The effect of the great moderation on the U.S. business cycle in a time-varying multivariate trend-cycle model

DREW CREAL
Department of Econometrics, Vrije Universiteit Amsterdam
De Boelelaan 1105, NL-1081 HV Amsterdam
`dcreal@feweb.vu.nl`

SIEM JAN KOOPMAN
Department of Econometrics, Vrije Universiteit Amsterdam
De Boelelaan 1105, NL-1081 HV Amsterdam
`s.j.koopman@feweb.vu.nl`

ERIC ZIVOT
Department of Economics, University of Washington
Box 353330, 302 Savery Hall, Seattle, WA 98195
`ezivot@u.washington.edu`

July 2008

Abstract

In this paper we investigate whether the dynamic properties of the U.S. business cycle have changed in the last fifty years. For this purpose we develop a flexible business cycle indicator that is constructed from a moderate set of macroeconomic time series. The coincident economic indicator is based on a multivariate trend-cycle decomposition model that accounts for time variation in macroeconomic volatility, known as the great moderation. 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 the result of 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 all the heteroskedasticity present in the data. The empirical results are based on a Bayesian analysis and show that time-varying volatility is only present in the a selection of idiosyncratic components while the coefficients driving the dynamic properties of the business cycle indicator have been stable over time in the last fifty years.

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

JEL Classification: C11; C32; E32

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, Koopman, and Rua (2006) and Harvey, Trimbur, and van Dijk (2007) are recent examples where the traditional structural time series approach based on a multivariate trend-cycle model is adopted. Closely related to the former methods are the indicators of Forni, Hallin, Lippi, and Reichlin (2000, 2001), Stock and Watson (2002a, 2002c) and Altissimo, Riccardo, Forni, Lippi, and Veronese (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* belief that periods of 1.5 to 8 years are of primary importance, interest centers on building an ideal filter. An ideal 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 filter, while Valle e Azevedo (2007) extends the ideal filter towards a multivariate setting.

Concurrent research documents sizeable changes in the volatility of U.S. macroeconomic time series; e.g., see Kim and Nelson (1999a), McConnell and Pérez-Quirós (2000), 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 contribution of this paper is the construction of a business cycle indicator that has bandpass filter properties and also accounts for time varying volatility. Our indicator is constructed from the multivariate unobserved components time series model of Valle e Azevedo, Koopman, and Rua (2006), where we extend the model to include stochastic volatility in both the common cycle and irregular components of the model. 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. We further adjust the stochastic cycle for phase shift and amplitude between series using the device of Rünstler (2004). The

innovations of the stochastic cycle and the irregulars are made subject to stochastic volatility processes which enable us to account for the heteroskedasticity present in the data. 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, Koopman, and Rua (2006) estimate the parameters of their model by maximum likelihood, we use Bayesian methods for inference. In particular, we use Markov chain Monte Carlo methods for the estimation of all parameters including the trend, cycle and irregular components. In this respect, our paper complements the work of Kim and Nelson (1998), Chauvet and Potter (2001), and Harvey, Trimbur, and van Dijk (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, Trimbur, and van Dijk (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 cyclical component, they do not do so in the idiosyncratic component nor does their cyclical component isolate business cycle frequencies.

We discuss the results of an empirical analysis based on eleven U.S. macroeconomic time series that are commonly used for constructing a business cycle indicator. The modeling framework allows time series that are coincident, leading or lagging the business cycle to be included in the analysis. The empirical results reveal that unemployment and inflation are lagging the business cycle while productivity, manufacturing and consumption are leading the cycle. The NBER recession dates are picked up by the estimated cycle of our model. 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. We may therefore conclude that the time-varying volatility in the macro-economic time series is associated with the high-frequency movements only. These empirical findings have not been observed earlier and may provide a justification for modeling the cycle component as a stationary process with time-invariant parameters.

The multivariate trend-cycle decomposition time series model with stochastic volatility and mixture distributions is presented in section 2. In section 3, we describe the data, our priors,

and the Markov chain Monte Carlo (MCMC) algorithm used to estimate the model. 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

Unobserved components models consist of stochastic components with dynamic specifications that have a direct interpretation. A classic model-based decomposition of a time series into trend, seasonal and irregular can be accomplished in this framework by specifying the trend component as a random walk process (low frequency), the seasonal component as a dynamic seasonal process (periodic frequency) and the irregular component as a white noise process (high frequency). For macroeconomic time series, which are often seasonally adjusted, a stationary component in its decomposition is of interest as well since it may detect dynamic characteristics associated with the business cycle. The dynamic properties of a business cycle are well documented. For example, much of the business cycle literature cite the seminal paper of Burns and Mitchell (1946) who argue that business cycle fluctuations should typically last between 1.5 years and 8 years. Autoregressive moving average processes with complex roots in the autoregressive polynomial can be designed that ensure the implied cyclical fluctuations to last within a certain band of years. A component with such specifications can then be interpreted as a business cycle component.

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 business cycle component. 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 for many equations and components in the model such that the model is robust to the impact of the great moderation.

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 equation of the multivariate model is given by

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

where τ_{it} and ε_{it} are the idiosyncratic trend and irregular components, respectively, for the ith variable. The cycle component $\psi_t^{(k)}$ is specified as a smooth cyclical process where k 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 the 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

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

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

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 $\sigma_{i,\varepsilon}^2$ and is normally distributed, that is

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

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 $\psi_t^{(k)}$ component is the kth-order stochastic cycle and modeled by

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

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

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

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

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 $\psi_t^{(k)}$ approximates a bandpass filter as the number of cycles k increases. As it is the case for the nonparametric filters of Baxter and King (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

5

its width. Bandpass filters eliminate 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 the high-frequency variation. In previous research carried out by, e.g., Harvey and Trimbur (2003) and Valle e Azevedo, Koopman, and Rua (2006), the frequency parameter λ has been fixed. Harvey, Trimbur, and van Dijk (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 cycle component estimates of the last two or three periods may be heavily revised when new data arrives. Valle e Azevedo, Koopman, and Rua (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, Riccardo, Forni, Lippi, and Veronese (2007) have adopted a similar approach to improve their estimates at the end of the sample.

