The business cycle of European countries. Bayesian clustering of country-individual IP growth series.

Sylvia Kaufmann*†

May 2003

Abstract

In the present paper, time series on industrial production growth of individual countries are used to investigate the following questions: (i) Is there a common growth cycle for the euro area countries? (ii) Did the synchronization change over time? (iii) Can we discriminate between a "European" and an "overseas" cycle? (iv) Which countries follow the "overseas" rather than the "European" cycle? To obtain the inference, I use an autoregressive panel data framework whereby the groups of co-moving countries are estimated adaptively along with the model parameters using Bayesian simulation methods.

JEL classification: C15, C33, E32

Key words: Business cycle, Bayesian clustering, Markov switching, Markov chain Monte Carlo, panel data.

1 Introduction

The paper deals with the question whether the growth rate of industrial production (IP) has followed the same or a similar business cycle pattern in euro area countries and in all European countries. Moreover, the approach taken here allows to assess the relation of the European countries with transatlantic or "overseas" countries, in particular Australia, Canada, Japan and the United States. The focus lies on three different observation periods, a long-term historical perspective (1978-2001), a medium-term (1990-2001) and a short-term perspective (1999-2001), which reveals whether the synchronization of IP growth has changed among the countries in the course of increasing European integration.

The individual country series are analyzed within a panel data framework which enables us to enlarge the focus from the euro area towards a European versus overseas perspective or/and to restrict the observation period to the recent three years of the common European currency. The model specification allows for a time-varying constant that switches according to a latent state indicator which itself follows a Markov switching process. Countries that follow the same or a similar switching pattern are grouped together

*Oesterreichische Nationalbank, Economic Studies Division, P.O. Box 61, A-1010 Vienna, e-mail sylvia.kaufmann@oenb.co.at, phone +431 40420-7222, fax +431 40420-7299

†I would like to thank Melanie Groschan for assembling the dataset. The views expressed in the paper are those of the author and do not necessarily reflect the views of the OeNB.

whereby the groupings, i.e. which countries can be pooled together, are not set a priori. Rather, for a given number of groups, the classification of countries is estimated along with the model parameters and the latent state indicator.

The first contribution of the paper lies in the methodological approach. The presence of two latent indicator variables renders maximum likelihood methods infeasible and therefore, the estimation is cast into a Bayesian framework. The posterior distributions of the model parameters and the inference on the state- and the group-indicators are obtained with the use of Markov chain Monte Carlo (MCMC) simulations. The sampling scheme builds on the one proposed in Frühwirth-Schnatter and Kaufmann (2002a) and is extended here to group-specific switching state-indicators. The work is related to Artis et al. (1999), who estimate a common growth cycle for nine European countries by means of a Markov switching vector autoregressive (MS-VAR) model. From a methodological point of view, the advantage of the panel framework used in the present paper is that the number of countries entering the analysis may be increased without having to cope with an exponentially growing number of parameters to estimate. Moreover, while in Artis et al. (1999) all countries are pooled into one group and underly the switches simultaneously, here, potentially, different groups of countries are allowed for, that do not switch all at the same time.

The second contribution of the paper consists in the two pieces of evidence obtained from the results on the euro area and the European countries on the one hand and on the European versus the overseas countries on the other hand. When the euro area (and the European countries) are analyzed on their own, it turns out that in the long-term and in the medium-term perspective, they may be pooled together into one group. When the investigation includes all European and the overseas countries, we are able to identify a "European" and an "overseas" cycle whereby the latter is primarily determined by Australia, Canada and the US.[1] This is broadly consistent with previous literature that analyzed the synchronization and the correlation structure of European economies. Artis and Zhang (1997, 1999) find that the contemporaneous correlation between European countries has increased during the ERM period while at the same time the correlation with the US cycle has decreased. Forni and Reichlin (2001) analyze regional output fluctuations in 9 European countries and find out that the European component explains almost 50% of output growth in most regions and that it is highly correlated among the regions, which indicates a high degree of synchronization of output fluctuations in Europe. Finally, in a recent study, Mitchell and Mouratidis (2002) find that average correlation of various business cycle measures among the euro area countries has increased since the 1980s and continues to rise. The exception appears to be the United Kingdom, which, according to Artis and Zhang (1997), follows more closely the US rather than the German cycle and, according to Forni and Reichlin (2001), has a larger national than European component in output fluctuations. The present results are partly at odds with this evidence as the UK follows the European countries more closely than the overseas countries, in the long-term historical perspective (1978-2001) as well as in the medium-term perspective (1990-2001). In the recent short-term perspective (1999-2001), however, it turns out that the UK has been moving more closely with the US and Canada than before. The results also document the increasing synchronization among the euro area countries due to the integration process. Whereas over the long-term perspective Finland

[1] Over the long-term and the medium-term perspective, Japan is following the European cycle, whereas during the last three years, it moved more closely with Canada and the US.

and Ireland are following the overseas cycle, they follow the European cycle when the perspective is restricted to the last decade (1990-2001) and to the past three years.

An additional byproduct of the estimation is obtained by using the inference on the posterior state probabilities to date business cycle turning points. They follow quite closely the dates for the 1990s published in the Monthly Bulletin of the ECB (2002) and lag the ones identified by EuroCOIN (see Altissimo et al., 2001) by two to three quarters.

The following section introduces the model and the questions that we might wish to answer after the investigation, section 3 discusses briefly the estimation method. Section 4 first describes the data and presents the results for the euro area countries. The results for all European and the overseas countries are found in section 5, and section 6 concludes. For the interested reader, two appendices describe in detail the assumptions on the prior distribution of the model parameters and the sampling scheme, respectively.

2 The model and the hypotheses

Let y_{it} represent the (quarterly or monthly) growth rate at date t of industrial production for country i, computed by taking the first difference of the logarithmic level. We then write:

$$\Delta y_{it} = \mu_i^G + \mu_i^R(I_{it} - 1) + \phi_1 \Delta y_{i,t-1} + \cdots + \phi_p \Delta y_{i,t-p} + \varepsilon_{it}, \tag{1}$$

with $\varepsilon_{it} \sim$ i.i.d$N(0, \sigma^2)$, $t = 1, \ldots, T$. For a single country, the model we fit to industrial production comes close to the one estimated in Hamilton (1989) for US GNP. We assume that the growth rate $\mu_{it} = \mu_i^G + \mu_i^R(I_{it} - 1)$ depends on a latent state variable I_{it}, which may take on either the value 0 or 1:[2]

$$\mu_{it} = \begin{cases} \mu_i^G - \mu_i^R & \text{iff} \quad I_{it} = 0 \\ \mu_i^G & \text{iff} \quad I_{it} = 1 \end{cases} . \tag{2}$$

The latent specification of I_{it} takes into account the fact that the state prevailing in each period t is usually not observable with certainty. Moreover, as the periods with a higher growth rate might have a different duration than the periods of lower growth rate, we specify I_{it} to follow a Markov switching process of order one, $P(I_{it} = l | I_{i,t-1} = j) = \eta_{jl}^i$, with the restriction $\sum_{l=0}^{1} \eta_{jl}^i = 1$, $j = 0, 1$.

When the observation period is long enough, equation (1) might be estimated for each of the N investigated countries separately. The various processes for I_{it} might then be compared to assess the synchronization of the business cycles across the countries. However, the variance of estimation for I_{it} (and for the model parameters, too) might be reduced, if countries that switch at the same time, i.e. follow the same or a similar business cycle pattern, would be pooled in a group (see e.g. Hoogstrate et al., 2000, and Frühwirth-Schnatter and Kaufmann, 2002). Also, the gain in estimation precision would be greater the shorter the observation period. The difficulty in following this procedure is to form the appropriate grouping of countries. If we do not have a priori certain information about it, we might wish to draw an inference on the appropriate grouping characterizing

[2] For the empirical analysis, it proved sufficient to assume two states driving the process. In principle, however, the present parameterization would allow a three-state specification with I_{it} potentially taking on additionally the value 2. The estimation assuming a higher number of states is possible; then, however, a "direct" parameterization would be appropriate: $\mu_{i,I_{it}} = \mu_{i,j}$ iff $I_{it} = j$, $j = 1, \ldots, J$.

the countries included in the panel. To this aim, an additional latent group-indicator S_i, $i = 1, \dots, N$, is defined that relates to group-specific parameters, whereby S_i can take on one out of K different values, $S_i = k$, $k = 1, \dots, K$, if we assume to have K distinct groups of countries in the panel. The model for μ_{it} given in (2) may thus be extended to:

$$\mu_{it} = \begin{cases} \mu_k^G - \mu_k^R & \text{iff} \quad S_i = k \text{ and } I_{kt} = 0 \\ \mu_k^G & \text{iff} \quad S_i = k \text{ and } I_{kt} = 1 \end{cases}, \quad k = 1, \dots, K \tag{3}$$

whereby the probabilities $P(S_i = k)$ are given by η_k^G, $k = 1, \dots, K$ with the restriction $\sum_{k=1}^{K} \eta_k^G = 1$.

Note that in model (1), we assume that the parameters of the autoregressive process, $\phi_1, \dots, \phi_p$ are group- and state-independent. Although the analysis carries over to group- and state-dependent parameters in general, we do not introduce it here for expositional convenience and also because the results obtained from a preliminary investigation favoured a specification with group- and state-independent autoregressive parameters. Likewise, we assume that the variance of the error process is not group-specific, $\sigma_i^2 = \sigma^2$.

With the present model specification, we might investigate the following questions:

- Is there a common growth cycle for the euro area countries?
  If this is the case, the estimation should yield a pooled group of the euro area countries, all following the same switching pattern of the state indicator, i.e. K should equal at most 1 when the euro area countries are investigated on their own. When all other countries are included in the panel, the euro area countries should again pool into the same group, even if more than one group could be identified ($K > 1$).

- Did the synchronization change over time?
  The first dimension of synchronization relates to the estimated pattern of the state indicators, if more than one group can be identified in the panel. A changing lead-lag behaviour between the state indicators (over various time horizons) would document changing synchronization. The second dimension relates to the estimated country groupings. A change in synchronization would show up in a changing composition of the estimated country groups.

- Can we discriminate between a "European" and an "overseas" cycle?
  Evidence in favour of a distinction would be reflected in two distinctively estimated state-indicators, whereby most European countries would fall into one group.

- Which countries follow the "overseas" rather than the "European" cycle?
  Over different time horizons, it might be possible that a country switches between groups. This question is also related to the one in changes in synchronization.

3 Estimation via MCMC simulations

To briefly describe the estimation procedure, we introduce the following notation. While y_{it} denotes the observation of country i in period t, y_i^t gathers all observations of country i up to period t, $y_i^t = \{y_{it}, y_{i,t-1}, \dots, y_{i1}\}$, $i = 1, \dots, N$. The variables Y_t and Y^T will denote accordingly all country observations in and up to period t, respectively, $Y_t =$

$\{y_{1t}, y_{2,t}, \ldots, y_{Nt}\}$ and $Y^T = \{Y_T, Y_{T-1}, \ldots, Y_1\}$. Likewise, the vectors $S^N = (S_1, \ldots, S_N)$ and $I^T = (I_1^T, \ldots, I_K^T)$, where $I_k^T = (I_{kT}, I_{k,T-1}, \ldots, I_{k1})$, $k = 1, \ldots, K$, collect the group and the state indicators, respectively. Moreover, for notational convenience, θ will denote all model parameters[3] and $\psi = (\theta, S^N, I^T)$ will be the augmented parameter vector which includes additionally the two latent indicators.

