Ricerche Economiche (1996) **50**, 135–162

Three tests for the existence of cycles in time series

FABIO CANOVA*†

Dipartimento di Economia, Universita' di Modena, I-41100 Modena, Italy

(Received 26 March 1993, accepted for publication 3 October 1994, final revision 23 December 1994)

Summary

Three tests for the presence of cycles in univariate time series are proposed. The asymptotic distribution of the tests is derived using the properties of the integrated periodogram and the small sample properties are examined using a Monte Carlo experiment. The tests are applied to U.S. data to detect the existence of significant seasonal and of other types of periodic fluctuations.

J.E.L. Classification: C15, C22.
Keywords: Spectrum, distance test, band pass test, seasonality, business cycles.

"His stories were good because he imagined them intensively, so intensively that he came to believe them . . ." Max Frish

1. Introduction

This paper describes three tests to assess the significance of cycles in economic time series.

There are several reasons for interest in such tests. Time series analysts often identify cycles with peaks in the spectral density of a time series (e.g. seasonal cycles in Sims, 1974). However, the question of the significance of these peaks has seldom been addressed. Business cycle practitioners are concerned with cyclical fluctuations in GNP and other variables, where cyclical fluctuations are measured as deviations from the trend of the process. Recent

† I would like to thank three anonymous referees, Steve Durlauf, Frank Diebold, Eric Ghysels, Bruce Hansen, Bruce Mizrach and Augustin Maravall, for helpful comments. A special thanks to Adrian Pagan for very useful suggestions.

Also at: Department of Economics, Universitat Pompeu Fabra, Balmes 132, E-08008 Barcelona, Spain; and CEPR.

0035–5054/96/020135 + 28 $18.00/0

literature on unit roots in macroeconomics has suggested that these types of cycles may not even exist (see, for example, Nelson & Plosser, 1982). In general, there is no insurance that deviations from the trend are not just contaminated noise and that no interesting cyclical fluctuations really exist. Similarly, recent literature on Bayesian learning (Nyarko, 1992) and on noisy traders in financial markets (Campbell & Kyle, 1993) propose models which generate irregularly spaced but significant cycles in economic activity and asset prices. Further, the recent political economy literature has argued that there are electoral cycles in government variables, for example, a periodicity of 4 years in the growth rate of government expenditure (Alesina & Roubini, 1992). Finally, a branch of the financial economics literature has examined the predictability of asset returns using particular speculative strategies in the short and in the long run (see, for example, Lo & MacKinley, 1988). The maintained hypothesis here is that efficiency implies martingale difference behaviour for these variables. Therefore, this literature is also interested in uncovering the presence of meaningful cycles in the data. Statistical tests which allow us to formally assess whether significant cycles exist are therefore useful to validate all these theories.

This paper presents a unified framework which can be used to assess the existence of cycles of any finite length in economic time series which are periodic or quasi-periodic with finite mean period. The tests are concerned with univariate time series, do not require *a priori* knowledge about which autocorrelations are important (as would be the case with time domain tests), are designed in the frequency domain and use the properties of the periodogram to derive the asymptotic distribution of the statistics of interest. The principle employed is general and the proposed procedures encompass tests for the existence of seasonal and cyclical fluctuations and of cycles of long but finite length as particular cases. The basic idea of the testing approach is that when cycles of mean duration r exist in the data, then their contribution to the total variance of the process is non-negligible. The three tests proposed here are alternative ways to measure how significant the contribution of these components to the variance of the process is. In the first case, the contribution of cycles of length r and of *all* its harmonics are considered. In the other two, the significance only of cycles of mean length r are considered.

The paper is organized as follows: Section 2 presents the definition of cycles employed and discusses the relationship of our concept with that of “hidden periodicity” recently employed by Hansen and Sargent (1993) and with the concept of cycles currently used in the macro literature (see Kydland & Prescott, 1990). Section 3 describes the three test statistics, their asymptotic properties and highlights the relationships among them. In Section 4 a Monte

Carlo study is performed to compute the small sample properties of the tests. I show that all three tests have reasonable size and power properties in small samples and that the only major source of misspecification occurs with the first test when there are peaks (or large masses) at harmonics of the frequency we care about. In Section 5, the methodology is applied to a number of post-Second World War U.S. macro series to detect the existence of seasonal, business and other interesting cyclical fluctuations. Section 6 provides some conclusions and discusses extensions to multivariate frameworks.

2. A definition of cycles

Let X_t be a general linear stochastic process with MA representation $X_t = g(\ell)e_t + \mu$, where $g(\ell) = g_0 + g_1\ell + g_2\ell^2, \ldots,$ is one-sided in non-negative powers of the lag operator ℓ, μ is the linearly deterministic component (possibly, a vector of initial conditions composed of sine and cosine functions) and e_t is an independently distributed white noise. It is typical to say that X_t exhibits cycles of mean period $r < \infty$ if the (non-normalized) spectral density of X_t, denoted by $h_x(\lambda)$, has a peak or a large mass at $\lambda_1 = (2\pi/r)$ and, possibly, one or more of its harmonics $\lambda_p = (2\pi p/r)$, $p = 2, \ldots, [r/2]$, where $[r/2]$ is the maximum integer less than or equal to $(r/2)$. This definition has been suggested by, for example, Granger and Newbold (1986), Sargent (1986), and has been used by Sims (1974) and Granger (1979) to identify the seasonal cycle in univariate time series.

One way to formalize the above notion of cycles is the following. Let the linearly deterministic (possibly purely periodic) component of X_t be modelled as an atom in the spectral measure. Then cycles of mean length $r < \infty$ exist if:

$$0 < \int_{-\pi}^{\pi} Q(\lambda)h_x(\lambda)d\lambda < \int_{-\pi}^{\pi} h_x(\lambda)d\lambda \tag{1}$$

where $Q(\lambda)$ is the transfer function of a linear filter $q(\ell)$ with the property that it has low power at some or all λ_p, $p = 1, \ldots, [r/2]$.

To intuitively understand why equation (1) captures the above notion of cycles suppose that the spectral density of X_t has a peak at $\lambda_1 = (2\pi/r)$, for example, X_t is a cosine function with period r plus white noise. Then if $Q(\lambda)$ has low power at λ_1, the contribution of the periodic component of length r to the total variance of X_t is removed and the "filtered variance" $\int_{-\pi}^{\pi} Q(\lambda)h_x(\lambda)d\lambda$ is smaller than the variance of the original series $\int_{-\pi}^{\pi} h_x(\lambda)d\lambda$. The idea of the tests is to determine whether the filtered variance is statistically different from the original one. One $Q(\lambda)$ function which can easily be used to illustrate the point we are interested in making is:

$$Q(\lambda)=|1-e^{-i\lambda r}|^2=2-2\cos(\lambda r) \tag{2}$$

in which case equation (1) reduces

$$0<\int_{-\pi}^{\pi} h_x(\lambda)d\lambda<2\int_{-\pi}^{\pi}\cos(\lambda r)h_x(\lambda)d\lambda \tag{3}$$

To gain further intuition on the meaning of equation (3) note that since $\cos(\lambda r)$ changes sign over $[-\pi,\pi]$, the expression on the right-hand side of equation (3) will be small (or even negative) unless the power of $h_x(\lambda)$ is concentrated in the region where $\cos(\lambda r)$ is large and positive. Hence, X_t has cycles of length r if $h_x(\lambda)$ has a sharp peak (or wide mass) in the neighbourhood of some or all λ_p. Note, however, that, because $\cos(\lambda r)$ is periodic mod$(2\pi/r)$, equation (3) does not distinguish between cycles at $(2\pi/r)$ or at one of its harmonics. In other words, equation (3) is consistent with X_t having only one peak (large mass) as well as having several peaks (large masses) at some or all λ_p.

The linear filter whose transfer function is given by equation (2) is the simple r-differencing operator. This filter is typically used in ARIMA modelling to take care of stochastic cycles of mean duration r and has a long history, in particular, when r is a seasonal period. In the time domain, equation (3) implies that cycles of mean period r exist if

$$\text{var}[(1-\ell^r)X_t]<\text{var}[X_t] \tag{4}$$

Noticing that $\text{var}[(1-\ell^r)X_t]=\text{var}(X_t)+\text{var}(X_{t-r})-2\,\text{cov}(X_t,X_{t-r})$, equation (4) implies that

$$\beta_r=\frac{\text{cov}(X_t,X_{t-r})}{\text{var}(X_{t-r})}>\frac{1}{2} \tag{5}$$

Hence, given the above choice of filter, cycles of length r exist if the regression coefficient of X_t on X_{t-r} exceeds 0·5, a result which corresponds to the notion that X_t displays cycles of period r if its correlogram shows high positive values at lag r [see Granger & Newbold (1986: p. 66) for a similar definition].

