

Wavelets and Economics

by Maria Joana Soares* and Luís Aguiar-Conraria**

The use of wavelet analysis is very common in a large variety of disciplines, such as signal and image processing, quantum mechanics, geophysics, medicine, biology, etc. In economics, however, wavelets are still a mysterious, but colorful, tool for time-series analysis. The pioneering work of Ramsey and Lampart [26] is unknown to the majority of economists. Among the exceptions to this rule, one can point to [4], [14], and [12]. See [6], for a recent survey of wavelet applications to economic data. Probably, wavelets are not more popular among economists, because wavelet multivariate analysis is still incipient. Recently, however, Gallegati [11] — using the maximum overlap discrete wavelet transform — and Crowley and Mayes [5] and Aguiar-Conraria and Soares [1] — using the continuous wavelet transform — showed how the cross-wavelet analysis could be fruitfully used to uncover time-frequency interactions between two economic time-series. Still, most surely, wavelets will not become very fashionable in economics until a concept analogous to the spectral partial-coherence is developed. On this regard, the proficient reader may be interested in our most recent working-paper [2].

We present a brief and self-contained introduction to the wavelet tools used, namely the continuous wavelet transform, the wavelet coherency and the wavelet phase-difference. Then we apply these tools to a real world economic problem — the study of the synchronization of

the Portuguese and Spanish economic cycles, in the last 5 decades. Decades that include the democratic transition in both countries (mid-1970s), the European Union membership of both countries (1986), and the adoption of a single currency, the Euro (1999).

TIME-FREQUENCY LOCALIZATION

In what follows, $L^2(\mathbb{R})$ denotes the set of square integrable functions, i.e. the set of functions defined on the real line and satisfying $\int_{-\infty}^{\infty} |x(t)|^2 dt < \infty$, with the usual inner product

$$\langle x, y \rangle = \int_{-\infty}^{+\infty} x(t)\overline{y(t)}dt$$

and associated norm $\|x\| = \langle x, x \rangle^{\frac{1}{2}}$. By influence of the signal process literature, this space is usually referred to as the space of finite energy signals, the energy of a signal x being simply its squared norm.

Given a function $x \in L^2(\mathbb{R})$, \hat{x} will denote its Fourier transform, here defined as:

$$\hat{x}(\omega) = \int_{-\infty}^{+\infty} x(t)e^{-i\omega t}dt. \quad (1)$$

NOTE 1.— With the above convention of the Fourier transform, ω is an *angular* (or radian) frequency. The relation to the usual Fourier frequency f is given by $f = \omega/2\pi$.

NOTE 2.— We use the symbol x to denote a general function, since this is a more common notation for time-se-

(1) As it is well known, for a function in $L^2(\mathbb{R})$, the above formula must be understood as the result of a limiting process, e.g. $\hat{x}(\omega) = \text{l.i.m.} \int_{-T}^T x(t)e^{-i\omega t}dt$, with l.i.m. denoting the limit in the mean, i.e. the limit in the L^2 sense; we will use this type of abuse of notation frequently in these notes.

* NIPE and Department of Mathematics and Applications, University of Minho.
E-mail: jsoares@math.uminho.pt

** NIPE and Economics Department, University of Minho.
E-mail: lfaguiar@eeg.uminho.pt




FIGURE 1.— A Gaussian function g (in red) and the real part of two functions $g_{\tau,\omega}$: $g_{0,4}$ (in black) and $g_{0,16}$ (in blue).

ries, which are our main objects of interest in this paper.

The spectral representation of a function given by its Fourier transform determines all the spectral components embedded in the function, but does not provide any information about when they are present. To overcome this problem, Denis Gabor, the Hungarian-born Nobel laureate in physics, proposed, in his fundamental paper on communication theory [10], the use of a modified version of the Fourier transform which became known as a windowed Fourier transform (or short time Fourier transform). The idea is simple: we first choose a window function g , i.e. a well localized function in time;^[2] by multiplying the function x by translated copies of g , we are able to select “local sections” of x , whose Fourier transforms are then computed. We thus obtain a function of two-variables, τ (the translation parameter) and ω (the angular frequency), given by

$$\mathcal{F}_{g^x}(\tau, \omega) = \int_{-\infty}^{+\infty} x(t)g(t - \tau)e^{-i\omega t} dt.$$

We can also view the above procedure in a different manner: starting with a basic window function g , a two-parameter family of functions $g_{\tau,\omega}$ is generated, via translation by τ and modulation by ω , $g_{\tau,\omega}(t) = g(t - \tau)e^{i\omega t}$, and the inner products of x with all the member of this family are then computed: $\mathcal{F}_{g^x}(\tau, \omega) = \langle x, g_{\tau,\omega} \rangle$. The principal limitation of this technique is that it gives us a fixed resolution over the entire time-frequency plane. In fact, the functions $g_{\tau,\omega}$, being obtained by simple translations in time and modulations (i.e. translations in frequency) of the window function g , all have the same “size” as g ; see Figure 1.

