BAYES ESTIMATES OF THE CYCLICAL COMPONENT IN TWENTIETH CENTURY US GROSS DOMESTIC PRODUCT

Econometric Institute Report EI 2004-45

By Andrew C. Harvey[a], Thomas M. Trimbur[a,b], and Herman K. van Dijk[c]

[a] *Cambridge University, Faculty of Economics and Politics*

[b] *U.S. Census Bureau, Washington DC*

[c] *Econometric Institute, Erasmus University*

Cyclical components in economic time series are analysed in a Bayesian framework, thereby allowing prior notions about periodicity to be used. The method is based on a general class of unobserved component models that encompasses a range of dynamics in the stochastic cycle. This allows for instance relatively smooth cycles to be extracted from time series. Posterior densities of parameters and estimated components are obtained using Markov chain Monte Carlo methods, which we develop for both univariate and multivariate models. Features such as time-varying amplitude may be studied by examining different functions of the posterior draws for the cyclical component and parameters. The empirical application illustrates the method for annual US real GDP over the last 130 years.

Keywords: Band pass filter, Business cycles, Gibbs sampler, Markov chain Monte Carlo, unobserved components.

JEL classification: C11, E32.

1 Introduction

Decomposing time series into trends and cycles is fundamental to a good deal of macroeconomic analysis. The Hodrick-Prescott (HP) filter is often used to detrend series but as shown in Harvey and Jaeger (1993) and Cogley and Nason (1995), using it inappropriately can result in the creation of spurious cycles. The same is true of the band pass filter recently proposed by Baxter and King (1999) for extracting cyclical movements over the range two to ten years; see Murray (2003).

Harvey and Jaeger (1993) argued that detrending is best accomplished by fitting a structural time series model consisting of trend, cycle and irregular unobserved components. In the classical approach, the model is estimated in state space form with the components extracted by the Kalman filter and associated smoother. However, fitting the model to

[1]A more extended version of this paper is available as: Cyclical components in economic time series: a Bayesian approach, Econometric Institute Report 2002-20 in revision.

Corresponding author : Thomas M. Trimbur. Mailing address - U.S. Census Bureau; 4700 Silver Hill Road; Washington, D.C. 20024; United States. Telephone - (00) 1 (301) 763-4973. Email - Thomas.M.Trimbur@census.gov.

series like GDP usually results in the irregular component nearly disappearing with the result that the cycle is quite noisy. The class of higher order stochastic cycles introduced in Harvey and Trimbur (2003) enables one to overcome this feature. Trimbur (2004) provides the exact characterization of this class of models, showing analytical expressions for the time and frequency domain properties.

Using higher order cycles in the unobserved components model for the series yields implicit filters that concentrate on extracting relatively more power from a narrower band of frequencies. More high frequency noise is forced into the irregular thereby yielding a smoother cycle. A perfectly sharp band pass filter, typically referred to as the "ideal" filter, emerges as a limiting case. In empirical work it has often been taken for granted that such a filter naturally provides the best basis for extracting cycles - this view is implicit in the terminology itself - and the "ideal" filter has been emulated in numerous applications with time series data. Examining the models that underpin different types of filters provides insight as to their relative appropriateness. In the general case, the gain tapers off gradually at the band edges, in contrast to the "ideal" filter, which exhibits discontinuity at the boundaries, a shape that is in some sense less natural, particularly in economic applications. The biggest concern, however, remains the potential for distortions of the type demonstrated in Murray (2003).

The model-based estimators of cyclical components, introduced in Harvey and Trimbur (2003), enable one to mitigate such distortions. These classical estimators are referred to as generalized Butterworth band-pass filters, given their close links with the Butterworth filters commonly used in engineering. The goal of this paper is to develop the complementary Bayesian analysis of models with higher order cycles so that parameter uncertainty is accounted for, with the estimated posterior densities providing compact and informative summaries of trend and cyclical dynamics. We show empirical application with the annual US real GDP series. Essentially, the paper introduces Bayesian band-pass filters for economic time series.

Prior knowledge on periodicity is readily available in the area of business cycle analysis. This paper introduces a Bayesian approach, whereby parameter uncertainty in accounted for and expectations about cyclical dynamics are updated optimally, producing a posterior probability density of the model parameters and component series. In a classical framework, fixed estimates for the parameters in structural time series models are usually obtained by maximum likelihood (ML) using the Kalman filter. Harvey and Trimbur (2003) found that this works well for series like investment where the cycle is pronounced. For some cases, however, it may be difficult to obtain plausible parameter estimates on the basis of the likelihood surface alone. In any case, the classical approach is unable to provide a satisfactory way to implement any prior knowledge that may be available about the period of oscillation.

In estimating economic cycles, prior information on the frequency range of interest has clearly played a role in a good deal of empirical work. Typically, however, such information has been implemented by imposing stringent conditions on the desired pattern of the gain function of the filter, for instance in Baxter and King (1999) where an approximation to an "ideal" filter with prespecified boundaries is proposed. The concerns associated with the use of "ideal" filters in economics were noted above. By using prior knowledge in a more consistent manner, a model-based framework is able to address these concerns.

