ÖSTERREICHISCHES INSTITUT
FÜR WIRTSCHAFTSFORSCHUNG

WORKING PAPERS

A Multi-country Approach to Analysing the Euro Area Output Gap

Florian Huber, Philipp Piribauer

560/2018

A Multi-country Approach to Analysing the Euro Area Output Gap

Florian Huber, Philipp Piribauer

WIFO Working Papers, No. 560

March 2018

Abstract

We develop a multivariate dynamic factor model that exploits euro area country-specific information on output and inflation for estimating an area-wide measure of the output gap. In the proposed multi-country framework we moreover allow for flexible stochastic volatility specifications for both the error variances and the innovations to the latent quantities in order to deal with potential changes in the commonalities of business cycle movements. By tracing the relative importance of the common euro area output gap component as a means to explaining movements in both output and inflation over time, the paper provides valuable insights in the evolution of the degree of synchronicity of the country-specific business cycles. In an out-of-sample forecasting exercise, the paper shows that the proposed approach performs well as compared to other well-known benchmark specifications.

E-mail address: philipp.piribauer@wifo.ac.at
2018/077/W/0

Medieninhaber (Verleger), Hersteller: Österreichisches Institut für Wirtschaftsforschung • 1030 Wien, Arsenal, Objekt 20 • Tel. (43 1) 798 26 01-0 • Fax (43 1) 798 93 86 • http://www.wifo.ac.at/ • Verlags- und Herstellungsort: Wien
Die Working Papers geben nicht notwendigerweise die Meinung des WIFO wieder
Kostenloser Download: http://www.wifo.ac.at/wwa/pubid/61004

A multi-country approach to analyzing the Euro Area output gap *

Florian Huber[1] and Philipp Piribauer[†2]

[1]Vienna University of Economics and Business (WU)
[2]Austrian Institute of Economic Research (WIFO)

Abstract

We develop a multivariate dynamic factor model that exploits euro area country-specific information on output and inflation for estimating an area-wide measure of the output gap. In the proposed multi-country framework we moreover allow for flexible stochastic volatility (SV) specifications for both the error variances and the innovations to the latent quantities in order to deal with potential changes in the commonalities of business cycle movements. By tracing the relative importance of the common euro area output gap component as a means to explaining movements in both output and inflation over time, the paper provides valuable insights in the evolution of the degree of synchronicity of the country-specific business cycles. In an out-of-sample forecasting exercise, the paper shows that the proposed approach performs well as compared to other well-known benchmark specifications.

Keywords: European Business Cycles, Dynamic factor model, Forecasting

JEL Codes: E32, C11, C32, C53

*The research carried out in this paper was supported by funds of the Oesterreichische Nationalbank (Oesterreichische Nationalbank, Anniversary Fund, project number: 17382).

†Corresponding author: Philipp Piribauer, Austrian Institute of Economic Research (WIFO), Arsenal 20, 1030 Vienna, Austria. E-mail: philipp.piribauer@wifo.ac.at.

1 Introduction

In this paper we develop a multi-country business cycle model for the euro area (EA). The proposed model assumes that country-specific business cycles are driven by a common latent factor and thus exploits cross-sectional information in the data. To control for changes in the degree of synchronicity, we furthermore assume that the innovations to the latent quantities as well as the measurement errors are time-varying and follow a stochastic volatility (SV) process. One methodological key innovation is the introduction of global-local shrinkage priors on the process variances to the state equations describing the law of motion of the logarithmic volatility components, effectively shrinking the system towards a homoskedastic specification, if necessary.

The model aims to connect the literature on output gap modeling (see, among many others, Kuttner, 1994; Orphanides and Van Norden, 2002; Basistha and Nelson, 2007; Planas et al., 2008) that focuses on estimating the output gap based on data for a single country/regional aggregate with the literature on dynamic factor models (Otrok and Whiteman, 1998; Kim and Nelson, 1999; Kose et al., 2003; Breitung and Eickmeier, 2015; Jarocinski and Lenza, 2015). We assume that output as well as inflation across selected EA member states share a common cyclical behavior, pointing towards an underlying area-wide latent gap process. Recent post-crisis evidence, however, also points towards significant deviations of countries that share a common set of macroeconomic fundamentals from this general pattern. To control for this, we assume that measurement errors as well as country-specific trend components display conditional heteroskedasticity, which provides sufficient flexibility to capture deviations from a common cyclical component in the presence of idiosyncratic shocks.

This increased flexibility, however, is costly in terms of additional parameters to estimate. We thus follow the recent literature on state space modeling (Frühwirth-Schnatter and Wagner, 2010; Belmonte et al., 2014; Kastner and Frühwirth-Schnatter, 2014; Bitto and Frühwirth-Schnatter, 2016; Feldkircher et al., 2017) and exploit a non-centered parameterization of the model (see Frühwirth-Schnatter and Wagner, 2010) to test whether SV is supported by the data. The non-centered parameterization allows treating the square root of the process innovation variances as standard regression coefficients, implying that standard shrinkage priors can be used. Here we follow Griffin and Brown (2010) and use a variant of the Normal-Gamma (NG) shrinkage prior that introduces a global shrinkage component that is applied to all process variances simultaneously, forcing all of them towards zero. Local shrinkage parameters are then used to drag sufficient posterior mass away from zero even in the presence of strong global shrinkage, allowing for non-zero process variances in the presence of strong global shrinkage.

Our measure of the output gap is closely linked to estimates reported in previous studies (Planas et al., 2008; Jarocinski and Lenza, 2015). To assess how much variance of the stationary component of a given time series is explained by the cycle, we compute the commonalities of the factor model over time. Using this measure we find that for output, the amount of variation explained is high with relatively little variation over time. For inflation, we generally observe lower levels of explained variation but the amount of time variation appears to be much higher, reaching a peak during the early 2000s and staying high afterwards. To assess the sensitivity of output and inflation across Europe, we perform a simple counterfactual exercise. Specifically, we shock the equation for the common output gap and inspect whether there exist country-specific differences.

The paper moreover evaluates the performance of the modeling approach in terms of forecasting, paying particular attention to how much the introduction of a common output gap improves predictive capabilities. Compared to a range of simpler alternatives that range from univariate benchmark models to models that use alternative ways to calculate the output gap, our model markedly improves output predictions. For inflation, the results appear to be rather mixed, with the model approach proposed in Stock and Watson (1999; 2007) performing best.

The remainder of the paper is structured as follows. Section 2 describes the econometric framework adopted. After providing an overview of the model, we discuss the Bayesian prior choice and briefly summarize the main steps involved in estimating the model. Section 3 presents the empirical application, starting with a summary of the dataset used and inspects various key features of our model. The section moreover studies the dynamic impacts of business cycle shocks to the country-specific output and inflation series. In a forecasting exercise, Section 4 compares the out-of-sample predictive performance of our model with other specifications. The final section summarizes and concludes the paper.

2 Econometric framework

2.1 A dynamic factor model for the euro area

Let denote y_{jt} and Δp_{jt} output and inflation for country $j = 1, \ldots, N$ in period $t = 1, \ldots, T$, respectively. We assume that output and inflation feature a country-specific non-stationary trend component τ_{jt}^{k}, for $k \in \{y, \Delta p\}$, and depend on a common cyclical component f_t,

$$y_{jt} = \tau_{jt}^{y} + \alpha_j^{y} f_t + \epsilon_{jt}^{y}, \tag{2.1}$$

$$\Delta p_{jt} = \tau_{jt}^{\Delta p} + \alpha_j^{\Delta p} f_t + \epsilon_{jt}^{\Delta p}, \tag{2.2}$$

with $\epsilon_{jt}^{k} \sim \mathcal{N}(0, e^{h_{jt}^{k}})$ being a set of independent heteroskedastic shocks. The sensitivity of output and inflation with respect to movements in the common component f_t are governed by a set of factor loadings α_j^{k}.