An example may further clarify how equation (5) captures the idea that there is a peak (or a large mass) in the spectral density at frequency λ. Let $X_t=-0{\cdot}8X_{t-1}+\varepsilon_t$ with $\varepsilon_t\sim(0,1)$. It is easy to check that $\beta_1 = -0{\cdot}8$, $\beta_2 = 0{\cdot}64$, $\beta_3 = -0{\cdot}48$, $\beta_4 = 0{\cdot}36$, etc. Using the rule that a cycle of length r exists if $\beta_r > 0{\cdot}5$, we find that X_t displays cycles of order 2, as intuition would suggest.

Given that many economic time series display unit root-like behaviour and Granger's (1966) "typical spectral shape", it is clear that equation (2) need not be the most obvious choice of $Q(\lambda)$ since it also removes power at the frequency zero and will therefore mistakenly give the impression that there are significant cycles of period r in the data even when only cycles with $r = \infty$ are present. However, because growth rates of variables are usually calculated before the spectrum is plotted and all the economic examples cited in the introduction are concerned with cycles in stationary series, this argument is of limited empirical importance as most economic time series have almost no power at the zero frequency after standard stationary inducing transformations are performed. If for some reason the time series still exhibits a peak of finite height at the zero frequency, all the previous arguments go through substituting the r-differencing operator with $q^*(\ell) = 1 + \ell - \ell^2 \ldots + \ell^{r-1}$) in which case equation (3) becomes:

$$0 < \int_{-\pi}^{\pi} \left| 1 - \sum_{j=1}^{r-1} (-1)^{j+1} e^{-i\lambda j} \right|^2 h_y(\lambda) d\lambda < \int_{-\pi}^{\pi} h_y(\lambda) d\lambda \qquad (6)$$

It is important to note that here we do not seek, in principle, any optimality property for $Q(\lambda)$, i.e. we do not require $Q(\lambda)$ to approximate any ideal filter. $Q(\lambda)$ is any arbitrary filter which has low power at some or all λ_p and makes equation (1) true.

As mentioned above, equation (1) is an appropriate definition of cycle only if $h_x(\lambda)$ exists everywhere on $[-\pi, \pi]$. Thus this paper is not interested in proposing yet another test for unit roots, in addressing questions concerning persistence, as in Cochrane (1988), or meaningful permanent components, as in Quah (1992). Instead, I focus on the problem of examining the significance of peaks (or wide masses) in the spectral density at some $0 < \lambda_p \leqslant \pi$, where λ_p is known *a priori* or provided by economic theory, which may be generated by, say, second order difference equations with complex roots. Because I start from the assumption that a researcher has some information regarding the location of interesting quasi-periodic components in the spectrum, the analysis of this paper differs from the ones based on Fisher's (1929) G-test where a researcher looks for cycles without specifying exactly the band of frequencies where they may appear.

A third important point that needs to be stressed is that equation (1) is a valid definition of cycles even when the data does not display purely periodic fluctuations. As a referee has pointed out, the concept of cycles employed in this paper relates to the power of the process over frequency bands (not to the magnitude of the peak) and it is sufficiently general to capture cycles which appear as stochastic with given mean periodicity (such as those present in crop series) or cycles which smoothly evolve over time (such as those evident in the electricity consumption series). One advantage of allowing for this level of generality is that the approach nicely relates to the Band Spectrum Regression setup initiated by Engle (1974) and avoids the problem of having to find consistent estimates of individual frequencies of the spectrum in designing test statistics.

Finally, since there is a large literature dealing with "spurious cycles" in time series (see Nelson & Kang, 1981; Cogley & Nason, 1995), it is not the purpose of this paper to discuss how incorrect stationary inducing transformations distort inferences about the presence of cycles. In the empirical section of the paper I do, however, briefly discuss how particular conclusions about the presence of cycles may be altered when alternative stationary inducing transformations are used.

It is useful to compare the notion of cycle proposed in this section with two other widely used concepts which have appeared in macroeconomic and time series literature. Tiao and Grupe (1980) and Hansen and Sargent (1993) have used the term "processes with hidden periodicities" to characterize Markov processes whose transition law is not time invariant but is strictly periodic with period r. Although there are similarities between their definition and the formalization of the notion of cycles employed in this paper, at least two differences should be noted. First, a process which has cycles according to equation (1) need not have a strictly periodic representation as an $r \times 1$ vector stochastic process. Second, because equation (1) is concerned with the variance of X_t, unless deterministic periodic components are modelled as point masses in the spectral density, the definition of cycles employed in this paper does not capture processes which are periodic in the mean (e.g. processes with seasonal initial conditions). Finally, note that both definitions are meaningful only when $r<\infty$ and $h_x(\lambda)$ exists everywhere on $[-\pi, \pi]$.

In the current macroeconomics literature, it is standard to consider the deviations from the trend of a process as representing cyclical fluctuations (see, for example, Kydland & Prescott, 1990). Because this notion does not coincide closely with the concept of cycle formalized in equation (1), it is useful to highlight the connection between the two concepts. According to the framework of this section, the concept of cyclicality used in the business cycle literature consists of the following three conditions: (i) $\int_{-\pi}^{\pi} F(\lambda)h_x(\lambda)d\lambda$ is large where $F(\lambda)$ is the transfer function of a high pass filter, i.e. a filter with $F(\lambda)=0$ for $0<|\lambda|<\bar{\lambda}$ and $F(\lambda)=1$ for $\bar{\lambda}<|\lambda|<\pi$ and $\bar{\lambda}$ is some predetermined frequency, (ii) $H(\lambda)=F(\lambda)h_x(\lambda)$ is significantly different from the spectrum of a white noise and (iii) for any chosen series (say W_t), $\int Co_{X,W}(\lambda)d\lambda$ is large, where $Co_{X,W}(\lambda)$ is the coherence of W_t and X_t at frequency λ. Note if X_t and W_t are two AR(1) processes with roots close to 1, they may satisfy all three of these conditions while they would not display cycles according to equation (1). In general, there is only a weak relationship between equation (1) and the above three conditions. I will come back to the relationship between the two concepts of cycle in the conclusions where I discuss how to extend equation (1) to multivariate frameworks. When a multivariate point of view is taken, the range of overlap between the two types of definitions becomes substantial.

3. The tests

To set up the tests I will use a slightly modified version of equation (1) with $I_{n,x}(\lambda)=(1/2\pi n)\,|\,\Sigma_{t=1}^{n} X_t e^{-i\lambda t}\,|^2$, the periodogram of X_t based on n observations at frequency λ (see, for example, Bloomfield, 1974: p. 78), in place of $h_x(\lambda)$. The rationale for using $I_{n,x}(\lambda)$ in place of $h_x(\lambda)$ is that an estimate of the former is easily obtained using the Fourier transform of the data. If e_t is independently and normally distributed, this substitution is innocuous since the periodogram is an asymptotically unbiased estimator of the spectrum (see Priestley, 1981: p. 425). If $h_x(\lambda)$ has bounded first derivatives, the order of magnitude of the bias is $O(\log(n)/n)$ which vanishes as $n\rightarrow\infty$. In addition, when $Q(\lambda)$ is a fixed bounded function independent of n, and since $\text{var}[I_{n,x}(\lambda)] \xrightarrow{n\rightarrow\infty} h_x^2(\lambda)$, for $\lambda\neq 0$, π and $\text{var}[I_{n,x}(\lambda)] \xrightarrow{n\rightarrow\infty} 2h_x^2(\lambda)$, for $\lambda=0$, π, then $\int Q(\lambda)I_{n,x}(\lambda)$ is a consistent estimator of $\int Q(\lambda)h_x(\lambda)$ (e.g. Priestley, 1981: p. 473).

For the rest of this section I make the following two assumptions:

Assumption 1: e_t is an independent Gaussian process with fixed variance σ_e^2.

Assumption 2: $\Sigma_j\,|g_j|\;|j|^{\zeta}<\infty$, $\zeta>0$.

These two assumptions require some discussion. The first is somewhat restrictive as it rules out conditionally heteroskedastic processes. For most applications this is not a strong limitation as the tests will be naturally applied to monthly and quarterly time series which hardly ever display any form of conditional heteroskedasticity. The condition imposed on the coefficients of the MA representation of X_t is stronger than required and ensures that the decay of the correlogram of X_t is sufficiently rapid (see Walker, 1965). It implies absolute and square integrability of the g_js and, therefore, second moment stationarity of X_t. However, it is more general than stationarity itself since it allows for the existence of cycles in non-stationary series which possess a smooth evolutionary spectrum (see Priestley, 1981: p. 828).

