Detrending, Stylized Facts and the Business Cycle

A. C. Harvey; A. Jaeger

Journal of Applied Econometrics, Vol. 8, No. 3 (Jul. - Sep., 1993), 231-247.

Stable URL:
http://links.jstor.org/sici?sici=0883-7252%28199307%2F09%298%3A3%3C231%3ADSFATB%3E2.0.CO%3B2-I

Journal of Applied Econometrics is currently published by John Wiley & Sons, Ltd..

JOURNAL OF APPLIED ECONOMETRICS, VOL. 8, 231–247 (1993)

DETRENDING, STYLIZED FACTS AND THE BUSINESS CYCLE

A. C. HARVEY

Department of Statistics, London School of Economics, Houghton Street, London WC2A 2AE, UK

AND

A. JAEGER

Austrian Institute of Economic Research, A-1103 Vienna, PO Box 91, Austria

SUMMARY

The stylized facts of macroeconomic time series can be presented by fitting structural time series models. Within this framework, we analyse the consequences of the widely used detrending technique popularised by Hodrick and Prescott (1980). It is shown that mechanical detrending based on the Hodrick–Prescott filter can lead investigators to report spurious cyclical behaviour, and this point is illustrated with empirical examples. Structural time-series models also allow investigators to deal explicitly with seasonal and irregular movements that may distort estimated cyclical components. Finally, the structural framework provides a basis for exposing the limitations of ARIMA methodology and models based on a deterministic trend with a single break.

1. INTRODUCTION

Establishing the 'stylized facts' associated with a set of time series is widely considered a crucial step in macroeconomic research (see e.g. Blanchard and Fischer, 1989 chapter 1). For such facts to be useful they should (1) be consistent with the stochastic properties of the data and (2) present meaningful information. Nevertheless, many stylized facts reported in the literature do not fulfil these criteria. In particular, information based on mechanically detrended series can easily give a spurious impression of cyclical behaviour. Analysis based on autoregressive-integrated-moving average (ARIMA) models can also be misleading if such models are chosen primarily on grounds of parsimony.

We argue here that structural time-series models provide the most useful framework within which to present stylized facts on time series. These models are explicitly based on the stochastic properties of the data. We illustrate how, when these models are fitted to various macroeconomic time series, they provide meaningful information and serve as a basis for exposing the limitations of other techniques. These arguments have, to some extent, been made before (Harvey, 1985, 1989; Clark, 1987). They are further elaborated here. In addition, we examine the consequences of the mechanical detrending method suggested by Hodrick and Prescott (1980), which has recently started to become popular in macroeconomics (see e.g. Danthine and Girardin, 1989; Backus and Kehoe, 1989; Kydland and Prescott, 1990; Brandner and Neusser, 1992). We show that the uncritical use of mechanical detrending can lead investigators to report spurious cyclical behaviour. This point has also been made by Cogley

0883–7252/93/030231–17$13.50

Received September 1991
Revised June 1992

(1990). We argue that the structural framework provides further insights and that trends and cycles should be fitted simultaneously to avoid such pitfalls.

The plan of the paper is as follows. In Section 2 we lay out the basic framework of structural time-series modelling in this context. Section 3 provides an analysis of the consequences of detrending using the Hodrick–Prescott filter approach. Section 4 considers modelling and detrending of several macroeconomic time series. Section 5 discusses several issues including seasonal adjustment, trends with deterministic break points, and spurious cross-correlations between inappropriately detrended series. Section 6 draws together the main conclusions.

2. THE TREND PLUS CYCLE MODEL

Structural time-series models are set up explicitly in terms of components that have a direct interpretation (see Harvey, 1989). In the present context we postulate the appropriate model to be

$$y_t = \mu_t + \psi_t + \varepsilon_t, \qquad t = 1, \ldots, T \tag{1}$$

where y_t is the observed series, μ_t is the trend, ψ_t is the cycle, and ε_t is the irregular component. The trend is a local linear trend defined as

$$\mu_t = \mu_{t-1} + \beta_{t-1} + \eta_t \qquad \eta_t \sim \text{NID}(0, \sigma_\eta^2) \tag{2}$$

$$\beta_t = \beta_{t-1} + \zeta_t \qquad \zeta_t \sim \text{NID}(0, \sigma_\zeta^2) \tag{3}$$

where β_t is the slope and the normal white-noise disturbances, η_t and ζ_t, are independent of each other. The stochastic cycle is generated as

$$\psi_t = \rho \cos \lambda_c \psi_{t-1} + \rho \sin \lambda_c \psi_{t-1}^* + \varkappa_t \tag{4}$$

$$\psi_t^* = -\rho \sin \lambda_c \psi_{t-1} + \rho \cos \lambda_c \psi_{t-1}^* + \varkappa_t^* \tag{5}$$

where ρ is a damping factor such that $0 \leqslant \rho \leqslant 1$, λ_c is the frequency of the cycle in radians, and $\varkappa_t$ and $\varkappa_t^*$ are both $\text{NID}(0, \sigma_\varkappa^2)$. The irregular component is $\text{NID}(0, \sigma_\varepsilon^2)$ and the disturbances in all three components are taken to be independent of each other.

The trend is equivalent to an ARIMA(0,2,1) process. However, if $\sigma_\zeta^2 = 0$, it reduces to a random walk with drift. If, furthermore, $\sigma_\eta^2 = 0$ it becomes deterministic, that is, $\mu_t = \mu_0 + \beta t$. When $\sigma_\eta^2 = 0$, but $\sigma_\zeta^2 > 0$, the trend is still a process integrated of order two, abbreviated $I(2)$, that is, stationary in second differences. A trend component with this feature tends to be relatively smooth. An important issue is therefore whether or not the constraint $\sigma_\eta^2 = 0$ should be imposed at the outset. We argue that there are series where it is unreasonable to assume a smooth trend *a priori* and therefore the question whether or not σ_η^2 is set to zero is an empirical one. The examples in Section 4 illustrate this point.

The cyclical component, ψ_t, is stationary if ρ is strictly less than one. It is equivalent to an ARMA(2,1) process in which both the MA and the AR parts are subject to restrictions (see Harvey, 1985, p. 219). The most important of these is that the AR parameters are constrained to lie within the region corresponding to complex roots. Since the purpose is to model the possible occurrence of stochastic cycles, imposing this constraint *a priori* is desirable.

Estimation of the hyperparameters, $(\sigma_\eta^2, \sigma_\zeta^2, \sigma_\varkappa^2, \rho, \lambda_c, \sigma_\varepsilon^2)$, can be carried out by maximum likelihood either in the time domain or the frequency domain. Once this has been done, estimates of the trend, cyclical, and irregular components are obtained from a smoothing algorithm. These calculations may be carried out very rapidly on a PC using the STAMP package.

The model in equation (1) can be extended to deal with seasonal data. Thus there is no need to use data that may have been distorted by a seasonal adjustment procedure. Furthermore, if we are interested in stylized facts relating to seasonal components, structural time-series models provide a ready tool to determine these components without imposing a deterministic structure on the seasonal pattern (see e.g. Barsky and Miron, 1989).

