Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Econometrics 120 (2004) 327–359
JOURNAL OF Econometrics
www.elsevier.com/locate/econbase

Forecasting and turning point predictions in a Bayesian panel VAR model

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

[a] *Departament d'Economia i Empresa, Universitat Pompeu Fabra, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain*
[b] *CEPR, London, UK*
[c] *Fundamentos del Análisis Económico, Universidad de Alicante, Campus de S. Vicente del Raspeig, 03690 Alicante, Spain*

Accepted 30 June 2003

Abstract

We provide methods for forecasting variables and predicting turning points in panel Bayesian VARs. We specify a flexible model, which accounts for both interdependencies in the cross section and time variations in the parameters. Posterior distributions for the parameters are obtained for hierarchical and for Minnesota-type priors. Formulas for multistep, multiunit point and average forecasts are provided. An application to the problem of forecasting the growth rate of output and of predicting turning points in the G-7 illustrates the approach. A comparison with alternative forecasting methods is also provided.

JEL classification: C11; C15; E32; E37

Keywords: Forecasting; Turning points; Bayesian methods; Panel VAR; Markov chain Monte Carlo methods

1. Introduction

Panel VAR models have become increasingly popular in macroeconomics to study the transmission of shocks across countries (Ballabriga et al., 1998), the propagation effects of monetary policy in the European Union (Gerlach and Smets, 1996) and the average differential response of developed and underdeveloped countries to domestic and external disturbances (Hoffmaister and Roldós, 1997; Rebucci, 1998). At the same time, recent developments in computer technology have permitted the estimation

* Corresponding author. Tel.: +49-69-1344-8721; fax: +49-69-1344-6575.
E-mail addresses: fabio.canova@upf.es (F. Canova), matteo.ciccarelli@ecb.int (M. Ciccarelli).

0304-4076/$ - see front matter
doi:10.1016/S0304-4076(03)00216-1

of increasingly complex multicountry VAR models in reasonable time, making them potentially usable for a variety of forecasting and policy purposes.

Despite this interest, the theory for panel VAR is somewhat underdeveloped. After the works of Chamberlain (1982, 1984), Holtz-Eakin et al. (1988) and Binder et al. (2000), who specify panel VAR models for micro-data, to the best of our knowledge only Pesaran and Smith (1996), Canova and Marcet (1997) and Hsiao et al. (1999) have considered problems connected with the specification and the estimation of (univariate) dynamic macro-panels. Garcia Ferrer et al. (1987), Zellner and Hong (1989), Zellner et al. (1991), on the other hand, have provided Bayesian shrinkage estimators and predictors for similar models. In general, a researcher focuses on the specification

$$y_{it} = A(L)y_{it-1} + \varepsilon_{it}, \tag{1}$$

where y_{it} is a $G \times 1$ vector, $i = 1,\ldots,N$; $A(L)$ is a matrix in the lag operator; $\varepsilon_{it} = \alpha_i + \delta_t + u_{it}$, where δ_t is a time effect; α_i is a unit specific effect and u_{it} a disturbance term. In some cases a specification with time varying slope coefficients and a fixed effect is used (see e.g. Holtz-Eakin et al., 1988). Two main restrictions characterize this specification. First, it assumes common slope coefficients. Second, it does not allow for interdependencies across units. With these restrictions, the interest is typically in estimating the *average* dynamics in response to shocks (the matrix $A(L)$).

Canova and Marcet, Pesaran and Smith and others, instead, use a univariate dynamic model of the form

$$y_{it} = \alpha_i + \rho_i y_{it-1} + x_{it}'\beta_i + v_t'\delta_i + \varepsilon_{it}, \tag{2}$$

where y_{it} is a scalar, x_{it} is a set of k exogenous unit specific regressors, v_t is a set of h exogenous regressors common to all units while ρ_i, β_i and δ_i are unit specific vectors of coefficients. Two restrictions are implicit also in this specification. First, no time variation is allowed in the parameters. Second, there are no interdependencies either among different variables within units or among the same variable across units.

In this paper we relax these restrictions and study the issues of specification, estimation and forecasting in a macro-panel VAR model. Our point of view is Bayesian. Such an approach has been widely used in the VAR literature since the works of Doan et al. (1984), Litterman (1986), and Sims and Zha (1998) and provides a convenient framework where one can allow for both interdependencies and meaningful time variations in the coefficients. The specification we consider has the general form

$$y_{it} = A_{it}(L)Y_{t-1} + \varepsilon_{it}, \tag{3}$$

where $Y_s (s < t)$ is a $GN \times 1$ vector (with G variables for each unit $i = 1,\ldots,N$). Because coefficients vary across units and along time, estimation of the parameters is impossible without imposing restrictions. Instead of constraining the coefficients to be the same across units, we assume that they are random and a prior distribution on $A_{it}(L)$ is introduced. We decompose the parameter vector into two components, one which is unit specific and one which is time specific. We specify a flexible prior on these two

components which parsimoniously accounts for interdependencies in the cross section and for time variations in the evolution of the parameters. The prior shares features with those of Lindley and Smith (1972), Doan et al. (1984) and Hsiao et al. (1999) and has a hierarchical structure, which allows for various degrees of ignorance in the researcher's information about the parameters.

Bayesian VARs are known to produce better forecasts than unrestricted VAR and, in many situations, ARIMA or structural models (Canova, 1995 for references). By allowing interdependencies and some degree of information pooling across units, we introduce an additional level of flexibility which may improve the forecasting ability of these models.

We describe in detail two situations of interest: one with fully hierarchical priors and one Minnesota-type priors. In the former case, a Markov Chain Monte Carlo method (the Gibbs sampler) is employed to calculate posterior distributions. Such an approach is useful in our setup since it exploits the recursive features of the posterior distribution. For the Minnesota-type prior, unknown parameters are estimated using the predictive density and posterior estimates are obtained by plugging-in our estimates in the relevant formulas in an empirical Bayes fashion. We provide recursive formulas for multistep, multiunit forecasts, consistent with the information available at each point in time using the posterior of the parameters or the predictive density of future observations. The latter is also used to compute turning point probabilities.

To illustrate the features of the proposed approach, we apply the methodology to the problem of predicting output growth and of forecasting turning points in output growth in the G-7 and computing the probability of a recession in the US. To evaluate the forecasting performance we provide an extensive comparison with other specifications suggested in the literature. The results indicate that our approach improves the forecasting performance of existing univariate and simple BVAR models, both at the one and at the four steps horizons. The improvements are of the order of 5–15% when the Theil-U is used and about 2–8% when the MAD is used. The forecasting performance of our specification is also preferable to the one of a BVAR model which mechanically extends the Litterman prior to the panel case. In terms of turning point predictions, the three versions of the panel approach we consider are able to recognize about 80% of turning points and they turn out to be the best for this task, along with Zellner's g-prior shrinkage approach. The simple extension of the Litterman's prior to the panel case does poorly along this dimension and it is the second worst among all the procedures employed. Finally, we show that the method is competitive with the best specifications in predicting the downturn in US economic activity occurred in 1990:3 when using the information available in 1988:4, a turning point which was missed by many commercial and government forecasting agencies. Depending on the specification, our approach suggests that downturn at that date occurs with 30–57% probability.

The rest of the paper is organized as follows. The next section gives the general model specification and the assumptions made. Section 3 provides general formulas for the posterior. Section 4 sets up the prior and discusses the computational issues involved. Section 5 describes formulas for multistep, multiunits forecasts. Section 6 contains the application. Section 7 concludes.

2. General specification

The statistical (reduced form) model we use is of the form

$$y_{it} = \sum_{j=1}^{N} \sum_{l=1}^{p} b_{it,l}^{j} y_{jt-l} + d_{it} v_t + u_{it}, \tag{4}$$

where $i = 1, \ldots, N$; $t = 1, \ldots, T$; y_{it} is a $G \times 1$ vector for each i, $b_{it,l}^{j}$ are $G \times G$ matrices, d_{it} is $G \times q$, v_t is a $q \times 1$ vector of exogenous variables, common to all units, and u_{it} is a $G \times 1$ vector of random disturbances. Here p is the number of lags, G the number of endogenous variables and q the number of exogenous variables (including a constant).

The generality of (4) comes from at least two features. First, the coefficients are allowed to vary both across units and across time. Second, there are interdependencies among units whenever $b_{it,l}^{j} \neq 0$ for $j \neq i$ and for any l. Both features constitute the main difference with the literature (Holtz-Eakin et al., 1988; Rebucci, 1998) that considers panel VAR models. It is easy to verify that if $d_{it} v_t = a_t$, $b_{it} = b_t \ \forall i$, $u_{it} = \psi_t f_i + \xi_{it} b_{it,l}^{j} = 0$, $j \neq i$, $\forall l$, our specification collapses to the one used by Holtz-Eakin et al. (1988). We rewrite (4) in a stacked regression manner

$$Y_t = W_t \gamma_t + U_t, \tag{5}$$

where $W_t = I_{NG} \otimes X_t'$; $X_t = (y_{t-1}', y_{t-2}', \ldots, y_{t-p}', v_t')'$; $\gamma_t = (\gamma_{1t}', \ldots, \gamma_{Nt}')'$ and $\gamma_{it} = (\beta_{it}^{1\prime}, \ldots, \beta_{it}^{G\prime})'$. Here $y_s (s < t)$ is an $NG \times 1$ vector, β_{it}^{g} are $k \times 1$ vectors, $k = NGp + q$, containing, stacked, the g rows of the coefficient matrices b_{it} and d_{it}, while Y_t and U_t are $NG \times 1$ vectors containing the endogenous variables and the random disturbances of the model.

If γ_{it} are different for each cross-sectional unit in different time periods, there is no way to obtain meaningful estimates of them using classical methods. One possibility is to view each coefficient vector as random with a given probability distribution. We make the following assumptions:

1. For each i, the $Gk \times 1$ vector γ_{it} has a time invariant and a time varying component:

$$\gamma_{it} = \alpha_i + \lambda_{it}. \tag{6}$$

2. For each i, the $Gk \times 1$ vector α_i is normally distributed

$$\alpha_i \sim \mathrm{N}(R_i \bar{\alpha}, \Delta_i), \tag{7}$$

where $R_i = I_G \otimes E_i$, $\Delta_i = V \otimes E_i \Omega_1 E_i$, the $G \times G$ matrix V and the $k \times k$ matrix Ω_1 are symmetric and positive semidefinite and E_i is a $k \times k$ matrix that commutes the k coefficients of unit i for each of the G equations with those of unit one. We assume $cov(\alpha_i, \alpha_j) = 0 \ \forall i \neq j$.

3. The mean vector $\bar{\alpha}$ is assumed to have a normal distribution

$$\bar{\alpha} \sim \mathrm{N}(\mu, \Psi). \tag{8}$$

4. For each i, $\lambda_{it} = R_i\lambda_t$, with λ_t independent of α_i. The $Gk \times 1$ vector λ_t evolves according to

$$\lambda_t = B\lambda_{t-1} + (I - B)\lambda_0 + e_t, \tag{9}$$

where $B = \rho * I_{Gk}$ and, conditional on U_t and W_t, $e_t \sim \mathrm{N}(0, V \otimes \Omega_{2t})$, $\Omega_{2t} = v_1\Omega_{2t-1} + v_2\Omega_{20}$ and Ω_{20} is a positive semidefinite, symmetric matrix. The initial conditions are such that $\lambda_0 \sim \mathrm{N}(\tilde{\lambda}_0, V \otimes \Omega_{20})$.

5. Conditional on W_t, the vector of random disturbances U_t has a normal distribution

$$U_t \sim \mathrm{N}(0, \Sigma_u). \tag{10}$$

where $\Sigma_u = \Sigma \otimes H$, Σ is $N \times N$ and H is $G \times G$, both positive definite and symmetric matrices.

Given the previous assumptions, the structure of the model can be summarized with the following a priori hierarchical scheme

$$\begin{aligned} &Y_t \mid F_t, \alpha, \lambda_t \sim \mathrm{N}(W_t\alpha + Z_t\lambda_t, \Sigma_u),\\ &\alpha \mid F_t \sim \mathrm{N}(S_N\bar{\alpha}, \Delta),\\ &\bar{\alpha} \mid F_t \sim \mathrm{N}(\mu, \Psi),\\ &\lambda_t \mid F_t \sim \mathrm{N}(\hat{\lambda}_{t|t-1}, \hat{\Omega}_{t|t-1}), \end{aligned} \tag{11}$$

where F_t is the information set at t (which includes Y_0, the presample information, and W_t); $S_N = diag\{R_i\}$; $Z_t = W_tS_N$; $\Delta = diag(\Delta_1, \ldots, \Delta_n)$, $\hat{\lambda}_{t|t-1} = B\hat{\lambda}_{t-1|t-1} + (I - B)\hat{\lambda}_0$; $\hat{\Omega}_{t|t-1} = B\hat{\Omega}_{t-1|t-1}B' + (I - B)(V \otimes \Omega_{20})(I - B)' + V \otimes \Omega_{2,t|t-1}$, and the notation $t|t-1$ indicates values at t predicted with information at $t-1$.

