Extracting a robust U.S. business cycle using a time-varying multivariate model-based bandpass filter

Drew Creal*
Booth School of Business, University of Chicago

Siem Jan Koopman
Department of Econometrics, Vrije Universiteit Amsterdam

Eric Zivot
Department of Economics, University of Washington

June 2009

Abstract

We develop a flexible business cycle indicator that accounts for potential time-variation in macroeconomic variables. The coincident economic indicator is based on a multivariate trend-cycle decomposition model and is constructed from a moderate set of U.S. macroeconomic time series. In particular, we consider an unobserved components time series model with a common cycle that is shared across different time series but adjusted for phase shift and amplitude. The extracted cycle can be interpreted as a model-based bandpass filter and is designed to emphasize the business cycle frequencies that are of interest to applied researchers and policymakers. Stochastic volatility processes and mixture distributions for the irregular components and the common cycle disturbances enable us to account for the heteroskedasticity present in the data. Forecasting results are presented for a set of different specifications. Point forecasts from the preferred model indicate a future recession with the uncertainty over the business cycle growing quickly as the forecast horizon increases.

Keywords: Bandpass filter; Markov chain Monte Carlo; Stochastic Volatility, Trend-cycle decomposition; Unobserved components time series model

JEL Classification: C11; C32; E32.

*Corresponding author: Drew Creal, University of Chicago Booth School of Business, 5807 S. Woodlawn Ave., Chicago, IL 60637, email: dcreal@chicagobooth.edu, tel: 773.834.5249

1 Introduction

The existence of common cyclical movements across macroeconomic time series has been a common theme of business cycle research since Burns and Mitchell (1946). Their contributions inspired future researchers to build business cycle indicators that would be valuable for economic policy and financial decision making. The techniques used to construct the indicators vary widely across the literature. The leading and coincident indicators of Stock and Watson (1989, 1991), Diebold and Rudebusch (1996), and Kim and Nelson (1998) are factor models containing a common business cycle component. Valle e Azevedo et al. (2006) and Harvey et al. (2007) are recent papers which adopt the traditional structural time series approach based on a multivariate trend-cycle model. Closely related to the former methods are the indicators of Forni et al. (2000), Forni et al. (2001), Stock and Watson (2002a, 2002c), and Altissimo et al. (2007), which are built from a combination of principal components analysis and factor models. Finally, some researchers prefer to apply nonparametric filters rather than explicitly formulating a model for the data. Given the *a priori* assumption that a business cycle exists and that its period is between 1.5 to 8 years, interest centers on building an ideal bandpass filter. An ideal bandpass filter has a spectral gain function that carves out exactly the movements in a frequency range specified by the researcher. Baxter and King (1999) and Christiano and Fitzgerald (2003) constructed univariate bandpass filters that approximate an ideal bandpass filter, while Valle e Azevedo (2007) extends their approximate ideal bandpass filter to a multivariate setting.

Concurrent research documents sizeable changes in the volatility of U.S. macroeconomic time series; e.g., see Kim and Nelson (1999), McConnell and Pérez-Quirós (2000), Blanchard and Simon (2001), Stock and Watson (2002b), and Sensier and van Dijk (2004). Most of the evidence from this literature suggests a sizeable reduction in volatility for many series; many of them used to construct business cycle indicators. With the exception of the beginning and the end of the series, the gain function of a lowpass filter such as the Hodrick and Prescott (1997) filter or a bandpass filter remains constant through time. Consequently, the estimates provided by these filters do not account for the great moderation. The methodological contribution of this paper is the construction of a business cycle indicator that has bandpass filter properties, accounts for time varying volatility and utilizes data sampled at different frequencies. Our indicator is constructed from the multivariate unobserved components time series model of Valle e Azevedo et al. (2006), where we extend the model to include stochastic volatility in both the common cycle and irregular components of the model. This enables us to account for the heteroskedasticity present in the data. The stochastic volatility in the model can account for both a gradual decline and a sudden change in volatility, see Blanchard and Simon (2001) and Stock and Watson (2002b) for the two respective cases. The trend components are flexible stochastic functions of random walk processes. The common cycle is a higher-order stochastic cycle formulated by Harvey and Trimbur (2003) which ensures that the extracted business cycle has bandpass filter properties. The adopted trend and cycle formulations in our model originate from the work of Gómez (1999, 2001) who made connections between lowpass/bandpass filters and unobserved components models. Following these authors, we interpret the methods in this paper as a model-based bandpass filter. As the business cycle indicator is a product of a bandpass filter, the paper takes as a working assumption the existence of the business cycle. The modeling framework also allows time series that are coincident, leading, or lagging the business cycle to be included in the analysis by employing the phase shift methodology of Rünstler (2004). Finally, we introduce mixture distributions for the innovations of the trend component and for the stochastic volatility processes. Although the coefficients of the mixture distributions are given known values, the specification remains sufficiently flexible and appears to be robust to other values for the coefficients and for aberrant observations.

Whereas Valle e Azevedo et al. (2006) estimate the parameters of their model by maximum likelihood, we use Bayesian methods for inference. In particular, we use Markov chain Monte Carlo (MCMC) methods for the estimation of all parameters including the trend, cycle, irregular, and stochastic volatility components. In this respect, our paper complements the work of Kim and Nelson (1998), Chauvet and Potter (2001), and Harvey et al. (2007) who use Bayesian methods for estimating business cycle indicators within a state space framework. Our work differs from theirs in several respects. Kim and Nelson (1998) and Harvey et al. (2007) do not account for time-varying volatility. Although Chauvet and Potter (2001) account for changes in volatility through a one-time structural break in the common cycle component, they do not do so in the idiosyncratic component nor does their cycle component isolate business cycle frequencies.

We discuss the results of an empirical analysis based on eleven U.S. macroeconomic time series from 1953 through September 2007. The empirical results reveal that unemployment and inflation are lagging the business cycle while productivity, manufacturing, and real consumption of nondurables plus services are the series leading the cycle by the most. Our model also includes an irregular component that captures the high-frequency movements in the time series. The empirical results suggest that the irregular component with a stochastic variance captures the majority of the time-varying volatility in the data. The estimated volatility associated with the cycle does not vary significantly over time. Although evidence of the great moderation is widely available, the persistence of the business cycle appears to be constant through time. The modeling framework enables the forecasting of each series and each component simultaneously.

In our empirical analysis, we find that point forecasts from our preferred specification indicate an oncoming recession for 2008 while point forecasts from other more restricted specifications point toward an expansion. The uncertainty over the business cycle does however grow quickly as the forecast horizon increases for all the specifications.

The multivariate trend-cycle model with stochastic volatility and mixture distributions is presented in section 2. In section 3, we describe the data and our priors for the parameters. Section 4 contains our empirical results including the new business cycle indicator and the estimated stochastic volatility components. Section 5 concludes.

2 A Model-Based Bandpass Filter with Stochastic Volatility

We adopt a multivariate class of unobserved components models in which all time series variables have separate trend and irregular components but where they share a common business cycle component. The common cyclical component is designed to extract a specific range of frequencies. Following Gómez (2001) and Harvey and Trimbur (2003), this approach to signal extraction is called a *model-based bandpass filter*. The model framework is sufficiently flexible for the inclusion of time series that either lead, lag or coincide with the business cycle. Furthermore, time-varying heteroskedasticity components are introduced such that the gain function implied by the model adapts over time to the time-varying volatility.

2.1 The model without stochastic volatility and mixture innovations

We develop a multivariate unobserved components time series model that decomposes an $M \times 1$ vector time series into trend, cycle, and irregular components. The tth observation for the ith macroeconomic variable is denoted by y_{it} for $i = 1, \ldots, M$ and $t = 1, \ldots, n$. The ith measurement equation of the model is given by

$$y_{it} = \tau_{it} + \delta_i \psi_t^{(q)} + \varepsilon_{it}, \qquad i = 1, \ldots, M, \qquad t = 1, \ldots, n, \tag{1}$$

where τ_{it} and ε_{it} are the idiosyncratic trend and irregular components, respectively, for the ith variable. The cycle component $\psi_t^{(q)}$ is specified as a smooth cyclical process where q is an integer for the level of smoothness. The cycle is common to all series and scaled for each series by the coefficient δ_i. The stochastic specifications of the three different components are given next.

The individual trend component is modeled as a smooth local linear trend process $\Delta^2 \tau_{it} = \zeta_{it}^*$ where Δ is the difference operator ($\Delta x_t = x_t - x_{t-1}$) and ζ_{it}^* is a disturbance term. Alternatively, we can represent this trend specification as

$$\tau_{i,t+1} = \tau_{it} + \beta_{it}, \tag{2}$$

$$\beta_{i,t+1} = \beta_{it} + \zeta_{it}, \qquad \zeta_{it} \sim \mathcal{NID}\left(0, \sigma^2_{i,\zeta}\right), \tag{3}$$

where β_{it} can be interpreted as the growth or slope term of the trend τ_{it} while it follows that $\zeta^*_{it} = \zeta_{i,t-2}$. The irregular component ε_{it} in (1) has mean zero, variance $\sigma^2_{i,\varepsilon}$ and is normally distributed, that is

$$\varepsilon_{it} \sim \mathcal{NID}\left(0, \sigma^2_{i,\varepsilon}\right). \tag{4}$$

The disturbances ζ_{it} and $\varepsilon_{i't'}$ are serially and mutually uncorrelated at all times $t, t' = 1, \ldots, n$ and for all variables $i, i' = 1, \ldots, M$.

The common cycle $\psi_t^{(q)}$ component is a qth-order stochastic cycle and modeled by

$$\begin{pmatrix} \psi_{t+1}^{(j)} \\ \psi_{t+1}^{+(j)} \end{pmatrix} = \rho \begin{bmatrix} \cos\lambda & \sin\lambda \\ -\sin\lambda & \cos\lambda \end{bmatrix} \begin{pmatrix} \psi_t^{(j)} \\ \psi_t^{+(j)} \end{pmatrix} + \begin{pmatrix} \psi_t^{(j-1)} \\ \psi_t^{+(j-1)} \end{pmatrix}, \qquad j = q, q-1, \ldots, 1, \tag{5}$$

where ρ is the damping parameter with restriction $0 < \rho < 1$ to ensure the stationarity of $\psi_t^{(q)}$ and λ is the frequency of the cycle measured in radians with $2\pi \,/\, \lambda$ as the period of the cycle. The dynamics of the stochastic cycle $\psi_t^{(q)}$ are determined by (5) for $j = q, q-1, \ldots, 1$ and with

$$\begin{pmatrix} \psi_t^{(0)} \\ \psi_t^{+(0)} \end{pmatrix} = \begin{pmatrix} \kappa_t \\ \kappa_t^+ \end{pmatrix}, \qquad \begin{pmatrix} \kappa_t \\ \kappa_t^+ \end{pmatrix} \sim \mathcal{NID}\left(0, \sigma^2_\kappa I_2\right), \tag{6}$$

where I_2 is the 2×2 identity matrix and σ^2_κ is the common variance for both disturbances κ_t and κ_t^+ for $t = 1, \ldots, n$. We refer to Gómez (2001) and Harvey and Trimbur (2003) for a detailed discussion of this specification of the cycle component $\psi_t^{(q)}$. We extend the specifications of the irregular and cyclical components below in section 2.2 by including time-varying variances.

