.903				0
0.5834		939.364				1
0.0497		941.431				2
0.0049		941.629				3
H_0: 'rank $= r$'	Max	$Max_{(T-mn)}$	95%	$Trace$	$Trace_{(T-mn)}$	95%
$r \leq 0$	70.92**	65.67**	21.0	75.45**	69.86**	29.7
$r \leq 1$	4.13	3.83	14.1	4.53	4.19	15.4
$r \leq 2$	0.40	0.37	3.8	0.40	0.37	3.8
standardised $\boldsymbol{\beta}'$ eigenvectors		q_t		p_t		y_t
$i = 1$		1.000		0.012		−0.331
$i = 2$		0.301		1.000		0.377
$i = 3$		0.939		−3.714		1.000
standardised $\boldsymbol{\alpha}$ coefficients		$i = 1$		$i = 2$		$i = 3$
q_t		−0.975		−0.014		−0.002
p_t		−0.013		−0.039		0.001
y_t		0.026		−0.007		−0.003

(NOTES: r denotes the rank of the cointegrating space, m the number of variables ($m = 3$), and n the lag order of the VAR ($n = 2$). ** denotes significance at the 1% level.)

where ξ_t is assumed to be an $\mathsf{IN}(0, \sigma_\xi^2)$ process and $\sigma_\xi^2 \equiv \sigma_{11} - \sigma_{12}\mathbf{\Sigma}_{22}^{-1}\sigma_{21}$, with

$$\mathbf{\Sigma} = \begin{bmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{21} & \mathbf{\Sigma}_{22}, \end{bmatrix}$$

the covariance matrix of the errors of the joint (VAR) model of $\mathbf{z}_t = [y_t : \mathbf{x}_t']'$.

If $\phi \neq 0$ and $\boldsymbol{\delta} \neq \mathbf{0}$ there exists a long–run relationship between the levels of y_t and $\mathbf{x}_t$ of the form:

$$y_t = \theta_0 + \theta_1 t + \boldsymbol{\theta}'\mathbf{x}_t + v_t, \tag{4}$$

where $\theta_0 \equiv -a_0/\phi$, $\theta_1 \equiv -a_1/\phi$, $\boldsymbol{\theta} \equiv -\boldsymbol{\delta}/\phi$ is the k-vector of long–run response parameters, and v_t a stationary process with mean zero. Note also that if $\phi < 0$ then the long–run relationship between the levels of y_t and $\mathbf{x}_t$ is stable.

Eq. (3) can be expressed as

$$\Delta y_t = a_0 + a_1 t + \phi\left(y_{t-1} - \boldsymbol{\theta}'\mathbf{x}_{t-1}\right) + \sum_{i=1}^{n_1-1} \psi_i \Delta y_{t-i} + \sum_{i=0}^{n_2-1} \boldsymbol{\varphi}_i \Delta \mathbf{x}_{t-i} + \xi_t, \tag{5}$$

Consequently, a test for $\phi = 0$ may be interpreted as a test for the existence of a long–run relationship.

The approach suggested by Pesaran et al. (1996) is to test for the *non*–existence of a long–run relationship between y_t and $\mathbf{x}_t$ by testing the joint hypothesis '$\phi = 0$ and $\boldsymbol{\delta} = \mathbf{0}$' in Eq. (3). This avoids the problem with Eq. (5) that when $\boldsymbol{\theta}$ is neither pre–specified nor estimated the elements of $\boldsymbol{\theta}$ in the case when $\phi = 0$ are not identified.

The asymptotic distributions of the Wald- and F- statistics are non–standard. Two sets of critical values are provided: one which is appropriate when all the variables are $I(1)$, and the other for the case when all the variables are $I(0)$, thus covering all possible classifications of the variables into $I(1)$, $I(0)$, or mutually cointegrated. In the sense that different orders of integration are allowed for, the approach is more general than the partial systems cointegration analysis of, for example, Boswijk (1992), Johansen (1992) and Urbain (1992).

Our findings of applying the Pesaran–Shin–Smith procedure to the quarterly data are summarised in Table 4. When four lags are included on all variables ($n_1 = n_2 = n = 4$), the Wald statistic is below the lower bound, implying the absence of a long-run relationship.

However, the trend appears redundant, and more importantly, the order of the model can be reduced to $n = 2$ without any apparent signs of mis–specification (a test for serial correlation is recorded, but the usual battery of tests failed to signal any problems) and then the Wald test exceeds the upper value. If the price variable is omitted, the rejection of the null for $n = 2$ is more decisive. Finally, we note the estimate of the long–run relationship implied by the model is $q - 0.48h - 0.35y$, which is similar to, and confirms, that obtained from the systems cointegration analysis.

4 The econometric models considered

4.1 A short–run model of energy demand

Figure 3 plots the quarterly time–series for q and h and suggests that part of the observed seasonal pattern in q can be explained by the seasonal pattern in h. Visual inspection indicates that to a first approximation the seasonal variation in q is roughly one half of that in h. Not surprisingly, given the visual correlation between the series, regressing q on h and a constant produces the relationship reported in Table 5.

Assuming q and h do not possess roots at the seasonal frequencies[2] but that they both contain zero–frequency roots, the regression reported in Table 5 appears to constitute a zero–frequency cointegrating vector, judged by the value of the DW statistic (see Sargan and Bhargava, 1983). The last entry in Table 2 formally tests for unit roots in the residual of this regression. The test for a zero–frequency root is close to the critical value reported in Hylleberg et al. (1990), although these critical values are not appropriate for an estimated residual, and would reject too often.

