

Spectral Analysis for Economic Time Series

N° 2003-07
November 2003

Alessandra IACOBUCCI
OFCE

Spectral Analysis for Economic Time Series *

Alessandra Iacobucci

OFCE, 69 quai d'Orsay, 75340 Paris Cedex 07 (France)

IDEFI, 250 rue Albert Einstein, 06560 Valbonne (France)

November 28, 2003

Abstract

The last ten years have witnessed an increasing interest of the econometrics community in spectral theory. In fact, decomposing the series evolution in periodic contributions allows a more insightful view of its structure and on its cyclical behavior at different time scales. In this paper I concisely broach the issues of cross-spectral analysis and filtering, dwelling in particular upon the windowed filter [15]. In order to show the usefulness of these tools, I present an application to real data, namely to US unemployment and inflation. I show how cross spectral analysis and filtering can be used to find correlation between them (i.e. the Phillips curve) in some specific frequency bands, even if it does not appear in raw data.

Keywords : spectral and cross-spectral methods, frequency selective filters, US Phillips curve.

JEL codes : C10, C14, E32.

*This paper is a chapter of a forthcoming book on new tools for economic dynamics analysis, edited by Springer-Verlag. I thank Francesco Saraceno and an anonymous referee for their precious comments and suggestions. E-mail: alessandra.iacobucci@sciences-po.fr

1 Introduction

The first appearance of spectral analysis in the study of macroeconomic time series dates from the middle 1960s, motivated by the requirement of a more insightful knowledge of the series structure and supported by the contemporaneous progress in spectral estimation and computation. The first works focused on the problem of seasonal adjustment procedures (see e.g. [20]) and on the general spectral structure of economic data [12]. Cross spectral methods were pointed out from the outset as being important in discovering and interpreting the relationships between economic variables [11, 13]. After the early years, the range of application of such analysis was extended to the study of other econometric issues, among which the controversial trend-cycle separation, the related problem of business cycles extraction and the analysis of co-movements among series, useful in the study of international business cycles. It has been clear from the beginning that spectral analysis is purely descriptive and cannot be straightforwardly used for forecasting; it is nevertheless a powerful tool for inspecting cyclical phenomena and highlighting lead-lag relations among series. It also provides a rigorous and versatile way to define formally and quantitatively each series components and, by means of filtering, it provides a reliable extraction method. In particular, cross spectral analysis allows a detailed study of the correlation among series.

In this synthetic overview I will focus on both filtering and cross spectral analysis, which are often two stages of the same procedure. As a matter of fact, besides the definition and extraction of the different components of a series – typically trend, business cycle and seasonalities – frequency filters can also be applied to perform a more targeted and efficient cross spectral analysis.

Time-frequency approaches — which represent the frequency content of a series, while keeping the time description parameter to give a three-dimensional time-dependent spectrum — will not be tackled in this paper. This is for essentially two reasons: first, they would require more than a simple section; second, and more importantly, because evolutionary spectral methods and wavelets are

suitable when dealing with very long time series, like those found in geophysics, astrophysics, neurosciences or finance. But their application to short series — the norm in macroeconomics — is difficult and may give unstable parameter-dependent results. For such series, traditional spectral analysis is probably more suitable.

The paper is organized as follows: the first section contains a concise description of spectral estimation and filtering issues¹ together with a recall of discrete Fourier analysis; in the second section I expose the cross spectral analysis procedure, with a very short account of the genuinely technical yet central issue of estimation; in the third section I show an application of the techniques to the US Phillips curve. Some remarks and the conclusion can be found in the fourth and last section.

2 Spectral Estimation and Filtering: a Brief Review

At a first glance, the overall behavior of time series may be decomposed in three main parts: long, medium and short run behavior. These three parts are respectively associated with slowly evolving secular movements (the trend), a faster oscillating part (the business cycles) and a rapidly varying, often irregular, component (the seasonality). As it is often the case when no testable *a priori* hypothesis on the data generating process (i.e. on the model) is available, this separation is very phenomenological.

Modern empirical macroeconomics employs an assortment of *ad hoc* detrending and smoothing techniques to extract the business cycle, like moving averages to eliminate the fast components, first-differences to cut out the long term movements, or even the simple subtraction of the linear trend, to cancel the slow drift variables. Though conceptually not wrong, these methods lack a formal decomposition of the series and are incapable of giving a definition of the business cycle based on some required and adjustable characteristics. This is why the Fourier de-

¹For an more extensive and detailed treatment the interested reader may refer, among others, to the celebrated book by Jenkins and Watts [16].

composition remains one of the most insightful ways of performing the separation of a signal into different purely periodic components.

Consider a finite series $u(j)$ of length $T = N\Delta t$, where N is the number of data and Δt the sampling periodicity; the frequency $\nu_k = k/(N\Delta t)$ and the time $t_j = j\Delta t$ are indexed by k and j respectively.