Specifically one may construct a Bayesian prior for the frequency parameter of the cyclical component as part of the model; the prior distribution is then updated, following Bayes rule, conditional on the information in the data. Prior information on other parameters may also be used, though this will typically be rather more vague. The resulting joint posterior density of parameters and components may then be used to provide a great deal of insight on many different aspects of the time series dynamics.

The Bayesian approach also has the practical advantage that it allows one to address potential finite sample problems. In particular, for any given model and series of limited length, the likelihood surface may exhibit some irregularities, with reasonable parameter estimates impossible to attain in an unrestricted classical framework. In Harvey and Trimbur (2003), difficulties were sometimes experienced in obtaining a sensible period for the cycle in real GDP; in such cases, estimation with a fixed period led to acceptable results. However, the use of suitable priors, within a Bayesian framework, enables one to address the problem more effectively in a consistent manner. Greater flexibility is provided, and generally, a Bayesian approach may be feasible for a broader range of applications with different model structures and sample sizes.

This paper presents a Markov chain Monte Carlo (MCMC) algorithm for a Bayesian analysis of stochastic cycles in time series. The method we develop is a new composite of the Gibbs sampler and Metropolis-Hastings algorithms. With the aid of state space modeling techniques, we set out an efficient procedure for computing the joint posterior density of parameters and components. This provides a great deal of information that may be used to study a variety of features of cyclical dynamics. For instance, one may characterize the complete distribution of the cyclical component series. As the algorithm allows for the computation of marginal posteriors of the trend and cyclical component at different points in time, based on the percentiles of the estimated distributions, bands showing the highest posterior density regions of the cycle and trend may be constructed.

The area of Bayesian analysis of dynamic econometric models has benefitted from developments in state space and MCMC methods in recent years. Some important contributions have been Carter and Kohn (1994), Fruhwirth-Schnatter (1994), deJong and Shephard (1995), and Durbin and Koopman (2002). The paper by Huerta and West (1999) sets out a Bayesian treatment of cyclical behaviour indirectly through autoregressive models. However, there has been no attempt to conduct a Bayesian analysis of cycles in an unobserved components framework.

Fruhwirth-Schnatter (1994) and Koop and van Dijk (2000) analyse trends and seasonals in various macroeconomic series. The treatment of cycles leads to a number of new technical issues which we address in this paper and the Bayesian approach allows us to present the results of model fitting in an informative way. As outlined in Harvey, Trimbur, and van Dijk (2002), the analysis may be naturally extended to multivariate unobserved components models. Again this has not been attempted in a Bayesian framework and it introduces a number of issues that require careful consideration.

The rest of the paper is arranged as follows. Section 2 describes the extension of the class of structural time series models to include higher order cyclical components. The Bayesian treatment is developed in section 3, while section 4 illustrates the method with an applications to an annual US macroeconomic series. Technical details on the state space form and the Markov chain Monte Carlo algorithms for the univariate and multivariate cases are set out in Harvey, Trimbur, and van Dijk (2002).

2 Cyclical and trend components in time series

We consider a class of unobserved component (UC) models in which the observations y_t, $t = 1, ..., T$, are made up of a nonstationary trend component μ_t, a cyclical component $\psi_{n,t}$, and an irregular term ε_t. Thus:

$$y_t = \mu_t + \psi_{n,t} + \varepsilon_t, \quad t = 1, ...T, \tag{1}$$

where the irregular is white noise, that is $\varepsilon_t \sim WN(0, \sigma_\varepsilon^2)$, the stochastic trend is an integrated random walk

$$\mu_t = \mu_{t-1} + \beta_{t-1}, \tag{2}$$

$$\beta_t = \beta_{t-1} + \zeta_t, \quad \zeta_t \sim WN(0, \sigma_\zeta^2) \tag{3}$$

while the $n - th$ order cycle is defined by

$$\begin{bmatrix} \psi_{1,t} \\ \psi_{1,t}^* \end{bmatrix} = \rho \begin{bmatrix} \cos\lambda_c & \sin\lambda_c \\ -\sin\lambda_c & \cos\lambda_c \end{bmatrix} \begin{bmatrix} \psi_{1,t-1} \\ \psi_{1,t-1}^* \end{bmatrix} + \begin{bmatrix} \kappa_t \\ \kappa_t^* \end{bmatrix}, \tag{4}$$

$$\begin{bmatrix} \kappa_t \\ \kappa_t^* \end{bmatrix} \sim WN\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_\kappa^2 & 0 \\ 0 & \sigma_\kappa^2 \end{bmatrix} \right)$$

$$\begin{bmatrix} \psi_{i,t} \\ \psi_{i,t}^* \end{bmatrix} = \rho \begin{bmatrix} \cos\lambda_c & \sin\lambda_c \\ -\sin\lambda_c & \cos\lambda_c \end{bmatrix} \begin{bmatrix} \psi_{i,t-1} \\ \psi_{i,t-1}^* \end{bmatrix} + \begin{bmatrix} \psi_{i-1,t-1} \\ \psi_{i-1,t-1}^* \end{bmatrix}, \quad i = 2, ..., n \tag{5}$$

The parameter λ_c denotes frequency in radians while ρ is a damping factor lying between zero and one; if it is equal to one, the cycle is nonstationary. The disturbances driving the trend and cycle are assumed to be uncorrelated with each other and with the irregular. The specification in (1) is well-suited for trend-cycle decompositions. The assumption in using model (1) directly is that the seasonal component has been removed. Alternatively, the raw (non-seasonally adjusted) data could be used, with the model augmented by adding a seasonal component.