The cycle component in the model is based on the similar stochastic cycle specification of Harvey and Koopman (1997). In a similar stochastic cycle model, the parameters ρ and λ are shared across series. This assumption reduces the number of parameters in the model and may be regarded as reasonable for extracting a business cycle indicator that is common across multiple time series. We emphasize that the methods employed in this paper assume *a priori* that a business cycle exists. In various contributions in the business cycle literature, it is discussed that estimated cycles obtained from lowpass and bandpass filters, but also from signal extraction methods based on a decomposition model such as (1) with a nonstationary trend, can potentially be spurious; see, for example, Nelson (1988) and Cogley and Nason (1995).

For the decomposition model (1) with trend (2), slope (3), irregular (4) and common cycle (5), Harvey and Trimbur (2003) demonstrate that the spectral gain function of the cycle component $\psi_t^{(q)}$ approximates a bandpass filter as the number of cycles q increases. As with the nonparametric filters of Baxter and King (1999), Gómez (1999) and Christiano and Fitzgerald (2003), the cycle is a function of a specific range of frequencies from the observed data. The frequency parameter λ corresponds roughly to the center of the gain function while the damping factor ρ determines its width. Bandpass filters reduce the high-frequency variation that is not considered part of the business cycle and may cause difficulty identifying a turning point. In an unobserved components model with higher-order cycle (5), the irregular component (4) receives most of the high-frequency variation. In previous research carried out by, e.g., Harvey and Trimbur (2003) and Valle e Azevedo et al. (2006), the frequency parameter λ has been fixed. Harvey et al. (2007) refer to fixing λ as one of their motivations to use a Bayesian approach for parameter estimation and signal extraction.

It is often argued that bandpass filters perform poorly at the end of the sample because the estimates of the cycle for the last two or three periods may be heavily revised when new data arrives. Valle e Azevedo et al. (2006) sought to alleviate this problem by combining monthly and quarterly data within the phase shift methodology of Rünstler (2004). Following their work, the model specification (1)–(5) is extended to include phase shifts in the common cycle. Adding the phase shift parameters allows the model to gather cyclical information from other series, particularly leading and lagging series that may be omitted from other indicators. The recent contribution of Altissimo et al. (2007) have adopted a similar approach to improve their estimates at the end of the sample.

The suggested modifications for the common cycle component require the measurement equation (1) for the ith variable to be replaced by

$$y_{it} = \tau_{it} + \delta_i \left\{ \cos(\xi_i \lambda)\, \psi_t^{(q)} + \sin(\xi_i \lambda)\, \psi_t^{+(q)} \right\} + \varepsilon_{it}, \qquad (7)$$

for $i = 1, \ldots, M$ and $t = 1, \ldots, n$, where the coefficient ξ_i determines the extent of the shift of the common cycle for the ith variable. When $\xi_i \geq 0$, the shift is forwards in time while it is backwards otherwise. The phase shift parameters are measured in calendar time according to the highest observed frequency. The parameters δ_i scale the remaining cycles by expanding or contracting the base cycles $\psi_t^{(q)}$ and $\psi_t^{+(q)}$ to fit each series. It is necessary for identification purposes to choose one of the time series' cycles as the base cycle and set $\delta_1 = 1$ and $\xi_1 = 0$ for this series. For convenience, we group the trend, slope, and cyclical components and denote them as $\alpha_{1:n} = \left[\{\tau_{it}, \beta_{it}\}_{i=1}^{M}, \psi_t^{(q)}, \psi_t^{+(q)} \right]_{t=1}^{n}$.

As the model imposes a common cycle (adjusted for phase shift and amplitude) across series, we consider it to be only an approximation of the true dynamics of the joint data generating process. This framework uses the common cycle to pool information from multiple series observed at different frequencies in order to improve estimates of the business cycle. Valle e Azevedo et al. (2006) argue that signal extraction based on the trend-cycle model (2)–(7) is effectively applying a multivariate model-based bandpass filter because it extracts a cycle with bandpass filter properties from multiple time series. The properties of this method can be evaluated by considering it as a filter and by inspecting the gain function of the cycle.

2.2 The model with stochastic volatility and mixture innovations

In this section, we extend the model (2)–(7) to account for recently documented changes in macroeconomic volatility. We therefore allow the variances of the irregular components $\sigma^2_{i,\varepsilon}$ in (4) to vary over time using independent stochastic processes. In particular, we consider the stochastically time-varying variance processes with mixture innovations given by

$$\sigma^2_{i,t,\varepsilon} = \exp(h_{i,t,\varepsilon}), \tag{8}$$

$$h_{i,t+1,\varepsilon} = h_{i,t,\varepsilon} + K_{i,t,\varepsilon}\omega_{i,t,\varepsilon}, \qquad \omega_{i,t,\varepsilon} \sim \mathcal{NID}\left(0,1\right), \tag{9}$$

for $i = 1, \ldots, M$ and $t = 1, \ldots, n$. The ith log-variance $h_{i,t,\varepsilon}$ is modeled by a random walk process whose innovations are a mixture of a standard Gaussian noise sequence and a stochastic indicator variable with known probabilities.

The initial value of the log-variance, that is $h_{i,1,\varepsilon}$, is treated as a diffuse prior (or non-informative) variable. In this model specification, the latent indicator variables $K_{i,t,\varepsilon}$ take two values $k^{(1)}_\varepsilon$ and $k^{(2)}_\varepsilon$, for all $i = 1, \ldots, M$ and $t = 1, \ldots, n$, with prior probabilities $p^{(1)}_\varepsilon$ and $p^{(2)}_\varepsilon = 1 - p^{(1)}_\varepsilon$, respectively. This specification of a stochastic volatility model based on mixture innovations was recently used by Giordani and Kohn (2008). The mixture framework can be designed to reflect the prior belief that changes in the variance structure of macroeconomic time series are reasonably rare. In this case, for example, one can take $k^{(1)}_\varepsilon = 0$ and $k^{(2)}_\varepsilon$ as a small positive value, say 0.1, with probabilities set to $p^{(1)}_\varepsilon = 0.95$ and $p^{(2)}_\varepsilon = 0.05$. The choices of these values can be set differently for each ith time series. Since the specification is flexible by design, we adopt the same k and p values for all indicator variables $K_{i,t,\varepsilon}$ in the log-variance processes $h_{i,t,\varepsilon}$ for $i = 1, \ldots, M$. The parameters k and p can be estimated for sufficiently long time series. However, macroeconomic time series are generally not long or informative enough to estimate these values. Previous authors using this approach have also noted this in their empirical work; see, e.g. Giordani and Kohn (2008). Therefore, we set the coefficients k and p to fixed values in our empirical analysis. These values are given in Table 1. We have found that different values for k and p do not influence our empirical results much. For example, we have estimated the model with $k^{(2)}_\varepsilon$ ranging from 0.0001 to 0.5 and $p^{(2)}_\varepsilon$ from 0.01 to 0.15. While the estimated mean of the business cycle indicator and the volatilities does not appear to be sensitive, the 0.95% highest posterior density intervals will get slightly larger as $k^{(2)}_\varepsilon$ increases.

<INSERT TABLE 1 HERE>

In a similar fashion, the common variance shared by the cyclical components (6) is also allowed to change stochastically over time via

$$\sigma^2_{t,\kappa} = \exp(h_{t,\kappa}), \tag{10}$$

$$h_{t+1,\kappa} = h_{t,\kappa} + K_{t,\kappa}\omega_{t,\kappa}, \qquad \omega_{t,\kappa} \sim \mathcal{NID}(0,1), \tag{11}$$

for $t = 1,\ldots,n$ where the log-variance $h_{t,\kappa}$ follows a random walk process that is independent of all random walk processes $h_{i,t,\varepsilon}$ in (9) for $i = 1,\ldots,M$. The indicator variables $K_{t,\kappa}$ for $t = 1,\ldots,n$ can take on the values $k^{(1)}_\kappa$ and $k^{(2)}_\kappa$ with prior probabilities $p^{(1)}_\kappa$ and $p^{(2)}_\kappa$, respectively, and are independent of the indicator variables $K_{i,t,\varepsilon}$ in (9). The values we use in our empirical work are also given in Table 1. The previous discussion on the coefficients of indicator variable $K_{i,t,\varepsilon}$ also applies to indicator $K_{t,\kappa}$. The disturbance sequence $\omega_{t,\kappa}$ is also independent of all other disturbance sequences in the model.

The stochastic log-variances with mixture innovations complete our specification of the filter for the empirical study regarding the U.S. business cycle presented in section 4. From a preliminary empirical study we have learned that for most macroeconomic time series the trend component (2) – (3) is estimated as a smooth trend function with a very small estimate for $\sigma^2_{i,\zeta}$, with $i = 1,\ldots,M$. Therefore, we adopt a similar strategy as for the log-variances by introducing mixture innovations for the disturbances ζ_{it} associated with the growth part of the trend component $\tau_{i,t}$. We replace the specification for the growth term of the trend $\tau_{i,t}$ in (3) by

$$\beta_{i,t+1} = \beta_{it} + K_{i,t,\zeta}\zeta_{it}, \qquad \zeta_{it} \sim \mathcal{NID}(0,1), \tag{12}$$

where the indicator variables $K_{i,t,\zeta}$ take on the values $k^{(1)}_\zeta$ and $k^{(2)}_\zeta$ with prior probabilities $p^{(1)}_\zeta$ and $p^{(2)}_\zeta$, respectively. This specification enables the model to account for moderate and infrequent changes in the trend. A similar specification for the trend function has been considered by Perron and Wada (2006). We take $k^{(1)}_\zeta = 0$ and $k^{(2)}_\zeta > 0$ with $p^{(1)}_\zeta = 0.95$ and $p^{(2)}_\zeta = 0.05$. The resulting trend can be interpreted as a stochastically evolving piece-wise linear trend. An alternative growth specification is (3) with variance $\sigma^2_{i,\zeta}$ evolving as in (10)–(11). However, specification (12) has proven to be effective for our purposes.

The full model is specified by the equations (2) and (4) – (12). We extend the set of latent variables $\alpha_{1:n}$ (for trend, slope and cycle) by two additional sets of latent variables. The first set is for the indicator variables associated with the M irregular log-variance innovations, the M slope innovations and for the common cycle log-variance innovations which we denote by

$\chi_{1:n} = \left[\{K_{i,t,\varepsilon}\}_{i=1}^{M^*}, \{K_{i,t,\zeta}\}_{i=1}^{M}, K_{t,\kappa} \right]_{t=1}^{n}$. The second set is for the log-variances of the M irregulars and the common cycle innovation which we denote by $\gamma_{1:n} = \left[\{h_{i,t,\varepsilon},\}_{i=1}^{M^*}, h_{t,\kappa} \right]_{t=1}^{n}$.

In applications with stochastic volatility models for time series of financial returns, the log-variance is typically specified as an autoregressive process. The random-walk specification for the $M+1$ log-variance processes in $\gamma_{1:n}$ of our model imposes smoothness on their evolution and reduces the overall number of parameters to estimate. In practice, only a subset, say $M^* < M$, of the measurement equations may require stochastic volatility for the irregular component ε_{it}, while a constant variance may be more appropriate for the remaining series. We discuss these issues in more detail in section 3.3 below.

