

QUANTITATIVE DESCRIPTION OF U.S. BUSINESS CYCLES USING MULTIVARIATE SINGULAR SPECTRAL ANALYSIS

ANDREAS GROTH^{a,b}, MICHAEL GHIL^{a,b,c,d}, STÉPHANE HALLEGATTE^{e,f} AND
PATRICE DUMAS^{b,e}

AbstractWe apply advanced methods of spectral analysis to U.S. macroeconomic data for a 52-year interval (1954–2005). The data set includes nine variables and we study the presence of cyclic behavior and the characteristics of this cyclicity. The methods used include singular spectral analysis and measuring synchronization between the nine variables. The individual business cycles are fairly irregular, but have a mean period of 5–6 years. They share several stylized facts, including a roughly sawtooth-shaped character, with recessions shorter than expansions, and near-synchronicity of gross domestic product (GDP), employment and personal income, while change in private inventories leads GDP.

KEYWORDS: advanced spectral methods, comovements, frequency domain, stylized facts, time domain.

1. INTRODUCTION

Business cycles, their causes and characteristics have been extensively studied since the beginnings of modern economic theory. Research to characterize their regularities and stylized facts has, therewith, a long history (Burns and Mitchell (1946); Kydland and Prescott (1998)). Besides a long-term upward drift, macroeconomic time series exhibit short-term fluctuations. A number of approaches have been proposed to separate the fluctuations from the trend (e.g., Baxter and King (1999); Canova (1998)), with the Hodrick-Prescott filter being the most commonly used tool to do so (Hodrick and Prescott (1997)). Since there is no fundamental theory — and hence no generally accepted definition — of the trend, the resulting residuals have to be analyzed very critically, in order to avoid spurious results due merely to the detrending procedure (Cogley and Nason (1995)). Evidence for these residuals exhibiting a certain recurrent or cyclical structure is plentiful, although there is no definitive theory to explain the underlying mechanism of cyclicity. Moreover, business cycles are understood as comovements of the deviations from the trend in several distinct macroeconomic

groth@lmd.ens.fr (A. Groth), ghil@lmd.ens.fr (M. Ghil), hallegatte@centre-cired.fr (S. Hallegatte), dumas@centre-cired.fr (P. Dumas)

^aLaboratoire de Météorologie Dynamique (CNRS and IPSL), Ecole Normale Supérieure, Paris, France

^bEnvironmental Research & Teaching Institute, Ecole Normale Supérieure, Paris, France

^cGeosciences Department, Ecole Normale Supérieure, Paris, France

^dDepartment of Atmospheric & Oceanic Sciences and Institute of Geophysics & Planetary Physics, University of California, Los Angeles, USA

^eCentre International de Recherche sur l'Environnement et le Développement, Nogent-sur-Marne, France

^fEcole Nationale de la Météorologie, Météo France, Toulouse, France

variables (Burns and Mitchell (1946); Lucas (1977)). It is imperative, therefore, to analyze business cycle properties as a multivariate process.

The purpose of this paper is to introduce two novel methodologies for the analysis of macroeconomic time series and to use these methodologies to help understand the origin and behavior of business cycles. The first of the two involves *singular spectral analysis (SSA)* and *multivariate SSA (M-SSA)*, both of which rely on the classical Karhunen-Loève spectral decomposition of time series and random fields (Karhunen (1946); Loève (1945, 1978)).

Broomhead and King (1986a,b) proposed to use SSA and M-SSA in the context of nonlinear dynamics for the purpose of reconstructing the attractor of a system from measured time series. These authors provided an extension and a more robust application of the Mañé-Takens idea of reconstructing dynamics from a single time series (Mañé (1981); Sauer, Yorke, and Casdagli (1991); Takens (1981)). Ghil, Vautard and associates (Ghil and Vautard (1991); Vautard and Ghil (1989); Vautard, Yiou, and Ghil (1992)) noticed that SSA can be used as a time-and-frequency domain method for time series analysis — independently from attractor reconstruction and including cases in which the latter may fail.

We rely here on M-SSA for the analysis of the spectral content and common characteristics of several time series that reflect the time evolution of a single economy. M-SSA combines two useful approaches of statistical analysis: (1) it determines major directions in the system's phase space that are spanned by the multivariate time series with the help of principal component analysis (PCA); and (2) it extracts major spectral components by using data-adaptive filters. In particular, M-SSA can separate distinct spectral components in a multivariate data set of limited length and in the presence of relatively high noise levels.

M-SSA has already proven its advantages in a variety of applications, such as climate dynamics, meteorology and oceanography, as well as the biomedical sciences. Ghil, Allen, Dettinger, Ide, Kondrashov, Mann, Robertson, Saunders, Tian, Varadi, and Yiou (2002) provide an overview and a comprehensive set of references to both the theory and applications of SSA and M-SSA; free software for implementation is provided by the SSA-MTM Toolkit at <http://www.atmos.ucla.edu/tcd/ssa>. M-SSA has also shown its ability to reduce the effect of noise in order to help predict future exchange rates (Lisi and Medio (1997)). The present paper, however, goes beyond mere denoising and describes in much greater detail how to use M-SSA in order to quantify the properties of the underlying macroeconomic system.

The second methodology we introduce here is that of studying phase synchronization (Pikovsky, Rosenblum, and Kurths (2003)). Using it, we analyze comovements, extract leads and lags, and derive turning points. This phase analysis enables us also to compare cycles that differ in length and amplitude, and to combine them into an average business cycle. This average cycle allows us to describe in a quantitative fashion typical macroeconomic behavior during expansions and recessions.

The paper is organized as follows. In Section 2 we introduce the SSA and

M-SSA methodology: starting from single-channel SSA, we discuss the extension to multi-channel time series. We summarize the properties of the methodology in terms of spectral decomposition, as well as of time-domain reconstruction. In Section 3, we analyze the detailed properties of the quarterly U.S. economic data by applying the decomposition–reconstruction procedure of Section 2; we also give recipes on how to assess the quality of the reconstruction and how to draw conclusions about the underlying system from the M-SSA results.

In Section 4, we introduce the phase synchronization methodology and apply it to systematically derive stylized facts, such as leads and lags, as well as the average behavior of business cycles. A summary and discussion follow in Section 5.

2. SPECTRAL DECOMPOSITION AND TIME-DOMAIN RECONSTRUCTION

2.1. *Data and pre-processing*

We apply M-SSA to U.S. macroeconomic data from the Bureau of Economic Analysis (BEA; see <http://www.bea.gov>). The nine variables analyzed are gross domestic product (GDP), investment, employment rate, consumption, total wage, change in private inventories, price, exports, and imports; all monetary variables are in constant 2005 dollars, while the employment rate is in percent. The time series of these nine variables are available on a quarterly basis, and cover 52 years, from the first quarter of 1954 to the last of 2005.

We first remove the trend of each time series separately, by using the [Hodrick and Prescott \(1997\)](#) filter with the parameter value $\lambda = 1600$, as recommended by these authors for quarterly data. Employment is the only one of the nine variables that does not exhibit an upward trend, being bounded between 0 and 100%; still we detrend it, in order to remove periods that are much longer than 20 years and thus are not relevant to our analysis. This approach to obtaining the (raw–trend) residuals allows us to have a common basis for comparison with previous studies. As stated already in the Introduction, we are fully aware of the danger of obtaining spurious cyclical behavior due to inadequate detrending ([Cogley and Nason \(1995\)](#)), and will discuss in Section 3.1 the ability of M-SSA to justify the results obtained in the present paper.

Next, we nondimensionalize each time series by dividing the residuals by the trend, i.e. we concentrate on relative values. Each time series of relative values is the normalized to possess standard deviation one, i.e. we divide the relative values by their standard deviation. Finally, we divide the normalized time series by $(DM)^{1/2}$ — where $D = 9$ is the number of variables (“channels”) and $M = 40$ is the window width — so that the sum of the partial variances equals one (see the next two subsections). The upper panel of Fig. 1 shows the results of this pre-processing. The U.S. recessions, as defined by the National Bureau of Economic Research (NBER), are indicated by shaded vertical bars in this figure.




FIGURE 1.— Pre-processing and singular spectrum analysis (SSA) of the nine time series of U.S. macroeconomic data; raw data from the BEA, 1954–2005. Upper panel: Detrended and standardized time series. Middle panel: Reconstruction of each series separately, using the first 10 components of its single-channel SSA. Lower panel: Joint reconstruction of the entire data set, using the first 16 components of its multi-channel SSA (M-SSA); the choice of 16 components is explained in Section 3.1. Both SSA and M-SSA use a window width of $M = 40$ quarters. The shaded vertical bars in all the panels indicate NBER-defined recessions, while the vertical lines in the lower panel mark the limits of the individual cycles, as derived from the methodology presented here; see Section 4.3 and Fig. 5 there.

2.2. Singular spectral analysis (SSA)

Before introducing the multivariate extension on which we shall concentrate in the present paper, we first briefly discuss the *univariate version of SSA* and present the main properties and capabilities of this analysis method.

The classical approach to describe cyclical behavior of a single time series is to decompose it into its spectral components by some version of Fourier analysis. This approach works well for fairly long time series $\{x(t) : t = 1 \dots N\}$, with N large, and relatively low noise levels; it works less well when N is small and the noise is large, which is often the case in economic, as well as geophysical time series.