To briefly summarise: the results of the unit–root testing procedures carried out in Section 2, and in this section, are difficult to interpret and lead to conflicting inferences concerning the time–series properties of the variables. h can not be $I(1)$ *a priori*, yet in line with the studies cited in Table 1 we find that q appears to be $I(1)$ *and* there appears to be a perfectly sensible relationship between the levels of q and h. Moreover, in the multivariate cointegration analysis reported in Section 3, we find that we can only reject the null of no cointegration between q, p and y, using the Johansen ML procedure, if either

Table 4: Test outcomes for the existence of a long–run relationship between energy demand and the explanatory variables

n	regressors	AR(5) p-value	Wald	critical values
4	$h, y, p; c, t, s$	0.081	3.18	16.26, 20.48
3	$h, y, p; c, s$	0.222	5.55	12.88, 17.51
2	$h, y, p; c, s$	0.096	18.98	12.88, 17.51
4	$h, y; c, t, s$	0.102	4.72	14.71, 17.62
3	$h, y; c, s$	0.258	8.88	11.38, 14.57
2	$h, y; c, s$	0.114	26.00	11.38, 14.57

(NOTES: The critical values of the Wald statistic are taken from Pesaran et al., 1996, Table A. c, t, s denote constant, trend and seasonals, respectively, and the third column is the p-value of a test for serial correlation up to fifth order.)

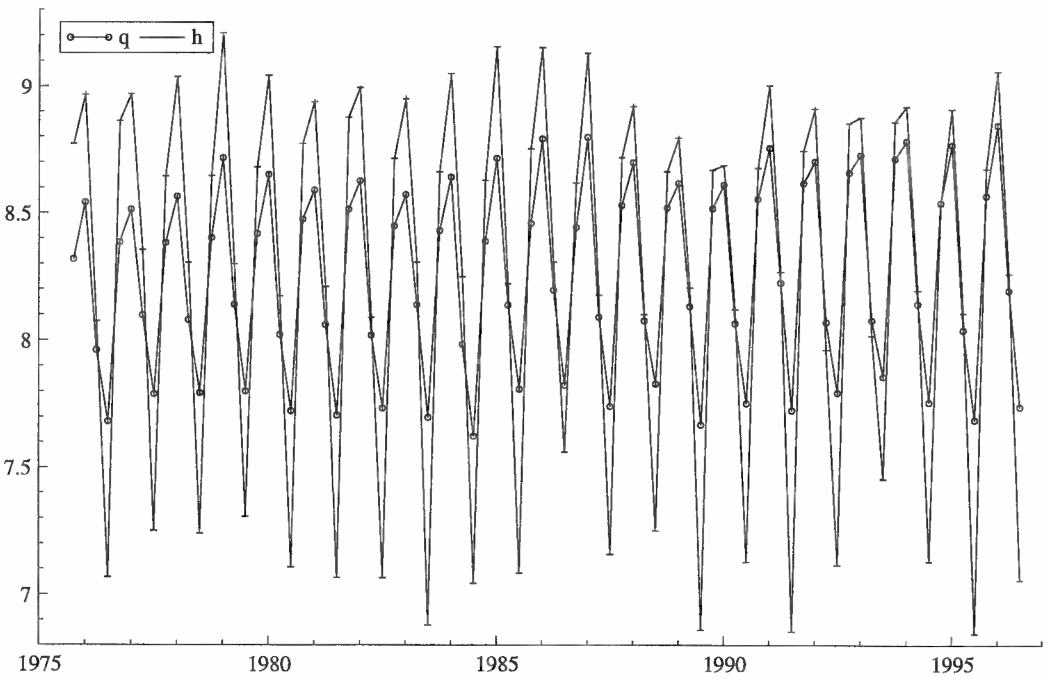

Figure 3: Time–series plots of q and h (matching means)

h is included unrestrictedly in the VAR (i.e., partialled out of $\Delta\mathbf{X}_t$ and $\mathbf{X}_{t-n}$ along with the $\Delta\mathbf{X}_{t-j}$, $j = 1, \ldots, n-1$, where $\mathbf{X}_t = [q_t, p_t, y_t]'$) or is restricted to the cointegrating space (see, e.g. Johansen and Juselius, 1990).

Nevertheless, our aim is not to establish the orders of integration (zero–frequency and seasonal) of the series once–and–for–all *per se*, but to use what we can learn about the properties of the series to build useful forecasting models. The results in Table 5, for example, suggest that a 'short–run' forecasting model that relates Δq to h (to eliminate the zero–frequency root in q) may perform poorly relative to a model that allows for a relationship in the levels of these two variables.

Starting from a general dynamic model for q and h, and adopting the sequential 'general–to–simple' model design procedure advocated by David Hendry and his co–authors (see, e.g. Hendry, 1995), we obtained the short–run model set out in Table 6. There, s_j is the seasonal dummy for quarter j (i.e. unity when t falls in quarter j, and zero otherwise); σ is the residual standard deviation; the diagnostic tests are of the form $\mathsf{F}_{\mathsf{j}}\,(k, T-l)$ which denotes an F-test against the alternative hypothesis j for: 4^{th}-order residual serial correlation (F_{ar}: see Godfrey, 1978), 4^{th}-order residual autoregressive conditional heteroscedasticity ($\mathsf{F}_{\mathsf{arch}}$: see Engle, 1982), heteroscedasticity ($\mathsf{F}_{\mathsf{het}}$: see White, 1980), omitted powers of $\hat{y}_t$, i.e. $\hat{y}_t^2, \hat{y}_t^3, \ldots$; ($\mathsf{F}_{\mathsf{reset}}$: see Ramsey, 1969); and a chi–square test for normality ($\chi^2_{\mathsf{nd}}(2)$: see Doornik and Hansen, 1994). * and ** denote significance at the 5% and 1% levels, respectively. The in–sample diagnostic tests are satisfactory apart from that for normality.

The short–run response of Δq to a change in Δh is 0.38, a little less than the implied long–run response of 0.45, which is close to the 'static regression' estimate of Table 5. Imposing the unit root in q and beginning again by testing down from a general specification results in a model with a worse fit and poorer forecasts, which is unsurprising given the magnitude of the t-value on q_{t-1}. The time trend suggests annual growth of approximately 1% per annum in q.

To assess the adequacy of the model for forecasting, 15 observations were held back at the model design stage. Ideally, if the sample were long enough we would generate sequences of, say, j-step forecasts, by rolling the origin forward through the sample, as in, e.g., Clements and Hendry (1997). However, we have to confine our attention to analysing forecasts based on a single origin, which unfortunately limits the generality that can be claimed for the results.