We augment the model by two additional equations that measure euro area output, y_{0t}, and inflation Δp_{0t},

$$y_{0t} = \tau_{0t}^{y} + f_t, \tag{2.3}$$

$$\Delta p_{0t} = \tau_{0t}^{\Delta p} + \alpha_0^{\Delta p} f_t + \epsilon_{0t}^{\Delta p}. \tag{2.4}$$

Note that Eq. (2.4) implies that f_t can be interpreted as the output gap, i.e. the deviation of output from trend output, measured through τ_{0t}^{y}.

The time-varying components of Eq. (2.1) to Eq. (2.4) are assumed to evolve according to a vector autoregressive model of order two,

$$\underbrace{\begin{bmatrix} \boldsymbol{\tau}_t^{y} \\ \boldsymbol{\tau}_t^{\Delta p} \\ f_t \end{bmatrix}}_{\boldsymbol{\tau}_t} = \underbrace{\begin{bmatrix} \boldsymbol{I}_{N+1} & \dots & 0 \\ \vdots & \boldsymbol{I}_{N+1} & \vdots \\ 0 & \dots & \phi_1 \end{bmatrix}}_{\boldsymbol{A}_1} \underbrace{\begin{bmatrix} \boldsymbol{\tau}_{t-1}^{y} \\ \boldsymbol{\tau}_{t-1}^{\Delta p} \\ f_{t-1} \end{bmatrix}}_{\boldsymbol{\tau}_{t-1}} + \underbrace{\begin{bmatrix} 0 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \phi_2 \end{bmatrix}}_{\boldsymbol{A}_2} \underbrace{\begin{bmatrix} \boldsymbol{\tau}_{t-2}^{y} \\ \boldsymbol{\tau}_{t-2}^{\Delta p} \\ f_{t-2} \end{bmatrix}}_{\boldsymbol{\tau}_{t-2}} + \underbrace{\begin{bmatrix} \boldsymbol{\eta}_t^{\tau_y} \\ \boldsymbol{\eta}_t^{\tau_{\Delta p}} \\ \eta_t^{f} \end{bmatrix}}_{\boldsymbol{\eta}_t}, \tag{2.5}$$

with $\boldsymbol{\tau}_t^k = (\tau_{0t}^k, \tau_{1t}^k, \dots, \tau_{Nt}^k)'$, and $\boldsymbol{\eta}_t^k = (\eta_{0t}^k, \eta_{1t}^k, \dots, \eta_{Nt}^k)'$.

For the AR(2) parameters ϕ_1 and ϕ_2 we follow Planas et al. (2008) and reparameterize the state equation in f_t as follows,

$$f_t = 2\ Q\ \cos(2\pi/\gamma) f_{t-1} - Q^2 f_{t-2} + \eta_t^f. \tag{2.6}$$

Hereby, Q determines the amplitude and γ the frequency of the cycle. This parameterization has the convenient property that prior information on the length as well as the intensity of the business cycle can be introduced in a relatively easy manner.

For the sake of simplicity, Eq. (2.5) may be rewritten in terms of a standard multivariate regression model

$$\boldsymbol{\tau}_t = \boldsymbol{A}\boldsymbol{X}_t + \boldsymbol{\eta}_t, \tag{2.7}$$

with $\boldsymbol{A} = (\boldsymbol{A}_1, \boldsymbol{A}_2)$ and $\boldsymbol{X}_t = (\boldsymbol{\tau}'_{t-1}, \boldsymbol{\tau}'_{t-2})'$.

Similarly to the measurement errors we follow Stock and Watson (1999; 2007) and assume that the shocks to the states are mutually orthogonal with time-varying variances,

$$\begin{bmatrix} \boldsymbol{\eta}_t^{\tau_y} \\ \boldsymbol{\eta}_t^{\tau_{\Delta p}} \\ \eta_t^f \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t). \tag{2.8}$$

$\boldsymbol{\Sigma}_t$ is a $K \times K$ diagonal variance-covariance matrix with $K = 2(N+1)+1$ and typical element $\sigma^2_{ii,t} = e^{s_{it}}$ for $i = 1, \ldots, K$.

We assume that the error variances of the observation equations in Eq. (2.1) to Eq. (2.4) as well as the shocks to the states in Eq. (2.5) follow a stationary stochastic volatility process,

$$h^k_{jt} = \mu^k_j + \varrho^k_j(h^k_{jt-1} - \mu^k_j) + \omega^k_{jt}, \qquad \omega^k_{jt} \sim \mathcal{N}(0, \vartheta^k_{hi}), \text{ for } j = 1, \ldots, N, \tag{2.9}$$

$$s_{it} = \mu^\tau_i + \varrho^\tau_i(s^\tau_{it-1} - \mu^\tau_i) + \omega^\tau_{it}, \qquad \omega^\tau_{it} \sim \mathcal{N}(0, \vartheta^\tau_{si}), \text{ for } i = 1, \ldots, K. \tag{2.10}$$

The autoregressive parameters are given by ϱ^k_j and ϱ^τ_i while the means of the log-volatility processes are given by μ^k_j and μ^τ_i. Finally, the state innovation variances are given by ϑ^k_{hi} and ϑ^τ_{si}. It worth noting that if a given ϑ^k_{hi} equals zero, the corresponding variance is constant.

The main implications of our model are as follows. First, we extract the cyclical component of EA output using a dynamic factor model with manifest variables given by output and inflation of the member countries. This, in turn, will enable us to assess to what extent output and inflation movements within the EA are driven by a common cyclical component that can be interpreted as an area-wide output gap. Second, we assume that the gap in inflation is proportional to the output gap plus a white noise error. This is mainly in line with recent applications of multivariate unobserved component models for inflation (see, for example, Stella and Stock 2013). Finally, the flexible stochastic volatility assumption on the measurement error variances and process variances captures the notion that trend output and inflation tend to display more uncertainty during turbulent times. This feature is particularly crucial for producing precise predictive densities.

2.2 Bayesian inference

The model outlined in the previous subsection is quite flexible, however, heavily parameterized. This calls for regularization in the form of Bayesian shrinkage. We start by outlining a general strategy to shrink our proposed factor model towards a simpler specification when it comes to deciding on what components should feature conditional heteroskedasticity. We moreover present our prior setup on the remaining free coefficients of the model described in the previous subsection.

In what follows we focus on how to flexibly shrink the variances in the state innovation variances in Eq. (2.5) to zero. Shrinkage to homoskedasticity in the observation equation is achieved in a similar manner. Following Frühwirth-Schnatter and Wagner (2010) and Kastner and Frühwirth-Schnatter (2014), squaring and taking logs of the jth equation of Eq. (2.7) and rewriting yields a non-centered parameterization of the state space model,

$$\tilde{\varepsilon}_{jt} = \mu_j^{\tau} + \sqrt{\vartheta_{sj}^{\tau}}\tilde{s}_{jt} + v_{jt}, \quad v_{jt} \sim \ln \chi(1) \tag{2.11}$$

$$\tilde{s}_{jt} = \varrho_j^{\tau}\tilde{s}_{jt-1} + w_{jt}, \quad w_{jt} \sim \mathcal{N}(0,1), \tag{2.12}$$

$$\tilde{s}_{jt} = \frac{s_{jt} - \mu_j^{\tau}}{\sqrt{\vartheta_{sj}^{\tau}}}, \tag{2.13}$$

with $\tilde{\varepsilon}_{jt} = \ln(\tau_{jt} - \boldsymbol{A}_{j\bullet}\boldsymbol{X}_t)^2$ and $\boldsymbol{A}_{j\bullet}$ selecting the jth row of the matrix $\boldsymbol{A}$. Equation (2.12) implies that the process variance ϑ_{sj}^{τ} as well as the unconditional mean μ_j^{τ} is moved into the observation equation. Conditional on the full history of the normalized log-volatilities and the mixture approximation to render Eq. (2.12) conditionally Gaussian (Kim et al. 1998), the process variances and parameters can be obtained by estimating an otherwise standard Bayesian linear regression model.

