Decomposition of the European GDP based on Singular Spectrum Analysis

Leon, Costas

28 July 2015

Decomposition of the European GDP based on Singular Spectrum Analysis[1]

Costas Leon[2]

Abstract
In this paper, the Singular Spectrum Analysis (SSA), a relatively new tool originated in natural sciences, for orthogonal decomposition of time series, is presented and applied in the European real, seasonally unadjusted quarterly GDP for the period 1995 - 2010. SSA is suitable for short and noisy time series, properties that characterize many macroeconomic time series. In this paper, I decompose the GDP in trend, cycle, seasonals and noise components. There are significant similarities but also some differences between the SSA-based filter and the other well-known macroeconomic filters. These differences are shown here by means of correlation matrices and spectral measures. Although SSA is a method that only very recently has been introduced in macroeconomics, its use in the natural sciences for more than three decades, makes it a serious candidate for tackling macroeconomic issues such as filtering, denoising and smoothing.

Keywords: Macroeconomics, economic fluctuations, business cycle, dynamical systems, spectral methods, singular spectrum analysis.

1. Introduction

The macroeconomics of business cycle requires a meaningful decomposition in signal and noise components. A very large volume of econometrics literature analyzes time series in the time domain whereas only a comparatively minor part of econometric methods and techniques, collectively called spectral methods, is conducted in the frequency domain. Given the "cyclical" properties of the business cycle component, a "natural" way of the analysis of the business cycle lies in the frequency domain because parameters of interest such as amplitude, phase, frequency and co-movements are indispensable parts of the frequency domain. In this context, Singular Spectrum Analysis (SSA) is a method, within the unobserved components structural time series analysis, that provides a unifying framework of decomposition of a time series into signals of interest, such as trend, cycle and seasonal variation, and noise. SSA does not provide directly the required parameters of interest (e.g. frequency or phase) but the decomposed signals can further be processed by usual spectral methods which will, finally, provide estimates of the required spectral parameters.

SSA is a relatively new method of time series analysis and combines elements of various scientific areas such as, among others, statistics and probability theory, dynamical systems and signal processing. It is based on the spectral decomposition of time series (Karhunen, 1946, and Loève, 1978) and on the Mañé (1981) and Takens (1981) embedding theorem. Applications of SSA typically include hydrology, geophysics, climatology, biology, physics, and very recently macroeconomics and financial economics. SSA is a model-free and data-adaptive method and depends only on one parameter, the window length (the embedding dimension). It gives remarkable results especially in short and noisy time series. The SSA was developed independently in the 1980's in the USA and UK under the name SSA and in the 1990's in Russia (St. Petersburg and Moscow) under the name Caterpillar-SSA.

SSA exists both in univariate and multivariate versions with several variations (eg. Basic SSA, sequential SSA, etc). Since this is a relatively new method, one expects more methodological advances and practical applications in the near future. The central concept of SSA is the partitioning of a vector space into orthogonal subspaces that have, within the context at hand, some meaningful interpretation, i.e. a decomposition of a time series in signal (such as trend, cycle, and seasonal variation) and noise.

In this paper, I present the application of the SSA method in the decomposition of the European GDP into trend, cycle, seasonal variation and noise. The cycle signals are statistically compared to those obtained from the well-known Hodrick - Prescott high-pass filter, the Christiano - Fitzgerald asymmetric frequency filter and the Baxter - King filter. The similarities and differences among the SSA and the other filters are presented. The comparison is made by means of correlation coefficients and spectral densities obtained by the maximum entropy method (Burg, 1975).

The paper is organized as follows. The next Section 2 exhibits the basic methodological aspects of the SSA. In Section 3, SSA is applied to the European real GDP which is decomposed into a trend, cycle, seasonal component and noise. Also, on the basis of this decomposition, the seasonally adjusted GDP is constructed. Further, the similarity of the cyclical component obtained from the SSA with those obtained from other filters is also presented. Finally, Section 4 concludes the paper.

---

[1] A first version of this paper had been submitted and presented in the Annual Meeting of the Swiss Society of Economics and Statistics, University of Zurich, 12-13 April 2012. The present paper, in July 2015, is a slightly modified version, focusing on some minor language and syntax issues of the first version.

[2] Independent researcher. Email: costas_leon@yahoo.com.

2. Methodological Aspects of the SSA

The SSA can be understood as the singular value decomposition of a matrix consisting of delayed vectors (the time series under consideration embedded into some integer dimension), called trajectory matrix, shown to be a Hankel matrix, and the regroupping of the obtained orthonormal vectors into sets of meaningful, to the application at hand, components (orthogonal subspaces) through the process of diagonal averaging. These sets in the macroeconomic analysis of the business cycle are the the trend, the cycle, the seasonals and the noise. This section describes the SSA algorithm used for the extraction of the signals. SSA comprises of two main stages: decomposition and reconstruction. In turn, the decomposition stage comprises of two steps: embedding and singular value decomposition. The reconstruction stage comprises of two steps, too: grouping and diagonal averaging (or Hankelization). The presentation below follows Beneki *et al.* (2011), Golyandina *et al.* (2001, Chapters 1 and 2), Hassani and Thomakos (*2010)* and Hassani and Zhigljavsky (2008).