The main idea of the continuous wavelet transform is again to compute the inner products of the function x with members of a two-parameter family of functions $\psi_{\tau,s}$. In this case, however, the functions $\psi_{\tau,s}$ are obtained

from a given window function ψ — the so-called *mother wavelet* — which is already oscillatory (and hence, in a certain way, can be seen as a function of a given frequency), by a dilation by a scaling factor s and a translation by τ ,

$$\psi_{\tau,s}(t) = |s|^{-1/2} \psi\left(\frac{t - \tau}{s}\right);$$

see Figure 2. For $|s| > 1$, the windows $\psi_{\tau,s}$ become larger (hence, correspond to functions with lower frequency) and when the scales satisfy $|s| < 1$, the windows become narrower (hence, become functions with higher frequency). The main advantage of the continuous wavelet transform, as opposed to the windowed Fourier transform, is now clear: it provides us a time-scale (or time-frequency) representation of a function with windows whose size automatically adjusts to frequencies.

WAVELET TOOLS

The Wavelet

The minimum requirement imposed on a function $\psi \in L^2(\mathbb{R})$ to qualify for being a *mother (admissible or analyzing) wavelet* is that it satisfies the following technical condition, usually referred to as the *admissibility condition* (AC):

$$0 < \int_{-\infty}^{+\infty} \frac{|\widehat{\psi}(\omega)|}{|\omega|} d\omega < \infty; \quad (1)$$

see [7, p.22]. In this case, the constant given by the value of the above integral,

$$C_\psi = \int_{-\infty}^{+\infty} \frac{|\widehat{\psi}(\omega)|}{|\omega|} d\omega,$$

is called the *admissibility constant*. The wavelet ψ is usually normalized to have unit energy, which we always assume here. We should point out that the square integrability of ψ is a very mild decay condition and that, in practice, much more stringent conditions are imposed. In

[2] Gabor, in his paper, used Gaussian functions as windows and the transform with this particular class of functions is now called a Gabor transform.

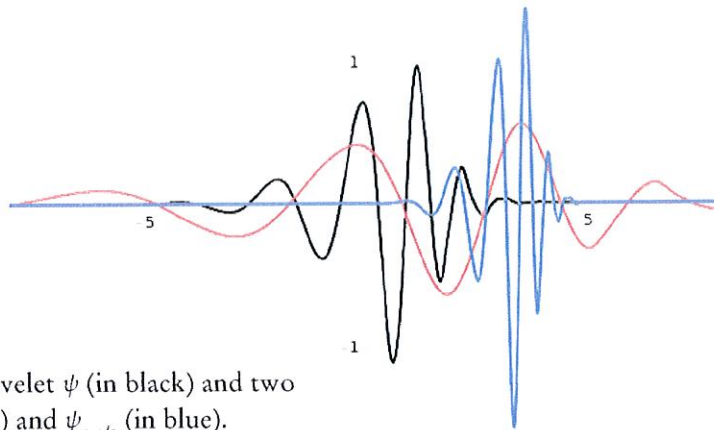


FIGURE 2.—A mother-wavelet ψ (in black) and two functions $\psi_{\tau,s}$: $\psi_{0,3}$ (in red) and $\psi_{3,1/2}$ (in blue).

fact, for the purpose of providing a useful time-frequency localization, the wavelet must be a reasonable well localized function, both in the time domain as well as in the frequency domain. For functions with sufficient decay, imposing the AC (1) is equivalent to requiring that

$$\widehat{\psi}(0) = \int_{-\infty}^{+\infty} \psi(t) dt = 0;$$

again, [7, p.24]. This implies that the function ψ has to wiggle up and down the t -axis, i.e. it must behave like a wave; this, together with the assumed decaying property, justifies the choice of the term wavelet (originally, in French, *ondelette*) to designate ψ .

The Continuous Wavelet Transform

As referred before, starting with a mother wavelet ψ , a family $\psi_{\tau,s}$ of “wavelet daughters” can be obtained by simply scaling ψ by s and translating it by τ

$$\psi_{\tau,s}(t) := \frac{1}{\sqrt{|s|}} \psi\left(\frac{t-\tau}{s}\right) \quad \tau, s \in \mathbb{R}, s \neq 0.$$

The parameter s is a scaling or dilation factor that controls the length of the wavelet (the factor $1/\sqrt{|s|}$ being introduced to guarantee preservation of the unit norm, $\|\psi_{\tau,s}\| = 1$) and τ is a location parameter that indicates where the wavelet is centered.

Given a function $x \in L^2(\mathbb{R})$, its *continuous wavelet transform* (CWT) with respect to the wavelet ψ is a function of two-variables, $W_{x;\psi}$, obtained by projecting x , in the L^2 sense, onto the over-complete family $\{\psi_{\tau,s}\}$:

$$W_{x;\psi}(\tau, s) = \langle x, \psi_{\tau,s} \rangle = \frac{1}{\sqrt{|s|}} \int_{-\infty}^{+\infty} x(t) \overline{\psi\left(\frac{t-\tau}{s}\right)} dt. \quad (2)$$

NOTE 3.—When the wavelet ψ is implicit from the context, we abbreviate the notation and simply write W_x for $W_{x;\psi}$.