Ghil, Vautard and several associates first proposed to apply the SSA methodology to handle the problem of describing cyclical behavior in short and noisy time series, for which standard methods derived from Fourier analysis do not work well (Ghil and Vautard (1991); Vautard and Ghil (1989); Vautard, Yiou, and Ghil (1992)). This methodology provides insight into the unknown or partially known dynamics of the underlying system that has generated the time series. Applying SSA to reconstruct the entire attractor of a nonlinear dynamical system from limited data, as originally proposed by Broomhead and King (1986a), may fail, however, even in relatively simple cases (Mees, Rapp, and Jennings (1987); Vautard and Ghil (1989)).

The key idea of Vautard and Ghil (1989) and Ghil and Vautard (1991) was to only reconstruct the “skeleton of the attractor”, i.e. the most robust, albeit unstable *limit cycles* embedded in it. Following Mañé (1981) and Takens (1981), the starting point of SSA is to embed the time series $\{x(t) : t = 1 \dots N\}$ into an M -dimensional *phase space* \mathbf{X} , by using M lagged copies

$$\mathbf{x}(t) = (x(t), x(t+1), \dots, x(t+M-1)),$$

with $t = 1 \dots N - M + 1$.

The SSA procedure starts by calculating the principal directions of the attractor in this embedding phase space \mathbf{X} from the vector sequence \mathbf{x} . The first step is then to compute the auto-covariance matrix \mathbf{C} of \mathbf{x} , whose elements $c_{i,j}$ are given by

$$(1) \quad c_{i,j} = \frac{1}{N - |i-j|} \sum_{t=1}^{N-|i-j|} x(t)x(t+|i-j|).$$

Vautard and Ghil (1989) observed that the $M \times M$ matrix \mathbf{C} has Toeplitz structure with constant diagonals; that is, its entries $c_{i,j}$ depend only on the lag $|i-j|$.

The eigenvalues λ_k and eigenvectors $\boldsymbol{\rho}_k$, $k = 1 \dots M$, of \mathbf{C} are obtained by solving

$$(2) \quad \mathbf{C}\boldsymbol{\rho}_k = \lambda_k\boldsymbol{\rho}_k.$$

The eigenvectors, which are pairwise orthonormal, span a new coordinate system in the M -dimensional embedding space \mathbf{X} , and each eigenvalue λ_k indicates the variance of \mathbf{x} in the corresponding direction $\boldsymbol{\rho}_k$. This computation helps us find major directions in \mathbf{x} that carry a large variance, and therefore describe major components of the system's dynamical behavior. Usually, this spectral decomposition of \mathbf{C} determines its directions of greatest variance successively, from the largest to the smallest, subject to the condition that each new direction be orthogonal to all the preceding ones.

By convention, the eigenvalues $\{\lambda_k, k = 1 \dots M\}$ are arranged in descending order, thus representing the directions $\boldsymbol{\rho}_k$ in \mathbf{x} in order of importance, from the largest to the smallest variance. Projecting the time series \mathbf{x} onto each eigenvector $\boldsymbol{\rho}_k$ yields the corresponding principal component (PC)

$$(3) \quad A_k(t) = \sum_{j=1}^M x(t+j-1)\rho_k(j), \quad k = 1 \dots M.$$

Note that the sum above is not defined close to the end of the time series, where $N - M \leq t \leq N$. It is customary, therefore, to consider the PCs as defined for only $N - M + 1$ indices, which could start at $t = M$, rather than at $t = 1$, and end at N , or start at $t = 1$ but end at $N - M + 1$; most commonly they are plotted centered for $M/2 \leq t \leq N - M/2$, with M even (Ghil, Allen, Dettinger, Ide, Kondrashov, Mann, Robertson, Saunders, Tian, Varadi, and Yiou (2002)).

We can now reconstruct that part of the time series that is associated with a particular eigenvector by convolving the PC with the corresponding eigenvector

$$(4) \quad r_k(t) = \frac{1}{M_t} \sum_{j=L_t}^{U_t} A_k(t-j+1)\rho_k(j), \quad k = 1 \dots M.$$

The values of the triplet of integers (M_t, L_t, U_t) for the central part of the time series, $M \leq t \leq N - M + 1$, are simply $(M, 1, M)$; for either end they are given in Ghil, Allen, Dettinger, Ide, Kondrashov, Mann, Robertson, Saunders, Tian, Varadi, and Yiou (2002). Thus the *reconstructed component* (RC) $r_k(t)$ associated with the variance λ_k has a complete set of N indices, but our confidence in its values decreases as we approach either end of the time series, since fewer values of the original time series are averaged over the window M to obtain these values of the RC.

Given any subset \mathcal{K} of eigenelements $\{(\lambda_k, \boldsymbol{\rho}_k) : k \in \mathcal{K}\}$, we obtain the corresponding reconstruction $r_{\mathcal{K}}(t)$ by summing the RCs r_k over $k \in \mathcal{K}$,

$$r_{\mathcal{K}}(t) = \sum_{k \in \mathcal{K}} r_k(t).$$

Typical choices of \mathcal{K} may involve (i) $\mathcal{K} = \{t : 1 \leq t \leq S\}$, where S is the *statistical dimension* of the time series, cf. Vautard and Ghil (1989), i.e., the number of

statistically significant components, commonly referred to as the “signal” (as opposed to the “noise”); (ii) a pair of successive components $(k_0, k_0 + 1)$ for which $\lambda_{k_0} \approx \lambda_{k_0+1}$, which might capture, as we shall see, a possibly cyclic mode of behavior of the system (see Section 3.1); or (iii) $\mathcal{K} = \{k : 1 \leq k \leq M\}$, i.e. the complete reconstruction of the time series. The latter case simply corresponds to the counterpart, for SSA, of the Parseval theorem for a Fourier series (or integral) or to wavelet reconstruction. The common, less mathematical notation for the reconstructed component r_k is RC- k , and for a sum of several, consecutive RCs — from index k to index k' — it is RCs $k - k'$.

The middle panel in Fig. 1 shows the reconstruction RCs 1–10, with the first 10 components, where SSA has been applied to each of the nine economic variables separately. The common embedding dimension for each of the nine is $M = 40$, and $S \leq 10$ for each of them.

The plot of these RCs supports the impression of a much smoother time series than the original, raw time series, for each of the economic variables. Indeed, the RCs can be considered as a filtered time series, where the filtering is *data adaptive*, being given by the projection onto the eigenvectors. From the viewpoint of signal processing, (3) can be viewed as a finite-impulse response (FIR) filter (Oppenheim and Schaffer (1989)), with ρ_k being an FIR filter of length M . When the so-called “scree diagram” of variance λ_k vs. order k (see Section 3.1) has a clear break in slope at $k = S$, cf. choice (i) for \mathcal{K} above, then this filter is optimal for the given time series, in the sense that it captures the largest amount of variance for the given signal part (Ghil, Allen, Dettinger, Ide, Kondrashov, Mann, Robertson, Saunders, Tian, Varadi, and Yiou (2002); Vautard and Ghil (1989)).

Next, the PCs $A_k(t)$ obtained in (3) are time-reversed, cf. (4), and the FIR filter is run again through them. After this second filter pass, the correct chronological order is restored by reversing the filtered result $r_k(t)$ once more. This procedure is called *forward-backward filtering*, and it is known to preserve the phase relations. The filtering by projection onto each eigenvector ρ_k is thus neutral in phase and acts only on the amplitudes. Hence, each RC $r_k(t)$ and the original time series $x(t)$ are in phase; this property is essential in allowing us to study phase relations between distinct variables in business cycles. In designing an appropriate band-pass filter, Baxter and King (1999) require, in particular, that this “filter should not introduce phase shifts.” Unlike their *a priori* band-pass filter of data-independent design, SSA is data adaptive, as the M filters are simply the eigenvectors of the auto-covariance matrix.

2.3. Multivariate SSA (*M*-SSA)

M-SSA extends univariate SSA to multivariate time series. In the context of Mañé–Takens-style attractor reconstruction for nonlinear, deterministic dynamical systems, it was proposed by Broomhead and King (1986b). Once again, complete reconstruction is problematic, especially for large systems and in the

presence of substantial amounts of noise. The more modest goal of reconstructing only the robust skeleton of the attractor led [Kimoto, Ghil, and Mo \(1991\)](#), [Keppen and Ghil \(1993\)](#), and [Plaut and Vautard \(1994\)](#) to formulate the algorithm described briefly herein; see also [Ghil, Allen, Dettinger, Ide, Kondrashov, Mann, Robertson, Saunders, Tian, Varadi, and Yiou \(2002\)](#). In this section we also present the advantages of applying M-SSA to the data set at hand.

Instead of a single time series $x(t)$, we now observe multiple quantities simultaneously. Let $\mathbf{x}(t) = \{x_d(t): d = 1 \dots D, t = 1 \dots N\}$ be now a D -channel, vector time series of length N . In the generalization of (1) we consider, beside all D auto-covariances $\mathbf{C}_{d,d}$, also all cross-covariances $\mathbf{C}_{d,d'}$ to yield the *grand covariance matrix* $\tilde{\mathbf{C}}$:

$$\tilde{\mathbf{C}} = \begin{pmatrix} \mathbf{C}_{1,1} & \mathbf{C}_{1,2} & \dots & \mathbf{C}_{1,D} \\ \mathbf{C}_{2,1} & \mathbf{C}_{2,2} & \dots & \mathbf{C}_{2,D} \\ \vdots & \vdots & \mathbf{C}_{d,d'} & \vdots \\ \mathbf{C}_{D,1} & \mathbf{C}_{D,2} & \dots & \mathbf{C}_{D,D} \end{pmatrix},$$

where $\tilde{\mathbf{C}}$ is of size $DM \times DM$ and the entries of the individual matrices $\mathbf{C}_{d,d'}$ are given by

$$(c_{i,j})_{d,d'} = \frac{1}{\tilde{N}} \sum_{t=\min\{1,1+i-j\}}^{\max\{N,N+i-j\}} x_d(t)x_{d'}(t+i-j).$$

The denominator \tilde{N} depends on the range of summation, namely $\tilde{N} = \min\{1, 1+i-j\} - \max\{N, N+i-j\} + 1$.

