Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections

Marta Bańbura, European Central Bank

Domenico Giannone, LUISS, ULB-ECARES and CEPR

Michele Lenza, European Central Bank and ULB-ECARES

August 28, 2014

Abstract

This paper describes an algorithm to compute the distribution of conditional forecasts, i.e. projections of a set of variables of interest on future paths of some other variables, in dynamic systems. The algorithm is based on Kalman filtering methods and is computationally viable for large models that can be cast in a linear state space representation. We build large vector autoregressions (VARs) and a large dynamic factor model (DFM) for a quarterly data set of 26 euro area macroeconomic and financial indicators. Both approaches deliver similar forecasts and scenario assessments. In addition, conditional forecasts shed light on the stability of the dynamic relationships in the euro area during the recent episodes of financial turmoil and indicate that only a small number of sources drive the bulk of the fluctuations in the euro area economy.

JEL Classification: C11, C13, C33, C53

Keywords: Vector Autoregression, Bayesian Shrinkage, Dynamic Factor Model, Conditional Forecast, Large Cross-Sections.

1 Introduction

Vector autoregressions (VARs) are very flexible and general models and provide a reliable empirical benchmark for alternative econometric representations such as dynamic stochastic general equilibrium (DSGE) models, which are more grounded in theory but, at the same time, impose more structure on the data (see, for example, Christiano, Eichenbaum, and Evans, 1999).

Recent literature has shown that VARs are viable tools also for large sets of data (see Bańbura, Giannone, and Reichlin, 2010). In this paper, we construct a large VAR for the euro area and we apply it to unconditional forecasting as well as for conditional forecasts and scenarios. These, along with structural analysis (assessing the effects of structural shocks), have been the main applications of VARs. Whereas large VARs have been used for unconditional forecasting and structural analysis,[1] limited attention has been devoted as yet to conditional forecasting. This is because popular algorithms for deriving conditional forecasts have been computationally challenging for large data sets. We overcome this problem by computing the conditional forecasts recursively using Kalman filtering techniques.

Conditional forecasts and, in particular, scenarios are projections of a set of variables of interest on future paths of some other variables. This is in contrast to unconditional forecasts, where no knowledge of the future path of any variables is assumed. The prior knowledge, albeit imperfect, of the future evolution of some economic variables may carry information for the outlook of other variables. For example, future fiscal packages would affect the future evolution of economic activity and, thus, might provide important off-model information. Moreover, it may be of interest to assess the impact of specific future events on a set of variables, i.e. to conduct scenario analysis. Notable examples of the latter are the stress tests recently conducted in the US and the euro area in order to assess the vulnerability of their banking systems. For recent examples of conditional forecasts, see Lenza, Pill, and Reichlin (2010); Giannone, Lenza, and Reichlin (2010); Jarociński and Smets (2008); Bloor and Matheson (2011); Giannone, Lenza, Pill, and Reichlin (2012); Stock and Watson (2012a); Giannone, Lenza, Momferatou, and Onorante (2014). Recently, Clark and McCracken (2014) propose

---

[1]See e.g. Koop (2013), Giannone, Lenza, Momferatou, and Onorante (2014), Giannone, Lenza, and Reichlin (2012), Paciello (2011),Giannone, Lenza, and Primiceri (2014).

and evaluate a range of tests of predictive ability for conditional forecasts from estimated models.

The scenario analysis described above and studied in this paper can be considered as "reduced form", in the sense that the forecasts are conditional on observables and the identification of structural shocks is not required.[2] Notice that, if needed, the structural shocks that are "compatible" with the scenario can be retrieved from the reduced form innovations with some identifying assumptions. An alternative approach consists in constructing scenarios by manipulating specific structural shocks so that the resulting paths of the observed variables are consistent with the conditioning information (see also Leeper and Zha, 2003; Adolfson, Laséen, Lindé, and Villani, 2005; Christoffel, Coenen, and Warne, 2007; Luciani, 2013). Along similar lines, Baumeister and Kilian (2013) construct scenarios for real price of oil from a VAR by conditioning on a sequence of appropriately derived structural shocks rather than on a pre-specified path for observables.

For VAR models, the conditional forecasts are typically computed by using the algorithm developed by Waggoner and Zha (1999). Roughly speaking, the methodology involves drawing (the entire) paths of reduced form shocks which are compatible with the conditioning path on the observables. Due to computational burden, this approach can easily become impractical or unfeasible for high dimensional data and long forecast horizons, even if the computationally more efficient version of Jarociński (2010) is employed. However, many problems in macroeconomics and finance can only be addressed by looking at the joint dynamic behavior of a large number of time series. For example, business cycle research, as in the NBER tradition, typically involves the analysis of many macroeconomic variables. Professional forecasters and policymakers look at a variety of different indicators to predict key variables of interest and to make their decisions. Investors analyze the joint behavior of many asset returns in order to choose their optimal portfolios. More in general, contemporary science relies more and more on the availability and exploitation of large data sets.

In this paper, building on an old insight by Clarida and Coyle (1984), we propose an algorithm based on Kalman filtering techniques to compute the conditional forecasts. Since the Kalman filter works recursively, i.e. period by period, this algorithm reduces significantly the

---

[2] For a discussion on the invariance of conditional forecast distribution to alternative identification assumptions for structural shocks see Waggoner and Zha (1999).

3

computational burden for longer forecast horizons and is particularly well suited for empirical approaches handling large data sets. Using a simulation smoother (see Carter and Kohn, 1994; de Jong and Shephard, 1995; Durbin and Koopman, 2002, for examples of simulation smoothers) allows for the computation of the full distribution of conditional forecasts. The algorithm applies to any model which can be cast in a linear state space representation. For the VAR framework, we compare the computational efficiency of different simulation smoothers and find that for large systems the simulation smoother of Durbin and Koopman (2002) can offer substantial computational gains with respect to the more popular algorithm of Carter and Kohn (1994).

The interest in issues which are best addressed by considering large information sets raises a trade-off between excessive simplicity of the models – misspecification due to omitted variables – and their excessive complexity – many free parameters leading to large estimation uncertainty. Recent developments in macroeconometrics have suggested two approaches to deal with the complexity of large sets of data, without losing their salient features: Bayesian VARs (BVARs) and dynamic factor models (DFMs).

The aforementioned flexibility of VARs comes at the cost of a high number of free parameters to be estimated. Specifically, for a generic VAR(p) model for a vector of n variables $y_t = (y_{1,t}, \ldots, y_{n,t})'$:

$$y_t = c + A_1 y_{t-1} + \cdots + A_p y_{t-p} + \varepsilon_t, \qquad \varepsilon_t \sim WN(0, \Sigma), \tag{1}$$

where $WN(0, \Sigma)$ refers to a white noise process with mean 0 and covariance matrix Σ, we count: i) pn^2 parameters in autoregressive matrices, $A_1, \ldots, A_p$, that are of dimension $n \times n$ each; ii) $n(n + 1)/2$ free parameters in the $n \times n$ covariance matrix of residuals Σ; iii) n parameters in the constant term c. The number of parameters proliferates as the number of variables in the model increases, making estimation unreliable or unfeasible. For example, when the number of variables in a VAR with 4 lags increases from 6, as in the original VAR model proposed by Sims (1980), to 20, 50 or 100, the total number of parameters to be estimated goes from 171 to, respectively, numbers in the order of 2, 10 and 50 thousands. Such a high number of parameters cannot be well estimated by ordinary least squares, for example, since the typical macroeconomic sample involves a limited number of data points (in the best case, $50 - 60$ years of data, i.e. $200 - 250$ data points with quarterly data). The

4

problem of parameter proliferation that prevents econometricians from conducting reliable inference with large dimensional systems is also known as the "curse of dimensionality".

A solution to the curse of dimensionality in the VAR framework consists in adopting Bayesian shrinkage. The idea of this method is to combine the likelihood coming from the complex and highly parameterised VAR model with a prior distribution for the parameters that is naïve but enforces parsimony. As a consequence, the estimates of the coefficients are "shrunk" toward the prior expectations, which are typically equal to 0.[3] This approach can also be interpreted as a penalised maximum likelihood method.

The shrinkage methods have been advocated by early proponents of VARs as a macroeconometric tool (Litterman, 1979; Sims, 1980; Doan, Litterman, and Sims, 1984) but they were typically used for low dimensional systems. Recently, it has been shown that the idea of shrinkage works also for high dimensional systems and provides results that are very similar to those obtained by using the DFMs (see De Mol, Giannone, and Reichlin, 2008; Bańbura, Giannone, and Reichlin, 2010; Giannone, Lenza, and Primiceri, 2014). This is not surprising since, as shown by De Mol, Giannone, and Reichlin (2008), when applied to collinear variables, as are typically macroeconomic variables, the forecasts produced by factor models and Bayesian shrinkage tend to get closer, as the size of the sample and of the cross-section get larger.

Factor models exploit the fact that macroeconomic and financial time series are characterised by strong cross-sectional correlation. Under the assumption that most of the fluctuations are driven by a relatively limited set of common sources, factor models offer a parsimonious representation by summarizing the information from a large number of data series in few common factors. DFMs further parameterize the dynamics of the factors, typically assuming a VAR process. The estimation of factor models generally requires that the data are stationary. Assuming that stationarity is achieved via taking first differences[4], the DFM is defined as follows:

$$\Delta y_t \;=\; \mu + \Lambda F_t + e_t,$$
$$F_t \;=\; \Phi_1 F_{t-1} + \ldots + \Phi_p F_{t-s} + u_t, \qquad u_t \sim WN(0, Q),$$

[3]For an extensive discussion of shrinkage in various contexts see e.g. Stock and Watson (2012b) and Ng (2013).

[4]Trending series are typically "logged" beforehand.

where $F_t = (F_{1,t}, \ldots, F_{r,t})'$ is an r-dimensional vector of common factors, with r typically much smaller than n and Λ is an $n \times r$ matrix of factor loadings. Since the number of common factors r is typically small, the estimation of the VAR describing the dynamics of the common factors does not pose any problem. The residual $e_t = (e_{1,t}, \ldots, e_{n,t})'$ is the idiosyncratic component. The most common approach is to assume that the idiosyncratic component is cross-sectionally uncorrelated. This assumption gives rise to the "exact" factor model, which highlights the fact that the cross-correlation between the variables is fully accounted for by the common factors. Interestingly, recent literature has shown that factor models can be estimated with large data sets, i.e. even in situations in which the cross-sectional dimension n is much larger than the sample size T. In addition, the estimates are asymptotically valid also when the data generating process is not the "exact" but rather an "approximate" factor model, in the sense that the idiosyncratic components are weakly cross-correlated (see Forni, Hallin, Lippi, and Reichlin, 2000; Stock and Watson, 2002b; Bai and Ng, 2002; Bai, 2003; Forni, Hallin, Lippi, and Reichlin, 2004; Doz, Giannone, and Reichlin, 2012). Stock and Watson (2011) provide an exhaustive survey of the literature.

Factor models are appealing also because many popular economic models can be cast in their format. The typical theoretical macro model, indeed, includes only a handful of shocks driving the key aggregate variables in the economy. The arbitrage pricing theory (APT) is built upon the existence of a set of common factors underlying all returns. Moreover, the distinction between common and idiosyncratic sources of fluctuations is often employed in international, regional and sectorial studies and represents a useful device to study macroeconomic implications of microeconomic behavior (see e.g. Kose, Otrok, and Whiteman, 2003; Foerster, Sarte, and Watson, 2011).

In our empirical application, we carry out a comprehensive comparison of the two modelling approaches, the VARs and the DFMs, on a large data set of quarterly euro area macroeconomic and financial variables. We consider two versions of the BVAR – with data in (log-)levels and in (log-)differences – and a dynamic factor model. First, we show that the three models produce quite accurate unconditional forecasts, compared to univariate benchmarks, and that the forecasts from the three approaches are very correlated. The latter finding lends empirical support to the theoretical argument that the approaches are tightly linked, complementing similar evidence already available for the US (see, for example De Mol, Giannone, and Reichlin,

2008; Giannone, Lenza, and Primiceri, 2014).

Then, we compare the two approaches also for what concerns scenarios and conditional forecasts. In particular, we study the economic developments associated to a scenario of an increase in world GDP as well as conditional forecasts based on the realised path of real GDP, consumer prices and the policy rate. We show that also the scenario analysis and the conditional forecasts computed for the three models provide similar insights. The fact that the results are not model specific is reassuring, since it indicates that the predictions of the models are reflecting genuine data features.

The results from the conditional forecasts yield support to two further conclusions. First, the fact that the conditional forecasts based on the three variables track, in general, quite closely the actual developments in most of the variables under analysis suggests that there are only a few "sources" that drive the bulk of the fluctuations in the euro area economy. Second, there appears to be some degree of stability in the economic relationships following the financial crisis as the the conditional forecasts for this period based on the parameters estimated with data until end of 2007 are relatively accurate, with the possible exception of some categories of loans and broad monetary aggregates (see Giannone, Lenza, and Reichlin, 2012, for an extensive discussion and interpretation of these results).

The structure of the paper is as follows. In section 2, we review the state-of-the-art techniques for the estimation and inference for DFMs and BVARs and we expound the close relationship linking the two approaches. In section 3, we describe a Kalman filter based methodology to compute conditional forecasts. In section 4, we present and discuss the empirical results. Section 5 concludes. The appendix contains some implementation details, comparison of computational performance of different algorithms and data descriptions.

