



SCHOOL OF ECONOMICS AND MANAGEMENT
FACULTY OF SOCIAL SCIENCES
AARHUS UNIVERSITY



CREATES
Center for Research in Econometric Analysis of Time Series

CREATES Research Paper 2011-30

Stochastic trends and seasonality in economic time

Stefano Grassi and Tommaso Proietti

School of Economics and Management
Aarhus University
Bartholins Allé 10, Building 1322, DK-8000 Aarhus C
Denmark

Stochastic trends and seasonality in economic time
series: new evidence from Bayesian stochastic
model specification search

Tommaso Proietti¹
University of Sydney

Stefano Grassi²
CREATES, Aarhus University

September 2, 2011

¹Address for Correspondence: Room 499 Merewether Building (H04), Discipline of Operations Management and Econometrics, The University of Sydney, NSW 2006. *E-mail*: t.proietti@econ.usyd.edu.au.

²CREATES, Aarhus University DK-8000 Aarhus C, Denmark. *E-mail*: sgrassi@creates.au.dk. Financial support from CREATES, funded by the Danish National Research Foundation, is gratefully acknowledged by Stefano Grassi.

Abstract

An important issue in modelling economic time series is whether key unobserved components representing trends, seasonality and calendar components, are deterministic or evolutive. We address it by applying a recently proposed Bayesian variable selection methodology to an encompassing linear mixed model that features, along with deterministic effects, additional random explanatory variables that account for the evolution of the underlying level, slope, seasonality and trading days. Variable selection is performed by estimating the posterior model probabilities using a suitable Gibbs sampling scheme.

The paper conducts an extensive empirical application on a large and representative set of monthly time series concerning industrial production and retail turnover. We find strong support for the presence of stochastic trends in the series, either in the form of a time-varying level, or, less frequently, of a stochastic slope, or both. Seasonality is a more stable component: only in 70% of the cases we were able to select at least one stochastic trigonometric cycle out of the six possible cycles. Most frequently the time variation is found in correspondence with the fundamental and the first harmonic cycles.

An interesting and intuitively plausible finding is that the probability of estimating time-varying components increases with the sample size available. However, even for very large sample sizes we were unable to find stochastically varying calendar effects.

Keywords: Nonstationarity. Variable selection. Linear Mixed Models.

J.E.L. Classification: E32, E37, C53

1 Introduction

Economic time series, recorded at monthly time intervals, exhibit trends, seasonality and the effects due to the aliasing of the weekly cycle in economic activity. Modeling and extracting these component has represented an important problem in time series analysis. See Zellner (1978), Zellner (1983) Nerlove et al. (1979), Harvey (1989) Hylleberg (1992), Peña et al. (2001), and Ghysels and Osborn (2001), among others.

Figure 1 displays three such series. The first is the French index of industrial production for total manufacturing; the middle series is the index of retail turnover for Germany, while the third is the UK retail turnover index. These series show trending behaviour and a strong seasonal pattern. It is definitely less straightforward to be able to spot the effect of trading days and moving festivals from the graph, but their contribution is also relevant. An interesting question is whether these components can be adequately represented by deterministic functions of time. For instance, the trend may be modelled by a time polynomial, and the seasonal component by a combination of sine and cosine functions with pre-specified frequencies. An alternative view is that these components are subject to random evolution, and thus we need more elaborate stochastic processes to model them.

The time series literature offers methods for discriminating the deterministic generation hypothesis against the stochastic one. One approach is performing the class of seasonal unit root tests proposed by Hylleberg et al. (1990), which is based on the finite autoregressive representation of the series and tests for the presence of roots with unit modulus and zero or seasonal phase in the autoregressive polynomial. An alternative approach is to carry out the stationarity tests proposed by Canova and Hansen (1995) and extended by Busetti and Harvey (2003).

In this paper we propose to investigate the issue as a model selection problem within a mixture model that encompasses both deterministic and stochastic generation hypotheses. For this purpose, we extend the stochastic model specification search proposed by Frühwirth-Schnatter and Wagner (2010), and applied by Proietti and Grassi (2012). The mixture model nests the different specifications for the components, with the elements of the mixture representing the evolution of a particular unobserved components, such as a stochastic level, a stochastic slope, a stochastic trigonometric cycle defined at the fundamental frequency and at the harmonics. By setting up a suitable Gibbs sampling scheme we can sample the indicators of the mixture, as well as the model parameters and underlying state, and obtain a Monte Carlo estimate of the posterior probability for the various different specifications. Deterministic components are obtained by imposing exclusion restrictions. Hence, discriminating between deterministic and stochastic components amounts to performing variable selection within a regression framework that is similar to that considered by George and McCulloch (1993).

The central contribution of this paper lies with the empirical analysis, as we apply

the methodology to a dataset consisting of 530 time series, with the aim of assessing the case for the presence of stochastic trends, seasonals and trading days effects in economic time series. For each of the series belonging to the dataset we perform model selection and evaluate the frequency by which time evolving components were selected. Since the available series are characterised by different lengths, we will be able to assess the role of the sample size in the probability of detecting time variation in the components. We find the evidence for the presence of stochastic trends overwhelming, whereas the probability of detecting stochastic variation in the seasonal cycles depends crucially on the length of the available series.

The paper is structured as follows. The reference model will be presented in Section 2. Section 3 discusses its relation with the literature and contextualises the various specifications. Section 4 discusses how stochastic model specification search can be applied for the selection of the components of the linear mixed models. This hinges on a convenient reparameterization of the standard deviations of the disturbances that drive the components. Section 5 discusses the state space representation of the non-centered model and Markov Chain Monte Carlo (MCMC) inference via Gibbs sampling for model selection and Bayesian estimation of the hyperparameters and the components. After a brief description of the dataset, section 6 presents the empirical results. In Section 7 we draw our conclusions.

2 An encompassing linear mixed model with trend and seasonal effects

Let y_t denote a time series observed at $t = 1, 2, \dots, n$. We focus on modelling y_t by a linear mixed model that accounts for a trend component, denoted μ_t , a seasonal component, S_t , a calendar component, C_t , and an irregular disturbance term, ϵ_t , specified as follows:

$$y_t = \mu_t + S_t + C_t + \epsilon_t, \quad t = 1, \dots, n, \quad (1)$$

The trend component has a deterministic linear part, and a random part, specified as follows:

$$\begin{aligned} \mu_t &= \mu_0 + q_0 t + \sigma_\eta \tilde{\mu}_t + \sigma_\zeta \tilde{A}_t, \\ \tilde{\mu}_t &= \tilde{\mu}_{t-1} + \tilde{\eta}_t, & \tilde{\eta}_t &\sim \text{NID}(0, 1), \\ \tilde{A}_t &= \tilde{A}_{t-1} + \tilde{q}_{t-1}, \\ \tilde{q}_t &= \tilde{q}_{t-1} + \tilde{\zeta}_t, & \tilde{\zeta}_t &\sim \text{NID}(0, 1), \end{aligned} \quad (2)$$

here $\tilde{\mu}_t$ is a random walk component with starting value $\tilde{\mu}_0 = 0$ and unit size; the parameter $\sigma_\eta \geq 0$ establishes the scale of this component. The process \tilde{A}_t is an integrated random walk (such that $\tilde{q}_0 = \tilde{A}_0 = 0$), driven by standard normal




Figure 1: Monthly time series: 1. France, Index of Industrial Production for Total Manufacturing; 2. Germany, Index of Retail Turnover, Total; 3. UK, Index of Retail Turnover, Total. Source: Eurostat, Europa Database.

disturbances, accounting for the random evolution of the slope; $\sigma_\zeta \geq 0$ is the scale parameter for the component.