2.1 Stage 1: Decomposition, Step 1: Embedding

Let $N > 2$. Consider a one-dimensional real-valued time series $F = (f_0, f_1, \ldots, f_{N\text{-}1})$ of length *N*, a positive integer *L* (*window length*) such that $1 < L < N$ and a mapping of the original series into a sequence of *L*-dimensional lagged vectors $\{X_i\}_{i=1}^{K}$, $K = N - L + 1$, by the formula:

$$X_i = (f_{i-1}, \ldots\ldots, f_{i+L\text{-}2})^T, \quad 1 \le i \le K.$$

The Hankel matrix **X**=$[X_1 ::: X_K]$ of size $L \times K$ is called the *L-trajectory matrix* (or simply, the *trajectory matrix*) of the series *F*. In linear algebra, a Hankel matrix is a matrix where all the elements along the diagonal *i+j=const.* are equal. In other words, the trajectory matrix is

$$\mathbf{X} = (x_{ij})_{i,j=1}^{L,K} = \begin{bmatrix} f_0 & f_1 & f_2 & \cdots & f_{K-1} \\ f_1 & f_2 & f_3 & \cdots & f_K \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ f_{L-1} & f_L & f_{L+1} & \cdots & f_{N-1} \end{bmatrix}$$

Note that if *N* and *L* are fixed, then there is a one-to–one mapping between the Hankel matrices of size $L \times K$ and the time series of length *N*.

2.2 Stage 1: Decomposition, Step 2: Singular Value Decomposition (SVD)

The SSA is based on a particular transformation known in linear algebra as singular value decomposition (SVD). The SVD of the trajectory matrix **X**=$[X_1 ::: X_K]$ is a decomposition of **X** in the form **X**=$\sum_{i=1}^{d} \sqrt{\lambda_i}\, U_i\, V_i^T$, where $V_i =$ **X**$^T U_i / \sqrt{\lambda_i}$, $d = max(i,\ such\ that\ \lambda_i > 0) = rank$ **X**, $\lambda_1, \ldots, \lambda_L$ are the eigenvalues of the $L \times L$ matrix **S** = **XX**^{T}, taken in the decreasing order of magnitude $(\lambda_1 \ge \lambda_2 \ldots \ge \lambda_L \ge 0)$ and $U_1, \ldots, U_L$ are the eigenvectors of the matrix **S** corresponding to these eigenvalues**.**

If we define **X**$_i$= $\sqrt{\lambda_i} U_i V_i^T$ $(i = 1, \ldots, d)$, then the SVD of the trajectory matrix can be written as a sum of rank-one orthogonal matrices:

$$\mathbf{X} = \mathbf{X}_1 + \ldots + \mathbf{X}_d \qquad (1)$$

where U_i are the orthonormal eigenvectors of **S** = **XX**^{T} (in the SSA terminology they are called empirical orthogonal functions) and V_i (in SSA terminology they are called principal components) can be regarded as the eigenvectors of the matrix **X**^{T}**X**. The collection $(\sqrt{\lambda_i}, U_i, V_i)$ is called the *i*-th eigentriple of the matrix **X**, $\sqrt{\lambda_i}$ $(i = 1, \ldots, d)$ are the singular values of the matrix **X** and U_i, V_i are the left and right singular vectors of **X**, respectively.

SVD is attractive because it ensures optimality. Among all the matrices $\mathbf{X}^{(r)}$ of rank $r<d$, the matrix $\sum_{i=1}^{r} X_i$ provides the best approximation to the trajectory matrix **X**, so that $\|\mathbf{X}-\mathbf{X}^{(\mathbf{r})}\|$ is minimum. Note that $\|X\|^2 = \sum_{i=1}^{d} \lambda_i$ and $\|X_i\|^2 = \lambda_i$ for $i = 1,\dots,d$. Thus, we can consider the ratio $\lambda_i / \sum_{i=1}^{d} \lambda_i$ as the contribution of the matrix $\mathbf{X}_i$ in the expansion (1) to the whole trajectory matrix **X**. Consequently, $\sum_{i=1}^{r} \lambda_i / \sum_{i=1}^{d} \lambda_i$, the sum of the first *r* ratios, is the contribution of the optimal approximation of the trajectory matrix by the matrices of rank *r*.

2.3 Stage 2: Reconstruction, Step 1: Grouping

The grouping step corresponds to splitting the elementary matrices $\mathbf{X}_i$ into several groups and summing the matrices within each group. Let $I = \{i_1,\dots,\mathrm{i_p}\}$ be a group of indices. Then the matrix $\mathbf{X}_I$ corresponding to the group *I* is defined as $\mathbf{X}_I = \mathbf{X}_{i1} + \dots + \mathbf{X}_{ip}$. These matrices are computed for $I = I_1, \dots, I_m$ and the expansion (1) leads to the decomposition

$$\mathbf{X} = \mathbf{X}_{I1} + \dots + \mathbf{X}_{Im} \tag{2}$$

The procedure of choosing the sets $I_1, \dots, I_m$ is called the eigentriple grouping.