Each type of stochastic cycle is defined in terms of a number of processes. The nth order stochastic cycle $\psi_{n,t}$ has periodic movements centered around a frequency of λ_c. The stochastic movements stem from the two disturbances, κ_t and κ_t^* in (4). Suppressing κ_t^* in the model specification yields a class of band pass filters that generalises the Butterworth class of filters; see Gomez (2001) and Harvey and Trimbur (2003). Our preference here is to work with the 'balanced form' of (4); for $n = 1$ this is identical to the stochastic cycle in (Harvey 1989, p. 39) and as shown in Trimbur (2004), analytical expressions for key properties are available for all n. The different spectral shapes for different orders give an illustration of the range of dynamics of the class of higher order cycles. With the cyclical parameters fixed, the spectrum becomes sharper as the order increases, concentrating around the central frequency. Further details may be found in Trimbur (2004).

State space methods play a key role in both classical and Bayesian treatments of the class of models (1). The state space formulation and dynamic characteristics of the higher order cycles are set out in Trimbur (2004). In the next section, we devise an efficient Bayesian approach, where components are estimated in a way that accounts for parameter uncertainty. This allows for an informative analysis of cyclical behavior with flexible expectations on periodicity.

3 Bayesian treatment

The three variance parameters and two cyclical parameters are arranged in the vector $\boldsymbol{\theta} = \{\sigma_{\zeta}^2, \sigma_{\kappa}^2, \sigma_{\varepsilon}^2, \rho, \lambda_c\}$. The model is assumed to have Gaussian disturbances throughout. Given a sample $\mathbf{y} = \{y_1, ..., y_T\}$, the likelihood function is specified by the model structure. Below we introduce a flexible set of priors for $\boldsymbol{\theta}$.

The goal is then to analyze the properties of the posterior distribution, $p(\boldsymbol{\theta}|\mathbf{y})$. Since this is not a member of a class of densities which has known analytical properties, a new Markov chain Monte Carlo routine is developed. The method produces draws from the posterior density of the parameters and components, giving smoothed estimates of the cycle as a by-product. The approach can be adapted to different kinds of priors, and in the next section we also set out the extension to multivariate model structures.

3.1 Priors and likelihood

We start by summarising the elicitation of priors. The direct interpretation of the cycle parameters makes it straightforward to design suitable priors; they are linked to economic intuition and previous experience of studying business cycles. The parameter λ_c represents a central frequency; the annual sample we investigate below includes cyclical swings that vary a great deal in their duration, and the expectation is that, on average, fluctuations in the cycle will have a period of around ten years. A standard peaked distribution for λ_c may be used by the researcher to reflect this expectation in a consistent and adaptable framework.

For annual macroeconomic data, we consider priors for λ_c centred around $2\pi/10$, based on the class of beta distributions; such priors are flexible, covering a variety of possibilities, and they are easy to work with analytically. We consider different degrees of concentration in the prior around the mean. The least informative prior covers a wide range of frequencies, while the sharpest density focuses attention narrowly around a period of ten years. Other priors, with different spreads and locations, may be implemented in this framework. If a particular shape of prior were desired, then another class of densities could be used and the algorithm would be constructed in the same way as shown below.

The informative priors we implement all reflect an emphasis on average business cycle periodicity of ten years.

For technical details on the class of priors we use, see appendix A in Harvey, Trimbur, and van Dijk (2002). The approach taken by Huerta and West (1999) uses priors for the autoregressive parameters implied by the reduced form of the first order cycle. We find it more useful to focus on λ_c. In any case adapting the Huerta and West approach to higher order cycles would not be straightforward. The parameter ρ is linked to the order of the cycle. In the first order case ρ is the rate of decay of the cycle, but for higher orders the interpretation of ρ changes somewhat so that different values are appropriate. However, since the precise form of the relationship between ρ and n is not clear, we use a uniform prior on ρ over the interval [0,1].

For the variance parameters, we use inverted gamma prior densities with shape and scale set to near zero, which serve as effectively noninformative. The class of inverted gamma priors are conditionally natural conjugate so that the conditional posteriors are also inverted gamma, and thus one has the advantage of direct simulation within the Gibbs sampler routine. Also, the flexibility of this class of density allows for implementing any prior knowledge that may be available for the variance parameters. Thus, for instance, if there is the expectation of stochastic trend dynamics, such an expectation may be expressed as a prior with positive mean and a degree of dispersion to reflect uncertainty in the exact value of the disturbance variance.

