Similarities and convergence in G-7 cycles

Fabio Canova[a], Matteo Ciccarelli[b],*, Eva Ortega[c]

[a]ICREA-Universitat Pompeu Fabra, CREI, and CEPR, Spain
[b]European Central Bank, DG Research, Kaiserstrasse 29, D-60311, Frankfurt am Main, Germany
[c]Banco de España, Madrid, Spain

Abstract

This paper examines the properties of G-7 cycles using a multicountry Bayesian panel VAR model with time variations, unit specific dynamics and cross-country interdependences. We demonstrate the presence of a significant world cycle and show that country specific indicators play a much smaller role. We detect differences across business cycle phases but, apart from an increase in synchronicity in the late 1990s, find little evidence of major structural changes. We also find no evidence of the existence of a Euro specific cycle or of its emergence in the late 1990s.
© 2006 Elsevier B.V. All rights reserved.

> Pricelings, resolve your private bickerings
> Yourselves. 'Tis folly to resort to kings.
> Fight your own battles, form your own opinions,
> And keep all strangers out of your dominions.
> <div align="right">La Fointaine</div>

1. Introduction

There is abundant evidence that economic activity in developed countries share a number of characteristics. For example, the real business cycle literature has demonstrated

*Corresponding author. Tel.: +49 69 13448721.
E-mail address: matteo.ciccarelli@ecb.int (M. Ciccarelli).

that macroeconomic fluctuations across industrialized countries are closely linked (see e.g. Backus et al., 1995; Baxter, 1995; Canova and Marrinan, 1998) while more structured time series analyses have shown that a large portion of regional and country specific fluctuations are common (see e.g. Gregory et al., 1997; Del Negro, 2002; Lumsdaine and Prasad, 2003; Kose et al., 2003) and that a number of linear and non-linear business cycle features are similar (Harding and Pagan, 2006). Although all existing evidence suggests the presence of a common source of fluctuations in developed countries, results are typically derived using restrictive or conventional assumptions about the nature of the dynamic relationships. In fact, except for Del Negro (2002) or Kose et al. (2003), the issue of whether commonalities are present or not is examined within an empirical framework which does not explicitly allow for cross-country and cross-variables interdependencies.

Recently, the stability of the cross-country business cycle relationships has come under scrutiny. For example, Helbling and Bayoumi (2003) suggest that the increase in cyclical synchronization observed after 2000 in advanced economies is the result of common shocks. This apparently represents a major shift relative to the 1980s where increased similarities in macroeconomic fluctuations were the result of improved trade relationships (see e.g. Canova and Dellas, 1992; Rose and Frankel, 2000 for a quantification of the magnitude of trade links). It also represents a somewhat different propagation mechanism from the one used to explain the transmission of US shocks to Asia (see Mackowiak, 2003) or Latin America (Canova, 2005). On the other hand, Stock and Watson (2003) document changes in the volatilities of G-7 business cycles in the 1990s and indicate that such changes may have also altered the correlation among international macroeconomic variables.

A third issue which has recently attracted attention is the behavior of European cycles in the period leading to the creation of the European Monetary Union, both within the area and in relationship with other developed countries. Casual observation indicates that European cycles have displayed more synchronized movements in the 1990s and, possibly, a larger transatlantic conformity. All in all, commentators suggest that national effects in Europe may be slowly but surely vanishing.

Knowing whether fluctuations in the industrialized world are similar, understanding their sources and characterizing their time variations is important for both academics and policymakers. From an academic point of view one is interested in knowing whether business cycle links are the result of cross-country interdependencies or common shocks. Hence, one may welcome studies empirically documenting similarities in economic fluctuations since the presence of a common cycle facilitates the study of the relationship between national and international policy decisions and the state of the world economy. Policymakers monitoring domestic or regional cycles are typically concerned with the effect of national idiosyncrasies and with the consequences that their actions have on the working of international markets. However, if variations in economic activity in countries with different institutions, economic structures or economic policies are driven by a common cause, international markets more than national policies are the key to understanding the comovements in economic activity. Moreover, national or regional policies designed to counteract world tendencies may be ineffective. Finally, structural time variations may undermine the usefulness of policies which may have been effective in the past.

This paper breaks ground in the area by explicitly addressing two interrelated questions. First, we would like to know whether there has been any tendency for G-7 cycles to

become more similar in the 1990s or if, on the contrary, they tend to be clustered along geographical, regional or other institutional characteristics. Second, we are curious as to whether there is any evidence that Euro cycles are different from those we observe in the rest of the G-7 or if they have become more different in the recent past. In answering these questions we also provide new evidence on the relative importance of world and country specific causes in driving G-7 cycles, on their evolution over time and over business cycle phases.

To study these questions we employ a panel VAR model of the type developed in Canova and Ciccarelli (2002). Their approach is useful in our context for at least two reasons. First, the econometric methodology is designed for large-scale dynamic models displaying unit specific dynamics and cross-country lagged interdependencies and flexibly allows for time variations in the correlation structure of cyclical fluctuations across variables and countries. Second, the parsimonious parametrization they propose endogenously produces an index structure where indicators of world and national specific cycles can be recursively constructed and dynamically span cross-country interdependencies. Therefore, the specification is particularly suited to study the interrelationships and the structural changes present in G-7 cycles and to analyze what drives the common and the idiosyncratic components of G-7 fluctuations.

Our investigation confirms some of the existing evidence. For example, as in Kose et al. (2003) or Lumsdaine and Prasad (2003), we find evidence of a significant world business cycle, despite different empirical techniques and data sets. However, we also provide new and important insights in the phenomenon. For instance, our results indicate that the world indicator accounts for about 30% of the fluctuations in sales, industrial production, output and employment of the seven most industrialized countries; that it captures the more persistent portions of G-7 fluctuations and that it has more information than simple average or principal component measures obtained using G-7 GDPs or IPs. We also show that country specific indicators are useful in explaining certain GDP and employment episodes over time, but fail to consistently track cyclical movements in the four variables over the entire sample. Perhaps more interestingly, we find that both world and country specific fluctuations are more synchronized in contractions than expansions. Expansions tend to have large idiosyncratic components, both across variables and countries, while declines in economic activity have common timing and similar dynamics, both within and across countries.

Regarding the questions of interest, we find no evidence of structural breaks in the country indicators in the 1990s. Hence, the often cited idea that national cycles are disappearing finds no support from our analysis. These indicators are as significant in explaining the differential growth rate of GDP across countries in the mid-1990s as they were in the mid-1980s and, if anything, slightly more important. We also find little support for the idea that Euro cycles are different or that a Euro area cycle is emerging in the 1990s. This result should be contrasted with Lumsdaine and Prasad (2003) and Artis et al. (2003), who instead detected the presence of a EU cycle using IP data, and with Del Negro and Otrok (2003), who find that within EU synchronization has increased in the late 1990s. Our analysis shows that the Euro signal is weaker when one considers a broader set of variables and that regional causes have minor explanatory power for G-7 fluctuations throughout the sample.

Taken together, these results imply that movements in the world indicator have been the stable and consistent explanation for the commonalities of the fluctuations in the G-7

economies over time and that structural breaks, in both the pattern of transmission across countries and in the sources of structural shocks, are probably absent.

We attempt to characterize the informational content of the indicators we construct using simple correlation analysis. We document that our world indicator captures a variety of influences going from the magnitude of world trade, to the stance of monetary policy in the G-7 and to the spending power of consumers. Hence one of the reasons why the world indicator has been a very stable driver of international business cycles over the period is that it captures a variety of sources of disturbances and endogenously allows their mix to change over time. We also show that oil shocks, US technology shocks and fiscal policy do not appear to be behind the movements in the world indicator. Furthermore, we document that, country indicators capture primarily the differential stance of monetary policy but, in some cases, are also related to local spending capability or US variables. Interestingly, their informational content is not related to the stance of local fiscal policy, except for Canada and Italy.

The rest of the paper is organized as follows: the next section presents the model specification, the technique used to construct the various indicators and the details of our empirical approach. Section 3 presents the results and Section 4 concludes.

2. The empirical model

2.1. The panel VAR model

The model we employ in our analysis has the form:

$$y_{it} = D_{it}(L)Y_{t-1} + c_{it} + e_{it}, \tag{1}$$

where $i = 1, \ldots, N$ refers to countries and $t = 1, \ldots, T$ to time. y_{it} is a $G \times 1$ vector for each country i and $Y_t = (y'_{1t}, y'_{2t}, \ldots, y'_{Nt})'$. $D_{it,j}$ are $G \times NG$ matrices for each lag $j = 1, \ldots, p$, c_{it} is a $G \times 1$ vector of intercepts and e_{it} is a $G \times 1$ vector of random disturbances.

Whenever the $NG \times NG$ matrix $D_t(L) = [D_{1t}(L), \ldots, D_{Nt}(L)]'$, is not block diagonal for some L, the model displays cross-unit lagged interdependencies. Eq. (1) displays two other important features. First, the coefficients are allowed to vary over time. Second, the dynamic relationships are allowed to be unit specific. All three ingredients are crucial when one wants to study similarities, the propagation and time variations in the structure of business cycles across countries. Moreover, they add considerable realism to the specification but are costly: the number of parameters is large (we have now $k = NGp + 1$ parameters in each equation) and there is only one time period in each unit to estimate them.

It is convenient to rewrite (1) in a simultaneous equations format:

$$Y_t = W_t \delta_t + E_t, \quad E_t \sim N(0, \Omega), \tag{2}$$

where $W_t = I_{NG} \otimes X'_t$; $X'_t = (Y'_{t-1}, Y'_{t-2}, \ldots, Y'_{t-p}, 1)$, $\delta_t = (\delta'_{1t}, \ldots, \delta'_{Nt})'$ and δ_{it} are $Gk \times 1$ vectors containing, stacked, the G rows of the matrix D_{it} and of c_{it}, while Y_t and E_t are $NG \times 1$ vectors containing the endogenous variables and the random disturbances.

Since δ_t varies with cross-sectional units in different time periods, it is impossible to estimate it using classical methods. However, even if δ_t is time invariant, its sheer dimensionality prevents its unconstrained estimation. Two shortcuts are typically employed in the literature: it is assumed that δ_t does not depend on the unit, apart from

a time invariant fixed effect, or that there are no interdependencies across units (see e.g. Holtz–Eakin et al., 1988; Binder et al., 2000). Since neither of these assumptions is appealing in our study, we factor δ_t as

$$\delta_t = \Xi_1 \lambda_t + \Xi_2 \alpha_t + \Xi_3 \rho_t + u_t, \tag{3}$$

where Ξ_1, Ξ_2, Ξ_3 are matrices of dimensions $NGk \times N_1 \ll N$, $NGk \times N$, $NGk \times G$, respectively, and $\lambda_t, \alpha_t, \rho_t$ are mutually orthogonal. λ_t captures movements in the coefficient vector which are common across units and variables—for most of this paper λ_t will be a scalar; in Section 3.3 it will be a 2×1 vector. α_t captures movements in the coefficient vector which are common within countries and therefore its dimension equals to N, the number of countries in the panel. ρ_t captures movements in the coefficient vector which are variable specific and its dimension is therefore equal to G, the number of variables in each country. Finally, u_t captures all the unmodelled features of the coefficient vector, which may have to do with lag specific, time specific or other idiosyncratic effects.

