

The Role of Oscillatory Modes in U.S. Business Cycles

Andreas Groth^{a,b}, Michael Ghil^{a,b,c,*}, Stéphane Hallegatte^{d,e}, Patrice Dumas^d

^a*Geosciences Department, Ecole Normale Supérieure, Paris, France*

^b*Environmental Research & Teaching Institute, Ecole Normale Supérieure, Paris, France*

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

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

^e*Ecole Nationale de la Météorologie, Météo France, Toulouse, France*

Abstract

We apply multivariate singular spectrum analysis to the study of U.S. business cycle dynamics. This method provides a robust way to identify and reconstruct oscillations, whether intermittent or modulated. We show such oscillations to be associated with comovements across the entire economy. The problem of spurious cycles generated by the use of detrending filters is addressed and we present a Monte Carlo test to extract significant oscillations. The behavior of the U.S. economy is shown to change significantly from one phase of the business cycle to another: the recession phase is dominated by a five-year mode, while the expansion phase exhibits more complex dynamics, with higher-frequency modes coming into play. We show that the variations so identified cannot be generated by random shocks alone, as assumed in ‘real’ business-cycle models, and that endogenous, deterministically generated variability has to be involved.

Keywords: Advanced spectral methods, Comovements, Frequency domain, Monte Carlo testing, Time domain

JEL classification: C15, C60, E32

1. Introduction

Dominated over many decades by a long-term upward drift (Solow, 1956), macroeconomic time series also exhibit smaller but still important shorter-term fluctuations often associated with business cycles. The causes and characteristics of these cycles have been extensively studied in modern economic theory (Burns and Mitchell, 1946; Kydland and Prescott, 1998, and references therein).

A number of approaches have been proposed to separate the shorter-term fluctuations from the long-term trend (Canova, 1998; Baxter and King, 1999). Morley and Piger (2012) recently attempted a classification of business cycle

*Corresponding author: Michael Ghil, ghil@lmd.ens.fr

analyses into (i) those that consider a cyclic sequence of expansions and contractions; and (ii) an output-gap view, in which business cycles merely correspond to transitory fluctuations superimposed on a permanent trend level.

The definition of business cycles depends, however, on the knowledge of a trend component that cannot be observed directly. Several filters have been developed (Graff, 2011) to extract such a trend component, of which the Hodrick-Prescott (HP) filter is the most commonly used one (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, to avoid spurious results due merely to the detrending procedure itself (Nelson and Kang, 1981; Harvey and Jaeger, 1993; Cogley and Nason, 1995).

Business cycles can also be understood as comovements of transitory fluctuations in several distinct macroeconomic 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 apply *multivariate singular spectrum analysis* (M-SSA)—the extension of *singular spectrum analysis* (SSA) to multivariate time series—to the analysis of business cycles. Both SSA and M-SSA rely on the classical Karhunen–Loève spectral decomposition of random processes (Karhunen, 1946; Loève, 1945, 1978). Broomhead and King (1986a,b) proposed to use both in the context of nonlinear dynamics as a more robust application of the Mañé–Takens idea of reconstructing dynamics from measured time series (Mañé, 1981; Takens, 1981; Sauer et al., 1991). Ghil, Vautard and associates (Vautard and Ghil, 1989; Ghil and Vautard, 1991; Vautard et al., 1992) noticed that SSA can be used as a time-and-frequency domain method for the analysis of time series, whether they are generated by a linear stochastic process, a nonlinear deterministic one or a superposition of the two.

We propose to use M-SSA for the analysis of business cycles in a completely multivariate fashion. This method combines two useful approaches of statistical analysis: (1) it determines—with the help of principal component analysis (PCA)—major directions in the system’s phase space that are populated by the multivariate time series; and (2) it extracts major spectral components by using data-adaptive filters. To get reliable information about significant oscillatory modes, we perform exhaustive statistical tests by means of Monte Carlo SSA (MC-SSA, Allen and Smith, 1996), which allow us to deal with the problem of spurious oscillations (Nelson and Kang, 1981; Cogley and Nason, 1995).

SSA and M-SSA have already proven their advantages in a variety of applications, such as climate dynamics, meteorology and oceanography, as well as the biomedical sciences. Ghil et al. (2002) provide an overview and a comprehensive set of references to their theory and applications. In economics, this approach has received little attention so far. Recent applications to the univariate SSA analysis of business cycles include de Carvalho et al. (2012), Sella and Marchionatti (2012), and Dumas et al. (2013). The present paper introduces the M-SSA methodology into the economic literature and demonstrates its advantages for the multivariate analysis of economic activity.

The paper is organized as follows. In Section 2, we introduce the method-

ology and summarize its properties in terms of spectral decomposition, as well as of time-domain reconstruction. In Section 3, we apply SSA to the U.S. gross domestic product (GDP) and M-SSA to the full data set. The reliability of the results is then discussed via Monte Carlo testing. Section 4 analyzes the cycle-to-cycle variability of the U.S. business cycles, and we draw conclusions about the underlying dynamics in Section 5.

2. Decomposition and reconstruction

2.1. Data and pre-processing

We study here U.S. macroeconomic data from the Bureau of Economic Analysis (BEA; see <http://www.bea.gov>). The nine variables analyzed are 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 percentage points. The quarterly time series cover 52 years, from the first quarter of 1954 to the third quarter of 2005.

Aligning ourselves with the output-gap view of business cycle analysis, we first remove the trend of each time series separately, by using the HP filter with the common parameter value $\lambda = 1600$. Employment is the only one of the nine variables that does not exhibit an upward trend; still, we do detrend it to consistently remove periods longer than 10 years, as done for the other variables.

The restriction to the interval 1954–2005 reduces end-to-end mismatches of the remaining transitory fluctuations and minimizes spectral leakage effects, i.e. the presence of spurious oscillations in the spectral analysis. On the macroeconomic side, the years 2007–2008 correspond to a well-known crisis, whose origin was financial, rather than economical. The time interval since that crisis exhibited several new characteristics, for which we do not yet have sufficiently long data sets to distinguish this interval from the previous half-a-century of data.

In contrast to our two-step approach—see also [Dumas et al. \(2013\)](#)—[de Carvalho et al. \(2012\)](#) have chosen to apply a single SSA analysis for the decomposition of the GDP into trend and fluctuations. Such an approach is indeed desirable as it appears to provide a more consistent separation into a permanent trend component and transitory fluctuations that are orthogonal to it (cf. [Vautard et al., 1992](#); [Ghil et al., 2002](#)). In the present macroeconomic context, however, the trend dominates the SSA’s variance-based phase-space decomposition and small fluctuations could be masked in such a single-sweep analysis ([Granger, 1966](#); [Sella and Marchionatti, 2012](#)). We will return to this issue in the estimation of the covariance matrix in Sec. 2.2.