Although the parameter constancy tests suggest some problems with the model (see Doornik and Hendry, 1997, for an explanation of these tests), the first plot in Figure 4 is perhaps more informative of how well the model predicts over this period. The forecasts are *ex post* in that they are formed using the true values of h. Note that we consider forecasts of q rather than Δq: Clements and Hendry (1993, 1995) discuss the role of the transformation of the variable on which forecast accuracy is to be assessed, and show that assessing forecast accuracy in terms of ability to predict differences of the data will tend to flatter poor models. The multi–step forecasts naturally show a slightly worsening performance as the horizon increases—here we continue to use the true values of h (i.e., $h_{T+j}, h_{T+j-1}, \ldots$) for forecasting q_{T+j}, but for horizons $j > 1$ explanatory variables involving lags of q are replaced by their forecasts.

4.2 A long–run model of energy demand

The short–run forecasting model in Section 4.1 accounts for the upward trend in q of just under 1% per annum (cf. Figure 1) by a linear time trend. A more satisfactory approach would attribute the long–run growth and any cyclical movements in q to economic fundamentals, such as y and p, and that is the task of this section.

The systems cointegration analysis suggests a role for y in the long run, but is less sanguine about the prospects of finding a significant role for p. In fact we were unable to find significant price terms in either the long or short run, and such terms do not appear in our preferred model given in Table 7, which was again obtained by a general–to–specific simplification procedure. The 'long–run' model is very similar to the 'short–run' model, except that a lagged income term y_{t-1} replaces the time trend and now accounts for the growth in q over time. Tying the long–run evolution of q to y is obviously more satisfactory than using a time trend, and improves the equation fit a little, but from comparing the forecasts between the short–run and the long–run model we learnt that the forecasts were not noticeably better.

The model is satisfactory within sample (apart from the test outcome for normality), but there again appears to be evidence of parameter instability over the forecast period.

Table 5: Regressing q on h by OLS, 1975(4) to 1996(3)

variable	coeff.	std.error	t-value	t-prob.	part.R^2
constant	5.277	0.103	51.19	0.00	0.970
h_t	0.485	0.017	29.07	0.00	0.912
$RSS = 1.008$	$\hat{\sigma} = 0.111$	$R^2 = 0.912$	$\bar{R}^2 = 0.910$	$DW = 1.90$	

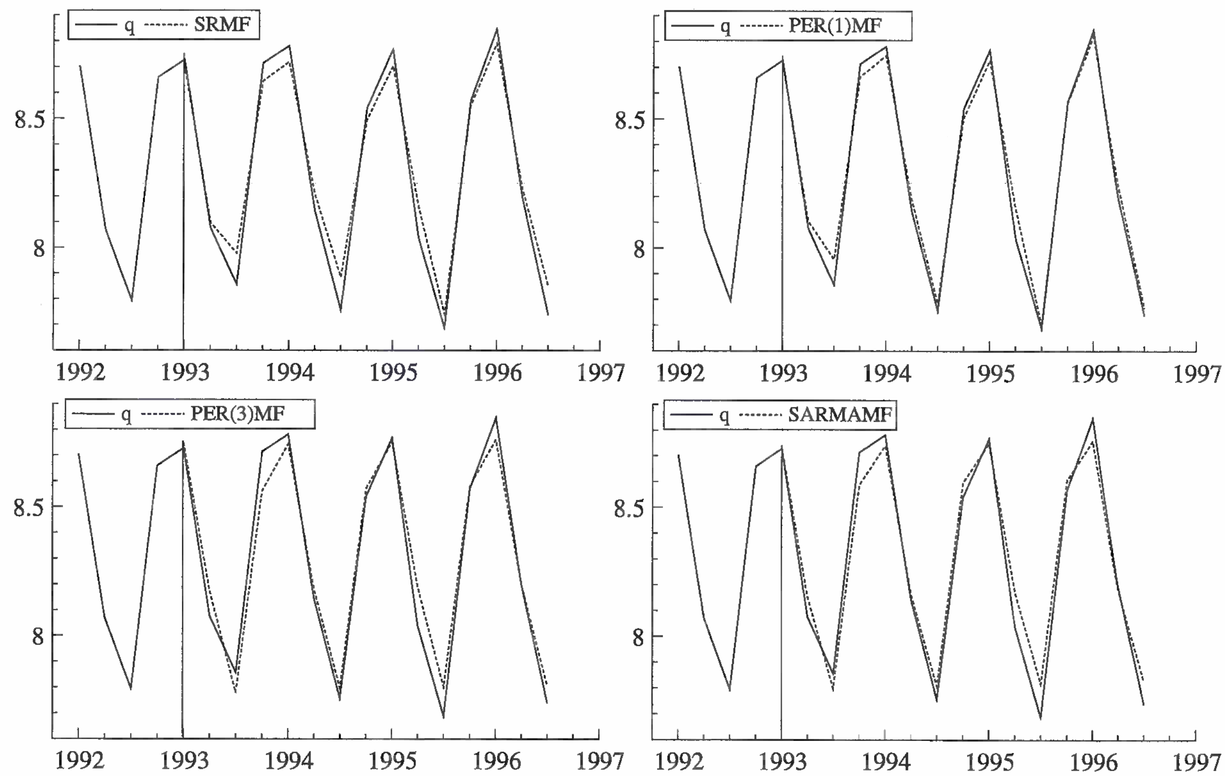

Figure 4: Multi–step forecasts from short–run model (Table 6), periodic model with *ex post* forecasts (Table 8), periodic model with *ex ante* forecasts, and SARMA model

Table 6: Short–run forecasting model for q

Modelling Δq_t by OLS, 1977(1) to 1996(3) less 15 forecasts. The forecast period is: 1993(1) to 1996(3)					
variable	coeff.	std.error	t-value	t-prob.	part.R^2
Δq_{t-1}	0.296	0.118	2.51	0.015	0.104
Δq_{t-3}	0.142	0.047	2.98	0.004	0.142
Δq_{t-4}	0.165	0.047	3.48	0.001	0.183
Δh_t	0.379	0.030	12.45	0.000	0.742
Δh_{t-1}	−0.154	0.062	−2.47	0.017	0.102
q_{t-1}	−0.976	0.149	−6.57	0.000	0.444
h_{t-1}	0.445	0.093	4.80	0.000	0.299
s_1	0.240	0.052	4.66	0.000	0.287
constant	5.161	0.912	5.66	0.000	0.372
trend	0.003	0.000	5.17	0.000	0.331
$RSS = 0.082$	$\hat{\sigma} = 0.039$	R^2=0.990		$\bar{R}^2 = 0.988$	$DW = 1.95$