2.4 Stage 2: Reconstruction, Step 2: Diagonal Averaging

The last step is, in a sense, opposite to the first step and transforms each matrix of the grouped decomposition (2) into a system of new (reconstructed) series of length *N*. This procedure is the so-called Hankelization or diagonal averaging. If z_{ij} stands for an element of a matrix **Z**, then the *k*-th term of the resulting time series is obtained by averaging z_{ij} over all *i, j* such that $i + j = k + 2$. The result of the Hankelization of a matrix **Z** is the Hankel matrix *H***Z**, which is the trajectory matrix corresponding to the time series obtained as a result of the diagonal averaging (see the formal description in Golyandina *et al., op.cit.*). Note that the Hankelization is an optimal procedure in the sense that the matrix *H***Z** is the closest to **Z** (with respect to the matrix norm) among all Hankel matrices of the corresponding size (Golyandina *et al., op.cit.*). In its turn, the Hankel matrix *H***Z** defines the series uniquely by relating the values in the diagonals to the values in the series. Diagonal averaging applied to a matrix $\mathbf{X}_{Ik}$ produces the series $F^{(k)} = \left(f_0^{(k)},\dots,f_{N-1}^{(k)}\right)$. Hence, the original series is decomposed into the sum of *m* series:

$$f_n = \sum_{k=1}^{m} \widetilde{f}_n^{(k)} \ .$$

3. Application: Decomposition of the European GDP

As an illustration, I apply the method to the European GDP (16 countries, at constant prices of the year 2005, for the period 1995:I - 2010:IV, i.e. 64 non-seasonally adjusted quarterly observations). Figure 1 displays the natural logarithm (designated as log) of the European GDP.

Figure 1: The log of the European GDP

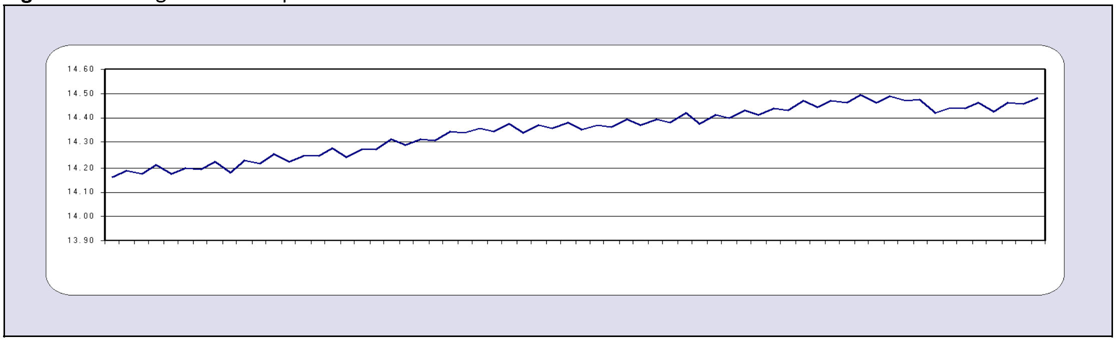

3.1 The choice of the window length

As mentioned in the Introduction, SSA depends on one parameter, called the window length L. L should be large enough in order to capture sufficiently the dynamics of the time series but not larger than $N/2$. Further, if any periodic components are known to be present in the time series, then L should be proportional to the longest wave length. Under the assumption of a local trend model (the signal in the state equation is hidden in noise in the observational equation), there is a relationship between the window length and the signal to noise ratio (S/N). As S/N approaches zero, then the window length should converge to $N/2$, and as S/N approaches infinity, then the window length should converge to $\sqrt{N}$ (see Hassani and Thomakos, 2010). Spectral density or other relevant measures showing peaks in the dominant frequencies or exogenous information may assist in the determination of periodic components.

3.2 The determination of the signal components

Once the window length has been determined, some other information might be valuable in order to proceed to the grouping step of the reconstruction stage. For example, in practice, a periodic component is identified by having two eigentriples with singular values close to each other (the exception is at frequency 0.5 which displays one eigentriple with saw-tooth singular vector). Therefore, a plot of the singular values $\sqrt{\lambda_i}$ against an index *i=1,…,L*, gives important information in the sense that, through this visual aid, one can easily discern the high and the low singular values. Since each singular value $\sqrt{\lambda_i}$ expresses the significance of the corresponding $\mathbf{X}_i$ to the total trajectory matrix $\mathbf{X}$, high singular values imply significance of the corresponding eigentriple as a determinant of the variance of the time series.

3.3 Decomposition and reconstruction

Based on the above considerations and having no further information about the S/N ratio (which could be measured, for example, by the Kalman filter), I choose the window length to be 32. This choice corresponds both to the minimum possible S/N ratio and to the maximum "typical" cyclical component of the business cycle (32 quarters). Figure 2 shows the 32 singular values.

Figure 2: Singular values

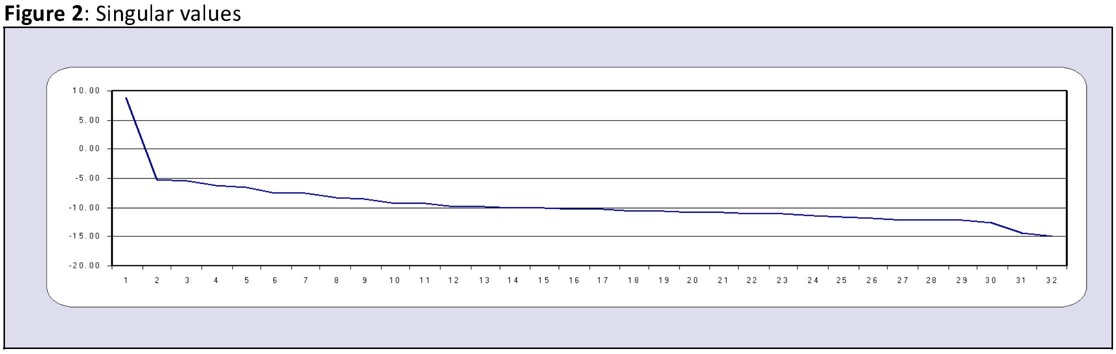

The highest singular value corresponds to the trend. Then, the following 8 singular values (2-9) decay almost linearly up to the value of (approximately) -10. They contain the cycle and the seasonal components. The last 23 singular values decay very slowly (with the exception of the last three). All these 23 singular values belong to noise.