The motivation to use stochastic volatility instead of a model with discrete breaks in the variances and to allow both irregular and cycle components to have time-varying variances originates from our goal to develop a flexible method for constructing a business cycle indicator. It is also based on earlier contributions in the literature on the great moderation. Some of this literature concentrates on finding a one-time break in U.S. real GDP estimated around 1984, see, for example, Kim and Nelson (1999), McConnell and Pérez-Quirós (2000). In other work, a larger number of time series are analyzed using a plethora of methods, see, for example, Chauvet and Potter (2001), Stock and Watson (2002b), Kim et al. (2004) and Ahmed et al. (2004). Some researchers are also concerned whether one or more breaks exist, see the discussion in Sensier and van Dijk (2004). However, the debate of a sudden versus a gradual decline in volatility addresses a key question in our study as well. Two distinct contributions in this debate are Blanchard and Simon (2001) who consider a slow reduction in volatility and Stock and Watson (2002b) who find a sudden change in volatility (both studies are based on U.S. macroeconomic time series). We believe that the stochastic volatility specification with mixture innovations is an effective way to accomodate changes or breaks in the volatility of each series. More importantly, it gives a flexible method to develop a robust indicator, which is our primary concern in this paper.

Another key question of interest is at what frequencies the break or breaks exist in each macroeconomic time series. Ahmed et al. (2004) analyzed this question in detail by using two different frequency-domain estimators. They found evidence of structural change at both business cycle and high frequencies. The stochastic volatility processes in the irregular and cycle components are intended to capture this possibility.

The number of breaks, relative timing, and underlying causes of the great moderation are still open research questions. Stock and Watson (2002b) reviewed the arguments proposed by many authors concerning the causes. They concluded, based partially on univariate stochastic volatility models, that most of the reduction is due to good luck, i.e. a reduction in the size of structural shocks hitting the economy. This conclusion is shared by Primiceri (2005), who analysed macroeconomic time series using a structural VAR with stochastic volatility. Although our model specification is different and the volatility components are associated with different dynamic features in the data, we also find significant empirical evidence of volatility reductions from the early 1980s onwards.

3 Design of the empirical study and estimation

In section 4 we present and discuss the empirical results based on the model described in sections 2.1 and 2.2, applied to a set of eleven U.S. macroeconomic time series. We carry out a Bayesian analysis based on Markov chain Monte Carlo (MCMC) methods which we discuss briefly in Appendix A. In particular we provide here the details of how we have constructed the data-set and how we have implemented the methods for the general model specification.

3.1 The data-set with missing values

Our analysis includes eleven time series that are commonly used to construct business cycle indicators. All series were taken from the FRED database at the Federal Reserve Bank of St. Louis. The monthly time series are industrial production, unemployment, manufacturing (PMI composite index), real retail sales, and retail sales and food services. The first three monthly series are measured from 1953:M4 to 2007:M9. The retail sales series is available for the period 1953:M4 to 2001:M4 while retail sales and food services is collected from 1992:M1 to 2007:M9. Real GDP, consumer price index inflation (all goods), real consumption of durables plus real fixed investment (these variables are added together and labeled as investment below), real consumption of nondurables plus services (these variables are added together and labeled as consumption below), productivity of the non-farm business sector, and hours of the non-farm business sector are available as quarterly time series for the period 1953:Q4-2007:Q3. The inflation series is constructed following Stock and Watson (2007) who average the three prior months of the monthly index, take logarithms, first differences, and multiply by 400. The unemployment rate has been multiplied by 100. All other series except unemployment are in logarithms and have been multiplied by 100.

The resulting panel of macro-economic time series is unbalanced and consists of a mixture of monthly and quarterly time series. Its treatment requires the handling of missing observations. For example, we include both monthly and quarterly data in our analysis as policy-making decisions typically occur between data releases at a higher than quarterly frequency. The state space framework handles this by treating the months in which quarterly data are unavailable as missing. Other data idiosyncracies can be handled analogously. For example, the monthly real retail sales series started prior to 1953 but was discontinued in 2001. A newly defined series, real retail sales and food services, began in 1992 and continues to the present. We treat this situation by assuming that both series share a common cyclical feature, namely the phase shift parameter ξ_i. Otherwise, the two series have separate trend, slope, and irregular components.

3.2 The prior distribution for the parameters

The full model for the 11 selected time series has a total of 25 parameters which we denote collectively as Θ. The prior distribution $p(\Theta)$ for the parameters are independent of each other except for the phase shifts ξ_i. Our base cycle is industrial production for which we define $\delta_1 = 1$ and $\xi_1 = 0$. An outline of the priors of the parameters is given below. Specific details on the hyperparameters within the prior can be found in an Appendix associated with this paper. We have located this additional material at `http://faculty.chicagobooth.edu/drew.creal/`.

Starting with the parameters of the common cyclical component, we chose a uniform prior for ρ on the interval $[0, 0.99)$ to ensure stationarity. Following Harvey et al. (2007), the frequency parameter λ has a beta distribution. We position the mode equal to $2\pi/60$ while the standard deviation of this prior is set to be wide for a beta distribution at 0.1. The mode implies a business cycle with a period of five years for monthly data. However, we stress that in practice the parameter λ does not correspond exactly to the center of the spectral gain function of the cycle component, see also the discussion in Valle e Azevedo et al. (2006).

Due to the periodicity of the sine and cosine functions of the cycle component, the phase shift parameters must be restricted within the interval $-\frac{1}{2}\pi < \xi_i\lambda < \frac{1}{2}\pi$ to remain identified. We specify a conditional prior $p(\xi_i|\lambda)$ for these parameters. It was selected as a truncated normal distribution with mean zero and standard deviation 2.5 while the left and right truncation points are set equal to $\pm\frac{1}{2}\pi\lambda^{-1}$, respectively. The scale parameters δ_i have normal distributions as priors which are centered at zero with standard deviations equal to 2.0. The priors on ξ_i and δ_i are relatively uninformative and are intended to see whether the information in the data enable the probability mass to move away from the prior mean of zero. An alternative strategy for eliciting informative priors might be to use information from an another source such as European data on similar series. We have not pursued this further. For the variances of the irregular components that do not have stochastic volatility, we adopt standard non-informative inverse gamma priors with the shape and scale parameters set to zero. This selection of priors is the same as in Harvey et al. (2007).

3.3 Final details of the model

For the empirical study in section 4, we consider the full model specification as given by (2) and (4) – (12). The stochastic volatility processes are introduced in section 2.2 for all M irregular components in the model. However, we may only need a subset of $M^* < M$ irregular components that have time-varying variances in the model under consideration. We started by estimating the parameters from the full model specification with stochastic volatility for all M irregulars and the cycle innovations. By inspection of the empirical results, we have concluded that a constant variance was more appropriate for the time series of hours, productivity and unemployment. Finally, we have selected $q = 2$ for the cycle in (5). We point out that it has also been the preferred specification in Harvey and Trimbur (2003) and Harvey et al. (2007).

4 Empirical study for the U.S. business cycle

In this section we present the empirical results of our study for the U.S. business cycle. Given the multivariate nature of the model, both the estimation output in tables and the graphical output is potentially large. We present and discuss below the empirical results based on the multivariate model with cycle shifts and mixture innovations to the slope and stochastic volatility components. Various other extensions of the model are considered and discussed at the end of this section. The main results are presented here while additional tables and graphical output are reported in an online Appendix at the website listed above. The results from a selection of alternative model specifications are reported there as well.

4.1 The business cycle indicator

Figure 1 depicts smoothed estimates of the four components of our base series of industrial production. The estimated business cycle indicator in the top left graph matches the NBER dates well with 2 additional small downturns, one in the late 1960's and one in the mid-1990's. Based on our data-set that covers a period until the end of the third quarter of 2007, the indicator provides evidence that the U.S. economy is on the brink of a recession. In a multivariate framework, the highest posterior density intervals are considerably smaller than for a univariate model. This is due both to the pooling of series which contain a common feature (the cycle) and the addition of stochastic volatility. The cycle component estimates with 95% highest posterior density intervals are pictured in Figure 2. Even though the highest posterior density intervals are relatively small, there still exists considerable uncertainty about the status of the business cycle over the second and third quarters of 2007. We observe from panel (iii) in Figure 1 that the estimated slope component corresponds closely to the growth rate of industrial production.

Smoothed estimates of the irregular component in panel (iv) show evidence of an extensive amount of heteroskedasticity. The irregular or idiosyncratic component is larger in the 1950's and early 1960's. Estimates of the stochastic volatility components presented in section 4.2 below will confirm this observation.

Table 2 compares the posterior mean and standard deviation of the key parameters of the model compared to their prior mean and standard deviation. The estimated phase shift parameters ξ_i imply that inflation lags industrial production by just over a quarter and unemployment lags it by 1 month. In turn, industrial production lags real GDP by 3 to 4 months. Interestingly, Valle e Azevedo et al. (2006) found the opposite relationship between these two variables for the Euro area. The variables productivity, real consumption of nondurables plus services, and manufacturing lead the cycle by the most. The series of real GDP, retail sales, and investment appear to be roughly coincident with one another.

<INSERT FIGURE 1 HERE>

<INSERT FIGURE 2 HERE>

<INSERT TABLE 2 HERE>

4.2 Stochastic volatility estimates

The primary goal of this paper is the development of a monthly business cycle indicator with bandpass-filter properties that accounts for time-varying volatility. However, it is also interesting to compare the estimated stochastic volatility components from our procedure to related findings in the empirical literature on the great moderation. Any conclusions drawn from the estimated components are of course dependent on the method used to extract them; a point that is made clearly by Canova (1998).

Chauvet and Potter (2001) built a factor model with a common component for four times series and provided evidence that there exists a break in the common cycle in 1984. Given two different definitions of the trend, Kim et al. (2004) also concluded that the break in real GDP volatility occurred in the cycle in 1984. The estimated volatility from our common cycle is pictured in Figure 3(i). It contains an increase from 1974 through 1984 with two peaks in 1976 and 1980. This timing agrees approximately with the oil price shocks in the 1970's as well as the U.S. monetary experiment from 1979-1984. These increases appear however insignificant relative to their 95% highest posterior density intervals. The overall dynamics of volatility in the common cycle is small compared to the changes in volatility of the irregular components, which are depicted in the remaining panels of Figures 3 and 4. Another key and clear finding is that volatility of inflation pictured in Figure 3(iv) appears to be increasing from 1995 onward. The increase of volatility in inflation is remarkable because the other series show a moderation or further reduction in volatility after 1995. The irregular component of real GDP in Figure 4(ii) indicates a moderate decline in volatility beginning in 1979 and ending in 1984. Other series such as manufacturing, real consumption of nondurables plus services, and real consumption of durables plus investment suggest a reduction in volatility but at different dates and with different dynamics than for real GDP. These graphs also imply that the decrease in volatility was mostly in the high frequency shocks hitting the economy.