As in the univariate case in (2) of Section 2.2, we now diagonalize the grand matrix $\tilde{\mathbf{C}}$ to yield its eigenvalues λ_k and eigenvectors $\tilde{\boldsymbol{\rho}}_k$,

$$\tilde{\mathbf{C}}\tilde{\boldsymbol{\rho}}_k = \lambda_k\tilde{\boldsymbol{\rho}}_k.$$

By solving this equation, we get DM pairs $(\lambda_k, \tilde{\boldsymbol{\rho}}_k)$, where each eigenvector $\tilde{\boldsymbol{\rho}}_k$ of length DM is composed of D consecutive segments $\boldsymbol{\rho}_k^d$, $d = 1 \dots D$, each of length M . The associated PCs are single-channel time series that are computed by projecting the multivariate time series (x_1, x_2, \dots, x_D) onto $\boldsymbol{\rho}_k^d$,

$$(5) \quad A_k(t) = \sum_{j=1}^M \sum_{d=1}^D x_d(t+j-1)\rho_k^d(j), \quad k = 1 \dots M.$$

In addition to the summation j over time as in the univariate SSA in (3), we have here a second summation with respect to the distinct channels d , representing a principal component analysis. In the geophysical applications reviewed by [Ghil, Allen, Dettinger, Ide, Kondrashov, Mann, Robertson, Saunders, Tian, Varadi, and Yiou \(2002\)](#), these channels often represent sample time series of the

same variable, e.g. temperature, at distinct spatial locations; here they represent distinct variables for the same economy.

Convolving the PCs with the eigenvectors,

$$(6) \quad r_k^d(t) = \frac{1}{M_t} \sum_{j=L_t}^{U_t} A_k(t-j+1)\rho_k^d(j), \quad k = 1 \dots M, d = 1 \dots D,$$

gives the RCs, as in the univariate case of (4). Like in single-channel SSA, this M-SSA result in (6) yields a set of M RCs, but it does so for each of the D time series $x_d(t)$. Depending on the information contained in the cross-covariances $\mathbf{C}_{d,d'}$, the RCs of different channels may or may not be correlated.

If the channels are actually independent, the submatrices $\mathbf{C}_{d,d'}$ equal zero for $d \neq d'$, and so M-SSA will yield the same results as single-channel SSA. If, however, $\mathbf{C}_{d,d'} \neq \mathbf{0}$ for at least some pairs (d, d') with $d \neq d'$, then M-SSA helps extract common spectral components from the multivariate data set, along with comovements of the channels. Past experience has shown, moreover, that M-SSA does not induce artificially more coherent behavior in the filtered data set, i.e. no spurious correlations between channels appear due to the M-SSA treatment of the data, when choosing the parameter carefully. In order to confirm that our results on coupling between channels were not affected by the use of M-SSA, we compare them to those obtained by using a null hypothesis of no coupling in Section 4.3.

To illustrate the advantages of applying M-SSA to the complete data set, with respect to separate application of SSA to each of the nine channels, we compare the results of the two approaches in Fig. 1. The overall picture of the M-SSA filtered time series shown in the lower panel exhibits a more coherent behavior of the nine economic variables. It is obvious from the comparison with the middle panel that this improves the extraction of common phase-space behavior and helps gain additional information about the underlying dynamical system.

According to the NBER definition, “a recession is a significant decline in economic activity spread across the economy, lasting more than a few months, normally visible in real GDP, real income, employment, industrial production, and wholesale-retail sales.” This definition clearly implies that comovements among several economic time series play an important role in the characteristics of business cycles.

3. SPECTRAL AND TEMPORAL CHARACTERISTICS OF BUSINESS CYCLES

3.1. Eigenvalues and oscillating pairs

A key advantage of M-SSA is its data-adaptive filtering property. As we saw in Sections 2.2 and 2.3, the filter defined by projection onto an eigenvector is not prescribed a priori, such as various band-pass filters — e.g., that of [Baxter and King \(1999\)](#) — but is determined by using the time series under study. Therefore,




FIGURE 2.— Spectrum of eigenvalues λ_k of M-SSA for the U.S. economic data in Fig. 1, using a window width of $M = 40$; scale on the ordinate is logarithmic. Only the largest 40 eigenvalues out of the total number of $DM = 9 \cdot 40 = 360$ are shown. The error bars in the figure are estimated by $\sigma_k = \lambda_k / (\hat{N}/2)^{1/2}$. The number of independent samples \hat{N} is estimated to be $\hat{N} = N / (\tau\kappa)$, where τ is the mean decorrelation time of all time series and $\kappa = 1.5$ is an empirical constant; this statistical significance test and several others are described in Ghil, Allen, Dettinger, Ide, Kondrashov, Mann, Robertson, Saunders, Tian, Varadi, and Yiou (2002) and available in the SSA-MTM Toolkit, <http://www.atmos.ucla.edu/tcd/ssa>; τ was introduced into this estimate by Unal and Ghil (1995).

it is important to understand the way such a filter adapts to the data and to control and assess the reconstruction process. We show here how to decompose the time series into the sum of its essential, dynamically relevant part and of the remaining, irrelevant fluctuations.

A first insight into the decomposition procedure is given by the eigenvalues. Roughly speaking, they indicate the importance of the corresponding eigenvectors; strictly speaking, they determine the variance of the time series in each of the mutually orthogonal directions given by the eigenvectors. The ideal picture would be that of the eigenvalues falling into two distinct sets, namely a few large eigenvalues and a lot of small ones. If such an obvious separation were to present itself, we would clearly choose the first set for reconstruction; Vautard and Ghil (1989) called the number of eigenvalues in the first set the statistical dimension S of the time series; see also Section 2.2 here. Unfortunately, in many real-world applications the separation is not that obvious.

Figure 2 depicts the 40 largest eigenvalues obtained by the M-SSA of the nine U.S. economic indicators in descending order. Although the first six eigenvalues are dominant and are followed by a small gap in the spectrum, they capture

only 60 percent of the time series' variance, and there is no obvious break in the slope at larger k , rather a smooth flattening of the tail as the eigenvalues continue to decrease. As already stated by [Vautard and Ghil \(1989\)](#), this tail can result from either red noise associated with many additional degrees of freedom or from more complex, but still deterministic dynamics.

An important feature of SSA and M-SSA is the identification of robust, possibly nonlinear oscillations that are reflected by pairs of eigenvalues whose corresponding eigenvectors are in phase quadrature, i.e. have a phase lag of one quarter period, like a sine and cosine of the same period. Such pairs are indeed the counterpart of sine-and-cosine pairs in classical Fourier analysis. The advantage of the data-adaptive character of (M-)SSA is that a purely periodic sawtooth, box-car or other markedly anharmonic function will require many sine-and-cosine pairs to reconstruct it by Fourier analysis, while one or at most two (M-)SSA pairs will suffice.

In [Fig. 2](#) the first eight eigenvalues clearly fall into pairs of nearly equal eigenvalues. Following [Plaut and Vautard \(1994\)](#), we consider them as an oscillatory pair when the (positive and negative) extrema of the cross-correlation function of the corresponding PCs, on either side of the origin (i.e. of lag zero), have absolute values that both exceed 0.6. If the oscillations is only intermittent, the cross correlations would decay to 0 and will be rejected. Applying this statistical significance criterion to our data set leads us to accept these first four pairs of eigenvalues as oscillating pairs.

The largest pair represents a 5.2-year cycle, where the period length is in agreement with the commonly accepted business cycle length of 5–6 years. To determine the period length, we calculate the cross-spectral density of the PC-pair, defined as the Fourier transform of the cross-correlation function. We have tested the robustness of the paired structure of the eigenvalues and corresponding RCs at different values of the window width $M = 30, 40, 50$ and 60, and it turns out that the leading pair always captures the 5–6-year business cycle. Moreover, for all these M -values, we always find a second pair whose period equals about one-half of 5.2 years; this pair clearly reflects the first harmonic of the dominant period. The other pairs vary somewhat in their frequency but they are still essential for the reconstruction of the smooth, statistically significant part of the time series, because they represent collectively a substantial part of its variance.

In their work about the “real facts” and monetary myths of business cycles, [Kydland and Prescott \(1998\)](#) discussed the origin of business cycles in terms of the [Slutzky \(1937\)](#) theory of random shocks. In the simplest real business cycle model, cyclicity originates exclusively from productivity shocks that can be modeled by a simple random walk. Such a random-walk model, however, can be easily rejected by the present results, because it does not exhibit robust, statistically significant oscillatory modes, as we find in our BEA data set.

3.2. Spectral characteristics

As discussed in Section 2.2, SSA can be seen as a forward-backward filtering procedure, with the eigenvectors as the data-adaptive filters. In the case of M-SSA, though, this is not exactly true, because a scalar (inner) product intervenes in Eq. (5) in the calculation of PCs, as opposed to regular multiplication of scalars in its SSA counterpart, i.e. (3).

Nevertheless, the eigenvectors are frequency-selective and extract individual spectral components of the multivariate time series and the individual RCs are band-limited in frequency. It is of interest, therefore, to see which part of the full spectrum is reconstructed from an individual RC or a subset of RCs. In the same way as each channel of the time series can be reconstructed by adding the corresponding RCs, the full spectrum of each channel can be reconstructed by adding the spectra of the corresponding RCs. By restricting our investigation to a subset of RCs, we are able to study in greater detail different aspects of the data set.

