TI 2010-017/4
Tinbergen Institute Discussion Paper

Models with Time-varying Mean and Variance: A Robust Analysis of U.S. Industrial Production

Charles S. Bos
Siem Jan Koopman

VU University Amsterdam, and Tinbergen Institute.

Tinbergen Institute

The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam
Roetersstraat 31
1018 WB Amsterdam
The Netherlands
Tel.: +31(0)20 551 3500
Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Most TI discussion papers can be downloaded at http://www.tinbergen.nl.

Models with Time-varying Mean and Variance: A Robust Analysis of U.S. Industrial Production

Charles S. Bos and *Siem Jan Koopman*

Department of Econometrics, VU University Amsterdam,
&
Tinbergen Institute, Amsterdam

January 21, 2010

Abstract

Many seasonal macroeconomic time series are subject to changes in their means and variances over a long time horizon. In this paper we propose a general treatment for the modelling of time-varying features in economic time series. We show that time series models with mean and variance functions depending on dynamic stochastic processes can be sufficiently robust against changes in their dynamic properties. We further show that the implementation of the treatment is relatively straightforward. An illustration is given for monthly U.S. Industrial Production. The empirical results including estimates of time-varying means and variances are discussed in detail.

Some Keywords: Common stochastic variance; Kalman filter; State space model; Unobserved Components Time Series Models.

JEL classification: C22, C51, C53, E23

1 Introduction

The analysis of macroeconomic time series requires a sufficiently large number of time series observations for identifying stable dynamic relationships between the economic variables. Seasonal economic time series become increasingly available at much longer time horizons. However, time series with long time horizons are also more likely to be subject to slow shifts and sudden breaks which can be due to changes in economic behaviour, in data collection, in variable definitions and, most importantly, changes in economic and fiscal policy decision making. Such changes are not always easy to capture in a model representation of the time series. As a result, a sequence of noisy observations can become part of the dataset and groups of observations may be less reliable since they have been subject to particular events that cannot be treated explicitly in the model. Therefore, the model specification needs to be modified to allow for data irregularities and breaks.

In this paper we present and discuss an approach that is able to account for changes in the mean and variance of time series. In this way, the observations will not be weighted equally in the analysis. A strong feature of our approach is that we do not require to formulate the changes in the time series explicitly. Furthermore, these changes in the time series will be accounted for as an integrated part of the analysis. This robust approach to time series modelling is illustrated by a time series analysis of monthly growth in U.S. industrial production (IP) between 1960 and 2009. This series is of interest as it is a key indicator of economic activity in the U.S.

The monthly IP time series is available at a long time span and has been subject to different economic settings. In the 1970s large variations in the series are related to several crises mainly caused by limited oil supply. In the 1990s smaller variations over time are observed due to tighter monetary policies and related control mechanisms, an effect which Stock and Watson (2002) terms the Great Moderation. The credit crunch in 2007-2008 and the resulting economic decline in 2008-2009 have led to various large negative shocks in the monthly growth of IP. These events illustrate the challenge of formulating an overall model for forecasting IP. The forecasting studies in Bruno and Lupi (2004) and Kawasaki and Franses (2004) focus on the suitability of specific linear and seasonal specifications. Heravi et al. (2004) compare the forecasting performance of linear models against the performance of a neural network. The study reported in Franses and Van Dijk (2005) considers non-linear models that account for seasonal unit roots, for smooth transitions in seasonal effects, and for deterministically or stochastically time-varying autoregressive coefficients.

The studies by Kim and Nelson (1999), Stock and Watson (1999) and Hendry and Clements (2000), and more recently Ewing and Thompson (2008), argue for introducing changes in the overall variance of the model. In this paper we confirm the importance of allowing for changes in the overall variance and propose a general method for its implementation. It is based on stochastically time-varying functions for the common variance of the model. We introduce a stochastic scaling mechanism for the time series of interest. As an illustration for modelling monthly IP growth, we consider a basic unobserved components time series model that includes a stationary growth process, a persistent seasonal component and an irregular. The model provides an effective base for the description of the salient features in IP growth. It is shown that an additional feature of the time series is characterised by a time-varying variance which is modelled by an unobserved stochastic component as well. The estimation method of Koopman and Bos (2004) can be adopted for state space models with a common stochastic time-varying variance. As part of this analysis, we estimate the model by maximum likelihood using an adapted version of the simulation techniques for the evaluation of the likelihood function by importance sampling methods, see Durbin and Koopman (2001).

The inclusion of time-varying variances in an unobserved components model has been considered earlier by e.g. Bos et al. (2000), Aguilar and West (2000), Cecchetti et al. (2007) and Stock and Watson (2007). In this paper we provide more empirical evidence of the importance of time-varying variances and we provide a framework in which several time-varying specifications for mean and variance functions can be treated simultaneously in a unified way. It also implies that formal tests for model specifications can be carried out as part of the presented methodology.

The paper is organised as follows. Section 2 introduces the industrial production time series data set and discusses various characteristics of the series. The basic modelling framework is presented in Section 3. The robustness of the model is obtained by the introduction of a common time-varying variance component in Section 3.2. The associated estimation methods are detailed in Section 3.3. The results of our empirical study for the U.S. industrial production time series are presented in Section 4, including a description of the extracted components in Section 4.2. In this section it is also seen how the observation weights for different periods vary due to the changing volatility in the model. Whether we reach stability of the recursive parameter estimates is discussed in Section 4.4. Conclusions are given in Section 5.

2 Industrial production data

The monthly time series of U.S. industrial production is obtained from the website of the Board of Governors of the Federal Reserve System[1] for the period between 1960:1 and 2009:9 (596 observations). Figure 1 presents the actual data (in panel (i)), in logs (panel (ii)), and the monthly differences of the logged data times 100 (panel (iii)). The latter series is interpreted as the monthly IP growth and will be the input for the analysis below. The time series plot of growth reveals that the variability of the series, and especially of the seasonality, is lower after the early 1980s.

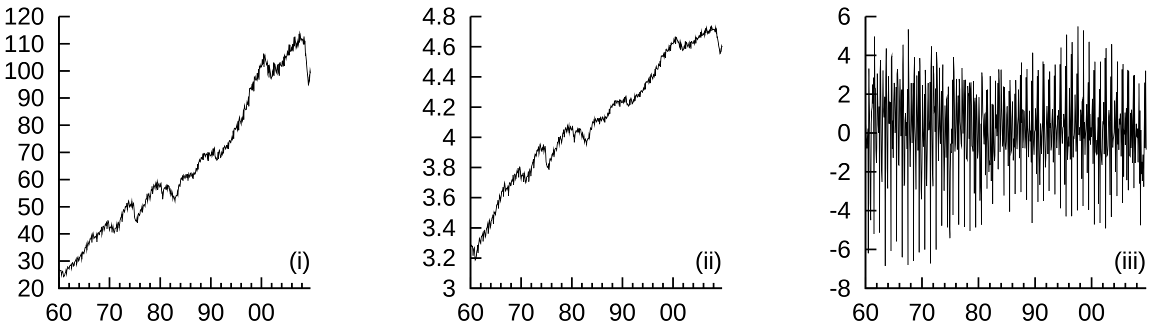

Figure 1: The U.S. monthly industrial production 1960:1 and 2009:9: (i) in levels, (ii) in logs and (iii) in percentage growth.

Figure 2 plots the correlation of the returns (in panel (i)) and of the squared returns (panel (ii)). Both panels indicate a strong yearly seasonal pattern in the data. Furthermore, even without accounting for the seasonality, there is a strong indication of non-seasonal correlation in the squared returns. This suggest that we need to account for a time-varying variance in the model. Further evidence of a time-varying variance in the time series is provided by the third panel of Figure 2, where the yearly average squared return is depicted. The observations in the years before 1980 exhibit up to twice the variability of the observations after 1980 in this dataset. An adequate model for this series should take this drop in the variance into account.