In line with our two-step decomposition, we next obtain trend residuals that we divide by the trend—i.e., we concentrate on relative values—and then transform to unit standard deviation. This normalization brings all the indicators to the same scale and gives equal weight to each of them in the M-SSA analysis. One could choose to give different weights to each time series to reflect a priori ideas on their relative importance. Our choice here is one of simplicity, and the

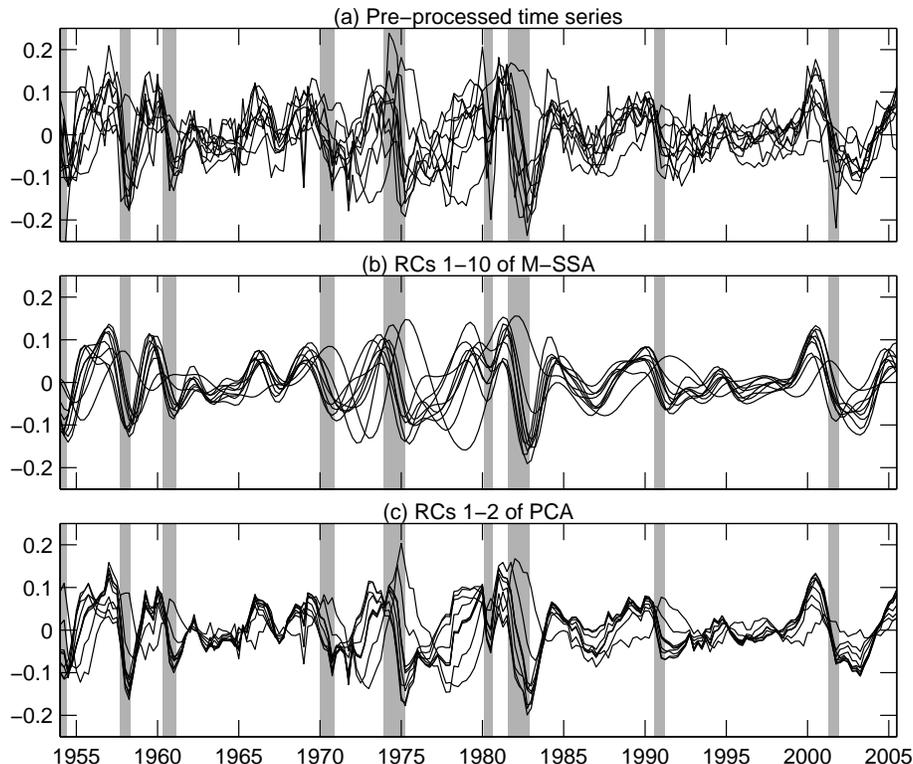


Figure 1: The nine time series of U.S. macroeconomic data used in this paper; raw data from the U.S. Bureau of Economic Analysis (BEA), 1954–2005. The figure illustrates the results of pre-processing and of applying either multivariate singular spectrum analysis (M-SSA) or principal component analysis (PCA); the shaded vertical bars in the three panels indicate NBER-defined recessions. (a) Detrended and standardized time series. (b,c) Reconstruction of the entire data set: (b) with the first 10 M-SSA components, using a window width of $M = 24$ quarters; and (c) with the first two PCA components. Both reconstructions capture 75% of the total variance, while the M-SSA reconstruction is smoother.

covariance matrix is transformed into a correlation matrix. Finally, we divide the normalized time series by $(DM)^{1/2}$ —with D the number of variables and M the window width—so that the sum of the partial variances equals one.

Figure 1a shows the results of this pre-processing. The U.S. recessions, as defined by the NBER, are indicated by shaded vertical bars.

2.2. Singular spectrum analysis (SSA)

In this section we discuss the univariate version of SSA and present its main properties, in particular, its ability to reconstruct cyclical dynamics.

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

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

with $t = 1, \dots, N-M+1$. The SSA procedure starts by calculating the principal directions of the embedded data set $\mathbf{x}(t)$.

Reconstructing the entire attractor of a nonlinear dynamical system from $\mathbf{x}(t)$, as originally proposed by [Broomhead and King \(1986a\)](#), may fail, however, even in relatively simple cases ([Mees et al., 1987](#); [Vautard and Ghil, 1989](#)). Ghil, Vautard and several associates first proposed, instead, to apply the SSA methodology to describe cyclical behavior in short and noisy time series, for which standard methods derived from Fourier analysis do not work well ([Vautard and Ghil, 1989](#); [Ghil and Vautard, 1991](#); [Vautard et al., 1992](#)). The key idea in their approach was to reconstruct the ‘skeleton of the attractor,’ i.e. the most robust, albeit unstable limit cycles embedded in it.

The next step in SSA is to compute the auto-covariance matrix \mathbf{C} of x . [Vautard and Ghil \(1989\)](#) proposed to estimate it by

$$c_{i,j} = \frac{1}{N-|i-j|} \sum_{t=1}^{N-|i-j|} x(t)x(t+|i-j|), \quad (2)$$

imposing a Toeplitz structure with constant diagonals: the entries $c_{i,j}$ of the matrix depend only on the lag $|i-j|$.

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

$$\mathbf{C}\boldsymbol{\rho}_k = \lambda_k\boldsymbol{\rho}_k. \quad (3)$$

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 in the corresponding direction $\boldsymbol{\rho}_k$. This computation helps us find, therefore, the major components of the system’s dynamical behavior.

The eigenvectors of such a Toeplitz matrix are necessarily either symmetric or anti-symmetric, and the method’s reliability in extracting oscillations is enhanced therewith by using this form of \mathbf{C} ([Allen and Robertson, 1996](#)). In the presence of strong non-stationarity, the *Toeplitz approach* yields a slightly larger bias in the reconstruction of low-frequency activity than the original *trajectory approach* of [Broomhead and King \(1986a\)](#).

The latter approach relies on a singular-value decomposition of the trajectory matrix \mathbf{x} and is more appropriate for the analysis of trends ([Ghil et al., 2002](#)). Our focus here is on the transitory fluctuations and we rely therefore on the Toeplitz approach for our analysis. [de Carvalho et al. \(2012\)](#) have found the trajectory approach to be also adequate in their one-step identification of the permanent trend and superimposed fluctuations in the U.S. business cycles.

By convention, the eigenvalues $\{\lambda_k, k = 1, \dots, M\}$ are arranged in descending order, from the largest to the smallest variance, yielding a so-called ‘scree diagram’ of eigenvalues λ_k vs. order k . In this diagram, one often looks for a clear break in the slope to distinguish ‘signal’ from ‘noise.’ Such a break,

however, occurs mostly when the noise is actually white, with no temporal correlations at all. The signal-to-noise separation test has, therefore, to be modified in the presence of non-vanishing correlations, as done in Sec. 3 below.

Projecting the embedded time series \mathbf{x} onto eigenvectors $\boldsymbol{\rho}_k$ yields the corresponding *principal components* (PCs),

$$A_k(t) = \sum_{j=1}^M x(t+j-1)\rho_k(j), \quad k = 1, \dots, M. \quad (4)$$

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 $t = M$ 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 et al., 2002).

Finally, we can reconstruct parts of the time series that are associated with a particular eigenvector by

$$r_k(t) = \frac{1}{M_t} \sum_{j=L_t}^{U_t} A_k(t-j+1)\rho_k(j), \quad k = 1, \dots, M, \quad (5)$$

cf. Ghil and Vautard (1991) and Vautard et al. (1992). 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 et al. (2002). Each *reconstructed component* (RC) $r_k(t)$ associated with the variance λ_k has a complete set of N indices, but with a reduced confidence in its values at either end of the time series.

Given any subset $k \in \mathcal{K}$ of eigenelements, we obtain the corresponding reconstruction $r_{\mathcal{K}}(t)$ by summing the RCs,

$$r_{\mathcal{K}}(t) = \sum_{k \in \mathcal{K}} r_k(t). \quad (6)$$

Typical choices of \mathcal{K} are (i) $\mathcal{K} = \{k : 1 \leq k \leq S\}$, where S is the *statistical dimension* of the time series, cf. Vautard and Ghil (1989), i.e., the number of statistically significant components; or (ii) a pair of components (k_0, k_1) for which $\lambda_{k_0} \approx \lambda_{k_1}$, and which may capture a cyclic mode (see Section 3). The whole set of RCs, $\mathcal{K} = \{k : 1 \leq k \leq M\}$, gives the complete reconstruction of the time series.