2 Models for large data sets

2.1 Dynamic factor models

The general representation of the dynamic factor model described in the introduction is:

$$\Delta y_t = \mu + \Lambda F_t + e_t, \tag{2}$$
$$F_t = \Phi_1 F_{t-1} + \ldots + \Phi_s F_{t-s} + u_t, \qquad u_t \sim WN(0, Q).$$

Following Doz, Giannone, and Reichlin (2012) the model can be estimated by means of quasi-maximum likelihood methods. In this context, the estimation of the model is performed by maximising a likelihood function, under the assumption that data are Gaussian and that the factor structure is exact, i.e. the idiosyncratic errors are cross-sectionally orthogonal: $u_t \sim$ i.i.d.$N(0, Q)$ and $e_t \sim$ i.i.d.$N(0, \Gamma_d)$, where Γ_d is a diagonal matrix.

Doz, Giannone, and Reichlin (2012) have shown that this estimation procedure provides consistent estimates for approximate dynamic factor models under general regularity conditions (convergence in probability of the covariance matrix of the data and data stationarity). Remarkably, consistency is achieved without any constraint on the number of variables, n, relative to the sample size, T, under the assumption of weak cross-sectional dependence of the idiosyncratic term, e_t, and of sufficient pervasiveness of the common factors.

As the factors are unobserved, the maximum likelihood estimators of the parameters Λ, Γ_d, $\Phi_1, \ldots, \Phi_s, Q$, which we collect in θ, are, in general, not available in closed form. They can be obtained either via a direct numerical maximisation of the likelihood, which can be computationally demanding,[5] or, as in Doz, Giannone, and Reichlin (2012), via the Expectation-Maximisation (EM) algorithm. The EM algorithm was proposed by Dempster, Laird, and Rubin (1977) as a general solution to problems with incomplete or latent data. In the case of the DFM, the algorithm alternates between the use of the Kalman smoother to estimate the common factors given a set of parameters (E-step), and multivariate regressions (corrected for the uncertainty in the estimation of the common factors) to estimate the parameters given the factors (M-step), see e.g. Watson and Engle (1983) or Shumway and Stoffer (1982).

---

[5]Jungbacker and Koopman (2008) show how to reduce the computational burden in case the number of observables is much larger than the number of factors.

The algorithm can be initialised using the sample principal components. In what follows, we assume that data are standardised to have sample mean equal to zero and variance equal to one.[6] Denote by d_j, $j = 1, \ldots n$, the eigenvalues of $\frac{1}{T} \sum_{t=1}^{T} \Delta y_t \Delta y_t'$ and by v_j, $j = 1, \ldots n$, the associated eigenvectors, i.e.

$$\left(\frac{1}{T} \sum_{t=1}^{T} \Delta y_t \Delta y_t' \right) v_j = v_j d_j, \quad j = 1, 2, \ldots n \,,$$

with $v_j' v_j = 1$, $v_j' v_k = 0$ for $j \neq k$ and $d_1 \geq d_2 \geq \ldots \geq d_n$. The sample principal components of Δy_t are defined as $z_{jt} = \frac{1}{\sqrt{d_j}} v_j' \Delta y_t$. The principal components are ordered accordingly to their ability to explain the variability in the data as the total variance explained by each principal component is equal to d_j. The principal components transform cross-sectionally correlated data, Δy_t, into linear combinations $z_t = (z_{1,t}, \ldots, z_{n,t})' = H \Delta y_t$ where $H = \left(\frac{1}{\sqrt{d_1}} v_1, \ldots, \frac{1}{\sqrt{d_n}} v_n \right)'$. These linear combinations are cross-sectionally uncorrelated, with unit variance, $\frac{1}{T} \sum_{t=1}^{T} z_t z_t' = I_n$.

The approximate factor structure is defined in terms of behavior of the eigenvalues of the population covariance matrix when the number of variables increases. Specifically, the first r eigenvalues of the population covariance matrix of Δy_t are assumed to grow with the dimension of the system, at a rate n. All the remaining eigenvalues remain, instead, bounded. It can be proved that these assumptions imply that the eigenvalues d_j of the sample covariance matrix will go to infinity at a rate n for $j = 1, \ldots, r$, where r is the number of common factors. On the other hand, $d_{r+1}, \ldots, d_n$ will grow at a rate given by $n/\sqrt{T}$ (see De Mol, Giannone, and Reichlin, 2008; Doz, Giannone, and Reichlin, 2011, 2012). Forni, Hallin, Lippi, and Reichlin (2000) and Stock and Watson (2002a,b) have shown that if data have an approximate factor structure[7], then the first r principal components can approximate well the space spanned by the unobserved common factors, when the sample size and the cross-sectional dimension are large.

---

[6]The zero mean assumption is without loss of generality, since it is equivalent to concentrating out the mean. Since maximum likelihood estimates are scale invariant, rescaling the data does not affect the estimates. On the other hand, homogeneity of the scale across variables is convenient, since the algorithm for maximizing the likelihood is more efficient from a computational standpoint. In addition, working with standardised data is useful since the initialisation of the algorithm is based on principal components, which are not scale invariant. Once the estimates are obtained, the factor loadings, $\hat{\Lambda}$, and the covariance matrix of the idiosyncratic components, $\hat{\Gamma}_d$, can be obtained by simple rescaling.

[7]As stressed above, this amounts to assuming that the idiosyncratic components are weakly cross-correlated.

The sample principal components offer thus good starting values for the common factors: $\hat{F}_t^{(0)} = z_t$.[8] The starting values for the parameters of the model, $\theta^{(0)}$, can then be estimated by means of OLS techniques, by treating the principal components as if they were the true factors. Once the parameters have been estimated, we can estimate a new set of factors by using the Kalman smoother: $\hat{F}_t^{(1)} = \mathrm{E}_{\theta^{(0)}}[F_t|\Delta y_1, \ldots, \Delta y_T]$. At this stage, we have the two-step procedure of Doz, Giannone, and Reichlin (2011). The quasi-maximum likelihood estimation via the EM algorithm consists essentially in iterating these steps until convergence. Details are reported in the appendix.

2.2 Bayesian vector autoregressions

For Gaussian data, the VAR model described in the introduction is:

$$y_t = c + A_1 y_{t-1} + \cdots + A_p y_{t-p} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \Sigma).$$

We consider conjugate priors belonging to the normal-inverse-Wishart family, where the prior for the covariance matrix of the residuals is inverse-Wishart and the prior for the autoregressive coefficients is normal. The priors are a version of the so-called Minnesota prior, originally due to Litterman (1979), which is centered on the assumption that each variable follows an independent random walk process, possibly with drift:

$$y_t = c + y_{t-1} + \varepsilon_t,$$

which is a parsimonious yet "reasonable approximation of the behavior of an economic variable".

For the prior on the covariance matrix of the errors, Σ, we set the degrees of freedom equal to $n + 2$, which is the minimum value that guarantees the existence of the prior mean, which we set as $\mathrm{E}[\Sigma] = \Psi$, where Ψ is diagonal.

---

[8]In fact, under the assumption that $\Phi_1 = \ldots = \Phi_s = 0$ and $\Gamma_d = \bar{\gamma} I_n$ (i.e. homoscedastic idiosyncratic components) the quasi-maximum likelihood solution is analytical, with the expected value for the factors proportional to the principal components of the data.

The prior moments for the VAR coefficients are as follows:

$$E\left[(A_s)_{ij}|\Sigma,\lambda,\Psi\right] = \begin{cases} 1 & \text{if } i=j \text{ and } s=1 \\ 0 & \text{otherwise} \end{cases}$$

$$cov\left[(A_s)_{ij},(A_r)_{hm}|\Sigma,\lambda,\Psi\right] = \begin{cases} \lambda^2\frac{1}{s^2}\frac{\Sigma_{ih}}{\Psi_{jj}} & \text{if } m=j \text{ and } r=s \\ 0 & \text{otherwise} \end{cases} .$$

Notice that the variance of this prior is lower for the coefficients associated with more distant lags, and that coefficients associated with the same variable and lag in different equations can be correlated. Finally, the key hyperparameter λ controls the scale of all the variances and covariances, and effectively determines the overall tightness of this prior. The terms Σ_{ij}/Ψ_{jj} account for the relative scale of the variables. The prior for the intercept, c, is diffuse.[9]

We include an additional prior, which implements a so-called "inexact differencing" of the data. More precisely, rewrite the VAR equation in an error correction form:

$$\Delta y_t = c + \Pi y_{t-1} + B_1\Delta y_{t-1} + \ldots + B_{\tilde{p}}\Delta y_{t-\tilde{p}} + \varepsilon_t.$$

where $\tilde{p} = p-1$, $B_s = -A_{s+1} - \ldots - A_p$, $s = 1,\ldots,\tilde{p}$ and $\Pi = A_1 + \ldots + A_p - I_n$.

A VAR in first differences implies the restriction $\Pi = 0$ (or $A_1 + \ldots + A_p = I_n$). We follow Doan, Litterman, and Sims (1984) and set a prior that shrinks Π to zero. Precisely, we set a prior centered at 1 for the sum of coefficients on own lags for each variable, and at 0 for the sum of coefficients on other variables' lags. This prior introduces correlation among the coefficients on each variable in each equation. The tightness of this prior on the "sum of coefficients" is controlled by the hyperparameter μ. As μ goes to infinity the prior becomes diffuse while, as it goes to 0, we approach the case of exact differencing, which implies the presence of a unit root in each equation.

Following Sims (1993) and Sims and Zha (1998), we complement such "inexact differencing" with an additional prior, known as "dummy-initial-observation" prior, that shrinks the forecast

---

[9]Koop (2013) considers non-conjugate priors which allow for exclusion of certain variables from some equations, however, he finds that these do not outperform simpler Minnesota priors in terms of forecast accuracy. Carriero, Clark, and Marcellino (2012) find that allowing for stochastic volatility helps to improve forecast accuracy. See Karlsson (2013) for a comprehensive overview of Bayesian methods for inference and forecasting with VAR models.

of each variable at the beginning of the sample toward a no-change forecast. The tightness of the prior is controlled by an additional hyperparameter δ.

The setting of the priors importantly depends on the hyperparameters λ, μ, δ and Ψ, which reflect the informativeness of the prior distributions for the model coefficients. These hyperparameters have been usually set on the basis of subjective considerations or rules-of-thumb. Instead, we closely follow the theoretically grounded approach proposed by Giannone, Lenza, and Primiceri (2014). This involves treating the hyperparameters as additional parameters, in the spirit of hierarchical modelling. As hyperpriors (i.e. prior distributions for the hyperparameters), we use proper but quite disperse distributions. The implementation details are reported in the appendix.

2.3 Bayesian vector autoregression and dynamic factor model

The connection between Bayesian shrinkage and dynamic factor models is better understood by focusing on the data that have been transformed to achieve stationarity, Δy_t, and that have been standardised to have mean zero and unit variance.

The VAR in differences can be represented by:

$$\Delta y_t = B_1 \Delta y_{t-1} + \ldots + B_{\tilde{p}} \Delta y_{t-\tilde{p}} + \varepsilon_t.$$

Imposing that the level of each variable y_t follows an independent random walk process, is equivalent to imposing that its difference, Δy_t, follows an independent white noise process.

Consequently, the prior on the autoregressive coefficients can be characterised by the following first and second moments:

$$
\begin{aligned}
E\left[(B_s)_{ij} | \Sigma, \lambda, \Psi\right] &= 0 \\
cov\left[(B_s)_{ij}, (B_r)_{hm} | \Sigma, \lambda, \Psi\right] &= \begin{cases} \lambda^2 \frac{1}{s^2} \frac{\Sigma_{ih}}{\Psi_{jj}} & \text{if } m = j \text{ and } r = s \\ 0 & \text{otherwise} \end{cases}.
\end{aligned}
$$

Since the variables are rescaled to have the same variance, the hyperparameter related to the scale can be set to be the same for all variables, i.e. $\Psi = \bar{\psi} I_n$.

The model can be rewritten in terms of the principal components described in section 2.1:

$$\Delta y_t = B_1 H^{-1} z_{t-1} + \ldots + B_{\tilde{p}} H^{-1} z_{t-\tilde{p}} + \varepsilon_t,$$

where $z_t = H \Delta y_t$ are the ordered principal components.

Interestingly, the prior set-up that imposes a uniform shrinkage on the parameters is equivalent to imposing a non-uniform degree of shrinkage on principal components:

$$E\left[\left(B_s H^{-1}\right)_{ij} | \Sigma, \lambda, \bar{\psi}\right] = 0$$

$$cov\left(\left(B_s H^{-1}\right)_{ij}, \left(B_r H^{-1}\right)_{hm} | \Sigma, \lambda, \bar{\psi}\right) = \begin{cases} (\lambda^2 d_j)\frac{1}{s^2}\frac{\Sigma_{ih}}{\bar{\psi}} & \text{if } m = j \text{ and } r = s \\ 0 & \text{otherwise} \end{cases}.$$

In fact, the prior variance for the coefficients on the j^{th} principal component turns out to be proportional to the variance explained by the latter (d_j).

As discussed in section 2.1, if the data are characterised by a factor structure then, as the number of variables and the sample size increase, d_j will go to infinity at a rate n for $j = 1, \ldots, r$, where r is the number of common factors. On the other hand, $d_{r+1}, \ldots, d_n$ will grow at a slower rate which cannot be faster than $n/\sqrt{T}$. As a consequence, if λ goes to zero at a rate that is faster than that for the smaller eigenvalues and slower than for the largest eigenvalues, i.e. $\lambda^2 = \kappa \frac{\sqrt{T}}{n} \frac{1}{T^\delta}$ with $0 < \delta < 1/2$ and κ an arbitrary constant, then $\lambda^2 d_j$ will go to infinity for $j = 1, \ldots, r$. Hence the prior on the coefficients associated with the first r principal components will become flat. Instead, for $j > r$, $\lambda^2 d_j$ will go to zero, i.e. the coefficients related to the principal components associated with the bounded eigenvalues will be shrunk to zero.