Next, we summarise the evaluation of the likelihood and posterior of our unobserved components model. The sample $\mathbf{y}$ represents the observed realization of the data generating process, which is a multivariate density $p(Y|\boldsymbol{\theta})$ over all possible realizations Y. As normality of the disturbances is assumed, the likelihood function, $L(\boldsymbol{\theta};\mathbf{y}) = p(\mathbf{y}|\boldsymbol{\theta})$, can be evaluated for any permissible value of $\boldsymbol{\theta}$ using the Kalman filter. This relies on the prediction error decomposition as described in Harvey (1989, p. 126). With an appropriate initialisation, the density $p(Y|\boldsymbol{\theta})$ is multivariate Gaussian so that computation of the likelihood for a given sample is straightforward.

The densities $\{p(\boldsymbol{\theta}), p(Y|\boldsymbol{\theta})\}$ give the full description of the model in the Bayesian context. Given the data $\mathbf{y}$, the prior-likelihood pair $\{p(\boldsymbol{\theta}), L(\boldsymbol{\theta};\mathbf{y})\}$ defines the analysis. The expectations reflected in $p(\boldsymbol{\theta})$ are updated using the information in the $\mathbf{y}$, through $L(\boldsymbol{\theta};\mathbf{y})$. By studying the characteristics of the posterior, the analysis addresses various questions about cyclical and trend dynamics.

3.2 Posterior

The posterior $p(\boldsymbol{\theta}|\mathbf{y})$ is proportional to the product of the prior and likelihood. However, the expression for the product $p(\boldsymbol{\theta})L(\boldsymbol{\theta};\mathbf{y})$ does not represent the kernel of a known distribution. The normalizing constant (equal to the marginal likelihood) required for evaluating the posterior ordinate is not known in terms of elementary functions (like $\sqrt{2\pi}$ in the case of the normal distribution).

A strategy is needed for analyzing the properties of $p(\boldsymbol{\theta}|\mathbf{y})$. With a five-dimensional parameter vector, MCMC methods offer an efficient way to sample (pseudo-random) parameter drawings from the posterior. This may be used to produce drawings of regular functions of the parameters. Thus for instance, finite sample results on posterior moments may be compared with ML estimates.

The strategy for posterior analysis is based on extending the parameter space to include the components and associated auxiliary processes in (1), which together form the basis for the state space form of the model. As described in Harvey, Trimbur, and van Dijk (2002), we design an MCMC routine that is able to capitalize on recent developments in state space modeling. Thus, the simulation smoothing techniques developed, for instance, in deJong and Shephard (1995) and Durbin and Koopman (2001), may be efficiently used in a Gibbs sampling setup. In this way, the algorithm is set up to produce drawings from the expanded density of both the parameters and unobserved components. This provides an efficient route for obtaining draws from $p(\boldsymbol{\theta}|\mathbf{y})$, and it also gives additional information that is useful for studying the trend and cycle.

3.3 Signal extraction

The MCMC method produces drawings from the joint density of the two unobserved components, the cycle and trend, over the sample period. These high-dimensional variates, conditional on the data, can be used to compute a Bayesian analogue of the classical smoother. As they form part of the Gibbs sampler, no additional effort is required to obtain them. Classical estimates correspond to the conditional means of the cyclical component, given the sample, assuming the parameters are fixed at estimated values. The *Bayesian smoother* incorporates parameter uncertainty (the parameter vector is integrated out) and accounts for prior knowledge about the period.

The Bayesian analysis produces draws from the joint posterior of the trend and cyclical components, $p(\mu_1, ..., \mu_T, \psi_{n,1}, ..., \psi_{n,T}|\mathbf{y})$. We note that the posterior mean is the optimal estimator for a quadratic loss. The estimated component series are obtained by averaging over the J state draws, that is

$$\widehat{\mu}_t = \frac{1}{J}\sum_{j=1}^{J}\mu_t^{(j)}, \qquad \widehat{\psi}_{n,t} = \frac{1}{J}\sum_{j=1}^{J}\psi_{n,t}^{(j)}, \quad t = 1, ..., T,$$

where $\mu_t^{(j)}$ denotes the jth draw for the trend at time t and similarly for the cyclical component. The standard deviation of the trend estimate at each time point is given by $\sqrt{\sum \mu_t^{2(j)}/J - \widehat{\mu}_t^2}$ and other higher-order moments, for the various components, may be computed in a similar way.

Drawings of regular functions of the trend, cycle, and other state vector elements over the sample are directly obtained. This enables properties of the time-varying cycle, such as amplitude, to be studied. The amplitude of the cycle at time t is estimated by

$$A_t = \frac{1}{J}\sum_{j=1}^{J}\sqrt{\psi_{n,t}^{2(j)} + \psi_{n,t}^{*2(j)}}, \quad t = 1, ..., T$$

Other features of interest of the unobserved components may be studied analogously.