Factoring δ_t as in (3) is advantageous in many respects. Computationally, it reduces the problem of estimating NGk coefficients into one of the estimating $1 + N + G$ or $2 + N + G$ factors driving the coefficients. Therefore, even when the number of interdependent units is large estimation is feasible, noise is averaged out and reliable estimates of the features of interest can be obtained. Practically, the factorization (3) transforms an overparametrized panel VAR into a parsimonious SUR model where the regressors are 'appropriate' averages of the right-hand side variables of the VAR. In fact, substituting (3) into (2) we have

$$Y_t = \mathcal{W}_t \lambda_t + \mathcal{A}_t \alpha_t + \mathcal{M}_t \rho_t + v_t, \tag{4}$$

where $\mathcal{W}_t = W_t \Xi_1, \mathcal{A}_t = W_t \Xi_2, \mathcal{M}_t = W_t \Xi_3$ capture, respectively, common, country specific and variable specific information present in the VAR and $v_t = E_t + W_t u_t$. Economically, the decomposition in (4) is convenient since it allows us to measure the relative importance of world and country specific influences for fluctuations in Y_t. In fact $WLI_t = \mathcal{W}_t \lambda_t$ plays the role of a world (or of a regional) indicator, while $CLI_t = \mathcal{A}_t \alpha_t$ plays the role of a vector of country specific indicators. Both coincident and leading versions of these indicators can be designed using either time t or time $t - 1$ information, since $\mathcal{W}_t$ and $\mathcal{A}_t$ only contain information present in the predetermined variables of the VAR, and recursively constructed, given a law of motion of λ_t and α_t. Note that WLI_t and CLI_t are correlated by construction—the same variables enter in both $\mathcal{W}_t$ and $\mathcal{A}_t$—but become uncorrelated as the number of countries becomes large.

To illustrate the structure of the matrices Ξ's and of $\mathcal{W}_t, \mathcal{A}_t, \mathcal{M}_t$, suppose there are $G = 2$ variables for each of $n = 2$ countries, $p = 1$ lags and an intercept:

$$\begin{bmatrix} y_t^1 \\ x_t^1 \\ y_t^2 \\ x_t^2 \end{bmatrix} = \begin{bmatrix} d_{1,1,t}^{1,y} & d_{2,1,t}^{1,y} & d_{1,2,t}^{1,y} & d_{2,2,t}^{1,y} \\ d_{1,1,t}^{1,x} & d_{2,1,t}^{1,x} & d_{1,2,t}^{1,x} & d_{2,2,t}^{1,x} \\ d_{1,1,t}^{2,y} & d_{2,1,t}^{2,y} & d_{1,2,t}^{2,y} & d_{2,2,t}^{2,y} \\ d_{1,1,t}^{2,x} & d_{2,1,t}^{2,x} & d_{1,2,t}^{2,x} & d_{2,2,t}^{2,x} \end{bmatrix} \begin{bmatrix} y_{t-1}^1 \\ x_{t-1}^1 \\ y_{t-1}^2 \\ x_{t-1}^2 \end{bmatrix} + \begin{bmatrix} c_1^y \\ c_1^x \\ c_2^y \\ c_2^x \end{bmatrix} + e_t. \tag{5}$$

Here $\delta_t = [d_{1,1,t}^{1,y}, d_{2,1,t}^{1,y}, d_{1,2,t}^{1,y}, d_{2,2,t}^{1,y}, c_1^y, d_{1,1,t}^{1,x}, d_{2,1,t}^{1,x}, d_{1,2,t}^{1,x}, d_{2,2,t}^{1,x}, c_1^x, d_{1,1,t}^{2,y}, d_{2,1,t}^{2,y}, d_{1,2,t}^{2,y}, d_{2,2,t}^{2,y}, c_2^y, d_{1,1,t}^{2,x}, d_{2,1,t}^{2,x}, d_{1,2,t}^{2,x}, d_{2,2,t}^{2,x}, c_2^x]'$ is a (20×1) vector containing the time varying coefficients of the model. Note that the typical element of δ_t, $\delta_{l,s,t}^{i,j}$, is indexed by the country i, the variable

j, the variable in an equation l (independent of the country), and the country in an equation s (independent of variable). If we are not interested in modelling all these aspects and call u_t all unaccounted features, one possible factorization of δ_t is

$$\delta_t = \Xi_1 \lambda_t + \Xi_2 \alpha_t + \Xi_3 \rho_t + u_t, \tag{6}$$

where for each t, λ_t is a scalar, α_t is a 2×1 vector, ρ_t is a 2×1 vector, Ξ_1 is a 20×1 vector of ones and

$$\Xi_2 \atop (20 \times 2) = \begin{bmatrix} \iota_1 & 0 \\ \iota_1 & 0 \\ 0 & \iota_2 \\ 0 & \iota_2 \end{bmatrix}, \qquad \Xi_3 \atop (20 \times 2) = \begin{bmatrix} \varkappa_1 & 0 \\ 0 & \varkappa_2 \\ \varkappa_1 & 0 \\ 0 & \varkappa_2 \end{bmatrix}$$

with $\iota_1 = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 \end{pmatrix}'$, $\iota_2 = \begin{pmatrix} 0 & 0 & 1 & 1 & 0 \end{pmatrix}'$, $\varkappa_1 = \begin{pmatrix} 1 & 0 & 1 & 0 & 0 \end{pmatrix}'$ and $\varkappa_2 = \begin{pmatrix} 0 & 1 & 0 & 1 & 0 \end{pmatrix}'$. Hence, the VAR (5) can be rewritten as

$$\begin{bmatrix} y_t^1 \\ x_t^1 \\ y_t^2 \\ x_t^2 \end{bmatrix} = \begin{bmatrix} \mathscr{W}_t \\ \mathscr{W}_t \\ \mathscr{W}_t \\ \mathscr{W}_t \end{bmatrix} \lambda_t + \begin{bmatrix} \mathscr{A}_{1,t} & 0 \\ \mathscr{A}_{1,t} & 0 \\ 0 & \mathscr{A}_{2,t} \\ 0 & \mathscr{A}_{2,t} \end{bmatrix} \alpha_t + \begin{bmatrix} \mathscr{M}_{1,t} & 0 \\ 0 & \mathscr{M}_{2,t} \\ \mathscr{M}_{1,t} & 0 \\ 0 & \mathscr{M}_{2,t} \end{bmatrix} \rho_t + v_t, \tag{7}$$

where $\mathscr{W}_t = y_{t-1}^1 + x_{t-1}^1 + y_{t-1}^2 + x_{t-1}^2 + 1$, $\mathscr{A}_{1,t} = y_{t-1}^1 + x_{t-1}^1$, $\mathscr{A}_{2,t} = y_{t-1}^2 + x_{t-1}^2$, $\mathscr{M}_{1,t} = y_{t-1}^1 + y_{t-1}^2$, $\mathscr{M}_{2,t} = x_{t-1}^1 + x_{t-1}^2$ and $v_t = e_t + W_t' u_t$. In the empirical application, all variables are measured in growth rates and therefore this type of averaging will indeed be appropriate. Note that if λ_t is large relative to α_t, y_t^1 and x_t^1 comove with y_t^2 and x_t^2. On the other hand, if λ_t is zero, y_t^1 and x_t^1 may drift apart from y_t^2 and x_t^2. Note also that in the general case when $p > 1$, lags can be weighted using a decay factor in the same spirit as Doan et al. (1984).

As the notation we have used makes it clear, the regressors in (4) are combinations of lags of the right-hand side variables of the VAR, while $\lambda_t, \alpha_t, \rho_t$ play the role of time varying loadings. Using averages as regressors is common in the factor model literature (see e.g. Stock and Watson, 1989; Forni and Reichlin, 1998) and in the signal extraction literature (see e.g. Sargent, 1989). However, there are five important differences between (4) and standard factor models. First, the indices we construct weight equally the information in all variables while in factor models the weights generally depend on the variability of the components.[1] Second, our indices dynamically span lagged interdependencies across units and variables while in standard factor models they statically span the space of the variables of the system. Third, our indices are directly observable while in factor models they are estimated. Fourth, the loadings are allowed to be time varying while, except for Del Negro and Otrok (2003), the loadings are constant in factor models. Finally, our averaging approach creates moving average terms of order p in the regressors of (4), even when y_{it} are serially independent. Therefore, contrary to what occurs in factor models, our indicators implicitly filter out from the right-hand side variables of the VAR

---

[1]It is important to stress that the equal weighting scheme comes from the prior specification and the unit in which the variables are measured. In some cases, weighting the variables by the relative size of the economies, as recently done by Pesaran et al. (2004), could also be appropriate.

high-frequency variability. The fact that the regressors of the SUR model emphasize the low-frequency movements in the variables of the VAR is important in forecasting in the medium run and in detecting turning points of GDP, as we will see in the next section.

2.2. The structure of time variations

To complete the empirical specification we need to describe the evolution of λ_t, α_t, ρ_t over time and the features of their (prior) distribution at time zero. Write (3) compactly as

$$\delta_t = \Xi\theta_t + u_t, \quad u_t \sim N(0, \Sigma \otimes V), \tag{8}$$

where $\Xi = [\Xi_1, \Xi_2, \Xi_3]$, $\theta_t = [\lambda_t, \alpha_t, \rho_t]'$, and V is a $k \times k$ matrix and let

$$\theta_t = \theta_{t-1} + \eta_t, \quad \eta_t \sim N(0, B_t). \tag{9}$$

We further assume that $\Sigma = \Omega$ and $V = \sigma^2 I_k$, σ^2 known; that $B_t = \gamma_1 B_{t-1} + \gamma_2 \bar{B}$, γ_1, γ_2 known; that $\bar{B} = diag(\bar{B}_1, \bar{B}_2, \bar{B}_3)$, and that Ξ_t, u_t and η_t are mutually independent. We treat $(\sigma^2, \gamma_1, \gamma_2)$ as parameters and their selection is described in Appendix A.

In (9) the factors evolve over time as random walks. Alternative specifications which allow for more complex dynamics or exchangeability across units are possible (see e.g. Canova and Ciccarelli, 2002). We stick to this simple setup since experimentation with more complicated structures did not produce qualitatively important improvements in our results. The spherical assumption on V reflects the fact that the factors have similar units, while setting $\Sigma = \Omega$ is standard in the literature (see e.g. Kadiyala and Karlsson, 1997). The variance of the innovations in θ_t is allowed to be time varying to account for heteroskedasticity and other generic volatility clustering that may appear in the coefficients of several, or all, series within and across units. Time invariant structures ($\gamma_1 = \gamma_2 = 0$), and homoskedastic variances ($\gamma_1 = 0$ and $\gamma_2 = 1$) are special cases of the assumed process. The block diagonality of $\bar{B}$ is needed to guarantee orthogonality of the factors, which is preserved a posteriori, and hence their identifiability. Finally, independence among the errors is standard.

To summarize, our empirical model has the state space structure:

$$\begin{aligned} Y_t &= (W_t \Xi)\theta_t + v_t, \\ \theta_t &= \theta_{t-1} + \eta_t, \end{aligned} \tag{10}$$

where $v_t \sim (0, \sigma_t \Omega)$ and $\sigma_t = (1 + \sigma^2 X_t' X_t)$. To compute posterior distributions for the unknowns we need prior densities for $(\Omega, \sigma^2, \bar{B}, \theta_0)$. Because we want to minimize the impact of our prior choices on the posterior distribution of the indicators, we specify rather loose but proper priors. Their exact form, the numerical approach used to compute posterior distributions and the details of the computations are in Appendix A. While model (10) can be estimated both with classical and Bayesian methods, we prefer the second since the exact small sample distribution of the objects of interest can be obtained even when T and N are small—as it is the case here—while classical estimates are justifiable only when either T or N or both go to infinity. Furthermore, classical estimators can be obtained only under restrictive assumptions about the nature of time variations

present in the coefficients and are typically more cumbersome to compute than the one employed here.

2.3. Verifying the hypotheses of interest

To evaluate the posterior support for the three main issues of interest, i.e. whether G-7 fluctuations are primarily driven by a world cycle; whether there has been any tendency for G-7 cycles to become more similar in the 1990s and whether Euro area cycles are different from those in the other G-7 countries or have become so in the 1990s, we employ two types of evidence. First, we examine the behavior of the posterior distribution of WLI_t and of CLI_t over time. Second, we compare the marginal likelihood of restricted and unrestricted specifications.