De Mol, Giannone, and Reichlin (2008) have shown that, if the data are generated accordingly to a factor model and the hyperparameter λ is set according to the rate described above, the point forecasts obtained by using shrinkage estimators converge to the unfeasible optimal forecasts that would be obtained if the common factors were observed.

3 Conditional forecasts for linear state space representations

3.1 Linear state space representation

Several univariate and multivariate time-series models may be cast in a linear state space representation. For the sake of notation, the generic linear state space representation is defined as[10]:

Measurement equation

$$Z_t = C_t S_t + v_t \tag{3}$$

Transition equation

$$S_{t+1} = G_t S_t + w_t \tag{4}$$

where $Z_t = (Z_{1,t}, Z_{2,t}, \ldots, Z_{k,t})'$ is a k-dimensional vector of observables, S_t an m-dimensional vector of potentially unobserved states, v_t and u_t two vectors of errors with: $v_t \sim$ i.d. $N(0, R_t)$, $w_t \sim$ i.d. $N(0, H_t)$ and $E[v_t w_s'] = 0 \ \forall \ t, s$. Finally, C_t and G_t are two, respectively, $k \times m$ and $m \times m$ matrices of potentially time-varying coefficients.

The dynamic factor model in (2) can be cast in the representation (3)-(4) with $Z_t := \Delta y_t$, $C_t := (\Lambda, \ 0_{n \times r(s-1)}, \ I_n)$, $R_t := \Gamma_d$ and

$$S_t := \begin{pmatrix} F_t \\ \vdots \\ F_{t-s+1} \\ \mu \end{pmatrix}, \quad G_t := \begin{pmatrix} \Phi_1 & \Phi_2 & \ldots & \Phi_s & 0_{r \times n} \\ I_r & 0_r & \ldots & 0_r & \vdots \\ \vdots & \ddots & \ldots & \vdots & \vdots \\ 0_r & \ldots & I_r & 0_r & 0_{r \times n} \\ 0_{n \times r} & \ldots & \ldots & 0_{n \times r} & I_n \end{pmatrix}, \quad H_t := \begin{pmatrix} Q & \ldots & 0_{r \times n} \\ \vdots & \ddots & \vdots \\ 0_{n \times r} & \ldots & 0_n \end{pmatrix}.$$

---

[10]See Harvey (1989) for a thorough treatment of state space techniques. To simplify notation we abstract from exogenous variables as they are not included in our empirical models.

For the VAR in (1), we have $Z_t := Y_t$, $C_t := (I_n, \ 0_{n \times np})$, $R_t := 0_n$ and

$$
S_t := \begin{pmatrix} Y_t \\ \vdots \\ Y_{t-p+1} \\ c \end{pmatrix}, \quad G_t := \begin{pmatrix} A_1 & A_2 & \ldots & A_p & I_n \\ I_n & 0_n & \ldots & 0_n & 0_n \\ \vdots & \ddots & \ldots & \vdots & \vdots \\ 0_n & \ldots & I_n & 0_n & 0_n \\ 0_n & \ldots & 0_n & 0_n & I_n \end{pmatrix}, \quad H_t := \begin{pmatrix} \Sigma & \ldots & 0_n \\ \vdots & \ddots & \vdots \\ 0_n & \ldots & 0_n \end{pmatrix}.
$$

For the implementation in differences the modifications are straightforward.

3.2 Conditional forecasts

Simulation smoothers (see Carter and Kohn, 1994; de Jong and Shephard, 1995; Durbin and Koopman, 2002, for example) can be used to generate a draw of the state vector S_t, $t = 1, \ldots, T$ conditional on the observations $\{Z_t, t = 1, \ldots, T\}$ and on (a draw of) the parameters, C_t, G_t, R_t, H_t, $t = 1, \ldots, T$:

$$
S_{t|T} \sim p(S_t | Z_t, C_t, G_t, R_t, H_t, t = 1, \ldots, T), \quad t = 1, \ldots, T.
$$

Let us now assume that for a subset of variables, $\mathcal{I}$, we are interested in obtaining conditional forecasts for $t > t_0$, conditional on their own past and on the past and future observations of the remaining variables, i.e. conditional on the information set $\Omega = \{Z_{l,t}, l \in \mathcal{I}, t \leq t_0, Z_{l,t}, l \notin \mathcal{I}, t = 1, \ldots, T\}$:

$$
Z_{i,t|\Omega} \sim p(Z_{i,t}|\Omega), \quad i \in \mathcal{I}, t > t_0. \tag{5}
$$

In order to obtain such conditional forecasts, we adopt the solution proposed for forecasting with ragged edge data sets using a Kalman filter methodology, see e.g. Giannone, Reichlin, and Small (2005). In fact, the variables for which we do not assume the knowledge of a future path can be considered as time series with missing data. The Kalman filter allows to easily deal with such time series. Going more in details, we follow a standard approach (see e.g. Durbin and Koopman, 2001, pp. 92-93) and apply the Kalman filter to a modified state space representation with Z_t, C_t and R_t replaced by $\bar{Z}_t$, $\bar{C}_t$ and $\bar{R}_t$ respectively. The latter are

derived from the former by removing the rows (and, for R_t, also columns) that correspond to the missing observations in Z_t.[11]

This insight is already sufficient in order to compute point conditional forecasts, as the Kalman smoother gives the expectation of the distribution in (5), conditional on the parameters. In addition, assuming that the posterior distribution of the model parameters conditional on the data is available, the following algorithm (described for the generic iteration j) may be used in order to draw from the distribution of the conditional forecasts:[12]

(i) Draw the parameters $\tilde{C}_t^{(j)}$, $\tilde{G}_t^{(j)}$, $\tilde{R}_t^{(j)}$ and $\tilde{H}_t^{(j)}$ from their posterior distribution.[13]

(ii) Draw the states $\tilde{S}_{t|\Omega}^{(j)}$ using a simulation smoother (Carter and Kohn, 1994; de Jong and Shephard, 1995; Durbin and Koopman, 2002) for the modified (for the missing data) state space representation with the parameters $\bar{\tilde{C}}_t^{(j)}$, $\tilde{G}_t^{(j)}$, $\bar{\tilde{R}}_t^{(j)}$ and $\tilde{H}_t^{(j)}$.

(iii) Draw the disturbances for the observation equation, $\tilde{v}_{i,t|\Omega}^{(j)}$, from a conditional multivariate distribution $p(v_{i,t}|\tilde{v}_{l,t|\Omega}^{(j)}, l \notin \mathcal{I})$, $i \in \mathcal{I}, t > t_0$.[14] In fact, for a VAR $\tilde{v}_{i,t|\Omega}^{(j)} \equiv 0$.

(iv) Compute $\tilde{Z}_{i,t|\Omega}^{(j)} = (\tilde{C}_t^{(j)})_{i\cdot} \tilde{S}_{t|\Omega}^{(j)} + \tilde{v}_{i,t|\Omega}^{(j)}$, $i \in \mathcal{I}, t > t_0$.

The algorithm can be modified in a straightforward manner for any pattern of "missing" observations in Z_t. Note that we can also easily condition on a linear combination of the observations. Suppose, in fact, that the aim is to condition on j linear combinations $W_t Z_t$

---

[11]Giannone, Reichlin, and Small (2005) propose an equivalent solution. Instead of removing rows (and columns) of C_t and R_t that correspond to missing observations, they replace R_t with $\bar{R}_t$ defined as follows:

$$(\bar{R}_t)_{ij} = \begin{cases} (R_t)_{ij} & \text{if } Z_{i,t} \text{ and } Z_{j,t} \text{ are available} \\ \infty & \text{otherwise} \end{cases}$$

where, in practice, ∞ is a large number.

[12]We will denote a draw of a random variable from a distribution by $\tilde{\ }$.

[13]In order to take the available future paths of selected variables into account when drawing the parameters, the initial values can be obtained from the "balanced" data set (up to t_0) and then steps 1-4 can be iterated, with the latest conditional forecasts treated as data when drawing the parameters in step (i).

[14]See e.g. Greene (2002) pp. 872 for conditional normal distributions. $\tilde{v}_{l,t|\Omega}^{(j)}, l \notin \mathcal{I}$ can be obtained from $Z_{l,t} - (\tilde{C}_t^{(j)})_{l\cdot} \tilde{S}_{t|\Omega}^{(j)}$ or, alternatively, a disturbance smoother (see de Jong and Shephard, 1995; Durbin and Koopman, 2002) can be used and the states can be derived indirectly from the disturbances.

where W_t is a sequence of matrices of dimension $j \times k$. Then we set

$$\check{Z}_t = \begin{pmatrix} W_t Z_t \\ Z_t \end{pmatrix}, \qquad \check{C}_t = \begin{pmatrix} W_t C_t \\ C_t \end{pmatrix}, \qquad \check{R}_t = \begin{pmatrix} W_t R_t W_t' & W_t R_t \\ R_t W_t' & R_t \end{pmatrix}.$$

Note that for a large VAR model with several lags, the size of the state vector, S_t, is much larger than the size of the vector of observables, Z_t. It turns out that in this case the simulation smoother of Durbin and Koopman (2002) along with the Kalman smoother implementation of de Jong (1988) offers large computational gains compared to the algorithm of Carter and Kohn (1994), see the appendix. This is the implementation we adopt in the empirical exercises.[15]

3.3 Comparison with the approach of Waggoner and Zha (1999)

The algorithm of Waggoner and Zha (1999) is a popular method to obtain conditional forecasts for VAR models. Roughly speaking, the methodology involves drawing directly vectors of ε_t, $t = t_0 + 1, \ldots, T$, in (1) which satisfy the required conditions.

For the VAR described above and the pattern of variable availability discussed in section 3.2:

$$\Omega = \{Y_{l,t}, l \in \mathcal{I}, t \leq t_0, Y_{l,t}, l \notin \mathcal{I}, t = 1, \ldots, T\},$$

this would involve an inversion of a $q \times q$ matrix, where $q = (n - \#(\mathcal{I}))(T - t0)$ denotes the number of restrictions, and, more importantly, a spectral decomposition of a $n(T-t0) \times n(T-t0)$ matrix (see e.g. Jarociński, 2010, for a detailed discussion).[16] Jarociński (2010) proposes a way to decrease the computational complexity which involves a singular value decomposition of a $q \times n(T-t0)$ matrix and avoids the matrix inversion. However, the complexity still heavily depends on the number of restrictions, which can be prohibitively large in case of a large data set and a long forecast horizon. In contrast, the application of the Kalman filter makes the problem "recursive", and the size of the largest matrix to be inverted is independent of the forecast horizon.

The appendix provides an assessment of computational performance of different algorithms for obtaining conditional forecast distributions for VARs and indeed shows major computational

---

[15]To further improve the computational performance for the VARs we include the intercept, c, as a constant in the transition equation (4) and remove it from the state vector, S_t, see the appendix.

[16]$\#(\mathcal{I})$ denotes the number of elements in $\mathcal{I}$.

gains of the Kalman filter based algorithm over the one of Waggoner and Zha (1999) and Jarociński (2010) when the number of restrictions increases.

3.4 Conditional forecasts and structural scenarios

The scenario analysis described above and considered in this paper is reduced form, in the sense that the scenarios are based on all the innovations that are compatible with the conditioning information. An alternative approach consists in constructing structural scenarios based only on certain structural shocks. This amounts to considering only some specific linear combination of the reduced form innovations. These scenarios can be implemented by modifying the covariance matrix of the innovations in the state space representation.

Let us consider the case of the VAR first. Let us denote by $\eta_t = (\eta_{1,t}, \ldots, \eta_{n,t})'$ the structural shocks, assumed to be orthonormal, $\mathrm{E}(\eta_t \eta_t') = I_n$. We have that $\varepsilon_t = \Gamma \eta_t = \gamma_1 \eta_{1,t} + \cdots + \gamma_n \eta_{n,t}$, where Γ contains the contemporaneous responses to the shocks and $\Gamma \Gamma' = \Sigma$. If one is interested in obtaining the scenario based only on the last $n - s$ shocks (i.e. $\eta_{1,t}, \ldots, \eta_{s,t}$ are assumed to be 0), it is sufficient to replace Σ with $\bar{\Sigma} = \gamma_{s+1} \gamma_{s+1}' + \ldots + \gamma_n \gamma_n'$.[17]

In the case of the factor model, structural shocks, $\eta_t = (\eta_{1,t}, \ldots, \eta_{r,t})'$, $\mathrm{E}(\eta_t \eta_t') = I_r$, are linear combinations of the shocks to the factors. We have that $u_t = \Gamma \eta_t = \gamma_1 \eta_{1,t} + \cdots + \gamma_r \eta_{r,t}$, where Γ contains the contemporaneous responses to the shocks and $\Gamma \Gamma' = Q$. If one is interested in obtaining the scenario based only on the last $r - s$ shocks, that it is sufficient to replace Q with $\bar{Q} = \gamma_{s+1} \gamma_{s+1}' + \ldots + \gamma_r \gamma_r'$.