The seasonal component results from the sum of six trigonometric cycles defined at the seasonal frequencies $\lambda_j = 2\pi j/12$, $j = 1, \dots, 6$. In particular, $S_t = \sum_{j=1}^6 S_{jt}$, with each S_{jt} made up of a deterministic and a random component: for $j = 1, \dots, 5$,

$$\begin{aligned}
 S_{jt} &= a_{j0} \cos \lambda_j t + b_{j0} \sin \lambda_j t + \sigma_j \left(\tilde{a}_{jt} \cos \lambda_j t + \tilde{b}_{jt} \sin \lambda_j t \right), \\
 \tilde{a}_{jt} &= \tilde{a}_{j,t-1} + \tilde{\omega}_{jt}, \quad \tilde{\omega}_{jt} \sim \text{NID}(0, 1), \quad t = 1, \dots, n, \\
 \tilde{b}_{jt} &= \tilde{b}_{j,t-1} + \tilde{\omega}_{jt}^*, \quad \tilde{\omega}_{jt}^* \sim \text{NID}(0, 1).
 \end{aligned} \tag{3}$$

with starting values $\tilde{a}_{j0} = \tilde{b}_{j0} = 0$ whereas, for $j = 6$,

$$\begin{aligned}
 S_{6t} &= a_{60}(-1)^t + \sigma_6 \tilde{a}_{6t}(-1)^t \\
 \tilde{a}_{6t} &= \tilde{a}_{6,t-1} + \tilde{\omega}_{6t}, \quad \tilde{\omega}_{6t} \sim \text{NID}(0, 1).
 \end{aligned} \tag{4}$$

The parameters a_{j0} and b_{j0} , $j = 1, \dots, 6$, determine the amplitude of the fixed trigonometric cycles, whereas σ_j regulate the contribution of the random component.

The calendar component, C_t , plays an important role for the class of economic time series that we investigate in the paper. In fact, production and sales are characterised by a strong weekly cycle, which is aliased since the data are recorded with reference to the months. The component accounts for trading days (TD) effects and for moving festivals. The former are related to the fact that the number of weekdays and weekend days is not the same across the months. Let D_{jt} denote the number of days of type j , $j = 1, \dots, 7$, occurring in month t , and define $x_{kt} = D_{jt} - D_{7t}$, $k = 1, \dots, 6$, which is a contrast between the number of days of a particular type (Mondays, Tuesdays, ..., Saturdays), and the number of Sundays occurring in the same month. A time varying trading day component can be modelled as a regression component with time-varying coefficients:

$$\begin{aligned} TD_t &= \sum_{k=1}^6 \phi_{k0} x_{kt} + \sigma_\nu \left(\sum_{k=1}^6 \tilde{\phi}_{kt} x_{kt} \right) \\ \tilde{\phi}_{kt} &= \tilde{\phi}_{k,t-1} + \tilde{\nu}_t, \quad \tilde{\nu}_t \sim \text{NID}(0, 1). \end{aligned} \quad (5)$$

The coefficients associated with the regressors x_{kt} evolve as independent random walks with starting value $\tilde{\phi}_{k0} = 0$. Obviously, if $\sigma_\nu = 0$ the trading days effect are time invariant.

As far as moving festivals are concerned, we focus on Easter and Labor Day (U.S. time series), and model their effects defining explanatory variable measuring the proportion of 7 days before Easter (x_{Et}) or Labor Day (x_{Lt}) that fall in month t and subtracting their monthly long run average, computed over the first 400 years of the Gregorian calendar (1583-1982). This treatment is quite standard in the literature; see Bell and Hillmer (1983), among others, and the references therein.

Finally, the irregular component is a Gaussian white noise process, $\epsilon_t \sim \text{NID}(0, \sigma_\epsilon^2)$.

3 Discussion

The linear mixed model proposed in the previous section is sufficiently general to accommodate both deterministic and stochastic trends, seasonals and calendar effects. The specification with σ_j constant across j and $\sigma_{TD} = 0$ is referred to as the basic structural model; see Harvey (1989). The representation used for the components is known as the non-centred (with respect to location and scale) representation; Frühwirth-Schnatter and Wagner (2010) and Strickland et al. (2007) discuss its advantages for Bayesian estimation of the model.

Notice that the trend component (2) can be written equivalently using the following recursions:

$$\begin{aligned} \mu_t &= \mu_{t-1} + q_{t-1} + \eta_t, & \eta_t &\sim \text{NID}(0, \sigma_\eta^2) \\ q_t &= q_{t-1} + \zeta_t, & \zeta_t &\sim \text{NID}(0, \sigma_\zeta^2) \end{aligned} \quad (6)$$

where q_t is the slope component and we assume that η_t and ζ_t are mutually uncorrelated and independent of ϵ_t and S_t (see Harvey (1989) and West and Harrison (1997)). The trend model is related to cubic spline smoothing (see Wecker and Ansley (1983)). If $\sigma_\zeta = 0$ the trend is a random walk with constant drift; if $\sigma_\eta = 0$ and $\sigma_\zeta > 0$ the trend is an integrated random walk; finally, if both $\sigma_\eta = \sigma_\zeta = 0$ the trend is linear deterministic.

The seasonal component consists of six cycles: the first is defined at the fundamental frequency, $\lambda_1 = \pi/6$ (corresponding to a period of 12 monthly observations) while the others are defined at the harmonic frequencies $\lambda_j = 2\pi j/12$, $j = 2, \dots, 6$, (corresponding, respectively, to periods of 6 months, i.e. two cycles in a year, 4 months, i.e. three cycles in a year, 3 months, i.e. four cycles in a year, 2.4, i.e. five cycles in a year, and 2 months).

Using trigonometric identities, the j -th seasonal cycle in 4 can be rewritten: $S_{jt} = \varphi_t \cos(\lambda_j t - \vartheta_t)$, where

$$\varphi_t = \sqrt{\left(a_{j0} + \sum_{k=0}^{t-1} \tilde{\omega}_{j,t-k}\right)^2 + \left(b_{j0} + \sum_{k=0}^{t-1} \tilde{\omega}_{j,t-k}^*\right)^2}$$

is the time varying amplitude and

$$\vartheta_t = \tan^{-1} \left(\frac{a_{j0} + \sum_{k=0}^{t-1} \tilde{\omega}_{j,t-k}}{b_{j0} + \sum_{k=0}^{t-1} \tilde{\omega}_{j,t-k}^*} \right)$$

represents the phase shift. If $\sigma_1 = \dots = \sigma_6 = 0$, the seasonal component is the sum of six perfectly deterministic cycles. The recursive representation of the j -th seasonal cycle is

$$\begin{bmatrix} S_{jt} \\ S_{jt}^* \end{bmatrix} = \begin{bmatrix} \cos \lambda_j & \sin \lambda_j \\ -\sin \lambda_j & \cos \lambda_j \end{bmatrix} \begin{bmatrix} S_{j,t-1} \\ S_{j,t-1}^* \end{bmatrix} + \begin{bmatrix} \varpi_{j,t} \\ \varpi_{j,t}^* \end{bmatrix}, j = 1, \dots, 5, \quad (7)$$

and $S_{6,t} = -S_{6,t-1} + \varpi_{6,t}$, where $\varpi_{jt} \sim \text{NID}(0, \sigma_j^2)$, $j = 1, \dots, 6$, $\varpi_{jt}^* \sim \text{NID}(0, \sigma_j^2)$, $j = 1, \dots, 5$. These recursions hold for $t = 1, \dots, n$, with starting values $S_{j,0} = a_{j0}$ and $S_{j,0}^* = b_{j0}$. This representation is the one usually adopted in the time series literature (see Harvey (1989) and West and Harrison (1997)).