1-step (*ex post*) forecast analysis 1993(1) to 1996(3). Parameter constancy forecast tests:

Forecast $\chi^2(15) = 41.141[0.00]^{**}$

Chow $F(15, 54) = 1.909[0.04]^{*}$

Model diagnostics:

$F_{ar}(4{,}50) = 0.275\ [0.89]$

$F_{arch}(4{,}46) = 0.581\ [0.68]$

$F_{het}(17{,}36) = 0.792\ [0.69]$

$F_{reset}(1{,}53) = 1.142\ [0.29]$

$\chi^2_{nd}(2) = 11.117[0.00]^{**}$

Table 7: Long–run forecasting model for q

Modelling Δq_t by OLS, 1977(1) to 1996(3) less 15 forecasts.

The forecast period is: 1993(1) to 1996(3)

variable	coeff.	std.error	t-value	t-prob.	part.R^2
Δq_{t-1}	0.329	0.119	2.77	0.008	0.124
Δq_{t-3}	0.145	0.047	3.09	0.003	0.150
Δq_{t-4}	0.162	0.047	3.47	0.001	0.182
Δh_t	0.381	0.030	12.69	0.000	0.749
Δh_{t-1}	−0.166	0.062	−2.67	0.010	0.117
q_{t-1}	−1.023	0.151	−6.78	0.000	0.460
h_{t-1}	0.471	0.094	5.03	0.000	0.319
y_{t-1}	0.360	0.067	5.40	0.000	0.350
s_1	0.229	0.051	4.53	0.000	0.275
constant	3.898	0.792	4.92	0.000	0.310
$RSS = 0.080$	$\hat{\sigma} = 0.038$	$R^2 = 0.990$		$\bar{R}^2 = 0.989$	$DW = 2.01$

1-step (*ex post*) forecast analysis 1993(1) to 1996(3). Parameter constancy forecast tests:

Forecast $\chi^2(15) = 35.428[0.00]^{**}$

Chow $F(15, 54) = 1.690$ [0.08]

Model diagnostics:

$F_{ar}(4{,}50) = 0.424$ [0.79]

$F_{arch}(4{,}46) = 0.927$ [0.46]

$F_{het}(17{,}36) = 0.852$ [0.63]

$F_{reset}(1{,}53) = 2.441$ [0.12]

$\chi^2_{nd}(2) = 7.025[0.03]^{*}$

4.3 A periodic (econometric) model of energy demand

Following on from important papers by Tiao and Grupe (1980) and Osborn (1988), much of the recent research on modelling seasonal processes has focused on periodic models. For example, Birchenhall et al. (1989) consider a periodic model of consumers expenditure, and it is natural to consider whether allowing periodic variation may yield improved forecasts of energy demand.

Periodic models have parameters that vary across the seasons (here, quarters). While such models often admit shorter lag lengths than non–periodic models, they are often of a high dimension. For instance, if there were no offset in overall lag lengths and all parameters were periodic, there would be four times as many parameters to estimate.

Our relatively short sample limits the generality of the periodic model that we entertain at the outset as the 'general' model. In fact, the most general model we allow is that given in Table 7, with all the parameters allowed to exhibit seasonal variation. We then consider each parameter in turn, and test whether it exhibits seasonal variation using a conventional F-test (with 3 degrees of freedom). The model is then re–estimated, imposing 'no–variation' for particular parameters where this is not rejected, and allowing periodic variation otherwise. A final step is to delete insignificant variables. The results are shown in Table 8, the associated multi–step forecasts in the second plot in Figure 4. In Table 8 $s_j x$, $j = 1, \ldots, 4$, denotes $s_j \times x$ for a variable x.

The periodic model has a better in–sample fit (the standard error of the regression is around 8% lower, as compared with the long–run model) and visually the forecasts are markedly better. A feature of the equation is that the coefficent on q_{t-1} appears to depend on the season. We can constrain the first and fourth quarter effects to be equal ($\chi^2(1) = 0.323$ [0.57]), but no further simplifications are data admissible (e.g., testing the first, second and fourth quarter coefficients for equality yielded $\chi^2(2) = 22.030$ [0.00]).

The parameter constancy forecast tests reported in Table 8 show that the hypothesis of parameter constancy is rejected at the 1% level for the forecast χ^2 test and at the 5% level for the Chow test, possibly indicating over–fitting in–sample. None of the in–sample diagnostics indicate mis–specification at the 1% level.

In order to mimic a more realistic situation for the energy analyst, we generated one–

Table 8: Periodic econometric forecasting model for q

Modelling Δq_t by OLS, 1977(1) to 1996(3) less 15 forecasts.

The forecast period is: 1993(1) to 1996(3)

variable	coeff.	std.error	t-value	t-prob.	part.R^2
Δh_t	0.441	0.039	11.42	0.000	0.711
$s_3 \Delta q_{t-1}$	1.163	0.223	5.22	0.000	0.339
$s_2 \Delta h_{t-1}$	−0.486	0.126	−3.87	0.000	0.220
$s_1 q_{t-1}$	−0.884	0.123	−7.19	0.000	0.494
$s_2 q_{t-1}$	−0.912	0.122	−7.49	0.000	0.514
$s_3 q_{t-1}$	−1.397	0.188	−7.43	0.000	0.511
$s_4 q_{t-1}$	−0.880	0.126	−7.00	0.000	0.480
h_{t-1}	0.445	0.070	6.33	0.000	0.431
y_{t-1}	0.391	0.057	6.87	0.000	0.471
s_3	4.384	1.503	2.92	0.005	0.138
constant	2.895	0.761	3.81	0.000	0.215

$RSS = 0.066$ $\hat{\sigma} = 0.035$ $R^2 = 0.992$ $\bar{R}^2 = 0.991$ $DW = 2.07$

1-step (*ex post*) forecast analysis 1993 (1) to 1996 (3). Parameter constancy forecast tests:

Forecast χ^2 (15) = 41.210 [0.00]**

Chow $F(15, 53) = 1.885$[0.05]*

Model diagnostics:

F_{ar}(4,49) = 0.653 [0.63]

F_{arch}(4,45) = 0.332 [0.86]

F_{het}(19,33) = 0.960 [0.52]

$F_{reset}(1, 52) = 5.461$[0.02]*

$\chi^2_{nd}(2) = 7.073$[0.03]*

step and multi–step ahead forecasts q_{T+1} based on forecast values of h_{T+1}. The results for the multi–step forecasts are reported in graphical form in the third plot in Figure 4, and are compared to those obtained from the other models in Table 9. The model for h_t was of the form

$$h_t = \mu_s + \gamma t + \varepsilon_t, \tag{6}$$

which proved to be a data admissible simplification of a general fourth–order PAR model (see below, Eq.(9)). Here μ_s varies across the seasons (equivalent to a constant and three seasonals), and t is a linear trend. The equation standard error was 0.125, the adjusted R^2 0.971. The calculation of true *ex ante* forecasts from the periodic econometric model necessitates an equation for y_t. In the event, a first–order autoregression in first differences proved adequate, with an equation standard error of 0.020, and an adjusted R^2 of 0.981.

5 The time–series models considered

We also estimated a number of univariate time–series models in order to provide a benchmark against which to evaluate the forecasts from the econometric models of q. Two classes of models were considered: seasonal autoregressive–moving average (SARMA) models from the time–series modelling tradition of Box and Jenkins (1970), and periodic autoregressive (PAR) models.

5.1 A SARMA model of energy demand

The general class of SARMA models can be written as

$$\phi(L)(1-L)\left(1-L^4\right)x_t = \mu + (1-\theta_1 L)\left(1-\theta_4 L^4\right)\epsilon_t, \tag{7}$$

where $\epsilon_t \sim \mathsf{IN}(0,\sigma_\epsilon^2)$, $|\theta_1| < 1$, $|\theta_4| < 1$. When $\phi(L) = 1$ this model is sometimes known as the 'airline' model, following the application to airline data of Box and Jenkins (1970). The rationale for the model is straightforward. The filter $\Delta_4 = (1 - L^4)$ captures the tendency for the value of the series in any quarter to be highly correlated with the value in the same quarter a year earlier, while $\Delta = (1 - L)$ relates to the non–seasonal part of the model and specifies a stochastic trend in the level of the series (with drift when $\mu \neq 0$).

The MA terms are usually found to be negative (and possibly quite close to unity), so that the impact of the unit roots tends to be moderated. This model has proved useful empirically (see Franses, 1996, pp.42–46, for references to empirical studies). Often there is a tension between the two zero–frequency roots implied by the AR polynomial ($\Delta\Delta_4$) and the outcome of unit root testing procedures, which may suggest only a single root at the zero frequency (see, e.g., Osborn, 1990; Hylleberg, Jørgensen, and Sørensen, 1993; Clements and Hendry, 1997).

For the sample period 1975:4–1992:4 we estimated the following model for q, where standard errors are in squared brackets:

$$\begin{aligned}(1-L)\left(1-L^4\right)q_t &= (1-\underset{[0.034]}{0.946}L)(1-\underset{[0.089]}{0.816}L^4)\hat{\epsilon}_t\end{aligned} \tag{8}$$

and $\hat{\sigma} = 0.069$, $\mathsf{R}^2 = 0.596$, $DW = 1.67$.

As expected, θ_1 and θ_4 are both close to unity. There is some evidence of serial correlation in the disturbances, but nevertheless the forecasts from Eq.(8) serve as a useful benchmark.

5.2 A PAR model of energy demand

Some success has been claimed for univariate periodic models for short–term forecasting. For example, Osborn and Smith (1989) found that periodic models offered some improvement in forecast accuracy for the components of seasonal UK consumption at short time horizons, although such a conclusion is not supported by the analysis of Clements and Smith (1997) based on an extended sample period. Franses (1996) cites empirical work and evidence from Monte Carlo studies attesting to the usefulness of periodic models for forecasting. Empirical evidence suggests that in many cases PAR models of low order, by contrast to periodic moving average (PMA) or periodic ARMA (PARMA) models, are sufficient to describe periodic time series (cf., McLeod, 1994).

Univariate periodic time–series models simply allow the slope or autoregressive parameters of the model to vary with the seasons, as well as the intercept (which is commonplace in non–periodic models of seasonal variables). Thus, the periodic autoregressive (PAR) model can be written as

$$y_t = \mu_s + \phi_{1s} y_{t-1} + \ldots + \phi_{ns} y_{t-n} + \epsilon_t, \tag{9}$$

where the intercept, μ_s, and the autoregressive parameters, $\phi_{1s}, \ldots, \phi_{ns}$, may vary with the season $s = 1, \ldots, 4$. It is implicit in the formulation in Eq.(9) that $\mu_s = \mu_1$ and $\phi_{js} = \phi_{j1}$ when t falls in season $s = 1$, etc. The disturbance term is assumed to be normally independently distributed with zero mean and variance σ^2. n is chosen so that the last assumption approximately holds in the empirical model. We can also consider 'restricted' PAR (RPAR) models with holes in the lag distribution.

The hypothesis of no periodic variation in the slope parameters is a simple F-test of the nested non–periodic (in the slopes) AR(n) model

$$y_t = \mu_s + \phi_1 y_{t-1} + \ldots + \phi_n y_{t-n} + \epsilon_t, \tag{10}$$

that is, that $\phi_{js} = \phi_j$, $s = 1, \ldots, 4$, $j = 1, \ldots, n$, which has an $\mathsf{F}_{3n,T-4n-4}$ distribution under the null.