Another observation is that the conditional forecast framework can be used to obtain impulse response functions for VARs with recursive identification schemes. In this case Γ is lower triangular and the impulse response function to a shock η_{it} can be obtained as:

$$\mathrm{IRF}_j^i = \mathrm{E}\left(Y_{t+j} | \epsilon_{1,t} = 0, \ldots, \epsilon_{i-1,t} = 0, \epsilon_{i,t} = \gamma_{ii}; Y_{t-1}, \ldots, Y_{t-p}\right) - \mathrm{E}\left(Y_{t+j} | Y_{t-1}, \ldots, Y_{t-p}\right),$$

(6)

$j = 0, \ldots$ The second term on the right-hand side of equation (6) is the unconditional forecast. The conditions $\epsilon_{i,t} = 0$ ($\epsilon_{i,t} = \gamma_{ii}$) are implemented by setting $Y_{i,t}$ equal to its unconditional

---

[17]One has to further insure that the reduced form innovations implied by the scenario are not "incompatible" with the assumptions on $\bar{\Sigma}$. For example, in a recursive identification scheme with Γ lower triangular, one cannot impose that $\eta_{1,t} = 0$ and $\epsilon_{1,t} \neq 0$.

forecast (plus γ_{ii}).[18] For $i = 1$, this is akin to estimating a generalised impulse response function (on this point, see Koop, Pesaran, and Potter, 1996; Pesaran and Shin, 1998).

4 Empirical Results

4.1 Data

Our data set includes 26 quarterly variables. Roughly, we include the most relevant real and nominal euro area aggregates and a set of international macroeconomic variables that proxies for global macroeconomic conditions (GDP and expenditure components, consumer and producer prices, labour market data, surveys, effective exchange rate, world economic activity, commodity prices), financial variables (short and long-term interest rates, stock prices), credit (both to households and firms) and monetary aggregates (M1 and M3).

The sample covers the period from 1995Q1 to 2012Q4. Most of the data come from the Area Wide Model data base (Fagan, Henry, and Mestre, 2005). The share price index is downloaded from Datastream and the US short-term interest rate from the IMF data base. Remaining variables can be downloaded from the ECB Statistical Data Warehouse. The data appendix at the end of the paper provides the details, including the data transformations applied prior to parameter estimation. For most of the variables that are not already expressed in rates we take the natural logarithm multiplied by 4. For the models specified in differences we further take first differences for all the variables. For the logged variables this corresponds to annualised one-period percentage changes (in decimal).[19]

4.2 Model specifications

We include 3 common factors in the DFM, which are meant to roughly capture the information on real developments, prices and interest rates. Further, we include 4 lags in the VAR of the common factors. The model is estimated by maximum likelihood using the EM algorithm. We also include 4 lags in the BVAR in differences and in order to be consistent with the

---

[18]On the connection between scenario analysis and impulse response functions to identified shocks see also Jardet, Monfort, and Pegoraro (2012).

[19]This reflects the setting of the prior for Ψ, see the appendix.

dynamics captured in the two approaches specified on variables in differences, in the BVAR in levels we include 5 lags. In the BVAR approaches, we have to choose the tightness of the prior distributions. As suggested in Giannone, Lenza, and Primiceri (2014), we follow a hierarchical approach and we treat the hyperparameters governing such tightness as random variables with relatively diffuse prior distributions.

4.3 Model validation: out-of-sample forecasting evaluation of unconditional forecasts

As a preliminary step, we gauge the accuracy of our empirical models in terms of out-of-sample unconditional forecasts. This preliminary step is particularly important because our models are specified on a large set of variables (26). This feature leads to a proliferation of parameters and, hence, potential instabilities due to estimation uncertainty might completely offset the gains obtained by limiting model misspecification due to variable omission. Assessing out-of-sample forecast accuracy, which reflects both estimation uncertainty and model misspecification, allows us to understand whether the benefits due to the generality of our models outweigh the costs.

We focus on point forecasts. For the DFM the forecasts are easily obtained with the Kalman filter. For the BVARs we use the modes of the posterior distribution of the parameters and the forecasts are obtained using the chain rule.

For each of the three models, we produce the forecasts recursively for three horizons (1, 2 and 4 quarters ahead). The evaluation period is 2004-2012. For each of the evaluation periods, $t = 2004Q1, \ldots, 2012Q4$, and forecast horizons, $h = 1, 2, 4$, the forecasts are produced using the data from 1995Q1 to $t - h$. At each step, the parameters (including the mode of the posterior distribution of the hyperparameters) are re-estimated.

For each variable, the target of our evaluation is defined as $m_{i,t+h}^h = \frac{100}{h}[y_{i,t+h} - y_{i,t}]$. For variables specified in logs, this is approximately the average annualised growth rate over the next h quarters (in percent), while for variables not transformed in logs, this is the average annualised quarterly change over the next h quarters .

We compare our model with a simple benchmark model, namely a random walk with drift for

the (log-)levels of the variables. This model is a particularly challenging benchmark over the monetary union sample and, in addition, it has also the appeal of being the "prior model" for the BVAR approaches. Hence, in case the BVAR models out-perform the benchmark model, this would suggest that they are able to extract valuable information from the sample.[20]

Table 1 reports the results of our analysis, for all models and variables. Results are cast in terms of ratios of the Mean Squared Forecast Errors (MSFE) of our three models with respect to the corresponding MSFE of the random walk benchmark model. Hence, values smaller than one indicate that our model outperforms the benchmark model.

INSERT TABLE 1

The outcomes of the evaluation show that, in general, the BVARs and the DFM are able to outperform the benchmark model, particularly for the short horizons. For the one- and two-quarter horizons, the models are more accurate than the random walk for most of the variables. Some failures to outperform the benchmark model are not particularly surprising, since it is well known that it is very hard to beat the random walk model for financial and commodity prices, in general. Also for consumer prices (HICP) and the GDP deflator there is a relatively ample documentation of the difficulties to beat the random walk due to the steady anchoring of inflation expectations in the monetary union sample.[21]

As argued in section 2.3 and in De Mol, Giannone, and Reichlin (2008), there is a tight connection between the BVAR and the DFM approaches. Indeed, the out-of-sample performance of the three different approaches is quite similar. This reflects the fact that the forecasts from the three approaches are very correlated. Figure 1 (panels a-c) reports the bivariate correlation with the BVAR in levels of the DFM (black bar with stripes) and the BVAR in differences (red solid bar) for all the variables and forecast horizons.

INSERT FIGURE 1

The figure reveals the strong collinearity of the forecasts across approaches, providing empirical

---

[20]We also compared our models to a battery of univariate autoregressive models, another class of popular benchmark models, with very similar outcomes.

[21]See, for example, Diron and Mojon (2005) and Giannone, Lenza, Momferatou, and Onorante (2014).

21

support for their theoretical connection highlighted in section 2.3 and in De Mol, Giannone, and Reichlin (2008).

4.4   Scenario analysis: an increase in world GDP

In this exercise we perform a scenario analysis to assess the effects associated with positive developments in the global economy, represented by a 0.1 percentage point stronger growth (on impact) in global GDP.

We compute the effects of the scenario by using our framework to produce conditional forecasts, as discussed in section 3.4. Precisely, we estimate our models on the whole sample and generate two forecasts: an unconditional forecast for $T + 1, \ldots, T + h$ given the sample $1, 2, \ldots, T$ (which provides a "baseline" scenario) and a conditional forecast in which the world GDP growth in $T + 1$ is set to the value of its own unconditional forecast plus 0.1 percentage points and all the remaining variables are left unconstrained (which we will refer to as the "shocked scenario"). The scenario results for all variables are computed by taking the difference between the conditional and the unconditional forecasts described above. This is equivalent to computing a generalised impulse response function to an increase in world GDP, see section 3.4. We explore the horizon of eight quarters.

Figure 2 shows the responses of some selected variables for the three models.[22] In particular, we report the distribution of the scenario effects computed in the context of the BVAR model in levels (shades of orange) and the point estimates of the effects in the other two modelling approaches (DFM: dashed blue line; BVAR in differences: black solid line). For the BVARs we use the algorithm described in section 3.2 and the point estimates of the effects for the BVAR in differences are defined as the medians of the distribution.[23] For the DFM we simply use the Kalman smoother to obtain the (conditional) forecasts given the maximum likelihood estimates of the parameters. All results are reported in terms of deviations of (log-)levels of the variables in the shocked scenario compared to the baseline.[24]

---

[22]The total set of responses is available upon request.

[23]We generate 25000 draws from the posterior distribution of the parameters and discard the first 5000. For each of the remaining draws we compute a point (conditional) forecast using the Kalman smoother.

[24]For the variables that are modelled in logs, this approximately corresponds to percentage deviation from the baseline for the levels.

INSERT FIGURE 2

The three approaches provide similar scenario assessments for all variables, at least qualitatively but, generally, also quantitatively. This result confirms the view that, for the variables commonly used in macroeconometric studies, dynamic factor models and Bayesian shrinkage (irrespective of data transformation) are both valid alternative methods to deal with the curse of dimensionality.

Going more in details of the scenario analysis, the top left panel reports the developments in global real GDP, which is 0.1 percent higher on impact (as assumed in the scenario assumption), keeps on increasing for the first year and then tends to drop back to the level prevailing before the initial increase.

The euro area real economy (GDP, exports, imports and unemployment) closely mirrors the developments in global GDP. Consumer prices are also higher, reaching a peak after about one year.

The short-term interest rate, which reflects systematic monetary policy, reacts positively to stabilize the economy and then drops back toward the initial level. According to the BVAR results, long-term interest rates are not particularly affected. This implies that the term-spread (defined as long-term interest rates minus short-term interest rates) decreases on impact, to finally revert to initial values. In this case, the DFM results are different from those of the BVARs.

Credit aggregates, which are traditionally very cyclical, follow the same path as GDP. Moreover, loans to households are coincident with GDP, while loans to firms lag behind. The narrow monetary aggregate M1 decreases on impact, reaching a trough after about 1-1.5 year. To understand this pattern, notice that M1 is negatively related to the short-term interest rate, indicating that its response to world demand is mainly driven by the liquidity effect. The effect on M3 is instead mostly driven by the increases in the short-term monetary assets included in the M3-M1 component, which completely offset the decrease in M1 (see Giannone, Lenza, and Reichlin, 2012, for an extensive discussion on the cyclical properties of credit and monetary aggregates in the euro area and their relationships with short and long-term interest rates).

23

An implicitly maintained assumption in this exercise is that the forecast paths we examine involve shocks small enough so as not to be subject to the Lucas critique.[25] Indeed, the reliability of the results rests on the fact that the perturbations we induce in the system by means of the scenario assumptions are not as big as to induce a substantial shift in the behavior of economic agents which could, in turn, change the economic structure and, hence, the estimated reduced form parameters.

4.5 Conditional forecasts

In this exercise we generate forecasts from the three models conditional on the realised paths for the following three variables: real GDP, HICP and the short-term interest rate.

The conditional forecasts are generated over the period 1997-2012. The first two years in the sample are used as initial conditions. The parameters are estimated over the sample 1995-2007. Thus the conditional forecasts for 1997-2007 can be considered as "in-sample" while those over 2008-2012 as "out-of-sample". The aim of this exercise (see Giannone, Lenza, and Reichlin, 2012; Stock and Watson, 2012a, for similar exercises) is twofold. First, the "in-sample" part (1997-2007) of the conditional forecasts can be compared with the observed developments in order to gauge whether knowing only the time series of real GDP, HICP and the short-term interest rate is sufficient in order to capture the salient features of the variables in our model. Second, by comparing the "out-of-sample" part (i.e. from 2008 onward) of the conditional forecasts with the observed developments, we can also assess whether the turmoil associated to the financial and the sovereign debt crises was reflected in a change in the structural economic relationships in the euro area. In fact, a change in the economic relationships would likely lead to relevant inaccuracies of the conditional forecasts based on parameters representing the pre-2007 economic relationships.

Figure 3 shows the conditional forecasts from the three models for the same selected variables shown in Figure 2.[26] As in the previous exercise, the distribution is generated using the BVAR in levels, using the algorithm described in section 3.2.[27] Blue dashed and black solid

---

[25]See, Kilian and Lewis (2011) and references therein for a discussion of this issue.

[26]Three additional variables replace the three variables shown in Figure 2 which in this exercise were used as conditions. A complete set of results is available upon request.

[27]We generate 25000 draws from the posterior distribution of the parameters and discard the first 5000. For

lines correspond, respectively, to the conditional point forecasts of the DFM, obtained via the Kalman smoother, and the BVAR in differences, obtained as the median of the distribution. In addition, the green line indicates the actual outcomes. For the interest rates and the unemployment rate we report the conditional forecasts for the levels. For the remaining variables the results are expressed in terms of annual rates of change.

INSERT FIGURE 3

Analogously to the previous exercise, the forecasts from the three models are similar for most of the variables, indicating that different methodologies capture similar cross-sectional and dynamic information. In addition, the conditional forecasts are close to the actual outcomes, in particular in the "in-sample" period. This fact suggests that 3 "dimensions" are sufficient to capture the developments in most of the economy[28] (Giannone, Reichlin, and Sala, 2004, reach a similar conclusion for the US economy).

Turning to the "out-of-sample" evidence, there is still a general similarity of the conditional forecasts across approaches. However, some differences appear between forecasts and observed developments for a few variables, indicating an instability in the relationships of these variables with the conditioning set. For example, notable differences appear in the developments in money and credit variables, whose actual developments were much more subdued than what would have been predicted based on the conditioning information.[29] For the variables where we have evidence of instability, we also notice some discrepancies in the forecasts across methods.