4 Bayesian stochastic specification search

A widely debated issue is whether trends, seasonals and trading day effects are deterministic or stochastically evolving over time; this translates into the following main specification issues with respect to the model set up in section 2:

- when $\sigma_\zeta = 0$, *cæteris paribus*, the trend changes are white noise around a constant drift;
- when $\sigma_\eta = \sigma_\zeta = 0$, *cæteris paribus*, the trend is linear deterministic.
- When $\sigma_j = 0, \forall j$, *cæteris paribus*, seasonality is represented by a set of perfectly periodic deterministic components.
- When $\sigma_\nu = 0$, *cæteris paribus*, the TD coefficients are time invariant.

In the econometric literature formal statistical tests are available for discriminating deterministic trends from stochastic ones. When seasonality is absent, unit root tests, see Dickey and Fuller (1979) and Phillips and Perron (1988), test the null of integration versus a stationary alternative see De Jong and Whiteman (1991), Koop (1992), Sims (1988), Sims and Uhlig (1991), Phillips (1991), Schotman and van Dijk (1991), Phillips and Perron (1994), among others, for the Bayesian approach to unit root testing; on the contrary, the tests proposed by Nyblom and Makelainen (1983) and Kwiatkowski et al. (1992) test trend stationarity against the alternative of integration. Unit root tests were extended to the seasonal case by Hylleberg et al. (1990), whereas the extension for stationarity tests was proposed by Canova and Hansen (1995), and Busetti and Harvey (2003). Other important references on whether seasonality is stochastically evolving over time include Hylleberg and Pagan (1997) and Koop and van Dijk (2000). The issue as to whether trading days affects are time varying has been addressed by Dagum et al. (1993), Dagum and Quenneville (1993), Bell and Martin (2004).

We can decide on the above main specification issues using the specification search methodology proposed by Frühwirth-Schnatter and Wagner (2010). The approach starts with the linear mixed model representation presented in section 2 and proceeds to the reparameterization of the hyperparameters representing standard deviations as regression parameters with unrestricted support, as it will be illustrated shortly.

It should be noticed that the linear mixed model is identified up to sign switches that operate on both the standard deviations and on the underlying stochastic components. Consider, for instance the trend component in equation (6): if we replace $\sigma_\eta \tilde{\mu}_t$ by the product $(-\sigma_\eta)(-\tilde{\mu}_t)$, i.e. we switch the sign to both the elements, we obtain an observationally equivalent representation, characterised by exactly the same likelihood. FS-W came up with the clever idea of replacing $\sigma_\eta \tilde{\mu}_t$ with $\beta_\mu \mu_t^*$, where, for $t = 1, \dots, n$,

$$\beta_\mu \mu_t^* = \begin{cases} \sigma_\eta \tilde{\mu}_t, & \text{with probability 0.5} \\ (-\sigma_\eta)(-\tilde{\mu}_t), & \text{with probability 0.5} \end{cases}$$

Hence, the sign switch is the outcome of a Bernoulli random experiment, with 50% success probability. According to this setting, the parameter β_μ can take any real value and it would be suitable to set up a normal prior for it centred in zero.

The same reasoning can be applied to the pairs $(-\sigma_\zeta)(-\tilde{A}_t)$ and $(\sigma_\zeta)(\tilde{A}_t)$, $(-\sigma_j) \left[- \left(\tilde{a}_{jt} \cos \lambda_j t + \tilde{b}_{jt} \sin \lambda_j t \right) \right]$ and $\sigma_j \left(\tilde{a}_{jt} \cos \lambda_j t + \tilde{b}_{jt} \sin \lambda_j t \right)$, and $(-\sigma_\nu) \left(- \sum_{k=1}^6 \tilde{\phi}_{kt} x_{kt} \right)$ and $\sigma_\nu \left(\sum_{k=1}^6 \tilde{\phi}_{kt} x_{kt} \right)$. The likelihood function is symmetric around zero along the parameter space and multimodal, if the true standard deviations are larger than zero, as resulting from the identifiability issue. This feature will be later exploited to judge whether the posterior of $\sigma_\eta, \sigma_\zeta, \sigma_j, j = 1, \dots, 6$, and σ_ν , is far away from or sufficiently close to zero.

The random switch process can be formalised by defining independent Bernoulli random variates with success probability 0.5, $\mathbf{B}_\mu, \mathbf{B}_A, \mathbf{B}_{sj}, j = 1, \dots, 6, \mathbf{B}_{TD}$, so that we can use the reparameterisation $\sigma_\eta \tilde{\mu}_t = \beta_\mu \mu_t^*$, where $\beta_\mu = (-1)^{\mathbf{B}_\mu} \sigma_\eta$, and $\mu_t^* = (-1)^{\mathbf{B}_\mu} \tilde{\mu}_t$; similarly, $\sigma_\zeta \tilde{A}_t = \beta_A A_t^*$, where $\beta_A = (-1)^{\mathbf{B}_A} \sigma_\zeta$, $A_t^* = (-1)^{\mathbf{B}_A} \tilde{A}_t$,

$$\sigma_j \left(\tilde{a}_{jt} \cos \lambda_j t + \tilde{b}_{jt} \sin \lambda_j t \right) = \beta_{sj} U_{jt}^*, \quad \beta_{sj} = (-1)^{\mathbf{B}_{sj}} \sigma_j, U_{jt}^* = (-1)^{\mathbf{B}_{sj}} \left(\tilde{a}_{jt} \cos \lambda_j t + \tilde{b}_{jt} \sin \lambda_j t \right),$$

for $j = 1, \dots, 6$, and

$$\sigma_\nu \left(\sum_k \phi_{kt} x_{kt} \right) = \beta_{TD} \Phi_t^*, \quad \beta_{TD} = (-1)^{\mathbf{B}_{TD}} \sigma_\nu, \Phi_t^* = (-1)^{\mathbf{B}_{TD}} \left(\sum_k \phi_{kt} x_{kt} \right).$$

As stated above, the reparameterisation aims at transforming a standard deviation into a regression parameter in a linear mixed model, so that the selection of an evolutive component is related to the inclusion of a particular explanatory variable. The different specifications for the trend and the seasonal components are obtained by imposing exclusion restrictions, so that discriminating between deterministic and stochastic components amounts to performing variable selection within the regression framework considered by George and McCulloch (1993).

Although in principle we could conduct variable selection for any of the explanatory variables in the model, for our purposes, it will suffice to carry it out on the slope term $q_0 t$, on the random walk and integrated random walk components μ_t^*, A_t^* , on the six stochastic terms U_{jt}^* and on $\left(\sum_{k=1}^6 \Phi_{kt}^* x_{kt} \right)$. We then introduce nine binary indicator variables $\gamma_\mu, \gamma_A, \gamma_{sj}, j = 1, \dots, 6, \gamma_{TD}$, taking value 1 if the random effects $\mu_t^*, A_t^*, U_{jt}^*, j = 1, \dots, 6, \left(\sum_{k=1}^6 \Phi_{kt}^* x_{kt} \right)$ are present and 0 otherwise, along with a binary indicator for the linear trend component, δ , taking values (0,1) according to whether the term $q_0 t$ is included in the model. The ten indicators can be further collected in the multinomial vector $\Upsilon = (\gamma_\mu, \gamma_A, \gamma_{sj}, j = 1, \dots, 6, \gamma_{TD}, \delta)$.