Assumptions 1–4 decompose the parameters vector for each i in 2 components: one is unit specific and constant over time; the other is common across units but varies with time. The prior possibility for time-variation increases the flexibility of the specification and provides a general mechanism to account for structural shifts without explicitly modelling the source of the shift. The fact that the time-varying parameter vector is common across units does not prevent unit-specific structural shifts, since γ_{it} can be re-written as

$$\gamma_{it} = (I - B)(\alpha_i + \lambda_{i0}) + (I - B)\gamma_{it-1} + e_{it}, \tag{12}$$

where unit specific variations of time occur through the common coefficient vector B.

Assumptions 2 and 3 can be used to recover the vector α or the mean coefficient vector $\bar{\alpha}$. In this sense, we can distinguish between "fixed" and "random" effects, following Lindley and Smith (1972). By fixed effects we mean the estimation of the vector γ_{it}, while the term random effects refers to the estimation of $\bar{\gamma}_t = \bar{\alpha} + \lambda_t$. For example, in a VAR without interdependencies (i.e. $b^j_{it,l} = 0$, $j \neq i$), we may be more interested in the relationships among the variables of the system for a "typical" unit, in which case interest centers in the estimation of the random effect $\bar{\gamma}_t$. If, instead, we are interested in the relationships across units, for example, wishing to find the effect of a shock in the g variable of unit j on the variables of unit i, we better estimate γ_{it}

for each i. In the context of forecasting, we may be concerned with point prediction using the average vector $\bar{\gamma}_t$ or in predicting future values of the variables of interest using information available for each unit. Assumption 2 allows for some degree of a priori pooling of cross-sectional information via an exchangeable prior on α_i. This may be useful in a panel when there are similarities in the characteristics of the vector of variables across units. In this case coefficients of other units may contain useful information for estimating the coefficients of unit i.

The structure underlying assumption 4 is similar to Canova (1993). There it is shown that (9) allows for nonlinearities in the moment structure (of both ARCH-M and Markov switching type) and nonnormalities in the time series under consideration. Note that the common component evolves over time with an heteroschedastic structure (for homoskedastic variations set $v_1 = 0$). Besides being useful to directly capture generic volatility clustering which are common across countries, time variations in the variance allow the model to quickly adapt when outliers or regime switches of short mean length are present.

The assumed Kronecker structure for the variance–covariance matrices is computationally convenient and allows us to nest interesting hypothesis. For instance, when $\Omega_1 = 0$, there is no heterogeneity in the cross-sectional dimension of the panel. If $B = I_{Gk}$, coefficients evolve over time as a random walk, while when $B = I_{Gk}$ and $v_1 = v_2 = 0$, the model reduces to a standard dynamic panel model with no time-variation in the coefficient vector. Moreover, when $V = 0$ neither heterogeneity nor time variation are present in the model. Finally, a single country VAR with fixed coefficients can be obtained by setting $b^j_{it,l} = 0$, $\forall j \neq i$, $\forall l$ and letting Ω_1, Ψ, Σ_t go to zero. Note that, with the Kronecker structure, the prior is fully symmetric in the sense that it is the same regardless of the variables and of the units we are considering. There are both notational and computation advantages in setting the prior this way. However, when the N units display scale and transmission differences, one may want to relax this assumption and let, e.g., $\Delta_i = V \otimes \Omega_{i1}$ while leaving the prior distributions for $\bar{\alpha}$ and λ_t are unchanged.

Stock and Watson (2002), Forni et al. (2000) and others have examined macro-panel models where either N or G or both are large. Their approach is to setup the problem so that it can be handled in the context of (dynamic) index models with classical methods. From (11) one can see that also our specification has an index structure, where the two indices we consider are a "common" one and a "time specific" one. Two major features differentiate our approach from theirs: first, the coefficients on our indices are allowed to vary over time; second, their inferential methods require asymptotic approximations, while our approach delivers exact estimates even when N or G are small.

3. Posterior estimates

3.1. Fixed effects model

From (11) the likelihood function is

$$L(Y_t|\gamma_t,\ F_t) = \mathrm{N}(W_t\alpha + Z_t\lambda_t,\ \Sigma_u)$$

and the prior, given information up to t, is

$$p(\gamma_t|F_t) = \mathrm{N}(\hat{\gamma}_{t-1},\ \hat{H}_{t-1}), \tag{13}$$

where $\hat{\gamma}_{t-1} = S_N\left(\mu + \hat{\lambda}_{t|t-1}\right)$ and $\hat{H}_{t-1} = (S_N \Psi S_N' + \Delta) + S_N \hat{\Omega}_{t|t-1} S_N'$. Standard calculations give us that the posterior $\pi_0(\gamma_t|F_t,\ Y_t)$ is normal with mean γ_t^* and variance H_t^* where

$$\begin{aligned}\gamma_t^* &= H_t^*\left(W_t'\Sigma_u^{-1}Y_t + \hat{H}_{t-1}^{-1}\hat{\gamma}_{t-1}\right)\\ &= \hat{\gamma}_{t-1} + \hat{H}_{t-1}W_t'\left[W_t\hat{H}_{t-1}W_t' + \Sigma_u\right]^{-1}(Y_t - W_t\hat{\gamma}_{t-1}),\end{aligned} \tag{14}$$

$$\begin{aligned}H_t^* &= \left[\hat{H}_{t-1}^{-1} + W_t'\Sigma_u^{-1}W_t\right]^{-1}\\ &= \hat{H}_{t-1} - \hat{H}_{t-1}W_t'\left[W_t\hat{H}_{t-1}W_t' + \Sigma_u\right]^{-1}W_t\hat{H}_{t-1}.\end{aligned} \tag{15}$$

In the second expression, posterior estimates are obtained recursively, given $\hat{\gamma}_0$ and $\hat{H}_0$.

In some cases one may want to obtain posterior distributions of α and λ_t separately. It is straightforward to show that

$$\begin{pmatrix}\alpha\\ Y_t\end{pmatrix}\Bigg| F_t \sim \mathrm{N}\left[\begin{pmatrix}S_N\mu\\ Z_t(\mu + \hat{\lambda}_{t|t-1})\end{pmatrix}, \begin{pmatrix}\phi_{11} & \phi_{12}\\ \phi_{21} & \phi_{22}\end{pmatrix}\right],$$

where $\phi_{11} = (S_N \Psi S_N' + \Delta)$; $\phi_{12} = \phi_{11}W_t'$; $\phi_{21} = W_t\phi_{11}$; $\phi_{22} = W_t\phi_{11}W_t' + Z_t\hat{\Omega}_{t|t-1}Z_t' + \Sigma_u$.

Using the properties of multivariate normal distributions, the conditional marginal $\pi_1(\alpha|F_t,\ Y_t)$ is normal with mean $\alpha^* = S_N\mu + \phi_{12}\phi_{22}^{-1}\left[Y_t - Z_t\left(\mu + \hat{\lambda}_{t|t-1}\right)\right]$ and variance $V_\alpha^* = \phi_{11} - \phi_{12}\phi_{22}^{-1}\phi_{21}$ and the conditional marginal $\pi_2(\lambda_t|Y_t,\ F_t)$ is normal with mean $\lambda_t^* = \hat{\lambda}_{t|t-1} + \hat{\Omega}_{t|t-1}Z_t'\phi_{22}^{-1}\left[Y_t - Z_t\left(\mu + \hat{\lambda}_{t|t-1}\right)\right]$ and variance $\Omega_t^* = \hat{\Omega}_{t|t-1} - \hat{\Omega}_{t|t-1}Z_t'\phi_{22}^{-1}Z_t\hat{\Omega}_{t|t-1}$.

3.2. Random effects model

When interest centers on the estimation of the mean vector $\bar{\gamma}_t$, we rewrite the model as

$$Y_t = Z_t\bar{\gamma}_t + \eta_t, \tag{16}$$

where $\bar{\gamma}_t = \bar{\alpha} + \lambda_t$ and $\eta_t = u_t + W_t v$. Standard manipulations give us that the posterior $\pi_3(\bar{\alpha}\,|\,Y_t,\ F_t) \sim \mathrm{N}(\bar{\alpha}^*, \Psi^*)$ where

$$\bar{\alpha}^* = \mu - \Psi Z_t'\left[Z_t\left(\Psi + \hat{\Omega}_{t|t-1}\right)Z_t' + \Sigma_u + W_t\,\Delta\,W_t'\right]^{-1}\left[Y_t - Z_t\left(\mu + \hat{\lambda}_{t|t-1}\right)\right], \tag{17}$$

$$\Psi^* = \Psi - \Psi Z_t'\left[Z_t\left(\Psi + \hat{\Omega}_{t|t-1}\right)Z_t' + \Sigma_u + W_t\,\Delta\,W_t'\right]^{-1}Z_t\Psi, \tag{18}$$

while the posterior $\pi_2(\lambda_t \mid Y_t,\ F_t) \sim \mathrm{N}(\lambda_t^*, \Omega_t^*)$ and λ_t^* and Ω_t^* are the same as before. Hence the posterior $\pi_4(\vec{\gamma}_t \mid Y_t, F_t) \sim \mathrm{N}(\vec{\gamma}_t^*, H_t^*)$ where

$$\vec{\gamma}_t^* = \left(\mu + \hat{\lambda}_{t|t-1}\right) + \left(\Psi + \hat{\Omega}_{t|t-1}\right) Z_t' \left[Z_t \left(\Psi + \hat{\Omega}_{t|t-1}\right) Z_t' + \Sigma_u + W_t\, \Delta W_t'\right]^{-1} \times \left[Y_t - Z_t \left(\mu + \hat{\lambda}_{t|t-1}\right)\right], \tag{19}$$

$$H_t^* = \left(\Psi + \hat{\Omega}_{t|t-1}\right) - \left(\Psi + \hat{\Omega}_{t|t-1}\right) Z_t' \left[Z_t \left(\Psi + \hat{\Omega}_{t|t-1}\right) Z_t' + \Sigma_u + W_t\ \Delta W_t'\right]^{-1} \times Z_t \left(\Psi + \hat{\Omega}_{t|t-1}\right). \tag{20}$$

4. Setting up the priors

The formulas derived in the previous section are operational only if the vector

$$\zeta = (vec(\mu), vec(\tilde{\lambda}_o), vech(\Sigma_u), vech(\Omega_{20}), vech(B), vech(\Psi), vech(\Delta))$$

is known.

When this is the case, to obtain marginal posteriors we need to integrate nuisance parameters out of the joint posterior density. This integration, in general, is difficult, even with brute force numerical methods, given the large number of parameters contained in ζ. In this section we describe two approaches which make the computation of the posterior feasible.

4.1. Informative priors

When the prior for the parameter vector is informative, the posterior distribution does not have an analytical closed form. Nevertheless, we can implement a hierarchical Bayes analysis using a sampling-based approach, such as the Gibbs sampler, (see e.g. Geman and Geman, 1984; Gelfand and Smith, 1990; Gelfand et al., 1990 among others).

The basic idea of the approach is to construct a (computable) Markov chain on a general state space such that the limiting distribution of the chain is the joint posterior of interest. Suppose we have a parameter vector ϑ with k components $(\vartheta_1, \vartheta_2, \ldots, \vartheta_k)$ and that the posterior distributions $\pi(\vartheta_j \mid \vartheta_s, s \neq j)$ are available. Then the algorithm works as follows. We start from arbitrary values for $\vartheta_1^{(o)}, \vartheta_2^{(o)}, \ldots, \vartheta_k^{(o)}$. Setting $i = 1$, we cycle through the conditional distributions sampling $\vartheta_1^{(1)}$ from $\pi\left(\vartheta_1 \mid \vartheta_2^{(o)}, \ldots, \vartheta_k^{(o)}\right)$, $\vartheta_2^{(1)}$ from $\pi\left(\vartheta_2 \mid \vartheta_1^{(1)}, \ldots, \vartheta_k^{(o)}\right)$, up to $\vartheta_k^{(1)}$ from $\pi\left(\vartheta_k \mid \vartheta_1^{(1)}, \ldots, \vartheta_{k-1}^{(1)}\right)$. Next, we set $i{=}i{+}1$ and repeat the cycle. After iterating on this cycle, say, M times, the sample value $\vartheta^{(M)} = \left(\vartheta_1^{(M)}, \vartheta_2^{(M)}, \ldots, \vartheta_k^{(M)}\right)$ can be regarded as a drawing from the true joint posterior density. Once this simulated sample has been obtained, any posterior moment of interest or marginal density can be estimated, using the ergodic theorem. Convergence to the desired distribution can be checked as suggested in Gelfand and Smith (1990).

In order to apply the Gibbs sampler to our panel VAR model, we need to specify prior information so that the conditional posterior distribution for components of the parameter vector can be obtained analytically. Recall that our hierarchical model is given by

$$Y_t = W_t\alpha + Z_t\lambda_t + u_t,$$

$$\alpha_i = R_i\bar{\alpha} + \varepsilon_i,$$

$$\bar{\alpha} = \mu + v,$$

$$\lambda_t = B\lambda_{t-1} + (1-B)\lambda_0 + e_t,$$

where $u_t \sim \mathrm{N}(0, \Sigma \otimes H)$; $\varepsilon_i \sim \mathrm{N}(0, V \otimes E_i\Omega_1 E_i)$; $v \sim \mathrm{N}(0, \Psi)$; $\lambda_o \sim \mathrm{N}(0, V \otimes \Omega_{20})$; $e_t \sim \mathrm{N}(0, V \otimes \Omega_{2t})$ and $\Omega_{2t} = v_1\Omega_{2t-1} + v_2\Omega_{20}$. We assume that the covariance matrices are independent, that V, Ψ, v_1, v_2, and μ are known and that $\Sigma \sim iW_N(\sigma_o, M_o)$, $H \sim iW_G(h_o, P_o)$, $\Omega_1 \sim iW_k(w_1, W_1)$, and $\Omega_{20} \sim iW_k(w_2, W_2)$, where the notation $\Phi \sim iW_p(v, Z)$ means that the symmetric positive definite matrix Φ follows a p-dimensional inverted Wishart distribution with v degrees of freedom and scale matrix Z. We also assume that for each of these distributions the degrees of freedom and the scale matrix are known. These assumptions are inconsequential and the analysis goes through, even when consistent estimates are substituted for the true ones.