Thus, the estimation of the model should not only yield an inference on the model parameters in θ, but also on the latent indicators S^N and I^T. If we knew S^N and I^T, roughly speaking we would be left with a regression model in (1) and standard methods (like GMM in the present case) could be used for estimation. Even if only one indicator were known, the estimation could be performed within the maximum likelihood framework, as the marginal likelihood $L(Y^T|\theta, S^N)$ or $L(Y^T|\theta, I^T)$ can alternatively be derived. However, if both indicators are unknown, the marginal likelihood $L(Y^T|\theta)$ is not feasible.[4] Therefore, the estimation is cast into a Bayesian framework and the posterior distribution of ψ is estimated using Markov chain Monte Carlo simulation methods based on an adapted version of the sampling scheme proposed in Frühwirth-Schnatter and Kaufmann (2002).

Note first, that for known values of S^N and I^T, the likelihood might be factorized in the following way:

$$L(Y^T|\psi) = \prod_{t=1}^{T}\prod_{i=1}^{N} f(y_{it}|y_i^{t-1}, \phi_1, \ldots, \phi_p, \sigma^2, I_{S_i,t}, S_i), \tag{4}$$

where $f(y_{it}|\cdot)$ denotes the density of the normal distribution:

$$f(y_{it}|y_i^{t-1}, \phi_1, \ldots, \phi_p, \sigma^2, I_{S_i,t}, S_i) = \\ \frac{1}{\sqrt{2\pi}\sigma} \exp\left\{-\frac{1}{2\sigma^2}\left(y_{it} - \mu_{S_i}^G - \mu_{S_i}^R(I_{S_i,t} - 1) - \sum_{j=1}^{p}\phi_j y_{i,t-j}\right)^2\right\}. \tag{5}$$

The densities of the K state indicators, $\pi(I_k^T|\eta^k)$, are mutually independent and depend a priori only on the transition distributions η^k, $k = 1, \ldots, K$. Thus,

$$\pi(I^T|\eta^1, \ldots, \eta^K) = \prod_{k=1}^{K}\pi(I_k^T|\eta^k) = \prod_{k=1}^{K}\prod_{j=0}^{1}(\eta_{jj}^k)^{N_{jj}^k}(1-\eta_{jj}^k)^{N_{1-j,j}^k}, \tag{6}$$

where $N_{jl}^k = \#(I_{kt} = l, I_{k,t-1} = j)$, $l, j = 0, 1$, is the number of times that state l followed state j for the indicator k, $k = 1, \ldots, K$. For the group indicator, in turn, we have

$$\pi(S^N|\eta_1^G, \ldots, \eta_K^G) = \prod_{i=1}^{N}\pi(S_i|\eta_1^G, \ldots, \eta_K^G) = \prod_{k=1}^{K}(\eta_k^G)^{(\#S_i=k)}. \tag{7}$$

The third layer within the Bayesian model setup comprises the specification of the prior distribution for the model parameter θ which, for the sake of brevity, is not describe here. The interested reader finds a detailed description of it in appendix A.

[3]That is: $\theta = (\mu_1^G, \ldots, \mu_K^G, \mu_1^R, \ldots, \mu_K^R, \phi_1, \ldots, \phi_p, \sigma^2, \eta_1^G, \eta_K^G, \eta^1, \ldots, \eta^K)$, where $\eta^k = (\eta_{00}^k, \eta_{01}^k, \eta_{10}^k, \eta_{11}^k)$, $k = 1, \ldots, K$.

[4]To derive $L(Y^T|\theta)$ we need to integrate out $L(Y^T|\theta, S^N)$ over S^N which itself depends on I^T. The same problem arises when we want to marginalize $L(Y^T|\theta, I^T)$. There is no way to circumvent the analytical problem.

The inference on the joint posterior distribution $\pi(\theta, S^N, I^T|Y^T)$ is obtained by successively simulating parameter values and values for the group and state indicators out of their conditional posterior distributions:

(i) $\pi(S^N|Y^T, \theta, I^T)$,

(ii) $\pi(I^T|Y^T, \theta, S^N)$,

(iii) $\pi(\theta|Y^T, S^N, I^T)$.

For given (sensible) starting values for θ and I^T, iterating several thousand times over the sampling steps (i)-(iii),[5] thereby replacing at each step the conditioning parameters by their respective actual simulated values, yields a sample out of the joint posterior distribution $\pi(\theta, S^N, I^T|Y^T)$. The simulated values may then be post-processed to estimate the properties of the posterior distribution, e.g. the mean and standard error may be inferred by computing the mean and the standard deviation of the simulated values. For practical implementation, step (iii) involves a further break-down of the parameter vector θ into appropriate sub-vectors for which the conditional posterior distributions can fully be derived and simulated straightforwardly. Appendix B describes in detail the sampling steps and derives the posterior distributions (i)-(iii).

4 The euro area countries

4.1 The data and the model selection procedure

All data are taken from the International Financial Statistics database. To perform the analysis for the long-term (1978-2001) and the medium-term perspective (1990-2001), we use quarterly data on seasonally adjusted industrial production of all (West-) European countries covering the period from the first quarter of 1978 to the last quarter of 2001. The recent period of the common European currency beginning with January 1999 and running through December 2001 is analyzed with monthly data, and to asses the relationship between the "European" and the "overseas" business cycle, we include additionally the industrial production series of Australia, Canada, Japan and the United States. For all countries, industrial production growth is computed by multiplying by 100 the first difference of the logarithmic level. Prior to the analysis, we standardize each time series to have unit standard deviation. Besides, the series undergo no other transformation, i.e. the series are included in the panel without weighting according to a country's size and without smoothing prior to the estimation, in contrast to Artis et al. (1999) and Mitchell and Mouratidis (2000). Some series, in particular Greece and Norway, display breaks which are accounted for by including dummy variables for the observations 1990Q3 and 1990Q4, and 1986Q2 and 1986Q3, respectively, in the estimation.

One advantage of using each country's industrial production series rather than a compiled euro area aggregate one, which usually weights each series relative to the entity under investigation, is that the country specific versus the regional/transatlantic optic may be compared directly. Also, the regional/transatlantic optic may be enlarged rather straightforwardly, without the need for reweighing each series. On the other hand, the standardization copes with the drawback one might perceive in the unweighed inclusion

[5] The programs are written in Matlab and the estimation time for the largest setting, i.e. when all countries are included in the panel, the observation period is 1978-2001 and $K = 3$, is approximately 15 minutes on a PC Pentium.

of each country's series in the panel. Although a small and a large country's performance enter with equal weights in the estimation of aggregate behaviour, the estimation yields an inference on groups of countries that display a common economic pattern over time that is not due to differences in business cycle volatility across countries. However, the estimation could also account for differences in volatility by assuming country-specific variances of the error process ($\sigma_i^2 \neq \sigma^2$). Here, the focus lies on the common pattern over time for groups of countries, so we abstain from this generalization.

Given (sensible) starting values for the model parameters and the state indicator, the model is estimated by iterating 10,000 times over the sampling steps described in detail in appendix B. The first 4,000 iterations are discarded to remove dependence from the starting values and the remaining simulations are used to infer the posterior distributions, e.g. the mean and the standard error may be estimated by the mean and the standard deviation of the simulated values or the values can be used to estimate the marginal probability distribution functions of the model parameters.

A preliminary estimation of model (1) over all time horizons for the panel of euro area countries included four autoregressive lags of industrial production and allowed for three groups, i.e. $p = 4$ and $K = 3$, respectively. It turned out that two autoregressive lags were sufficient as higher order lags were not significantly different from zero. Moreover, the estimation revealed that all countries may be pooled into one group, $K = 1$, over the long-term and the medium-term perspective and, hence, that one state indicator is enough to capture the common pattern of the euro area countries' industrial production. Over the short-term perspective, however, the countries may be classified into two groups. The selection of K is based on the fact that when allowing for more than one group, the sampler is not able to discriminate between them distinctively which is reflected in a uniform posterior group probability distribution for the countries and in equal group-specific parameter estimates.[6]

Table 1 summarizes the posterior inference for the parameters of interest of the chosen model specification for the euro area countries. Over all time horizons, the two states specification is significant as μ_1^R and μ_2^R are significantly different from zero. The two-groups specification for the short-term perspective seems to be a borderline case. The mean of μ_2^G is not included in the confidence interval for μ_1^G whereas it is marginally the case the other way round. Having an additional look a the posterior group probabilities (see figure 5) turned the balance in favour of $K = 2$ rather than $K = 1$.

4.2 The long-term historical perspective

To interpret the results of the one-group specification for the euro area countries, we have a look at the posterior state probabilities depicted in figure 1 which are estimated by taking the average over all simulated paths I^T. The figure reveals that $I_t = 0$ relates to periods of economic recessions at the beginning of the 1980s, the 1990s and during 2001, and reflects additionally periods of growth slowdown in the mid 1980s and in the course of 1990s. The regime switches are quite distinct and, due to the fact of capturing recessions and slowdowns, nearly equally persistent for periods of economic recovery and for periods

[6]For example, when K is set $K = 3$ whereas in fact it should be $K = 1$, the estimated posterior group probabilities for each country would be approximately 1/3 for each group and the group-specific parameters μ_1^G, μ_2^G and μ_3^G would not significantly be different from each other, i.e. their confidence interval would largely overlap.

Table 1: Mean group-specific parameter estimates of the euro area model (with confidence interval). The confidence intervals are estimated by the shortest interval containing 95% of the 6,000 simulated parameter values. η_{00} and η_{11} refer to the persistence of state $I_t = 0$ and $I_t = 1$, respectively.

	$I_t = 1$		$I_t = 0$			
	μ_1^G	μ_2^G	$\mu_1^G - \mu_1^R$	$\mu_2^G - \mu_2^R$	μ_1^R	μ_2^R
1978Q1-2001Q4	0.70		-0.09		0.79	
	(0.57 0.82)		(-0.19 0.02)		(0.67 0.92)	
η_{00}	0.79					
conf.int.	(0.65 0.92)					
η_{11}	0.77					
conf.int.	(0.63 0.92)					
1990Q1-2001Q4	0.66		-0.21		0.87	
	(0.52 0.81)		(-0.36 -0.08)		(0.70 1.02)	
η_{00}	0.77					
conf.int.	(0.58 0.94)					
η_{11}	0.79					
conf.int.	(0.62 0.95)					
1999M1-2001M12	0.38	1.16	-0.63	-0.30	1.01	1.46
	(0.02 0.74)	(0.32 3.03)	(-2.14 0.12)	(-1.48 1.80)	(0.03 2.19)	(0.05 3.34)
no. of countries	5	7				
η_{00}	0.61	0.74				
conf.int.	(0.22 0.99)	(0.37 1.00)				
η_{11}	0.74	0.67				
conf.int.	(0.37 1.00)	(0.27 1.00)				

Figure 1: Posterior probability $P(I_t = 0|Y^T, S^N, \theta)$ for the grouped euro area countries.

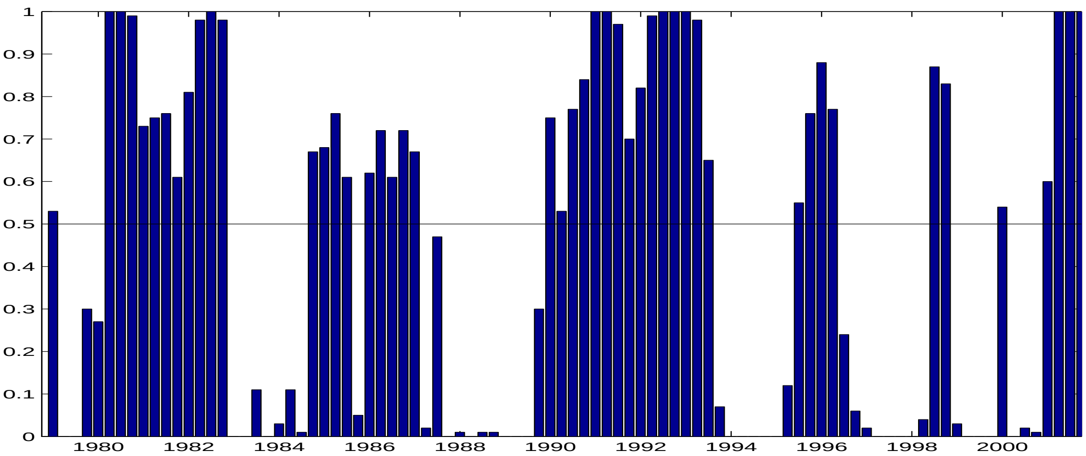

Figure 2: State specific mean growth rates of the euro area countries.