This implies that standard shrinkage priors can specified on $\sqrt{\theta_{sj}}$. Here we adopt the flexible shrinkage prior proposed in Griffin and Brown (2010) and recently adopted within the framework of univariate state space models in Bitto and Frühwirth-Schnatter (2016),

$$\vartheta_{sj}^{\tau} \sim \mathcal{G}\left(1/2, 1/(2B_{sj})\right) \Leftrightarrow \sqrt{\vartheta_{sj}^{\tau}} \sim \mathcal{N}(0, B_{sj}), \tag{2.14}$$

with B_{sj} being a shrinkage hyperparameter with

$$B_{sj} \sim \mathcal{G}(\kappa_s, \kappa_s \xi_s/2), \quad \xi \sim \mathcal{G}(d_0, d_1). \tag{2.15}$$

ξ_s is a so-called global shrinkage parameter that pushes $\sqrt{\boldsymbol{\vartheta}_s^{\tau}} = (\sqrt{\vartheta_{s1}^{\tau}}, \dots, \sqrt{\vartheta_{sK}^{\tau}})'$ to zero. Moreover, κ_s and d_0, d_1 are hyperparameters specified by the researcher. Intuitively speaking, the global shrinkage parameter exerts shrinkage towards the origin while B_{sj} serves to pull certain elements of $\sqrt{\boldsymbol{\vartheta}_s^{\tau}}$ away from zero even if ξ_s is large (i.e. heavy global shrinkage is introduced). Notice that ξ_s introduces prior dependence and can be thought of as a common factor that efficiently pools information across coefficients.

The same prior choice is also adopted for the process innovation variances in the log volatility equations for the measurement errors (see Eq. (2.9)), i.e.

$$\sqrt{\vartheta_{hi}^{k}} \sim \mathcal{N}(0, B_{hj}^{k}), B_{hj}^{k} \sim \mathcal{G}(\kappa_h, \kappa_h \xi_h/2), \quad \xi_h \sim \mathcal{G}(e_0, e_1). \tag{2.16}$$

Notice that the common parameter κ_h pools information on error variances in the log-volatilities across all output and inflation equations, effectively introducing global shrinkage across variable types.

Bitto and Frühwirth-Schnatter (2016) label this prior a double Gamma prior (if placed on the variances). Consistent with the literature we set $\kappa_s = \kappa_h = 0.1$ and $d_0 = d_1 = e_0 = e_1 = 0.01$. This choice is consistent with heavy shrinkage on all process variances while maintaining heavy tails in the underlying marginal prior.

Following Planas et al. (2008), we specify a Beta distributed prior on Q,

$$Q \sim \mathcal{B}(a_Q, b_Q), \tag{2.17}$$

with a_Q and b_Q denoting hyperparameters. For γ we adopt also adopt a Beta prior with

$$\frac{\gamma - \gamma_L}{\gamma_H - \gamma_L} \sim \mathcal{B}(a_\gamma, b_\gamma). \tag{2.18}$$

This prior restricts the support of γ by specifying a minimum wave length γ_L, which is set equal to two, and a maximum length γ_H set equal to T. The parameters a_γ, b_γ are fixed hyperparameters.

For the remaining parameters of Eqs. (2.9) - (2.10) we follow Kastner and Frühwirth-Schnatter (2014) and use an uninformative Gaussian prior on the unconditional mean, i.e. $\mu_j^k \sim \mathcal{N}(0, 10^2)$, μ_j^τ for all i, j, k as well as a Beta prior on the persistence parameter $\varrho_j^k \sim \mathcal{B}(25, 5)$ and $\varrho_i^\tau \sim \mathcal{B}(25, 5)$. On the factor loadings α_j^k we use a sequence of independent Gaussian priors,

$$\alpha_j^k \sim \mathcal{N}(0, 1).$$

Finally, we specify the priors on the initial state $\boldsymbol{\tau}_0$ and the log-volatilities to be fairly uninformative with each element being normally distributed with zero mean and a variance 10^2.

Estimation is carried out using a Markov chain Monte Carlo (MCMC) algorithm described in section A. The algorithm is repeated 50,000 times with the first 25,000 draws discarded as burn-in. Convergence and mixing of most model parameters appear to be satisfactory. However, we find a substantial degree of autocorrelation for the factor loadings in selected countries. To assess the sensitivity of our findings, we thus re-estimated the model a moderate number of times based on different initial values. The corresponding findings appear to be remarkably robust.

3 Empirical application

3.1 Data overview and model specification

For the empirical application, we use quarterly data for economic output and inflation from 1985Q1 to 2013Q3. Our country sample comprises Austria (AT), Belgium (BE), Finland (FI), France (FR), Germany (DE), Greece (GR), Italy (IT), Netherlands (NL), Portugal (PT) and Spain (ES). Economic output and inflation is measured in terms of (the logarithm of) real seasonally adjusted gross domestic product and the rate of consumer price inflation, respectively.

3.2 Key features of the model

In this section we present some key features of the proposed model. We start by discussing the estimated output gap along other competing measures of the output gap in Fig. 1. The black line in Fig. 1 presents the posterior median of the estimated output gap for the euro area resulting from the model framework sketched above (DFM-SV). In addition, the orange line in Fig. 1 shows the output gap using a standard HP-filter (Hodrick and Prescott 1997), whereas the blue line depicts the output gap based on recent work by Hamilton (2017) as a straightforward alternative to the HP-filter.[1]

[1] Specifically, the approach proposed by Hamilton (2017) is based on standard regressions of future output on a constant and lags of the four most recent observations of output.

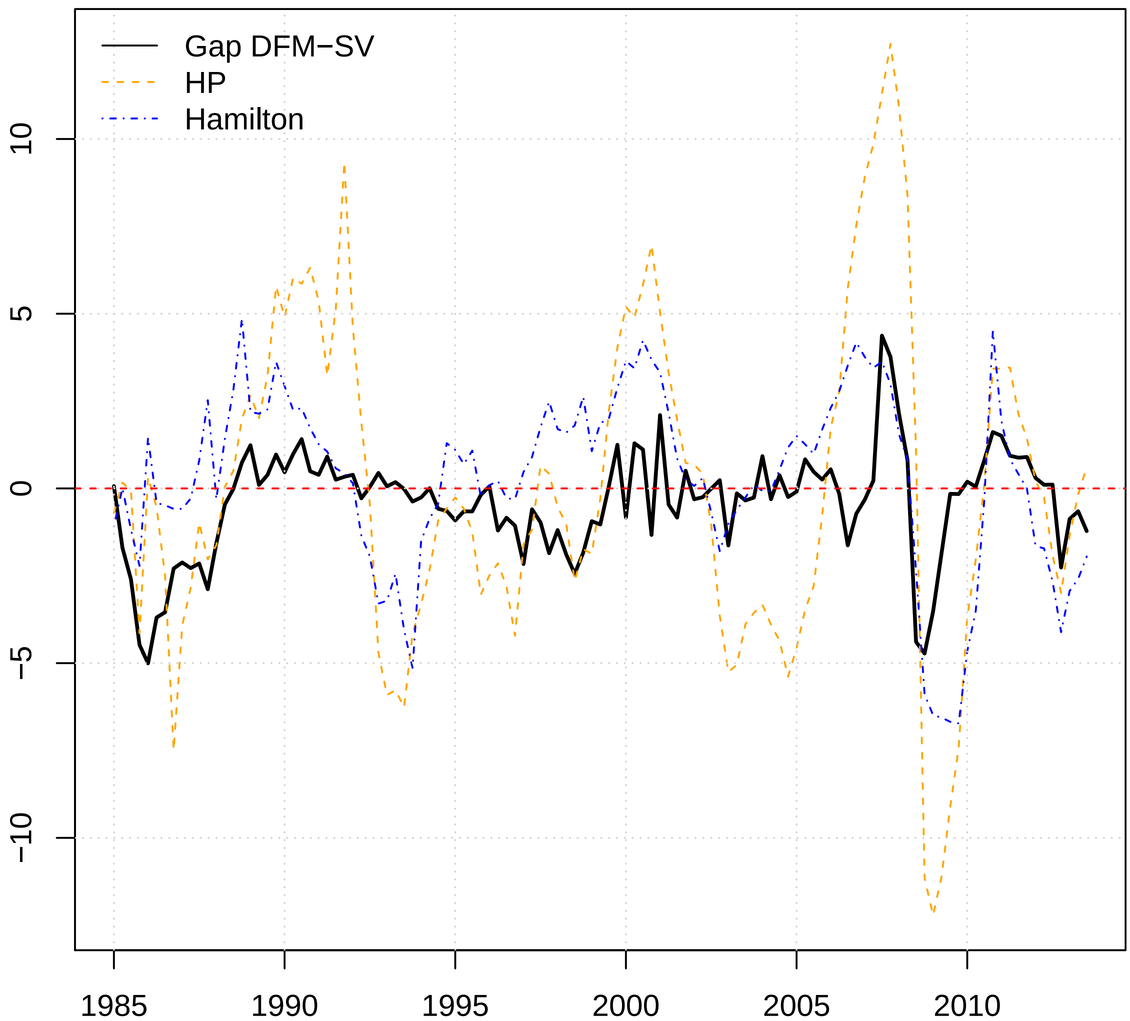

Fig. 1: Estimated euro area output gap

Considering the results in the end of the 1980s indicates that the aggregate growth rate of output has been below its potential for quite some time. Moreover, especially the economic turmoils due to the financial crisis as well as the recent Euro crisis periods appear particularly pronounced.