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.

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

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

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 $\psi_t^{(k)}$ and $\psi_t^{(+k)}$ 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 $\alpha_{1:n} = \left[\{\tau_{it}, \beta_{it}\}_{i=1}^M, \psi_t, \psi_t^+ \right]_{t=1}^n$.

Valle e Azevedo, Koopman, and Rua (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. As the model imposes a common cycle (adjusted for phase shift and amplitude) across series, it is however not an attempt to approximate the true joint data generating process of all series. The properties of this method are instead best evaluated by taking it as a filter and by inspecting the gain function of the cycle, see section 2.3 below.

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 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 $K_{i,t,\varepsilon}$ takes 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 is recently introduced 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 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 $i = 1,\ldots,M$. All indicator variables are mutually and serially independent. Also, the disturbance sequences $\omega_{i,t,\varepsilon}$ are mutually and serially independent.

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

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

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

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 $k_\kappa^{(1)}$ and $k_\kappa^{(2)}$ with prior probabilities $p_\kappa^{(1)}$ and $p_\kappa^{(2)}$, respectively. and are independent of the indicator variables $K_{i,t,\varepsilon}$ in (9). The disturbance sequence $\omega_{t,\kappa}$ is also independent of all other disturbance sequences in the model.

The specifications for the stochastic log-variances with mixture innovations complete our model 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 $\sigma_{i,\zeta}^2$, with $i = 1, \ldots, M$. In other words, changes in the trend are reasonably rare which makes the variance $\sigma_{i,\zeta}^2$ hard to estimate accurately. Therefore, as an alternative approach, we adopt a similar strategy as for the log-variances by introducing mixture innovations for the trend component $\tau_{i,t}$. For this purpose we replace the specification for the slope or 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 $k_\zeta^{(1)}$ and $k_\zeta^{(2)}$ with prior probabilities $p_\zeta^{(1)}$ and $p_\zeta^{(2)}$, respectively. This specification enables the model to account for possible large but infrequent changes in the trend.

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 of, say, M^* of the M 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 3.5 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 aim to develop a flexible model to construct 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;

e.g., see Kim and Nelson (1999a), McConnell and Pérez-Quirós (2000), and Blanchard and Simon (2001). Other work, including Chauvet and Potter (2001), Stock and Watson (2002b), Kim, Nelson, and Piger (2004), Ahmed, Levin, and Wilson (2004), and Sensier and van Dijk (2004), analyzes a larger number of time series using a plethora of methods. Researchers are also concerned whether one or more breaks exist, e.g. Sensier and van Dijk (2004). The findings in this work do not yet suggest a consistent, definitive pattern. Some real and nominal series break only once while other real and nominal series may break more often. If each series breaks only once, it is not clear that this date is the same for all the series. Whether the break in each series was sharp versus gradual is still debated. We regard the specification of stochastic volatility with mixture innovations as a flexible means for developing a robust indicator, which is our primary interest in this study.

Another key question of interest is at what frequencies the break or breaks exist in each macroeconomic time series. Ahmed, Levin, and Wilson (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. Our flexible model specification is consistent with this view as well.

2.3 Time-varying gain and weight functions

The properties of a bandpass filter are evaluated in the frequency domain through the gain function. When constructing a trend or a cycle indicator from a single time series using a filter, the gain function describes which frequencies of the original time series are being used to construct the indicator. With the exception of the beginning and the end of the series, the gain function of a lowpass filter such as the one of Hodrick and Prescott (1997) or a bandpass filter remains constant through time. Consequently, the estimates provided by these filters do not account for the great moderation.

In a model-based approach, the trend and cycle indicators are constructed by the Kalman filter and associated smoothing algorithms, see Harvey (1989) and Durbin and Koopman (2001). These indicators are effectively the minimum mean square linear estimators of the trend and cycle

components based on all observations $y_{11}, \ldots, y_{MT}$. In case of the trend-cycle decomposition model (1) – (6) with $M = 1$ and $\delta_1 = 1$, the implicit gain functions for the trend and cycle indicators are determined by the *signal-to-noise* ratios $\sigma_{1,\zeta}^2 / \sigma_{1,\varepsilon}^2$ and $\sigma_\kappa^2 / \sigma_{1,\varepsilon}^2$, see Harvey (1989). For the multivariate model ($M > 1$), the gain functions are intricate functions of all such signal-to-noise ratios. In case of the full model specification with stochastic volatility, the gain function for the cycle indicator will implicitly vary over time because the signal-to-noise ratios in the model vary. It is of interest to investigate these changes in some detail. A method for computing the gain functions for our model specification is described in section 3.4.

3   Design of the empirical study and estimation

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), consumption of durables, investment, 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. Although consumption and investment are components of real GDP, they are included because their dynamics help in identifying the business cycle. 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 are in logarithms and have been multiplied by 100.

The resulting panel of macro-economic time series is unbalanced and consists of a mix 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 efficiently 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. Visual inspection of the manufacturing index indicates the existence of four large outliers in M12:1960, M1:1982, M12:1995, and M1:1996, respectively. We also treat these observations as missing.

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 of the prior settings for each of the parameters can be found in an Appendix associated with this paper. We have located this additional material at `http://staff.feweb.vu.nl/dcreal/`.

