Analysing Productivity Cycles in the Euro area, US and UK Using Wavelet Analysis

Patrick M. Crowley*, Douglas Maraun† and David Mayes‡

November 2005

Abstract

Using recent advances in time-varying spectral methods, this research analyses the growth cycles of productivity for three economies: that of the euro area, the US and the UK. The methodology uses a discrete wavelet transform (MODWT), a continuous wavelet transform (CWT) and also Hilbert wavelet pairs in order to analyze bivariate time series in terms of conventional frequency domain measures from spectral analysis. The findings are that there appears to be cycles at various frequencies, apart from the usual longer cycle that the economics literature focuses on, and there is apparently little significant correlation or coherence between productivity cycles between the three, signifying that country-idiosyncratic factors are much more likely to drive productivity than international factors.

Keywords: Time-varying spectral analysis, coherence, phase, productivity growth, EMU, wavelet analysis, Hilbert transform.

JEL Classification: C19, C63, C65, E32, E39, E58, F40

Prepared for the CEPR/EABCN Seminar on "Productivity and the Business Cycle: Evicence for Europe and the US", held at the Bank of Finland, Helsinki, November 28-29, 2005.

Acknowledgement: Maraun acknowledges financial support from Deutsche Forschungsgemeinschaft, Sonderforschungsbereich 555.

EXTREMELY PRELIMINARY - NOT TO BE QUOTED WITHOUT AUTHORS PERMISSION

*Corresponding author: College of Business, Texas A&M University, Corpus Christi, TX, USA. Email: pcrowley@cob.tamucc.edu

†Douglas Maraun, Nonlinear Dynamics Group Physics Institute, University of Potsdam, Germany. Email: maraun@agnld.uni-potsdam.de

‡David Mayes: Monetary Policy and Research Department, Bank of Finland, 00101, Helsinki, Finland. Email: david.mayes@bof.fi

1 Introduction

One of the (unfortunate) consequences of the real business cycle literature has been the increasing tendency for economists to think only in terms of the paradigm of the steady state buffetted by "shocks" to describe the source of the business cycle[1]. As productivity has been treated almost exclusively as a shock in these models, the predominant view in economics is that productivity is driven by an unpredictable process of technological innovation (see Schumpeter (1934)) rather than a data-generating process which has relatively reliable properties (akin to the way in which economists think about other economic variables). In the real business cycle view of the world, these supply side or productivity shocks impact the macroeconomy and cause deviations in the steady-state path of economic growth, along with other types of shocks. To quote Chatterjee (2000), "Modern theories of business cycles attribute cyclical fluctuations to cumulative shocks and disturbances that continuously buffet the economy. In other words, without shocks there are no cycles" (page 1). To be fair, real business cycle models incorporate supply side shocks by using fluctuations in the Solow residual, itself a measure of the deviations of labour productivity growth from its "model-predicted" growth, indicating an improvement or worsening in the economy's technological capability (via the capital/labour ratio). These variations are interpreted as productivity shocks and are the shocks that drive most real business cycle models.

But why should productivity be treated as a "shock" without some empirical evidence that indeed productivity contains some unpredictable elements of sizeable magnitude, elements that clearly dominate any perhaps lower frequency cyclical properties? Certainly, elementary logic would tend to suggest that cyclical behaviour in productivity measures would constitute a sizeable part of the movements in productivity, as productivity is measured as a "derived" variable in empirical analysis - that is, it (here labour productivity) is not measured directly but rather is measured by using both GDP and either the number of employees or the number of hours worked. Given that both the variables from which productivity is derived exhibit cycles as stylized facts, it should hardly be surprising to find that cycles in productivity are present at various frequencies. Yet this approach has received little attention from economists in general, and has little coverage in the productivity literature at all. In this setup, how might "shocks" be characterised? Given regularities

[1] Indeed it is shocking to note the extreme prevalance of the term "shock" in economics - the word has taken on an extremely broad definition in economics and is used to describe not only sudden changes in variables, but also anticipated and unanticipated tendencies, changes in trend, changes in expectations, policy changes, and a whole host of other phenomena.

in cycles at different frequencies, shocks might occur if the cycle is suddenly shortened or lengthened, perhaps due to underlying technological innovations or creative destruction.

The paper uses time and scale resolved wavelet techniques, both from a spectral (i.e. frequency domain) analysis point ov view (the continuous wavelet transform) as well as from a scale decomposition approach in the time domain (the discrete wavelet transform) to analyse productivity cycles and the relationship between these cycles in the Euro area, the US and the the UK. Do these cycles and accompanying shocks appear to be similar between the Euro area, the US and the UK? As the real business cycle literature (Backus, Kehoe, and Kydland (1992), Backus and Kehoe (1992) and Backus, Kehoe, and Kydland (1995)) suggests that technological innovations will occur relatively simultaneously across industrialised countries, this suggests that the correlation between any productivity cycles at longer frequencies should be positive. Obviously if country-idiosyncratic labour market factors are important then there will likely be a zero or negative correlation between any productivity cycles though. More recent research replicating the Backus et al results using a GMM methodology by Ambler, Cardia, and Zimmerman (2004) suggests that empirically, productivity correlations are greater than output correlations, which suggests significantly positive correlations between productivity cycles should be evident in the data at business cycle frequencies. Clearly this issue is as yet unresolved and it is in this regard that this paper hopefully makes a contribution to the existing literature.

The paper is divided into into five parts. In the next section, the productivity cycles are reviewed, with an emphasis on the co-movement of shocks or cycles. In the third section, the maximal overlap discrete wavelet transform is introduced and applied to the productivity data, and the fourth section follows suit with the introduction of the continuous wavelet transform and application to the same productivity data. The fifth section introduces a variation on the maximal overlap discrete wavelet transform that allows spectral type measures to be obtained, and then this is applied again to the productivity data. A final section brings together the three methodologies and makes some concluding remarks.

2 Productivity cycles

This paper deals with two issues, the first is the characterisation of the evolution of productivity in a more complex cyclical framework through the use of wavelet decomposition. The second is an examination of the similarities of this decomposition across the US, UK and euro area. In very elementary terms it makes sense to assume that productivity may be subject to quite complex patterns since it is such an endogenous variable, being affected by changes in inputs, technology, organisation and output.

It is already clear that there are considerable differences in the nature of cyclical variation among the euro area countries, the UK, the US and other OECD countries (Crowley and Lee (2005) and Crowley (2005)). Although these cycles have been becoming more synchronous at all frequencies, there are clearly factors which lead them to differ both at low frequencies (trends) and at frequencies higher than the typical business cycle. Increasing synchronicity does not mean that the cycles are necessarily in phase, particularly at higher frequencies. Pure time series analysis makes it fairly difficult to decide the extent to which the shape of these cycles is inherently different across countries, from whether it is the pattern of shocks, or unanticipated external factors, that lead to the differences.

[To be completed]

3 Maximal Overlap Discrete Wavelet Transforms (MODWT)

3.1 Methodology

Wavelets are a useful tool for analysing time series, and probably represent the biggest breakthrough in time-frequency methods in several decades. They are used in many disciplines which rely on time series for validation of theories or hypotheses[2], and yet they are still largely still ignored by economists[3]. This is indeed puzzling, as wavelet analysis (which is a subset of methods which are commonly termed time-frequency analysis) has the ability to reveal cycles in data at different frequencies through time - which is something that can be very revealing for empirical work in economics. Although to date, most of the applications to economic and financial data have been made by scholars in the finance area, there are some applications using wavelets in economics, and these are reviewed in Crowley (2005).

What are wavelets? Wavelets are, by definition, small waves. That is, they begin at a finite point in time and die out at a later finite point in time. As such they must, whatever their shape, have a defined number of oscillations and last through a certain period of time or space. Clearly these small wavelike functions are ideally suited to locally approximating variables in time or space as they have the ability to be manipulated by

[2] Disciplines such as signal processing, physics, meteorology, astronomy, medicine, engineering and biology.

[3] There are some exceptions to this: James Ramsay (New York University, USA) and Ramazan Gencay (Simon Fraser University, Canada) are notable in this regard.

being either "stretched" or "squeezed" so as to mimic the series under investigation. Thus the cycles within a data series can be extracted creating a new series which just incorporates the cycles in the original data at that particular frequency. As wavelet methodology is relatively foreign to most economists, a brief review of discrete wavelet transforms is now presented.

The main feature of wavelet analysis is that it enables the researcher to separate out a variable or signal into its constituent multiresolution components. In order to retain tractability (- many wavelets have an extremely complicated functional form), assume we are dealing with symmlets, then the father and mother pair can be given respectively by the pair of functions:

$$\phi_{j,k} = 2^{-\frac{j}{2}}\phi(\frac{t-2^j k}{2^j}) \tag{1}$$

$$\psi_{j,k} = 2^{-\frac{j}{2}}\psi(\frac{t-2^j k}{2^j}) \tag{2}$$

where j indexes the scale, and k indexes the translation. It is not hard to show that any variable $x(t)$ can be built up as a sequence of projections onto father and mother wavelets indexed by both j, the scale, and k, the number of translations of the wavelet for any given scale, where if k is dyadic we obtain the basic discrete wavelet transform (DWT). As shown in Bruce and Gao (1996), if the wavelet coefficients are approximately given by the integrals:

$$s_{J,k} \approx \int x(t)\phi_{J,k}(t)dt \tag{3}$$

$$d_{j,k} \approx \int x(t)\psi_{j,k}(t)dt \tag{4}$$

$j = 1, 2, ...J$ such that J is the maximum scale sustainable with the data to hand then a multiresolution representation of $x(t)$ is given by:

$$x(t) = \sum_k s_{J,k}\phi_{J,k}(t) + \sum_k d_{J,k}\psi_{J,k}(t) + \sum_k d_{J-1,k}\psi_{J-1,k}(t) + ... + \sum_k d_{1,k}\psi_{1,k}(t) \tag{5}$$

where the basis functions $\phi_{J,k}(t)$ and $\psi_{J,k}(t)$ are assumed to be orthogonal, that is:

$$\begin{array}{c} \int \phi_{J,k}(t)\phi_{J,k'}(t) = \delta_{k,k'} \\ \int \psi_{J,k}(t)\phi_{J,k'}(t) = 0 \\ \int \psi_{J,k}(t)\psi_{J',k'}(t) = \delta_{k,k'}\delta_{j,j'} \end{array} \tag{6}$$

where $\delta_{i,j} = 1$ if $i = j$ and $\delta_{i,j} = 0$ if $i \neq j$. The multiresolution decomposition (MRD) of the variable or signal $x(t)$ is then summarised as:

$$\{S_J, D_J, D_{J-1}, ...D_1\} \tag{7}$$

where

For ease of exposition, the informal description above assumes a continuous signal, which in signal processing is usually the case, but in economics although variables we use for analysis represent continuous "real time signals", they are invariably sampled at pre-ordained points in time. The continuous version of the wavelet transform (known as the CWT) assumes an underlying continuous signal, whereas a discrete wavelet transform (DWT) assumes a variable or signal consisting of observations sampled at evenly-spaced points in time.