In the following we refer to the common notation for the reconstructed component r_k as RC k , and for a sum of several, consecutive RCs in Eq. (6) from index k to index k' as RCs k - k' .

From the viewpoint of signal processing, the RCs can be considered as filtered time series, with the eigenvectors being a set of *data adaptive* filters. Both Eqs. (4) and (5) can be interpreted as a finite-impulse response (FIR) filter (Oppenheim and Schaffer, 1989), with $\boldsymbol{\rho}_k$ being an FIR filter of length M . The PCs obtained in Eq. (4) are time-reversed in Eq. (5), and the FIR filter is run

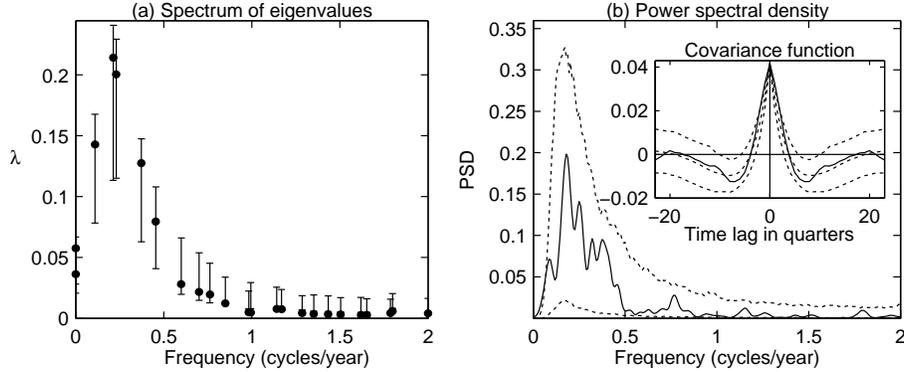


Figure 2: Univariate spectral analysis of U.S. gross domestic product (GDP). (a) Eigenvalue spectrum of λ_k (filled circles) vs. dominant frequency of the associated eigenvector ρ_k , with window width $M = 24$ quarters; the error bars indicate the significance levels (cf. Sec. 3.1). (b) Power spectral density (PSD) estimate (solid line) using Welch’s averaged periodogram method, with a Hamming window of length 128 quarters and 75% overlap (Priestley, 1991); the dashed lines indicate the significance levels. Inset: Covariance estimates (solid line) and their significance levels (dashed lines). The upper and lower significance levels in both panels and in the inset are derived from the 2.5% and 97.5% percentiles of 1000 surrogate time series from an AR(1) null hypothesis; see Sec. 3.1.

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. Hence, each RC k and the original time series $x(t)$ are in phase and the filtering acts only on the amplitude.

In designing an appropriate band-pass filter, Baxter and King (1999) require, in particular, that this “filter should not introduce phase shifts.” Unlike their band-pass filter, with its data-independent weights chosen *a priori*, SSA is data adaptive. The M filters are the eigenvectors of the auto-covariance matrix and provide an optimal spectral decomposition of the time series, i.e., a maximum of the variance is captured by a minimal number of spectral components.

Following Vautard et al. (1992), we assign to each pair (λ_k, ρ_k) a frequency, given by the maximum of the Fourier transform of ρ_k . Plotting each eigenvalue vs. its dominant frequency provides a complementary perspective on SSA in terms of an analogy with classical spectral analysis.

This analogy becomes more obvious when analyzing the trend residuals of GDP (Figure 2). In the SSA analysis of panel (a), we observe a maximum in the spectrum of eigenvalues at the usually reported mean business cycle length of 5–6 years, which agrees with the classical estimation of the power spectral density (PSD) in panel (b). For various PSD estimation algorithms that we have tested, we observe a maximum around the same period; at this point, though, the trend residuals are still subject to the Nelson and Kang (1981) criticism of spurious cycles. Therefore, we have to perform additional significance tests

before relying on the results, cf. Sec. 3.1.

2.3. Multivariate SSA (M-SSA)

M-SSA provides an extension of SSA to multivariate time series (Broomhead and King, 1986b; Kimoto et al., 1991; Keppenne and Ghil, 1993; Plaut and Vautard, 1994). Let $\mathbf{x}(t) = \{x_d(t): d = 1, \dots, D; t = 1, \dots, N\}$ be a vector time series of length N , with D channels. In generalizing (2), we use the D auto-covariances $\mathbf{C}_{d,d}$, as well as the $D \times (D - 1)$ cross-covariances $\mathbf{C}_{d,d'}$ to form a *grand covariance matrix* $\tilde{\mathbf{C}}$:

$$\tilde{\mathbf{C}} = \begin{pmatrix} \mathbf{C}_{1,1} & \mathbf{C}_{1,2} & \cdots & \mathbf{C}_{1,D} \\ \mathbf{C}_{2,1} & \mathbf{C}_{2,2} & \cdots & \mathbf{C}_{2,D} \\ \vdots & \vdots & \mathbf{C}_{d,d'} & \vdots \\ \mathbf{C}_{D,1} & \mathbf{C}_{D,2} & \cdots & \mathbf{C}_{D,D} \end{pmatrix}. \quad (7)$$

Here $\tilde{\mathbf{C}}$ is a $DM \times DM$ matrix and the entries of the individual matrices $\mathbf{C}_{d,d'}$ can be estimated as

$$(c_{i,j})_{d,d'} = \frac{1}{\tilde{N}} \sum_{t=\max\{1,1+i-j\}}^{\min\{N,N+i-j\}} x_d(t)x_{d'}(t+i-j). \quad (8)$$

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

As before we diagonalize the grand matrix $\tilde{\mathbf{C}}$ to yield its eigenvalues λ_k and eigenvectors $\tilde{\boldsymbol{\rho}}_k$,

$$\tilde{\mathbf{C}}\tilde{\boldsymbol{\rho}}_k = \lambda_k\tilde{\boldsymbol{\rho}}_k, \quad k = 1, \dots, DM. \quad (9)$$

In contrast to SSA, the M-SSA eigenvectors $\tilde{\boldsymbol{\rho}}_k$ have now length DM , and are composed of D consecutive segments $\boldsymbol{\rho}_k^d$, $d = 1, \dots, D$ of length M . These segments can be likewise interpreted as frequency-selective FIR filters, combined here into one multivariate filter $\tilde{\boldsymbol{\rho}}_k$.

The associated PCs are single-channel time series that are computed by projecting the multivariate time series $\mathbf{x}(t)$ onto $\tilde{\boldsymbol{\rho}}_k$,

$$A_k(t) = \sum_{d=1}^D \sum_{j=1}^M x_d(t+j-1)\rho_k^d(j), \quad k = 1, \dots, DM. \quad (10)$$

In addition to the summation j over time, as in Eq. (4), we get a second summation d over the channels. This summation involves a classical PCA. In particular, setting $M = 1$ reduces M-SSA to PCA in D variables.

Finally, one can reconstruct parts of each time series $x_d(t)$ associated with its corresponding eigenvector segment $\boldsymbol{\rho}_k^d$ by (Plaut and Vautard, 1994)

$$r_k^d(t) = \frac{1}{M_t} \sum_{j=L_t}^{U_t} A_k(t-j+1)\rho_k^d(j), \quad k = 1, \dots, DM; d = 1, \dots, D. \quad (11)$$

This formula provides a set of DM RCs for each of the D time series and—depending on the information contained in the cross-covariances $\mathbf{C}_{d,d'}$ —the RCs of different time series may or may not be correlated. In this way, M-SSA helps extract common spectral components from the multivariate data set, along with comovements of the time series.