Given this prior information, the posterior density of the parameter vector

$$\vartheta_T = \left(vec(\alpha), vech(\Sigma), vech(H), vec(\bar{\alpha}), vech(\Omega_1), vec\{\lambda_t\}_{t=0}^{T}, vech(\Omega_{20})\right)$$

is given by

$$\pi(\vartheta_T | Y_T, F_T) \propto f(Y_T | \vartheta_T, F_T)\, p(\vartheta_T | F_T), \tag{21}$$

where $Y_T = (Y_1, \ldots, Y_T)$ is the sample data and $p(\vartheta | F_T)$ is the prior information available at T.

To obtain marginal posteriors, we iterate on the conditional distributions of the parameters, which can easily be obtained from the conditional posterior (21). To deal with the presence of time varying parameters we adapt the results of Carter and Kohn (1994) and Chib and Greenberg (1996). In fact, conditional on $\{\lambda_t\}_{t=0}^{T}$, the distribution of the remaining parameters can be derived without difficulty. Let ψ_{-x} be the vector ϑ containing all the parameters but x. Then, the conditional distributions for parameters other than $\{\lambda_t\}$ are:

$$\Omega_1 \,|\, \psi_{-\Omega_1}, Y_T, F_T \sim iW_k(w_1 + NG, \hat{W}_1),$$

$$\Omega_{20} \,|\, \psi_{-\Omega_{20}}, Y_T, F_T \sim iW_k(w_2 + TG, \hat{W}_2),$$

$$\Sigma \,|\, \psi_{-\Sigma}, Y_T, F_T \sim iW_N(\sigma_o + GT, \hat{M}_o),$$

$$H \mid \psi_{-H}, Y_T, F_T \sim iW_G(h_o + NT, \hat{P}_o),$$

$$\alpha \mid \psi_{-\alpha}, Y_T, F_T \sim \mathrm{N}(\hat{\alpha}, \hat{V}_\alpha),$$

$$\bar{\alpha} \mid \psi_{-\bar{\alpha}}, Y_T, F_T \sim \mathrm{N}(\alpha^*, \hat{V}^*), \tag{22}$$

where the expressions for $\hat{W}_1, \hat{W}_2, \hat{M}_o, \hat{P}_o, \hat{\alpha}, \hat{V}_\alpha, \alpha^*, \hat{V}^*$ are given in the appendix.

Following Chib and Greenberg (1996), the parameter vector λ_t can be included in the Gibbs sampler via the distribution $\pi(\lambda_o, \ldots, \lambda_T \mid Y_T, F_T, \psi_T)$ where $\psi_t \equiv \vartheta_{-\{\lambda_t\}_t}$. We can rewrite such a distribution as

$$\pi(\lambda_T \mid Y_T, F_T, \psi_T) \times \pi(\lambda_{T-1} \mid Y_T, F_T, \psi_{T-1}, \lambda_T) \times \cdots$$
$$\times \pi(\lambda_o \mid Y_T, F_T, \psi_0, \lambda_1, \ldots, \lambda_T). \tag{23}$$

A draw from the joint distribution can be obtained by drawing $\tilde{\lambda}_T$ from $\pi(\lambda_T \mid Y_T, F_T, \psi_T)$; then $\tilde{\lambda}_{T-1}$ from $\pi(\lambda_{T-1} \mid Y_T, F_T, \psi_{T-1}, \tilde{\lambda}_T)$ and so on. Let $\lambda^s = (\lambda_s, \ldots, \lambda_T)$ and $Y^s = (Y_s, \ldots, Y_T)$ for $s \leqslant T$. The density of the typical term in (23) is

$$\pi(\lambda_t \mid Y_T, F_T, \psi_t, \lambda^{t+1})$$
$$\propto \pi(\lambda_t \mid Y^t, F_t, \psi_t)\pi(\lambda_{t+1} \mid Y_t, F_t, \psi_{t-1}, \lambda_t) f(Y^{t+1}, \lambda^{t+1} \mid Y_t, F_t, \lambda_t, \lambda_{t+1})$$
$$\propto \pi(\lambda_t \mid Y^t, F_t, \psi_t)\pi(\lambda_{t+1} \mid F_t, \psi_{t-1}, \lambda_t). \tag{24}$$

The last row of (24) follows from the fact that, conditional on λ_{t+1}, the joint density of (Y^{t+1}, λ^{t+1}) is independent of λ_t and, conditional on λ_t, λ_{t+1} is independent of Y_t. The second density in (24) is Gaussian with moments $B\lambda_t + (I - B)\lambda_0$ and $V \otimes \Omega_{2t}$. The first was derived in Section 3, and it is Gaussian with mean $\hat{\lambda}_{t|t} = \hat{\lambda}_{t|t-1} + \hat{\Omega}_{t|t-1} Z_t' \phi_{22}^{-1} (Y_t - Z_t\mu - Z_t\hat{\lambda}_{t|t-1})$ and variance $\hat{\Omega}_{t|t} = \hat{\Omega}_{t|t-1} - \hat{\Omega}_{t|t-1} Z_t' \phi_{22}^{-1} Z_t \hat{\Omega}_{t|t-1}$. Hence, $\pi(\lambda_t | Y_T, F_t, \psi_t, \lambda^{t+1}) \sim \mathrm{N}(\hat{\lambda}_t, \hat{\Omega}_t)$ where $\hat{\lambda}_t = \hat{\lambda}_{t|t} + M_t(\lambda_{t+1} - B\hat{\lambda}_{t|t} - (I - B)\lambda_0)$; $\hat{\Omega}_t = \hat{\Omega}_{t|t} - M_t \Omega^*_{t+1|,t} M_t'$ and $M_t = \rho\hat{\Omega}_{t|t}\hat{\Omega}^{*-1}_{t+1|t}$, with $\hat{\Omega}^*_{t+1|t} = [\hat{\Omega}_{t+1|t} - (I - B)(V \otimes \Omega_{20})(I - B)']$.

To be concrete the following algorithm can be used to sample $\{\lambda_t\}$: first, starting from given initial conditions, we run the Kalman filter to recursively get $\hat{\lambda}_t$ and $\hat{\Omega}_t$; then we simulate $\tilde{\lambda}_T$ from a normal with mean $\hat{\lambda}_{T|T}$ and variance $\hat{\Omega}_{T|T}$; $\tilde{\lambda}_{T-1}$ from $\mathrm{N}(\hat{\lambda}_{T-1}, \hat{\Omega}_{T-1})$, and so on until $\tilde{\lambda}_o$ is simulated from $\mathrm{N}(\hat{\lambda}_o, \hat{\Omega}_o)$ where, for each t,

$$\hat{\lambda}_t = \hat{\lambda}_{t|t} + M_t(\tilde{\lambda}_{t+1} - \rho\hat{\lambda}_{t|t} - (1 - \rho)\hat{\lambda}_{0|0}),$$

and

$$\hat{\Omega}_t = \hat{\Omega}_{t|t} - M_t \hat{\Omega}^*_{t+1|,t} M_t'.$$

One special case of this setup deserves some attention. Assume informative priors on all the parameters except on H, whose prior is now diffuse. Then our framework resembles the Normal-Diffuse prior of Kadiyala and Karlsson (1997) where posterior dependence among the coefficients of different equations obtains even when there is prior independence. There are two additional major difference with the specification

used by these authors: first, we assume that both the mean and the variance of γ_t are random variables—they take the mean and the variance of γ_t to be fixed. Second, we do not restrict Σ_u to be diagonal and therefore allow complicated interactions among variables within and across countries.

Canova and Ciccarelli (1999) describe in detail two special cases of the general setup of this subsection: (i) no information on the location of the mean of the unit specific effect ($\Psi^{-1}=0$) and (ii) no information on the time varying component of the coefficients at a particular point in time ($\hat{\Omega}_{t|t-1}^{-1}=0$). They show that a diffuse prior on $\bar{\alpha}$ does not allow to update the prior information we have on λ_t and that, in the latter case, the posterior mean and variance for γ_t are the same as those obtained when only prior information on α is used.

Finally, it is worth reminding that the structure employed imposes symmetry restrictions, which are desirable in an unrestricted VAR system. Clearly, these restrictions may be inappropriate for structural or restricted VAR systems and alternative specifications, along the lines of Sims and Zha (1998), should be used.

4.2. Minnesota-type prior

Given the computational complexity involved in calculating posterior Gibbs sampling estimates for large-scale problems, one may be interested in knowing whether shortcuts, which do not require iterative procedures, may be used.

Here, we adapt the so-called Minnesota prior to a panel VAR framework. The Minnesota prior, described in Litterman (1986), Doan et al. (1984), Ingram and Whiteman (1994), Ballabriga et al. (1998) among others, is a way to account for the near non-stationarity of many macroeconomic time series and, at the same time, to weakly reduce the dimensionality of a VAR model. Given that the intertemporal dependence of the variables is believed to be strong, the prior mean of the VAR coefficients on the first own lag is set equal to one and the mean of remaining coefficients is equal to zero. The covariance matrix of the coefficients is diagonal (so we have prior and posterior independence between equations) and the elements are specified in a way that coefficients of higher order lags are likely to be close to zero (the prior variance decreases when the lag length increases). Moreover, since most of the variations in the VAR variables is accounted for by own lags, coefficients of variables other than the dependent one are assigned a smaller relative variance. The prior on the constant term, other deterministic and exogenous variables, is diffuse. Finally, the variance–covariance matrix of the error term is assumed to be fixed and known.

For a panel VAR setup we introduce the following modifications. The covariance-matrices $\Omega_{20}, \Psi, \Delta$, are assumed to have the same structure. Take, for example, $\Delta = diag(\Delta_1, \ldots, \Delta_n)$, where $\Delta_i = V \otimes E_i \Omega_1 E_i$. We assume that $V = H$ (see Eq. (10)) and that Ω_1 is diagonal with elements:

$$\sigma_{g_i j_s}^2 = \left(\frac{\theta_{1\alpha} \theta_3^{\delta(g_i j_s)}}{l^{\theta_2}} \frac{1}{\sigma_{j_s}} \right)^2 \quad g,\ j = 1, \ldots, G \quad i,\ s = 1, \ldots, N \quad l = 1, \ldots, p,$$

where $\delta(g_i, j_s) = 0$ if $i = s$ and 1 otherwise and

$$\sigma_{gm}^2 = (\theta_{1\alpha}\theta_4)^2, \quad m = 1, \ldots, q.$$

Here, g_i represents equation g of unit i, j_s the endogenous variable j of unit s, l the lag, m exogenous or deterministic variables.

The hyperparameter $\theta_{1\alpha}$ controls the tightness of beliefs for the vector α; θ_2 the rate at which the prior variance decays with the lag; θ_3 the degree of uncertainty for the coefficients of the variables of unit s in the equations of unit i; θ_4 the degree of uncertainty of the coefficients of the exogenous variables and σ_{j_s} are the diagonal elements of the matrix Σ_u used as scale factors to account for differences in units of measurement. Notice that we do not have prior independence between equations: our prior information specifies that, for example, the coefficient on lag 1 of the GNP equation for the US may have some relationship with the same coefficient in the PRICE equation for US. Moreover, we have not specified a hyperparameter which controls the overall tightness of beliefs because the randomness of the coefficients depends on α_i and λ_t and we parametrize the uncertainty in each of them separately. Finally, there is no distinction between own versus other countries variables (see Sims and Zha, 1998). This would not be the case if V_i were country specific. The structures for Ψ and Ω_o are identical with $\theta_{1\alpha}$ being replaced by $\theta_{1\tilde{\alpha}}$ and $\theta_{1\lambda}$, respectively.

To complete the specification, the elements of the matrix H and the σ's are estimated from the data to tune up the prior to the application.

The prior time-varying features of the model are determined by specifying the matrices B, Ω_{2t}. We assume that B is diagonal and that each of the $k \times k$ diagonal blocks B_g satisfies: $B_g = diag(\theta_5)$. Furthermore, we let $\Omega_{2t} = \theta_6\Omega_{2o}$. Here, θ_5 controls the evolution of the law of motion of λ_t and $\theta_6 = v_1$ control the time variations in λ_t.

Finally, we assume that the $k \times 1$ vectors μ_g and $\tilde{\lambda}_{og}$ have the following structures:

$$\mu_g = \begin{bmatrix} 0 \\ \vdots \\ \theta_7 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \tilde{\lambda}_{og} = \begin{bmatrix} 0 \\ \vdots \\ 1-\theta_7 \\ 0 \\ \vdots \\ 0 \end{bmatrix},$$

where μ_g and $\tilde{\lambda}_{og}$ are the gth-elements of the mean vectors μ and $\tilde{\lambda}_o$ and θ_7 controls the prior mean on the first own lag coefficient of the dependent variable in equation g for unit i.