(a) 1978Q1-2001Q4 (b) 1990Q1-2001Q4

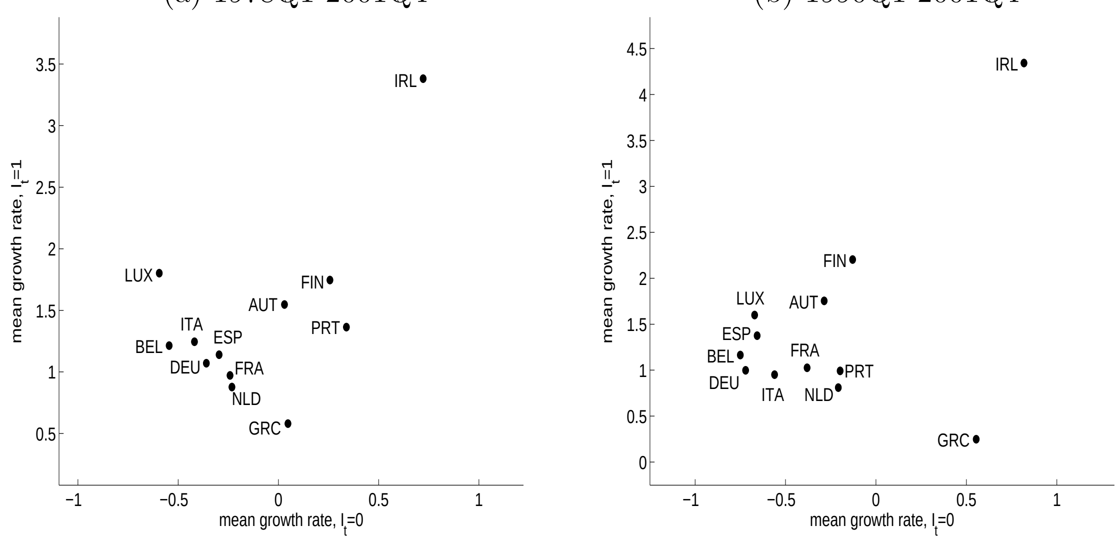

of economic slowdown. Table 1 documents that the mean persistence of the states $I_t = 0$ and $I_t = 1$, η_{00} and η_{11}, is 0.79 and 0.77, respectively. Note that due to the standardization, the coefficients μ_1^G and $\mu_1^G - \mu_1^R$ are not directly interpretable as quarterly growth rates. However, based on the estimate of the common state indicator, one might compute for each country the average state-dependent quarterly growth rates over the observation period. These are depicted in the left-hand scatter plot of figure 2. Although the countries followed a common economic pattern over the last two decades, there have been some remarkable differences in their growth performance. Ireland, followed by Finland, Austria and Portugal, represents the country experiencing the most pronounced catching up process. On the other hand, the three largest euro area countries, Germany, France and Italy, along with Spain, Belgium and the Netherlands displayed a more traditional growth pattern over the business cycle, around +4% a year in economic expansion and between -1% and -2% in slowdown periods.

Table 2: Business cycle dating of the euro area countries. The bottom lines reproduce the dating of the European Central Bank (2002) based on the euro area GDP cycle extracted by the band-pass filter of Baxter and King (1999) and the dating based on the EuroCOIN indicator (see Altissimo et al., 2001, and www.cepr.org)

sample period	P	T	P	T	P	T	P	T	P	T	P	T	P
1978-2001	80:1	82:4	84:3	87:1	89:4			93:3	95:2	96:2	98:2	98:4	00:4
1990-2001						91:3	91:4	93:2	95:2	96:2	98:2	98:4	00:4
ECB's dating							92:1	93:3	95:1	97:1	98:1	99:1	00:3
EuroCOIN					89:1			92:4	94:4	95:4	97:4	98:4	99:4

The posterior state probabilities depicted in figure 1 might be used to date the business cycle turning points for the euro area, defining a peak (P) the quarter t if $P(I_t = 0|Y^T, \theta, S^N) < 0.5$ and $P(I_t + 1 = 0|Y^T, \theta, S^N) > 0.5$, $P(I_t + 2 = 0|Y^T, \theta, S^N) > 0.5$, and correspondingly the quarter t to be a trough (T) if $P(I_t = 0|Y^T, \theta, S^N) > 0.5$ and $P(I_t + 1 = 0|Y^T, \theta, S^N) < 0.5$, $P(I_t + 2 = 0|Y^T, \theta, S^N) < 0.5$. The first line in table 2 summarizes the turning points. The first line in the bottom panel includes the dating of the European Central Bank (ECB, 2002) obtained by extracting the euro area wide GDP cycle for the 1990s from the Baxter and King (1999) band-pass filter. Both dating series are quite in accordance with each other, in general peaks and troughs identified with industrial production lag and precede the ones identified with euro area GDP by one quarter, respectively, with the exception of the trough in 1996:2 which leads the ECB's dating by three quarters. The second line in the bottom panel reproduces the turning points identified with EuroCOIN (Altissimo et al., 2001). The three full cycles during the 1990s broadly comove, EuroCOIN leads the state indicator by two to three quarters, however. Note finally, that the most recent downturn is identified by EuroCOIN nearly a year before the ECB and the present state indicator identify it.

The relative position of each country to the common cycle identified in figure 1 and table 2 can be assessed by estimating the model for each country separately. The posterior state probabilities estimated by a univariate (country-specific) Markov switching specification are depicted in figure 3 where the shaded areas refer to the common cycle. First of all, note that two significantly different states are identified for Germany and Finland only. For all other countries, the two-states specification does not seem as significant as for these two countries. In particular for Greece and for Ireland, the state indicator seems to relate more to long-term growth rather than to business cycle periods. As seen before, the catching up process has been especially strong for Ireland. The figure, however, documents the gain obtainable from pooling all countries' series. Despite the fact that generally two separate regimes cannot clearly be discriminated on an individual country basis, the information contained in the pooled data series is valuable to do so.

4.3 The nineties

To get an insight on potential changes during the 1990s we restrict the sample period to run from the first quarter of 1990 to the last quarter of 2001. As already mentioned in subsection 4.1, the countries behave similar enough to be pooled into one group. There-

Figure 3: Posterior probabilities $P(I_t = 0|Y^T, S^N, \theta)$ estimated for each country separately. The shaded area refers to the common periods of economic slowdown.

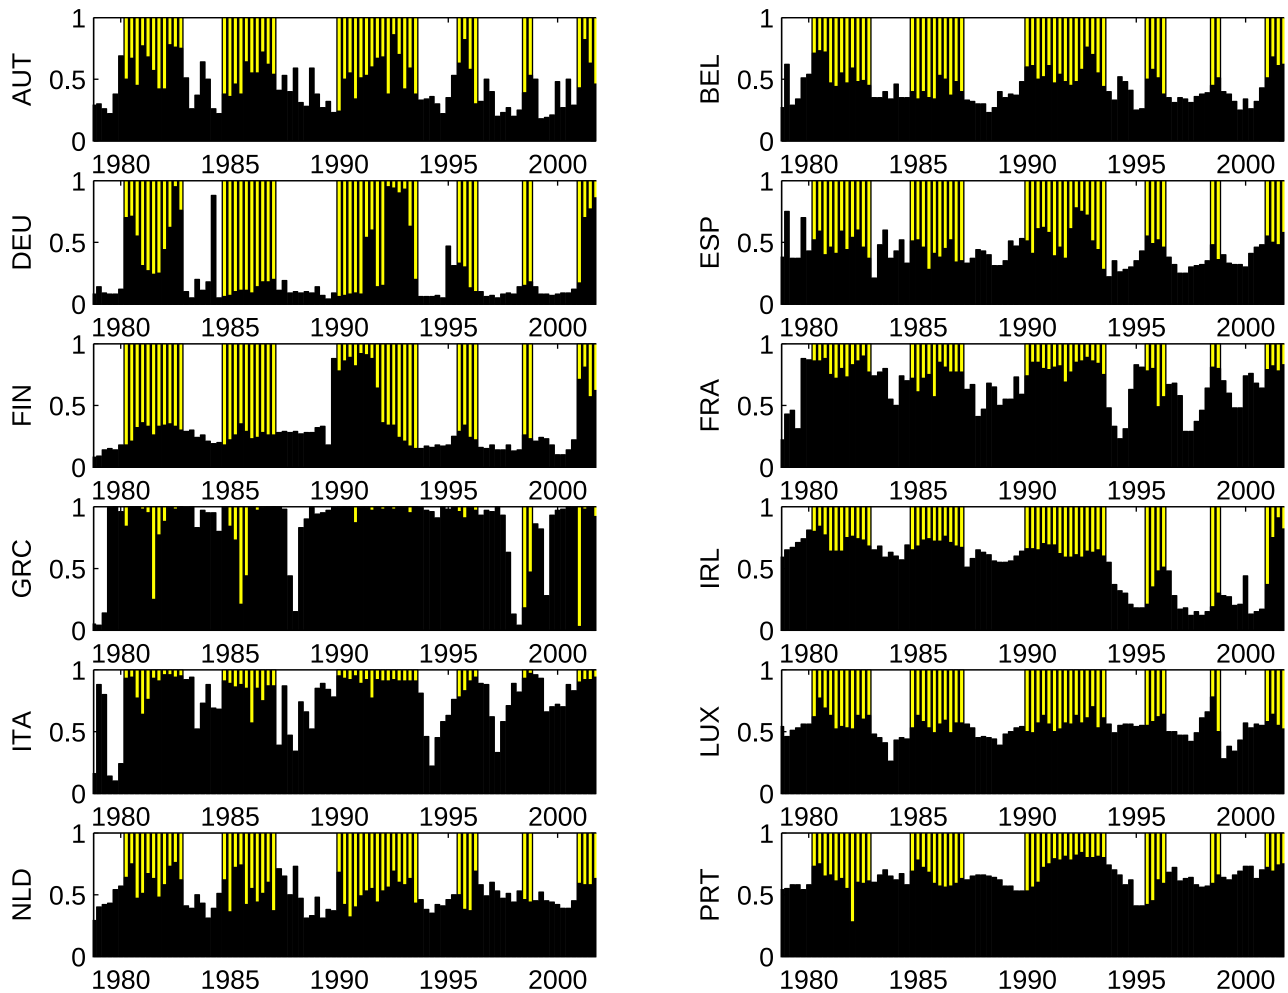

fore, the specification remains the same as over the long-term perspective, $K = 1$ and $p = 2$. The pattern of the posterior state probabilities (see figure 4) reveals that the euro area countries experienced three peak-to-peak cycles during the 1990s, the turning points of which are given on the second line of table 2. Interestingly, although the observation sample has been restricted, the identification of the turning points appears quite robust and remains in strong accordance with the ECB's dating. Again, peaks and troughs identified with industrial production growth lag and precede the ones identified with band-pass filter detrended GDP.