For 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, Trimbur, and van Dijk (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, that is 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.

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 is set 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.

3.3 Parameter estimation using Markov chain Monte Carlo

We estimate the model using Bayesian methods and have developed a Markov chain Monte Carlo (MCMC) algorithm; see, e.g. Kim and Nelson (1999b) and Robert and Casella (2004) for

an overview of MCMC. We describe the MCMC algorithm only briefly here. A more detailed description is presented in the Appendix 5 of this paper.

The posterior distribution of our 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 $\psi_t^{(k)}$) and $\chi_{1:n}$ contains the indicator variables for the volatility and slope innovations. We have adopted data-augmentation methods to expand the posterior for including the latent state variables. The first term on the right hand side is the data-augmented likelihood, the second and third terms are the transition density of the state space model, and the last term is the prior described above.

Given initial values for the parameters $\Theta^{(0)}$, log-variances $\gamma_{1:n}^{(0)}$, and indicator variables $\chi_{1:n}^{(0)}$, the MCMC algorithm consists of making a series of N draws using four blocks where N is a large number. In the jth round (with $j = 1, \ldots, N$), we first draw $\alpha_{1:n}^{(j)}$ conditional on the parameters $\Theta^{(j-1)}$, the log-variances $\gamma_{1:n}^{(j-1)}$, and the indicators $\chi_{1:n}^{(j-1)}$ using the simulation smoothing algorithm of Durbin and Koopman (2002). This Monte Carlo method is for linear Gaussian state space model and is based on the Kalman filter and an associated smoothing algorithm. Next, we draw the indicator variables $\chi_{1:n}^{(j)}$ using the reduced conditional sampling algorithms for mixture models developed independently by Gerlach, Carter, and Kohn (2000) and Doucet and Andrieu (2001). Conditional on $\Theta^{(j-1)}$, $\alpha_{1:n}^{(j)}$, and $\chi_{1:n}^{(j)}$, we draw new log-variances using the mixture of normals approximation to the stochastic volatility model developed by Kim, Shephard, and Chib (1998). We use the more accurate 10 component mixture from Omori, Chib, Shephard, and Nakajima (2007) which improves upon the original 7 component mixture in the former paper. Given the latent variables in $\gamma_{1:n}^{(j)}$ and $\chi_{1:n}^{(j)}$, we draw the parameters $\Theta^{(j)}$ using a combination of so-called Gibbs sampling and Metropolis-Hastings steps. 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.

Calculation of the exact likelihood function for the MH acceptance probabilities as well as draws from the simulation smoother require that careful attention is paid to the initialization of the Kalman filter for both the stationary and nonstationary components of the model. In particular, the initial distribution of the cycle depends on its order and is based on the unconditional mean and variance of the cycle process; details and derivations are given by Trimbur (2006). For the nonstationary components, we follow the exact initiliazation methods described in Koopman (1997). Textbook treatments of the Kalman filter and associating methods can be found in Anderson and Moore (1979), Harvey (1989) and Durbin and Koopman (2001).

12

3.4 Estimation of time-varying gain functions

In a linear Gaussian state space model, the estimates of the components (in $\alpha_{1:n}$) obtained by the Kalman filter and associated smoothing methods are linear functions of all observations. Koopman and Harvey (2003) demonstrate how to compute the observation weights for this linear function conditional on the model parameters and, in our case, the latent variables in $\chi_{1:n}$ and $\gamma_{1:n}$. These weights can then be transformed into the gain function for a given component of the model. In a frequentist setting, the weights and gains are computed at the maximum likelihood parameter estimate while in a Bayesian framework they are functions of the posterior. Consequently, in a Bayesian analysis they are computed by averaging over successive iterations of the MCMC algorithm. In particular, for each jth round of the MCMC algorithm, we can compute the weights and gains in the first block before we draw $\alpha_{1:n}^{(j)}$ using the simulation smoother. In section 4, we demonstrate how the gain function remains centered over the business cycle frequencies for each time period and that it adjusts over time according to the variations in volatility.

3.5 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 $M^* \leq M$ irregular components that have time-varying variances in the model under consideration. We have started to estimate 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.

The stochastic volatility processes have innovations with a mixture of an indicator variable and a standardized Gaussian noise component. A similar specification is used for the innovations of the slope component of the trend. In all cases, the indicators have two possible values with certain probabilities. These values and probabilities are pre-fixed and given in Table 1.

Finally, we have selected $k = 2$ for the cycle in (5) and have found that this setting produces a sufficiently smooth cycle. We notice that it has also been the preferred specification in Harvey, Trimbur, and van Dijk (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 graph-

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 possible 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 $K_{i,t,\zeta}$.*