3. THE HODRICK–PRESCOTT FILTER

Nelson and Kang (1981) have drawn attention to the distortions that can arise from fitting deterministic trends to series actually driven by stochastic trends. Similarly, it has long been known that the mechanical use of moving average filters can create a wide range of undesirable effects in the data (see Fishman, 1969). The filter adopted by Hodrick and Prescott (1980), hereafter denoted HP filter, has a long tradition as a method of fitting a smooth curve through a set of points. It may be rationalized as the optimal estimator of the trend component in a structural time-series model

$$y_t = \mu_t + \varepsilon_t \qquad t = 1, \ldots, T \tag{6}$$

where μ_t is defined by equation (2) and (3) but with σ_η^2 set equal to zero.[1]

Of course, the reason for estimating the trend in the present context is to detrend the data. The optimal filter which gives the detrended observations, y_t^d, is, for large samples and t not near the beginning or end of the series

$$y_t^d = \left[\frac{(1-L)^2(1-L^{-1})^2}{q_\zeta + (1-L)^2(1-L^{-1})^2}\right] y_t \qquad q_\zeta > 0 \tag{7}$$

where $q_\zeta = \sigma_\zeta^2/\sigma_\varepsilon^2$ and L is the lag operator. This expression can be obtained as the optimal estimator of ε_t in equation (6) by means of the standard signal extraction formulae which, as shown by Bell (1984), apply to non-stationary, as well as stationary, time series.[2]

If equation (6) was believed to be the true model, q_ζ could be estimated by maximum likelihood as outlined in the previous section. However, the whole reason for applying the HP filter is the belief that detrended data consist of something more than white noise. Thus, a value of q_ζ is imposed, rather than estimated. We will denote this value of q_ζ by $\bar{q}_\zeta$. From the point of view of structural time-series modelling, HP filtering is therefore equivalent to postulating model (1) and imposing the restrictions $\sigma_\zeta^2/\sigma_\varepsilon^2 = \bar{q}_\zeta$, $\sigma_\eta^2 = 0$, and $\psi_t = 0$. The HP estimate of the cyclical component is then simply given by the smoothed irregular component.

Given a particular model for y_t, the process followed by the HP detrended series, y_t^{HP}, and hence its dynamic properties, may be determined by substituting for y_t in equation (7). Suppose that we specify y_t as an ARIMA(p,d,q) process, possibly representing the reduced form of a structural time-series model such as equation (1). That is,

$$(1-L)^d y_t = \varphi^{-1}(L)\theta(L)\xi_t, \qquad \xi_t \sim \mathrm{NID}(0, \sigma^2) \tag{8}$$

where $\varphi(L)$ and $\theta(L)$ denote the AR and MA polynomials in the lag operator. Then

$$y_t^{\mathrm{HP}} = \left[\frac{(1-L)^{2-d}(1-L^{-1})^2}{\bar{q}_\zeta + (1-L)^2(1-L^{-1})^2}\right] \frac{\theta(L)}{\varphi(L)}\, \xi_t \tag{9}$$

---

[1] The continuous time version of this model can be used to rationalize a cubic spline fitted to data which may not be regularly spaced. See Wecker and Ansley (1983).

[2] The exact solution is given by the smoother obtained from the state space model. A fast algorithm is given in Koopman (1991). Such a smoother is valid even if the trend is deterministic, that is, $q_\zeta = 0$.

As noted by King and Rebelo (1989), y_t^{HP} is a stationary process for $d \leqslant 4$.

The autocovariance generating function (a.c.g.f.) of y_t^{HP} is given by

$$g_{HP}(L) = \left[\frac{(1-L)^{4-d}(1-L^{-1})^{4-d}}{[\bar{q}_\zeta + (1-L)^2(1-L^{-1})^2]^2}\right] g_y(L) \tag{10}$$

where $g_y(L)$ is the a.c.g.f. of $\Delta^d y_t$ and $\Delta = 1 - L$.

Note that if y_t is specified in terms of unobserved components, $g_y(L)$ may be obtained directly without solving for the ARIMA reduced form. Setting $L = \exp(i\lambda)$ gives the spectrum of y_t^{HP}, $f_{HP}(\lambda)$.[3] The spectrum may be calculated straightforwardly and it provides particularly useful information if we wish to study the possible creation of spurious cycles.

Specifying y_t to be a structural time-series model of the form (1) gives insight into the effects of detrending, since the contribution of each of the unobserved components to the overall spectrum, $f_{HP}(\lambda)$, can be assessed. To this end, rewrite model (1) in the single-equation form

$$y_t = \frac{\zeta_{t-1}}{\Delta^2} + \frac{\eta_t}{\Delta} + \psi_t + \varepsilon_t \tag{11}$$

The first term is integrated of order two, abbreviated as $I(2)$, the second term is $I(1)$, and the last two terms are stationary or $I(0)$. From model (10) we have

$$f_{HP}(\lambda) = \tau(\lambda)\{\sigma_\zeta^2 + 2(1-\cos\lambda)\sigma_\eta^2 + 4(1-\cos\lambda)^2[g_\psi(\lambda) + \sigma_\varepsilon^2]\} \tag{12}$$

where

$$\tau(\lambda) = \frac{1}{2\pi}\frac{4(1-\cos\lambda)^2}{[\bar{q}_\zeta + 4(1-\cos\lambda)^2]^2}$$

and $g_\psi(\lambda)$ is the spectral generating function of ψ_t.

More generally, suppose we have any unobserved components model with $I(2)$, $I(1)$, and $I(0)$ components. Then the transfer function for an $I(d)$ component is

$$\tau_d(\lambda) = 2^{(2-d)}(1-\cos\lambda)^{2-d}\tau(\lambda) \qquad d = 0, 1, 2 \tag{13}$$

The transfer function tells us the effect of the filter on the spectrum of the dth difference of an $I(d)$ component. Note that $\tau_2(\lambda) = \tau(\lambda)$.

For the macroec' nomic series they study, Kydland and Prescott (1990) set $\bar{q}_\zeta = 0{\cdot}000625$. Figures 1(a) to 1(c) show the transfer functions for $I(0)$, $I(1)$, and $I(2)$ components assuming this value for $\bar{q}_\zeta$ plotted against a period over the range $0 \leqslant 2\pi/\lambda \leqslant 4$. The filter for $I(0)$ components removes the very low frequency components, but frequencies corresponding to periods of less than 20 are virtually unaffected. Multiplying $\tau(\lambda)$ by $2(1-\cos(\lambda))$, on the other hand, produces a transfer function $\tau_1(\lambda)$, with a peak at

$$\lambda_{max} = \arccos[1 - \sqrt{0{\cdot}75\bar{q}_\zeta}] \tag{14}$$

which for $\bar{q}_\zeta = 0{\cdot}000625$ corresponds to a period of about 30. Thus applying the standard HP filter to a random walk produces detrended observations which have the characteristics of a business cycle for quarterly observations. Such cyclical behaviour is spurious and is a classic example of the Yule–Slutzky effect.

Spurious cycles can also emanate from the $I(2)$ component. The transfer function in Figure 1(c) has a peak at a frequency corresponding to a period of about 40. The nature of any spurious cyclical behaviour in the detrended observations depends on the relative importance

[3] Neglecting factors of proportionality.