As it is the case for the long-term perspective, the common cyclical pattern in fact hides the still diverging growth performance among the euro area countries. Figure 2, panel (b), depicts the average country- and state-specific growth rates during the 1990s. Ireland still experienced a strong growth period, with average quarterly growth rates above 4% and nearly 1% when $I_t = 1$ and $I_t = 0$, respectively. Finland and Austria again follow with higher growth rates than the other countries, both countries come closer to the countries with a traditional positive and negative growth rate business cycle pattern,

Figure 4: Posterior probability $P(I_t = 0|Y^T, S^N, \theta)$ for the grouped euro area countries, sample period 1990Q1-2001Q4.

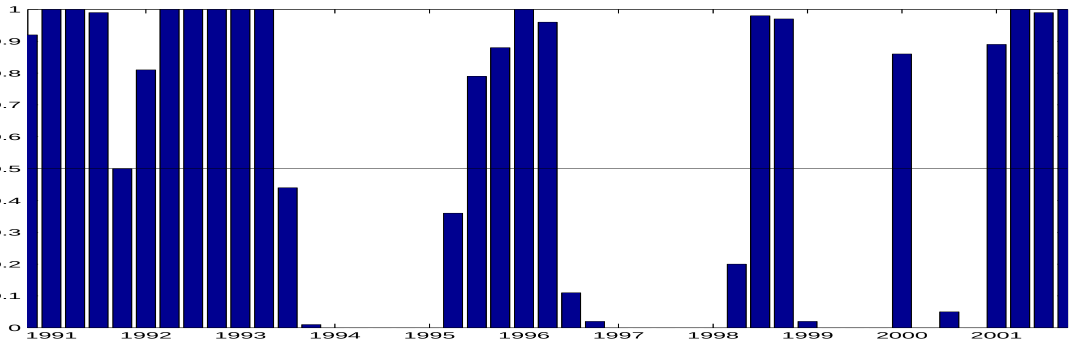

however. Greece, experiencing only a slow catching up process, builds the other exception of the euro area group.

4.4 The recent past

To estimate model (1) for the period during which the single common European currency was effectively in place, we use monthly data on industrial production covering the years from 1999 through 2001. The preferred model specification sets $K = 2$ and $p = 2$. According to the posterior group probabilities depicted in figure 5 (estimated by averaging over all simulated S^N), the first group of countries consists of Finland, Ireland, Luxembourg, the Netherlands and Portugal, while the largest euro area countries along with Austria, Belgium and Greece fall into the second group. The pattern of the posterior group-specific state probabilities in figure 6 and the mean state persistence, η_{00} and η_{11}, recorded in table 1 suggest that the countries of the first group experienced a longer recovery period through 2000 than the largest euro area countries, before the economic slowdown taking its pace in 2001 finally affected all countries. However, with the exception of Ireland and Finland, which both recorded relatively strong positive and negative growth rates, the growth performance of the first group's countries lies in the range of the largest euro area countries (see figure 7).

Figure 5: Posterior group probability $P(S_i = k|Y^T, I^T, \theta)$ for each euro area country, sample period 1999M1-2001M12.

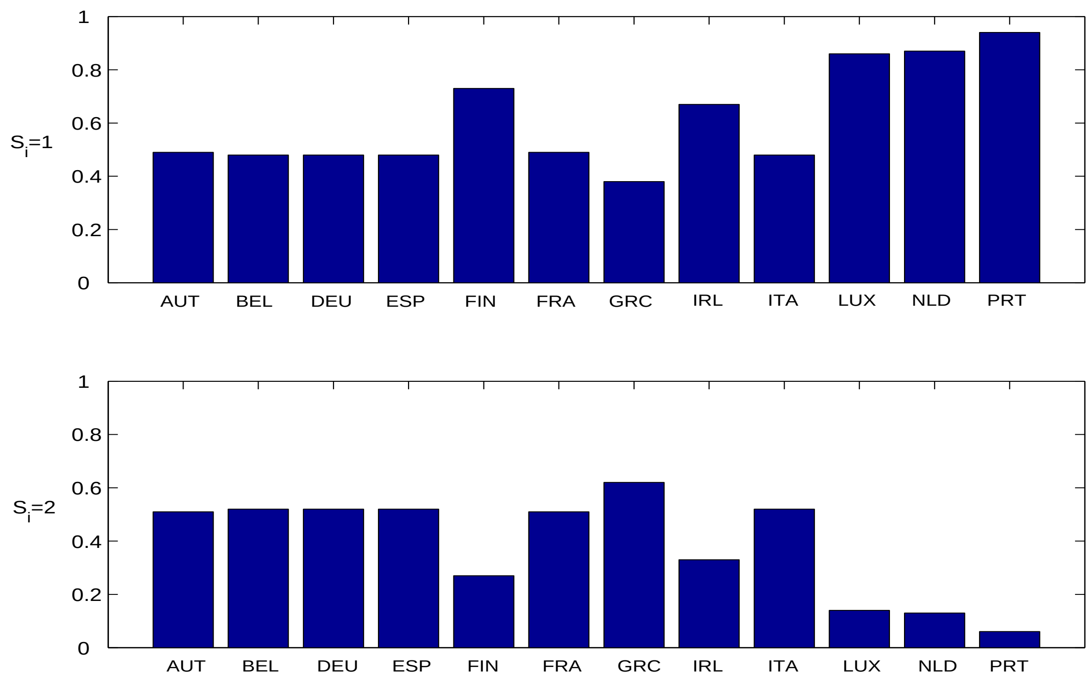

Figure 6: Posterior probability $P(I_t = 0|Y^T, S^N, \theta)$ for the grouped euro area countries, sample period 1999M1-2001M12.

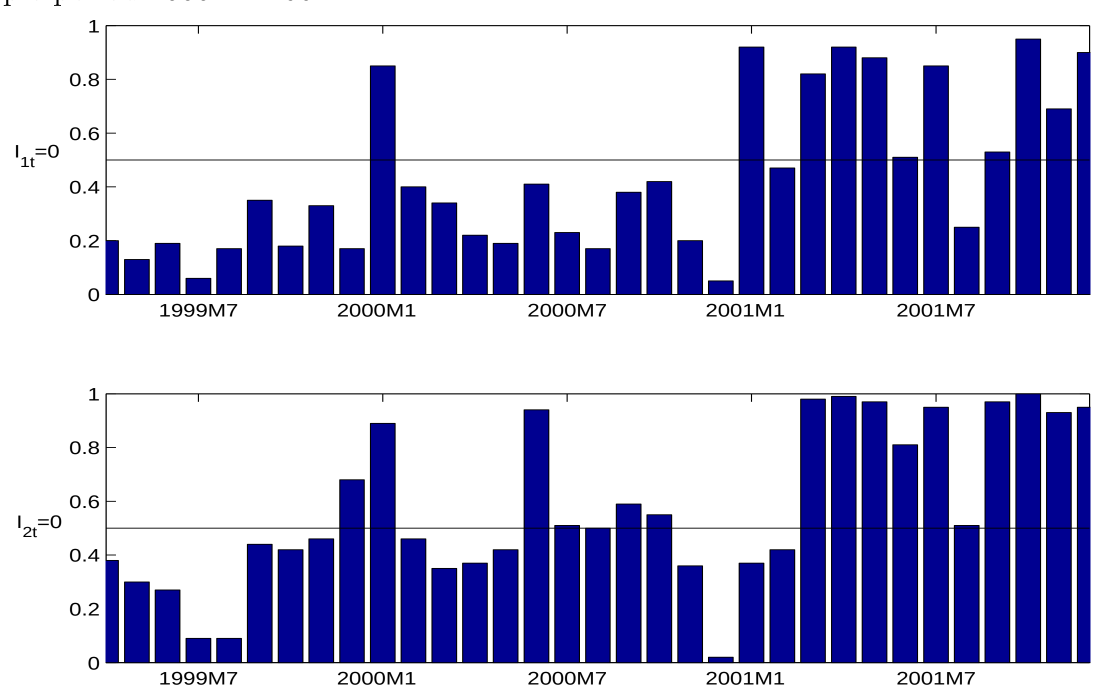

Figure 7: State specific (monthly) mean growth rates of the euro area countries. The first group is identified by the dotted marker, the second group by the squared marker.

(c) 1999M1-2001M12

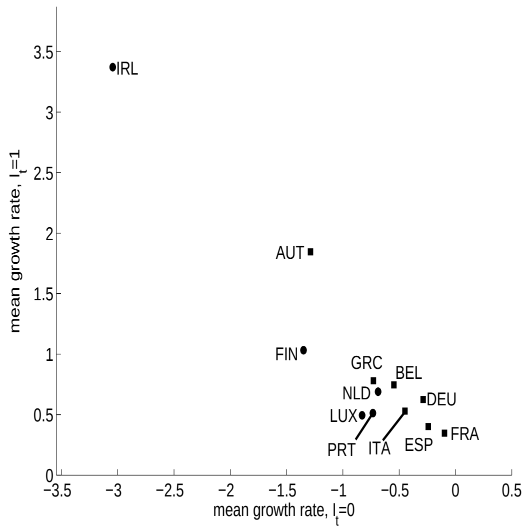

5 The "European" and the "overseas" cycles

Before investigating the business cycle pattern of the European countries in relation to major overseas countries, i.e. Australia, Canada, Japan and the United States, it is worth mentioning that performing the analysis for all West-European countries, EU countries plus Norway and Switzerland, revealed that the main result, that countries may be pooled for the long term and the medium term analysis and may be classified into two groups in the short term perspective, turned out to be robust.[7] Moreover, under the European long-term perspective, the estimated posterior state probabilities identified more distinctively the recessionary periods at the beginning of the 1980s, the 1990s and in 2001, without including the slowdown periods in the middle of the 1980s and 1990s as before (see table 3).[8]

The estimation of model (1) for 17 European and 4 overseas countries identified for all time horizons two groups with distinct business cycle timing patterns whereby a preliminary analysis also confirmed that two autoregressive lags were sufficient to model the data. Table 4 summarizes the estimation for the group- and state-specific parameters. Over all time horizons, the two states are significant in each group and also significantly different from each other. Moreover, in particular over the long- and the medium-term perspective, the mean persistence of $I_t = 1$ is now higher than the mean persistence of

[7]Under this setting, the two nordic countries, Norway and Sweden, and Greece joined Finland, Ireland, Luxembourg, the Netherland and Portugal (the countries defining the first group in the euro area setting) in the first group, while Denmark and the UK were classified into the second group with the largest euro area countries along with Austria and Belgium.

[8]As the results, overall, do not differ significantly from the previous ones, we do not display them here in order to save space.

Table 3: Business cycle dating of European countries in comparison with euro area dating.

sample period	P	T	P	T	P	T	P	T	P	T	P	T	P	T	P
European countries															
1978-2001	80:1	80:4	81:4	82:4			90:4	91:2	92:1	93:2					01:1
1990-2001							90:3	91:3	92:1	93:2	95:3	96:2	98:2	98:4	00:4
Euro area															
1978-2001	80:1			82:4	84:3	87:1	89:4			93:3	95:2	96:2	98:2	98:4	00:4
1990-2001								91:3	91:4	93:2	95:2	96:2	98:2	98:4	00:4

$I_t = 0$. The changing number of countries falling into the two groups[9] already hints towards changing business cycle synchronization among the countries.