The interpretation of the MRD using the DWT is of interest in terms of understanding the frequency at which activity in the time series occurs. For example with a quarterly time series (as we have here), table 1 shows the interpretation of the different scale crystals:

Scale crystals	Quarterly frequency resolution
d1	1-2
d2	2-4
d3	4-8=1-2yrs
d4	8-16=2-4yrs
d5	16-32=4-8ys
d6	64-128=8-16yrs
d7	128-256=16-32yrs
d8	etc

Table 1: Frequency interpretation of MRD scale levels

Although extremely popular due to its intuitive approach, the classic DWT suffers from two drawbacks: dyadic length requirements and the fact that the DWT is non-shift invariant. In order to address these two drawbacks, the maximal-overlap DWT (MODWT)[4]

[4]As Percival and Walden (2000) note, the MODWT is also commonly referred to by various names in the wavelet literature. Equivalent labels for this transform are non-decimated DWT, time-invariant DWT, undecimated DWT, translation-invariant DWT and stationary DWT. The term "maximal overlap" comes from its relationship with the literature on the Allan variance (the variation of time-keeping by atomic clocks) - see Greenhall (1991).

gives up the orthogonality property of the DWT to gain other features, such as given in Percival and Mofjeld (1997) as:

- the ability to handle any sample size regardless of whether dyadic or not;
- increased resolution at coarser scales as the MODWT oversamples the data;
- translation-invariance - in other words the MODWT crystal coefficients do not change if the time series is shifted in a "circular" fashion; and
- the MODWT produces a more asymptotically efficient wavelet variance estimator than the DWT.

Both Gençay, Selçuk, and Whicher (2001) and Percival and Walden (2000) give a description of the matrix algebra involved in the MODWT, but for our purposes the MODWT can be described in intuitive terms as simply moving a wavelet function along a series, data point by data point to obtain a detail crystal, rather than moving the wavelet function along to the next datapoints not already convolved with the data (- which would constitute a DWT). Put another way, in terms of applying wavelet filters to the data, the MODWT simply skips downsampling after filtering the data, whereas the DWT downsamples. This means that the size of each crystal is the same as the length of the data under analysis with the MODWT, whereas the number of datapoints in each crystal halves for each crystal of a higher order with a DWT.

3.2 Empirical results

3.2.1 MODWT results

The data used in this paper are described in Appendix A. First we plot the data as used in the analysis - this is shown in figure below.

The MODWT using the nearly-asymmetric Debauchies wavelet of length 8 was used to multi-scale decompose the productivity data using a total of 6 scales (- this therefore extracts fluctuations of up to 16 years in length). The results are phase corrected, and then shown graphically in stack plots below as figures 2 to 4.

The MODWT stackplots are interesting in 4 respects:

i) Crystal d1 appears in all cases to exhibit some volatility, but since the 1970s, the volatility at this frequency has tended to be low;

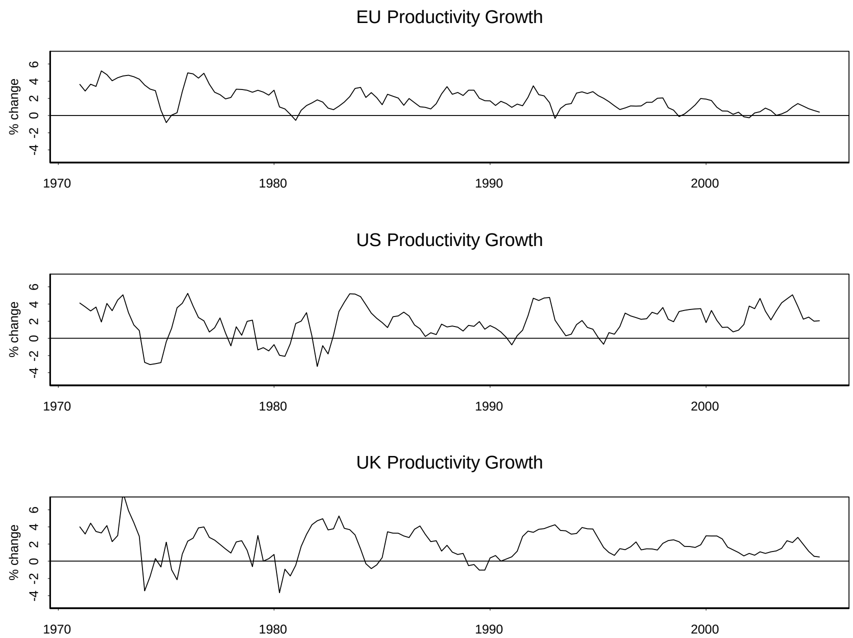

Figure 1: Percentage change in year-over-year quarterly productivity data

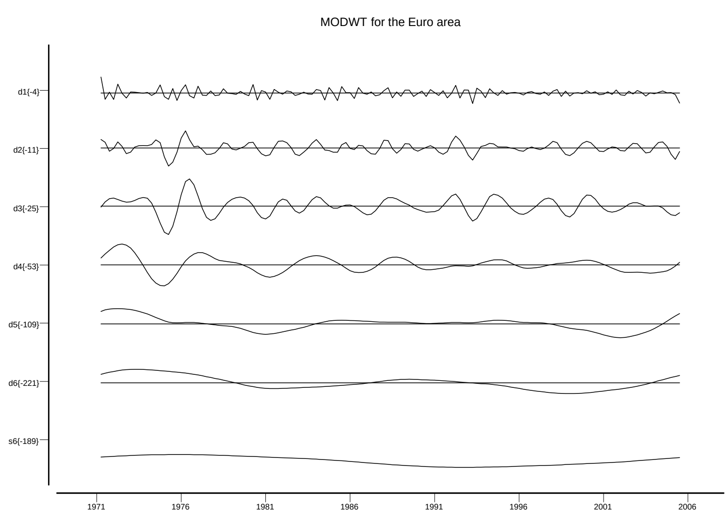

Figure 2: MODWT for Euro area productivity

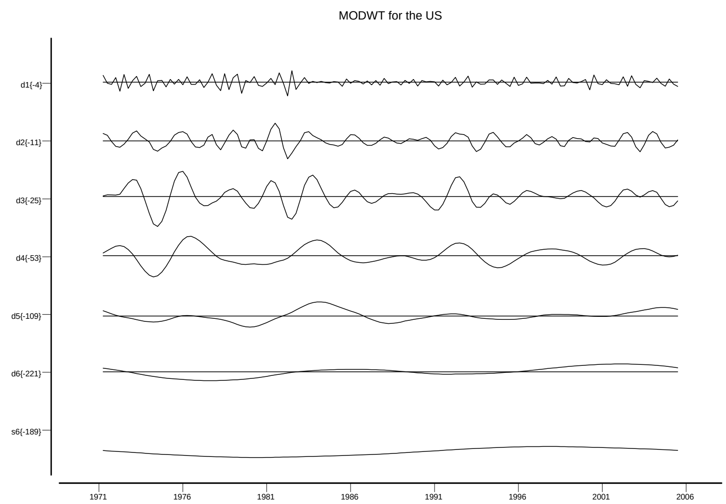

Figure 3: MODWT for US productivity

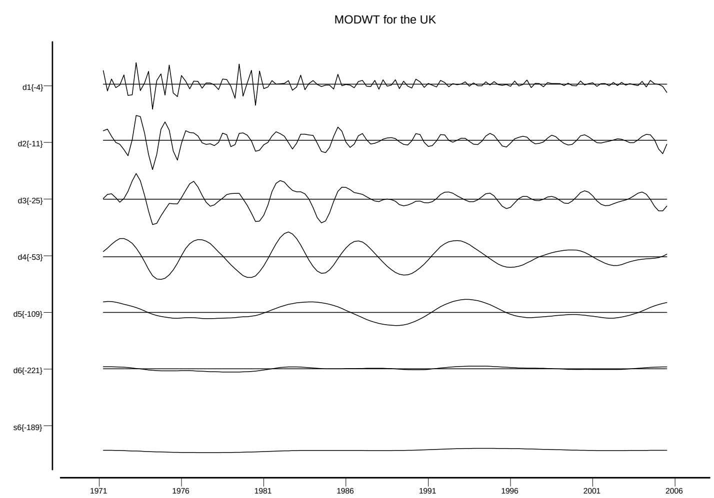

Figure 4: MODWT for UK productivity

ii) In all 3 cases, crystals d2 and d3 appear to have roughly the same pattern, suggesting that just one cycle is at work at these frequencies (of 2-8 quarters);

iii) In all 3 cases crystal d4 (2-4yrs) appears to have a fairly regular cycle, although this crystal appears to be stronger for the US and UK than it does for the Euro area. Further d5 appears to follow a similar pattern here for the Euro area and the US, suggesting a 2-8yrs cycle, which accords with business cycle frequencies;

iv) In all cases their appears to be little activity in the d6 crystal (8-16yrs) although in all cases there appears to be a residual cycle in the s6 crystal (at a frequency of approximately 20yrs, given the half cycles that we observe in the wavelet smooth). The wavelet smooth does appear to match the general perceptions about productivity trends in each of the three cases considered here.

Another way of presenting this data is to present it in terms of the energy decomposition, where the energy of a crystal for scale j, E_j, is given by:

$$E_j^d = \frac{1}{E}\sum_{k=1}^{\frac{n}{2^j}} d_{j,k}^2 \tag{8}$$

where d refers to the detail crystals and E is the total energy of the series. Orthogonal wavelets are energy (variance) preserving, so that:

$$E = E_j^s + \sum_{i=1}^{j} E_i^d \tag{9}$$

where E_j^s is the energy of the smooth. As already noted, crystals d4, d5, d6 and s6 contain most of the series energy. Note that as the detail crystals have mean zero by construction, the energy distribution of the detail crystals amounts to a various decomposition of the series by frequency band. This is shown by means of a histogram, a box plot and a pie chart for each of the series in figures 5 to 7. The figures show in all cases that the wavelet smooth contains most energy (which is hardly surprising as this crystal does not have zero mean), but of the detail crystals d3 (1-2 year cycles) has most energy for the Euro area and the US, followed by the d4 crystal (2-4 years), but for the UK, the d4 crystal energy has most energy followed by the d3 crystal.