Since time series plots of the posterior of WLI_t and CLI_t are difficult to read, we summarize their information reporting the median of the distribution and a 68% posterior central band—the latter corresponding to one standard deviation around the mean in classical frameworks. Visual evidence in favor or against the first and the third hypotheses is obtained by examining whether the 68% posterior band of the relevant indicator includes zero or not at most or all dates and whether the posterior distribution of the difference between Euro and non-Euro indicators are centered around zero. Visual evidence in favor of time variations/structural breaks can be obtained when the 68% posterior bands includes zero at some dates but not at others. We complement this visual information by computing the percentage of the variations in the endogenous variables due to the world indicator and examining the relative importance of world and national indicators in specific historical episodes.

The marginal likelihood of model i is $f(Y|M_i) = \int L(y|\psi_i, M_i)p(\psi_i|M_i)\,d\psi_i$, where ψ_i are the parameters of the model i. Model i is preferred to model i' if its marginal likelihood is higher or, equivalently, if the Bayes factor $BF(i,i') = f(Y|M_i)/f(Y|M_{i'})$ substantially exceeds 1. To evaluate the questions of interest we consider three models. A benchmark one, M_0, whose specification includes a world, seven country specific and four variable specific factors; a restricted version of this benchmark, M_1, which excludes country specific factors; and an extended version, M_2, where the world indicator WLI_t has two distinct components, one for Euro countries and one for non-Euro ones. Thus, in examining the similarities of cycles across countries, we will compare the marginal likelihood of M_0 and M_1 over the full sample; while in examining the differences of Euro cycles from non-Euro cycles, the marginal likelihood of M_0 is compared with the one M_2 over the full and the 1990s. Marginal likelihoods are computed here as in Chib (1995), letting both θ_t and σ_t be vectors of latent variables.

3. The results

3.1. The data and some forecasting features of the model

For each country we use quarterly growth rates of seasonally adjusted real GDP, employment, sales and industrial production as our basic variables. We choose these four since they are among the variables used by the NBER when deciding the state of the business cycle in the US, by the CEPR when deciding the state of the business cycle in the Euro area and by the Economic Cycle Research Institute (ECRI) when measuring business

cycles around the world.[2] The sample maximizes the amount of common data and covers the period 1979:1–2002:2.

Real GDP data is measured in constant 1995 prices, except for Canada (the base year is 1997). The Japanese series starts only in 1980 and has been extended backwards using the real GDP series measured in constant 1990 prices. Similarly, the Canadian real GDP was extended backwards from 1981 using real GDP data with base-year 1992. The source of all data is the Quarterly National Accounts of OECD, except for Germany whose real GDP comes from the Bundesbank database. We prefer this series since it explicitly takes into account the effects of German unification. Employment is measured by the civilian employment index, with base year 1995, and is from the OECD Main Economic Indicators. Sales are measured by the retail sales volume index, with base year 1995, and come from OECD Main Economic Indicators, except for the US, where the source is the Department of Commerce. The industrial production index also has 1995 as base year and comes from OECD Main Economic Indicators. Whenever these series are provided in non-seasonal adjusted form, we seasonally adjusted them using the TRAMO-SEATS program. In the panel VAR we use these four series in growth rates. While information about their long-run properties may be lost, our choice avoids distortions in the construction of the indicators due to the different size of various aggregates across countries. This is important since the procedure cannot distinguish if a 2% growth is generated in the countries with a large level (say, the US) or with a smaller one (say, Canada).

Before we examine the hypotheses of interest, we first demonstrate the ability of our model to capture the dynamics of the data. We do so presenting statistics which intuitively and visually summarize both the in-sample and the out-of-sample features of our model.

In Table 1 we report the posterior mean and the posterior standard deviation of the out-of-sample RMSE of our Panel VAR model for the seven countries at three horizons (1, 4 and 8 quarters) relative to three benchmarks: a random walk model, a univariate $AR(p_1)$ model, and a AR-factor model where p_2 lags of two factors, extracted from the cross-section of data, are added to a basic AR model, and where both p_1 and p_2 are chosen using the SIC criteria. Standard errors are computed randomizing on the parameters of our Panel VAR model while keeping fixed RMSE estimates for the other models. The evaluation sample is 1990:1–2002:2. Clearly, our specification is preferable in a forecasting sense if the mean statistics are less than one and if standard errors are such that a value of one has low or negligible posterior probability.

Table 1 shows that our model is superior to the competitors in forecasting the four variables on average and for the majority of the countries. Improvement are of the order of 15%–20% on average and statistically significant at all the three horizons. Gains are noticeable for employment growth and for GDP growth while for the other two series the performance of the three models is statistically similar. Interestingly, the AR-factor model improves over a basic AR model at the one step horizon but not at longer horizons. Hence, while the information pooling produced by a factor model is useful for short-term predictions, it is the information contained in cross-sectional lagged interdependencies and the flexible structure of time variations (both of which are absent the AR-factor model) that help most in forecasting in the medium run.

---

[2] The NBER uses real personal income instead of GDP, but such series are not available on a quarterly basis for all the G-7 countries. The CEPR also looks at investment and at some other national variables like the unemployment rate when dating cycles.

Table 1
Posterior means and standard deviations of RMSE relative to benchmarks. Evaluation sample: 1990:1–2002:2

|	Factor model								AR								Random walk							
	GDPG		EG		IPG		SALEG		GDPG		EG		IPG		SALEG		GDPG		EG		IPG		SALEG	
	Mean	Stde	Mean	Stde	Mean	Stde	Mean	Stde	Mean	Stde	Mean	Stde	Mean	Stde	Mean	Stde	Mean	Stde	Mean	Stde	Mean	Stde	Mean	Stde
One step ahead																								
US	1.00	0.06	0.78	0.09	1.06	0.05	1.07	0.05	1.09	0.07	0.69	0.08	0.87	0.04	1.12	0.07	0.89	0.06	0.61	0.07	1.11	0.07	1.06	0.07
JP	0.88	0.03	0.93	0.08	0.92	0.02	0.88	0.02	0.89	0.03	0.82	0.07	0.61	0.01	1.09	0.04	0.67	0.02	0.71	0.06	1.06	0.04	0.97	0.04
GE	0.84	0.04	0.80	0.07	0.91	0.02	0.96	0.02	0.92	0.04	0.77	0.07	0.62	0.01	0.87	0.04	0.70	0.03	0.65	0.06	0.89	0.04	0.77	0.03
UK	1.01	0.08	0.89	0.07	1.32	0.07	1.12	0.06	0.98	0.07	0.79	0.06	1.28	0.07	1.17	0.08	1.06	0.08	0.74	0.06	1.12	0.07	0.97	0.06
FR	0.92	0.06	0.88	0.09	1.06	0.03	1.09	0.03	0.90	0.05	0.95	0.09	0.56	0.01	0.81	0.04	0.71	0.04	0.95	0.09	0.83	0.04	0.72	0.04
IT	0.91	0.05	0.83	0.04	0.91	0.02	0.95	0.02	0.85	0.05	0.84	0.04	0.53	0.01	0.93	0.03	0.62	0.04	0.85	0.04	0.94	0.03	0.84	0.03
CA	0.89	0.06	0.86	0.05	0.98	0.03	0.97	0.03	0.99	0.07	0.83	0.08	0.70	0.02	1.08	0.05	0.99	0.07	0.80	0.08	0.98	0.05	0.86	0.08
Median	0.92	0.03	0.85	0.04	1.05	0.04	0.98	0.02	0.93	0.03	0.81	0.04	0.99	0.04	0.98	0.02	0.73	0.03	0.75	0.04	0.87	0.04	0.62	0.01
Four steps ahead																								
US	0.99	0.08	0.78	0.09	0.80	0.06	1.01	0.05	1.06	0.08	0.77	0.09	0.80	0.06	1.10	0.05	0.75	0.06	0.58	0.07	0.68	0.05	0.80	0.04
JP	0.88	0.03	0.80	0.08	0.75	0.03	0.94	0.02	0.84	0.03	0.78	0.07	0.80	0.03	0.98	0.02	0.64	0.02	0.74	0.07	0.58	0.02	0.79	0.01
GE	0.73	0.04	0.62	0.05	0.80	0.03	0.78	0.02	0.87	0.05	0.80	0.07	0.86	0.03	0.84	0.04	0.60	0.04	0.57	0.07	0.63	0.02	0.58	0.04
UK	0.99	0.09	0.70	0.07	0.73	0.06	1.32	0.08	1.01	0.10	0.70	0.07	0.83	0.07	1.33	0.08	0.88	0.08	0.61	0.07	0.68	0.06	1.04	0.07
FR	0.78	0.05	0.62	0.06	0.79	0.04	0.99	0.02	0.82	0.05	0.58	0.06	0.82	0.04	0.99	0.02	0.58	0.04	0.60	0.06	0.60	0.03	0.85	0.02
IT	0.96	0.06	0.79	0.04	0.97	0.04	0.78	0.02	0.92	0.06	0.80	0.05	0.90	0.04	0.84	0.02	0.63	0.04	0.65	0.04	0.58	0.02	0.61	0.01
CA	0.72	0.07	0.72	0.02	0.82	0.04	0.93	0.03	0.83	0.08	0.79	0.04	0.88	0.04	0.96	0.03	0.65	0.06	0.71	0.07	0.68	0.03	0.66	0.02
Median	0.84	0.03	0.72	0.04	0.80	0.03	0.95	0.02	0.88	0.03	0.76	0.04	0.84	0.03	0.98	0.01	0.64	0.02	0.63	0.03	0.61	0.02	0.78	0.02
Eight steps ahead																								
US	0.99	0.08	0.83	0.08	0.92	0.07	1.11	0.06	1.03	0.08	0.93	0.08	0.88	0.06	1.11	0.06	0.58	0.07	0.60	0.05	0.75	0.06	0.72	0.04
JP	0.82	0.03	0.61	0.06	0.75	0.03	0.89	0.02	0.85	0.03	0.60	0.06	0.83	0.03	0.92	0.02	0.74	0.07	0.53	0.05	0.62	0.02	0.72	0.01
GE	0.72	0.04	0.57	0.08	0.80	0.04	0.86	0.02	0.89	0.05	0.93	0.08	0.90	0.04	0.91	0.02	0.57	0.05	0.54	0.05	0.60	0.03	0.50	0.01
UK	1.00	0.08	0.56	0.06	0.90	0.07	1.11	0.08	1.12	0.11	0.72	0.06	1.03	0.09	1.12	0.09	0.61	0.08	0.53	0.05	0.73	0.06	0.90	0.06
FR	0.80	0.05	0.56	0.04	0.88	0.03	0.93	0.02	0.83	0.05	0.48	0.04	0.93	0.03	0.88	0.02	0.60	0.06	0.45	0.04	0.66	0.02	0.70	0.01
IT	0.73	0.05	0.72	0.04	0.83	0.04	0.82	0.01	0.84	0.05	0.76	0.04	0.91	0.04	0.79	0.02	0.65	0.04	0.66	0.03	0.60	0.03	0.54	0.01
CA	0.84	0.07	0.80	0.09	0.87	0.05	0.99	0.03	0.92	0.08	0.86	0.09	0.94	0.05	0.96	0.03	0.71	0.07	0.57	0.06	0.64	0.04	0.59	0.02
Median	0.82	0.05	0.61	0.06	0.87	0.04	0.93	0.02	0.90	0.04	0.77	0.04	0.91	0.03	0.96	0.02	0.64	0.03	0.55	0.03	0.64	0.02	0.69	0.02
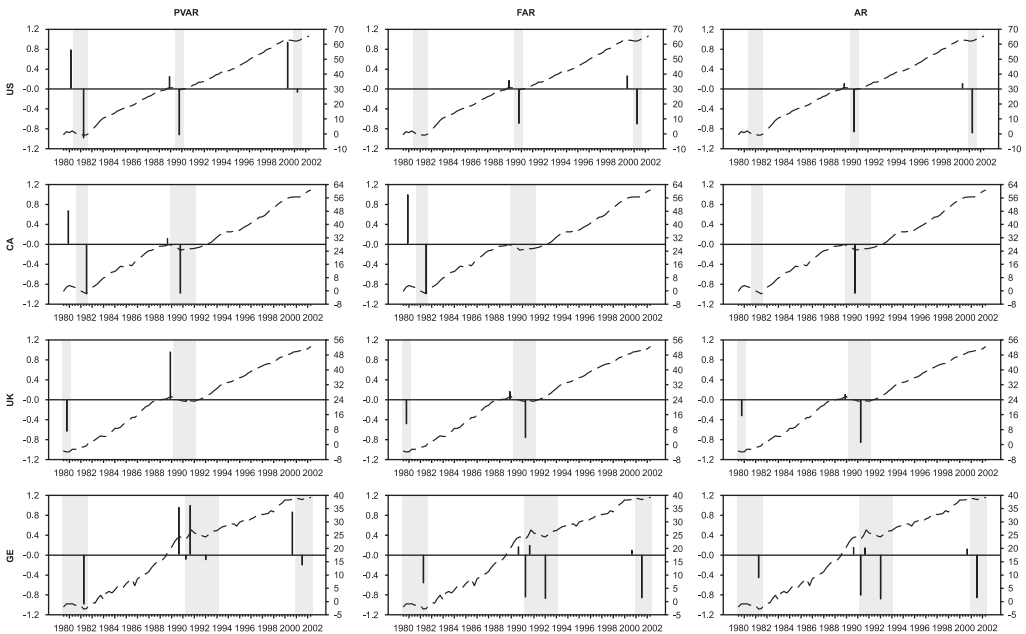
Fig. 1. GDP levels, recession phases and turning point probabilities.