3.4 Annual US GDP

Annual time series are available over a fairly long period of time and this allows one to investigate issues concerning long-term changes in the business cycle. We examine annual US real GDP data from 1870 to 1998 compiled from the OECD publications *Monitoring the World Economy* and *The World Economy: a Millennial Perspective.* The enormous swing from the beginning of the Great Depression to the end of World War II constitutes a much longer and more pronounced cycle than is found in post-war data. The priors we consider for λ_c have a mean of $2\pi/10$; thus we allow for relatively long periods with an average of around ten years. In addition we assume that the variance of the cyclical disturbances from 1929 to 1946 is ten times what it is elsewhere. The state space model has no difficulty handling such an extension. (A similar device could have been adopted for the frequency, that is the period from 1929 to 1946 could have been assumed to be double what it is elsewhere).

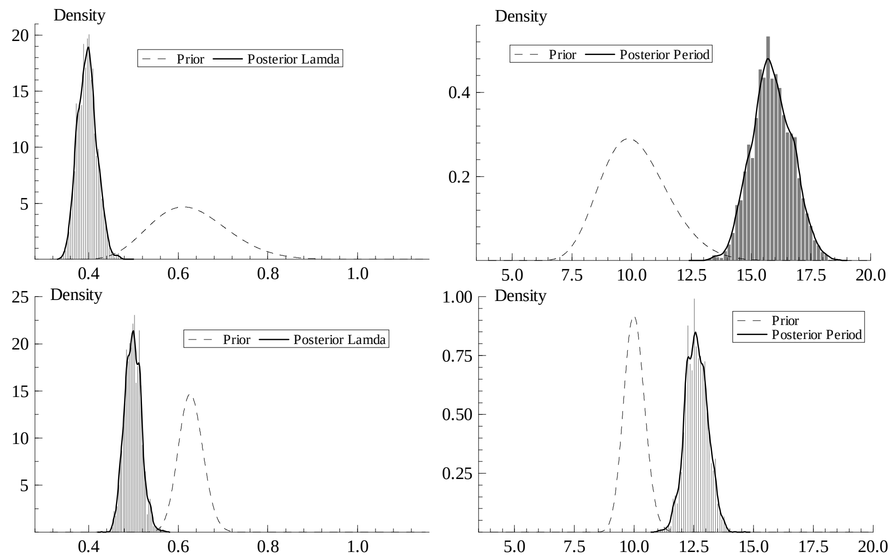

Figure 1: Marginal posterior densities of frequency and period with intermediate and very informative priors on λ_c for annual US real GDP (logarithms) from 1870 to 1998.

The results for $n = 1$ with the least informative prior on λ_c show that the posterior distribution of period concentrates around 17.5 years, owing to the dominating influence of the Great Depression and the subsequent World War II recovery. It may be preferable, therefore, to use a sharper prior. The posterior densities for frequency and period are shown in figure 1 for intermediate and very informative priors, and the resulting trend and cycle for the very informative prior is shown in figures 2 and 3. Note that the data are in logarithms.

The series of trend and cycle estimates over the sample period are computed as posterior means, obtained by averaging over the J state draws as described in sub-section 3.3. The HPD (Highest Posterior Density) regions for the estimated series are shown as well.

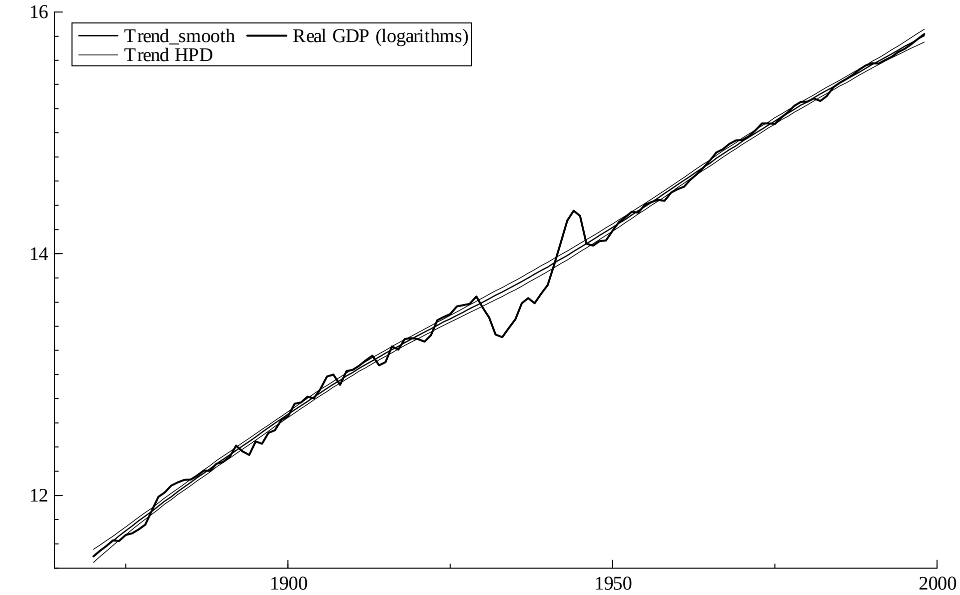

Figure 2: Estimated trend in annual US real GDP (logarithms) from 1870 to 1998 for $n = 1$ with most informative prior on λ_c.