[1]Source: Federal Reserve, Board of Governors d.d. November 2009, Industrial Production – Market Group – Total Index – Not Seasonally Adjusted. Series `b500001_ipnsa`.

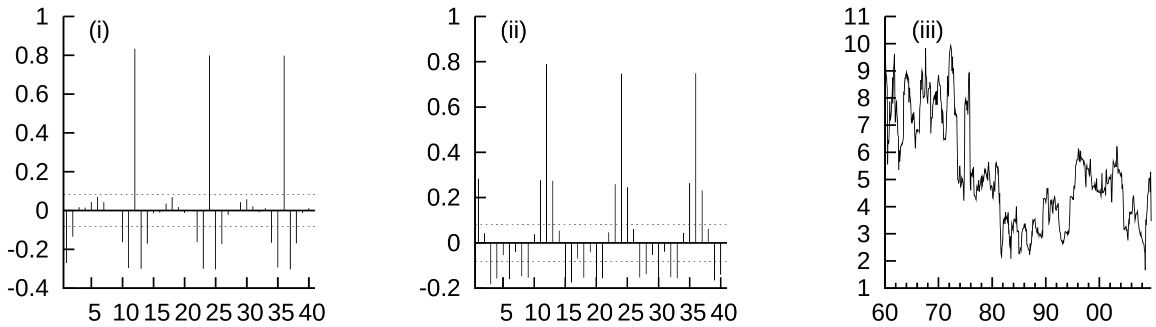

Figure 2: Autocorrelation of U.S. monthly industrial production between 1960:1 and 2009:9: (i) for the returns, and (ii) for the squared returns, with (iii) the yearly averaged squared return.

3 Decomposition model with stochastic volatility

3.1 The mean equation

To analyse the monthly time series of IP growth, we consider the unobserved components time series model consisting of a constant growth term c, a stochastic seasonal component γ_t, a business cycle component ψ_t and an irregular component ε_t. The model is given by

$$y_t = c + \gamma_t + \psi_t + \varepsilon_t, \qquad \varepsilon_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2_{\text{Irr}}), \qquad t = 1, \ldots, n, \tag{1}$$

where y_t is the observed IP growth at time t. The notation in (1) implies that the irregular is independently and normally distributed with mean zero and variance σ^2_{Irr}.

The seasonal component γ_t is modelled by the trigonometric seasonal process

$$\gamma_t = \sum_{j=1}^{\lfloor s/2 \rfloor} \gamma^+_{j,t}, \qquad \text{with} \begin{pmatrix} \gamma^+_{j,t+1} \\ \gamma^*_{j,t+1} \end{pmatrix} = \begin{pmatrix} \cos\lambda_j & \sin\lambda_j \\ -\sin\lambda_j & \cos\lambda_j \end{pmatrix} \begin{pmatrix} \gamma^+_{j,t} \\ \gamma^*_{j,t} \end{pmatrix} + \begin{pmatrix} \omega^+_{j,t} \\ \omega^*_{j,t} \end{pmatrix}, \tag{2}$$

where s is the seasonal length (in the case of monthly data, $s = 12$) and $\lfloor x \rfloor$ is the function which truncates x to the nearest integer value $\leq x$. The seasonal frequencies are given by $\lambda_j = 2\pi j/s$ and the seasonal disturbances are modelled by

$$\begin{pmatrix} \omega^+_{j,t} \\ \omega^*_{j,t} \end{pmatrix} \overset{\text{i.i.d.}}{\sim} \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \sigma^2_{\text{Seas}} I_2 \right), \qquad \text{for } j = 1, \ldots, \lfloor s/2 \rfloor,$$

where I_2 is the 2×2 identity matrix.

Economic time series are also often subject to business cycle features and to capture these we also include a stochastic cycle component in the mean equation. The stationary stochastic cycle process is given by

$$\begin{pmatrix} \psi^+_{t+1} \\ \psi^*_{t+1} \end{pmatrix} = \varphi_{\text{Cyc}} \begin{pmatrix} \cos\lambda_{\text{Cyc}} & \sin\lambda_{\text{Cyc}} \\ -\sin\lambda_{\text{Cyc}} & \cos\lambda_{\text{Cyc}} \end{pmatrix} \begin{pmatrix} \psi^+_t \\ \psi^*_t \end{pmatrix} + \begin{pmatrix} \kappa^+_t \\ \kappa^*_t \end{pmatrix}, \qquad \begin{pmatrix} \kappa^+_t \\ \kappa^*_t \end{pmatrix} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2_{\text{Cyc}} I_2). \tag{3}$$

with cycle frequency $\lambda_{\text{Cyc}} = 2\pi/p_{\text{Cyc}}$. The cycle period p_{Cyc} typically ranges from 3 to 8 years (36 to 96 months) and the discounting factor $0 < \varphi_{\text{Cyc}} < 1$ ensures the stationary behaviour of

ψ_t. For this component the cycle variance σ^2_{Cyc}, the discounting factor φ_{Cyc} and the cycle period p_{Cyc} will be estimated. Harvey (1989, Sec. 2.5.3) shows that the cycle process (3) can be reduced to a stationary autoregressive moving average process. If $\lambda_{\text{Cyc}} \downarrow 0$ or $p_{\text{Cyc}} \to \infty$, the component simplifies to a simple autoregressive process of order one.

The disturbance series $\varepsilon_t, \kappa_t^+, \kappa_t^*$ and $\omega_{j,t}^+, \omega_{j,t}^* (j = 1, \ldots, \lfloor s/2 \rfloor)$ are mutually independent. The model (1)–(3) is a special case of an unobserved components time series model. A detailed discussion of this class of models is given in Harvey (1989).

We provide empirical evidence in Section 4 that the unobserved components model (1) allows for most of the variation and serial correlation in the time series of IP growth. The seasonal dynamics in the time series are estimated simultaneously with the stationary dynamics represented by the cycle and irregular components. The mean equation of the model describes the typical features in seasonal macroeconomic time series.

3.2 The variance equation

The three variances in the model (1)–(3) can be specified as

$$\sigma^2_{\text{Irr}} = q_{\text{Irr}}\sigma^2, \qquad \sigma^2_{\text{Seas}} = q_{\text{Seas}}\sigma^2, \qquad \sigma^2_{\text{Cyc}} = q_{\text{Cyc}}\sigma^2, \tag{4}$$

where the common variance σ^2 is strictly positive while the scaled variances q_{Irr}, q_{Seas} and q_{Cyc} are strictly non-negative. The common variance can be interpreted as the scaling coefficient of the time series y_t.

In the empirical analysis we will treat the common variance σ^2 either as a fixed unknown coefficient, or we consider it as a (unknown) time-varying parameter. In the latter case, σ^2 is replaced by σ_t^2 which is a (possibly stochastic) function of time. As a consequence, the variances σ^2_{Irr}, σ^2_{Seas} and σ^2_{Cyc} also become time-varying as they are set equal to σ_t^2 times their respective q values in (4). We consider the following time-varying functions for the common variance:

Case 1: Break in variance A single break of the common variance can be specified as the step-function

$$\sigma_t^2 = \begin{cases} \sigma_I^2 & \text{if } t \leq \lfloor \tau T \rfloor, \\ \sigma_{II}^2 & \text{if } t > \lfloor \tau T \rfloor, \end{cases} \qquad 0 < \tau < 1, \tag{5}$$

where both σ_I^2 and σ_{II}^2 are strictly positive and τ is a coefficient governing the timing of the break.

