TI 2021-056/III
Tinbergen Institute Discussion Paper

Vector Autoregressions with Dynamic Factor Coefficients and Conditionally Heteroskedastic Errors

Paolo Gorgi[1]
Siem Jan Koopman[1,2]
Julia Schaumburg[1]

[1] Vrije Universiteit Amsterdam and Tinbergen Institute, The Netherlands
[2] CREATES, University of Aarhus, Denmark

Vector Autoregressions with Dynamic Factor Coefficients and Conditionally Heteroskedastic Errors[1]

Paolo Gorgi$^{(a)}$ *Siem Jan Koopman*$^{(a,b)}$ *Julia Schaumburg*$^{(a)}$

$^{(a)}$ Vrije Universiteit Amsterdam and Tinbergen Institute, The Netherlands

$^{(b)}$ CREATES, University of Aarhus, Denmark

June 2021

Abstract

We introduce a new and general methodology for analyzing vector autoregressive models with time-varying coefficient matrices and conditionally heteroskedastic disturbances. Our proposed method is able to jointly treat a dynamic latent factor model for the autoregressive coefficient matrices and a multivariate dynamic volatility model for the variance matrix of the disturbance vector. Since the likelihood function is available in closed-form through a simple extension of the Kalman filter equations, all unknown parameters in this flexible model can be easily estimated by the method of maximum likelihood. The proposed approach is appealing since it is simple to implement and computationally fast. Furthermore, it presents an alternative to Bayesian methods which are regularly employed in the empirical literature. A simulation study shows the reliability and robustness of the method against potential misspecifications of the volatility in the disturbance vector. We further provide an empirical illustration in which we analyze possibly time-varying relationships between U.S. industrial production, inflation, and bond spread. We empirically identify a time-varying linkage between economic and financial variables which are effectively described by a common dynamic factor. The impulse response analysis points towards substantial differences in the effects of financial shocks on output and inflation during crisis and non-crisis periods.

Keywords: time-varying parameters, vector autoregressive model, dynamic factor model, Kalman filter, generalized autoregressive conditional heteroskedasticity, orthogonal impulse response functions.

JEL classification: C32, E31

[1]Emails: `p.gorgi@vu.nl` (P Gorgi), `s.j.koopman@vu.nl` (S J Koopman), `j.schaumburg@vu.nl` (J Schaumburg).
Schaumburg thanks the Netherlands Organization for Scientific Research (NWO, grant VENI451-15-022) for financial support.

1 Introduction

The development of dynamic stochastic general equilibrium models to study structural interactions between economic and financial variables has gained much interest in economic research; Gertler and Kiyotaki (2011) and Brunnermeier et al. (2013) provide overviews of this extensive literature. In empirical studies, the reduced form of these relations are typically formulated in terms of a vector autoregressive (VAR) model for a selection of macroeconomic and financial variables. The VAR model has been introduced in the seminal paper of Sims (1980). The econometrics of the VAR model has been reviewed, among others, by Stock and Watson (2001). The VAR analysis has become popular in empirical work due to its convenience in evaluating the potential impact of shocks by means of analyzing impulse response functions; see, for example, Inoue and Kilian (2013). Furthermore, in many empirical studies it has been concluded that VAR models form a good basis for obtaining accurate forecasts; see, for example, Chauvet and Potter (2013) and Clements and Galvão (2013).

The empirical macroeconometric literature has increasingly focused on capturing time-variation in VAR coefficient matrices, see, for instance, Primiceri (2005), Canova and Ciccarelli (2004), Hubrich and Tetlow (2015), Prieto et al. (2016), Galvão and Owyang (2018). Furthermore, as argued by Justiniano and Primiceri (2008), among others, a crucial feature in macroeconomic and financial variables is the presence of volatility changes. These studies have provided overwhelming evidence of parameter instability in the context of VAR models. These instabilities may originate from model misspecifications due to the existence of nonlinear relationships which are not appropriately addressed in reduced form VAR representations of structural interactions between variables. Also, time-varying features may be implied from structural changes in economic policy and persistent changes in the volatility of financial and commodity markets. The econometric challenge is to properly account for this time-variation in parameters of VAR models in order to enhance estimation accuracy but also to obtain reliable impulse response functions.

