Common Drifting Volatility in Large Bayesian VARs*

Andrea Carriero
Queen Mary, University of London
a.carriero@qmul.ac.uk

Todd E. Clark
Federal Reserve Bank of Cleveland
todd.clark@clev.frb.org

Massimiliano Marcellino
European University Institute, Bocconi University and CEPR
massimiliano.marcellino@eui.eu

This version: March 2, 2012

Abstract

The estimation of large Vector Autoregressions with stochastic volatility using standard methods is computationally very demanding. In this paper we propose to model conditional volatilities as driven by a single common unobserved factor. This is justified by the observation that the pattern of estimated volatilities in empirical analyses is often very similar across variables. Using a combination of a standard natural conjugate prior for the VAR coefficients, and an independent prior on a common stochastic volatility factor, we derive the posterior densities for the parameters of the resulting BVAR with common stochastic volatility (BVAR-CSV). Under the chosen prior the conditional posterior of the VAR coefficients features a Kroneker structure that allows for fast estimation, even in a large system. Using US and UK data, we show that, compared to a model with constant volatilities, our proposed common volatility model significantly improves model fit and forecast accuracy. The gains are comparable to or as great as the gains achieved with a conventional stochastic volatility specification that allows independent volatility processes for each variable. But our common volatility specification greatly speeds computations.

Keywords: Bayesian VARs, stochastic volatility, forecasting, prior specification.

J.E.L. Classification: C11, C13, C33, C53.

*The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System. We would like to thank Haroon Mumtaz and Jonathan Wright for helpful comments on a previous draft. The usual disclaimers apply.

1 Introduction

Several recent papers have shown that two key ingredients for the empirical success of Vector Autoregressions are the use of a rather large information set and the inclusion of drifting volatilities in the model. Banbura, Giannone, and Reichlin (2010), Carriero, Kapetanios, and Marcellino (2011), and Koop (2012) show that a system of 15-20 variables performs better than smaller systems in point forecasting and structural analysis. With small models, studies such as Clark (2011), Cogley and Sargent (2005), and Primiceri (2005) show how the inclusion of drifting volatility is key for understanding the dynamics of macroeconomic variables and for density forecasting. Koop and Korobilis (2012) show that a computational shortcut for allowing time-varying volatility (roughly speaking, using a form of exponential smoothing of volatility) improves the accuracy of point and density forecasts from larger VARs.

However, introducing stochastic volatility within a Vector Autoregressions poses serious computational burdens, and typically all the empirical implementations of such models have been limited to a handful of variables (3 to 5). The computational burden is driven by the use of Markov Chain Monte Carlo (MCMC) estimation methods needed to accommodate stochastic volatility (the same applies to Bayesian estimation of other models of time-varying volatilities, including Markov Switching and GARCH). In particular, as noted in such studies as Sims and Zha (1998), the challenge with larger VAR models is that drawing the VAR coefficients from the conditional posterior involves computing a (variance) matrix with the number of rows and columns equal to the number of variables squared times the number of lags (plus one if a constant is included). The size of this matrix increases with the square of the number of variables in the model, making CPU time requirements highly nonlinear in the number of variables.

In this paper we propose a computationally effective way to model stochastic volatility, to greatly speed up computations for smaller VAR models and make estimation tractable for larger models. The proposed method hinges on the observation that the pattern of estimated volatilities in empirical analyses is often very similar across variables. We propose to model conditional volatilities as driven by a single common unobserved factor. Our volatility model corresponds to the stochastic discount factor model described in Jacquier, Polson, and Rossi (1995). While Jacquier, Polson, and Rossi (1995) had in mind using the model in an asset return context, we incorporate the volatility model in a VAR. Using a combination of (1) a standard natural conjugate prior for the VAR coefficients and (2) an independent prior on a common stochastic volatility factor, we derive the posterior densities for the parameters of the resulting BVAR with common stochastic volatility (BVAR-CSV). Under the chosen prior the conditional posterior of the VAR coefficients features a Kroneker structure that allows for fast estimation. Hence, the BVAR-CSV can be also estimated with a larger set of endogenous variables.

Our proposed volatility model treats the commonality as multiplicative. We need both the single factor and the multiplicative structure in order to be able to define a prior and factor out volatility in such a way as to exploit the Kroneker structure that is needed to speed up the VAR computations. Prior work by Pajor (2006) considered the same basic model of volatility for the errors of a VAR(1) process, in just a few variables, without the VAR prior we incorporate to speed up computations. Still other work in such studies as Osiewalski and Pajor (2009) and references therein has considered common volatility within GARCH-type specifications. Some other papers introduce the commonality in volatility as additive. For example, in an asset return context, Chib, et al. (2002, 2006) and Jacquier, et al. (1995) employ a factor structure multivariate stochastic volatility model. In a macro context, in a setup similar to that used in some finance research, Del Negro and Otrok (2008) develop a factor model with stochastic volatility. Viewed this way, the factor structure multivariate stochastic volatility model or factor model with stochastic volatility is somewhat different from the one proposed here: in the BVAR-CSV we have a VAR that captures cross-variable correlations in conditional means and captures a common factor in just volatility; in these other models, the factor captures both cross-variable correlations in conditional means and drives commonality in volatility.

To establish the value of our proposed model, we compare CPU time requirements, volatility estimates, and forecast accuracy (both point and density) across VAR models of different sizes and specifications. The model specifications include: a VAR with constant volatilities; a VAR with stochastic volatility that treats the volatilities of each variable as independent, as pioneered in Cogley and Sargent (2005) and Primiceri (2005); and our proposed VAR with common stochastic volatility. More specifically, using VARs for US data, we first document the efficiency gains associated with imposing common volatility. We then compare alternative estimates of volatility, for both 4-variable and 8-variable systems, and show that there is substantial evidence of common volatility. We then proceed to examine real-time forecasts from 4-variable and 8-variable macroeconomic models for the US, finding that the imposition of common stochastic volatility consistently improves the accuracy of real-time point forecasts (RMSEs) and density forecasts (log predictive scores). We also compare final-vintage forecasts from 15-variable models for the US data and again find that common stochastic volatility improves forecast accuracy.

Finally, as a robustness check, we repeat much of the analysis using UK data, obtaining broadly similar results. Most notably, despite evidence of more heterogeneity in the volatility patterns across variables for the UK than for the US, we find the BVAR with common stochastic volatility significantly improves the accuracy of forecasts.[1] Actually, the gains are comparable to those for the US when using a BVAR as the benchmark, and are even larger with a simple AR model for each variable as the benchmark. Furthermore, the gains apply to both point and density forecasts.

We interpret these results as evidence that the BVAR-CSV model efficiently summarizes the information in a rather large dataset and successfully accounts for changing volatility, in a way that is much more computationally efficient than in the conventional approach that treats the volatility of each variable as independent.

The structure of the paper is as follows. Section 2 presents the model, discusses the priors, derives the posteriors (with additional details in the Appendix), and briefly describes the other BVAR models to which we compare the results from our proposed BVAR-CSV model. Section 3 discusses the MCMC implementation. Section 4 presents our US-based evidence, including computational time for the estimates of alternative models and full-sample volatility estimates and presents the forecasting exercise for the 4-, 8- and 15-variable BVAR-CSV. Section 5 examines the robustness of our key findings using data for the UK. Section 6 summarizes the main results and concludes.

2 The BVAR-CSV model

2.1 Model Specification

Let y_t denote the $n \times 1$ vector of model variables and p the number of lags. Define the following: $\Pi_0 =$ an $n \times 1$ vector of intercepts; $\Pi(L) = \Pi_1 - \Pi_2 L - \cdots - \Pi_p L^{p-1}$; $A =$ a lower triangular matrix with ones on the diagonal and coefficients a_{ij} in row i and column j (for $i = 2, \ldots, n$, $j = 1, \ldots, i-1$), where $a_i, i = 2, \ldots, n$ denotes the vector of coefficients in row i; and $S = \text{diag}(1, s_2, \ldots, s_n)$.

The VAR(p) with common stochastic volatility takes the form

$$
\begin{aligned}
y_t &= \Pi_0 + \Pi(L) y_{t-1} + v_t, &(1)\\
v_t &= \lambda_t^{0.5} A^{-1} S^{1/2} \epsilon_t, \ \ \epsilon_t \sim N(0, I_n), &(2)\\
\log(\lambda_t) &= \log(\lambda_{t-1}) + \nu_t, \ \ \nu_t \sim \text{iid } N(0, \phi). &(3)
\end{aligned}
$$

[1] As detailed below, in light of the more limited availability of real-time data for the UK than the US, our UK results are based on final vintage data, not real-time data.

As is standard in macroeconomic VARs with stochastic volatility, the log variance λ_t follows a random walk process, with innovations having a variance of ϕ. Here, there is a single volatility process that is common to all variables, and drives the time variation in the entire variance covariance matrix of the VAR errors. As we will see, empirically this assumption yields sizable forecasting gains with respect to a specification with constant volatility. Moreover, it leads to major computational gains with respect to a model with n independent stochastic volatilities, with in general no major losses and often gains in forecasting accuracy. The scaling matrix S allows the variances of each variable to differ by a factor that is constant over time. The setup of S reflects an identifying normalization that the first variable's loading on common volatility is 1. Similarly, the matrix A rescales the covariances.

Under the above specification, the residual variance–covariance matrix for period t is $\text{var}(v_t) = \Sigma_t \equiv \lambda_t A^{-1} S A^{-1\prime}$. To simplify some notation, let $\tilde{A} = S^{-1/2}A$. Then the inverse of the reduced-form variance-covariance matrix simplifies to:

$$V_t^{-1} = \frac{1}{\lambda_t}\tilde{A}'\tilde{A}. \tag{4}$$

2.2 Priors

The parameters of the model consist of the following: $\Pi \equiv k \times n$ matrix of coefficients contained in $(\Pi_0,\ \Pi_1,\ ..., \Pi_p)$; A (non-zero and non-unit elements), composed of vectors a_i, $i = 2, \ldots, n$; s_i, $i = 2, \ldots, n$; ϕ; and λ_0. The model also includes the latent states λ_t, $t = 1, \ldots, T$. Below, we use Λ to refer to the history of variances from 1 to T.

We use $N(a, b)$ to denote a normal distribution (either univariate or multivariate) with mean a and variance b. We use $IG(a, b)$ to denote an inverse gamma distribution with scale term a and degrees of freedom b.

We specify priors for the parameter blocks of the model, as follows (implementation details are given below).

$$\begin{aligned}
\text{vec}(\Pi)|A, S &\sim N(\text{vec}(\underline{\mu}_\Pi), \underline{\Omega}_\Pi) && (5)\\
a_i &\sim N(\underline{\mu}_{a,i}, \underline{\Omega}_{a,i}),\ i = 2, \ldots, n && (6)\\
s_i &\sim IG(d_s \cdot \underline{s}_i, d_s),\ i = 2, \ldots, n && (7)\\
\phi &\sim IG(d_\phi \cdot \underline{\phi}, d_\phi) && (8)\\
\log \lambda_0 &\sim N(\underline{\mu}_\lambda, \underline{\Omega}_\lambda) && (9)
\end{aligned}$$

To make estimation with large models tractable, the prior variance for vec(Π) needs to be specified with a factorization that permits a Kroneker structure. To be able to exploit a Kroneker structure and achieve computational gains, we need not only a single common, multiplicative volatility factor but also a prior that permits factorization. Specifically, we use a prior conditional on $\tilde{A} = S^{-1/2}A$, of the following form:

$$\underline{\Omega}_\Pi = (\tilde{A}'\tilde{A})^{-1} \otimes \underline{\Omega}_0, \tag{10}$$

where $\underline{\Omega}_0$ incorporates the kind of symmetric coefficient shrinkage typical of the natural conjugate Normal-Wishart prior. Under the usual Minnesota-style specification of the Normal-Wishart prior for $\underline{\Omega}_0$, the prior variance takes account of volatility (and relative volatilities of different variables) by using variance estimates from some training sample. Note that the use of a prior for the coefficients conditional on volatility is in alignment with the natural conjugate Normal-Wishart prior, but it does depart from the setup of Clark (2011) and Clark and Davig (2011), in which, for a VAR with independent stochastic volatilities, the coefficient prior was unconditional.

The prior used here, combined with the assumption of a single volatility factor, implies that the posterior distribution of the VAR coefficients, conditional on $\tilde{A}$ and Λ, will have a variance featuring a Kroneker structure. As a result the computations required to draw from such a distribution via MC sampling are of order $n^3 + k^3$ rather than of order n^3k^3.[2] While such advantage can be considered minor with a small system, it becomes crucial in estimating larger VARs.

2.3 Coefficient posteriors

The parameters Π, a_i, s_i, and ϕ have closed form conditional posterior distributions which we present here. Draws from these conditionals will constitute Gibbs sampling steps in our MCMC algorithm. Drawing from the process λ_t instead will involve a Metropolis step and is discussed below. We define some additional notation incorporated in the computation of certain moments:

$$v_t = y_t - \Pi_0 - \Pi(L)y_{t-1}, \tag{11}$$

$$\tilde{v}_t = Av_t, \tag{12}$$

$$\nu_t = \log(\lambda_t) - \log(\lambda_{t-1}), \tag{13}$$

[2] Direct inversion of $\underline{\Omega}_\Pi$ would require n^3k^3 elementary operations (using Gaussian elimination). If instead $\underline{\Omega}_\Pi$ has a Kronecker structure, then its inverse can be obtained by inverting $(\tilde{A}'\tilde{A})^{-1}$ and $\underline{\Omega}_0$ separately. As these matrices are of dimension n and k respectively, their inversion requires $n^3 + k^3$ elementary operations (plus the operations necessary to compute the Kronecker product, which being of order n^2k^2 are negligible).

and:

$$w_t = n^{-1}\tilde{v}_t' S^{-1}\tilde{v}_t. \tag{14}$$

In the Appendix we show that the conditional posterior distributions of Π, a_i, s_i, and ϕ take the following forms:

$$\begin{aligned}
\text{vec}(\Pi)|A, S, \phi, \Lambda, y &\sim N(\text{vec}(\bar{\mu}_\Pi), \bar{\Omega}_\Pi) & (15)\\
a_i|\Pi, S, \phi, \Lambda, y &\sim N(\bar{\mu}_{a,i}, \bar{\Omega}_{a,i}),\ i = 2, \ldots, n & (16)\\
s_i|\Pi, A, \phi, \Lambda, y &\sim IG\left(d_s \cdot \underline{s}_i + \sum_{t=1}^{T}(\tilde{v}_{i,t}^2/\lambda_t),\ d_s + T\right),\ i = 2, \ldots, n & (17)\\
\phi|\Pi, A, S, \Lambda, y &\sim IG\left(d_\phi \cdot \underline{\phi} + \sum_{t=1}^{T}\nu_t^2,\ d_\phi + T\right), & (18)
\end{aligned}$$

where y is a nT-dimensional vector containing all the data.

The mean and variance of the conditional posterior normal distribution for vec(Π) take the following forms:

$$\begin{aligned}
\text{vec}(\bar{\mu}_\Pi) &= \bar{\Omega}_\Pi\left\{\text{vec}\left(\sum_{t=1}^{T} X_t y_t' \Sigma_t^{-1}\right) + \underline{\Omega}_\Pi^{-1}\text{vec}(\underline{\mu}_\Pi)\right\} & (19)\\
\bar{\Omega}_\Pi &= \left(\tilde{A}'\tilde{A}\right)^{-1} \otimes \left(\underline{\Omega}_0^{-1} + \sum_{t=1}^{T}(\frac{1}{\lambda_t}X_t X_t')\right)^{-1}. & (20)
\end{aligned}$$

Again, the key to the computational advantage of this model is the Kroneker structure of the conditional posterior variance. Achieving this Kroneker structure requires both a single, multiplicative volatility factor and the conditional prior described above.

In practice, the posterior mean of the coefficient matrix can be written in an equivalent form that may often be more computationally efficient. This equivalent form is obtained by defining data vectors normalized by the standard deviation of volatility, to permit rewriting the VAR in terms of conditionally homoskedastic variables: specifically, let $\tilde{y}_t = \lambda_t^{-0.5} y_t$ and $\tilde{X}_t = \lambda_t^{-0.5} X_t$. Then, the posterior mean of the matrix of coefficients can be equivalently written as

$$\bar{\mu}_\Pi = \left(\sum_{t=1}^{T}\tilde{X}_t\tilde{X}_t' + \underline{\Omega}_0^{-1}\right)^{-1}\left(\underline{\Omega}_0^{-1}\underline{\mu}_\Pi + \sum_{t=1}^{T}\tilde{y}_t\tilde{X}_t'\right), \tag{21}$$

or, using full-data matrices,

$$\bar{\mu}_\Pi = \left(\tilde{X}'\tilde{X} + \underline{\Omega}_0^{-1}\right)^{-1}\left(\underline{\Omega}_0^{-1}\underline{\mu}_\Pi + \tilde{X}'\tilde{y}\right). \tag{22}$$

As detailed in Cogley and Sargent (2005), the mean and variance of the posterior normal distribution for the rows of A are obtained from moments associated with regressions, for

$i = 2, \ldots, n$, of $v_{i,t}/(s_i\lambda_t)^{0.5}$, on $v_{j,t}/(s_i\lambda_t)^{0.5}$, where $j = 1, \ldots, i-1$. Treating each equation i separately, let $Z_i'Z_i$ denote the second moment matrix of the variables on the right-hand side of the regression, and $Z_i'z_i$ denote the product of the right-hand side with the dependent variable. Then, for each i, the posterior mean and variance of the normal distribution are as follows:

$$\bar{\mu}_{a,i} = \bar{\Omega}_{a,i}(Z_i'z_i + \underline{\Omega}_{a,i}^{-1}\underline{\mu}_{a,i}) \quad (23)$$

$$\bar{\Omega}_{a,i} = (Z_i'Z_i + \underline{\Omega}_{a,i}^{-1})^{-1}. \quad (24)$$

2.4 Volatility

Our treatment of volatility follows the approach of Cogley and Sargent (2005), in a univariate setting, based on Jacquier, Polson, and Rossi (1994). Exploiting the Markov property of the volatility process one can write:

$$f(\lambda_t|\lambda_{-t}, u^T, \phi, y) = f(\lambda_t|\lambda_{t-1}, \lambda_{t+1}, u^T, \phi), \quad (25)$$

where λ_{-t} denotes the volatilities at all dates but t and u^T denotes the full history of $u_t = AS^{-1/2}\epsilon_t$. Jacquier, Polson, and Rossi (1994) derive the conditional posterior kernel for this process:

$$f(\lambda_t|\lambda_{t-1}, \lambda_{t+1}, u^T, \phi, y) \sim \lambda_t^{-1.5} \exp\left(\frac{-w_t}{2\lambda_t}\right) \exp\left(\frac{-(\log\lambda_t - \mu_t)}{2\sigma_c^2}\right), \quad (26)$$

where the parameters μ_t and σ_c^2 are the conditional mean and variance of $\log\lambda_t$ given λ_{t-1} and λ_{t+1}. With the random walk process, for periods 2 through $T-1$, the conditional mean and variance are $\mu_t = (\log\lambda_{t-1} + \log\lambda_{t+1})/2$ and $\sigma_c^2 = \phi/2$, respectively (the conditional mean and variance are a bit different for periods 1 and T). Draws from the process λ_t need to be simulated using a Metropolis step, spelled out in Cogley and Sargent (2005).

2.5 Other models for comparison

To establish the merits of our proposed model, we will consider estimates from a VAR with independent stochastic volatilities for each variable (denoted BVAR-SV) and a VAR with constant volatilities (denoted BVAR).

The BVAR-SV model takes the form

$$y_t = \Pi_0 + \Pi(L)y_{t-1} + v_t, \quad (27)$$

$$v_t = A^{-1}\Lambda_t^{0.5}\epsilon_t,\ \epsilon_t \sim N(0, I_n),\ \Lambda_t = \mathrm{diag}(\lambda_{1,t}, \ldots, \lambda_{n,t}), \quad (28)$$

$$\log(\lambda_{i,t}) = \log(\lambda_{i,t-1}) + \nu_{i,t},\ \nu_{i,t} \sim N(0, \phi_i),\ i = 1, n,$$

With this model, the residual variance–covariance matrix for period t is $\text{var}(v_t) \equiv \Sigma_t = A^{-1}\Lambda_t A^{-1\prime}$.

In the interest of brevity, we don't spell out all of the priors and posteriors for the model. However, as detailed in Clark (2011) and Clark and Davig (2011), the prior for the VAR coefficients is unconditional, rather than conditional as in the BVAR-CSV. From a computational perspective, the key difference between the BVAR-SV and BVAR-CSV models is that the posterior variance for the (VAR) coefficients of the BVAR-SV model does not have the overall Kroneker structure of the posterior variance for the coefficients of the BVAR-CSV model (given in equation (20)). For the BVAR-SV specification, the posterior mean (the vector of coefficients) and variance are:

$$\text{vec}(\bar{\mu}_\Pi) = \bar{\Omega}_\Pi \left\{ \text{vec}\left(\sum_{t=1}^{T} X_t y_t' \Sigma_t^{-1} \right) + \underline{\Omega}_\Pi^{-1} \text{vec}(\underline{\mu}_\Pi) \right\} \tag{29}$$

$$\bar{\Omega}_\Pi^{-1} = \underline{\Omega}_\Pi^{-1} + \sum_{t=1}^{T} (\Sigma_t^{-1} \otimes X_t X_t'). \tag{30}$$

The BVAR takes the form

$$y_t = \Pi_0 + \Pi(L) y_{t-1} + v_t, \; v_t \sim N(0, \Sigma). \tag{31}$$

For this model, we use the Normal-diffuse prior and posterior detailed in such studies as Kadiyala and Karlsson (1997).