The discrete Fourier transform (DFT) $U(k)$ of $u(j)$ and its inverse (IDFT) for finite series are

$$U(k) = \frac{1}{N} \sum_{j=0}^{N-1} u(j) e^{-i2\pi jk/N}, \quad (1)$$

$$u(j) = \sum_{k=-\lfloor N/2 \rfloor}^{\lfloor (N-1)/2 \rfloor} U(k) e^{i2\pi jk/N}, \quad (2)$$

where $\lfloor \cdot \rfloor$ denotes the largest integer smaller or equal than the operand, $k \in [-\lfloor N/2 \rfloor, \lfloor (N-1)/2 \rfloor]$ and $j = 0, \dots, N-1$. Of course, the discretization of the signal (i.e. its sampling with some finite period Δt) implies a limitation of its spectrum to the band $\nu \in [-(2\Delta t)^{-1}, (2\Delta t)^{-1}[$, where $(2\Delta t)^{-1}$ is the *Nyquist frequency*, as frequencies outside that range are folded inside by the sampling (an effect known as *aliasing* [7]). On the other hand, the finiteness of the signal in time implies a discretization of the spectrum, the interval between two successive values being $2/N$.

Equation (1) can only be an approximation of the corresponding real quantity, since it provides only for a finite set of discrete frequencies. The quantity $P_u(k) = |U(k)|^2$ is the signal (power) spectrum and its “natural” estimator would be Schuster’s *periodogram*

$$\begin{aligned} P_u(k) &= \Delta t \sum_{J=-(N-1)}^{N-1} \gamma_{uu}(J) e^{-i2\pi Jk/N} \\ &= \Delta t \sum_{J=-(N-1)}^{N-1} \gamma_{uu}(J) \cos \frac{2\pi Jk}{N}, \end{aligned} \quad (3)$$

where $\gamma_{uu}(J) = \gamma_{uu}(-J) = N^{-1} \sum_{j=-(N-J)}^{N-J} (u(j) - \bar{u})(u(j+J) - \bar{u})$ is the standard sample estimation at lag J of the autocovariance function.

The periodogram is a real quantity – since the series is real and the autocovariance is an even function – and is an asymptotically unbiased estimator of the theoretical spectrum. Yet, in the case of finite series, it is *non-consistent* since the power estimate at the individual frequency fluctuates with N , making difficult its interpretation. To build a spectral estimator which is more stable – i.e. has a smaller variance – than $P_u(k)$, we turn to the technique of *windowing* (see [8, 16, 21] among others). This technique is employed both in time and in frequency domain to smoothen all abrupt variations and to minimize the spurious fluctuations generated every time a series is truncated. The result of windowing is the *smoothed spectrum*

$$\hat{S}_u(k) = \Delta t \sum_{J=-(N-1)}^{N-1} w_M(J) \gamma_{uu}(J) \cos \frac{2\pi Jk}{N}, \quad (4)$$

where the autocorrelation function is weighted by the *lag window* $w(j)$ of width M [1]. It can be shown that this is equivalent to splitting the series in N/M sub-series of length M , computing their spectra and taking their mean.

Since $P_u(k)$ and $\gamma_u(J)$ are related by DFT (equation 3), equation (4) can also be written as

$$\hat{S}_u(k) = \Delta t \sum_{k'=-\lfloor N/2 \rfloor}^{\lfloor (N-1)/2 \rfloor} P_u(k') W_{M'}(k - k'), \quad (5)$$

that is, the *convolution* of the periodogram $P_u(k)$ with the Fourier transform of $w_M(j)$, the *spectral window* $W_{M'}(k)$ of width $M' = M^{-1}$. Thus the smoothed spectrum at k is nothing but the periodogram seen through a window opened on a convenient interval around k . Equations (4) and (5) represent two perfectly equivalent ways to compute the smoothed spectrum. Usually, the multiplication, as in equation (4), is chosen because it is easier to compute. Nevertheless, sometimes the convolution may be more convenient as we shall see in the section devoted the windowed filter.

The choice of the lag window width M is performed by choosing a "reasonably" narrow window, i.e. a small initial value of M , and then widening it until a good spectral stability is obtained, i.e. until the spectral density remains roughly

unchanged as M increases. Widening the lag window $w_M(j)$ corresponds to narrowing the band covered by its Fourier transform, the spectral window $W_M(k)$. This is why the procedure is called *window-closing* [16]. This method allows to learn progressively about the shape of the spectrum. The initial choice of a wide bandwidth usually masks some details of the spectrum. By decreasing the bandwidth, more significant details can be explored. The choice of M is rather tricky since it has to be large enough to let all the fundamental details of the spectrum appear, but not too large, to prevent the generation of spurious peaks.

Windows can be chosen among those already existing in the literature (rectangular, triangular, Bartlett, Parzen, Tuckey, Blackman, Hamming,...) or can be built *ad hoc* for the specific problem treated. The research of the optimal window involves a compromise between *accuracy* and *stability* of the estimator (see [8, 16, 21] among others). Moreover, windows are used both in time and in frequency domain, according to the researcher needs. Both lag windows and spectral-windows can be used either as multiplying window, like the lag window in equation (4), or as convolving windows, like the spectral window in equation (5). Since a convolution in the time domain becomes a multiplication in the frequency domain and *vice versa*, a multiplying (convolving) lag window becomes by Fourier transform a convolving (multiplying) spectral window.