Since the seminal article of Primiceri (2005), it has become common practice to use Bayesian Markov chain Monte Carlo (MCMC) methods for the analysis of time-varying VAR models. In this and related studies, Gibbs sampling is adopted for repeatedly sampling features of the model: sampling one feature or parameter in each step of the chain, conditionally on the other features and parameters, until convergence. Since Primiceri (2005), there is a growing literature on Bayesian treatments of VAR models with stochastically time-varying parameters, for both autoregressive coefficients and covariances; see, for instance, Canova and Ciccarelli (2004, 2009). It is widely acknowledged that Bayesian methods require computationally intensive algorithms that are not necessarily trivial to implement. Eickmeier et al. (2015) and Abbate et al. (2016) consider a factor-augmented vector autoregressive model with time-varying parameters. Their estimation approaches are more classical but also somewhat restrictive: they rely on equation by equation methods and variance changes are driven by lagged factors only.

In this study we propose a new and simple approach to VAR modeling with time-varying parameters. We consider an unobserved factor process to specify the time-variation in the VAR coefficients and a generalised autoregressive conditionally hetereoskedastic (GARCH) specification for the modeling of volatility in the disturbance vector. An appealing feature of our proposed model is its simplicity in terms of interpretation as well as implementation. We show that the paramters in the model can be estimated straightforwardly by the method of maximum likelihood. We show how a closed form expression of the log-likelihood function is obtained by extending the Kalman filter equations. We also derive impulse response functions that allow us to account for the time-varying factors and the uncertainty in the estimated factors.

We illustrate the effectiveness and reliability of the method through a simulation study. In particular, the experiment highlights how the estimation of the unobserved factors is robust against potential mis-specification of the volatility of the error term. In an empirical application, we study financial-macro linkages in a three-dimensional VAR for two U.S. macroeconomic variables, industrial production and inflation, and the corporate bond spread as financial variable. The sample includes monthly data from January 1970 until January 2019. The time-variation of the VAR coefficients is driven by unobserved factors that have a structural economic interpretation. Since estimation is likelihood-based, standard model selection criteria such as BIC can be used to determine the optimal number of factors and the lag order. We find that the spillover strength of financial shocks to the two macro variables is driven by a common dynamic factor. Additional factors capture time-variation in the persistence of each variable and the linkage between the two macro variables. Impulse response analysis shows that the impacts of financial shocks differ substantially depending on whether the economy is in crisis or not. Furthermore, the financial crisis 2008/09 is found to be different from other crises.

The remainder of the paper is structured as follows. Section 2 introduces the factor VAR model with conditionally heteroskedastic errors and the estimation approach. Section 3 illustrates how to derive impulse response functions. Section 4 presents two simulation studies where the reliability of the proposed method is tested. Section 5 presents the empirical application that describes time-variation in macro-financial linkages. Section 6 concludes.

2 Vector Autoregressive Model with Time-Varying Parameters

In our study we consider the vector autoregressive (VAR) model for a time series of N-dimensional vectors $\{y_t\}_{t\in\mathbb{Z}}$. In particular, we specify the VAR model of order p, as denoted by VAR(p), through the equation

$$y_t = A_1 y_{t-1} + \ldots + A_p y_{t-p} + u_t, \qquad t \in \mathbb{Z},$$

where A_i, for $i = 1, \ldots, p$, is the $N \times N$ matrix of autoregressive coefficients, $\{u_t\}$ is an independent and identically distributed (iid) sequence of N-dimensional normal random disturbance vectors with mean vector 0 and covariance matrix H. For convenience of exposition, we rewrite the VAR(p) model equation as

$$y_t = \Phi Y_{t-1:p} + u_t, \tag{1}$$

where $Y_{t-1:p} = (y'_{t-1}, \ldots, y'_{t-p})'$ is an $(Np \times 1)$-dimensional vector in which the lagged values of y_t are stacked, and where $\Phi = (A_1, \ldots, A_p)$ is the $N \times Np$ coefficient matrix containing the corresponding coefficient matrices for the lags $1, \ldots, p$. We refer to Hamilton (1994) and Lütkepohl (2005) for a review on the stochastic properties of the VAR(p) model, the estimation of the autoregressive coefficient matrix Φ and covariance matrix H, and the forecasting and impulse response analysis for the VAR(p) model.