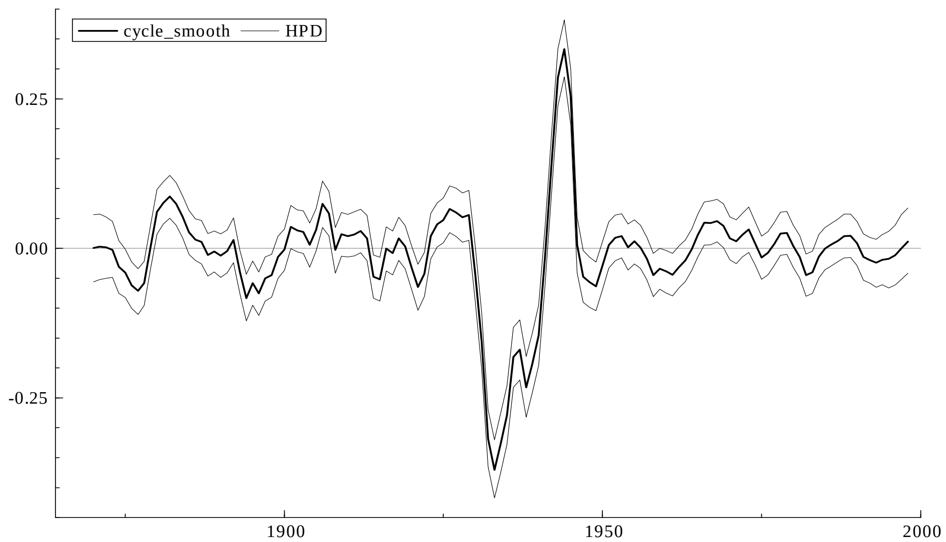

Figure 3: Estimated cycle in annual US real GDP (logarithms) from 1870 to 1998 for $n = 1$ with most informative prior on λ_c.

The HPD regions in figures 2 and 3 are obtained by taking the 2.5 and 97.5 percentiles, and we will refer to the series of HPD regions as HPD bands[1]. These resemble classical 95% confidence intervals but their interpretation is distinct; the bands in figures 2 and 3 give exact finite sample measures of uncertainty. The Bayesian smoother refers to the actual distribution of the trend and cycle conditional on the dataset. The associated HPD bands incorporate the posterior uncertainty in components and parameters. To our knowledge, no one has given, so far, finite sample confidence bands around the trend and cycle components.

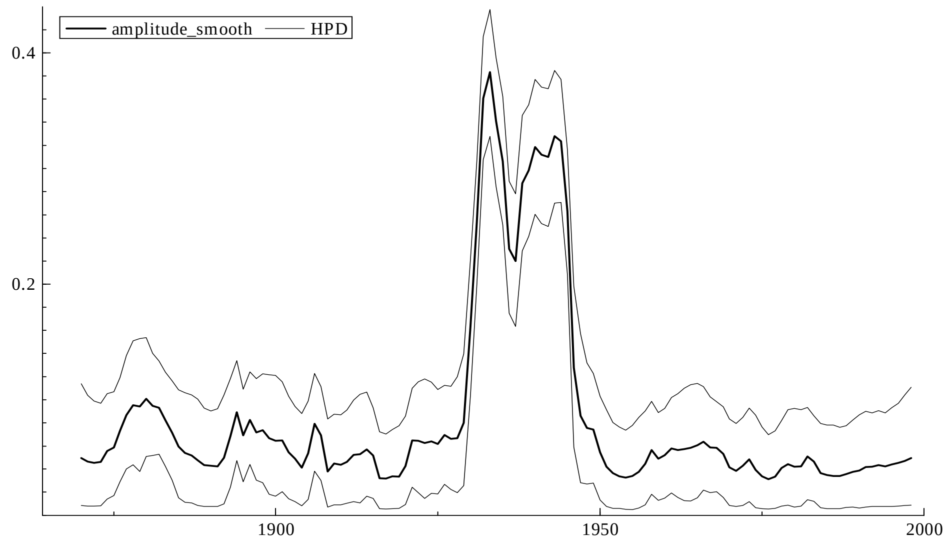

Figure 4: Evolving amplitude of the cyclical component in annual US real GDP (logarithms) from 1870 to 1998 for $n = 1$ with most informative prior on λ_c.

An interesting question is whether business cycles have generally declined in intensity since World War II. Figure 4 displays the evolving amplitude. The graph is dominated by the Great Depression and World War II, but there is some indication that the average amplitude of the cycle is smaller in the post war period as compared with the period before 1929. This contradicts the findings of Backus and Kehoe (1992) from a study based on detrending with the Hodrick-Prescott filter[2].

[1]The precise definition of the 95% highest posterior density (HPD) region is the minimum interval that contains 95% of the probability mass. For symmetric distributions this gives a result identical to the 2.5 to 97.5 percentile interval as defined in the paper. For posteriors with a limited degree of asymmetry, the HPD region is approximately given by the percentile-based definition.