Overall, the estimated peaks and troughs of the estimated output gap component by DFM-SV shows marked similarities to the gap estimates of recent work by Planas et al. (2008) or Jarocinski and Lenza (2015). However, some differences between DFM-SV and the alternative specifications depicted in Fig. 1 are still visible. Particularly in the 1990s, the estimated peaks and troughs of the output gap produced by DFM-SV appear to be lower as compared to Hamilton or HP. The figure moreover shows that the estimated gap component in our multi-country framework appears to react more quickly in the wake of the financial crisis. We conjecture that the faster reaction of the output gap based on our proposed specification can be traced back to the fact that information stemming from other countries' output reactions, with some countries leading the others.

The estimated posterior median of the stochastic volatility component of the euro area output gap is presented it the top panel of Fig. 2 along with the lower 16th and upper 84th percentile of the credible interval (orange lines). Figure Fig. 2 reports a significant increase in volatility during periods of economic stress, especially during the crisis associated with the burst of the dot-com bubble, the 9/11 terrorist attacks and the period of the global financial crisis. Notice that volatility remains at an elevated level during the crisis of the Euro area in 2011 and afterwards.

An additional indication for the importance of accounting for time variation in the error variances is depicted in the bottom panel of Figure 2. The figure reports the signed square root of the error variance. Since the sign of the square roots of ϑ_{hi}^{k} and ϑ_{si}^{τ} in equations (2.9) and (2.10) are not identified, the non-identification of the respective signs may be exploited by randomly switching the sign of ϑ_{hi}^{k} and ϑ_{si}^{τ} and assessing the corresponding posterior distribution. If the resulting posterior density is centered on zero, we obtain only limited evidence in favor of a heteroskedastic specification for the error variances. By contrast, a bimodal posterior distribution points towards time-variation in the corresponding volatility component. The bottom panel in Figure 2 clearly shows a bimodal posterior, providing a simple yet effective visual assessment whether heteroscedasticity is needed.

In order to assess the importance of accounting for a common component in our multi-country framework over time, we moreover compute variance decompositions for the country-specific variations in output and inflation. Figure 3, for example, shows the share

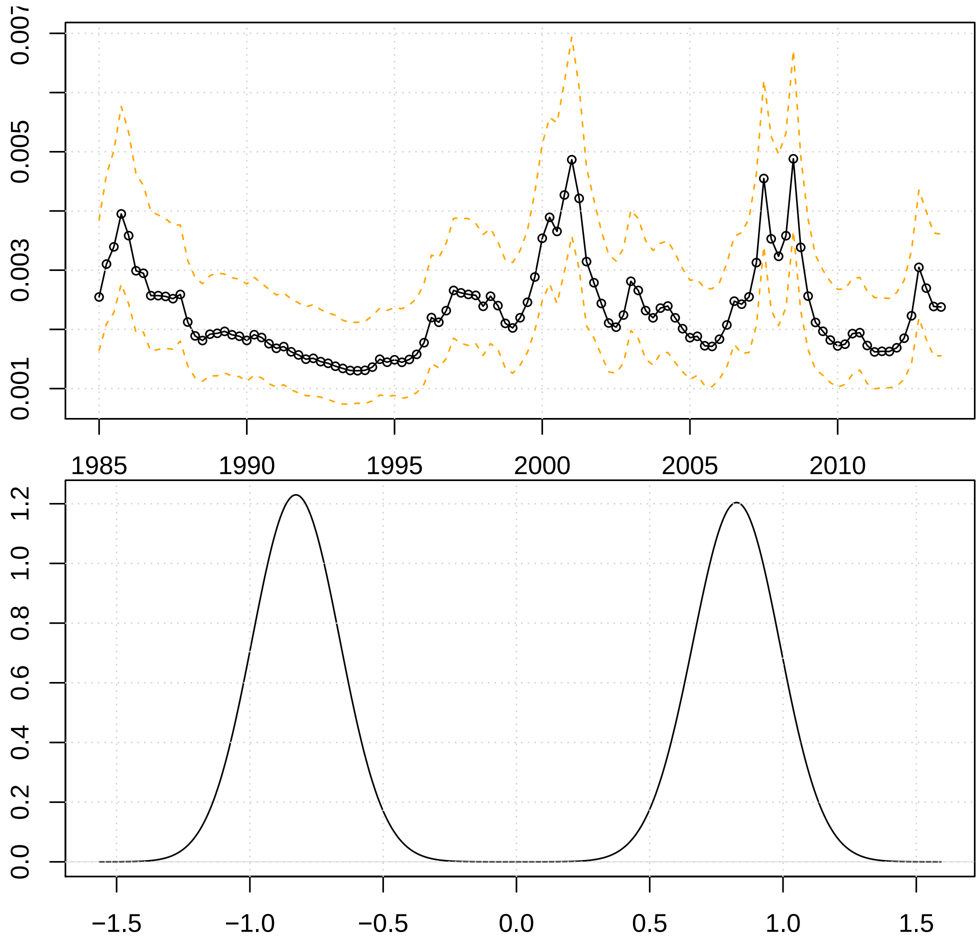

Fig. 2: Stochastic volatility of the EA output gap (top panel) and signed square root of error variance (bottom panel)

of the variance explained by the common gap component in the respective country-specific output equation. The respective decompositions for the inflation series are depicted in Fig. 4. Since the country-specific trend components given in Eqs. (2.1) and (2.2) are non-stationary and thus increase over time in an unbounded manner, we compute the share attributable to the common component in terms of the stationary part of the model (i.e. by considering the variance of the measurement errors as well as the gap component).

The proportion of variance explained for output given in Fig. 3 appears to be rather large across all output series considered. Concerning the overall variance explained for economic output, some countries under scrutiny, however, exhibit notable smaller shares as compared to the remaining economies. These countries include Germany, Greece, Spain and Italy, pointing towards a more diverging behavior in output trajectories. With some exceptions, the share of variance explained is typically above a 50% threshold for the entire time span. Some notable exceptions include Germany, Greece, Spain, or Italy, where overall shares appear markedly lower as compared to the other countries. Similar to the stochastic volatility component depicted in Figure 2, Figure 3 also shows pronounced increases in the variance decompositions in the early 2000s, during the financial crisis as well as in the wake of the recent Euro crisis. This, again, provides considerable evidence that during economic downturns, the cross-correlation across countries increases and business cycle synchronization becomes stronger.

As compared to the variance decompositions for economic output, the respective decompositions for the country-specific inflation series appear more heterogeneous, both across countries and time. Especially for Belgium, Germany, France and Italy, the share explained by the common component appears particularly high across the estimation period. An additional interesting aspect of figure is the sharp decline in variation explained in the beginning of the 1990s for Germany. This period captures the German reunification process, with the sharp increase in explanatory power of the common component pointing towards a better synchronization of price dynamics across Europe. In this period, similar patterns are visible due to the economic turmoils in Italy.