2.1 Time-varying factor-driven coefficient matrices

We introduce a time-varying autoregressive coefficient matrix in the VAR(p) model specification by having equation (1) replaced by

$$y_t = \Phi_t Y_{t-1:p} + u_t, \tag{2}$$

where the $N \times Np$ matrix Φ_t has the same dimension as Φ and varies with time-index t. We assume that the time-variation of Φ_t relies on an $(r \times 1)$-dimensional vector of unobserved factors $f_t = (f_{t,1}, \ldots, f_{t,r})'$ as

$$\Phi_t = \Phi^c + \Phi^f_1 f_{t,1} + \cdots + \Phi^f_r f_{t,r}, \tag{3}$$

where $N \times Np$ matrices Φ^c and Φ^f_i, for $i = 1, \ldots, r$, are the “factor loading” coefficient matrices and $f_{t,i}$ is a time-varying scalar (the ith element of vector f_t), for $i = 1, \ldots, r$. We specify the unobserved factor as another vector autoregressive process which we formulate as

$$f_{t+1} = \varphi f_t + \eta_t, \tag{4}$$

where $\varphi = \text{diag}\{\varphi_1, \dots, \varphi_r\}$ is a $r \times r$ diagonal matrix of autoregressive coefficients $\varphi_1, \dots, \varphi_r$, and where $\{\eta_t\}$ is an iid sequence of r-dimensional disturbance vectors of normal random variables with mean vector 0 and covariance matrix Σ_η. We assume that the factor f_t is a stationary process and therefore we impose $|\varphi_i| < 1$ for $i = 1, \dots, r$. For identification purposes, we restrict the covariance matrix of η_t to be $\Sigma_\eta = \text{I}_r - \varphi\,\varphi'$, where I_r is an $r \times r$ identity matrix. This restriction implies that the r factors in f_t are orthogonal and have a unit unconditional variance, we have $\text{Var}(f_t) = \text{I}_r$. Furthermore, we assume that the disturbance vector series u_t and η_t are uncorrelated, that is $\mathbb{E}(u_t\,\eta_s') = 0$ for all $t, s \in \mathbb{Z}$.

The proposed time-varying VAR process as specified in the equations (2), (3) and (4) can be referred to as the dynamic factor VAR (DFVAR) model. The model specification is very general. It nests the standard VAR model by setting $r = 0$ such that equation (3) reduces to $\Phi_t = \Phi^c$. On the other hand, by having $r = p \times N^2$ and setting the loading matrices Φ_i^f, for $i = 1, \dots, r$, as unique selection matrices for all elements in Φ_t, each entry of Φ_t can be specified as a distinct unobserved autoregressive process. Hence, the general specification of Primiceri (2005) can be considered in our modeling framework. Many other choices of the loading matrices allow different entries of Φ_t to depend on one dynamic factor or on a subset of dynamic factors. For instance, in the empirical application of Section 5, we specify the time-varying linkage between macro and financial variables through one common dynamic factor. The specification (3) also allows for the factor-based time-varying VAR model as suggested in Canova (2007, Section 10.4).

The loading matrices Φ_i^f, for $i = 1, \dots, r$, are not typically specified as full matrices of coefficients since restrictions are needed for identification. Such restristictions are reminiscent of dynamic factor analysis; see, for example, Stock and Watson (2011). We typically specify the loading matrices Φ_i^f as selection matrices or as sparse and parsimoniously designed matrices. In the remainder of this section, we discuss stability conditions of the DFVAR model and the estimation of parameters. In the next section, the model is extended to include conditional heteroskedasticity in the disturbance vector u_t.

2.2 Stationarity conditions for the DFVAR process

We establish the stationarity conditions for the DFVAR model. For this purpose, we define the stochastic $Np \times Np$ matrix

$$\tilde{\Phi}_t = \begin{bmatrix} \Phi_t \\ \mathrm{I}_{N(p-1)} \quad \mathbf{0}_{N(p-1),N} \end{bmatrix}, \tag{5}$$

where I_k is the $k \times k$ identity matrix, for any $k \in \mathbb{N}$, and $\mathbf{0}_{k,m}$ is the $k \times m$ matrix of zeros, for any $m \in \mathbb{N}$. Furthermore, we denote the matrix norm operator by $\|\cdot\|$. The theorem below delivers sufficient conditions for the existence of a strictly stationary solution for the DFVAR model.