Summing up, our prior information can be represented with a nine-dimensional vector of hyperparameters $\Theta = (\theta_{1\alpha}, \theta_{1\lambda}, \theta_{1\tilde{\alpha}}, \theta_2, \theta_3, \theta_4, \theta_5, \theta_6, \theta_7)$. Estimates of Θ can be obtained by maximizing the predictive density of the model as in Doan et al. (1984). Posterior distributions for the parameters are obtained by plugging-in the resulting

estimates for $\mu, \tilde{\lambda}_o, \Omega_{2o}, \Sigma_u, B, \Psi, \Delta$ in the formulas of Section 3 in an empirical Bayes fashion (see e.g. Berger, 1985).

Compared with Ballabriga et al. (1998), who used a Minnesota prior on a panel VAR model for the Spanish, German and French economies, our specification separates the prior information for the time and the individual component (they have one parameter in place of $\theta_{1\alpha}, \theta_{1\lambda}, \theta_{1\tilde{\alpha}}$); it introduces a further level of uncertainty by specifying a prior for $\tilde{\alpha}$ and allows for a priori pooling of the information present in the cross section of the panel. None of these features is present in their specification.

5. Forecasting

Once posterior estimates are obtained, forecasts can be computed. In order to obtain multistep forecasting formulas for a panel VAR and to compute turning points probabilities, we rewrite (4) in a companion VAR(1) form

$$Y_{it} = \sum_{j=1}^{N} B_{it}^{j} Y_{jt-1} + D_{it} z_t + U_{it}, \tag{25}$$

where Y_{it} and U_{it} are $Gp \times 1$ vectors, B_{it}^{j} is a $Gp \times Gp$ matrix and D_{it} is a $Gp \times q$ matrix.

Stacking for i, and repeatedly substituting we have

$$Y_t = \left[\prod_{r=0}^{h-1} B_{t-r}\right] Y_{t-h} + \sum_{s=0}^{h-1} \left[\prod_{r=0}^{s-1} B_{t-r}\right] D_{t-s} z_{t-s} + \sum_{s=0}^{h-1} \left[\prod_{r=0}^{s-1} B_{t-r}\right] U_{t-s} \tag{26}$$

or

$$y_t = J\left[\prod_{r=0}^{h-1} B_{t-r}\right] Y_{t-h} + \sum_{s=0}^{h-1} \Phi_{st} D_{t-s} z_{t-s} + \sum_{s=0}^{h-1} \Phi_{st} u_{t-s}, \tag{27}$$

where $\Phi_{st} = \Pi_{r=0}^{s-1} B_{t-r}$, and $J = I_N \otimes J_1$, $J_1 = [I_G \ 0]$ and J is a selection matrix such that $JY_t = y_t, JU_t = u_t$ and $J'JU_t = U_t$. The expression in (27) can be used to compute the h-steps ahead forecast of the NG-dimensional vector Y_t.

First, we compute a "point" forecast for y_{t+h}. The forecast function is given by

$$y_t(h) = J\left[\prod_{r=0}^{h-1} B_{t+h-r}\right] Y_t + \sum_{s=0}^{h-1} \Phi_{st+h} D_{t+h-s} z_{t+h-s} \tag{28}$$

or, recursively

$$y_t(h) = J\tilde{B}_{t+h} Y_t(h-1) + \tilde{D}_{t+h} z_{t+h},$$

where $\tilde{D}_{t+h}$ is the $NG \times q$ matrix $[d_{1t}\ d_{2t} \cdots d_{Nt}]'$ and $\tilde{B}_{t+h} = diag(B_{1t}, B_{2t}, \ldots, B_{nt})$ with $B_{it} = (B_{it}^{1}, B_{it}^{2}, \ldots, B_{it}^{N})$. One way to obtain a h-step ahead forecasts is to use the posterior mean of $\tilde{B}_{t+h}$ and $\tilde{D}_{t+h}$ and the mean of the predictive density for z_{t+h}, conditional on the information at time t. Estimates for the posterior mean of the coefficients can be

obtained from the recursive formulas for λ_t (and, consequently, for γ_t) using expressions like (9) and by drawing from distributions like (23) in a recursive fashion. Call this estimates $\hat{B}_{t+h|t}$ and $\hat{D}_{t+h|t}$. The forecast error is $y_{t+h} - \hat{y}_t(h) = \sum_{s=0}^{h-1} \Phi_{st+h} u_{t+h-s} + [y_t(h) - \hat{y}_t(h)]$. To measure the forecasting performance we compute the mean square error (MSE) or the mean absolute error (MAD) of the estimated forecast which are given by

$$MSE(\hat{y}_t(h)) = \sum_{s=0}^{h-1} \Phi_{st+h} \Sigma_u \Phi'_{st+h} + MSE[y_t(h) - \hat{y}_t(h)],$$

$$MAD(\hat{y}_t(h)) = \sum_{s=0}^{h-1} |u_{t+h-s}| + MAD[y_t(h) - \hat{y}_t(h)].$$

The first term on the RHS of each equation can be obtained using posterior mean estimates of B_{t+h-r} and of U_t, conditional on the information at time t, while for the second term an approximation can be computed along the lines of Lütkepohl (1990, pp. 86–89). Clearly, if a researcher is interested in point forecasts using the average value of the parameters, then the previous formulas apply using for $\hat{B}_{t+h|t}$ and $\hat{D}_{t+h|t}$ the posteriors derived in Section 3.2.

In many situations, it may be more appealing to compute "average" forecasts h-step ahead using the predictive density

$$f(Y_{t+h}|F_t) = \int f(Y_{t+h}|F_t, \vartheta) p(\vartheta|F_t),$$

where $f(Y_{t+h}|F_t, \vartheta)$ is the conditional density of the future observation vector given ϑ, and $p(\vartheta|F_t)$ is the posterior pdf of ϑ at time t. To compute forecasts for Y_{t+h} we can sample from the predictive density numerically. For each $i=1,\ldots,M$ we draw $\vartheta^{(i)}$ from the posterior distribution and simulate the vector $Y_{t+h}^{(i)}$ from the density $f(Y_{t+h}|F_t, \vartheta^{(i)})$. $\{Y_{t+h}^{(i)}\}_{i=1}^{M}$ constitutes a sample, from which we can compute the necessary moments. The value of the forecast is then the ergodic average $\hat{Y}_{t+h} = M^{-1} \sum_{i=1}^{M} Y_{t+h}^{(i)}$ and its numerical variance can be estimated using $var(\hat{Y}_{t+h}) = M^{-1}[Q_o + \sum_{s=1}^{r}(1 - s/(r+1))(Q_s + Q'_s)]$ where $Q_s = M^{-1} \sum_{i=s+1}^{M} [Y_{t+h}^{(i)} - \hat{Y}_{t+h}][Y_{t+h}^{(i)} - \hat{Y}_{t+h}]'$.

Since the computation of the impulse response function for orthogonalized shocks is a simple corollary of the calculation of forecasts, the approach we provide here to calculate point and average forecasts can also be used to compute impulse responses. In fact, given the information up to time t, computing impulse response at $t+h$ is equivalent to calculating the difference between the conditional forecasts at $t+h$, given that at $t+1$ there has been a one unit impulse in one of the orthogonal shocks, and the unconditional forecast, i.e. with the value of the vector that would have occurred without shocks (see Koop, 1992 for an application to structural VAR models). This idea is exploited in a recent paper by Waggoner and Zha (1998). The authors, using a version of (27), develop two Bayesian methods for computing probability distributions of conditional forecasts. The last term in (27) represents the dynamic impact of structural shocks which affect future realizations of variables through the impulse

response matrix Φ_{st}. With conditions or constraints imposed on this last term we can produce what they call *conditional forecasts*.

In order to compute structural impulse responses and their error bands we must work with a structural VAR, e.g. impose some restrictions on the contemporaneous coefficient matrix. A prior (flat or informative) can then be assigned to the non-zero elements of this matrix, as suggested by Sims and Zha (1998). The extension of their approach to panel data is however not straightforward and we postpone this issue to future work.

Turning point predictions can also be computed from the predictive density of future observations (see in Zellner et al., 1991). We define turning points as follows:

Definition 5.1. A downward turn for unit i at time $t+h+1$ occurs if S_{it+h}, the growth rate of the reference variable (typically, *GNP*), satisfies for all h $S_{it+h-2}, S_{it+h-1} < S_{it+h} > S_{it+h+1}$. An upward turn for unit i at time $t+h+1$ occurs if the growth rate of the reference variable satisfies $S_{it+h-2},\ S_{it+h-1} > S_{it+h} < S_{it+h+1}$.

Similarly, we define a non-downward turn and a non-upward turn:

Definition 5.2. A non-downward turn for unit i at time $t+h+1$ occurs if S_{it+h} satisfies for all h $S_{it+h-2},\ S_{it+h-1} < S_{it+h} \leqslant S_{it+h+1}$. A non-upward turn for unit i at time $t+h+1$ occurs if the growth rate of the reference variable satisfies $S_{it+h-2},\ S_{it+h-1} > S_{it+h} \geqslant S_{it+h+1}$.

Although there are other definitions in the literature (see e.g. Lahiri and Moore, 1991) this is the most used one and it suffices for our purposes. Let $\tilde{f}(Y_{i,t+h}|F_t) = \int_{Y_{p,t+h}} f(Y_{t+h}|F_t)\, dY_{p,t+h}$ be the marginal predictive density for the variables of unit i after integrating the remaining p variables and let $\mathscr{K}(S^1_{it+h}|F_t) = \int \cdots \int f(S^1_{it+h} \cdots S^G_{it+h}|F_t)\, \mathrm{d}S^2_{it+h} \cdots \mathrm{d}S^G_{it+h}$ be the marginal predictive density for the growth rate of the reference variable, which we order first in the list, in unit i.

Take now the simplest case of $h=0$. To compute the probability of a turning point we have to calculate S^1_{it+1}. Given the marginal predictive density $\mathscr{K}$, the probability of a downturn in unit i is

$$P_{Dt} = P_r(S^1_{it+1} < S^1_{it} | S^1_{it-2}, S^1_{it-1} < S^1_{it}, F_t)$$

$$= \int_{-\infty}^{S^1_{it}} \mathscr{K}(S^1_{it+1}|S^1_{it-2}, S^1_{it-1}, S^1_{it}, F_t)\, \mathrm{d}S^1_{it} \tag{29}$$

and the probability of an upturn is

$$P_{Ut} = Pr(S^1_{it+1} > S^1_{it} | S^1_{it-2}, S^1_{it-1} > S^1_{it}, F_t)$$

$$= \int_{S^1_{it}}^{\infty} \mathscr{K}(S^1_{it+1}|S^1_{it-2}, S^1_{it-1}, S^1_{it}, F_t)\, \mathrm{d}S^1_{it}. \tag{30}$$

Using a numerical sample from the predictive density satisfying $S^1_{it-2}, S^1_{it-1} < S^1_{it}$, we can approximate these probabilities using the frequencies of realizations which are less

than or greater than S_{it}. With a symmetric loss function, minimization of the expected loss leads to predict the occurrence of turning point at $t+1$ if $P_{Dt} > 0.5$ or $P_{Ut} > 0.5$.

For $h \neq 0$ the probability of a turning point can be computed using the joint predictive density for all future observations, i.e. in the case of a downturn,

$$\begin{aligned} P_{Dt+h} &= Pr(S^1_{it+h+1} < S^1_{it+h} > S^1_{it+h-2}, S^1_{it+h-1}|F_t) \\ &= \int_{-\infty}^{S^1_{it+h}} \int_{S^1_{it+h}}^{\infty} \int_{S^1_{it+h}}^{\infty} \mathcal{K}(S^1_{it+h+1} < S^1_{it+h} > S^1_{it+h-2}, S^1_{it+h-1}|F_t) \\ &\quad \times \mathrm{d}S^1_{it+h}\, \mathrm{d}S^1_{it+h-1}\, \mathrm{d}S^1_{it+h-2}. \end{aligned} \tag{31}$$

Given the available panel data structure we may also be interested in computing the probability that a turning point occurs jointly for $m \leqslant N$ units of panel. For example, we would like to compute the probability that at $t+1$ there will be a recession in European countries. Let $\tilde{\mathcal{K}}(S^1_{t+h}|F_t)$ be the joint predictive density of the reference variable for the m units of interest. Then the probability of a downturn is

$$\begin{aligned} P^m_{Dt} &= Pr(S^1_{it+1} < S^1_{it}\, i = 1, \ldots, m | S^1_{it-2}, S^1_{it-1} < S^1_{it}, F_t) \\ &= \int_{-\infty}^{S^1_{it}} \cdots \int_{-\infty}^{S^1_{mt}} \tilde{\mathcal{K}}(S^1_{t+1}|S^1_{t-2}, S^1_{t-1} < S^1_t, F_t)\, \mathrm{d}S^1_{1t} \cdots \mathrm{d}S^1_{mt}. \end{aligned} \tag{32}$$

6. An application

In this section we apply the methodology to the problem of forecasting growth rates and predicting turning points in the G-7 countries. For each country we consider three national variables (GNP, real stock returns and real money growth) and a world one (the median real stock return in OECD countries) which is assumed to be exogenous in each equation. Figs. 1–3 plot the series. Hence there are 21 variables in the panel VAR. These variables are chosen after a rough specification search over about 10 variables because they appear to have the highest in-sample pairwise and multiple correlation with output growth. Among the variables we tried are also the nominal interest rate, the slope of the term structure and inflation. Data are sampled quarterly from 1973,1 to 1993,4 and taken from IMF statistics. The sample 1973,1–1988,4 is used to estimate the parameters and the sample 1989,1–1993,4 to evaluate the forecasting ability and to predict turning points.