It is well known that RMSE improvements can be obtained by biasing the model's forecasts as long as their variance is reduced. One way of accomplishing this result is to make forecasts very flat. It is therefore possible that good RMSE forecast miss important ups and downs of the series and fail to capture the direction of the actual changes.

To show that our model does not display this feature we present in Fig. 1 the record of turning point probabilities for US, UK, Canada and Germany GDPs in levels. The left column reports the outcomes of our Panel VAR model, the central one those of the AR-factor model and the right column those of a basic AR model. Probabilities of turning points are obtained calculating the percentage of times each model generates the pattern predicted by the following simple rule: there is an upturn at t if $gdp_{t-2} < gdp_{t-1} < gdp_t > gdp_{t+1} > gdp_{t+2}$ and there is a downturn if $gdp_{t-2} > gdp_{t-1} > gdp_t < gdp_{t+1} < gdp_{t+2}$. While the rule is simple and does not make provision for a minimum length of a cycle or for consecutive turning point signals as, for example, the Bry and Boschan (1971) procedure, it suffices to give an idea of the ability of the model to replicate important non-linear functions of the data.

For readability, we present probabilities of downturns with positive sign and probabilities of upturns with a negative sign. Superimposed on the graph are the level of the series and the ECRI downturn (peak to trough) phases are shaded.[3] While our model is designed to capture the features of cycles in the growth rates, it can also reproduce quite well turning points in the level of GDP of several of the G-7 countries. Fig. 1 shows that the model generates high probabilities of turning points in correspondence of the

---

[3]False alarms, i.e., positive probabilities of downturns or upturns not verified in the data, are not reported to avoid confusion. They were nevertheless always lower than 40% for all models considered.

actual ups and downs of the series. In fact, it does not miss any of the recession phases chosen by ECRI for the four countries, a comforting result, given the simplicity of the dating rule. In comparison, the basic AR model misses upturn/downturn signals at the beginning of the 1980s, in both the US and Canada and the AR-factor model faces similar shortcomings.

3.2. Are G-7 cycles (more) similar?

To answer this question we first plot the median and the 68% posterior central band for the world indicator WLI_t and for the seven country specific indicators CLI_{it} in Fig. 2, all in growth rates.

There are several features of Fig. 2 that are worth commenting upon. First, the world indicator captures major events occurred in the last 20 years: it displays the double dip experienced by growth rate of output of several countries at the beginning of the 1980s and a sustained growth thereafter; it captures both the US recession of the beginning of the 1990s and the European one a couple of years later; it shows a positive trend growth in the 1990s, it drops below zero in the first quarter of 2001 and wiggles around zero afterwards. According to this indicator, the 1990 recession is the deepest and, apparently, the longest of all while the (common) growth in the 1980s was larger, on average, than the one experienced in 1990s—Japan's poor performance in the 1990s being partially responsible for this outcome.

Second, posterior uncertainty, as measured by the size of the 68% band, depends on the state of the world economy. Bands are tighter at the beginning, in the middle and at the
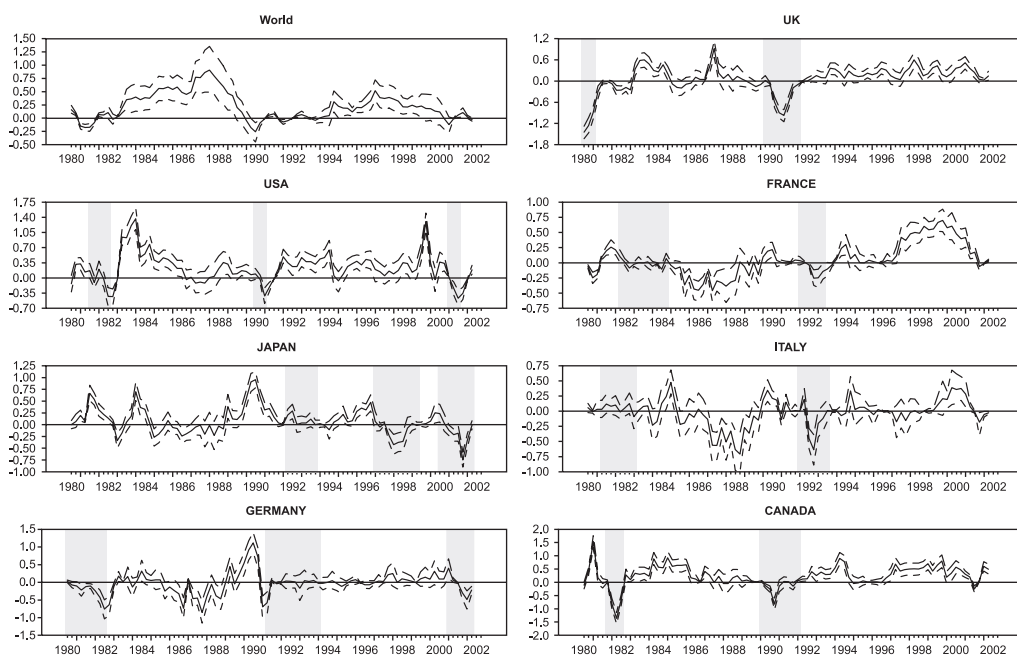

Fig. 2. World and country indicators.

end of the sample and these happened to be periods when the world indicator was either negative or very close to zero. Since posterior uncertainty in λ_t is roughly independent of the state of the economy, one must conclude that the timing and the size of comovements across variables and countries is more similar in contraction than in expansions.

The pattern of comovements observed after 2001 has drawn the attention of several researchers (see e.g. Doyle and Faust, 2002; Peersman, 2005; Helbling and Bayoumi, 2003). In particular, the high synchronicity of the downturn in the industrialized world has been contrasted with the much slower transmission processes experienced early on and commentators have suggested that the sources of international business cycle must have changed. While our analysis confirms the importance of world influences in determining the extent of such a slowdown (although the world indicator does not seem entirely responsible for the similarities in the comovements of G-7 GDPs after that date), it also suggests that the contribution of world disturbances is not unusual when compared with other recessionary episodes of the last 20 years. What appears to have drawn the attention of researchers is not so much a (permanent) structural change in the relationship but the switch in the uncertainty present in recession and expansion phases.

Third, at each time t, the world indicator has little posterior probability mass in the symmetric region centered around zero. Therefore, there is posterior evidence that it significantly contributes to the comovements of the 28 variables of the system. A simple way to numerically measure this contribution is to check how much of the variance of each of the four series is explained by the posterior median of the world indicator.[4]

We find that, on average across countries, the median value of the world indicator explains about 32% of the fluctuations in the growth rate of GDP, about 54% of the fluctuations in employment growth, 17% of the fluctuations of IP growth and 15% of the fluctuations of sales growth. Similarly, we find that 30% of the fluctuations of the four variables within countries are explained by the indicator, with France (44%) and Canada (24%) being the two opposite extremes. The numbers for GDP are surprisingly similar to those reported by Kose et al. (2003) for G-7 countries (36%), despite the fact that the sample and the frequency of the data they use is different. They are also remarkably stable. In fact, considering the 1980s and the 1990s separately, we find that the world indicator explains 32% of GDP fluctuations in the first and 36% in the second sample. Hence, the changes experienced by the world economy over the last 20 years have not altered the importance that common influences have in shaping international business cycles.

Fourth, the 68% posterior band of some of the country indicators (e.g. Germany and Italy) includes zero at most dates. This means that, barring exceptional circumstances, their explanatory power for domestic fluctuations over the last 20 years is minor. To put this observation in another way, German and Italian variables display important fluctuations but their source does not appear to be distinctively national; rather, world wide influences are the reason behind this commonality. Note that, consistent with common perceptions, US fluctuations displays important country specific features in the beginning of the 1980 and in the end of the 1990s and that the Japanese indicator displays a downward trend since the early 1990s. Also for country indicators, posterior uncertainty

---

[4]Since world and country indicators are correlated by construction, some orthogonalization is needed to produce such number and the ordering of the orthogonalization could be important. Luckily, in our system, the maximum correlation of WLI_t with CLI_{it}, $i = 1, \ldots, 7$, is only about 0.15. Hence, the orthogonalization turns out to be unimportant. Here we report averages across the two possible orderings of the two components.

significantly differ across business cycle phases. Hence, asymmetries in business cycle phases are also present at national level, with synchronicity being stronger in downturns than upturns.

To visually evaluate the contribution of country indicators to cyclical fluctuations in the G-7 economies we plot in Fig. 3 the growth rate of GDP (GDPG), the median value of WLI_t (w_factor) and the median value of CLI_t for each country (*country*_factor). Sizeable differences between the first two indicators emerge when country specific idiosyncrasies are important. In the US, Japan and Canada, country specific factors explain movements in the growth rate of domestic GDP for the period 1983–1990 and, to a smaller extent, in the late 1990s. Confirming conventional wisdom, national reasons are behind the poor GDP growth performance of Japan in the late 1990s. The country specific indicator adds little explanatory power to GDP growth movements in Germany: roughly speaking, it captures high-frequency GDP growth movements but little else. In Italy and France, it significantly deviates from the world indicator in the middle of the 1980s and again in the late 1990s—in the earlier episode on the negative side for both countries, in the more recent one on the positive side for France and on the negative side for Italy. The 1980s pattern is probably due to the repeated disruptions of European Monetary System occurred during the period, while in the 1990s stronger convergence following the Maastricht could be the reason for the positive deviation of country factors from world ones. In the UK, country specific disturbances appear to be relevant in capturing the downward trend in GDP growth experienced over the 1986–1992 period.

Fig. 3 also shows that the world indicator captures low-frequency comovements in G-7 countries, while country specific indicators replicate more high-frequency-type fluctuations. In fact, the AR(1) coefficient of the posterior median of the world indicator is 0.92, while the AR(1) coefficient of the posterior median of the country indicators varies from 0.57 for Germany to 0.81 for Canada, with Italy and France close to the lower end.