Theorem 2.1. *Let $|\varphi_i| < 1$, for $i = 1, \ldots, r$, and let the following Lyapunov coefficient be strictly negative*

$$\gamma_m = \frac{1}{m} \mathbb{E} \log \|\tilde{\Phi}_{t-1}\tilde{\Phi}_{t-2} \cdots \tilde{\Phi}_{t-m}\| < 0,$$

for some $m \in \mathbb{N}$. Then, the DFVAR process as specified by equations (2) and (3), with time-index $t \in \mathbb{Z}$, admits a unique stationary solution.

Proof. We consider the following Markov representation of the DFVAR process

$$Y_{t:(p-1)} = \tilde{\Phi}_t Y_{t-1:p} + \tilde{u}_t,$$

where $\tilde{u}_t = (u_t', \mathbf{0}_{1,N(p-1)})'$. Then the result follows immediately by an application of Theorem 1.1 of Bougerol and Picard (1992). In particular, we note that $\{f_t\}_{t\in\mathbb{Z}}$ is a stationary process given the assumption that $|\varphi_i| < 1$, for $i = 1, \ldots, r$. Hence, we have that $\{\tilde{\Phi}_t\}_{t\in\mathbb{Z}}$ is a stationary sequence of matrices. Furthermore, it is immediate to see that $\mathbb{E}\|\tilde{\Phi}_t\| < \infty$ since each element of the matrix $\tilde{\Phi}_t$ can be expressed as a linear combination of the Gaussian process f_t, which has finite moments of any order. Therefore, given that $\gamma_m < 0$ for some $m \in \mathbb{N}$, we conclude that all the assumptions of Theorem 1.1 of Bougerol and Picard (1992) are satisfied. □

We can verify the stationarity condition of $\gamma_m < 0$ in practical settings through simulations. We start by choosing a large value for m. Then, we replace the expectation $\mathbb{E}$ in Theorem 2.1 with the sample average, over a sufficient number of Monte Carlo draws for y_t from the DFVAR model. The time series length is typically set to a relatively large value.

2.3 Estimation of the parameters in the DFVAR model

Assume that we observe a sample of size $T \in \mathbb{N}$, we have $\{y_t\}_{t=1}^{T}$, and our aim is to estimate the parameters in the DFVAR model. The parameter vector is denoted by ψ and collects all parameters in the matrices Φ^c, $\Phi_1^f, \ldots, \Phi_r^f$, H, φ and Σ_η. The parameter vector ψ can be estimated by Maximum Likelihood (ML) straightforwardly by state space methods. The first step is to express the model in state space form. The Kalman filter equations can then be employed to obtain the likelihood function via the prediction error decomposition.

We define $\tilde{y}_t = y_t - \Phi^c Y_{t-1:p}$ and rewrite the DFVAR model equation for $\tilde{y}_t$ as

$$\begin{aligned}
\tilde{y}_t &= \left[\Phi_1^f f_{t,1} + \ldots + \Phi_r^f f_{t,r}\right] Y_{t-1:p} + u_t \\
&= \Phi_*^f \left(f_t \otimes \mathrm{I}_{Np}\right) Y_{t-1:p} + u_t \\
&= \left(Y_{t-1:p}' \otimes \Phi_*^f\right) \mathrm{vec}\left(f_t \otimes \mathrm{I}_{Np}\right) + u_t \\
&= \left(Y_{t-1:p}' \otimes \Phi_*^f\right) J\, f_t + u_t, \qquad (6)
\end{aligned}$$

where $\otimes$ denotes the Kronecker matrix product and $\mathrm{vec}(\cdot)$ denotes the column vectorization operator, with $N \times Npr$ matrix $\Phi_*^f = \left[\Phi_1^f, \ldots, \Phi_r^f\right]$ collecting the loading matrices for the time-varying coefficient matrix Φ_t and $N^2p^2r \times r$ selection matrix J as given by