In the first two subsections I derive tests using two specific forms of $Q(\lambda r)$. Although the derivation of the tests for general $Q(\lambda r)$ functions does not present particular problems, in empirical applications one is faced with the problem of choosing a particular $Q(\lambda r)$ and different researchers may choose different functions. By selecting one particular family of $Q(\lambda r)$ functions, I attempt to avoid problems connected with a possible non-comparability of the results and this may be a preferable alternative to seeking a higher level of generality. In the first subsection $Q(\lambda r)$ is the transfer function of the r-differencing filter. In the next subsection, $Q(\lambda r)$ is the transfer function of a band-pass filter.

3.1. A DISTANCE-TYPE TEST

The idea of the test is simple. We want to know whether the difference between the two quantities in equation (3) is significant relative to their variances. Taking discrete approximations to the integrals in equation (3) at $\lambda_k=(2k\pi/n)$, $k=0, \ldots, [n/2]$, $n\geqslant N$, for some N, using $I_{n,x}(\lambda)$ in place of $h_x(\lambda)$ and exploiting the symmetry of $I_{n,x}(\lambda)$ around $\lambda=0$, equation (3) becomes

$$2\sum_k \cos(\lambda_k r) I_{n,x}(\lambda_k) > \sum_k I_{n,x}(\lambda_k) > 0 \tag{7}$$

The following lemma characterizes the asymptotic properties of the quantities in equation (7):

Lemma 1. *Let*

$$A_{1n}=\frac{M}{n}\sum_k W_1(\lambda-\lambda_k) I_{n,x}(\lambda_k)$$

and

$$A_{2n}=\frac{M}{n}\sum_{k} W_2(\lambda-\lambda_k)I_{n,x}(\lambda_k)$$

where

$$W_1(\lambda-\lambda_k)\ \mathrm{o}\ \frac{4\pi}{M}$$

$$W_2(\lambda-\lambda_k)\ \mathrm{o}\ \frac{4\pi\cos(\lambda_k r)}{M}$$

and where $M=M(n)$ is a parameter regulating the width of W_1 and W_2. Then

$$\lim_{n\ )\ x}\sqrt{v_1}A_{1n}\xrightarrow{D}\mathcal{N}(H_1, H_1^2)$$
$$\lim_{n\ )\ x}\sqrt{v_2}A_{2n}\xrightarrow{D}\mathcal{N}(H_2, H_2^2) \tag{8}$$

where

$$v_1=\frac{n}{\pi\,\Sigma_k\, h(\lambda_k)^2}$$

$$v_2=\frac{n}{\pi\,\Sigma_k \cos(\lambda_k r)^2 h(\lambda_k)^2}$$

$$H_1=\frac{4\pi}{n}\sum_{k} h_{n,x}(\lambda_k)$$

$$H_2=\frac{4\pi}{n}\sum_{k}\cos(\lambda_k r)h_{n,x}(\lambda_k)$$

Lemma 1 follows from the normality of e_t, the asymptotic independence of the normalized periodogram estimates and the fact that χ^2 variates with a large number of degrees of freedom behave like normal random variables. The only complication emerges because linear combinations of χ^2 variates with unequal weights are not necessarily χ^2. Using the trick discussed in Fuller (1981: pp. 295–296) it is possible to overcome this problem.

From lemma 1 it is clear that A_{1n} and $2A_{2n}$ are different but not independent kernel estimators of the same quantity with kernels given by $W_1(\lambda_k-\lambda)$ and $W_2(\lambda_k-\lambda)$. The test I propose is based on the idea that the difference between these two ways of estimating the same quantity is small under the null hypothesis (and large under the alternative) in the metric given by the variances of A_{1n} and $2A_{2n}$. Let

$$J_{1n} = \frac{1}{\sqrt{v_1}} \frac{A_{1n} - EA_{1n}}{\sqrt{\mathrm{var}(A_{1n})}}$$

and

$$J_{2n} = \frac{1}{\sqrt{v_2}} \frac{A_{2n} - EA_{2n}}{\sqrt{\mathrm{var}(A_{2n})}}$$

and consider the quadratic form:

$$B_n = (2J_{2n} - J_{1n})(\mathrm{var}(2J_{2n} - J_{1n}))^{-1}(2J_{2n} - J_{1n}) \tag{9}$$

Corollary 1. $\lim_{n\ )\ x} B_n = B \sim \chi^2(1)$.

Corollary 1 follows from the evaluation of the quantities of interest in the limit. Under the null hypothesis that no cycles of mean length r exist in the data, B_n will not significantly exceed a predetermined value Z_α at α% confidence level.

The implementation of the test is particularly simple since one needs only to specify a band of frequencies and therefore the mean duration r of the cycles he wants to detect. Typically, a researcher has *a priori* beliefs about a band of interesting frequencies he wants to investigate (for example, seasonals or business cycles). At other times, no *a priori* information is available so that one may plot the spectrum of the series and observe a peak at frequency λ_p. In this case, one may choose the r which corresponds to the frequency where the peak appears and proceed under the assumption that the peak is due to cycles of mean duration r, but be aware of the possibility that the observed peak may have been generated by spillovers due to cycles present at neighbouring frequencies. In other cases a researcher may spot a large mass in the periodogram in a band around frequency λ_p. One choice which apparently gives good results is to select r corresponding to the central frequency where the large mass appears. This implies that, although there are time variations in the length of the cycle over the years, the mean duration is exactly r. Note that with the current choice of $Q(\lambda)$ the number of frequencies which are suppressed by the filter is automatically determined by the resolution of $I_{n,x}(\lambda)$ and the sample size.

One case where the test is very simple to apply is when one attempts to detect non-exactly periodic but stable seasonal patterns. Seasonality appears if there is a peak (or a large mass) at some $\lambda_p = (2\pi p/s)$, $p = 1, \ldots, [s/2]$, where s is the number of seasons in the year. In this case, the transfer function of the s-differencing operator is the appropriate $Q(\lambda)$ function to use and we say that X_t exhibits (stochastic) seasonal behaviour if $\mathrm{var}[(1-\ell^s)X_t] < \mathrm{var}[X_t]$.

3.2. A TEST BASED ON BAND-PASS SERIES

The second test is also based on the implication that if a peak (or a large mass) at λ_p is significant, then the contribution of cycles of mean length (r/p) to the total variance of the process should be non-negligible. The major differences between the test presented in this subsection and the previous one are in the $Q(\lambda)$ function used and, consequently, in the null hypothesis being tested, and in the asymptotic distribution used to detect deviations from the null.

Let $\tilde{X}_t = X_t - f(X_t)$ where

$$f(X_t) = \int_{\Omega'} e^{i\lambda t} dZ_x(\lambda)$$

$$\Omega' = \left[\frac{2\pi k}{r} - \varepsilon, \frac{2\pi k}{r} + \varepsilon\right]$$

$Z_x(\lambda)$ is the spectral measure of X_t and the integral is of the Fourier–Stieltjes variety. $\tilde{X}_t$ is the filtered series and $f(\cdot)$ a band-pass filter which wipes out the power of X_t on Ω'. Let M be a parameter controlling the number of periodogram ordinates in a 2ε neighbourhood of λ_k and let

$$W_3(\lambda_k) \equiv \frac{h_x(\lambda_k)}{h_{\tilde{x}}(\lambda_k)}$$

Under H_0, $W_3 \approx 1$. Define

$$K_n = \sum_k \frac{I_{n,x}(\lambda_k)}{h_x(\lambda_k)}$$

and

$$\tilde{K}_n = \sum_k \frac{I_{n,x}(\lambda_k)}{h_{\tilde{x}}(\lambda_k)}$$

Lemma 2. *Under* H_0,

$$\lim_{n\,)\,x} 2\sqrt{\tilde{K}_n} \rightarrow \mathcal{N}\left(\sqrt{4\left(\left[\frac{n}{2}\right]-1\right.}, 1\right)$$

Under H_1,

$$\lim_{n\,)\,x} 2\sqrt{\tilde{K}_n} \rightarrow \mathcal{N}(\sqrt{2v_4-1}, 1)$$

where

$$v_4 = \frac{2n}{M\,\Sigma_k\, W_3^2(\lambda_k)^2}$$

Lemma 2 follows from the fact that under H_0, $2\,8(\tilde{K}_n)$ has the same asymptotic χ^2 distribution as $2\,8(K_n)$, while under the alternative the parameter of central tendency of the distribution is shifted to the right. Therefore, if cycles of mean length (r/k) give an important contribution to the total variance of the process, the tail of the distribution will contain a mass larger than expected. It is worth emphasizing that while the r-differencing filter used in Section 3.1 eliminates power at each λ_k, $k=1, \ldots, [r/2]$, while introducing extraneous power at frequencies in between, the filter used here sets to zero the power at one λ_k only and leaves unchanged the power outside a 2ε band centred around this frequency.