In Figure 1, we compare the pre-processed time series in panel (a) with the M-SSA reconstruction in panel (b), using the ten leading RCs, and with PCA reconstruction in panel (c), using the two leading RCs of the data set; the latter results from Eq. (11) with $M = 1$. Both the M-SSA and PCA reconstructions capture 75% of the total variance and extract coherent behavior manifest in the nine economic variables. In contrast to PCA, the M-SSA results are much smoother, having removed irregular temporal fluctuations. It is especially the inclusion of temporal correlations that makes M-SSA superior to PCA in the extraction of dynamical behavior.

3. Oscillatory behavior and its statistical significance

The trend residuals in Figure 1 exhibit obviously more structure than pure white noise; we need, therefore, a stringent test to decide whether the visually apparent cyclical behavior can be attributed to random fluctuations or to a more regular oscillatory behavior, of possibly intrinsic origin. Cogley and Nason (1995), among others, have discussed in detail the effects of detrending, in particular the possibility that it might give rise to spurious cycles; their discussion was placed in the context of the detrending effect on a standard real business cycle (RBC) model, with no oscillatory dynamics.

We follow Cogley and Nason (1995) and test against a first-order autoregressive process, AR(1), to verify the statistical significance of oscillations; this test is already well-established in the geosciences (Allen and Smith, 1996; Ghil et al., 2002). Since AR(1) processes exhibit maximum variance at zero frequency, detrending with the HP filter may yield a possibly spurious maximum in the PSD at other frequencies, e.g. around the commonly reported business cycle length of 5-6 years.

3.1. Univariate time series: the GDP

We first focus on GDP alone, the most widely studied macroeconomic indicator. A Monte Carlo-type test consists in first fitting an AR(1) process $X(t)$ to the scalar time series $x(t)$ of interest,

$$X(t) = aX(t-1) + \sigma_0 \epsilon(t), \quad (12)$$

with $\epsilon(t)$ being Gaussian white noise of variance $\sigma = 1$, and then comparing the spectral properties of many realizations of $X(t)$ with that of $x(t)$.

We estimate the regression coefficient in Eq. (12) to be $a = 0.82$, and the variance $\sigma_0 = 0.04$, with the influence of the HP filter taken into account. That is, we choose the parameter a to minimize the mean-square distance between

the GDP covariance function (solid line in the inset of Figure 2b) and the HP-filtered AR(1) covariance function (dotted line), while to estimate σ_0 we use the unbiased estimator proposed by [Allen and Smith \(1996\)](#) for short time series. Given the model parameters, we create a set of 1000 surrogate realizations of length N from the AR(1) model, lowpass filter each with the HP filter, and normalize it to the same variance as $x(t)$. An additional division by the trend is not necessary, since the innovations in the AR(1) process have constant variance.

In Figure 2b, we compare the PSD estimate of the GDP residuals (solid line) with that of the surrogate time series. Frequency-dependent significance levels at the 2.5% and 97.5% quantiles (dashed lines) also show high power around five years, and the PSD estimate of the GDP falls entirely between them.

Other PSD estimates, such as the maximum entropy or the multi-taper method (see [Ghil et al., 2002](#), and references therein), confirm this finding and yield the conclusion that GDP residuals alone cannot be distinguished from the null hypothesis of an HP-filtered AR(1) process. The high PSD values around five years could be due to the detrending of an otherwise stable model with exogenous excitation, in complete agreement with the findings of [Cogley and Nason \(1995\)](#) and [Nelson and Kang \(1981\)](#).

The same lack of statistical significance holds for the lag-covariance function, as expected from the Wiener-Khinchin theorem, according to which the PSD and the lag-covariance function of a time series are related by the Fourier transform ([Blackman and Tukey, 1958](#)). The swing below zero for the surrogate time series is thus likewise due to the HP filter’s effect; see the inset in Figure 2b. The preliminary conclusion is that we cannot falsify the null hypothesis of an AR(1) process for the GDP residuals, as analyzed separately here.

This conclusion can also be confirmed by MC-SSA, which tests whether an eigenvalue λ_k captures more partial variance in the direction of the corresponding eigenvector $\boldsymbol{\rho}_k$ than present in the null hypothesis. To derive significance levels, the covariance matrix \mathbf{C}_S for each surrogate time series $\mathbf{x}_S(t)$ is projected onto the set of eigenvectors \mathbf{E} of the original time series via

$$\boldsymbol{\Lambda}_S = \mathbf{E}^\top \mathbf{C}_S \mathbf{E}; \tag{13}$$

here the eigenvectors $\boldsymbol{\rho}_k$ are the columns of \mathbf{E} , and $(\cdot)^\top$ denotes the transpose of a matrix.

Equation (13) is not the eigendecomposition of \mathbf{C}_S , and $\boldsymbol{\Lambda}_S$ is not necessarily diagonal, as it would be for \mathbf{C} . Instead, $\boldsymbol{\Lambda}_S$ provides a measure of the discrepancy between \mathbf{C}_S and \mathbf{C} , and by computing quantiles of the diagonal elements’ distribution from a set of $\boldsymbol{\Lambda}_S$, we derive significance levels for each eigenvalue λ_k ([Allen and Smith, 1996](#)).

The resulting significance levels for the SSA analysis of GDP in Figure 2a are indicated as vertical bars and we see that all eigenvalues fall within these error bars. Hence, the observed spectrum of eigenvalues cannot falsify the null hypothesis either. In the following subsection, we show that additional information from other macroeconomic indicators helps reject the null hypothesis.

Table 1: NULL-HYPOTHESIS PARAMETERS

Variable	AR(1) parameters		Time lag behind GDP
	a_d	σ_d	Δ_d (quarters)
GDP	0.82	0.040	0
Investment	0.92	0.028	1
Employment	0.96	0.021	0
Consumption	0.88	0.034	0
Total wage	0.95	0.024	0
Δ (Inventories)*	0.61	0.055	0
Price	0.99	0.012	7
Exports	0.95	0.023	1
Imports	0.83	0.039	1

* Changes in inventories

3.2. Multivariate time series

To test significance in M-SSA results, comovements should be taken into account in formulating the null hypothesis. Vector AR models, however, may support oscillations even for order one; when present, these are referred to as principal oscillation patterns (von Storch et al., 1995; Penland and Matrosova, 2001). We keep, therefore, the idea of fitting univariate AR(1) processes

$$X_d(t) = a_d X_d(t-1) + \sigma_d \epsilon_d(t), \quad (14)$$

but build characteristics of the cross-correlations into the null hypothesis. The estimated parameters for each indicator are listed in Table 1.

The cross-correlation information is included by coupling the noise residuals $\epsilon_d(t)$ at a certain temporal lag Δ_d . Relative to GDP, denoted by $x_1(t)$, we chose Δ_d so as to maximize the correlations between $x_1(t)$ and $x_d(t + \Delta_d)$. Doing so is especially necessary for the price, for which we observe a correlation maximum at seven quarters (cf. Figure 3g).

The covariance matrix \mathbf{R} for innovation processes $\epsilon_d(t)$ has elements

$$R_{d,d'} = \frac{1}{\tilde{N}_{d,d'}} \sum_{t=\max\{1,1+\Delta_{d'}-\Delta_d\}}^{\min\{N,N+\Delta_{d'}-\Delta_d\}} x_d(t)x_{d'}(t+\Delta_{d'}-\Delta_d); \quad (15)$$

the denominator $\tilde{N}_{d,d'}$ depends on the range of summation, namely $\tilde{N}_{d,d'} = \min\{N, N + \Delta_{d'} - \Delta_d\} - \max\{1, 1 + \Delta_{d'} - \Delta_d\} + 1$. Cholesky decomposition yields $\mathbf{R} = \mathbf{L}^\top \mathbf{L}$ and we derive correlated innovation processes from

$$(\epsilon_1(t), \dots, \epsilon_d(t + \Delta_d), \dots, \epsilon_D(t + \Delta_D))^\top = \mathbf{L}^\top (\xi_1(t), \dots, \xi_d(t), \dots, \xi_D(t))^\top, \quad (16)$$

with the ξ_d 's being independent white-noise processes. Finally, we pass the realizations through the HP filter and normalize it to the same standard deviation as the data set.