We can reparameterise Eq.(9) as

$$y_t - \alpha_s y_{t-1} = \mu_s + \sum_{j=1}^{n-1} \beta_{js}(y_{t-j} - \alpha_{s-j} y_{t-j-1}) + \epsilon_t \tag{11}$$

$$(1 - \alpha_s L) y_t = \mu_s + \sum_{j=1}^{n-1} \beta_{js}(1 - \alpha_{s-j} L) y_{t-j} + \epsilon_t, \tag{12}$$

where $\alpha_{s-4i} = \alpha_s$, $i = 1, 2, \ldots$, and test whether $\alpha_1\alpha_2\alpha_3\alpha_4 = 1$, in which case the periodic filter $(1 - \alpha_s L)$ removes the stochastic trend, and the variable is said to be periodically integrated. An interesting feature of such series is that they can not be decomposed into seasonal and stochastic trend components—the two are inextricably linked (see, e.g. Franses, 1996, Ch. 8).

The order of the unrestricted PAR model we estimated for q was $n = 4$. We could set the lag 2 and lag 3 terms to zero: $\mathsf{F}_{8,45} = 0.793\,[0.61]$. Testing for periodic variation within the RPAR model $H_0 : \phi_{js} = \phi_j$, $s = 1, \ldots, 4$, $j = 1, 4$, rejected the null: $\mathsf{F}_{6,53} = 5.562\,[0.00]$. We were able to impose the restriction that $\alpha_1\alpha_2\alpha_3\alpha_4 = 1$ in the unrestricted PAR model Eq.(11), and we then deleted insignificant regressors at conventional critical values. This led to a model in the form of Eq.(11) but with $\mu_s = 0$, $s = 1, 2, 4$, $\beta_{12} = \beta_{22} = \beta_{23} = \beta_{33} = 0$, which we call the RPIAR model. The individual coefficient estimates and standard errors are not very informative, and are not reported, but for the RPIAR model we obtained $\hat{\sigma} = 0.057$, $\mathsf{R}^2 = 0.978$, and $DW = 2.18$, and for the RPAR model a similar fit: $\hat{\sigma} = 0.058$, $\mathsf{R}^2 = 0.979$, and $DW = 1.99$.

Table 9 shows a comparison of the forecast errors. On the basis of RMS forecast errors, the SARMA model is the best of the time series models. Compared with the econometric models, however, they generally forecast less accurately. For the three periodic models analysed, the deterioration of forecast precision when producing true *ex ante* forecasts is also evident.

Table 9: Comparison of the root mean square (RMS) forecast errors

RMSFE	econometric models					time–series models		
	s–r	l–r	PER (1)	PER (2)	PER (3)	SARMA	RPAR	RPIAR
1-step	0.0646	0.0591	0.0585	0.0715	0.0715	0.0737	0.0779	0.0862
multi–step	0.0758	0.0713	0.0501	0.0736	0.0784	0.0755	0.0904	0.0728

(NOTES: (1) refers to the *ex post* periodic forecasting model (i.e. where the true values of h are used), (2) to the periodic model where the values of h are forecast (i.e. with y known), and (3) to the *ex ante* periodic forecasting model (i.e. where both h and y are forecast).)

6 Conclusions

A large number of empirical papers in energy economics have sought to obtain estimates of the long–run elasticities of the determinants of the demand for energy, primarily using either the Johansen ML systems approach, or the Engle–Granger static regression approach. Neither is particularly satisfactorily, for the reasons explained in the paper, but nevertheless we find that the Johansen estimates are broadly similar to those obtained using an approach due to Pesaran, Shin, and Smith (1996), which tests for a long–run relationship between the variables of interest, without the pre–supposition that all the variables are $I(1)$. Formally this is attractive, since it turns out that we can be fairly ambiguous about the time series properties of the temperature variable, which is often found to be vital in studies of energy demand that seek to estimate a cointegrating relationship. We find a long–run relationship irrespective of what we assume about this variable.

However, incorporating an estimate of the long–run relationship adds little to the short–run forecast performance of the model, over and above using a linear trend, and seasonal univariate time series models (such as the 'airline' model and periodic models) provide competitive forecasts, particularly when we allow for the additional uncertainty from forecasting the explanatory variables of the econometric models. Allowing for seasonally varying responses of demand to its principle determinants does establish the primacy of such models over the univariate time series rivals, but we caution against reading too much into this, given the fairly short sample period and the single epoch over which forecasts are evaluated.

Notes

[1]Note that the Pesaran–Shin–Smith approach is appropriate when there is only one 'fundamental' long–run relationship between the variables, and that relationship enters the conditional model for the chosen dependent variable. Thus the approach places more 'structure' on the problem than the VAR–based approach of Johansen(Johansen, 1988), in terms of notions such as causality and exogeneity—see Pesaran et al. (1996) for details.

[2]If they do, the *Theorem* in Engle et al. (1989, p.50) shows that the zero–frequency cointegrating vector can still be consistently estimated by first pre–filtering both series with $S(L) = 1 + L + L^2 + L^3$. This is approximately what the annual data does, and in the absence of seasonal roots it is unsurprising that the systems cointegration analyses on the quarterly and annual series in Section 3 yield similar results.

References

Bentzen, J. (1994). An empirical analysis of gasoline demand in Denmark using cointegration techniques. *Energy Economics*, *16*(2), 139–143.

Bentzen, J. and Engsted, T. (1993). Short- and long–run elasticities in energy demand: a cointegration approach. *Energy Economics*, *15*(1), 9–16.

Birchenhall, C. R., Bladen-Hovell, R. C., Chui, A. P. L., Osborn, D. R., and Smith, J. P. (1989). A seasonal model of consumption. *The Economic Journal*, *99*, 837–843.

Boswijk, H. P. (1992). *Cointegration, Identification and Exogeneity.* Vol. no. 37 of *Tinbergen Institute Research Series.* Thesis Publishers, Amsterdam.

Box, G. E. P. and Jenkins, G. M. (1970). *Time Series Analysis, Forecasting and Control.* Holden–Day, San Francisco.

Campbell, J. Y. and Perron, P. (1991). Pitfalls and opportunities: what macroeconomists should know about unit roots. In Blanchard, O. J. and Fischer, S. (Eds.), *NBER Economics Annual 1991.* MIT Press.

Cavanaugh, C. L., Elliott, G., and Stock, J. H. (1995). Inference in models with nearly integrated regressors. *Econometric Theory*, *11*, 1131–1147.