Inversion of CWT

The importance of the admissibility condition (1) is due to the fact that its fulfilment guarantees that the energy of

the original function x is preserved by the wavelet transform, i.e., the following Parseval-type relation holds:

$$\int_{-\infty}^{+\infty} |x(t)|^2 dt = \frac{1}{C_\psi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} |W_x(\tau, s)|^2 \frac{d\tau ds}{s^2}.$$

In other words, the operator defined by

$$\mathcal{W}_\psi : L^2(\mathbb{R}) \longrightarrow L^2\left(\mathbb{R} \times \mathbb{R} \setminus \{0\}, \frac{d\tau ds}{s^2}\right)$$

$$x \mapsto \frac{1}{\sqrt{C_\psi}} W_{x;\psi}$$

is an isometry. Also, if ψ satisfies (1), it is possible to recover x from its wavelet transform. In fact, due to the high redundancy of this transform (observe that a function of one variable is mapped into a bivariate function), many reconstruction formulas are available. For example, when the wavelet ψ and x are real-valued, it is possible to reconstruct x by using the formula

$$x(t) = \frac{2}{C_\psi} \int_0^{+\infty} \left[\int_{-\infty}^{+\infty} W_x(\tau, s) \psi_{\tau,s}(t) d\tau \right] \frac{ds}{s^2},$$

showing that no information is lost if we restrict the computation of the transform only to positive values of the scaling parameter s , which is a usual requirement, in practice; see e.g. [7].

Wavelet Power and Wavelet Phase

When the wavelet ψ is a complex-valued function, the wavelet transform $W_{x;\psi}$ is also complex-valued and can, therefore, be expressed in polar form as

$$W_x(\tau, s) = |W_x(\tau, s)| e^{i\phi_x(\tau, s)}, \quad \phi_x \in (-\pi, \pi].$$

The square of the amplitude, $|W_x(\tau, s)|^2$ is called the *wavelet power* and the angle $\phi_x(s, \tau)$ is known as the (*wavelet*) *phase*.

For real-valued wavelet functions, the imaginary part is constantly zero and the phase is, therefore, uninformative. Hence, in order to obtain phase information about a time-series, it is necessary to make use of complex wavelets.

Analytic Wavelets

When a complex wavelet ψ is to be used, it is convenient to choose it as *analytic* (or *progressive*), where by this we mean that its Fourier transform is supported on the positive real-axis only, i.e. $\widehat{\psi}(\omega) = 0$ for $\omega \leq 0$.⁽³⁾ In fact, when ψ is analytic and x is real, reconstruction formulas involving only positive values of the scale parameter s are still available; in particular, if the wavelet satisfies $0 < |K_\psi| < \infty$ where

$$K_\psi = \int_0^\infty \frac{\overline{\widehat{\psi}(\omega)}}{\omega} d\omega,$$

then one can use the following reconstruction formula, known as the *Morlet formula*, which is particularly useful for numerical applications:

$$x(t) = 2\Re\left(\frac{1}{K_\psi} \int_0^\infty W_x(t, s) \frac{ds}{s^{3/2}}\right),$$

where $\Re(\cdot)$ denotes real part; see, e.g. [8] or [17]. For other useful features of analytic wavelets, we refer the reader to [27], [25], [20], [21] and also [24].

NOTE 4. — In what follows, we assume that all the wavelets considered are analytic and hence, that the wavelet transform is computed only for positive values of the scaling parameter s . For this reason, in all the formulas that would normally involve the quantity $|s|$, this will simply be replaced by s .

The Morlet Wavelets

The admissibility condition (1) is a very weak condition. In fact, it can be shown easily that the set of wavelet functions, $\{\psi \in L^2(\mathbb{R}) : \psi \text{ satisfies AC(1)}\}$, is dense in $L^2(\mathbb{R})$; see, e.g. [23, p. 5]. In practice, however, the choice of which wavelet to use is an important aspect to be taken into account, and this will be dictated by the kind of application one has in mind.

To study the synchronism between different time-series, it is important to select a wavelet whose corresponding transform contains information on both amplitude and phase, and hence, a complex-valued analytic wavelet is a natural choice.

The most popular analytic wavelets used in practice belong to the so-called *Morlet wavelet* family. This is a one-parameter family of functions, first introduced in [15], and given by

$$\psi_{\omega_0}(t) = \pi^{-1/4} e^{i\omega_0 t} e^{-t^2/2}.$$

Strictly speaking, the above functions are not true wavelets, since they fail to satisfy the admissibility condition.⁽⁴⁾ In fact, since the Fourier transform of the Morlet wavelet $\widehat{\psi}_{\omega_0}$ is given by $\widehat{\psi}_{\omega_0}(\omega) = \sqrt{2\pi}^{1/4} e^{-\frac{1}{2}(\omega-\omega_0)^2}$, one has $\widehat{\psi}_{\omega_0}(0) = \sqrt{2\pi}^{1/4} e^{-\omega_0^2/2} \neq 0$. However, for sufficiently large ω_0 , e.g. $\omega_0 > 5$, the values of $\widehat{\psi}_{\omega_0}(\omega)$ for $\omega \leq 0$ are so small that, for numerical purposes, ψ_{ω_0} can be considered as an analytic wavelet; see [9].