Considering all the possible values of Υ , there are $K = 2^{10} = 1024$ possible models in competition, which are nested in the specification:

$$\begin{aligned}
y_t &= \mu_t + S_t + C_t + \epsilon_t, & \epsilon_t &\sim \text{NID}(0, \sigma_\epsilon^2), \\
\mu_t &= \mu_0 + \delta q_0 t + \gamma_\mu \beta_\mu \mu_t^* + \gamma_A \beta_A A_t^*, \\
\mu_t^* &= \mu_{t-1}^* + \tilde{\eta}_t, & \tilde{\eta}_t &\sim \text{NID}(0, 1), \\
A_t^* &= A_{t-1}^* + \tilde{q}_{t-1}, \\
\tilde{q}_t &= \tilde{q}_{t-1} + \tilde{\zeta}_t, & \tilde{\zeta}_t &\sim \text{NID}(0, 1), \\
S_t &= \sum_{j=1}^5 (a_{j0} \cos \lambda_j t + b_{j0} \sin \lambda_j t) + a_{60} (-1)^t + \sum_{j=1}^6 \gamma_{sj} \beta_{sj} U_{jt}^*, \\
U_{jt}^* &= A_{jt}^* \cos \lambda_j t + B_{jt}^* \sin \lambda_j t, \quad j = 1, \dots, 5, & U_{6t}^* &= A_{6t}^* \cos \pi t, \\
A_{jt}^* &= A_{j,t-1}^* + \tilde{\omega}_{jt}, & \tilde{\omega}_{jt} &\sim \text{NID}(0, 1), \\
B_{jt}^* &= B_{j,t-1}^* + \tilde{\omega}_{jt}^*, & \tilde{\omega}_{jt}^* &\sim \text{NID}(0, 1), \\
C_t &= \sum_{k=1}^6 \phi_{k0} x_{kt} + \gamma_{TD} \beta_{TD} \left(\sum_{k=1}^6 \Phi_{kt}^* x_{kt} \right) + \phi_E x_{Et} + \phi_L x_{Lt}, \\
\Phi_{kt}^* &= \Phi_{k,t-1}^* + \tilde{\nu}_t, & \tilde{\nu}_t &\sim \text{NID}(0, 1).
\end{aligned} \tag{8}$$

where we have defined $A_{jt}^* = (-1)^{\mathbf{B}_{sj}} \tilde{a}_{jt}$, $B_{jt}^* = (-1)^{\mathbf{B}_{sj}} \tilde{b}_{jt}$, $\Phi_{kt}^* = (-1)^{\mathbf{B}_{TD}} \phi_{kt}^*$.

All the specifications will include the constant term, the set of 11 sine and cosine terms at the seasonal frequencies, the six trading days regressors and the moving festivals regressors, so that the most elementary model is a model with a constant level, deterministic seasonals and fixed calendar effects; this corresponds to the most elementary model and has $\Upsilon = (0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)$. When $\Upsilon = (0, 0, 0, 0, 0, 0, 0, 0, 0, 1)$, the model features a deterministic linear trend and a perfectly deterministic seasonal component (assuming there is no moving festival):

$$y_t = \mu_0 + q_0 t + \sum_{j=1}^5 (a_{j0} \cos \lambda_j t + b_{j0} \sin \lambda_j t) + a_{60} (-1)^t + \sum_{k=1}^6 \phi_{k0} x_{kt} + \epsilon_t, \tag{9}$$

In turn, $\Upsilon = (1, 0, 1, 0, 0, 0, 0, 0, 0, 1)$ corresponds to

$$y_t = \mu_0 + q_0 t + \sigma_\eta \tilde{\mu}_t + \sum_{j=1}^5 (a_{j0} \cos \lambda_j t + b_{j0} \sin \lambda_j t) + a_{60} (-1)^t + \sigma_1 \left(\tilde{a}_{1t} \cos \lambda_1 t + \tilde{b}_{1t} \sin \lambda_1 t \right) + \sum_{k=1}^6 \phi_{k0} x_{kt} + \epsilon_t, \tag{10}$$

The different models will be labelled by

$$M_k, \quad k = 1 + \sum_{u=1}^U 2^{U-u} \Upsilon_u,$$

where Υ_u is the u -th element of the vector Υ , $u = 1, \dots, U$. For instance, $\Upsilon = (1, 0, 1, 0, 0, 0, 0, 0, 0, 1)$ is model M_{641} .

5 Statistical Treatment

Model selection entails the computation of the posterior model probabilities $\pi(M_k|y) \propto \pi(M_k)\pi(y|M_k)$, where y denotes the collection of time series values $\{y_t, t = 1, \dots, n\}$. The evaluation of the marginal likelihood $\pi(y|M_k)$ for each model is computationally intensive; in fact, it would be unfeasible to compute the posterior model probabilities for each of the 1024 specifications and select the specification characterised by the largest. It is feasible instead to draw samples from the posterior distribution of Υ given the data by Markov Chain Monte Carlo methods, as by a suitable design of the priors the full conditional posterior distribution of the multinomial vector Υ is available in closed form. A suitable Gibbs sampling (GS) scheme can in fact be devised which enables Υ to be sampled along with the model parameters and states. After the GS scheme has converged, we estimate $\pi(\Upsilon|y)$, by the proportion of times a particular specification was drawn.

Depending on the value of Υ , the models nested in (10) admit the following state space representation:

$$\begin{aligned} y_t &= x'_{\delta,t}\rho_\delta + z'_{\gamma,t}\alpha_{\gamma,t} + \epsilon_t, \quad \epsilon_t \sim \text{NID}(0, \sigma_\epsilon^2), \quad t = 1, \dots, n, \\ \alpha_{\gamma,t} &= T_\gamma\alpha_{\gamma,t-1} + R_\gamma u_{\gamma,t}, \quad u_{\gamma,t} \sim \text{NID}(0, I), \end{aligned} \quad (11)$$

where $\alpha_{\gamma,0} = 0$, and

$$\begin{aligned} x_{\delta,t} &= (1, \delta t, \cos \lambda_1 t, \sin \lambda_1 t, \dots, \cos \pi t, x_{1t}, \dots, x_{6t}, x_{Et}, x_{Lt})' \\ \rho_\delta &= (\mu_0, q_0, a_{10}, b_{10}, \dots, a_{60}, \phi_1, \dots, \phi_6, \phi_E, \phi_L)', \\ z_{\gamma,t} &= (\gamma_\mu \beta_\mu, \gamma_A \beta_A, 0, \gamma_{s1} \beta_{s1} \cos \lambda_1 t, \gamma_{s1} \beta_{s1} \sin \lambda_1 t, \dots, \gamma_{s6} \beta_{s6} \cos \pi t, \\ &\quad \gamma_{TD} \beta_{TD} x_{1t}, \dots, \gamma_{TD} \beta_{TD} x_{6t})', \\ \alpha_{\gamma,t} &= (\mu_t^*, A_t^*, \tilde{q}_t, A_{1t}^*, B_{1t}^*, \dots, A_{6t}^*, \Phi_{1t}^*, \dots, \Phi_{6t}^*), \end{aligned}$$

$$T_\gamma = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & I_{12} \end{pmatrix}, \quad R_\gamma = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & I_{12} \end{pmatrix}.$$

We will denote by α the collection of the latent states $\{\alpha_{\gamma,t}, t = 0, 1, \dots, n\}$, and by ψ_Υ the appropriate subset of the parameters $(\mu_0, q_0, a_{10}, b_{10}, \dots, a_{60}, \phi_{10}, \dots, \phi_{60}, \beta_\mu, \beta_A, \beta_{s1}, \dots, \beta_{s6}, \beta_{TD})$ that enter the model for a particular value of Υ .