ical 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 a separate online Appendix of this paper available at http://staff.feweb.vu.nl/dcreal/. 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 upto the end of 2007, the cycle indicator at the end of the sample provides evidence that the U.S. economy is on the brink of a recession at the end of 2007. In a multivariate framework, the highest posterior density intervals are considerably smaller than for a univariate model. The cycle component estimates with 95% highest posterior density intervals are pictured in Figure 2, where it can be seen that there still remains considerable uncertainty of the business cycle over the last 2 to 3 quarters in 2007. We observe from Figure 1(iii) that the estimated slope component corresponds closely to the growth rate of industrial production. Smoothed estimates of the irregular component in (iv) show evidence of a considerable 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 4.3 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, Koopman, and Rua (2006) found the opposite relationship between these two
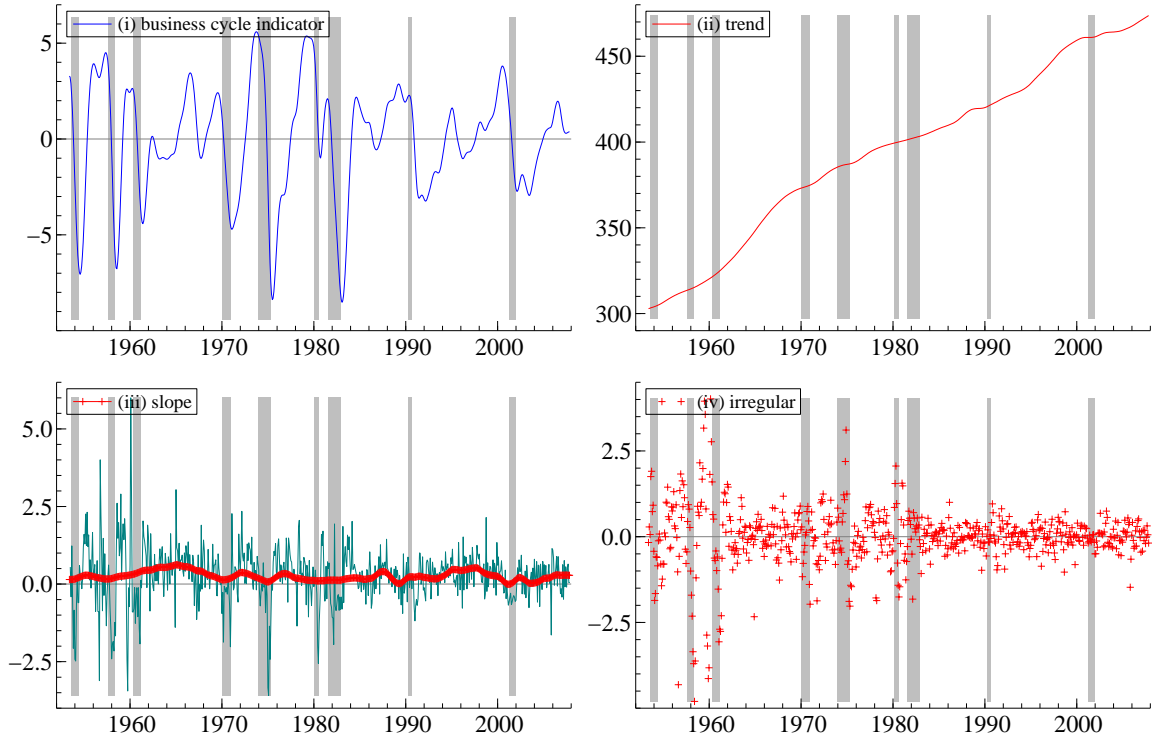
Figure 1: *(i) Business cycle indicator; (ii) smoothed estimates of the trend in industrial production; (iii) smoothed estimates of the slope and the growth rate of industrial production; (iv) smoothed estimates of the irregular component.* NBER *recession dates are represented by the vertical bands.*

variables for the Euro area. This may be explained by the different structures and dynamics in the U.S. and Euro area economies. Productivity and consumption of durables lead real GDP while manufacturing, retail sales, and investment appear to be roughly coincident with real GDP.

4.2 Time-varying gain functions

The gain function from a multivariate filter is a high-dimensional object and therefore cannot be as conveniently graphed as in a univariate framework. This makes interpretation of the gain function slightly more complicated. In Figure 3, we present the gain functions of the cycle in the model for two different quarters, Q1 1982 and Q1 1997. They are computed during the MCMC algorithm from the linear Gaussian model conditional on the log-variances for that date, the latent indicators, and the parameters of the model Θ. We observe several characteristics of interest for our analysis. Both gain functions in 1982 and 1997 eliminate the high frequency variation from each series because the gain function returns to zero by frequency 0.30. The
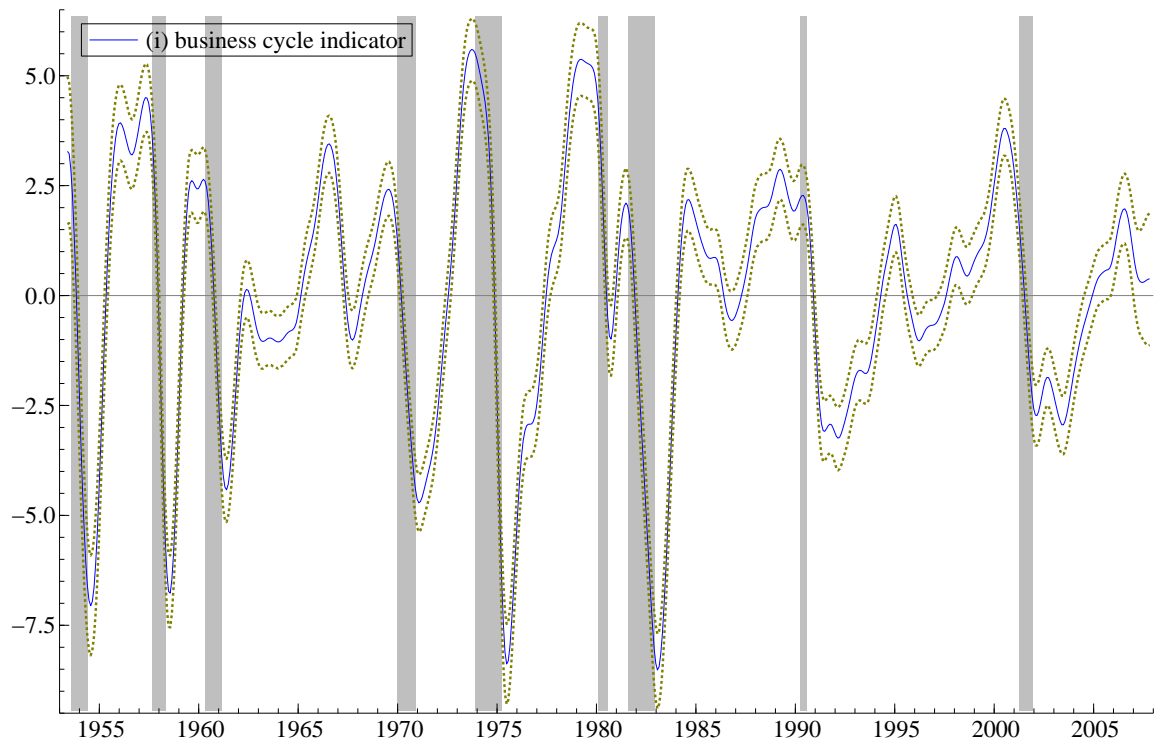
Figure 2: *Business cycle indicator with 95% HPDI.* NBER *recession dates are represented by the vertical bands.*

