INTERPRETATION OF THE EFFECTS OF FILTERING INTEGRATED TIME SERIES

João Valle e Azevedo

September 2007

Please address correspondence to
João Valle e Azevedo
Economics and Research Department
Banco de Portugal, Av. Almirante Reis no. 71, 1150-012 Lisboa, Portugal;
Tel.: 351 21 3130163, Email: jvazevedo@bportugal.pt

Interpretation of the Effects of Filtering Integrated Time Series

João Valle e Azevedo[*]

Banco de Portugal & Universidade NOVA de Lisboa

September 21, 2007

Abstract

We resort to a rigorous definition of spectrum of an integrated time series in order to characterise the implications of applying linear filters to such series. We conclude that in the presence of integrated series the transfer function of the filters has exactly the same interpretation as in the covariance stationary case, contrary to what many authors suggest. This disagreement leads to different conclusions regarding the link of the original fluctuations with the transformed fluctuations in the time series data, embodied in various unjustified criticisms to the application of detrending filters. Despite this, and given the frequency domain characteristics of filtered macroeconomic integrated series, we acknowledge that the choice of a particular detrending filter is far from being a neutral task.

JEL Classification: E32; C22

Keywords: Unit roots, Band-pass filters, Pseudo-spectrum

1   Introduction

The failure of classical Fourier analysis in the presence of unit roots has led to widespread misunderstandings regarding the effects of applying linear filters to integrated time series. While some authors completely disregard spectral theory under non-stationarity, others resort to the

so-called pseudo-spectrum without providing a rigorous justification for its use. If we want to fully understand the consequences of applying transformations (linear filters) to integrated time series, we should be able to decide which view, if any, is correct. Hopefully, we will be able to address an apparently simple question. What are the consequences, in the frequency domain, of differencing a random walk? Or the consequences of applying trend-extraction filters, or more specifically band-pass filters, to integrated time series? These questions are recurrent in the context of business cycle analysis. What are we isolating when applying the Hodrick-Prescott (HP) filter (see Hodrick and Prescott, 1997), the Butterworth filters proposed by Gomez (2001) the Baxter and King filter (see Baxter and King, 1999) or the Christiano and Fitzgerald filter (see Christiano and Fitzgerald, 2003) to integrated time series? Are the properties of the resulting cycle components just artefacts? Can we give a rigorous description of its properties or a rigorous interpretation of the effects of such transformations? As we shall review, these questions have a trivial answer if the original series are covariance-stationary.

Our sole contribution is assembling some results in the literature as well as our own, and use them to show that the usual interpretation of the effects of filtering, summarised by the transfer function of the filter, carries over if the input series is integrated. This occurs essentially because, although the second moments of an integrated process are time-dependent (or infinite depending on the specification of initial conditions), the distribution of the (infinite) variance (that is, the spectrum) over frequencies is not time-dependent. This phenomenon occurs with various infinite variance processes. The standard representation theorems do not hold, since the variance is not a constant and the processes exhibit persistent time dependence, but the spectrum exhibits time invariance. Understanding this apparent paradox is the purpose of this paper. We will then be able to discuss various criticisms to the application of filters in integrated macroeconomic time series. We will see that most of them are unwarranted. Nonetheless, we will show that even though a rigorous interpretation of the effects of filtering is possible, choosing or designing a specific filter is far from being a neutral task.

The remainder of the paper is organised as follows: In section 2, the usual spectrum of an integrated series is defined. A justification for its use is provided and the bridge with the spectral representation theorem for stationary processes is made. In section 3 we discuss why the previous results are important to understand the consequences of filtering. We discuss various criticisms in the literature, related to the application of band-pass filters to integrated time series and argue that most are unjustified. Section 4 gives a note of caution, on why choosing a specific filter can be a sensitive exercise. Section 5 concludes.

2   The Spectrum

2.1   Stationary case

It is well known that a covariance-stationary time series can be decomposed into an infinite weighted sum of periodic orthogonal components. This is summarised in the spectral representation theorem. Let $\{x_t\}$ be a covariance-stationary sequence with mean zero and define the spectrum of $\{x_t\}$ as:

$$S_x(\omega) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} e^{-i\omega k} \gamma(k), \quad -\pi \leq \omega \leq \pi \tag{1}$$

where $i^2 = -1$, ω denotes the frequency measured in radians and $\gamma(k)$ is the autocovariance function of $\{x_t\}$ at lag k. It is well known (see, e.g., Brockwell and Davis, 1991) that there exists a right-continuous orthogonal increment process $\{z(\omega), -\pi \leq \omega \leq \pi\}$ such that:

$$i)\ E[(z(\omega) - z(-\pi))(z(\omega) - z(-\pi))^*] = F_X(\omega) = \int_{-\pi}^{\omega} S_x(\omega) d\omega, \quad -\pi \leq \omega \leq \pi$$

where $*$ denotes the complex conjugate and:

$$ii)\ x_t = \int_{-\pi}^{\pi} e^{i\omega t} dz(\omega)\ \ a.e.$$

So, x_t can be decomposed into an infinite weighted sum of orthogonal fluctuations, each with frequency ω. $S_x(\omega)$ can be interpreted as the decomposition of the variance of x_t in terms of these fluctuations. $S_x(\omega)$ contains the same information as the second order moments characterised by $\gamma(k)$, $k = 0, \pm 1, \pm 2, ...$. $S_x(\omega)$ and $\gamma(k)$ form a pair of Fourier transforms in that:

$$\gamma(k) = \int_{-\pi}^{\pi} e^{i\omega k} dF_x(\omega) = \int_{-\pi}^{\pi} e^{i\omega k} S_x(\omega) d\omega, k = 0, \pm 1, \pm 2, ... \tag{2}$$