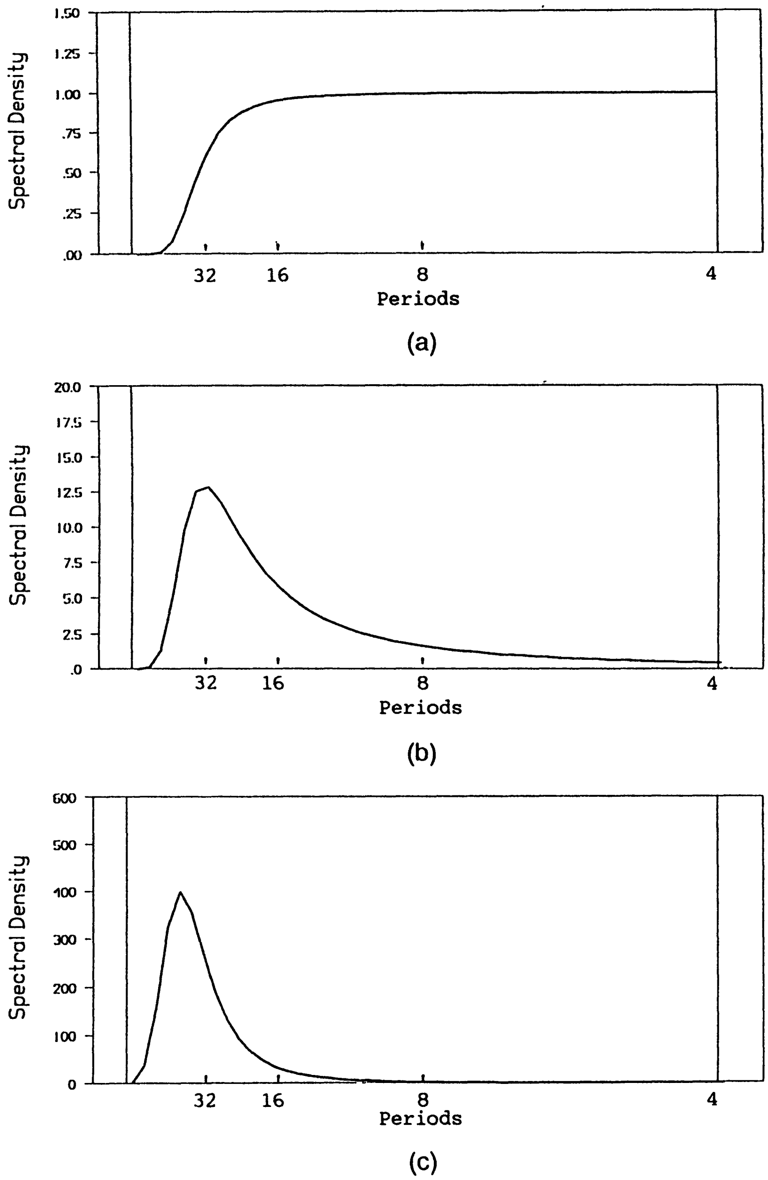

Figure 1. Transfer functions of HP filter. (a) $I(0)$ component; (b) $I(1)$ component; (c) $I(2)$ component

of $I(1)$ and $I(2)$ components. For data generated by (1) the peaks created in the spectrum are of similar height if $\sigma_\zeta^2/\sigma_\eta^2 \approx 25$. In this case they tend to merge together, and the overall effect is of a transfer function with a single peak corresponding to a period between 30 and 40.

4. MACROECONOMIC TIME SERIES

We now examine the stylized facts that can be produced by different techniques when applied to quarterly macroeconomic time series (all in logarithms). The series are US, real GNP, Austrian real GDP, the implicit deflator for US GNP, and the nominal value of the US monetary base.[4] All four series were seasonally adjusted by some variant of the Census X-11 program.

The HP filter was always applied with $\bar{q}_{\zeta}=0{\cdot}000625$. Attempts to estimate this ratio by applying maximum likelihood to model (6) usually produced a very high value of q_{ζ}, leading to the trend effectively picking up most of the movements in the stationary part of the series. Thus unless model (1) is a reasonable model for the series in question, q_{ζ} must be fixed in order to obtain sensible results.

Details of the results of fitting structural time-series model (1) are shown in Table I. Estimation was carried out in the frequency domain. Estimation in the time domain gave similar results and so these are not reported. Goodness of fit can be assessed by means of the estimated prediction error variance (σ^2), or, equivalently, by R_D^2 which is the coefficient of determination with respect to first differences. The Box–Ljung statistic, $Q(P)$, is based on the first P residual autocorrelations. The degrees of freedom for the resulting χ^2 statistic should be taken to be $P+1$ minus the number of estimated hyperparameters (see Harvey, 1989, chapter 5). Tests for normality and heteroscedasticity were also carried out. They are not reported in Table I and are only mentioned in the text if they were significant.

Estimating model (1) for real US GNP gives $\tilde{\sigma}_{\eta}^2=0$. Thus the result of unrestricted estimation is a relatively smooth trend. Furthermore, since $\tilde{\sigma}_{\varepsilon}^2=0$, the series effectively decomposes into a smooth trend plus a cycle (see Figure 2(a)). This is not surprising since $\tilde{\sigma}_{\zeta}^2$ is small and, coupled with the zero for $\tilde{\sigma}_{\eta}^2$, this means there is little contamination from non-stationary components in the series. Application of the HP filter yields a detrended series which is difficult to distinguish from the cycle extracted from the structural model (see Figure 2(b)). Thus applying the HP filter to real US GNP data is practically equivalent to estimating the structural time-series model (1). The striking coincidence between the estimated business cycle component and the HP cycle suggests that the HP filter is tailor-made for extracting the business cycle component from US GNP.

The close similarity between estimated and HP cycle reported for US GNP may not necessarily obtain for output series from other countries. To illustrate this point, we estimated model (1) using real Austrian GDP. Attempting to estimate the full model, (1), leads to the cyclical component virtually disappearing and we are led to a local linear trend model. On the other hand, if we impose $\sigma_{\eta}^2=0$ on (1) we obtain a smooth trend and a cycle. A graph of

Table I. Estimates of structural time-series models

Series	Time range	Restrictions	σ_{ζ}^2	σ_{η}^2	σ_{κ}^2	ρ	$2\pi/\lambda_c$	σ_{ε}^2	σ^2	R_D^2	$Q(P)$
US GNP	1954:1–89:4	None	8	0	625	0·92	22·2	0	937	0·05	8·01
Austrian GDP	1964:1–88:4	None	9	578	0	—	—	244	1126	0·05	13·63
		$\sigma_{\eta}^2=0$	21	—	36	0·97	13·0	438	1071	0·09	7·46
US Prices	1954:1–89:4	None	28	94	0	—	—	0	161	0·64	5·78
		$\sigma_{\eta}^2=0$	19	—	79	0·94	27·5	3	160	0·65	4·27
US monetary	1959:1–89:4	None	40	63	3	0·98	5·6	0	151	0·64	7·89
base		$\sigma_{\eta}^2=0$	47	—	25	0·73	6·0	0	153	0·64	10·68

Notes: All variance estimates have been multiplied by 10^7. $2\pi/\lambda_c$ is period (in quarters). $Q(P)$ is Box–Ljung statistic based on first P residual autocorrelations. $P=12$ for US Series and $P=10$ for Austrian GDP.