functions are centered over the business cycle frequences with their peaks varying from 0.08 to 0.14. This is despite the fact that the posterior estimate of λ in Table 2 is 0.151, which demonstrates the known fact that λ does not correspond exactly to the middle of the gain function for this model. Finally, it is interesting to see that the gain function in 1997 has adjusted to include more high frequency movements from each series. This typically leads to more short-lasting cyclical movements in the time series. For example, the estimated cycle period in 1982 is 6.5 years while in 1997 it is 4.4 years, approximately.

4.3 Stochastic volatility estimates

A major goal of this paper is the development of a monthly business cycle indicator with bandpass-filter properties that accounts for time-varying volatility. We have not intended to build a model to explain the joint data-generating process of all 11 time series. However, it remains interesting to compare the estimated stochastic volatility components from our procedure to the 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).

16

	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	-0.255	0.010	0.00	2.50	-1.013	0.185
manufacturing	0.00	2.00	0.352	0.014	0.00	2.50	3.923	0.228
inflation	0.00	2.00	0.318	0.034	0.00	2.50	-3.883	0.614
retail	0.00	2.00	0.679	0.033	0.00	2.50	2.983	0.343
retail/food	0.00	2.00	0.506	0.092	–	–	–	–
productivity	0.00	2.00	0.312	0.024	0.00	2.50	8.538	0.584
real GDP	0.00	2.00	0.548	0.020	0.00	2.50	3.417	0.259
hours	0.00	2.00	0.546	0.021	0.00	2.50	0.621	0.264
consumption	0.00	2.00	1.245	0.061	0.00	2.50	4.682	0.370
investment	0.00	2.00	2.662	0.103	0.00	2.50	2.713	0.260
	Prior ρ		Posterior ρ		Prior λ		Posterior λ	
	0.50	0.08	0.974	0.004	0.09	0.10	0.151	0.012

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

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, Nelson, and Piger (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 4(i). It contains a marked 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 4 and 5. The irregular component of real GDP in Figure 5(ii) indicates a moderate decline in volatility beginning in 1979 and ending in 1984. Other series such as manufacturing, consumption of durables, and 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. This agrees with the interpretation given in Stock and Watson (2002b) and Primiceri (2005) that most of the volatility changes are in the idiosyncratic component. Finally, we observe that the volatility of quarterly CPI inflation pictured in Figure 4(iv) appears to be increasing over the last two years.

The differences between our estimates and those of Chauvet and Potter (2001) and Kim, Nelson, and Piger (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
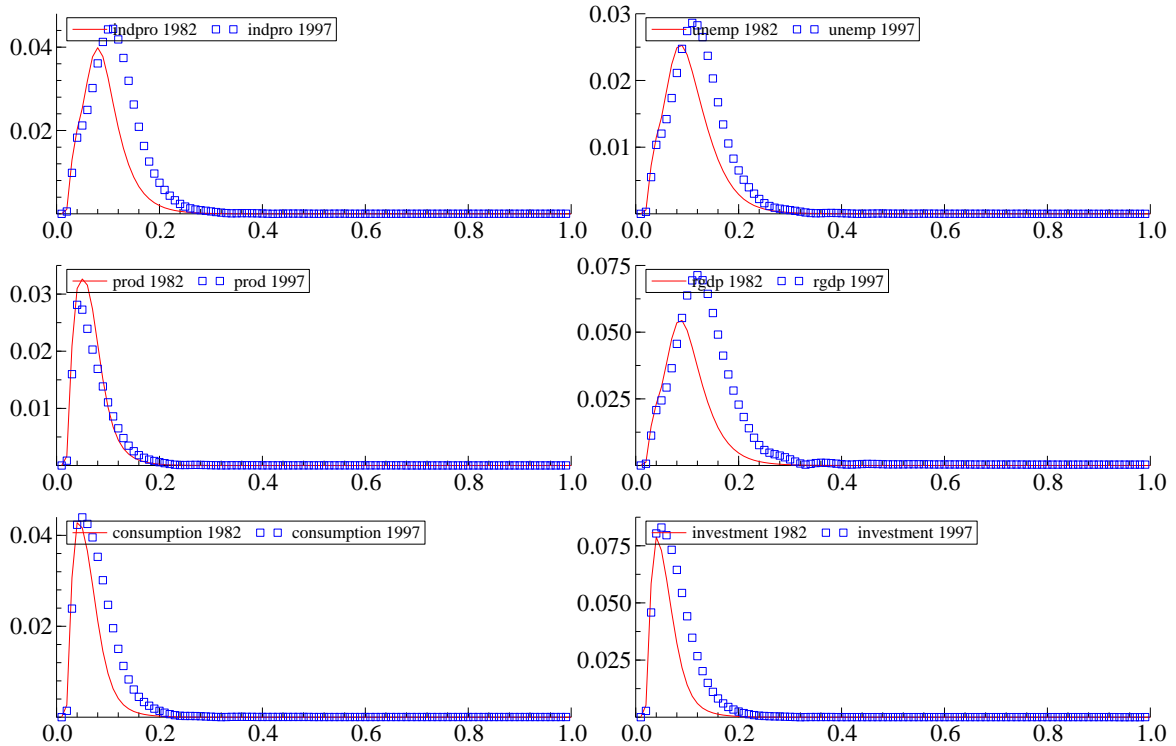
Figure 3: *Gain functions from Q1 1982 (straight lines) and Q1 1997 (boxes) for clockwise from top left: industrial production, unemployment, real GDP, investment, consumption, and productivity.*

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.