If we apply a time-invariant linear filter $h(L) = \sum_{j=-\infty}^{\infty} h_j L^j$ where $L^j x_t = x_{t-j}$ and such that $\sum_{j=-\infty}^{\infty} |h_j| < \infty$ to the sequence $\{x_t\}$ we obtain a filtered sequence $y_t = \sum_{k=-\infty}^{\infty} h_j x_{t-j}$. It is easy to

verify that the spectrum of $\{y_t\}$ is given by:

$$S_y(\omega) = |h(e^{-i\omega})|^2 S_x(\omega) \tag{3}$$

where $h(e^{-i\omega})$ is denoted as the transfer function of the filter. The interpretation of a time series in the frequency domain, and the analysis of the consequences of filtering, is straightforward given the spectral representation theorem. However, the conditions for it to hold are rather restrictive, the assumption of covariance stationarity being the crucial one. How can we interpret the consequences of filtering an integrated time series? Can we extend somehow the relation in (3)?

Although it is straightforward to *know* the consequences of filtering an integrated time series (i.e., to know exactly what the spectrum of the filtered series is once it becomes stationary), the link with the original integrated series is hard to make if no reasonable definition of spectrum of an integrated series is available. We will argue that a more general interpretation of the spectrum is available, that is, we argue that the spectral representation theorem is a restrictive way of interpreting a time series in the frequency domain. Obviously, if we extend the definition of spectrum in order to encompass integrated series, the spectrum of a covariance stationary series should still be consistent with the spectral representation theorem. That is, under both definitions the spectrum of a covariance stationary process coincides.

2.2 Integrated case

Spectral analysis of integrated time series is a difficult task given the failure of classical Fourier analysis in that case. There have been several approaches to the spectral analysis of non-stationary processes. Hatanaka and Suzuki (1967) build a spectral theory of non-stationary processes by focusing on finite subsets of the sequence $\{x_t\}$ that have finite second moments. The defined pseudo-spectrum is time-varying and closely related to the evolutionary spectra of Priestley (1981). Another approach, widely followed, is to define the power distribution of an integrated series as the limit of the spectrum of a stationary process when the smallest autoregressive roots converge to 1 (e.g., Harvey, 1993; Den Haan and Sumner, 2004; Young, Pedregal and Tych, 1999). The pseudo-spectrum of a general ARIMA process is defined as (see Harvey, 1993):

$$S_x(\omega) = \frac{\sigma_\varepsilon^2}{2\pi} \frac{|\phi^{-1}(e^{-i\omega})|^2 |\theta(e^{-i\omega})|^2}{|1 - e^{-i\omega}|^{2s}} = \frac{\sigma_\varepsilon^2}{2\pi} \frac{|\psi(e^{-i\omega})|^2}{|1 - e^{-i\omega}|^{2s}}, \omega \neq 0 \tag{4}$$

where x_t satisfies:

$$\phi(L)(1 - L)^s x_t = \theta(L)\varepsilon_t, \ \forall t$$

σ_ε^2 is the variance of the white-noise innovations ε_t, we assume the roots of $\phi(L)$ lie outside the unit circle and are different from those of $\theta(L)$, $\psi(L) = \phi(L)^{-1}\theta(L)$ and $s \geqslant 0$ is the order of integration of the series. This limit is a time-invariant continuous function at all frequencies except at those associated with autoregressive roots with unit modulus. At those frequencies the pseudo-spectrum is not well-defined, there is a pole, a singularity associated with the unit roots[1]. The spectrum of an integrated series is thus a functional form extension of the spectrum of stationary processes. Within this approach, the analysis of the consequences of filtering can be made, but the results hold by definition, as we shall review in section 3. No rigorous justification for the use of the pseudo-spectrum is presented. It is assumed that this function represents indeed a distribution of variance. Without further results it is merely an ad-hoc extension of the spectrum of a stationary series. If the spectral representation theorem does not hold, alternative definitions and the exploration of its properties should be done.

But recently, in an important paper, Bujosa, Bujosa and García-Ferrer (2002) extend the classical spectral analysis by developing an extended Fourier transform to the field of fractions of polynomials. A pseudo-autocovariance generating function is defined to account for the presence of unit roots and the corresponding extended Fourier transform is defined as the pseudo-spectrum. Within this approach, which leads to exactly the same functional forms as the limit of the stationary spectrum in (4), the pseudo-spectrum collapses to the standard spectrum when no non-stationary roots are present, since the extended Fourier transform is just the classical Fourier transform in that case. This extended Fourier transform is just the ratio of the Fourier transform of $\psi(L)$ to that of the non-invertible polynomial $(1 - L)^s$, associated with the unit roots. This leads to (4).

This complements and extends an earlier definition of the spectrum of non-stationary processes, provided by Solo (1992). Solo (1992) has shown that certain continuous-time, stationary increment processes, possess many of the frequency domain properties of stationary processes. Crucially, although their variance is infinite or time-varying, they have a time-invariant spectrum, defined there as the limit of the expected periodogram. This definition is motivated by the following observation, that we adapt to the discrete-time case. If $\{x_t\}$ is a stationary time series with mean μ and the autocovariance function $\gamma(.)$ is absolutely summable, it can be shown (see,

---

[1]Since we assumed the roots of $\phi(L)$ lie outside the unit circle, we are only considering the existence of a pole at zero frequency. This assumption can straightforwardly be relaxed in order to include singularities at frequencies other than zero, e.g., due to non-stationary seasonal components.