The popularity of the Morlet wavelets is due to their interesting properties. First, for numerical purposes, as we have just seen, they can be treated as analytic wavelets. Second, since the wavelet ψ_{ω_0} is the product of a complex sinusoidal of angular frequency ω_0 , $e^{i\omega_0 t}$, by a Gaussian envelope, $e^{-t^2/2}$, it makes perfect sense to associate the angular frequency ω_0 — i.e. the usual Fourier frequency $f_0 = \omega_0/(2\pi)$ — to this function; in this case, the wavelets at scale s can be associated with frequencies $f_s = \omega_0/2\pi s$; in particular, for the very common choice of $\omega_0 = 6$, we have $f_s \approx 1/s$ and hence the period (or wavelength) is $p_s \approx s$, which greatly facilitates the interpretation of the wavelet analysis as a time-frequency analysis. Finally, the wavelet ψ_{ω_0} is a function with optimal joint time-frequency concentration, in the sense that it attains the minimum possible value of uncertainty associated with the Heisenberg uncertainty principle.

All our numerical results were obtained with the Morlet wavelet ψ_{ω_0} for the particular choice $\omega_0 = 6$.

Computational Aspects

In practice, when one is dealing with a discrete time-series $x = \{x_k : k = 0, \dots, T-1\}$ of T observations with a uniform time step δt , the integral in (2) has to be discretized and is, therefore, replaced by a summation over the T time steps; also, it is convenient, for computational efficiency, to compute the transform for T values of the parameter τ , $\tau = n\delta t$; $n = 0, \dots, T-1$. Naturally, the wavelet transform is computed only for a selected set of scale values $s \in \{s_m : m = 0, \dots, F-1\}$ (corresponding to some frequencies $f_m; m = 0, \dots, F-1$). Hence, our computed wavelet spectrum of the discrete time-series x will simply be a $F \times T$ matrix $W_x = (w_{mn})$ whose (m, n) element is given by

⁽³⁾ Functions with a positive frequency spectrum were introduced in signal analysis by D. Gabor in [10]. He called them “analytic signals”, because they can be extended analytically to the upper-half complex plane.

⁽⁴⁾ In order to fulfill the admissibility condition, a correction term has to be added, as:

$$\psi_{\omega_0}(t) = \pi^{-1/4} (e^{i\omega_0 t} - e^{-\omega_0^2 t}) e^{-t^2/2}.$$

$$w_{mn} = \frac{\delta t}{\sqrt{s_m}} \sum_{k=0}^{T-1} x_k \bar{\psi} \left((k-n) \frac{\delta t}{s_m} \right).$$

Although it is possible to calculate the wavelet transform using the above formula for each value of m and n , one can also identify the computation for all the values of n simultaneously as a simple convolution of two sequences; in this case, one can follow the standard procedure and calculate this convolution as a simple product in the Fourier domain, using the Fast Fourier Transform algorithm to go forth and back from the time and spectral domains.

As with other types of transforms, the CWT applied to a finite length time-series inevitably suffers from border distortions; this is due to the fact that the values of the transform at the beginning and the end of the time-series are incorrectly computed, in the sense that they involve missing values of the series which are then artificially prescribed. Since the “size” of the wavelets $\psi_{\tau,s}$ increases with s , these edge-effects also increase with s . The region in which the transform suffers from these edge effects is called the *cone-of-influence* (COI). In this area of the time–frequency plane the results are less reliable and have to be interpreted carefully.

We also compute the so-called *wavelet ridges* which are simply the “local maxima” of the wavelet power matrix $|W_x|$. These are computed in the following manner: in each column, every element is compared with the neighbors located up to a specified distance, and the values which are larger than a given factor of the “global maximum” of $|W_x|$ (i.e. its largest value) are selected.

NOTE 5. — Although, numerically, we compute the wavelet transform in a discrete grid of the time-scale plane, the time and scale discretizations are so fine that we still refer to this as the continuous wavelet transform. The so-called *discrete wavelet transform* (DWT) often used in practice, but which we do not consider in this paper, corresponds to a very specific choice of s and τ — the dyadic grid $s = 2^j$, $\tau = 2^{jk}$; $j, k \in \mathbb{Z}$ — and makes the transform non-redundant. This, however, imposes far more stringent conditions on the choice of the mother wavelet ψ ; see, e.g. [7].

Example: time-frequency localization of the CWT

We have argued that the main advantage of wavelet analysis over spectral analysis is the possibility of tracing transitional changes across time. To illustrate this, we now consider an example, with simulated data, taken from [2]. We generate 100 years of monthly data, according to the following data generating process:

$$y_k = \cos\left(\frac{2\pi}{10}t_k\right) + \cos\left(\frac{2\pi}{p_k}t_k\right) + \varepsilon_k, \quad k = 1, \dots, 1200,$$

with $t_k = k/12$, $k = 1, \dots, 1200$, and where

$$p_k = \begin{cases} 5, & \text{if } 480 \leq k \leq 720 \\ 3, & \text{other values of } k \end{cases}$$

and ε_k is a white noise. The above time-series is the sum of two periodic components with a random error term. The first periodic component represents a 10-year cycle, while the second periodic component shows some transient dynamics. It represents a 3-year cycle that, between the fourth and sixth decades, changes to a 5-year cycle. Figure 3 displays some results related with this example.