[4] The US series are taken from Citibase data bank. Austrian GDP data are taken from the data bank of the Austrian Institute of Economic Research.

the cycle shows it to be relatively small (see Figure 3(a)); it rarely deviates from the trend by more that 2 per cent. However, the cycles do coincide with the Austrian experience of somewhat muted booms and recessions. The cyclical component obtained from HP filtering, shown in Figure 3(b), is more volatile and quite erratic because it includes the irregular movements in the series which appear to be substantial.

The cyclical model obtained by imposing a smooth trend by setting $\sigma_\eta^2 = 0$ has a slightly better fit than the local linear trend model. The explanation lies in the fact that the local linear trend model emerges as a limiting case of the smooth trend and cycle model as $\rho \to 0$ and $\lambda_c \to 0$. Thus, when σ_κ^2 is quite small, as it is here, it is difficult to pick out the cycle in an

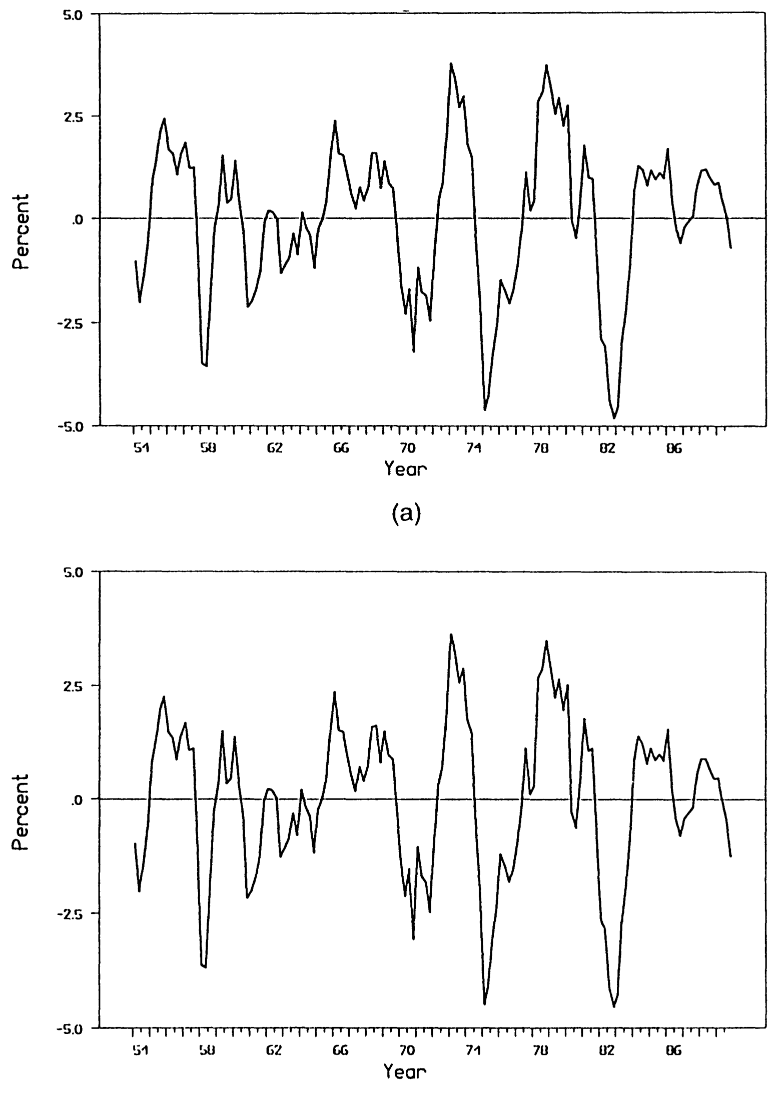

Figure 2. Business cycles in US GNP. (a) Estimated cyclical component; (b) Hodrick–Prescott cycle

unrestricted model since the likelihood function is very flat. The fact that the cycle model would be rejected on grounds of parsimony does not mean that it does not provide a valid description of the data. Furthermore, if we feel *a priori* that the underlying trend should be smooth then the cycle model is to be preferred over the more parsimonious local linear trend.

The two examples considered so far are based on real output series. Next we look at a price series and a nominal money stock series. Unrestricted estimation of model (1) for the implicit US GNP deflator leads quite clearly to a random walk plus noise model in first differences. This very simple model is also consistent with the correlogram of the second differences which is $-0\cdot47$ at lag 1 and $-0\cdot07$, $0\cdot04$, $0\cdot05$, $-0\cdot01$ for lags 2 to 5. Thus Box–Jenkins

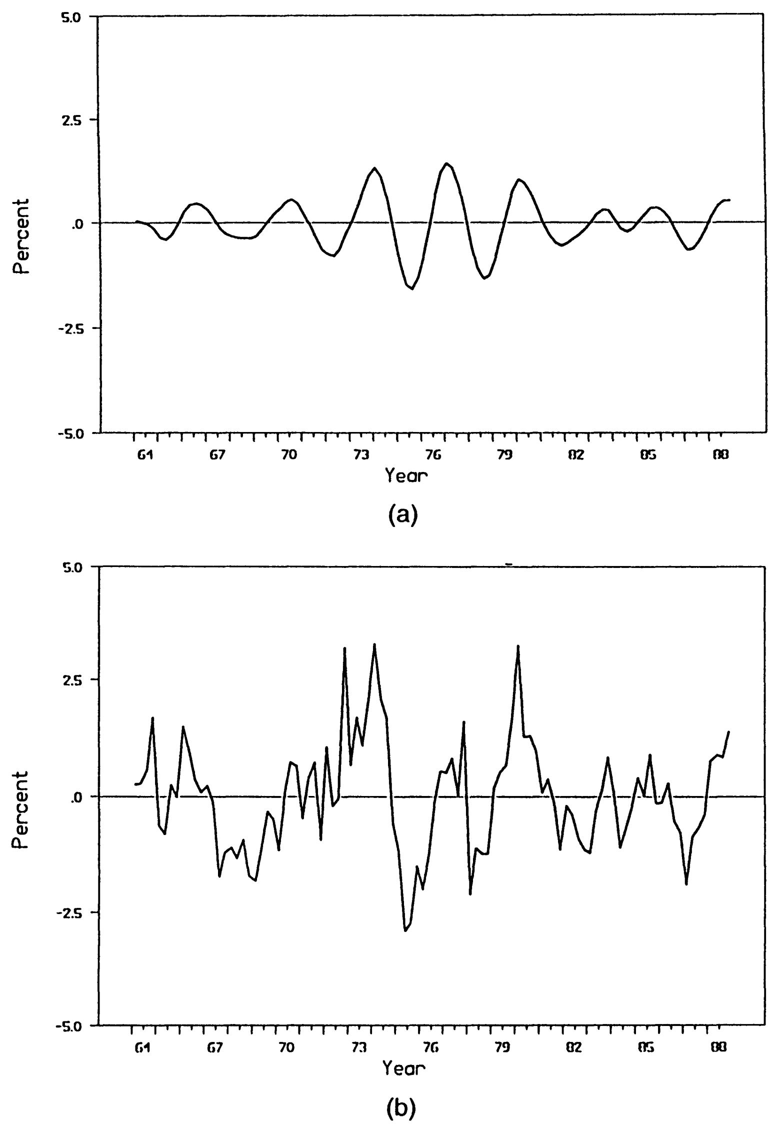

Figure 3. Business cycles in Austrian GDP. (a) Estimated cyclical component; (b) Hodrick–Prescott cycle