e.g., Brockwell and Davis, 1991, p.343) that:

$$E[I_{T,x}(0) - T\mu^2] \to 2\pi S_x(0) \text{ as } T \to \infty \tag{5}$$

$$E[I_{T,x}(\omega)] \to 2\pi S_x(\omega) \text{ as } T \to \infty, \omega \neq 0$$

where $S_x(\omega)$ is the spectrum of x_t and $I_{T,x}(\omega)$ is the periodogram for a sample of size T. That is, when T grows the periodogram converges to the distribution of variance as revealed by the spectral representation theorem. Solo (1992) argues that the result in (5) is a less restrictive inversion relation than that implied by (1), hence justifying the limit of the expected periodogram as a more general definition of spectrum. It would be comforting if the definition of Solo (1992) lead to functional forms as in (4). Unfortunately it does not. Valle e Azevedo (2007) provides an exact limit for the expected periodogram when the (discrete-time) time series contains one unit root ($s = 1$), under general conditions[2,3]. The limit (divided by 2π) is the following, assuming as above that x_t follows $(1 - L)x_t = \psi(L)\varepsilon_t, \forall t$:

$$S_x^*(\omega) = \frac{\sigma_\varepsilon^2}{2\pi} \frac{(|\psi(e^{-i\omega})|^2 + |\psi(1)|^2)}{|1 - e^{-i\omega}|^2}, \omega \neq 0 \tag{6}$$

which differs from (4) due to the term $|\psi(1)|^2$ in the numerator. Using this definition of spectrum instead of that in (4) would in general lead to different conclusions regarding the effects of applying linear filters to integrated processes. But these differences would not be dramatic after a simple normalisation, given the fact that the inverse of $|1 - e^{-i\omega}|^2$ dominates the behaviour of both $S_x(\omega)$ and $S_x^*(\omega)$ at frequencies close to the pole located at zero frequency. Consider the simplest case, when $\{x_t\}$ verifies $x_t - x_{t-1} = \varepsilon_t, \forall t$ where $\{\varepsilon_t\}$ is a white noise sequence such that $E[\varepsilon_t] = 0$ and $Var[\varepsilon_t] = \sigma_\varepsilon^2$. We have, since $\psi(e^{-i\omega}) = \psi(1) = 1$:

$$S_x^*(\omega) = \frac{\sigma_\varepsilon^2}{\pi|1 - e^{-i\omega}|^2}, \omega \neq 0$$

which shows that the pseudo-spectrum, defined as in (6), is just proportional to the inverse of the Fourier transform of the differencing operator $(1 - L)$. However, if we apply the first difference filter to $\{x_t\}$ the spectrum of $(1 - L)x_t = \varepsilon_t$ is given by $\sigma_\varepsilon^2/2\pi$. To perfectly maintain the relation

---

[2]Put $(1 - L)x_t = \psi(L)\varepsilon_t = \sum\limits_{j=-\infty}^{\infty} \psi_j L^j$. The condition $\sum\limits_{j=-\infty}^{\infty} |\psi_j||j|^\alpha < \infty$ for some (small) $\alpha > 0$ is imposed to guarantee this result.

[3]The limit of the (appropriately normalised) expected periodogram when the time series process contains more than one unit root is still unexplored.

$S_\varepsilon(\omega) = |1 - e^{-i\omega}|^2 S_x(\omega)$ as in (3) we would need to define the pseudo-spectrum of x_t as:

$$S_x(\omega) = \frac{\sigma_\varepsilon^2}{2\pi|1 - e^{-i\omega}|^2}, \omega \neq 0$$

which is exactly what we obtain in (4) and seems a neutral normalisation of the (non-integrable) power distribution (if it were defined by (6)) of x_t. In this case, the first difference filter maintains the usual interpretation, summarised by the function $|1 - e^{-i\omega}|^2$. It attenuates low frequencies and amplifies high frequencies, thus producing a "noisier" output series. The stationary case interpretation would also still be valid for any linear filter containing at least one unit root. But in general this normalisation is only neutral (or leads to values of the distribution close to those in (4)) at frequencies close to zero, which nonetheless dominate clearly the behaviour of an integrated time series. Formally, it can be shown (see Hurvich and Ray, 1995) that for a unit root process:

$$E[\frac{1}{2\pi} I_{T,x}(\omega_j)/S_x(\omega_j)] \to 2 \text{ as } T \to \infty, \ \omega_j = 2\pi j/T \tag{7}$$

where $S_x(\omega_j)$ is defined as in (4) and regarded as the power distribution of x_t, as is typical in the long-memory literature (see also, e.g., Velasco, 1999). It should be noted that j is held fixed, whereas the result in (6) is valid for any fixed $\omega \neq 0$. It is easy to reconcile the two results. Heuristically, once T grows, ω_j approaches 0 and hence $|\psi(e^{-i\omega_j})|^2$ approaches $|\psi(1)|^2$. Therefore $\frac{1}{2\pi} I_{T,x}(\omega_j)$ approaches $2S_x(\omega_j)$, with $S_x(\omega_j)$ defined as in (4). In the stationary case the limit in (7) is just 1.

In short, the rigorous definition of pseudo-spectrum due to Bujosa, Bujosa and García-Ferrer (2002) is not completely consistent with that of Solo (1992) for processes containing one unit root, due to the "slight" breakdown of the result in (5) in that case. Nonetheless, the unreasonableness of disregarding any spectral characterisation of unit root processes is clear.