4.4 Robustness of our findings

In this section, we investigate the robustness of our main findings discussed in the previous sections. First, we consider the multivariate trend-cycle model without stochastic volatility components. Second, the main results give evidence of considerable changes in the volatility of most of the series considered. However, we also inspect the model to detect possible changes that are not accounted for by time-varying volatility.

Model without stochastic volatility and mixture innovations

In section 2.1 we have introduced the model without stochastically time-varying log-variances and mixture innovations. Some elaborations of the MCMC algorithm for the linear Gaussian
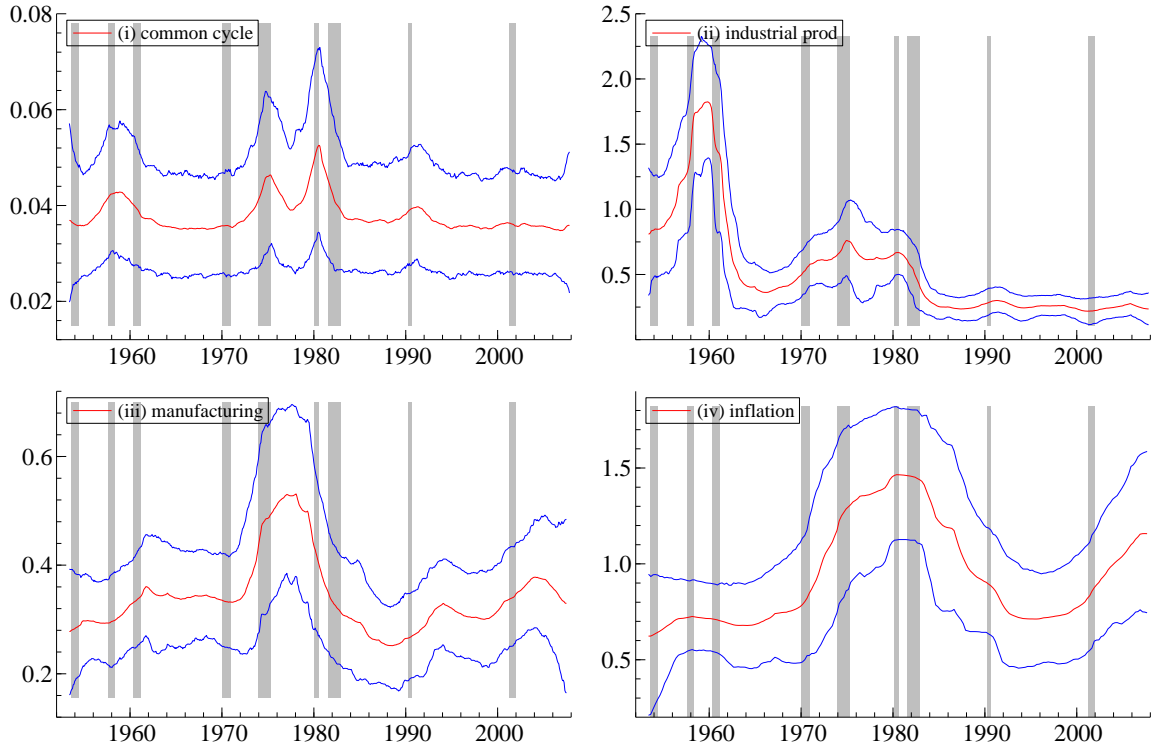
Figure 4: (i) Smoothed estimates of the stochastic volatility $\sigma_{t,\kappa}$ of the common cycle with 95% HPDI. Smoothed estimates of the stochastic volatility $\sigma_{i,t,\varepsilon}$ of the irregular components with 95% HPDI for (ii) industrial production; (iii) manufacturing; and (iv) inflation. NBER recession dates are represented by the vertical bands.

model (2) – (7) reduces since we do not need to integrate over $\chi_{1:n}$ and $\gamma_{1:n}$. The top left graph in Figure 6 provides estimates of the cycle and slope from a model without stochastic volatility. The amplitude of the estimated cycle is larger while the peaks and troughs continue to match the NBER recession dates. The largest difference between this cycle and the estimated cycle from the model with stochastic volatility comes in 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 the final months of 2007.