Case 2: GARCH A well-known time-varying variance specification is proposed by Bollerslev (1986) and is usually referred to as the generalised autoregressive conditional heteroskedasticity (GARCH) model. In combination with the unobserved componets model (1), we use the specification

$$\sigma_{t+1}^2 = \omega + \delta\sigma_t^2 + \alpha u_t^2, \qquad t = 1, \ldots, n, \tag{6}$$

where u_t is the mean/variance-corrected observation and ω, α and δ are fixed coefficients with $0 < \omega, 0 < \alpha < 1, 0 \leq \delta < 1$ and $\alpha + \delta < 1$. The formal definition of u_t is given below. The

GARCH recursion is started with $\sigma_1^2 \equiv \omega/(1-\alpha-\delta)$, which is effectively the unconditional expectation of σ_t^2.

The mean/variance-corrected observation u_t in (6) determines the variance σ_t^2. In case of the stochastic mean equation in our unobserved components model, it may not be clear immediately how to correct y_t. Since the GARCH model determines the variance σ_t^2 only from past u_t's, we will also determine u_t using past observations only. More specifically, we define u_t as the rescaled standardized one-step ahead prediction error which is given by

$$u_t \equiv \frac{\nu_t}{\sqrt{F_t}}\sigma_t, \qquad \nu_t = y_t - \mathrm{E}(y_t|y_1,\ldots,y_{t-1}) \qquad F_t = \mathrm{var}(y_t|y_1,\ldots,y_{t-1}).$$

The prediction error ν_t and variance F_t are evaluated by the Kalman filter, see Durbin and Koopman (2001). Detailed discussions of incorporating a GARCH specification into an unobserved components time series model are given by Harvey et al. (1992) and Broto and Ruiz (2006).

Case 3: Stochastic Variance The time-varying common variance can alternatively be specified as a stochastic volatility model, see Harvey et al. (1994) and Jacquier et al. (1994). The common log-variance evolves over time as a stationary autoregressive process, that is

$$\sigma_t^2 = \sigma_h^2 \exp(h_t), \qquad h_{t+1} = \phi h_t + \xi_t, \qquad \xi_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_{\mathrm{SV}}^2), \tag{7}$$

where $h_1 \equiv 0$ and with $0 < \phi < 1$. The disturbance series ξ_t is independent of all disturbances in the mean equation. The parameter σ_h^2 is used to govern the overall level of the variance process.

When the model contains a common time-varying variance σ_t^2, a restriction is needed for identification. In all cases we set the initial common variance equal to 1. This implies either $\sigma^2 \equiv 1$ (for the constant variance case), $\sigma_I^2 \equiv 1$ (for the variance break), $\omega \equiv 1-\delta-\alpha$ and $\sigma_1^2 \equiv 1$ (for GARCH) or $\sigma_h^2 \equiv 1$ (for SV).

3.3 Estimation procedure for the models

Both the mean and variance equations consist of parameters that need to be estimated. In case $\sigma_t^2 = \sigma^2 = 1$ is fixed, we only have parameters in the mean equation including $q_{\mathrm{Irr}}, q_{\mathrm{Seas}}$, q_{Cyc}, λ_c and φ_{Cyc}. We estimate these parameters by the method of maximum likelihood in which the common variance is concentrated out of the likelihood function together with the unknown constant c. The prediction error decomposition allows the loglikelihood function to be evaluated by means of the Kalman filter, see Harvey (1989). We have implemented this approach of estimation using Ox (Doornik, 2009) and SsfPack (Koopman et al., 2008). Once the parameters are estimated, the unobserved components can be estimated using the Kalman filter and associated smoothing algorithms.

In case the common variance is subject to a break (case 1), the same methods can be used for estimation and signal extraction since the Kalman filter is sufficiently flexible to deal with different variances for different time periods. The GARCH specification for the common variance (case 2) is also a deterministic function of mean-corrected observations that are evaluated by the Kalman filter. The Kalman filter can still be used for likelihood evaluation but the filter cannot be regarded as a linear filter since parameters in the model depend directly on (squared) observations. Therefore the well-known optimal properties of the Kalman filter are lost. Also, in this case the Kalman filter evaluates a quasi-likelihood function. A more detailed discussion is provided by Harvey et al. (1992).

In case the stochastic volatility specification is adopted for the common variance (case 3), the model is clearly not a linear Gaussian model, as the likelihood function must account for the uncertainty in the common variance. A closed-form expression for the likelihood function is not available and therefore we need to rely on numerical methods. Methods based on Monte Carlo integration have been regarded as the most feasible approach to likelihood evaluation for non-linear unobserved component time series models. We adopt the methodology that is described in detail by Koopman and Bos (2004) and is based on importance sampling. In this approach we repeatedly apply the Kalman filter to evaluate the likelihood function conditional on different random draws for the common variance sequence $\sigma_t^2, t = 1, \ldots, n$. The draws are obtained from an approximating linear Gaussian model for $\log \sigma_t^2$. The likelihoods from the Kalman filter are weighted by normalised importance sampling corrections. This Monte Carlo estimate of likelihood function is a consistent estimator of the true likelihood function of the model but is clearly subject to Monte Carlo error. However, the method can be used for maximum likelihood estimation since the likelihood is guaranteed to be smooth when a fixed random seed is used for each Monte Carlo likelihood evaluation.

4 Empirical results

4.1 Estimation results for main class of models

The first aim of our empirical study is to investigate how well the different model specifications are able to estimate the monthly growth of U.S. industrial production (IP) introduced in Section 2. The seasonal features of IP growth are clearly presented in the third panel of Figure 1. We therefore consider first a model that contains only seasonal and irregular components in the state space framework (1), that is our model without the cycle component ψ_t of (3) and with a fixed common variance $\sigma_t^2 = 1$.

Table 1, in the first set of columns labelled *Seas*, reports estimation results for this model. The results indicate that the standard deviation of the seasonal component is small but it is clearly different from zero. Hence the seasonal component is changing slowly over time. Test statistics include the normality test of Doornik and Hansen (1994, based on the 3rd and 4th moment of the residuals, asymptotically following a χ_2^2 density under the null of normality), the Box-Ljung Q-statistic for residual correlation (Ljung and Box, 1978, using $a = \lfloor \sqrt{n} \rfloor = 24$ lags, with asymptotically a χ_{a+1-p}^2 density, with p the number of parameters in the model), Engle's (1982) test for ARCH-type behaviour in the squared residuals (leading to a χ_1^2 density, using one lag of squared returns) and a general heteroskedasticity test H($n/3$) comparing variability in the first and last $n/3$ observations (resulting in an $F_{n/3,n/3}$ density). All four statistics are reported with their p-values in Table 1. For this first model, the diagnostics indicate a strong rejection of the null hypotheses of normality, of no-correlation, and constant variance.

A first step to improve our initial results is to include the cyclical component (3) in the model. The results are presented in the second set of columns with label *Cyc-Seas*. The cycle component leads to a lower residual variance and it also leads to a strong increase in the likelihood. The discounting factor $\varphi_{\text{Cyc}} \approx .90$ provides the expected persistence in U.S. growth of industrial production, with a cycle period p_{Cyc} of close to 50 months (over 4 years). The test statistics indicate that residual correlation disappears, but some evidence of non-normality and heteroskedasticity remains.

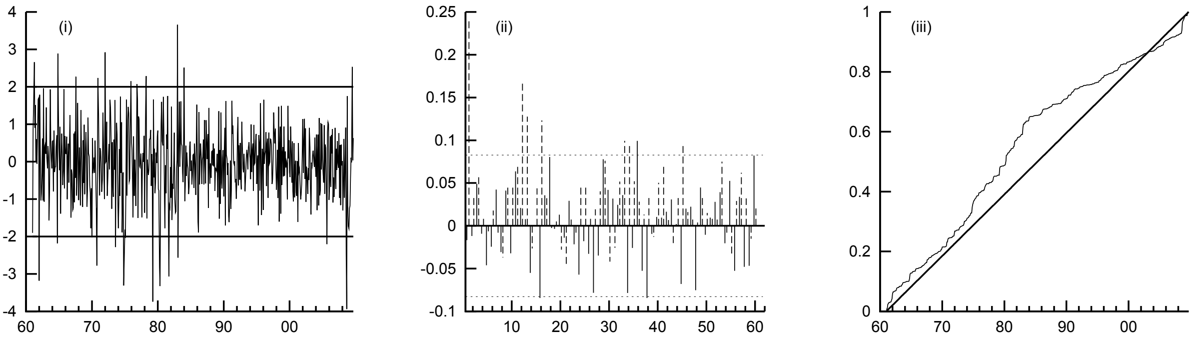

Figure 3: Prediction residuals of the estimated model *Cyc-Seas* with cyclical, seasonal and irregular components: (i) plot of standardised residuals; (ii) autocorrelogram of residuals (solid lines) and squared residuals; (iii) scaled cumulative sum of squared residuals.