3 Implementation

3.1 Specifics on priors: BVAR-CSV model

For our proposed BVAR-CSV model, we set the prior moments of the VAR coefficients along the lines of the common Minnesota prior, without cross-variable shrinkage:

$$\underline{\mu}_\Pi = 0, \text{ such that } E[\Pi_l^{(ij)}] = 0 \; \forall \; i, j, l \tag{32}$$

$$\underline{\Omega}_0 \text{ such that the entry corresponding to } \Pi_l^{(ij)} = \begin{cases} \frac{\theta^2}{l^2} \frac{\sigma_1^2}{\sigma_j^2} & for \; l > 0 \\ \varepsilon^2 \sigma_1^2 & for \; l = 0 \end{cases}. \tag{33}$$

With all of the variables of our VAR models transformed for stationarity (in particular, we use growth rates of GDP, the price level, etc.), we set the prior mean of all the VAR coefficients to 0.[3] The variance matrix $\underline{\Omega}_0$ is defined to be consistent with the usual Minnesota

[3]Our proposed BVAR-CSV specification can also be directly applied to models in levels with unit root priors, with the appropriate modification of the prior means on the coefficients. Including priors on sums

prior variance, which is a diagonal matrix. Note that σ_1^2, the prior variance associated with innovations to equation 1, enters as it does to reflect the normalization of S, in which all variances are normalized by σ_1^2. With a bit of algebra, omitted for brevity, by plugging in $A = I_n$ and $S_{ii} = \sigma_i^2/\sigma_1^2$, the prior $\underline{\Omega}_\Pi = (\tilde{A}'\tilde{A})^{-1} \otimes \underline{\Omega}_0$ can be shown to equal the conventional Minnesota prior given below for the BVAR-SV model. The shrinkage parameter θ measures the tightness of the prior: when $\theta \to 0$ the prior is imposed exactly and the data do not influence the estimates, while as $\theta \to \infty$ the prior becomes loose and results will approach standard GLS estimates. We set $\theta = 0.2$ and $\varepsilon = 1000$. The term $1/l^2$ determines the rate at which the prior variance decreases with increasing lag length. To set the scale parameters σ_i^2 we follow common practice (see e.g. Litterman, 1986; Sims and Zha, 1998) and fix them to the variance of the residuals from a univariate AR(4) model for the variables, computed for the estimation sample.

Following Cogley and Sargent (2005), we use an uninformative prior for the elements in the matrix A:

$$\underline{\mu}_{a,i} = 0, \ \underline{\Omega}_{a,i} = 1000^2 \cdot I_{i-1}. \tag{34}$$

In line with other studies such as Cogley and Sargent (2005), we make the priors on the volatility-related parameters loosely informative. Specifically, the prior scale and shape parameters for the elements $\underline{s}_i$ in S and for ϕ are:

$$\underline{s}_i = \hat{s}_{i,OLS}, \ d_s = 3, \tag{35}$$

$$\underline{\phi} = 0.035, \ d_\phi = 3. \tag{36}$$

Finally the prior moments for the initial value of the volatility process are:

$$\underline{\mu}_\lambda = \log \hat{\lambda}_{0,OLS}, \ \underline{\Omega}_\lambda = 4. \tag{37}$$

In the prior for S, the mean $\hat{s}_{i,OLS}$ is set on the basis of residual variances obtained from AR models fit with the estimation sample (in line with common practice). For each variable, we estimate an AR(4) model. For each $j = 2, \ldots, n$, we regress the residual from the AR model for j on the residuals associated with variables 1 through $j-1$ and compute the error variance (this step serves to filter out covariance as reflected in the A matrix). Letting $\hat{\sigma}_{i,0}^2$ denote these error variances, we set the prior mean on the relative volatilities at $\hat{s}_{i,OLS} = \hat{\sigma}_{i,0}^2/\hat{\sigma}_{1,0}^2$ for $i = 2, \ldots, n$. In the prior for log volatility in period 0, we follow the same steps in obtaining residual variances $\hat{\sigma}_{i,0}^2$, but with data from a training sample of the

of coefficients and initial observations as in such studies as Sims and Zha (1998) is also possible, subject to appropriate adjustment for the conditional heteroskedasticity of y_t and X_t.

40 observations preceding the estimation sample.[4] We set the prior mean of log volatility in period 0 at $\log \hat{\lambda}_{0,OLS} = \log(n^{-1} \sum_{i=1}^{n} \hat{\sigma}_{i,0}^2)$.

3.2 Specifics on priors: BVAR-SV and BVAR models

For the BVAR-SV model, we use a conventional Minnesota prior, without cross-variable shrinkage:

$$\bar{\mu}_\Pi \text{ such that } E[\Pi_l^{(ij)}] = 0 \; \forall \; i, j, l \tag{38}$$

$$\underline{\Omega}_\Pi \text{ such that } V[\Pi_l^{(ij)}] = \begin{cases} \frac{\theta^2}{l^2} \frac{\sigma_i^2}{\sigma_j^{-2}} & for \; l > 0 \\ \varepsilon^2 \sigma_i^2 & for \; l = 0 \end{cases}. \tag{39}$$

Consistent with our prior for the BVAR-CSV model, we set $\theta = 0.2$ and $\varepsilon = 1000$, and we set the scale parameters σ_i^2 at estimates of residual variances from AR(4) models from the estimation sample.

In the prior for the volatility-related components of the model, we follow an approach similar to that for the BVAR-CSV model. Broadly, our approach to setting volatility-related priors is similar to that used in such studies as Clark (2011), Cogley and Sargent (2005), and Primiceri (2005). The prior for A is uninformative, as described above. For the prior on each ϕ_i, we use a mean of 0.035 and 3 degrees of freedom. For the initial value of the volatility of each equation i, we use

$$\underline{\mu}_{\lambda,i} = \log \hat{\lambda}_{i,0,OLS}, \; \underline{\Omega}_\lambda = 4. \tag{40}$$

To obtain $\log \hat{\lambda}_{i,0,OLS}$, we use a training sample of 40 observations preceding the estimation sample to fit AR(4) models for each variable and, for each $j = 2, \ldots, n$, we regress the residual from the AR model for j on the residuals associated with variables 1 through $j-1$ and compute the error variance (this step serves to filter out covariance as reflected in the A matrix). Letting $\hat{\sigma}_{i,0}^2$ denote these error variances, we set the prior mean of log volatility in period 0 at $\log \hat{\lambda}_{i,0,OLS} = \log \hat{\sigma}_{i,0}^2$.[5]

[4] In the real-time forecasting analysis, for the vintages in which a training sample of 40 observations is not available, the prior is set using the training sample estimates available from the most recent vintage with 40 training sample observations.

[5] In the real-time forecasting analysis, for the vintages in which a training sample of 40 observations is not available, the prior is set using the training sample estimates available from the most recent vintage with 40 training sample observations.

3.3 MCMC Algorithm

We estimate the BVAR-CSV model with a five-step Metropolis-within-Gibbs MCMC algorithm.[6] The Metropolis step is used for volatility estimation, following Cogley and Sargent (2005), among others. The other steps rely on Gibbs samplers. In order to facilitate the description of some of the steps, we rewrite the VAR as in Cogley and Sargent (2005) and Primiceri (2005):

$$A(y_t - \Pi_0 - \Pi(L)y_{t-1}) \equiv \tilde{v}_t = \lambda_t^{0.5} S^{1/2} \epsilon_t. \tag{41}$$

Step 1: Draw the matrix of VAR coefficients Π conditional on A, S, ϕ, and Λ, using the conditional (normal) distribution for the posterior given in equation (15).

Step 2: Draw the coefficients in A conditional on Π, S, ϕ, and Λ, using the conditional (normal) distribution for the posterior given in (16). This step follows the approach detailed in Cogley and Sargent (2005), except that, in our model, the VAR coefficients Π are constant over time.

Step 3: Draw the elements of S conditional on Π, A, ϕ, and Λ, using the conditional (IG) distribution for the posterior given above in (17)

Using equation (41), for each equation $i = 2, \ldots, n$, we have that $\tilde{v}_{i,t}/\lambda_t^{0.5} = s_i^{1/2}\epsilon_{i,t}$. We can then draw s_i using a posterior distribution that incorporates information from the sample variance of $\tilde{v}_{i,t}/\lambda_t^{0.5}$.

Step 4: Draw the time series of volatility λ_t conditional on Π, A, S, and ϕ, using a Metropolis step. From equation (41) it follows that

$$w_t = n^{-1}\tilde{v}_t' S^{-1} \tilde{v}_t = n^{-1}\lambda_t \epsilon_t' \epsilon_t. \tag{42}$$

Taking the log yields

$$\log w_t^2 = \log \lambda_t + \log(n^{-1}\epsilon_t'\epsilon_t). \tag{43}$$

As suggested in Jacquier, Polson, and Rossi (1995), the estimation of the time series of λ_t can proceed as in the univariate approach of Jacquier, Polson, and Rossi (1994). Our particular implementation of the algorithm is taken from Cogley and Sargent (2005).

Step 5: Draw the variance ϕ, conditional on Π, A, S, and Λ, using the conditional (IG) distribution for the posterior given in (18)

We estimate the BVAR-SV model with a similar algorithm, modified to drop the step for sampling S and to draw time series of volatilities of all variables, not just common

[6]While not detailed in the interest of brevity, we follow Cogley and Sargent (2005) in including in the algorithm checks for explosive autoregressive draws, rejecting explosive draws and re-drawing to achieve a stable draw.

volatility. We estimate the BVAR with a simple Gibbs sampling algorithm, corresponding to the Normal-diffuse algorithm described in Kadiyala and Karlsson (1997).

In all cases, we obtain forecast distributions by sampling as appropriate from the posterior distribution. For example, in the case of the BVAR-CSV model, for each set of draws of parameters, we: (1) simulate volatility time paths over the forecast interval using the random walk structure of log volatility; (2) draw shocks to each variable over the forecast interval with variances equal to the draw of Σ_{t+h}; and (3) use the VAR structure of the model to obtain paths of each variable. We form point forecasts as means of the draws of simulated forecasts and density forecasts from the simulated distribution of forecasts. Conditional on the model, the posterior distribution reflects all sources of uncertainty (latent states, parameters, and shocks over forecast interval).

4 Empirical results with US data

4.1 Data and design of the forecast exercise

In most of our analysis, we consider models of a maximum of eight variables, at the quarterly frequency: growth of output, growth of personal consumption expenditures (PCE), growth of business fixed investment (in equipment, software, and structures, denoted BFI), growth of payroll employment, the unemployment rate, inflation, the 10-year Treasury bond yield, and the federal funds rate. This particular set of variables was chosen in part on the basis of the availability of real-time data for forecast evaluation. Consistent with such studies as Clark (2011), we also consider a four-variable model, in output, unemployment, inflation, and the funds rate. We also examine forecasts from a 15-variable model, similar to the medium-sized model of Koop (2012), using his data.

For the 4- and 8-variable models, we consider both full-sample estimates and real-time estimates. Our full-sample estimates are based on current vintage data taken from the FAME database of the Federal Reserve Board. The quarterly data on unemployment and the interest rates are constructed as simple within-quarter averages of the source monthly data (in keeping with the practice of, e.g., Blue Chip and the Federal Reserve). Growth and inflation rates are measured as annualized log changes (from $t-1$ to t).

For the 15-variable model, we report only forecasts based on current vintage data, using data from Koop (2012). The set of variables is listed in Tables 9 and 10 (please see Koop's paper for additional details). The data are transformed as detailed in Koop (2012). The forecast evaluation period runs from 1985:Q1 through 2008:Q4, and the forecasting models are estimated using data starting in 1965:Q1. We report results for forecasts at horizons of 1, 2, 4, and 8 quarters ahead.

In the real-time forecast analysis of models with 4 or 8 variables, output is measured as GDP or GNP, depending on data vintage. Inflation is measured with the GDP or GNP deflator or price index. Quarterly real-time data on GDP or GNP, PCE, BFI, payroll employment, and the GDP or GNP price series are taken from the Federal Reserve Bank of Philadelphia's Real-Time Data Set for Macroeconomists (RTDSM). For simplicity, hereafter "GDP" and "GDP price index" refer to the output and price series, even though the measures are based on GNP and a fixed weight deflator for much of the sample. In the case of unemployment, the Treasury yield, and the fed funds rate, for which real-time revisions are small to essentially non–existent, we simply abstract from real-time aspects of the data, and we use current vintage data.

The full forecast evaluation period runs from 1985:Q1 through 2010:Q4, which involves real-time data vintages from 1985:Q1 through 2011:Q2. As described in Croushore and Stark (2001), the vintages of the RTDSM are dated to reflect the information available around the middle of each quarter. Normally, in a given vintage t, the available NIPA data run through period $t-1$. For each forecast origin t starting with 1985:Q1, we use the real-time data vintage t to estimate the forecast models and construct forecasts for periods t and beyond. The starting point of the model estimation sample is always 1965:Q1.

The results on real-time forecast accuracy cover forecast horizons of 1 quarter ($h = 1Q$), 2 quarters ($h = 2Q$), 1 year ($h = 1Y$), and 2 years ($h = 2Y$) ahead. In light of the time $t-1$ information actually incorporated in the VARs used for forecasting at t, the 1-quarter ahead forecast is a current quarter (t) forecast, while the 2-quarter ahead forecast is a next quarter ($t+1$) forecast. In keeping with Federal Reserve practice, the 1– and 2–year ahead forecasts for growth in GDP, PCE, BFI, and payroll employment and for inflation are 4–quarter rates of change (the 1–year ahead forecast is the percent change from period t through $t+3$; the 2–year ahead forecast is the percent change from period $t+4$ through $t+7$). The 1– and 2–year ahead forecasts for unemployment and the interest rates are quarterly levels in periods $t+3$ and $t+7$, respectively.

As discussed in such sources as Croushore (2005), Romer and Romer (2000), and Sims (2002), evaluating the accuracy of real-time forecasts requires a difficult decision on what to take as the actual data in calculating forecast errors. The GDP data available today for, say, 1985, represent the best available estimates of output in 1985. However, output as defined and measured today is quite different from output as defined and measured in 1970. For example, today we have available chain-weighted GDP; in the 1980s, output was measured with fixed-weight GNP. Forecasters in 1985 could not have foreseen such changes and the potential impact on measured output. Accordingly, we follow studies such as Clark (2011), Faust and Wright (2009), and Romer and Romer (2000) and use the second available estimates of GDP/GNP, PCE, BFI, payroll employment, and the GDP/GNP deflator as actuals in evaluating forecast accuracy. In the case of h–step ahead (for h = 1Q, 2Q, 1Y, and 2Y) forecasts made for period $t + h$ with vintage t data ending in period $t - 1$, the second available estimate is normally taken from the vintage $t + h + 2$ data set. In light of the abstraction from real-time revisions in unemployment and the interest rates, for these series the real-time data correspond to the final vintage data.

Finally, note that, throughout our analysis, we include four lags in all of our models. With Bayesian methods that naturally provide shrinkage, many prior studies have used the same approach of setting the lag length at the data frequency (e.g., Banbura, Giannone, and Reichlin (2010), Clark (2011), Del Negro and Schorfheide (2004), Koop (2012), and Sims (1993)).

4.2 Results on MCMC convergence properties and CPU time requirements

We begin with documenting the convergence properties of our MCMC algorithm for the BVAR-CSV model compared to existing algorithms for the BVAR-SV and BVAR models and with comparing CPU time requirements for each type of model.

Table 1 reports summary statistics for the distributions of inefficiency factors (IF) for the posterior estimates of all groups of model parameters. We consider 4-variable and 8-variable BVARS with independent and common volatility, using skip intervals of 10, 20, or 30 draws, intended to yield reasonable mixing properties (sufficiently low IF's). As noted above, all BVARs have four lags. The IF is the inverse of the relative numerical efficiency measure of Geweke (1992), and is estimated for each individual parameter as $1+2\sum_{k=1}^{\infty}\rho_k$, where ρ_k is the k-th order autocorrelation of the chain of retained draws. The estimates use the Newey and West (1987) kernel and a bandwidth of 4 percent of the sample of draws.

These convergence measures reveal two broad patterns: the IF's tend to rise as the model size increases from 4 to 8 variables, and the IF's are about the same for the BVAR-CSV as for the BVAR-SV. More specifically, the table indicates that for the 4-variable BVAR-SV, the highest IF is for the set of parameters ϕ_i, $i = 1, \ldots, n$, the innovation variances in the random walk models for the log variances $\lambda_{i,t}$. For the 4-variable BVAR-CSV, the highest IF is instead for the scaling matrix S (which allows the variances of each variable to differ by a factor that is constant over time). For both types of BVAR the IFs are substantially reduced when the skip interval increases from 10 to 20, and all the values are anyway lower than 20, which is typically regarded as satisfactory (see e.g. Primiceri (2005)). For the 8-variable BVAR-CSV the IF is again highest for S. In this case a skip interval of 20 would produce IFs lower than 25, but using a skip interval of 30 gets the average IF's for the elements of S down to 15 or less.

Based on this evidence, in all subsequent analysis in the paper, our results are based on 5000 retained draws, obtained from a larger sample of draws in which we set the skip interval as follows: BVAR-CSV, skip interval of 20 in 4-variable models and 30 in larger models; BVAR-SV, skip interval of 20 in 4-variable models and 30 in larger models; BVAR, skip interval of 2 in all cases. In all cases, we initialize the MCMC chain with 5000 draws, which are discarded.

As to the CPU time requirements for the different models, Table 2 shows that they increase substantially when increasing the number of variables and/or adding stochastic volatility to the BVAR.[7] As noted above, a key determinant of the CPU time requirements is the size of the posterior variance matrix that must be computed for sampling the VAR coefficients; the size of the matrix is a function of the square of the number of variables in the model. The CPU time for models with independent stochastic volatilities can be considerable. For our quarterly data sample of 1965:Q1-2011:Q2, it takes about 84 minutes to estimate the 4-variable BVAR-SV and 880 minutes (14.7 hours) to estimate the 8-variable BVAR-SV. The time requirement for the 8-variable BVAR-SV makes it infeasible to consider the model in a real-time forecast evaluation. Moreover, this time requirement likely deters other researchers and practitioners from using the independent stochastic volatility specification with models of more than a few variables (a deterrence evident in the fact that existing studies have not considered more than a handful of variables).

Introducing stochastic volatility through our common volatility specification yields significant computational gains relative to the independent volatility specification. With 4 variables, the BVAR-CSV estimation takes about 18 minutes, compared to almost 84 for the BVAR-SV. With 8 variables, the BVAR-CSV estimation takes nearly 47 minutes, compared to 879.5 minutes (14.7 hours) for the BVAR-SV. As noted earlier in the paper, these computational gains stem from the Kroneker structure of the coefficient variance matrix that results from having a single multiplicative volatility factor and the coefficient prior developed above. With these computational gains, we can readily consider stochastic volatility in the form of common volatility in our real-time forecasting analysis, for models of 4, 8, or

[7]We estimated the models with 2.93 GHZ processors, using the RATS software package.

15 variables.[8]

4.3 Full-sample results

Having established the computational advantages of our proposed common volatility specification, we turn now to a comparison of volatility estimates from the common volatility model (BVAR-CSV) versus a model that allows independent volatility processes for each variable (BVAR-SV). We consider both 4-variable and 8-variable models.

Figure 1 reports the volatility estimates for the 4-variable BVAR-SV, which, despite the independence across variables, are fairly similar in shape across variables, with higher volatility in the 1970s, a marked decrease starting in the early 1980s, in line with the literature on the Great Moderation, and a new increase with the onset of the financial crisis. Figure 2 shows the same estimates for BVAR-CSV, which are of course equal across variables apart from a scaling factor. These common volatility estimates follow paths quite similar to those obtained from the BVAR-SV model.

Figures 3 and 4 present corresponding estimates for the 8-variable specifications. The shape of Figure 3's volatility estimates from the BVAR-SV model that allows independent volatilities across variables are again similar across variables. The similarity is reflected in high correlations (ranging from 0.58 to 0.97) of each volatility estimate with the first principal component computed from the posterior median volatility estimates of each variable (the principal component is reported in Figure 5).[9] The common volatility estimates from the BVAR-CSV model follow paths similar to the BVAR-SV estimates. Figure 5 shows that the common volatility estimate closely resembles the first principal component computed from the posterior median volatility estimates obtained with the BVAR-SV model; the correlation between the common volatility estimate and the principal component is 0.99.

Based on these results, it seems that, in applications to at least standard macroeconomic VARs in US data, our common stochastic volatility specification can effectively capture time variation in conditional volatilities. Of course, in real time, reliable estimation of volatility may prove to be more challenging, in part because, at the end of the sample, only one-sided filtering is possible, and in part because of data revisions. Accordingly, in Figures 6-8 we compare time series of volatility estimates from five different real-time data vintages. In the

[8] While we don't include the result in the table because the estimation sample isn't the same, estimating the BVAR-CSV with 15 variables takes about 144 minutes.

[9] To compute the principal component, we take the posterior median estimates of volatility from the BVAR-SV model, standardize them, and compute the principal component as described in such studies as Stock and Watson (2002).

4-variable case, we consider estimates from both the BVAR-CSV and BVAR-SV models. In the 8-variable case, in light of the computational burden of the BVAR-SV model, we only consider results for the BVAR-CSV specification.

Three main messages emerge from the real-time estimates in Figures 6-8. First, the commonality in the volatility estimates for the 4 four variables in the BVAR-SV is confirmed for each vintage (Figure 6). For GDP growth and inflation, data revisions can shift the estimated volatility path but typically have little effect on the contours of the volatility estimate. The larger shifts in the volatility paths tend to be associated with benchmark or large annual revisions of the NIPA accounts. In the case of the unemployment and federal funds rates, volatility estimates tend to change less across vintages, presumably because the underlying data are not revised. Second, the BVAR-CSV volatility estimates for the 4-variable model are also quite similar across vintages (Figure 7). Finally, applied to the 8-variable model, the BVAR-CSV specification yields volatility estimates that follow very similar patterns across vintages (Figure 8). Again, contours are very similar across vintages, although data revisions can move the levels of volatility across data vintages.

To assess more generally how the competing models fit the full sample of data, we follow studies such as Geweke and Amisano (2010) in using 1-step ahead predictive likelihoods. The predictive likelihood is closely related to the marginal likelihood: the marginal likelihood can be expressed as the product of a sequence of 1-step ahead predictive likelihoods. In our model setting, the predictive likelihood has the advantage of being simple to compute. For model M_i, the log predictive likelihood is defined as

$$\log \mathrm{PL}(\mathrm{M}_i) = \sum_{t=t_0}^{T} \log p(y_t^o | y^{(t-1)}, M_i), \tag{44}$$

where y_t^o denotes the observed outcome for the data vector y in period t, $y^{(t-1)}$ denotes the history of data up to period $t-1$, and the predictive density is multivariate normal. Finally, in computing the log predictive likelihood, we sum the log values over the period 1980:Q1 through 2011:Q2.