$$J = \begin{bmatrix} \mathrm{I}_r \otimes e_1 \\ \vdots \\ \mathrm{I}_r \otimes e_{Np} \end{bmatrix},$$

where e_j, for $j = 1, \ldots, Np$, is the jth column of the identity matrix I_{Np}. In the development towards equation (6), we use the well-known property that $\mathrm{vec}(ABC) = (C' \otimes A)\mathrm{vec}(B)$, for any appropriate set of matrices A, B, C; see Magnus and Neudecker (2019, Section 18.11). The matrix J effectively carries out the vectorization of the matrix $f_t \otimes \mathrm{I}_{Np}$. Finally, given the expression in (6), we can obtain the linear Gaussian state space representation of the DFVAR model. It consists of two equations: the measurement and transition equations which are given respectively by

$$\tilde{y}_t = Z_t f_t + u_t, \qquad f_{t+1} = \varphi f_t + \eta_t, \qquad (7)$$

with $N\times r$ matrix $Z_t = \left(Y_{t-1:p}' \otimes \Phi_f\right) J$, vector f_t playing the role of the $r\times 1$ state vector, $r\times r$ time-invariant transition matrix φ, and with the properties for the disturbance vectors u_t and η_t as discussed above. We notice that the system matrix Z_t has relatively low dimensions and can be constructed in a computationally fast way given that matrix J is a selection matrix consisting of zeroes and ones. Furthermore, matrix Z_t is constant conditional on the σ-field generated by past observations $\mathcal{F}_{t-1} = \sigma(y_s, 0 < s \leq t-1)$. Therefore, we can apply the Kalman filter to derive the *conditional* mean and variance of $\tilde{y}_t$.

The one-step ahead prediction error for the DFVAR model is defined and given by

$$\begin{aligned} v_t &= y_t - \mathbb{E}(y_t|\mathcal{F}_{t-1};\psi) \\ &= y_t - \Phi^c Y_{t-1:p} - \mathbb{E}(\tilde{y}_t|\mathcal{F}_{t-1};\psi) \\ &= \tilde{y}_t - Z_t\mathbb{E}(f_t|\mathcal{F}_{t-1};\psi) \ = \ \tilde{y}_t - Z_t\, a_t \end{aligned} \tag{8}$$

where $a_t = \mathbb{E}(f_t|\mathcal{F}_{t-1};\psi)$ for $t = p+1,\ldots,T$. The actual computation of the prediction error v_t can only start at $t = p+1$ since only then the data vector $Y_{t-1:p}$ is complete. It follows that the variance of the prediction error is defined and given by

$$\begin{aligned} F_t &= \text{Var}(\tilde{y}_t - Z_t\, a_t|\mathcal{F}_{t-1};\psi) \\ &= Z_t\text{Var}(f_t - a_t|\mathcal{F}_{t-1};\psi) + \text{Var}(u_t|\mathcal{F}_{t-1};\psi) \\ &= Z_t P_t Z_t' + \text{Var}(u_t;\psi) \ = \ Z_t P_t Z_t' + H, \end{aligned} \tag{9}$$

where $P_t = \text{Var}(f_t|\mathcal{F}_{t-1};\psi)$ for $t = p+1,\ldots,T$. When the model is correctly specified, the sequence $\{v_{p+1},\ldots,v_T\}$ is serially uncorrelated. For a given vector ψ and $t = p+1,\ldots,T$, the Kalman filter update equations for a_{t+1} and P_{t+1} are given by

$$a_{t+1} = \varphi a_t + K_t v_t, \qquad P_{t+1} = \varphi P_t(\varphi - K_t Z_t)' + \Sigma_\eta, \tag{10}$$

where the Kalman gain matrix is defined as $K_t = \varphi P_t Z_t' F_t^{-1}$; see Durbin and Koopman (2012, Section 4.2) for derivations of the Kalman filter. Given the unconditional properties of f_t, we can initialize the Kalman filter with $a_{p+1} = 0$ and $P_{p+1} = \text{I}_r$.