The differences between our estimates and those of Chauvet and Potter (2001) and Kim et al. (2004) are likely due to two sources. Our definition of the cycle intentionally eliminates most of the high-frequency variation from the cycle and separates it into the irregular component. The former papers do not differentiate between business cycle and high-frequency movements. It is also important to note that our procedure forces a common cycle among the series. Periods when this common cycle does not hold exactly may result in a larger irregular component.

<INSERT FIGURE 3 HERE>

<INSERT FIGURE 4 HERE>

4.3 Robustness of our findings

In this section, we investigate the robustness of our main findings discussed in the previous sections. First, we inspect the model to detect possible changes that are not accounted for by time-varying volatility. Then, we estimate a multivariate trend-cycle model without stochastic volatility components followed by a multivariate model with stochastic volatility but without phase shifts. Next, we consider a model with a common trend for real GDP and consumption. Finally, we consider the robustness of the cycle to small changes in the variables used in our analysis.

<INSERT TABLE 3 HERE>

To facilitate model comparisons, we report and discuss further below the log-predictive score (LPS) of Good (1952) as given by

$$LPS = -\frac{1}{n}\sum_{t=1}^{n} \log p(y_t|y_1,\ldots,y_{t-1};\hat{\Theta}), \tag{13}$$

where $\hat{\Theta}$ is the posterior median of the parameter vector Θ and $p(y_t|y_1,\ldots,y_{t-1};\hat{\Theta})$ is the contribution of y_t to the log-likelihood function. The LPS ranks models according to their predictive ability with smaller values being preferred. Values of the LPS for seven different models are reported in Table 3. We have computed estimates of the log-likelihood function in (13) using a particle filter that is implemented as a combination of the auxiliary particle filter of Pitt and Shephard (1999) and the mixture Kalman filter of Chen and Liu (2000). Creal (2009) provides a recent survey of the particle filtering literature with additional references and details of the implementation as adopted here.

<INSERT TABLE 4 HERE>

4.3.1 The model with breaks in first quarter of 1984

It is possible that additional changes in the business cycle have occurred which have not been accounted for in our model. For example, changes in the persistence of macroeconomic time series have also been reported in the literature. To investigate the possible instability of other parameters, we reestimate the model from sections 4.1 – 4.2 conditional on a known break in ρ, λ, and ξ_i, in Q1 1984. Estimation of all parameters and unobserved components continues to be performed jointly. In other words, we do not need to separate the data into sub-periods and do not estimate parameters for the two periods separately.

Columns 2 and 3 of Table 4 contain estimates of the parameters in the model with a break; additional results are available in the online Appendix. A clear and interesting result is that the value of the business cycle period $2\pi/\lambda$ is estimated to increase from 38.81 to 55.22 after 1984. This finding is consistent with the fact that on average periods between business cycles in more recent years are longer than in the 1960s and 1970s. The persistence parameter ρ does not appear to change between the two sub-periods. The estimated values for ρ in both periods correspond to the value reported in Table 2 for the main model. The estimation results from the model with a break do indicate substantial changes in a few of the phase shift parameters. For example, productivity, real GDP, and manufacturing lead industrial production by several more months after 1984. There is also considerably more uncertainty associated with the post-1984 parameter estimates. The standard deviations of the marginal distributions for each of the parameters are significantly larger. This raises the possibility that the relative position of the cycles may be different before and after 1984. Furthermore, the relationships between the series and the common cycle are potentially more unstable afterwards. This is an interesting finding that we leave for future research. The LPS values for this model and for a model with only a break in ρ, in 1984, are reported in Table 3. The values are slightly smaller than the LPS value for the main model. It may indicate that the model with breaks are preferred, although the differences are small.

The estimated cycle from this specification is shown in panel (i) of Figure 5. Although some of the parameters in the model with breaks appear to change in the second half of the sample, it does not appear to have a substantial effect on the estimated business cycle indicator. The indicator shares the same features as the indicator produced from the main model in Figure 2 without the structural break. Estimates of the stochastic volatility components from this specification are also similar to those found in the Figures 3 and 4 and are available in the online Appendix.

<INSERT FIGURE 5 HERE>

4.3.2 The model without stochastic volatility

In section 2.1 we have introduced the model without stochastic volatility. Some parts of the MCMC algorithm for the model (2) – (7) simplify since we do not need to integrate over $\chi_{1:n}$ and $\gamma_{1:n}$. Panel (ii) of Figure 5 provides the estimated cycle component from a model without stochastic volatility. The amplitude of the estimated cycle is larger and the 95% highest probability density intervals are substantially wider. The peaks and troughs continue to match the NBER recession dates reasonably well. The most important difference between this cycle and the estimated cycle from the model with stochastic volatility comes in the last several quarters at the end of the sample. The model without stochastic volatility implies that the economy is at or even slightly below trend during September 2007. As a result, the forecasts from this model are different than for the model with stochastic volatility, see the discussion in section 4.4.

A selection of parameter estimates, for ξ_i, ρ, and λ, for the current model are reported in Table 4 with the remaining values available in the online Appendix. The parameters of this model are estimated to be slightly different from the model with stochastic volatility reported in Table 2. Unemployment and inflation lag industrial production and real GDP by roughly two more months. The estimated persistence of the business cycle ρ is also smaller. However, the order in which the cycles are estimated (in terms of the phase shifts) remains the same in both specifications. The LPS reported in Table 3 indicates that the model without stochastic volatility performs the poorest of the models considered in our study.

4.3.3 The model with a common trend specification

The aim of our method is to extract a business cycle indicator which is a function of the movements of a time series within a specific range of frequencies. For this purpose we have kept the model flexible and included a separate trend component for each individual series in our model specification (2) – (7) for the mean. Since we consider eleven U.S. macroeconomic time series, common dynamics in the trend components may exist. The existence of common trends in multivariate time series imply cointegration as argued by Stock and Watson (1988). Since our primary focus is on the cycle component, we do not investigate the multivariate trend specification in much detail. However, empirical evidence of cointegration between income and consumption is widespread and we therefore consider a common trend for U.S. time series of real GDP and real consumption of nondurables plus services. The trend specification for these two series is given by

$$\tau_{i',t} = \tau_t^c, \qquad \tau_{i'',t} = a + b\tau_t^c, \tag{14}$$

where a and b are unknown coefficients and where the indices i' and i'' correspond to the indices of the income and consumption series. The two trend specifications in (14) replace the specification in (2) for $i = i', i''$ whereas the common trend τ_t^c is specified as one of the independent trends in (2). The number of trends has reduced from 11 to 10 in this specification. Extending the MCMC algorithm to estimate the parameters a and b requires an additional Gibbs sampling step to draw these parameters from their conditional distribution.

The effect of a common trend for income and consumption on the estimated business cycle can be viewed in panel (iv) of Figure 5. When it is compared with the original cycle estimate in Figure 2, we find that differences are small but do exist. In particular, the significant peaks and troughs in the years 1985–1990 and 2002–2005 of our cycle in Figure 2 are not confirmed in the cycle obtained from the model with common trends. We do however note that NBER recession dates have also not been recorded in these years. The LPS value for the model with cointegration in Table 3 is higher than the LPS values for some other model specifications. It indicates that several models without cointegration are preferred.

4.3.4 The model without phase shifts and one-step ahead prediction errors

We have also considered the model with stochastic volatility but without phase shifts. Panel (iii) of Figure 5 presents the estimated business cycle indicator and 95% highest posterior density intervals from this model. This estimated indicator is reasonably close to earlier cycle estimates. To further evaluate the different model specifications, we compare the standardized one-step ahead prediction errors from our main model of section 4.1 to those of the model without phase shifts. The prediction errors are computed using a particle filter.

<INSERT FIGURE 6 HERE>

The sample autocorrelation functions (ACF) of the standardized one-step ahead prediction errors for a selection of the eleven series are given in Figure 6. We report results for productivity, real consumption of nondurables plus services, manufacturing, and real GDP because these series are the leading indicators from our previous analysis, see the values of ξ_i in Table 2. The four panels on the left-hand side of Figure 6 contain the sample ACFs of the prediction errors for the four selected series from our main model. Although there are still some significant autocorrelations at the first lag for consumption and manufacturing, our preferred specification produces prediction residuals that are reasonable given our restrictive modeling framework with a single common cycle. The sample ACFs for the prediction errors from the model without phase shifts are presented in the four panels on the right-hand side of Figure 6. These graphs indicate that the prediction errors still contain some of the dynamics in the series. This finding provides some evidence that the addition of phase shifts provides an adequate simultaneous modeling framework for the eleven time series. The LPS for this model in Table 3 is considerably higher than our main model but is preferred to the model without stochastic volatility.

4.3.5 Sensitivity to the selection of variables

The estimated business cycle indicator is a function of the eleven time series that are included in our analysis. Here we evaluate how robust the indicator is to small changes in the selection of different series. For this purpose we exclude a single variable from the analysis (one at a time) and re-estimate the parameters of the model based on a selection of ten time series. Since real GDP and industrial production are key series in the analysis of a business cycle, we keep these two series in the analysis at any time. The aim of this exercise is to see whether there is an individual series that has a heavy impact on our results. Figure 7 presents the resulting nine different estimated business cycle indicators from our experiment (dropping one series at a time) together with the estimated indicator from our preferred model from section 4.1. We learn from this exercise that our estimated business cycle is largely robust to these small deviations in the selection of series. The exception is one estimated indicator which is substantially different at the end of the sample. Interestingly, this estimated indicator is constructed from the model that omits the productivity series. In our main analysis reported in Table 2, the estimates of the phase shifts ξ_i indicate that productivity leads industrial production and real GDP by the most. This experiment demonstrates that pooling more series can help to stabilize an estimator of the business cycle but that researchers should still be careful in selecting time series for the analysis

<INSERT FIGURE 7 HERE>

4.4 Forecasting

Since our approach to signal extraction and bandpass filtering is model-based, it remains possible to forecast the future business cycle using standard state space methods. For the purpose of forecast comparisons, we have estimated four different models using data through September 2007 and then computed the predictive density for one-year from September 2007 through September 2008. Figure 8 presents the smoothed estimates of the cycle component with 95% highest posterior density intervals from January 1999 to September 2007 together with the one year forecasts from October 2007 onward (after the dotted vertical line). The first three models we consider for Figure 8 are the main model from section 4.1 with stochastic volatility and mixture innovations in panel (i), the model with stochastic volatility but without phase shifts in panel (ii) and the model with phase shifts but without stochastic volatility in panel (iii). Finally, we also have computed the one year forecasts from a univariate version of the model (1) – (5) without stochastic volatility. This model was applied to quarterly real GDP and the forecast is presented in panel (iv) of Figure 8. The main model (i) and the model without phase shifts (ii) both forecast a downturn in the economy. The mean of the predictive distribution signals a potential recession in 2008. The 95% highest posterior density intervals however show a considerable amount of uncertainty which increases quickly as the forecast horizon grows to one year.