To implement this test we need two inputs: the mean duration (r/k) of the cycles we are interested in testing for and a band-pass filter which wipes out the power of X_t on Ω'. One such filter could be a standard MA or ARMA filter.

3.3. A TEST BASED ON BAND SPECTRUM VARIANCES

The final test described in this section also examines the behaviour of the spectrum in the band around some λ_k. But contrary to the two previous tests, I compare the magnitude of the *average* mass appearing inside a band of frequencies and outside of it. This test may be useful to detect the presence of cycles with slow time varying characteristics since they seldom generate peaks, although they tend to produce wide masses in a frequency band. To set up the test, consider the following two quantities:

$$C_{1n} = \frac{\Sigma_{\lambda_k \mathrm{v} \Gamma}\, I_{n,x}(\lambda_k)}{2\|\Gamma\|}$$

$$C_{2n} = \frac{\Sigma_{\lambda_k \mathrm{v} \Omega-\Gamma}\, I_{n,x}(\lambda_k)}{2\|\Omega-\Gamma\|} \qquad (10)$$

where $\Omega=[-\pi, \pi]$, and $\|\cdot\|$ represents the number of periodogram ordinates in the interval. C_{1n} and C_{2n} measure the average power of X_t inside the band Γ centred around the frequency λ_k, for some k, and outside the band Γ, respectively. If cycles of mean length (r/k) are the only ones existing in X_t, then the null hypothesis of the test is $C_{1n}=C_{2n}$, i.e. the average amount of power inside the band Γ is identical to the average amount outside of it. Notice that when X_t is a white noise, $C_{1n}=C_{2n}$. However, there are other processes which do not have a flat spectrum and may still satisfy the null hypothesis. For example, a process whose spectrum decays linearly with the frequency will have the required property. However, because these processes generate cycles of infinite duration (if any), they do not satisfy the conditions we imposed on equation (1) to be a meaningful definition of cycle. Under the alternative $C_{1n}>C_{2n}$, i.e. cycles of mean length (r/k) exist in the data.

Next I derive the distribution of the statistic $D_n=(C_{1n}/C_{2n})$ under the null hypothesis and under the alternative.

Lemma 3. *Under H_0 and as n) x ,*

$$G=2\|\Omega-\Gamma\|D_n\to\chi^2(2\|\Gamma\|)$$

Under H_1 and as n) x ,

$$G\to\frac{\chi^2(v_3)E(G)}{v_3}$$

where

$$v_3=\frac{32n\|\Gamma\|}{\Sigma_k\, h^2_{n,x}(\lambda_k)}$$

Lemma 3 employs the fact that under the null C_{1n} and C_{2n} are weighted averages of χ^2 random variables with equal weights while under the alternative the weights are frequency dependent. I normalize D_n by $2\|\Omega-\Gamma\|$ because the non-normalized quantity has a degenerate distribution as $n\to\infty$. Note that, in general, it is hard to predict the direction of the shift of the distribution under the alternative. However, for n large enough, $2\|\Gamma\|<v_3$.

To implement this test we also need two inputs: the mean duration (r/k) of the cycles we want to test for and the size of Γ.

3.4. DISCUSSION

Although all three tests are designed to assess the significance of cycles of mean length r, they reflect several differences. As already noted, the filter used in the distance-type-test knocks out power in the neighbourhood of each $(2\pi p/r)$, $p=1, \ldots, [r/2]$ and adds power in the neighbourhood of each $(2(p-1)\pi/r)$. There are two implications of this fact. First, if the series truly has cycles of mean duration r and we specify $r_1=\mu r$, $\mu=1, \ldots, [n/2]$, the null hypothesis will be steadily rejected. In other words, because the first test examines the size of the mass belonging to several frequency bands (the ones centred around cycles of mean duration r_1 and all of its harmonics), it is unable to distinguish whether it is the contribution of one frequency band or another which is significant. Second, a misspecified value for r may result in negative values of A_{2n}. Since lemma 1 requires $2A_{2n}-A_{1n}$ to be positive and since A_{1n} is positive everywhere on λ, the sign of A_{2n} provides a pre-test procedure to detect an inappropriate specification for r. For example, if X_t is a white noise, A_{2n} is non-positive for all r. Since the other two tests are concerned with only one particular frequency band, they do not display the same type of "aliasing" problem faced by the first test and can be used to clarify which of the harmonics of the basic frequency significantly contributes to the total variance of the process.

All three tests may face some problems when multiple peaks (or masses of similar size) are present in the spectrum. It is convenient to distinguish between two situations: one where a second peak (or a second large mass) occurs somewhere in the spectrum. Another where a second peak (or a second large mass) occurs within the frequency band we are concerned with. With the first two tests, if a peak of *larger* magnitude exists outside the frequency band we are interested in, we expect the null hypothesis to be rejected very infrequently since the contribution of cycles of mean duration r to the total variance of the process will be small. If, by any chance, this second peak occurs in a frequency band which is intermediate between λ_k and its harmonics, the first test will be able to detect this misspecification since A_{2n} will be negative. In general, if a researcher suspects that a time series displays multiple peaks of uneven size, it is advisable to plot the spectrum of the series, locate the peaks, roughly compare their magnitudes and run the test for the largest peak. With the third test, this problem is likely to be minor as the averaging procedure that the test employs will substantially reduce the effect of a secondary peak (or a secondary large mass) in the spectrum.

If two peaks of similar magnitude exist in the spectrum, it may be convenient to conduct the tests conditional on the removal of the other peak. That is, as in Bloomfield (1976: pp. 22–23) one could sequentially remove one of the peaks and test for the significance of the other. Although theoretically important, the presence of multiple peaks of similar magnitude in economic time series is a very rare event. One may object that the lack of detection of multiple peaks is due to the leakage from low order terms. It is well known in fact that most existing monthly and quarterly time series hardly ever deviate from Granger's (1966) "typical spectral shape" unless seasonals are present. Therefore, when growth rates and seasonally adjusted data are considered, as is standard in the macro literature, the relevance of the leakage problem for practical applications is probably very limited.

If a secondary peak emerges within the frequency band we are interested in or spills over from frequencies which are very close to the band, all three tests may have distorted size and power properties since the spillover effect from the neighbouring frequencies may be large. However, one should keep in mind that this is not a problem specific to the tests described in this paper and it is common to all spectral techniques examining the behaviour of a time series in a band of frequencies. As is well known, this problem is less important as sample size increases since the resolution of the spectrum improves and the spillover effect from neighbouring frequencies is reduced. For fixed sample size and to minimize distortions, it is recommendable to select the width of the band as in Andrews (1991) to minimize distortions and examine the robustness of the results by slightly varying the width of the band.

Finally, even though the alternative is the same in all three tests (i.e. the contribution of cycles of mean duration r to the variance of the process is large), the tests examine different null hypotheses. In the distance test the null hypothesis is that the total power appearing in frequency bands corresponding to cycle of mean duration r and all its harmonics is small. In the band-pass test the null hypothesis is that the contribution of the frequency band centred around cycle of mean duration r is small. In the average variance test the null hypothesis is that the average contributions to the total variance of the process of frequencies inside and outside a band are identical. Because of these differences, one should expect the tests to differ in their size properties, especially in small samples.

The first and third tests proposed here share features with what Priestley calls the "Bartlett homogeneity test" (Priestley, 1981: p. 487). That test was designed to check if independent estimates of the variance of the same quantity are significantly different. The statistic used, however, is slightly different from the ones employed here. The third test has also some relationship with Fisher (1929) and Whittle (1952) tests for jumps in the integrated spectrum. The major difference is that while the statistic they use takes the max periodogram ordinate to the sum of periodogram ordinates over the entire range of frequencies, here I take the average periodogram ordinate over a band to an average of periodogram ordinates over the remaining range of frequencies. The reason for choosing averages as opposed to the maximum is that there are many interesting situations where peaks may not be very sharp and yet there is a large mass concentrated around a particular frequency (e.g. when there are time varying seasonals or business cycle fluctuations). In this case the Fisher–Whittle test may fail to find meaningful economic cycles which are irregularly concentrated around a particular frequency while the averaging procedure employed here allows the test to detect the presence of periodic components if the mass appearing in a band is significant.