[2]Backus and Kehoe (1992) are concerned with the volatility of macroeconomic time series in different historical periods, but in fixing the HP smoothing constant (the inverse of the signal-noise ratio), they effectively make an assumption about the very thing they are trying to determine! Ravn and Uhlig (2002) show that the result is overturned if a more plausible value of the HP smoothing constant is used. A model-based approach, focusing on the relative variance of cycle and trend components, gives a coherent answer to the question, though in doing so it raises the whole issue of whether one should be doing HP detrending in the first place.

There is some difficulty in estimating plausible cycles in the annual series for higher order models; with noninformative priors the cyclical variance falls to near zero so that the trend accounts for most of the variation in the series. Therefore, for $n > 1$ more informative priors on σ_ζ^2 were used to ensure that the trend does not change too rapidly. Posterior means for different values of n are shown in table 1 for the sharp prior; with less informative priors on λ_c the posterior mass of the period shifts to higher values while the results for the other parameters remain similar. Considerably more noise is removed for $n > 1$, as can be seen from the higher estimates for σ_ε^2.

Order n	σ_ζ^2	σ_κ^2	σ_ε^2	ρ	λ_c	$2\pi/\lambda_c$
1	60.7	5416	3413	0.918	0.507	12.4
2	18.8	417	7847	0.897	0.525	11.9
3	20.7	64.4	8296	0.872	0.549	11.5
4	12.9	3.11	10,469	0.883	0.542	11.6

Table 1 : Posterior means for annual US real GDP from 1870 to 1998 for different values of n with an informative prior on λ_c, centred at $2\pi/10$. For $n = 2, 3$, the shape and scale of the inverted gamma prior on σ_ζ^2 were set to 20 and 2 x 10^{-5}, respectively. For $n = 4$, the shape and scale of the inverted gamma prior on σ_ζ^2 were set to 100 and 10^{-4}, and additionally, a moderately informative prior on σ_ε^2 was used, with shape and scale equal to 10 and 10^{-5}. The period in years is $2\pi/\lambda_c$. All variance parameters are multiplied by 10^7.

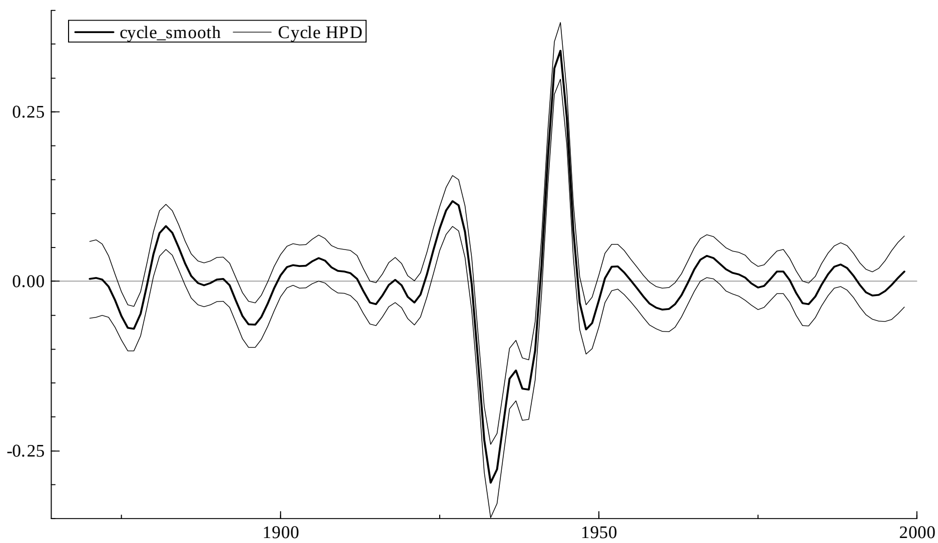

Figure 5: Estimated cycle in annual US real GDP (logarithms) from 1870 to 1998 for $n = 2$ with most informative prior on λ_c.

The biggest change occurs in moving from the first order to the second order cycle. The estimated cyclical component for $n = 2$ is displayed in figure 5. The increased smoothness relative to the first order case means that the turning points become more clearly defined.

4 Conclusion

A structural time series model provides a consistent framework for extracting trends and cycles. This article has investigated the Bayesian treatment of such a model, paying particular attention to the cyclical component and the way in which prior information on periodicity can be used. Markov chain Monte Carlo routines are successfully designed for univariate and multivariate models, including those with higher order cycles of the kind introduced recently by Harvey and Trimbur (2003). Smooth cycles were successfully extracted from the annual US GDP time series. These cycles have a simple interpretation in terms of the percentage by which they exceed or fall below the long-term level. The series of smoothed Bayesian estimates give more information about the relative position of the business cycle over time, along with measures of uncertainty.

Two key incentives for exploring a Bayesian approach are, on the practical side, the possibility of addressing a broader range of problems including those with limited sample sizes, and from a theoretical perspective, the more consistent and flexible use of prior knowledge on periodicity. Bayesian analysis in an unobserved components framework provides a direct and informative approach to studying cyclical dynamics. Cyclical turning points and features such as changing amplitude are assessed while accounting for parameter uncertainty. Multivariate applications provide the capacity to combine the information in sets of economic indicators and to study underlying business cycle components and relationships across different series.