Are national cycles disappearing? Figs. 2 and 3 indicate that there are significant time variations in the posterior distributions of both the world and the country specific indicators primarily linked with the varying level of uncertainty present in recessions and expansions. However, one could also notice that country specific indicators may have become slightly more synchronized in the last decade—probably as a result of greater synchronization in the ups and downs of domestic variables. From Fig. 4 we can see that, if we exclude Japan, the degree of synchronization across countries (as measured by correlations between country indicators over subsamples) has increased slightly on average from the 80s to the 90s. While US, Canada and Germany have become less synchronized with the others, UK, France and Italy have substantially increased their synchronization. Moreover, the average correlation between the world indicator and the country indicators has also increased.

These observations are important for policy purposes—the greater synchronization means that the information contained in GDP may suffice to characterize the state of the local economy—and intriguing from an academic point of view—is the slight increase in domestic synchronization the result of policy actions, of domestic specialization, of increased similarities in the sectorial shocks hitting the economies or simply of reduced uncertainty connected with a specific growth episode? However, our analysis fails to detect any significant break in the structure of comovements across the two parts of the sample. The 1980s are not different from the 1990s as far as sources of business cycles are concerned: national cycles play a role in both decades and the hypothesis that they are
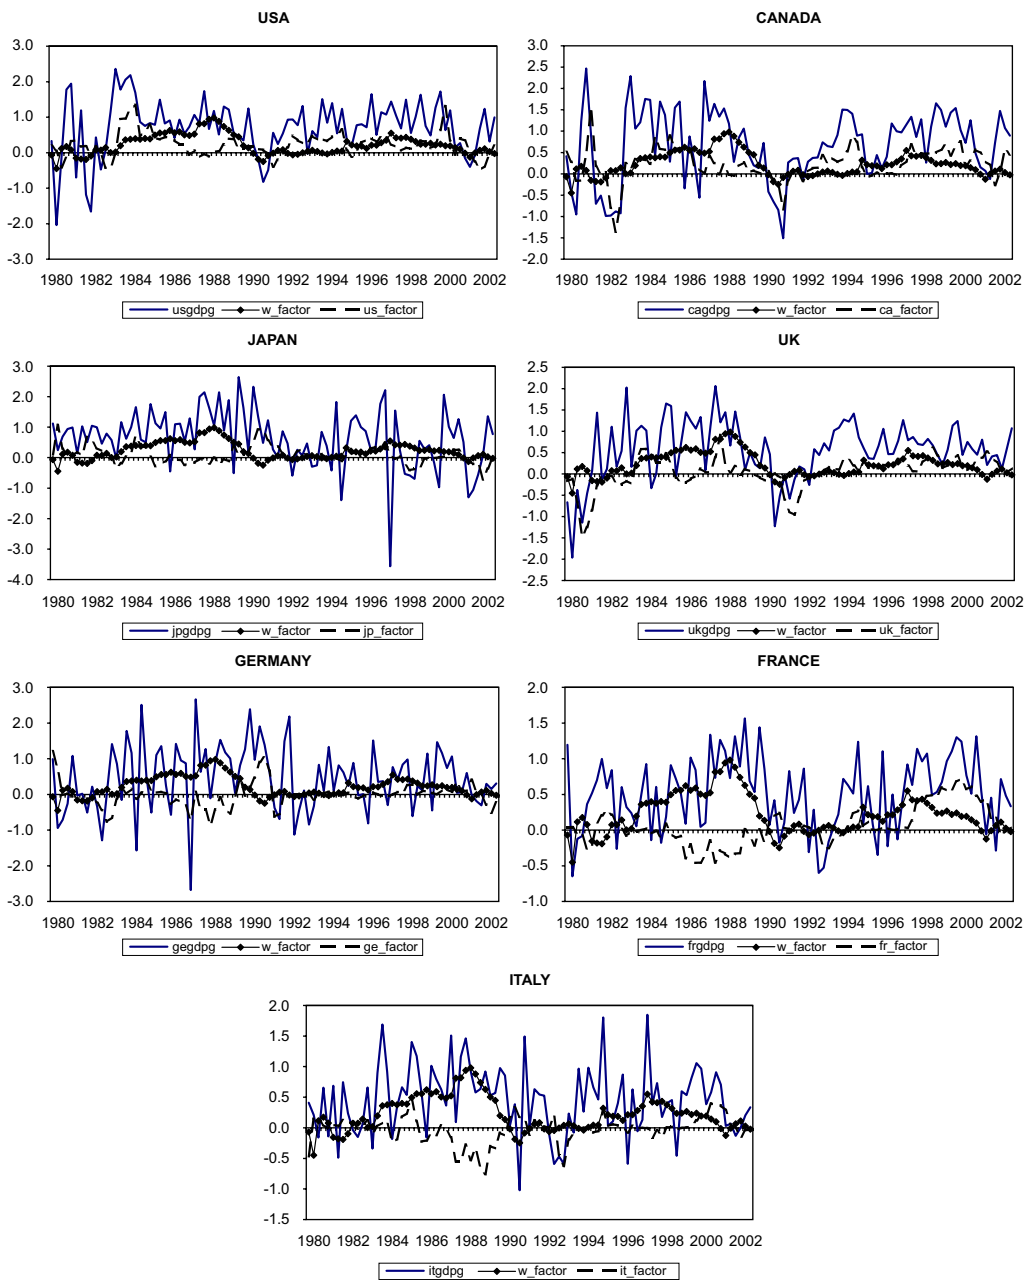
Fig. 3. Actual GDP growth and posterior medians of the indicators.

vanishing has hardly any support in the data. In fact, when formally comparing the benchmark model and a restricted one without country specific factors, the data always prefer the benchmark specification and by a large amount for the full sample (the log of the Bayes factor is 471.5). On the other hand, the performance of the restricted model is
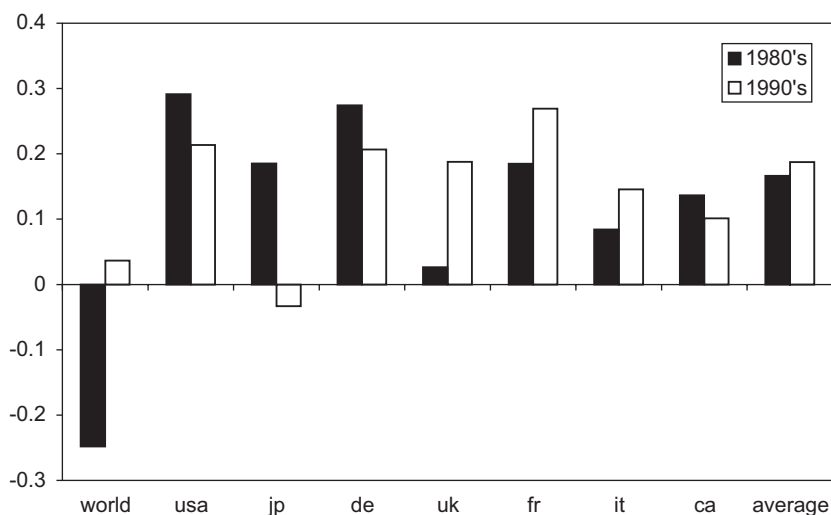
Fig. 4. Average correlations between a country (world) indicator with all the others.

relatively better over the sample 1992–2003 (log of the Bayes factor is 130.1) than in the sample 1980–1991 (log of the Bayes factor is 271.0).

This conclusion is not necessarily in contrast with any idea concerning the globalization of markets (see also Perri and Heathcothe, 2001). Conventional wisdom suggests that more globalization should bring a stronger synchronicity in the business cycles of the industrialized world as a result of an increase in trade. However, the possibility of easy trade may also make local production more specialized (along the lines of standard static or dynamic comparative advantage). Specialization may therefore bring less synchronicity in production. There is no reason to expect one of the two effects to dominate, and one can envision scenarios where increased globalization leaves the synchronicity of G-7 cycles unaffected.

3.3. Are European cycles different?

Country specific indicators may turn out to be significant in explaining national cycles if there is a regional element in the fluctuations which has been neglected in our analysis so far. For example, if Euro cycles tend to deviate from say, rest-of-the-world (ROW) cycles, the presence of a single world indicator may spuriously produce significant country specific indicators because of omitted variable biases. Similarly, the importance of national factors could be artificially reduced if the regional factor has become more highly correlated with the world factor. Should one worry about these problems? The analyses of Lumsdaine and Prasad (2003); Del Negro and Otrok (2003); Stock and Watson (2003), who all find regional EU factors to be important, and the fact that France is the country which display the most significant country specific indicator indicate that such a risk may be present.

We therefore repeated our exercise allowing for 2×1 vector of world indices: one for Euro area countries (France, Italy and Germany) and one for the Anglo-Saxons (USA, Canada, UK). The rest of the model is unchanged. We decided to leave Japan out of the two groups since the synchronicity of Japan's economy with the other G-6 cycles

dramatically declines from the mid-1980s when other East Asian countries (Korea, Taiwan, China) became the major trading partners of Japan. We also experimented including Japan in a ROW cycle and switching UK into the European group, despite evidence of diversity of Euro area and UK cycles (see e.g. Harding and Pagan, 2006; Artis et al., 2003). We settled on the above specification since the Bayes factor always preferred it and, in both cases, by a large margin.

Fig. 5 presents the time path of the median and 68% posterior central band for the two regional indicators. Three features of the figure stand out. First, the posterior band of the Euro area indicator almost always includes zero. In fact, if we exclude the 1986–1989 period, the posterior 68% band always sits over zero, roughly, 90% of the times. Hence, there is little posterior evidence that a Euro area indicator is needed to capture the dynamics of the data. Second, there is also little posterior evidence that Euro area cycles differ from the Anglo-Saxon ones: the overlap of the two posterior distributions at each t is substantial. Moreover, with the exclusion of the 1983–1985 period, the evolution of the median of the posterior of the two indicators is similar, in both the 1980s and the 1990s. Fourth, while for the full sample, a model with two regional indicators is preferable according to marginal likelihood calculations (log of Bayes factor is 35.9), for the 1992–2003 sample, the difference in the log of the marginal likelihoods is negligible (equal to 1.6). Given Fig. 5, this is not surprising: even though one extra degree of freedom is gained, the predictive ability does not necessarily improve splitting the world factor into two, and this is true especially in the 1990s, where posterior uncertainty seems to be somewhat larger than in the 1980s. Taken together these observations provide no posterior support for the idea that Euro area cycles have become more important in the later part of the sample. Nevertheless, one should remember that our analysis is based on the three largest Euro area countries. It could well be that considering other members of the Euro area, for example including Portugal or Greece, will make the Euro indicator significantly
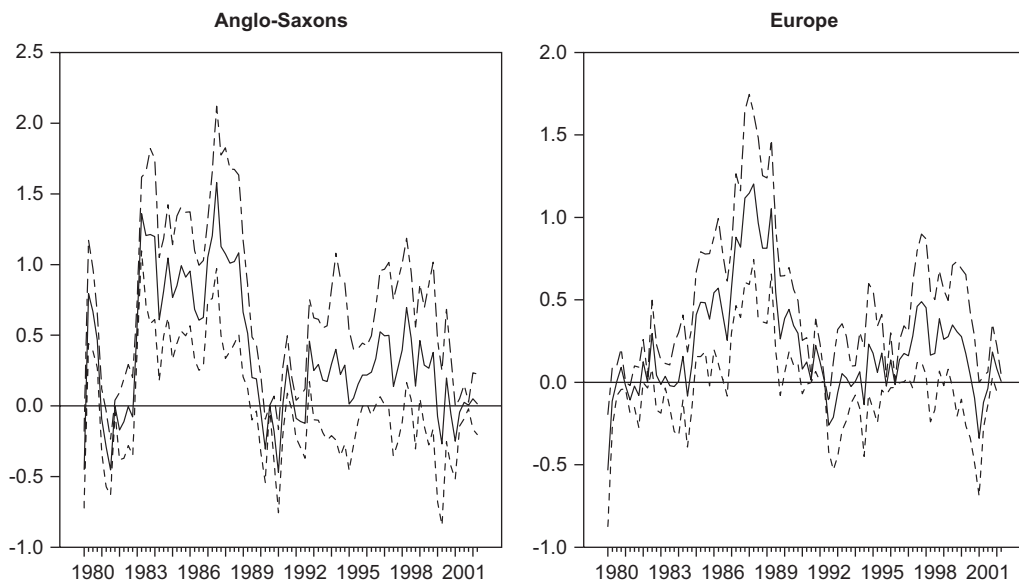

Fig. 5. Anglo Saxons and EMU indicators.

different from the ROW in the second subsample. Finally, notice that also with this specification, posterior uncertainty tends to be reduced during contraction phases. Hence, the larger posterior uncertainty present in the two indicators (relative to a world one) for most of the sample is not due to the lack of information in the two aggregates.