The Kalman filter effectively carries out the prediction error decomposition based on the linear Gaussian state space representation of the DFVAR model. It follows that the conditional log-likelihood function $\ell(\psi)$

can then be computed by

$$\ell(\psi) = \sum_{t=p+1}^{T} \ell_t(\psi), \qquad \ell_t(\psi) = -\frac{N}{2}\log 2\pi - \frac{1}{2}\log|F_t| - \frac{1}{2}v_t' F_t^{-1} v_t, \tag{11}$$

where $|\cdot|$ denotes the determinant of a matrix. The ML estimator is defined as the maximizer of the log-likelihood function which can be obtained using standard numerical optimization routines. The gravity of this estimation challenge depends on the number of parameters in ψ and on how informative the observations are about ψ. The econometrician can consider different strategies of exploring ways to establish a satisfactory model specification. In our empirical study in Section 5, we show for a concrete and relevant analysis that parameter estimation does not cause much concern. We conclude by emphasizing that the computation of the log-likelihood function $\ell(\psi)$ in (11), for the DFVAR model with a given ψ, is straightforward, simple and fast.

2.4 Conditional heteroskedastic disturbances: the CH-DFVAR model

In empirical analyses of macroeconomic and financial time series data, it is often concluded that it is crucial to account for time-varying scales (or variances). To address these concerns in empirical studies, we allow for conditional heteroskedasticity for the disturbance vector u_t in equation (2) of the DFVAR model. For this purpose we extend the model to include time-variation in the conditional covariance matrix of the disturbance vector u_t. In effect, we replace the $N \times N$ covariance matrix H of u_t by a conditionally time-varying covariance matrix H_t. In case we have a univariate model with a single disturbance variance, we could have considered the generalized autoregressive conditional heteroskedastic (GARCH) model for the time-varying variance; see Bollerslev (1986). There are various multivariate extensions of the GARCH model available to us, in order to treat a time-varying covariance matrix. Here, we consider the multivariate GARCH specification of (Engle and Kroner, 1995) as it is a convenient extension for our purposes and in the context of the DFVAR model. In particular, this specification allows us to update the covariance matrix H_t by the prediction error and, hence, within the Kalman filter. Koopman et al. (2010) have considered a similar adjustment for the Kalman filter in the context of GARCH and a different modeling framework.

The DFVAR model with conditional heteroskedastic disturbances (CH-DFVAR) remains specified as $y_t = \Phi_t Y_{t-1:p} + u_t$ but with disturbance vector u_t that is assumed to have a time-varying covariance matrix H_t that is specified by the dynamic matrix equation

$$H_t = \Omega + BH_{t-1}B' + Av_{t-1}v_{t-1}'A', \qquad t = 1, \dots, T, \tag{12}$$

where $N \times N$ matrix Ω is symmetric and positive-definite, $N \times N$ matrices A and B are coefficient matrices, and $v_t = y_t - \mathrm{E}(y_t|\mathcal{F}_{t-1}; \psi)$ is defined as in equation (8). This specification is proposed by (Engle and Kroner, 1995) and is usually referred to as the BEKK model. The updating equation for H_t is simple and it ensures that H_t remains a symmetric and positive-definite matrix, for all $t = 1, \dots, T$. A more parsimonious specification of the conditional variance is the scalar BEKK model as given by

$$H_t = \Omega + \beta^2 H_{t-1} + \alpha^2 v_{t-1}v_{t-1}', \qquad t = 1, \dots, T, \tag{13}$$

where β and α are scalar coefficients. We will consider this scalar BEKK model specification in the simulation study as well as in the empirical study.

2.5 Estimation of the parameters in CH-DFVAR model

The BEKK model for H_t enables us to retain a closed form of the log-likelihood function. Hence, our approach to parameter estimation remains unaltered. We only require an extension of the Kalman filter to incorporate the BEKK updating for H_t. The main recursive equations of the Kalman filter, for $t = p + 1, \dots, T$, only need minor adjustments. The details are as follows.

1. For time-index t, and for given values of a_t, P_t and H_t, we compute

$$v_t = \tilde{y}_t - Z_t a_t, \qquad F_t = Z_t P_t Z_t + H_t,$$

   as in equations (8) and (9), respectively, but with H replaced H_t.