3.2.2 Correlation analysis

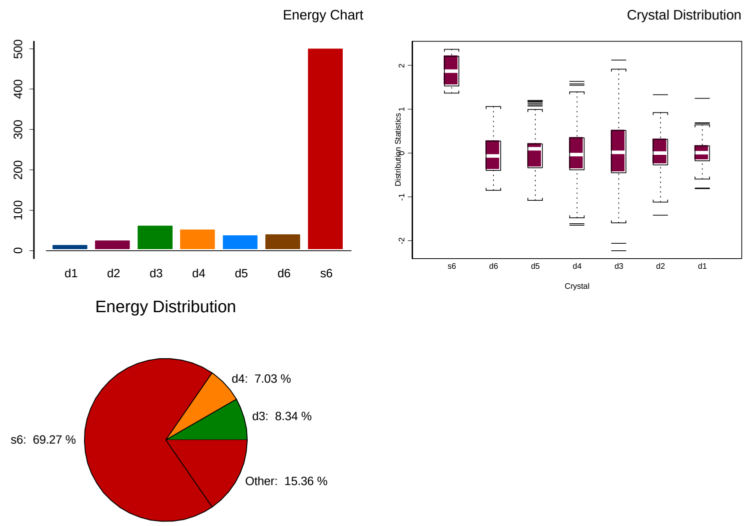

Figure 5: Euro area crystal energy distribution

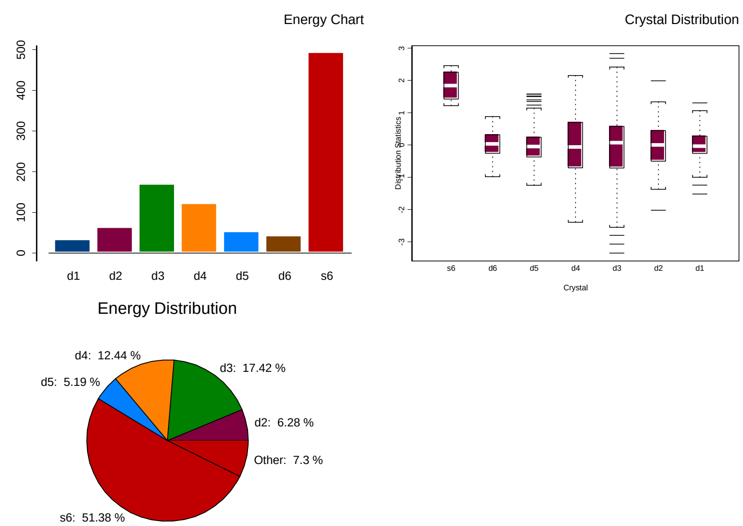

Figure 6: US crystal energy distribution

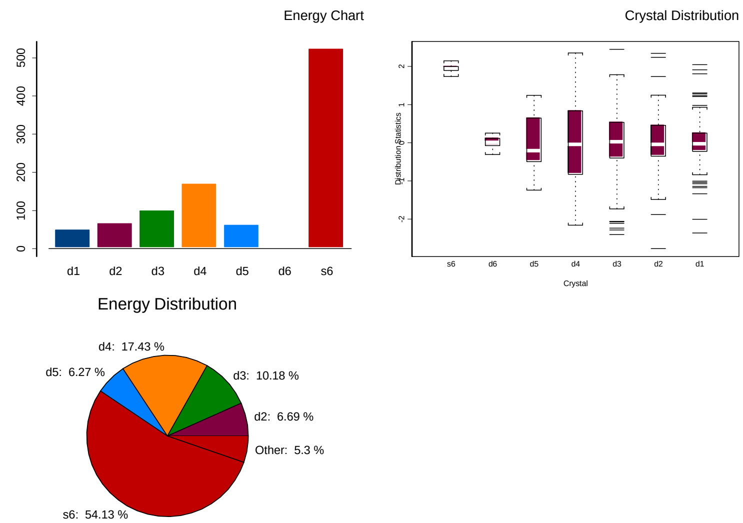

Figure 7: UK crystal energy distribution

Given that wavelet analysis can decompose a series into sets of crystals at various scales, it is not such a big leap to then take each scale crystal and use it as a basis for decomposing the variance of a given series into variances and covariances at different scales. Once covariance by scale has been obtained, the wavelet variances and covariances can be used together to obtain scale correlations between series. Confidence intervals can be derived for the correlation coefficients by scale (these are also derived in Whitcher, Guttorp, and Percival (2000)). Further details can be found in Constantine and Percival (2003) which is originally based on Whitcher, Guttorp, and Percival (2000) (with full-blown mathematical background provided in Whitcher, Guttorp, and Percival (1999)). Other more technical sources for this material are Percival and Walden (2000) and Gençay, Selçuk, and Whicher (2001). Figure 8 shows the wavelet correlations by scale for each of the series against one of the other series. Because of disposing with the boundary coefficients so as to ensure that the correlation estimates are unbiased, only 5 scales could now be resolved. The grey bars represent 95% confidence intervals. In none of the cases was the correlation coefficient between detail crystals reported as significantly positive. Appendix B gives the values for these correlations. The highest correlation was recorded for the UK against the US for the d5 crystal (4-8 year cycles) at 0.64. The highest correlation for the EU vs the US was recorded for the d4 crystal at 0.30. The wavelet smooth correlations are point estimates

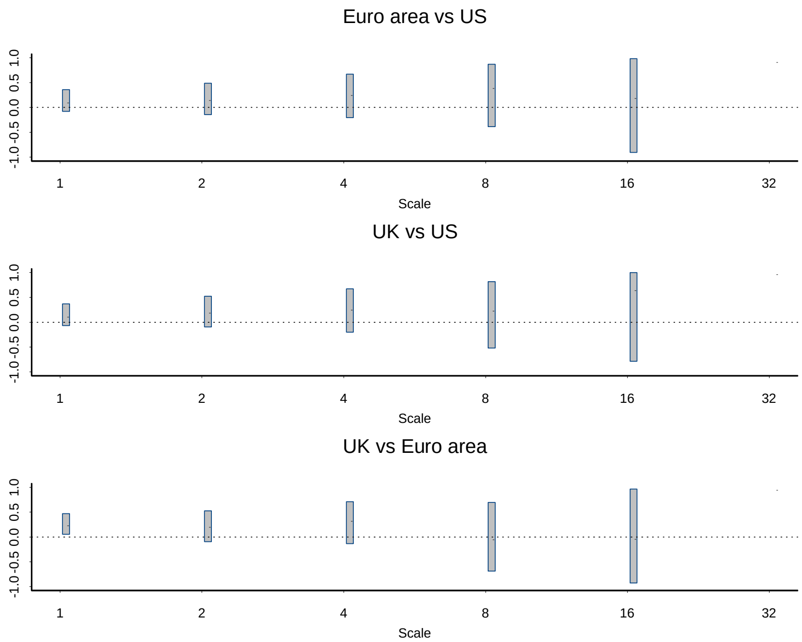

Figure 8: Wavelet correlations with 95% confidence intervals

as confidence intervals cannot be calculated for these cycles as they can theoretically be infinite. Cross-correlations can also be plotted for each of these variables to see if any lags are apparent in the data. The cross-correlation plots are shown in figures 9 to 11.

The results are interesting, because they highlight the leads and lags that occur in productivity cycles at different frequencies. In figure 9 the Euro area clearly lags behind the US for the d2, d3, d4 and d5 crystals, with maximum correlations occurring for these crystals at 0.41 (3Q), 0.58 (3Q), 0.42 (2Q) and 0.6 (6Q) respectively. In figure 10 the contemporaneous correlations are nearly all the highest correlations, although there is up to a 2 quarter lag for the largest correlations to occur in the d2, d3 and d5 crystals. The most surprising result perhaps is shown in figure 11 where the UK clearly leads the Euro area in its productivity cycles at all longer cycles and the largest correlation here is for the d4 crystal, at 0.66 (6Q lead).

Perhaps the most noteworthy result though in this analysis is a general one: productivity cycles are strongest at the 1-4 year frequency cycles, which do not generally accord with business cycle frequencies. Clearly there are longer cycles at work in the data as well, but if productivity is a driving force behind business cycle fluctuations, then they clearly act across frequencies rather than at the same frequency.

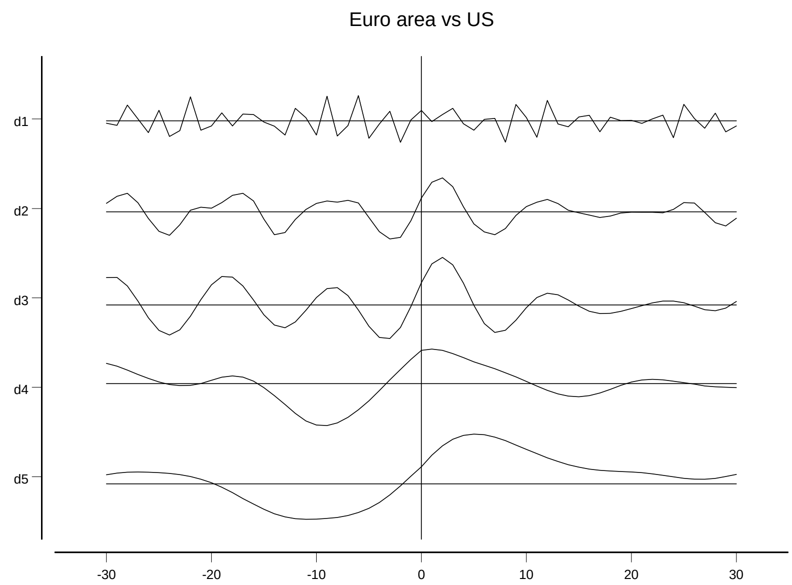

Figure 9: Cross-correlation plot for Euro area vs US productivity

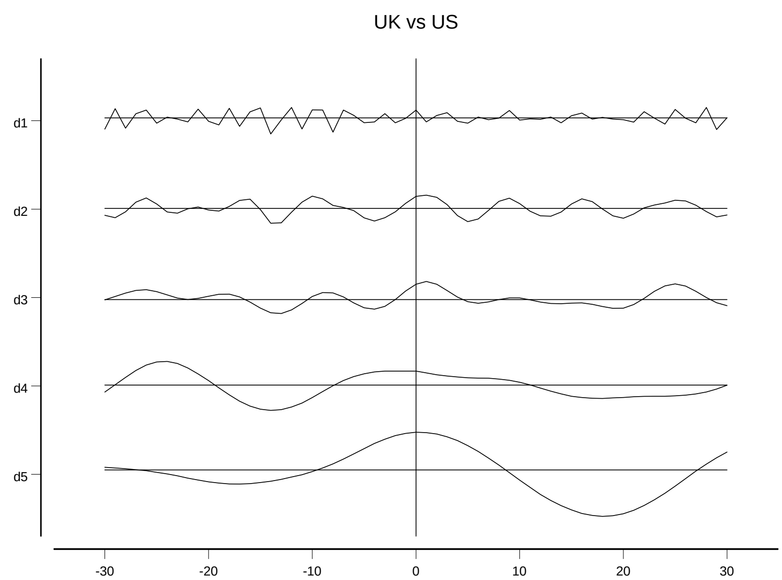

Figure 10: Cross-correlation plot for UK vs US productivity

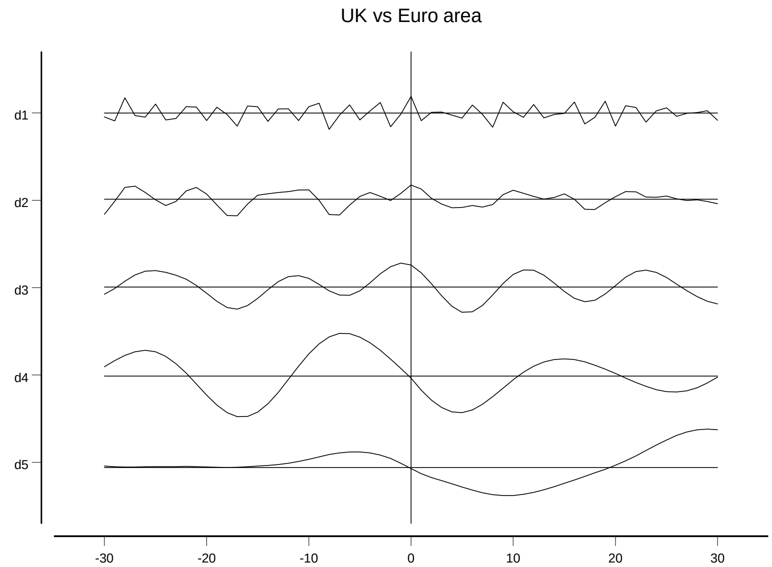

Figure 11: Cross correlation plot for UK vs Euro area productivity

4 Continuous Wavelet Transforms (CWTs)

4.1 Methodology

Wavelet analysis is neither strictly in the time domain nor the frequency domain: it straddles both - the links between these domains is explored in some detail by Priestley (1996). It is therefore quite natural that wavelet applications have been closely related to the frequency domain and can produce measures associated with spectral analysis. Perhaps the best introduction into the theoretical side of this literature can be found in Lau and Weng (1995), Holschneider (1995) and Chiann and Morettin (1998), while Torrence and Compo (1998) probably provides the most illuminating examples of empirical applications to time series from meteorology and the atmospheric sciences.