The log predictive likelihood (LPL) estimates reported in Table 3 show that our proposed common volatility specification significantly improves the fit of a BVAR. In the four-variable case, the LPL of the BVAR-CSV model is about 87 points higher than the LPL of the constant volatility BVAR (in log units, a difference of just a few points would imply a meaningful difference in fit and, in turn, model probabilities). In the eight-variable case, the BVAR-CSV also fits the data much better than the BVAR, with a LPL difference of about 81 points. In the four-variable case, extending the volatility specification to permit independent volatilities for each variable offers some further improvement in fit: the LPL of the BVAR-SV is about 19 points higher than the BVAR-CSV.[10] However, as we emphasized above, this improvement comes at considerable cost, in terms of CPU time. Our proposed BVAR-CSV specification yields much of the gain in model fit to be achieved by allowing stochastic volatility, but at much lower CPU time cost.

4.4 Real-time forecast results

In this subsection we compare the relative performance of the 4-variable BVAR-SV model and the 4- and 8- variable BVAR-CSV models, starting with point forecasts and moving next to density forecasts. We also include univariate AR(4) models in the context, since they are known to be tough benchmarks, but our main focus is the relative performance of BVARs with no, common or independent volatility.[11] As mentioned, the evaluation sample is 1985Q1-2010Q4, we consider four forecast horizons, and the exercise is conducted in a real time manner, using recursive estimation with real time data vintages.

Table 4 reports the root mean squared error (RMSE) of each model relative to that of the BVAR, and the absolute RMSE for the BVAR for the 4-variable case (including GDP growth, unemployment, GDP inflation and the Fed funds rate). Hence, entries less than 1 indicate that the indicated model has a lower RMSE than the BVAR. Table A4 in the Appendix contains the same results but using AR models as benchmarks. To provide a rough gauge of whether the RMSE ratios are significantly different from 1, we use the Diebold and Mariano (1995) t-statistic for equal MSE, applied to the forecast of each model relative to the benchmark. Our use of the Diebold-Mariano test with forecasts that are, in many cases, nested is a deliberate choice. Monte Carlo evidence in Clark and McCracken (2011a,b) indicates that, with nested models, the Diebold-Mariano test compared against normal critical values can be viewed as a somewhat conservative (conservative in the sense of tending to have size modestly below nominal size) test for equal accuracy in the finite sample. As most of the alternative models can be seen as nesting the benchmark, we treat the tests as one-sided, and only reject the benchmark in favor of the null (i.e., we don't consider rejections of the alternative model in favor of the benchmark). Differences in accuracy that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-

[10] We don't report LPL results for the 8-variable BVAR-SV specification because the CPU time requirements for the model rule out using it for forecast evaluation and for computing the LPL.

[11] The AR(4) models are estimated with the same approach we have described for the BVAR, with the shrinkage hyperparameter θ set at 1.0.

values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h-1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Three main comments can be made based on the figures in Table 4 (and A4). First, adding independent stochastic volatility to the BVAR model with no volatility systematically improves the forecasts, and in general the gains are statistically significant. Second, constraining the volatility to be common in general further improves the forecasts. The BVAR-CSV produces lower RMSEs than the BVAR-SV in 12 out of 16 cases, with the BVAR-SV doing slightly better only for short term forecasts for the interest rate. While the advantages of the BVAR-CSV model over the BVAR-SV specification are small or modest, they are consistent. Third, the AR benchmark produces the lowest RMSEs for GDP growth and inflation. However, the MSEs differences are not statistically significant from those of the various BVARs. Instead, the BVARs with volatility are better for inflation and the Fed funds rate, and the gains with respect to the AR are statistically significant.

Table 5 provides corresponding results for the 8-variable case, adding consumption, investment, employment, and the Treasury yield to the variable set. However, in light of the computational requirements of the BVAR-SV specification with 8 variables, our forecasting results for the larger set do not include this model. For the included BVAR and BVAR-CSV models, Table 5 shows two main results. First, the larger BVAR is better than the 4-variable BVAR in 11 out of 16 cases. This is in line with several findings in the literature showing that a larger information set generally yields more accurate forecasts — see, e.g., Banbura, Giannone, and Reichlin (2010), Carriero, Clark and Marcellino (2011), Carriero, Kapetanios, and Marcellino (2011), and Koop (2012). More precisely, compared to the small model, the large system consistently yields more accurate point forecasts of GDP growth and unemployment, while the large model is beaten at short horizons for GDP inflation and the Fed funds rate. A second result is that, compared to a BVAR with constant volatility, the BVAR with common stochastic volatility significantly improves the accuracy of point forecasts. Compared to the BVAR, our proposed BVAR-CSV model lowers the RMSE in 75% of the cases (24 out of 32), and in many cases the gains are statistically significant. Admittedly, while the BVAR-CSV doesn't fare quite as well against the AR benchmark, beating the AR models in only 40% of the cases (but at least statistically significantly in most of these cases), BVARs generally have a difficult time beating AR models in data since 1985.

The RMSE, while informative and commonly used for forecast comparisons, is based on the point forecasts only and therefore ignores the rest of the forecast density. Of course the introduction of drifting volatility in a VAR makes it particularly well suited for density forecasting; for a 4-variable model, Clark (2011) shows that adding independent stochastic volatilities to a VAR significantly improves density forecasts. The overall calibration of the density forecasts can be measured with log predictive density scores, motivated and described in such sources as Geweke and Amisano (2010). At each forecast origin, we compute the log predictive score using the quadratic approximation of Adolfson, et al. (2007).[12] Specifically, we compute the log score with:

$$s_t(y^o_{t+h}) = -0.5\left(n\log(2\pi) + \log|V_{t+h|t}| + \left(y^o_{t+h} - \bar{y}_{t+h|t}\right)' V^{-1}_{t+h|t}\left(y^o_{t+h} - \bar{y}_{t+h|t}\right)\right), \quad (45)$$

where y^o_{t+h} denotes the observed outcome, $\bar{y}_{t+h|t}$ denotes the posterior mean of the forecast distribution, and $V_{t+h|t}$ denotes the posterior variance of the forecast distribution.

Table 6 reports differences in log scores with respect to the BVAR for the 4-variable case, such that entries greater than 0 indicate that the model has a better average log score (better density forecast) than the benchmark BVAR model. Table A6 in the Appendix contains the same results but using the AR as benchmark. To provide a rough gauge of the statistical significance of differences in average log scores, we use the Amisano and Giacomini (2007) t-test of equal means, applied to the log score for each model relative to the benchmark BVAR forecast. We view the tests as a rough gauge because, with nested models, the asymptotic validity of the Amisano and Giacomini (2007) test requires that, as forecasting moves forward in time, the models be estimated with a rolling, rather than expanding, sample of data. As most of the alternative models can be seen as nesting the benchmark, we treat the tests as one-sided, and only reject the benchmark in favor of the null (i.e., we don't consider rejections of the alternative model in favor of the benchmark). Differences in average scores that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h-1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

The results in Table 6 yield four main conclusions. First, consistent with Clark (2011), compared to a BVAR with constant volatilities, adding independent stochastic volatility to a BVAR model almost always improves the density forecasts, the only exceptions being 1- and 2-year ahead forecasts for growth and unemployment, and 2-year forecasts for the Fed fund rate. Second, as in the RMSE comparison, constraining the volatility to be common

[12] In some limited checks, we obtained qualitatively similar results with some other approaches to computing the predictive score. In our application, though, the quadratic approximation is easier to use.

in general further improves the forecasts. The BVAR-CSV beats the BVAR-SV model in 12 out of 16 cases, with the BVAR-SV model doing slightly better only for short term forecasts for the interest rate. Third, the relative performance of the AR model deteriorates with respect to the RMSE comparison. It is now better than the BVAR in only 6 out of 16 cases, and of the BVAR-SV in only 5 out of 16 cases. Moreover, the AR gains are never statistically significant. Fourth, the system scores are much better for the BVARs than the AR model forecasts, presumably because the VAR forecasts better account for covariance among the forecasts.

Table 7 provides corresponding results for the 8-variable case. Again, in light of the computational requirements of the BVAR-SV specification with 8 variables, our forecasting results for the larger set do not include this model. Focusing first on the original four variables (GDP growth, unemployment, GDP inflation and the Fed funds rate), the 8-variable model often improves on the density forecast accuracy of the 4-variable model. As in the RMSE comparison, the larger BVAR is generally better for growth and unemployment, worse for inflation, and better for the Fed funds rate but only at medium and long horizons. Compared to the BVAR benchmark, our proposed model with common stochastic volatility generally improves the accuracy of density forecasts, although more so at short horizons than long horizons. In particular, the BVAR-CSV is clearly better than the simple BVAR for $h = 1, 2$, improving average scores in 14 out of 16 cases. But at the one and two year horizons, log scores from the BVAR-CSV model are worse than those from the BVAR in 12 out of 16 cases. Finally, by the density metric, the BVAR-CSV specification beats the AR models in slightly less than 50% of the cases (15 out of 32, up from 13 out of 32 when using the RMSE, and often concentrated for h=1,2), and in most of these cases the gains are statistically significant.

In light of the seemingly high volatility of the period surrounding the financial crisis and recession of 2007-2009, it is natural to ask how our competing models fared over this period. In the interest of brevity, we simply summarize results here, in lieu of providing additional tables or charts. To assess forecast performance over the crisis period, the 1-step ahead log predictive score provides a broad indicator of performance. We have considered cumulative sums of these scores (of real-time forecasts) for the 2005:Q1-2010:Q4 period. This measure shows that, as might be expected, the crisis period produces some big jumps in absolute forecast performance; basically, some of the crisis out-turns were really tail events, most noticeably for unemployment and to a lesser extent for GDP growth. Particularly for unemployment, forecast accuracy dropped during the crisis, for all models. Generally, through this 2005-2010 period, as in the 1985-2004 period, the BVAR models with stochastic volatility are more accurate than the constant volatility BVAR benchmark.

In summary, the estimation analysis confirms that estimated volatilities are often similar across variables when estimated unrestrictedly, and therefore provides support for our common volatility specification. The forecast results are also favourable, in the sense that the typical ranking according to both the RMSE and the log predictive score is BVAR-CSV, BVAR-SV and BVAR. The ranking is clear-cut in the 4-variable case and when using RMSE, while the performance is more mixed in the 8-variable case and when using the log predictive score as the evaluation criterion. However, even in these cases the BVAR-CSV is preferred in the majority of cases, and in particular for shorter forecast horizons.

4.5 Current-vintage forecasting results for 15-variable model

In light of recent evidence that medium-scale BVARs often yield forecasts more accurate than small BVARs (e.g., Banbura, Giannone, and Reichlin (2010) and Koop (2012)), we also consider forecasts from 15-variable models. In this exercise, we use Koop's data (2012) and variable transformations, for the variables listed in Tables 8 and 9. In the interest of brevity, we compare just a BVAR-CSV and simple BVAR specification.[13] Tables 8 and 9 provide the results on RMSEs and average log scores, respectively.

Consistent with our real time results for smaller models, the RMSE results in Table 8 show that adding common stochastic volatility to the 15-variable model yields fairly consistent, small improvements in the accuracy of point forecasts, particularly at shorter horizons. More specifically, at forecast horizons of 1 or 2 quarters, the RMSE ratios are below 1 for almost all variables. In contrast, at a forecast horizon of 8 quarters, the RMSE ratios are slightly above 1 for most variables.

Similarly, the average score results in Table 9 show that including common stochastic volatility in a model typically improves forecast accuracy at shorter forecast horizons. At the 1-quarter horizon, the BVAR-CSV model yields a better score for all but one variable (with rough statistical significance in 1/2 of the cases). The BVAR-CSV also yields better scores for most variables at the 2-quarter horizon. But at the 8-quarter horizon, the BVAR yields slightly better — slightly enough that none of the differences appear to be statistically significant — scores than the BVAR-CSV for all but two variables.

[13] For the simple BVAR, we generate forecasts with a simple Normal-Wishart prior and posterior, simulating 5000 forecast draws.

5 Additional evidence: the case of the UK

The results we have obtained are favourable to the BVAR-CSV but one may wonder whether they are country specific. To provide additional evidence on the robustness of the good performance of the BVAR-CSV, in this section we repeat the key components of the US-based analysis of the previous section with data for the UK: the comparison of BVAR-CSV and BVAR-SV volatility estimates and forecast performance, for an 8-variable macroeconomic model.

5.1 Data and design of the forecast exercise

As for the US, we consider models of eight variables, at the quarterly frequency, with variables selected to match the US case as closely as possible. Specifically, the variables under analysis include growth of real GDP, growth of real household consumption expenditure, growth of real gross fixed investment, growth of employment, the unemployment rate, inflation as measured by the GDP deflator, a yield on bellwether (10-year when issued) government bonds, and the 3-month interbank rate.[14] In light of the sharp mean shifts that occurred in the nominal variables (inflation and the interest rates) between the start of the sample and the early 1990s, we include these nominal variables in the model as differences (i.e., as the change in inflation and the change in each interest rate). After forming forecasts of the changes of these variables, we cumulate to obtain forecasts of the levels of inflation and the interest rate.

In light of the more limited availability of real-time data for the UK, all of our results use current vintage data taken from the FAME database of the Federal Reserve Board. The quarterly data on the interest rates are constructed as simple within-quarter averages of the source monthly data. Growth and inflation rates are measured as annualized log changes (from $t-1$ to t).

The full forecast evaluation period runs from 1985:Q1 through 2011:Q2. For each forecast origin t starting with 1985:Q1, we use data through quarter $t-1$ to estimate the forecast models and construct forecasts for periods t and beyond. The starting point of the model estimation sample is always 1972:Q3, the earliest possible for our data on the included variables.

The results on forecast accuracy cover forecast horizons of 1 quarter ($h = 1Q$), 2 quarters ($h = 2Q$), 1 year ($h = 1Y$), and 2 years ($h = 2Y$) ahead. As in the US analysis, the 1– and 2–year ahead forecasts for growth in GDP, consumption, investment, and employment and for

[14]We constructed the GDP deflator as the ratio of nominal to real (chain-weight) GDP.

inflation are 4–quarter rates of change. The 1– and 2–year ahead forecasts for unemployment and the interest rates are quarterly levels in periods $t+3$ and $t+7$, respectively (given a forecast horizon of $t-1$).

5.2 Full sample results

Figures 9 and 10 present volatility estimates from the 8-variable BVAR-SV and BVAR-CSV models. Figure 9's volatility estimates from the BVAR-SV model that allows independent volatilities across variables are broadly similar across variables. However, compared to the US estimates (Figure 3), the UK estimates show somewhat more heterogeneity. Volatility generally trends down for inflation and interest rates, trends up for unemployment and employment, and shows the familiar volatility moderation in the period 1985-2008 for the growth rates of GDP, investment and consumption. Despite some heterogeneity, each volatility estimate is significantly correlated (with correlations ranging from 0.40 to 0.96 for all variables except employment, for which the correlation is -0.30) with the first principal component computed from the posterior median volatility estimates of each variable (the principal component is reported in Figure 11).[15]

The common volatility estimates from the BVAR-CSV model shown in Figure 10 follow paths broadly similar to the BVAR-SV estimates, for most variables. Figure 11 shows that the common volatility estimate closely resembles the first principal component computed from the posterior median volatility estimates obtained with the BVAR-SV model; the correlation between the common volatility estimate and the principal component is 0.99. Both the common volatility estimate and the principal component of the individual volatility estimates from the BVAR-SV model decreases rather monotonically from the early 1970s till about 2005, and increases mildly after that.

While our proposed model seems to reasonably capture variation over time in conditional volatilities, how much does that matter for the full sample model fit? The log predictive likelihood (LPL) estimates reported in Table 10 show that our proposed common volatility specification significantly improves the fit of the 8-variable BVAR. The LPL of the BVAR-CSV model is about 70 points higher than the LPL of the constant volatility BVAR.

Overall, the UK results, like the US results, suggest our proposed common stochastic volatility specification can effectively capture time variation in conditional volatilities. However, in the UK evidence, there is more heterogeneity across variables. Accordingly, it could

[15] To compute the principal component, we take the posterior median estimates of volatility from the BVAR-SV model, standardize them, and compute the principal component as described in such studies as Stock and Watson (2002).

be that the forecasting gains from our proposed BVAR-CSV model could be more limited than for the US. We will assess whether this is the case in the next subsection.

5.3 Forecast results

In this subsection we compare the relative performance of 8-variable BVARs with or without common stochastic volatility (all with four lags), starting with point forecasts and moving next to density forecasts. As for the US, we also include univariate AR(4) models in the comparison. The model priors are the ones described in section 3. As mentioned above, the evaluation sample is 1985Q1-2011Q2, we consider four forecast horizons, and the exercise is conducted in a pseudo-real time manner, using recursive estimation but a single data vintage (the most recent available).

Table 11 reports the root mean squared error (RMSE) of each model relative to that of the BVAR, and the RMSE level for the BVAR. Hence, entries less than 1 indicate that the model has a lower RMSE than the BVAR. Table A11 in the Appendix contains the same results but using the AR model as benchmark.[16]

Consistent with the findings of a range of studies of data for the US, our point forecast results for the UK suggest that it is difficult to rank the AR and the constant volatility BVAR, since the former has a lower RMSE than the latter in 15 out of 32 cases (8 variables and 4 forecast horizons). An AR model beats the BVAR for GDP growth, unemployment, inflation and bond yields, but by a statistically significant margin in only a few cases. The BVAR beats the AR for consumption, investment, employment and the bank rate, but rarely with statistical significance (indicated in Appendix Table A11, in which the AR model is the benchmark).

Adding common stochastic volatility to the BVAR model improves the forecasts in 20 out of 32 cases with a BVAR benchmark (Table 11), and in 21 out of 32 cases with the AR benchmark (Appendix Table A11). Against the BVAR benchmark, the BVAR-CSV's payoff in forecast accuracy is almost uniform at horizons up to and including 1 year. Against the AR benchmark, the BVAR-CSV consistently improves the accuracy of forecasts of real

[16] As for the US, we use the Diebold-Mariano t-statistic for equal MSE, applied to the forecast of each model relative to the benchmark, to provide a rough gauge of whether the RMSE ratios are significantly different from 1. The tests are one-sided, only rejecting the alternative model in favor of the benchmark. Differences in accuracy that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h-1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

variables, but not nominal variables. In general, in keeping with the US results, the larger improvements in RMSEs are often statistically significant.

Moving now to the evaluation of the density forecasts, Table 12 reports differences in log scores with respect to the BVAR, such that entries greater than 0 indicate that the BVAR-CSV has a better average log score (better density forecast) than the benchmark BVAR model. Table A12 in the Appendix presents comparable results using the AR model as benchmark.[17] Both tables provide the levels of average log scores for the benchmark model.

Against either the constant volatility BVAR or AR benchmark, our proposed common stochastic volatility BVAR yields significant gains in density forecast accuracy, especially at shorter horizons. Compared to the BVAR, the BVAR-CSV model yields a better score in 23 out of 32 cases. These figures are comparable to those for the US and, as for the US, the gains are larger and concentrated at the one- and two-quarter horizons, when the BVAR-CSV is systematically better. The performance deteriorates for one- and two-year ahead density forecasts of growth, consumption and investment. Not surprisingly, these are the variables whose BVAR-SV estimates of volatility (estimates obtained by treating the estimates of volatility as independent) are more different from the common stochastic volatility estimate. Compared to AR model forecasts, the BVAR-CSV model yields a better score in 19 out of 32 cases, again with gains that are concentrated at shorter forecast horizons. However, when all 8 variables are considered jointly, the BVAR-CSV (significantly) beats the AR model at all horizons, presumably because the BVAR forecasts better account for the covariances among variables.

In summary, notwithstanding the higher heterogeneity in independent estimates of volatility for each variable, the forecasting gains from the BVAR-CSV are generally confirmed for the UK. Broadly, the gains for the UK are comparable to those for the US, for both point and density forecasts.

[17] As for the US, to provide a rough gauge of the statistical significance of differences in average log scores, we use the Amisano-Giacomini (2007) t-test of equal means, applied to the log score for each model relative to the benchmark BVAR forecast. Again, we treat the tests as one-sided. Differences in average scores that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h-1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

6 Conclusions

In this paper we propose to model conditional volatilities as driven by a single common unobserved factor. Using a combination of a standard natural conjugate prior for the VAR coefficients, and an independent prior on a common stochastic volatility factor, we derive the posterior densities for the parameters of the resulting BVAR with common stochastic volatility (BVAR-CSV). Under the chosen prior the conditional posterior of the VAR coefficients features a Kroneker structure that allows for fast estimation.

Empirically, we start with systems composed of 4 and 8 US variables, and we show that there is substantial evidence of common volatility. We then examine the accuracy of real-time forecasts from VARs with constant volatility, independent stochastic volatilities, and our proposed common stochastic volatility. We find that compared to a model with constant volatilities, our proposed common volatility model significantly improves model fit and forecast accuracy. The gains are comparable to or as great as the gains achieved with a conventional stochastic volatility specification that allows independent volatility processes for each variable. But our common volatility specification greatly speeds computations.

As a robustness check, we repeat the volatility and forecasting analysis using comparable UK data. Notwithstanding slightly higher heterogeneity in the estimated volatility across variables than for the US, the BVAR-CSV still delivers improved accuracy of both point and density forecasts.

We interpret these results as evidence that the BVAR-CSV efficiently summarizes the information in a possibly large dataset and accounts for changing volatility, while helping to significantly reduce computation costs relative to a model with independent stochastic volatilities. For these reasons this class of models should have a wide range of applicability for forecasting and possibly also for policy simulation exercises.