Clements, M. P. and Hendry, D. F. (1993). On the limitations of comparing mean squared forecast errors. *Journal of Forecasting*, *12*, 617–637. With discussion.

Clements, M. P. and Hendry, D. F. (1995). Forecasting in cointegrated systems. *Journal of Applied Econometrics*, *10*, 127–146.

Clements, M. P. and Hendry, D. F. (1997). An empirical study of seasonal unit roots in forecasting. *International Journal of Forecasting.* Forthcoming.

Clements, M. P. and Smith, J. (1997). Forecasting seasonal UK consumption components. Mimeo, Department of Economics, University of Warwick.

Doornik, J. A. and Hansen, H. (1994). A practical test of multivariate normality. Discussion paper, Nuffield College.

Doornik, J. A. and Hendry, D. F. (1996). *PcGive 9: An Interactive Econometric Modelling System.* International Thomson Publishing, London.

Doornik, J. A. and Hendry, D. F. (1997). *Modelling Dynamic Systems Using PcFiml 9 for Windows.* International Thomson Business Press, London.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflations. *Econometrica*, *50*, 987–1007.

Engle, R. F. and Granger, C. W. J. (1987). Cointegration and error correction: representation, estimation and testing. *Econometrica*, *55*, 251–276.

Engle, R. F., Granger, C. W. J., and Hallman, J. J. (1989). Merging short- and long–run forecasts: An application of seasonal cointegration to monthly electricity sales forecasting. *Journal of Econometrics*, *40*, 45–62.

Engsted, T. and Bentzen, J. (1993). Expectations, adjustment costs, and energy demand. *Resources and Energy Economics*, *15*, 371–385.

Fouquet, R. (1995). The impact of VAT introduction on UK residential energy demand. *Energy Economics*, *17*(3), 237–247.

Franses, P. H. (1996). *Periodicity and Stochastic Trends in Economic Time Series.* Oxford University Press, Oxford.

Godfrey, L. G. (1978). Testing for higher order serial correlation in regression equations when the regressors include lagged dependent variables. *Econometrica*, *46*, 1303–1313.

Hargreaves, C. P. (1994). A review of methods of estimating cointegrating relationships. In Hargreaves, C. P. (Ed.), *Nonstationary Time Series Analysis and Cointegration*, pp. 87–131. Oxford University Press, Oxford.

Hendry, D. F. (1995). *Dynamic Econometrics.* Oxford University Press, Oxford.

Hunt, L. and Manning, N. (1989). Energy price- and income–elasticities of demand: some estimates for the UK using the cointegration procedure. *Scottish Journal of Political Economy*, *36*, 183–193.

Hunt, L. C. and Lynk, E. L. (1992). Industrial Energy Demand in the UK: a Co–Integration Approach. In Hawdon, D. (Ed.), *Energy Demand: Evidence and Expectations*, pp. 143–162. Surrey University Press, Guildford, UK.

Hunt, L. C. and Witt, R. (1995). An Analysis of UK Energy Demand Using Multivariate Cointegration. Surrey Energy Economics Discussion Paper Series (SEEDS) No. 86, Department of Economics, University of Surrey, Guildford, UK.

Hylleberg, S. (1992). *Modelling Seasonality.* Oxford University Press, Oxford.

Hylleberg, S., Engle, R. F., Granger, C. W. J., and Yoo, B. S. (1990). Seasonal integration and co–integration. *Journal of Econometrics*, *44*, 215–228.

Hylleberg, S., Jørgensen, C., and Sørensen, N. K. (1993). Seasonality in macroeconomic time series. *Empirical Economics*, *18*, 321–325.

Johansen, S. (1988). Statistical analysis of cointegration vectors. *Journal of Economic Dynamics and Control*, *12*, 231–254.

Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in Gaussian Vector Autoregressive Models. *Econometrica*, *59*(6), 1551–1580.

Johansen, S. (1992). Cointegration in partial systems and the efficiency of single equation analysis. *Journal of Econometrics*, *52*, 389–402.

Johansen, S. and Juselius, K. (1990). Maximum likelihood estimation and inference on cointegration—with applications to the demand for money. *Oxford Bulletin of Economics and Statistics*, *52*, 169–210.

Kremers, J. J. M., Ericsson, N. R., and Dolado, J. (1992). The power of co–integration tests. *Oxford Bulletin of Economics and Statistics*, *54*, 325–348.

Madlener, R. (1996a). Econometric analysis of residential energy demand: A survey. *The Journal of Energy Literature*, *II*(2), 3–32.

Madlener, R. (1996b). On the Use of Multivariate Cointegration Analysis in Residential Energy Demand Analysis. Economics Series No. 30, Institute for Advanced Studies, Vienna. Paper presented at the Econometric Society European Meeting (ESEM), August 25–29, 1996, Istanbul.

McLeod, A. I. (1994). Diagnostic checking of periodic autoregression models with application. *Journal of Time Series Analysis*, *15*, 221–233.

Osborn, D. R. (1988). Seasonality and habit persistence in a life cycle model of consumption. *Journal of Applied Econometrics*, *3*, 255–266.

Osborn, D. R. (1990). A survey of seasonality in UK macroeconomic variables. *International Journal of Forecasting*, *6*, 327–336.

Osborn, D. R. and Smith, J. P. (1989). The performance of periodic autoregressive models in forecasting seasonal UK consumption. *Journal of Business & Economic Statistics*, *7*, 117–127.

Pesaran, M. H., Shin, Y., and Smith, R. J. (1996). Testing for the existence of a long–run relationship. Forthcoming.

Ramsey, J. B. (1969). Tests for specification errors in classical linear least squares regression analysis. *Journal of the Royal Statistical Society B*, *31*, 350–371.

Reimers, H. E. (1992). Comparisons of tests for multivariate cointegration. *Statistical Papers*, *33*, 335–359.

Sargan, J. D. and Bhargava, A. (1983). Testing residuals from least squares regression for being generated by the Gaussian random walk. *Econometrica*, *51*, 153–174.

Taylor, A. M. R. (1997). On the Practical Problems of Computing Seasonal Unit Root Tests. Department of Economics and Related Studies, University of York. Forthcoming in the *International Journal of Forecasting.*

Tiao, G. C. and Grupe, M. R. (1980). Hidden periodic autoregressive–moving average models in time series data. *Biometrika*, *67*, 365–373.