Spectral analysis is perhaps the most commonly known frequency domain tool used by economists (see Collard (1999), Camba Mendez and Kapetanios (2001), Valle e Azevedo (2002), Kim and In (2003), Süssmuth (2002) and Hughes Hallett and Richter (2004) for some examples), and therefore needs no detailed introduction here. In brief though, a representation of a covariance stationary process in terms of its frequency components can be made using Cramer's representation, as follows:

$$x_t = \mu + \int_{-\pi}^{\pi} e^{i\omega t} z(\omega) d\omega \tag{10}$$

where $i = \sqrt{-1}$, μ is the mean of the process, ω is measured in radians and $z(\omega)d\omega$ represents a complex orthogonal increment processes with variance $f_x(\omega)$, where it can be shown that:

$$f_x(\omega) = \frac{1}{2\pi} \left(\gamma(0) + 2 \sum_{\tau=1}^{\infty} \gamma(\tau) \cos(\omega\tau) \right) \tag{11}$$

where $\gamma(\tau)$ is the autocorrelation function. $f_x(\omega)$ is also known as the spectrum of a series as it defines a series of orthogonal periodic functions which essentially represent a decomposition of the variance into an infinite sum of waves of different frequencies. Given a large value of $f_x(\omega_i)$, say at a particular value of ω_i, $\widehat{\omega}_i$, this implies that frequency $\widehat{\omega}_i$ is a particularly important component of the series.

To conduct wavelet analysis in discrete terms, we would choose an orthogonal basis and then convolve the data with a wavelet filter to produce a set of coefficients (or crystals) which can be transformed back into the original series. In contrast to the DWT, we analyse a continuous set of scales (and thus choose a non-orthogonal basis accepting highly redundant results). So given a time series $x(t)$ and an analysing wavelet function $\psi(\theta)$, then the continuous wavelet transformation (CWT) is given by:

$$W(s,t) = \int_{-\infty}^{\infty} \frac{d\tau}{s^{\frac{1}{2}}} \psi^* \left(\frac{\tau - t}{s} \right) x(\tau) \tag{12}$$

For an easier computation making use of FFT algorithms this can be rewritten in Fourier space. For a discrete numerical evaluation we get:

$$W_k(s) = \sum_{k=0}^{N} s^{\frac{1}{2}} \widehat{x}_t \widehat{\psi}^* (s\omega_k) e^{i\omega_k t \partial t} \tag{13}$$

where $\widehat{x}_k$ is the discrete Fourier transform of x_t:

$$\widehat{x}_k = \frac{1}{(N+1)} \sum_{k=0}^{N} x_t \exp \left\{ \frac{-2\pi i k t}{N+1} \right\} \tag{14}$$

where $\widehat{x}_k$,represents the Fourier coefficients. In this research we use a Morlet wavelet, which is defined as:

$$\psi(\theta) = e^{i\omega\pi} e^{-\frac{\pi^2}{2}} \tag{15}$$

This is a symmetric wavelet, and is widely used in CWT analysis in the wavelet literature. Given our analysis above, it is also then possible to calculate conventional spectral measures, such as the spectral power:

$$WPS(t,s) = E\{W(t,s)W(t,s)^*\} \tag{16}$$

With two variables, x and y, it is posible to also derive and empirically estimate the cross wavelet power spectrum:

$$WCS^{xy}(t,s) = E\{W^x(t,s)W^y(t,s)^*\} \tag{17}$$

which gives rise to other multivariate spectral measures such as the coherence:

$$WCO^{xy}(t,s) = \frac{|WCS^{xy}(t,s)|}{[WPS^x(t,s)WPS^y(t,s)^*]^{\frac{1}{2}}} \tag{18}$$

which can also be measured as the magnitude squared coherence[5], which is simply $[WCO(t,s)_{xy}]^2$. As wavelet analysis essentially identifies cycles in the data, if such cycles are detected then the phasing, $\Phi(s)$, between the cycles can also be calculated from:

$$WCS^{xy}(t,s) = |WCS^{xy}(t,s)|\, e^{i\Phi(s)} \tag{19}$$

This measure is of particular interest to economists, as it shows the degree of synchronization between cycles at different frequencies. By using these measures in conjunction with wavelet analysis, we can also assess how synchronization changes through time.

4.2 Empirical results

4.2.1 Power spectra

Wavelet power spectra essentially measure the strength of cycles at various frequencies. Figure 12 shows the log power spectrum for Euro area productivity growth against an AR1 ("red noise") background spectra with a colour scale[6] and zero padding, where a basic Morlet wavelet is used for the analysis. The vertical axis measures the scale in years, and significance levels have been estimated using monte-carlo methods for 90% and 95% levels

[5] Andrew Hughes-Hallett recently referred to mean squared coherence as the R^2 of the frequency domain. We prefer to think of coherence as a measure of similarity of frequency content as the notion of an R^2 in the frequency domain can lead to some confusion in interpretation.

[6] Here we normalize the lowest value in the spectrum to one, and use a log scale for plotting contours.

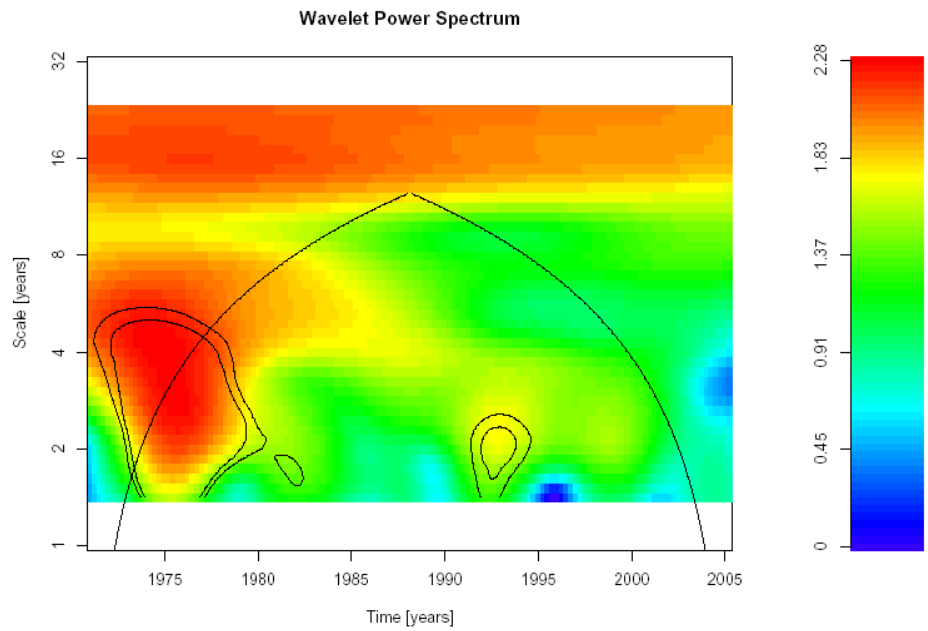

Figure 12: Power spectrum for Euro area productivity

and are plotted as contours, thus signifying movements in the series that are significantly different from an AR1 process. Only in the 1970s were there significant movements at higher frequencies, and once again at higher frequencies in the early 1990s, but there does appear to be a long cycle at work, although clearly this has been weakening throughout the period under consideration. The arch drawn in the plot shows the "cone of influence"[7], so points outside the one are to be interpreted as being less reliable than those placed within the cone. Figure 13 for the US shows significant movements also in the 1970s and early 1980s, but longer cycles in productivity have clearly been lengthening or weakening, as these do not show up clearly on this plot. Figure 14 for the UK again shows significant movements at higher frequencies through to the mid-1980s, and no apparent strength in the longer cycles, or evidence of them.

4.2.2 Cross spectral analysis

Multivariate spectral analysis essentially combines the spectra to study the frequency content of pairs of series at particular frequencies, and also the phasing of any cycles located at those frequencies. No background spectra assumption is necessary here, as cross-spectral

[7]This indicates the central area of the graph where the full length wavelets are applied to the data, so are free of any bias resulting from the use of boundary coefficients to enable wavelet application.

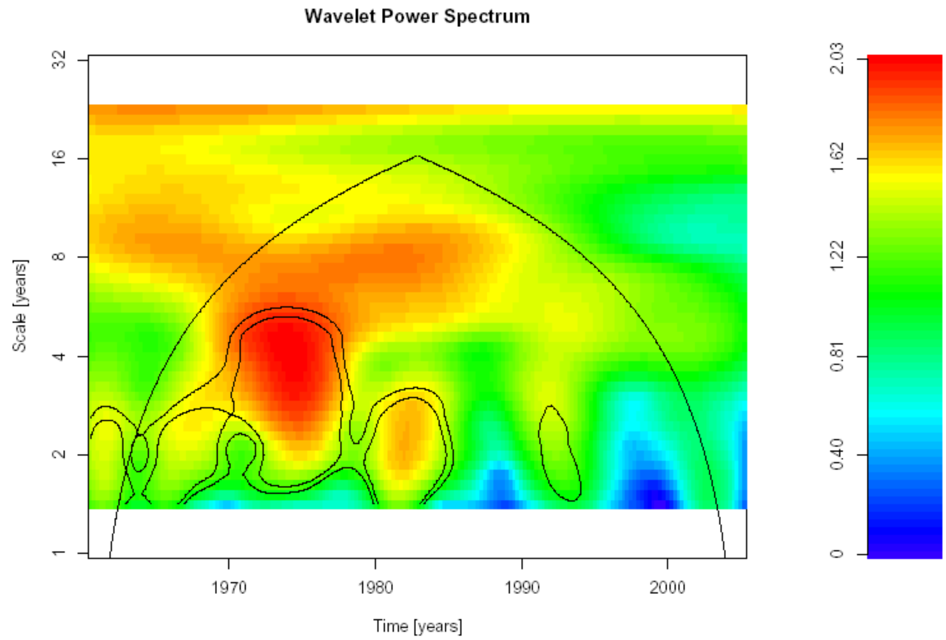

Figure 13: Power spectrum for US productivity

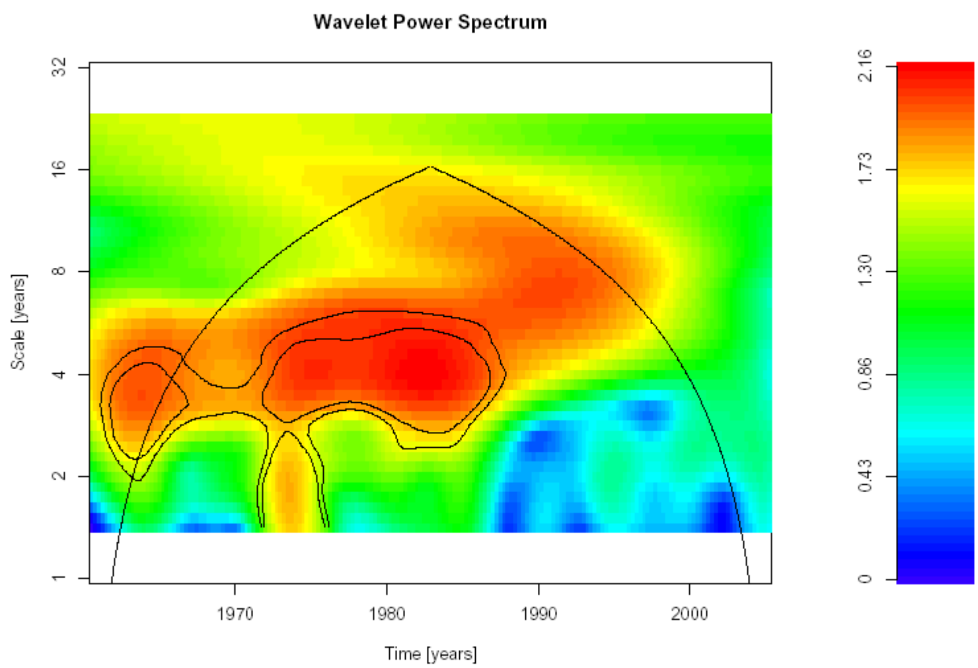

Figure 14: Power spectrum for UK productivity

analysis essentially tests for significant coherence against all linear gaussian processes. As null hypothesis for zero coherency, we choose two non-coherent white noise spectra. This is a valid background spectrum to test coherency against for all linear gaussian processes (see (see Brockwell and Davis (1998)[8]). This basically holds because coherency is a measure normalized to the process-characteristic single power spectra.