We estimate here the individual RC spectra by fitting an autoregressive (AR) process of order 16 with the Burg algorithm. Since the RCs are narrow-band time series, the order of such an AR process is typically much smaller than that required to fit the entire time series, and even smaller than the one given by the Akaike criterion as a good compromise between spectral resolution and statistical robustness; see Penland, Ghil, and Weickmann (1991).

Figure 3 shows the “stack spectra” for each of our nine economic indicators. We see that RCs 1-2 are indeed restricted to a narrow band around 5.2 years. Beside this, there are several spectral components that are successively covered by other RC pairs. By reconstructing the original time series using RCs 1–16, almost the entire spectral information below one-half year is included. When higher-frequency components are removed, as in the lower panel of Fig. 1, the remaining RCs 1–16 yield a smooth temporal behavior.

3.3. Temporal characteristics

The existence of an oscillatory pair indicates that the trajectory of the system that produced the time series is attracted by a limit cycle in phase space; in this case, the cycle has a period of 5.2 years. But this limit cycle does not explain the full dynamical behavior, as we can see from the existence of further significant eigenvalues.

In order to better understand the role of oscillatory pairs in U.S. business cycles, we study the partial variance of a subset \mathcal{K} of PCs, relative to that of all PCs, as a function of discrete time t :

$$(7) \quad V_{\mathcal{K}}(t) = \frac{\sum_{k \in \mathcal{K}} A_k(t)^2}{DM} \Bigg/ \sum_{k=1} A_k(t)^2;$$

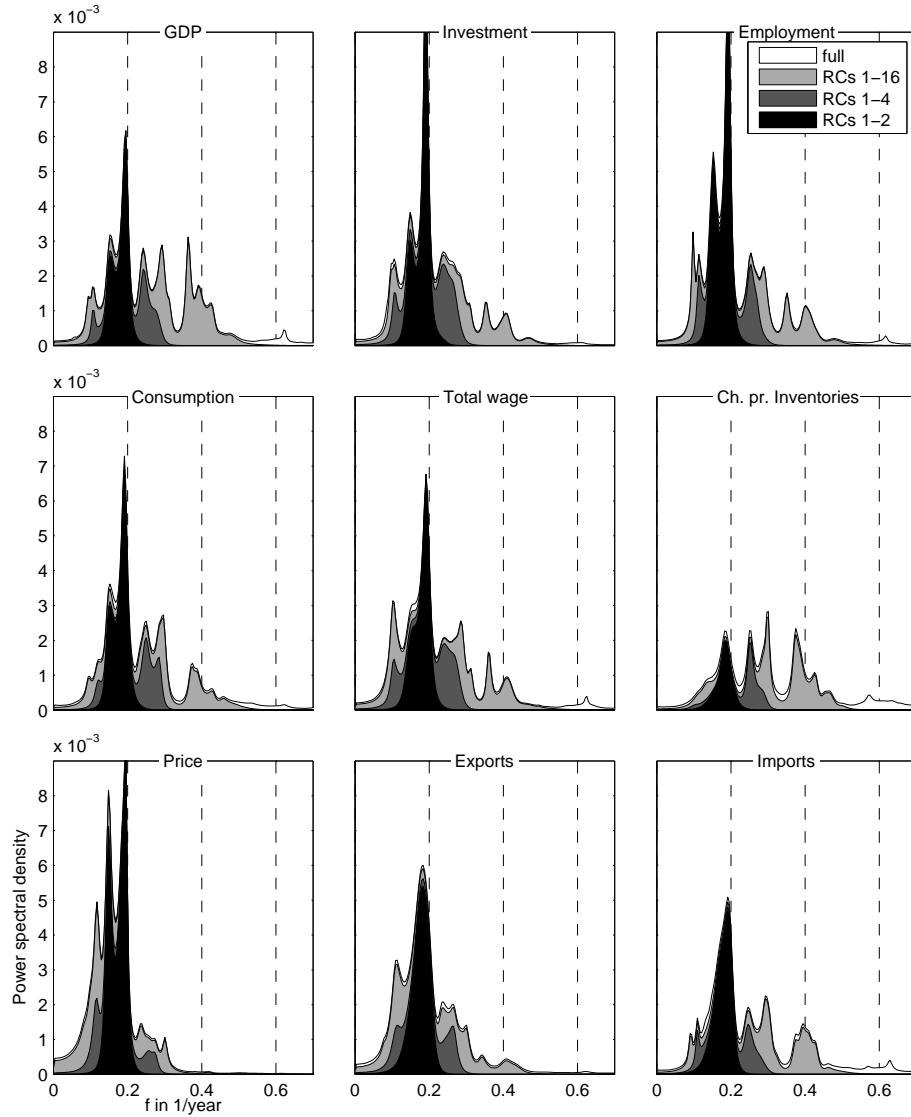


FIGURE 3.— Spectral densities of the nine U.S. macroeconomic indicators. The different shadings (see legend in upper-right corner) indicate the “stacking” — i.e., the separate contributions — of the power spectra of different subsets of reconstructed components (RCs). The 5.2-year peak captured by RCs 1-2 clearly dominates all the economic indicators.

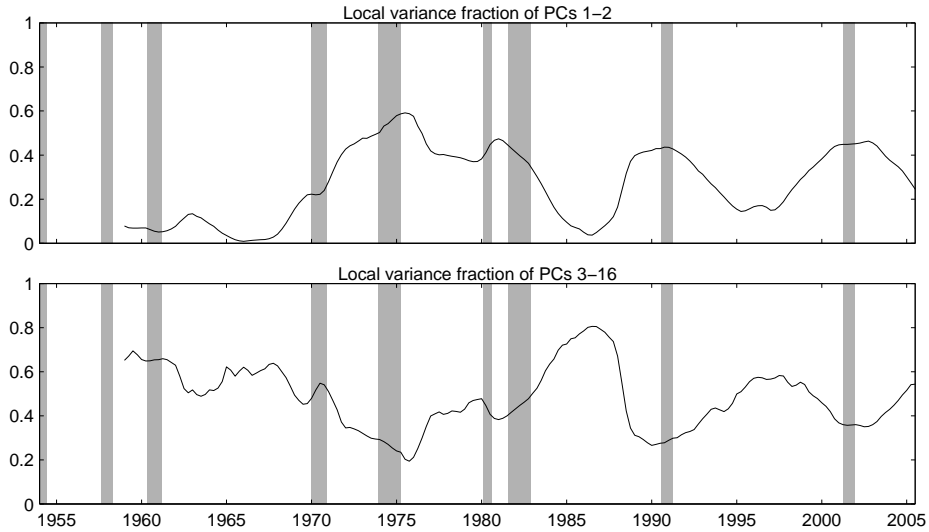


FIGURE 4.— Local variance fraction $V(t)$ for RCs 1-2 (upper panel) and RCs 3-16 (lower panel). The shaded vertical bars in both panels indicate NBER-defined recessions.

Plaut and Vautard (1994) called this index the *local variance fraction*. The PCs are used in (7), rather than the RCs, since it is the former that give us the projection of the multivariate time series onto the eigenvector pairs; as already noted in connection with (3), we plot the PCs centered, i.e. starting at $M/2$.

Figure 4 shows this index, along with the NBER dating of recessions, for PCs 1-2 (upper panel) and PCs 3-16 (lower panel). Starting after 1970, it is quite remarkable that the fraction of the 5.2-year cycle (upper panel) is high during recessions and low during expansions. During intervals of growth, the trajectory is less attracted by the limit cycle captured by PCs 1-2 and reveals more complex behavior, which is described by PCs 3-16. This result indicates that the complexity of the U.S. economy depends on the phase of the business cycle it is in, with more complex behavior during growth; it also suggests that recessions could be more easily explained than expansions.

4. COMOVEMENTS AND THE AVERAGE BUSINESS CYCLE

Now that we have extracted the key cyclical aspects of the U.S. economy with the help of M-SSA, we go into greater depth and detail, by applying the synchronization methodology. In this section, we consider comovements between the different economic indicators and analyze leads and lags between them, as well as the strength of the coupling between these distinct aspects of the econ-

omy. Moreover, we justify the statistical significance of the results by means of surrogate-data tests.

4.1. Analytical signal and phase synchronization

Since each business cycle differs from the others in amplitude and duration, we are not able to compare them by classical methods. The quantification of comovements by means of traditional cross-correlations suffers from this lack of exact periodicity; we are thus led to describe the behavior of stylized facts, such as expansion and recession, by using the *instantaneous-phase* concept. Heuristically, the phase of a cyclic phenomenon corresponds to the angle of rotation of the system's trajectory around a central point, taken as the origin in the system's phase space.

To introduce a well-defined phase for each single time series $x(t)$, we use the *analytical signal* approach. Applying the Hilbert transform $\tilde{x}(t) = \mathcal{H}\{x(t)\}$ to the time series (Hilbert (1953)), we extend the series into the complex plane as $x(t) + i\tilde{x}(t)$; the instantaneous phase of this analytical signal equals the angle $\phi(t)$ given by

$$(8) \quad \phi(t) = \arg \frac{x(t)}{\tilde{x}(t)}.$$

The extension $x(t) + i\tilde{x}(t)$ is estimated by applying a discrete Fourier transform to $x(t)$, setting all negative frequencies to zero, and transforming it back into the time domain. The mathematical properties of an analytic signal, as well as its estimation are discussed by Marple (1999). For each sine component $x(t) = \sin(t)$, the Hilbert transform returns its cosine counterpart $\tilde{x}(t) = -\cos(t)$; (8) yields the correct phase, but with the opposite sign. Sometimes one uses, therefore, $-\mathcal{H}$ as a more convenient transform; since we only need to consider phase differences in our investigation, we shall keep here \mathcal{H} as our transform, with the plus sign.