3 Interpreting filtered integrated time series

3.1 Consequences of filtering

The lack of a rigorous definition of spectrum of an integrated time series is not crucial to understand its properties. But if we want to interpret the frequency domain effects of applying linear filters to such series, this definition is essential. Otherwise, important questions would be left unanswered. What would we be isolating when applying band-pass filters to integrated

time series? What would be the link with the original time series fluctuations? These questions have been puzzling for a long time due to the insufficiencies of spectral theory in the presence of unit roots. This has led to numerous misinterpretations of the role of filtering in the behaviour of integrated time series, as we shall discuss below. Detractors of the use of commonly used filters often cite the fact that the spectral representation theorem does not hold in presence of integrated series. Those in support of their use fail to fully justify the frequency domain interpretation given to an integrated process. Given that we have at hand such interpretation, embodied in a definition of spectrum that encompasses integrated processes, we can proceed to discuss the effects of filtering an integrated time series.

It is straightforward to verify that if we apply an absolutely summable linear filter $h(L)$ to an integrated time series x_t, and this series becomes stationary, the spectrum of the filtered series y_t is given by:

$$S_y(\omega) = |h(e^{-i\omega})|^2 S_x(\omega) \tag{8}$$

where $S_x(\omega)$ is the pseudo-spectrum, defined as in (4), of the integrated series x_t. That is, an extension of the relation stated in (3) holds. This relation would always hold regardless of the validity of the pseudo-spectrum as a frequency domain description of x_t. This occurs because once the unit roots in the filter polynomial $h(L)$ cancel those in the auto-regressive polynomial of x_t, $\phi(L)(1-L)^s$, what we get is a transformed series given by $y_t = h_*(L)\phi^{-1}(L)\theta(L)\varepsilon_t$, where $h_*(L) = h(L)/(1-L)^s$. The spectrum of this series is obviously $|h(e^{-i\omega})|^2 S_x(\omega)$, even for $\omega = 0$. But this relation, without any further justification for the use of the pseudo-spectrum, is irrelevant for the analysis of the consequences of filtering. However, we can now state that the function $|h(e^{-i\omega})|^2$ has exactly the same interpretation as in the covariance stationary case. It represents the way fluctuations of x_t with frequency ω are amplified or mitigated by the filter.

3.2   Band-pass filters applied to integrated series

In the remainder of the paper we will focus in the analysis of band-pass filters, filters designed to isolate specific fluctuations in the original time series. We give special attention to the HP filter and to finite sample approximations of ideal band-pass filters, given the emphasis put in these filters in the literature. We define an ideal band-pass filter for the frequencies band $]\omega_l, \omega_h[\subseteq [0, \pi]^4$, as the filter that would perfectly eliminate the fluctuations with frequencies

---

[4]We restrict hereafter the analysis to the interval $[0, \pi]$, due to the symmetry of the spectrum around $\omega = 0$.

outside this band and preserve without distortion the remaining fluctuations. The transfer function of an ideal band-pass filter $B(L)$ is thus given by $B(e^{-i\omega}) = 1$ for $\omega \in]\omega_l, \omega_h[$ and 0 otherwise. The ideal filter can be expressed as $B(L) = \sum\limits_{j=-\infty}^{\infty} B_j L^j$. The filter weights are well known and given by:

$$B_o = \frac{\omega_h - \omega_l}{\pi}, \qquad B_j = \frac{\sin[\omega_h j] - \sin[\omega_l j]}{\pi j}, |j| \geq 1 \qquad (9)$$

It can be shown that $B(L)$ removes up to two unit roots, which means that $B(L)$ can be factored as $B(L) = (1 - L)^2 B_1(L) = (1 - L)B_*(L)$ where $B_*(1) = 0$. Furthermore, we can easily conclude that by applying an ideal band-pass filter to the integrated time series x_t (with order of integration $s = 1$), no fluctuations associated with frequencies outside the band of interest $]\omega_l, \omega_h[\subseteq [0, \pi]$ will be present in the filtered series. That is, for $\omega \neq 0$ and $\omega \in]0, \pi]/]\omega_l, \omega_h[$ we have $S_y(\omega) = 0$ by (8) since $B(e^{-i\omega}) = 0$. This is true regardless of the interpretation given to the pseudo-spectrum since (8) holds by definition. Also, we can conclude that $S_y(0) = 0$, even if the pseudo-spectrum is not well defined for $\omega = 0$. This occurs because $y_t = B_*(L)(1 - L)x_t = B_*(L)z_t$. Now, z_t is stationary and its spectrum well defined for all ω. Since $B_*(1) = 0$ the result follows. Finally, for $\omega \in]\omega_l, \omega_h[$ and since $B(e^{-i\omega}) = 1$, we can conclude by (8) that $S_y(\omega) = S_x(\omega)$. In short, the ideal band pass filter retains its usual interpretation in the presence of one unit root. Further unit roots would be needed in the filter if we wanted to isolate some range of frequencies in a time series with order of integration greater than 1 and enforce $S_y(0) = 0$.

Since the ideal filter in (9) is not applicable in practice, approximations have been developed. Baxter and King (1999) construct symmetric moving averages (BK filters) aimed at approximating an ideal band-pass filter. Christiano and Fitzgerald (2003) developed an optimal (in the mean squared error sense) band-pass filter (CF filter) that can be applied to every observation in the sample (since symmetry in the filter weights is not imposed). The HP filter (see Hodrick and Prescott, 1997) has long been used to eliminate low frequencies in the data. The fact that this filter was indeed a high-pass filter (a filter that eliminates only low frequencies and retains without distortion high frequencies) was pointed by King and Rebelo (1993). Using the HP filter, it is easy to construct a band-pass filter, by applying successively an HP high-pass filter and the complementary of another HP high-pass filter (a low-pass filter). The choice of the smoothing parameters can also be reconciled with the choice of a band of frequencies (see Pedersen, 2001). Butterworth filters, which can be seen as generalisations of the HP filter, provide better approximations (at least in infinite samples) to the ideal filter, the filter that would perfectly isolate the desired frequencies (see Gomez, 2001). The BK and CF filters are explicitly designed to

account for the removal of unit roots while the HP and Butterworth filters fortunately possess that property. This means, together with the interpretation given to the pseudo-spectrum, that they can in fact be seen as band-pass filters when applied to integrated series (with sufficiently small order of integration)[5].