2.1 Filtering

The filtering operation can be performed either in time or in frequency domain since both approaches are equivalent by

$$\begin{aligned} v(j) &= h(j) \circledast u(j) \equiv \sum_{n=0}^{N-1} h(n) u(j-n)_{\text{mod } N} , \\ &= \sum_{k=-N/2}^{(N-1)/2} H(k) U(k) e^{-i2\pi jk/N} . \end{aligned} \quad (6)$$

The previous relation is nothing but the finite discrete version of the *convolution theorem* [16], where the linear convolution has been substituted by the *circu-*

lar convolution (\circledast) of length equal to the number of data N . Thus filtering simply consists in multiplying $U(k)$ by the filter frequency response $H(k)$ or, equivalently, in convolving the signal $u(j)$ with the filter time response $h(j)$, obtained from $H(k)$ by IDFT. In particular, the *band-pass filter* selects a frequency range, so that $H(k) = 1$ for $k_l \leq |k| \leq k_h$ (*pass-band*) and zero elsewhere (*stop-band*). Of course, the *low-pass filter* has $k_l = 0$ and selects all frequencies lower than k_h , while the *high-pass filter* has $k_h = N/2$, correspondent to the Nyquist frequency $\nu_N = 1/(2\Delta t)$, and selects all frequencies higher than k_l . Notice that the filter $H(k)$ is not causal in the time domain because it requires future values as well as past ones (see equation (6)). Asymmetrical (one-sided) filters using only past values may seem interesting because they allow forecasting [5, 18]; but, unless special care is taken in designing them — e.g allowing for a complex time response function — they are dangerous to use because they induce frequency-dependent phase shifts and may thus change the causality relations among different frequency components [15, 21]. This would make cross-correlation analysis useless.

A filter which is real in time domain ($h(j) = h^*(j)$), is symmetric in frequency domain ($H(k) = H(-k)$) and *vice versa*. Therefore, if we want real signals to remain real after filtering, both time *and* frequency response functions have to be real and symmetric, to avoid time and phase shifts. Indeed, it is easy to see that if the filter $H(k)$ is a complex function different frequencies undergo different phase shift and timing relations among components are destroyed (*dispersive filter*).

In the circular convolution the finite signal is replaced by its periodic version $u(N+j) = u(j \bmod N)$, the maximum period length being implicitly assumed by the Fourier transform to be $T = N\Delta t$. This amounts to assuming that the only frequencies present in the signal are integer multiples of T^{-1} , which is in general false and affects the analysis. Indeed the “forced” periodicity introduces an artificial discontinuity at the edges of the time series, that is reflected by spurious oscillations in the series DFT, the so-called *Gibbs phenomenon*. These oscillations are due to the form of the DFT of a rectangular window of width $T = N\Delta t$, which

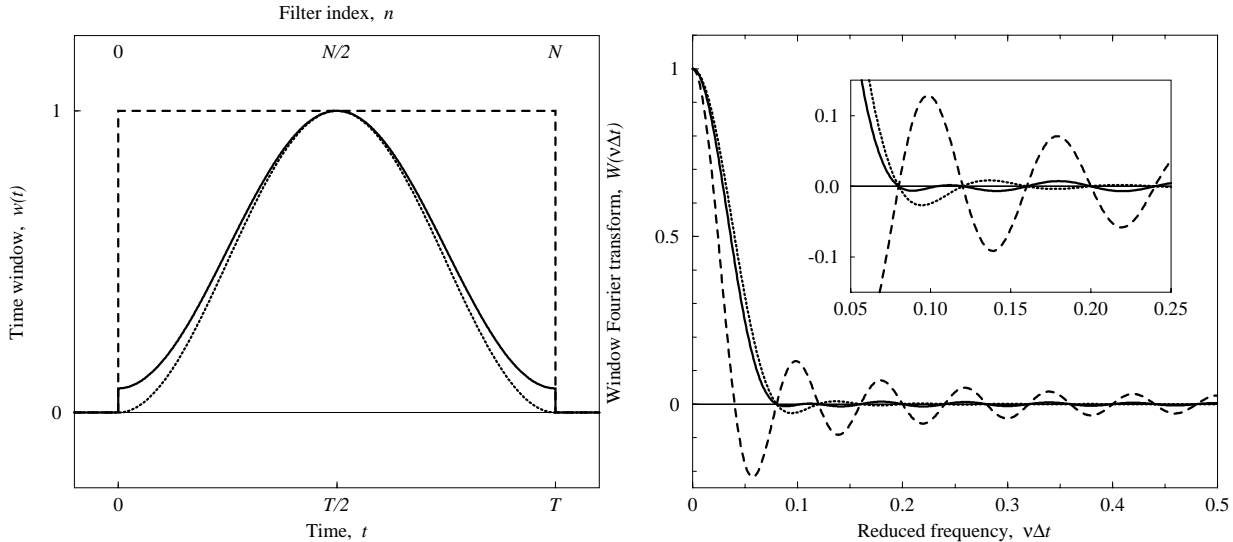


Figure 1: **Window Functions and Their Frequency Response.** *The rectangular (dashed line), Hanning (dotted line) and Hamming (full line) time windows (left panel) and their respective Fourier transforms (right panel). For the latter, the number of points $N = 15$ has been chosen rather small to emphasize the differences. Note in the zoom (right panel, inset) the reduced side lobe amplitude and leakage of the Hanning and Hamming windows with respect to the rectangular one, the Hamming window performing better in the first side lobe.*

goes like $\sin(\pi\nu T)/(\pi\nu T)$ (see Figure 1).

The only way to prevent this effect, would be to choose T (or equivalently N) as a multiple of the largest period that is likely to occur. Unfortunately, this is feasible only if we have some idea of the frequencies involved in the process and would in any case entail some loss of data at one or both sample ends. As for the cutoff frequencies $\nu_l = k_l/(N\Delta t)$ and $\nu_h = k_h/(N\Delta t)$, given the value of N , they must be chosen to be multiples of T^{-1} , otherwise the filter does not completely remove the zero frequency component (i.e. the signal mean) and cannot help in eliminating unit roots (see below).