<INSERT FIGURE 8 HERE>

The forecasts in panel (i) and (iii) of Figure 8 allow us to compare a multivariate model with stochastic volatility to a corresponding model without stochastic volatility as in Valle e Azevedo et al. (2006). In this case, we see a reduction in size of the 95% highest posterior density intervals. However, the indicator is estimated to be below the trend in September 2008 and the one-year ahead point forecasts are for a moderate expansion. The point forecasts from the univariate model applied to real GDP are presented in panel (iv) of Figure 8. The univariate forecasts appear similar but have much larger uncertainty. This finding may be due to the lack of pooling of data and to the omission of stochastic volatility. The forecasting results presented here give some indication of the value of pooling information from multiple series as well as the impact of modeling volatility.

5 Conclusion

We have constructed a business cycle indicator that explicitly accounts for the time variation in macroeconomic volatility commonly known as the great moderation. Our indicator is built from a multivariate unobserved components time series model with a common stochastic cycle that is adjusted for phase shift and amplitude and that is shared across series. A novelty is the introduction of stochastic volatility processes (with innovations from mixture distributions) for irregular and common cycle disturbances. The empirical results reveal that the cycles in inflation and unemployment lag the cycles in other variables while productivity, real consumption of nondurables plus services, and manufacturing (in this order) are leading the U.S. business cycle. The 95% highest posterior density intervals of the estimated volatilities indicate significant changes over time of the volatitities in the irregular (high frequency) components of real retail sales, real GDP, real consumption of durables plus real fixed investment, and manufacturing. Although parameter estimates may have changed when considering other model specifications, the estimated business cycle indicator is relatively robust to such changes. In particular, we have provided some evidence that the duration, the persistence and the phase shifts associated with the business cycle are subject to structural breaks before and after the first quarter of 1984. We have shown that the cycle phase shifts play an important role in our simultaneous model with a common cycle in capturing the dynamics for each individual series. The business cycle forecasts of our preferred specification are precise; we have been able to predict a potential recession in 2008 based on our data-set through October 2007.

Acknowledgements

The authors have benefited from the comments of participants in seminar presentations at the 14th International Conference on Computing in Economics and Finance in Paris, the Econometric Institute of Erasmus University Rotterdam, the 5^{th} Eurostat Colloquium on "Modern Tools for Business Cycle Analysis" in Luxembourg and the European Central Bank in Frankfurt. The first author would like to thank the hospitality of the Vrije Universiteit Amsterdam, where most of this research was conducted. He would also like to acknowledge financial support from the Grover and Creta Ensley Fellowship, which funded part of this research while he was a graduate student at the University of Washington. All the computations reported in this paper were carried out using the OxMetrix 5.0 programming environment of Doornik (2007) along with functions from SsfPack 3.0 of Koopman, Shephard and Doornik (1999, 2008) and SvPack 2.0 of Shephard (1998).

References

Ahmed, S., A. Levin, and B. A. Wilson (2004). Recent U.S. macroeconomic stability: good policies, good practices, or good luck? *The Review of Economics and Statistics 86*(3), 824–832.

Altissimo, F., C. Riccardo, M. Forni, M. Lippi, and G. Veronese (2007). New Eurocoin: tracking economic growth in real time. Unpublished manuscript, Universitá di Roma 'La Sapienza'.

Anderson, B. D. O. and J. B. Moore (1979). *Optimal Filtering*. New Jersey: Prentice-Hall.

Baxter, M. and R. King (1999). Measuring business cycles: approximate band-pass filters for economic time series. *The Review of Economics and Statistics 81*, 575–593.

Blanchard, O. and J. Simon (2001). The long and large decline in U.S. output volatility. *Brookings Papers on Economic Activity 1*, 135–164.

Burns, A. and W. C. Mitchell (1946). Measuring business cycles. *National Bureau of Economic Research*. New York, NY.

Canova, F. (1998). Detrending and business cycle facts. *Journal of Monetary Economics 41*, 475–512.

Chauvet, M. and S. Potter (2001). Recent changes in the U.S. business cycle. *The Manchester School 69*(5), 481–508.

Chen, R. and J. S. Liu (2000). Mixture Kalman filters. *Journal of the Royal Statistical Society, Series B 62*(3), 493–508.

Christiano, L. and T. Fitzgerald (2003). The band-pass filter. *International Economic Review 44*, 435–465.

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*, 253–278.

Creal, D. D. (2009). A survey of sequential Monte Carlo methods for economics and finance. Unpublished manuscript, Department of Econometrics, Vrije Universiteit Amsterdam.

Diebold, F. X. and G. Rudebusch (1996). Measuring business cycles: a modern perspective. *The Review of Economics and Statistics 78*, 67–77.

Doornik, J. A. (2007). *An Object-Oriented Matrix Language Ox 5*. London: Timberlake Consultants Press.

Durbin, J. and S. J. Koopman (2001). *Time Series Analysis by State Space Methods*. Oxford, UK: Oxford University Press.

Forni, M., M. Hallin, M. Lippi, and L. Reichlin (2000). The generalized factor model: identification and estimation. *The Review of Economics and Statistics 82*(4), 540–554.

Forni, M., M. Hallin, M. Lippi, and L. Reichlin (2001, May). Coincident and leading indicators for the Euro area. *The Economic Journal 111*, C62–C85.

Giordani, P. and R. Kohn (2008). Efficient Bayesian inference for multiple change-point and mixture innovation models. *Journal of Business and Economic Statistics 26*(1), 66–77.

Gómez, V. (1999). Three equivalent methods for filtering finite nonstationary time series. *Journal of Business and Economic Statistics 17*, 109–116.

Gómez, V. (2001). The use of Butterworth filters for trend and cycle estimation in economic time series. *Journal of Business and Economic Statistics 19*, 365–373.

Good, I. J. (1952). Rational decisions. *Journal of the Royal Statistical Society, Series B 14*, 107–114.

Harvey, A. C. (1989). *Forecasting, structural time series models and the Kalman filter*. Cambridge, UK: Cambridge University Press.

Harvey, A. C. and S. J. Koopman (1997). Multivariate structural time series models. C. Heij, H. Schumacher, B. Hanzon, C. Praagman, Eds. *System Dynamics in Economics and Financial Models*. New York, NY: John Wiley Press.

Harvey, A. C. and T. Trimbur (2003). Generalized model-based filters for extracting trends and cycles in economic time series. *The Review of Economics and Statistics 85*(2), 244–255.

Harvey, A. C., T. Trimbur, and H. K. van Dijk (2007). Trends and cycles in economic time series: a Bayesian approach. *Journal of Econometrics 140*(2), 618–649.

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

Kim, C.-J. and C. R. Nelson (1998). Business cycle turning points, a new coincident index, and tests of duration dependence in a dynamic factor model with regime switching. *The Review of Economics and Statistics 80*(2), 188–201.

Kim, C.-J. and C. R. Nelson (1999). Has the U. S. economy become more stable? A Bayesian approach based on a Markov-switching model of the business cycle. *The Review of Economics and Statistics 81*(4), 608–616.

Kim, C.-J., C. R. Nelson, and J. Piger (2004). The less-volatile U.S. economy: a Bayesian investigation of timing, breadth, and potential explanations. *Journal of Business and Economic Statistics 22*(1), 80–93.

Koopman, S. J., N. Shephard, and J. A. Doornik (1999). Statistical algorithms for models in state space using SsfPack 2.2. *The Econometrics Journal 2*, 107–166.

Koopman, S. J., N. Shephard, and J. A. Doornik (2008). *SsfPack 3.0: Statistical algorithms for models in state space form*. London: Timberlake Consultants.

McConnell, M. M. and G. Pérez-Quirós (2000). Output fluctuations in the United States: what has changed since the early 1980's? *American Economic Review 90*, 1464–1476.

Nelson, C. R. (1988). Spurious trend and cycle in the state space decomposition of a time series with a unit root. *Journal of Economic Dynamics and Control 12*, 475–488.

Perron, P. and T. Wada (2006). Let's take a break: Trends and cycles in U.S. real GDP. Unpublished manuscript, Department of Economics, Boston University.

Pitt, M. K. and N. Shephard (1999). Filtering via simulation: auxiliary particle filters. *Journal of the American Statistical Association 94*(446), 590–599.

Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies 72*, 821–852.

Robert, C. P. and G. Casella (2004). *Monte Carlo Statistical Methods* (Second ed.). New York: Springer Press.

Rünstler, G. (2004). Modelling phase shifts among stochastic cycles. *The Econometrics Journal 7*, 232–248.

Sensier, M. and D. van Dijk (2004). Testing for volatility changes in U.S. macroeconomic time series. *The Review of Economics and Statistics 86*(3), 833–839.

Shephard, N. (1998). SvPack 2.0: Software for stochastic volatility models. Nuffield College, Oxford.

Stock, J. H. and M. W. Watson (1988). Testing for common trends. *Journal of the American Statistical Association 83*, 1097–1107.

Stock, J. H. and M. W. Watson (1989). New indexes of coincident and leading economic indicators. in O. Blanchard and S. Fischer, Eds. *NBER Macroeconomics Annual 1989*. Cambridge, MA: MIT Press.

Stock, J. H. and M. W. Watson (1991). A probability model of the coincident economic indicators. in K. Lahiri and G. H. Moore, Eds. *Leading Economics Indicators: New Approaches* *and Forecasting Records*. Cambridge, UK: Cambridge University Press.

Stock, J. H. and M. W. Watson (2002a). Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association 97*, 1167–79.

Stock, J. H. and M. W. Watson (2002b). Has the business cycle changed and why? in M. Gertler and K. Rogoff, Eds. *NBER Macroeconomics Annual 2002*. Cambridge, MA: MIT Press.

Stock, J. H. and M. W. Watson (2002c). Macroeconomic forecasting using diffusion indexes. *Journal of Business and Economic Statistics 20*, 147–162.

Stock, J. H. and M. W. Watson (2007). Why has U.S. inflation become harder to forecast? *Journal of Money, Credit, and Banking 39*(1), 3–34.

Valle e Azevedo, J. (2007). A multivariate band-pass filter. Unpublished manuscript, Economics and Research Department, Banco de Portugal.

Valle e Azevedo, J., S. J. Koopman, and A. Rua (2006). Tracking the business cycle of the Euro area: a multivariate model-based bandpass filter. *Journal of Business and Economic Statistics 24*(3), 278–290.

A Appendix: implementation details of Bayesian estimation

A.1 Parameter estimation using Markov chain Monte Carlo

The posterior distribution of the model can be derived using Bayes' rule as

$$p\left(\alpha_{1:n},\gamma_{1:n},\chi_{1:n},\Theta|y_{1:n}\right)\propto p\left(y_{1:n}|\alpha_{1:n},\gamma_{1:n},\chi_{1:n},\Theta\right)p\left(\alpha_{1:n}|\gamma_{1:n},\Theta\right)p\left(\gamma_{1:n}|\chi_{1:n},\Theta\right)p(\chi_{1:n})p\left(\Theta\right),$$

where $\gamma_{1:n}$ includes the stochastic volatility processes in (9) and (11), $\alpha_{1:n}$ includes the unobserved components of the model (trend τ_{it}, slope β_{it} and cycle $\psi_t^{(q)}$) and $\chi_{1:n}$ contains the indicator variables for the volatility and slope innovations. We estimate the model using a Markov chain Monte Carlo (MCMC) algorithm; see, Robert and Casella (2004) for an overview of MCMC.