3.3 Discussion of criticisms to the use of band-pass filters

The comprehension of the consequences of filtering a time series (either stationary or integrated) is crucial to this analysis. Many criticisms to the HP filter and to other band-pass filters forget such consequences. For instance, differences between second moments of filtered and unfiltered series cannot be taken as evidence of a weak performance of a filter as do King and Rebelo (1993), Guay and St-Amant (1997) and Canova (1998). The differences are a consequence of filtering, since the spectrum of the filtered series changes and the relation in (2) holds. The differences would arise even if an ideal band-pass filter were used. The analyst interested in the fluctuations approximately isolated by the filter should not be concerned with these consequences. As he should not be concerned if the filter does not isolate the cycle component in an unobserved-components model with a stochastic trend, as Murray (2003) shows. This type of criticism seems redundant except if there are arguments against the usefulness of looking at a portion of the spectrum in each particular application.

A more involved type of criticism concerns the alleged spurious ciclicity induced by band-pass filters when applied to persistent time series. Pedersen (2001) has fully addressed this issue in the case when the original series is persistent but still covariance-stationary. However, he does not resort to the definition of pseudo-spectrum, essential to rigorously extend the argument to integrated series. But let us review this type of criticism. Harvey and Jaeger (1993), Cogley and Nason (1995) and Guay and St-Amant (1997) criticise the HP filter for inducing business cycle periodicity in integrated or near-integrated time series, those with the "typical spectral shape" of Granger (1966). Figure 1 shows an example of this spectral shape for a covariace-stationary series. The spectrum decreases monotonically towards π, with the power clearly concentrated in the low frequencies. A peak in the spectrum of HP filtered series, which is absent in the spectrum of the original series, is taken as evidence of Slutzky-effect (or spurious ciclicity). This peak occurs because the low frequencies clearly dominate the spectrum of a persistent time series. Once we apply a band-pass filter, and thus eliminate low frequencies, a peak in the

---

[5]If some of the unit roots from the input series were not removed, the filtered series would not consist only of fluctuations with the desired frequencies since the spectrum of the filtered series would still not be well defined at zero frequency.

spectrum of the filtered series will be detected near the lowest frequency of the band of interest. Pedersen (2001) shows that these critics rest on an "inadequate definition of the Slutzky-effect – a definition which has the unfortunate consequence that even an ideal high-pass filter induces a Slutzky effect". In fact, neither an ideal band-pass filter nor the commonly used band-pass filters amplify or weaken certain frequencies so as to produce a peak in the spectrum of the filtered series. They only isolate certain fluctuations that *are* in the data. If we define, as would be reasonable, the Slutzky-effect as a cycle in the transfer function of a filter, then the HP filter and other band-pass filters do not induce a Slutzky effect. There is clearly no cycle in the transfer function of the HP filter (see figure 2).

Crucially, the previous argument carries over to the analysis of integrated series, since we can interpret the effects of filtering integrated time series in the same way as we do for stationary processes. Many authors completely neglect this fact and analyse the consequences of filtering by first "transforming" the process into a stationary form. This completely hides or distorts relevant characteristics of the original process. E.g., Harvey and Jaeger (1993) and Cogley and Nason (1995) show that there is a peak in the power transfer function of a subcomponent of the HP filter. This subcomponent is isolated to interpret the effects of the filter in $I(1)$ time series. Let the transfer function of the HP cyclical filter be decomposed as:

$$C(\omega) = (1 - e^{-i\omega}).C_1(\omega), \ \ C_1(\omega) = \frac{\lambda e^{i2\omega}(1 - e^{-i\omega})^3}{\lambda e^{i2\omega}(1 - e^{-i\omega})^4 + 1}$$

where λ is the value of the smoothing parameter. According to this decomposition, applying the HP filter to an $I(1)$ time series is equivalent to: First, filter the non-stationary time series with the first difference filter and then filter the remaining stationary component with the asymmetric filter determined by $C_1(\omega)$. The modulus of the transfer function of this second filter is plotted in figure 3. It has a clear peak at business cycle frequencies that Harvey and Jaeger (1993) classify as "a classical example of the Yule-Slutzky-effect". The definition of spectrum for an integrated process tells us that we must look to the total transfer function $C(\omega)$, not only to the subcomponent $C_1(\omega)$. That is, we cannot overlook the consequences of the first-difference filter $(1 - L)$. This introduces a complete discontinuity in the analysis of the effects of filtering once the unit root case is considered. We believe this challenges even casual observation[6].

---

[6]Rigour apart, a band-pass filtered AR(1) with coefficient 0.99 is indistinguishable from a band-pass filtered pure random walk with the same innovations. The same can be said of the spectrum of the filtered series, which can be precisely determined regardless of the reasonableness of the pseudo-spectrum as a frequency domain description of an integrated process.

We can thus state that the effects of applying a band-pass filter to an integrated series are similar to those of applying the band-pass filter to a highly persistent (but stationary) time series. The spectrum of the filtered series will have a clear peak at the lowest frequency of the band of interest, given the typical shape of the pseudo-spectrum of an integrated time series (see figure 4).