Estimates of some of the parameters, ξ_i, ρ, and λ, for the current model are reported in Table 3 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 the previous section. Unemployment and inflation lag industrial production and real GDP by roughly two more months. The estimated persistence of the businesss cycle ρ is also smaller. However, we note that the order in which the cycles are estimated (in terms of the phase shifts) remains the same in both specifications.
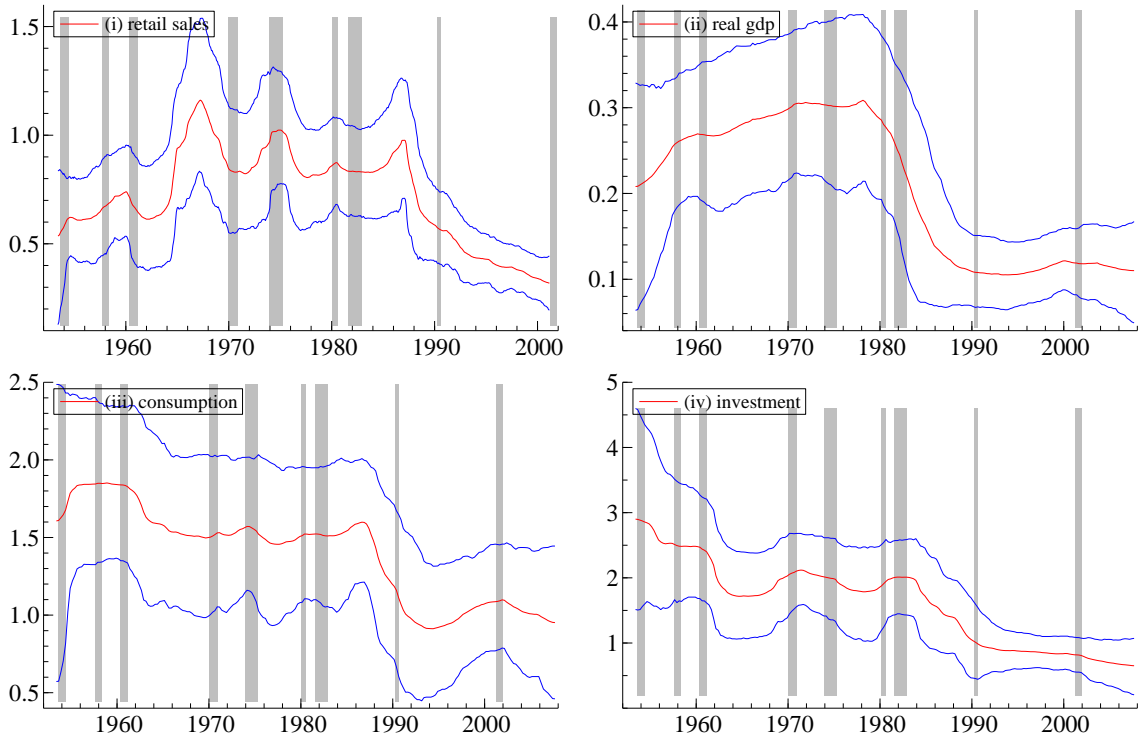
19

Figure 5: *Smoothed estimates of the stochastic volatility* $\sigma_{i,t,\varepsilon}$ *of the irregular components with 95% HPDI for (i) real retail sales; (ii) real GDP; (iii) consumption of durables; and (iv) investment.* NBER *recession dates are represented by the vertical bands.*

The model with breaks in Q1 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 of the model, we reestimate the model from sections 4.1 – 4.3 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 3 contain estimates of the parameters in the model with a break; additional results are available in the online Appendix. 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. Interestingly, the value of λ is estimated to decrease from 0.160 to 0.112 after 1984. This is consistent with the fact that on average periods between business cycles are now longer than in the past. Estimates

	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	-2.401	0.292	-1.058	0.164	-2.329	0.560
manufacturing	4.132	0.373	3.284	0.201	8.015	0.798
inflation	-5.731	0.958	-4.416	0.602	-1.247	1.702
retail	3.796	0.461	2.362	0.325	5.271	1.154
productivity	9.323	0.751	7.386	0.528	10.150	1.866
real GDP	3.732	0.395	2.796	0.233	6.561	0.822
hours	-0.111	0.313	0.497	0.247	-0.079	0.929
consumption	4.724	0.503	4.176	0.366	5.195	1.307
investment	3.086	0.322	2.406	0.286	4.206	0.585
ρ	0.873	0.009	0.981	0.003	0.971	0.005
λ	0.120	0.009	0.160	0.003	0.112	0.007

Table 3: *Posterior means and standard deviations for ρ, λ, 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.

from the model also 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 estimated cycle and slope from this specification are shown in panels (iii) and (iv) of Figure 6. 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 affect the estimated business cycle indicator. The indicator shares the same features as the indicator produced from the main model in Figure 1 without the structural break. Estimates of the stochastic volatility components from this specification are also similar to those found in the Figures 4 and 5 and are available in the online Appendix.