The time coefficients of an *ideal band-pass filter* are

$$h(j) = \frac{\sin(2\pi\nu_h j \Delta t) - \sin(2\pi\nu_l j \Delta t)}{\pi j}, \quad j = 1, \dots, \infty \quad (7)$$

$$h_0 = 2(\nu_h - \nu_l) \Delta t,$$

and are obtainable only in the case of infinite series, since this could be the only way to precisely select the frequency band $[\nu_l, \nu_h]$. Indeed the filter must distinguish between frequencies ν_h (or ν_l) and $\nu_h + d\nu$ (or $\nu_l - d\nu$) when $d\nu \rightarrow 0$,

that is, $N \rightarrow \infty$. This is why this filter is called “ideal”. In the case of finite series of length N only the first N coefficients can be calculated

$$\begin{aligned} h(j) &= \frac{\sin(2\pi k_h j/N) - \sin(2\pi k_l j/N)}{\pi j}, & j = 1, \dots, N-1 \\ h_0 &= \frac{2(k_h - k_l)}{N}. \end{aligned} \quad (8)$$

The same result is obtained by multiplying the coefficients (7) by the coefficients of a rectangular lag window of width N . As we saw above, the effect of this truncation is the Gibbs phenomenon, i.e. the appearance of spurious oscillations in the frequency response (see Figure 1). This causes the so-called *leakage*: the component at one frequency “contaminates” the neighboring components by modifying their amplitude. Thus, frequency components which are contiguous to the band limits, are allowed to leak into the band. Again the application of an appropriate window is the most straightforward way to bypass this problem and obtain a smoother response, as shown below, in the section devoted to the windowed filter.

Since the Fourier theory and the definition of the spectrum only apply to stationary time series, it is necessary to detect non-periodic components prior to the analysis of a series spectrum. First and foremost, it should be established whether the series has a trend, and, if so, whether the trend is *stochastic* or *deterministic*. Unfortunately there is no direct method to distinguish between the two categories in the case of raw data with no underlying model, so that the choice may often depend on the researcher’s insight (see, e.g. [6]). If the trend is deterministic, e.g. a polynomial function of time, the Fourier basis decomposition is not unique, since the polynomial term and the periodic one are not orthogonal (a polynomial term contains all possible frequency components). Therefore, the operations of detrending and filtering do not commute and the trend must be preliminarily removed. It is also *necessary* to remove the artificial discontinuity introduced at the edges of the interval by the combination of the trend and the periodicity induced by the Fourier representation. In the case of a linear deterministic trend — that should be established beforehand by looking at the correlation coefficient

of the signal with time —, the subtraction of the ordinary least-squares linear fit from the original series is performed, more or less explicitly, by some filtering procedures [2, 5, 15].

If the trend is stochastic, and the observed signal is an $I(p)$ process, i.e. the result of p integrations of a stationary process, it has a spectrum that goes as ν^{-p} for small ν . Thus, a filter whose frequency response function goes like ν^p makes the filtered series stationary. In particular, a ν^2 -like response is sufficient for the elimination of two unit roots. The typical way of treating $I(p)$ signals would be to apply p times the first-difference operator to remove the p unit roots. The main drawback of this procedure is that the difference operator is an asymmetric filter, thus it has a complex response $H_{L^p}(k) = i^p e^{ip\pi k/(N\Delta t)} \left(2 \sin \frac{\pi k}{N\Delta t}\right)^p e^{i2\pi k j/N}$ which introduces a frequency-dependent phase shift. Moreover it amplifies all frequencies larger than one third of the Nyquist frequency (see among others [8, 15]). This means that p applications of this filter will cause a dramatic amplification of high-frequency components and thus of noise. Moreover, the filter response varies almost linearly for small frequencies, so that low-frequency components are strongly attenuated. It is then very hard to obtain an ideal filter after differencing, especially when dealing with series with a Granger-shaped [11] spectrum, in which much of the power occurs at very low frequencies, like those common in macroeconomics.

2.2 The Windowed Filter

Good approximations of the ideal filter — “good” referring to some optimization criteria, like the (weighted) difference between the desired and the effective response [21] — are the Hodrick-Prescott [9] and the Baxter-King [2] filters². These procedures make stationary at least $I(2)$ processes. In particular, the HP-filter can eliminate up to four unit roots. As for these filters, which are widely known,

²Christiano and Fitzgerald [5] have also designed a band-pass filter, which is more complicated than the previous ones but, in my opinion, also more questionable (see [15]).

the reader is referred to the original papers [9, 2]. Here I will focus on a filter recently proposed by Iacobucci and Noullez [15], which is obtained by the technique of windowing. This filter performs better than the others, since it has minimum leakage, a significantly flatter frequency response function in the pass-band and involves no loss of data.

As we have previously seen, the filter obtained by truncation has two main drawbacks: large amplitudes and a slow decay of the spurious lobes in the response function (see Figure 1). These can be ascribed, as previously said, to the discontinuous shape of the above-mentioned lag window, whose $\sin(\pi\nu T)/(\pi\nu T)$ -profile Fourier transform disturbs the ideal frequency response. It seems then natural to try to adjust the shape of the rectangular window to obtain a gain that goes to zero faster. For this purpose, the “adjusted” window should be chosen to go to zero continuously with its highest possible order derivatives, at both ends of the observation interval [15].