5.1 The long term historical perspective

The characterization of the two groups is obtained from figure 9 which depicts the posterior group probabilities of the countries. The classification is very clear with all group probabilities being above 0.9 (0.8 for Finland) for one of the two groups. The countries falling into the second group over all time horizons, in particular Canada and the United States along with Australia, will determine what we call the "overseas" cycle. Under the long-term perspective, Finland and Ireland are following more closely the overseas rather than the "European" cycle, defined by the rest of the European countries falling into the first group. This might reflect overseas oriented trade relations of these two countries. On the other hand, the UK and Japan clearly fall into the group of European countries, while the situation will be the opposite in the short term perspective. The posterior group-specific state probabilities are graphed in figure 8. It nicely depicts that the state $I_t = 0$ relates to periods of recessionary tendencies, whereby until the mid 1990s the countries in the second group, following the "overseas" cycle, lead the countries following the "European" cycle by one quarter to over half a year (see the corresponding dates of the business cycle turning points given in table 5). Due to the lasting recovery experienced by the countries of the second group, the leading pattern has disappeared during the 1990s.

Figure 10, panel (a), depicts the mean state-specific growth rate of each country. From there we can see that Ireland and Finland have been growing quite strongly over the last two decades, whereas Canada, Australia and the United States have not been growing (much) faster than the largest European countries during periods of economic recovery. Canada and the United States even recorded stronger negative growth rates during periods of economic slowdown than most European countries did; but these periods turned out to be less frequent, however, in particular during the 1990s.

[9]No monthly industrial production series are available for Switzerland and Australia, therefore they are excluded from the panel over the short-term perspective.

Table 4: Mean group-specific parameter estimates of the grouped European and overseas countries (with confidence interval). The confidence intervals are estimated by the shortest 95% interval of the 6,000 simulated parameter values. η_{00} and η_{11} refer to the persistence of state $I_t = 0$ and $I_t = 1$, respectively.

	$I_t = 1$		$I_t = 0$			
	μ_1^G	μ_2^G	$\mu_1^G - \mu_1^R$	$\mu_2^G - \mu_2^R$	μ_1^R	μ_2^R
1978Q1-2001Q4	0.50	0.70	-0.33	-0.39	0.83	1.09
	(0.40 0.60)	(0.56 0.84)	(-0.50 -0.15)	(-0.70 -0.08)	(0.70 0.98)	(0.75 1.38)
no. of countries	16	5				
η_{00}	0.71	0.79				
conf.int.	(0.52 0.89)	(0.61 0.95)				
η_{11}	0.87	0.90				
conf.int.	(0.76 0.96)	(0.82 0.98)				
1990Q1-2001Q4	0.53	0.69	-0.24	-0.83	0.77	1.51
	(0.38 0.68)	(0.46 0.92)	(-0.50 -0.04)	(-1.45 -0.13)	(0.57 1.00)	(0.77 2.21)
no. of countries	18	3				
η_{00}	0.74	0.76				
conf.int.	(0.53 0.95)	(0.51 1.00)				
η_{11}	0.80	0.93				
conf.int.	(0.61 0.98)	(0.79 1.00)				
1999M1-2001M12	0.43	0.81	-0.29	-0.96	0.72	1.77
	(0.32 0.57)	(0.57 1.09)	(-0.46 -0.13)	(-1.21 -0.72)	(0.55 0.89)	(1.42 2.17)
no. of countries	15	4				
η_{00}	0.68	0.94				
conf.int.	(0.39 0.95)	(0.83 1.00)				
η_{11}	0.76	0.90				
conf.int.	(0.56 0.94)	(0.77 1.00)				

Table 5: Business cycle dating of the grouped European and overseas countries, sample period 1978Q1-2001Q4: The first group consists of the European countries (except Finland and Ireland), Japan and Australia, the second one of the United States, Canada and Finland and Ireland.

	P	T	P	T	P	T	P	T	P	T	P
"European" cycle	80:1	80:4	81:4	82:4	86:3	87:1	90:4	93:2	98:2	98:4	00:4
"Overseas" cycle	79:3	80:2	81:3	82:4	85:4	86:2	89:3	91:2			00:4

Figure 8: Posterior probability $P(I_t = 0|Y^T, S^N, \theta)$ for the "European" and the "overseas" state indicator , sample period 1978Q1-2001Q4.

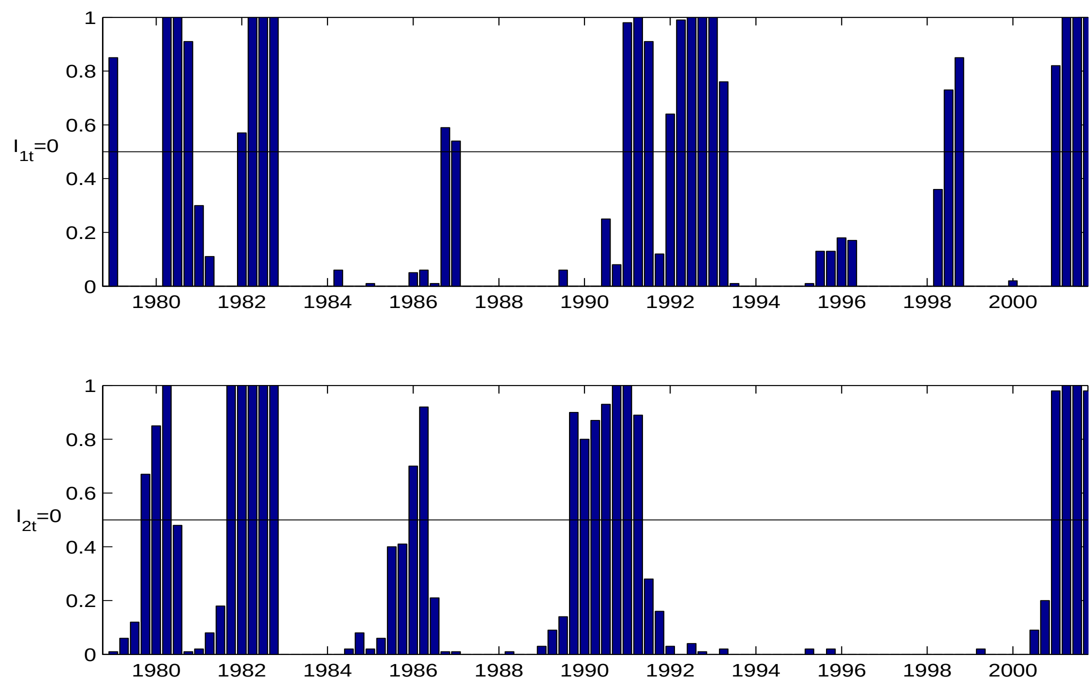

5.2 The nineties

Restricting the sample period to 1990Q1-2001Q4 yields a change in the country classification. During this decade, Ireland and Finland followed more closely the "European" cycle (see figure 11). Note that, still, Japan and the United Kingdom fall again into the first group. This reflects the convergence process that took place among the European countries in the course of increased economic and financial integration. The dating of the cycles in table 6 and the group-specific state probabilities in figure 12 reveal the already identified three full cycles for the euro area and all European countries, and the strong growth period of the overseas countries during the 1990s. Interestingly, the downturn that affected all countries in 2001 is identified to have already begun in the second half of 2000 for the overseas countries. Figure 10, panel (b) shows that the Nordic countries Finland and Sweden along with Ireland were the countries with the strongest growth during the recovery periods. Note that also Norway has been growing throughout the sample period. It is again the case that Canada and the US have not been growing much faster than the

Figure 9: Posterior group probability $P(S_i = k|Y^T, I^T, \theta)$ for each country, sample period 1978Q1-2001Q4.

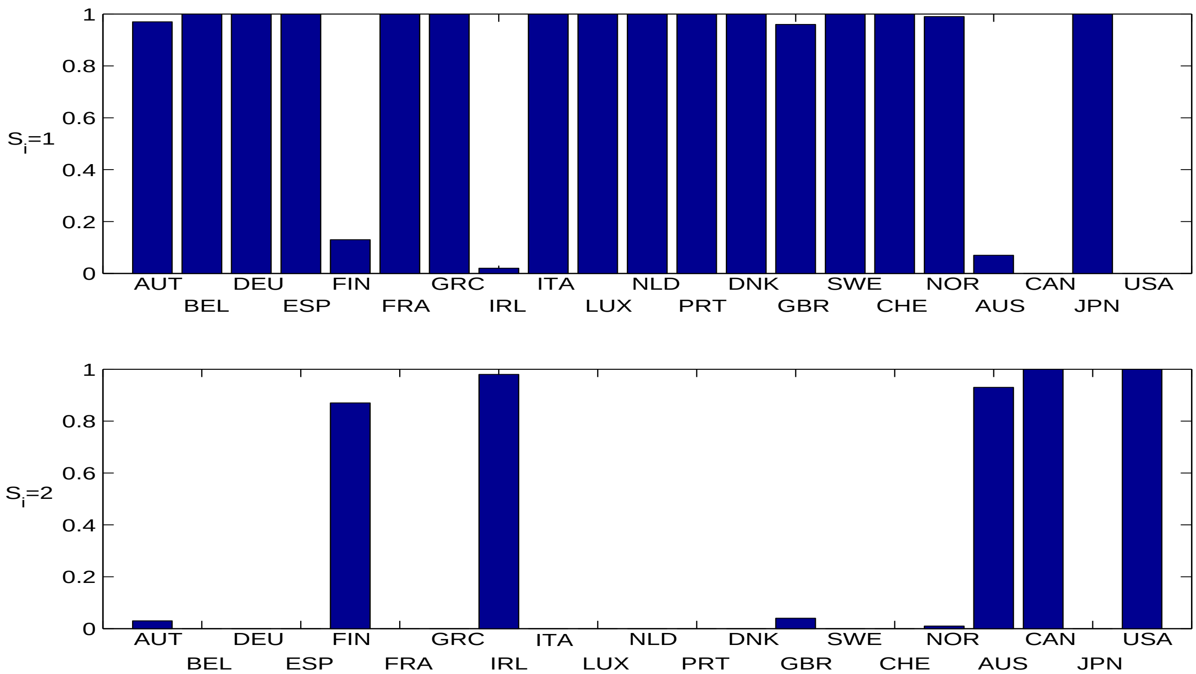

three largest European countries. However, their relative better economic performance was achieved by the long-lasting recovery throughout the 1990s.

Table 6: Business cycle dating of the grouped European and overseas countries: The first group consists of the European countries and Japan, the second one of the United States, Canada and Australia.

	P	T	P	T	P	T	P	T	P
"European" cycle	90:4	91:3	92:1	93:2	95:2	96:2	98:1	98:4	00:4
"Overseas" cycle		91:1							00:3

5.3 The recent past

Finally, the observation period is restricted to the years 1999 through 2001. According to figure 14, the classification into two groups is again very distinct and reveals that during the last three years, Japan and the UK followed the "overseas" cycle more closely than the "European" cycle. Moreover, the posterior group-specific state probabilities in figure 13 show that the overseas countries experienced a downturn in economic activity already in the second half of 2000, the European countries being affected later at the beginning of 2001. Finally, figure 15 reveals that Ireland and Finland, along with Denmark, represent the countries with the strongest growth rates during the first years of the common European currency. Canada, the US and the UK, on the other hand, experienced less

Figure 10: State specific mean growth rates of the European and the overseas countries.