5   Conclusion

In this paper, we propose the construction of a business cycle indicator that explicitly accounts for the time variation in macroeconomic volatility commonly known as the great moderation. Our indicator is constructed from a multivariate unobserved components time series model with a
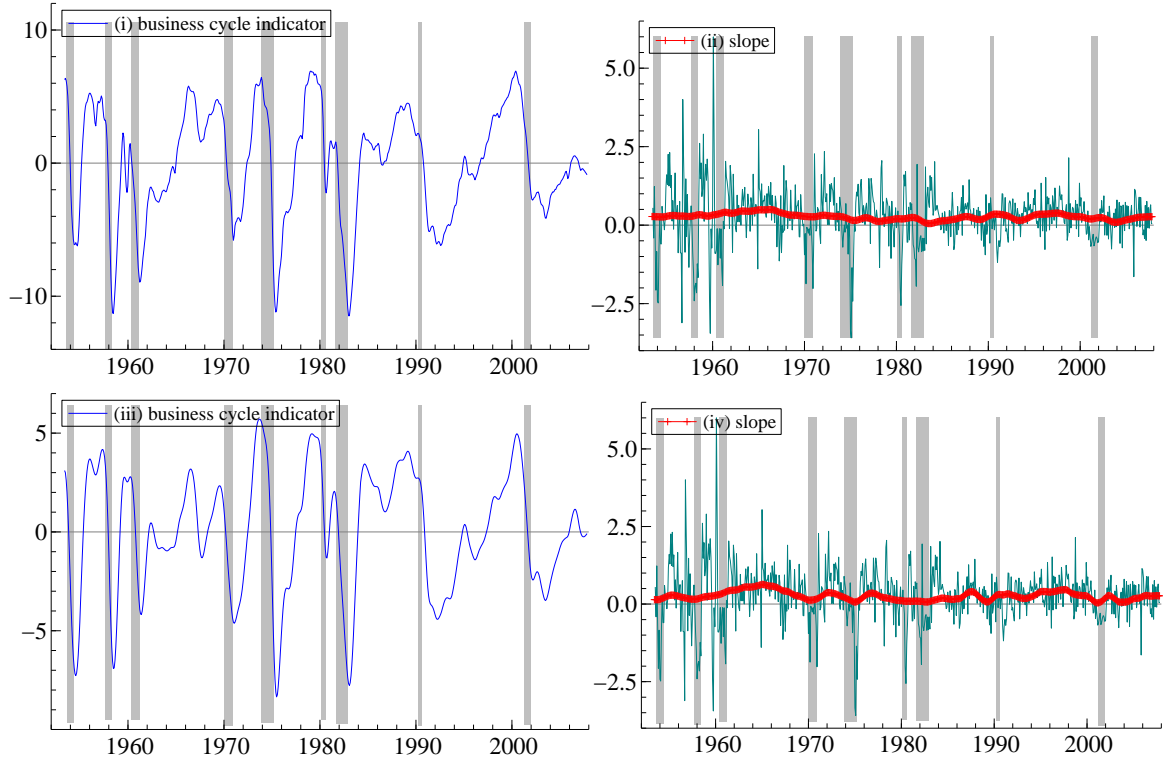
Figure 6: *(i) Business cycle indicator from a model with no stochastic volatility; (ii) smoothed estimates of the slope from a model with no stochastic volatility; (iii) Business cycle indicator from a model with a known break in Q1 1984; (iv) smoothed estimates of the slope from a model with a known break in Q1 1984.* NBER *recession dates are represented by the vertical bands.*

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. All parameters are estimated simultaneously using the MCMC algorithm that consists of Gibbs sampling and Metropolis-Hastings steps. We interpret our approach as a model-based bandpass filter method because the extracted cycle emphasizes the business cycle frequencies that are of interest to applied researchers and policymakers. The methods are applied to a panel of eleven macroeconomic time series which are indicative for the U.S. business cycle. The empirical results reveal that unemployment is counter-cyclical to all other macroeconomic time series, the cycle in inflation lags those in other variables while productivity, consumption and manufacturing (in this order) are leading the U.S. business cycle. The estimated volatility patterns are significantly changing over time for the irregular (high frequency) components of real retail sales, real GDP, durable consumption and investment. However, the disturbances associated with the common business cycle appear not to be subject to stochastic volatility processes and therefore we regard the

dynamic behaviour of the business cycle as stable over time. These empirical findings have so far not been documented in the context of business cycle studies. We have found that our main empirical findings are robust towards other model specifications. However, some evidence is given that the duration and the persistence of the business cycle are subject to structural breaks before and after the first quarter of 1984. Nevertheless, such breaks do not alter the other empirical findings when using our flexible modeling framework.

Acknowledgements

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, 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.

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

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). *Ox: An Object-Oriented Matrix Programming Language*. London: Timberlake Consultants Press.

Doucet, A. and C. Andrieu (2001). Iterative algorithms for state estimation of jump Markov linear systems. *IEEE Transactions on Signal Processing 49*, 1216–1227.

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

Durbin, J. and S. J. Koopman (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika 89*, 603–616.

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.

Gerlach, R., C. Carter, and R. Kohn (2000). Efficient Bayesian inference for dynamic mixture models. *Journal of the American Statistical Association 95*, 819–828.

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.

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 (1999a). 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. and C. R. Nelson (1999b). *State Space Models with Regime Switching Classical and Gibbs Sampling Approaches with Applications*. Cambridge, MA: MIT Press.

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.

Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility: likelihood inference and comparison with ARCH models. *The Review of Economic Studies 65*(3), 361–393.

Koopman, S. J. (1997). Exact initial Kalman filtering and smoothing for non-stationary time series models. *Journal of the American Statistical Association 92*(440), 1630–1638.

Koopman, S. J. and A. C. Harvey (2003). Computing observation weights for signal extraction and filtering. *Journal of Economic Dynamics and Control 27*(7), 1317–1333.

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.

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.

Omori, Y., S. Chib, N. Shephard, and J. Nakajima (2007). Stochastic volatility with leverage: fast and efficient likelihood inference. *Journal of Econometrics 140*(2), 425–449.

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 (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.

Trimbur, T. (2006). Properties of higher-order stochastic cycles. *Journal of Time Series Analysis 27*(1), 1–17.

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.

Appendix: some details of the 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 $j = 1, ..., N$ by an algorithm consisting of four main steps. The fourth step requires the sampling of parameters from different distributions.

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.

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, Carter, Kohn (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 developed by Kim, Shephard, and Chib (1998). In practice, we use the 10 component mixture recently provided in Omori, Chib, Shephard, and Nakajima (2007).

Setp 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 $\psi_t^{(k)}, \psi_t^{(k)+}$, 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 $\sigma_{i,\zeta}^2$*: Given the slopes β_{it}, the inverse gamma prior on these parameters mean that each variance $\sigma_{i,\zeta}^2$ can be drawn from its full conditional distribution via a Gibbs sampling step. Draws from the full conditional distributions which are inverse gamma distributions are standard; e.g., see Kim and Nelson (1999 pp. 175).

(d) *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% 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. See Robert and Casella (2004) for more details on the Metropolis-Hastings algorithm.