In independent work Durlauf (1991) designed a spectral-based test for the martingale hypothesis which is similar to the third test presented here. His formulation builds on work by Grenander and Rosenblatt (1957) and is more general than mine since it allows for non-normal and weakly dependent disturbances and for a generally specified alternative hypothesis. Durlauf's (1991) procedure has advantages and disadvantages. Because of its level of generality, his approach is free from data-mining activities which may affect the distribution of the test statistic under the null (see Hansen, 1990). However, there are many situations when a researcher has *a priori* knowledge about the possible location of interesting cycles in the data (e.g. seasonal or political cycles). In this case his tests may be less powerful in testing against a specific alternative than those described here. In addition, since Durlauf's (1991) tests are designed to assess general deviations from the white noise assumption, they cannot be used to examine questions such as: Is the *total* power at seasonal frequencies significantly different from a white noise? The distance-type test presented here, on the other hand, can be used for this purpose. Moreover, while it is very easy to verify the assumptions underlying the third test proposed in Section 3, it is much more complicated to examine whether Durlauf's mixing conditions are satisfied and this hampers the range of applicability of his test. Finally, while Durlauf's procedure is valid also for $r=\infty$, the tests designed here are appropriate only when $r<\infty$.

The third test presented in this paper is also very close in spirit to the *F*-test proposed by Engle (1974) to examine the hypothesis that band spectrum regression estimates are the same as regression estimates obtained using all the frequencies in the spectrum. The major differences between the two testing procedures are in the assumptions imposed on the disturbance term e_t and in the generic setup used (static regression framework in Engle, univariate time series models here).

4. The small sample power of the tests

This section describes the results of a Monte Carlo study designed to assess the small sample properties of the three tests described in Section 3. I use five different data generating processes (DGP):

$$\begin{aligned} \text{DGP(I)}\ & x_t = 5\cdot0 + bx_{t-1} + e_t \\ \text{DGP(II)}\ & x_t = 5\cdot0 + bx_{t-4} + e_t \\ \text{DGP(III)}\ & x_t = 5\cdot0 + bx_{t-4} + cx_{t-7} + dx_{t-20} + e_t \\ \text{DGP(IV)}\ & x_t = 5\cdot0 + e_t + fe_{t-1} \\ \text{DGP(V)}\ & x_t = e_t \end{aligned} \tag{11}$$

where e_t is an i.i.d. $\mathcal{N}(0, 1)$ random variable. Initial conditions x_{-s}, $s=0, 1, \ldots, 20$ and e_0 are set equal to zero. In DGP(I) b is equal to $(-0\cdot9, -0\cdot2)$ while in DGP(II) b is equal to $\pm0\cdot9$, $\pm0\cdot2$). In DGP(III) the values for the triplet (b, c, d) are either $[-0\cdot68, 0\cdot16, -0\cdot34]$ or $[0\cdot80, -0\cdot22, 0\cdot30]$. In DGP(IV) $f=(0\cdot8, 0\cdot2)$.

The first DGP covers the case of cycles with a periodicity of two while DGP(II) covers the case of cycles of four and eight periods and power at all the harmonics. When $b=\pm0\cdot2$, both DGPs generate samples with very small peaks in the spectral density so that cycles are insignificant. Hence we should expect all three tests to reject the null approximately 5% of the time, regardless of the value of r chosen. When $b=\pm0\cdot90$, there are sharp peaks in the spectral density so that the tests will have good power if they reject frequently when r is correctly specified and have good size if they reject infrequently when r is misspecified. DGP(III) covers the case of high order dynamics with significant cycles of 14 and 20 periods, respectively, and power at most of their harmonics. Once again we expect the tests to reject frequently when r is correctly specified and to reject infrequently when r is misspecified. The fourth experiment is designed to examine the size of the tests when the DGP generates time series with a short memory and not very significant serial correlation. With this DGP the tests should reject infrequently, regardless of the value of r chosen. The last experiment similarly examines the size of the tests when the underlying DGP is a white noise. Here to implement the tests I search for the highest peak in the periodogram and test for its significance.

For each parameter setting two sample sizes, $N=60$ and $N=154$, are considered. While the first sample size is arbitrary, the second sample size is typical of those series used in Section 5. The number of ordinates taken in the discrete Fourier transform is 200 for both sample sizes. The window Γ is chosen to contain three periodogram ordinates when $N=60$ and five periodogram ordinates when $N=154$, as suggested by Andrews (1991), and the band over which the power of the spectral density is wiped out in the second test contains, depending on the sample size, either three or five ordinates. The number of replications in each case is equal to 1000.

The results of the experiment appear in Table 1 where I tabulate the percentage of rejections of the null hypothesis at the 5% level over replications. There are two numbers in each cell, the first is the percentage for $N = 60$, the second the percentage for $N = 154$. Note that since the samples are short, tests performed choosing large r should be interpreted with caution.

The results of the table are encouraging. Test 1 performs well and type I error is around its nominal size of 5% in almost all cases. The non-negativity constraint on A_{2n} provides a useful sufficient condition to check for misspecifications of r. As expected, the test has distorted size properties when the selected b generates significant cycles at harmonics of the frequency corresponding to the cycles we are testing for. Note also that the performance of the test is good, regardless of the sample size used.

The performance of the second test is reasonable. The size of the test is slightly larger than its nominal size even for the largest sample, and relative to the other two tests the power is lower. For this test, both the size and the power properties depend on the sharpness of the peaks appearing in the periodogram, which in turn depend on the number of observations and the magnitude of b. Borderline cases of processes with low correlation at lag r or specifications where the sample size is small produce a percentage of type II errors sufficiently large. With low parameter values, increasing the sample size does not always improve the performance of the test. The reason is that when a process has low serial correlation the peaks are "unclean", regardless of the sample size, and the power of the test is weakened when the contribution of these peaks spill over into frequencies which are outside the band we are examining.

Test 3 is quite accurate. In general, its performance improves with the sample size and is best at frequencies away from 0. When the sample size is small the test has difficulties in correctly assessing the contribution of low frequencies to the total variance of the process. Also, the size of the test is distorted when the true frequency band and the band we test for contain some common frequencies. For example, if the DGP generates cycles of 20 periods and we select a frequency band centred around cycles of 24 periods, large size distortions emerge. However, shrinking the width of the window Γ makes most of the problems disappear.

In conclusion, whenever the DGP has parameter values which induce sharp peaks in the periodogram of X_t, all three tests are equally accurate: the size of the tests is close to its nominal size even in small samples and the tests reject frequently when the null is false. When the sample size is particularly small and the DGP has parameter values which generate peaks in the periodogram of borderline magnitude, the first and the third tests outperform the second, with the distance test being subject only to a misspecification at the harmonics of the basic frequency.

TABLE 1 *Size and power properties of the tests. Percentage of rejections of the null hypothesis over 1000 replications*

(a) Data generating mechanism: $X_t = 5{\cdot}0 + bX_{t-1} + e_t$

Value of b	*Value of* r	*Test 1*	*Test 2*	*Test 3*
−0·90	2	92·8/94·5	89·9/84·7	91·0/93·4
	4	87·3/90·4	11·7/7·2	8·0/6·0
−0·20	2	7·9/5·9	12·3/9·9	17·3/13·8
	4	8·6/6·2	5·1/5·1	8·4/7·7

(b) Data generating mechanism: $X_t = 5{\cdot}0 + bX_{t-4} + e_t$

Value of b	*Value of* r	*Test 1*	*Test 2*	*Test 3*
−0·90	8	94·5/95·1	92·8/93·9	94·2/94·7
	4	85·3/86·5	5·1/5·9	4·8/5·0
+0·90	4	95·1/94·7	59·1/66·6	83·6/92·8
	24	65·3/81·2	7·6/6·9	6·2/6·1
−0·20	8	6·6/5·2	14·2/11·7	8·8/7·3
	4	5·1/5·1	7·3/7·2	6·1/5·7
+0·20	4	6·8/5·6	7·9/6·6	9·8/8·0
	24	5·3/5·1	10·8/9·7	7·0/6·3

(c) Data generating mechanism: $X_t = 5{\cdot}0 + bX_{t-1} + cX_{t-7} + dX_{t-20} + e_t$

Value of b, c, d	*Value of* r	*Test 1*	*Test 2*	*Test 3*
−0·68, 0·16, −0·34	14	92·6/94·5	90·3/92·8	88·6/91·7
	7	87·3/81·4	9·6/8·7	6·3/5·8
0·80, −0·22, 0·30	20	94·6/95·2	89·9/93·6	90·4/92·0
	8	5·8/5·3	8·1/7·3	6·6/6·1

(d) Data generating mechanism: $X_t = e_t + fe_{t-1}$

Value of f	*Value of* r	*Test 1*	*Test 2*	*Test 3*
0·80	8	5·2/5·1	5·9/5·1	6·3/5·7
	8	5·3/5·1	6·5/5·8	6·1/5·5
0·20	4	5·1/5·0	6·3/5·7	6·1/5·4
	8	4·8/4·8	6·4/5·8	5·8/5·2

(e) Data generating mechanism: $X_t = e_t$

	Value of r	*Test 1*	*Test 2*	*Test 3*
		5·6/4·9	5·8/5·1	7·2/6·1

Note: In each cell the first number refers to $N = 60$, the second to $N = 154$. For all cases e_t is i.i.d. $\mathcal{N}(0, 1)$.