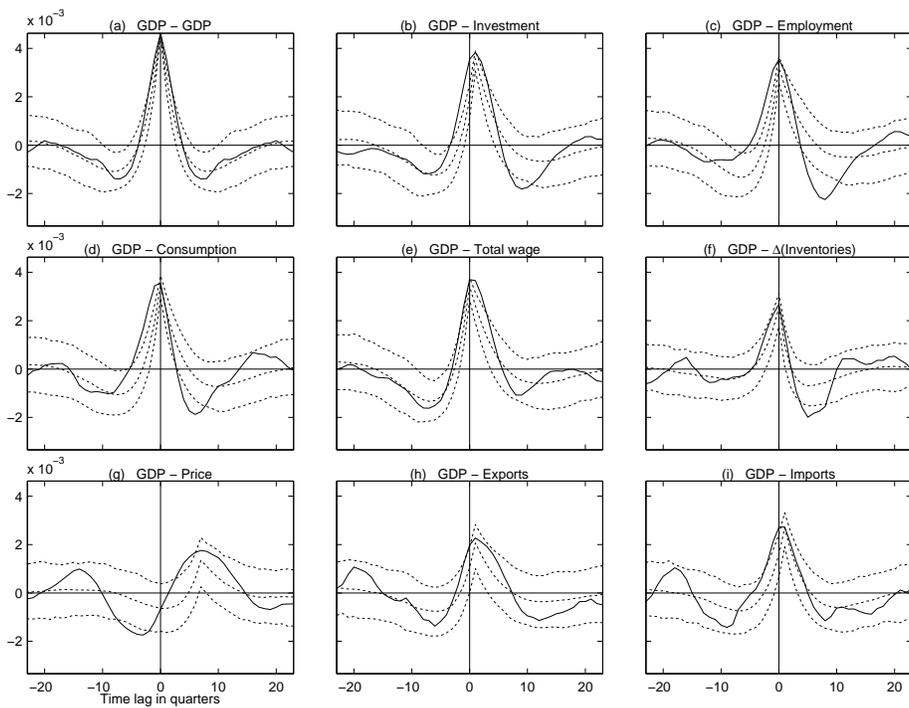


Figure 3: Auto- and cross-covariance functions of the nine U.S. economic indicators with respect to GDP (solid lines). Dashed lines indicate the significance levels 2.5% and 97.5%, as well as the median from the realization of 1000 surrogate time series. Panels (a)–(i) are labeled directly in the figure; $\Delta(\text{Inventories})$ in the legend of panel (f) indicates changes in inventories.

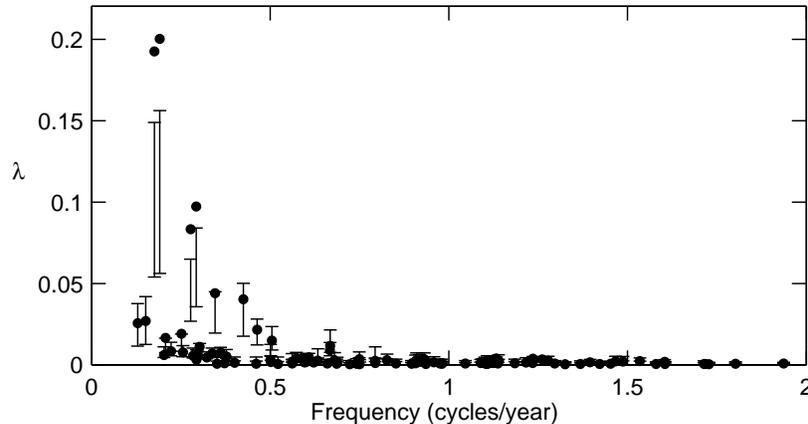


Figure 4: Spectrum of M-SSA eigenvalues (non-dimensional, filled circles) using all nine U.S. indicators, with $M = 24$. The error bars indicate the significance levels, derived from the 2.5% and 97.5% percentiles of 1000 multivariate surrogate time series.

Figure 3 compares the covariance of the data with that of the null hypothesis. Lead-lag relations among economic indicators are reproduced by the null-hypothesis model, and the data covariance function lies almost, but not quite, within the variability of the null hypothesis. As in the univariate case, the HP filter introduces a swing below zero and which would lead once again to spurious cycles.

Projecting the grand covariance matrix $\tilde{\mathbf{C}}_S$ for each of the surrogate realizations onto the data eigenvectors, $\tilde{\mathbf{\Lambda}}_S = \tilde{\mathbf{E}}^T \tilde{\mathbf{C}}_S \tilde{\mathbf{E}}$, yields again a measure of the discrepancy between $\tilde{\mathbf{C}}_S$ and $\tilde{\mathbf{C}}$, which allows us to derive significance levels for λ_k from the distribution of the diagonal elements of $\tilde{\mathbf{\Lambda}}_S$ (Figure 4).

As in the case of GDP alone, we observe higher significance levels near a five-year period, but this time the data eigenvalues clearly exceed the upper significance level. Hence the high variance associated with the leading pairs of eigenvalues can no longer be explained by spurious cycles induced by inappropriate detrending.

We have further tested the consistency of the present results by using different values of the window width, namely $M = 20, 30, 40$ and 50 . It turns out that the leading pair describes a significant oscillation of five-year period. The three-year oscillation in the second pair is somewhat less robust, but probably still deserves further examination in future work.

We have performed additional, exhaustive tests—as proposed by [Allen and Smith \(1996\)](#)—to cope with the problem of overestimating large eigenvalues in SSA and have found further evidence that the five-year oscillatory pair is indeed statistically significant at the 95% level. We focus, therefore, in the next section on this oscillatory mode and investigate its role in business cycle dynamics.

4. Business cycle dynamics: phase dependence and evolution

The presence of oscillatory pairs indicates recurrent behavior in the system’s dynamics. Such more-or-less regular recurrences are typically produced by the presence of an attracting or only weakly unstable limit cycle in the dynamics (Vautard and Ghil, 1989; Vautard et al., 1992; Ghil et al., 2002). On the other hand, vector AR(1) processes can possess oscillatory solutions as well, as mentioned at the beginning of Section 3.2. In practice, however, it might be hard to distinguish between purely deterministic, but chaotic oscillators and stochastically driven ones.

From the point of view of economic theory, the distinction is not that important: in fact, deterministic mechanisms—whether linear or nonlinear—are the only ones that give rise to cyclic behavior in a vector AR(1) process or in a vector random process generated by a system of linear (Arnold, 1974) or nonlinear (Arnold, 1998) stochastic differential equations as well (Schuss, 1980); for illustration purposes, we discuss the case of a vector AR(1) process in the Appendix. The stochastic forcing, if present, only contributes truly irregular fluctuations. It is, therefore, only the deterministic part of the dynamics that is of genuine interest in discussing cyclicity in macroeconomics, whether stochastic forcing is present or not.

The first two eigenvalues capture 40% of the total variance, and RCs 1-2 give already a good approximation of the GDP, cf. Figure 5a. To better understand the role of this five-year oscillatory mode in the processes of expansion and recession, we study the temporal evolution of its variance.