Among a certain number of possible windows, Iacobucci and Noullez choose the Hamming window

$$w^{\text{Ham}}(j) = 0.54 + 0.46 \cos\left(\frac{2\pi j}{N}\right), \quad (9)$$

which is obtained by a combination of the Hanning window $w^{\text{Han}}(j) = \frac{1}{2} + \frac{1}{2} \cos\left(\frac{2\pi j}{N}\right)$ and the rectangular window, to minimize the amplitude of the side lobes (see Figure 1). Its Fourier transform

$$W^{\text{Ham}}(k) = \left[0.08 + \frac{0.92}{1 - k^2}\right] \frac{\sin(\pi k)}{\pi k} \quad (10)$$

decreases like $(\nu T)^{-1}$ for large ν , but with a much smaller amplitude than the rectangular window. Moreover, it has non-zero components only at $k = 0$ and $k = \pm 1$.

The windowed filter algorithm is the following:

- subtract, if needed, the least-squares line to remove the artificial discontinuity introduced at the edge of the series by the Fourier representation;

— compute the discrete Fourier transform of $u(j)$

$$U(k) = \frac{1}{N} \sum_{j=0}^{N-1} u(j) e^{-i2\pi jk/N}, \quad k = 0, \dots, \lfloor N/2 \rfloor,$$

where $U(k) = U^*(-k) = U(-k)$;

— compute the DFT of the Hamming-windowed filtered series

$$\begin{aligned} V(k) &= [W(k) * H(k)] U(k) = \sum_{k'=-\lfloor N/2 \rfloor}^{\lfloor N/2 \rfloor} W(k') H(k - k') U(k) \\ &= [0.23H(k - 1) + 0.54H(k) + 0.23H(k + 1)] U(k), \end{aligned}$$

where $k = 0, \dots, \lfloor N/2 \rfloor$ and $H(k)$ is defined by the frequency range as

$$H(k) = H(k)^{\text{ideal}} \equiv \begin{cases} 1 & \text{if } \nu_l N \Delta t \leq |k| \leq \nu_h N \Delta t \\ 0 & \text{otherwise} \end{cases};$$

— compute the inverse transform

$$v(j) = \left[V(0) + \sum_{k=1}^{\lfloor N/2 \rfloor} (V(k) e^{i2\pi jk/N} + V(k)^* e^{-i2\pi jk/N}) \right], \quad j = 0, \dots, N-1.$$

Notice that windowing is performed in the frequency domain by convolution of the window Fourier transform with the ideal filter response. This is computationally more convenient than time domain multiplication, since the Hamming window Fourier transform has only three non-zero components, as I have already stressed. This procedure ensures both the best possible behavior in the upper part of the spectrum and the complete removal the signal mean. In the application I propose in Section 4, I make use of this filter because of its many advisable properties compared to the others, namely a flat, well-behaved response function and the fact that it involves no loss of data.

3 Cross Spectral Analysis: the Bivariate Extension

While univariate spectral analysis allows the detection of movements inside each series, by means of bivariate spectral analysis it is possible to describe pairs of time

series in frequency domain, by decomposing their covariance in frequency components. In other words, cross spectral analysis can be considered as the frequency domain equivalent of correlation analysis. The definition of the (smoothed) *cross spectrum*, analogously to that of the (smoothed) spectrum (see equation (4)), is obtained by substituting the cross covariance function for the autocovariance function. Thus, if we have two time series $u_1(j)$ and $u_2(j)$ and their crosscovariance $\gamma_{12}(J) = \gamma_{21}(-J)$, the cross spectrum is

$$\hat{S}_{12}(k) = \Delta t \sum_{J=-(N-1)}^{N-1} w(J) \gamma_{12}(J) e^{-i2\pi Jk/N} = \hat{C}_{12}(k) - i\hat{Q}_{12}(k) \quad (11)$$

and is in general *complex*, since the crosscovariance is not an even function. The real part $\hat{C}_{12}(k)$ is the *cospectrum* and the imaginary part $\hat{Q}_{12}(k)$ the *quadrature spectrum*. Keeping the time-frequency analogy, I introduce the typical cross spectral quantities and indicate in parentheses the time domain equivalent:

— the *coherency spectrum* (correlation coefficient)

$$\hat{K}_{12}(k) = \frac{|\hat{S}_{12}(k)|}{\sqrt{\hat{S}_1(k)\hat{S}_2(k)}} = \frac{\sqrt{\hat{C}_{12}(k)^2 + \hat{Q}_{12}(k)^2}}{\sqrt{\hat{S}_1(k)\hat{S}_2(k)}}, \quad (12)$$

which measures the degree to which one series can be represented as a linear function of the other (sometimes its square is used, whose time domain equivalent is the R^2);

— the *phase spectrum* (time-lag)

$$\hat{\Phi}_{12}(k) = \arctan\left(-\frac{\hat{Q}_{12}(k)}{\hat{C}_{12}(k)}\right), \quad (13)$$

which measures the phase difference between the frequency components of the two series: the number of leads ($\hat{\Phi}_{12}(k) > 0$) or lags ($\hat{\Phi}_{12}(k) < 0$) of $u_1(k)$ on $u_2(k)$ in sampling intervals at frequency ν_k is given by the so-called *standardized phase* $(2\pi\nu_k)^{-1}\hat{\Phi}_{12}(k)$;

— the *gain* (regression coefficient)

$$\hat{G}_{12}(k) = \frac{|\hat{S}_{12}(k)|}{\hat{S}_1(k)}, \quad (14)$$




Figure 2: **US series.** *Raw series used in the application and the corresponding detrended ones: unemployment (raw: thin line, detrended: thick line) and inflation (raw: thin dashed line, detrended: thick dashed line). The data are monthly and cover the period Jan60-Dec01.*

which indicates the extent to which the spectrum of $u_1(k)$ has been modified to approximate the corresponding frequency component of $u_2(k)$.