For the remaining countries under consideration, the common gap component appears to explain far less, especially in the beginning of the sample. Due to the convergence process of the country-specific inflation series in the euro area in the late 1990s, Fig. 4 also shows pronounced increases in the explained variation in this period. In the late 2000s, however, a diverging behavior in the inflation series is clearly visible, translating in sharp decreases in the variance decompositions.

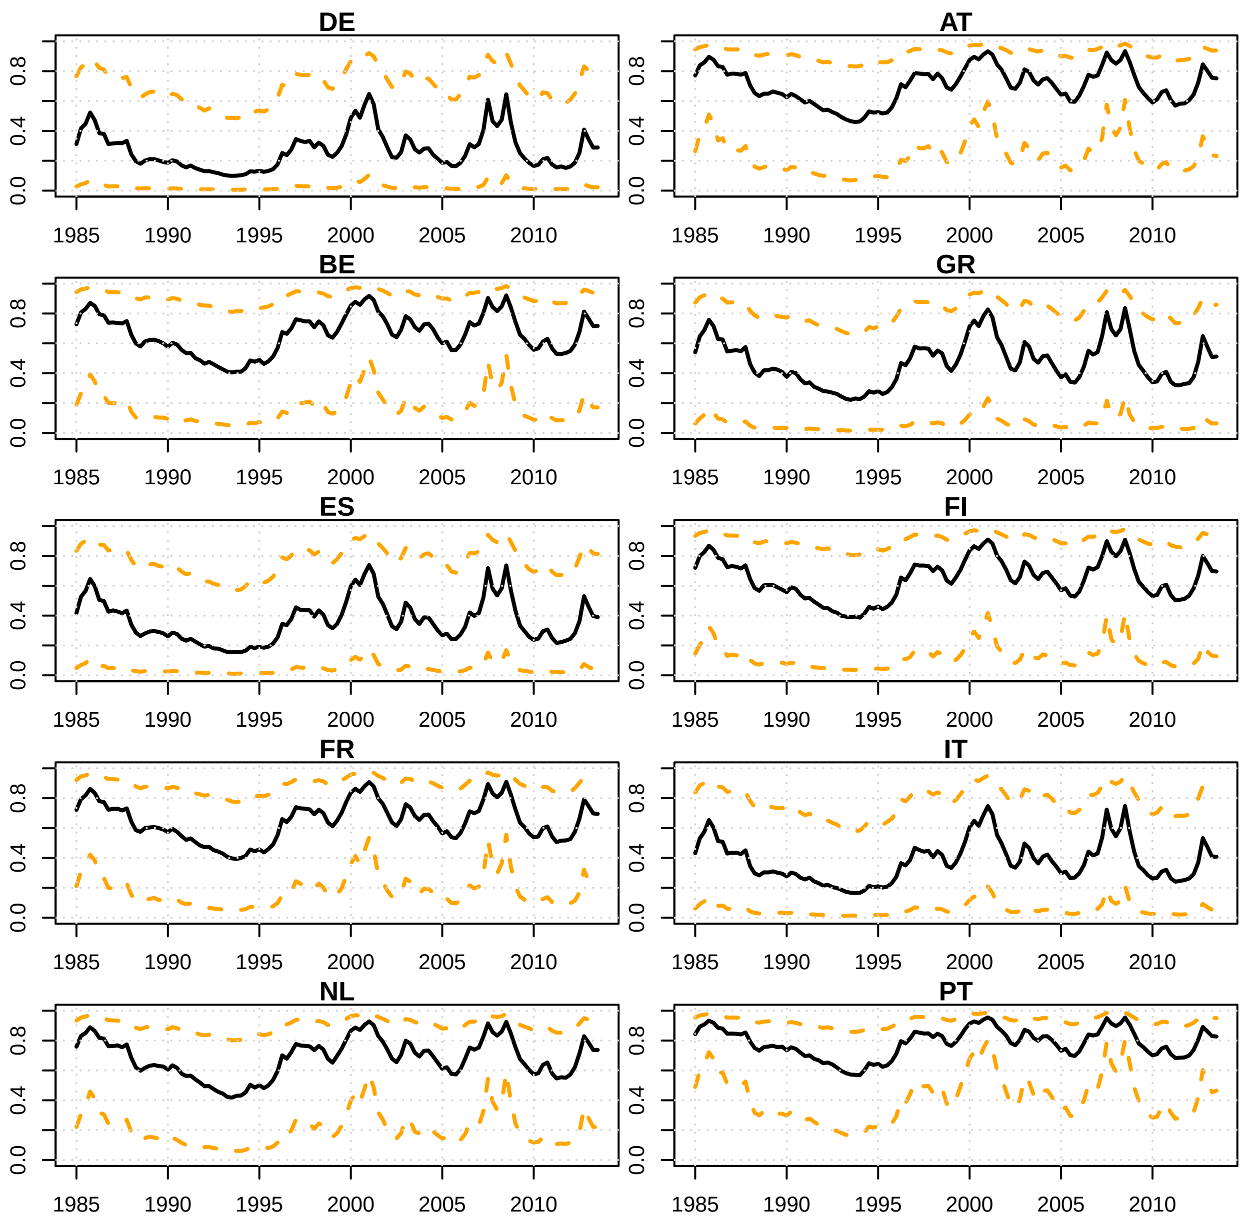

Fig. 3: Variance decomposition for output

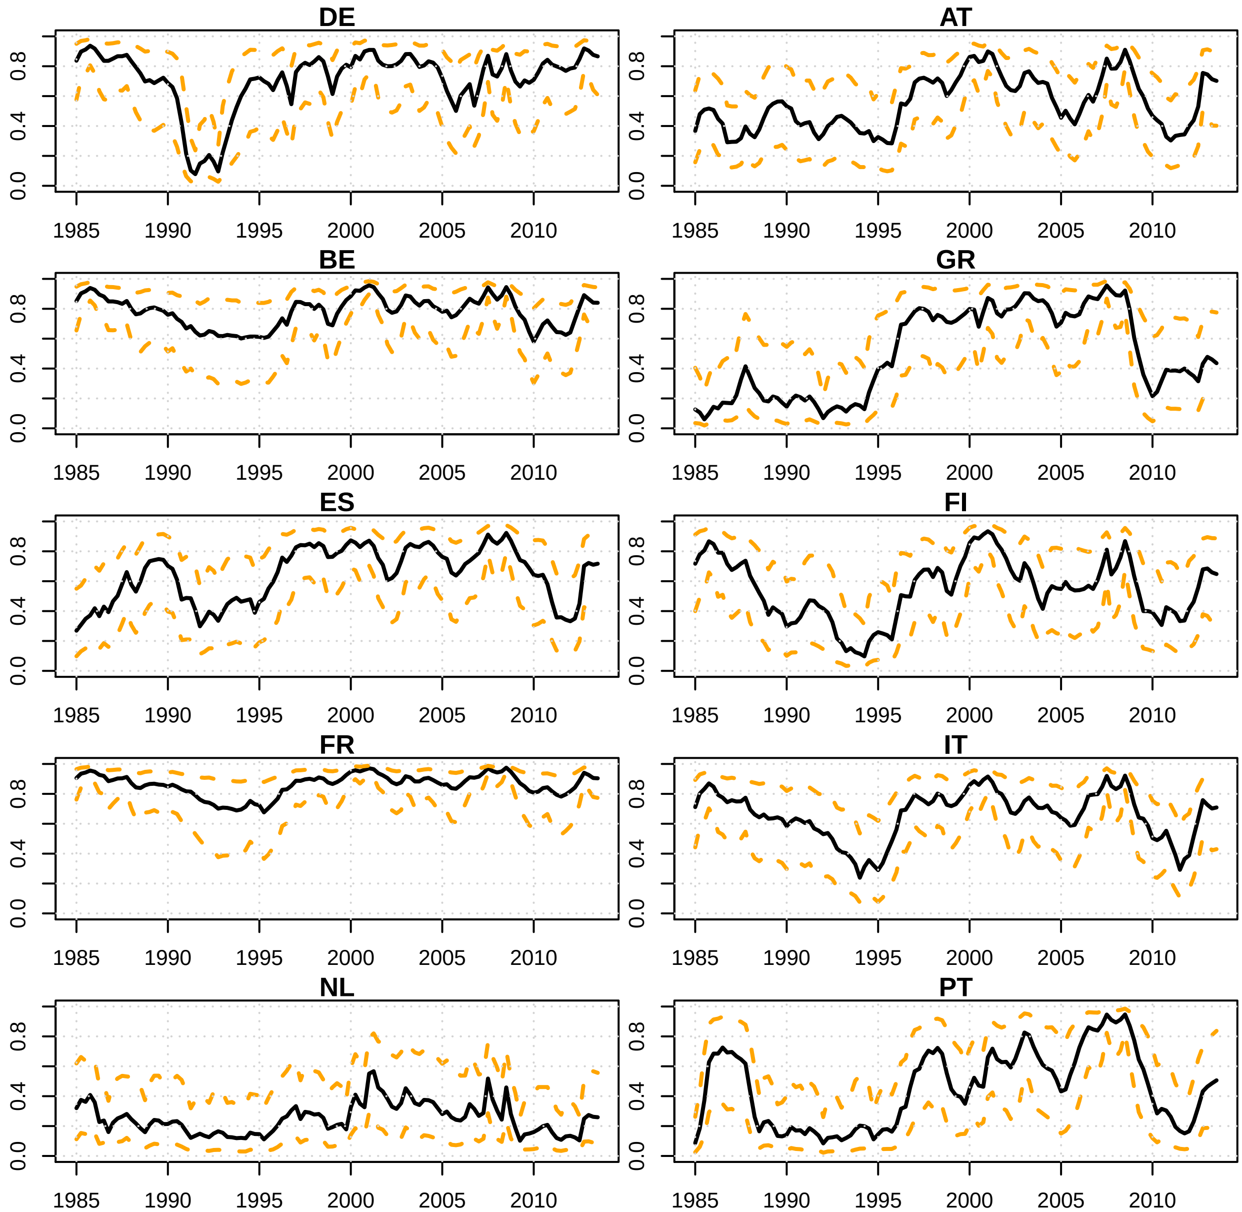

Fig. 4: Variance decomposition for inflation

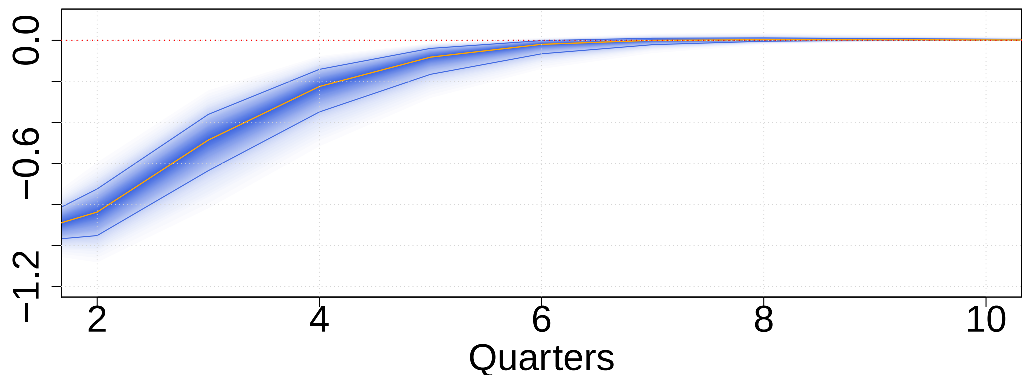

Fig. 5: Responses of the euro area output gap to a one standard error business cycle shock

3.3 Dynamic responses of output and inflation to an area-wide business cycle shock

This subsection aims at studying the dynamic effect of business cycle shocks to output and inflation across the euro area. Figure 5 depicts the posterior distribution of the dynamic response of the common output gap to a (negative) one standard deviation business cycle shock. The orange line in the figure shows the median responses over time along with lower 16th and upper 84th percentiles of the posterior distribution (in blue).

We find a negative and immediate impact on the common gap component. This effect appears to die out after around five to six quarters. Notice that Fig. 5 only measures the response of the latent gap component. Polcy makers, however, might be interested in how changes in the common cycle impact prices and output within each country considered. To this end, the left panel in Fig. 6 presents boxplots of the posterior distribution of the maximum output responses with red whiskers indicating 16th and upper 84th percentiles of the posterior distribution. Maximum absolute responses to the inflation series are shown in the right panel in Fig. 6.