5. Some applications

In this section I apply the three tests presented in Section 3 to the problem of detecting the presence of seasonals, of political cycles and of certain types of business cycle fluctuations. In this exercise I use quarterly seasonally non-adjusted U.S. data on the log of fixed investments (residential, non-residential, structures and inventories), and seasonally adjusted data for the log GNP, the log of the monetary base, the log of federal government expenditure and the log of the consumer price index for the period 46,1–85,4 (59,1–85,4 for the monetary base). The sources of the data appear in Barsky and Miron (1989).

Table 2 reports the results of testing the following hypotheses.

- Does seasonality exist in investment and inventory data? Is there still some form of seasonality after we extract deterministic seasonal components?
- Do GNP and Inflation exhibit significant business cycle fluctuations?
- Is there a "presidential election" cycle, i.e. is there a tendency for government expenditure and for the monetary base to follow a four-year pattern?

With quarterly data, and given the symmetry of the spectrum around $\lambda = 0$, business cycle frequencies are those with period between 16 and 24 quarters and seasonals appear at the annual (four quarters) and semi-annual (two quarters) frequency. Therefore, the second and the third tests will have two entries for each investment and inventory series, one for the yearly frequency and one for the semi-annual frequency. Also, because all of the series appear to contain a unit root, I log first difference all the data before the tests are performed (see panel A). Diagnostic testing on the residuals of an AR(5) regression on the log first difference of the data indicate that they are well behaved. In particular, I do not find any evidence of conditional heteroskedasticity in any of the series examined.

The results indicate that a yearly cycle is very significant for all investments and for the inventory series examined. All three tests suggest the presence of cycles at seasonal frequencies although the significance of the first test is marginal for non-residential fixed investments and inventories. Also, while for fixed residential and non-residential investments seasonals are reasonably modelled as purely periodic, investments in structures and inventories display significant seasonals even after a deterministic seasonal component is extracted. Note also that for the case of fixed non-residential investment the distance-type test detects the presence of significant seasonality after purely periodic patterns are taken into account while the other two do not. Hence, even though

TABLE 2 *Tests for cycles in U.S. data, sample 1946,1–1985,4*

	Test 1 P-*value*	*Test 2* P-*value*	*Test 3* P-*value*
Panel A: Log first differenced data			
		Seasonality	
Non-residential fixed investments	0·06	0·00/0·08	0·00/0·00
(with dummies)	0·00	0·34/0·56	0·06/0·06
Non-residential structures	0·00	0·00/0·03	0·00/0·00
(with dummies)	NA	0·19/0·18	0·00/0·00
Residential fixed investments	0·00	0·00/0·05	0·00/0·00
(with dummies)	NA	0·29/0·47	0·19/0·99
Inventories	0·06	0·39/0·66	0·00/0·01
(with dummies)	NA	0·59/0·81	0·00/0·31
		Business cycles	
GNP $r=16$	NA	0·76	0·35
GNP $r=24$	NA	0·92	0·40
Inflation $r=16$	NA	0·83	0·66
Inflation $r=24$	NA	0·88	0·81
		Political cycles	
Government expenditure $r=16$	NA	0·26	0·33
Monetary base $r=16$	0·58	0·22	0·07
Panel B: Hodrick and Prescott filtered data			
		Seasonality	
Non-residential fixed investments (with dummies)	NA	0·00/0·79	0·00/0·98
Non-residential structures (with dummies)	NA	0·00/0·08	0·00/0·05
Residential fixed investments (with dummies)	NA	0·55/0·03	0·00/0·98
Inventories (with dummies)	NA	0·73/0·31	0·00/0·18
		Business cycles	
GNP $r=16$	NA	0·18	0·08
GNP $r=24$	NA	0·12	0·03
Inflation $r=16$	NA	0·34	0·04
Inflation $r=24$	NA	0·08	0·00
		Political cycles	
Government expenditure $r=16$	0·24	0·33	0·04
Monetary base $r=16$	0·14	0·08	0·02

Note: In the case of seasonality the cells for the second and the third tests report the *P*-value at $(\pi/2)$ first and the *P*-value at π second. NA indicates that $2A_{2n}<0$.

none of the peaks at seasonal frequencies is significant after deterministic patterns are taken into account, the total mass appearing in bands centred around seasonal frequencies is significant relative to the variance of the process.

The table also indicates that GNP does not possess significant business cycle fluctuations (according to test 1, the total mass appearing in the band covering cycles of 16–24 quarters *and* its harmonics is small), and it does not have a sharp peak or a large mass anywhere in the region corresponding to cycles of 16–24 quarters (the second and the third test statistics are insignificant). Similarly, the inflation rate has no significant business cycle fluctuations according to all three tests.

As far as political cycles are concerned, government expenditure does not display any significant 4-year periodicity even though a small mass appears in the region of the periodogram centred around four-year cycles. Note that because in the third test the band covers frequencies corresponding to cycles ranging from 14 to 18 quarters, it is unlikely that this type of periodicity in government expenditure is of crucial importance in accounting for fluctuations in the U.S. economy. For the monetary base the results are different. There appears to be a peak in the spectral density of the series in the band centred around 4-year cycles, but it is of insignificant magnitude. However, the average mass appearing in the region centred around cycles of 4 years is significantly larger than the mass appearing outside the band.

Next, I briefly discuss the question of the sensitivity of the results to alternative ways of rendering the series stationary. As mentioned in Section 2, the paper is not concerned with the question of spurious cycles induced by erroneous transformations. However, it may be interesting to know from an economic point of view, if any of the results we have just discussed are sensitive to the preliminary transformation employed to eliminate possible unit roots from the data.

There are many ways to remove unit roots from time series which are alternatives to taking the log first order difference of the data. One transformation, typically used in the business cycle literature, is the Hodrick and Prescott (1980) filter. The Hodrick and Prescott (1980) filter decomposes a series into permanent and transitory components where the permanent component τ_t is obtained by minimizing $\Sigma_{t=1}^{T}(Y_t-\tau_t)^2+\lambda \quad \Sigma_{t=3}^{T}[(\tau_t-\tau_{t-1})-(\tau_{t-1}-\tau_{t-2})]^2$, where Y_t is the original series and λ is a smoothing parameter which, for quarterly data, is routinely set to 1600. The transitory component is obtained as $X_t=Y_t-\tau_t$. King and Rebelo (1993) have shown that the Hodrick and Prescott (1980) filter effectively removes up to four unit roots and is therefore a legitimate candidate to check the sensitivity of the results described in Table 2, panel A.

As I have argued elsewhere (Canova, 1993), the transfer function of the Hodrick and Prescott filter differs from the transfer function of the first order differencing filter so that we should expect the results to change with the stationary inducing transformation employed. Roughly speaking, log first order differencing eliminates most of the power in a large frequency band which goes from frequency 0 (infinite cycles) to frequency ($\pi/6$) (cycles of about 3 years) while the Hodrick and Prescott filter eliminates most of the power of cycles belonging to a band that goes from frequency 0 to frequency ($\pi/14$) (cycles of about 7 years). Therefore, while tests for seasonality should be largely unaffected by the preliminary transformation used, tests for business cycle or political cycle periodicity may give different results. To facilitate the comparison I report in Figure A plots of the log spectrum of log GNP, inflation, log of government expenditure and log of the monetary base using each of the two preliminary detrending transformations. Shaded regions correspond to frequencies in the band covering 16 to 24 quarter cycles. Spectra are computed smoothing the periodograms using tent windows containing nine periodogram ordinates. Table 2, panel B, where the results of testing the three hypotheses with Hodrick and Prescott detrending data are reported, confirms our expectations. For example, while tests 2 and 3 do not detect any significant 16–24 quarter cycles in inflation when the data is first order differenced, they do find some evidence of 24 quarter cycles when the data are made stationary with the Hodrick and Prescott filter. Also, when the Hodrick and Prescott filter is used, the third test suggests that the average mass of government expenditure and of the monetary base in frequency bands centred around cycles of 16 quarters is significantly larger than the mass outside it, while the opposite is true when the data is first order differenced. Finally, note that even for those cases where we do not expect significant changes in the results, e.g. in testing for significant seasonality in inventory and investments, we do find that the magnitude of the seasonal peak at the semi-annual frequency is somewhat altered when Hodrick and Prescott filtering is used, and in general, that the size of the mass at seasonal frequencies relative to the total mass depends on the type of first stage transformation used.