There is, however, a hitch: in general, the Hilbert transform is numerically ill posed, i.e. $\mathcal{H}\{x(t)\}$ depends sensitively on small errors in $x(t)$. Intuitively, for a broad-band signal one might get a multitude of phases that are hard to separate from each other, rather than a single phase. For an amplitude-modulated carrier wave $x(t) = u(t) \sin(2\pi ft + \phi)$, though, where the $u(t)$ modulation contains no frequency above the carrier frequency f , the Bedrosian (1963) theorem states that $\tilde{x}(t) = u(t)\mathcal{H}\{\sin(2\pi ft + \phi)\} = -u(t)\cos(2\pi ft + \phi)$. We thus still get from (8) a unique, correct phase (modulo the sign).

Since the RCs are narrow-band signals, the Bedrosian theorem justifies the application of the Hilbert transform to obtain a well-defined phase. Feliks, M.Ghil, and Robertson (2009) first suggested that a good method for obtaining narrow-band components of an arbitrary scalar or vector signal, $x(t)$ or $\mathbf{x}(t)$, is precisely to apply SSA or M-SSA, respectively, as a prefilter. We follow this suggestion here, and find that visual inspection of RCs 1–16 and of their Hilbert transforms shows indeed well-defined oscillations about the origin; it is reasonable,

therefore, to consider the instantaneous phases of the separate channels, in this data-adaptively filtered form.

Having defined the analytic-signal concept and the instantaneous phase, we now introduce here the tools of phase synchronization analysis, in order to study comovements of economic variables that characterize business cycles — independently of their exact duration and amplitude. During the last two decades, considerable attention has been paid to the interaction of deterministically chaotic oscillators; their behavior is considerably more complex than that of linear oscillators, such as the harmonic oscillator. [Rosenblum, Pikovsky, and Kurths \(1996\)](#), in particular, showed that, even for weak coupling of such oscillators, their phases become locked, while their amplitudes may remain uncorrelated. This phenomenon is referred to as *phase synchronization*.

As the coupling between the oscillators becomes stronger, it also becomes visible in the amplitude relation, and one observes so-called complete synchronization. The connection between synchronization and coupling for chaotic oscillators is now fairly well established in dynamical systems theory; it has also been widely accepted in time series analysis, cf. [Pikovsky, Rosenblum, and Kurths \(2003\)](#), who provide a comprehensive overview of the field. Still, we have to be careful in using the term of “synchronization,” which describes a property of interactions between self-sustained oscillators.

From the analysis of short and noisy time series, such as the U.S. economic indicators of interest here, we are not able to draw definitive conclusions about the underlying interaction mechanism. Indeed, the method applied here does not distinguish between phase synchronization and the more classical phase locking between some external forcing and one or more driven oscillators (e.g., [Ghil and Childress \(1987\)](#); [Winfree \(1980\)](#)) or some other simple transfer system ([Winterhalder, Schelter, Kurths, Schulze-Bonhage, and Timmer \(2006\)](#)). In the following — while using the basic tools of this methodology, enhanced by SSA prefiltering — we avoid the term “synchronization” and refer merely to lead-and-lag relations.

4.2. Comovements, leads and lags

To quantify comovements between two time series we consider their phase difference. If the cyclical behavior of the two time series is related, we expect their phases to be locked, i.e. their phase difference to be nearly constant. To the contrary, if the two series are unrelated, we expect a random distribution of their phase differences. Hence, the coupling strength between channels d and d' can be quantified by the resultant length $R_{d,d'}$ of their instantaneous phase differences $\phi_d(t) - \phi_{d'}(t)$, viewed as the exponents of the complex numbers $e^{i(\phi_d(t) - \phi_{d'}(t))}$,

$$(9) \quad R_{d,d'} = \left| \frac{1}{N} \sum_{t=1}^N e^{i(\phi_d(t) - \phi_{d'}(t))} \right|.$$

TABLE I
COMOVEMENTS IN THE U.S. ECONOMIC INDICATORS

Time series	Phase difference ^a in degrees	$\Delta\phi$ ($2\sigma_\phi$) in months	Coupling strength ^b R (R_α)
GDP	0 (± 0)	0 (± 0)	1.00 (0.46)
Investment	-19 (± 5)	-3 (± 1)	0.91 (0.48)
Employment	10 (± 8)	2 (± 1)	0.81 (0.46)
Consumption	19 (± 6)	3 (± 1)	0.87 (0.48)
Total wage	-21 (± 5)	-4 (± 1)	0.93 (0.47)
Δ (Inventories)	36 (± 8)	6 (± 1)	0.81 (0.46)
Price	-114 (± 14)	-19 (± 2)	0.51 (0.46)
Exports	-33 (± 13)	-6 (± 2)	0.57 (0.45)
Imports	-9 (± 8)	-2 (± 1)	0.78 (0.44)

^aPhase differences (leads and lags) $\Delta\phi_{d,d'}$ with respect to GDP. Positive $\Delta\phi$ means that the indicator d leads GDP ($d' = 1$), while a negative $\Delta\phi$ means it lags GDP; conversion between degrees (second column) and months (third column) is based on a business cycle lasting 5.2 years. The circular standard deviation σ_ϕ of $\Delta\phi$ is indicated in parentheses.

^bCoupling strength $R_{d,d'}$ with respect to GDP. The significance level $R_{\alpha=0.99}$ of R is derived from 10 000 realizations of surrogate data; see Section 4.3 for details.

Indeed, if the phases are totally locked then $\phi_d(t) - \phi_{d'}(t) \equiv \text{const.}$; by choosing this constant, without loss of generality, to be zero, the sum in (9) will be unity. On the other hand, if the phases are randomly distributed on the circle, and $N \rightarrow \infty$, then the sum on the right-hand side of the equation will be zero. The phase lag between the two channels is given by the circular mean of this phase difference

$$(10) \quad \Delta\phi_{d,d'} = \arg \left(\frac{1}{N} \sum_{t=1}^N e^{i(\phi_d(t) - \phi_{d'}(t))} \right).$$

The results of applying the definitions in (9) and (10) to our economic data are shown in Table I, where we have chosen the GDP as our reference time series and consider all leads and lags, as well as the coupling strengths, with respect to GDP. The second column in the table lists the leads and lags in terms of degrees of angle. To infer the corresponding values in terms of months, in the third column of the table, we assume an average business cycle length of 5.2 years, as derived from the M-SSA results of Sections 3.1 and 3.2.

The circular standard deviation σ_ϕ of the phase difference $\Delta\phi_{d,d'}$ between the channels d and d' is directly related to the coupling strength $R = R_{d,d'}$, $\sigma_\phi = (-2/\hat{N} \log R)^{1/2}$ (Fisher (1993)), with $\hat{N} = N/(\tau\kappa)$, the number of independent samples (cf. Fig. 2 and its caption). The coupling strength is shown in the fourth column, along with its significance level (in parentheses). We will discuss this property in Section 4.3 below.

The results in Table I are in good agreement with widely accepted stylized facts of the business cycle (e.g., Kydland and Prescott (1998); Zarnowitz (1985)). In particular, they support the fact that the price is countercyclical, i.e. that it lags

GDP by almost one-third of the period ($\Delta\phi = -114^\circ$, or a lag of 19 months), and that it is only weakly coupled to GDP ($R = 0.51$ with $R_\alpha = 0.46$, i.e. only slightly above the 99% significance level). Exports are also weakly coupled ($R = 0.57$) and lag GDP by about 6 months. Imports, on the other hand, are more strongly coupled ($R = 0.78$) and show almost no phase shift. Indeed, imports depend directly on national demand, which is linked to national income, investment and GDP, while exports depend largely on the demand from other countries and their economic situation. Even though the economic situation of other countries may be partly synchronized with the U.S. cycle, the coupling between U.S. exports and GDP is weaker than for other variables.

All other indicators are strong related to GDP, with $R \geq 0.81$. Thus consumption leads GDP by one quarter, which confirms its well-known role in the U.S. of stimulating the economy. Change in private inventories also leads by about half a year. Total wage and investments both lag GDP by about one quarter. Special attention in this analysis has to be paid to the phasing of employment since, unlike GDP, it has no exponential trend (see also Section 2.1). The detrending of GDP yields a shift in its peaks and troughs, which modifies the estimation of the phase lag. Later on, in Section 4.4, where we consider the amplitude along with the phase, we will see that GDP and employment track each other quite closely.

When interpreting these results, one has to consider that the phase differences are not fixed over time: Table I provides only an average value, while the corresponding circular standard deviation gives a rough measure of uncertainty and variability.

4.3. *Turning points and surrogate data*

In order to justify the results of Section 4.2, it is necessary to validate them against an alternative, or “null” hypothesis. This means finding an alternative point of view, which rejects phase coupling, and repeating the whole analysis. Since we only have a single trial of our macroeconomic time series, and we cannot repeat the experiment — i.e., rerun the whole U.S. economy for 52 years without phase coupling — we have to construct a *surrogate data set* from the original one, in which we are sure that the phase coupling is absent. A straightforward approach would be to simply use white noise, in which the phases are well known to be statistically independent. As already discussed in Section 3.1, however, the cyclical behavior is not purely random in nature, and so we have to use a different approach to constructing the surrogate data. Our method preserves typical business cycle patterns in each channel, while only destroying the phase relationships between the channels.