While the absence of a Euro area cycle may appear surprising, it is worth stressing that our result is not unique. For example, Kose et al. (2003) find no evidence of a second world factor, in general, and of a EU factor, in particular. How should one then interpret the evidence? Our favorite interpretation is the following: while Euro area aggregates display common fluctuations across variables and economies, their source is not distinctively European. Euro area and Anglo-Saxon fluctuations are similar in timing, size and amplitude because they are driven by the same source of disturbances. Since regional causes play a minor role in explaining comovements in cyclical fluctuations, Euro area policymakers should closely monitor the state of the international business cycle and de-emphasize Euro area cycles.

4. What drives our indicators?

So far the analysis has been primarily statistical. However, to go beyond the simple documentation of the time series properties of the world and country indicators and study how they relate to the more structural evidence presented, e.g., in Perri and Heathcote (2001), it is necessary to study the informational content of the indicators we constructed.

To start with, we examine how the world indicator relates to simple and easily computable measures of common statistical fluctuations such as arithmetic averages of GDP, IP, employment or sales growth in the G-7 or their principal components. If the behavior of the world indicator we construct can be reproduced with such simple measures, our more complicated setup can be easily dismissed. We find that arithmetic averages share some informational content with our indicator but the overlap is far from perfect. In fact, the point estimate of the bivariate correlation between the posterior median of the world indicator and the simple average at each point in time of the GDPs, IPs, employments and sales growth rates are 0.52, 0.45, 0.62, 0.44, respectively.

Similarly, we find that our indicator is correlated with the principal component of GDPs, IPs, employment and sales growth (point estimates 0.59, 0.46, 0.59, 0.32). Fig. 6 graphically provides evidence of this association using GDP and employment growth: the principal component is typically much more volatile than our indicator, while simple arithmetic average measures occasionally miss important cyclical movements in the data. Our world indicator is the smoothest of all and mimics the local trend present in the data better than these standard measures.

To study the informational content of the indicators one could proceed as in the recent factor-VAR literature (see e.g. Bernanke et al., 2003) and measure the contribution of structural shocks to their fluctuations. Since structural shocks are difficult to identify in our international context, we prefer to compute a number of simple correlations (Table 2). These correlations are non-structural and therefore are only suggestive of the possible sources of the cyclical movements captured by our indicators.

We have correlated the world indicator with the growth rate of the US real personal non-agricultural income, the growth rate of world commodity price index, the NYSE stock return index, the growth rate of the world market crude petroleum price, the growth rate of the world goods trade, the quarterly US technology shock extracted by Gali et al., 2003,
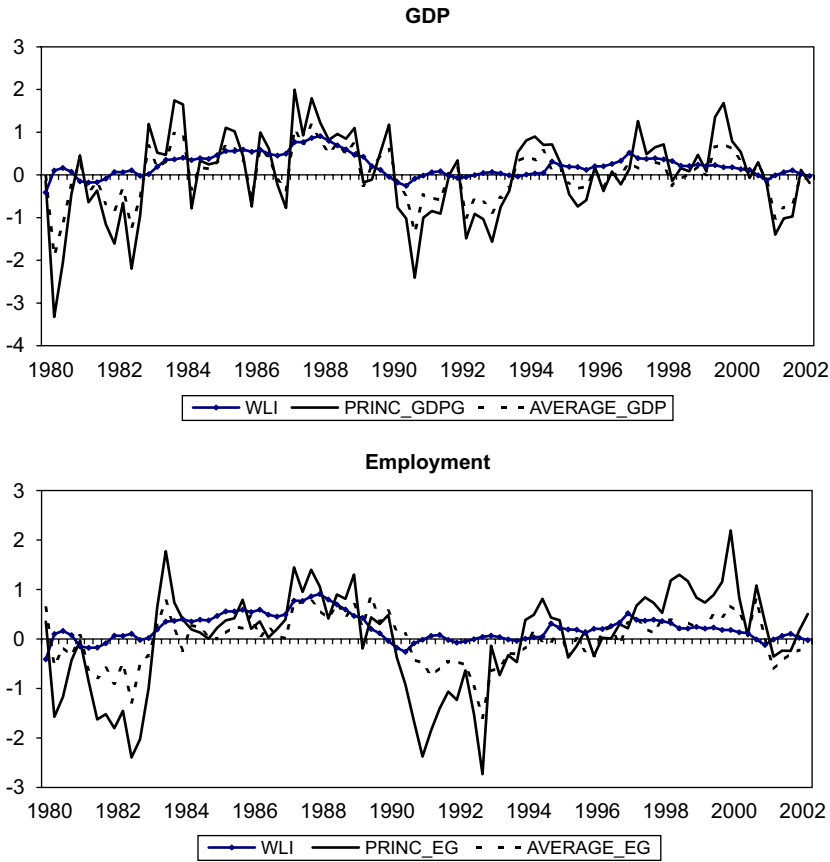
GDP

Employment

Fig. 6. World indicator, principal component and averages.

the average spread between US and other G-7 long-term interest rates, both nominal and real, the average growth rates of real private consumption to GDP ratio, real effective exchange rates, M3 and average G-7 government deficit.[5]

Is the world indicator a stand-in for oil shocks, technology shocks or other types of supply side disturbances? It does not seem to be the case. We find that the correlations between the world indicator and commodity prices, oil prices and US technology shocks are small and insignificant (point estimates -0.02, -0.10 and -0.06). Similarly, the correlation with real exchange rates is low and insignificant. Therefore, while oil shocks may be an important source of national disturbances, and technology shocks in the US may explain an important portion of US fluctuations, it is necessary to go beyond these disturbances to explain and interpret existing world business cycles.

---

[5]The series are seasonally adjusted and come from the OECD Quarterly National Accounts and Main Economic Indicators databases, except for the general government deficits, oil price (world market crude petroleum price) and world commodity price index, which were taken from the IMF International Financial Statistics database. For acronyms and definitions, see the appendix.

Table 2
Contemporaneous correlations with *WLI*

Average GDPG	Average IPG	Average EG	Average SALEG
0.52**	0.45**	0.62**	0.44**
per capita GDP	per capita IP	per capita E	per capita SALE
0.59**	0.46**	0.59**	0.32**
CPG	OIL	GLVTECH	REE
−0.02	−0.1	−0.05	0.06
USSTOCK	US DEF	M3G	AVGLIR
0.18*	−0.02	0.20*	−0.26**
WTRADE	CY	USPIN	AVGLID
0.29**	−0.20*	0.27**	−0.45**

Notes: * means significantly different from zero at 10% significance level, ** at 5% significance level. For the acronyms used see Table 4.

Does the world indicator capture fluctuations originating in financial or monetary markets or the stance of fiscal policy? The answer is mixed. The stance of the fiscal policy does not seem to be related to the world indicator; its correlation with average G-7 government deficit is −0.02 and not significantly different from zero. However, the correlation with US stock returns and average growth of M3 is significant (point estimates 0.18 and 0.20, respectively), while the world indicator appears to be most significantly correlated with the average long-term real and nominal interest rate prevailing in the G-7 (point estimates −0.27 and −0.45, respectively).

Besides monetary and financial developments, what else is behind the movements of the world indicator? It appears that the world trade (point estimate 0.29), the average consumption/output ratio (point estimate −0.20) and US personal income (point estimate 0.27) are the variables with the most significant correlations. Interestingly, as in Cochrane (1994), we find that the consumption/output ratio is almost as good as any other variables in explaining common movements in the G-7.

In sum, our world indicator captures a number of influences (trade, monetary policy, financial conditions, spending capacity) that analysts and academics have indicated to be important to understand the dynamics of business cycles in the industrialized world. One explanation for the remarkable stability in the explanatory power of this index may be the fact that it robustly and flexibly captures different sources of fluctuations, allowing for time variations and adaptive changes in the weights.

There are another couple of other interesting facts which our investigation has discovered. First, there is a strong negative relationship between the average explanatory power of the world indicator for national cycles and volatility of GDP (see Fig. 7). Such a relationship is at times studied in growth literature to evaluate the desirability of stable growth. In our context, this pattern implies that synchronicity with the rest of the G-7 improves (worsens) as the volatility of domestic GDP fluctuations is reduced (increases). Second, the world indicator has slightly larger explanatory power for employment than output (and larger correlation with the average employment growth than with average
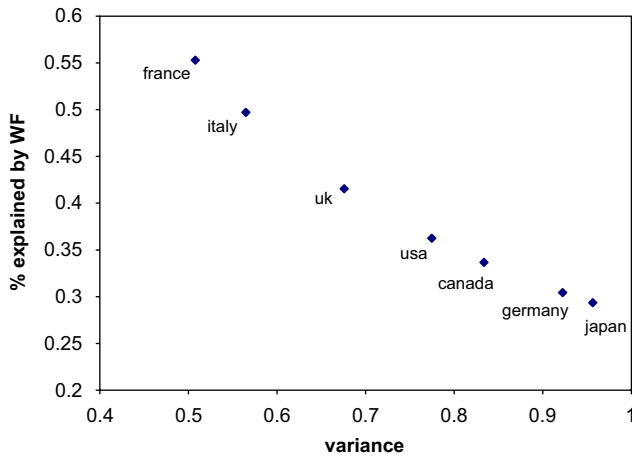
Fig. 7. Syncronicity and GDP volatality.

GDP growth). This result squares well with those derived by the international RBC literature where employment correlations larger than GDP correlations are interpreted as suggestive of the presence of an important world cycle.

We have also attempted to identify the informational content of country specific indicators. We have correlated them with the individual series used in the model (the growth rate of real GDP, industrial production, employment and real retail sales) as well as with other domestic variables like the three month interbank nominal interest rate, the nominal yield on 10 year bonds, the real short-term and long-term interest rates, the growth rate of M3, the general government deficit, the consumption output ratio and the real exchange rate and with a few other international variables (the US personal income, the commodity price index, oil prices, world trade, the US technology disturbances, US stock returns). We found that all the country indicators are very significantly correlated to local interest rates (long and short term, either nominal or real), with the only exception of Germany, while France and Japan national indicators are also significantly correlated to their respective money growth rates.

World trade is found to be the international variable more significantly correlated with country indicators (only insignificant for France and Germany), followed by US personal income and the US stock exchange. Besides monetary policy indicators, the other domestic variable with which correlations are found very significant is the consumption to GDP ratio, with the exceptions of Italy and Germany. Interestingly, national indicators are unrelated to the stance of fiscal policy, except in Canada and Italy.

5. Conclusions and directions for future work

This paper studies similarities and convergence of G-7 business cycles using a panel VAR model with cross-country interdependencies and time variations. The framework of analysis is advantageous in several respects. First, the structure allows for multiple types of contemporaneous and lagged comovements and for time variations in the correlation

structure of cyclical fluctuations across variables and countries. Second, the parsimonious parametrization we use endogenously produces an observable index structure where indicators of world and national cycles can be recursively constructed and dynamically span cross-country interdependencies. Third, the specification allows us to verify the posterior support for the hypotheses of interest and to analyze what drives the common and the idiosyncratic components of G-7 fluctuations.