4   Choosing a specific filter is not neutral

As we have seen, it is almost impossible to identify a dominant stationary component in a typical macroeconomic time series simply by looking to the pseudo-spectrum. There is not, in most cases, a clear peak in the pseudo-spectrum. It is definitely possible to have a (usually slight) peak if, e.g., the first differences of the series can be described by an AR process with complex roots. Also, if the true data generating process of the integrated series has a cycle component like the ones in Harvey and Jaeger (1993), there will be a small peak in the pseudo-spectrum at intermediate frequencies. All this means that (regular) business cycle fluctuations are hardly identified through the pseudo-spectrum. Sargent (1987) and more recently Pagan (1999), point this fact. If this was not the case we could comfortably define, in a purely statistical sense, business cycle fluctuations as those fluctuations with frequecies around the peak in the spectrum of a set of macroeconomic time series.

In view of the above, choosing to isolate a specific band of fluctuations in a macroeconomic time series (e.g., the famous [6, 32] quarters band in quarterly data) or choosing a particular smoothing parameter in the case of the HP filter is essentially an arbitrary exercise, with potentially important consequences. Suppose for instance that a central bank uses a Taylor rule in determining interest rates, and resorts to a band-pass filter in order to estimate the output gap. The bigger the upper period in the band of interest, the more volatile will be the output gap series, since more power is being assigned to the cycle component.

Furthermore, the pseudo-spectrum is so clearly increasing as we approach the zero frequency that a cut determined by a band-pass filter will deliver a very clear peak near the lowest frequency of the band of interest, regardless of the existence of intermediate peaks. We argued before that this is not a distortion of the power distribution of fluctuations that *are* in the data. But it is nonetheless an important property of the filtered series.

Should we be concerned with the above facts if we know that more or less regular fluctuations affect output and other economic variables? To make things worse, similar spectral patterns

would arise in the filtered series if the original series were a pure random walk. By applying the HP filter or another band-pass filter to a random walk process, we are again defining a trend and allowing specific frequencies to remain in the cycle component. But does it make sense to filter a random walk? Band-pass filtering is certainly "a sensible way to look at the data" (Kaiser and Maravall, 1999), just as is ideal band-pass filtering. There is no "truth" in any isolated cycle component. Prior information must be added through the specification of business cycle fluctuations. An important element of definition must therefore be assumed. As we have seen before, the peak in the spectrum of filtered series, which can be interpreted as the period of the cycle component, is mostly determined by the lowest frequency of the band of interest (or by the value of the smoothing parameter λ in the case of the HP filter). Given this important element of definition we share the view of Kaiser and Maravall (1999): "The analyst should first decide the length of the period around which he wishes to measure economic activity." The variance of the cycle component will therefore be mostly justified by fluctuations around the critical length. A similar decision is embodied in the now standard definition of business cycle fluctuations as "fluctuations with a specified range of periodicities" (Baxter and King, 1999)[7].

5   Conclusions

In this paper, as opposed to common practice in the literature, we have resorted to a rigorous definition of spectrum of an integrated time series in order to characterise the implications of appying linear filters to such series. We show that the fact that the spectral representation theorem does not hold is not fundamental to interpret rigorously the effects of filtering integrated data. A major conclusion is that we can indeed interpret filtered integrated data in the same way as we do for stationary processes, as has been suggested by other authors. However, given the frequency domain characteristics of typical macroeconomic integrated series, we have acknowledged that the choice of a particular detrending filter is far from being a neutral task.

References

[1] Baxter, M. and King, R. (1999). Measuring business cycles: approximate band-pass filters for economic time series. *Review of Economics and Statistics*, 81:575–93.

---

[7]E.g., the [6, 32] quarters band also implies a peak in the spectrum of the filtered series at frequency $2\pi/32$, that is, the variance of the cycle component will be concentrated around that frequency. However, noisy or high frequency components with period less than 6 quarters are also disregarded.

[2] Bujosa, A., Bujosa, M. and García-Ferrer (2002). A Note on Pseudo-spectra and Pseudo-Covariance Generating Functions of ARMA processes. mimeo

[3] Brockwell, P.J. and Davis, R. (1991). *Time Series: Theory and methods* (2nd ed.). Springer

[4] Canova, F. (1998). Detrending and Business Cycle Facts. *Journal of Monetary Economics*, 41:475-512.

[5] Cogley, T. and Nason, J. (1995). Effects of the Hodrick-Prescott Filter on Trend and Difference stationary time series. *Journal of Economic Dynamics and Control*, 19(1):253-278.

[6] Christiano, L. and Fitzgerald, T. (2003). The band-pass filter. *International Economic Review*, 44:435–65.

[7] Gomez, V. (2001). The use of Butterworth filters for trend and cycle estimation in economic time series. *Journal of Business and Economic Statistics*,19:365–73.

[8] Granger, C.W.J. (1966).The Typical Spectral Shape of an Economic Variable, *Econometrica*, 34(1):150-161

[9] Guay, A. and St Amant, P. (1997). Do the Hodrick-Prescott and Baxter-King Filters Provide a Good Approximation of Business Cycles? . Bank of Canada, Working paper 97-53

[10] Den Haan, W.J., and Sumner, S.W. (2004). The Comovements between Real Activity and Prices in the G7. *European Economic Review*, 48:1333-1347

[11] Harvey, A.C. (1993). *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press

[12] Harvey, A. C. and Jaeger, A. (1993). Detrending, Stylized Facts, and the Business cycle. *Journal of Applied Econometrics*, 8:231-47

[13] Hatanaka, M., and M. Suzuki (1967). A Theory of the Pseudospectrum and its Application to Nonstationary Dynamic Econometric Models, in M. Shubik (ed.), *Essays in Mathematical Economics in Honor of Oskar Morgenstern*, Princeton, Princeton University Press.