5 Conclusions

We have modelled the dynamic interactions among a large set of macroeconomic and financial indicators in the euro area by means of large dynamic factor models and large Bayesian vector autoregressions.

---

each of the remaining draws we compute a draw of the conditional forecast.

[28]Notable exceptions are wages, GDP deflator, government consumption and the effective exchange rate (not shown). For these variables, the conditional forecast distributions cover a relatively wide range of values and the central forecasts are often quite far from the outcomes.

[29]Giannone, Lenza, and Reichlin (2012) extensively discuss and interpret the anomalies in the developments in credit and money markets during the crisis.

We find that both classes of models are reliable tools for analyzing large data sets, since they produce accurate unconditional forecasts and meaningful scenarios.

Interestingly, the predictions of the two model classes are not only equally reliable, but are also very similar, in general. The fact that the results are not model specific is reassuring since it indicates that the predictions of the models are reflecting genuine data features.

The robustness and reliability of dynamic factor models and Bayesian vector autoregressions for analyzing large macroeconomic data sets has been already established for the United States in relation to forecasting and impulse response function analysis (see e.g. Bańbura, Giannone, and Reichlin, 2010; Giannone, Lenza, and Primiceri, 2014). We document that the same holds true for the euro area and for conditional forecasts.

In addition, we have shown how to implement scenario analysis and, in general, to compute conditional forecasts in the context of large data sets. The procedure is computationally feasible, produces meaningful results and interesting insights. The methodology has been already used in a number of papers including Giannone, Lenza, Momferatou, and Onorante (2014), Giannone, Lenza, and Reichlin (2010, 2012), Giannone, Lenza, Pill, and Reichlin (2012), Lenza, Pill, and Reichlin (2010) and Luciani (2013).

A   Simulation smoothers

A.1   Implementation

As mentioned above, for the VARs, we consider a version of the transition equation (4), in which the intercept is included as a constant:

$$S_{t+1} = \bar{c}_t + G_t S_t + w_t, \tag{7}$$

where $\bar{c}_t = (c' \; 0_{1 \times n(p-1)})'$ whereas S_t, w_t, G_t and H_t are obtained from the corresponding terms in equation (4) by removing the last n rows (and columns).

Given the parameter set C_t, G_t, R_t, H_t and $\bar{c}_t$, the algorithm of Carter and Kohn (1994) derives draws from the conditional distribution of the state vector, $\tilde{S}_{t|T}, t = 1, 2, \ldots, T$, from the following recursions:

$$
\begin{aligned}
S_{bt} &= S_{t|t} + P_{t|t} G_t' (G_t P_{t|t} G_t' + H_t)^{-1} (\tilde{S}_{t+1|T} - G_t S_{t|t} - \bar{c}_t), \\
P_{bt} &= P_{t|t} - P_{t|t} G_t' (G_t P_{t|t} G_t' + H_t)^{-1} G_t P_{t|t}, \\
\tilde{S}_{t|T} &= S_{bt} + \tilde{\xi}_t, \qquad\qquad \xi_t \sim N(0, P_{bt}),
\end{aligned}
$$

with $\tilde{S}_{T|T} = S_{T|T} + \tilde{\xi}_T$, $\xi_T \sim N(0, P_{T|T})$. $S_{t|t} = \mathrm{E}\left[S_t | Z_1, \ldots, Z_t\right]$ and $P_{t|t} = \mathrm{Var}\left[S_t | Z_1, \ldots, Z_t\right]$ are obtained from the Kalman filter (see below). In the case of a VAR, the algorithm involves a (pseudo) inversion of matrices of size $np \times np$ as well as a singular value decomposition of matrices of the same size.

The simulation smoother of Durbin and Koopman (2002) can be implemented via the following steps:

(i) Draw the disturbances $\tilde{v}_t$ and $\tilde{w}_t$, $t = 1, 2, \ldots, T$, from the unconditional distribution of v_t and w_t, i.e. $N(0, R_t)$ and $N(0, H_t)$, respectively.

(ii) Generate $\tilde{Z}_t$ and $\tilde{S}_t$, $t = 1, 2, \ldots, T$, using the state space representation given by (3) and (7) and $\tilde{v}_t$ and $\tilde{w}_t$ from the previous step.

(iii) A draw from the conditional distribution of the state vector can be obtained as $\tilde{S}_{t|T} = \mathrm{E}\left[S_t | \tilde{Z}_1^*, \ldots, \tilde{Z}_T^*\right] + \tilde{S}_t$, where $\tilde{Z}_t^* = Z_t - \tilde{Z}_t$, $t = 1, 2, \ldots, T$.

Let $S_{t|s}^* = \mathrm{E}\left[\mathrm{S_t}|\tilde{\mathrm{Z}}_1^*, \ldots, \tilde{\mathrm{Z}}_\mathrm{s}^*\right]$ and $P_{t|s}^* = \mathrm{Var}\left[\mathrm{S_t}|\tilde{\mathrm{Z}}_1^*, \ldots, \tilde{\mathrm{Z}}_\mathrm{s}^*\right]$. We obtain $S_{t|T}^*$, $t = 1, 2, \ldots, T$, using the following implementation of the Kalman filter:

$$E_t = Z_t^* - C_t S_{t|t-1}^*, \qquad\qquad F_t = C_t P_{t|t-1}^* C_t' + R_t,$$

$$K_t = P_{t|t-1}^* C_t' F_t^{-1}, \qquad\qquad L_t = I - K_t C_t,$$

$$S_{t|t}^* = S_{t|t-1}^* + K_t E_t, \qquad\qquad P_{t|t}^* = P_{t|t-1}^* L_t'$$

$$S_{t+1|t}^* = G_t S_{t|t}^* + \bar{c}_t, \qquad\qquad P_{t+1|t}^* = G_t P_{t|t}^* G_t' + H_t$$

and smoother:

$$r_{t-1} = C_t' F_t^{-1} E_t + L_t' r_t, \qquad\qquad S_{t|T}^* = S_{t|t-1}^* + P_{t|t-1}^* r_{t-1},$$

with $r_T = 0$. This approach involves the inversion of matrix F_t, which has the (row and column) size of $n - \#(\mathcal{I}) \leq n$. It is anyway needed for the run of the Kalman filter and can be stored. By contrast, the implementation of the Kalman smoother as in e.g. Hamilton (1994) for a VAR requires the inversion of $np \times np$ matrices $P_{t+1|t}$. Thus, the approach of Durbin and Koopman (2002) can offer sizable computational gains for large p.

In the empirical exercises we run the simulation smoother only for the part of the sample with missing data, that is we truncate the data keeping only Z_t, $t > t_0$ and we use $Z_{t_0}, Z_{t_0-1}, \ldots$ to derive the initial conditions, S_0 (e.g. in case of the VAR in levels $S_0 = (Y_{t_0}' \ \ldots \ Y_{t_0-p+1}')'$). Then we set $S_{1|0} = G_0 S_0 + \bar{c}_0$ and $P_{1|0} = H_0$. For the simulation smoother of Durbin and Koopman (2002) we set $S_0^* = 0$ and $\bar{c}_t = 0_{np \times 1}$.

A.2   Computational time

Table A compares the average computational time of a draw of conditional forecast for the 26 variable VAR considered in section 4 for the following algorithms: the simulation smoother of Carter and Kohn (1994) (*CK*); the simulation smoother of Durbin and Koopman (2002) with a "traditional" Kalman smoother implementation (see e.g. Hamilton, 1994, pp. 394-396) (*DK,H*); the simulation smoother of Durbin and Koopman (2002) with the Kalman smoother implementation of de Jong (1988) (*DK,dJ*); the algorithm of Waggoner and Zha (1999) with the implementation of Jarociński (2010) (*WZ,J*)[30]. We consider the cases of different number of

---

[30] We would like to thank Marek Jarociński for sharing his Matlab code for this procedure.

lags ($p = 2$, $p = 3$ or $p = 5$), different number of variables in the conditioning set ($n - \#(\mathcal{I}) = 5$, $n - \#(\mathcal{I}) = 15$ or $n - \#(\mathcal{I}) = 25$) and different number of conditioning periods ($T - t_0 = 5$, $T - t_0 = 20$ or $T - t_0 = 60$). In each case the time for the fastest algorithm is marked in boldface, whereas the time for the slowest is put in italics. Average time in seconds over 1000 repetitions is reported.[31]

Table A: Average time of a draw of conditional forecast

	CK	DK,H	DK,dJ	WZ,J	CK	DK,H	DK,dJ	WZ,J	CK	DK,H	DK,dJ	WZ,J
	$n - \#(\mathcal{I}) = 5$				$n - \#(\mathcal{I}) = 15$				$n - \#(\mathcal{I}) = 25$			
	$p = 2$											
$T - t_0 = 5$	*0.04*	0.02	**0.01**	0.01	*0.04*	0.02	**0.01**	0.02	0.03	0.02	0.01	*0.05*
$T - t_0 = 20$	*0.16*	0.09	**0.04**	0.13	0.16	0.09	**0.04**	*0.51*	0.13	0.08	**0.03**	*1.51*
$T - t_0 = 60$	0.42	0.28	**0.06**	*2.20*	0.38	0.26	**0.07**	*10.61*	0.40	0.26	**0.08**	*39.25*
	$p = 3$											
$T - t_0 = 5$	*0.08*	0.05	0.01	**0.01**	*0.08*	0.05	**0.01**	0.02	*0.08*	0.05	**0.01**	0.05
$T - t_0 = 20$	*0.34*	0.23	**0.04**	0.15	0.32	0.20	**0.04**	*0.52*	0.35	0.20	**0.05**	*1.59*
$T - t_0 = 60$	1.03	0.69	**0.10**	*2.25*	0.96	0.62	**0.13**	*10.82*	1.11	0.64	**0.14**	*39.75*
	$p = 5$											
$T - t_0 = 5$	*0.26*	0.16	0.03	**0.02**	*0.25*	0.15	0.03	**0.03**	*0.27*	0.16	**0.03**	0.06
$T - t_0 = 20$	*1.14*	0.73	**0.11**	0.21	*1.05*	0.65	**0.12**	0.60	1.19	0.70	**0.12**	*1.62*
$T - t_0 = 60$	*3.46*	2.20	**0.33**	2.40	2.49	1.56	**0.34**	*10.34*	3.69	2.15	**0.39**	*39.67*

Note: Table provides average time in seconds over 1000 repetitions of a draw of conditional forecast for the 26 variable VAR considered in section 4. The following algorithms are considered: CK - the simulation smoother of Carter and Kohn (1994); DK,H - the simulation smoother of Durbin and Koopman (2002) with a "traditional" Kalman smoother implementation (see e.g. Hamilton, 1994, pp. 394-396); DK,dJ - the simulation smoother of Durbin and Koopman (2002) with the Kalman smoother implementation of de Jong (1988); WZ,J - the algorithm of Waggoner and Zha (1999) with the implementation of Jarociński (2010). $n - \#(\mathcal{I})$ refers to the number of variables in the conditioning set, $T - t_0$ to the number of conditioning periods and p is the number of lags in the VAR. The time of fastest algorithm in each case is marked in boldface, whereas the time for the slowest is put in italics.

The results indeed show large computational gains of the algorithms based on the Kalman filter over the approach of Waggoner and Zha (1999) and Jarociński (2010) when the number of restrictions (a combination of the number of conditioning variables and conditioning periods) increases. In the extreme case of 60 conditioning periods and 25 conditioning variables the latter algorithm needs 11 hours for 1000 draws compared to 1-7 minutes in case of the simulation smoother of Durbin and Koopman (2002) with the Kalman smoother implementation of de Jong (1988). This algorithm fares by far the best among the approaches based on the Kalman filter, with larger improvements relative to the algorithm of Carter and Kohn (1994) as the number of lags included in the VAR increases. As expected, the computational time of

---

[31] The computations were performed with Matlab R2008a on a computer with Intel CORE Duo ˜ 2925 Mhz processor and 4.1 GB physical memory.

the Waggoner and Zha (1999) and Jarociński (2010) algorithm is unaffected by the number of lags in the VAR while the performance of the algorithms based on the Kalman filter is not much affected by the number of conditioning variables.

B Estimation

B.1 Dynamic factor model

If the factors were observed the joint likelihood of the data and the factors would be easily maximised and the estimates of the parameters would correspond to ordinary least squares outcomes. Specifically Λ and Γ_d would be obtained by regressing Δy_t on F_t while the autoregressive parameters $\Phi_1, \ldots, \Phi_s$ and the covariance matrix Q would be obtained by regressing F_t on its lags, $F_{t-1}, \ldots, F_{t-s}$.

As the factors are unobserved the likelihood of the data cannot be maximised explicitly. As an alternative to numerical optimisation methods, the EM algorithm alternates between computing the expectation of the joint likelihood of the data and the factors given the parameter estimates from the previous step (E-step) and deriving new estimates by maximising the expected likelihood (M-step). An interesting property is that at each step the likelihood of the data increases, insuring that a convergence to a local maximum is reached.

Maximising the expected likelihood given the parameters at the j^{th} iteration is achieved through substituting the sufficient statistics with their expectation. This amounts to replacing the unobserved factors with their expected value $\hat{F}_t^{(j)} = \mathrm{E}_{\theta^{(j-1)}}[F_t | y_1, \ldots, y_T]$, and correcting for estimation uncertainty which is measured as

$$\hat{V}_{l,t}^{(j)} = \mathrm{E}_{\theta^{(j-1)}}\left[(F_t - \hat{F}_t^{(j)})(F_{t-l} - \hat{F}_{t-l}^{(j)})' | y_1, \ldots, y_T\right].$$

Those quantities can be computed recursively using the Kalman smoother.