(a) 1978Q1-2001Q4 (b) 1990Q1-2001Q4

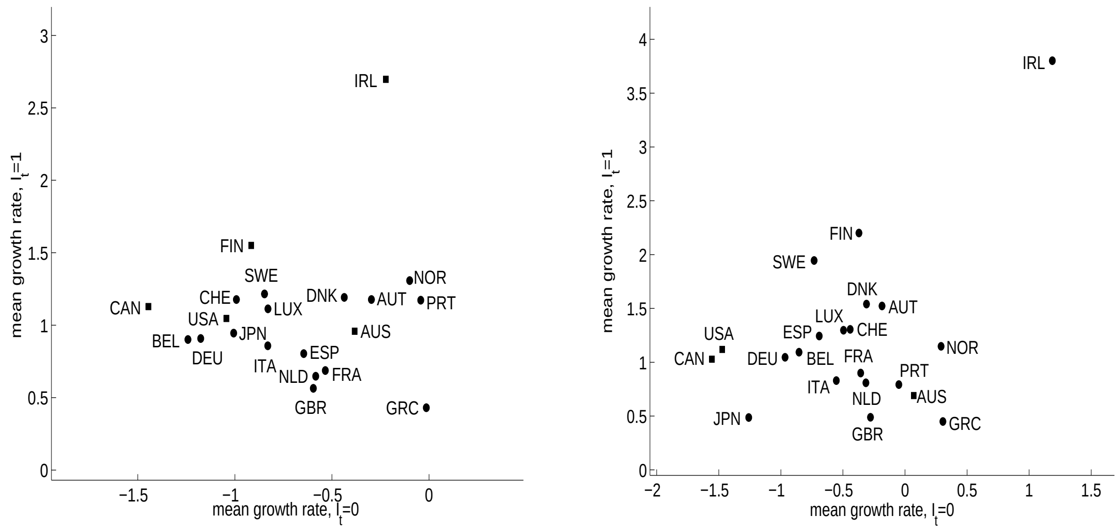

decline in industrial production than most euro area countries during the recent downturn period.

Figure 11: Posterior group probability $P(S_i = k|Y^T, I^T, \theta)$ for each country, sample period 1990Q1-2001Q4.

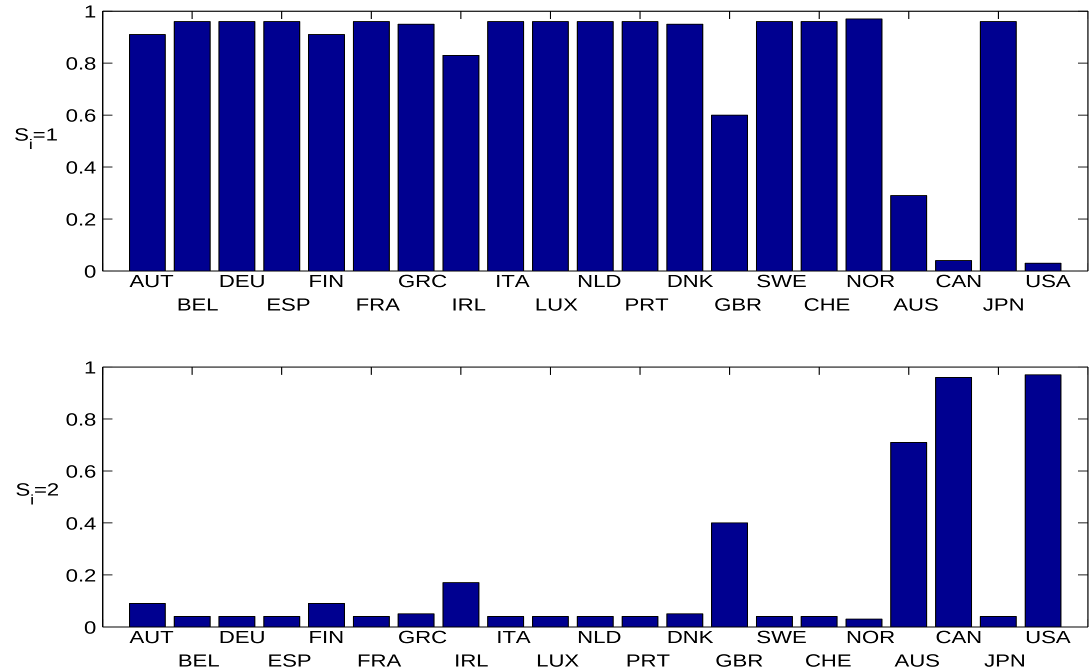

Figure 12: Posterior probability $P(I_t = 0|Y^T, S^N, \theta)$ for the "European" and the "overseas" state indicator , sample period 1990Q1-2001Q4.

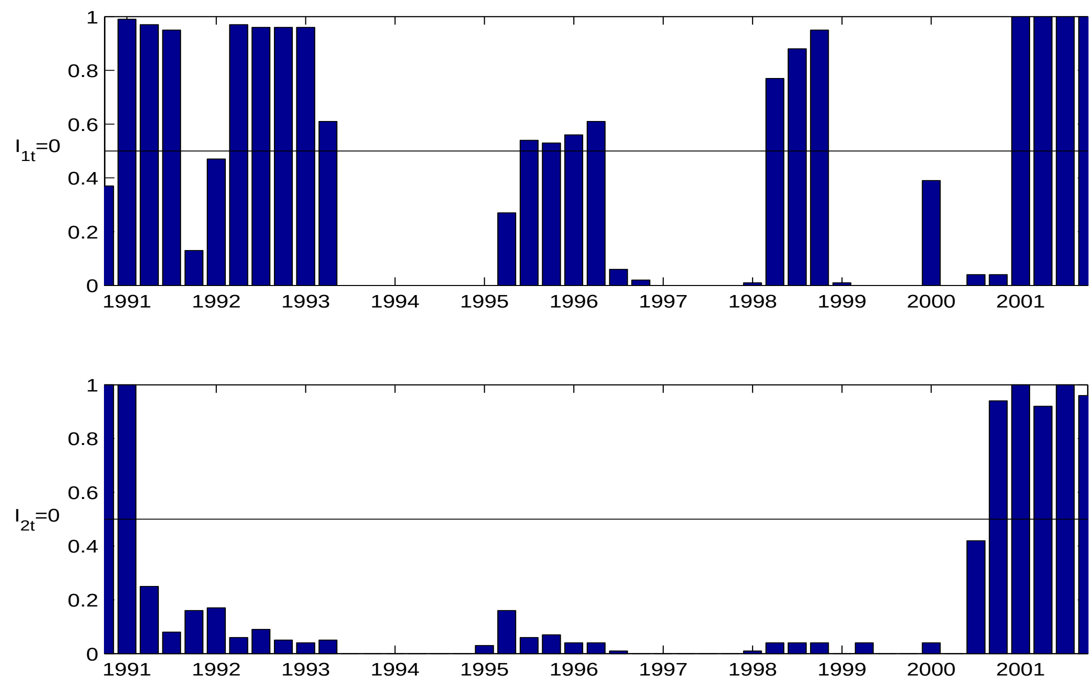

Figure 13: Posterior probability $P(I_t = 0|Y^T, S^N, \theta)$ for the "European" and the "overseas" state indicator , sample period 1999M1-2001M12.

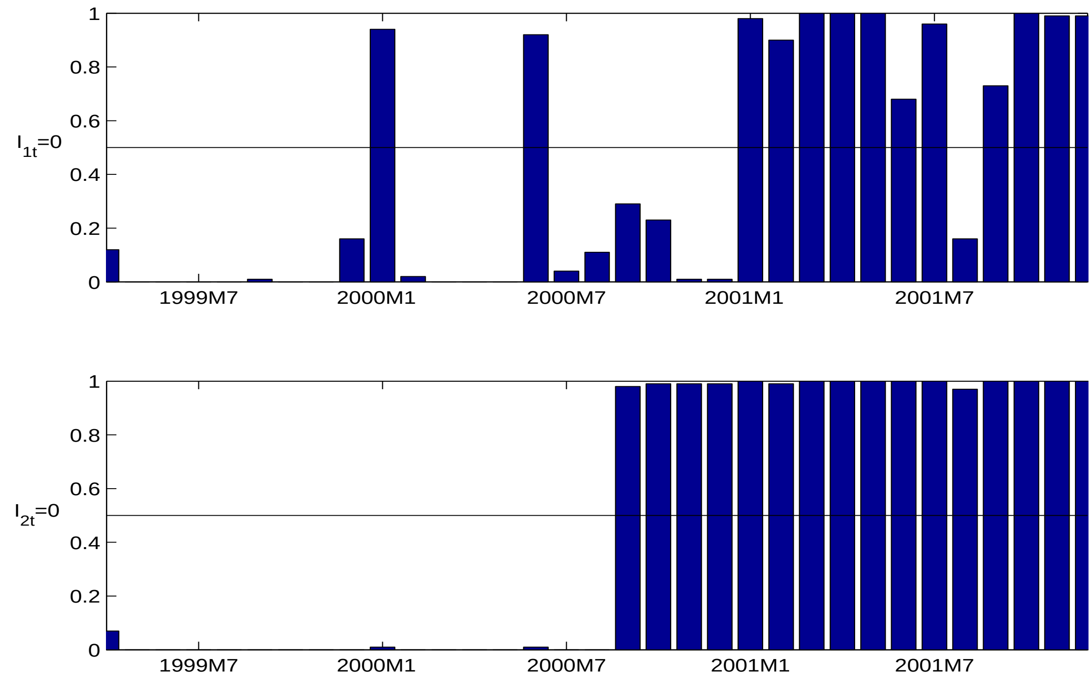

Figure 14: Posterior group probability $P(S_i = k|Y^T, I^T, \theta)$ for each country, sample period 1999M1-2001M12.

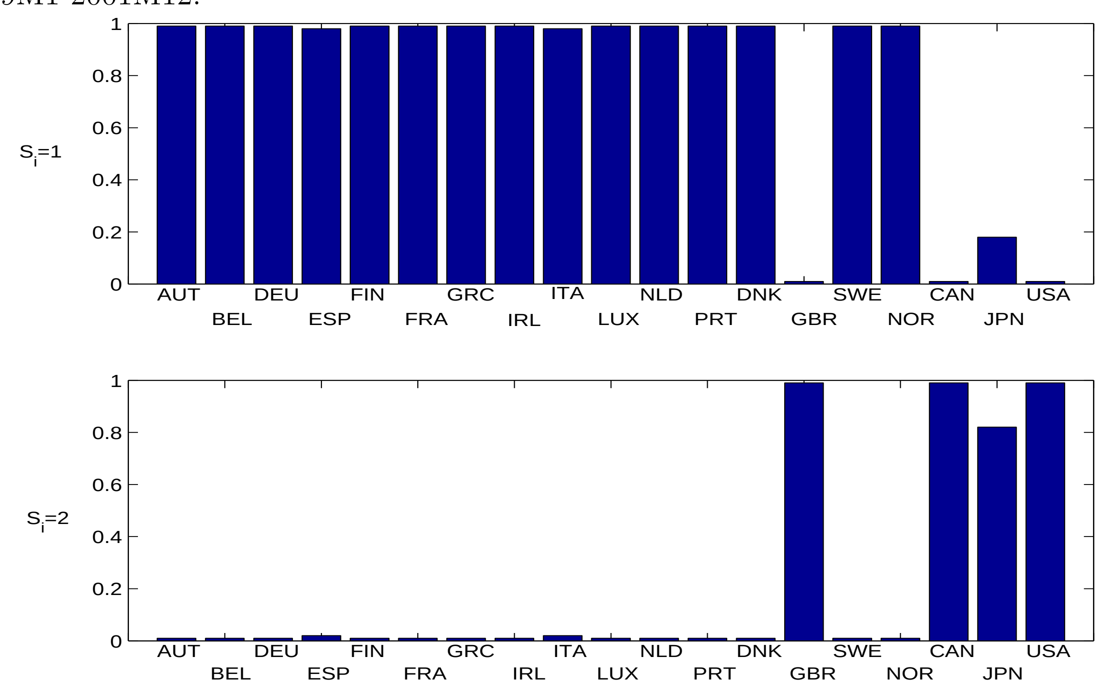

Figure 15: State specific (monthly) mean growth rates of the European and the overseas countries. The first group is identified by the dotted marker, the second group by the squared marker.