[14] Hodrick, R. J. and Prescott, E. C. (1997). Postwar U.S. business cycles: an empirical investigation. *Journal of Money, Credit and Banking*, 29:1-16.

[15] Hurvich, C.M. and Ray, B.K., (1995). Estimation of the memory parameter for nonstationary or noninvertible fractionally integrated processes. *Journal of Time Series Analysis*, 16:17-42.

[16] Kaiser, R. and Maravall, A. (1999). Estimation of the Business cycle: A Modified Hodrick-Prescott Filter, Banco de España -Servicio de Estudios, Working Paper 9912

[17] King, R. and Rebelo, S. (1993). Low Frequency Filtering and Real Business Cycles. *Journal of Economic Dynamics & Control*, 17:207-231.

[18] Murray, C. J. (2003). Cyclical Properties of Baxter-King Filtered Time Series. *Review of Economics and Statistics*, 85: 472-476.

[19] Pagan, A. (1999). The Getting of Macroeconomic Wisdom, mimeo, University of Melbourne.

[20] Pedersen, T. M.(2001). The Hodrick–Prescott filter, the Slutzky effect, and the distortionary effect of filters, *Journal of Economic Dynamics and Control*, 25:1081–1101.

[21] Priestley, M.B. (1981), Spectral Analysis of Time Series, Academic Press, London

[22] Sargent, Thomas J. (1987). *Macroeconomic Theory -2nd Ed.*, Academic Press

[23] Solo, V.(1992). Intrinsic random functions and the paradox of 1/ f noise. *SIAM Journal of Applied Mathematics,* 52(1):270-291

[24] Velasco, C., (1999). Non-stationary log-periodogram regression. *Journal of Econometrics*, 91:325-371.

[25] Valle e Azevedo, J. (2007). Exact Limit of the Expected Periodogram in the Unit Root case, mimeo, Stanford University

[26] Young, P. C., Pedregal, D. and Tych, W. (1999). Dynamic Harmonic Regression. *Journal of Forecasting*, 18:369-394.
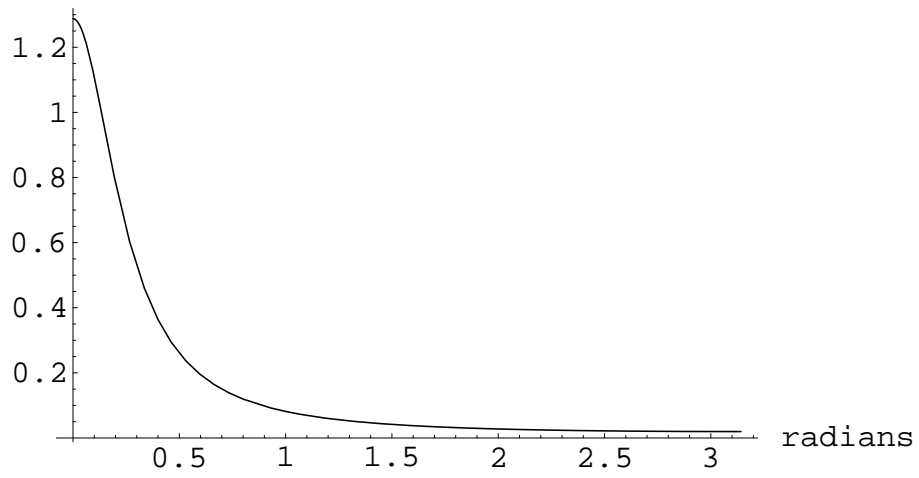
Figure 1: Spectrum of a time series with the typical "Granger" shape.
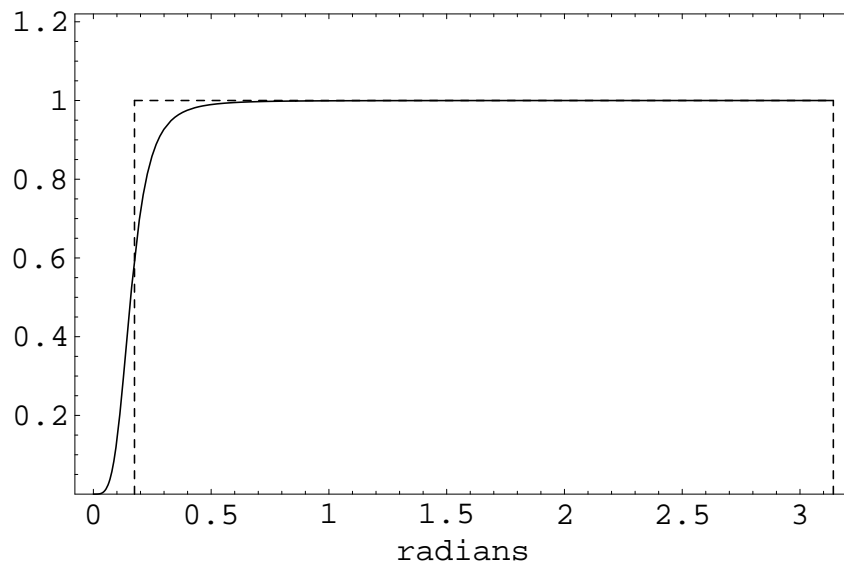
Figure 2: Transfer function of the HP cyclical filter $C(\omega)$.