7 Appendix: Some derivations

In this Appendix we derive the conditional posterior distributions used in the MCMC scheme. Recall the model:

$$y_t = \Pi_0 + \Pi(L) y_{t-1} + v_t, \tag{46}$$

$$v_t = \lambda_t^{0.5} A^{-1} S^{1/2} \epsilon_t, \ \epsilon_t \sim N(0, I_n), \tag{47}$$

$$\log(\lambda_t) = \log(\lambda_{t-1}) + \nu_t, \ \nu_t \sim \text{iid } N(0, \phi). \tag{48}$$

The conjectured posteriors are:

$$\begin{aligned}
\mathrm{vec}(\Pi)|A, S, \phi, \Lambda, y &\sim N(\mathrm{vec}(\bar{\mu}_{\Pi}), \bar{\Omega}_{\Pi}) & (49)\\
a_i|\Pi, S, \phi, \Lambda, y &\sim N(\bar{\mu}_{a,i}, \bar{\Omega}_{a,i}),\ i = 2, \ldots, n & (50)\\
s_i|\Pi, A, \phi, \Lambda, y &\sim IG\left(d_s \cdot \underline{s}_i + \sum_{t=1}^{T}(\tilde{v}_{i,t}^2/\lambda_t),\ d_s + T\right),\ i = 2, \ldots, n & (51)\\
\phi|\Pi, A, S, \Lambda, y &\sim IG\left(d_\phi \cdot \underline{\phi} + \sum_{t=1}^{T}\nu_t^2,\ d_\phi + T\right). & (52)
\end{aligned}$$

7.1 Likelihood of the VAR

Let us now use another representation for the VAR in (46). The VAR is:

$$y_t = \Pi_0 + \Pi_1 y_{t-1} + ... + \Pi_p y_{t-p} + v_t \tag{53}$$

By defining $\Pi = (\Pi_0,\ \Pi_1,\ ..., \Pi_p)'$ of dimension $N \times k$ (where $k = 1 + Np$) and $x_t = (1, y'_{t-1}, y'_{t-2}, ..., y'_{t-p})'$ of dimension $k \times 1$ we have:

$$y_t = \Pi' x_t + v_t. \tag{54}$$

Now consider the equations for all observations $t = 1, ..., T$. By stacking them by columns and then transposing the system we get:

$$Y = X\Pi + v, \tag{55}$$

where Y is a $T \times N$ data-matrix with rows y'_t, X is a $T \times k$ data-matrix with rows $x'_t = (1, y'_{t-1}, y'_{t-2}, ..., y'_{t-p})$ and v is a $T \times N$ data-matrix with rows v'_t. Consider now the equations for variable j in the system above. We have:

$$\underset{T\times 1}{Y_j} = X \underset{k\times 1}{\Pi_j} + \underset{T\times 1}{v_j}, \tag{56}$$

where Y_j is the j-th column of Y, and Π_j the j-th column of Π. By stacking these equations by column for $j = 1, ...N$ we get:

$$\mathrm{vec}(Y) = \mathrm{vec}(X\Pi I) + \mathrm{vec}(v). \tag{57}$$

Setting $y = \mathrm{vec}(Y)$, $Z = (I \otimes X)$,$v = \mathrm{vec}(v)$ we can write:

$$y = Z\mathrm{vec}(\Pi) + v. \tag{58}$$

Under our specification, the residual variance–covariance matrix for period t is $Var(v_t) = \Sigma_t \equiv \lambda_t A^{-1} S A^{-1\prime} = \lambda_t \left(\tilde{A}'\tilde{A}\right)^{-1}$. Define a diagonal matrix having the whole history of λ in the main diagonal:

$$\Lambda = diag(\lambda_1, ..., \lambda_T). \tag{59}$$

The variance of v is given by[18]:

$$Q = \left(\tilde{A}'\tilde{A}\right)^{-1} \otimes \Lambda. \tag{60}$$

The error term in (58) has the following conditional distribution:

$$v|\Pi, A, S, \phi, \Lambda, Z \sim N(0, Q). \tag{61}$$

It follows that the likelihood of (58) is:

$$p(y|\Pi, A, S, \phi, \Lambda, Z) = 2\pi^{-\frac{Tn}{2}} |Q|^{-\frac{1}{2}} \exp(-(y - Z\text{vec}(\Pi))' \, Q^{-1}(y - Z\text{vec}(\Pi)) \, /2) \tag{62}$$

7.2 Drawing $\Pi|A, S, \phi, \Lambda, y$

Conditionally on the other parameters of the model, the posterior of Π can be obtained by using standard results for the $N-IW$ prior. The assumed prior distribution is $\text{vec}(\Pi)|A, S \sim$

[18] This can be derived as follows:

$$Q = E[vv'] = \begin{bmatrix} E[\mathbf{v}_1\mathbf{v}_1'] & ... & E[\mathbf{v}_N\mathbf{v}_1'] \\ ... & \underset{T\times T}{E[\mathbf{v}_i\mathbf{v}_j']} & ... \\ E[\mathbf{v}_1\mathbf{v}_N'] & ... & E[\mathbf{v}_N\mathbf{v}_N'] \end{bmatrix}$$

The generic term $E[\mathbf{v}_i\mathbf{v}_j']$ in the above matrix is equal to:

$$\begin{aligned} E[\mathbf{v}_i\mathbf{v}_j'] &= \begin{bmatrix} E[v_1^{(i)}v_1^{(j)\prime}] & & E[v_T^{(i)}v_1^{(j)\prime}] \\ & ... & \\ E[v_1^{(i)}v_T^{(j)\prime}] & & E[v_T^{(i)}v_T^{(j)\prime}] \end{bmatrix} = \begin{bmatrix} [\Sigma_1]_{ij} & & \\ & ... & \\ & & [\Sigma_T]_{ij} \end{bmatrix} \\ &= \begin{bmatrix} \lambda_1 \left[\left(\tilde{A}'\tilde{A}\right)^{-1}\right]_{ij} & & \\ & ... & \\ & & \lambda_T \left[\left(\tilde{A}'\tilde{A}\right)^{-1}\right]_{ij} \end{bmatrix} = \left[\left(\tilde{A}'\tilde{A}\right)^{-1}\right]_{ij} \cdot \Lambda \end{aligned}$$

Therefore we have:

$$Q = \begin{bmatrix} \left[\left(\tilde{A}'\tilde{A}\right)^{-1}\right]_{11} \cdot \Lambda & ... & \left[\left(\tilde{A}'\tilde{A}\right)^{-1}\right]_{N1} \cdot \Lambda \\ ... & \left[\left(\tilde{A}'\tilde{A}\right)^{-1}\right]_{ij} \cdot \Lambda & ... \\ \left[\left(\tilde{A}'\tilde{A}\right)^{-1}\right]_{1N} \cdot \Lambda & ... & \left[\left(\tilde{A}'\tilde{A}\right)^{-1}\right]_{NN} \cdot \Lambda \end{bmatrix} = \left(\tilde{A}'\tilde{A}\right)^{-1} \otimes \Lambda$$

$N(\text{vec}(\underline{\mu}_\Pi), \underline{\Omega}_\Pi)$, therefore the conditional posterior kernel of vec(Π) is given by:

$$p(\text{vec}(\Pi)|A, S, \phi, \Lambda, y) \propto p(y|\text{vec}(\Pi), A, S, \phi, \Lambda)p(\text{vec}(\Pi)|A, S). \tag{63}$$

Under the knowledge of A, S, ϕ, Λ, this is the normal kernel usually found when combining a normal likelihood with a normal prior. As shown in Geweke (2005), the conditional posterior mean and variance of vec(Π) are:

$$\text{vec}(\overline{\mu}_\Pi) = \overline{\Omega}_\Pi(\underline{\Omega}_\Pi^{-1}\underline{\mu}_\Pi + Z'Q^{-1}y) \tag{64}$$

$$\overline{\Omega}_\Pi = (\underline{\Omega}_\Pi^{-1} + Z'Q^{-1}Z)^{-1}. \tag{65}$$

Thanks to the Kroneker structure of the prior, we can write these moments as follows:

$$\text{vec}(\bar{\mu}_\Pi) = \bar{\Omega}_\Pi\left\{\text{vec}\left(\sum_{t=1}^{T} X_t y_t' \Sigma_t^{-1}\right) + \underline{\Omega}_\Pi^{-1}\text{vec}(\underline{\mu}_\Pi)\right\} \tag{66}$$

$$\bar{\Omega}_\Pi = \left(\tilde{A}'\tilde{A}\right)^{-1} \otimes \left(\underline{\Omega}_0^{-1} + \sum_{t=1}^{T}(\frac{1}{\lambda_t}X_t X_t')\right)^{-1}. \tag{67}$$

Note that — once written in this way — the inversion of $\overline{\Omega}_\Pi$ which will be needed in our MCMC sampling scheme does involve computations of order $n^3 + k^3$ rather than of order n^3k^3. This is the same simplification that typically happens in the derivation of the posteriors from the natural conjugate N-IW prior.

We now show how to go from (64) and (65) to (66) and (67). Consider first the variance expression in (65). The term $Z'Q^{-1}Z$ can be written as follows:

$$\underset{n^2p\times Tn}{Z}{}' \underset{Tn\times Tn}{Q^{-1}} \underset{Tn\times n^2p}{Z} = \begin{pmatrix} I_n \otimes X_1 & ... & I_n \otimes X_T \end{pmatrix} \left(\tilde{A}'\tilde{A} \otimes \Lambda^{-1}\right) \begin{pmatrix} I_n \otimes X_1' \\ ... \\ I_n \otimes X_T' \end{pmatrix}$$

$$= \sum_{t=1}^{T} \frac{1}{\lambda_t}(I_n \otimes X_t)\left(\tilde{A}'\tilde{A} \otimes 1\right)(I_n \otimes X_t').$$

Therefore we can write:

$$\begin{aligned} \overline{\Omega}_\Pi^{-1} &= \underline{\Omega}_\Pi^{-1} + \sum_{t=1}^{T}\frac{1}{\lambda_t}(I_n \otimes X_t)\left(\tilde{A}'\tilde{A} \otimes 1\right)(I_n \otimes X_t') \\ &= \left(\tilde{A}'\tilde{A}\right) \otimes \underline{\Omega}_0^{-1} + \sum_{t=1}^{T}\left(\tilde{A}'\tilde{A}\right) \otimes (\frac{1}{\lambda_t}X_t X_t') \\ &= \left(\tilde{A}'\tilde{A}\right) \otimes \underline{\Omega}_0^{-1} + \left(\tilde{A}'\tilde{A}\right) \otimes \sum_{t=1}^{T}(\frac{1}{\lambda_t}X_t X_t') \\ &= \left(\tilde{A}'\tilde{A}\right) \otimes \left(\underline{\Omega}_0^{-1} + \sum_{t=1}^{T}(\frac{1}{\lambda_t}X_t X_t')\right), \end{aligned} \tag{68}$$

which implies (67). Now consider the mean in (64). The term $Z'Q^{-1}y$ can be written as follows:

$$\begin{aligned}
\underset{n^2p\times Tn}{Z}{}' \underset{Tn\times Tn}{Q^{-1}} \underset{Tn\times 1}{y} &= \sum_{t=1}^{T} \frac{1}{\lambda_t} \underset{n^2p\times n}{\left(I_n \otimes \underset{np\times 1}{X_t}\right)} \left(\underset{n\times n}{\tilde{A}'\tilde{A}} \otimes 1\right) y_t \\
&= \sum_{t=1}^{T} \left(\frac{1}{\lambda_t} \underset{n\times n}{\widetilde{A}'\widetilde{A}} \otimes \underset{np\times 1}{X_t}\right) y_t \\
&= \sum_{t=1}^{T} \underset{n^2p\times n}{\left(\Sigma_t^{-1} \otimes X_t\right)} \underset{n\times 1}{y_t} \\
&= \sum_{t=1}^{T} \left(\Sigma_t^{-1} \otimes X_t\right) vec(y_t') \\
&= \sum_{t=1}^{T} vec\left(X_t y_t' \Sigma_t^{-1}\right) \\
&= vec\left(\sum_{t=1}^{T} X_t y_t' \Sigma_t^{-1}\right).
\end{aligned} \tag{69}$$

7.3 Drawing $a_i|\Pi, S, \phi, \Lambda, y$

For the covariance elements we can use the derivation in Cogley and Sargent's (2005) Appendix B.2.4, in particular equations 76 through 78. The resulting conditional posterior is:

$$a_i|\Pi, S, \phi, \Lambda, y \sim N(\bar{\mu}_{a,i}, \bar{\Omega}_{a,i}), \;\; i = 2, \ldots, n. \tag{70}$$

7.4 Drawing $s_i|\Pi, A, \phi, \Lambda, y$

For the scaling variances s_i, we conjecture that the conditional posterior is:

$$s_i|\Pi, A, \phi, \Lambda, y \sim IG(d_s \cdot \underline{s}_i + \sum_{t=1}^{T} (\tilde{v}_{i,t}^2/\lambda_t), \;\; d_s + T), \;\; i = 2, \ldots, n. \tag{71}$$

This can be derived as follows. Recall the model in (58), which has the likelihood given in (62). The matrix Q^{-1} can be written as:

$$\begin{aligned}
Q^{-1} &= (\tilde{A}'\tilde{A}) \otimes \Lambda^{-1} = (A'S^{-1}A) \otimes \Lambda^{-1} \\
&= (A' \otimes I_T)(I_N \otimes \Lambda^{-1/2})(S^{-1} \otimes I_T)(I_N \otimes \Lambda^{-1/2})(A \otimes I_T) \\
&= A^{*'} \Lambda^* (S^{-1} \otimes I_T) \Lambda^* A^*,
\end{aligned} \tag{72}$$

where we defined $\Lambda^* = I_N \otimes \Lambda^{-1/2}$ and $A^* = A \otimes I$. Define the rescaled residuals $v^* = \Lambda^* A^*(y - Z\text{vec}(\Pi))$. We can write:

$$\begin{aligned}
& (y - Z\text{vec}(\Pi))' Q^{-1}(y - Z\text{vec}(\Pi)) \\
= & \ (y - Z\text{vec}(\Pi))' A^{*'} \Lambda^* (S^{-1} \otimes I_T) \Lambda^* A^* (y - Z\text{vec}(\Pi)) \\
= & \ v^{*\prime}(S^{-1} \otimes I_T) v^* = \sum_i \sum_j v_i^{*\prime} v_j^* S_{ij}^{-1} = tr(RS^{-1}),
\end{aligned} \tag{73}$$

where R is the matrix of rescaled residual cross products with generic element $[r_{ij}] = v_i^{*\prime} v_j^* = (y_i - Z_i\text{vec}(\Pi))' A^{*\prime} \Lambda^{*\prime} \Lambda^* A^*(y_j - Z_j\text{vec}(\Pi))$. Using (73), the likelihood in (62) can be written as:

$$\begin{aligned}
p(y|\Pi, A, S, \phi, \Lambda, Z) &= 2\pi^{-\frac{Tn}{2}} |A^{*'} \Lambda^* (S^{-1} \otimes I_T) \Lambda^* A^*|^{-\frac{1}{2}} \\
& \quad \cdot \exp(-(y - Z\text{vec}(\Pi))' \ Q^{-1}(y - Z\text{vec}(\Pi))/2) \\
&\propto |S|^{\frac{-T}{2}} \exp(-tr(RS^{-1})/2 \ ),
\end{aligned} \tag{74}$$

which is the kernel of a Wishart distribution for S^{-1}, or of an Inverse Wishart distribution for S. The conjugate prior for this distribution is:

$$S^{-1} \sim W(R_0^{-1}, \ d_s), \tag{75}$$

with prior density:

$$\begin{aligned}
p(S^{-1}) &= 2^{-d_s n/2} \cdot \pi^{-(n-1)n/4} \ |R_0^{-1}|^{d_s/2} \cdot \left(\prod_{i=1}^{p} \Gamma[(d_s + i - 1)/2] \right)^{-1} \\
& \quad \cdot |S^{-1}|^{(d_s - 1 - n)/2} \cdot \exp(-tr(R_0 S^{-1})/2).
\end{aligned} \tag{76}$$

The prior used in the paper is an inverse gamma for s_i, which can be interpreted as a special case of $S \sim IW(R_0, \ d_s)$ where S is diagonal, in particular we set $R_0 = diag(d_s \cdot \underline{s}_i)$. The posterior kernel is therefore:

$$\begin{aligned}
p(S^{-1}|y, \Pi, A, \phi, \Lambda) &\propto p(y|\Pi, A, S, \phi, \Lambda) * p(S^{-1}) \\
&\propto \left[2\pi^{-\frac{Tn}{2}} |S|^{-\frac{T}{2}} \exp(-tr(RS^{-1})/2 \ ) \right] \cdot \\
& \quad \left[\begin{array}{c} 2^{-d_s n/2} \cdot \pi^{-(n-1)n/4} \ |R_0^{-1}|^{d_s/2} \cdot \left(\prod_{i=1}^{p} \Gamma[(d_s + i - 1)/2] \right)^{-1} \\ \cdot |S^{-1}|^{(d_s - 1 - p)/2} \cdot \exp(-tr(R_0^{-1} S^{-1})/2) \end{array} \right] \\
&\propto |S|^{-\frac{T}{2}} |S^{-1}|^{(d_s - 1 - p)/2} \cdot \exp(-tr(RS^{-1})/2 \ ) \cdot \exp(-tr(R_0 S^{-1})/2) \\
&\propto |S^{-1}|^{\frac{T + d_s - 1 - n}{2}} \cdot \exp(-tr(R + R_0) S^{-1})/2),
\end{aligned} \tag{77}$$

which is the kernel of $S^{-1}|\Pi, A, \phi, \Lambda, y \sim W((R+R_0)^{-1},\ T+d_s)$, which implies that $S|\Pi, A, \phi, \Lambda, y \sim IW(R+R_0,\ T+d_s)$. Recalling that the generic element of R is r_{ij} $= v_i^{*\prime} v_j^*$, the sum of squares of the rescaled residuals, and that we are imposing diagonality on R, we see that $R = \sum_{t=1}^{T}(\tilde{v}_{i,t}^2/\lambda_t)$ and the posterior scale matrix is given by $diag(d_s \cdot \underline{s}_i) + \sum_{t=1}^{T}(\tilde{v}_{i,t}^2/\lambda_t)$.

7.5 Drawing $\phi|\Pi, A, S, \Lambda, y$

For the variance of the volatility term,

$$\phi|\Pi, A, S, \Lambda, y \sim IG(d_\phi \cdot \underline{\phi} + \sum_{t=1}^{T} \nu_t^2,\ d_\phi + T)$$

we refer the reader to Cogley and Sargent's (2005) equation 69.

7.6 Drawing $\lambda_t|\Pi, A, S, \phi, y$

Finally, the expression for the conditional kernel

$$f(\lambda_t|\lambda_{t-1}, \lambda_{t+1}, u^T, \phi, y) \sim \lambda_t^{-1.5} \exp\left(\frac{-w_t}{2\lambda_t}\right) \exp\left(\frac{-(\log \lambda_t - \mu_t)}{2\sigma_c^2}\right),$$

follows from Cogley and Sargent's (2005) equation 80, once we condition on ϕ.

References

[1] Adolfson, M., Linde, J., and Villani, M., 2007. Forecasting Performance of an Open Economy DSGE Model, *Econometric Reviews* 26, 289-328.

[2] Amisano, G., and Giacomini, R., 2007. Comparing Density Forecasts via Weighted Likelihood Ratio Tests, *Journal of Business and Economic Statistics* 25, 177-190.

[3] Banbura, M., Giannone, D., and Reichlin, L., 2010. Large Bayesian Vector Autoregressions, *Journal of Applied Econometrics* 25(1), 71-92.

[4] Bernanke B., Boivin, J., and Eliasz, P.S., 2005. Measuring the Effects of Monetary Policy: A Factor-augmented Vector Autoregressive (FAVAR) Approach, *Quarterly Journal of Economics* 120(1), 387-422

[5] Carriero, A., Kapetanios, G., and Marcellino M., 2011. Forecasting Large Datasets with Bayesian Reduced Rank Multivariate Models, *Journal of Applied Econometrics* 26(5), 735-761.

[6] Chib, S., Nardari, F., and Shephard, N., 2006. Analysis of High Dimensional Multivariate Stochastic Volatility Models, *Journal of Econometrics* 134(2), 341-371.

[7] Chib, S., Nardari, F., and Shephard, N., 2002. Markov Chain Monte Carlo Methods for Stochastic Volatility Models, *Journal of Econometrics* 108(2), 281-316.

[8] Clark, T., 2011. Real-Time Density Forecasts from BVARs with Stochastic Volatility, *Journal of Business and Economic Statistics* 29, 327-341.

[9] Clark, T., and Davig, T., 2011 Decomposing the declining volatility of long-term inflation expectations, *Journal of Economic Dynamics and Control*, 35(7), 981-999.

[10] Clark, T., and McCracken, M., 2011a. Nested Forecast Model Comparisons: A New Approach to Testing Equal Accuracy, manuscript, Federal Reserve Bank of St. Louis.

[11] Clark, T., and McCracken, M., 2011b. Testing for Unconditional Predictive Ability, in *Oxford Handbook of Economic Forecasting*, Michael P. Clements and David F. Hendry, eds., Oxford: Oxford University Press.

[12] Cogley, T., and Sargent, T., 2005. Drifts and Volatilities: Monetary Policies and Outcomes in the post-WWII US, *Review of Economic Dynamics* 8 (2), 262-302.

[13] Croushore, D., 2006. Forecasting with Real-Time Macroeconomic Data, in *Handbook of Economic Forecasting*, eds. G. Elliott, C. Granger, and A. Timmermann, Amsterdam: North Holland.

[14] Croushore, D., and Stark, T., 2001. A Real-Time Data Set for Macroeconomists, *Journal of Econometrics* 105, 111-30.

[15] Diebold, F., and Mariano, R., 1995. Comparing Predictive Accuracy, *Journal of Business and Economic Statistics* 13, 253-263.

[16] Del Negro, M., and Otrok, C., 2008. Dynamic Factor Models with Time-Varying Parameters: Measuring Changes in International Business Cycles, manuscript, Federal Reserve Bank of New York

[17] Del Negro, M., and Schorfheide, F., 2004. Priors from General Equilibrium Models for VARs, *International Economic Review* 45(2), 643-673.

[18] Faust, J., and Wright, J., 2009. Comparing Greenbook and Reduced Form Forecasts using a Large Realtime Dataset, *Journal of Business and Economic Statistics* 27, 468-479.

[19] Geweke, J., and Amisano, G., 2010. Comparing and Evaluating Bayesian Predictive Distributions of Asset Returns, *International Journal of Forecasting* 26, 216-230.

[20] Geweke, J., 2005. *Contemporary Bayesian Econometrics and Statistics*, Wiley, Hoboken, New Jersey.

[21] Geweke, J., 1992. Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments, in: Bernardo, J., Berger, J., Dawid, A., Smith, A. (Eds.), Bayesian Statistics 4. Clarendon Press, Oxford, pp. 641-649.

[22] Jacquier, E., Polson, N., and Rossi, P., 1995. Stochastic Volatility: Univariate and Multivariate Extensions, Working Paper 19-95, Rodney L. White Center for Financial Research.