We compare the forecasting performance of our panel VAR specifications with those obtained with other models suggested in the literature. As a benchmark we first run two versions of a tri-variable VAR(2) model for each country separately. The first one is an unrestricted (VAR). The second a weakly restricted VAR(BVAR) where we use a standard Litterman-prior with a mean of one on the first lag, a general tightness of 0.15, no decay in the lags and a weight of 0.5 on the lags of other variables. Since these two models do not exploit cross-sectional information nor do they allow for time variation, they are the natural benchmark to measure the improvements obtained by specifications which allow any of these two features in the model.

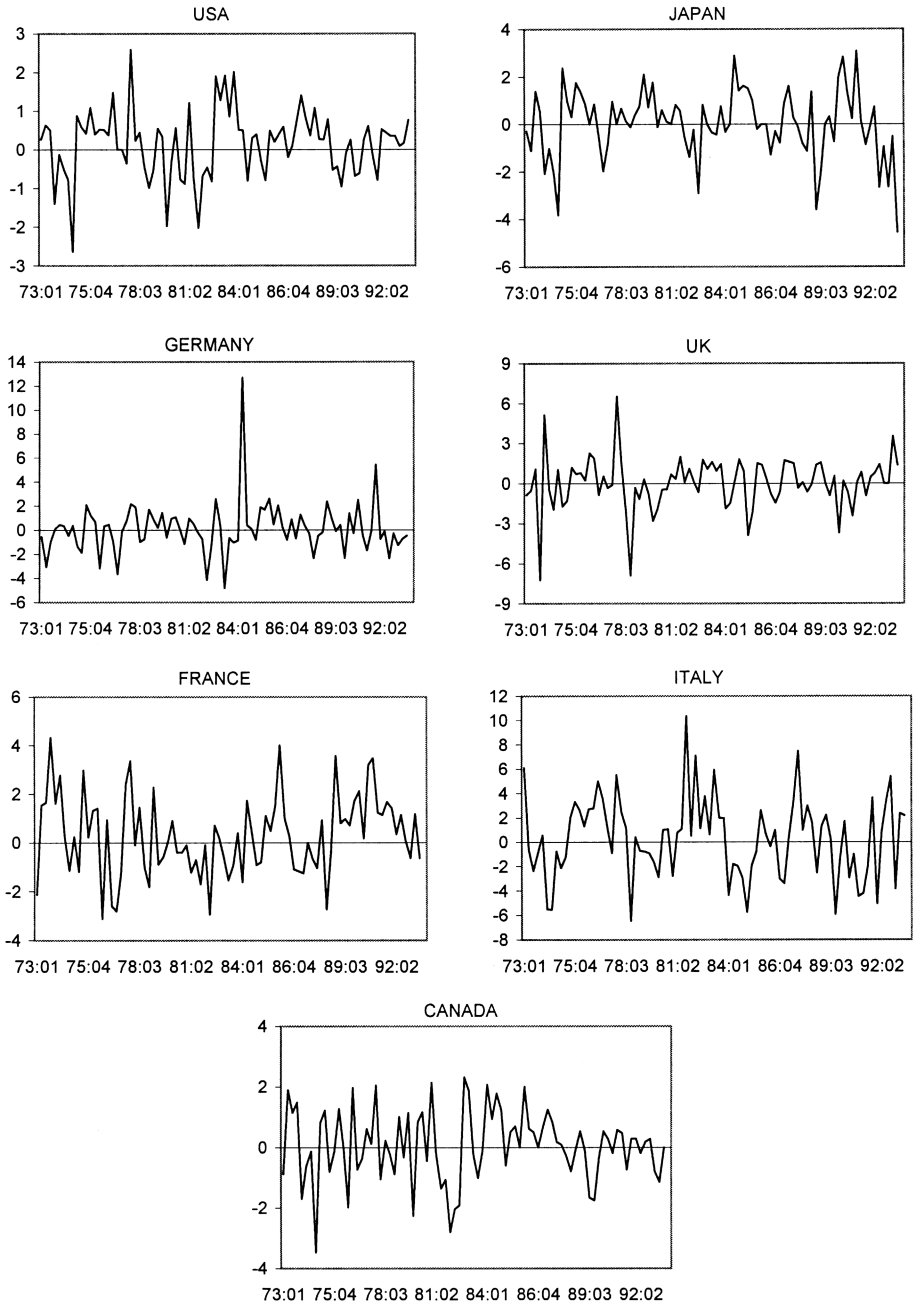

Fig. 1. Growth rate of GNP.

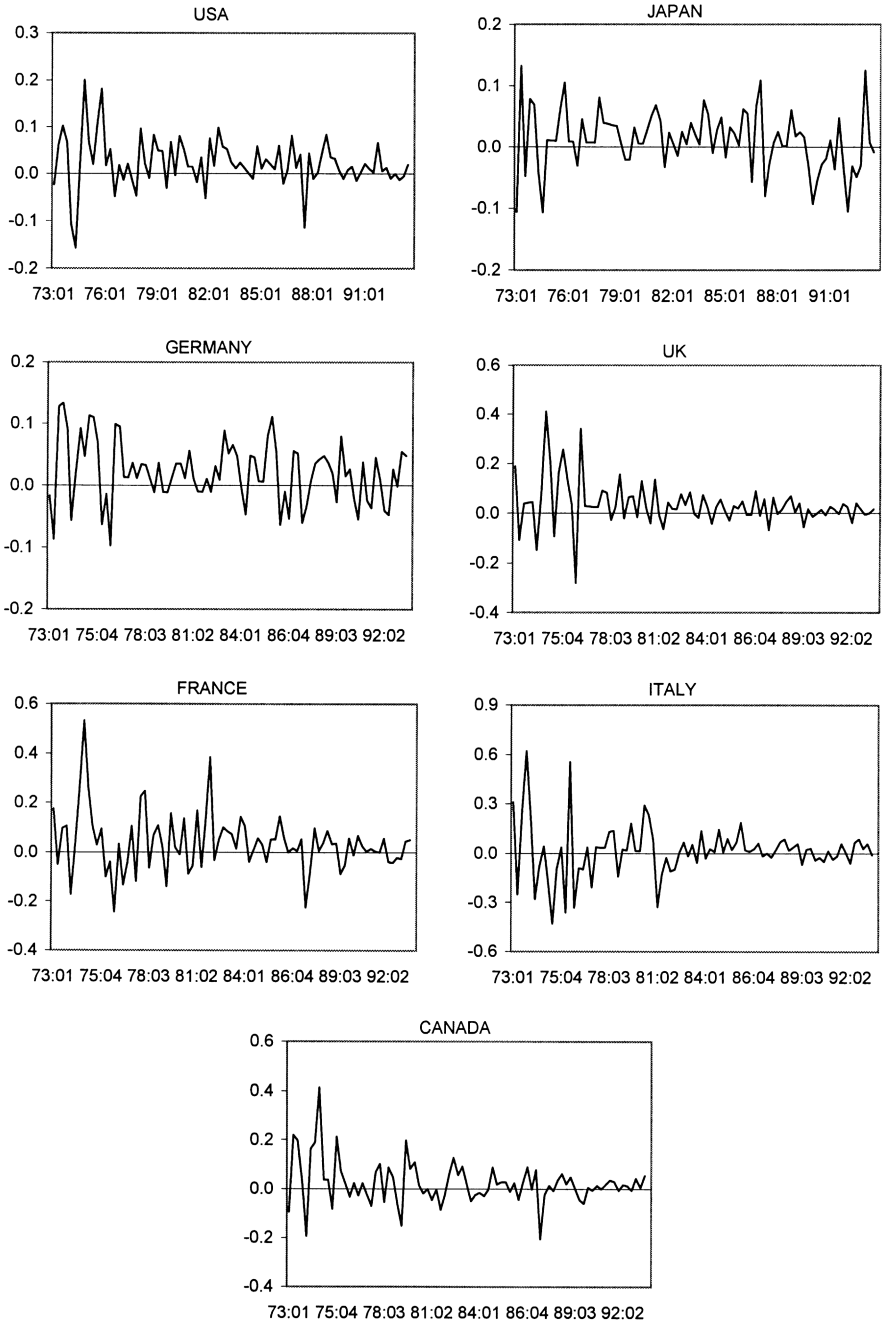

Fig. 2. Real stock return.

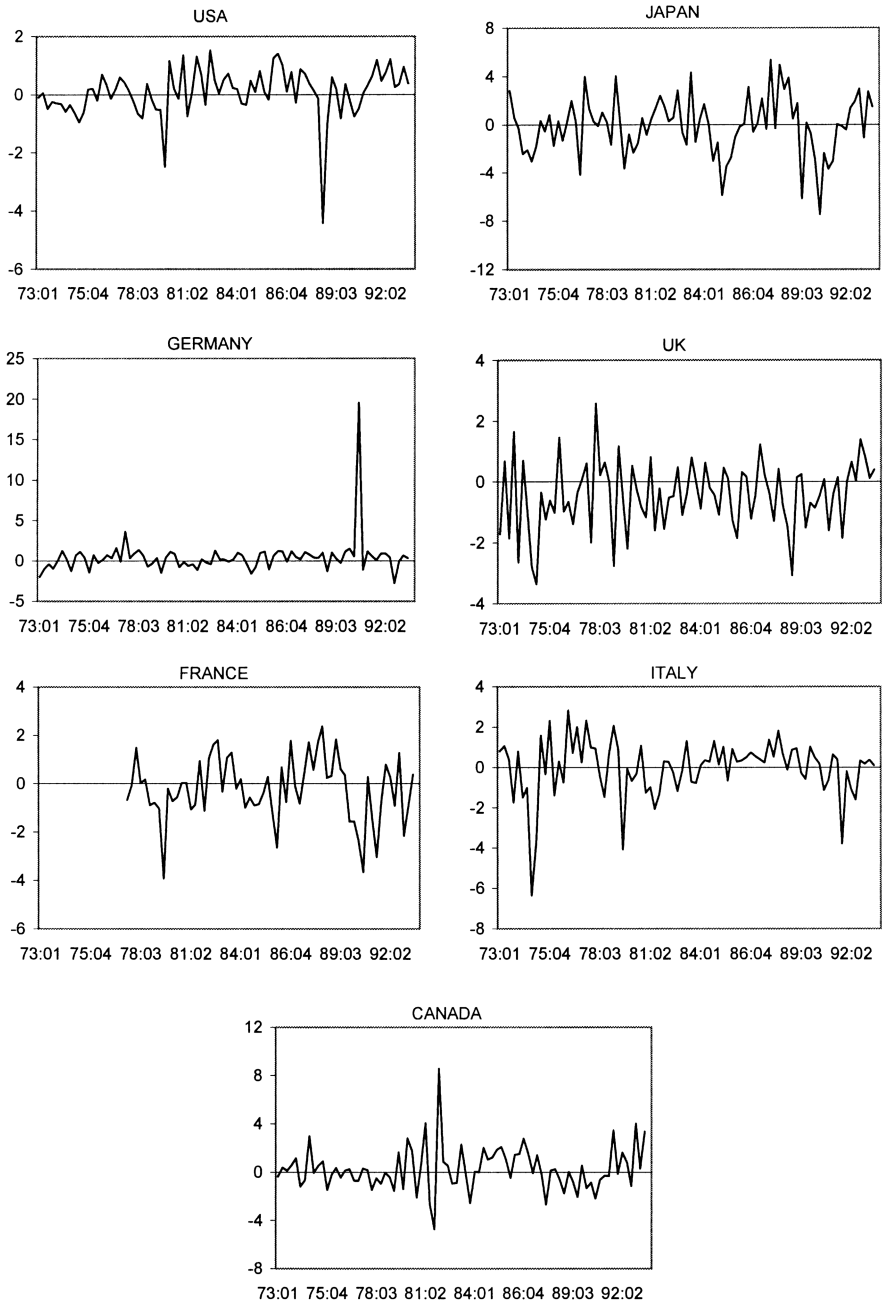

Fig. 3. Real money growth.

Also for comparison, we run a single equation AR(3) for GNP growth for each single country, augmented with two lags of real stock returns, 1 lag of real money balances and one lag of the median world real stock return. This is the specification used by Garcia Ferrer et al. (1987), Zellner and Hong (1989) and Zellner et al. (1991) to forecast annual growth rates of output in 18 countries. With the extended sample and the higher frequency of the data we have available, we confirm their results for all of the G-7 countries. This model represents a restricted version of the previous unrestricted VAR where insignificant lags are purged from the specification. The forecasting power of this model is measured when parameters are estimated with OLS and with the three shrinkage procedures: a ridge estimator (RIDGE), an estimator obtained assuming an exchangeable prior on the coefficients (as in Garcia Ferrer et al., 1987) (EXCHANGEABLE) and an estimator obtained using a g-prior (as in Zellner and Hong, 1989) (G-PRIOR). The two latter estimators attempt to improve upon OLS by combining the information coming from all the units in the cross section. They differ in the way they combine this information. Since none of these estimators allows for time variations in the coefficients we also consider a specification where the coefficient vector of the AR(3) model smoothly evolves over time. The evolution of the coefficients is controlled by five parameters: a general tightness θ_1, a lag decay parameter θ_2, a tightness on world variables θ_3, a parameter controlling the persistence in the coefficients θ_4 and a parameter controlling for time variations in the variance θ_5. Optimal values of these parameters are obtained maximizing the in-sample predictive density with a simplex algorithm and appear in Table 1.