Plaut and Vautard (1994) introduced the concept of *local variance fraction* $V_{\mathcal{K}}(t)$,

$$V_{\mathcal{K}}(t) = \frac{\sum_{k \in \mathcal{K}} A_k(t)^2}{DM \sum_{k=1} A_k(t)^2}, \quad (17)$$

which quantifies the fraction of the total variance that is described by a subset \mathcal{K} of PCs in a window of length M . We consider the PCs as centered, i.e. starting at $M/2$, cf. Section 2.2.

Figure 5 shows this index, along with the NBER-defined recessions, for the leading PCs 1-2 in panel (b) and for PCs 3-150 in panel (c). The sum of PCs 1-150 capture 99% of the total variance; see dash-dotted line in panel (b). Starting after 1980, it is quite remarkable that the fraction of the five-year oscillatory mode in PCs 1-2 is high during recessions and low during expansions. It shows that during the recessions, the trajectory of the data set stays closer to a suspected five-year limit cycle—like the one in the Non-Equilibrium Dynamic Model (NEDyM) of Hallegatte and Ghil (2008) or in other endogenous business cycle models (Chiarella et al., 2005)—while this trajectory reveals more complex behavior during expansions.

During the 1970s, PCs 1-2 capture roughly 50% of the variance or more over the full decade, while from 1980 on, PCs 1-2 play a significant role only during recessions. This result of our analysis suggests a change in the system’s dynamics

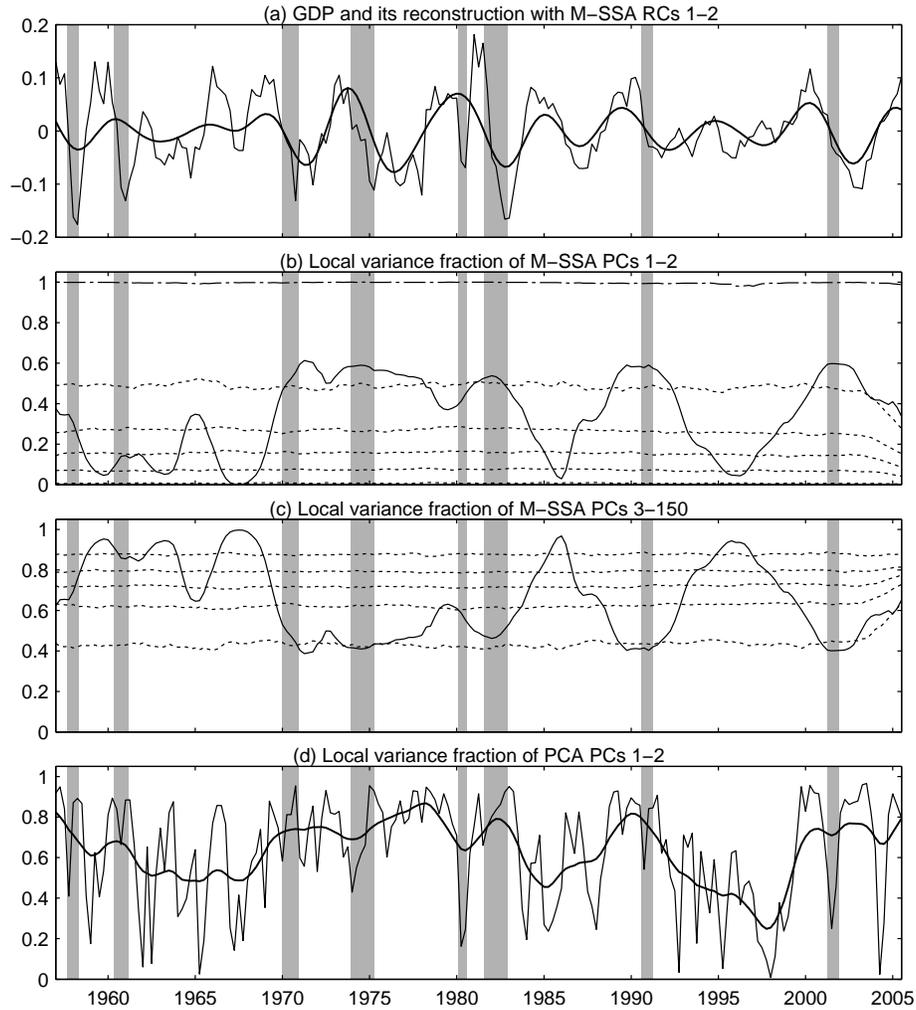


Figure 5: (a) Pre-processed GDP data set (light solid line) and its reconstruction with RCs 1–2 of our M-SSA analysis (heavy solid line). (b,c,d) Local variance fraction $V(t)$: (b) for M-SSA PCs 1–2 (solid line) and PCs 1–150 (near-total variance, dash-dotted line); (c) for M-SSA PCs 3–150 (solid line); and (d) for PCs 1–2 of a PCA analysis (light solid line), as well as after smoothing with a two-year moving average (heavy solid line). The dashed lines in panels (b) and (c) give the 2.5%, 25%, 50%, 75%, and 97.5% percentiles based on 1000 surrogate time series. The shaded vertical bars indicate NBER-defined recessions.

in the 1980s, a change that agrees in timing with the “great moderation,” during which volatility in GDP growth diminished markedly (Kim and Nelson, 1999; McConnell and Perez-Quiros, 2000; Stock and Watson, 2002; Kim et al., 2004).

There has been considerable debate on the cause of this shift, as well as on the expected duration of the U.S. economy’s new mode of functioning; in particular it has been proposed that this moderate behavior terminated in 2007, i.e. before and during the “great recession” of 2008–2009. In any case, our results are at least consistent with the hypothesis of structural changes in the 1980s, and our M-SSA methodology can help provide sophisticated analysis tools to determine whether and when the great moderation ended, once additional BEA data become available.

To assess the significance of the local variance results, we compare their variability with that of the null hypothesis in Eq. (14). To wit, we project each surrogate realization onto the data eigenvectors $\tilde{\rho}_k$ to obtain surrogate PCs in the same way as for the data set in Eq. (10). As in the significance test of the eigenvalues, the resulting surrogate PCs are not orthogonal, since their covariance matrix $\tilde{\Lambda}_S$ is not diagonal.

For each surrogate PC we calculate the local variance fraction in the same way as for the data set in Eq. (17), and derive time-dependent significance levels from the distribution of $V_{\mathcal{K}}(t)$ (Figs. 5b,c, dashed lines). Since the AR(1) processes are stationary, these levels are supposed to be constant; this stationarity is seen in fact in Figure 5, except near the end of the time series, i.e. starting at $t \simeq N - M$.

In contrast to the approximate constancy of $V_{\mathcal{K}}(t)$ in the null hypothesis, the five-year oscillatory mode in the data set exhibits much greater variance during the recessions, when it does exceed the 97.5% significance level. The variance in PCs 3–150 is also larger than can be explained by the null hypothesis, with $V_{\mathcal{K}}(t)$ values that are significantly larger than the 97.5% percentile during expansions and smaller than the 2.5% percentile during recessions, respectively.

We have further examined the variability of $V_{\mathcal{K}}(t)$ during the whole 1954–2005 interval by using other quantities, such as standard deviation and interquartile range (not shown here). All these estimates confirm that the U.S. macroeconomic indicators exhibit larger variability than can be explained by the random fluctuations of the null hypothesis.

It is interesting to note that a similar phenomenon can also be identified by applying PCA to the data (Figure 5d). Although, at first glance, the local variance fraction of the leading two PCs of PCA (light solid line) is rather irregular, with no apparent link to the business cycle, smoothing with a two-year moving-average filter (heavy solid line) does indeed produce a behavior comparable to that in Figure 5b. It would, however, be difficult to guess that from the raw PCA results: the moving-average filtering was only inspired by the M-SSA results in panels (b) and (c), which did not require any additional post-processing.