As shown in Fig. 6 (a), the maximum posterior responses for output appear rather heterogeneous among the countries in the sample. For five out of ten economies, we find pronounced movements in output in reaction to a shift in the common cycle. The strongest negative reactions can be found in France, Austria, and Finland. Interestingly, we find no evidence that Germany reacts to a common business cycle shock when the maximum is considered. Notice, however, that this could also be purely driven by the selection of the absolute maximum response which might be linked to a particular impulse response horizon.

Turning to the maximum responses of inflation in panel (b) of Fig. 6 reveals that all countries face lower levels of inflation. The underlying transmission mechanism indicates that if economic agents face a downturn in real activity, companies lower prices in order to increase sales, effectively mitigating the drop in demand. Consistent with the reactions of economic output, we observe comparatively stronger price reactions in Austria, France and Finland. The rather heterogeneous response pattern corroborates and extends findings in Peersman (2004) and Barigozzi et al. (2014), who report asymmetric responses of macroeconomic quantities to common monetary policy shocks in the euro area.

4 Forecasting evidence

In this section, we aim to assess whether using our proposed model specification (labeled here DFM-SV) pays off in terms of predictive capabilities by using an out-of-sample forecast exercise in other to shed light on the predictive importance of specific components of the specification. Forecasts of DFM-SV are computed for the period ranging from 1989Q1 to 2013Q3. The competing benchmark specifications are as follows:

- A variant without a euro area gap component, labeled as UC-SV. This model serves to assess the merits of including a common cyclical component.

(a) Output (b) Inflation

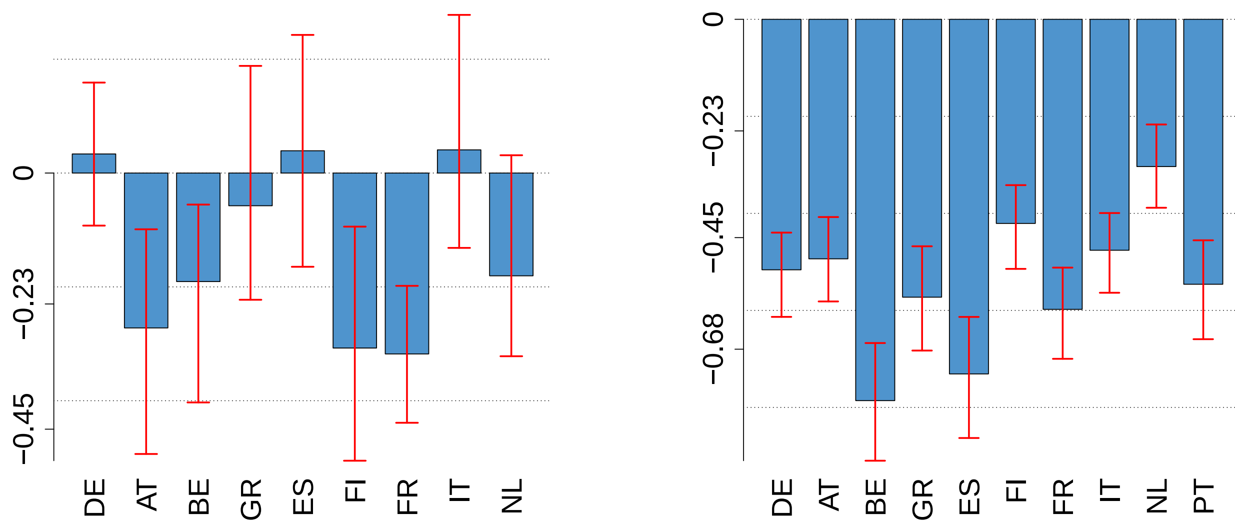

Fig. 6: Negative of the maximum absolute impact of a negative one standard error business cycle shock on (a) Output and (b) Inflation

- A specification labeled as UC-SV-Cycle augments the former by a (potentially over-parameterized) trend-cyle decomposition. This implies that for each time series, we estimate a model with a stochastic trend and a cyclical AR(2) component.