Figure 15 shows the squared coherency and phase plots for the Euro area against US productivity, with contours delineated once again for 90% and 95% significance levels[9]. Interestingly there was a high level of coherence between the Euro area and US productivities at most frequencies during the 1970s (although with a lag as can be seen in the phase diagram) which tapered to an 8 year level of coherency, but by the mid-1980s this had disappeared as well. There was significant levels of coherence at higher frequencies in the early 1990s although ahead of the US levels (likely German reunification driven) and then more recently there has been significant coherence at around a 6 year frequency (lagged) and currently at roughly a 12 year frequency (anticyclical though).

Figure 16 shows the same plots for the UK against US productivity, although with a longer dataset, as data was available from the 1960s in this instance. Two significant areas of coherence are apparent, one in the mid to late 1970s and the late 1990s and early 2000s (both of which were lagged against the US) at roughly the same frequencies of 2-6 years. Interestingly at lower frequencies coherence is quite low, perhaps indicating different long term productivity cycles at work.

Figure 17 contains the plots for UK against Euro area productivity. Only two significant areas of coherence is shown here, one in the early 1970s (with the UK leading against the Euro area) for frequencies ranging from around 3 to 8 years, and one at very high frequencies in the first half of the 1980s (and again leading the Euro area). Apart from this, no coherence at lower frequencies appears in the frequency ranges shown here[10].

---

[8] Page 438, remark 2. The Fourier Coherency between two processes x and y does not change when you apply different linear filters for each of the processes.

[9] The 90% significance level is 0.768 and for the 95% significance level it is 0.824

[10] The results for the cross spectral analysis should be interpreted with some caution, as the significant areas are all short compared to the corresponding scale and thus might be just coincidence due to multiple testing effects (see Maraun and Kurths (2004)).

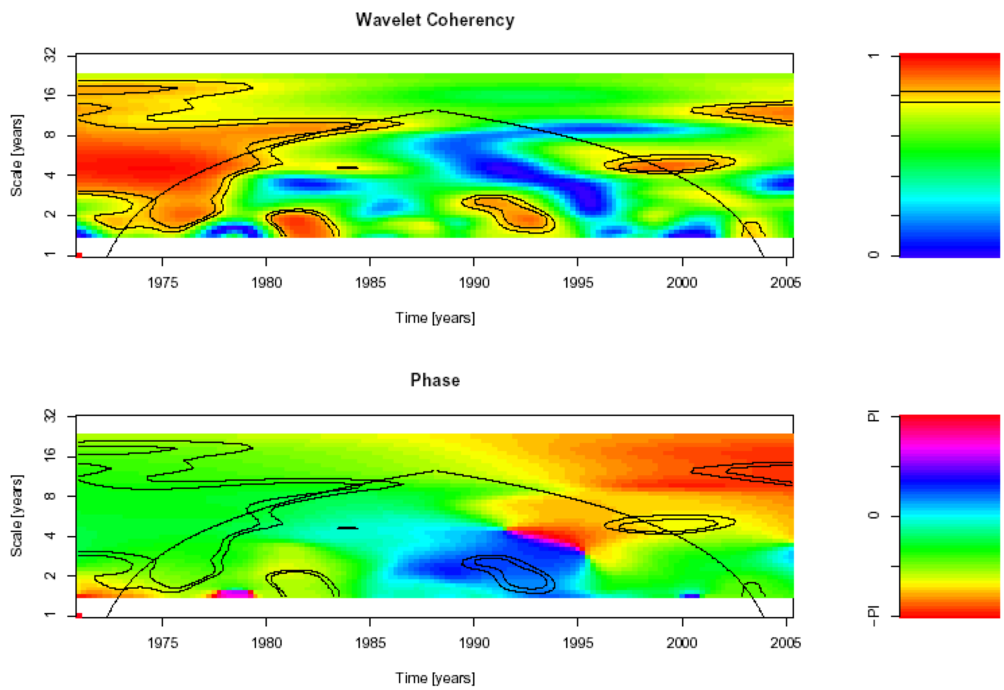

Figure 15: CWT for Euro area vs US productvity

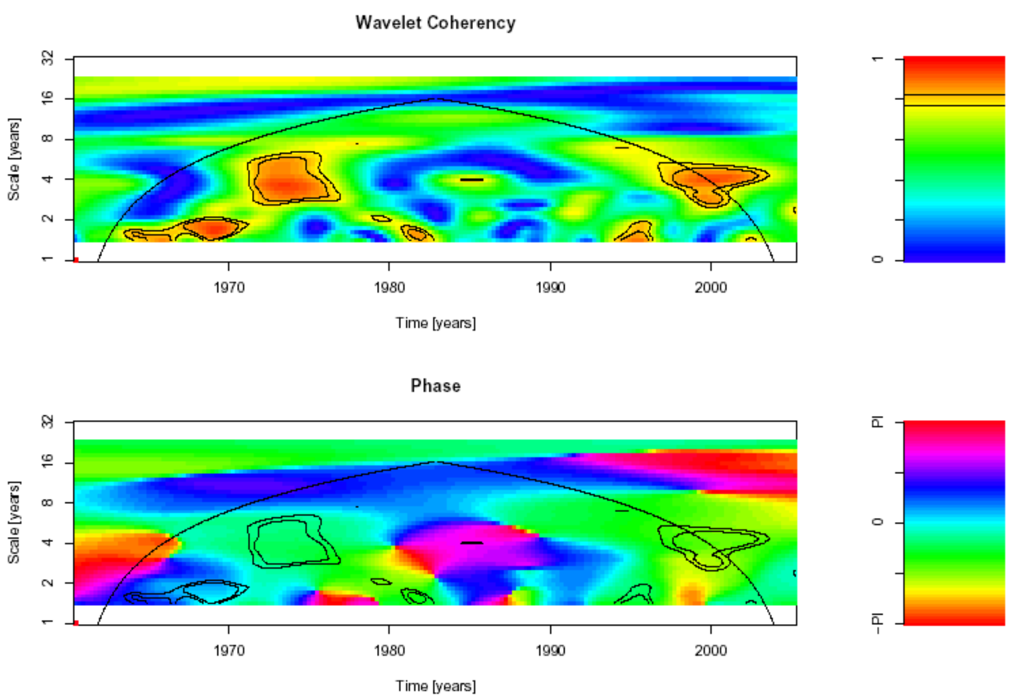

Figure 16: CWT for UK vs US productivity

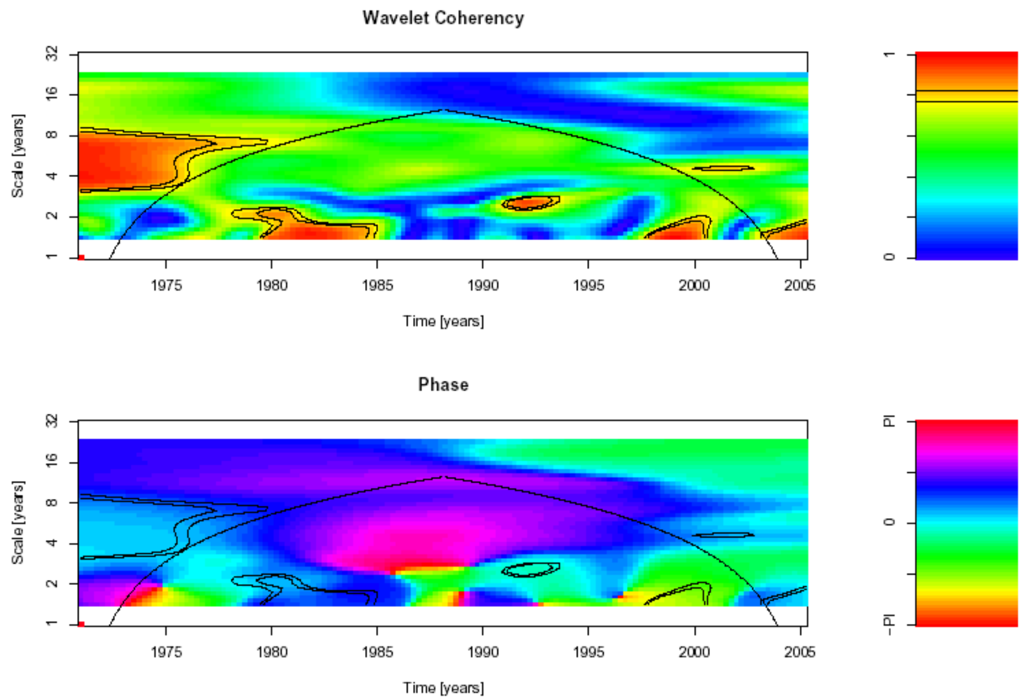

Figure 17: CWT for UK vs Euro area productivity

5 Multivariate spectral analysis using Hilbert wavelet pairs (HWP)

5.1 Methodology

The last technique employed in this paper uses a special version of the discrete wavelet transform by applying pairs of wavelets, not as continuous functions but as filters. The pairs of wavelets are identical in continuous time terms, except that one wavelet is displaced slightly compared to the other wavelet. The idea of applying 2 sets of wavelet filters to a series originated in work done by Kingsbury (2000) and is known as a dual-tree or complex wavelet transform (DTWT). The wavelets here though have specific properties - that is that each filter contains a wavelet but the wavelet filter coefficients are different as one of the wavelet filters possesses coefficients that lead the other wavelet by a certain amount. The amount of this phasing difference accords to what is known as the Hilbert transform[11]. The advantage of this methodology is that it provides a discrete wavelet transform analogue for spectral analysis, as it produces the usual spectral measures that have been defined above.

[11] A basic introduction to the Hilbert transform can be found in Bendat and Piersol (1986)

In a sense this is combination of the DTWT and the CWT spectral methodology, but allows multivariate coherency and phase measures to be defined in the time-frequency domain. In Craigmile and Whitcher (2004), the basis for using the Hilbert wavelet pairs (HWPs) is defined analytically and asymptotic theory used so that statistical inference can be applied.