We split the vector time series $\mathbf{x}(t)$ of nine economic indicators into its individual cycles and shuffle the phases, separately for each time series $x_d(t)$. To find the individual business cycles, we describe the state of the economy by an




FIGURE 5.— Differences $\Delta z(t)$ between the instantaneous-phase vector $\phi(t)$ and its value at the reference time $t_{\text{ref}} = 1990:3$; see (11). The shaded areas indicate NBER dates of U.S. recessions, as in Figs. 1 and 4, while the horizontal line placed at $\Delta z(t) = 1.2$ indicates our threshold for selecting the minima of $\Delta z(t)$.

instantaneous-phase vector $\phi(t)$ that contains the phases $\phi_d(t)$ of each $x_d(t)$:

$$\phi(t) = (\phi_1(t), \phi_2(t), \dots, \phi_D(t)).$$

To split the time series into its individual cycles, we start with a reference time t_{ref} : this becomes our first “turning point,” where a new cycle begins. We arbitrarily chose the NBER peak date of $t_{\text{ref}} = 1990:3$, where the notation indicates the third quarter of the calendar year 1993; this means that $\phi(t_{\text{ref}})$ represents a typical phase vector of the economy during a peak. To find other epochs when the economy passes through a peak, we look for a return of $\phi(t)$ into the neighborhood of $\phi(t_{\text{ref}})$, i.e., we seek local minima in the average phase distance between $\phi(t)$ and $\phi(t_{\text{ref}})$,

$$(11) \quad \Delta z(t) = \left[\frac{1}{D} \sum_{d=1}^D \left\{ \arg e^{i(\phi_d(t) - \phi_d(t_{\text{ref}}))} \right\}^2 \right]^{1/2}.$$

The evolution of $\Delta z(t)$ is plotted in Fig. 5, and it ranges from 0 to π . Beside the global minimum at $t = t_{\text{ref}}$, where $\Delta z(t) = 0$, we find other distinct local minima. They match very well the NBER dates of peaks, i.e. the beginning of the recessions, shown as shaded areas in the figure. Unsurprisingly, there are also some discrepancies: these arise from the fact that the NBER takes into account a broad set of macroeconomic aggregates to define recessions, and uses a qualitative methodology with no fixed rules. Nevertheless, the agreement is evident and justifies describing the state of the U.S. economy with the help of the instantaneous-phase vector $\phi(t)$ of (11).

We choose those minima that lie below a threshold indicated by the horizontal line in Fig. 5. This choice represents, of course, a certain degree of arbitrariness,

like in every statistical study based on a relatively small and noisy data set. Various thresholds yield different numbers of turning points, which correspond more or less well to the NBER dates of peaks. While changing the results slightly, the average business cycle behavior presented next, in Section 4.4, preserves its general characteristics.

Furthermore, we found that the results of our procedure for splitting the record into separate cycles depend on the choice of the macroeconomic aggregates: just relying on GDP does not characterize the health of the economy. Taking more and more aggregates into account, our peaks and troughs became closer to those of the NBER (not shown here). It is thus essential to consider several aggregates to describe and understand the economy’s dynamical behavior.

The turning points derived by this procedure are indicated in the lower panel of Fig. 1 as vertical lines. With these results in hand, we can shuffle the individual cycles for each economic indicator d and recalculate the coupling indices $R_{d,d'}$ for each pair of indicators, in particular for $d' = 1$, the GDP. For every time series in Table I, we repeat this shuffling procedure 10 000 times and derive a significance level for the coupling index $R = R_{d,1}$, based on the 99% quantile. Above this quantile, $R > R_{\alpha=0.99}$, the null hypothesis of no phase coupling has to be rejected.

We see that the R -values of almost all coupling indices in the table lie well above the corresponding significance levels $R_{\alpha=0.99}$, meaning that the cyclical behavior of ups and downs of almost all eight remaining indicators is strongly related to the GDP. Only the coupling of price with GDP is fairly weak, although it still cannot be rejected at the 99% level. Remarkable is also the fact, already discussed in the preceding Section 4.2, that exports are less closely related to GDP than imports.

4.4. Average business cycle behavior

In the previous section we have discussed the separation of the macroeconomic time series into individual business cycles. We now proceed to extract typical features that can be found in all these cycles. Every historical period is affected by different events — such as technological developments, political changes, geopolitical events, or exogenous shocks like natural disasters — while endogenous mechanisms may also play a significant role (e.g., Chiarella, Flaschel, and Franke (2005); Day (1982); Goodwin (1967); Grandmont (1985); Hallegatte and Ghil (2008); Hallegatte, Ghil, Dumas, and Hourcade (2008); Ireland (2003); Schmitt-Grohé (2000)).

As a consequence of exogenous factors, no time period is directly comparable to another, and stylized facts may change over time to some extent (e.g., Stock and Watson (2002)). Understanding and explaining business cycles requires one to distinguish between the effect on a given cycle of a specific event, such as an oil shock, and the generic, endogenous part. For instance, it is not clear whether the downturn of the early 1970s is purely a consequence of a geopolitical event

and of the subsequent rise in oil prices, or whether this evolution is also linked to an endogenous mechanism and, therefore, to the essence of the business cycle. Averaging the individual cycles reduces exogenous effects, so that endogenous mechanisms become more clearly visible.

The strong variation in cycle length — between 3 years (1967–1970) and 9 (1982–1991), say (cf. lower panel of Fig. 1) — renders the use of simple averaging, as in [Sichel \(1994\)](#) or [Kontolemis \(1997\)](#), rather problematic. Those analyses of the “typical” business cycle have considered, therefore, mainly a few quarters during, before and after the NBER-determined recessions. We present here a rather different approach to the analysis of the average business cycle, allowing us to study the full length of the cycle.

In order to make individual cycles comparable to each other, it is necessary to stretch or shrink them to the same duration of 5.2 years, say. One must therewith interpolate unknown values for any cycle that differs in length from 5.2 years. Doing so is not possible for the original time series, because it represents a broadband signal and exhibits a lot of high-frequency fluctuations. On the other hand, in the M-SSA-filtered time series of RCs 1–16, the high-frequency fluctuations have been removed, leaving only smooth behavior. The conditions of the sampling theorem ([Shannon \(1948\)](#)) are thus fulfilled, and so we can fit a curve between the quarterly data points by using standard interpolation algorithms.

To stretch or shrink all the cycles to the same length, we redistribute their equidistant data points into an interval of 5.2 years and interpolate unknown values with a month-by-month resolution by linear interpolation. Unknown values between two known data points $x(t_1)$ and $x(t_2)$ are thus calculated by $x(t) = x(t_1) + [(t - t_1)/(t_2 - t_1)](x(t_2) - x(t_1))$, for any $t_1 < t < t_2$. Finally, we determine the mean value over all the cycles, for each month. The average business cycle behavior so derived is illustrated in Fig. 6. Cubic-spline interpolation (not shown) yields visually indistinguishable results.

To return from relative to absolute values, one has to reverse the pre-processing transformations described in Section 2.1. In order to reverse the subtraction of the trend and the division by it, we use the mean value of the trend for the last five years of the 52-year time interval in the data.

In Fig. 6, the recession phase of the cycle is quite short in comparison with the expansion phase: the growth period in the GDP is of about 44 months, while the recession is of about 18 months. Our result differs slightly from the NBER statement; the difference arises, amongst other reasons, from their considering original and not detrended indicators.

Our average cycle in Fig. 6 confirms several stylized facts from the literature. In particular, we find a higher-growth phase in the first year, followed by lower growth in the second and third years of the expansion. This is consistent with an uneven growth rate during expansion ([Emery and Koenig \(1992\)](#); [Kontolemis \(1997\)](#); [Sichel \(1994\)](#)). [Sichel \(1994\)](#) explains the presence of three phases by the role of inventory dynamics, which is responsible for the initial high-growth period in his analysis. In our average business cycle, the high-growth phase is also

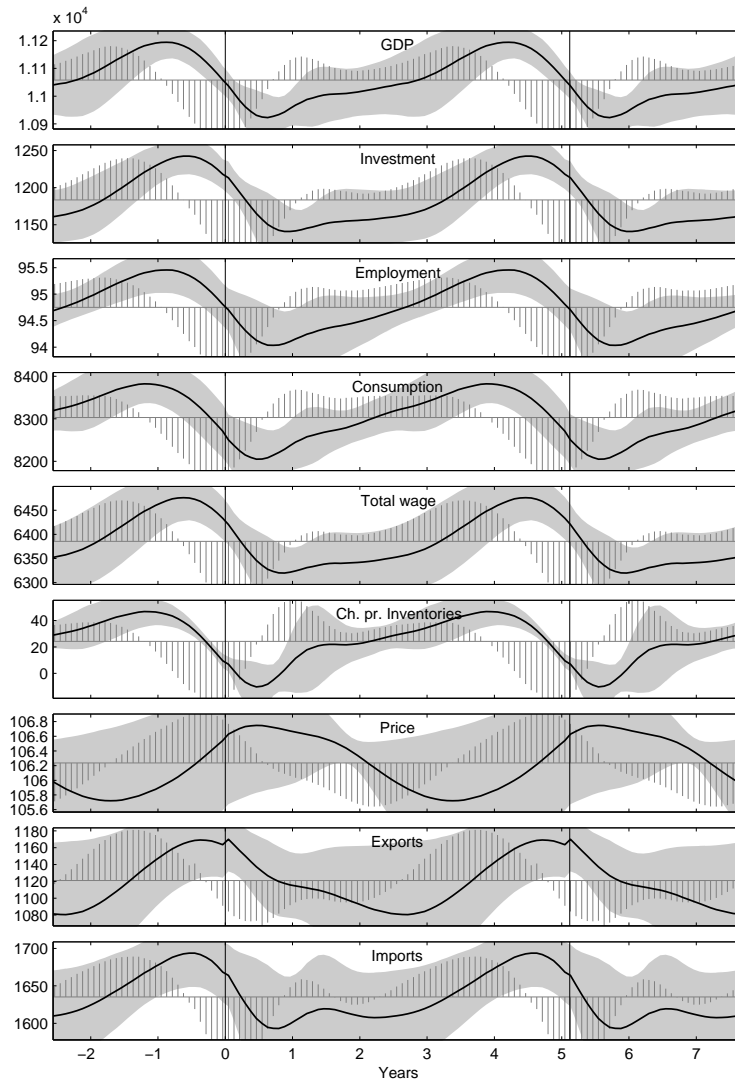


FIGURE 6.— The average business cycle, with the cycle length set to 5.2 years. Heavy solid lines show the average over the individual cycles, while the gray areas indicate the corresponding standard deviation, and the shorter vertical lines indicate the one-year growth rate. For better visualization, about half the cycle is repeated on either side of the full period.

explained by the particularly rapid growth of change in private inventories. This component grows rapidly in the first year of the expansion, and then decelerates substantially in the second year.