[23] Jacquier, E., Polson, N., and Rossi, P., 1994, Bayesian Analysis of Stochastic Volatility Models, *Journal of Business and Economic Statistics* 12, 371-418.

[24] Harvey, D., Leybourne, S. and Newbold, P., 1997. Testing the Equality of Prediction Mean Squared Errors, *International Journal of Forecasting* 13, 281-291.

[25] Kadiyala, K., and Karlsson, S., 1997. Numerical Methods for Estimation and Inference in Bayesian VAR-Models, *Journal of Applied Econometrics* 12(2), 99-132.

[26] Koop, G., 2012. Forecasting with Medium and Large Bayesian VARs, *Journal of Applied Econometrics*, forthcoming.

[27] Koop, G., and Korobilis, D., 2012. Large Time-Varying Parameter VARs, manuscript, University of Strathclyde.

[28] Litterman, R., 1986. Forecasting with Bayesian Vector Autoregressions-Five Years of Experience, *Journal of Business and Economic Statistics* 4, 25-38.

[29] Newey, W., and West, K., 1987. A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55, 703-08.

[30] Osiewalski, J., and Pajor, A., 2009. Bayesian Analysis for Hybrid MSF-SBEKK Models of Multivariate Volatility, *Central European Journal of Economic Modelling and Econometrics* 1, 179-202.

[31] Pajor, A., 2006. Bayesian Analysis of the Conditional Correlation Between Stock Index Returns with Multivariate Stochastic Volatility Models, *Acta Physica Polonica B* 37, 3093-3103.

[32] Primiceri, G., 2005. Time Varying Structural Vector Autoregressions and Monetary Policy, *Review of Economic Studies* 72, 821-852.

[33] Romer, C., and Romer, D., 2000. Federal Reserve Information and the Behavior of Interest Rates, *American Economic Review* 90, 429-457.

[34] Sims, C., 2002. The Role of Models and Probabilities in the Monetary Policy Process, *Brookings Papers on Economic Activity* 2, 1-40.

[35] Sims, C., 1993. A Nine-Variable Probabilistic Macroeconomic Forecasting Model, in *Business Cycles, Indicators and Forecasting*, James H. Stock and Mark W. Watson, editors, University of Chicago Press, 179-212.

[36] Sims, C., and Zha, T., 1998. Bayesian Methods for Dynamic Multivariate Models, *International Economic Review* 39(4), 949–68.

[37] Stock, J.H., and Watson, M.H., 2002. Macroeconomic Forecasting Using Diffusion Indexes, *Journal of Business and Economic Statistics* 20(2), 147-162.

Table 1. Summary of Inefficiency Factors for Various Model Specifications

parameter block	# parameters	median	mean	min	max
4 variables, independent st. vol.: skip interval of 10					
Π	68	1.40	1.47	0.79	2.74
A	6	1.65	1.87	1.09	3.00
ϕ	4	11.54	11.13	7.80	13.65
Λ	744	4.81	4.96	0.80	14.93
4 variables, independent st. vol.: skip interval of 20					
Π	68	1.09	1.19	0.54	2.03
A	6	1.04	1.30	0.77	2.32
ϕ	4	8.12	7.92	4.12	11.34
Λ	744	2.82	2.92	0.68	8.39
4 variables, common st. vol.: skip interval of 10					
Π	68	1.52	1.64	0.73	5.28
A	6	1.79	2.10	1.07	4.25
S	3	16.31	16.24	13.69	18.72
ϕ	1	8.55	8.55	8.55	8.55
Λ	186	10.70	11.24	7.31	18.66
4 variables, common st. vol.: skip interval of 20					
Π	68	1.21	1.27	0.73	3.70
A	6	1.29	1.25	0.95	1.47
S	3	11.03	10.35	8.84	11.19
ϕ	1	3.15	3.15	3.15	3.15
Λ	186	5.65	5.80	3.24	10.53

Notes:
1. For each individual parameter, the inefficiency factor is estimated as $1 + 2\sum_{k=1}^{\infty} \rho_k$, where ρ_k is the k-th order autocorrelation of the chain of retained draws. The estimates use the Newey-West kernel and a bandwidth of 4 percent of the sample of draws.
2. The table provides summary statistics for the inefficiency factors computed for groups of model parameters.

Table 1, Continued. Summary of Inefficiency Factors for Various Model Specifications

parameter block	# parameters	median	mean	min	max
8 variables, independent st. vol.: skip interval of 10					
Π	264	1.47	1.56	0.65	3.94
A	28	2.00	2.31	1.17	5.36
ϕ	8	14.01	13.39	5.42	21.01
Λ	1488	8.74	9.61	1.42	27.89
8 variables, independent st. vol.: skip interval of 20					
Π	264	1.18	1.23	0.63	2.62
A	28	1.44	1.62	0.87	3.42
ϕ	8	7.85	8.13	4.91	12.85
Λ	1488	4.71	4.93	0.77	12.55
8 variables, independent st. vol.: skip interval of 30					
Π	264	1.12	1.15	0.58	2.01
A	28	1.19	1.38	0.70	2.96
ϕ	8	5.06	5.22	3.55	8.22
Λ	1488	3.26	3.41	0.69	8.43
8 variables, common st. vol.: skip interval of 10					
Π	264	1.69	1.94	0.67	5.53
A	28	1.85	1.98	0.90	4.01
S	7	39.11	38.77	35.82	41.74
ϕ	1	8.75	8.75	8.75	8.75
Λ	186	25.84	26.26	17.43	35.92
8 variables, common st. vol.: skip interval of 20					
Π	264	1.38	1.48	0.67	4.47
A	28	1.36	1.52	0.95	2.75
S	7	22.97	23.13	21.85	24.63
ϕ	1	3.50	3.50	3.50	3.50
Λ	186	14.80	14.76	10.79	18.52
8 variables, common st. vol.: skip interval of 30					
Π	264	1.21	1.27	0.50	3.32
A	28	1.31	1.41	0.79	2.73
S	7	15.39	15.32	13.16	17.13
ϕ	1	1.92	1.92	1.92	1.92
Λ	186	9.35	9.56	6.62	12.61

Notes:
1. For each individual parameter, the inefficiency factor is estimated as $1 + 2\sum_{k=1}^{\infty} \rho_k$, where ρ_k is the k-th order autocorrelation of the chain of retained draws. The estimates use the Newey-West kernel and a bandwidth of 4 percent of the sample of draws.
2. The table provides summary statistics for the inefficiency factors computed for groups of model parameters.

Table 2. CPU time requirements for different models

model	skip interval (between draws)	CPU time (minutes)
BVAR(4), 4 variables, independent stochastic volatility	20	83.6
BVAR(4), 4 variables, common stochastic volatility	20	18.1
BVAR(4), 8 variables, independent stochastic volatility	20	1291
BVAR(4), 8 variables, common stochastic volatility	20	46.8

Notes: Each model is estimated to generate a sample of 5000 retained draws, by skipping every k'th draw of $k\times$ 5000 draws generated after a burn-in sample of 5000 draws. The skip intervals were chosen to deliver desirable mixing and convergence properties of the MCMC chains. The reported CPU run times are averages across 10 different sets of model estimates, based on different MCMC chains.

Table 3. Log predictive likelihoods, 1980:Q1-2011:Q2

model	log PL
BVAR(4), 4 variables	-656.578
BVAR(4), 4 variables, independent stochastic volatility	-550.363
BVAR(4), 4 variables, common stochastic volatility	-569.269
BVAR(4), 8 variables	-1545.288
BVAR(4), 8 variables, common stochastic volatility	-1464.062

Notes: The table reports log predictive likelihoods, formed as the sum of 1-step ahead likelihoods, over the period 1980:Q1 through 2011:Q2. The estimates are based on final vintage data, not real-time data.

Table 4. Real-Time Forecast RMSEs, 4-variable BVARs, 1985:Q1-2010:Q4

(RMSEs for BVAR benchmark, RMSE ratios in all others)

	$h = 1Q$	$h = 2Q$	$h = 1Y$	$h = 2Y$
AR				
GDP growth	0.777 ***	0.779 ***	0.754 **	1.042
Unemployment	0.752 **	0.729 *	0.742 *	0.867
GDP inflation	0.984	0.931 **	0.875 ***	0.770 ***
Fed funds rate	1.025	1.158	1.188	1.154
BVAR				
GDP growth	2.653	2.823	2.206	1.730
Unemployment	0.258	0.521	1.042	1.590
GDP inflation	1.118	1.235	1.031	1.780
Fed funds rate	0.492	0.864	1.479	2.465
BVAR with independent stochastic volatilities				
GDP growth	0.908 ***	0.908 ***	0.899 **	1.005
Unemployment	0.948 ***	0.932 **	0.929 *	0.975
GDP inflation	0.939 ***	0.913 ***	0.838 ***	0.791 ***
Fed funds rate	0.905 ***	0.936 *	0.953	0.945 *
BVAR with common stochastic volatility				
GDP growth	0.881 ***	0.881 ***	0.867 **	1.036
Unemployment	0.877 ***	0.868 **	0.882 *	0.960
GDP inflation	0.930 ***	0.875 ***	0.778 ***	0.725 ***
Fed funds rate	0.984	0.987	0.957	0.926 **

Notes: For the forecasts from AR models, the BVAR with independent stochastic volatilities, and the BVAR with common stochastic volatility, entries less than 1 indicate the model has a lower RMSE than the benchmark. To provide a rough gauge of whether the RMSE ratios are significantly different from 1, we use the Diebold-Mariano t-statistic for equal MSE, applied to the forecast of each model relative to the benchmark. Differences in accuracy that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h-1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Table 5. Real-Time Forecast RMSEs, 8-variable BVARs, 1985:Q1-2010:Q4

(RMSEs for BVAR benchmark, RMSE ratios in all others)

	$h = 1Q$	$h = 2Q$	$h = 1Y$	$h = 2Y$
AR				
GDP growth	0.872 **	0.847 ***	0.843 **	1.037
Consumption	0.848 ***	0.874 ***	0.753 ***	0.941 ***
BFI	1.081	1.050	1.085	1.157
Employment	0.854 ***	0.810 ***	0.814 **	0.943
Unemployment	0.824 **	0.805 *	0.816	0.894
GDP inflation	0.951 *	0.899 ***	0.852 **	0.797 ***
Treasury yield	0.950	1.044	1.232	1.420
Fed funds rate	0.863	1.063	1.249	1.292
BVAR				
GDP growth	2.364	2.598	1.976	1.738
Consumption	2.484	2.425	1.847	1.713
BFI	8.630	9.537	6.732	6.971
Employment	1.118	1.723	1.631	2.010
Unemployment	0.236	0.471	0.947	1.541
GDP inflation	1.157	1.279	1.058	1.721
Treasury yield	0.445	0.685	0.923	1.108
Fed funds rate	0.584	0.941	1.407	2.202
BVAR with common stochastic volatility				
GDP growth	0.960 *	0.940 **	0.931 *	1.028
Consumption	0.964 **	0.971 *	0.942 *	1.038
BFI	0.991	0.993	1.000	1.013
Employment	0.867 ***	0.870 ***	0.872 **	0.957
Unemployment	0.931 **	0.921 *	0.923 *	0.968
GDP inflation	0.956 ***	0.904 ***	0.831 ***	0.766 ***
Treasury yield	0.991	1.032	1.031	0.979
Fed funds rate	1.002	1.028	0.993	0.960

Notes: For the forecasts from AR models and the BVAR with common stochastic volatility, entries less than 1 indicate the model has a lower RMSE than the benchmark. To provide a rough gauge of whether the RMSE ratios are significantly different from 1, we use the Diebold-Mariano t-statistic for equal MSE, applied to the forecast of each model relative to the benchmark. Differences in accuracy that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h-1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Table 6. Average log predictive scores, 4-variable BVARs, 1985:Q1-2010:Q4

(avg. score for BVAR benchmark, differences in scores in all others)

	$h = 1Q$	$h = 2Q$	$h = 1Y$	$h = 2Y$
AR				
All variables	-8.059	-8.872	-8.004	-8.205
GDP growth	0.067 *	0.071	0.270	-0.074
Unemployment	0.185 **	0.309	0.751	1.145
GDP inflation	0.006	0.038 **	0.072 **	0.203 ***
Fed funds rate	-0.020	-0.133	-0.142	-0.138
BVAR				
All variables	-4.916	-6.384	-7.410	-9.570
GDP growth	-2.439	-2.498	-2.333	-2.052
Unemployment	-0.084	-0.865	-1.990	-2.946
GDP inflation	-1.549	-1.663	-1.480	-2.004
Fed funds rate	-1.127	-1.526	-1.941	-2.368
BVAR with independent stochastic volatilities				
All variables	0.810 ***	0.690 **	0.633	-0.166
GDP growth	0.149 ***	0.080	-0.062	-0.180
Unemployment	0.187 ***	0.147	-0.098	-0.639
GDP inflation	0.089 ***	0.109 ***	0.186 ***	0.196 ***
Fed funds rate	0.504 ***	0.261 **	0.010	-0.101
BVAR with common stochastic volatility				
All variables	0.678 ***	0.739 ***	0.704 **	0.165
GDP growth	0.196 ***	0.132 *	-0.070	-0.173
Unemployment	0.230 ***	0.207 **	0.076	-0.314
GDP inflation	0.090 ***	0.124 ***	0.222 ***	0.266 ***
Fed funds rate	0.267 ***	0.191 ***	0.088	0.000

Notes: For the forecasts from AR models, the BVAR with independent stochastic volatilities, and the BVAR with common stochastic volatility, entries greater than 0 indicate the model has a better average log score (better density forecast) than the benchmark model. To provide a rough gauge of the statistical significance of differences in average log scores, we use the Amisano-Giacomini t-test of equal means, applied to the log score for each model relative to the benchmark of the local level-SV forecast. Differences in average scores that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h - 1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Table 7. Average log predictive scores, 8-variable BVARs, 1985:Q1-2010:Q4

(avg. score for benchmark BVAR, differences in scores in all others)

	$h = 1Q$	$h = 2Q$	$h = 1Y$	$h = 2Y$
AR				
All variables	-0.795	-1.096	-0.867	-0.392
GDP growth	-0.011	0.017	0.097	-0.059
Consumption	0.116 ***	0.095 ***	0.310 **	0.047 *
BFI	-0.081	-0.042	-0.013	-0.107
Employment	0.051 *	0.144	0.410 *	0.191
Unemployment	0.109	0.182	0.468	0.789
GDP inflation	0.041 **	0.069 ***	0.096 ***	0.167 ***
Treasury yield	0.036	-0.063	-0.200	-0.292
Fed funds rate	-0.014	-0.098	-0.154	-0.214
BVAR				
All variables	-12.180	-14.161	-14.547	-17.384
GDP growth	-2.362	-2.443	-2.160	-2.067
Consumption	-2.334	-2.328	-2.131	-1.996
BFI	-3.583	-3.717	-3.428	-3.442
Employment	-1.564	-2.013	-2.190	-2.308
Unemployment	-0.007	-0.738	-1.708	-2.590
GDP inflation	-1.583	-1.694	-1.503	-1.969
Treasury yield	-0.654	-1.055	-1.360	-1.619
Fed funds rate	-1.134	-1.561	-1.930	-2.291
BVAR with common stochastic volatility				
All variables	0.449 ***	0.368 **	-0.072	-0.590
GDP growth	0.100 **	0.074	-0.120	-0.118
Consumption	0.025	0.012	-0.035	-0.142
BFI	0.029	-0.034	-0.137	-0.190
Employment	0.162 ***	0.111 **	0.104	-0.107
Unemployment	0.115 ***	0.056	-0.111	-0.272
GDP inflation	0.032 *	0.064 ***	0.113 ***	0.158 ***
Treasury yield	0.044 ***	-0.006	-0.017	-0.022
Fed funds rate	0.113 ***	0.067 ***	0.018	-0.014

Notes: For the forecasts from AR models and the BVAR with common stochastic volatility, entries greater than 0 indicate the model has a better average log score (better density forecast) than the benchmark model. To provide a rough gauge of the statistical significance of differences in average log scores, we use the Amisano-Giacomini t-test of equal means, applied to the log score for each model relative to the benchmark of the local level-SV forecast. Differences in average scores that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h-1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Table 8. Forecast RMSEs, 15-variable BVARs,
Koop's (2012) Data, 1985:Q1-2008:Q4
(RMSEs for BVAR benchmark, RMSE ratios for BVAR-CSV)

	$h = 1Q$	$h = 2Q$	$h = 4Q$	$h = 8Q$
BVAR				
GDP	1.966	2.330	2.408	2.327
CPI	1.960	2.172	2.257	2.284
Fed funds rate	0.510	0.604	0.525	0.515
Consumption	1.901	2.137	2.203	2.170
Industrial production	3.218	4.173	4.330	4.239
Capacity Utilization	0.816	1.570	2.590	3.441
Unemployment rate	0.164	0.204	0.238	0.234
Housing starts	29.436	42.943	62.535	77.052
PPI for finished goods	4.308	4.639	4.739	4.836
PCE price index	1.421	1.538	1.606	1.604
Real average hourly earnigs	1.731	1.759	1.859	1.835
S&P stock price index, industrials	27.485	29.206	29.434	27.643
10-year Treasury bond yield	0.461	0.476	0.436	0.399
Effective exchange rate	13.295	13.459	13.084	12.516
Payroll employment	0.884	1.391	1.900	2.014
BVAR with common stochastic volatility				
GDP	0.989	0.977 **	0.985	1.008
CPI	1.009	1.008	1.001	0.997
Fed funds rate	0.908 ***	0.954 **	0.991	1.013
Consumption	1.000	0.979 *	0.996	1.010
Industrial production	0.982	0.997	0.985	1.010
Capacity Utilization	0.981 *	0.983	0.975 ***	0.977
Unemployment rate	0.977 *	0.983 **	0.980 *	1.010
Housing starts	1.006	0.998	1.007	1.028
PPI for finished goods	0.994	1.002	1.003	0.997
PCE price index	0.995	0.999	0.994	0.995
Real average hourly earnigs	0.977	0.997	0.994	1.007
S&P stock price index, industrials	0.989	0.979 *	0.995	1.016
10-year Treasury bond yield	0.977 **	0.991	1.007	1.001
Effective exchange rate	1.001	1.002	1.003	1.007
Payroll employment	0.963 **	0.966 **	0.957 *	0.992

Notes: For the forecasts from the BVAR with common stochastic volatility, entries less than 1 indicate the model has a lower RMSE than the benchmark BVAR. To provide a rough gauge of whether the RMSE ratios are significantly different from 1, we use the Diebold-Mariano t-statistic for equal MSE, applied to the forecast of each model relative to the benchmark. Differences in accuracy that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h - 1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Table 9. Average log predictive scores, 15-variable BVARs, Koop's (2012) Data, 1985:Q1-2008:Q4

(avg. score for benchmark BVAR, differences in scores for BVAR-CSV)

	$h = 1Q$	$h = 2Q$	$h = 1Y$	$h = 2Y$
BVAR				
All variables	-30.327	-32.283	-33.723	-34.536
GDP	-2.233	-2.357	-2.433	-2.442
CPI	-2.187	-2.304	-2.302	-2.311
Fed funds rate	-0.948	-1.116	-1.166	-1.202
Consumption	-2.094	-2.219	-2.270	-2.283
Industrial production	-2.689	-2.902	-2.999	-3.022
Capacity Utilization	-1.290	-1.879	-2.397	-2.737
Unemployment rate	0.302	0.095	-0.071	-0.097
Housing starts	-4.829	-5.222	-5.618	-5.836
PPI for finished goods	-2.973	-3.065	-3.056	-3.082
PCE price index	-1.825	-1.926	-1.951	-1.947
Real average hourly earnigs	-1.987	-2.006	-2.059	-2.061
S&P stock price index, industrials	-4.772	-4.832	-4.838	-4.758
10-year Treasury bond yield	-0.643	-0.689	-0.651	-0.627
Effective exchange rate	-4.116	-4.084	-4.045	-3.957
Payroll employment	-1.460	-1.819	-2.116	-2.201
BVAR with common stochastic volatility				
All variables	0.593 ***	0.411 ***	0.108	-0.331
GDP	0.145 ***	0.101 **	0.077 *	0.083 *
CPI	0.003	-0.035	-0.055	-0.079
Fed funds rate	0.073 ***	0.072 ***	0.036	0.004
Consumption	0.046 ***	0.034 **	-0.005	-0.015
Industrial production	0.072 ***	0.034	-0.010	-0.021
Capacity Utilization	0.081 ***	0.029	-0.033	-0.117
Unemployment rate	0.038 ***	0.034 *	-0.007	-0.025
Housing starts	0.009	-0.017	-0.083	-0.110
PPI for finished goods	0.069	0.045	0.002	-0.028
PCE price index	0.018	-0.015	-0.021	-0.045
Real average hourly earnigs	0.005	-0.036	-0.033	-0.050
S&P stock price index, industrials	0.037	0.013	0.000	-0.060
10-year Treasury bond yield	-0.006	-0.013	-0.053	-0.090
Effective exchange rate	0.128 ***	0.084 *	0.062	-0.014
Payroll employment	0.079 ***	0.067 **	-0.010	-0.061

Notes: For the forecasts from the BVAR with common stochastic volatility, entries greater than 0 indicate the model has a better average log score (better density forecast) than the benchmark BVAR model. To provide a rough gauge of the statistical significance of differences in average log scores, we use the Amisano-Giacomini t-test of equal means, applied to the log score for each model relative to the benchmark of the local level-SV forecast. Differences in average scores that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h - 1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Table 10. Log predictive likelihoods for UK models, 1985:Q1-2011:Q2

model	log PL
BVAR(4), 8 variables	-1541.702
BVAR(4), 8 variables, common stochastic volatility	-1471.830

Notes: The table reports log predictive likelihoods, formed as the sum of 1-step ahead likelihoods, over the period 1985:Q1 through 2011:Q2.