The change in the dynamics of the series is nearly impossible to spot in Figure 3 (a), which contains a simple representation of the time-series. Furthermore, if we use the traditional spectral analysis, the information on the transient dynamics is completely lost, as we can see in Figure 3 (d). The power spectral density estimate is able to capture both the 3-year and the 10-year cycles, but it completely fails to capture the 5-year cycle that occurred in the fifth and sixth decades.

Comparing with Figure 3 (c), we observe that spectral analysis gives us essentially the same information as the global wavelet power spectrum (GWPS), which is an average, across all times, of the wavelet power spectrum. On the other hand, Figure 3 (b) shows the wavelet power spectrum itself. On the horizontal axis, we have the time dimension (in years) and the vertical axis gives us the periods.⁽⁵⁾ The intensity of power is given by the color. The color code for power ranges from blue (low power) to red (high power), with regions with warm colors thus representing areas of high power. The cone-of-influence is shown with a thick grey line. The white lines show the local maxima (or ridges) of the wavelet power spectrum, giving us a more precise estimate of the cycle period. We observe a white line on period 10 across all times, meaning that there is a permanent cycle with this period. We are also able to spot the 3-year cycle that occurs up to year 40 and, again, between years 60 and 100. Finally, we are also able to identify a yellow/orange region between the years 40 and 60, with the white stripes indicating a cycle of period five. This means that a cycle of roughly 5-year periodicity, relatively important in explaining the total variance of the time-series and taking place between years 40 and 60, was hidden by the Fourier power spectrum estimate.

⁽⁵⁾ Note that, since we are using a ψ_6 Morlet wavelet, the periods are almost identical to the scales.

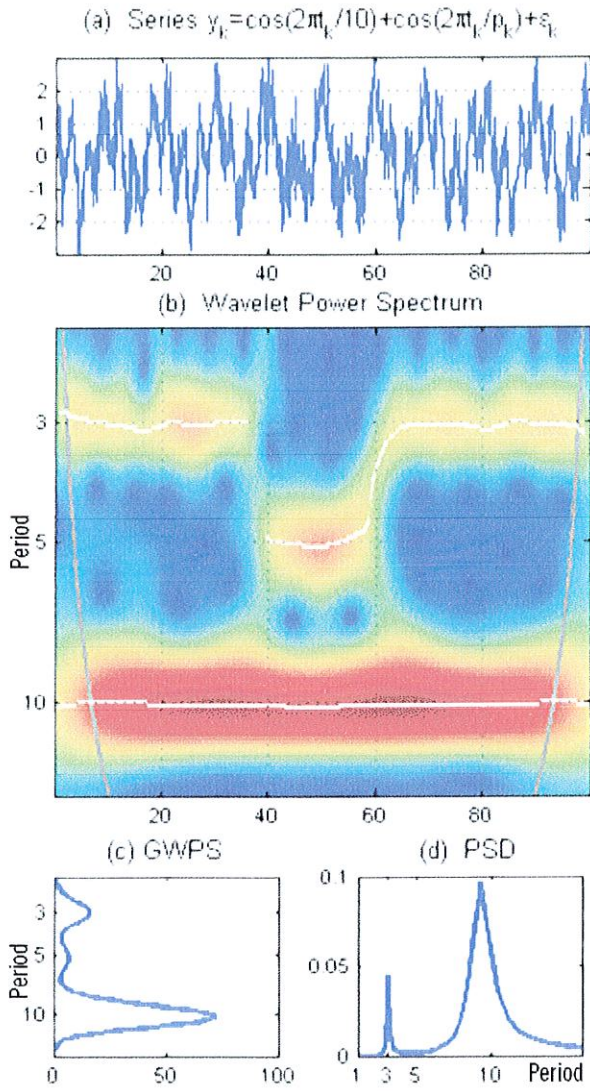


FIGURE 3.—(a) Series $y_k = \cos(2\pi t_k/10) + \cos(2\pi t_k/p_k) + \varepsilon_k$ (b) Wavelet power spectrum of y — The cone-of-influence is shown with a grey line; the color code for power ranges from blue (low power) to red (high power); the white lines show the local maxima of the wavelet power spectrum. (c) Global wavelet power spectrum, i.e. average wavelet power (over all times) for each frequency. (d) Fourier power spectral density.

Figure 3 (b) clearly illustrates the big advantage of wavelet analysis over spectral analysis. While the Fourier transform is silent about changes that happen across time, with wavelets we are able to estimate the power spectrum as a function of time and, therefore, we do not lose the time dimension. The wavelet power spectrum is able to capture not only the 3-year and 10-year cycles, but also

^[6] More correctly, we have $\phi_{xy} = \phi_x - \phi_y \pmod{2\pi}$.

to capture the change that occurred between years 40 and 60.