5. Concluding remarks

In this article, we applied multivariate singular spectrum analysis (M-SSA) to study business cycles dynamics in a consistent and multivariate way. Our M-SSA results allowed us to reconcile and combine the NBER definition of recessions with quantitative analysis. This methodology extends, in particular, the recently proposed SSA business cycle analysis of a single macroeconomic indicator (de Carvalho et al., 2012; Sella and Marchionatti, 2012; Dumas et al., 2013). The present M-SSA analysis uses information on nine U.S. macroeconomic indicators from the BEA for 52 years (1954–2005).

This extended business cycle analysis leads to three major conclusions: (i) the presence of genuine periodicity in macroeconomic behavior and its deterministic causes; (ii) the essential role of comovements of economic aggregates in the proper definition of business cycles; and (iii) the dependence of economic ‘volatility’ on the phase of the business cycle. We describe these conclusions in greater detail below.

Genuine periodicity and its deterministic causes. 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 Slutsky (1937) theory of random shocks. In the simplest RBC models, cyclicity originates exclusively from productivity shocks that can be modeled by a simple random walk. Cogley and Nason (1995) have, moreover, argued that the spurious appearance of business cycle dynamics can be generated by the HP filter even if none is present, even in a random walk.

In agreement with the findings of Cogley and Nason (1995), a simple univariate analysis of GDP does not reveal any significant oscillatory modes (see Figure 2). Our multivariate analysis, however, uses a larger amount of information about macroeconomic behavior, and allows us to identify a five-year oscillatory mode with high statistical confidence (see Figure 4). This mode cannot be explained by artificial effects due to detrending by the HP filter, and a random-walk-driven model of business cycles has to be questioned in the light of the results obtained in our paper.

A major result of our study thus points to the presence of deterministic, endogenous effects in the business cycles of the U.S. economy and leads to the conclusion that business cycles cannot be explained by exogenous shocks alone. This conclusion does not exclude the importance of random effects in the economy, as discussed at the beginning of Section 4 and in the Appendix: it only emphasizes the role of the deterministic ones in giving rise to cyclic behavior.

Comovements of macroeconomic aggregates. The role of the additional information provided by the M-SSA analysis emphasizes the need to understand business cycles as a phenomenon that is not limited to GDP variations, but involves all aspects of the economy; it is reflected, therefore, in the comovements of several macroeconomic aggregates. In the present study, we have performed a quantitative analysis of the BEA data set that is consistent with the NBER definition of the business cycle, inasmuch as it is entirely multivariate and takes into

account the lead-and-lag relationships between the various indicators present in the data (see Table 1 and Figure 3). These lead-and-lag relationships, in turn, reflect certain widely acknowledged “stylized facts” of economic cycles (Kydland and Prescott, 1998; Zarnowitz, 1985).

State-dependent fluctuations. M-SSA also allowed us to provide further insight into the underlying macroeconomic dynamics, and especially into the crucial question of the complex interplay between endogenous dynamics and exogenous shocks. We showed that the U.S. economy changes its behavior from one phase of the business cycle to another: the recession phase is dominated by the five-year mode, while the expansion phase exhibits more complex dynamics, with higher-frequency modes coming into play (see Figure 5). This type of behavior cannot be explained by the random fluctuations that drive a simple stationary RBC model, in the absence of endogenous oscillatory dynamics.

It thus appears that the dynamics of the U.S. economy can indeed be decomposed into a five-year cycle and more complex, higher-frequency behavior superimposed on this cycle. The amplitude of the latter, irregular component is higher during expansions, i.e. the business cycle is more volatile during expansions than during recessions.

This phase-dependent volatility is consistent with the response to natural disasters predicted by Hallegatte and Ghil (2008) in an endogenous business cycle (EnBC) model. These authors have shown that exogenous shocks, whether positive or negative, are likely to have a bigger impact in the presence of EnBCs.

Their modeling framework (Hallegatte et al., 2008), while highly simplified, has Keynesian features and their EnBC model’s predictions can be interpreted in terms of production being closer to capacity during expansions. On the other hand, the lower volatility during recessions in our analysis here is consistent with the predicted reduction in sensitivity to exogenous shocks, due to underutilized production capacity in the low phases of an EnBC. This so-called “vulnerability paradox” was also highlighted by Ghil et al. (2011) and by Dumas et al. (2013).

Such a variable-volatility pattern may seem at odds with the findings of French and Sichel (1993). But this apparent discrepancy can be explained by the fact that our analysis and theirs are not defining the fluctuations in the same manner. French and Sichel (1993) modeled the variance of the residuals on a long-term trend, without decomposing these residuals into cyclical and non-cyclical behavior, as we do here, and found higher variance during epochs of recession. In the present paper, we study the fluctuations superimposed on the sum of the long-term trend, plus a possible cyclical component. It is the variance of the fluctuations so defined that is largest during expansions.

The next step in our research program is to investigate whether the change in the economy’s dynamical behavior between boom and bust also leads to different types of response to exogenous shocks. This question is fundamental in attempting to evaluate the efficiency of economic policy in different phases of the business cycle.

Acknowledgments

It is a pleasure to thank J.-C. Hourcade for informative and stimulating discussions and J. H. Stock for detailed and constructive comments on an earlier version of the paper. This work was supported by a CNRS post-doctoral fellowship to AG, the Groupement d'Intérêt Scientifique (GIS) Réseau de Recherche sur le Développement Soutenable (R2DS) of the Région Ile-de-France, and the Chaire Développement Durable of the Ecole Polytechnique.

Appendix. Cyclicity in deterministic and stochastic models

It is well known, as already mentioned at the beginning of Section 3.2, that vector AR(1) processes can possess oscillatory solutions, due to the presence of pairs of complex conjugate eigenvalues $(\lambda_k, \lambda_{k+1}) = (\lambda_k^{(r)} \pm \lambda_k^{(i)})$ in the spectrum of the matrix $\mathbf{A} = (a_{ij})$ that characterizes such a process,

$$\mathbf{X}(t) = \mathbf{A} \mathbf{X}(t-1) + \mathbf{\Sigma} \epsilon(t); \quad (\text{A.1})$$

here $\mathbf{\Sigma}$ is a covariance matrix multiplying the noise vector ϵ . For a stationary AR(1) process, all the real parts $\lambda_k^{(r)}$ of the eigenvalues of \mathbf{A} must be negative, and the damped oscillations are maintained at a statistically constant amplitude by the noise ϵ .

In practice, such a stochastically driven oscillator might be hard to distinguish from a purely deterministic, possibly chaotic one. But the term $\mathbf{A} \mathbf{X}(t-1)$ in the former case, on the right-hand side of Eq. (A.1), still captures the manifestation of a coupled pair of deterministic feedbacks, one positive, the other negative, whether linear or nonlinear, noise-driven or not, as the ultimate cause of any complex conjugate pair of eigenvalues $(\lambda_k^{(r)} \pm \lambda_k^{(i)})$.

References

- Allen, M. R., Robertson, A. W., 1996. Distinguishing modulated oscillations from coloured noise in multivariate datasets. *Climate Dynamics* 12 (11), 775–784.
- Allen, M. R., Smith, L. A., 1996. Monte Carlo SSA: Detecting irregular oscillations in the presence of colored noise. *Journal of Climate* 9, 3373–3404.
- Arnold, L., 1974. *Stochastic Differential Equations: Theory and Applications*. J. Wiley.
- Arnold, L., 1998. *Random Dynamical Systems*. Springer.
- Baxter, M., King, R. G., 1999. Measuring business cycles: Approximate band-pass filters for economic time series. *Review of Economics and Statistics* 81 (4), 575–593.