- Recent work by Hamilton (2017) advocates a simple alternative as a means to decompose cyclical and trend components of econometric time series. The approach relies on simple forecasts using a constant and the four most recent observations of the quantity under consideration. This benchmarked, labeled as Hamilton, replaces the common component by using an estimate of the gap based on the approach discussed in Hamilton (2017).

- Standard benchmark specifications include simple random walk specifications (RW) or first order autoregressive processes (AR(1)), both estimated with stochastic volatility in the errors.

	DFM-SV	UC-SV	UC-SV-Cycle	Hamilton	AR(1)	RW
			Output			
AT	323.0	323.7	320.4	317.3	330.1	298.2
BE	338.6	325.7	331.5	324.0	247.8	208.8
DE	334.5	329.5	323.0	323.8	330.3	311.0
GR	272.1	276.2	271.3	272.6	247.5	176.3
ES	320.2	322.5	308.1	316.2	294.0	327.0
FI	285.0	291.2	280.0	287.8	291.1	269.2
FR	358.6	359.6	347.9	351.2	314.7	339.5
IT	355.4	346.4	335.0	343.5	309.7	309.4
NL	319.3	333.5	325.1	327.7	333.8	307.8
PT	322.3	304.1	312.0	311.0	309.4	175.8
			Inflation			
AT	403.1	434.2	447.4	327.4	391.0	420.5
BE	412.9	422.9	421.4	376.1	400.0	395.4
DE	383.1	434.4	440.9	330.5	404.5	330.1
GR	385.2	383.6	384.2	335.2	349.2	335.2
ES	400.3	391.2	409.7	366.5	395.0	374.0
FI	440.5	438.0	441.7	396.4	434.2	428.3
FR	453.7	448.7	451.0	316.2	405.3	356.5
IT	443.3	443.7	450.1	285.7	440.5	429.7
NL	430.2	429.2	438.1	423.1	379.8	362.0
PT	397.7	393.7	388.3	247.9	263.0	267.5
			Joint performance			
Output	3,228.9	3,212.4	3,154.2	3,175.2	3,008.5	2,722.9
Inflation	4,150.0	4,219.4	4,272.9	3,405.0	3,862.4	3,699.1

Table 1: Marginal log-predictive scores

The country-specific out-of-sample forecast performance for the competing model specifications under scrutiny, measured in terms of marginal log-predictive scores, are reported in Table 1. DFM-SV appears to outperform the alternative specifications in countries such as Belgium, Germany, Italy, or Portugal. On the contrary, the specification without an euro area output gap (UC-SV) appears to slightly outperform for Greece, Spain, Finland, and France. However, compared to the alternative specifications, the table reveals non-negligible advantages in the overall out-of-sample predictive performance of the proposed model (DFM-SV) in terms of predicting economic output. Specifically, DFM-SV displays the strongest forecasting performance, closely tracked by UC-SV and the model based on the Hamilton approach ranked third. Standard benchmarks such as the simple AR(1) model or the random walk appear to work well in selected countries but are generally beaten by models that take into account cross-country information.

With some few exceptions, the specification UC-SV-Cycle hardly manages to produce more precise forecasts for output as compared to the more parsimonious UC-SV and DFM-SV models. The specification using the trend-cycle decomposition advocated by Hamilton (2017) produces marked increases in the predictive out-of-sample performance in economic output as compared to UC-SV-Cycle for most countries. However, similarly to UC-SV-Cycle, this approach also fails to outperform our proposed DFM-SV model.

Summing up, for output we find that our multivariate state space model clearly outperforms all alternative specifications considered. Standard benchmarks such as first-order autoregressive processes or random walks also appear to severely underperform as compared to DFM-SV. Notable exceptions are output predictions in Austria or the Netherlands, where the first-order autoregressive specifications appear to produce the most precise density predictions.

For inflation, we observe a slightly weaker performance of our proposed model with the UC-SV-Cycle specification performing best for almost all countries under scrutiny. Exceptions appear to be inflation forecasts for France and Portugal, where the DFM-SV specification yields the best predictive performance. For Belgium, we observe that UC-SV performs best. Considering joint predictive performance for inflation corroborates the findings based on marginal log scores, showing that UC-SV-Cycle outperforms all competing alternatives. This finding is consistent with Stock and Watson (1999; 2007) who show that using similar models like the UC-SV and the UC-SV-Cycle yields precise inflation predictions.

(a) Output

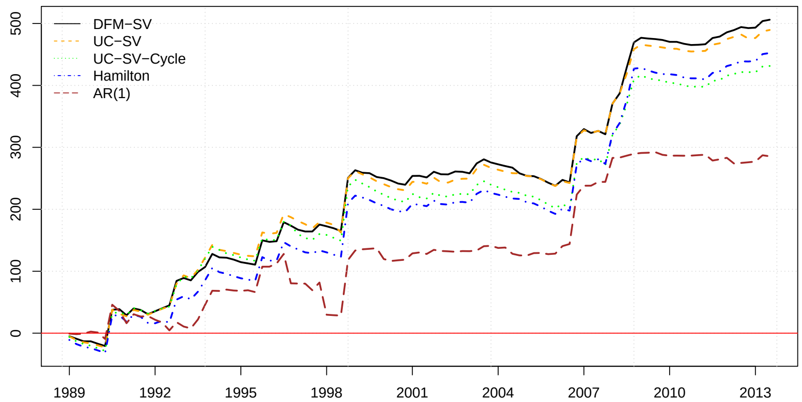

(b) Inflation

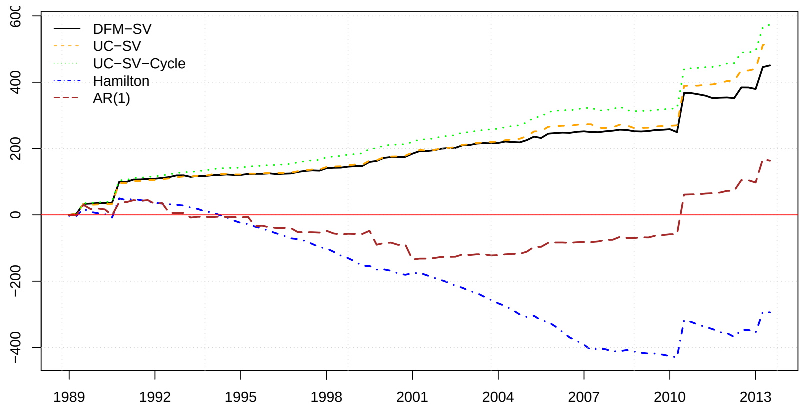

Fig. 7: Marginal log-predictive scores over time

Figure 7 depicts the predictive performance measured in terms of marginal log-predictive scores relative to the random walk (red line) over time. The top panel of the figure presents the joint performance for output, whereas the bottom panel shows the overall predictive performance for inflation. In line with the summary metrics presented in Table 1, Figure 7 shows a pronounced outperformance of DFM-SV in terms of forecasting economic output especially in the recent decade. Interestingly, albeit the out-of-sample performances between DFM-SV and UC-SV appeared very similar in the past, a marked decoupling of the two series is particularly pronounced in the aftermath of the financial crisis. In the recent decade DFM-SV appeared to outperform the other specifications in terms of forecasting real activity, indicating the importance of accounting for a common gap component in the aftermath of the economic turmoil. For the joint predictive performance of inflation, the opposite is the case. Especially in the last decade of the sample, UC-SV appeared to outperform the specification including the common factor (DFM-SV). However, in terms of predictive performance, both specification are outperformed by UC-CV-Cycle for almost the entire forecast period.

5 Concluding remarks

In this paper we estimate a Bayesian multivariate unoberved components for output and inflation in the spirit of Stock and Watson (2002) for the euro area countries. The multi-country framework allows to explicitly account for country-specific trajectories in the econometric series, augmenting and extending the framework proposed in Stella and Stock (2013) along several dimensions. To account for both common and country-specific factors, the proposed model specification explicitly accounts for a latent common gap component for economic output in the euro area. In order to alleviate the potential problem of over-parameterization of the model, the proposed estimation strategy moreover involves recent regularizations in terms of Bayesian shrinkage.