The analysis of these three quantities together with the (auto) spectra of each series and with the amplitude of their cross spectrum gives us an overall view of the frequency interaction of the two series. As anticipated at the beginning, filtering procedures are often coupled to cross spectral analysis, either as preliminary or as a consequential step. In fact, it is sometimes evident from spectral peaks investigation that most of the power is contained in one or more bands. In particular, many macroeconomic time series (in level) have the typical Granger-shaped spectrum [12]. Such peaks may leak into nearby components and corrupt spectral and cross spectral investigation in low-power bands. That is why it may be advantageous to “pre-filter” the data. On the other hand, filtering can also be required afterwards when the spectral power concentration occurs in the coherency spectrum. This would involve that only some bands are important for the “interaction” between the series, all the remaining frequency components being useless.

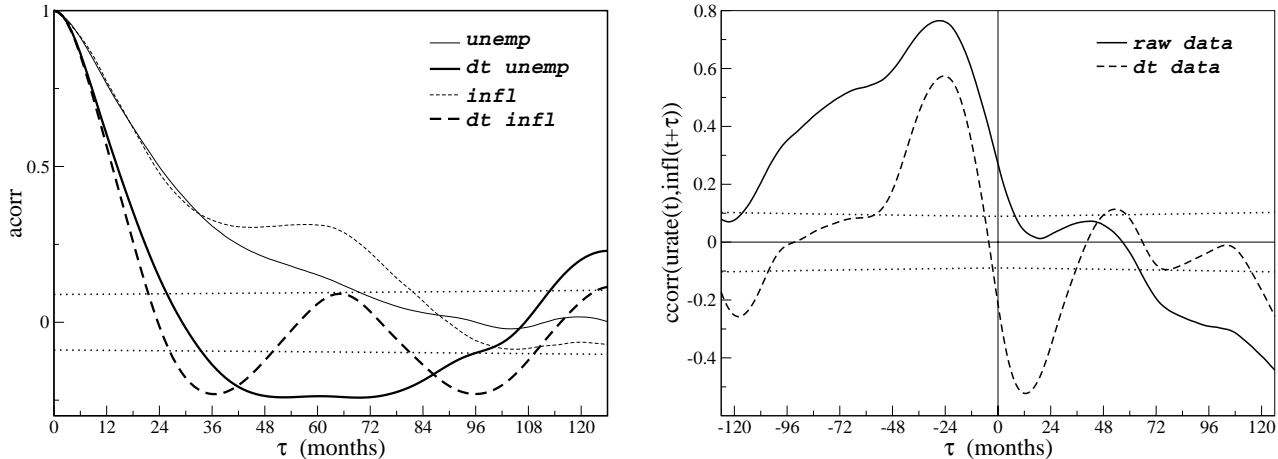


Figure 3: **Auto and Cross Correlograms.** The dotted lines represent the approximated two standard error bound from the null value, computed as $\pm 2/\sqrt{N - \tau}$.

4 An Application to the US Phillips Curve

In this section the methods just described are applied to the analysis of the US Phillips curve in the frequency domain³.

We start by looking at the raw data (Figure 2): neither unemployment nor inflation show any obvious trend. Nevertheless, since the data cover a period of 42 years, we could expect the existence of a low-frequency trend, unobservable by simple visual inspection. Moreover, as the data are not seasonally adjusted, we risk to find an effect on correlation we are not interested in. I thus perform a filtering operation by means of the windowed filter described above which eliminates all periodicities smaller than one year and higher than 21 years, which, as we saw, may be fictitiously introduced by the Fourier approach. The filtering operation on this particular band has the effect of detrending and smoothing our 42-years original series and hereafter I will refer to the resulting series as to the *detrended* series. Figure 3 reports raw and detrended data autocorrelation functions for both unemployment and inflation and their cross-correlation function. Only about $N/4$ lags (10.5 yr) are reported, as suggested by [4], in order to have

³This section builds on [10]. The issue of the Phillips curve historical behavior at different frequencies is also broached in [14, 17].

enough lagged products at the highest lag, so that a reasonably accurate average is obtained. It may be seen that autocorrelation functions (left panel) drop to zero more quickly in the detrended than in the raw series. Furthermore a sort of oscillating behavior emerges in the case of detrended inflation, while it was absent in the raw case. We also observe significant negative autocorrelation values for both series which appear only after the detrending operation. These values correspond approximatively to lags between 3 and 7.5 years for unemployment and between 2.5 and 4 years and 7 and 9 years for inflation, the intermediate values being not significantly different from zero. Finally and more crucially, we notice that following the detrending operation a negative cross correlation (right panel) emerges in the short-to-medium run (0 to 36 months lag), which was absent in the raw case. This justifies our fears about the “hiding” effect of low-frequency high-power spectral components on short-to-medium term correlation visibility. The negative cross correlation between the detrended series means that: (a) in the (wide) band of frequency $\nu \in [21^{-1}, 1] \text{ yr}^{-1}$ that we extracted it does exist a contemporaneous Phillips curve, as shown by the negative cross correlation at lag zero, which did not appear by the sole visual inspection of raw data; (b) there is a retarded negative effect of unemployment on inflation in this frequency range, which reaches its maximum at about 1 year, meaning that a rise in inflation will follow of an year a fall in unemployment, in other words a delayed Phillips curve.