It is easily seen that the expected sufficient statistics are as follows:

$$\mathrm{E}_{\theta^{(j-1)}}\left[\Delta y_t \Delta y_t' | y_1, \ldots, y_T\right] = \Delta y_t \Delta y_t',$$

$$\mathrm{E}_{\theta^{(j-1)}}\left[\Delta y_t F_t' | y_1, \ldots, y_T\right] = \Delta y_t \hat{F}_t^{(j)'}$$

and

$$\mathrm{E}_{\theta^{(j-1)}} \left[F_t F'_{t-l} | y_1, \ldots, y_T \right] = \hat{F}_t^{(j)} \hat{F}_{t-l}^{(j)'} + \hat{V}_{l,t}^{(j)}.$$

As a consequence the M-step consists of the following equations, where for simplicity we consider the case $s = 1$[32]:

$$\hat{\Lambda}^{(j)} = \left(\sum_{t=1}^{T} \Delta y_t \hat{F}_t^{(j)'} \right) \left(\sum_{t=1}^{T} \hat{F}_t^{(j)} \hat{F}_t^{(j)'} + \hat{V}_{0,t}^{(j)} \right)^{-1},$$

$$\hat{\Gamma}_d^{(j)} = \frac{1}{T} \mathrm{diag} \left[\left(\sum_{t=1}^{T} \Delta y_t \Delta y_t' \right) - \hat{\Lambda}^{(j)} \left(\sum_{t=1}^{T} \hat{F}_t^{(j)} \Delta y_t' \right) \right],$$

$$\hat{\Phi}_1^{(j)} = \left(\sum_{t=2}^{T} F_t^{(j)} \hat{F}_{t-1}^{(j)'} + \hat{V}_{1,t}^{(j)} \right) \left(\sum_{t=2}^{T} \hat{F}_{t-1}^{(j)} \hat{F}_{t-1}^{(j)'} + \hat{V}_{0,t-1}^{(j)} \right)^{-1}$$

and

$$\hat{Q}^{(j)} = \frac{1}{T-1} \left(\left(\sum_{t=2}^{T} F_t^{(j)} F_t^{(j)'} + \hat{V}_{0,t}^{(j)} \right) - \hat{A}^{(j)} \left(\sum_{t=2}^{T} \hat{F}_{t-1}^{(j)} \hat{F}_t^{(j)'} + \hat{V}_{1,t}^{(j)} \right) \right).$$

Principal components represent a good starting point for the EM algorithm. The initial estimates of the factor loadings are obtained by regressing Δy_t on the principal components:

$$\hat{\Lambda}^{(0)} = \left(\sum_{t=1}^{T} \Delta y_t \hat{F}_t^{(0)'} \right) \left(\sum_{t=1}^{T} \hat{F}_t^{(0)} \hat{F}_t^{(0)'} \right)^{-1}.$$

The variance of the idiosyncratic residuals is hence given by:

$$\hat{\Gamma}_d^{(0)} = \frac{1}{T} \mathrm{diag} \left[\left(\sum_{t=1}^{T} \Delta y_t \Delta y_t' \right) - \hat{\Lambda}^{(0)} \left(\sum_{t=1}^{T} \hat{F}_t^{(0)} \Delta y_t' \right) \right].$$

The total variance of the residual is given by: $\mathrm{trace}(\hat{\Gamma}_d^{(0)}) = d_{r+1} + \ldots + d_n$.

Turning to the estimation of the VAR model for the common factors, the OLS estimates, treating the factors as known, can be obtained as follows:

$$\hat{\Phi}_1^{(0)} = \left(\sum_{t=2}^{T} F_t^{(0)} \hat{F}_{t-1}^{(0)'} \right) \left(\sum_{t=2}^{T} \hat{F}_{t-1}^{(0)} \hat{F}_{t-1}^{(0)'} \right)^{-1}$$

---

[32]Extending it for more general situations is straightforward. Bańbura and Modugno (2014) show how to modify the following formulas in case some of the observations in Δy_t are missing.

and

$$\hat{Q}^{(0)} = \frac{1}{T-1} \left(\left(\sum_{t=2}^{T} F_t^{(0)} F_t^{(0)'} \right) - \hat{\Phi}_1^{(0)} \left(\sum_{t=2}^{T} \hat{F}_{t-1}^{(0)} \hat{F}_t^{(0)'} \right) \right).$$

It is important to stress that this algorithm is more efficient when the scale of all the variables is similar. Hence, although the QML estimates are scale invariant, it is useful to standardise variables beforehand. The scale can be re-attributed accordingly once the likelihood is maximised. Standardisation is also useful for assuring a good initialisation since principal components are not scale invariant.

B.2 Bayesian vector autoregression

In this section we summarise the procedures derived by Giannone, Lenza, and Primiceri (2014). Consider the VAR model of section 2.2:

$$
\begin{aligned}
y_t &= c + A_1 y_{t-1} + \cdots + A_p y_{t-p} + \varepsilon_t, \quad t = 1, \ldots, T, \\
\varepsilon_t &\sim N(0, \Sigma),
\end{aligned}
$$

and rewrite it as

$$
\begin{aligned}
Y &= X\beta + \epsilon, \\
\epsilon &\sim N(0, \Sigma \otimes I_{T-p}),
\end{aligned}
$$

where $y \equiv (y_{p+1}, \ldots, y_T)'$, $Y \equiv vec(y)$, $x_t \equiv \left(1, y_{t-1}', \ldots, y_{t-p}'\right)'$, $x \equiv (x_{p+1}, \ldots, x_T)'$, $X \equiv I_n \otimes x$, $\varepsilon \equiv (\varepsilon_{p+1}, \ldots, \varepsilon_T)'$, $\epsilon \equiv vec(\varepsilon)$, $B \equiv (c, A_1, \ldots, A_p)'$ and $\beta \equiv vec(B)$. Finally, denote the number of regressors for each equation by $k \equiv np + 1$.

For expositional convenience we will focus first on the implementation of the Minnesota prior. Later in the section, we will describe how to implement the sum-of-coefficient and the dummy-initial-observation prior.

The Minnesota prior on (β, Σ) is given by the following normal-inverse-Wishart distribution:

$$
\begin{aligned}
\Sigma | \Psi &\sim IW(\Psi, d), \\
\beta | \Sigma, \Psi, \lambda &\sim N(b, \Sigma \otimes \Omega_{\Psi, \lambda}).
\end{aligned}
$$

The posterior is given by:

$$\Sigma | \Psi, \lambda, Y \quad \sim \quad IW\left(\Psi + \hat{\varepsilon}'\hat{\varepsilon} + \left(\hat{B} - \hat{b}\right)' \Omega_{\Psi,\lambda}^{-1} \left(\hat{B} - \hat{b}\right), T - p + d\right),$$

$$\beta | \Sigma, \Psi, \lambda, Y \quad \sim \quad N\left(\hat{\beta}, \Sigma \otimes \left(x'x + \Omega_{\Psi,\lambda}^{-1}\right)^{-1}\right),$$

where $\hat{B} \equiv \left(x'x + \Omega_{\Psi,\lambda}^{-1}\right)^{-1}\left(x'y + \Omega_{\Psi,\lambda}^{-1}\hat{b}\right)$, $\hat{\beta} \equiv vec\left(\hat{B}\right)$, $\hat{\varepsilon} \equiv y - x\hat{B}$, $\hat{\epsilon} \equiv vec\left(\hat{\varepsilon}\right)$, and $\hat{b}$ is a $k \times n$ matrix obtained by reshaping the vector b in such a way that each column corresponds to the prior mean of the coefficients of each equation (i.e. $b \equiv vec(\hat{b})$).

We follow Giannone, Lenza, and Primiceri (2014) and set an almost flat, but proper, hyperprior. For λ we choose Gamma distribution with mode equal to 0.2 and standard deviation equal to 0.4. Our prior on Ψ is an inverse-Gamma with scale and shape equal to $(0.02)^2$.

The posterior for the hyperparameters is $p(\Psi, \lambda | Y) \propto p(Y | \Psi, \lambda) p(\Psi, \lambda)$, where $p(Y | \Psi, \lambda)$ is the marginal likelihood, which takes the following form (see Giannone, Lenza, and Primiceri, 2014):

$$
\begin{aligned}
p\left(Y | \Psi, \lambda\right) \;=\; & \left(\frac{1}{\pi}\right)^{\frac{n(T-p)}{2}} \frac{\Gamma_n\left(\frac{T-p+d}{2}\right)}{\Gamma_n\left(\frac{d}{2}\right)} \times \\
& \left|\Omega_{\Psi,\lambda}\right|^{-\frac{n}{2}} \cdot \left|\Psi\right|^{\frac{d}{2}} \cdot \left|x'x + \Omega_{\Psi,\lambda}^{-1}\right|^{-\frac{n}{2}} \times \\
& \left|\Psi + \hat{\varepsilon}'\hat{\varepsilon} + \left(\hat{B} - \hat{b}\right)' \Omega_{\Psi,\lambda}^{-1}\left(\hat{B} - \hat{b}\right)\right|^{-\frac{T-p+d}{2}}.
\end{aligned}
$$

Draws from the joint posterior of the parameters and hyperparameters can be easily derived by using the following algorithm. Since the marginal likelihood conditional on the hyperparameters is available in closed form, the hyperparameters can be drawn using the Metropolis-Hastings algorithm. For any draw of the hyperparameters Ψ and λ, the covariance matrix of the residuals Σ and the autoregressive parameters β can be drawn from their distributions conditional on Ψ and λ.

B.2.1 Implementing the "sum-of-coefficients" and the "dummy-initial-observation" priors

The sum-of-coefficient prior is implemented by using the following dummy observations:

$$Y_\mu = \text{diag}(\bar{y}_0)/\mu; \quad X_\mu = [\mathbf{0}_n, \mathbf{1}'_p \otimes \text{diag}(\bar{y}_0)]/\mu,$$

where $\bar{y}_0$ is the average of the first p initial observations, $\mathbf{1}_p$ is a $p \times 1$ vector of ones and $\mathbf{0}_n$ is a $n \times 1$ vector of zeros.

Similarly, the dummy-initial-observation prior is implemented by using the following dummy observations:

$$Y_\delta = \bar{y}_0'/\delta; \quad X_\delta = [1, \mathbf{1}_p' \otimes \bar{y}_0']/\delta.$$

These dummy observations are added to the data and the procedure described above is performed on the augmented data set $Y_{\mu,\delta}^* = \left(Y' \; Y_\mu' \; Y_\delta'\right)'$ and $X_\mu^* = \left(X' \; X_\mu' \; X_\delta'\right)'$. The only correction that has to taken into account concerns the marginal likelihood which should be computed on the original data only. As derived in Giannone, Lenza, and Primiceri (2014), this is equivalent to taking the ratio between the marginal likelihood of the augmented data set relative to the marginal likelihood of the dummy observations:

$$p(Y|\Psi, \lambda, \mu, \delta) = p(Y_{\mu,\delta}^*|\Psi, \lambda)/p(Y_\mu, Y_\delta|\Psi, \lambda).$$

The prior distributions for μ and δ are Gamma distributions with mode and standard deviation equal to 1.

C Description of the data set

No	Short name	Description	Source	Transformation (BVAR in levels)
1	Global GDP	World gross domestic product	AWM	4×log-levels
2	Real GDP	Real gross domestic product, euro area	AWM	4×log-levels
3	Real consumption	Real private consumption, euro area	AWM	4×log-levels
4	Government consumption	Real government consumption, euro area	AWM	4×log-levels
5	Real investment	Real gross investment, euro area	AWM	4×log-levels
6	Real exports	Real exports of goods and services, intra and extra euro area	AWM	4×log-levels
7	Real imports	Real imports of goods and services, intra and extra euro area	AWM	4×log-levels
8	Employment	Total employment (persons), euro area	AWM	4×log-levels
9	Unemployment rate	Unemployment rate (as a ratio to the civilian workforce), euro area	AWM	Raw
10	Economic sentiment	Economic sentiment indicator, survey of the European Commission, euro area	Eurostat	Raw/100
11	Oil price	Price of oil in US dollars	AWM	4×log-levels
12	N.-o. comm. prices	Non-oil commodity prices in US dollars	AWM	4×log-levels
13	HICP	Harmonised index of consumer prices, euro area	AWM	4×log-levels
14	PPI ex. const.	Producer price index, domestic sales, total industry excluding construction, euro area	ECB	4×log-levels
15	GDP deflator	GDP deflator, euro area	AWM	4×log-levels
16	Imports deflator	Imports of goods and services deflator, intra and extra euro area	AWM	4×log-levels
17	Nominal wages	Compensation per employee, euro area	AWM	4×log-levels
18	US short-term i. r.	US short-term interest rate, 3-month dep. LIBOR	IMF (IFS)	Raw/100
19	Short-term i. r.	Short-term interest rate, 3-month EURIBOR	AWM	Raw/100
20	Long-term i. r.	Long-term interest rate, euro area 10-year government benchmark bond yield	AWM	Raw/100
21	M1	Monetary aggregate M1, index of notional stocks, euro area	ECB	4×log-levels
22	M3	Monetary aggregate M3, index of notional stocks, euro area	ECB	4×log-levels
23	Loans to households	Loans to households, sum of consumer loans, loans for house purchases and other loans, index of notional stocks, euro area	ECB	4×log-levels
24	Loans to firms	Loans to non-financial corporations, index of notional stocks, euro area	ECB	4×log-levels
25	Effective exchange rate	Nominal effective exchange rate of the euro (against main 20 trading partners)	AWM	4×log-levels
26	Stock prices	Dow Jones Euro Stoxx price index	DataStream	4×log-levels

Note: In the BVAR in differences and in the DFM specification we take the first difference of the variables transformed as in the BVAR in levels. AWM refers to the 13[th] update of the Area Wide Model data base (Fagan, Henry, and Mestre, 2005).