A.2 Details of our implemented MCMC algorithm

The MCMC algorithm begins with initial values for the parameters $\Theta^{(0)}$, indicator variables $\chi_{1:n}^{(0)}$, and log-variances $\gamma_{1:n}^{(0)}$. It continues by making draws for $j=1,...,N$ using an algorithm consisting of four main steps. The fourth step requires the sampling of parameters from different distributions. In practice, we use a burn-in of $10,000$ draws and then make $60,000$ draws, skipping every 20 to retain a total of $3,000$ draws for inference.

Step 1. Drawing $\alpha_{1:n}^{(j)}$ conditional on $\Theta^{(j-1)}$, $\gamma_{1:n}^{(j-1)}$, $\chi_{1:n}^{(j-1)}$, and the data.

Conditional on $\Theta^{(j-1)}$, $\gamma_{1:n}^{(j-1)}$, $\chi_{1:n}^{(j-1)}$, and the data, the model reduces to a conditionally linear, Gaussian state space model. The states $\alpha_{1:n}^{(j)}$ are drawn using the simulation smoothing algorithm of Durbin and Koopman (2002). The simulation smoothing algorithm runs the Kalman filter forward, the Kalman smoothing algorithm backward, and another run forward to produce the required sequence of draws. We follow Trimbur (2006) and Koopman (1997) to initialize the stationary and nonstationary components of the model, respectively. Textbook treatments of the Kalman filter and smoother can be found in Anderson and Moore (1979), Harvey (1989) and Durbin and Koopman (2001).

Step 2. Drawing indicators $\chi_{1:n}^{(j)}$ conditional on $\Theta^{(j-1)}$, $\alpha_{1:n}^{(j)}$, and the data.

Conditional on $\alpha_{1:n}^{(j)}$, $\Theta^{(j-1)}$, and the data, we have to draw $M+M^*+1$ series of indicator variables that are independent of one another. We construct conditionally linear Gaussian state space models from the data. This allows us to apply the reduced conditional sampling algorithm of either Gerlach et al. (2000) or Doucet and Andrieu (2001). See also Giordani and Kohn (2008) for more details.

Step 3. Drawing log-variances $\gamma_{1:n}^{(j)}$ conditional on $\Theta^{(j-1)}$, $\alpha_{1:n}^{(j)}$, $\chi_{1:n}^{(j)}$, and the data.

Conditional on $\chi_{1:n}^{(j)}$,$\alpha_{1:n}^{(j)}$, $\Theta^{(j-1)}$, and the data, we have to draw $M^* + 1$ series of log-variances that are independent of one another. We construct $M^* + 1$ stochastic volatility models from the data and $\alpha_{1:n}^{(j)}$. This allows us to apply the mixture of normals approximation to the SV model provided in Omori et al. (2007).

Step 4. Drawing $\Theta^{(j)}$ conditional on $\gamma_{1:n}^{(j)}$, $\alpha_{1:n}^{(j)}$, and the data.

(a) *Drawing δ_i*: Given ψ_i, ψ_i^+, and τ_{it}, the other parameters of the model, and the data, the priors on δ_i are conjugate allowing these parameters to be drawn from their full conditional distributions via a Gibbs sampling step. The full conditional distributions are normal distributions and are standard; e.g., see Kim and Nelson (1999, pp. 173).

(b) *Drawing $\sigma_{i,\varepsilon}^2$ for those series with no stochastic volatility*: Given the cycles $\psi_t^{(q)}, \psi_t^{+(q)}$, the trend τ_{it}, other parameters of the model, and the data, the non-informative inverse gamma prior on these parameters mean that each variance $\sigma_{i,\varepsilon}^2$ can be drawn from its full conditional distribution via a Gibbs sampling step. The full conditional distributions are inverse gamma distributions and are standard; e.g., see Kim and Nelson (1999, pp. 175).

(c) *Drawing ρ, ξ_i, and λ*: Our priors on these parameters are not conjugate requiring Metropolis-Hastings steps. We use a standard random-walk Metropolis algorithm where we grouped all the parameters ρ, ξ_i, and λ together in one block labeled Θ_{-c}. The other parameters of the model in Θ that remain constant are labeled Θ_c. The covariance matrix on the random walk was estimated over several initial runs. We then tuned the scales on the random walk to have roughly a 35-40% acceptance rate. For example, given $\Theta_{-c}^{(j-1)}$, we draw a new candidate value Θ_{-c}^* from a multivariate normal distribution centered at $\Theta_{-c}^{(j-1)}$ and accept this candidate with probability

$$\min\left(\frac{p(y_{1:n}|\Theta_c, \Theta_{-c}^*, \gamma_{1:n}^{(j)}, \chi_{1:n}^{(j)})p(\Theta_{-c}^*)}{p(y_{1:n}|\Theta_c, \Theta^{(j-1)}, \gamma_{1:n}^{(j)}, \chi_{1:n}^{(j)})p(\Theta_{-c}^{(j-1)})}, 1\right),$$

where $p(y_{1:n}|\Theta_c, \Theta_{-c}^*, \gamma_{1:n}^{(j)}, \chi_{1:n}^{(j)})$ can be computed by the Kalman filter conditional on the current log-variances $\gamma_{1:n}^{(j)}$ and the remaining parameters of the model Θ_c.

Indicator	$k^{(1)}$	$k^{(2)}$	$p^{(1)}$	$p^{(2)}$
$K_{i,t,\varepsilon}$	0	5×10^{-2}	0.97	0.03
$K_{t,\kappa}$	0	5×10^{-2}	0.97	0.03
$K_{i,t,\zeta}$	5×10^{-5}	1×10^{-3}	0.95	0.05

Table 1: *Fixed values ($k^{(1)}$ and $k^{(2)}$) and their probabilities ($p^{(1)}$ and $p^{(2)}$, respectively) for the mixture indicator variables $K_{i,t,\varepsilon}$, $K_{t,\kappa}$ and $K_{i,t,\zeta}$ used in the empirical work.*

	Prior δ_i		Posterior δ_i		Prior ξ_i		Posterior ξ_i	
	mean	st. dev.	mean	st. dev.	mean	st. dev.	mean	st. dev.
unemployment	0.00	2.00	-0.263	0.010	0.00	2.50	-0.958	0.182
manufacturing	0.00	2.00	0.368	0.014	0.00	2.50	4.001	0.242
inflation	0.00	2.00	0.332	0.035	0.00	2.50	-3.807	0.604
retail	0.00	2.00	0.699	0.036	0.00	2.50	2.969	0.354
retail/food	0.00	2.00	0.556	0.095	–	–	–	–
productivity	0.00	2.00	0.326	0.025	0.00	2.50	8.391	0.536
real GDP	0.00	2.00	0.577	0.024	0.00	2.50	3.409	0.250
hours	0.00	2.00	0.550	0.023	0.00	2.50	0.648	0.261
consumption	0.00	2.00	0.142	0.015	0.00	2.50	6.365	0.741
investment	0.00	2.00	2.239	0.091	0.00	2.50	3.349	0.231
	Prior ρ		Posterior ρ		Prior $2\pi/\lambda$		Posterior $2\pi/\lambda$	
	0.50	0.08	0.969	0.004	63.56	67.02	41.07	1.25

Table 2: *Prior and posterior means and standard deviations for ρ, $2\pi/\lambda$, and for each series δ_i and ξ_i. Industrial production has $\delta = 1$ and $\xi = 0$.*

	Main	All-SV	No-SV	Break in ρ	Break in ρ, λ, ξ_i	Coint	No phase shifts
LPS	6.40	6.43	10.11	6.39	6.36	7.82	8.62

Table 3: *Log-predictive scores (LPS) for a selection of models.* Main *is the model of section 4.1*, Break in ρ and Break in ρ, λ, ξ_i *are from section* 4.3.1, No-SV *is from section* 4.3.2, Coint *is from section* 4.3.3, No phase shifts *is from section* 4.3.4. *We also report the LPS for the model with SV on all series as* All-SV.

	No-SV		Model with a known break			
	ξ_i		ξ_i 1953:Q1-1983:Q4		ξ_i 1984:Q1-2007:Q3	
	mean	st. dev.	mean	st. dev.	mean	st. dev.
unemployment	-2.366	0.312	-1.082	0.161	-2.043	0.560
manufacturing	4.282	0.445	3.409	0.232	7.781	0.826
inflation	-5.970	0.920	-4.535	0.601	-0.729	1.594
retail	3.739	0.461	2.470	0.363	4.893	1.154
productivity	9.301	0.758	7.393	0.508	10.35	1.771
real GDP	4.032	0.412	2.752	0.238	6.215	0.763
hours	-0.171	0.331	0.566	0.229	0.062	0.747
consumption	5.968	0.862	5.923	0.392	4.575	1.854
investment	3.609	0.341	2.766	0.212	5.422	0.621
ρ	0.869	0.009	0.972	0.004	0.969	0.007
$2\pi/\lambda$	51.94	3.783	38.81	1.170	55.22	4.360

Table 4: *Posterior means and standard deviations for ρ, $2\pi/\lambda$, and ξ_i for two different models. Column 1 is the model with no stochastic volatility. Columns 2 & 3 are for the model with a structural break in 1984:Q1. The priors are the same as in* Table 2.

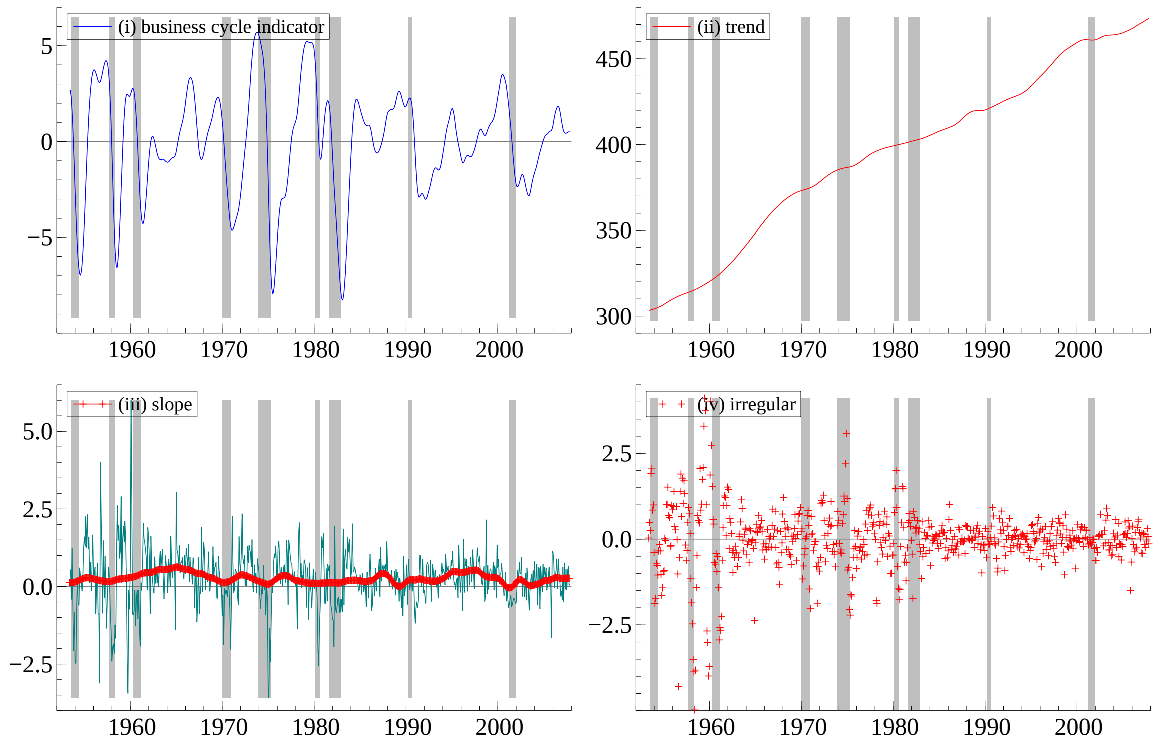

Figure 1: *(i) Business cycle indicator; (ii) smoothed estimates of the trend in industrial production; (iii) smoothed estimates of the slope and the growth rate of industrial production; (iv) smoothed estimates of the irregular component.* NBER *recession dates are represented by the vertical bands.*