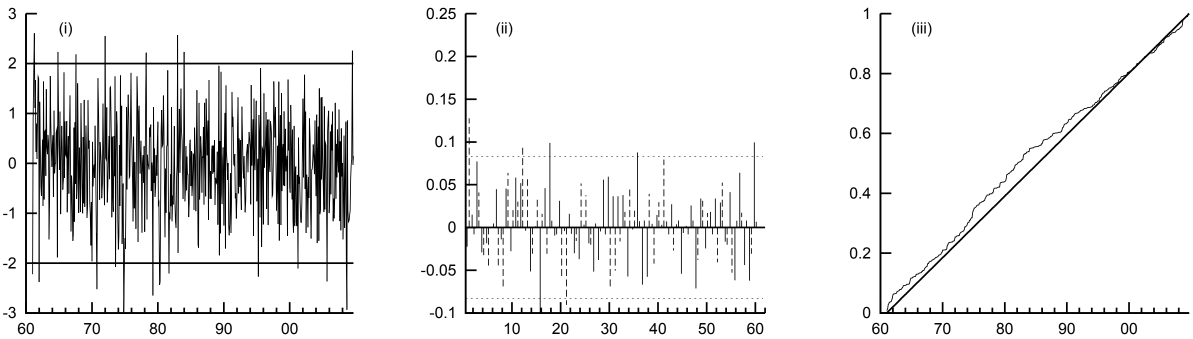

Figure 4: Residuals of the *Cyc-Seas-SV* model with cyclical, seasonal, irregular and common SV components: (i) plot of standardised residuals; (ii) autocorrelogram of residuals (solid lines) and squared residuals; (iii) scaled cumulative sum of squared residuals.

Figure 3 presents the standardised residuals of the model (panel (i)), with the autocorrelations of the residuals and of the squared residuals (panel (ii)) and the cumulative sum of the squared residuals. The residuals themselves show several outliers in the first half of the sample, with possibly too many 'inliers' in the second part. The squared residuals depict significant first order autocorrelation, and also the plot of the cumulative squared residuals indicates that around 1974

Table 1: Estimation results for U.S. industrial growth $\Delta \log(\mathrm{IP})$, 1960-2008, 585 observations, basic models.

Parameter	Seas		Cyc-Seas		Cyc-Seas-SV		Cyc-Seas-GARCH		Cyc-Seas-$\Delta\sigma$	
σ_{Irr}	0.888	[0.83, 0.95]	0.729	[0.66, 0.80]	0.729	[0.62, 0.85]	1.318	[0.85, 2.05]	0.843	[0.76, 0.94]
σ_{Seas}	0.029	[0.02, 0.04]	0.033	[0.03, 0.04]	0.033	[0.03, 0.04]	0.055	[0.03, 0.09]	0.041	[0.03, 0.05]
σ_{Cyc}			0.453	[0.35, 0.59]	0.358	[0.26, 0.49]	0.644	[0.40, 1.03]	0.509	[0.39, 0.66]
φ_{Cyc}			0.899	[0.59, 0.96]	0.929	[0.75, 0.97]	0.932	[0.76, 0.97]	0.910	[0.69, 0.96]
p_{Cyc}			49.399	[26.84, 92.44]	55.446	[33.24, 93.43]	55.815	[32.65, 96.48]	51.752	[28.47, 95.49]
σ_{SV}					0.250	[0.15, 0.41]				
ϕ_{SV}					0.912	[0.78, 0.97]				
δ_{GARCH}							0.941	[0.89, 0.97]		
α_{GARCH}							0.056	[0.03, 0.11]		
$\Delta \log \sigma^2_{1984}$									−0.692	[-0.92, -0.46]
LL	−855.81		−826.39		−795.83		−811.61		−809.59	
Normality-DH(2)	28.62	(0.00)	26.40	(0.00)	0.44	(0.80)	18.94	(0.00)	18.94	(0.00)
Box-Ljung Q($\sqrt{n} - p$)	116.88	(0.00)	23.14	(0.28)	28.25	(0.06)	27.59	(0.07)	25.46	(0.15)
ARCH(1)	82.71	(0.00)	35.66	(0.00)	9.70	(0.00)	22.17	(0.00)	39.78	(0.00)
H($n/3, n/3$)	0.60	(0.00)	0.58	(0.00)	0.80	(0.12)	1.03	(0.86)	1.13	(0.39)

Parameter estimates with 95% confidence intervals of five models for U.S. monthly industrial growth 1960-2009. The models consist of a combination of seasonal plus irregular components (*Seas*), cyclical (*Cyc-Seas*) component and common stochastic variance (*Cyc-Seas-SV*), GARCH (*Cyc-Seas-GARCH*) or variance-with-change (*Cyc-Seas-*$\Delta\sigma$) component respectively. The loglikelihood (LL) value and various diagnostic test statistics (with p-values) are also reported.

there was a first clear increase in variability in the series, with a drop in variability occurring around 1984.

Instead of adapting manually for atypical observations (by including structural changes at pre-specified time periods), we modify the specification of the model further. The most general option proposed in the previous section was to introduce a common stochastic variance component (7), allowing the variance to change over time. In the columns with label *Cyc-Seas-SV* the relevant results are presented. The first indication that the model is preferred for this data set over the *Cyc-Seas* specification is the likelihood score that is found: An increase of 30.6 points is a clear indication of a better fit. The cycle becomes relatively more pronounced, with a smaller estimated variance and a larger estimated discounting factor. These estimates indicate a smoother and more persistent cycle component in the time series. The estimated autoregressive coefficient of the log-variance component h_t is estimated as $\phi_{\mathrm{SV}} = 0.91$, indicating a persistence in the time-varying variance process. However the estimated standard deviation $\sigma_{\mathrm{SV}} = 0.25$ allows for some amount of flexibility in the variance process such that it can move upwards and downwards over time.

The test statistics indicate that the normality, the (lack of) autocorrelation, and the equality of variance between the first and last third of the sample of residuals[2] are acceptable, but the test for ARCH effects is still rejected. Figure 4 displays similar panels of estimated components for the present model as Figure 3 did for the components of the model *Cyc-Seas*. The first and last panel indicate that volatility of the standardized residuals seems constant, with possibly a large residual around 1975.

The middle panel of Figure 4 displays the autocorrelation of residuals and squared residuals. For the squared residuals, at both the first and twelfth lag the ACF still fall slightly outside the confidence bounds. However, in comparison with the corresponding panel in Figure 3, the autocorrelations have become much smaller.