CROSS-WAVELET ANALYSIS

In many applications, one is interested in detecting and quantifying relationships between two non-stationary time-series. The concepts of cross-wavelet power, wavelet coherency and wavelet phase-difference are natural generalizations of the basic wavelet analysis tools that enable us to appropriately deal with the time-frequency dependencies between two time-series.

Cross-Wavelet Transform, Cross-Wavelet Power and Phase-Difference

The *cross-wavelet transform* (XWT) of two time-series, x and y , first introduced by Hudgins, Friehe and Mayer [18], is simply defined as

$$W_{xy}(\tau, s) = W_x(\tau, s) \overline{W_y(\tau, s)},$$

where W_x and W_y are the wavelet transforms of x and y , respectively. The modulus of the XWT, $|W_{xy}(\tau, s)|$ is known as the *cross-wavelet power*.

As for the wavelet transform, if the wavelet ψ is complex-valued, the cross-wavelet transform is also complex-valued and can be written as

$$W_{xy}(\tau, s) = |W_{xy}(\tau, s)| e^{i\phi_{xy}(\tau, s)},$$

where

$$\phi_{xy}(\tau, s) = \phi_x(\tau, s) - \phi_y(\tau, s),$$

with ϕ_x and ϕ_y denoting the phases of x and y respectively,^[6] is the *phase-difference* of x and y (also called the *phase-lead* of x over y). A phase-difference of zero indicates that the two time-series move together at the specified (τ, s) value; if $\phi_{xy} \in (0, \pi/2)$, then the series move in-phase, but the time-series x leads over y ; if $\phi_{xy} \in (-\pi/2, 0)$, the series also move in-phase, but, in this case, is the series y that is leading; a phase-difference of π indicates an anti-phase relation; if $\phi_{xy} \in (\pi/2, \pi)$, then the series are out-of-phase, and y is leading; finally, if $\phi_{xy} \in (-\pi, -\pi/2)$, the series are out-of-phase and x is leading.

Complex Wavelet Coherency

In analogy with the concept of coherency used in Fourier analysis, given two time-series x and y one can define their *complex wavelet coherency*, ρ_{xy} , by:

$$\rho_{xy}(\tau, s) = \frac{\mathcal{S}(W_{xy}(\tau, s))}{[\mathcal{S}(|W_x(\tau, s)|^2) \mathcal{S}(|W_y(\tau, s)|^2)]^{1/2}}, \quad (3)$$

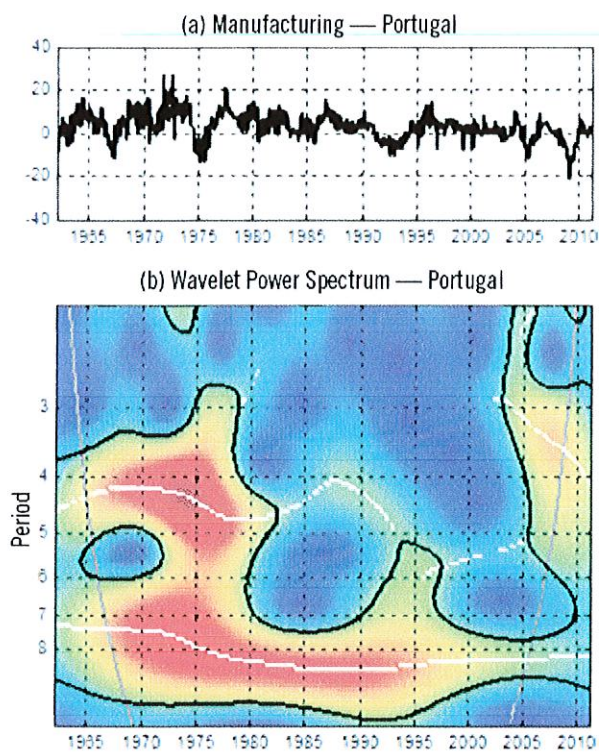


FIGURE 4.— Total Manufacturing for Portugal (a) and corresponding wavelet power spectrum (b). Color codes are as in Fig. 3. The thick black contour represents the 5% significance level.

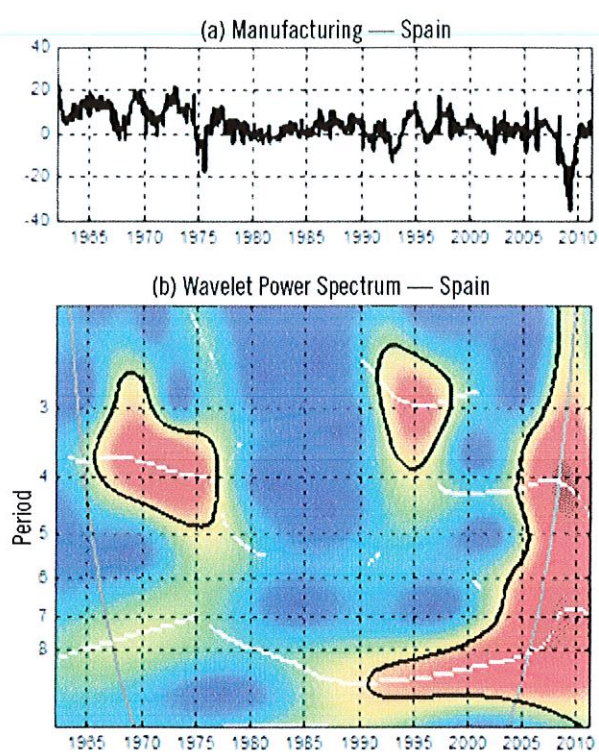


FIGURE 5.— Total Manufacturing for Spain (a) and corresponding wavelet power spectrum (b). Color codes are as in Fig. 3. The thick black contour represents the 5% significance level.