In technical terms, the maximal overlap discrete Hilbert wavelet transform (MODHWT) is implemented using a pair of mother and father wavelet filters such that the two sets only differ in their phase, and not in their gain functions. A good example of a Hilbert pair of functions would the sine and cosine functions, that differ by only a quarter phase (or what is known as a half sample in the signal processing literature). As wavelets are not defined in trigonometric terms, Selesnick (2002) provides a way of obtaining near HWPs, which involves making the gain of two low pass (father) filters, denoted $A_0(f)$ and $B_0(f)$ relate in the following way:

$$B_0(f) = A_0(f)\exp^{-i\theta(f)} \tag{20}$$

where $A_0(f)$ and $B_0(f)$ form a wavelet pair as long as $\theta(f) = \pi f$, so that they have a half sample delay between them - the same can clearly be done for the high pass or mother wavelets too. To characterise a HWP in practice, two parameters are required, denoted K and M, where K is the number of zero wavelet moments (which directly relates to the smoothness of the wavelet) and M represents the degree of approximation to the half sample delay (- as M increases this approximation improves). Care needs to be taken applying the HWP to high frequencies, as at high frequencies the relationship between the two filters is no longer characterised by the Hilbert transform, so should not be used[12]. To implement the MODHWT or its decimated equivalent the DHWT, define the high pass and low pass filters as:

$$\widetilde{h}_l = \widetilde{a}_{1,l} + i\widetilde{b}_{1,l} \tag{21}$$

$$\widetilde{g}_l = \widetilde{a}_{0,l} + i\widetilde{b}_{0,l} \tag{22}$$

where equation 21 represents the mother wavelet filter and equation 22 the father wavelet filter, with both sometimes known as "complex wavelets" because of the form of the equation representing the wavelet filters. Convolution occurs with the data using the wavelet filters defined above to give the crystal coefficients. As the wavelet filters are simultaneously moved along the series, phasing can also be studied by looking at the differences in

[12]This is shown in the appendix in terms of a frequency function for the two filters.

crystal coefficients through time. There is also an analogous packet table available for the MODHWT as well.

In order to conduct time-varying spectral analysis, define $\{(W_t^X, W_t^Y)^T : t \in \mathbb{Z}\}$ as the MODHWT detail crystals from two series X_t and Y_t with a total of T observations in each series (- each crystal will also have T observations in this MODDHWT version of the analysis). The time-varying cross spectrum of X_t and Y_t can then be defined as:

$$S_{XY}(\lambda_j, t) = E\left[W_{j,t}^X W_{j,t}^Y\right] \tag{23}$$

and then corresponding amplitude and phase spectra can be extracted as with conventional frequency domain analysis as per the previous section. More details of this approach and an illustration using atmospheric monsoon data can be found in Craigmile and Whitcher (2004).

5.2 Empirical results

In terms of implementation of the Hilbert wavelet pairs, the same data is used as above, with an HWP(2,4) choice of wavelet pairs using a moving average window of 16 quarters. Periodic boundary conditions are applied, but unfortunately no confidence intervals are currently available for this methodology[13]. Figure 18 shows the output from the MODHWT analysis for Euro area vs US productivity growth in terms of squared coherence. The 95% significance level is reproduced as a horizontal line - levels of coherency above this level should therefore be interpreted as significant. In terms of the plot in figure the results seem not too disimilar from those of the CWT analysis, although the results for cycles longer than 8 years are clearly quite different. Nevertheless coherence at the 2-8 year frequencies seem to be higher during the 1970s, with little coherence in the 1980s followed by a return to reasonably high levels of coherence in the 1990s at the 4-8 year cycle frequency. Figure 19 which contains the phase plots, shows that productivity cycles tend to lag those of the US at all frequencies.

Figure 20 shows coherence for the UK against US productivity. Once again coherence was high in the 1970s at longer frequencies (2-16 years), but declined in the 1980s only to rebound in the 1990s. One interesting aspect of these plots is that high frequencies up to two years, have particularly low coherence levels (- which is largely in keeping with the

[13] This was implemented using the waveslim package in R language. We acknowledge the assistance of Brandon Whitcher in this regard.

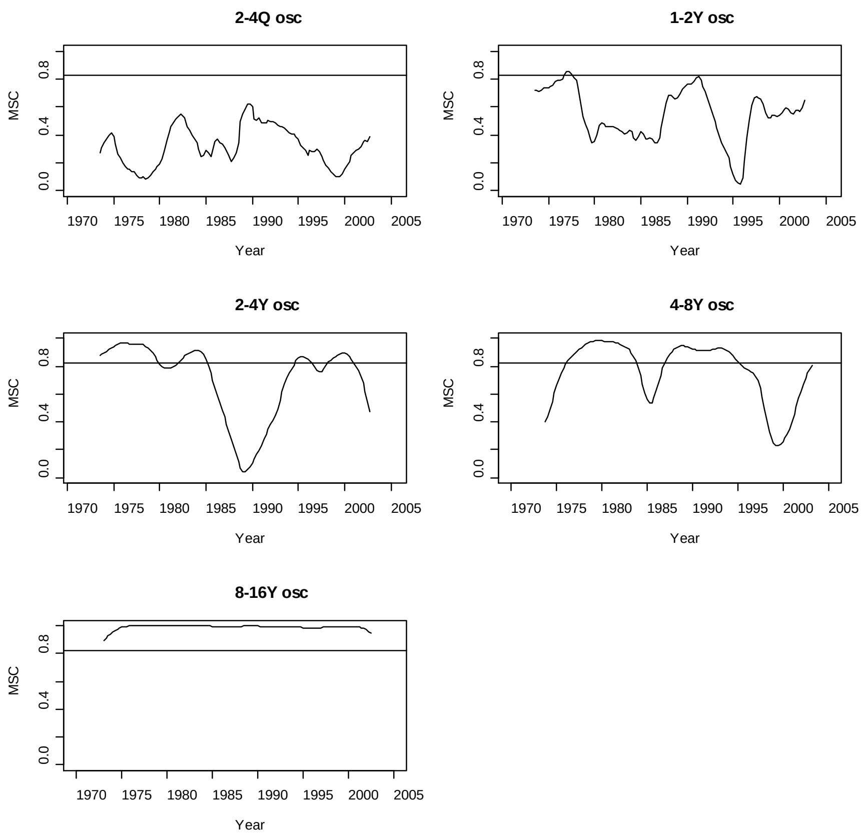

Figure 18: HWP Squared Coherency plots for EU vs US productivity

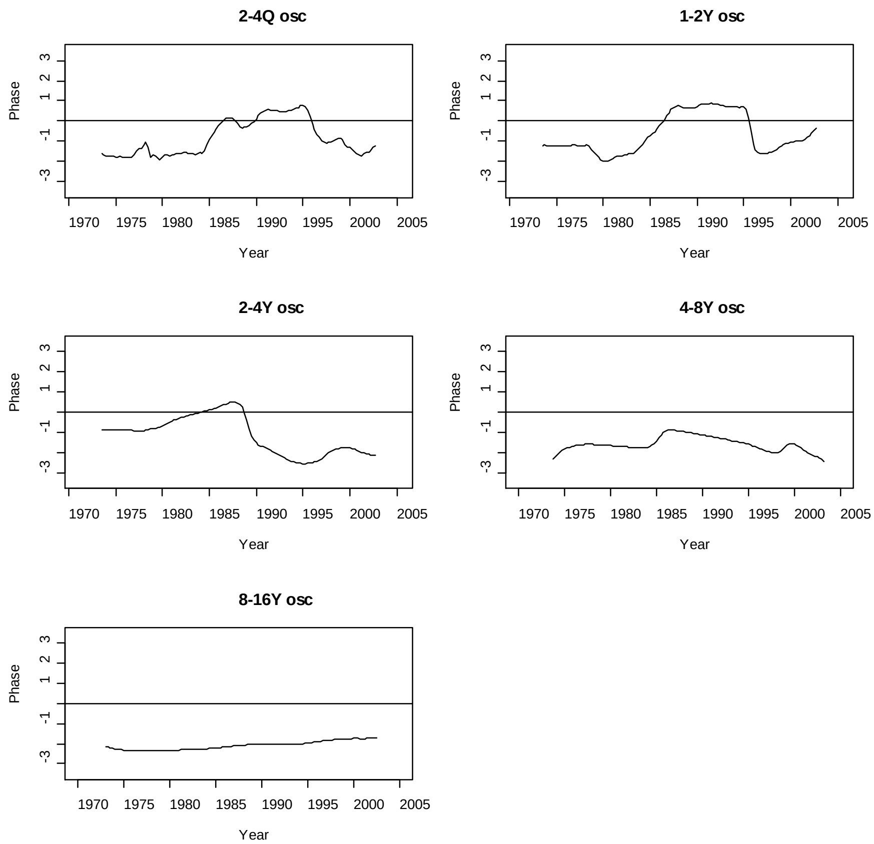

Figure 19: HWP Phase plots for Euro area vs US productivity

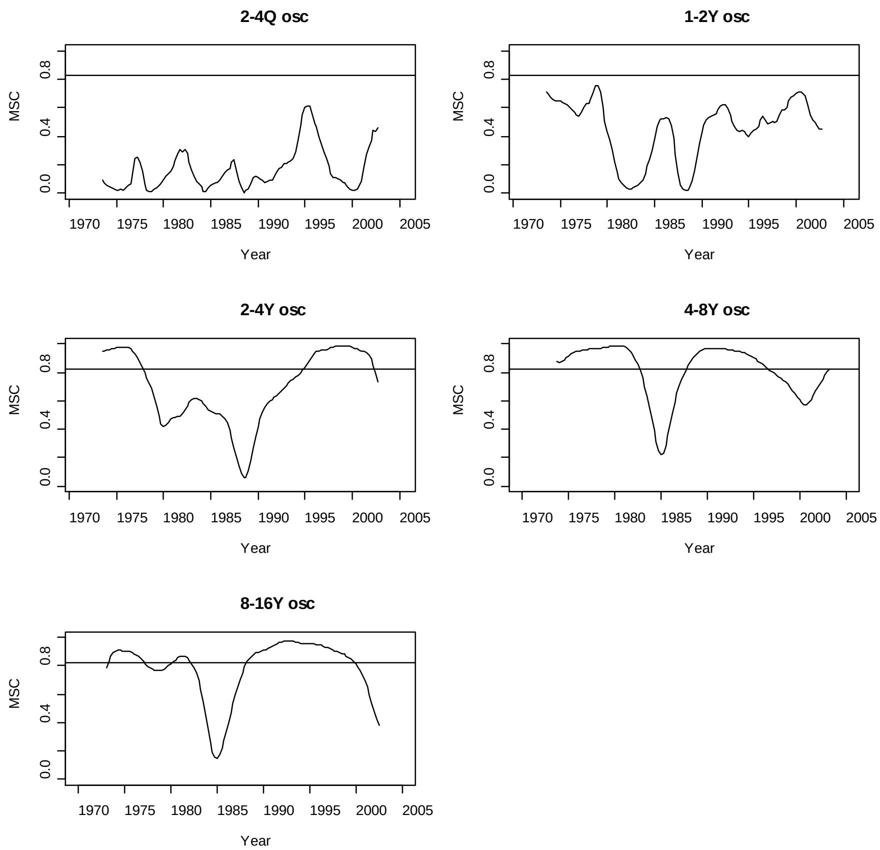

Figure 20: HWP Squared Coherency plots for UK vs US productivity

CWT analysis), and that coherence at longer cycles is less evident than with the Euro area plots. In terms of the phase plots shown in figure 21, there appears to be a phase jump in the 1980s at most frequencies, but for the most part phases indicate that cycles were ahead of those for the US until the 1980s and then the phase jump has led to productivity cycles lagging behind those of the US, although the 8-16 year cycle appears to be now roughly in line with that of the US.