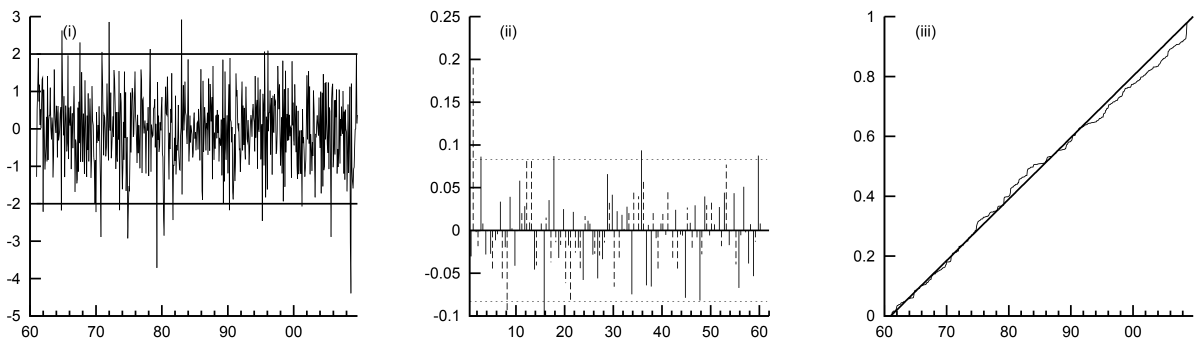

Figure 5: Residuals of the *Cyc-Seas-GARCH* model with cyclical, seasonal, irregular and GARCH components: (i) plot of standardised residuals; (ii) autocorrelogram of residuals (solid lines) and squared residuals; (iii) scaled cumulative sum of squared residuals.

Alternatively, we can allow the variance to change over time via a GARCH specification. The

[2]The residuals are extracted by iteratively running the importance smoother for the likelihood evaluation and construction of the unobserved state, keeping only the one step ahead out-of-sample predictions. Alternative results, using a particle filter (Pitt and Shephard, 1999; Christophe and Doucet, 2002), are similar.

results for this option are presented in the fourth set of columns in Table 1 (labelled *Cyc-Seas-GARCH*). The final likelihood is in between the likelihoods of the model with SV and the *Cyc-Seas* model. The estimates for the GARCH coefficients in the model indicate that $\delta + \alpha \approx 1$, implying that the so-called integrated GARCH model might be preferred for this data set. These estimation results indicate a strong autocorrelation in the variance series and, therefore, a very slow mean-reverting variance series.

The diagnostic test statistics indicate that the *Cyc-Seas-GARCH* model passes the test of equal variance of the standardised residuals in the first and last third of the sample. The test for serial correlation in the residuals is passed with a p-value of 7%, but both the ARCH and the normality tests are rejected. It may be due to the fact that the GARCH model adapts slowly to the altered variability of the series. This can also be seen from Figure 5. The residual plot in the first panel shows that some of the standardised residuals are too large for assuming normality. Possibly the time-varying variance needs to adapt faster to account for the increasing variance at the end of the sample. This can result in a remaining first-order autocorrelation of the squared residuals, and in turn leads to the rejection of the ARCH test-statistic. However, the third panel of the figure indicates that the GARCH model is adequate for describing the overall level of time-varying variance. For example, the scaled cumulative squared residuals do not deviate much from the 45 degrees line.

In Table 1 we also report the estimation results for a model in which a simple one-time break in the variance component is specified. The date of the break has to be known in advance, whereas the other earlier models have variance series that adapt automatically. In this case, the break in variance is set to January 1984. The estimated variance break is highly significant and it even leads to a slightly higher likelihood compared to the value obtained from the model with GARCH.

4.2 Signal extraction of mean and variance components

The main differences between our models are the manner in which the variance adapts over time. In case of GARCH, the time-varying variance is specified as a deterministic function of past observations. The model with a common SV component has a stochastically time-varying variance that is unobserved. Once the model parameters are estimated, we can estimate the variance using all observations (smoothing) or using past observations (filtering). The filtered SV estimates can be compared with the variance estimates from the GARCH specification since these estimates also depend only on past observations.

Figure 6 displays the standard deviations of the irregular component, $\sigma_{\text{Irr}}\sigma_t$ of the *Cyc-Seas*, filtered *Cyc-Seas-SV*, *Cyc-Seas-GARCH* and *Cyc-Seas-*$\Delta\sigma$ models. For the SV model, both a filtered and a smoothed estimate are presented in the bottom panel. We see that both SV and GARCH models adapt to the changing situations of the mid-seventies and the lower variability starting after 1985. However, the GARCH estimates adapt more slowly than the (filtered) SV estimates. Taking future observations into account, by means of the smoothed SV estimate in the bottom panel, leads to an even clearer and more swiftly reacting signal for the time-varying variance. This is coherent with the finding in the literature that the GARCH volatility estimates

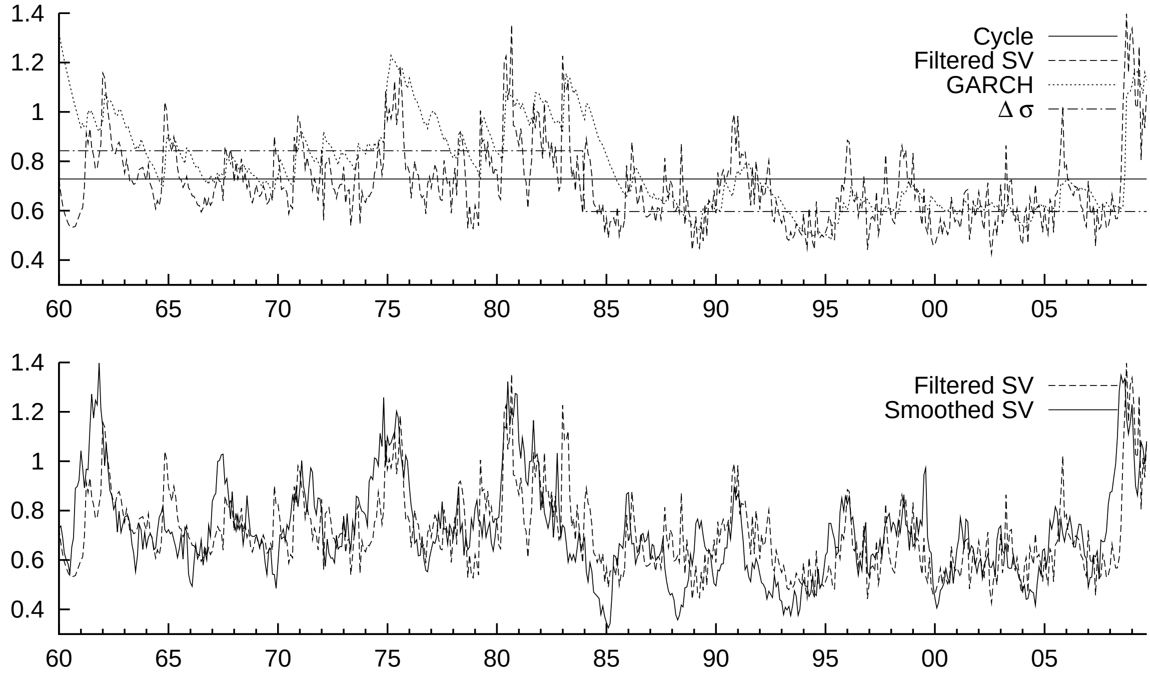

Figure 6: Filtered standard deviations of the irregular component according to *Cyc-Seas*, *Cyc-Seas-SV*, *Cyc-Seas-GARCH* and *Cyc-Seas-*$\Delta\sigma$ models (top panel), and comparison between the filtered and the smoothed standard deviations using the *Cyc-Seas-SV* model.

tend to show higher persistence than SV estimates, see Carnero et al. (2004).

Figure 6 presents the estimates of standard deviations of the different models. We clearly observe the drop of the variance around 1984. This motivated the model *Cyc-Seas-*$\Delta\sigma$ with two different variances, before and after 1984. Although this model is simple, we can only find an appropriate location for the break in the variance when the complete sample of observations is available. Therefore, the model is not appropriate for forecasting. We cannot predict a new break in the variance. For example, the variance estimates for the GARCH and SV models increase strongly after 2008:7. The model with a single variance shift cannot accommodate changing patterns in the variance.