(c) 1999M1-2001M12

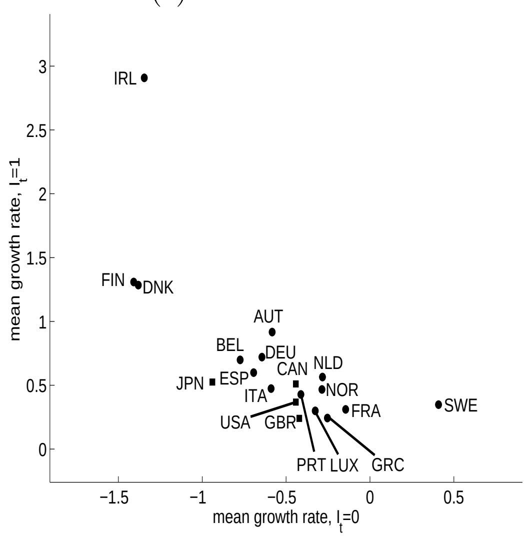

6 Conclusion

The paper presents a model for analyzing the common growth pattern of country-individual industrial production series. Within an autoregressive panel framework, the growth rate (the constant) is allowed to switch between two states according to a latent state indicator which is thought to capture the unobservable expansionary or recessionary state of the economy. Moreover, for a given number of groups, the appropriate grouping of the countries in terms of a "similar" growth rate pattern, is estimated along with the model parameters rather than setting it a priori on the basis of some ad hoc defined country characteristics.

For the analysis, we use quarterly and monthly data for all (Western) European countries and the major overseas countries, in particular Australia, Canada, Japan and the US. With quarterly data, the focus is set on a long-term (1978-2001) and on a medium-term (1990-2001) historical perspective. The monthly data are used to analyze the recent period of the common European currency (1999-2001). The advantage of using a panel data framework lies first in the possibility of enlarging the number of countries to be analyzed within a single model without having to cope with an exponentially increasing number of parameters to be estimated as it would be the case when estimating a vector autoregressive system. Another advantage lies in the possibility of analyzing a shorter time period as we can exploit the information contained in the cross-section.

The results obtained from the estimation of the model give us answers to the following questions. *Is there a common growth cycle for the euro area countries?* Yes; when the euro area and also the European countries are analyzed on their own, they fall into one group under either the long-term or the medium-term perspective. However, during the recent three years, two groups of countries can be discriminated. The first group consists of countries recording a longer-lasting expansionary period than the second group, whereby the largest euro area countries along with Spain and the UK, and along with Austria, Belgium and Denmark fall into this second group (and the Nordic countries with Greece, Ireland, Luxembourg, the Netherlands and Portugal form the first group). When we include additionally the overseas countries, the European countries fall into one group in general, with the exception of Finland and Ireland under the long-term perspective and the UK during the recent three years. In these cases, the countries follow more closely the "overseas" cycle mainly characterized by Australia, Canada and the US.

Can we discriminate between a "European" and an "overseas" cycle? Yes; when the analysis includes Australia, Canada and the US, these three countries characterize one group over all time horizons, while most European countries fall into the other group.

Did the synchronization change over time? Yes; the increased integration of European countries is reflected in the changing classification of Finland and Ireland. While under the long-term perspective they follow more closely the "overseas" cycle, they follow more closely the European cycle during the nineties and during the period of the common European currency. Moreover, there is evidence for a changing synchronization between the European and the overseas cycle. Until the early nineties, overseas downturns were leading European downturns by half a year to about one year. During the nineties, however, overseas countries went through a long-lasting expansionary period while European countries experienced three full growth cycles. Several factors might have led to this disentangled cyclical pattern of the two regions, such as the German re-unification, a potential stronger European exposure to the Asian and the Russian crises, and a slow pace of appropriate (labour market and/or fiscal and economic) policy reforms needed to foster growth in an increasingly integrated European environment.

Which countries follow the "overseas" rather than the "European" cycle? The results document the increased integration of European countries and in particular the integration of the euro-area countries. Since the nineties, Finland and Ireland are following more closely the European cycle, while under the longer historical perspective, they were following the overseas cycle. In contrast to previous studies, the UK, and also Japan, follow more closely the European rather than the overseas cycle. It is only during the recent three years, that both countries are classified into one group with Canada and the US. Finally, the downturn in these countries is identified to have begun already in the second half of 2000, before affecting all countries in 2001.

References

Altissimo, F., A. Bassanetti, R. Cristadoro, M. Forni, M. Lippi, L. Reichlin, and G. Veronese (2001). EuroCOIN: A real time coincident indicator of the euro area business cycle. Discussion Paper 3108, CEPR.

Artis, M. J., H.-M. Krolzig, and J. Toro (1999). The European business cycle. Discussion Paper 2242, CEPR.

Artis, M. J. and W. Zhang (1995). International business cycles and the ERM: Is there a European business cycle? *International Journal of Finance and Economics 2*, 1–16.

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

Carter, C. K. and R. Kohn (1994). On Gibbs sampling for state space models. *Biometrika 84*, 541–553.

Chib, S. (1996). Calculating posterior distributions and modal estimates in Markov mixture models. *Journal of Econometrics 75*, 79–97.

European Central Bank (2002). Characteristics of the euro area business cycle in the 1990s. Monthly bulletin, july, 39-49.

Forni, M. and L. Reichlin (2001). Federal policies and local economies; Europe and the US. *European Economic Review*.

Frühwirth-Schnatter, S. (1994). Data augmentation and dynamic linear models. *Journal of Time Series Analysis 15*, 183–202.

Frühwirth-Schnatter, S. (2001). MCMC estimation of classical and dynamic switching and mixture models. *Journal of the American Statistical Association 96*, 194–209.

Frühwirth-Schnatter, S. and S. Kaufmann (2002a). Bayesian clustering of many short time series. mimeo, Vienna University of Economics and Business Administration, and Oesterreichische Nationalbank.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica 57*, 357–384.

Hoogstrate, A. J., F. C. Palm, and G. A. Pfann (2000). Pooling in dynamic panel-data models: An application to forecasting GDP growth rates. *Journal of Business & Economic Statistics 18*, 274–283.

Kaufmann, S. (2002). Is there an asymmetric effect of monetary poliy over time? A Bayesian analysis using Austrian data. *Empirical Economics 99*, 277–298.

Mitchell, J. and K. Mouratidis (2002). Is there a common Euro-zone business cycle? National Institute of Economic and Social Research.

Shephard, N. (1994). Partial non-Gaussian state space. *Biometrika 81*, 115–131.

A The prior distribution of θ

The vector θ gathers all model parameters, i.e. the group-specific regression parameters $(\mu_1^G, \dots, \mu_K^G)$ and $(\mu_1^R, \dots, \mu_K^R)$, the autoregressive parameters $(\phi_1, \dots, \phi_p)$, the variance σ^2, the group probabilities $(\eta_1^G, \dots, \eta_K^G)$ and the group-specific transition probabilities $(\eta^1, \dots, \eta^K)$, where $\eta^k = (\eta_{00}^k, \eta_{01}^k, \eta_{10}^k, \eta_{11}^k)$.

A priori, we assume that the parameter vectors are independent of each other:

$$\pi(\theta) = \pi(\mu_1^G, \dots, \mu_K^G, \mu_1^R, \dots, \mu_K^R)\pi(\phi_1, \dots, \phi_p)\pi(\sigma^2)\pi(\eta_1^G, \dots, \eta_K^G)\pi(\eta^1, \dots, \eta^K).$$

Specifically, the prior distributions are parameterized in the following way:

- The pairs of group-specific regression parameters, (μ_k^G, μ_k^R), $k = 1, \dots, K$, are independent of each other and have a normal distribution,

  $$N(m_0, M_0), \tag{8}$$

  where

  $$m_0 = \begin{pmatrix} b_0 \\ 0 \end{pmatrix}, \quad M_0 = \begin{pmatrix} B_0 & B_0 \\ B_0 & 2B_0 \end{pmatrix}.$$

  The specific feature of the prior comes from the parameterization of the state-specific constant in model (1). Remind that state 1 relates to $I_t = 0$ with parameter $\mu_k^G - \mu_k^R$ and state 2 relates to $I_t = 1$ with parameter μ_k^G. To apply the permutation sampler described in detail in the next section, the prior distribution of the state-specific parameters needs to be symmetric and invariant with respect to state permutation. Therefore, the prior distribution on the pair of coefficients $(\mu_k^G - \mu_k^R, \mu_k^G)$ is assumed to be normal,

  $$\begin{pmatrix} \mu_k^G - \mu_k^R \\ \mu_k^G \end{pmatrix} \sim N\left(\begin{pmatrix} b_0 \\ b_0 \end{pmatrix}, \begin{pmatrix} B_0 & 0 \\ 0 & B_0 \end{pmatrix}\right).$$

  It is then easy to derive that the prior specification on $(\mu_k^G - \mu_k^R, \mu_k^G)$ implies the one given in (8) for (μ_k^G, μ_k^R).

- The autoregressive parameters $(\phi_1, \dots, \phi_p)$ are independently normally distributed $N(0, \kappa I_p)$ where I_p denotes the identity matrix of dimension p and κ is a positive constant. Note that for the present investigation, it is not necessary to truncate the prior distribution to ensure a stationary autoregressive process. Throughout, the sampler itself yields autoregressive values that confine to the stationarity region.

- A natural prior for the variance σ^2 is the inverse gamma distribution, $\sigma^2 \sim IG(g_0, G_0)$.

- For the group probabilities $(\eta_1^G, \dots, \eta_K^G)$ we assume a Dirichlet prior distribution, $D(e_{1,0}^G, \dots, e_{K,0}^G)$.

- The distribution of the group-specific transition distributions, $(\eta^1, \dots, \eta^K)$, are assumed to be independent of each other and to have independent Dirichlet distributions *a priori*, $\eta^k \sim \prod_{j=0}^{1} D(e_{j0,0}, e_{j1,0})$, $\forall k$.

When estimating the model, one can not be totally uninformative, i.e. use improper prior distributions, in particular about the group- and state-specific parameters. If the number of groups or states is set too large, it is quite possible that none country falls into the superfluous group or that one of the states is never simulated. In this case, non-informative/improper prior distributions would lead to improper posterior distributions. Therefore, we have to be informative; nevertheless, the hyperparameters chosen put only little information on the prior distribution. Specifically, $b_0 = 0.5$ and $B_0 = 4$; $\kappa = 1$; $g_0 = 1$ and $G_0 = 1$; $e^G_{k,0} = 4$, $k = 1, \ldots, K$; $(e_{00,0}, e_{01,0}) = (2, 1)$, $(e_{10,0}, e_{11,0}) = (1, 2)$.

B The sampling scheme

The Markov chain Monte Carlo simulations are based on the permutation sampler proposed in Frühwirth-Schnatter and Kaufmann (2002a) which is extended to the present case of group-specific state indicators. Restate the model first:

$$\Delta y_{it} = \mu^G_{S_i} + \mu^R_{S_i}(I_{S_i,t} - 1) + \phi_1 \Delta y_{i,t-1} + \cdots + \phi_p \Delta y_{i,t-p} + \varepsilon_{it},$$

with $\varepsilon_{it} \sim$ i.i.d. $N(0, \sigma^2)$.[10]

To obtain the inference on the joint posterior distribution $\pi(\theta, S^N, I^T|Y^T)$, the sampling scheme involves the following steps:

(i) $\pi(S^N|Y^T, \theta, I^T)$,
(ii) $\pi(I^T|Y^T, \theta, S^N)$,
(iii) $\pi(\theta|Y^T, S^N, I^T)$.