Lastly, in figure 22, coherence plots for the UK against Euro area productivity are shown. The 2-4 frequency cycles appear to have fairly high coherence, but the 4-8 year cycle has extremely high coherence (confirming the CWT analysis here) up until early 1990s. All lower frequency cycles appear to have lower coherence currently. In terms of the phase plots, the 2-4 and 4-8 year frequency cycles seem to consistently ahead of those for the Euro area. This observation also confirms the results obtained in the previous section. For

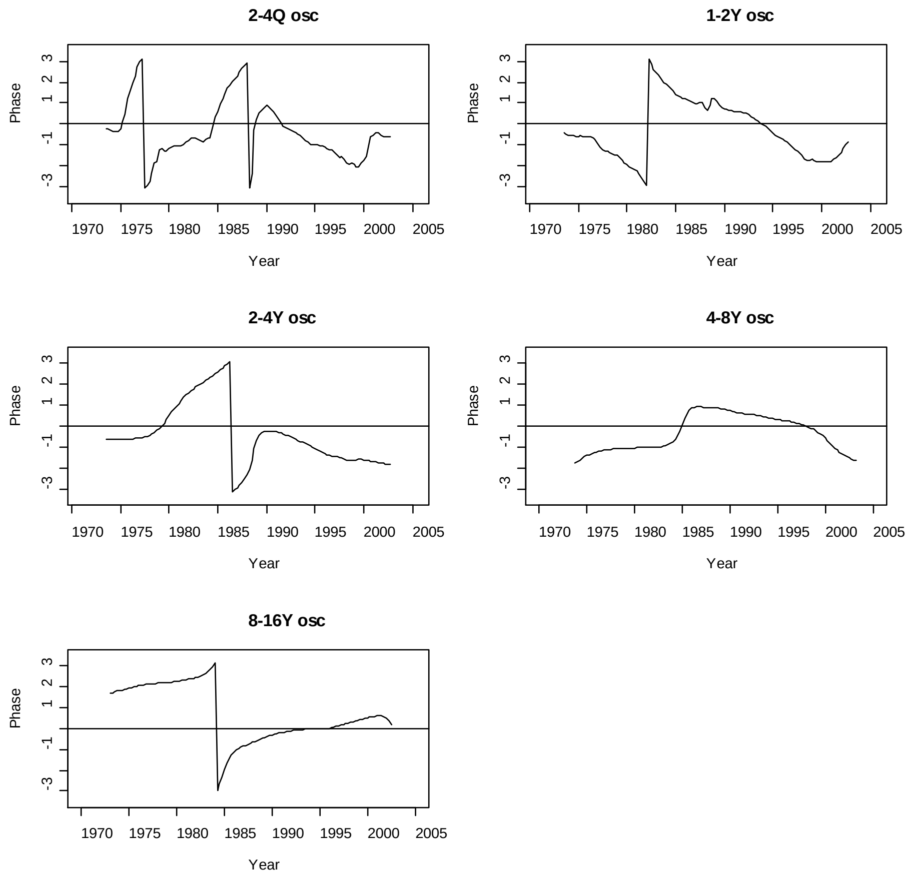

Figure 21: HWP Phase plots for UK vs US productivity

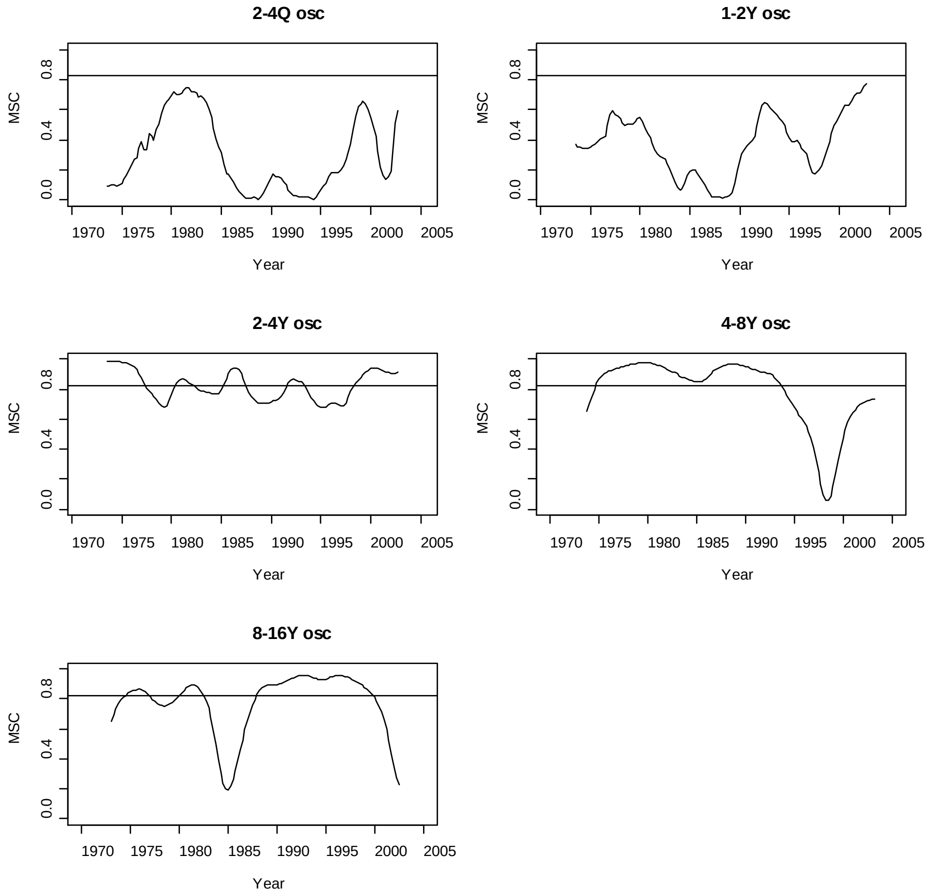

Figure 22: HWP Squared Coherency plots for UK vs Euro area productivity

the longer cycle with frequency greater than 8 years, the phasing appears to have moved from a lagged phase to a lead phase, although as coherence is currently not high, perhaps not too much should be read into this.

Several things stand out from this analysis:

i) productivity growth cycles below the 2 year frequency do not seem to be significant in terms of their coherence for any of the comparisons;

ii) the productivity cycles of the Euro area tend to lag behind those of the US at all cycles and the UK currently at longer cycles;

iii) there appears to have been a large dip in coherency during the mid-1980s, although this shows up at different frequencies depending on the comparison (for the Euro area

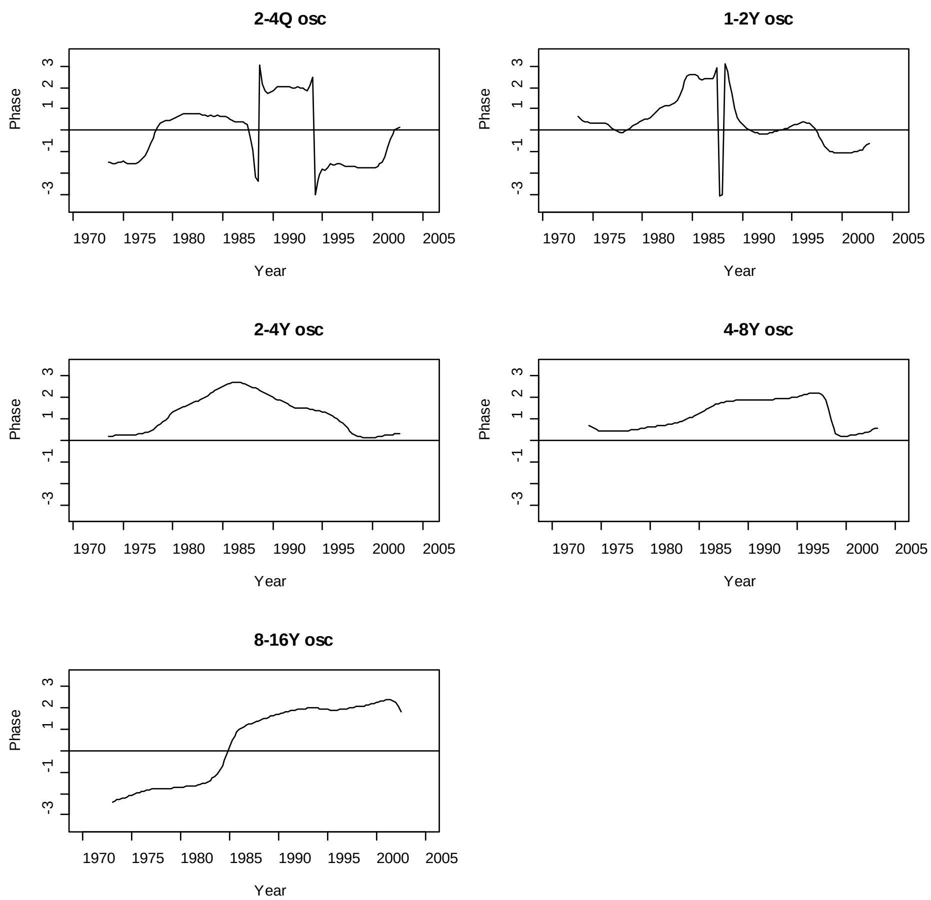

Figure 23: HWP Phase plots for UK vs Euro area productivity

vs the US it shows up in 2-8 year oscillations, for the UK vs the US in 2-16 year oscillations and for the UK vs the Euro area for 8-16 year oscillations only); and

iii) there appears to be a significantly coherent long cycle between the Euro area and the US, although detecting this cycle and its frequency is clearly difficult given the data limitations in this exercise.

6 Concluding Comments

6.1 Comparing methodologies

This paper attempted to extract and compare growth cycles in productivity in three major economic areas, those of the Euro area, the US and the UK. To do this, three different variations of wavelet analysis were applied to the same data. The results clearly showed some differences between the methodologies, which is hardly surprising as these three methodologies adopt different ways of comparing cycles. The MODWT method does give wavelet correlations, and while these correlations are not time-varying in their simplest form, other methodologies can be applied (as shown by Crowley and Lee (2005)), and indeed the MODWT does have the advantage of allowing the researcher to extract cycles and use this data for further analysis. The CWT methodology did largely corroborate the findings of the MODWT on levels of energy and the frequencies at which these occurred - this was clearly at shorter cycles than the business cycle. The CWT and HWP methodologies both produce spectral-type measures, but in one sense these measures are of limited usefulness for economic analysis, as they compare frequency content and not movement in itself. Nevertheless with significant levels of coherence, it is likely that there are similar driving processes at work. When coherence is significant, then also the phasing is relevant, as spectral techniques can properly identify the phasing between cycles - caution should be used in interpreting the validity of the phasing when coherency is not high.

Some precautionary comments should also be made when interpreting the output from wavelet analysis. First, boundaries are always a problem. In the MODWT and when using the HWP we use a periodic assumption to estimate boundary crystal coefficients, whereas with the CWT a reflective assumption is common. Hence areas outside the "cone of influence" in the CWT, or near to the end of the crystal series should not be viewed with the same level of certainty as points within the cone of influence. Likewise for the MODWT or the HWP methodology, points at the very end or beginning of the plots are subject to the assumptions made to continue the series in either direction so as to extract crystal coefficients. Second, there is much less certainty about very long cycles, given the amount of data we have (34 years) and the maximum length of cycle that we seek to resolve (16 years) - this was shown by the contradiction in results between the CWT and HWP methodologies for the Euro area vs US at these longer cycle frequencies. Only by adding additional data or using a long annual dataset can this be resolved, as clearly there are resolution issues here.

6.2 Economic results

In this paper there are five findings as described below:

a) cyclical behaviour in productivity growth is established as a stylized fact, and that while the frequency range differs by country or countries, the strongest component of this growth occurs at higher frequencies than that of the traditionally measured business cycle;

b) in the cases under investigation there is also likely a cycle which is at a much lower frequency than the business cycle also at work driving productivity growth;

c) over the period 1971-2005, while there are positive correlations for cycles near or at the business cycle frequency, none of the positive correlations between cycles for the Euro area, the US and the UK are significant (- our results do not confirm those of Ambler, Cardia, and Zimmerman (2004));

d) coherence at business cycle frequencies was high and significant in the 1970s, but that during the 1980s there was a dip in this coherence, followed by some rebound, but not consistently across the board; and

e) Euro area productivity cycles tend to lag those of the US, UK cycles lag only slightly against the US, and UK cycles tend to lead those of the Euro area. These results are not consistent across all cycle frequencies though.