The prior has the following conditional independence structure:

$$\pi(\Upsilon, \psi, \sigma_\epsilon^2, \alpha) = \pi(\Upsilon)\pi(\sigma_\epsilon^2)\pi(\psi|\Upsilon, \sigma_\epsilon^2)\pi(\alpha|\Upsilon),$$

where individual factors are given as follows.

- We assume that the models $M_k, k = 1, \dots, K$, are equally likely a priori, that is $\pi(\Upsilon) = 2^{-U}$.

- For σ_ϵ^2 we adopt a hierarchical inverse Gamma (IG) prior: $\sigma_\epsilon^2 \sim \text{IG}(c_0, C_0)$, where $C_0 \sim \text{G}(g_0, G_0)$, $\text{G}(\cdot)$ denoting the Gamma distribution, $c_0 = 2.5$, $g_0 = 5$, and $G_0 = g_0/[0.75\text{Var}(y_t)(c_0 - 1)]$. The hierarchical prior is intended to make the posteriors less sensitive to the choice of the hyperparameters of the IG distribution.
- Denoting the i -th element of the parameter vector ψ_Υ by $\psi_{\Upsilon i}$, $i = 1 \dots, p$, we set $\pi(\psi_\Upsilon|\Upsilon, \sigma_\epsilon^2) = \prod_{i=1}^p \pi(\psi_i|\sigma_\epsilon^2)$, where all the priors are conjugate. For instance, $q_0|\sigma_\epsilon^2 \sim \text{N}(0, d_0\sigma_\epsilon^2)$, etc. A distinctive feature of the methodology proposed by Frühwirth-Schnatter and Wagner (2010) is the adoption of Gaussian priors, centered at zero, for the parameters $\beta_\mu, \beta_A, \beta_{sj}, \beta_{TD}$:

$$\beta_\mu|\sigma_\epsilon^2 \sim \text{N}(0, \kappa_\mu\sigma_\epsilon^2), \beta_A|\sigma_\epsilon^2 \sim \text{N}(0, \kappa_A\sigma_\epsilon^2),$$

$$\beta_{sj}|\sigma_\epsilon^2 \sim \text{N}(0, \kappa_j\sigma_\epsilon^2), j = 1, \dots, 6, \beta_{TD}|\sigma_\epsilon^2 \sim \text{N}(0, \kappa_{TD}\sigma_\epsilon^2).$$

Not only this allows conjugate analysis, but FS-W show that inference will benefit substantially from the use of a normal prior for e.g. $\beta_\mu = \pm\sigma_\eta$, $\beta_\mu|\sigma_\epsilon^2 \sim \text{N}(0, \kappa_\mu\sigma_\epsilon^2)$, instead of the usual inverse Gamma prior for the variance parameter σ_η^2 .

For the constant term and the coefficients $a_{j0}, j = 1, \dots, 6, b_{j0}, j = 1, \dots, 5, \phi_{k0}, k = 1, \dots, 6$ we adopt the uninformative priors $\pi(\mu_0|\sigma_\epsilon^2) \propto 1$.

- The prior distribution for α is given directly by the Gaussian dynamic model (11):

$$\pi(\alpha|\Upsilon) = \pi(\alpha_{\gamma 0}) \prod_{t=1}^n \pi(\alpha_{\gamma t}|\alpha_{\gamma, t-1}),$$

with $\alpha_{\gamma t}|\alpha_{\gamma, t-1} \sim \text{N}(T_\gamma\alpha_{\gamma, t-1}, R_\gamma R'_\gamma)$ and $\alpha_{\gamma, 0} = 0$.

The GS scheme can be sketched as follows. After specifying a set of initial values $\Upsilon^{(0)}, \sigma_\epsilon^{2(0)}, \alpha^{(0)}, \psi^{(0)}$, we iterate for $i = 1, 2, \dots, M$, the following operations:

- Draw $\Upsilon^{(i)} \sim \pi(\Upsilon|\alpha^{(i-1)}, y)$
- Draw $\sigma_\epsilon^{2(i)} \sim \pi(\sigma_\epsilon^2|\Upsilon^{(i)}, \psi^{(i-1)}, \alpha^{(i-1)}, y)$
- Draw $\psi^{(i)} \sim \pi(\psi|\Upsilon^{(i)}, \sigma_\epsilon^{2(i)}, \alpha^{(i-1)}, y)$
- Draw $\alpha^{(i)} \sim \pi(\alpha|\Upsilon^{(i)}, \sigma_\epsilon^{2(i)}, \psi^{(i)}, y)$

The above complete conditional densities are available, up to a normalizing constant, from the form of the likelihood and the prior.

For the sake of notation, let us write the linear mixed model as $y = Z_{\Upsilon}\psi_{\Upsilon} + \epsilon$, where y and ϵ are vectors stacking the values $\{y_t\}$ and $\{\epsilon_t\}$, respectively, and the generic row of matrix Z_{Υ} contains the relevant subset of the explanatory variables.

Step a. is carried out by sampling the indicators with probabilities proportional to the conditional likelihood of the regression model, as

$$\begin{aligned}\pi(\Upsilon|\alpha, y) &\propto \pi(\Upsilon)\pi(y|\Upsilon, \alpha) \\ &\propto \pi(y|\Upsilon, \alpha),\end{aligned}$$

which is available in closed form (see below).

Under the normal-inverse Gamma conjugate prior for $(\psi_{\Upsilon}, \sigma_{\epsilon}^2)$

$$\sigma_{\epsilon}^2 \sim \text{IG}(c_0, C_0), \quad \psi_{\Upsilon}|\sigma_{\epsilon}^2 \sim \text{N}(0, \sigma_{\epsilon}^2 D_{\Upsilon}),$$

where D_{Υ} is a diagonal matrix with elements κ_{μ}, κ_A , etc., steps b. and c. are carried out by sampling from the posteriors

$$\begin{aligned}\sigma_{\epsilon}^2|\Upsilon, \alpha, y &\sim \text{IG}(c_{T^*}, C_{T^*}) \\ \psi_{\Upsilon}|\Upsilon, \sigma_{\epsilon}^2, \alpha, y &\sim \text{N}(m, \sigma_{\epsilon}^2 S)\end{aligned}$$

where

$$\begin{aligned}S &= (Z'_{\Upsilon}Z_{\Upsilon} + D_{\Upsilon}^{-1})^{-1}, & m &= SZ'_{\Upsilon}y \\ c_{T^*} &= c_0 + T^*/2, & C_{T^*} &= C_0 + \frac{1}{2}(y'y - m'S^{-1}m).\end{aligned}$$

Finally,

$$\pi(y|\Upsilon, \alpha) \propto \frac{|S|^{0.5}}{|D_{\Upsilon}|^{0.5}} \frac{\Gamma(c_{T^*})}{\Gamma(c_0)} \frac{C_0^{c_0}}{C_{T^*}^{c_{T^*}}},$$

see e.g. Geweke (2005), where $\Gamma(\cdot)$ denotes the Gamma function.

The sample from the posterior distribution of the latent states, conditional on the model and its parameters, in step d., is obtained by the conditional simulation smoother proposed by Durbin and Koopman (2002).