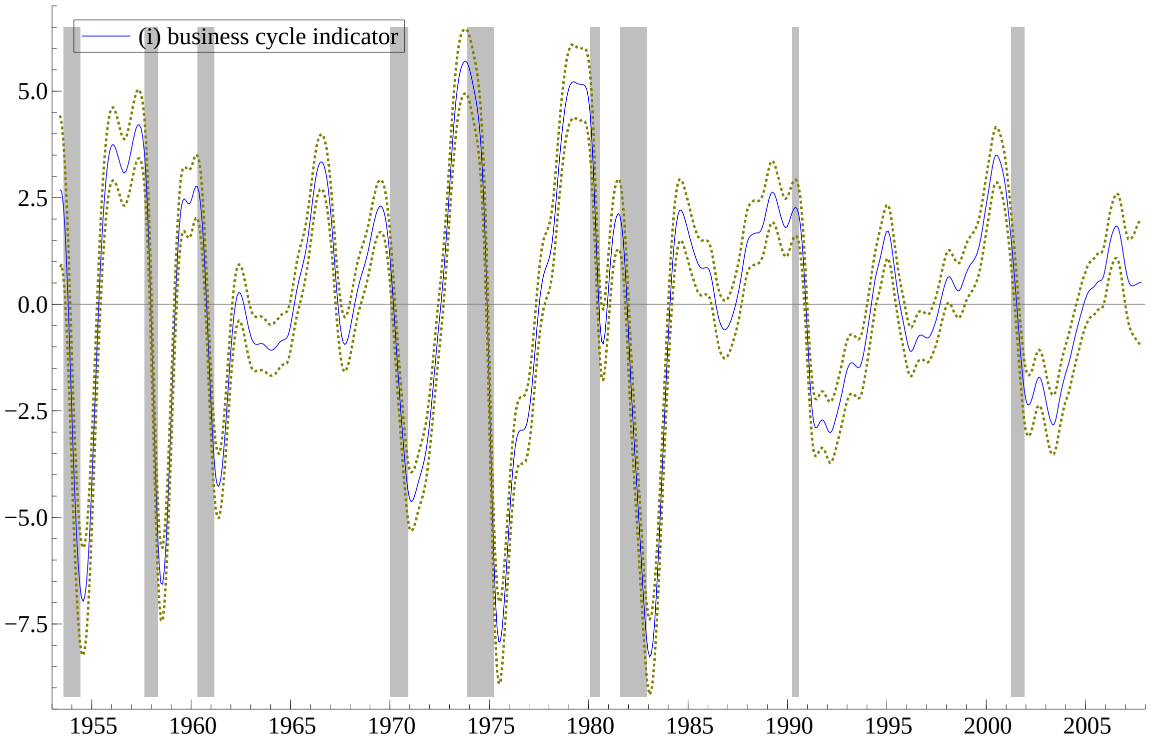

Figure 2: *Business cycle indicator with 95% highest probability density interval.* NBER *recession dates are represented by the vertical bands.*

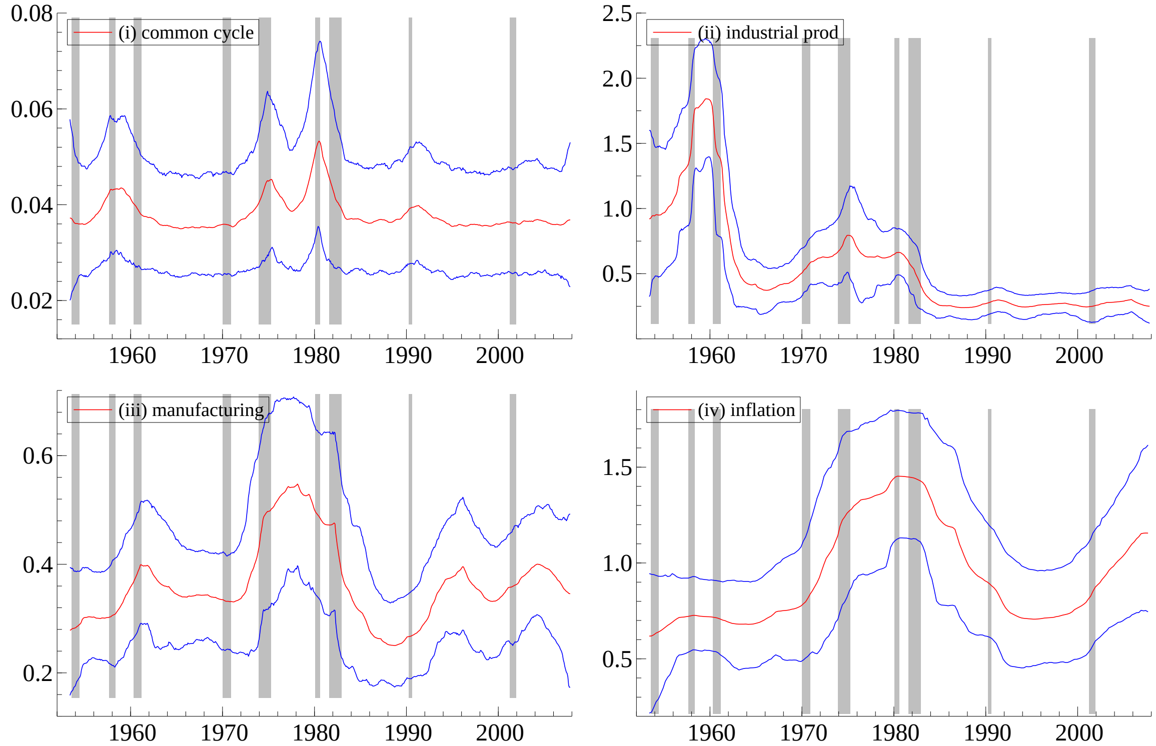

Figure 3: *(i) Smoothed estimates of the stochastic volatility* $\sigma_{t,\kappa}$ *of the common cycle with 95% highest probability density interval. Smoothed estimates of the stochastic volatility* $\sigma_{i,t,\varepsilon}$ *of the irregular components with 95% highest probability density interval for (ii) industrial production; (iii) manufacturing; and (iv) inflation.* NBER *recession dates are represented by the vertical bands.*

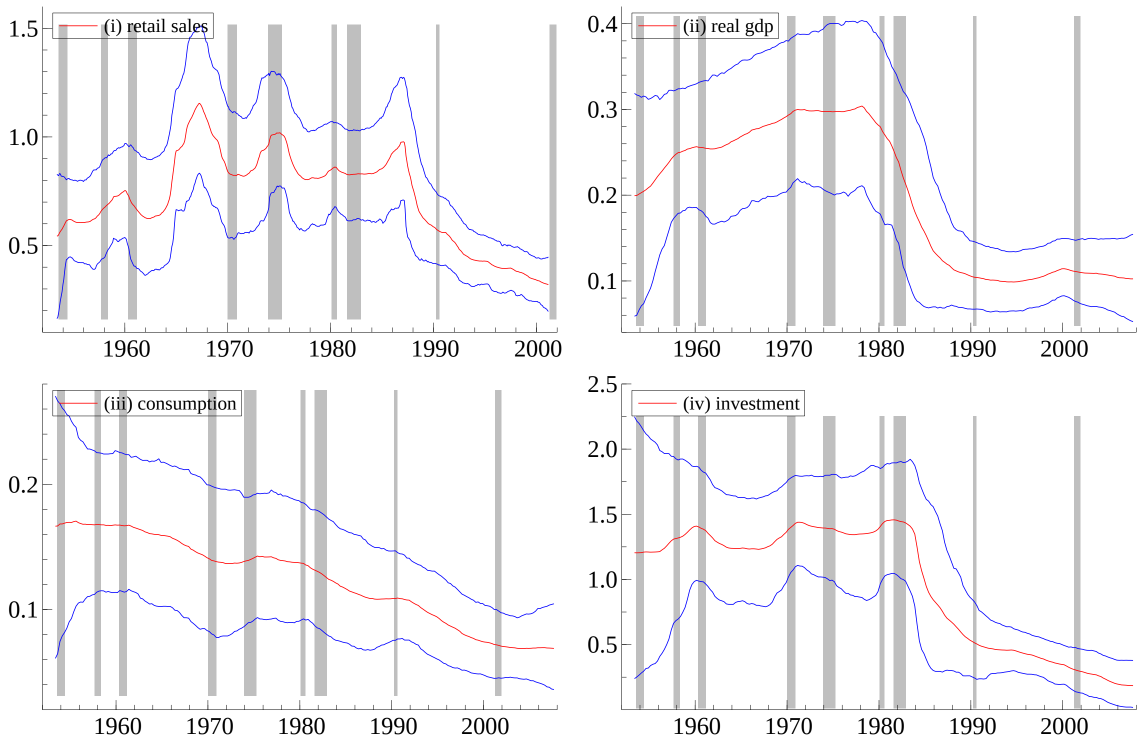

Figure 4: *Smoothed estimates of the stochastic volatility* $\sigma_{i,t,\varepsilon}$ *of the irregular components with 95% highest probability density interval for (i) real retail sales; (ii) real GDP; (iii) consumption of durables; and (iv) investment.* NBER *recession dates are represented by the vertical bands.*