As a final term of comparison, we use a version of the panel VAR specification suggested by Ballabriga et al. (1998) (PBVAR). This model specification does not use the information coming from the cross section—every variable is treated in the same way regardless of the country where is from—but allows for time variations in the coefficients of the model. The model has the same structure as Doan et al. (1984) and assumes that the coefficient vector β_t has an AR(1) structure of the form $\beta_t = M\beta_{t-1} + u_t$ where u_t, conditional on the information available, is normal with mean zero and variance Σ_u. The matrices β_0, M, and Σ_u depend on seven hyperparameters: five parameters control the structure of Σ_{u_0} (a general tightness (θ_1), a tightness on variables of the same country (θ_3), a tightness on the variables of other countries (θ_4), a geometric lag decay parameter (θ_2), and a tightness on world variables (θ_5)); a parameter describes the structure of M (θ_6); and a parameter controls the prior mean

Table 1
Estimated hyperparameters: univariate TVC

	US	Japan	Germany	UK	France	Italy	Canada
General tightness (θ_1)	0.022	0.022	0.096	0.091	0.087	0.122	0.022
Lag decay (θ_2)	2.0	1.0	1.0	0.5	1.0	0.5	2.0
World tightness (θ_3)	10000	100	100	10000	10000	10000	10000
AR coefficient (θ_4)	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Time variations (θ_5)	1.0E − 05	1.0E − 05	1.0E − 05	1.0E − 05	1.0E − 04	1.0E − 05	1.0E − 05

on the first lag of β_0 (θ_7). Table 2 reports the optimal values obtained by maximizing the in-sample predictive density with a simplex algorithm.

We produce forecasts from two versions of our panel VAR model: one with a modified Minnesota-prior (PANEL1) and one with a fully hierarchical heteroschedastic specification (PANEL2). In the first one, the nine prior parameters are selected to maximize the predictive density using a simplex method. Their optimal values are in Table 3. For both PBVAR and PANEL1 forecasts are computed using the posterior mean of the coefficients, after we have plugged-in estimates of the prior parameters in the relevant formula. For PANEL2 posterior estimates are computed numerically using the Gibbs sampler and forecasts are directly obtained from these estimates.

For both models we set $H = V$. For the PANEL1 specification we compute the scale factors V and the matrix Σ_u as follows. We estimate a trivariate VAR for each country and take the average of the estimated variance–covariance matrix of the residuals across countries as a measure of V. Furthermore, for each of the three variable we estimate a 7-variable VAR (the same variable across countries) and store the variance–covariance matrices of the residuals. An estimate of Σ_u is

$$\hat{\Sigma}_u = \sum_{j=1}^{3} \begin{pmatrix} \sigma_1 & 0 & \cdots & 0 \\ 0 & \sigma_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_7 \end{pmatrix}_j \otimes \begin{pmatrix} 0 & 0 & 0 \\ 0 & v_j & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

Table 2
Estimated hyperparameters: PBVAR

General tightness (θ_1)	0.01
Lag decay (θ_2)	13.96
Own country tightness (θ_3)	3.5 − e005
Other countries tightness (θ_4)	7.3 − e004
World variable tightness (θ_5)	5.0e − 007
AR coefficient (θ_6)	0.95
Prior mean on the first lag (θ_7)	0.11048

Table 3
Estimated hyperparameters: PANEL1

Tightness for α ($\theta_{1\alpha}$)	0.1207
Tightness for λ ($\theta_{1\lambda}$)	0.1300
Tightness for $\bar{\alpha}$ ($\theta_{\bar{\alpha}}$)	0.0004
Lag decay (θ_2)	1.9156
Tightness on other countries (θ_3)	0.0046
Tightness on world variables (θ_4)	4.7804
Law of motion of λ (θ_5)	0.1211
Time variation (θ_6)	0.4295
Prior mean on the first lag (θ_7)	0.0754

where the first matrix contains on the diagonal the estimated standard deviations obtained by running the three 7-variate VARs; while the second matrix contains just one element different from zero, the (j,j) element, which is obtained from the diagonal of the matrix V. For the PANEL2 specifications we need to choose the scale and the degrees of freedom in the various Wishart distribution. We still estimate V as before. Following Kadiyala and Karlsson (1997) we set $\sigma_0=N+2+(T-p)*G$, $\omega_1=k+2+N+g$, $\omega_2=k+2+(t-p)*G$, while the scale matrices M_0, W_1 and W_2 are such that $\Sigma_u, \Delta, \Sigma_\varepsilon$ have the same structure as in the PANEL1 specification. Furthermore, v_1 and v_2 are selected with a rough specification search and Ω_{i1} is drawn from an inverted Wishart with $k+G$ degrees of freedom and scale matrix $W_1+(A_iE_i-\bar{A})'V^{-1}(A_iE_i-\bar{A})$.

We compare the forecasting ability of various models using both the Theil-U Statistics and the mean absolute deviation (MAD), 1 and 4 periods ahead, which are reported in Table 4. Since no measure of uncertainty is available for these statistics (so that differences across methods may be due to chance) we complement this information by examining the p-value of a Diebold and Mariano (1995) sign-test for predictive accuracy, using our PANEL2 specification as a benchmark, and by reporting, in Fig. 4, simulated distributions for median and mean values of the MSE for the PANEL2 specification, constructed using the output of the Gibbs sampling. In the case of the Diebold and Mariano test the null hypothesis is that the loss differential between the forecasts obtained with the PANEL2 model and with any other candidate is an i.i.d. random variable. In the second case, the forecasts of the PANEL2 model will be different from those of a rival model if the MSE produced by the candidate is outside a specific range (interquartile, interdecile, etc.) the investigator chooses.

To examine the performance of various models as business cycle indicators we compute turning points predictions one period ahead. Following Zellner et al. (1991), we compute the total number of turning points, the number of downturns and no-downturns, and the number of upturns and no-upturns in the sample (across all countries) and for each procedure we report the number of correct cases in Table 5. Finally, for each model, we compute the probability of a downward turn in the growth rate of US output over the period 1989:1–1991:4, given the information available in 1988:4. According to the official NBER classification the long expansion of the 1980s terminated in 1990:3 and it was followed by a brief and shallow recession. The probabilities for the 10 models for each of the 12 periods we consider are presented in Table 6.

The forecasting performance of univariate OLS, ridge and exchangeable procedures are very similar. The minimum and maximum values of the Theil-U across countries at one and four steps for the latter two are slightly smaller, but the mean and the median at both steps are practically identical. On the other hand, a univariate model where the parameters are shrunk with a g-prior is somewhat better than OLS in all the dimensions using the Theil-U at both steps. The improvements obtained using time variations in univariate specifications are somewhat limited: there are some gains for Japan, but on average a time varying coefficients model is somewhat worse than a model with exchangeable and g-priors.

Unrestricted VAR models are not very successful in forecasting growth rates of output because of the large number of parameters to be estimated. This is noticeable in particular for Japan, Germany and the UK where the Theil-Us are significantly worse

Table 4
Statistics

Method	Step	US	Japan	Germany	UK	France	Italy	Canada	Median	Mean
Theil-U										
VAR	1	1.06	0.88	0.91	0.94[a]	1.00[a]	0.73	0.95[a]	0.94	0.92
	4	0.73	0.95	0.56	0.81	1.32	0.96[a]	0.72	0.81	0.86
BVAR	1	0.83[a]	0.89	0.69	0.91[a]	0.90[a]	0.80	0.85[a]	0.85[a]	0.84
	4	0.75[a]	0.89	0.65[a]	0.79	1.16	1.00	0.70	0.89	0.85
OLS	1	1.21	0.86	0.88	0.86[a]	0.90[a]	0.79	0.91	0.88	0.90
	4	0.77	0.90[a]	1.07[a]	0.76[a]	0.98[a]	1.03[a]	0.67	0.90[a]	0.88
Ridge	1	1.17	0.83	0.89	0.85	0.89[a]	0.79	0.89	0.89[a]	0.90
	4	0.76[a]	0.88	1.06[a]	0.75[a]	0.99[a]	1.01[a]	0.68	0.88	0.87
Exchangeable	1	1.18	0.84	0.90	0.85	0.89[a]	0.78	0.89	0.89[a]	0.90
	4	0.76[a]	0.90	1.09[a]	0.75[a]	0.99[a]	1.01[a]	0.68	0.90	0.88
g-prior	1	1.06	0.86	0.69	0.78	1.00	0.72	0.92	0.86	0.86
	4	0.83[a]	1.07	0.77[a]	0.75[a]	1.12[a]	1.02[a]	0.70	0.83[a]	0.89
TVC	1	1.07[a]	0.85[a]	0.97	0.94[a]	0.84[a]	0.83	0.90[a]	0.90[a]	0.91
	4	0.74[a]	0.70[a]	1.25[a]	0.77	0.88[a]	1.09	0.88	0.88[a]	0.90
PBVAR	1	0.82	0.85	0.68	0.76	0.98[a]	0.73	0.85	0.82	0.81
	4	0.86[a]	0.91	0.77[a]	0.75	1.08[a]	1.03[a]	0.66	0.86[a]	0.87
Panel 1	1	0.81	0.88	0.67	0.75	1.02	0.70	0.88	0.81	0.81
	4	0.86	0.90	0.76	0.74	1.07	1.03	0.66	0.86	0.86
Panel 2	1	0.85	0.80	0.63	0.76	0.95	0.75	0.88	0.80	0.80
	4	0.85	0.87	0.78	0.76	1.07	0.99	0.66	0.85	0.85
MAD										
VAR	1	0.46[a]	1.71	1.74	1.35	1.26[a]	2.91[a]	0.65[a]	1.35	1.44
	4	0.35[a]	1.55	1.18	1.33	1.66	2.74[a]	0.56[a]	1.33	1.34
BVAR	1	0.46[a]	1.62	1.48	1.32	1.15[a]	3.22[a]	0.58[a]	1.32	1.40
	4	0.40[a]	1.39	1.25	1.28	1.42[a]	2.98[a]	0.51	1.28	1.40
OLS	1	0.56[a]	1.59	1.51	1.37[a]	1.06	3.17[a]	0.57[a]	1.37[a]	1.40
	4	0.34[a]	1.54	1.58	1.28[a]	1.14	3.19[a]	0.54	1.28[a]	1.37
Ridge	1	0.54[a]	1.50	1.68	1.31[a]	1.07	3.14[a]	0.56[a]	1.31[a]	1.40
	4	0.36[a]	1.46	1.72[a]	1.25[a]	1.17[a]	3.09[a]	0.53	1.25[a]	1.37
Exchangeable	1	0.54[a]	1.52	1.68	1.32[a]	1.06	3.14[a]	0.56[a]	1.32[a]	1.40
	4	0.35[a]	1.48	1.73[a]	1.26[a]	1.17	3.09[a]	0.53	1.26[a]	1.37
g-prior	1	0.53[a]	1.63	1.33	1.18	1.26[a]	2.89[a]	0.54[a]	1.26[a]	1.34
	4	0.41[a]	1.60	1.35	1.18[a]	1.34	3.12[a]	0.51[a]	1.34	1.36
TVC	1	0.51[a]	1.52	1.53[a]	1.41	1.04	3.48[a]	0.59[a]	1.41	1.44
	4	0.36[a]	1.24	1.76[a]	1.19	1.10[a]	3.28[a]	0.68	1.24	1.37
PBVAR	1	0.46[a]	1.47	1.29[a]	1.17[a]	1.27	2.85[a]	0.53[a]	1.27	1.29
	4	0.44[a]	1.48	1.27[a]	1.12[a]	1.31	3.14	0.51	1.27[a]	1.32
Panel 1	1	0.46	1.53	1.24	1.08	1.37	2.82	0.54	1.24	1.29
	4	0.44	1.48	1.27	1.11	1.31	3.14	0.50	1.27	1.32

Table 4 (*continued*)

Method	Step	US	Japan	Germany	UK	France	Italy	Canada	Median	Mean
Panel 2	1	0.49	1.44	1.23	1.18	1.20	2.93	0.56	1.20	1.25
	4	0.43	1.43	1.33	1.21	1.26	2.99	0.52	1.26	1.31

Notes: VAR is a VAR(2) model for output growth, real stock returns and real money growth, BVAR is the same model with a Minnesota prior. OLS refer to a model where the parameters are estimated with OLS, Ridge to a Ridge correction, Exchangeable to a model with an exchangeable prior, g-prior to Zellner's g-prior specification and TVC to a univariate time varying coefficient model. PBVAR is a 21 VAR model with a Minnesota prior and time variation, Panel 1 is a panel VAR model with all seven countries with a modified Minnesota prior, Panel 2 is the same model with a hierarchical prior and heteroschedasticity.