methodology would almost certainly select an equivalent model, namely ARIMA(0,2,1). Nevertheless, setting $\sigma_\eta^2 = 0$ does give a cycle and the model has essentially the same fit. It would clearly be rejected on grounds of parsimony, but it is consistent with the data and so cannot be dismissed, just as we could not dismiss the smooth trend/cycle model for Austrian GDP. However, while it may be reasonable to argue that a real series, such as GDP, contains a smooth trend, such an argument is less convincing for prices. Abrupt upwards or downwards movements in the price level can easily arise from indirect tax changes or oil price shocks, suggesting that the underlying trend is not smooth and contains an $I(1)$ component.

Applying the HP filter to the US price series yields what Kydland and Prescott (1990)

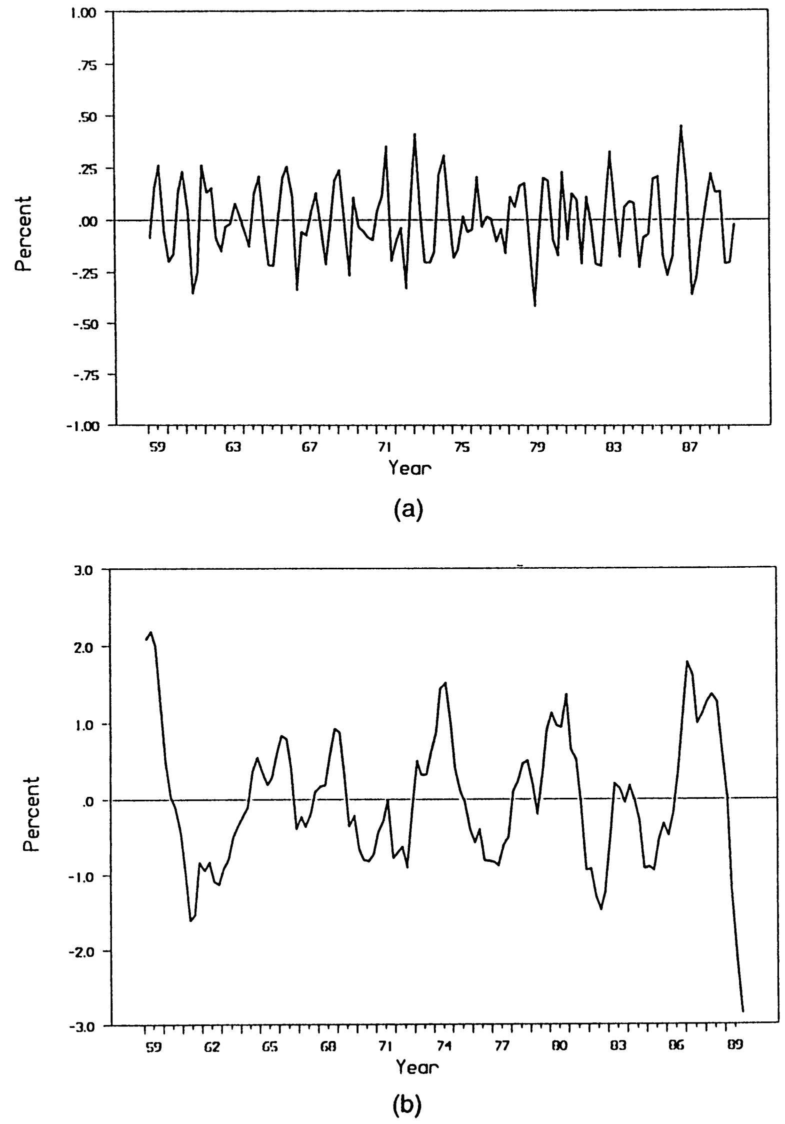

Figure 4. Business cycles in US monetary base. (a) Estimated cyclical component; (b) Hodrick–Prescott cycle

identify as cyclical movements. While we cannot rule out the possibility that the price level contains cycles, we note that the transfer function for our preferred model, the random walk plus noise in first differences, has a peak since it is a combination of the $\tau_2(\lambda)$ and $\tau_1(\lambda)$ filters shown in Figures 1(b) and 1(c). It is therefore possible that the price cycle identified by Kydland and Prescott (1990) is spurious.

For the US monetary base series, the unrestricted structural model is a local linear trend with a very small cycle. Setting $\sigma_\eta^2 = 0$ gives a model with basically the same fit and a cycle with a somewhat larger amplitude (see Figure 4(a)). The HP filter procedure imposes a smaller variance on the trend component and gives rise to a large cycle (see Figure 4(b)). This provides an illustration of how HP filtering may change substantially the volatility and periodicity properties of an estimated cyclical component.

5. FURTHER ISSUES

5.1. Seasonality

In common with most studies, the results reported above used seasonally adjusted data. Such data may not always have desirable properties, particularly if the seasonality pattern changes in some way, and is not of a kind that a standard adjustment method, such as the Bureau of the Census X-11, handles well. Data on real GDP for the United Kingdom provide a good example.[5] With the seasonally adjusted data and the restriction $\sigma_\eta^2 = 0$ imposed we estimated the cyclical component given in Figure 5(b). This cycle is not well defined and does not coincide particularly well with the known booms and recessions in the UK. On the other hand, seasonally unadjusted data produce much better results when a seasonal component[6] is added to model (1) (compare Figure 5(a) with 5(b)). The estimated seasonality pattern given in Figure 5(c) changes quite noticeably and the adjustment procedure presumably creates distortions in the series in attempting to cope with it.

5.2. ARIMA Methodology and Smooth Trends

ARIMA methodology usually results in the stylized fact that real output series are $I(1)$. Informal Box–Jenkins identification as well as formal root tests support this notion. For example, the first five autocorrelations of the first differences of real US GNP are 0·29, 0·19, −0·02, −0·10, −0·01. These autocorrelations show no need for second differencing of the data. A standard augmented Dickey–Fuller test rejects the null hypothesis of a second unit root in US GNP quite clearly. The relevant t-statistic is around −6·0, the precise value depending on the number of lags included. Thus, an ARIMA model of order (0,1,2) with a constant might be a reasonable selection. If we restrict attention to pure autoregressions and test down from a high number of lags an ARIMA(1,1,0) model with constant is obtained.