6. Conclusions

This paper describes three tests to assess the significance of cycles in univariate time series. The tests are based on the frequency domain features of the series, do not require parametric assumptions and employ the properties of the integrated spectrum to derive the asymptotic distribution of the tests. The paper shows

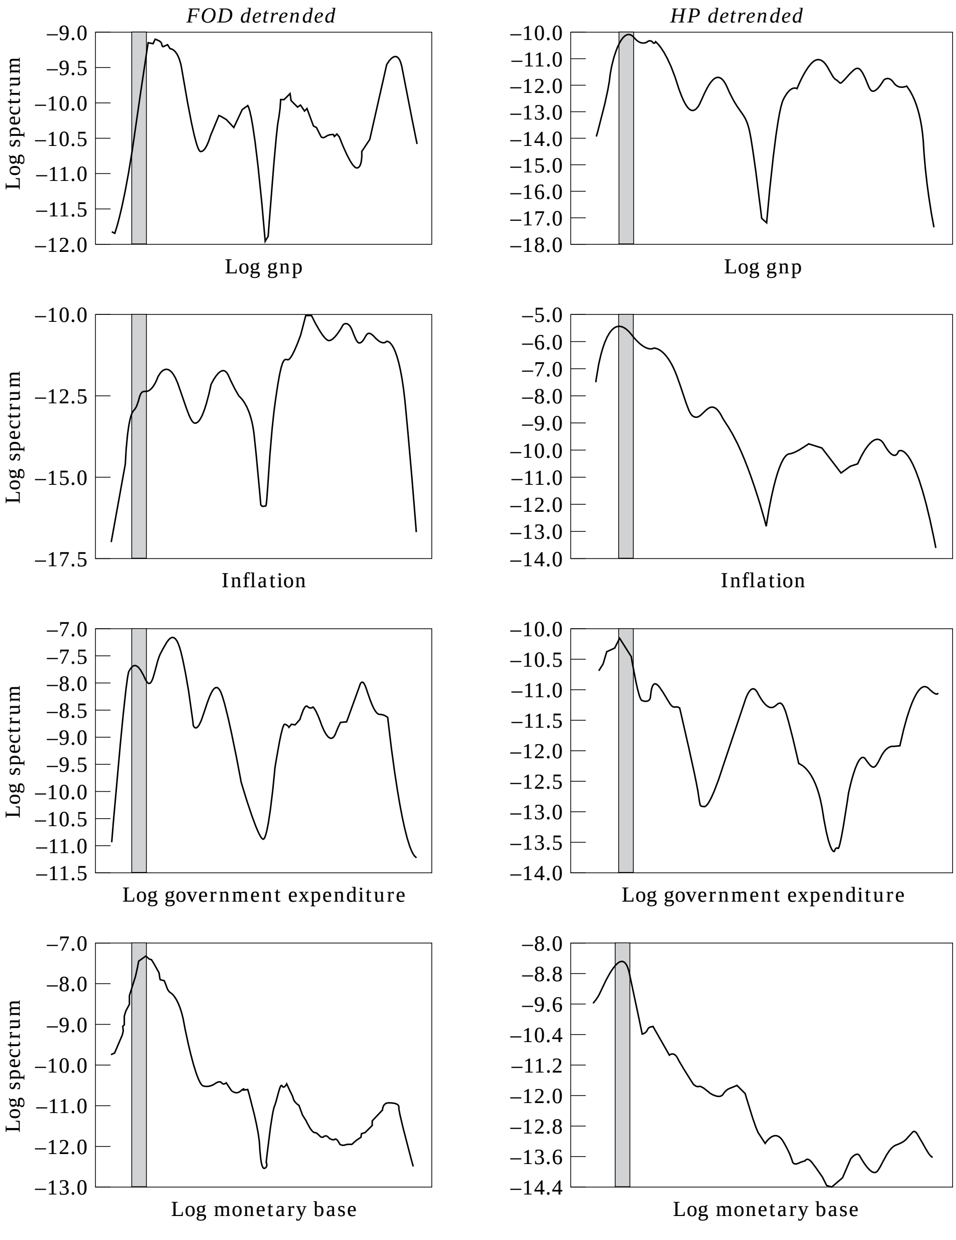

FIGURE A. Time series plots.

that all three tests are simple to implement, have good small sample size and power properties, and can be used to assess the significance of seasonal and various type of business cycle fluctuations.

Although the tests are designed for univariate time series, they can be extended to multivariate frameworks where propositions concerning the seasonal and cyclical behaviour of a multitude of time series can be formulated and tested. These extensions are pretty simple to handle. For example, if $h(\lambda)$ is the spectral density matrix of a bivariate process, then equation (1) applied to this vector process means that X_t displays significant cycles of period $r<\infty$ if each of the two components has a peak or large mass in the spectral density at some λ_p and if the coherence between the two components at frequency λ_p is high. The three tests can then be extended in a straightforward way to a vector of conditions concerning the spectral density of each of the components of X_t and of the coherence of pairs of components of X_t. Finally, note that once equation (1) of cycles is extended to multivariate frameworks, it approximately coincides with the one used by business cycle macroeconomists.

References

Alesina, A. & Roubini, N. (1992). Politics cycles in OECD countries. *Review of Economic Studies*, **59**, 663–688.

Anderson, T.W. (1971). *Statistical Analysis of Time Series*. New York, NY: Wiley & Sons.

Andrews, D. (1991). Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica*, **59**, 817–858.

Barsky, R. & Miron, J. (1989). The seasonal and the business cycle. *Journal of Political Economy*, **97**, 503–534.

Bloomfield, P. (1976). *Fourier Analysis of Time Series: An Introduction*. New York, NY: Wiley & Sons.

Campbell, J. & Kyle, A. (1993). Smart money, noise trading and stock price behaviour. *Review of Economics Studies*, **60**, 1–34.

Canova, F. (1993). *Detrending and business cycle facts*. CEPR, Working Paper 782, London, U.K.

Cochrane, J. (1988). How big is the random walk in output. *Journal of Political Economy*, **96**, 893–920.

Cogley, T. & Nason, J. (1995). Effects of the Hodrick and Prescott filter on integrated time series. *Journal of Economic Dynamics and Control*, **19**, 253–278.

Durlauf, S. (1991). Spectral based testing of the martingale hypothesis. *Journal of Econometrics*, **50**, 355–376.

Engle, R. (1974). Band spectrum regression. *International Economic Review*, **15**, 1–14.

Fisher, R. (1929). Tests of significance in harmonic analysis. *Proceedings of the Royal Society, Series A*, **125**, 54–59.

Fuller, W. (1981). *An Introduction to Statistical Time Series*. New York, NY: Wiley & Sons.

Granger, C. (1966). The typical spectral shape of an economic variable. *Econometrica*, **34**, 150–161.

Granger, C. (1979). Seasonality: causation, interpretation and implications. In A. Zellner, Ed., *Seasonal Analysis of Economic Time Series*. Washington, DC: U.S. Department of Commerce.

Granger, C. & Newbold, P. (1986). *Forecasting Economic Time Series*. New York, NY: Academic Press.