Finally, the draw of the parameters $\beta_{\mu}, \beta_A, \beta_{sj}, j = 1, \dots, 6, \beta_{TD}$ are obtained by performing a final random sign permutation. This is achieved by drawing independently Bernoulli random variables $\mathbf{B}_{\mu}, \mathbf{B}_A, \mathbf{B}_{sj}, j = 1, \dots, 6, \mathbf{B}_{TD}$ with probability 0.5, and recording $(-1)^{\mathbf{B}_{\mu}}(\sigma_{\eta}, \tilde{\mu}_t), (-1)^{\mathbf{B}_A}(\sigma_{\zeta}, \tilde{A}_t, a_t)$, etc.

A key assumption is that σ_{ϵ}^2 is strictly greater than zero, i.e. the irregular component is always present.

6 Empirical Results

Our application deals with data set consisting of 530 monthly time series for 10 Euro Area countries, the UK, and the US, referring to the index of industrial production

and retail turnover. This is a large and representative sample, with 379 series referring to the index of industrial production (IPI) and 151 to the index of retail turnover (RT). The breakdown of the series by country and their sample period is available in Table 1. For the IPI we consider series from Sectors B (Mining and quarrying), C (Manufacturing), D (Energy), and B–D, and the series for the manufacturing sectors are from those identified by two digits of the NACE statistical classifications of economic activities (sectors C1–C31). For the US we consider the 63 series for Market and Industry Group and the 32 series for Special Aggregates and Selected Detail (see http://www.federalreserve.gov/releases/g17/table1_2.htm for more details). For retail turnover, we focus on the series available with code starting with G47 (Retail trade, except of motor vehicles and motorcycles). The sources of the series are Eurostat (<http://epp.eurostat.ec.europa.eu/portal/page/portal/eurostat/home/>), the Federal Reserve and the US Census Bureau. All the series are analysed in logarithms.

We set the scale parameters $d_0 = \kappa_\mu = \kappa_A = \kappa_\mu = \kappa_j = \kappa_{TD} = 100$ for the priors (our experience is that the results are very insensitive to the choice of these priors: although the priors have to be proper for the purposes of model selection, taking 10 or 100 or 1000 did not make a difference), and run the GS scheme outlined in the previous section. For each series y_{it} , $i = 1, \dots, N$ we record the 10 modal models, denoted Υ_{ik} , $k = 1, \dots, 10$, visited by the GS scheme, as well as the number of times they were visited. All the results are based on 40,000 MCMC draws (after a burn-in sample of 20,000 draws). Let c_{ik} denote the number of times model Υ_{ik} was selected and let c_i be the total number of draws (which is actually invariant with i). Then, $\hat{\pi}_{ij} = \frac{c_{ij}}{c_i}$ estimates the posterior probability of model Υ_j for the i -th series. Limiting ourselves to the first 10 model models is not at all restrictive as the median percent of draws absorbed by them across the 530 series amounts to 99.32%. Figure 2 shows the histogram of $\sum_{k=1}^{10} \hat{\pi}_{ik} = \frac{\sum_{j=1}^{10} c_{ij}}{c_i}$, i.e. of the total probability attached to the 10 modal specification, which turned out to be highly concentrated around 100%.

Table 2 reports the first three models that were visited more frequently by the Gibbs sampler for the time series considered in the introduction. As far as the IPI for France is concerned, the evidence is overwhelmingly in favour of a stochastic level and seasonality. An important source of model uncertainty is about the presence of a nonzero drift q_0 . The retail turnover series for Germany feature a stochastic level, possibly with a constant drift, and a time-varying seasonal cycle at the fundamental frequency; the remaining harmonics are predominantly deterministic, except the fourth (which is responsible for a cycle with period of three months). The UK retail turnover series behaves rather differently. First and foremost, the model posterior probabilities are more diffuse, the modal model being drawn only 15.66% of the times. There are four stochastic cycles at the fundamental and the first three harmonic frequencies, and the slope component is evolutive.

Figure 3 displays the distribution of the 5300 vectors Υ_{ij} , $i = 1, \dots, 530$, $j =$

Table 1: Breakdown of the series by country, sample period, number of time series.

<i>Index of industrial production</i>		
Country	Sample period	Number of series
Austria	1996.1-2010.12	28
Belgium	1995.1-2010.12	27
Finland	1990.1-2010.12	20
France	1990.1-2010.12	28
Germany	1991.1-2010.12	28
Greece	2000.1-2010.12	29
Italy	1990.1-2010.12	27
Netherlands	1990.1-2010.12	22
Portugal	1995.1-2010.12	20
Spain	1980.1-2010.12	28
UK	1990.1-2010.12	28
US	1947.1-2010.12	94
<i>Index of retail turnover</i>		
Country	Sample period	Number of series
Austria	1999.1-2010.12	6
Belgium	1998.1-2010.12	15
Finland	1995.1-2010.12	14
France	1994.1-2010.12	14
Germany	1994.1-2010.12	15
Greece	1995.1-2010.12	13
Italy	2000.1-2010.12	14
Netherlands	1996.1-2010.12	9
Portugal	1995.1-2010.12	9
Spain	2000.1-2010.12	14
UK	2000.1-2010.12	14
US	1992.1-2010.12	14



Figure 2: Distribution of $\sum_{j=1}^{10} \hat{\pi}_{ij}$ for the 530 series in our dataset.

1, ..., 10, in the model space: the horizontal axis refers to $M_k, k = 1, \dots, 1024$, and each bar is proportional to the average number of times the corresponding specification was visited per series. The graph shows that the specification M_{513} , featuring a stochastic level, no slope, deterministic seasonality and constant trading days effects, was the most visited, followed by M_{514} , which is as the same specification as before, but with a constant non zero slope. The next two modal models feature stochastic seasonality in the form of a fundamental stochastic cycle (model 641), also in conjunction to the first four harmonic cycles.

To obtain an estimate of the marginal probabilities that a particular component is present in series i we compute

$$\hat{\pi}_i = \frac{\sum_{j=1}^{10} c_{ij} \Upsilon_{ij}}{\sum_{j=1}^{10} c_{ij}};$$

this is actually the probability that each component is present in the first ten modal specifications that were recorded, as it amounts to summing up the values $c_{ij} / \sum_{j=1}^{10} c_{ij}$ over those specifications which contain the component, i.e. for which

Table 2: Estimation results for selected time series.
France, Index of industrial production

M_k	γ_μ	γ_A	γ_{s1}	γ_{s2}	γ_{s3}	γ_{s4}	γ_{s5}	γ_{s6}	γ_{TD}	δ	$100 \times \hat{\pi}_{ij}$
729	1	0	1	1	0	1	1	0	0	0	55.69
730	1	0	1	1	0	1	1	0	0	1	40.77
665	1	0	1	0	0	1	1	0	0	0	0.92

Germany, Retail turnover

M_k	γ_μ	γ_A	γ_{s1}	γ_{s2}	γ_{s3}	γ_{s4}	γ_{s5}	γ_{s6}	γ_{TD}	δ	$\hat{\pi}_{ij}$
641	1	0	1	0	0	0	0	0	0	0	58.21
642	1	0	1	0	0	0	0	0	0	1	29.88
657	1	0	1	0	0	1	0	0	0	0	7.18

UK, Retail turnover

M_k	γ_μ	γ_A	γ_{s1}	γ_{s2}	γ_{s3}	γ_{s4}	γ_{s5}	γ_{s6}	γ_{TD}	δ	$\hat{\pi}_{ij}$
497	0	1	1	1	1	1	0	0	0	0	15.66
881	1	1	0	1	1	1	0	0	0	0	7.88
498	0	1	1	1	1	1	0	0	0	1	7.38

the elements of Υ is one. However, as we have illustrated, the modal specifications absorb the quasi totality of the GS draws, so that the vector $\hat{\pi}_i$ can be thought of as an approximation to the true marginal probabilities given the data.