ACKNOWLEDGMENTS

Earlier versions of this paper was presented at the Cambridge PhD Workshop in Econometrics and at the NAKE conference at the Netherlands Central Bank. Trimbur wishes to thank the Cambridge Commonwealth Trust and the Richard Kahn Fund for financial support, and is grateful to the Tinbergen Institute for its hospitality and financial support in autumn 2001. Harvey thanks the Economic and Social Research Council (ESRC) for support as part of a project on Dynamic Common Factor Models for Regional Time Series, grant number L138 25 1008. We would like to thank Bill Bell and Simon Godsill for helpful comments.

REFERENCES

BACKUS, D. K., AND P.J. Kehoe (1992): "International evidence on the historical properties of business cycles," *American Economic Review* 82, 864-88.

BAXTER, M., AND KING, R. G. (1999): "Measuring business cycles: approximate band-pass filters for economic time series," *Review of Economics and Statistics* 81, 575-93.

CARTER, C. K., AND R. KOHN (1994): "On Gibbs sampling for state space models," *Biometrika* 81, 541-53.

COGLEY, T., AND J. M. NASON (1995): "Effects of the Hodrick-Prescott filter on trend and difference stationary time series: implications for business cycle research," *Journal of Economic Dynamics and Control* 19, 253-78.

DEJONG, P., AND N. SHEPHARD (1995): "The simulation smoother for time series models," *Biometrika* 82, 339-50.

DOORNIK, J. A. (1999): *Ox: An Object-Oriented Matrix Programming Language.* Timberlake Consultants Ltd., London.

DURBIN, J., AND Koopman, S.J., (2002): "A simple and efficient simulation smoother," Biometrika 89, 603-16.

FRUHWIRTH-SCHNATTER, S. (1994): "Data augmentation and dynamic linear models," *Journal of Time Series Analysis* 15, 183-202.

GOMEZ, V. (2001): "The use of Butterworth filters for trend and cycle estimation in economic time series," *Journal of Business and Economic Statistics* 19, 365-73.

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

HARVEY, A. C. AND A. JAEGER (1993): "Detrending, stylised facts and the business cycle," *Journal of Applied Econometrics* 8, 231-47.

HARVEY, A. C., AND S. J. KOOPMAN (1997): "Multivariate structural time series models," in *System dynamics in economic and financial models,* ed. by C. Heij et al. Wiley and Sons, Chichester.

HARVEY, A. C., AND S. J. KOOPMAN (2000): "Signal extraction and the formulation of unobserved component models," *Econometrics Journal* 3, 84-107.

HARVEY, A. C., AND T. M. TRIMBUR (2003): "General model-based filters for extracting trends and cycles in economic time series," *Review of Economics and Statistics* 85, 244-55.

HARVEY, A. C., TRIMBUR, T. M., AND H. K. VAN DIJK (2002): "Cyclical components in economic time series: a Bayesian approach," *Econometric Institute Research Report,* Erasmus University Rotterdam.

HODRICK, R. J., AND E.C. PRESCOTT (1997): "Postwar US business cycles: an empirical investigation," *Journal of Money, Credit and Banking* 24, 1-16.

HUERTA, G., AND M. WEST (1999): "Priors and component structures in autoregressive time series models," *Journal of the Royal Statistical Society*, Series B, 61, 881-99.

KASS, R. E., AND A.E. RAFTERY (1995): "Bayes factors," *Journal of the American Statistical Association* 24, 773-95.

KITAGAWA, G., AND W. GERSCH (1996) *Smoothness Priors Analysis of Time Series.* Berlin: Springer-Verlag.

KOHN, R., ANSLEY, C., AND C.H. WONG (1992): "Nonparametric spline regression with autoregressive moving average errors," *Biometrika* 79, 335-46.

KOOP, G. AND H.K. VAN DIJK (2000): "Testing for integration using evolving trend and seasonals models: a Bayesian approach," *Journal of Econometrics* 97, 261-91.

KOOPMAN, S. J., AND A. C. HARVEY (2003): "Computing observation weights for signal extraction and filtering," *Journal of Economic Dynamics and Control* 27, 1317-33.

KOOPMAN, S. J., HARVEY, A. C., DOORNIK, J. A., AND N. SHEPHARD (2000): *STAMP 6.0 Structural Time Series Analysis Modeller and Predictor.* Timberlake Consultants Ltd., London.

KOOPMAN, S. J., SHEPHARD, N., AND J. A. DOORNIK (1999): "Statistical algorithms for models in state space using SsfPack 2.2," *Econometrics Journal* 2**,** 113-66**.**

MURRAY, C. J. (2003): "Cyclical properties of Baxter-King filtered time series," *Review of Economics and Statistics* 85, 471-6.

TRIMBUR, T. M. (2004): "Properites of higher order stochastic cycles," *Journal of Time Series Analysis,* forthcoming.

YOUNG, P. (1984): *Recursive Estimation and Time Series Analysis.* Springer-Verlag, Berlin.