TABLES AND FIGURES

Table 1: Ratio of MSFE relative to random walk benchmark

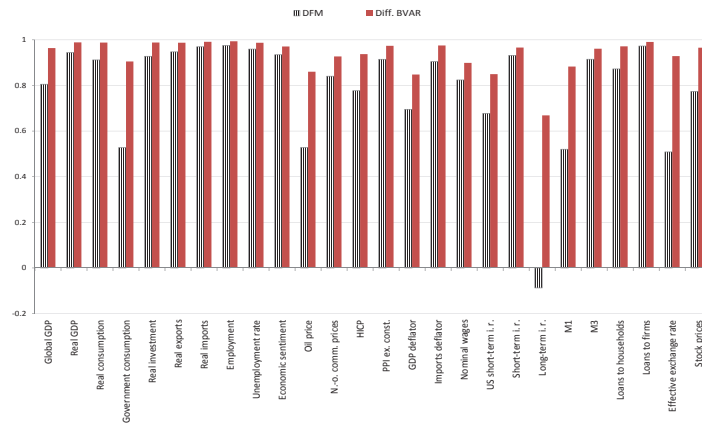| Variables	BVAR in levels			BVAR in differences			Dynamic Factor model		
	H=1	H=2	H=4	H=1	H=2	H=4	H=1	H=2	H=4
Global GDP	**0.61**	**0.67**	**0.76**	**0.65**	**0.69**	**0.69**	**0.60**	**0.62**	**0.63**
Real GDP	**0.59**	**0.78**	**1.00**	**0.53**	**0.73**	**0.88**	**0.56**	**0.71**	**0.84**
Real consumption	**0.39**	**0.21**	**0.27**	**0.41**	**0.28**	**0.36**	**0.55**	**0.41**	**0.64**
Government consumption	**0.54**	**0.43**	**0.31**	**0.59**	**0.44**	**0.37**	**0.85**	**0.85**	**0.91**
Real investment	**0.51**	**0.55**	**0.60**	**0.46**	**0.51**	**0.56**	**0.43**	**0.46**	**0.51**
Real exports	**0.71**	**0.98**	1.28	**0.66**	**0.91**	1.04	**0.72**	**0.99**	1.03
Real imports	**0.51**	**0.63**	**0.73**	**0.46**	**0.59**	**0.67**	**0.54**	**0.65**	**0.73**
Employment	**0.16**	**0.21**	**0.35**	**0.16**	**0.20**	**0.33**	**0.17**	**0.19**	**0.33**
Unemployment rate	**0.39**	**0.50**	**0.72**	**0.35**	**0.46**	**0.66**	**0.30**	**0.40**	**0.61**
Economic sentiment	**0.48**	**0.67**	**0.73**	**0.58**	**0.88**	**0.96**	**0.73**	**0.97**	1.17
Oil price	1.08	1.34	1.52	1.28	1.65	1.79	1.93	2.89	3.44
N.-o. comm. prices	1.08	1.33	1.40	1.16	1.48	1.61	1.23	1.55	1.63
HICP	1.00	1.42	2.54	**0.96**	1.28	1.79	1.55	2.28	3.25
PPI ex. const.	**0.72**	1.04	1.56	**0.69**	1.07	1.67	**0.97**	1.58	2.26
GDP deflator	**0.79**	**0.98**	1.71	1.06	1.17	1.33	1.25	1.38	1.81
Imports deflator	**0.75**	1.09	1.58	**0.74**	1.18	1.84	**0.91**	1.53	2.08
Nominal wages	1.17	1.39	2.24	1.02	**0.98**	1.18	**0.87**	**0.96**	1.44
US short-term i. r.	**0.90**	1.05	1.14	**0.80**	**0.88**	**0.82**	**0.87**	**0.88**	**0.81**
Short-term i. r.	**0.67**	**0.94**	1.46	**0.57**	**0.92**	1.64	**0.72**	1.18	1.75
Long-term i. r.	**0.90**	1.01	1.58	1.05	1.30	1.64	1.08	1.26	1.62
M1	**0.57**	**0.60**	**0.90**	**0.60**	**0.69**	1.23	**0.89**	1.22	1.74
M3	**0.46**	**0.47**	**0.66**	**0.53**	**0.54**	**0.65**	**0.54**	**0.49**	**0.55**
Loans to households	**0.09**	**0.11**	**0.18**	**0.12**	**0.18**	**0.30**	**0.25**	**0.31**	**0.50**
Loans to firms	**0.06**	**0.09**	**0.26**	**0.07**	**0.10**	**0.25**	**0.15**	**0.20**	**0.41**
Effective exchange rate	1.33	1.50	2.13	1.38	1.39	1.59	1.29	1.43	1.76
Stock prices	**0.83**	**0.93**	1.21	**0.88**	1.02	1.05	**0.90**	1.12	1.50

Note: The table reports the ratio of Mean Squared Forecast Errors (MSFE) of the BVAR in levels, BVAR in differences and the DFM over the MSFE of the random walk with drift for the (log-)levels (the model that would prevail if we assumed a dogmatic prior). The ratios are reported for the horizons of one, two and four quarters ahead. Values smaller than one (in bold) indicate that the MSFE of a specific model is lower than the corresponding MSFE of the random walk model.

Figure 1:

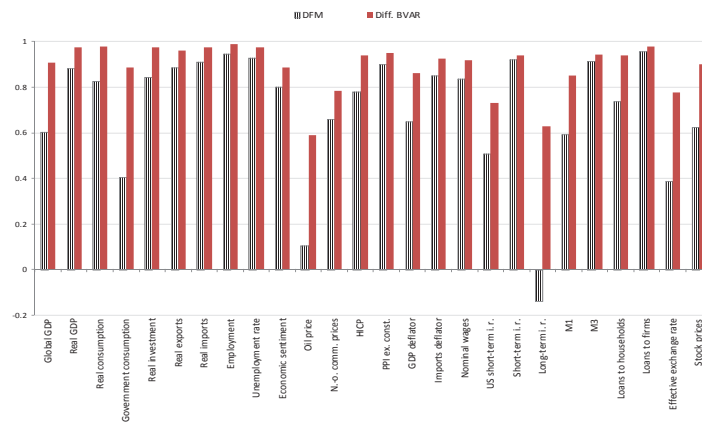Correlation of DFM and BVAR in differences forecasts with BVAR in levels forecasts

(a) One quarter ahead
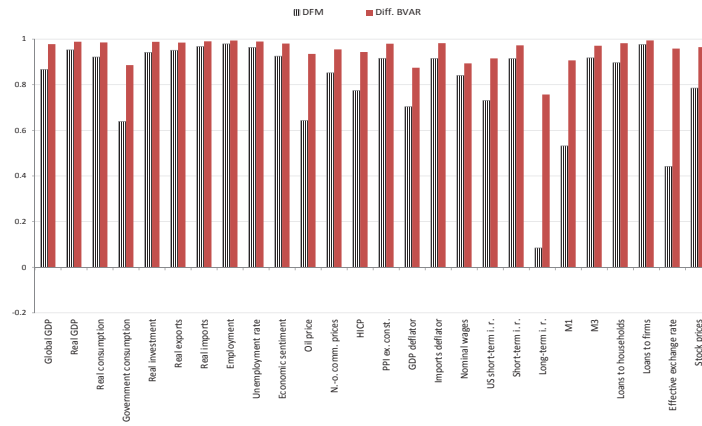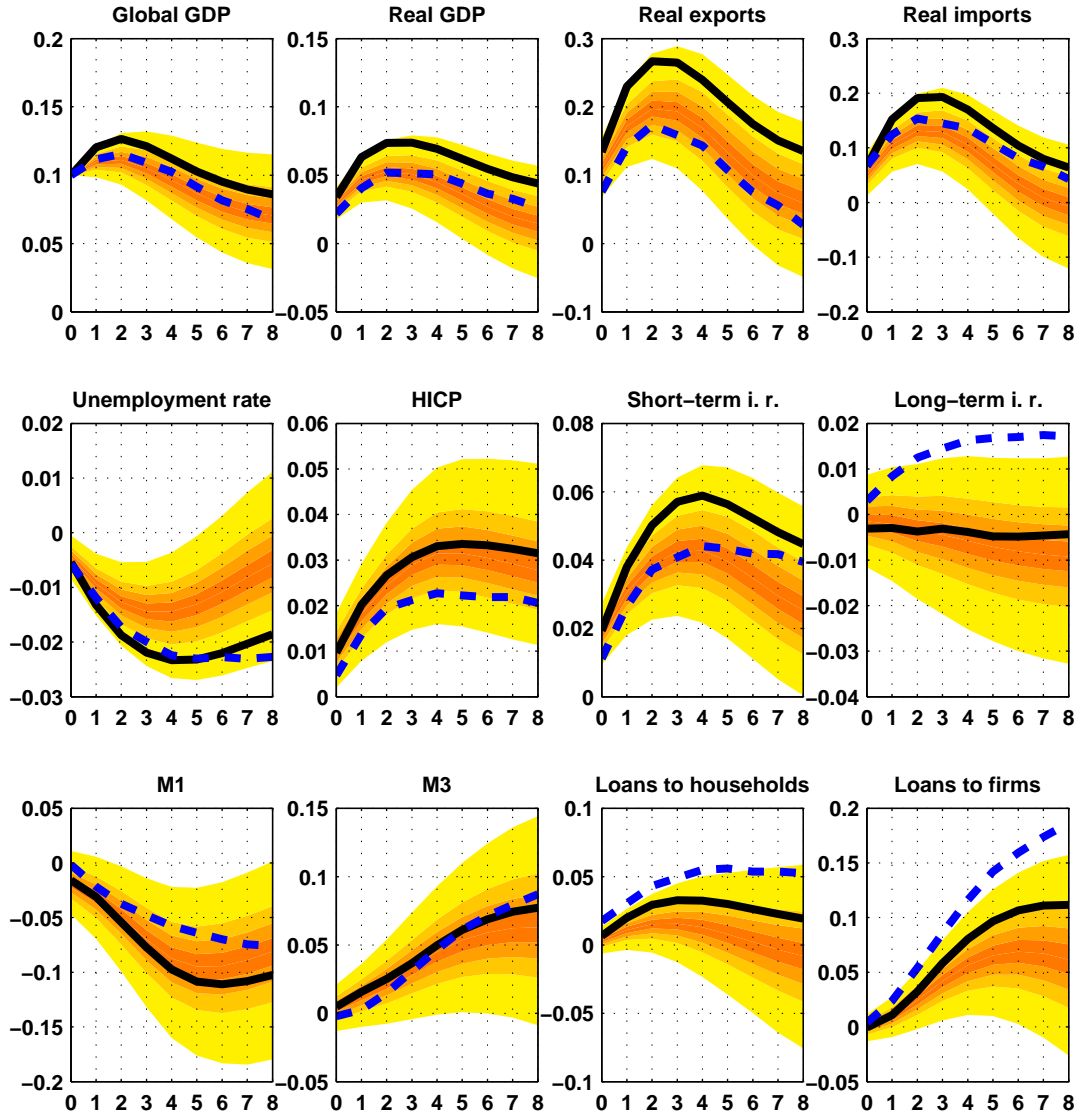

(b) Two quarters ahead


(c) Four quarters ahead

Note: For each variable on the horizontal axis, we report the correlation between the forecasts from the DFM and BVAR in levels (bars with white stripes) and between the forecasts from the BVAR in differences and BVAR in levels (bars with red solid fill).

Figure 2: Scenario analysis: an increase in world GDP

Note: Shades of orange: distribution of the scenario responses in the BVAR in levels, excluding the lower and higher 5% quantiles. Dashed blue line: point estimate of the scenario responses in the DFM model. Solid black line: point estimate of the scenario responses in the BVAR in differences, which is computed as the median of the distribution of the scenario responses in this model. The scenario responses are shown in terms of percentage deviation (of the levels of the variables) in the shocked scenario relative to the baseline scenario (except for the unemployment rate and the interest rates, for which we show deviations).