[a]Null hypothesis that the loss differential between the forecasts of the model and of a PANEL2 is iid, where the loss functions are either the MSE and the MAD has a p-value smaller or equal than 0.10. The test for the mean is not available.

than those obtained with univariate specifications at the one step horizon. However, unrestricted VAR models outperform all univariate specifications at the four step horizon. Hence, the presence of interdependencies across variables helps in predicting the evolution of the growth rate of output in the medium run. BVAR are significantly better than VAR and univariate approaches at the one step horizon: in the median the gains are of the order of 5–6% over univariate specifications and of more than 10% over the unrestricted VAR. At the four step horizon BVARs turn out to be inferior to unrestricted VARs, and comparable to univariate shrinkage procedures. This is to be expected since to improve the performance at short horizons BVARs reduce both the memory and the interdependencies of the system, which are useful when medium-long run forecasts have to be made.

Adding time variation to the coefficients and interdependencies across countries substantially improves the forecasting performance of multivariate models both at short and at medium horizons. For example, the median (mean) Theil-U at one step goes from 0.85 (0.84) with a simple BVAR to 0.80 (0.80) with the PANEL2 version of the model. There are gains also relative to the PBVAR specification but smaller in magnitude. Relatively speaking, the information contained in the cross section is crucial for the UK (the Theil-U is lower by as much as 15%) and important for Japan, Germany and Italy. For the other three nations country specific models appear to be sufficient to predict output growth. The improvements are noticeable also at longer horizons: the distribution of the Theil-U across countries at the four step horizon is similar to the one obtained with an unrestricted VAR, which is the best among the benchmark models.

The size of the forecasts errors are large for every procedure. The median size of the scaled MSE (forecast error divided by the quantity being forecasted) is always greater than 1.0. In certain countries these numbers are even more dramatic. For example, when forecasting German output growth the size of the scaled error is larger than 3.0 for five approaches. Similarly, the scaled MSE obtained in forecasting UK output growth exceeds 3.0 with seven of the 10 procedures. Zellner's g-prior, the PBVAR and our Panel model tend to provide the smallest errors (on average and in the median)

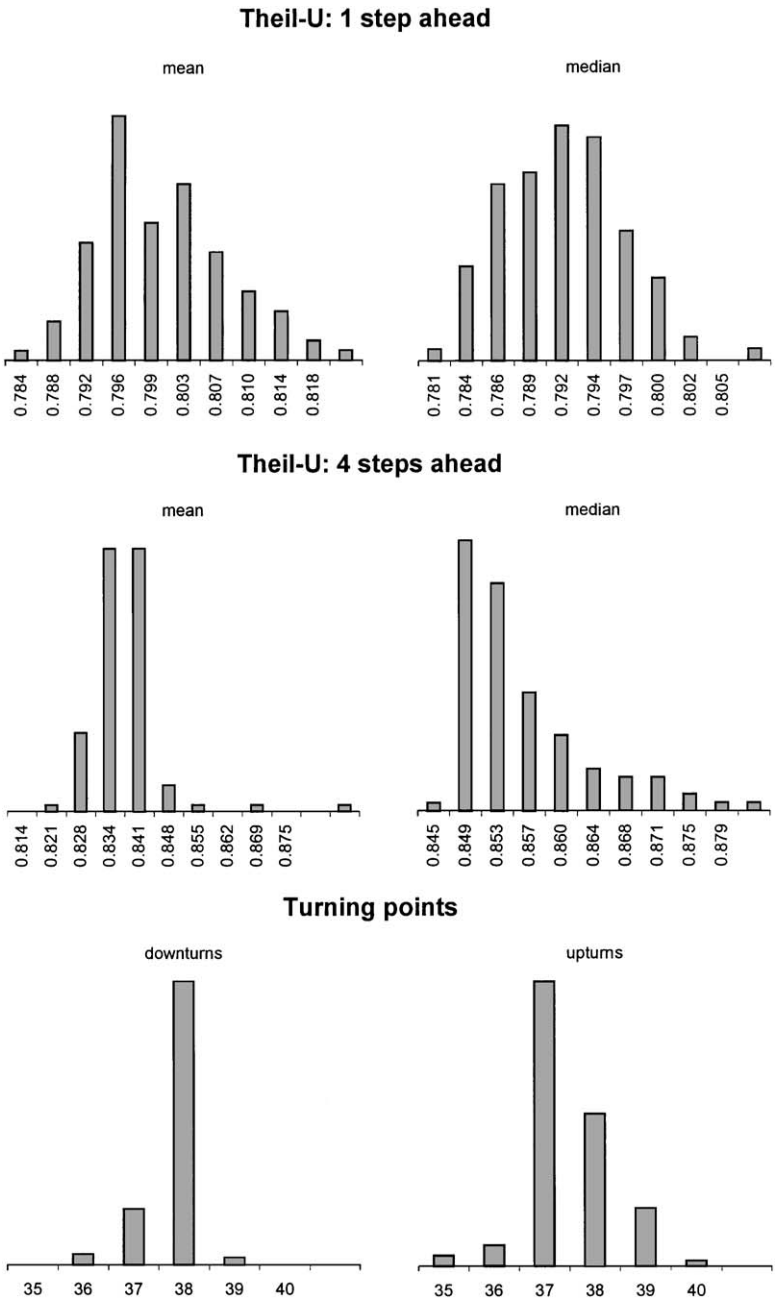

Fig. 4. Empirical distributions of main statistics.

Table 5
Turning points forecasts

Method	Turning points	DT & NDT	UT & NUT
TRUE	96	47	49
VAR	65	32	33
BVAR	72	34	38
OLS	74	37	37
Ridge	72	37	35
Exchangeable	72	37	35
g-prior	75	37	38
TVC	72	34	38
PBVAR	68	32	36
Panel 1	73	36	37
Panel 2	76	38	38

Notes: VAR is a VAR(2) model for output growth, real stock returns and real money growth, BVAR is the same model with a Minnesota prior. OLS refer to a model where the parameters are estimated with OLS, Ridge to a Ridge correction, Exchangeable to a model with an exchangeable prior and g-prior to Zellner's g-prior specification. PBVAR is a 21 VAR model with a Minnesota prior and time variation, Panel 1 is a panel VAR model with all seven countries with a modified Minnesota prior and Panel 2 is the same model with a hierarchical prior and heteroschedasticity. DT means downturn, NDT means non-downturn, UT means upturn and NUT means non-upturn.

Table 6
Probabilities of a downturn in US GDP growth

Quarter	VAR	BVAR	OLS	RIDGE	EXCH	g-PRIOR	TVC	PBVAR	PANEL1	PANEL2
89:1	0.000	0.000	0.000	0.000	0.000	0.000	0.025	0.000	0.000	0.000
89:2[a]	0.000	0.005	0.005	0.010	0.000	0.270	0.781	0.420	0.410	0.625
89:3	0.020	0.010	0.005	0.010	0.200	0.250	0.089	0.010	0.250	0.000
89:4	0.780	0.590	0.625	0.815	0.370	0.280	0.000	0.070	0.210	0.000
90:1	0.200	0.375	0.365	0.160	0.070	0.050	0.000	0.070	0.230	0.000
90:2	0.000	0.005	0.000	0.000	0.070	0.080	0.983	0.040	0.220	0.517
90:3[a]	0.645	0.660	0.700	0.660	0.320	0.360	0.017	0.820	0.300	0.572
90:4	0.005	0.010	0.030	0.015	0.280	0.380	0.000	0.040	0.250	0.000
91:1	0.000	0.005	0.000	0.003	0.230	0.050	0.000	0.130	0.240	0.000
91:2	0.000	0.000	0.000	0.000	0.170	0.060	0.073	0.000	0.250	0.572
91:3[a]	0.005	0.015	0.000	0.000	0.180	0.490	0.881	0.790	0.230	0.633
91:4	0.015	0.005	0.005	0.035	0.250	0.350	0.033	0.080	0.240	0.000

Notes: [a]Downturn in output growth occurred at that date.

both at one and four steps ahead. Therefore, measured by this yardstick, the PANEL2 model is, at least, competitive with the best univariate one.

Are these differences significant? Fig. 4 indicates that the distribution of the mean and median Theil-U statistics for the PANEL2 model are very much concentrated around

the point estimates we present in Table 4, both at the one and four steps ahead. Therefore, for example, the interquartile range of these distributions do not include the value of the Theil-U at that step for any of the competitors but PBVAR model. The Diebold and Mariano test indicates that the PBVAR, which is the closer competitor, produce forecasts which are qualitatively different from those of the PANEL2 model for the UK at one step ahead and for US, Germany, France and Italy at the four step ahead. Note also that while at the one step ahead the forecasts obtained by the g-prior and exchangeable models are not qualitatively different from those of the PANEL2 model, at four step ahead differences are noticeable for, at least, five countries. Hence, the PANEL2 model produce forecasts which are different, both qualitatively and quantitatively, from other models.

The results obtained using the MAD are somewhat similar but four features deserve comments. First, all shrinkage procedures are better than OLS at the one step horizon. The same is true at four steps horizons except for the case of g-prior, which is significantly worse. Second, unrestricted and simple BVAR display a somewhat mediocre performance at both horizons: the distribution of the MAD across countries is more concentrated but the mean and the median are above those obtained with univariate approaches. Third, the improvements obtained with panel VAR approaches are largest for US, Japan, Germany and France and the PANEL2 version of the model produces the best distribution of MAD at the one step horizon. Fourth, at the four step horizon the improvements obtained over univariate shrinkage procedures are small or at times negligible. The Diebold and Mariano test suggests that the forecasts produced by the PANEL2 model are qualitatively different from those of the other models when a MAD loss function is used. The differences are evident both for single countries and for median measures.

It is worth discussing the relative merits of the three models with cross-country interdependencies. Our refinement of the Litterman's prior, which employs both cross-sectional and time series a priori restrictions, has similar forecasting performance as the PBVAR model at both horizons. However, there are three features which are worth discussing. First, the maximized value of predictive density of the PANEL1 model is significantly higher than the one of the PBVAR model (-36.90 vs. -985.35) suggesting that the former fits the data for the in-sample period better. Second, θ_6, the time variation parameter in the process for λ is significant. Third, while in the PBVAR, the coefficient vector evolves with a persistence equal to 0.95 but with very small variance, in our PANEL1 specification the time varying component of the coefficients is close to be a white noise. This difference can be explained by examining the role of the parameters regulating the cross-sectional prior (i.e. the tightness on α and $\bar{\alpha}$). These parameters force a high degree of coherence across countries in the time invariant component and leave the time varying component to evolve randomly. In the PBVAR this distinction is not possible and to produce coefficients which are almost constant over time it is necessary to have close to a random walk dynamics coupled with a small variance. Using Eq. (12), one can see in fact that coefficients of the PANEL1 model are approximately constant over time and are tightly linked to each other because of the restrictions imposed on α_i. The omission of the fixed effect component in the PBVAR is therefore likely to bias upward estimates of the persistence parameter.

The performance of a fully hierarchical model with heteroschedasticity is somewhat better than the two versions of the Minnesota prior we consider. However, one should weight this improved performance against the computation costs. While forecasts for the fully hierarchical prior required several hours of computer time on a Pentium IV-700, the forecasts produced with a PBVAR or PANEL1 were obtained in a matter of minutes. We are currently experimenting with an alternative specification which drastically reduces the dimensionality of the time component. Preliminary results suggest that the computational costs are reduced by as much as 95% while the quality of the forecasts is only marginally affected.

In sum, using interdependencies and cross-sectional restrictions in the coefficients helps in improving forecasts at short-medium horizon. Given the poor results obtained by the univariate-TVC specification and the good performance of g-prior specification, it appears that it is the cross-sectional information that helps most in the exercise. Nevertheless, because of the substantial differences in the process for the growth rate of output across countries, improvements are not uniform and in some cases considering only domestic variables seems appropriate. Also, time variation in the coefficients is important only when coupled with heteroschedasticity in the variance probably because the structure of the economies is changing over time and outliers are present.

How good are various approaches in predicting turning points? Out of 96 total actual turning points in the sample, univariate approaches recognize between 72 and 75. Differences emerge when we try to predict upturns and non-upturns and for this type of turning points, Zellner's-g and TVC models are the best. Unrestricted VARs fare very poorly in this dimension and recognize about 10% less turning points than the best ones. The performance of the BVAR model is comparable to the one of Ridge and Exchangeable approaches but, contrary to them, it predicts upturns and non-upturns better than downturns and non-downturns. The performance of the PBVAR model is surprisingly poor: it is the second worst in recognizing the total number of turning points and is comparable to unrestricted VARs in predicting downturns and non-downturns. The PANEL2 model produces 76 correct turning point forecasts and recognizes the same number of upturns and downturns (38). From Fig. 4, we can see that the distribution of turning points selected by the PANEL2 model is very concentrated around the median and the interquartile range of the distribution of downturns includes only the value of 38 and of upturns the values 37 and 38. Hence, with the exception of the g-prior specification, the record of turning point recognition obtained with the model is significantly different than all the others.

Two other conclusions can be drawn from Table 5. First, different models are better in recognizing different types of turning points. If predicting downturns (and non-downturns) is more important than predicting upturns (and non-upturns) our results suggest that VAR, BVAR and PBVAR should not be used. Second, while in terms of linear forecasting statistics there was a clear ranking of procedures, with more complicated ones doing a better job, when we look at non-linear forecasting statistics, simple univariate approaches, and OLS in particular, are almost as good as other more refined approaches.