The first point above is important as it shows that if productivity shocks drive business cycles then they likely do so across the frequency ranges that are pre-specified with discrete wavelet analysis. But this is nevertheless consistent with real business cycle theory, as productivity shocks would presumably occur at a higher frequency than the business cycle itself. It is in this sense though that there is both good news and bad news for real business cycle research in our results. The good news is that we find the strongest component in this data is consistently at higher frequency cycles than the traditionally measured business cycle. This likely constitutes the "shocks" that economic theorists specify in their models, and there is clearly potential for using these higher frequency crystals in future research as a representation of these productivity "shocks". The bad news for real business cycle research is that our research strongly suggests that there is also a long cycle in productivity out there, which may be an important factor in driving business cycles as well, although our results using the present dataset are somewhat inconclusive. Nevertheless, this longer cycle could hardly be described as a "shock". Clearly our future research should use long datasets to try and isolate and study this longer cycle.

Points c) and d) above are also important in terms of the real business cycle research program. These "shocks", if they are characterised by reasonably short cycles, are not highly correlated or similar in terms of frequency content between the Euro area, the US and the UK. Interestingly, the weaker cycles at the business cycle frequencies seemed to have much larger correlations and coherencies between the three economies, although we are not sure of the reason why here. Clearly our results tend to suggest that in the data at least, these "shocks" appear largely as economy- idiosyncratic rather than as common shocks. Perhaps this has to do with the way in which innovations are introduced, or instititional features of each economy that cause these "shocks" to appear as relatively idiosyncratic.

Lastly, when there is coherence between cycles, our research confirms the US productivity cycles leading those of Europe at nearly all frequencies. There is a more complicated relationship between the UK and the Euro area, but nevertheless, our findings also indicate a weaker leading relationship in productivity cycles for the UK.

References

Ambler, S., E. Cardia, and C. Zimmerman (2004). International business cycles: What are the facts? *Journal of Monetary Economics 51*, 257–276.

Backus, D. and P. Kehoe (1992). International evidence on the historical properties of business cycles. *American Economic Review 82*, 864–888.

Backus, D., P. Kehoe, and F. Kydland (1992). International real business cycles. *Journal of Political Economy 101*, 745–775.

Backus, D., P. Kehoe, and F. Kydland (1995). International business cycles: Theory and evidence. In F. Cooley (Ed.), *Frontiers of Businss Cycle Research*, pp. 331–356. Princeton, NJ, USA: Princeton University Press.

Bendat, J. and A. Piersol (1986). *Random Data: Analysis and Measurement Procedures*. New York, USA: Wiley.

Brockwell, P. and R. Davis (1998). *Time Series: Theory and Methods*. New York, NY, USA: Springer Verlag.

Bruce, A. and H.-Y. Gao (1996). *Applied Wavelet Analysis with S-PLUS*. New York, NY, USA: Springer-Verlag.

Camba Mendez, G. and G. Kapetanios (2001). Spectral based methods to identify common trends and common cycles. Working Paper 62, ECB, Frankfurt, Germany.

Chatterjee, S. (2000, March-April). From cycles to shocks: Progress in business-cycle theory. *Business Review of the Federal Reserve Bank of Philadelphia*, 1–11.

Chiann, C. and P. Morettin (1998). A wavelet analysis for time series. *Nonparametric Statistics 10*, 1–46.

Collard, F. (1999). Spectral and persistence properties of cyclical growth. *Journal of Economic Dynamics and Control 23*, 463–488.

Constantine, W. and D. Percival (2003, October). S+Wavelets 2.0.

Craigmile, P. and B. Whitcher (2004). Multivariate spectral analysis using hilbert wavelet pairs. *International Journal of Wavelets, Multiresolution and Information Processing 2*(4), 567–587.

Crowley, P. (2005, March). An intuitive guide to wavelets for economists. Discussion paper 01-05, Bank of Finland, Helsinki, Finland.

Crowley, P. and J. Lee (2005). Decomposing the co-movement of the business cycle: A time-frequency analysis of growth cycles in the euro area. Bank of Finland Discussion Paper 12/05, Helsinki, Finland.

Fagan, G., J. Henry, and R. Mestre (2001, January). An area-wide model (AWM) for the euro area. Working Paper 42, ECB, Frankfurt, Germany.

Gençay, R., F. Selçuk, and B. Whicher (2001). *An Introduction to Wavelets and Other Filtering Methods in Finance and Economics*. San Diego, CA, USA: Academic Press.

Greenhall, C. (1991). Recipes for degrees of freedom of frequency stability estimators. *IEEE Transactions on Instrumentation and Measurement 40*, 994–9.

Holschneider, M. (1995). *Wavlets: An Analysis Tool*. Oxford, UK: Oxford University Press.

Hughes Hallett, A. and C. Richter (2004). Spectral analysis as a tool for financial policy: An analysis of the short-end of the british term structure. *Computational Economics 23*, 271–288.

Kim, S. and F. In (2003). The relationship between financial variables and real economic activity: Evidence from spectral and wavelet analyses. *Studies in Nonlinear Dynamics and Econometrics 7*(4).

Kingsbury, N. (2000, September). A dual-tree complex wavelet transform with improved orthogonality and symmetry properties. In *Proceedings of the IEEE Conference on Image Processing*, Vancouver, Canada.

Lau, K.-M. and H.-Y. Weng (1995). Climate singal detection using wavelet transform: How to make a time series sing. *Bulletin of the American Meteorological Society 76*, 2391–2402.

Maraun, D. and J. Kurths (2004). Cross wavelet analysis. significance testing and pitfalls. *Nonlinear Proceedings in Geophysics 11*(4), 505–514.

Percival, D. and H. Mofjeld (1997). Analysis of subtidal coastal sea level flusctuations using wavelets. *Journal of the American Statistical Association 92*, 868–80.

Percival, D. and A. Walden (2000). *Wavelet Methods for Time Series Analysis*. Cambridge, UK: Cambridge University Press.

Priestley, M. (1996). Wavelets and time-dependent spectral analysis. *Journal of Time Series Analysis 17*, 85–103.

Schumpeter, J. (1934). *The Theory of Economic Development.* Cambridge, MA, USA: Harvard University Press.

Selesnick, I. (2002). The design of approximate hilbert transform pairs of wavelet bases. *IEEE Transactions on Signal Processing 50*(5), 1144–1152.

Süssmuth, B. (2002, January). National and supranational business cycles (1960-2000): A multivariate description of central g7 and euro15 NIPA aggregates. CESifo Working Paper 658(5).

Torrence, C. and G. Compo (1998). A practical guide to wavelet analysis. *Bulletin of the American Meteorological Society 79*(1), 61–78.

Valle e Azevedo, J. (2002, April). Business cycles: Cyclical comovement within the european union in the period 1960-1999. a frequency domain approach. WP 5-02, Banco do Portugal, Lisbon, Portugal.

Whitcher, B., P. Guttorp, and D. Percival (1999). Mathematical background for wavelet estimators of cross-covariance and cross-correlation. Technical Report TSE No. 038, NRCSE.

Whitcher, B., P. Guttorp, and D. Percival (2000). Wavelet analysis of covariance with application to atmospheric time series. *Journal of Geophysical Research 105*(D11), 14,941–14,962.

Appendices

A Data

All data are quarterly labour productivity data from 1970Q1 to 2005Q2. The data are transformed into year-over-year percentage change form (so run from 1971Q1-2005Q2), which constitutes data series of 138 observations. The data are sourced as follows:

i) Euro area data - ECB area wide model (see Fagan, Henry, and Mestre (2001))

ii) US - Dept of Labor (www.bls.gov)

iii) UK - National Statistics Office (www.statistics.gov.uk)

B Wavelet correlations

	Euro area vs US	UK vs US	UK vs Euro area
d1	0.01	-0.01	-0.02
d2	-0.11	0.09	0.09
d3	-0.02	0.15	0.37
d4	0.30	0.25	0.12
d5	0.09	0.64	0.07
s5	0.93	0.98	0.96

C Hilbert wavelet filter properties

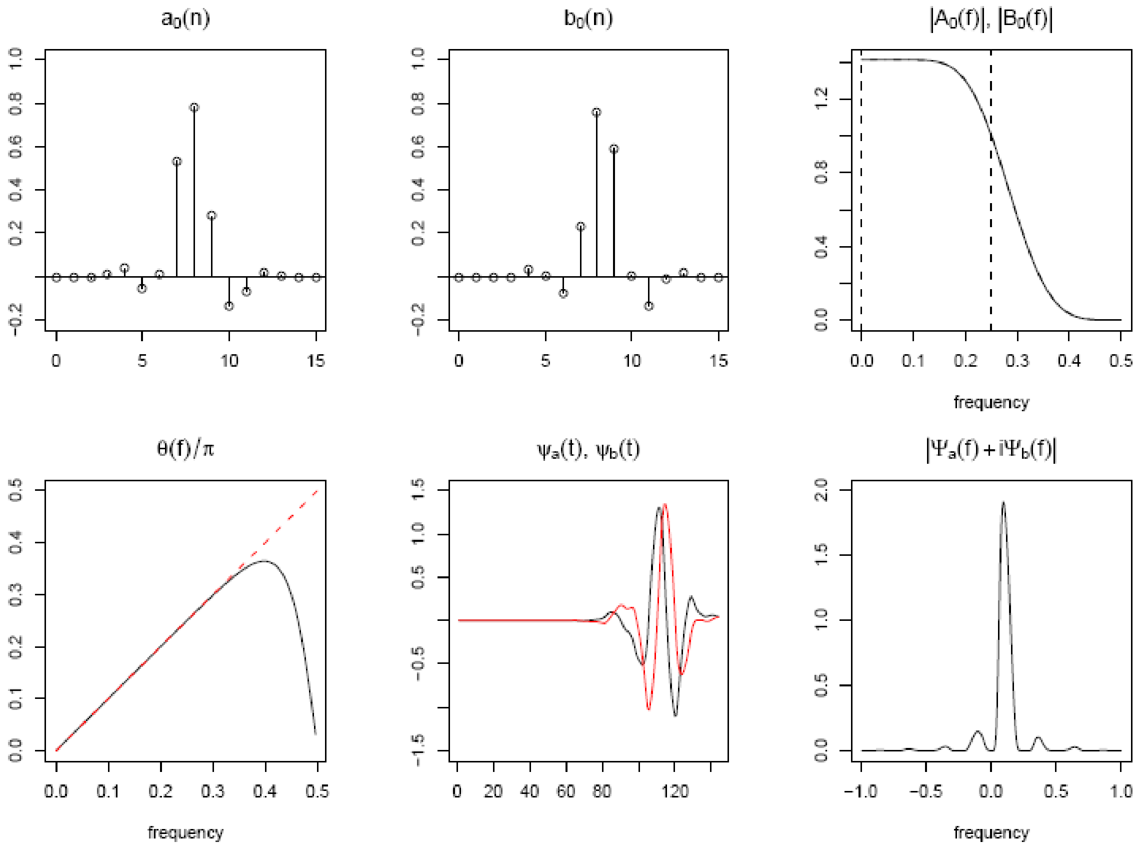

Source: Craigmile and Whitcher (2004)