16
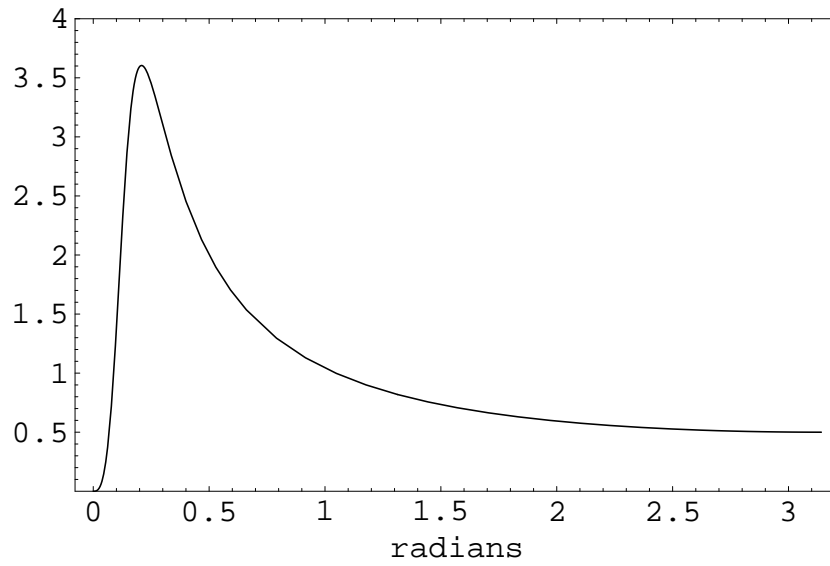
Figure 3: Modulus of the transfer function of the subcomponent $C_1(\omega)$ of the HP cyclical filter
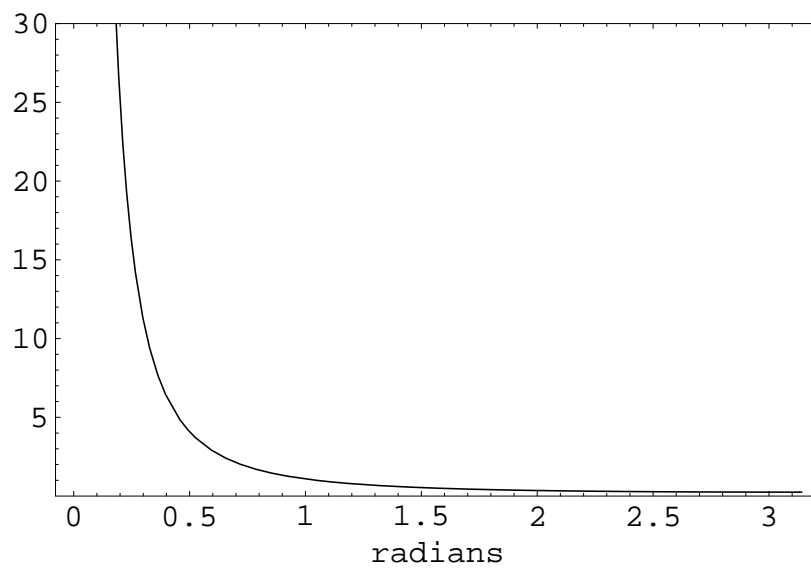
Figure 4: Typical pseudo-spectrum of an integrated macroeconomic time series.

WORKING PAPERS

2000

1/00 UNEMPLOYMENT DURATION: COMPETING AND DEFECTIVE RISKS
— *John T. Addison, Pedro Portugal*

2/00 THE ESTIMATION OF RISK PREMIUM IMPLICIT IN OIL PRICES
— *Jorge Barros Luís*

3/00 EVALUATING CORE INFLATION INDICATORS
— *Carlos Robalo Marques, Pedro Duarte Neves, Luís Morais Sarmento*

4/00 LABOR MARKETS AND KALEIDOSCOPIC COMPARATIVE ADVANTAGE
— *Daniel A. Traça*

5/00 WHY SHOULD CENTRAL BANKS AVOID THE USE OF THE UNDERLYING INFLATION INDICATOR?
— *Carlos Robalo Marques, Pedro Duarte Neves, Afonso Gonçalves da Silva*

6/00 USING THE ASYMMETRIC TRIMMED MEAN AS A CORE INFLATION INDICATOR
— *Carlos Robalo Marques, João Machado Mota*

2001

1/01 THE SURVIVAL OF NEW DOMESTIC AND FOREIGN OWNED FIRMS
— *José Mata, Pedro Portugal*

2/01 GAPS AND TRIANGLES
— *Bernardino Adão, Isabel Correia, Pedro Teles*

3/01 A NEW REPRESENTATION FOR THE FOREIGN CURRENCY RISK PREMIUM
— *Bernardino Adão, Fátima Silva*

4/01 ENTRY MISTAKES WITH STRATEGIC PRICING
— *Bernardino Adão*

5/01 FINANCING IN THE EUROSYSTEM: FIXED VERSUS VARIABLE RATE TENDERS
— *Margarida Catalão-Lopes*

6/01 AGGREGATION, PERSISTENCE AND VOLATILITY IN A MACROMODEL
— *Karim Abadir, Gabriel Talmain*

7/01 SOME FACTS ABOUT THE CYCLICAL CONVERGENCE IN THE EURO ZONE
— *Frederico Belo*

8/01 TENURE, BUSINESS CYCLE AND THE WAGE-SETTING PROCESS
— *Leandro Arozamena, Mário Centeno*

9/01 USING THE FIRST PRINCIPAL COMPONENT AS A CORE INFLATION INDICATOR
— *José Ferreira Machado, Carlos Robalo Marques, Pedro Duarte Neves, Afonso Gonçalves da Silva*

10/01 IDENTIFICATION WITH AVERAGED DATA AND IMPLICATIONS FOR HEDONIC REGRESSION STUDIES
— *José A.F. Machado, João M.C. Santos Silva*