where \mathcal{S} denotes a smoothing operator in both time and scale; smoothing is necessary, because, otherwise, coherency would be identically one at all scales and times. Time and scale smoothing can be achieved, e.g. by convolution with appropriate windows; see [3] or [16], for details.

The absolute value of the complex coherency is called the *wavelet coherency* and is denoted by $R_{xy}(\tau, s)$. As in the case of the usual Fourier coherency, wavelet coherency satisfies the inequality $0 \leq R_{xy}(\tau, s) \leq 1$, whenever the ratio (3) is well defined. At points (τ, s) for which $\mathcal{S}(|W_x(\tau, s)|^2) \mathcal{S}(|W_y(\tau, s)|^2) = 0$, we will define $R_{xy}(\tau, s) = 0$.

As referred by Liu [22], the advantage of these “wavelet-based” quantities is that they may vary in time and can detect transient associations between studied time-series.

SIGNIFICANCE TESTS

As with other time-series methods, it is important to assess the statistical significance of the results obtained

by wavelet analysis. The seminal paper by Torrence and Compo [28] is one of the first works to discuss significance testing for wavelet and cross-wavelet power. However, more work needs to be done on this area. For reasonable general processes, like an ARMA process, one has to rely on bootstrap techniques or Monte Carlo Simulations. To our knowledge, there is no good way of assessing the statistical significance of the phase-difference. In fact, Ge [13] argues that one should not use significance tests for the wavelet phase-difference. Instead, its analysis should be complemented by inspection of the coherence significance.

BUSINESS CYCLE SYNCHRONIZATION BETWEEN PORTUGAL AND SPAIN

In this section we describe the application of the continuous wavelet tools — more specifically, wavelet coherency, phase, and the phase-difference — to the study of the synchronization of the economic cycles of Portugal and Spain, before and after these two countries joined




FIGURE 6.—(a) Wavelet coherency — coherency ranges from blue (low coherency) to red (high coherency); the black thick contour designates the 5% significance level. (b) Phases and phase-difference at the 3–8 year frequency band — The green line represents the phase for Spain, the blue line the phase for Portugal and the red line represents the phase-difference between Spain and Portugal.

the Euro Zone. The data used consist of the total manufacturing year-to-year growth rates for Portugal and Spain. We gathered monthly data from January of 1962 until February of 2011. The results are displayed in Figures 4–6. Figures 4 (a) and (b) represent the time-series for Portugal and the corresponding wavelet power spectrum and Figures 5 (a) and (b) the same quantities, but for Spain. The color codes are as in Figure 3. The black thick contour indicates the 5% significance level.

Wavelet coherency is shown in Figure 6 (a). Figure 6 (b) displays the average values of the phases and phase-difference at the 3–8 year frequency band: the green line represents the phase for Spain, the blue line the phase for Portugal and the red line the phase-difference between Spain and Portugal. Figure 4 (b) tells us that for the entire period of analysis Portugal has a very active business cycle around the 8-year frequency. A cycle with a similar period is also present in Spain — Figure 5 (b) — how-

ever, it is statistically significant only after 1990. Cycles associated with shorter periods are also statistically significant in both countries until 1980. It is interesting to note that at the end of the sample, probably because of the world financial and economic crisis, both countries display a large wavelet power spectrum. This is particularly true in the case of Spain, one might add. We recall that, although the precise details are different, these two countries, during the decade of 1960 and in the first half of the decade of 1970, both had proto-fascist regimes ([19]) and, in the second half of that decade, they both turned into democratic regimes.

In 1982, Portugal had a severe Current Account crisis that led to an IMF intervention in 1983. The two countries applied together to be part of the European Economic Community, which they joined in the first of January of 1986. It is very interesting to note how this historical evolution of the countries is reflected in the evolution of the synchronization of their economic cycles and how this can be read with the wavelet tools.

In Figure 6 (a), we see that until early 1980's, the two time-series are highly coherent and in Figure 6 (b), we observe that their phases, at business cycle frequencies 3–8 years period cycles, were well aligned in this period of time, with a slight lead from Portugal. From the early 1980s to about 1986, coinciding with the period immediately after the IMF intervention, we clearly see a de-synchronization between the two countries business cycles. Between 1986 and 1995, the two countries became more synchronized again, in particular at lower frequencies, as we can see in Figure 6 (a), but the phases were not aligned anymore. Instead, the phase-difference between Spain and Portugal (red line) tells us that the Portuguese business cycle was lagging the Spanish one. After 1999, when both countries joined the Euro, the phase-difference started approaching zero. After 2002, the phase-difference became almost zero, suggesting that the business cycles became aligned again. After 2004, we also observe a region of high coherency, which reinforces our previous conclusion. Therefore, coinciding with the adoption of a common currency, the business cycles became more synchronized.