3.3.1 The long-run trend

Based on Figure 1, the long run trend is reconstructed on the basis of the first eigenvector, Figure 3.

Figure 3: Eigenvector 1 and the reconstructed trend

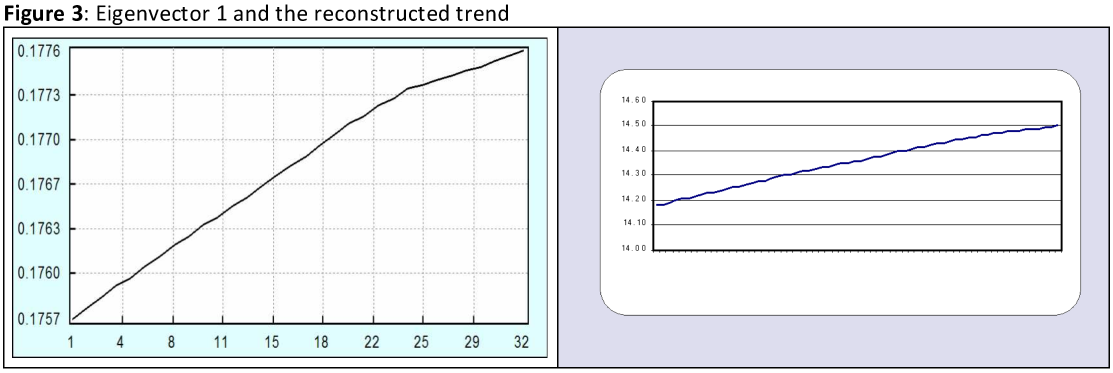

Note: The left part is the eigenvector. The shaded area is the reconstructed trend.

3.3.2 The cycle

The cycle consists of the eigentriples 3,4,5,8,9. Figure 4 shows these eigenvectors and the reconstructed cycle.