Neither of the above models is consistent with the structural time-series model (1) which has an ARIMA(2,2,4) reduced form. However, since σ_ζ^2 is relatively small, the $I(2)$ component may be difficult to detect by ARIMA methodology given typical sample sizes. To verify this conjecture we conducted two small Monte Carlo experiments. The data-generating process for

[5] The series run from 1960:1 1987:4. The seasonally adjusted series is taken from the OECD Main Economic Indicator data bank whereas the seasonally unadjusted series from the OECD Quarterly National data bank.
[6] The estimated seasonal component is modelled using the trigonometric formulation described in Harvey (1989, pp. 42–3).

both experiments is the estimated structural time series model for real US GNP reported in Table I. Table II reports the sample autocorrelations up to lag eight for the first differences of the generated series using sample sizes 100 and 500, respectively. Table III reports the empirical size of augmented Dickey–Fuller tests at the 5 per cent level for the null hypothesis that the first difference of the generated data has a unit root. The numbers of lags included in the Dickey–Fuller regression is fixed at 4, 8, and 16. The experiments are based on 500 replications. The results for $T = 100$ in Table II, confirm that much longer time series would be needed than commonly available to detect small but persistent changes in growth rates using ARIMA methodology. As regards the results for unit root tests reported in Table III, the findings of Schwert (1989) and Pantula (1991) are clearly applicable. They demonstrate that

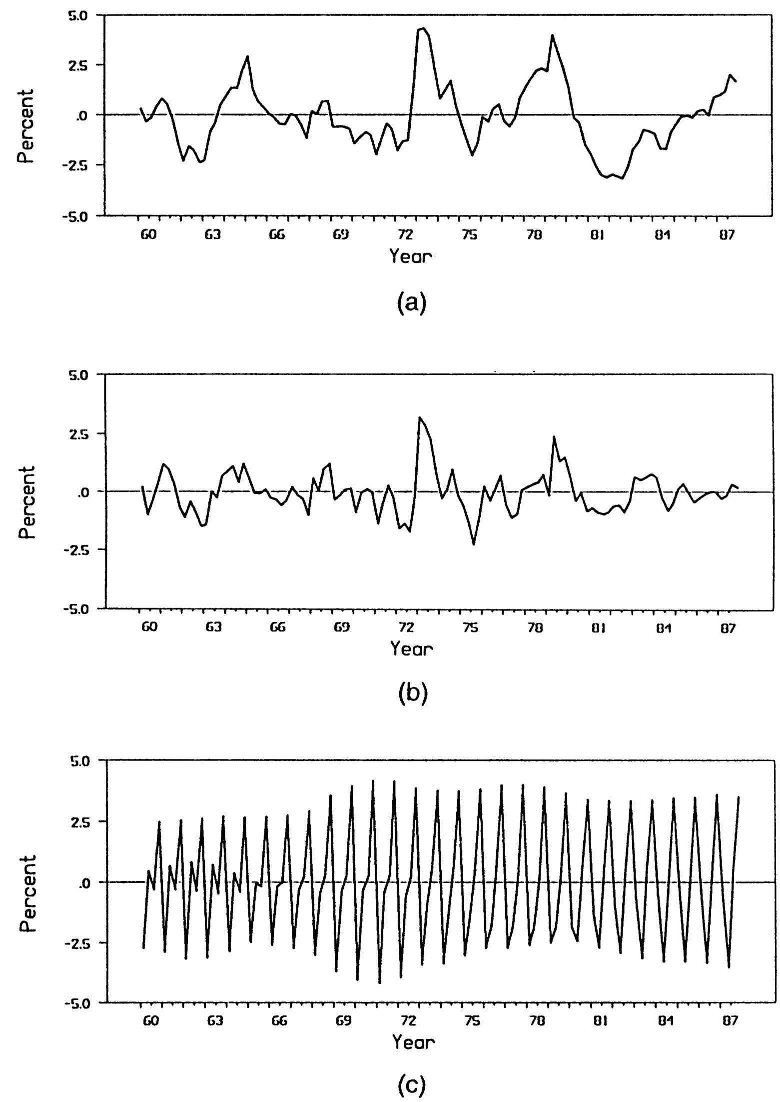

Figure 5. Business cycles in UK GDP. (a) Estimated cycle in seasonally unadjusted data; (b) estimated cycle in seasonally adjusted data (c) estimated seasonal component

if, after first differencing, we have an MA process which is close to being non-invertible, standard unit root tests will tend to reject the null hypothesis that a unit root is present with much higher frequency than the nominal test size. This tendency appears to be even more pronounced in a situation where a smooth local linear trend model (6) is appropriate since the reduced-form ARIMA (0,2,2) model will then have two roots close to the unit circle. For example, the results in Table III show that with a sample size of 100 and the number of lags in the Dickey–Fuller regression fixed at 8, the empirical size of the test is 0·74, exceeding the nominal 5 per cent size of the test substantially.

For purposes of short-term forecasting a parsimonious ARIMA model, such as ARIMA (1,1,0), may well be perfectly adequate compared with a trend plus cycle model. But as a descriptive device it may have little meaning and may even be misleading. For example, it may lead to a rejection of cyclical behaviour when such behaviour is, in fact, quite consistent with the data (see Harvey, 1985). Perhaps more important is the concept of 'persistence' associated with the identification of $I(1)$ models. A trend plus cycle model of the form (1) with $\sigma_\eta^2 = 0$ has stationary components with no persistence and a smooth $I(2)$ trend with infinite persistence. But since the trend is reflecting slow long-term changes in growth rates, perhaps

Table II. Autocorrelations of first differences of $I(2)$ process

	Autocorrelation at lag							
Sample	1	2	3	4	5	6	7	8
100	0·41	0·30	0·19	0·09	0·01	−0·05	−0·09	−0·11
	(0·12)	(0·12)	(0·11)	(0·12)	(0·13)	(0·13)	(0·14)	(0·14)
500	0·58	0·50	0·42	0·34	0·28	0·23	0·18	0·18
	(0·11)	(0·13)	(0·15)	(0·17)	(0·18)	(0·20)	(0·20)	(0·20)

Notes: This table presents the results of a Monte Carlo experiment. The data generating process is given by the estimated structural time-series model for US GNP. The table reports the mean of the autocorrelations of the first differences, and, in parentheses, standard deviations of the estimates. The results are based on 500 replications.

Table III. Empirical size of 5 per cent ADF-test for unit root in first differences of $I(2)$ process

	Empirical size		
Sample size	$k = 4$	$k = 8$	$k = 16$
100	0.90	0.74	0.25
500	0.99	0.94	0.38

Notes: This table presents the results of a Monte Carlo experiment. The data generating process is given by the estimated structural time-series model for US GNP. The table reports the empirical size of the 5 per cent level augmented Dickey–Fuller test of the null hypothesis of a unit root in the first differences. k denotes the numbers of lags included in the Dickey–Fuller regressions. The results are based on 500 replications.

arising from demographic changes, innovations in technology, changes in savings behaviour, or increasing integration of capital and goods markets, the shocks which drive the smooth trend may have no connection with short-term economic policy. Following the extensive literature on the productivity slowdown phenomenon, we may well argue that understanding the reasons for persistent changes in growth rates is one of the key problems in macroeconomics.