Figure 3: Conditional Forecasts
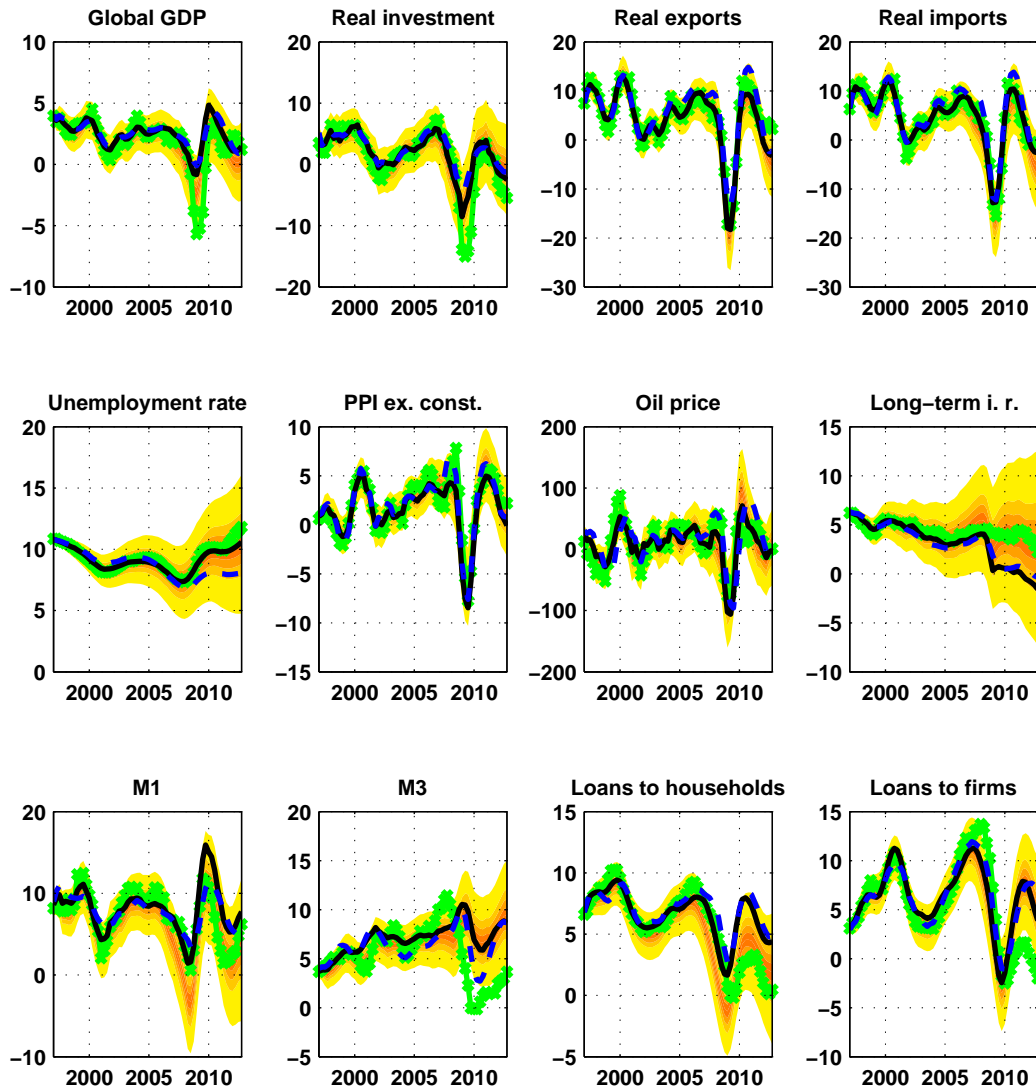
Note: Shades of orange: distribution of the conditional forecasts in the BVAR in levels, excluding the lower and higher 5% quantiles. Dashed blue line: point estimate of the conditional forecasts in the DFM model. Solid black line: point estimate of the conditional forecasts in the BVAR in differences, which is computed as the median of the distribution of the conditional forecasts in this model. Green line with crosses: actual values. The variables are all reported in terms of annual percentage changes, except for the unemployment rate and the long-term interest rate, which are in levels. Compared to Figure 2, we report total investment, producer price index and the oil price in place of real GDP, HICP and the short-term interest rate, which are our conditioning assumptions.

39

References

ADOLFSON, M., S. LASÉEN, J. LINDÉ, AND M. VILLANI (2005): "Are constant interest rate forecasts modest interventions? Evidence from an estimated open economy DSGE model of the euro area," *International Finance*, 8, 509–544.

BAI, J. (2003): "Inferential theory for factor models of large dimensions," *Econometrica*, 71(1), 135–171.

BAI, J., AND S. NG (2002): "Determining the number of factors in approximate factor models," *Econometrica*, 70(1), 191–221.

BAŃBURA, M., D. GIANNONE, AND L. REICHLIN (2010): "Large Bayesian vector auto regressions," *Journal of Applied Econometrics*, 25(1), 71–92.

BAŃBURA, M., AND M. MODUGNO (2014): "Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data," *Journal of Applied Econometrics*, 29(1), 133–160.

BAUMEISTER, C., AND L. KILIAN (2013): "Real-time analysis of oil price risks using forecast scenarios," *IMF Economic Review*, 62(1), 119–145.

BLOOR, C., AND T. MATHESON (2011): "Real-time conditional forecasts with Bayesian VARs: An application to New Zealand," *The North American Journal of Economics and Finance*, 22(1), 26–42.

CARRIERO, A., T. E. CLARK, AND M. MARCELLINO (2012): "Common drifting volatility in large Bayesian VARs," Working Paper 1206, Federal Reserve Bank of Cleveland.

CARTER, C., AND P. KOHN (1994): "On Gibbs sampling for state space models," *Biometrica*, 81, 541–553.

CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (1999): "Monetary policy shocks: What have we learned and to what end?," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1, chap. 2, pp. 65–148. Elsevier.

CHRISTOFFEL, K., G. COENEN, AND A. WARNE (2007): "Conditional versus unconditional forecasting with the new area-wide model of the euro area," Mimeo, European Central Bank.

CLARIDA, R. H., AND D. COYLE (1984): "Conditional projection by means of Kalman filtering," NBER Technical Working Papers 0036, National Bureau of Economic Research, Inc.

CLARK, T. E., AND M. W. MCCRACKEN (2014): "Evaluating conditional forecasts from vector autoregressions," Mimeo, Federal Reserve Bank of St. Louis.

DE JONG, P. (1988): "A cross-validation filter for time series models," *Biometrika*, 75(3), 594–600.

DE JONG, P., AND N. SHEPHARD (1995): "The simulation smoother for time series models," *Biometrika*, 2, 339–350.

DE MOL, C., D. GIANNONE, AND L. REICHLIN (2008): "Forecasting using a large number of predictors: Is Bayesian shrinkage a valid alternative to principal components?," *Journal of Econometrics*, 146(2), 318–328.

DEMPSTER, A., N. LAIRD, AND D. RUBIN (1977): "Maximum likelihood estimation from incomplete data," *Journal of the Royal Statistical Society*, 14, 1–38.

DIRON, M., AND B. MOJON (2005): "Forecasting the central banks inflation objective is a good rule of thumb," Working Paper Series 0564, European Central Bank.

DOAN, T., R. LITTERMAN, AND C. A. SIMS (1984): "Forecasting and conditional projection using realistic prior distributions," *Econometric Reviews*, 3, 1–100.

DOZ, C., D. GIANNONE, AND L. REICHLIN (2011): "A two-step estimator for large approximate dynamic factor models based on Kalman filtering," *Journal of Econometrics*, 164(1), 188–205.

——— (2012): "A quasi-maximum likelihood approach for large, approximate dynamic factor models," *The Review of Economics and Statistics*, 94(4), 1014–1024.

DURBIN, J., AND S. KOOPMAN (2002): "A simple and efficient simulation smoother for state space time series analysis," *Biometrika*, 89(3), 603–615.

DURBIN, J., AND S. J. KOOPMAN (2001): *Time Series Analysis by State Space Methods*. Oxford University Press.

FAGAN, G., J. HENRY, AND R. MESTRE (2005): "An area-wide model for the euro area," *Economic Modelling*, 22(1), 39–59.

FOERSTER, A. T., P.-D. G. SARTE, AND M. W. WATSON (2011): "Sectoral versus aggregate shocks: a structural factor analysis of industrial production," *Journal of Political Economy*, 119(1), 1 – 38.

FORNI, M., M. HALLIN, M. LIPPI, AND L. REICHLIN (2000): "The generalized dynamic factor model: identification and estimation," *Review of Economics and Statistics*, 82, 540–554.

——— (2004): "The generalized dynamic factor model: consistency and rates," *Journal of Econometrics*, 119, 231–245.

GIANNONE, D., M. LENZA, D. MOMFERATOU, AND L. ONORANTE (2014): "Short-term inflation projections: a Bayesian vector autoregressive approach," *International Journal of Forecasting*, 30(3), 635–644.

GIANNONE, D., M. LENZA, H. PILL, AND L. REICHLIN (2012): "The ECB and the interbank market," *Economic Journal*, 122(564), F467–F486.

GIANNONE, D., M. LENZA, AND G. E. PRIMICERI (2014): "Prior selection for vector autoregressions," *Review of Economics and Statistics*, forthcoming.

GIANNONE, D., M. LENZA, AND L. REICHLIN (2010): "Business cycles in the euro area," in *Europe and the Euro*, NBER Chapters, pp. 141–167. National Bureau of Economic Research, Inc.

——— (2012): "Money, credit, monetary policy and the business cycle in the euro area," CEPR Discussion Papers 8944, C.E.P.R. Discussion Papers.

GIANNONE, D., L. REICHLIN, AND L. SALA (2004): "Monetary policy in real time," in *NBER Macroeconomics Annual*, ed. by M. Gertler, and K. Rogoff, pp. 161–200. MIT Press.

GIANNONE, D., L. REICHLIN, AND D. SMALL (2005): "Nowcasting GDP and inflation: the real-time informational content of macroeconomic data releases," Finance and Economics Discussion Series 2005-42, Board of Governors of the Federal Reserve System (U.S.).

GREENE, W. H. (2002): *Econometric Analysis*. Prentice Hall.

HAMILTON, J. D. (1994): *Time Series Analysis*. Princeton University Press.

HARVEY, A. (1989): *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.

JARDET, C., A. MONFORT, AND F. PEGORARO (2012): "New information response functions and applications to monetary policy," Mimeo, Banque de France.

JAROCIŃSKI, M. (2010): "Conditional forecasts and uncertainty about forecast revisions in vector autoregressions," *Economics Letters*, 108(3), 257–259.

JAROCIŃSKI, M., AND F. R. SMETS (2008): "House prices and the stance of monetary policy," *Review*, (Jul), 339–366.

JUNGBACKER, B., AND S. J. KOOPMAN (2008): "Likelihood-based analysis for dynamic factor models," Tinbergen Institute Discussion Papers 08-007/4, Tinbergen Institute.

KARLSSON, S. (2013): "Forecasting with Bayesian vector autoregressions," in *Handbook of Economic Forecasting, Volume 2B*, ed. by G. Elliott, and T. Timmermann. North Holland, Elsevier.

KILIAN, L., AND L. T. LEWIS (2011): "Does the Fed respond to oil price shocks?," *Economic Journal*, 121(555), 1047–1072.

KOOP, G., M. H. PESARAN, AND S. M. POTTER (1996): "Impulse response analysis in nonlinear multivariate models," *Journal of Econometrics*, 74(1), 119–147.

KOOP, G. M. (2013): "Forecasting with medium and large Bayesian VARs," *Journal of Applied Econometrics*, 28(2), 177–203.

KOSE, M. A., C. OTROK, AND C. H. WHITEMAN (2003): "International business cycles: world, region, and country-specific factors," *American Economic Review*, 93, 1216–1239.

LEEPER, E. M., AND T. ZHA (2003): "Modest policy interventions," *Journal of Monetary Economics*, 50(8), 1673–1700.

LENZA, M., H. PILL, AND L. REICHLIN (2010): "Monetary policy in exceptional times," *Economic Policy*, 25, 295–339.

LITTERMAN, R. B. (1979): "Techniques of forecasting using vector autoregressions," Discussion paper.

LUCIANI, M. (2013): "Monetary policy and the housing market: a structural factor analysis," *Journal of Applied Econometrics*, forthcoming.

NG, S. (2013): "Variable selection in predictive regressions," in *Handbook of Economic Forecasting, Volume 2B*, ed. by G. Elliott, and T. Timmermann. North Holland, Elsevier.

PACIELLO, L. (2011): "Does inflation adjust faster to aggregate technology shocks than to monetary policy shocks?," *Journal of Money, Credit and Banking*, 43(8), 1663–1684.

PESARAN, H., AND Y. SHIN (1998): "Generalized impulse response analysis in linear multivariate models," *Economics Letters*, 58(1), 17 – 29.

SHUMWAY, R., AND D. STOFFER (1982): "An approach to time series smoothing and forecasting using the EM algorithm," *Journal of Time Series Analysis*, 3, 253–264.

SIMS, C. A. (1980): "Macroeconomics and reality," *Econometrica*, 48(1), pp. 1–48.

——— (1993): "A nine-variable probabilistic macroeconomic forecasting model," in *Business Cycles, Indicators and Forecasting*, NBER Chapters, pp. 179–212. National Bureau of Economic Research, Inc.

SIMS, C. A., AND T. ZHA (1998): "Bayesian methods for dynamic multivariate models," *International Economic Review*, 39(4), 949–68.

STOCK, J. H., AND M. W. WATSON (2002a): "Forecasting using principal components from a large number of predictors," *Journal of the American Statistical Association*, 97, 147–162.

——— (2002b): "Macroeconomic forecasting using diffusion indexes.," *Journal of Business and Economics Statistics*, 20, 147–162.

——— (2011): "Dynamic factor models," in *The Oxford Handbook of Economic Forecasting*, ed. by M. P. Clements, and D. F. Hendry. Oxford University Press.

———— (2012a): "Disentangling the channels of the 2007-2009 recession," NBER Working Papers 18094, National Bureau of Economic Research, Inc.

———— (2012b): "Generalized shrinkage methods for forecasting using many predictors," *Journal of Business & Economic Statistics*, 30(4), 481–493.

WAGGONER, D. F., AND T. ZHA (1999): "Conditional forecasts in dynamic multivariate models," *The Review of Economics and Statistics*, 81(4), 639–651.

WATSON, M. W., AND R. F. ENGLE (1983): "Alternative algorithms for the estimation of dynamic factor, mimic and varying coefficient regression models," *Journal of Econometrics*, 23, 385–400.