Figure 4: Eigenvectors 3,4,5,8,9 and the reconstructed cycle

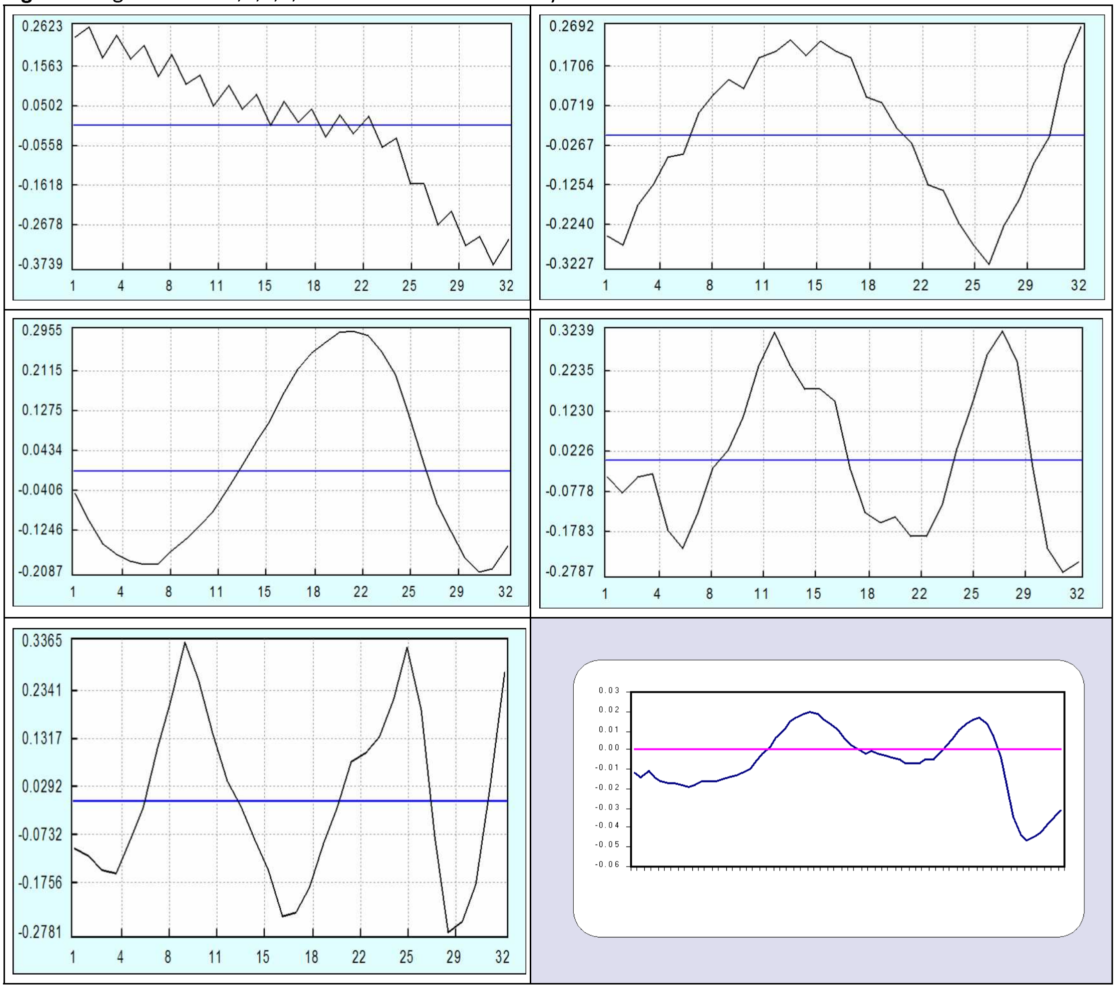

Note: All graphs, except the last one, are eigenvectors. The last graph with the shaded area is the reconstructed cycle.

3.3.3 The seasonal component

The seasonal component is reconstructed according to the eigentriples 2,6,7. Figure 5 shows the reconstructed seasonal component.

Figure 5: Eigenvectors 2,6,7 and the reconstructed seasonal component

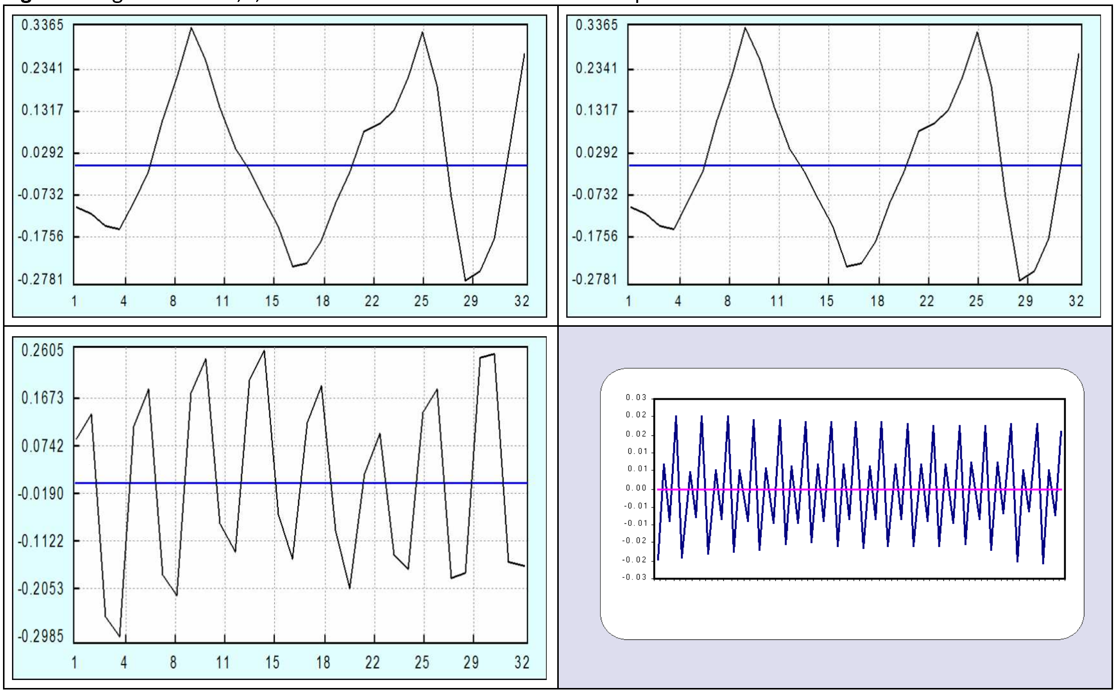

Note: All graphs, except the last one, are eigenvectors. The last graph with the shaded area is the reconstructed seasonal component.

3.3.4 The noise

The noise component is reconstructed on the basis of the remaining 23 eigentriples (10-32). Figure 6 shows the reconstructed noise.