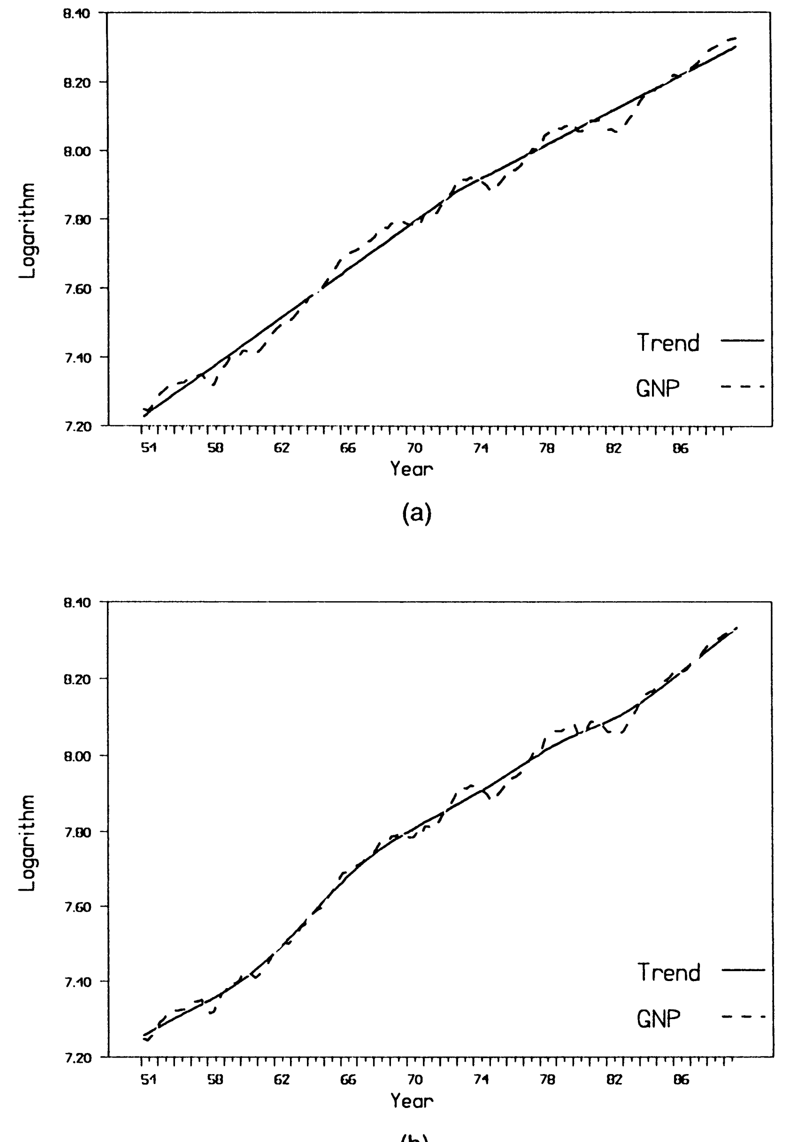

Figure 6. Segmented trends in US GNP. (a) Deterministic trend with break in 1973 : 1; (b) estimated trend component

5.3. Segmented Trends

It is sometimes argued that the trend component in economic time series is deterministic, but with a break at one or more points. We do not find the argument for such a trend particularly persuasive but if the data were really generated in this way, it is worth noting that a smooth trend within a structural time-series model would adapt to it. Thus the structural time series model would still give a good indication of the appropriate stylized facts. Indeed it is interesting to note that the trend component we estimated for US GNP shows a slowdown in the underlying growth rate in the late 1960s (see Figure 6(b)) and not in the first quarter of 1973 as maintained by Perron (1989) (see Figure 6(a)).[7] The imposition of exogenously determined breakpoints could therefore be potentially misleading and subject to many of the pitfalls associated with fitting deterministic trends to the series as a whole.

Making segmented trends more flexible by allowing several endogenously determined breaks also has a limited appeal. Such an approach is unnecessarily complicated and the conclusions could be very sensitive to the method used to choose the breaks. Structural models are not only likely to be more robust, but are also easier to fit.

5.4. Spurious Cross-correlations Between Detrended Series

The illustrative examples in Section 4 cast serious doubt on the validity of the cycles in the detrended price and monetary base series obtained using the HP filter. For US data, Kydland and Prescott (1990) draw wide-ranging conclusions about macroeconomic behaviour based on such data by examining sample cross-correlations. In particular, they argue that mainstream macroeconomic theory is inconsistent with a negative contemporaneous correlation of about $-0{\cdot}50$ for US data between HP detrended prices and real GNP.

In this section we use some of the results developed in Section 3 to study the possibility of spurious sample cross-correlations between spurious cycles. From the point of view of the structural time-series model (1), arbitrary cross-correlations can arise if one or both of the cyclical components is absent and the shocks of the trend components are correlated across series. In the following, we focus our attention on the analytically tractable case where spurious HP cycles are imposed on two series and the two series are independent by construction. First, note that the asymptotic distribution of the sample cross-correlations between two independent stationary series is asymptotically normal (AN) and given by (see e.g. Brockwell and Davis, 1987, p. 400)

$$\hat{r}_{12}(h) \sim \mathrm{AN}\left(0, T^{-1}\left(1+2\sum_{j=1}^{\infty} r_1(j)r_2(j)\right)\right) \tag{15}$$

where $\hat{r}_{12}(h)$ is the sample cross-correlation at lag h between two series with sample size T and $r_1(j)$ and $r_2(j)$ are the autocorrelations of the two sta tionary processes at lag j, respectively. The standard deviation of $\hat{r}_{12}(h)$ can be used to evaluate the probability of finding large spurious sample cross correlations between spurious cycles imposed on independent series. To evaluate the standard deviation of the sample cross-correlations we need the autocorrelations of the spurious HP cycles. As a benchmark case, assume we have two independent random walk processes

$$(1-L)y_{i,t} = \xi_{i,t}, \qquad \xi_{i,t} \sim \mathrm{NID}(0, \sigma_i^2) \tag{16}$$

[7] In fact, Nordhaus (1972) published a paper entitled 'The recent productivity slowdown' before the assumed breakpoint.

where $i = 1, 2$ and the $\xi_{i,t}$ are uncorrelated with each other. From equation (12), the spectra of the two HP filtered random walk processes are

$$f_{c,i}(\lambda) = \frac{8(1-\cos\lambda)^3}{[\bar{q}_\zeta + 4(1-\cos\lambda)^2]^2}\frac{\sigma_i^2}{2\pi} \tag{17}$$

The autocovariances of the HP-filtered processes may be calculated by taking the inverse Fourier transform of equation (17)

$$c(j) = \frac{\sigma_i^2}{2\pi}\int_{-\pi}^{\pi}\frac{\cos(\lambda j)8(1-\cos\lambda)^3}{[\bar{q}_\zeta + 4(1-\cos\lambda)^2]^2}\,d\lambda \qquad j = 0, 1, \ldots \tag{18}$$

and the autocorrelations are therefore given as $r(j) = c(j)/c(0)$ for $j = 1, \ldots$. Setting $\sigma_i^2 = 1\cdot 0$, the autocorrelations can be calculated by numerical integration up to some maximum lag $j_{\max}$.