- Blackman, R. B., Tukey, J. W., 1958. *The Measurement of Power Spectra*. Dover.
- Broomhead, D. S., King, G. P., 1986a. Extracting qualitative dynamics from experimental data. *Physica D* 20 (2-3), 217–236.
- Broomhead, D. S., King, G. P., 1986b. On the qualitative analysis of experimental dynamical systems. In: Sarkar, S. (Ed.), *Nonlinear Phenomena and Chaos*. Adam Hilger, Bristol, England, pp. 113–144.
- Burns, A. F., Mitchell, W. C., 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., Flaschel, P., Franke, R., 2005. *Foundations for a Disequilibrium Theory of the Business Cycle*. Cambridge: Cambridge University Press.
- Cogley, T., Nason, J. M., 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.
- de Carvalho, M., Rodrigues, P. C., Rua, A., Jan. 2012. Tracking the US business cycle with a singular spectrum analysis. *Economics Letters* 114 (1), 32–35.
- Dumas, P., Ghil, M., Groth, A., Hallegatte, S., 2013. Dynamic coupling of the climate and macroeconomic systems. *Mathematics of the Social Sciences*, in press.
- French, M. W., Sichel, D. E., Jan. 1993. Cyclical patterns in the variance of economic activity. *Journal of Business & Economic Statistics* 11 (1), 113–119.
- Ghil, M., Allen, M. R., Dettinger, M. D., Ide, K., Kondrashov, D., Mann, M. E., Robertson, A. W., Saunders, A., Tian, Y., Varadi, F., Yiou, P., 2002. Advanced spectral methods for climatic time series. *Reviews of Geophysics* 40 (1), 1–41.
- Ghil, M., Vautard, R., 1991. Interdecadal oscillations and the warming trend in global temperature time series. *Nature* 350 (6316), 324–327.
- Ghil, M., Yiou, P., Hallegatte, S., Malamud, B. D., Naveau, P., Soloviev, A., Friederichs, P., Keilis-Borok, V., Kondrashov, D., Kossobokov, V., Mestre, O., Nicolis, C., Rust, H., Shebalin, P., Vrac, M., Witt, A., Zaliapin, I., 2011. Extreme events: Dynamics, statistics and prediction. *Nonlinear Processes in Geophysics* 18, 295–350.
- Graff, M., 2011. International business cycles: How do they relate to Switzerland? Tech. Rep. 291, KOF Swiss Economic Institute, ETH Zurich Working Paper.

- Granger, C. W. J., 1966. The typical spectral shape of an economic variable. *Econometrica* 34 (1), 150–161.
- Hallegatte, S., Ghil, M., 2008. Natural disasters impacting a macroeconomic model with endogenous dynamics. *Ecological Economics* 68 (1–2), 582–592.
- Hallegatte, S., Ghil, M., Dumas, P., Hourcade, J.-C., 2008. Business cycles, bifurcations and chaos in a neo-classical model with investment dynamics. *Journal of Economic Behavior & Organization* 67 (1), 57–77.
- Harvey, A. C., Jaeger, A., 1993. Detrending, stylized facts and the business cycle. *Journal of Applied Econometrics* 8 (3), 231–247.
- Hodrick, R. J., Prescott, E. C., 1997. Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit and Banking* 29 (1), 1–16.
- Karhunen, K., 1946. Zur Spektraltheorie stochastischer Prozesse. *Annales Academiæ Scientiarum Fennicæ. Ser. A1, Math. Phys.* 34.
- Keppenne, C. L., Ghil, M., 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.
- Kim, C.-J., Nelson, C. R., 1999. Has the U.S. economy become more stable? a Bayesian approach based on a Markov-switching model of the business cycle. *Review of Economics and Statistics* 81, 608–616.
- Kim, C.-J., Nelson, C. R., Piger, J., Jan. 2004. The less-volatile U.S. economy: A Bayesian investigation of timing, breadth, and potential explanations. *Journal of Business & Economic Statistics* 22 (1), 80–93.
- Kimoto, M., Ghil, M., Mo, K.-C., 1991. Spatial structure of the extratropical 40-day oscillation. In: Eighth Conference on Atmospheric and Oceanic Waves and Stability (Denver, Colo.). American Meteorological Society, Boston, Mass., pp. 115–116.
- Kydland, F. E., Prescott, E. C., 1998. Business cycles: Real facts and a monetary myth. In: James E. Hartley, Kevin D. Hoover, K. D. S. (Ed.), *Real Business Cycles: A Reader*. Routledge, pp. 231–247.
- Loève, M., 1945. Fonctions aléatoires de second ordre. *Comptes Rendus de l'Académie des sciences Paris* 220, 380.
- Loève, M., 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*. Springer-Verlag, Berlin, pp. 230–242.
- McConnell, M. M., Perez-Quiros, G., 2000. Output fluctuations in the United States: What has changed since the early 1980's? *American Economic Review* 90 (5), 1464–1476.
- Mees, A. I., Rapp, P. E., Jennings, L. S., 1987. Singular-value decomposition and embedding dimension. *Physical Review A* 36 (1), 340–346.
- Morley, J., Piger, J., 2012. The asymmetric business cycle. *Review of Economics and Statistics* 94 (1), 208–221.
- Nelson, C. R., Kang, H., 1981. Spurious periodicity in inappropriately detrended time series. *Econometrica* 49 (3), 741–751.
- Oppenheim, A. V., Schaffer, R. W., 1989. *Discrete-time signal processing*. Prentice-Hall, Inc. Upper Saddle River, NJ.
- Penland, C., Matrosova, L., 2001. Expected and actual errors of linear inverse model forecasts. *Monthly Weather Review* 129, 1740–1745.
- Plaut, G., Vautard, R., 1994. Spells of low-frequency oscillations and weather regimes in the Northern Hemisphere. *Journal of the Atmospheric Sciences* 51 (2), 210–236.
- Priestley, M., 1991. *Spectral Analysis and Time Series*. Academic Press.
- Sauer, T., Yorke, J. A., Casdagli, M., 1991. Embedology. *Journal of Statistical Physics* 65 (3-4), 579–616.
- Schuss, Z., 1980. *Theory and Applications of Stochastic Differential Equations*. J. Wiley.
- Sella, L., Marchionatti, R., 2012. On the cyclical variability of economic growth in Italy, 1881–1913: a critical note. *Cliometrica*, 1–22.
- Slutzky, E., 1937. The summation of random causes as a source of cyclic processes. *Econometrica* 5, 105–146.
- Solow, R., 1956. A contribution to the theory of economic growth. *Journal of Economics* 70(1), 65–94.
- Stock, J. H., Watson, M. W., 2002. Has the business cycle changed and why? In: Gertler, M., Rogoff, K. S. (Eds.), *NBER Macroeconomics Annual*. MIT Press, Cambridge, Mass., pp. 159–218.
- Takens, F., 1981. Detecting strange attractors in turbulence. In: *Dynamical Systems and Turbulence*. Vol. 898 of *Lecture Notes in Mathematics*. Springer, Berlin, pp. 366–381.

- Vautard, R., Ghil, M., 1989. Singular spectrum analysis in nonlinear dynamics, with applications to paleoclimatic time series. *Physica D* 35 (3), 395–424.
- Vautard, R., Yiou, P., Ghil, M., 1992. Singular-spectrum analysis: A toolkit for short, noisy chaotic signals. *Physica D* 58 (1-4), 95–126.
- von Storch, H., Bürger, G., Schnur, R., von Storch, J.-S., 1995. Principal oscillation patterns : a review. *Journal of Climate* 8 (3), 377–400.
- 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.