Figure 6: Eigenvectors 10-32 and the reconstructed noise

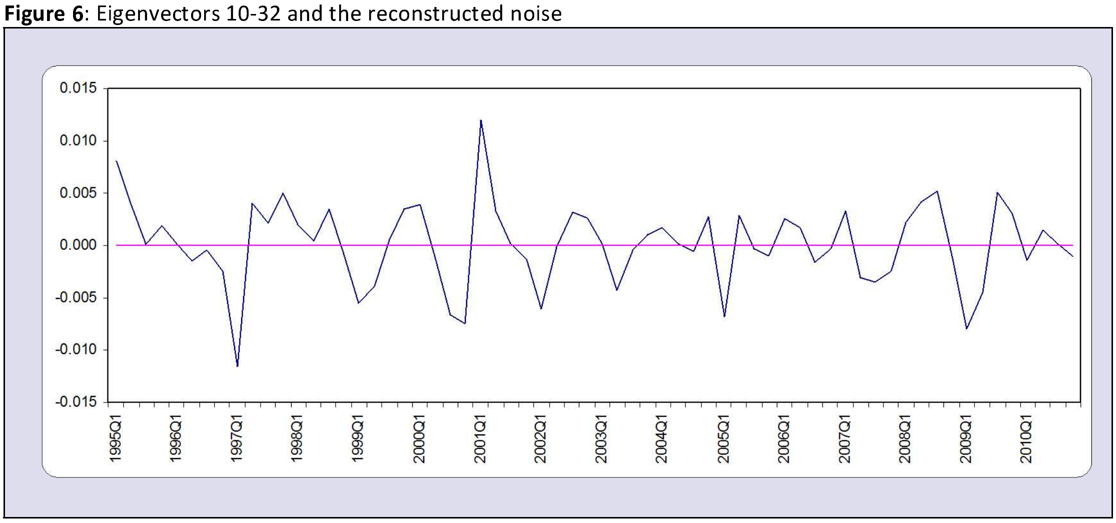

3.3.5 The reconstructed European GDP signal

Based on the decomposition of the European GDP into the three subsignals of interest (trend, cycle and seasonals), I now proceed to the reconstruction of GDP signal (i.e. signal = trend + cycle + seasonals = actual GDP - noise). The GDP signal is displayed in Figure 7.

Figure 7: The reconstructed European GDP signal

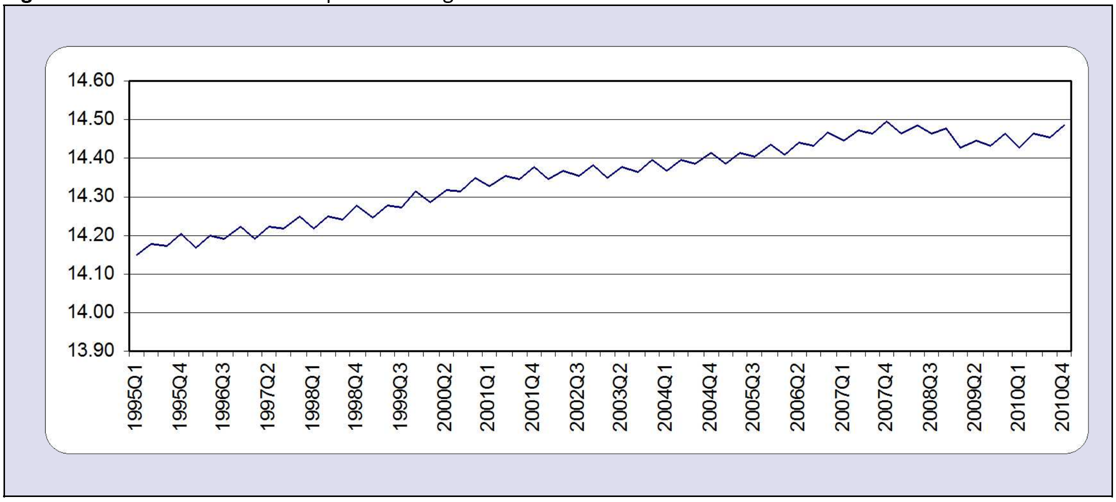

3.3.6 The seasonally adjusted GDP

Since the signal in Figure 7 contains the seasonal component, seasonal adjustment can easily be made by reconstructing a time series which contains only the trend and the cycle. This could be considered as the seasonally adjusted time series (without the noise). The seasonally adjusted GDP signal is displayed in Figure 8.

Figure 8: The seasonally adjusted GDP signal

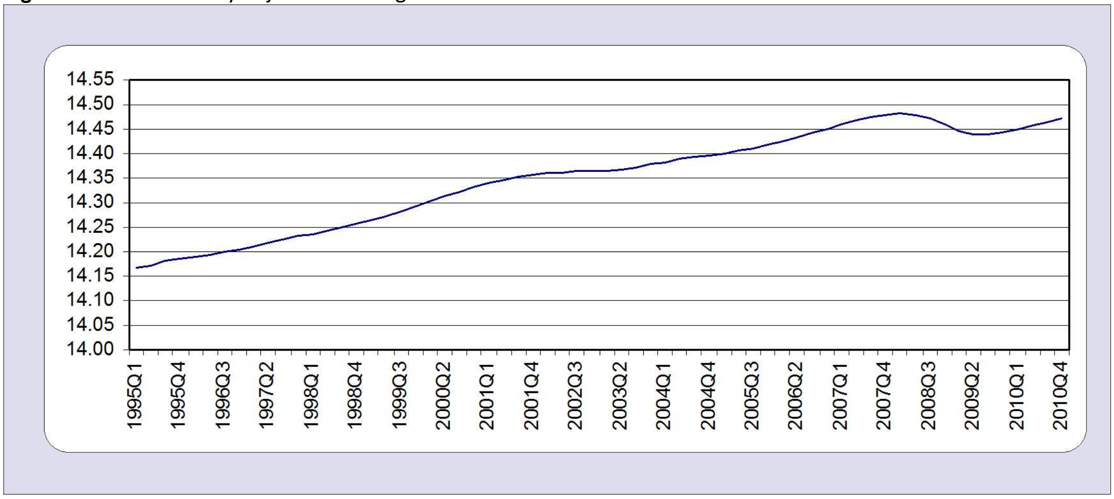

3.3.7 The orthogonality of the decomposition

The decomposition is optimal when the sub-signals and noise are statistically "separable". This implies that the correlation matrix among these sub-signals and the noise component should be diagonal, with the main diagonal being one and the off-diagonal elements being, ideally, zero. The appropriate statistic here is $t = r\sqrt{\frac{N-2}{1-r^2}}$.