We address two interrelated questions. First, we study whether there has been any tendency for G-7 cycles to become more similar in the 1990s or if, on the contrary, they tend to be clustered along geographical, regional or other institutional characteristics. Second, we examine whether there is any evidence that Euro cycles are different from those we observe in the rest of the G-7 or if they have become so in the recent past. In answering these questions we also provide some new evidence on the relative importance of world and country specific cycles, on their evolution over time and over business cycle phases.

Our investigation confirms some of the existing evidence. For example, as in Kose et al. (2003) or Lumsdaine and Prasad (2003), we find evidence of a significant world business cycle, despite different empirical techniques and data sets. However, it also provides new insights in the phenomenon. For instance, our results indicate that the common (world) indicator accounts for about 30% of the fluctuations in sales, industrial production, output and employment of the seven most industrialized countries; that it captures the more persistent portions of G-7 fluctuations and that it has more information than simple average or principal component measures obtained using G-7 GDPs or IPs. On the other hand, country specific indicators are useful in explaining certain GDP and employment episodes across time, but fail to track cyclical movements in the four variables over the entire sample. Perhaps more interestingly, we find that both world and country specific fluctuations are much more synchronized in contractions than expansions. Expansions tend to have idiosyncratic components, both across variables and countries, while declines in economic activity have common timing and similar dynamics, both within and across countries.

Regarding the questions of interest, we do not find evidence of structural breaks in country indicators in the 1990s. Hence, the often cited idea that national cycles are disappearing finds no support from our analysis. These indicators are as significant in explaining the differential growth rate of GDP across countries in the mid-1990s as they were in the mid-1980s and, if anything, slightly more important. We also find little support for the idea that Euro cycles are different from those of the ROW or that a Euro area cycle is emerging in the 1990s. This result should be contrasted with Lumsdaine and Prasad (2003) and Artis et al. (2003) who instead detected the presence of a EU cycle using IP data. Our analysis shows that the Euro signal is weaker when one considers a broader set of variables and that regional causes have minor explanatory power for G-7 fluctuations throughout the sample.

These set of results taken together imply that movements in the world indicator have been the stable and consistent reason for the commonalities of the fluctuations in the G-7 economies over time and that structural breaks in both the pattern of transmission across countries and in the sources of structural shocks are probably absent.

We document that the world indicator captures a variety of influences going from the magnitude of world trade, to the stance of monetary policy in the G-7 and to the spending power of consumers. Interestingly, oil shocks, US technology shocks and fiscal policy do not appear to be behind the fluctuations of the world indicator.

Our results are important from a policy point of view for several reasons. First, since variations in economic activity in countries with different institutions, economic structures and/or economic policies are driven by a common cause, international markets more than domestic policies appear to be the key to understanding the comovements in economic activity. This also means that policy institutions should probably de-emphasize national and regional cycles and instead focus on the identification of the market(s) or channel(s) that foster cross-country transmission. Second, since national or regional variables drive the national but not the world indicator, policies designed to counteract the tendencies dictated by world conditions may be ineffective. In addition, the presence of significant time variations indicates that reliance on policy actions which have been effective in the past is doomed to failure. Third, since cyclical time variations imply important asymmetries in the shape and the dynamics of international cycles, reliance on linear models in policy analyses may miss important and pervasive features of the data.

There are several interesting questions that our paper has left unanswered. For example, one would like to know more about the decreased posterior variability present in both the world and the national indicators in the late 1990s. What are the causes of this decrease? Is it the size of the shocks which has declined, the synchronization that has increased? If it is the latter, what can we say about the relative importance of contemporaneous vs. lagged transmission? Similarly, one would like to know why there is such a negative relationship between international synchronicity and volatility of GDP. Which way does causality go? What factors drive this correlation? Furthermore, since the model tracks reasonably well our four macroeconomic variables, one may want to see whether this ability translates also in useful predictions of the future state of the world economy at various horizons. The results reported in Table 1 are promising. The exercises conducted in Canova and Ciccarelli (2002) suggest that this could be the case using information available up to one or two years in advance, but more evidence is clearly needed. We plan to take up all these questions in future work.

Acknowledgments

Appendix A. Estimation

A.1. Prior information

We let $\bar{B}_i = b_i I$, $i = 1, 2, 3$, where b_i is a parameter which controls the tightness of factor i in the coefficients, and $p(\Omega^{-1}, \sigma^2, b_i, \theta_0) = p(\Omega^{-1})p(\sigma^2)p(\theta_0)\prod_i p(b_i)$ with

$$
\begin{aligned}
p(\Omega^{-1}) &= \text{Wi}(z_1, Q_1), \\
p(\sigma^2) &= \text{IG}\left(\frac{\zeta}{2}, \frac{\zeta s^2}{2}\right), \\
p(b_i) &= \text{IG}\left(\frac{\varpi_0}{2}, \frac{\delta_0}{2}\right), \\
p(\theta_0 \mid \mathscr{F}_{-1}) &= \text{N}(\bar{\theta}_0, \bar{R}_0),
\end{aligned}
\tag{11}
$$

where N stands for Normal, Wi for Wishart and IG for inverse gamma distributions, and $\mathscr{F}_{-1}$ denotes the information available at time -1. The prior for θ_0 and the law of motion for the coefficient factors imply the prior for θ_t is $p(\theta_t \mid \mathscr{F}_{t-1}) = \text{N}(\bar{\theta}_{t-1|t-1}, \bar{R}_{t-1|t-1} + B_t)$.

We collect the hyperparameters of the prior in the vector

$$
\mu = (z_1, \zeta, s^2, \varpi_0, \delta_0, \gamma_1, \gamma_2, vech(Q_1), \bar{\theta}_0, vech(\bar{R}_0)),
$$

where $vech(\cdot)$ denotes the column-wise vectorization of a symmetric matrix. We assume that the elements of μ are either known or can be estimated in the data, for example, splitting the sample into two pieces, using the first part ("training" sample) to estimate the μ and the second to estimate posterior distributions and to conduct inference. We have experimented with both informative and non-informative priors and report results obtained with the latter set of priors. The following table presents the hyperparameters values used in the two cases:

	Prior hyperparameters								
	ζ	s^2	z_1	Q_1	ϖ_0	δ_0	γ_1	γ_2	$\bar{\theta}_0$	$\bar{R}_0$
Informative	1.0	$\hat{\sigma}^2$	$NG + 20$	$\hat{Q}_1$	10	1	1	0	$\hat{\theta}_0$	I_J
Non-informative	0.0	$\hat{\sigma}^2$	$NG + 1$	$\hat{Q}_1$	0	0	1	0	$\hat{\theta}_0$	I_J

Here $\hat{\sigma}^2$ is the average of the estimated variances of NG AR(p) models, $\hat{Q}_1$ is the estimated variance–covariance of the time invariant version of (1), $\hat{\theta}_0$ is initialized with a sequential OLS on (1), over the sample 1975–1980, and J is the dimension of θ_t. The values of the remaining hyperparameters have been chosen using previous experience.

A.2. Posterior distributions

To calculate the posterior distribution of the unknowns $\psi = (\Omega^{-1}, b_i, \sigma^2, \{\theta_t\}_{t=1}^T)$, we combine the prior with the likelihood of the data, which is proportional to

$$L \propto \left(\prod_{t=1}^T \sigma_t^{-NG/2}\right)|\Omega|^{-T/2} \exp\left[-\frac{1}{2}\sum_t (Y_t - W_t \Xi \theta_t)'(\sigma_t \Omega)^{-1}(Y_t - W_t \Xi \theta_t)\right], \quad (12)$$

where $Y^T = (Y_1, \ldots, Y_T)$ denotes the data, and $\sigma_t = (1 + \sigma^2 X_t' X_t)$. Using Bayes rule, $p(\psi \mid Y^T) = (p(\psi)L(Y^T \mid \psi)/p(Y^T)) \propto p(\psi)L(Y^T \mid \psi)$. Given $p(\psi \mid Y^T)$, the posterior distribution for the components of ψ, $p(\Omega \mid Y^T)$, $p(b_i \mid Y^T)$, $p(\sigma^2 \mid Y^T)$ and $p(\{\theta_t\}_{t=1}^T \mid Y^T)$, can be obtained by integrating out nuisance parameters from $p(\psi \mid Y^T)$. Once these distributions are obtained, location and dispersion measures for ψ and for any interesting continuous function of them can be obtained.

For the model considered in this paper, it is impossible to compute $p(\psi \mid Y^T)$ analytically. However, we can numerically simulate a sample from it using Monte Carlo techniques. A method which is particularly useful in our context is the Gibbs sampler since it only requires knowledge of the conditional posterior distribution of ψ. However, while the conditional posteriors of Ω^{-1}, b_i and $\{\theta_t\}_{t=1}^T$ are available in closed form, the conditional posterior of σ^2 is not and a metropolis step within the Gibbs sampler is needed.

Denoting $\psi_{-\kappa}$ the vector ψ excluding the parameter κ, the conditional distributions of interest are

$$\theta_t \mid Y^T, \psi_{-\theta_t} \sim N(\bar{\theta}_{t|T}, \bar{R}_{t|T}), \quad t \leq T,$$

$$\Omega^{-1} \mid Y^T, \psi_{-\Omega} \sim \text{Wi}\left(z_1 + T, \left[\frac{\sum_t (Y_t - W_t \Xi \theta_t)(Y_t - W_t \Xi \theta_t)'}{\sigma_t} + Q_1^{-1}\right]^{-1}\right),$$

$$b_i \mid Y^T, \psi_{-b_i} \sim \text{IG}\left(\frac{\varpi^i}{2}, \frac{\sum_t (\theta_t^i - \theta_{t-1}^i)'(\theta_t^i - \theta_{t-1}^i) + \delta_0}{2\xi_t}\right),$$

$$\sigma^2 \mid Y^T, \psi_{-\sigma^2} \propto L(Y^T \mid \psi) \times p(\sigma^2), \quad (13)$$

where $\bar{\theta}_{t|T}$ and $\bar{R}_{t|T}$ are the one-period-ahead forecasts of θ_t and the variance–covariance matrix of the forecast error, respectively, calculated with the Kalman smoother, as described in Chib and Greenberg (1995), and $\varpi^1 = T + \varpi_0$, $\varpi^2 = Tg + \varpi_0$ and $\varpi^3 = TN + \varpi_0$.

The posterior for σ^2 is simulated using a random walk metropolis algorithm, where, at each iteration l, we generate candidate draws according to $(\sigma^2)^* = (\sigma^2)^{(l-1)} + z$, where z is a normal random variable with mean zero and variance chosen to ensure that the acceptance rate is approximately 0.2–0.4.

Under regularity conditions (see Geweke, 2000), cycling through the conditional distributions in (13) will produce in the limit draws from the joint posterior of interest. From these, marginal distributions can be computed averaging over draws nuisance dimensions. In particular, using the draws, the posterior distributions of λ_t and α_t can be estimated using kernel methods and, in turns, the posterior distributions of WLI_t and CLI_t can be obtained. For example, the posterior mean of WLI_t can be approximated by $(1/H)\sum_h \mathscr{W}_t \lambda_t^h$ and a credible 68% interval is obtained ordering the draws of WLI_t^h for each t and taking the 16th and the 84th percentile of the distribution.

Because we are not directly sampling from the posterior, it is important to monitor that the Markov chain induced by the sampler converges to the ergotic (posterior) distribution.