Grenander, U. & Rosenblatt, M. (1957). *Statistical Analysis of Stationary Time Series*. New York, NY: Wiley & Sons.
Hansen, B. (1990). Lagrange multiplier tests for parameter instability in non-linear models. Unpublished manuscript, University of Rochester, Rochester, NY, U.S.A.
Hansen, L. & Sargent, T. (1993). Seasonality and approximation error in rational expectation models , *Journal of Econometrics*, **55**, 21–56.
Hastings, N.A.J. & Peacock, J.B. (1983). *Statistical Distributions*. New York, NY: Wiley & Sons.
Hodrick, R. & Prescott, E. (1980). *PostWar US business cycles: an empirical investigation*. Discussion Paper 450, Carnegie Mellon University, Pittsburgh, PA, U.S.A.
King, R. & Rebelo, S. (1993). Low filtering and business cycles. *Journal of Economic Dynamics and Control*, **17**, 207–231.
Kydland, F. & Prescott, E. (1990). Business cycles: real facts and a monetary myth. *Federal Reserve Bank of Minneapolis, Quarterly Review*, Spring, 3–18.
Lo, A. & MacKinlay, C. (1988). Stock market prices do not follow a random walk: evidence from a simple specification test. *Review of Financial Studies*, **1**, 41–66.
Nelson, C. & Kang, M. (1981). Spurious periodicity in inappropriately detrended time series. *Econometrica*, **49**, 741–751.
Nelson, C. & Plosser, C. (1982). Trends and random walks in macroeconomic time series. *Journal of Monetary Economics*, **10**, 139–167.
Nyarko, Y. (1992). Learning in misspecified models and the possibility of cycles. *Journal of Economic Theory*, **55**, 416–427.
Priestley, M.B. (1981). *Spectral Analysis and Time Series*. New York, NY: Academic Press.
Quah, D. (1992). The relative importance of permanent and transitory components: identification and some theoretical bounds. *Econometrica*, **60**, 107–118.
Sargent, T. (1986). *Macroeconomic Theory*, 2nd edition. New York, NY: Academic Press.
Sims, C. (1974). Seasonality in regression. *Journal of the American Statistical Association*, **69**, 618– 626.
Tiao, G. & Grupe, M. (1980). Hidden periodic autoregressive-moving average models in time series data. *Biometrika*, **67**, 365–373.
Walker, P. (1965). Some asymptotic results for the periodogram of stationary time series. *Journal of the Australian Mathematical Society*, **5**, 107–128.
Whittle, P. (1952). Tests for fit in time series. *Biometrika*, **39**, 309.

Appendix

Proof of Lemma 1: For linear processes satisfying assumptions 1 and 2, $I_{n,x}(\lambda_k)=0{\cdot}5h(\lambda_k)(I_{n,e}(\lambda_k)/\sigma_e^2)+R_n(\lambda_k)$ uniformly in λ_k, where for some $\gamma>0$, $E[R_n(\lambda_k)^2]=O(1/n^{2\gamma})$ (Priestley, 1981: p. 424). Therefore, $2(I_{n,x}(\lambda_k)/h(\lambda_k))$ are, asymptotically, independently distributed random variables for each λ_k and have the same distribution of $(I_{n,e}(\lambda_k)/\sigma_e^2)$. By normality of e_t we have: $(I_{n,x}(\lambda_k)/h_x(\lambda_k))\sim\chi^2(1)$ if $k=0$, $[n/2]$ and $(I_{n,x}(\lambda_k)/h_x(\lambda_k))\sim 0{\cdot}5\chi^2(2)$ otherwise. For large enough n, A_{1n} and A_{2n} are linear combinations of independent variables each proportional to a χ^2 distribution. Excluding $k=0$ and $k=[n/2]$, the weights in the summations are

$$\frac{2\pi h_x(\lambda_k)}{M} \quad \text{and} \quad \frac{2\pi\cos(\lambda_k r)h_x(\lambda_k)}{M}$$

respectively. Because the weights are unequal over the range of the summation, A_{1n} and A_{2n} are no longer χ^2 distributed. Following Fuller (1981: p. 296), we can approximate the distribution of A_{1n} by $E(A_{1n})\chi^2(v_1)/v_1$ and that of A_{2n} by $E(A_{2n})\chi^2(v_2)/v_2$, where v_1 and v_2 are equivalent degrees of freedom given by

$$\frac{2n}{M\,\Sigma_k\, q(k/M)^2}$$

where $q(k/M)$ are the weights in each expression. Since as $n\to\infty$ both v_1 and $v_2\to\infty$,

$$J_{1n}=\frac{1}{\sqrt{v_1}}\frac{A_{1n}-EA_{1n}}{\sqrt{\mathrm{var}(A_{1n})}} \quad \text{and} \quad J_{2n}=\frac{1}{\sqrt{v_2}}\frac{A_{2n}-EA_{2n}}{\sqrt{\mathrm{var}(A_{2n})}}$$

have asymptotic normal distributions with zero mean and unit variance where

$$E(A_{1n})=4\pi/n\sum_k h_{n,x}(\lambda_k)$$

$$E(A_{2n})=4\pi/n\sum_k \cos(\lambda_k r)h_{n,x}(\lambda_k),$$

$$\mathrm{var}(A_{1n})=16\pi^2/n^2\sum h^2_{n,x}(\lambda_k) \text{ and}$$

$$\mathrm{var}(A)_{2n}=16\pi^2/n^2\sum \cos^2(\lambda_k r)h^2_{n,x}(\lambda_k)$$

(see, for example, Anderson, 1971: pp. 539, 545).

Proof of Corollary 1: The corollary follows from the fact that J_{1n} and J_{2n} are asymptotically $\mathcal{N}(0,1)$ variates.

Proof of Lemma 2: From lemma 1,

$$E[I_{n,x}(\lambda_k)]=h_x(\lambda_k)+O\left(\frac{\log(n)}{n}\right) \quad \text{and}$$

and

$$\mathrm{var}[I_{n,x}(\lambda_k)]=h_{x^2}(\lambda_k)+O\left(\frac{1}{n}\right)$$

Under H_0, $h_x(\lambda_k) \approx h_{\tilde{x}}(\lambda_k)$ at all λ_k where $h_{\tilde{x}}$ is the spectral density of $\tilde{X}_t$. Therefore, for $k \neq 0$ and $[n/2]$, $\tilde{K}_n \sim 0{\cdot}5\chi^2(2([n/2]-2))$. As $n \to \infty$, $2([n/2]-2) \to \infty$ so that asymptotically $2 8(\tilde{K}_n)$ has approximately the same distribution as $\mathcal{N}(8 4([n/2]-1)-1, 1)$ (see Hastings & Peacock, 1983: p. 50). Under H_1:

$$\lim_{n\to\infty} 2\sqrt{\tilde{K}_n} = \lim_{n\to\infty} 2\sqrt{K_n\, h_{\tilde{x}}(\lambda_k)/h_x(\lambda_k)} = 2\sqrt{\Sigma_k\, W_3(\lambda_k)\chi^2(2)}$$

where $W_3(\lambda) = 1$ for λ outside $(\lambda_k \pm \varepsilon)$. $\tilde{K}_n$ is a weighted average of χ^2s with unit weight outside the band centred around λ_k and weight given by $W_3(\lambda_k)$ inside the band. This weighted sum can be approximated by $E(\tilde{K}_n)(\chi^2(v_4)/v_4)$ where $v_4 = (2n/M\, \Sigma_k\, W_3(\lambda_k)^2)$ in the band around $(2\pi k/r)$. Since $v_4 > 2([n/2]-2)$, if cycles of mean length (r/k) exist in the data, the value of $2 8(\tilde{K}_n)$ will exceed the Z_α value determined by the asymptotically normal approximation computed under H_0.

Proof of Lemma 3: From lemma 1, when x_t is a white noise, $I_{n,x}(\lambda_k) \sim 0{\cdot}5\chi^2(2)\sigma_x^2$, for k different from 0 and $[n/2]$. Therefore, for all such λ_k, C_{1n} and C_{2n} are weighted averages of χ^2 with equal weights. Hence,

$$2\|\Gamma\|C_{1n} \sim 0{\cdot}5\chi^2(2\|\Gamma\|)\sigma_x^2$$

and

$$2\|\Omega-\Gamma\|C_{2n} \sim 0{\cdot}5\chi^2(2\|\Omega-\Gamma\|)\sigma_x^2$$

and

$$\lim_{n\to\infty} D_n \sim F(2\|\Gamma\|, 2\|\Omega-\Gamma\|)$$

since C_{1n} and C_{2n} are asymptotically independent by construction. Since $\|\Omega\|$ can be chosen to grow with the sample size, $2\|\Omega-\Gamma\| \to \infty$ as $n \to \infty$ and $G_n = 2\|\Omega-\Gamma\|D_n \xrightarrow{n\to\infty} \chi^2(2\|\Gamma\|.)$ Under the alternative,

$$\frac{2I_{n,x}(\lambda_k)}{h_x(\lambda_k)} \sim \chi^2(2)$$

for each λ_k in Γ so that the distribution of C_{1n} is no longer a χ^2. Using the procedure described in lemma 1, C_{1n} can be approximated with a χ^2 distribution. Therefore G_n is approximately asymptotically distributed as $(\chi^2(v_1)E(G)/v_3)$ where

$$v_3 = \frac{32n\|\Gamma\|}{\Sigma_k\, h_{n,x}^2(\lambda_k)}$$