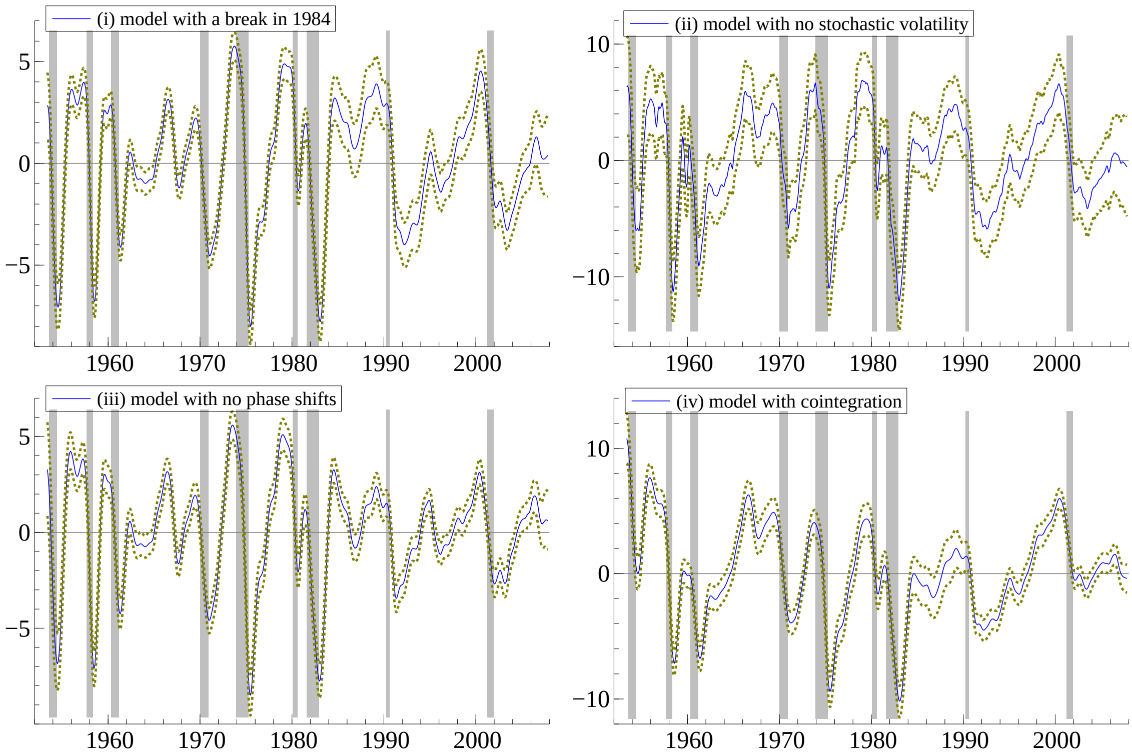

Figure 5: *Estimates of the business cycle indicator from four alternative models: (i) model with a known break in* ρ *and* λ *in 1984; (ii) model with no stochastic volatility; (iii) model with no phase shifts; (iv) model with a common trend for real GDP and real consumption of nondurables + services.* NBER *recession dates are represented by the vertical bands.*

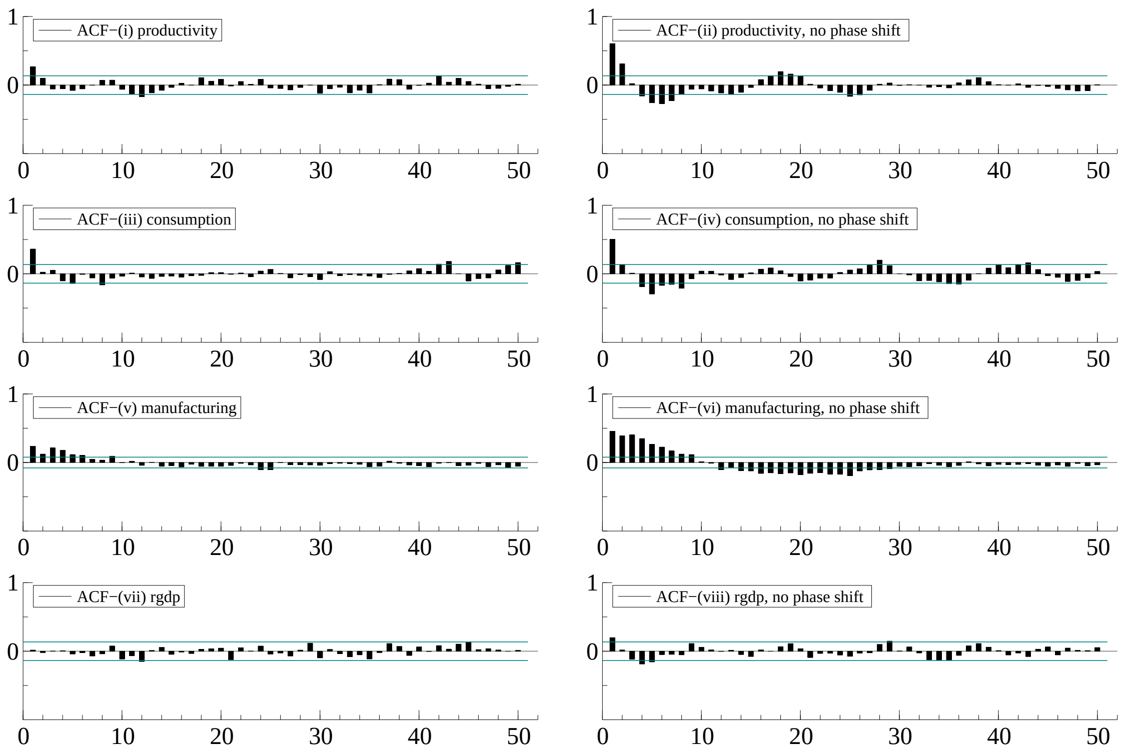

Figure 6: *ACF of the one-step ahead prediction errors from the main model (left column) and the model with no phase shifts (right column) for (i-ii) productivity; (iii-iv) real consumption of nondurables + services; (v-vi) manufacturing; (vii-viii) real GDP.*

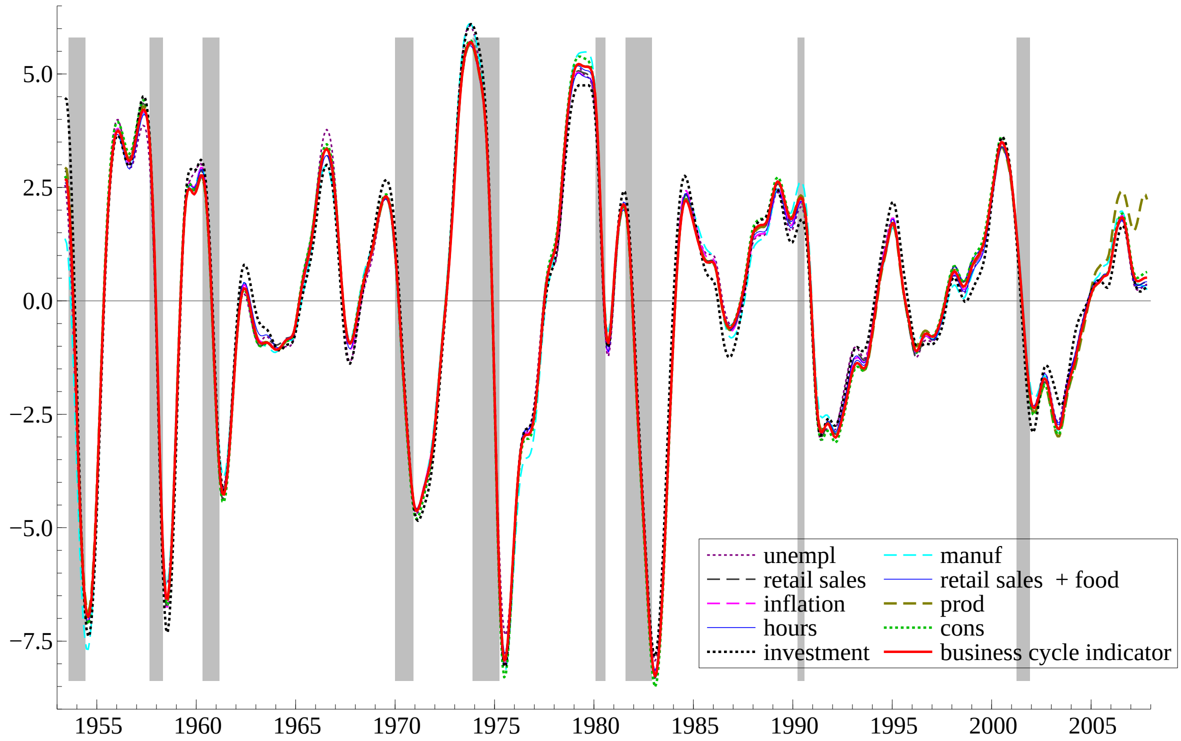

Figure 7: *Estimates of the business cycle indicator computed by dropping one series at a time.* NBER *recession dates are represented by the vertical bands.*

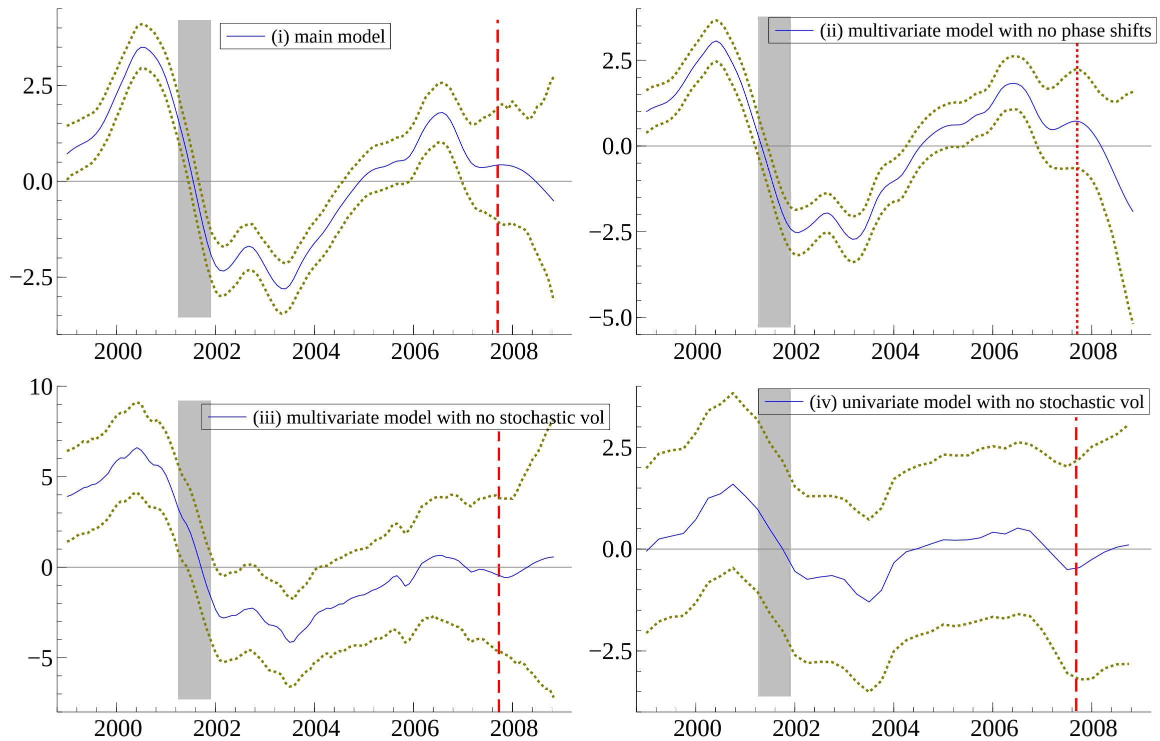

Figure 8: *Smoothed cycle estimates from 1999 to September 2007 and one-year forecasts from September 2007 onward: (i) main model; (ii) model with stochastic volatility but no phase shifts; (iii) multivariate model with no stochastic volatility; (iv) univariate model for real GDP with no stochastic volatility.* NBER *recession dates are represented by the vertical bands.*