Given that our suggested specifications are good in forecasting on average, we would like to know if they are also good in predicting a specific episode of interest, i.e., the

downturn in real activity occurred in the US in 1990:3. This is an interesting episode because many commercial and government models, which were forecasting pretty well in the 1980s, failed to find any relevant signs that a downturn and a short recession were forthcoming (see e.g. Stock and Watson, 1993). Interestingly, all procedures except the univariate TVC model predict that there is a nonnegligible probability of a downturn in economic activity at 1990:3. For OLS and ridge methods this probability is larger than the threshold of 0.5, which we use to date a downturn. Single country VAR, with and without a Bayesian prior, are comparable to the best univariate procedures (probability 0.64 and 0.66 respectively) and significantly improve over Exchangeable and g-prior specifications. The PBVAR specification is overwhelmingly predicting a downward turn in 1990:3 (probability is 0.82) and does not produce any false alarm in the neighborhood of this date. The performance of the PANEL2 approach is somewhat mixed: the probability of a downturn in 1990:3 is above the 0.5 threshold but smaller than the best univariate and single country VAR models. However, the approach gives a strong warning signal also a quarter before the downturn occur, which is indicative of the good anticipatory features of the model. In comparison, the anticipatory signal given by the univariate TVC model is not corroborated at the time of actual downturn. Note also that while four out of five univariate approaches produce a false alarm in 1989:4, probably due to the stock market crash of the fall of 1989, the probabilities produced by PBVAR and PANEL models at dates other than 1990:3 are small. The latter four models in the table also produce a high probability of a downturn in 1991:3, a date where a downturn materialized. Finally, the downturn in 1989:2 is missed by all approaches but the univariate TVC and the PANEL2 models.

7. Conclusions

The task of this paper was to describe the issues of specification, estimation and forecasting in a macro-panel VAR model with interdependencies. The point of view used is Bayesian. Such an approach has been widely used in the VAR literature since the works of Doan et al. (1984), Litterman (1986), and Sims and Zha (1998), and provides a convenient framework where one can allow for both interdependencies and meaningful time variations in the coefficients. We decompose the parameter vector into two components, one which is unit specific and the other which is time specific. We specify a flexible prior on these two components which parsimoniously accounts for interdependencies in the cross section and for variations in the evolution of the parameters over time. The prior shares features with those of Lindley and Smith (1972), Doan et al. (1984) and Hsiao et al. (1999) and it is specified to have a hierarchical structure, which allows for various degrees of ignorance in the researcher's information about the parameters.

Bayesian VARs are known produce better forecasts than unrestricted VAR and, in many situations, ARIMA or structural models (Canova, 1995 for references). By allowing interdependencies and some degree of information pooling across units in the model specification, we introduce an additional level of flexibility which may improve the forecasting ability of these models.

In the case of fully hierarchical priors, the Gibbs sampler is employed to calculate posterior distributions and to construct forecasts. Such an approach is useful in our setup since it exploits the recursive features of the posterior distribution. We also consider a version of the Minnesota prior. In this case we employ the predictive density of the model to estimate unknown parameters and plug-in our estimates in the relevant formulas in an empirical Bayes fashion.

To illustrate the approach, we apply the methodology to the problem of predicting output growth, of forecasting turning points in output growth in the G-7 and computing the probability of a recession in the US. We show that our panel VAR approach is competitive and improves over existing univariate and simple BVAR models using either the Theil-U or the MAD criteria both at the one step and at the four steps horizons. The improvements are of the order of 5–15% with the Theil-U and about 2–8% with the MAD. The forecasting performance of our specification is also better than the one of a BVAR model which mechanically extends the Minnesota prior to the panel case. In terms of turning point prediction, the three Panel VARs are able to recognize about 80% of turning points in the sample and they turn out to be the best for this task, along with Zellner's g-prior shrinkage approach. The simple extension of the Litterman's prior to the panel case does poorly along this dimension and, among all the procedures employed, is the second worst. Finally, all the procedures produce a significant probability of a downturn at 90:3, the date selected by the NBER committee to terminate the long expansion of the 1980s. However, the other two downturn dates in our sample are correctly recognized only by one of our Panel VAR specifications.

We consider the current work the first step in developing a coherent theory for Bayesian Panel VAR models which take into consideration the nature of interdependencies, the similarities in the model across units and the existence of time variation in the coefficients. Extensions of the theory include the set-up of leading indicator index models; the formulation of interesting hypothesis on the nature of the interdependencies, the similarities across units and the variations over time; and the development of tools to undertake structural identification in these models. The work of Sims and Zha (1998) is the starting point for developments in this latter case.

Acknowledgements

We would like to thank Arnold Zellner, the Associate Editor, two anonymous referees, Albert Marcet, Alessandro Rebucci and the participants of seminars at Universitat Pompeu Fabra, Universidad de Alicante, Universidade do Minho, European Central Bank, the Bundesbank, the European Meetings of the Econometric Society, the ASSET meetings, the International Symposia in forecasting and the CEPR/BdI conference "Monitoring the Euro Area Business Cycles" for comments and suggestions. Ciccarelli's research benefitted from funding by IVIE, DGIMCYT (BEC2002-03097) and Generalitat Valenciana (CTIDIB/2002/175). All remaining errors are of the authors.

Appendix A. Definition of the matrices for the Gibbs sampler

$$\hat{W}_1 = W_1 + \sum_i (A_i E_i - \bar{A})' V^{-1} (A_i E_i - \bar{A}),$$

$$\hat{W}_2 = W_2 + \sum_t (\Lambda_t - \rho\Lambda_{t-1} - (1-\rho)\Lambda_0)'(\delta_t V_1)^{-1}(\Lambda_t - \rho\Lambda_{t-1} - (1-\rho)\Lambda_0),$$

$$\hat{M}_0 = M_0 + \sum_t (\mathbf{Y}_t - \mathbf{\Gamma}_t \mathbf{W}_t')H^{-1}(\mathbf{Y}_t - \mathbf{\Gamma}_t \mathbf{W}_t')',$$

$$\hat{P}_0 = P_0 + \sum_t{}' (\mathbf{Y}_t - \mathbf{\Gamma}_t \mathbf{W}_t')'\Sigma^{-1}(\mathbf{Y}_t - \mathbf{\Gamma}_t \mathbf{W}_t'),$$

$$\hat{\alpha} = \hat{V}_\alpha \left(\sum_t W_t'(\Sigma \otimes H)^{-1}(Y_t - Z_t\lambda_t) + \Delta^{-1} S_N \bar{\alpha} \right),$$

$$\hat{V}_\alpha = \left(\sum_t W_t'(\Sigma \otimes H)^{-1} W_t + \Delta^{-1} \right)^{-1},$$

$$\alpha^* = \hat{V}^* \left((V \otimes \Omega_1)^{-1} \sum_i R_i \alpha_i + \Psi^{-1}\mu \right),$$

$$\hat{V}^* = (N(V \otimes \Omega_1)^{-1} + \Psi^{-1})^{-1},$$

where $\delta_t = v_2^t + v_1(1 - v_2^t)/(1 - v_2)$, $\mathbf{Y}_t$ is the $N \times G$ matrix such that $vecr(\mathbf{Y}_t) = Y_t$, $\mathbf{\Gamma}_t - [vecr(\mathbf{\Gamma}_{1t}), \ldots, vecr(\mathbf{\Gamma}_{Nt})]'$ is $N \times Gk$ and $\mathbf{W}_t = (I_G \otimes X_t')$. Here, $vecr()$ is the row vectorization of a matrix; $\mathbf{\Gamma}_{it} = A_i + \Lambda_t E_i$ is a $G \times k$ matrix and $\alpha = vecr(A)$; $\lambda_t = vecr(\Lambda_t)$.

References

Ballabriga, F.C., Sebastian, M., Vallés, J., 1998. European asymmetries. Journal of International Economics 48, 233–253.

Berger, J.O., 1985. Statistical Decision Theory and Bayesian Analysis, 2nd Edition. Springer, New York.

Binder, M., Hsiao, C., Pesaran, M.H., 2000. Estimation and inference in short panel vector autoregressions with unit roots and cointegration. DAE Working Paper No. 0003, University of Cambridge.

Canova, F., 1993. Modelling and forecasting exchange rates with a Bayesian time-varying coefficient model. Journal of Economic Dynamics and Control 17, 233–261.

Canova, F., 1995. VAR: specification, estimation, testing, forecasting. In: Pesaran, H., Wickens, M.R., (Eds.), Handbook of Applied Econometrics, Vol. 1. Macroeconomics. Handbooks in Economics. Blackwell, Oxford and Malden, MA. pp. 73–138.

Canova, F., Ciccarelli, M., 1999. Forecasting and turning point predictions in a Bayesian panel VAR model. Working Paper No. 443, Universitat Pompeu Fabra, available at http://www.econ.upf.es.

Canova, F., Marcet, A., 1997. The poor stay poor: non-convergence across countries and regions. UPF Working Paper No. 137.

Carter, C.K., Kohn, R., 1994. On Gibbs sampling for state space models. Biometrika 81, 541–553.

Chamberlain, G., 1982. Multivariate regression model for panel data. Journal of Econometrics 18, 5–46.

Chamberlain, G., 1984. Panel data. In: Griliches, Z., Intriligator, M.D. (Eds.), Handbook of Econometrics, Vol. 2. North-Holland, Amsterdam, pp. 1247–1318.

Chib, S., Greenberg, E., 1996. Markov Chain Monte Carlo simulation methods in econometrics. Econometric Theory 12, 409–431.

Diebold, F., Mariano, R., 1995. Comparing predictive accuracy. Journal of Business and Economic Statistics 13, 253–263.

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

Forni, M., Hallin, M., Lippi, M., Reichlin, L., 2000. The generalized dynamic-factor model: identification and estimation. Review of Economics and Statistics 82, 540–554.

Garcia Ferrer, A., Highfield, F., Palm, F., Zellner, A., 1987. Macroeconomic forecasting using pooled international data. Journal of Business and Economic Statistics 5, 53–67.

Gelfand, A.E., Smith, A.F.M., 1990. Sampling-based approaches to calculating marginal densities. Journal of the American Statistical Association 85, 398–409.

Gelfand, A.E., Hills, S.E., Racine-Poon, A., Smith, A.F.M., 1990. Illustration of Bayesian inference in normal data models using Gibbs sampling. Journal of the American Statistical Association 85, 972–985.

Geman, S., Geman, D., 1984. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. IEEE Transactions on Pattern Analysis and Machine Intelligence 6, 721–741.

Gerlach, S., Smets, F., 1996. The monetary transmission mechanism: evidence from the G-7 countries. mimeo, Bank of International Settlements, Basle.

Hoffmaister, W., Roldós, J.F., 1997. Are business cycles different in Asia and Latin America? IMF Working paper No. 97/79.

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

Hsiao, C., Pesaran, M.H., Tahmiscioglu, A.K., 1999. Bayes estimation of short run coefficients in dynamic panel data models. In: Hsiao et al. (Eds.), Analysis of Panels and Limited Dependent Variable Models: In Honor of G.S. Maddala. Cambridge University Press, Cambridge, pp. 268–296.

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

Koop, G., 1992. Aggregate shocks and macroeconomic fluctuations: a Bayesian approach. Journal of Applied Econometrics 7, 395–411.

Ingram, B., Whiteman, C., 1994. Supplanting the Minnesota prior: forecasting macroeconomics series using real business cycles model prior. Journal of Monetary Economics 34, 497–510.

Lahiri, K., Moore, J., 1991. Leading economic indicators: new approaches and forecasting records. University of Chicago Press, Chicago, IL.

Lindley, D.V., Smith, A.F.M., 1972. Bayes estimates for the linear model (with discussion). Journal of the Royal Statistical Society B 34, 1–41.

Litterman, R.B., 1986. Forecasting with Bayesian Vector Autoregressions—five years of experience. Journal of Business and Economic Statistics 4, 25–38.

Lütkepohl, H., 1990. Introduction to Multiple Time Series Analysis. Springer, Berlin.

Pesaran, M.H., Smith, R., 1996. Estimating long run relationships from dynamic heterogeneous panels. Journal of Econometrics 68, 79–113.

Rebucci, A., 1998. External shocks, macroeconomic policy and growth: a panel VAR approach, Working Paper no. 40, Global Economic Institutions, available at http://www.cepr.org/GEI/geiwp36.htm.

Sims, C., Zha, T., 1998. Bayesian methods for dynamic multivariate models. International Economic Review 39, 949–979.

Stock, J., Watson, M., 1993. A procedure for predicting recessions with leading indicators: econometric issues and recent experience. In: Stock, J., Watson, M. (Eds.), Business Cycles, Indicators, and Forecasting, NBER Studies in Business Cycles, Vol. 28. University of Chicago Press, Chicago and London, pp. 95–153.

Stock, J., Watson, M., 2002. Macroeconomic forecasting using diffusion indexes. Journal of Business and Economic Statistics 20, 147–162.

Waggoner, D.F., Zha, T., 1998. Conditional forecasts in dynamic multivariate models. Working Paper No. 98-22, Federal Reserve Bank of Atlanta.

Zellner, A., Hong, C., 1989. Forecasting international growth rates using Bayesian shrinkage and other procedures. Journal of Econometrics 40, 183–202.

Zellner, A., Hong, C., Min, C., 1991. Forecasting turning point in international output growth rates using Bayesian exponentially weighted autoregression, time varying parameter, and pooling techniques. Journal of Econometrics 49, 275–304.