In contrast to Sichel’s three-phase model, we observe a further acceleration in GDP at the end of the expansion, which could be explained by a more sustained growth of investments. This difference in GDP behavior could also be affected by the different approaches to averaging: we have stretched all cycles to the same length, whereas Sichel (1994) did not apply such a temporal correction. Consumption and employment, though, are found to grow at a constant rate during expansion. Strikingly, the average business cycle in the figure here shows that employment is almost in phase with GDP when considering the amplitude, too, and not just the phase difference, as in Table. I.

During recessions, all variables exhibit strongly negative rates of change (short vertical “sticks” in the figure), while the recovery exhibits a sharp “rebound,” with a rapid increase in growth rate. Following the terminology of Sichel (1993), we do find *steepness* in production and employment. Our results, on the other hand, do not confirm *deepness*, since peaks and troughs in the average business cycle of Fig. 6 are of about the same amplitude. A definitive quantitative statement to this effect, however, is beyond the scope of the present paper, and the question is left for further investigation.

5. CONCLUDING REMARKS

In this article, we proposed two novel methodologies — singular spectral analysis (SSA; see Sections 2 and 3) and phase synchronization (see Section 4) — to extract a consistent multivariate business cycle from time series of macroeconomic indicators. We applied these method to nine U.S. indicators available from the BEA for 1954–2005 (see Section 2.1 and Fig. 1 there) and showed that our combined methodologies are able to extract an average business cycle (see Section 4.4 and Fig. 6 there) and its stylized facts from these macroeconomic time series.

SSA and its multivariate version, M-SSA, allow us to identify oscillatory modes of behavior (see Figs. 2 and 3) with high statistical confidence, cf. Ghil, Allen, Dettinger, Ide, Kondrashov, Mann, Robertson, Saunders, Tian, Varadi, and Yiou (2002) and the SSA-MTM Toolkit at <http://www.atmos.ucla.edu/tcd/ssa>. The modes we find are dominated by a mean cycle period of 5.2 years. This period agrees, on the one hand, with previous data-based studies that yielded the accepted business cycle length of 5–6 years; on the other, it agrees with the period of the simple endogenous business cycle model of Hallegatte, Ghil, Dumas, and Hourcade (2008) and Hallegatte and Ghil (2008). M-SSA also allows us to provide further insight into the underlying macroeconomic dynamics, with the recession phase being dominated by the 5.2-year mode, while the expansion exhibits more complex dynamics, with higher-frequency modes at play (see Fig. 4).

To analyze in greater depth the characteristic features of the business cycles extracted by M-SSA, we applied the phase synchronization methodology, cf. [Pikovsky, Rosenblum, and Kurths \(2003\)](#). The stylized facts so obtained are consistent with previous work on business cycles, yielding high correlation between GDP, employment, and wages. [Sichel \(1994\)](#) posited a three-phase pattern of the business cycle, with relatively rapid growth in the first year of the expansion phase, followed by slower growth in its second and third years, and then only by the recession phase. We identified here a further phase of acceleration at the end of expansion (see again [Fig. 6](#)), thus proposing a four-phase cycle. [Emery and Koenig \(1992\)](#) and [Kontolemis \(1997\)](#) have also discussed the existence of more than three phases.

The [Slutzky \(1937\)](#) theory of purely random shocks generating business cycles implies the absence of significant oscillatory modes in macroeconomic indicators. Our results on the unambiguous presence of such modes allow us to reject this theory (see error bars in [Fig. 2](#) and spectral peaks in [Fig. 3](#)).

The next step in our research program is to investigate whether the regular stylized facts of our average business cycle arise (i) from stylized facts in productivity shocks, which could translate into some oscillatory structure of economic fluctuations (e.g., [King and Rebelo \(2000\)](#)); (ii) from the response of the economic system to random productivity shocks (e.g., [Ireland \(2003\)](#)); or (iii) from endogenous instabilities in the economic system (e.g., [Chiarella, Flaschel, and Franke \(2005\)](#); [Day \(1982\)](#); [Goodwin \(1951\)](#); [Grandmont \(1985\)](#); [Hallegatte, Ghil, Dumas, and Hourcade \(2008\)](#); [Harrod \(1939\)](#); [Hicks \(1950\)](#); [Kaldor \(1940\)](#); [Kalecki \(1937\)](#); [Samuelson \(1939\)](#)). One possible way of advancing this program is to refine and calibrate the nonequilibrium dynamical model (NEDyM) of [Hallegatte, Ghil, Dumas, and Hourcade \(2008\)](#), which already has the proper period of 5–6 years, and several — but not all — stylized facts that are correct.

To achieve a suitable calibration of this model’s parameters on economic data, we expect to rely on the “data assimilation” approach, whose use is common by now in the geosciences ([Bengtsson, Ghil, and Källén \(1981\)](#); [Ghil and Malanotte-Rizzoli \(1991\)](#); [Kondrashov, Sun, and Ghil \(2008\)](#)). Data assimilation has been used in the econometric context ([Harvey \(1989\)](#)), too, but is only starting to be applied to macroeconomic models ([Lemoine and Pelgrin \(2004\)](#)). Preliminary results (not reported here) were obtained by applying data assimilation methods — such as Kalman filtering and extended Kalman filtering ([Ghil \(1997\)](#); [Kalman \(1960\)](#)) — to the NEDyM model, by using at first synthetic data produced by the model itself. This so-called “identical-twin” approach ([Bengtsson, Ghil, and Källén \(1981\)](#)) produced results that were encouraging but not yet conclusive.

We plan to use in subsequent work the average business cycle obtained here, instead of synthetic data, for model calibration. Doing so could allow us to bridge the gap between real-cycle theory and endogenous cycle theory, by acknowledging that both endogenous instabilities and exogenous shocks play a role in economic fluctuations.

To conclude, we have seen that it is essential to understand the economy as

a complex system, whose dynamical behavior cannot be described by a single macroeconomic indicator, such as the GDP. In the present study, we have described and analyzed the data, therefore, in an entirely multivariate fashion. First, the business cycle behavior was extracted by applying M-SSA. Second, to create an average business cycle, we considered the lead-and-lag relations between all nine indicators of the U.S. economy. In fact, as we increased the number of aggregates taken into account, the peaks and troughs of our purely objective methodology approximated better and better those obtained more subjectively, by relying on the vast experience of leading economists, at the NBER. In future work, we plan to extend this analysis in order to study relationships between the economies of different countries, so as to determine the degree of synchronization among them.

REFERENCES

- BAXTER, M., AND R. G. KING (1999): "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series," *Review of Economics and Statistics*, 81(4), 575–593.
- BEDROSIAN, E. (1963): "A product theorem for Hilbert transforms," *Proc. IEEE*, 51(5), 868–869.
- BENGTSSON, L., M. GHIL, AND E. KÄLLÉN (eds.) (1981): *Dynamic Meteorology: Data Assimilation Methods*. Springer-Verlag, New York/Heidelberg/Berlin.
- BROOMHEAD, D. S., AND G. P. KING (1986a): "Extracting qualitative dynamics from experimental data," *Physica D*, 20(2-3), 217–236.
- (1986b): "On the qualitative analysis of experimental dynamical systems," in *Nonlinear Phenomena and Chaos*, ed. by S. Sarkar, pp. 113–144. Adam Hilger, Bristol, England.
- BURNS, A. F., AND W. C. MITCHELL (1946): *Measuring Business Cycles*. NBER, New York City, New York.
- CANOVA, F. (1998): "Detrending and business cycle facts," *Journal of Monetary Economics*, 41(3), 475–512.
- CHIARELLA, C., P. FLASCHEL, AND R. FRANKE (2005): *Foundations for a Disequilibrium Theory of the Business Cycle*. Cambridge: Cambridge University Press.
- COGLEY, T., AND J. M. NASON (1995): "Effects of the Hodrick-Prescott filter on trend and difference stationary time series: Implications for business cycle research," *Journal of Economic Dynamics and Control*, 19(1-2), 253–278.
- DAY, R. (1982): "Irregular growth cycles," *American Economic Review*, 72, 406–414.
- EMERY, K., AND E. KOENIG (1992): "Forecasting turning points. Is a two-state characterization of the business cycle appropriate?," *Economic Letters*, 39, 431–435.
- FELIKS, Y., M. GHIL, AND A. W. ROBERTSON (2009): "Oscillatory climate modes in the Eastern Mediterranean: Synchronization with the NAO and Southern Oscillation," *Journal of Climate*, submitted.
- FISHER, N. I. (1993): *Statistical Analysis of Circular Data*. Cambridge University Press.
- GHIL, M. (1997): "Advances in sequential estimation for atmospheric and oceanic flows," *Journal of the Meteorological Society of Japan*, 75(1B), 289–304.
- GHIL, M., M. R. ALLEN, M. D. DETTINGER, K. IDE, D. KONDRASHOV, M. E. MANN, A. W. ROBERTSON, A. SAUNDERS, Y. TIAN, F. VARADI, AND P. YIOU (2002): "Advanced spectral methods for climatic time series," *Reviews of Geophysics*, 40(1), 1–41.
- GHIL, M., AND S. CHILDRESS (1987): *Topics in Geophysical Fluid Dynamics: Atmospheric Dynamics, Dynamo Theory and Climate Dynamics*. Springer-Verlag.
- GHIL, M., AND P. MALANOTTE-RIZZOLI (1991): "Data assimilation in meteorology and oceanography," *Advances in Geophysics*, 33, 141–266.
- GHIL, M., AND R. VAUTARD (1991): "Interdecadal oscillations and the warming trend in global temperature time series," *Nature*, 350(6316), 324–327.