The distribution of the $\hat{\pi}_i, i = 1, \dots, 430$, can be effectively represented graphically using a biplot; see Gower and Hand (1996) and Greenacre (2010). The latter is a two-dimensional display that is based on the singular value decomposition of the matrix Π , obtained by stacking the row vectors $\hat{\pi}_i$.

In the graph, obtained using the BiplotGUI package, described in la Grange et al. (2009), the individual series are represented as points and the columns are represented as calibrated axes, as advocated by Gower and Hand (1996). Points are marked by a circle if they refer to the IPI series, whereas the RT series are marked by a square. The interpretation of the biplot is such that (the best rank 2 approximation to) the individual probabilities $\hat{\pi}_{ik}$ are obtained from the orthogonal projection of the point representing the series on the calibrated axis representing the k -th component. Moreover, the Euclidean distances among the points are an approximation to the Mahalanobis distances between the vectors $\hat{\pi}_i$, so that series that follow similar models are represented close in the plane. The calibrated axes are defined by the eigenvectors corresponding to the two largest eigenvalues of the covariance matrix of Π . The orientation of the axes can be gauged from the position of the labels. For instance slope and level span essentially the same subspace, which we can label the trend subspace, but move along opposite directions, which is a consequence of the negative correlation between the corresponding columns of the




Figure 3: Distribution of results in the model space.

matrix Π . Also, the trend subspace is almost orthogonal to the space spanned by the seasonal components.

The cluster of points in the lower left part of the graph refers to series for which the level is stochastic, the slope is fixed (the probability of selecting this component is well below the average, and seasonality is stable, i.e. the projection of those points along the Seas1-6 axes is low. This is the most numerous cluster. The points to the left will display stochastic seasonality as well. On the contrary, the set of point on the top left corner are characterised by low probabilities for stochastic level and seasonal cycles, but the probability that the slope is stochastic is high.

The vectors $\hat{\pi}_i$ can be further aggregated across the series to yield

$$\hat{\pi} = \sum_{i=1}^N \hat{\pi}_i w_i, \quad w_i = \sum_{j=1}^{10} c_{ij} / \sum_{i=1}^N \sum_{j=1}^{10} c_{ij}.$$

The elements of this vector are presented in table 3 for the 379 industrial production series (IPI), the 151 retail turnover series (RT) and for all the 530 series. The main result is that the marginal probability of having a stochastic level component is very high, and it is much higher for the industrial production series.




Figure 4: Principal components biplot of $\hat{\pi}_i, i = 1, \dots, 430,$. Circles represent industrial production series, and squares represent RT series. The orientation of the calibrated axes is provided by the position of the labels.

The probability of having a stochastic slope is higher for the retail series, and it is very low for the IPI series. Thus, we find that the trend component in the IPI series has a different characterisation than the RT series. The marginal probability of a stochastically time-varying seasonal trigonometric cycle is always less than 50% and tends to decrease with j , as order of the harmonic cycles increases. Time-varying trading day effects are never detected.

We conclude the presentation of the empirical results by discussing the joint frequencies by which given stochastic components are detected. In particular, we focus on the joint frequency distribution for the indicators γ_μ, γ_A and $I(\sum_j \gamma_j > 0)$; the latter is the indicator for the presence of at least one stochastic cycle at anyone of the seasonal frequencies. This is presented in table 4 for the complete dataset (Total,

530 series), the industrial production series (IPI, 379 series), the retail turnover series (RT, 151 series), the subset of very long time (VL) series, consisting of the US industrial production series (for which more than sixty years of data are available), the subset of long time series (L), featuring 195 series having more less than 30 but more that 18 years of data; (all the series except the US RT series belong to the IPI group), the subset of medium sized series (M), featuring more than 12 and no more than 18 years of data (149 series are in this group), and finally the subset of short time series, with at most 12 years of data (this subsets comprises 96 time series).

Table 3: Probability of identifying a particular component.

Component of Υ	IPI	RT	Total
Stochastic Level γ_μ	0.93	0.60	0.84
Stochastic Slope γ_A	0.18	0.43	0.25
Stochastic Seas1 γ_{s1}	0.39	0.46	0.41
Stochastic Seas2 γ_{s2}	0.42	0.47	0.43
Stochastic Seas3 γ_{s3}	0.41	0.44	0.42
Stochastic Seas4 γ_{s4}	0.34	0.47	0.37
Stochastic Seas5 γ_{s5}	0.24	0.34	0.27
Stochastic Seas6 γ_{s6}	0.18	0.23	0.19
Time-Varying Calendar γ_{TD}	0.00	0.00	0.00
Drift δ	0.34	0.38	0.35

The table reports the proportion of the MCMC draws referring to the 10 modal specifications that featured a particular combination of the indicators. A stochastic trend (either $\gamma_\mu = 1$ or $\gamma_A = 1$, or both) is detected in most occurrences; only in the case of short time series a completely deterministic trend was found in 4% of the draws. The modal representation in that case is (1,0,0), i.e. features a stochastic level, but deterministic slope and seasonals. If we consider the entire dataset, the modal representation (48.29%) features a stochastic level and at least one stochastic seasonal cycle. An interesting finding is that the frequency by which stochastic slope and seasonality are detected depends inversely on the sample size. The percentage of specifications featuring all three random components (corresponding to the triple (1,1,1) of the indicators) is 36.57 for the U.S. industrial production series, which are very long. This percentage decreases quite rapidly as the sample sizes decreases. The different results for the IPI and RT subsets may be the consequence of the different sample sizes of the series making up the two groups, the RT series being much shorter.

Table 4: Joint frequency distribution of the three indicators γ_μ, γ_A and $I(\sum_j \gamma_j > 0)$ for the complete dataset (Total), the industrial production series subset (IPI), the retail turnover series subset (RT), the subsets consisting of very long time series (VL), long time series (L), medium sized (M), and short time series (S).

γ_μ	γ_A	$I(\sum_j \gamma_j > 0)$	Total	IPI	RT	VL	L	M	S
0	0	0	0.26	0.10	0.68	0.00	0.00	0.00	1.52
0	0	1	0.69	0.33	1.61	0.00	0.25	0.64	2.68
0	1	0	2.35	1.03	5.69	0.00	0.59	2.45	8.17
0	1	1	12.74	5.25	31.69	7.79	7.57	22.54	11.56
1	0	0	25.77	27.74	20.80	13.98	18.08	31.63	41.89
1	0	1	48.29	54.02	33.78	36.89	69.57	39.68	32.87
1	1	0	0.99	1.23	0.37	4.77	0.02	0.22	0.49
1	1	1	8.91	10.30	5.38	36.57	3.92	2.84	0.80

7 Conclusions

We find strong support for the presence of a stochastic trend in the series, either in the form of a time-varying level, or, more rarely, of a stochastic slope, or both. We estimate the probability of detecting a stochastic trend close to 1. There is however a difference in the trend model for the industrial production series and the retail turnover, as for the latter a stochastic slope is more likely to be found, whereas for the former the slope is either fixed or zero in most of the cases.

Seasonality is a more stable component: only in 70% of the cases we were able to select at least one stochastic trigonometric cycle out of the six possible cycles. Most frequently the time variation is found in correspondence with the fundamental and the first harmonic frequencies. As we move to higher order harmonics, it is less probable to find time variation.