Urbain, J. P. (1992). On weak exogeneity in error–correction models. Research paper 9103, C.R.E.D.E.L. University of Liège. Reprinted in *Oxford Bulletin of Economics and Statistics, 54:187–207.*

Vaage, K. (1993). The Dynamics of Residential Electricity Demand: Empirical Evidence from Norway. Working Paper No. 0193, Department of Economics, University of Bergen, Norway.

White, H. (1980). A heteroskedastic–consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica*, *48*, 817–838.

Appendix 1: Data description

The UK quarterly data set used is for the time period 1975Q4 to 1996Q3, the annual data cover the years 1976 to 1995. All variables are expressed in logarithms, where applicable [ ] denotes the appropriate filename on the CSO (or better: ONS) tape.

q denotes the log of final domestic energy consumption (not temperature corrected). The values for the period 1992Q2 to 1996Q3 have been converted from the new reporting unit 1,000 tonnes of oil equivalent (toe) to million therms (conversion rate 1 toe = 396.8 therms). (DATA SOURCE: *Energy Trends*, various issues)

y is log of real personal disposable income (indexed), not seasonally adjusted and revalued by the implied consumers' expenditure deflator (1990=1). (DATA SOURCES: CSO tape [CECQ], Monthly Digest, various issues)

p denotes the retail price index (RPI) for fuel and light, relative to the seasonally adjusted GDP deflator at market prices. From 1994Q2 the series includes 8% VAT for coal and coke, electricity, and heating oils, respectively. (DATA SOURCES: CSO tape [CHBC], Monthly Digest, various issues)

h stands for log of heating degree days (HDD), a series constructed from published monthly HDD data for 18 regions in the UK (corresponding to 18 measurement points, base temperature 15.5 °C), annual population data by county for England, Wales and Scotland, as well as for Northern Ireland (with estimates used for 1995 and projections for 1996, respectively), and a region–county match table kindly supplied by the Department of Trade and Industry (DTI). (DATA SOURCES: (i) population data: OPCS and Gen. Reg. Office for Scotland, as provided by DTI; Northern Ireland Abstract of Statistics (until 1981 Digest of Statistics Northern Ireland). 1995 and 1996 data have been kindly provided by ONS, the General Register Office for Scotland, and the Statistical Directorate of the Welsh Office, respectively; (ii) monthly HDD data: 1975–94 data kindly supplied by DTI; 1995 and 1996 data obtained from the Department of the Environment).

The complete data set used is available from the authors upon request.

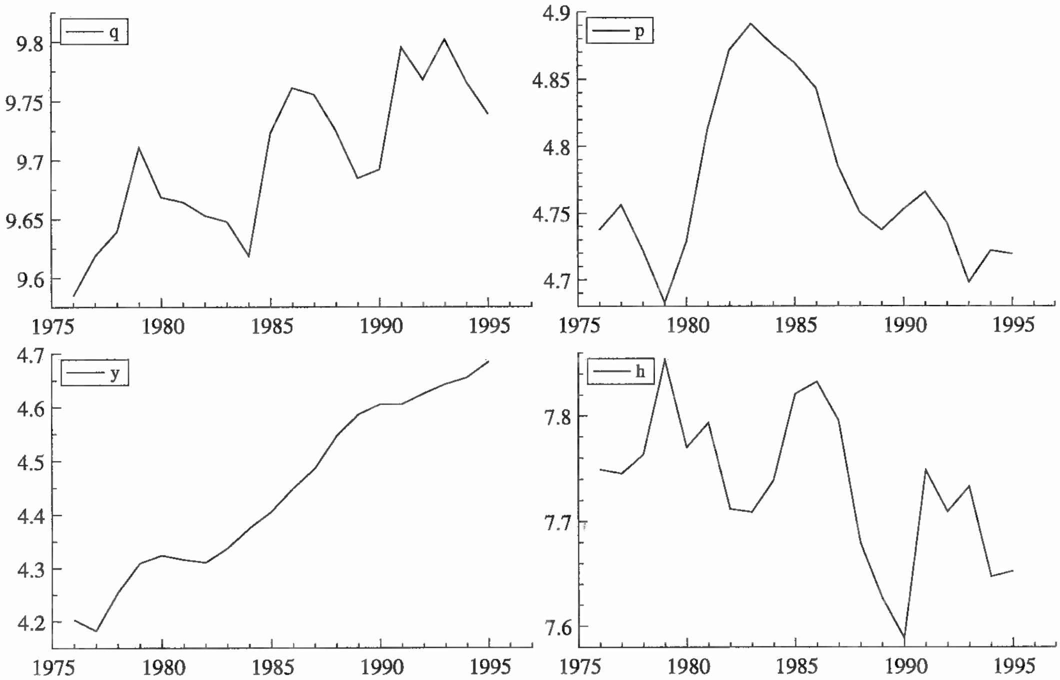

Figure 5: Annual data plots (all series in logs)

Appendix 2: Cointegration test results (annual data)

Table 10: Cointegration statistics (annual data)

μ_i		l_i				rank r
		196.341				0
0.8222		211.886				1
0.6336		220.921				2
0.0395		221.283				3
H_0: 'rank $= r$'	Max	$\text{Max}_{(T-mn)}$	95%	Trace	$\text{Trace}_{(T-mn)}$	95%
$r \leq 0$	31.09**	25.91**	21.0	49.89**	41.57**	29.7
$r \leq 1$	18.07*	15.06*	14.1	18.80*	15.66*	15.4
$r \leq 2$	0.73	0.60	3.8	0.73	0.60	3.8

standardised β' eigenvectors	q_t	p_t	y_t
$i = 1$	1.000	0.057	−0.495
$i = 2$	2.202	1.000	0.498
$i = 3$	3.107	−30.310	1.000
standardised α coefficients	$i = 1$	$i = 2$	$i = 3$
q_t	−0.821	−0.023	−0.001
p_t	0.020	−0.090	0.003
y_t	0.139	−0.014	−0.002

(NOTE: ** denotes significance at the 1% level, * at the 5% level.)