- GOODWIN, R. (1951): "The non-linear accelerator and the persistence of business cycles," *Econometrica*, 19, 1–17.
- (1967): "A growth cycle," in *Socialism, Capitalism and Economic Growth*, ed. by C. Feinstein, pp. 54–58. Cambridge University Press.
- GRANDMONT, J.-M. (1985): "On endogenous competitive business cycles," *Econometrica*, 5, 995–1045.
- HALLEGATTE, S., AND M. GHIL (2008): "Natural disasters impacting a macroeconomic model with endogenous dynamics," *Ecological Economics*, 68(1–2), 582–592.
- HALLEGATTE, S., M. GHIL, P. DUMAS, AND J.-C. HOURCADE (2008): "Business cycles, bifurcations and chaos in a neo-classical model with investment dynamics," *Journal of Economic Behavior & Organization*, 67(1), 57–77.
- HARROD, R. (1939): "An essay on dynamic economic theory," *Economic Journal*, 49, 1433.
- HARVEY, A. C. (1989): *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- HICKS, J. (1950): "The Cycle in Outline," in *A Contribution to the Theory of the Trade Cycle*, chap. 8, pp. 95–107. Oxford University Press, Oxford.
- HILBERT, D. (1953): *Grundzüge einer allgemeinen Theorie der linearen Integralgleichungen*. Chelsea Publ. Co.
- HODRICK, R. J., AND E. C. PRESCOTT (1997): "Postwar U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit and Banking*, 29(1), 1–16.
- IRELAND, P. (2003): "Endogenous money or sticky prices?," *Journal of Monetary Economics*, 50, 1623–1648.
- KALDOR, N. (1940): "A model of the trade cycle," *Economic Journal*, 50, 78–92.
- KALECKI, M. (1937): "A theory of the business cycle," *Review of Economic Studies*, 4, 77–97.
- KALMAN, R. E. (1960): "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering, Transactions of the ASME, Series D*, 82, 5–45.
- KARHUNEN, K. (1946): "Zur Spektraltheorie stochastischer Prozesse," *Ann. Acad. Sci. Fenn. Ser. A1, Math. Phys.*, 34.
- KEPPENNE, C. L., AND M. GHIL (1993): "Adaptive filtering and prediction of noisy multivariate signals: An application to subannual variability in atmospheric angular momentum," *International Journal of Bifurcation and Chaos*, 3, 625–634.
- KIMOTO, M., M. GHIL, AND K.-C. MO (1991): "Spatial structure of the extratropical 40-day oscillation," in *Eighth Conf. Atmos. & Oceanic Waves and Stability (Denver, Colo.)*, pp. 115–116. American Meteorological Society, Boston, Mass.
- KING, R., AND S. REBELO (2000): "Resuscitating real business cycles," in *Handbook of Macroeconomics*, ed. by J. Taylor, and M. Woodford, pp. 927–1007. North-Holland, Amsterdam.
- KONDRASHOV, D., C. SUN, AND M. GHIL (2008): "Data Assimilation for a Coupled Ocean-Atmosphere Model. Part II: Parameter Estimation," *Monthly Weather Review*, , In press(2008).
- KONTOLEMIS, Z. (1997): "Does growth vary over the business cycle? Some evidence from the G7 countries," *Economica*, 64, 441–460.
- KYDLAND, F. E., AND E. C. PRESCOTT (1998): "Business cycles: Real Facts and a Monetary Myth," in *Real Business Cycles: A Reader*, ed. by K. D. S. James E. Hartley, Kevin D. Hoover, pp. 231–247. Routledge.
- LEMOINE, M., AND F. PELGRIN (2004): "Introduction aux modèles espace-état et au filtre de Kalman," *Revue de l'OFCE*, 86.
- LISI, F., AND A. MEDIO (1997): "Is a random walk the best exchange rate predictor?," *International Journal of Forecasting*, 13(2), 255–267.
- LOÈVE, M. (1945): "Fonctions aléatoires de second ordre," *Comptes Rendus de l'Académie des sciences Paris*, 220, 380.
- (1978): *Probability Theory, Vol. II, 4th ed.*, vol. 46 of *Graduate Texts in Mathematics*. Springer-Verlag.
- LUCAS, R. E. (1977): "Understanding business cycles," *Carnegie-Rochester Conference Series on Public Policy*, 5(1), 7–29.
- MAÑÉ, R. (1981): "On the dimension of the compact invariant sets of certain non-linear maps,"

- in *Dynamical Systems and Turbulence*, vol. 898 of *Lecture Notes in Mathematics*, pp. 230–242. Springer-Verlag, Berlin.
- MARPLE, JR., L. (1999): “Computing the discrete-time analytic signal via FFT,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 47(9), 2600–2603.
- MEES, A. I., P. E. RAPP, AND L. S. JENNINGS (1987): “Singular-value decomposition and embedding dimension,” *Physical Review A*, 36(1), 340–346.
- OPPENHEIM, A. V., AND R. W. SCHAFER (1989): *Discrete-time signal processing*. Prentice-Hall, Inc. Upper Saddle River, NJ.
- PENLAND, C. M., M. GHIL, AND K. M. WEICKMANN (1991): “Adaptive filtering and maximum entropy spectra, with application to changes in atmospheric angular momentum,” *Journal of Geophysical Research*, 96, 22659–22671.
- PIKOVSKY, A., M. ROSENBLUM, AND J. KURTHS (2003): *Synchronization - A universal concept in nonlinear sciences*. Cambridge University Press.
- PLAUT, G., AND R. VAUTARD (1994): “Spells of Low-Frequency Oscillations and Weather Regimes in the Northern Hemisphere,” *Journal of the Atmospheric Sciences*, 51(2), 210–236.
- ROSENBLUM, M., A. PIKOVSKY, AND K. KURTHS (1996): “Phase Synchronization of Chaotic Oscillators,” *Physical Review Letters*, 76(11), 1804–1807.
- SAMUELSON, P. (1939): “A Synthesis of the Principle of Acceleration and the Multiplier,” *Journal of Political Economy*, 47, 786–797.
- SAUER, T., J. A. YORKE, AND M. CASDAGLI (1991): “Embedology,” *Journal of Statistical Physics*, 65(3-4), 579–616.
- SCHMITT-GROHÉ, S. (2000): “Endogenous business cycles and the dynamics of output, hours, and consumption,” *The American Economic Review*, 90(5), 1136–1159.
- SHANNON, C. E. (1948): “A mathematical theory of communication,” *Bell System Technical Journal*, 27(3,4), 379–423, 623–656.
- SICHEL, D. E. (1993): “Business Cycle Asymmetry: A Deeper Look,” *Economic Inquiry*, 31(2), 224–236.
- (1994): “Inventories and the Three Phases of the Business Cycle,” *Journal of Business & Economic Statistics*, 12(3), 269–277.
- SLUTZKY, E. (1937): “The summation of random causes as a source of cyclic processes,” *Econometrica*, 5, 105–146.
- STOCK, J. H., AND M. W. WATSON (2002): “Has the Business Cycle Changed and Why?,” NBER Working Papers 9127, National Bureau of Economic Research, Inc.
- TAKENS, F. (1981): “Detecting strange attractors in turbulence,” in *Dynamical Systems and Turbulence*, vol. 898 of *Lecture Notes in Mathematics*, pp. 366–381. Springer, Berlin.
- UNAL, Y. S., AND M. GHIL (1995): “Interannual and interdecadal oscillation patterns in sea level,” *Climate Dynamics*, 11(5), 255–278.
- VAUTARD, R., AND M. GHIL (1989): “Singular spectrum analysis in nonlinear dynamics, with applications to paleoclimatic time series,” *Physica D*, 35(3), 395–424.
- VAUTARD, R., P. YIOU, AND M. GHIL (1992): “Singular-spectrum analysis: A toolkit for short, noisy chaotic signals,” *Physica D*, 58(1-4), 95–126.
- WINFREE, A. T. (1980): *The Geometry of Biological Time*. Springer-Verlag, New York.
- WINTERHALDER, M., B. SCHELTER, J. KURTHS, A. SCHULZE-BONHAGE, AND J. TIMMER (2006): “Sensitivity and specificity of coherence and phase synchronization analysis,” *Physics Letters A*, 356(1), 26–34.
- ZARNOWITZ, V. (1985): “Recent Work on Business Cycles in Historical Perspective: A Review of Theories and Evidence,” *Journal of Economic Literature*, 23(2), 523–580.