An interesting intuitive finding is that the probability of estimating time-varying components increases with the sample size available. This is particularly true of seasonality. However, even for very large sample sizes (such as those available for the US industrial production) we were unable to find stochastically varying calendar effects.

References

- Bell, W. R. and Hillmer, S. C. (1983). Modeling time series with calendar variation. *Journal of Business and Economic Statistics*, 78:526–534.
- Bell, W. R. and Martin, D. E. K. (2004). Modeling time-varying trading day effects in monthly time series. In *ASA Proceedings of the Joint Statistical Meetings: Alexandria*. American Statistical Association.
- Busetti, F. and Harvey, A. (2003). Seasonality tests. *Journal of Business and Economic Statistics*, 21:420–436.
- Canova, F. and Hansen, B. (1995). Are seasonal patterns constant over time? a test for seasonal stability. *Journal of Business and Economic Statistics*, 13:237–252.
- Dagum, E. and Quenneville, B. (1993). Dynamic linear models for time series components. *Journal of Econometrics*, 55:333–351.
- Dagum, E., Quenneville, B., and Sutradhar, B. (1993). Trading-day variations multiple regression models with random parameters. *International Statistical Review*, 60:57–73.
- De Jong, D. N. and Whiteman, C. H. (1991). The case for trend-stationarity is stronger than we thought. *Journal of Applied Econometrics*, 6:413 – 421.
- Dickey, D. and Fuller, W. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74:427–431.
- Durbin, J. and Koopman, S. (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika*, 89:603–615.
- Frühwirth-Schnatter, S. and Wagner, H. (2010). Stochastic model specification search for gaussian and partial non-gaussian state space models. *Journal of Econometrics*, 154(1):85–100.
- George, E. I. and McCulloch, R. (1993). Variable selection via gibbs sampling. *Journal of the American Statistical Association*, 88:881– 889.
- Geweke, J. (2005). *Contemporary Bayesian Econometrics and Statistics*. Wiley Series in Probability and Statistics.
- Ghysels, E. and Osborn, D. (2001). *The econometric analysis of seasonal time series*. Cambridge: Cambridge University Press.

- Harvey, A. (1989). *Forecasting, Structural Time Series and the Kalman Filter*. Cambridge: Cambridge University Press.
- Hylleberg, S. (1992). *Modelling Seasonality*. Oxford: Oxford University Press.
- Hylleberg, S., Engle, R. F., Granger, C. W. J., and Yoo, B. S. (1990). Seasonal integration and cointegration. *Journal of Econometrics*, 44:215–238.
- Hylleberg, S. and Pagan, A. (1997). Seasonal integration and the evolving seasonals model. *International Journal of Forecasting*, 13:329–340.
- Koop, G. (1992). “objective’ bayesian unit root tests. *Journal of Applied Econometrics*, 7:65 – 82.
- Koop, G. and van Dijk, H. (2000). Testing for integration using evolving trend and seasonals models: A bayesian approach. *Journal of Econometrics*, 97:261–291.
- Kwiatkowski, D., Phillips, P. C. B., Schmidt, P., and Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54:159 – 178.
- Nerlove, M., Grether, D. M., and Carvalho, J. L. (1979). *Analysis of economic time series: a synthesis*. New York: Academic Press.
- Nyblom, J. and Makelainen, T. (1983). Comparisons of tests for the presence of random walk coefficients in a simple linear model. *Journal of the American Statistical Association*, 78:865 – 864.
- Peña, D., Tiao, G., and Tsay, R. (2001). *A Course in Time Series Analysis*. New York: J. Wiley and Sons.
- Phillips, P. (1991). To criticize the critics: An objective bayesian analysis of stochastic trends. *Journal of Applied Econometrics*, 6:333 – 364.
- Phillips, P. and Perron, P. (1988). Posterior odds testing for a unit root with data-based model selection. *Biometrika*, 75:335 – 346.
- Phillips, P. and Perron, P. (1994). Testing for a unit root in time series regression. *Econometric Theory*, 10:774 – 808.
- Proietti, T. and Grassi, S. (2012). Bayesian stochastic model specification search for seasonal and calendar effects. In Harvey, A. C., Koopman, S. J., and Shephard, N., editors, *Economic Time Series: Modeling and Seasonality*. Chapman & Hall/CRC Press.

- Schotman, P. C. and van Dijk, H. (1991). On bayesian routes to unit roots. *Journal of Applied Econometrics*, 6:387 – 401.
- Sims, C. (1988). Bayesian skepticism on unit root econometrics. *Journal of Economic Dynamic and Control*, 12:463–474.
- Sims, C. and Uhlig, H. (1991). Understanding unit rooters: A helicopter tour. *Econometrica*, 59:1591 – 1599.
- Strickland, C. M., Martin, G., and Forbes, C. (2007). Parameterisation and efficient mcmc estimation of non-gaussian state space models. *Computational Statistical and Data Analysis*, 97(52):2911–2930.
- Wecker, W. and Ansley, C. (1983). The signal extraction approach to nonlinear regression and spline smoothing. *Journal of the American Statistical Association*, 78:81–89.
- West, M. and Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models*. New York, Springer-Verlag.
- Zellner, A. (1978). *Seasonal Analysis of Economic Time Series*. US Dept. of Commerce-Bureau of the Census.
- Zellner, A. (1983). *Applied Time Series Analysis of Economic Data*. US Dept. of Commerce-Bureau of the Census.

Research Papers 2011

- 2011-17: Søren Johansen and Theis Lange: Some econometric results for the Blanchard-Watson bubble model
- 2011-18: Tom Engsted and Thomas Q. Pedersen: Bias-correction in vector autoregressive models: A simulation study
- 2011-19: Kim Christensen, Roel Oomen and Mark Podolskij: Fact or friction: Jumps at ultra high frequency
- 2011-20: Charlotte Christiansen: Predicting Severe Simultaneous Recessions Using Yield Spreads as Leading Indicators
- 2011-21: Bent Jesper Christensen, Olaf Posch and Michel van der Wel: Estimating Dynamic Equilibrium Models using Macro and Financial Data
- 2011-22: Antonis Papapantoleon, John Schoenmakers and David Skovmand: Efficient and accurate log-Lévi approximations to Lévi driven LIBOR models
- 2011-23: Torben G. Andersen, Dobrislav Dobrev and Ernst Schaumburg: A Functional Filtering and Neighborhood Truncation Approach to Integrated Quarticity Estimation
- 2011-24: Cristina Amado and Timo Teräsvirta: Conditional Correlation Models of Autoregressive Conditional Heteroskedasticity with Nonstationary GARCH Equations
- 2011-25: Stephen T. Ziliak: Field Experiments in Economics: Comment on an article by Levitt and List
- 2011-26: Rasmus Tangsgaard Varneskov and Pierre Perron: Combining Long Memory and Level Shifts in Modeling and Forecasting of Persistent Time Series
- 2011-27: Anders Bredahl Kock and Timo Teräsvirta: Forecasting Macroeconomic Variables using Neural Network Models and Three Automated Model Selection Techniques
- 2011-28: Anders Bredahl Kock and Timo Teräsvirta: Forecasting performance of three automated modelling techniques during the economic crisis 2007-2009
- 2011-29: Yushu Li: Wavelet Based Outlier Correction for Power Controlled Turning Point Detection in Surveillance Systems
- 2011-30: Stefano Grassi and Tommaso Proietti: Stochastic trends and seasonality in economic time