Under the null hypothesis that ρ (the theoretical correlation coefficient) is zero, the t statistic, with r the sample correlation coefficient and N =64 (the sample size), is distributed as the usual t (Student) statistic with *N-2* degrees of freedom. In the above cases, the null hypothesis cannot be rejected at the conventional levels of significance (5% and 10%). Thus, the proposed decomposition is optimal in the sense that a high degree of separability has been achieved. Table 1 presents the correlation matrix among all the sub-signals and the noise.

Table 1: Correlation matrix among the sub-signals and noise

	Trend	Cycle	Seasonals	Noise
Trend	1.000	0.075	0.042	0.066
Cycle	0.075	1.000	0.027	0.024
Seasonals	0.042	0.027	1.000	0.047
Noise	0.066	0.024	0.047	1.000

3.4 Comparison to other known filters

3.4.1 Descriptive statistics and correlation coefficients

In the following, I compare the SSA cyclical signal to the cyclical signals obtained by the Hodrick - Prescott filter (HP) with the smoothing parameter equal to 1600, the Christiano - Fitzgerald (CF) asymmetric frequency filter under the assumption that the real European GDP follows a random walk and that the business cycle is between 6 and 32 quarters, and the Baxter-King (BK) frequency filter under the same assumptions about the cycle as the CF filter. It needs to be noted that, because of the truncation of the coefficients in the BK filter, 24 observations (12 in the beginning and 12 in the end of the sample) are not defined. For the HP, the CF and BK filters, the initial time series has been first seasonally adjusted by the X-12 seasonal adjustment method. Table 2 shows the minimum, the maximum, the standard deviation and the Bera - Jarque (BJ) statistic for normality, and Table 3 shows the correlation matrix among all these four filters. Figure 9 depicts these four cycles over the whole period 1995 - 2010. The normality of the cyclical components cannot be rejected in any of the filters at 5% significance level since the Bera-Jarque (BJ) statistic < 5.99, the critical value at the 5% significance level, in all four cases. On the basis of the correlation matrix in the Table 3, the SSA filter resembles more to the BK filter (cor. coef. = 0.846), then to the HP filter (cor. coef.=0.791) and then with the CF filter (cor. coef.=0.771). It is observed that differences are rather slight although the SSA filter shows deeper troughs and milder peaks in comparison to the other three filters.

Table 2: Descriptive statistics for the cyclical component among the four filters

	SSA	HP	CF	BK
Minimum	-0.047	-0.025	-0.028	-0.010
Maximum	0.019	0.029	0.028	0.029
Std. deviation	0.017	0.012	0.011	0.010
BJ statistic	0.019	0.028	0.422	4.225

Table 3: Correlation matrix among the four cyclical components

	SSA	HP	CF	BK
SSA	1.000	0.791	0.771	0.846
HP	0.791	1.000	0.931	0.933
CF	0.771	0.931	1.000	0.982
BK	0.846	0.933	0.982	1.000

Figure 9: The cyclical component obtained from the four filters

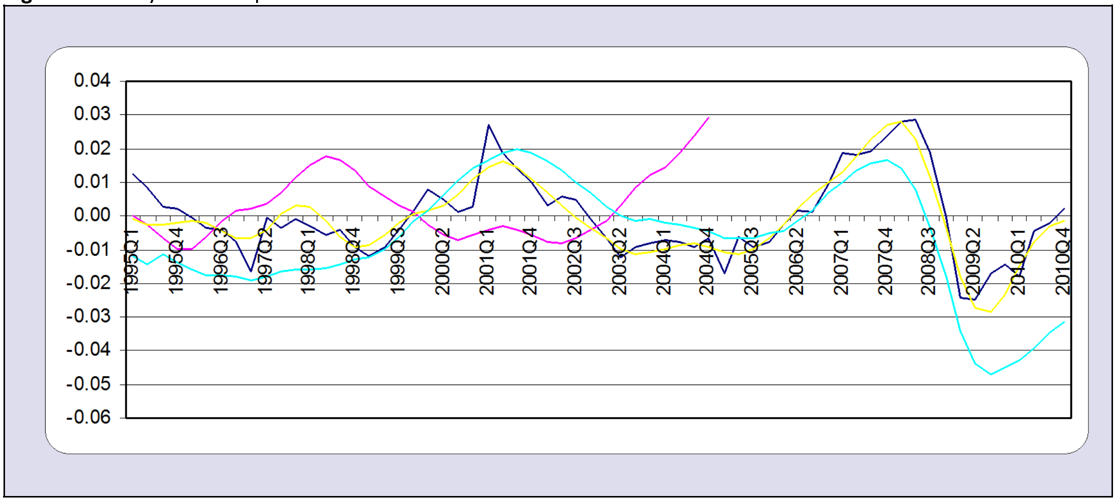

Note: SSA (light blue), CF (yellow), HP (dark blue), BK (pink).

3.4.2 Spectral measures of the filters

The frequency of the cycle can be measured in the frequency domain by means of some spectral density function. I apply the maximum entropy method (Burg, *op. cit.*) which estimates the power spectrum based on partial knowledge of the autocorrelation function of a stationary time series (here the cycle). After several experimentations with different AR structures, an AR model of order 4 has been used in the spectral density estimation. The determination of the spectral density for the BK filter did not yield reliable results and therefore it is not shown. As it is observed in Figure 10, the SSA filter has a peak in a lower frequency (0.078 cycles), in comparison to both HP (0.156 cycles) and CF filters (0.156 cycles, the same as the HP filter). This implies that the cycle obtained by the SSA filter has higher wave length in comparison to the other two filters.