NOTE 6.—The pictures and the numerical results given in the paper were obtained using a matlab toolbox developed by the authors, the ASToolbox, freely available at <http://sites.google.com/site/aguiarconraria/joanasoares-wavelets>.

REFERENCES

- [1] Aguiar-Conraria and M. J. Soares. Business cycle synchronization and the Euro: a wavelet analysis. *Journal of Macroeconomics*, 33(3):477–489, 2011.
- [2] L. Aguiar-Conraria and M. J. Soares. The Continuous Wavelet Transform: a Primer. Working Paper 16, NIPE/ University of Minho, 2011.
- [3] B. Cazelles, M. Chavez, G. C. de Magny, J.-F. Guégan, and S. Hales. Time-dependent spectral analysis of epidemiological time-series with wavelets. *Journal of the Royal Society Interface*, 4:625–636, 2007.
- [4] J. Connor and R. Rossiter. Wavelet transforms and commodity prices. *Studies in Nonlinear Dynamics and Econometrics*, 9(1):Article 6, 2005.
- [5] P. Crowley and D. Mayes. How fused is the Euro area core?: An evaluation of growth cycle co-movement and synchronization using wavelet analysis. *J. Bus. Cycle. Meas. Anal.*, 4:63–95, 2008.
- [6] P. M. Crowley. A guide to wavelets for economists. *Journal of Economic Surveys*, 21:207–267, 2007.
- [7] I. Daubechies. *Ten Lectures on Wavelets, volume 61 of CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM, Philadelphia, 1992.
- [8] M. Farge. Wavelet transforms and their applications to turbulence. *Annu. Rev. Fluid Mech.*, 24:395–457, 1992.
- [9] E. Foufoula-Georgiou and P. Kumar, editors. *Wavelets in Geophysics*, volume 4 of *Wavelet Analysis and its Applications*. Academic Press, 1994.
- [10] D. Gabor. Theory of communication. *J. Inst. Electr. Eng.*, 93:429–457, 1946.
- [11] M. Gallegati. Wavelet analysis of stock returns and aggregate economic activity. *Computational Statistics and Data Analysis*, 52:3061–3074, 2008.
- [12] M. Gallegati and M. Gallegati. Wavelet variance analysis of output in G-7 countries. *Studies in Nonlinear Dynamics and Econometrics*, 11(3):Article 6, 2007.
- [13] Z. Ge. Significance tests for the wavelet cross spectrum and wavelet linear coherence. *Ann. Geophys.*, 26:3819–3829, 2008.
- [14] R. Gençay, F. Selçuk, and B. Withcher. Multiscale systematic risk. *Journal of International Money and Finance*, 24:55–70, 2005.
- [15] P. Goupillaud, A. Grossman, and J. Morlet. Cycle-octave and related transforms in seismic signal analysis. *Geoexploration*, 23:85–102, 1984.
- [16] A. Grinsted, J. C. Moore, and S. Jevrejeva. Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlinear Proc. Geoph.*, 11:561–566, 2004.
- [17] M. Holschneider. *Wavelets: An Analysis Tool*. Clarendon Press, Oxford, 1995.
- [18] L. Hudgins, C. A. Friehe, and M. E. Mayer. Wavelet transforms and atmospheric turbulence. *Phys. Rev. Lett.*, 71(20):3279–3282, 1993.
- [19] A. A. Kallis. The ‘regime-model’ of Fascism: A typology. *European History Quarterly*, 30(1):77–104, 2000.
- [20] J. M. Lilly and S. C. Olhede. On the analytic wavelet transform. *IEEE Transactions on Signal Processing*, 1(11):1–15, 2007.
- [21] J. M. Lilly and S. C. Olhede. Higher-order properties of analytic wavelets. *IEEE T. Signal Proces.*, 57(1):146–160, 2009.
- [22] P. C. Liu. *Wavelets in Geophysics*, volume 4 of *Wavelet Analysis and its Applications*, chapter Wavelet spectrum analysis and ocean wind waves, pages 151–166. Academic Press, 1994.
- [23] A. K. Louis, P. Maaß, and A. Rieder. *Wavelets: Theory and Applications*. John Wiley & Sons, Chichester, 1998.
- [24] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, New York, 1998.
- [25] S. C. Olhede and A. T. Walden. Generalized Morse wavelets. *IEEE T. Signal Proces.*, 50(11):2661–2670, 2002.
- [26] J. B. Ramsey and C. Lampart. Decomposition of economic relationships by timescale using wavelets: Money and income. *Macroeconomics Dynamics*, 2:49–71, 1998.
- [27] I. W. Selesnick, R. G. Baraniuk, and N. G. Kingsbury. The dual-tree complex wavelet transform. *IEEE Signal Process. Mag.*, 22:123–151, 2005.
- [28] C. Torrence and G. P. Compo. A practical guide to wavelet analysis. *Bull. Am. Meteorol. Soc.*, 79:61–78, 1998.