Table 11. Forecast RMSEs, 8-variable BVARs, UK data, 1985:Q1-2011:Q2

(RMSEs for BVAR benchmark, RMSE ratios in all others)

	$h = 1Q$	$h = 2Q$	$h = 1Y$	$h = 2Y$
AR				
GDP growth	0.933 *	0.966	0.997	1.057
Consumption	1.008	1.026	1.038	1.064
Investment	1.062	1.009	1.045	1.066
Employment	1.073	1.023	1.048	1.051
Unemployment	0.925	0.910	0.898	0.907
GDP inflation	0.927 *	0.949	0.840 ***	0.741 **
Bond yield	0.961	0.981	0.928	0.830 *
Bank rate	1.009	1.036	1.047	1.019
BVAR				
GDP growth	2.731	2.732	2.244	2.279
Consumption	3.322	3.173	2.527	2.569
Investment	10.520	10.659	6.149	6.368
Employment	1.102	1.218	1.037	1.417
Unemployment	0.204	0.372	0.778	1.526
GDP inflation	2.875	2.834	2.249	2.713
Bond yield	0.451	0.704	1.002	1.598
Bank rate	0.764	1.197	1.842	3.028
BVAR with common stochastic volatility				
GDP growth	0.885 **	0.944 **	0.942 *	1.043
Consumption	0.951 *	0.994	0.995	1.066
Investment	0.968	1.013	1.020	1.055
Employment	0.967	0.959	0.948	1.038
Unemployment	0.881 ***	0.849 **	0.857	0.888
GDP inflation	1.013	1.022	1.050	1.109
Bond yield	0.996	0.987	0.930 **	0.833 ***
Bank rate	1.033	1.017	0.977	0.960

Notes: For the forecasts from AR models and the BVAR with common stochastic volatility, entries less than 1 indicate the model has a lower RMSE than the benchmark. To provide a rough gauge of whether the RMSE ratios are significantly different from 1, we use the Diebold-Mariano t-statistic for equal MSE, applied to the forecast of each model relative to the benchmark. Differences in accuracy that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h - 1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Table 12. Average log predictive scores, 8-variable BVARs, UK data, 1985:Q1-2011:Q2

(avg. score for benchmark BVAR, differences in scores in all others)

	$h = 1Q$	$h = 2Q$	$h = 1Y$	$h = 2Y$
AR				
All variables	-0.809	-0.734	-1.757	-1.422
GDP growth	0.056 ***	0.064 ***	0.007	-0.039
Consumption	0.036 ***	0.042 ***	-0.017	-0.052
Investment	-0.072	-0.018	-0.054	-0.077
Employment	-0.079	-0.030	-0.024	0.110
Unemployment	0.035	0.047	0.162	0.357
GDP inflation	0.066 ***	0.124 ***	0.206 ***	0.290 ***
Bond yield	0.045 ***	0.019	0.039 *	0.078 ***
Bank rate	-0.007	-0.027	-0.041	0.005
BVAR				
All variables	-14.581	-16.247	-15.273	-18.050
GDP growth	-2.633	-2.668	-2.306	-2.332
Consumption	-2.803	-2.806	-2.393	-2.410
Investment	-3.758	-3.790	-3.258	-3.300
Employment	-1.545	-1.618	-1.496	-1.960
Unemployment	0.187	-0.399	-1.237	-2.114
GDP inflation	-2.745	-2.868	-2.739	-3.272
Bond yield	-0.837	-1.261	-1.645	-2.076
Bank rate	-1.340	-1.799	-2.224	-2.687
BVAR with common stochastic volatility				
All variables	0.637 ***	0.612 ***	-0.051	-0.007
GDP growth	0.300 ***	0.210	-0.349	-0.447
Consumption	0.156 ***	0.129 ***	-0.095	-0.156
Investment	0.023	-0.005	-0.095	-0.152
Employment	0.028	0.051	0.085	0.072
Unemployment	0.098 **	0.120 *	0.172	0.322
GDP inflation	0.109 ***	0.143 ***	0.171 ***	0.170 ***
Bond yield	0.079 ***	0.057 ***	0.056 *	0.072
Bank rate	0.058 ***	0.056 **	0.041	0.019

Notes: For the forecasts from AR models and the BVAR with common stochastic volatility, entries greater than 0 indicate the model has a better average log score (better density forecast) than the benchmark model. To provide a rough gauge of the statistical significance of differences in average log scores, we use the Amisano-Giacomini t-test of equal means, applied to the log score for each model relative to the benchmark of the local level-SV forecast. Differences in average scores that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h - 1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Appendix Table A4. Real-Time Forecast RMSEs, 4-variable BVARs, 1985:Q1-2010:Q4

(RMSEs for benchmark AR, RMSE ratios in all others)

	$h = 1Q$	$h = 2Q$	$h = 1Y$	$h = 2Y$
AR				
GDP growth	2.062	2.200	1.665	1.804
Unemployment	0.194	0.380	0.773	1.379
GDP inflation	1.100	1.150	0.902	1.371
Fed funds rate	0.504	1.000	1.757	2.845
BVAR				
GDP growth	1.287	1.283	1.325	0.959 **
Unemployment	1.329	1.371	1.349	1.153
GDP inflation	1.016	1.074	1.143	1.298
Fed funds rate	0.976	0.863 **	0.842 ***	0.866 ***
BVAR with independent stochastic volatilities				
GDP growth	1.168	1.165	1.191	0.964
Unemployment	1.260	1.278	1.253	1.124
GDP inflation	0.954 **	0.981	0.958	1.027
Fed funds rate	0.883 *	0.809 **	0.802 ***	0.818 ***
BVAR with common stochastic volatility				
GDP growth	1.134	1.130	1.149	0.994
Unemployment	1.165	1.190	1.189	1.108
GDP inflation	0.945 **	0.939 **	0.890 **	0.941
Fed funds rate	0.960	0.852 *	0.805 **	0.802 ***

Notes: For the forecasts from BVAR models, entries less than 1 indicate the model has a lower RMSE than the benchmark. To provide a rough gauge of whether the RMSE ratios are significantly different from 1, we use the Diebold-Mariano t-statistic for equal MSE, applied to the forecast of each model relative to the benchmark. Differences in accuracy that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h-1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Appendix Table A5. Real-Time Forecast RMSEs, 8-variable BVARs, 1985:Q1-2010:Q4

(RMSEs for benchmark AR, RMSE ratios in all others)

	$h = 1Q$	$h = 2Q$	$h = 1Y$	$h = 2Y$
AR				
GDP growth	2.062	2.200	1.665	1.804
Consumption	2.106	2.120	1.390	1.611
BFI	9.333	10.013	7.302	8.068
Employment	0.954	1.396	1.329	1.895
Unemployment	0.194	0.380	0.773	1.379
GDP inflation	1.100	1.150	0.902	1.371
Treasury yield	0.422	0.715	1.137	1.573
Fed funds rate	0.504	1.000	1.757	2.845
BVAR				
GDP growth	1.146	1.181	1.187	0.964 *
Consumption	1.180	1.144	1.329	1.063
BFI	0.925 **	0.952	0.922	0.864 **
Employment	1.171	1.235	1.228	1.061
Unemployment	1.214	1.242	1.225	1.118
GDP inflation	1.052	1.112	1.173	1.255
Treasury yield	1.053	0.958	0.812 **	0.704 ***
Fed funds rate	1.158	0.941	0.801 **	0.774 **
BVAR with common stochastic volatility				
GDP growth	1.101	1.110	1.106	0.991
Consumption	1.138	1.111	1.252	1.103
BFI	0.916 ***	0.946 *	0.922	0.875 **
Employment	1.016	1.074	1.071	1.015
Unemployment	1.130	1.143	1.130	1.082
GDP inflation	1.005	1.005	0.975	0.961
Treasury yield	1.044	0.989	0.837 *	0.689 ***
Fed funds rate	1.161	0.968	0.795 **	0.743 **

Notes: For the forecasts from BVAR models, entries less than 1 indicate the model has a lower RMSE than the benchmark. To provide a rough gauge of whether the RMSE ratios are significantly different from 1, we use the Diebold-Mariano t-statistic for equal MSE, applied to the forecast of each model relative to the benchmark. Differences in accuracy that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h-1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Appendix Table A6. Average log predictive scores, 4-variable BVARs, 1985:Q1-2010:Q4

(avg. score for benchmark AR, differences in scores in all others)

	$h = 1Q$	$h = 2Q$	$h = 1Y$	$h = 2Y$
AR				
All variables	-12.975	-15.256	-15.414	-17.775
GDP growth	-2.372	-2.426	-2.063	-2.126
Unemployment	0.101	-0.556	-1.239	-1.801
GDP inflation	-1.542	-1.625	-1.407	-1.801
Fed funds rate	-1.147	-1.659	-2.084	-2.505
BVAR				
All variables	8.059 ***	8.872 ***	8.004 ***	8.205 ***
GDP growth	-0.067	-0.071	-0.270	0.074
Unemployment	-0.185	-0.309	-0.751	-1.145
GDP inflation	-0.006	-0.038	-0.072	-0.203
Fed funds rate	0.020	0.133 ***	0.142 ***	0.138 ***
BVAR with independent stochastic volatilities				
All variables	8.869 ***	9.562 ***	8.638 ***	8.039 ***
GDP growth	0.082 *	0.008	-0.332	-0.106
Unemployment	0.002	-0.161	-0.849	-1.784
GDP inflation	0.082 ***	0.071 ***	0.114 **	-0.007
Fed funds rate	0.524 ***	0.395 ***	0.153	0.037
BVAR with common stochastic volatility				
All variables	8.737 ***	9.612 ***	8.709 ***	8.370 ***
GDP growth	0.129 **	0.061	-0.339	-0.098
Unemployment	0.045	-0.101	-0.675	-1.459
GDP inflation	0.084 ***	0.086 ***	0.149 ***	0.063 **
Fed funds rate	0.287 ***	0.325 ***	0.230 ***	0.138 **

Notes: For the forecasts from BVAR models, entries greater than 0 indicate the model has a better average log score (better density forecast) than the benchmark model. To provide a rough gauge of the statistical significance of differences in average log scores, we use the Amisano-Giacomini t-test of equal means, applied to the log score for each model relative to the benchmark of the local level-SV forecast. Differences in average scores that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h-1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Appendix Table A7. Average log predictive scores, 8-variable BVARs, 1985:Q1-2010:Q4

(avg. score for benchmark AR, differences in scores in all others)

	$h = 1Q$	$h = 2Q$	$h = 1Y$	$h = 2Y$
AR				
All variables	-12.975	-15.256	-15.414	-17.775
GDP growth	-2.372	-2.426	-2.063	-2.126
Consumption	-2.218	-2.234	-1.821	-1.948
BFI	-3.664	-3.760	-3.441	-3.548
Employment	-1.513	-1.869	-1.781	-2.117
Unemployment	0.101	-0.556	-1.239	-1.801
GDP inflation	-1.542	-1.625	-1.407	-1.801
Treasury yield	-0.618	-1.118	-1.561	-1.912
Fed funds rate	-1.147	-1.659	-2.084	-2.505
BVAR				
All variables	0.795 ***	1.096 **	0.867	0.392
GDP growth	0.011	-0.017	-0.097	0.059 *
Consumption	-0.116	-0.095	-0.310	-0.047
BFI	0.081 **	0.042	0.013	0.107
Employment	-0.051	-0.144	-0.410	-0.191
Unemployment	-0.109	-0.182	-0.468	-0.789
GDP inflation	-0.041	-0.069	-0.096	-0.167
Treasury yield	-0.036	0.063	0.200 ***	0.292 ***
Fed funds rate	0.014	0.098 **	0.154 ***	0.214 ***
BVAR with common stochastic volatility				
All variables	1.244 ***	1.464 ***	0.795	-0.199
GDP growth	0.111 *	0.058	-0.217	-0.059
Consumption	-0.091	-0.082	-0.345	-0.190
BFI	0.110 **	0.008	-0.124	-0.084
Employment	0.112 ***	-0.033	-0.306	-0.297
Unemployment	0.006	-0.126	-0.580	-1.060
GDP inflation	-0.009	-0.004	0.017	-0.009
Treasury yield	0.008	0.057	0.183 **	0.270 ***
Fed funds rate	0.127 ***	0.165 ***	0.172 ***	0.200 **

Notes: For the forecasts from BVAR models, entries greater than 0 indicate the model has a better average log score (better density forecast) than the benchmark model. To provide a rough gauge of the statistical significance of differences in average log scores, we use the Amisano-Giacomini t-test of equal means, applied to the log score for each model relative to the benchmark of the local level-SV forecast. Differences in average scores that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h-1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Appendix Table A11. Forecast RMSEs, 8-variable BVARs, UK data, 1985:Q1-2011:Q2

(RMSEs for AR benchmark, RMSE ratios in all others)

	$h = 1Q$	$h = 2Q$	$h = 1Y$	$h = 2Y$
AR				
GDP growth	2.547	2.640	2.238	2.408
Consumption	3.349	3.257	2.623	2.733
Investment	11.167	10.755	6.424	6.786
Employment	1.183	1.245	1.086	1.488
Unemployment	0.189	0.339	0.698	1.383
GDP inflation	2.666	2.689	1.889	2.010
Bond yield	0.434	0.691	0.930	1.327
Bank rate	0.771	1.240	1.929	3.087
BVAR				
GDP growth	1.072	1.035	1.003	0.946
Consumption	0.992	0.974	0.963	0.940
Investment	0.942 *	0.991	0.957	0.938
Employment	0.932 **	0.978	0.955	0.952
Unemployment	1.081	1.099	1.113	1.103
GDP inflation	1.078	1.054	1.191	1.349
Bond yield	1.041	1.019	1.077	1.205
Bank rate	0.991	0.965	0.955	0.981
BVAR with common stochastic volatility				
GDP growth	0.949	0.977	0.945	0.987
Consumption	0.944	0.969	0.959	1.002
Investment	0.912 **	1.004	0.976	0.990
Employment	0.901 ***	0.937 *	0.905 **	0.988
Unemployment	0.953 *	0.933 *	0.955	0.980
GDP inflation	1.092	1.077	1.250	1.497
Bond yield	1.036	1.006	1.002	1.004
Bank rate	1.023	0.982	0.933	0.942

Notes: For the forecasts from the BVAR and the BVAR with common stochastic volatility, entries less than 1 indicate the model has a lower RMSE than the benchmark. To provide a rough gauge of whether the RMSE ratios are significantly different from 1, we use the Diebold-Mariano t-statistic for equal MSE, applied to the forecast of each model relative to the benchmark. Differences in accuracy that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h-1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

Appendix Table A12. Average log predictive scores, 8-variable BVARs, UK data, 1985:Q1-2011:Q2

(avg. score for benchmark AR, differences in scores in all others)

	$h = 1Q$	$h = 2Q$	$h = 1Y$	$h = 2Y$
AR				
All variables	-15.390	-16.982	-17.030	-19.472
GDP growth	-2.577	-2.605	-2.299	-2.371
Consumption	-2.767	-2.764	-2.411	-2.461
Investment	-3.830	-3.808	-3.312	-3.377
Employment	-1.624	-1.648	-1.520	-1.850
Unemployment	0.222	-0.352	-1.075	-1.757
GDP inflation	-2.678	-2.744	-2.533	-2.981
Bond yield	-0.792	-1.242	-1.605	-1.997
Bank rate	-1.346	-1.826	-2.265	-2.682
BVAR				
All variables	0.809 ***	0.734 ***	1.757 ***	1.422 ***
GDP growth	-0.056	-0.064	-0.007	0.039
Consumption	-0.036	-0.042	0.017	0.052
Investment	0.072 **	0.018	0.054	0.077
Employment	0.079 *	0.030	0.024	-0.110
Unemployment	-0.035	-0.047	-0.162	-0.357
GDP inflation	-0.066	-0.124	-0.206	-0.290
Bond yield	-0.045	-0.019	-0.039	-0.078
Bank rate	0.007	0.027	0.041	-0.005
BVAR with common stochastic volatility				
All variables	1.446 ***	1.346 ***	1.705 ***	1.415 ***
GDP growth	0.244 ***	0.147	-0.356	-0.409
Consumption	0.121 ***	0.087 **	-0.077	-0.105
Investment	0.095 **	0.012	-0.040	-0.075
Employment	0.107 **	0.081 **	0.109 **	-0.038
Unemployment	0.064 *	0.073	0.011	-0.035
GDP inflation	0.043 **	0.019	-0.035	-0.120
Bond yield	0.034 **	0.037 *	0.017	-0.007
Bank rate	0.065 ***	0.083 **	0.082 **	0.014

Notes: For the forecasts from the BVAR and the BVAR with common stochastic volatility, entries greater than 0 indicate the model has a better average log score (better density forecast) than the benchmark model. To provide a rough gauge of the statistical significance of differences in average log scores, we use the Amisano-Giacomini t-test of equal means, applied to the log score for each model relative to the benchmark of the local level-SV forecast. Differences in average scores that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h - 1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997).

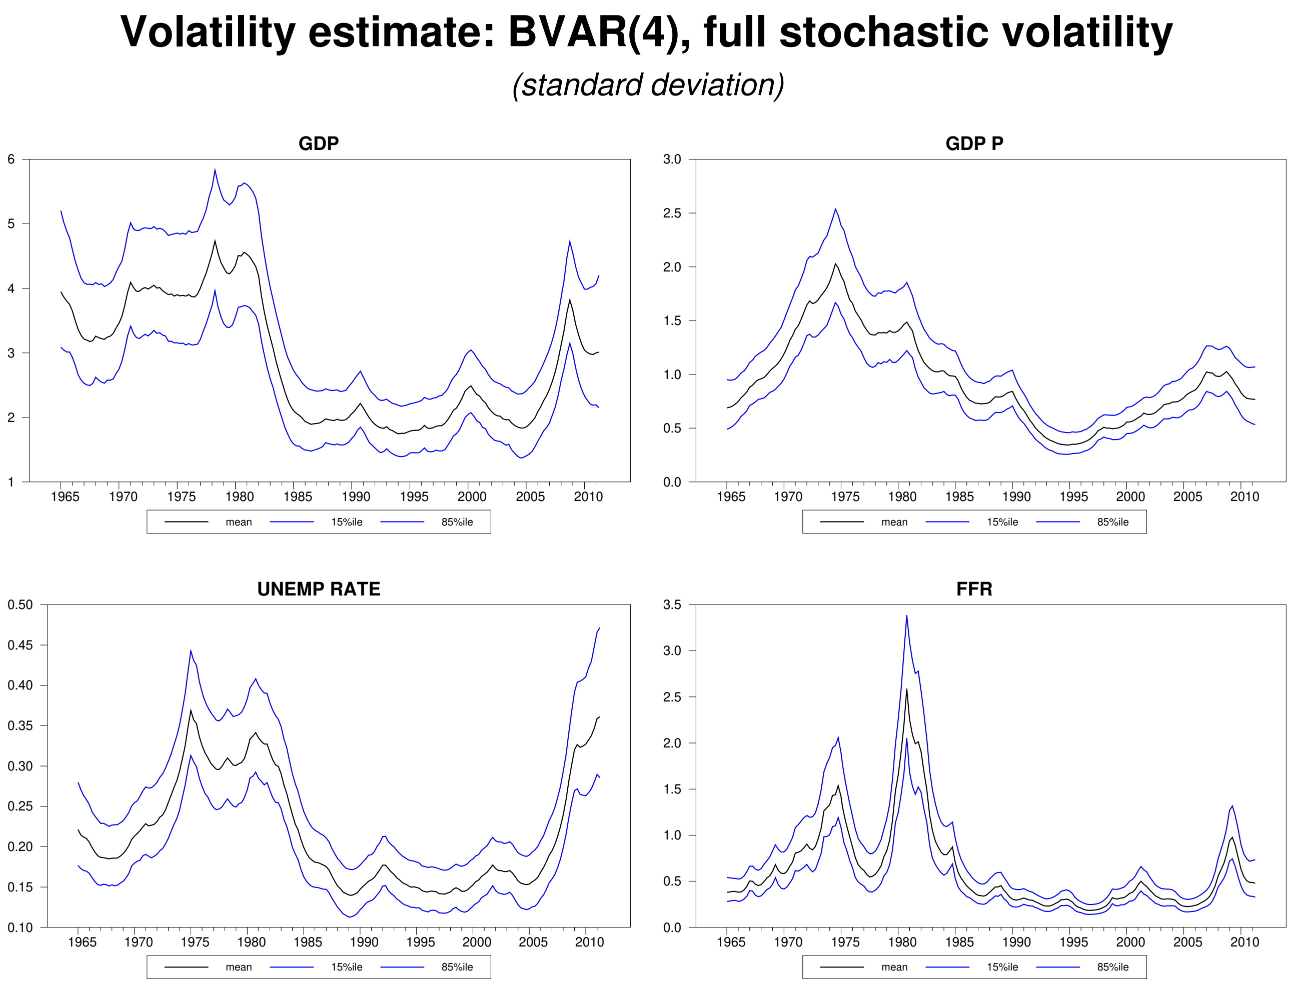

Figure 1: Volatility estimates from 4-variable BVAR with independent stochastic volatilities, final vintage data