Figure 10: Maximum entropy spectra for SSA-HP and SSA-CF filters

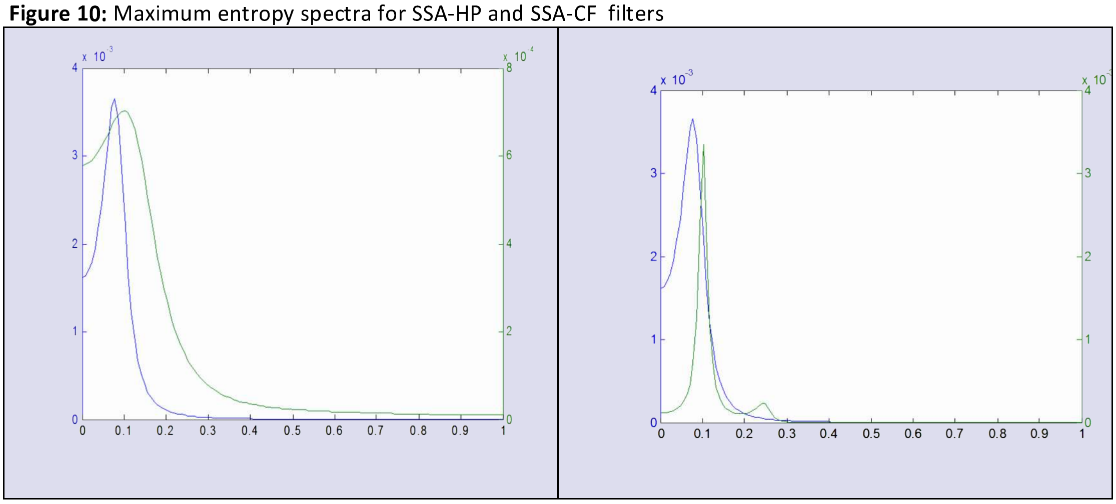

Note: Left: SSA filter (blue), HP filter (green). Right: SSA filter (blue), CF (green).

4. Conclusion

In this paper, I have presented the SSA method and its application in the decomposition of the European real GDP with quarterly data for the period 1995 - 2010. This method decomposes the initial time series vector space into a sum of signal and noise orthogonal subspaces in a way that these subspaces are meaningful in empirical macroeconomics. SSA depends only on one parameter, the window length, and can be considered as a unified framework within which the signals of interest (trend, cycle and seasonal variation) are obtained as special cases. As far as the cycle is concerned, the application of the method to the European real GDP shows similarities (strong correlation coefficient among the cycles obtained by all four filters) but also some differences to other usual filters. The SSA filter with a window length equal to 32 shows deeper troughs and milder peaks than the other three filters, as well as a longer wave length in the cyclical component of the European real GDP in comparison to HP and CF filters. Possible applications of SSA may include seasonal adjustment procedures, denoising, filtering and smoothing of short and noisy economic time series. SSA can also be used as a first stage in empirical macroeconomics in the sense that it provides a cleaner signal of the cycles in the context of business cycle synchronization. Once the signals are obtained from the SSA, they may be further processed by the usual time or frequency domain methods. Given that SSA has been very recently introduced in macroeconomics, some further issues should be addressed in a future research. For example, a more precise determination of the window length (a relation between window length and signal to noise ratio) or the robustness across different time domain or frequency domain estimators, such as a comparison of the spectral densities when they are obtained by different spectral estimators such as maximum entropy or smoothed periodogram techniques. Further, the issue of the obtained noise is still open. In all experimentations with SSA, the obtained noise component was not purely white and some signal, expressed in the form of statistically significant autocorrelations, was hidden into it.

References

Beneki, B., Eeckels, B., Leon, C. (2011), "Signal Extraction and Forecasting of the UK Tourism Income Time Series. A Singular Spectrum Approach." *Journal of Forecasting*, 9 March 2011, DOI:10:1002/for.1220.

Burg, J.P. Maximum Entropy Spectral Analysis, PhD Thesis, Stanford University, 1975.

Golyandina, N., Nekrutkin, V. and Zhigljavsky, A. Analysis of Time Series Structure: SSA and Related Techniques, Chapman and Hall/CRC, 2001.

Hassani, H., Thomakos, D. (2010). "A Review on Singular Spectrum Analysis for Economic and Financial Time Series", *Statistics and Its Interface*, Forthcoming.

Karhunen, K. Zur Spektraltheorie Stochastischer Prozesse, Ann. Acad. Sci. Fenn. Ser. A1, Math. Phys., 34, 1946.

Loève, M.: Probability Theory, Vol. II, 4th ed., Springer-Verlag, 1978.

Mañé, R., (1981), "On the dimension of the compact invariant sets of certain nonlinear maps", in Dynamical Systems and Turbulence, Eds. D. A. Rand and L. S. Young, Springer-Verlag, New York, 230–242.

Takens, F., (1981), "Detecting strange attractors in turbulence". In Dynamical Systems and Turbulence, D. A. Rand and L.-S. Young (Eds.), Springer-Verlag, New York, pp. 366–381.