Line 1 in Table IV reports the asymptotic standard deviations for the chosen benchmark case. Sample sizes T are 25, 100, and 500; $\bar{q}_\zeta$ is fixed at 0·000625; and the first 100 autocorrelations are used to approximate the infinite sum for the asymptotic variance defined in equation (15). If the sample size T is 100, the standard deviation for the sample cross-correlations $\hat{r}_{12}(h)$ is 0·20. Thus, given a normal distribution there is about a 30 per cent chance of finding spurious cross-correlations exceeding 0·20 in absolute value. To reduce the chance of finding spurious cross-correlations to about 5 per cent, cross-correlations have to exceed 0·40 in absolute value. If the sample size is as low as $T = 25$, the standard deviation increases to 0·41.[8] Even if the sample size is as large as 500, there is still a chance of about 5 per cent that the sample cross-correlations will exceed 0.18 in absolute value. If the two independent processes are specified as doubly integrated random walks, $(1-L)^2 y_{i,t} = \xi_{i,t}$, appropriately modified versions of equations (17) and (18) give the standard deviations reported in line 2 of Table IV. For $T = 100$, the standard deviation is 0·34 and so values of sample cross-correlations which are quite high in absolute value may easily arise under this specification for the two independent processes. These examples illustrate that the danger of finding large sample cross-correlations between independent but spurious HP cycles is not negligible. Furthermore, they strongly indicate that research on stylized business cycle facts should report standard errors in addition to point estimates of cross-correlations.[9]

Table IV. Asymptotic standard deviation of sample cross-correlations

	Standard deviation		
Process	$T = 25$	$T = 100$	$T = 500$
$(1-L)y_{i,t} = \xi_{i,t}$	0·41	0·20	0·09
$(1-L)^2 y_{i,t} = \xi_{i,t}$	0·67	0·34	0·15

Notes: This table reports the asymptotic standard deviations for the sample cross-correlations between two independent spurious HP cycles. $\bar{q}_\zeta$ is fixed at 0·000625.

[8] Monte Carlo experiments indicate that the asymptotic distribution in equation (15) approximates the actual small sample distribution well for sample sizes as low as $T = 25$.

[9] As an exception, Brandner and Neusser (1992) suggest the rule of thumb that cross-correlations between detrended series exceeding $2/\sqrt{T}$ in absolute value are significant at the 5 per cent level. From equation (15), however, this rule of thumb is misleading because it implicitly assumes that at least one of the detrended series is white noise.

6. CONCLUSIONS

Given the nature of macroeconomic time series, it is almost impossible to unambiguously obtain stylized facts from a single series. Instead we must be content with the less ambitious objective of extracting sets of stylized facts that are consistent with the data. It will often be possible to obtain several sets of stylized facts for a series and these may have very different implications. In such cases it is necessary to look for corroborating evidence from other sources.

We have argued in this article that because structural time-series models are formulated in terms of components that have a direct interpretation, they are a particularly useful way of presenting stylized facts. Furthermore, they provide a framework for assessing the limitations of stylized facts obtained by other methods. Our principal conclusions are as follows:

(1) ARIMA models fitted on the basis of parsimony may be uninformative and are sometimes misleading. A process integrated of order 2, or $I(2)$, is unlikely to be chosen in small samples using correlogram and standard unit root tests. The net result are simple $I(1)$ representations which are not consistent with a smooth trend plus cycle representation. If the latter representation is believed to be appropriate, measures of persistence associated with $I(1)$ models have little meaning.

(2) Pure autoregressive models are even more unlikely than ARIMA models to be consistent with trend plus cycle models. Furthermore, they have virtually no hope of adequately modelling the kind of changing seasonality that is to be found in the UK GDP series. These points need to be borne in mind when making inferences from vector autoregressions.

(3) The Hodrick–Prescott filter may create spurious cycles and/or distort unrestricted estimates of the cyclical component. This property of the Hodrick–Prescott filter may lead to misleading conclusions being drawn on the relationship between short-term movements in macroeconomic time series. A proper presentation of the stylized facts associated with a trend plus cycle view needs to be done within the framework of a model that fits both components at the same time.

ACKNOWLEDGEMENTS

We would like to thank Robert Kunst, Klaus Neusser, George Tiao, and Mark Watson for helpful comments on the original draft of this paper.

REFERENCES

Backus, D. K., and P. J. Kehoe (1989), 'International evidence on the historical properties of business cycles', Working Paper 402R, Federal Reserve Bank of Minneapolis.

Barsky, R. B., and J. A. Miron (1989), 'The seasonal cycle and the business cycle', *Journal of Political Economy*, **97**, 503–35.

Blanchard, O. J., and S. Fischer (1989), *Lectures in Macroeconomics*, MIT Press, Cambridge.

Bell, W. (1984), 'Signal extraction for nonstationary time series', *The Annals of Statistics*, **12**, 646–664.

Brandner, P., and K. Neusser (1992), 'Business cycles in open economies: Stylized facts for Austria and Germany', *Weltwirtschaftliches Archiv*, **128**, 67–87.

Brockwell, P. J., and R. A. Davis (1987), *Time Series: Theory and Methods*, Springer, New York.

Clark, P. (1987), 'The cyclical component of U.S. economic activity', *Quarterly Journal of Economics*, **102**, 797–814.

Cogley, T. (1990), 'Spurious business cycle phenomena in HP filtered time series', Mimeo, University of Washington.

Danthine, J. P., and M. Girardin (1989), 'Business cycles in Switzerland. A comparative study', *European Economic Review*, **33**, 31–50.

Fishman, G. S. (1969), *Spectral Methods in Econometrics*, Harvard University Press, Cambridge.

Harvey, A. C. (1985), 'Trends and cycles in macroeconomic time series', *Journal of Business and Economic Statistics*, **3**, 216–27.

Harvey, A. C. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press, Cambridge.

Hodrick, R. J., and E. C. Prescott (1980), 'Postwar U.S. business cycles: An empirical investigation', Discussion Paper No. 451, Carnegie-Mellon University.

King, R. G., and S. T. Rebelo (1989), 'Low frequency filtering and real business cycles', Working Paper No. 205, University of Rochester.

Koopman, S. (1991), 'Efficient smoothing algorithms for time series models', Mimeo, London School of Economics.

Kydland, F. E., and E. C. Prescott (1990), 'Business cycles: Real facts and a monetary myth', *Federal Reserve Bank of Minneapolis Quarterly Review*, **14**, 3–18.

Nelson C. R., and H. Kang (1981), 'Spurious periodicity in inappropriately detrended time series', *Econometrica*, **49**, 741–51.

Nordhaus, W. D. (1972), 'The recent productivity slowdown', *Brookings Papers on Economic Activity*, **3**, 493–546.

Pantula, S. G. (1991), 'Asymptotic distributions of unit-root tests when the process is nearly stationary', *Journal of Business and Economic Statistics*, **9**, 63–71.

Perron, P. (1989), 'The great crash, the oil shock, and the unit root hypothesis', *Econometrica*, **57**, 1361–1401.

Schwert, G. W. (1989), 'Tests for unit roots: A Monte Carlo investigation', *Journal of Business and Economic Statistics*, **7**, 147–160.

Wecker, W. E., and C. F. Ansley (1983), 'The signal extraction approach to nonlinear regression and splint smoothing', *Journal of the American Statistical Association*, **78**, 81–9.