Obviously, we are not able to establish a lead-lag relation by the simple study of the cross-correlation function. In fact, if we look at its negative lags part, we find a positive correlation between inflation and retarded unemployment, which reaches its maximum at a 2-year lag. This implies a positively sloped Phillips curve. We would need additional information to establish which of the two is leading the other. Cross spectral analysis (Figure 4) can answer this purpose and more. Notice that the residual spectral power at frequencies lower than $(21 \text{ yr})^{-1}$ is an effect of the smoothing. The same happens to the other quantities, thus their value outside the band should be disregarded. The parameter M was set to 140




Figure 4: **Cross spectral analysis.** *Main quantities relative to US unemployment and inflation: auto and cospectrum (top left), standardized phase (top right), coherency (bottom left) and gain (bottom right). For a greater legibility, on abscissae I report the period, which is the inverse of the frequency and is expressed in years.*

after the preliminary window closing procedure. We notice (top left panel) that the spectrum of inflation is higher than that of unemployment, confirming the former's higher variance. Moreover, inflation shows two non-harmonic peaks at the periodicities ν^{-1} of 14 and 6 years, while the only unemployment spectral peak is at 10.5 years. The cospectrum shows a concentration of the two series covariance approximatively in the periodicity band $[21, 6]$ yr. In point of the standardized phase (top right panel), the plot shows a leading behavior of unemployment on inflation in the *same* band ($[21, 6]$ yr). This is consistent with the findings of the cross correlation inspection, which showed a negative correlation between unemployment and lagged inflation (Figure 3). The leads (lags) of unemployment

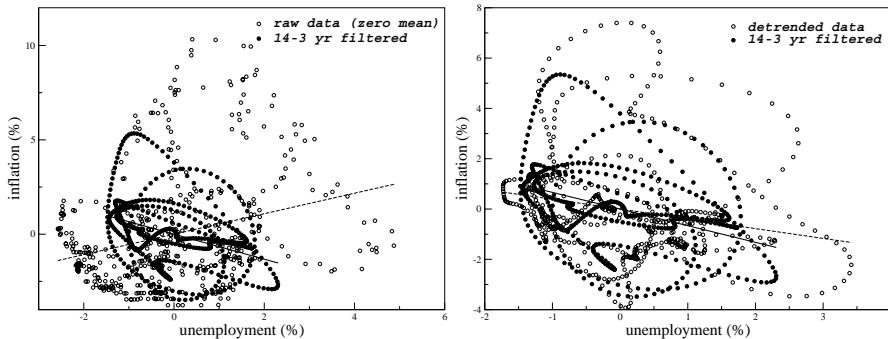


Figure 5: **Phillips curves.** Raw data (left panel), detrended (right panel) and 14 – 3 yr filtered US Phillips curves, with their corresponding OLS line. The raw data curve has been translated to the origin by subtracting from both series their means.

(inflation) components vary from a maximum of about 2 months and a half for the $\nu^{-1} = 21$ yr component to zero (i.e. coincident) for the $\nu^{-1} = 6$ yr component. For periodicities in the band $[6, 3]$ yr, we remark a negative phase, which would imply a leading behavior of inflation on unemployment. Nevertheless, the cospectrum has very low values in the second half of this band, where the phase is more significantly different from zero. We may thus conclude that these components do not account for much of the series covariance and that the prevailing trend is the former, i.e. unemployment leading inflation. The coherency plot (bottom left panel) shows a maximum frequency domain correlation at 5.25 yr, suggesting that filtering around this frequency would yield a more pronounced Phillips relation. The same information is given by the gain function (bottom right panel), whose maximum is found at the same frequency as the coherency.

Regarding the overall behavior of coherency, gain and phase, there is a transition at $\nu^{-1} = 3$ yr. In fact, while their behavior for $\nu^{-1} < 3$ yr can be considered reliable for cross spectral analysis, their non negligible value for $\nu^{-1} > 3$ yr might be the effect of some divergence (little denominators), since the cospectrum and the individual spectra above this value are nearly zero. To see more clearly at high frequencies, that part of the spectrum should be studied separately, that is we should extract the frequency components corresponding to the band $[3, 1]$ yr or $[3, 0.5]$ yr and perform the cross spectral analysis again. The interested reader

is referred to [10].

Turning to the Phillips curve analysis, Figure 5 shows the raw data, the detrended and the 14 – 3 yr filtered curve. The first thing we notice is that the raw curve OLS line has a positive slope of 0.54 ($t_{Student} = 6.30$), with a low correlation coefficient of 0.27. If we look at the detrended curve, the slope becomes negative (-0.39 with $t_{Student} = -4.9$), with a lower correlation coefficient of -0.21 . This is in agreement with the information given by the cross correlation (Figure 3). Finally, I filtered both unemployment and inflation on the band [14, 3] yr, containing the coherency maximum and indeed we find a negative slope of -0.66 with the highest correlation coefficient $\rho = -0.38$ ($t_{Student} = -9.32$). Thus, cross spectral analysis guided us in finding the band where we can detect a stronger Phillips relation.