The standard deviations are not the only components that can be extracted from the model. Figure 7 displays the estimated smoothed seasonal and cycle components in the *Cyc-Seas-SV* model. They display the main characteristics of the dynamics in U.S. growth. The cycle component exhibits the typical business cycle features. The severe recessions in 1974-1975 (oil crises) 2008-2009 (financial and credit crises) are clearly present in this component. The estimated seasonal component in the second panel shows the smooth changes in the seasonality over time which may be due to the changing structure of the U.S. economy.

Time-varying variances have an impact on signal extraction: observations in a low variance period are given (relatively) more weight and observations from a more volatile period are given less weight. For example, the IP series in 1984 becomes less volatile than in the years before 1984. The effect of this on signal extraction is illustrated in Figure 8. The weights are displayed that

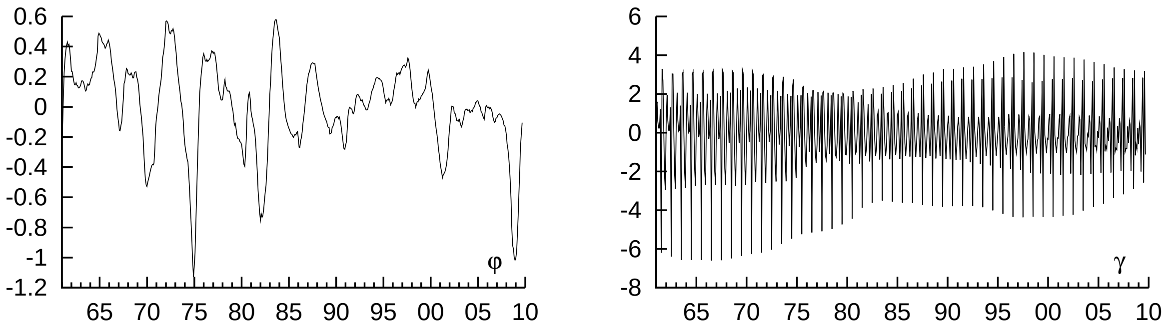

Figure 7: Smoothed estimation results for the industrial production model *Cyc-Seas-SV*: (i) estimated cycle, (ii) estimated seasonal.

are given to observations for the signal extraction of the component $\gamma_t + \rho_t$ (smooth estimate) in January 1984 based on a model with a fixed and with a stochastic variance. Koopman and Harvey (2003) provides the algorithm to compute the observation weights.

In case of a model with a fixed variance, the weights (as displayed in the first panel of Figure 8) display a symmetric pattern. The observations receive relatively the same weights in the pre- and post-1984 periods. In case of a model with a common stochastic variance, most weights are clearly higher for the post-1984 observations compared to the weights for the pre-1984 observations. These differences are due to the time-varying common variance in the model. The additional flexibility of the SV model provides the ingredients that can lead to the more successful modelling of a time series over a long horizon.

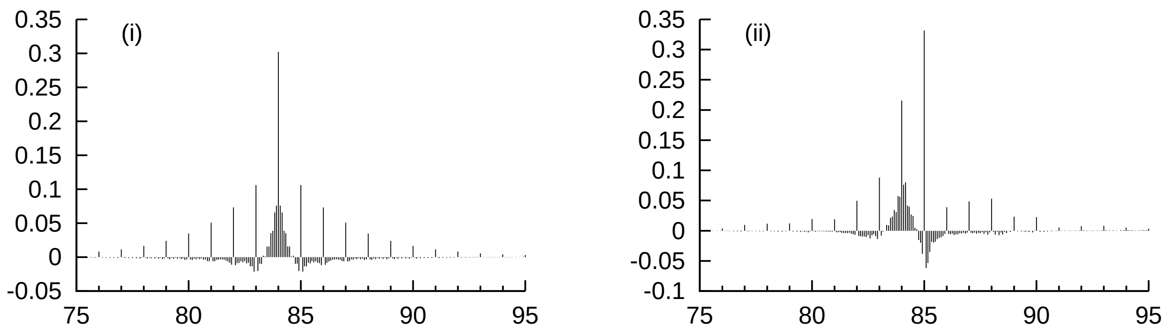

Figure 8: Weights for the Cyc-Seas (i) and Cyc-Seas-SV (ii) models, around the period 1984:1.

4.3 Estimation results for alternative specifications

To investigate whether other specifications within our model class can lead to improved or more parsimonious descriptions of the dynamic features in the data, we have considered the following model extensions.

1. *Cyc-Seas-ISV*: In Table 1, the parameter $\phi_{\rm SV}$ was restricted into the stationary region, but its estimate was found close to 1. In the first set of columns of Table 2 it is shown that the unit

Table 2: Estimation results for U.S. industrial growth $\Delta\log(\text{IP})$, combination of models.

Parameter	ISV		SV-$\Delta\sigma$		Cyc-Seas- SV-SVSeasTr		SV-SVSeasDum	
σ_{Irr}	1.106	[0.75, 1.62]	0.843	[0.73, 0.97]	0.728	[0.62, 0.85]	1.337	[1.26, 1.41]
σ_{Seas}	0.048	[0.03, 0.07]	0.038	[0.03, 0.05]	0.033	[0.03, 0.04]	0.063	[0.06, 0.07]
σ_{Cyc}	0.561	[0.36, 0.88]	0.418	[0.30, 0.58]	0.354	[0.26, 0.49]	0.824	[0.68, 1.00]
φ_{Cyc}	0.920	[0.71, 0.96]	0.921	[0.69, 0.97]	0.930	[0.76, 0.97]	0.856	[0.76, 0.91]
p_{Cyc}	57.381	[29.66,112.89]	55.800	[30.70,102.85]	54.754	[32.89,92.09]	41.032	[23.43,73.08]
σ_{SV}	0.089	[0.04, 0.18]	0.313	[0.20, 0.50]	0.247	[0.15, 0.41]	0.186	[0.11, 0.31]
ϕ_{SV}	1		0.797	[0.57, 0.92]	0.910	[0.77, 0.97]	0.921	[0.85, 0.96]
$\Delta\log\sigma^2_{1984}$			-0.157	[-0.31,-0.01]				
$A_{\text{SV sin}}$					0.077	[-0.04, 0.20]		
$B_{\text{SV sin}}$					0.581	[0.35, 0.81]		
$g_{\text{SV}i}$							See Figure 9	
LL	-796.90		-788.27		-795.01		-777.60	
Normality-DH(2)	2.73	(0.26)	1.07	(0.59)	0.47	(0.79)	0.10	(0.95)
Box-Ljung Q($\sqrt{n}-p$)	30.78	(0.04)	30.46	(0.02)	28.01	(0.03)	31.90	(0.00)
ARCH(1)	17.30	(0.00)	11.56	(0.00)	9.89	(0.00)	15.27	(0.00)
H($n/3, n/3$)	1.14	(0.36)	1.14	(0.36)	0.81	(0.15)	0.74	(0.04)

Parameter estimates with 95% confidence intervals of alternative combined models for U.S. monthly industrial growth 1960-2009. The models combine earlier components, with common stochastic variance restricted to integrated SV (ISV), with a combination of a break and SV (SV-$\Delta\sigma$) or with an additional seasonal component in the SV part (either trigonometric, in *SVSeasTr*, with amplitude A and shift B, or using seasonal dummies in *SVSeasDum*). The loglikelihood (LL) value and various diagnostic test statistics (with p-values) are also reported.

root restriction leading to the integrated stochastic volatility (ISV) model results in virtually the same loglikelihood value, and a likelihood ratio test between the ISV and standard SV model would indeed not reject the hypothesis that $\phi_{\rm SV} = 1$. The Box-Ljung test statistic for remaining autocorrelation in the residuals however now rejects the null hypothesis at the 5% level.

2. *Cyc-Seas-SV-*$\Delta\sigma$: A combination of the stochastic volatility model with the volatility break. Instead of fixing $\phi_{\rm SV}$ to one, the common SV model could be extended by allowing for an additional volatility break. This leads to a strong further increase in the loglikelihood as a result of combining the models; the single break parameter leads to an improvement of the loglikelihood of 7.6 points. Both the ARCH test and the test for residual correlation reject the null however. The remaining SV effect seems to be less important than before, with larger confidence bounds around the parameters $\phi_{\rm SV}$ and $\sigma_{\rm SV}$.