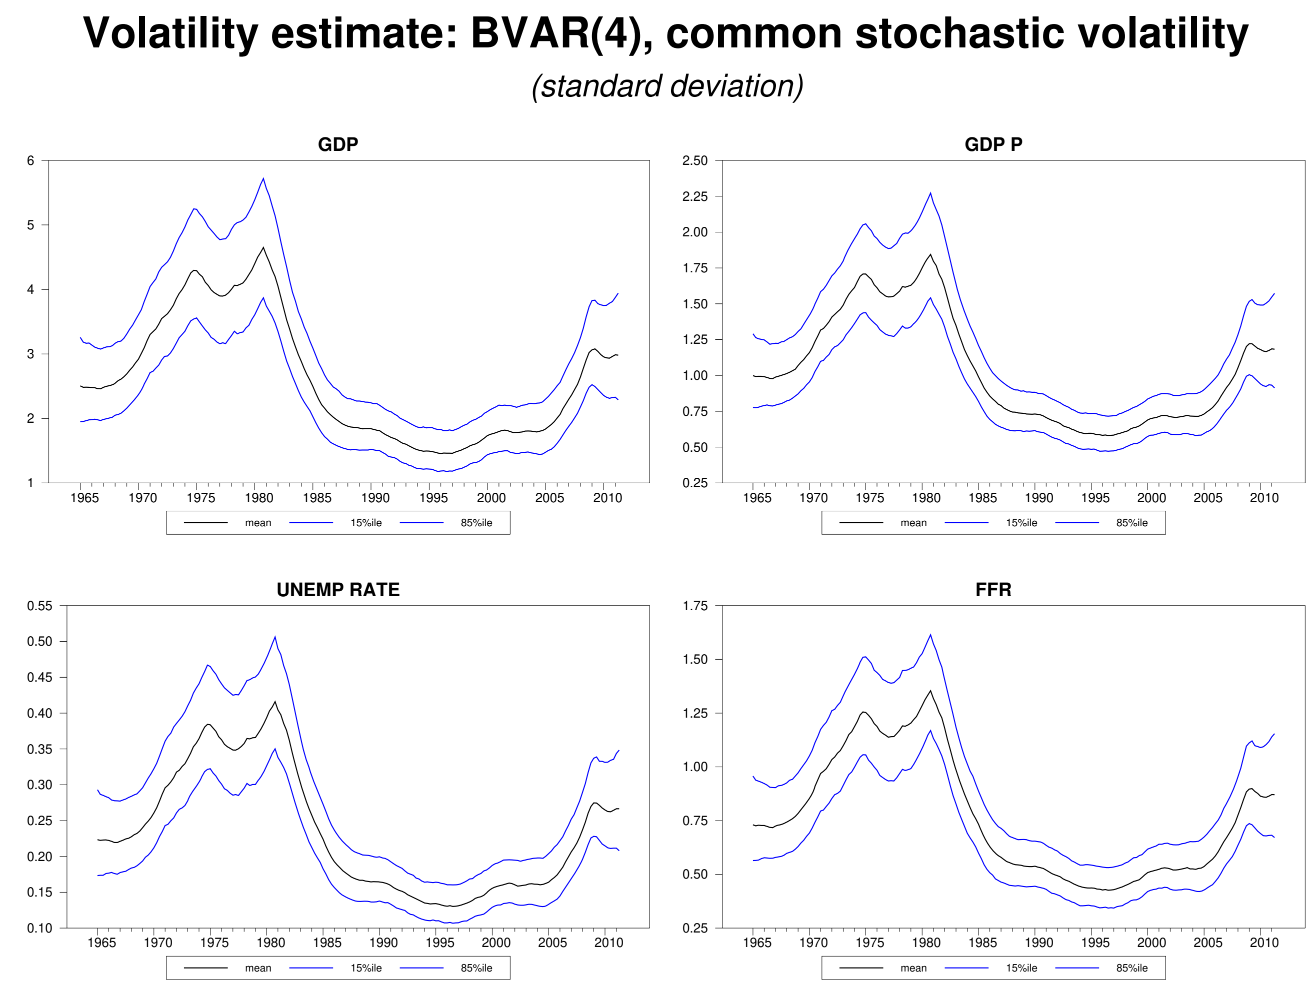

Figure 2: Volatility estimates from 4-variable BVAR with common stochastic volatility, final vintage data

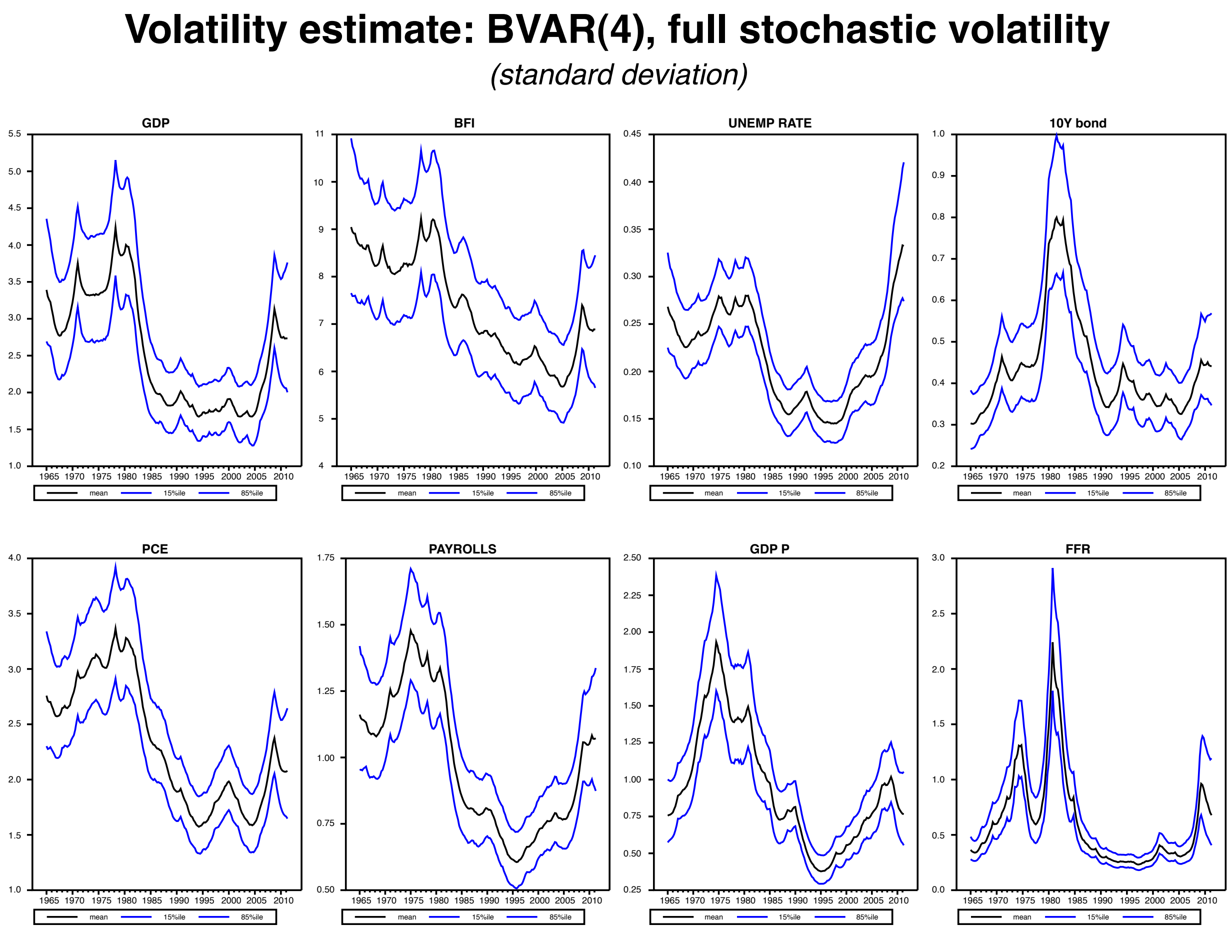

Figure 3: Volatility estimates from 8-variable BVAR with independent stochastic volatilities, final vintage data

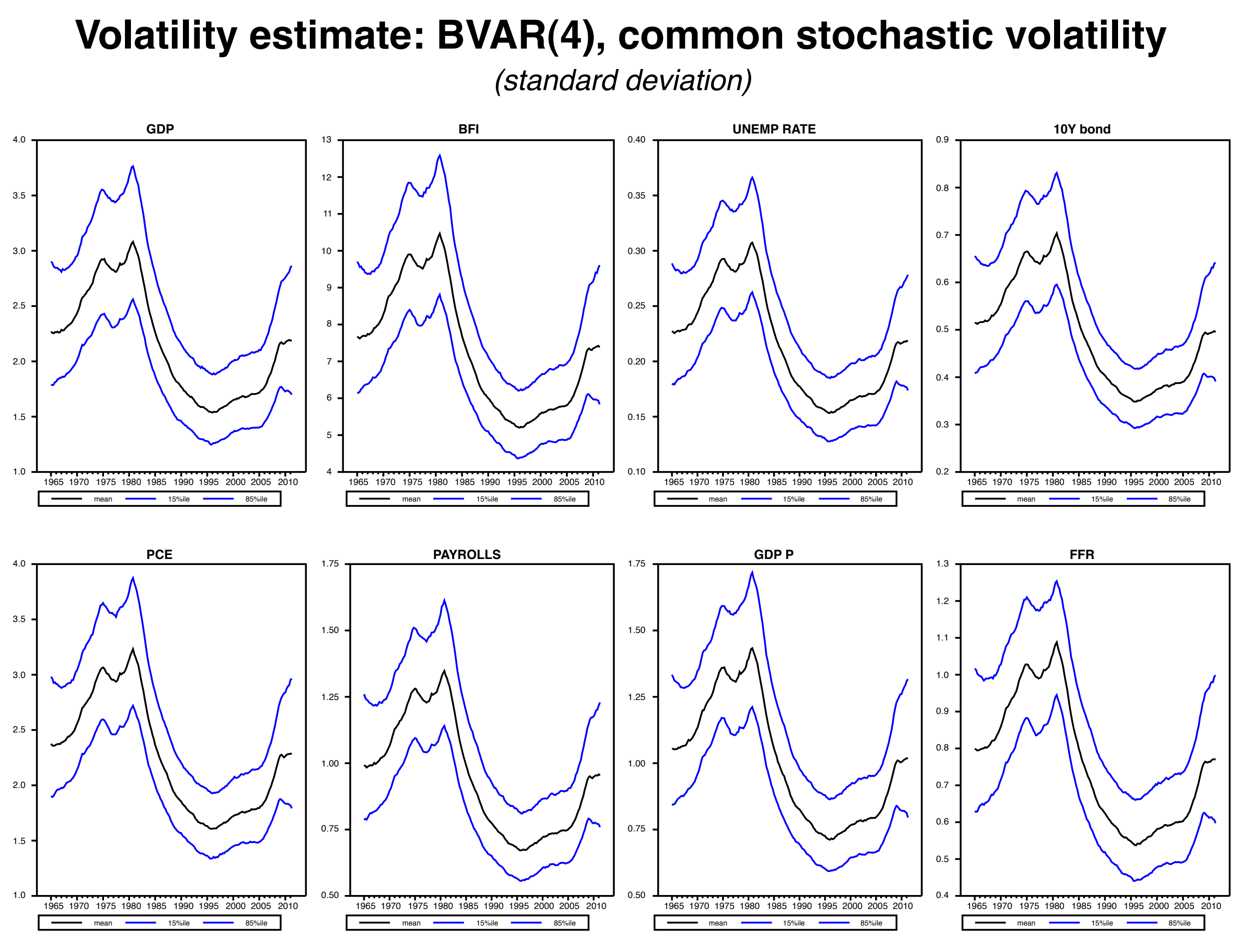

Figure 4: Volatility estimates from 8-variable BVAR with common stochastic volatility, final vintage data

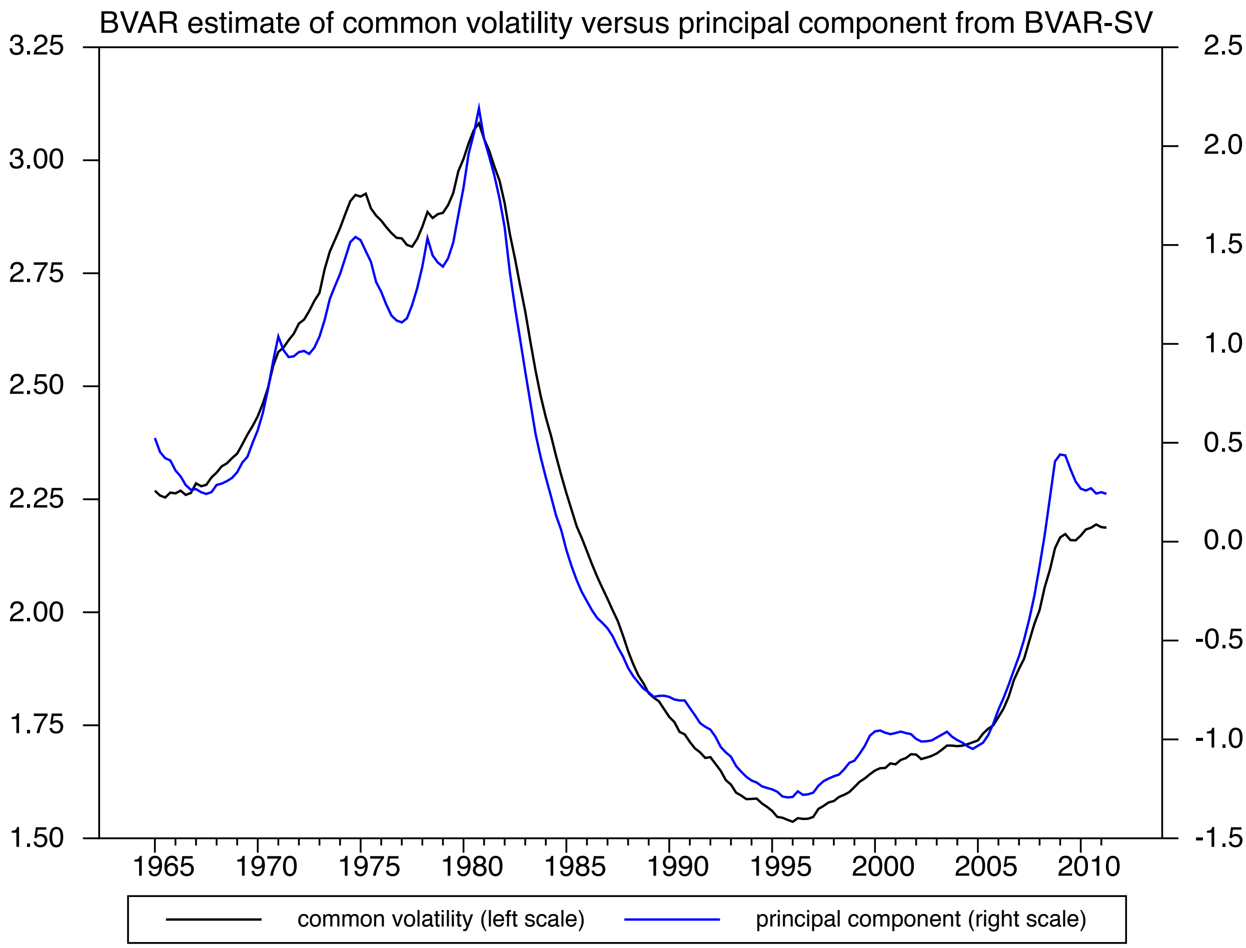

Figure 5: Principal component of BVAR-SV estimates of volatility versus common factor of volatility in 8-variable BVAR-CSV, final vintage data

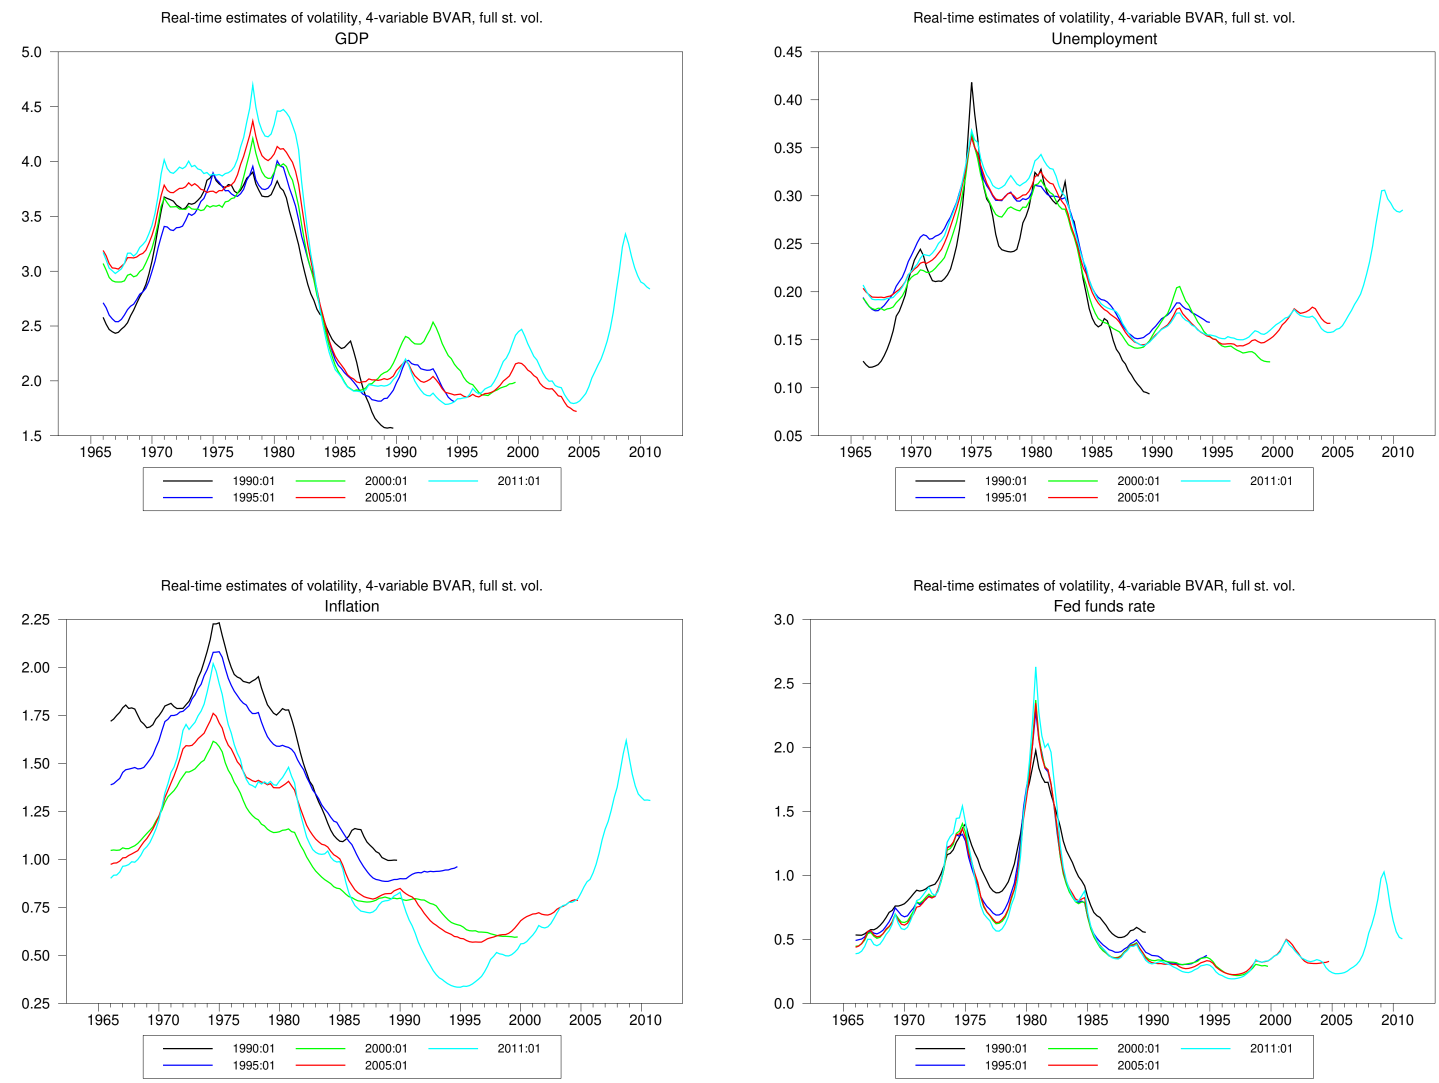

Figure 6: Volatility estimates from 4-variable BVAR with independent stochastic volatilities, real-time data

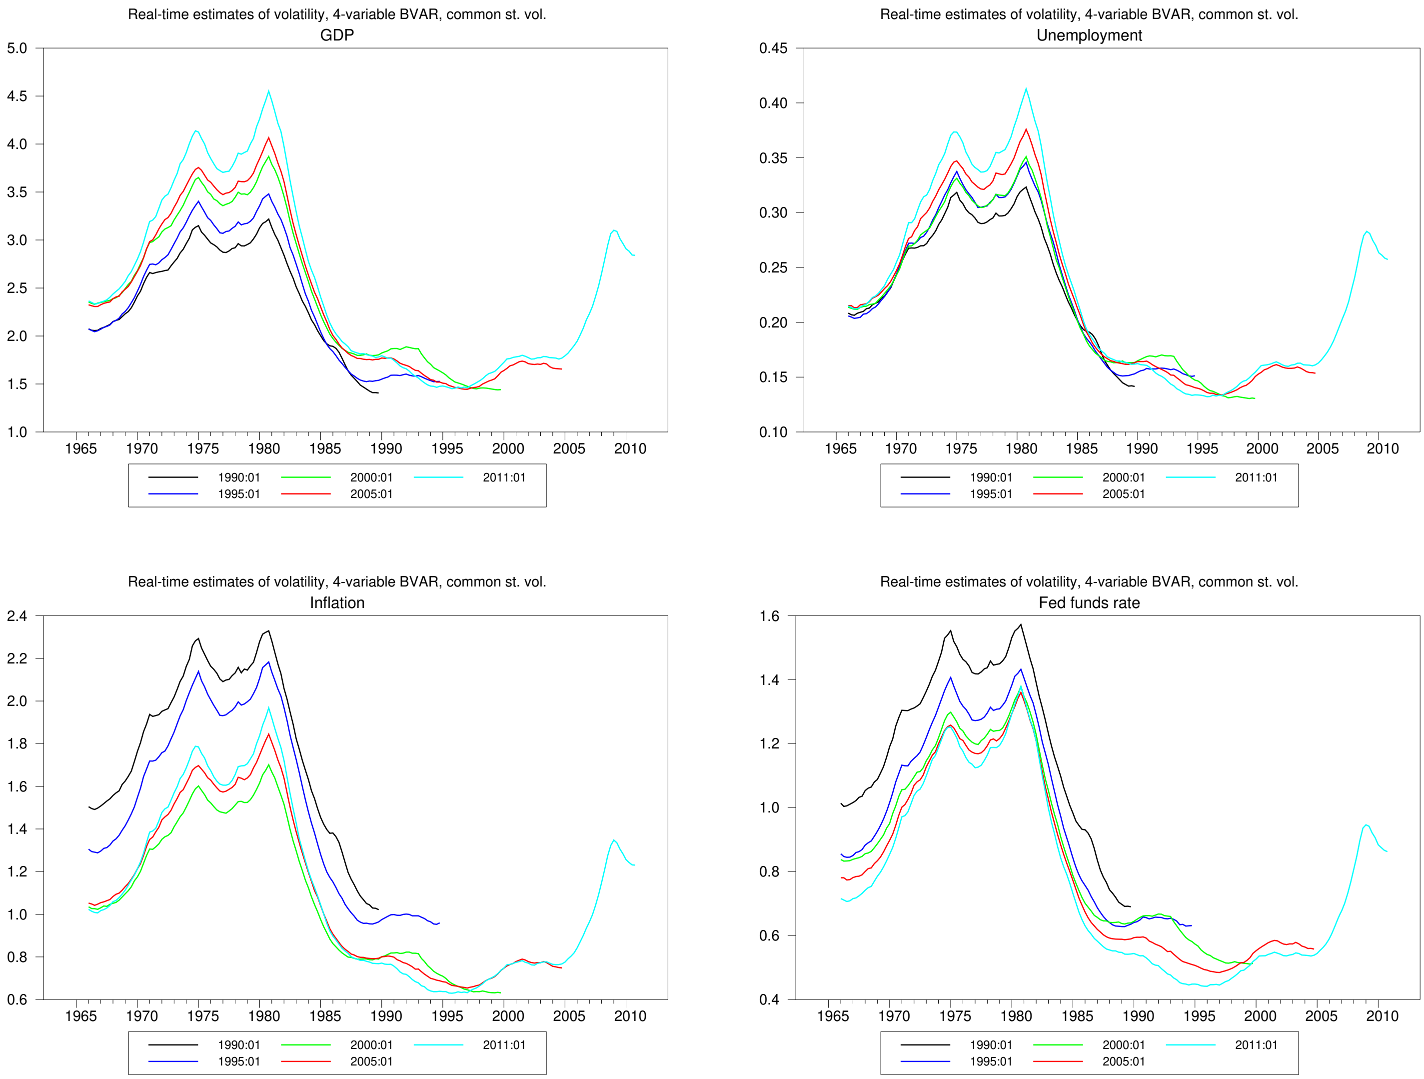

Figure 7: Volatility estimates from 4-variable BVAR with common stochastic volatility, real-time data