5 Conclusion

This paper highlighted the main features of spectral analysis and their practical application. After a general theoretical introduction, I approached the issue of filtering for the extraction of particular components, mostly those related to the business cycle. In fact one of the advantages of the method is that it allows a quantitative definition of the cycle, and the extraction of long, medium or short term components, according to the researcher's wish. Then, I sketched the theory and practice of cross spectral analysis introducing some typical concepts, like coherency and phase spectrum, which may provide some essential information, complementary to that given by time domain methods. Finally, I applied these tools to the study of US Phillips curve. Thanks to the combined analysis, I managed to show that a Phillips relation arises between inflation and unemployment, at the typical business cycle components ([14, 6] yr), even if there is no hint of it in raw data. Moreover, by means of phase spectrum analysis, I showed that unemployment leads inflation, the latter being delayed of about one year.

To conclude I would like to spend some words in favor of the *fact-without theory* technique of filtering, which, after a period of great glamour, has lately come under attack. It goes without saying that as all other methods in time series analysis it has limits which have to be known and thoroughly explored to ensure a proper utilization; and, as any other method or model, it can not be expected to be universal.

A major limit of this approach is that it is impossible to say anything about the evolution in time of the frequency content of the series, since the spectrum depends on the frequency but not on time. Thus, if a particular frequency component remains “switched on” only during a subsample, it is impossible to detect this interval by means of the inspection of sole series spectrum. To keep the information about time-localization some other kind of analysis is required, like time-frequency or wavelets, which, as we said, are tricky to apply to short series and for this reason may give unstable results.

Nevertheless, some other criticisms seem to be off the mark. For example, it has been often argued that spectral analysis does not allow to disentangle the data generating process of a series from its spectrum (see, e.g. [3, 19]) — typically to separate the business cycle from the trend in a model where the latter is nonstationary, e.g. an $I(1)$ process [3, 19]. But in general, this would require a separation *within* the individual frequency components, something that this purely descriptive method is not meant to do. It would be like dismissing correlation analysis because it fails to detect causality between two variables.

In other words, the tool is far from being perfect, but it has been too hastily dismissed by some quarters with unwarranted arguments.

References

- [1] Bartlett, M. S., *An Introduction to Stochastic Processes with Special Reference to Methods and Applications*, Cambridge University Press, Cambridge (1953).
- [2] Baxter, M. and R. G. King, "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series", *The Review of Economics and Statistics* vol. 8, no. 4 (November 1999), pp. 575–93.
- [3] Benati, L., "Band-Pass Filtering, Cointegration and Business Cycle Analysis", Bank of England, working paper (2001).
- [4] Box, G. E. P., G. M. Jenkins and G. C. Reinsel, *Time Series Analysis. Forecasting and Control*, third edition, Prentice Hall, Upper Saddle River, New Jersey (1994).
- [5] Christiano, L. and T. J. Fitzgerald, "The Band-Pass Filter", *International Economic Review*, vol. 44, no. 2, May 2003.
- [6] Hamilton, J. D., "Time Series Analysis", Princeton University Press, Princeton, New Jersey (1994).
- [7] Hamming, R. W., *Numerical Methods for Scientists and Engineers*, second edition, Dover Publications, Inc., New York, 1973.
- [8] Hamming, R. W., *Digital Filters*, third edition, Dover Publications, Inc., New York, 1998.
- [9] Hodrick, R. J. and E. C. Prescott, "Postwar US Business Cycles: An Empirical Investigation", *Journal of Money, Credit, and Banking*, vol. 29-1 (1997), pp. 1–16.
- [10] Gaffard, J. L. and A. Iacobucci, "The Phillips Curve: Old Theories and New Statistics", mimeo.

- [11] Granger, C. W. J. and M. Hatanaka, *Spectral Analysis of Economic Time Series*, Princeton University Press, Princeton, New Jersey (1964).
- [12] Granger, C. W. J, "The Typical Spectral Shape of an Economic Variable", *Econometrica*, vol. 34(1), (1966), pp. 150–161.
- [13] Granger, C. W. J, "Investigating Casual Relations by Econometric Models and Cross-Spectral Methods", *Econometrica*, vol. 37(3), (1969), pp. 424–438.
- [14] Haldane, A. and D. Quah, "UK Phillips Curves and Monetary Policy", *Journal of Monetary Economics*, Special Issue: *The Return of the Phillips Curve*, vol. 44, no. 2, (1999), pp. 259-278.
- [15] Iacobucci, A. and A. Noullez, "Frequency Filters for Short Length Time Series", Working Paper IDEFI–IDEE 2002, n. 1.
- [16] Jenkins, G. M. and D. G. Watts, *Spectral Analysis and Its Applications*, Holden-Day, San Francisco, (1969).
- [17] Lee, J., "The Phillips Curve Behavior Over Different Horizons", *Journal of Economics and Finance*, 19 (1995), pp. 51-69.
- [18] Mitchell, J. and K. Mouratidis, "Is There a Common Euro-Zone Business Cycle?", presented at the colloquium on *Modern Tools for Business Cycles Analysis*, EUROSTAT, Luxembourg, 28-29 November 2002.
- [19] Murray, C. J., "Cyclical Properties of Baxter-King Filtered Time Series", *The Review of Economics and Statistics*, May 2003, 85(2): 472-476.
- [20] Nerlove, M., "Spectral Analysis of Seasonal Adjustment Procedures", *Econometrica*, vol. 32(3), (1964), pp. 241–286.
- [21] Oppenheim, A. V. and R. W. Schaffer, *Discrete-Time Signal Processing*, second edition, Prentice-Hall, New Jersey, 1999.