Simulating S^N out of $\pi(S^N|Y^T, \theta, I^T)$. Note first that the posterior distribution may be factorized

$$\pi(S^N|Y^T, \theta, I^T) \propto \prod_{i=1}^{N} \prod_{t=p+1}^{T} f(y_{it}|y_i^{t-1}, \mu^G_{S_i}, \mu^R_{S_i}, \phi_1, \ldots, \phi_p, \sigma^2, I_{S_i,t}) \pi(S_i|\eta^G).$$

Therefore, the group indicators S^N are independent and may be simulated from the discrete distribution

$$\pi(S_i = k|y_i, \theta, I^T_k) \propto \prod_{t=p+1}^{T} f(y_{it}|y_i^{t-1}, \mu^G_k, \mu^R_k, \phi_1, \ldots, \phi_p, \sigma^2, I_{kt}) \cdot \eta^G_k,$$

where the observation density $f(y_{it}|\cdot)$ is normal with mean $\mu^G_k + \mu^R_k(I_{kt} - 1) + \sum_{j=1}^{p} \phi_j \Delta y_{i,t-j}$ and variance σ^2.

Denote by p_{ik} the kth cumulated sum of the normalized probability distribution,

$$p_{ik} = \sum_{1}^{k} \frac{\pi(S_i = k|y_i, \theta, I^T_k)}{\sum_{k=1}^{K} \pi(S_i = k|y_i, \theta, I^T_k)}.$$

Then, each group indicator S_i, $i = 1, \ldots, N$, may be simulated independently from a uniform distribution. Given a draw U, $S_i = k$ according to $k = 1 + \sum_{j=1}^{K} 1(U > p_{ij})$.

Simulating I^T out of $\pi(I^T|Y^T, \theta, S^N)$. Given the group indicator S^N, the group-specific state indicators are independent of each other. Therefore, the posterior distribution can

[10] The posterior distributions are derived explicitly for the present model. If the notation is generalized, the sampler also applies to models involving more right-hand side variables, however (see Frühwirth-Schnatter and Kaufmann, 2002a).

be factorized by

$$\pi(I^T|Y^T,\theta,S^N) = \prod_{k=1}^{K} \pi(I_k^T|Y^T,\theta,S^N) = \prod_{k=1}^{K} \pi(I_{kT}|Y_T,\theta,S^N) \prod_{t=0}^{T-1} \pi(I_{kt}|Y^t, I_{k,t+1},\theta,S^N).$$

whereby the typical element $\pi(I_{kt}|Y^t, I_{k,t+1},\theta,S^N)$ is proportional to

$$\pi(I_{kt}|Y^t, I_{k,t+1},\theta,S^N) \propto \pi(I_{kt}|Y^t,\theta,S^N) \cdot \eta^k_{I_{kt},I_{k,t+1}}. \tag{9}$$

The first factor $\pi(I_{kt}|Y^t,\theta,S^N)$ is obtained from a forward filter, whereby

$$\pi(I_{kt}|Y^t,\theta,S^N) \propto \prod_{S_i=k} f(y_{it}|y_i^{t-1},\theta,S_i,I_{kt}) \cdot \pi(I_{kt}|Y^t,\theta,S^N).$$

The observation density $f(y_{it}|\cdot)$ is normal with mean and variance given above, the element $\pi(I_{kt}|Y^t,\theta,S^N)$ is given by extrapolation:

$$\pi(I_{kt}|Y^t,\theta,S^N) \propto \sum_{I_{k,t-1}=0}^{1} \pi(I_{k,t-1}|Y^{t-1},\theta,S^N) \cdot \eta^k_{I_{k,t-1},I_{kt}}.$$

The filter must be initialized at time 0, one possibility is to use the ergodic probabilities of the kth state indicator, $\pi(I_{k0}) = \rho^k$, or to start with a user-specific distribution, e.g. a uniform distribution, $\pi(I_{k0}) = (1/2, 1/2)$.

Once the filter densities $\pi(I_{kt}|Y^t,\theta,S^N)$, for $t = 1,\dots,T$ are computed, beginning in T, each group-specific state indicator I_{kT}, $k = 1,\dots,K$, may be sampled out of the discrete distribution $\pi(I_{kT}|Y_T,\theta,S^N)$. Then, the recursion in (9) is used to simulate I_{kt} for $t = T-1,\dots,0$, $\forall k$. This multi-move method has been independently proposed in the literature by Carter and Kohn (1994), Frühwirth-Schnatter (1994), and Shephard (1994) for latent state space variables and has been explicitly derived for latent Markov switching processes in Chib (1996).

Given S^N and I^T, the vector θ is blocked appropriately to simulate the model parameters out of their full conditional distributions.

Simulating the group probabilities $\eta^G = (\eta_1^G,\dots,\eta_K^G)$ *from* $\pi(\eta^G|Y^T,S^N,I^T)$. Given S^N, the posterior distribution in fact reduces to $\pi(\eta^G|S^N)$ and is given by the Dirichlet distribution

$$\pi(\eta^G|S^N) = D(e_1^G,\dots,e_K^G),$$

with $e_k^G = e_{k,0}^G + \sum_{i=1}^N 1(S_i = k)$.

Simulating the transition probabilities $\eta = (\eta^1,\dots,\eta^K)$ *from* $\pi(\eta|Y^T,S^N,I^T)$. In analogy, the posterior distribution reduces to $\pi(\eta|I^T)$ and is a mixture of independent Dirichlet distributions,

$$\pi(\eta|I^T) = \prod_{k=1}^{K}\prod_{j=0}^{1} D(e_{j0}^k, e_{j1}^k)$$

with $e_{jl}^k = e_{jl,0} + \sum_{t=1}^T 1(I_{kt} = l, I_{k,t-1} = j)$.

Simulating the model parameters from $\pi(\mu_1^G,\dots,\mu_K^G,\mu_1^R,\dots,\mu_K^R,\phi_1,\dots,\phi_p|Y^T,S^N,I^T,\sigma^2)$. Given the sampled values S^N and I^T, the model in (1) can be written compactly as

$$\Delta y_{it} = Z_{it}\beta + \varepsilon_{it}$$

where $\beta = (\mu_1^G, \dots, \mu_K^G, \mu_1^R, \dots, \mu_K^R, \phi_1, \dots, \phi_p)$ and

$$Z_{it} = \left(D_i^{(1)}, \dots, D_i^{(K)}, D_i^{(1)}(I_{1t} - 1), \dots, D_i^{(K)}(I_{Kt} - 1), \Delta y_{i,t-1}, \dots, \Delta y_{i,t-p}\right).$$

The dummy variables are defined as $D_i^{(k)} = 1$ iff $S_i = k$ and $D_i^{(k)} = 0$ otherwise. All parameters are then simulated jointly from a normal posterior distribution, $N(b, B)$, with variance B and mean b given by, respectively,

$$\begin{aligned} B &= \left(\sum_{i=1}^{N} \sum_{t=p+1}^{T} Z'_{it} Z_{it}/\sigma^2 + B_0^{*-1}\right)^{-1} \\ b &= B\left(\sum_{i=1}^{N} \sum_{t=p+1}^{T} Z'_{it} y_{it}/\sigma^2 + B_0^{*-1} b_0^*\right). \end{aligned}$$

The joint normal prior distribution $N(b_0^*, B_0^*)$ is constructed from the prior distribution of the parameter blocks (see appendix A for the definition of m_0 and M_0):

$$b_0^* = \begin{pmatrix} m_0 \otimes \imath_K \\ 0_{p\times 1} \end{pmatrix}, \quad B_0^* = \begin{bmatrix} M_0 \otimes I_K & 0 \\ 0 & \kappa I_p \end{bmatrix},$$

where $\otimes$ denotes the Kronecker product, $\imath_K$ a $K \times 1$ vector of ones, $0_{p\times1}$ a $p \times 1$ vector of zeros and I_K the identity matrix of dimension K.

Simulating the variance σ^2 from $\pi(\sigma^2|Y^T, S^N, I^T, \beta)$. The posterior distribution is given by the inverted Gamma

$$\sigma^2|Y^T, S^N, I^T, \beta \sim IG(g, G)$$

with

$$\begin{aligned} g &= g_0 + N(T-p)/2 \\ G &= G_0 + \frac{1}{2} \sum_{i=1}^{N} \sum_{t=p+1}^{T} (\Delta y_{it} - Z_{it}\beta)^2. \end{aligned}$$

Note that, throughout, the simulations are not constrained, i.e. we sample out of the unconstrained posterior distribution. However, to obtain an identified model, restrictions should be set on the group- and state-specific simulated values. To discriminate between the groups, we might set a restriction on the constant in state 2 ($I_t = 1$):

$$\mu_1^G < \dots < \mu_K^G, \qquad \text{(R1)}$$

and to discriminate between the states (in each group) we might use the restriction:

$$\mu_k^G > \mu_k^G - \mu_k^R, \quad \forall k \quad \text{(R2)}$$

which in fact implies $\mu_k^R > 0$. It is obvious that a number of other restrictions could discriminate between the groups and the states, like e.g. the restrictions $\mu_1^R < \dots < \mu_K^R$ or $\eta_{00}^k < \eta_{11}^k$, $\forall k$, respectively. If the appropriate group- and state-identifying restrictions are known a priori, the Gibbs sampler described above is completed by a permutation step, in which all simulated group- and state-specific values are reordered appropriately to fulfill the restrictions. This amounts to first interchange the state-specific values in each group k, $k = 1, \ldots, K$, if (R2) is violated:

$$\begin{aligned} \mu_k^G &:= \mu_k^G - \mu_k^R \\ \mu_k^R &:= -\mu_k^R \\ \eta_{jl}^k &:= \eta_{1-j,1-l}^k, \; j, l = 0, 1 \\ I_{kt} &:= 1 - I_{kt}, \; t = 0, \ldots, T. \end{aligned}$$

Then, if the group identifying restriction (R1) is violated, the group-specific parameters are reordered according to the permutation $\rho(1), \ldots, \rho(K)$ that fulfills (R1):[11]

$$\begin{aligned} (\mu_1^G, \ldots, \mu_K^G) &:= (\mu_{\rho(1)}^G, \ldots, \mu_{\rho(K)}^G) \\ (\mu_1^R, \ldots, \mu_K^R) &:= (\mu_{\rho(1)}^R, \ldots, \mu_{\rho(K)}^R) \\ (\eta^1, \ldots, \eta^K) &:= (\eta^{\rho(1)}, \ldots, \eta^{\rho(K)}) \\ (I_1^T, \ldots, I_K^T) &:= (I_{\rho(1)}^T, \ldots, I_{\rho(K)}^T) \\ (S_1, \ldots, S_N) &:= (\rho(S_1), \ldots, \rho(S_N)) \\ (\eta_1^G, \ldots, \eta_K^G) &:= (\eta_{\rho(1)}^G, \ldots, \eta_{\rho(K)}^G). \end{aligned}$$

If, on the other hand, a state- and a group-identifying restriction are not known a priori, one might permute the simulated values randomly at each iteration to force the sampler to explore the whole unconstrained posterior distribution. The simulated values might then be post-processed and displayed in scatter plots to find an appropriate identifying restriction (see Frühwirth-Schnatter, 2001, and Kaufmann, 2002). Indeed, it turned out that the restrictions (R1) and (R2) are the appropriate ones to obtain the inference in the present paper.

[11] If e.g. $K = 3$ and the simulated values fulfill $\mu_1^G > \mu_2^G > \mu_3^G$ instead of (R1), the appropriate permutation would be $(\rho(1), \ldots, \rho(K)) = (3, 2, 1)$.