The estimated trends and gap components appear to match the timing of the economic peaks and troughs very well. In a forecasting exercise, the paper moreover compares the forecast performance of the proposed model specification with other well-known benchmark specifications. The out-of-sample forecast exercise shows that accounting for a common euro area output gap component produces particularly precise forecasts for the economic output series under scrutiny.

Acknowledgements

The research carried out in this paper was supported by funds of the Oesterreichische Nationalbank (Oesterreichische Nationalbank, Anniversary Fund, project number: 17382).

References

Barigozzi M, Conti AM and Luciani M (2014) Do euro area countries respond asymmetrically to the common monetary policy? *Oxford bulletin of economics and statistics* 76(5), 693–714

Basistha A and Nelson CR (2007) New measures of the output gap based on the forward-looking new Keynesian Phillips curve. *Journal of Monetary Economics* 54(2), 498–511

Belmonte MA, Koop G and Korobilis D (2014) Hierarchical shrinkage in time-varying parameter models. *Journal of Forecasting* 33(1), 80–94

Bitto A and Frühwirth-Schnatter S (2016) Achieving shrinkage in a time-varying parameter model framework. *arXiv preprint arXiv:1611.01310*

Breitung J and Eickmeier S (2015) Analyzing business cycle asymmetries in a multi-level factor model. *Economics Letters* 127, 31–34

Carter CK and Kohn R (1994) On Gibbs sampling for state space models. *Biometrika* 81(3), 541–553

Feldkircher M, Huber F and Kastner G (2017) Sophisticated and small versus simple and sizeable: When does it pay off to introduce drifting coefficients in Bayesian VARs? *arXiv preprint arXiv:1711.00564*

Frühwirth-Schnatter S (1994) Applied state space modelling of non-Gaussian time series using integration-based Kalman filtering. *Statistics and Computing* 4(4), 259–269

Frühwirth-Schnatter S and Wagner H (2010) Stochastic model specification search for Gaussian and partial non-Gaussian state space models. *Journal of Econometrics* 154(1), 85–100

Griffin JE and Brown PJ (2010) Inference with normal-gamma prior distributions in regression problems. *Bayesian Analysis* 5(1), 171–188

Hamilton JD (2017) Why you should never use the Hodrick-Prescott filter

Hodrick RJ and Prescott EC (1997) Postwar US business cycles: An empirical investigation. *Journal of Money, Credit, and Banking* , 1–16

Jarocinski M and Lenza M (2015) Output gap and inflation forecasts in a Bayesian dynamic factor model of the euro area. *manuscript, European Central Bank*

Kastner G and Frühwirth-Schnatter S (2014) Ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC estimation of stochastic volatility models. *Computational Statistics & Data Analysis* 76, 408–423

Kim CJ and Nelson CR (1999) *State-space models with regime switching: classical and Gibbs-sampling approaches with applications*, volume 2. MIT press Cambridge

Kim S, Shephard N and Chib S (1998) Stochastic volatility: likelihood inference and comparison with ARCH models. *The review of economic studies* 65(3), 361–393

Kose MA, Otrok C and Whiteman CH (2003) International business cycles: World, region, and country-specific factors. *American Economic Review* 93(4), 1216–1239

Kuttner KN (1994) Estimating potential output as a latent variable. *Journal of business & economic statistics* 12(3), 361–368

Orphanides A and Van Norden S (2002) The unreliability of output-gap estimates in real time. *The Review of Economics and Statistics* 84(4), 569–583

Otrok C and Whiteman CH (1998) Bayesian leading indicators: measuring and predicting economic conditions in Iowa. *International Economic Review* 39(4), 997–1014

Peersman G (2004) The transmission of monetary policy in the Euro Area: are the effects different across countries? *Oxford Bulletin of Economics and Statistics* 66(3), 285–308

Planas C, Rossi A and Fiorentini G (2008) Bayesian analysis of the output gap. *Journal of*

Business & Economic Statistics 26(1), 18–32

Stella A and Stock J (2013) A state-dependent model for inflation forecasting. *FRB International Finance Discussion Paper No. 1062*

Stock JH and Watson MW (1999) Forecasting inflation. *Journal of Monetary Economics* 44(2), 293–335

Stock JH and Watson MW (2002) Macroeconomic forecasting using diffusion indexes. *Journal of Business & Economic Statistics* 20(2), 147–62

Stock JH and Watson MW (2007) Why has US inflation become harder to forecast? *Journal of Money, Credit and banking* 39(s1), 3–33

A Full conditional posterior distributions

It is worth noting that the joint posterior distribution of the model parameters and the set of latent states is intractable. Fortunately, the full conditional posterior distributions for most quantities are of a simple form and thus amenable to standard Gibbs updating.

In order to obtain a draw from the joint posterior we design a straightforward Markov chain Monte Carlo (MCMC) algorithm that cycles through the following steps:

(1) Simulate the full history of $\{\boldsymbol{\tau}_t\}_{t=1}^T$ using a forward filtering backward sampling algorithm (Carter and Kohn 1994, Frühwirth-Schnatter 1994).

(2) Draw the sequence of log-volatilities $\{h_{jt}^k\}_{t=1}^T, \{s_{it}\}_{t=1}^T$ for all i, j, k as well as the parameters in the corresponding state equations independently using the algorithm proposed in Kastner and Frühwirth-Schnatter (2014).

(3) Conditional on the latent states, we simulate the loadings α_j^k by estimating $N+1$ independent regression models with heteroskedastic innovations

(4) The parameters Q and γ are updated in a block by using a standard random walk Metropolis Hastings algorithm.

(5) Update B_{sj} and its counterpart B_{hi}^k by sampling from an generalized inverted Gaussian (GIG) distribution.

(6) Sample ξ_s and ξ_h from a Gamma distributed conditional posterior distribution.

Steps (1) to (4) are standard and easily executed. Steps (5) and (6) deserve more attention. In the empirical application we repeat this algorithm 30,000 times and discard the first 15,000 as burn-in.

The full conditional of B_{sj} follows a GIG distribution that is obtained by combining the conditional density $p(\sqrt{\vartheta_{sj}^{\tau}}|B_{sj})$ with the conditional prior $p(B_{sj}|\xi_s)$,[2]

$$B_{sj}|\bullet \sim \mathcal{GIG}(\kappa_s - 1/2, \vartheta_{sj}^{\tau}, \xi_s\kappa_s), \text{ for } j = 1, \ldots, K, \tag{A.1}$$

where $\bullet$ denotes conditioning on all remaining quantities of the model.

Likewise, the full conditional posterior of B_{hj}^{k} is given by

$$B_{hi}^{k}|\bullet \sim \mathcal{GIG}(\kappa_h - 1/2, \vartheta_{hi}^{k}, \xi_s\kappa_s), \text{ for } i = 1, \ldots, N; \quad k \in \{y, \Delta p\}. \tag{A.2}$$

To obtain the full conditional posterior distribution for the global scaling parameters, we combine the joint density $\prod_{j=1}^{K} p(B_{sj}|\xi_s)$ with the prior $p(\xi_s)$. This yields a Gamma distributed conditional posterior distribution,

$$\xi_s|\bullet \sim \mathcal{G}\left(d_0 + \kappa_s K, d_1 + \kappa_s/2 \sum_{j=1}^{K} B_{sj}\right). \tag{A.3}$$

Similarly to the conditional posterior of ξ_s, ξ_h also follows a Gamma distribution

$$\xi_h|\bullet \sim \mathcal{G}\left(e_0 + (2N+1)\kappa_h, d_1 + \kappa_h/2\left[\sum_{j=1}^{N} B_{hj}^{y} + \sum_{j=0}^{N} B_{hj}^{\Delta p}\right]\right). \tag{A.4}$$

[2]The GIG distribution has a density which is proportional to $p(x) \propto x^{\nu-1} exp(-\{\chi/x + \psi x\}/2)$.