3. *Cyc-Seas-SV-SVSeasTr*, *Cyc-Seas-SV-SVSeasD*: In the original data description, it was found that the seasonalitity of this data series was of great importance. One could expect that the volatility likewise contains a seasonal pattern. The last two sets of columns propose to add an extra seasonal component within the stochastic variance component, as in

   $$\sigma_t^2 = \exp(h_t + g_t).$$

   The first set of columns, labelled *SVSeasTr*, consider a deterministic trigonometric pattern,

   $$g_t = A \sin\left[2\pi\left(B + \frac{t}{s}\right)\right], \qquad s = 12,$$

   with amplitude A and shift B are treated as two unknown coefficients that need to be estimated. The last set of columns, labelled *SVSeasD*, consider a set of seasonal dummies,

   $$g_1 = 0, \qquad g_2, \ldots, g_{12} \text{ unknown}, \qquad g_t = g_{t-12}, \qquad t = 13, \ldots, T.$$

   in which the eleven dummy coefficients $g_2, \ldots, g_{12}$ need to be estimated.

   The estimation results for the trigonometric model, *Cyc-Seas-SV-SVSeasTr* do not present tangible improvements over the original SV model. Also, the amplitude A is estimated as not significant. The alternative specification using seasonal dummies leads to an improvement of the log-likelihood by 18.2 points using 11 additional parameters; a likelihood-ratio test would still indicate that the seasonal dummies do not lead to a significant increase of the fit. However, some individual volatility dummies $g_{{\rm SV},i}$ are estimated as significant; the estimated seasonal coefficients are presented in Figure 9.

The alternative specifications do not lead to clear improvements in the fit over our more parsimonious and robust Cycle-Seasonal-Stochastic Volatility model specified in Sections 3.1-3.2. It appears that a break in the variance has occurred. Allowing for this break does not change our earlier findings significantly. The Cyc-Seas-SV model can be considered as a reasonable and robust model for fitting the most important characteristics of the data set.

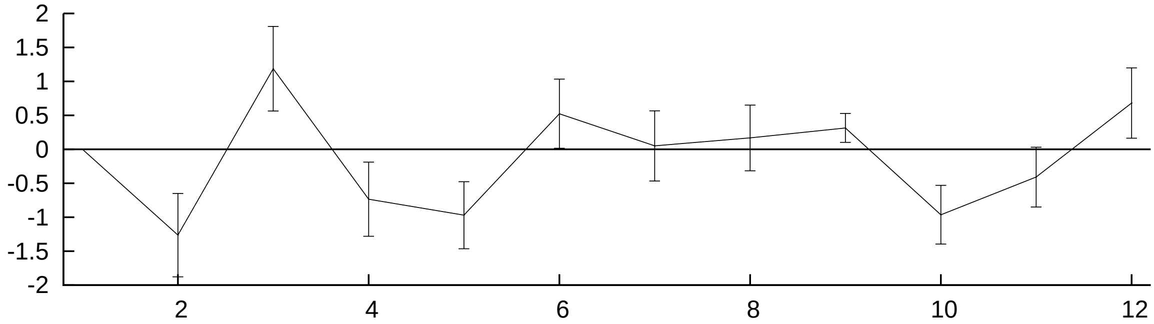

Figure 9: Monthly log-variance offsets $g_{\mathrm{SV},i}$ with 95% confidence bounds for the *Cyc-Seas-SV-SVSeasD* model.

4.4 Recursive parameter estimation

To investigate the robustness of our empirical results, we have re-estimated the parameters in the different model specifications for an increasing window of observations. From these recursive estimation results we can learn how parameter estimates vary across different samples.

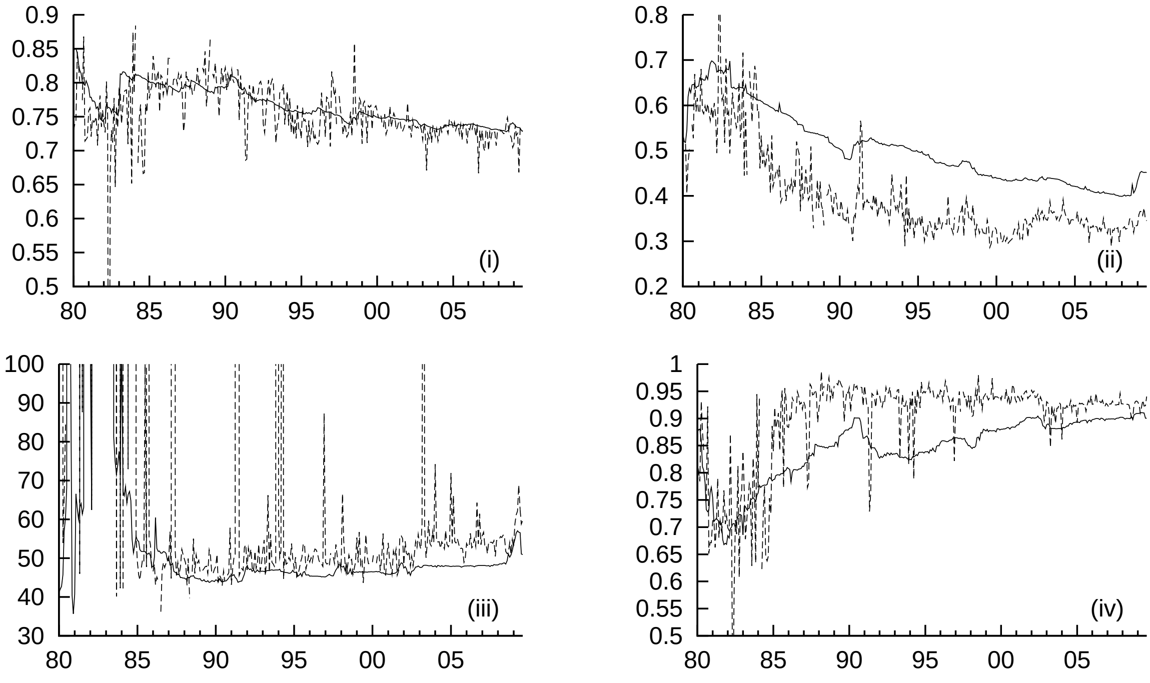

Figure 10: Recursive parameter estimates $\sigma_{\mathrm{Irr}}, \sigma_{\mathrm{Cyc}}, p_{\mathrm{Cyc}}$ and φ_{Cyc} in panels (i) – (iv), for the *Cyc-Seas* (solid line) and *Cyc-Seas-SV* (dotted line) models.

Figure 10 displays the recursive estimates of four parameters in the mean equation of the model. Estimation starts with the sample 1960:1-1980:1, and repetitively an observation is added to the sample until the full sample of 1960:1-2009:9 is reached. The first panel displays the estimate of σ_{Irr} for the *Cyc-Seas* (solid line) and $\sigma_{\mathrm{Irr}}\overline{\sigma}$ for the *Cyc-Seas-SV* (dotted line) models. For the latter

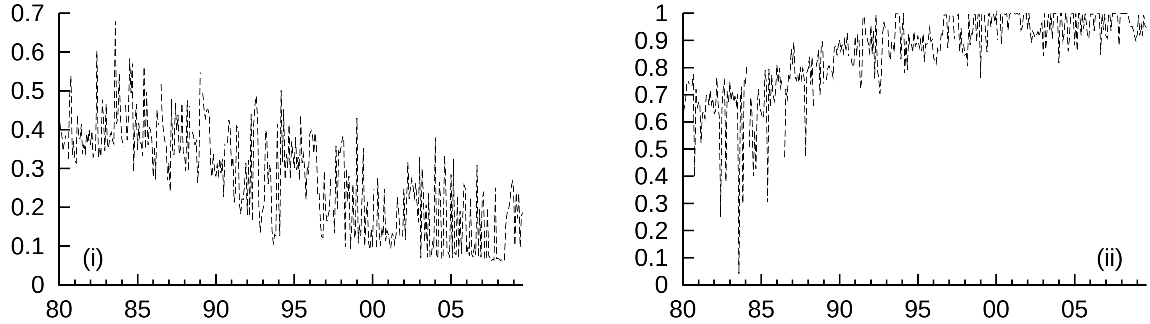

Figure 11: Recursive parameter estimates σ_{SV} and ϕ_{SV} in panels (i) – (ii), for the *Cyc-Seas-SV* model.

model, σ_{Irr} was multiplied by $\overline{\sigma}_t$, the average time varying standard deviation.