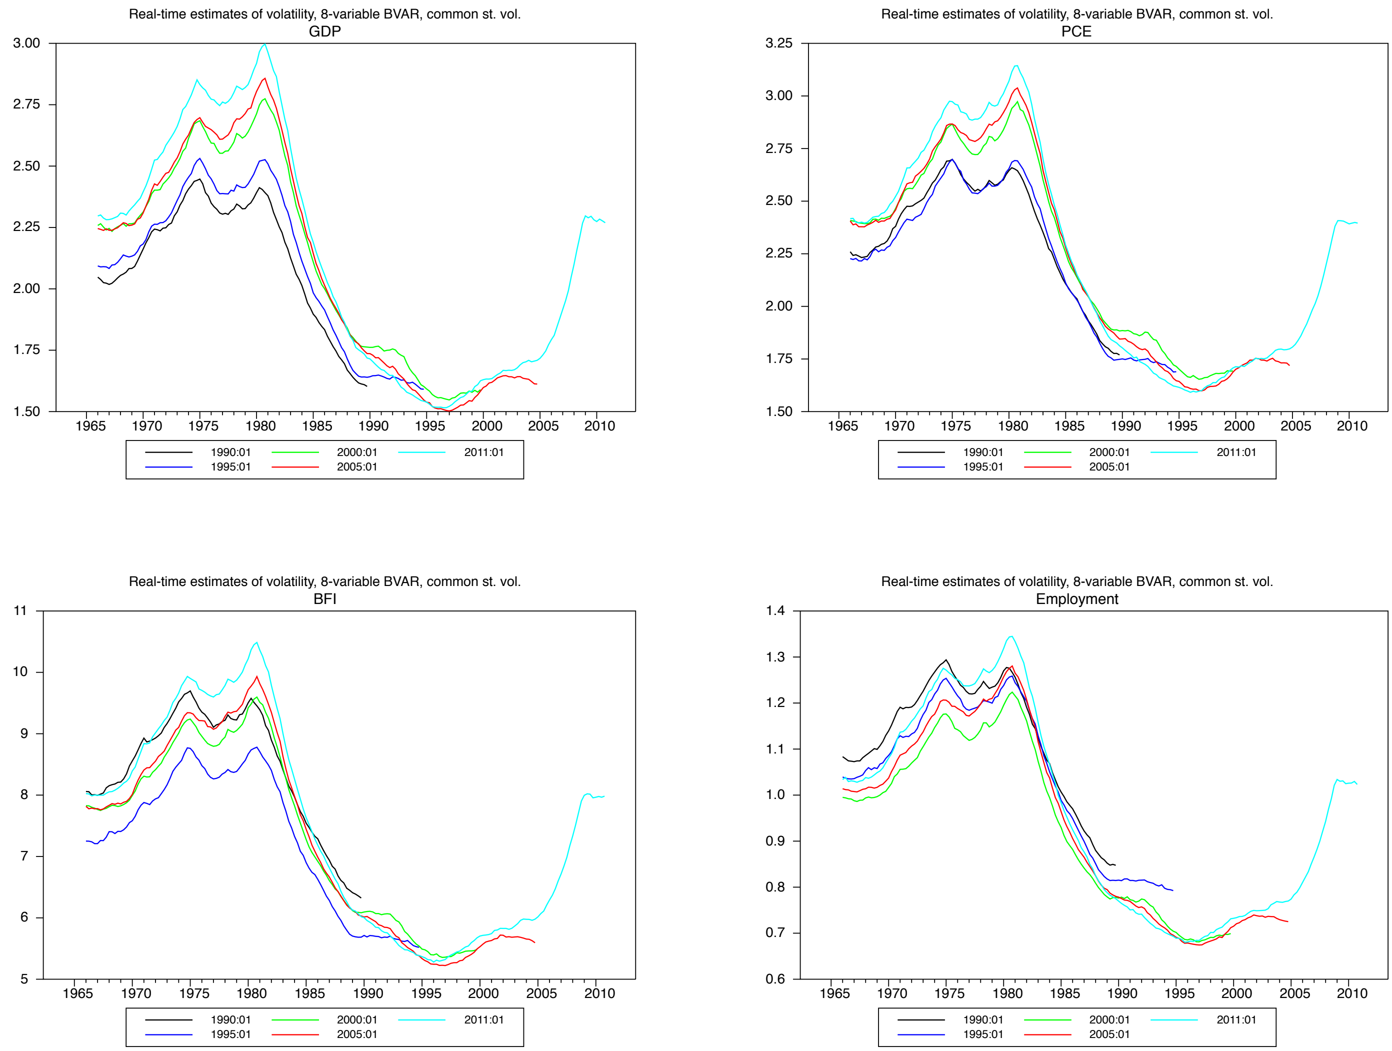

Figure 8: Volatility estimates from 8-variable BVAR with common stochastic volatility, real-time data

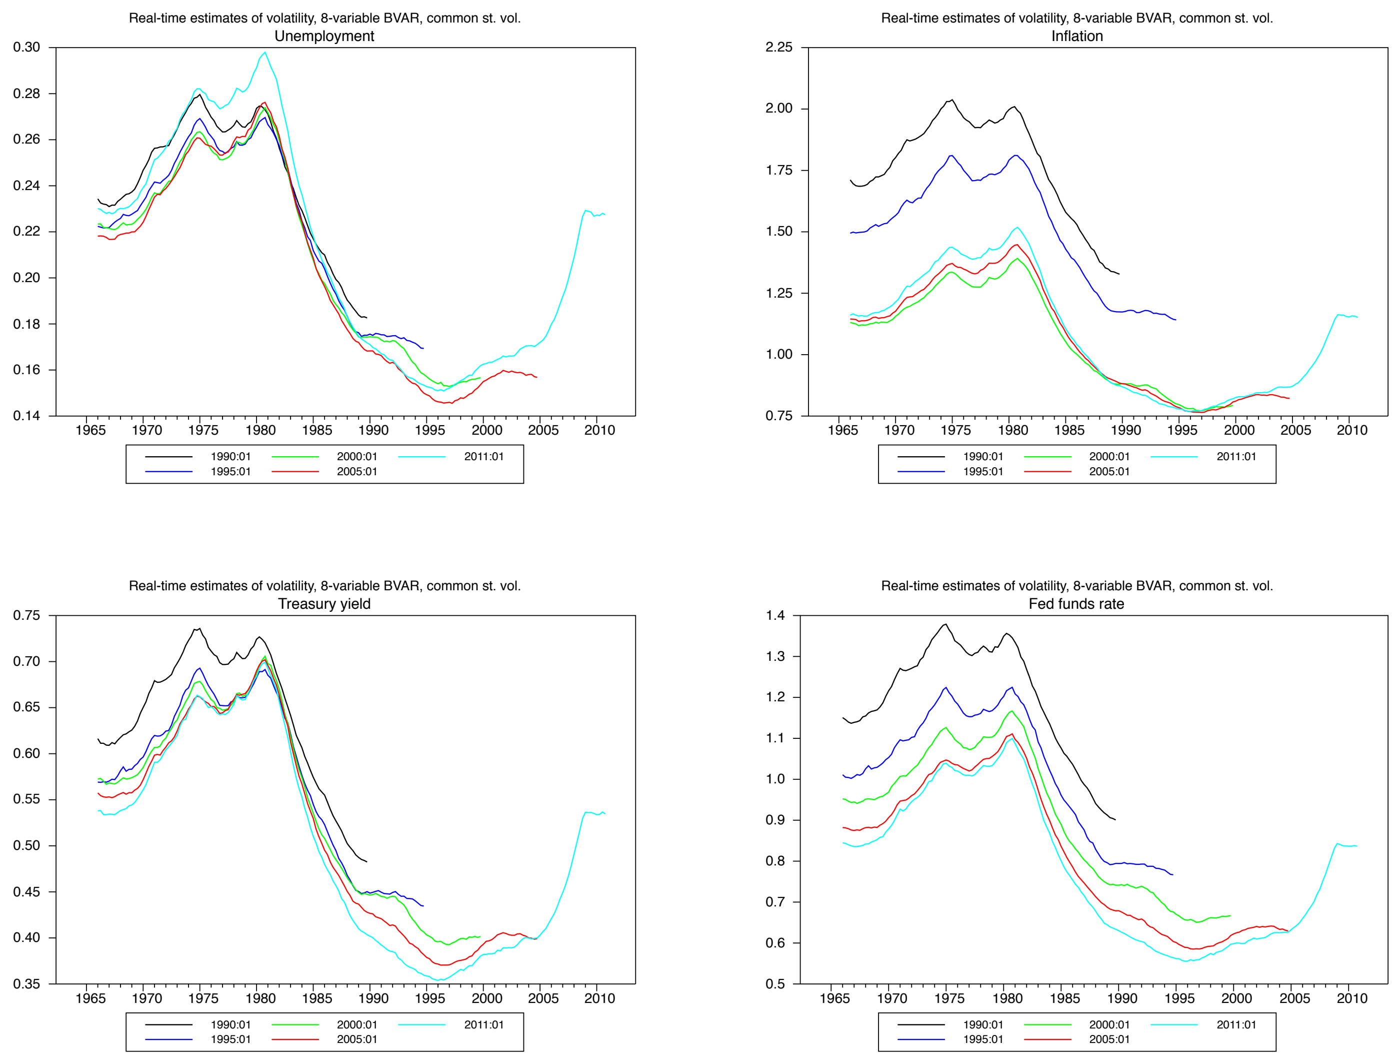

Figure 8: Volatility estimates from 8-variable BVAR with common stochastic volatility, real-time data, continued

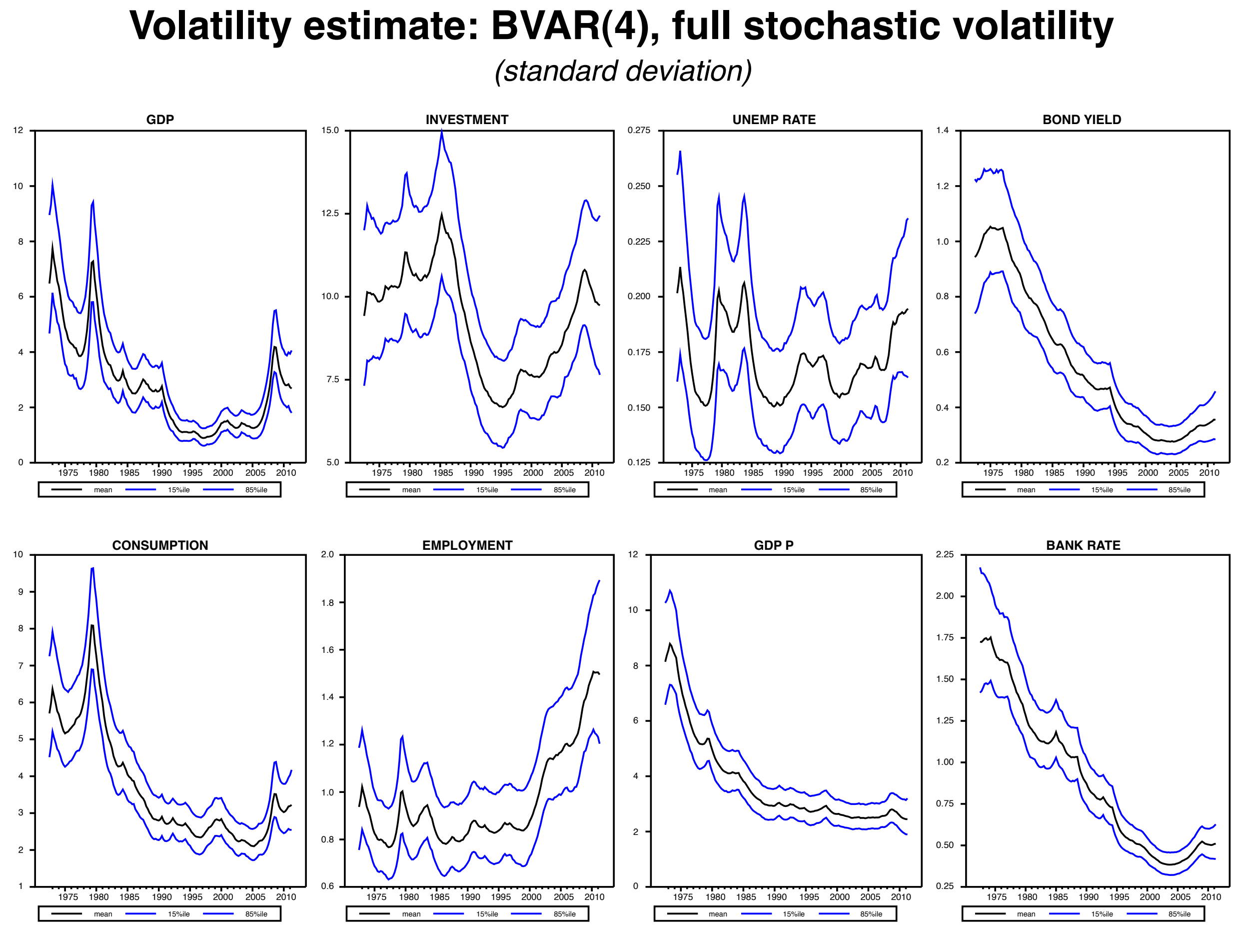

Figure 9: Volatility estimates from 8-variable BVAR with independent stochastic volatilities, UK data

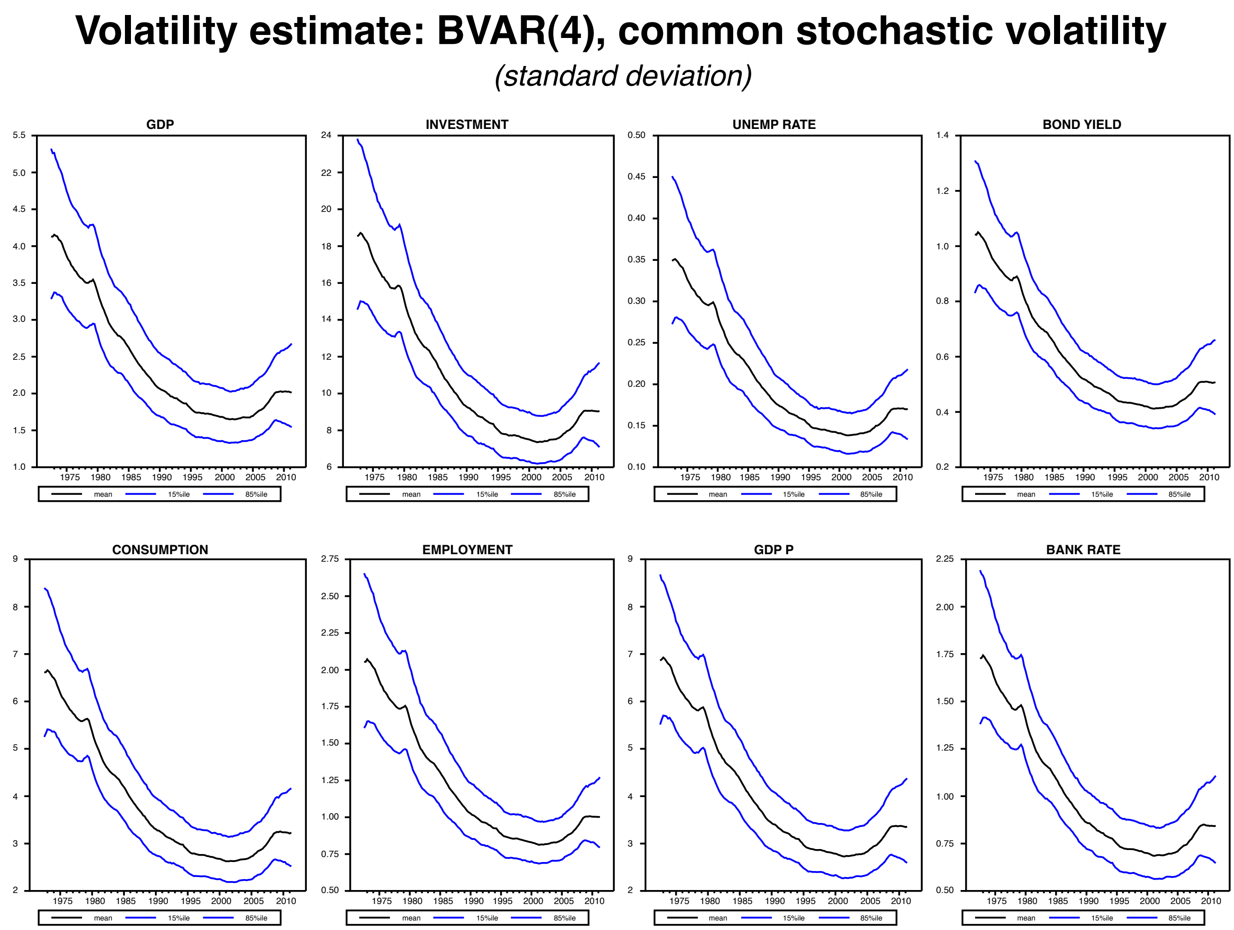

Figure 10: Volatility estimates from 8-variable BVAR with common stochastic volatility, UK data

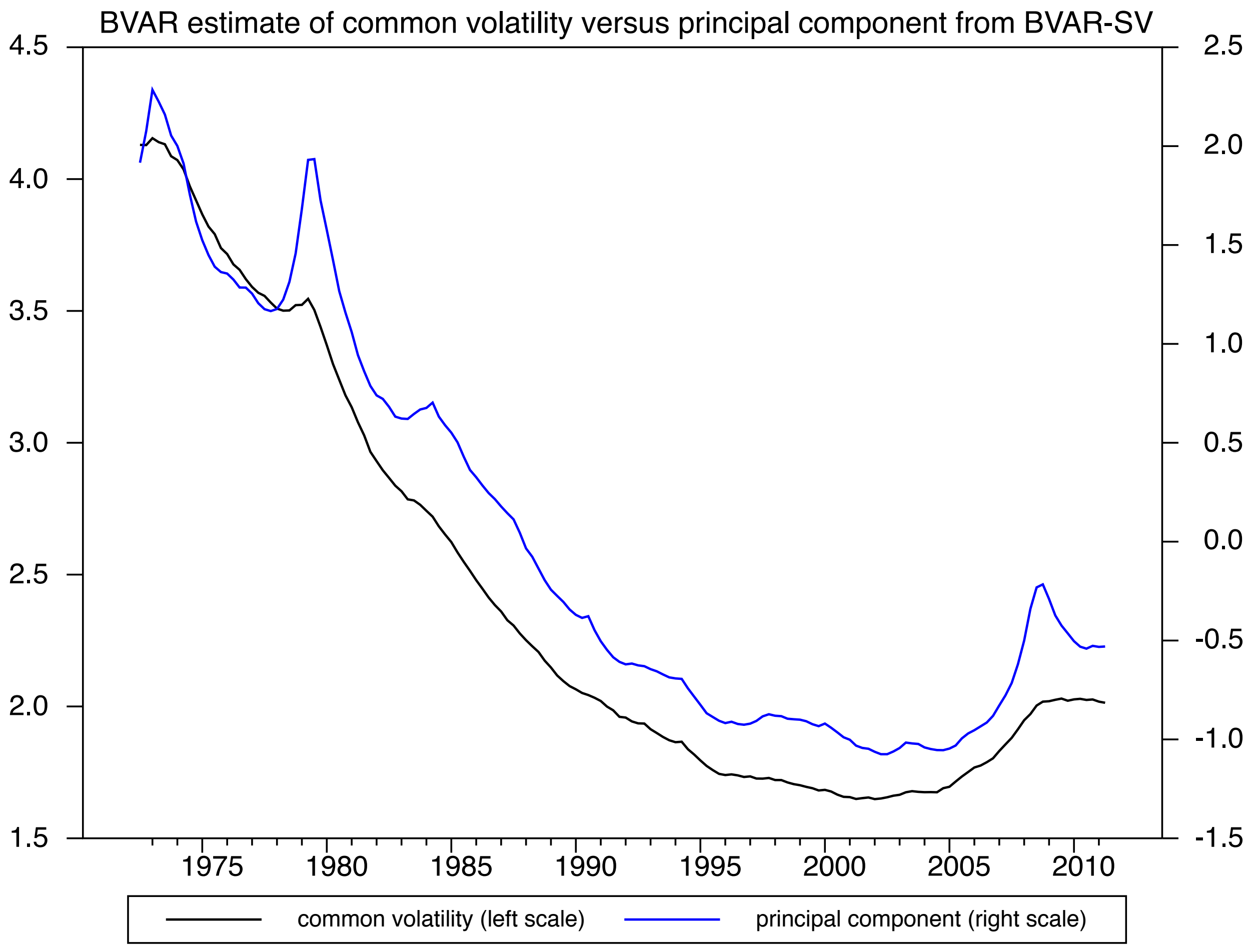

Figure 11: Principal component of BVAR-SV estimates of volatility versus common factor of volatility in 8-variable BVAR-CSV, UK data