We have checked convergence in several ways: increasing the length of the chain, splitting the chain into two after a burn-in period and calculating whether the mean and the variances are similar; checking if cumulative means settle at some value. The result we present are based on chains with 24 000 draws: 600 blocks of 40 draws were made and the last draw for each block is retained after the discarding the first 4000. This means that a total of 500 draws is used at each t to conduct posterior inference.

Appendix B. Specification searches

We have conducted a number of preliminary checks on the model to examine whether all the features included in the specification are really needed to capture the dynamics of the variables under consideration. In particular, we have examined whether the factorization (3) is exact, whether the presence of international lagged interdependencies and of time variations in the coefficients are necessary to capture the dynamics of the variables, and whether we can exclude the variable indicators (ρ_t) from the factorization of δ_t. Verification of the first three hypotheses is important because the specification can be considerably simplified if the factorization is exact, interdependencies are absent and time variations unimportant.

To check whether the factorization is exact we have examined the marginal likelihoods of the benchmark model (M_0) and of a model restricted to have $\sigma^2 = 0$ (denoted by M_3). In addition, we have visually examined whether the posterior density of σ^2 is more concentrated around zero
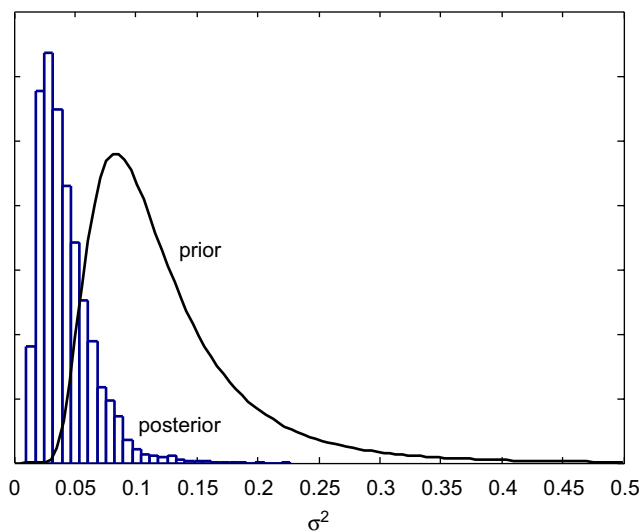

Fig. 8. Prior and posterior densities of σ^2.

Table 3
Log marginal likelihood of models

Sample	M_0	M_1	M_2	M_3	M_4	M_5	M_6
1980–2002	116.4	−355.1	151.3	158.6	−53.7	−582.8	−182.1
1980–1992	−68.2	−339	−13.8	55.4	−145.3	−457.2	−216.6
1993–2002	166.9	30.8	168.5	169.5	126.1	−71.5	58.9

than is the prior (see Fig. 8). Both the Bayes factor and Fig. 8 favor the model where the factorizations is exact. For example, the log of the Bayes factor of the two specifications is -42.2 while the prior and posterior mean of σ^2 are, respectively, 0.125 and 0.044.

To measure the importance of interdependencies we compare the marginal likelihood of model M_3 and the marginal likelihood of a model where interdependencies are excluded (denoted by M_4). It turns out that the log of the Bayes factor exceeds 200, overwhelmingly indicating that lagged interdependencies matter at least as contemporaneous ones in generating common fluctuations in the G-7. Hence, there is important information in the lags of the variables which is neglected, e.g., when standard static factor approaches are used.

The presence of time variations in the factors is examined comparing the marginal likelihood of model M_3 and the marginal likelihood of a model where $\gamma_1 = \gamma_2 = 0$ (denoted by M_5). The log of the Bayes factor exceeds 600 therefore indicating that time variations in the coefficients are an important feature of our data. A deeper look into the specification reveals other interesting facts. First, time variations become less significant in the second part of the sample (the log of Bayes factor is 'only' around 200). Second, though important, time variations, as measured by the size of B_t are relatively small, in close agreement with Del Negro and Otrok (2003). Third, the world factor λ_t displays more significant movements over the sample than the other two components. In fact, the posterior distributions of b_1, b_2 and b_3 are centered around 0.01, 0.0062 and 0.0045, respectively (the prior mean of all three parameters is 0.125).

Finally, we have also examined the performance of M_3 against a model where the coefficient vector excludes variable specific factors (denoted by M_6). In this case the specification with three factors is preferred by a large amount (the log of the Bayes factor is 340.1). Interestingly, the marginal likelihood of M_6 dramatically increases over the 1990s. Hence, over this sample, variable specific dynamics seem to have become less crucial to explain the business cycle characteristics of G7 countries. This result is comforting since sales, employment, industrial production were partially chosen because coincident with GDP not only in the US (as the NBER practice indicates) but also in the other G-7 countries. Table 3 reports a summary by sample periods of the logarithm of $f(Y|M_i)$ for all models described above.

Appendix C. Data

The sources of the data and their acronyms and definition used in the paper are listed in Table 4.

Table 4
Data used: definitions and sources

Acronym	Definition	Source
uspin	US real personal non-agricultural income, SA, quarterly growth rates	Bureau of Economic Analysis
cpg	World commodity PRICE index, SA, quarterly growth rates	IMF, International Financial Statistics
usstock	NYSE stock prices index, SA, quarterly growth rates	Datastream
oil	World market crude petroleum PRICE, SA, quarterly growth rates	IMF, International Financial Statistics

Table 4 (*continued*)

Acronym	Definition	Source
wtrade	World goods trade, SA, quarterly growth rates	OECD, Main Economic Indicators
gdpg	Real GDP, SA, quarterly growth rates	OECD, Quarterly National Accounts
ipg	Real industrial production, SA, quarterly growth rates	OECD, Quarterly National Accounts
eg	Employment, SA, quarterly growth rates	OECD, Quarterly National Accounts
saleg	Real retail sales, SA, quarterly growth rates	OECD, Quarterly National Accounts
cy	Real private consumption over real GDP, SA, levels	OECD, Quarterly National Accounts
ree	Real effective exchange rate, SA, quarterly growth rates	OECD, Main Economic Indicators
glvtech	Quarterly US technology shock	Gali et al. (2003) 723–743, Fig. 3
si	Three month interbank nominal interest rate (various sources)	G7 National Central Banks
sri	si minus 4 times quarterly CPI inflation rate (various sources)	G7 National Central Banks
li	Nominal yield on 10 year bond (various sources)	G7 National Central Banks
lir	li minus 4 times quarterly CPI inflation rate (various sources)	G7 National Central Banks
avglid	Avg of G7 (excl US) long-term nominal interest rate spreads	G7 National Central Banks
avglir	Avg of G7 real li, i.e., li-annualized quarterly CPI inflation	G7 National Central Banks
m3g	M3 (various sources), SA, quarterly growth rates	G7 National Central Banks
def	General government deficit, current prices, SA, levels	IMF, International Financial Statistics

References

Artis, M., Marcellino, M., Proietti, T., 2003. Dating the euro area business cycle. CEPR Working Paper No. 3696.

Backus, D., Kehoe, P., Kydland, F., 1995. International business cycles: theory and evidence. In: Cooley, T.F. (Ed.), Frontiers of Business Cycle Research. Princeton University Press, Princeton NJ.

Baxter, M., 1995. International trade and business cycle. In: Grossman, G., Rogoff, K. (Eds.), Handbook of International Economics. North-Holland, Amsterdam, pp. 1801–1864.

Bernanke, B., Boivin, L., Eliasz, P., 2003. Measuring the effects of monetary policy. NBER Working Paper 10220.

Binder, M., Hsiao, C., Pesaran, H., 2000. Estimation and inference in short panel vector autoregressions with unit roots and cointegration. Manuscript, University of Maryland.

Bry, G., Boschan, C., 1971. Cyclical analysis of time series: selected procedures and computer programs. NBER, New York.

Canova, F., 2005. The transmission of US shocks to Latin America. Journal of Applied Econometrics 20, 229–251.

Canova, F., Ciccarelli, M., 2002. Bayesian panel VAR: specification, estimation, forecasting and leading indicators. CEPR Working Paper No. 4033.

Canova, F., Dellas, H., 1992. Trade interdependence and the international business cycle. Journal of International Economics 34, 23–47.

Canova, F., Marrinan, J., 1998. Sources and propagations of international output cycles: common shocks or transmission? Journal of International Economics 46, 133–166.

Chib, S., 1995. Marginal likelihood from the Gibbs output. Journal of the American Statistical Association 90, 1313–1321.

Chib, S., Greenberg, E., 1995. Hierarchical analysis of SUR models with extensions to correlated serial errors and time-varying parameter models. Journal of Econometrics 68, 409–431.

Cochrane, J., 1994. A simple test of consumption insurance. Journal of Political Economy 99, 957–976.

Del Negro, M., 2002. Asymmetric shocks among US states. Journal of International Economics 56, 273–297.

Del Negro, M., Otrok, C., 2003. Dynamic factor models with time-varying parameters: measuring the evolution of the European business cycle. Manuscript, Federal Reserve Bank of Atlanta.

Doan, T., Litterman, R., Sims, C., 1984. Forecasting and conditional projections using realist prior distributions. Econometric Review 3, 1–100.

Doyle, B.M., Faust, J., 2002. An investigation of the co-movements among growth rates of the G-7 countries. Federal Reserve Bulletin (October), 427–437.

Forni, M., Reichlin, L., 1998. Let's get real: a factor analytic approach to disaggregated business cycle dynamics. Review of Economic Studies 65, 453–473.

Gali, J., López-Salido, D., Vallés, J., 2003. Technology shocks and monetary policy: assessing the Fed's performance. Journal of Monetary Economics 50, 723–743.

Geweke, J., 2000. Simulation based Bayesian inference for economic time series. In: Mariano, R., Schuermann, T., Weeks, M. (Eds.), Simulation Based Inference in Econometrics: Methods and Applications. Cambridge University Press, Cambridge.

Gregory, A., Head, A., Raynauld, J., 1997. Measuring world business cycles. International Economic Review 38, 677–702.

Harding, D., Pagan, A., 2006. Synchronization of cycles. Journal of Econometrics 132, 59–79.

Helbling, T., Bayoumi, T., 2003. Are they all in the same boat? The 2000–2001 growth slowdown and the G-7 business cycle linkages. IMF Working Paper, 03-46.

Holtz–Eakin, D., Newey, W., Rosen, H., 1988. Estimating vector autoregressions with panel data. Econometrica 56, 1371–1395.

Kadiyala, K.R., Karlsson, S., 1997. Numerical methods for estimation and inference in Bayesian-VAR models. Journal of Applied Econometrics 12, 99–132.

Kose, M., Otrok, C., Whiteman, C., 2003. International business cycles: world, region and country specific factors. American Economic Review 93, 1216–1239.

Lumsdaine, R., Prasad, E., 2003. Identifying the common component in international business cycles. Economic Journal 113, 101–127.

Mackowiak, B., 2003. External shocks, US monetary policy, macroeconomic fluctuations in emerging markets. Manuscript, Humboldt University.

Peersman, G., 2005. What caused the early millennium slowdown? Evidence based on vector autoregressions. Journal of Applied Econometrics 20, 185–207.

Perri, F., Heathcothe, J., 2001. Financial globalization and real regionalization. Journal of Economic Theory 119, 207–243.

Pesaran, M.H., Shuermann, T., Weiner, S.M., 2004. Modelling regional interdependencies using a global error-correcting macroeconometric model. Journal of Business and Economic Statistics 22, 129–162.

Rose, A., Frankel, J., 2000. Estimating the effects of currency unions on trade and output. NBER Working Paper No. 7557.

Sargent, T., 1989. Two models of measurement and the investment accelerator. Journal of Political Economy 97, 251–287.

Stock, J., Watson, M., 1989. New indices of coincident and leading indicators. In: Blanchard, O., Fisher, S. (Eds.), NBER Macroeconomic Annual, vol. 4.

Stock, J., Watson, M., 2003. Understanding changes in international business cycle dynamics. NBER Working Paper No. 9859.