The levels and shapes of the estimate of σ_{Irr} are fully comparable for both models. The main difference is found in the estimates for the model including the common stochastic variance being more volatile; the added uncertainty in the estimation of having to extract the variance sequence as well leads to much more pronounced parameter uncertainty, especially for smaller sample sizes. For σ_{Seas} (not displayed in the figure) and σ_{Cyc} (in panel (ii)), the overall shape of the recursively estimated parameter plot is similar between the two models, though for longer samples the season and cycle standard deviation is lower for the model including SV.

The parameter governing the period, p_{Cyc}, displays a different effect. For smaller samples, the data effectively is not informative enough to estimate the cycle period, and estimates tend to infinity. This would imply that for samples shorter than approximately 1960:1–1985:1, an AR(1) component instead of the cycle would suffice. For these smaller sample sizes, the cycle (or: AR(1)) autocorrelation parameter φ_{Cyc} is estimated around 0.75. For longer sample sizes, and more pronounced cycles, the parameter increases to 0.9–0.95, especially for the model including the SV.

Figure 11 displays the resulting parameters for the SV component. It is relatively hard to estimate these parameters, as seen from the large variability of successive estimates. σ_{SV} steadily drops to a level of around 0.2 at the end of the sample, when the autocorrelation of the volatility ϕ_{SV} moves up towards the limit of 1. As seen before, it seems that the restriction of having $\phi_{\mathrm{SV}} \equiv 1$ cannot be rejected for larger sample sizes.

When a similar exercise of reestimating the model is performed using the *Cyc-Seas-GARCH* specification, the shape of the basic parameters is similar. The GARCH parameters however show how pre-1994 only an ARCH-model is found in the data, with $\alpha \approx 0.3$. In 1994 there is a series of years where the data cannot decide on the best parameter values, and starting 1995 it settles on a GARCH specification, with $\delta \approx 0.95, \alpha \approx 0.05$. This sudden alteration in the parameter estimates is a sign of a lack of robustness for this specification.

5 Concluding remarks

The introduction of time-varying variances in a model can lead to clear improvements in the fit of a time series. In this paper we have given evidence of such improvements in the context of an unobserved components time series model for the analysis of monthly growth in U.S. Industrial Production over a long horizon. We have shown that the inclusion of a stochastically time-varying common variance component can lead to substantial improvements in the fit of the time series. The adaptation of the variance enables the use of observations in volatile and non-volatile periods although observations from different periods will be weighted differently. At a more volatile period, the model assigns relatively more weight to observations that originate from a low volatile period. The implementation of this methodology is straightforward within our model-based approach to time series analysis. The empirical illustration has shown the effectiveness of our approach. In particular, we have shown that our basic decomposition model is appropriate over a long time period that includes the major economic crises (1974-1975 and 2008-2009) and the great moderation of volatility after 1985.

References

Aguilar, O. and M. West (2000). Bayesian dynamic factor models and portfolio allocation. *Journal of Business and Economic Statistics 18*(3), 338–357.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics 31*(3), 307–327.

Bos, C. S., R. J. Mahieu, and H. K. Van Dijk (2000). Daily exchange rate behaviour and hedging of currency risk. *Journal of Applied Econometrics 15*(6), 671–696.

Broto, C. and E. Ruiz (2006). Unobserved component models with asymmetric conditional variances. *Computational Statistics and Data Analysis 50*(9), 2146–2166.

Bruno, G. and C. Lupi (2004). Forecasting industrial production and the early detection of turning points. *Empirical Economics 29*(3), 647–671.

Carnero, M. A., D. Peña, and E. Ruiz (2004). Persistence and kurtosis in GARCH and stochastic volatility models. *Journal of Financial Econometrics 2*(2), 319–342.

Cecchetti, S. G., P. Hooper, B. C. Kasman, K. L. Schoenholtz, and M. W. Watson (2007). Understanding the evolving inflation process. U.S. Monetary Policy Forum 2007.

Christophe, A. and A. Doucet (2002). Particle filtering for partially observed Gaussian state space models. *Journal of the Royal Statistical Society, Series B 64*(4), 827–836.

Doornik, J. A. (2009). *Ox6: An Object-Oriented Matrix Programming Language.* London: Timberlake Consultants Ltd.

Doornik, J. A. and H. Hansen (1994). An omnibus test for univariate and multivariate normality. Technical report, Nuffield College, Oxford, UK.

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

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflations. *Econometrica 50*, 987–1008.

Ewing, B. T. and M. A. Thompson (2008). Industrial production, volatility, and the supply chain. *International Journal of Production Economics 115*(2), 553–559.

Franses, P. H. and D. Van Dijk (2005). The forecasting performance of various models for seasonality and nonlinearity for quarterly industrial production. *International Journal of Forecasting 21*(1), 87–102.

Harvey, A. C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter.* Cambridge: Cambridge University Press.

Harvey, A. C., E. Ruiz, and E. Sentana (1992). Unobserved component time series with ARCH disturbances. *Journal of Econometrics 52*(1-2), 129–157.

Harvey, A. C., E. Ruiz, and N. Shephard (1994). Multivariate stochastic variance models. *Review of Economic Studies 61*(2), 247–264.

Hendry, D. F. and M. P. Clements (2000). Economic forecasting in the face of structural breaks. In S. Holly and M. Weale (Eds.), *Econometric Modelling: Techniques and Applications*, Chapter 2, pp. 3–37. Cambridge University Press.

Heravi, S., D. R. Osborn, and C. R. Birchenhall (2004). Linear versus neural network forecasts for European industrial production series. *International Journal of Forecasting 20*(3), 435–446.

Jacquier, E., N. G. Polson, and P. E. Rossi (1994). Bayesian analysis of stochastic volatility models (with discussion). *Journal of Business and Economic Statistics 12*(4), 371–417.

Kawasaki, Y. and P. H. Franses (2004). Do seasonal unit roots matter for forecasting monthly industrial production? *Journal of Forecasting 23*(2), 77–88.

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

Koopman, S. J. and C. S. Bos (2004). State space models with a common stochastic variance. *Journal of Business and Economic Statistics 22*(3), 346–357.

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 (2008). *SsfPack 3.0: Statistical Algorithms for Models in State Space Models*. London: Timberlake Consultants Ltd.

Ljung, G. M. and G. E. P. Box (1978). On a measure of lack of fit in time series models. *Biometrika 65*(2), 67–72.

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

Stock, J. H. and M. W. Watson (1999). Business cycle fluctuations in U.S. macroeconomic time series. In T. J. B. and W. Michael (Eds.), *Handbook of Macroeconomics*, Volume 1a, Chapter 1, pp. 3–64. Amsterdam: Elsevier Science Publishers.

Stock, J. H. and M. W. Watson (2002). *Has the Business Cycle Changed and Why?*, Volume 17, pp. 159–218. The University of Chicago Press.

Stock, J. H. and M. W. Watson (2007). Why has U.S. inflation become harder to forecast? *Journal of Money, Credit, and Banking 39*(s1), 3–33.