2. We apply the Kalman updating equations (10) to obtain a_{t+1} and P_{t+1}.

3. We update the conditional covariance matrix using the scalar BEKK filtering equation

$$H_{t+1} = \Omega + \beta^2 H_t + \alpha^2 v_t v_t'.$$

The covariance matrix updating is initialized with $H_{p+1} = \Omega \,/\, (1 - \beta^2 - \alpha^2)$. The log-likelihood function is computed as in (11). The parameters for H are replaced by those for Ω, α and β in the parameter vector ψ.

3 Impulse Response Functions

For a given DFVAR model, with the unknown parameters replaced by their ML estimates, we discuss how to obtain the orthogonalized impulse response function (IRF). The orthogonalized IRF is discussed as early as in Sims (1980) and further explored in Cooley and LeRoy (1985) and Pagan (1987); the default method for orthogonalization is the Cholesky decompostion which has been used in influential macroeconomic studies such as Christiano et al. (1996).

The IRF allows us to assess how a unit shock in the error term at time τ is expected to propagate to the observed variables over time. Under stationarity conditions, the DFVAR model admits a vector moving average (VMA) process of an infinity length, that is a VMA(∞). In case of the "future" observation $y_{\tau+h}$, $h \in \mathbb{N}$, we have the representation

$$y_{\tau+h} = \sum_{i=0}^{\infty} \Psi_{\tau-i}(h) u_{\tau+h-i},$$

where $\Psi_{\tau}(h) = \mathrm{I}_N$ and $\Psi_{\tau-i}(h)$ is defined as the $N \times N$ matrix containing the submatrix given by the first N rows and N columns of the $Np \times Np$ matrix $\prod_{j=1}^{i} \tilde{\Phi}_{\tau+h-j+1}$ with matrix $\tilde{\Phi}_{\tau}$ being defined in (5). As it is implied by equation (3) that Φ_t is a function of f_t, matrix $\Psi_{\tau-i}(h)$ is a function of $f_{\tau+h-i+1}, \ldots, f_{\tau+h}$ for $i = 1, 2, \ldots$. The orthogonal IRF is based on the Cholesky decomposition of the covariance matrix H_τ as given by $H_\tau = Q_\tau Q_\tau'$ where Q_τ is an $N \times N$ lower-triangular matrix, with strictly-positive diagonal values. We obtain the VMA(∞) representation with orthogonal errors by

$$y_{\tau+h} = \sum_{i=0}^{\infty} \Psi_{\tau-i}(h) Q_{\tau+h-i} Q_{\tau+h-i}^{-1} u_{\tau+h-i} = \sum_{i=0}^{\infty} \Psi_{\tau-i}^{*}(h) \varepsilon_{\tau+h-i},$$

where $\Psi_{\tau-i}^{*}(h) = \Psi_{\tau-i}(h) Q_{\tau+h-i}$ and $\varepsilon_{\tau+h-i} = Q_{\tau+h-i}^{-1} u_{\tau+h-i}$ such that the covariance matrix for $\varepsilon_{\tau+h-i}$ equals the identity matrix I_N. It follows that the impact on $y_{\tau+h}$ of a unit shock of the k-th element of ε_τ is given by

$$\frac{\partial y_{\tau+h}}{\partial \varepsilon_{k,\tau}} = \Psi_{\tau-h}^{*}(h) e_k, \tag{14}$$

where e_k denotes the k-th column of the identity matrix I_N. Equation (14) provides the definition of the orthogonal IRF, for a fixed τ and as a function of $h \in \mathbb{N}$ and $k = 1, \dots, N$.

In practice, the IRF in equation (14) cannot be used directly since in our DFVAR model the time-varying coefficient matrices $\{\Psi^*_{\tau-h}(h)\}_{h\in\mathbb{N}}$ are subject to the random shocks $\eta_1, \dots, \eta_{\tau+h}$ which is apparent from equations (3) and (4). Hence we only obtain the IRF by taking the expectation of (14) conditional on the observed data before the shock occurs $\mathcal{F}_{\tau-1}$, that is

$$\text{IRF}_\tau(h, k) = \mathbb{E}\big(\Psi^*_{\tau-h}(h)|\mathcal{F}_{\tau-1}\big)e_k = \mathbb{E}\big(\Psi_{\tau-h}(h)|\mathcal{F}_{\tau-1}\big)Q_\tau e_k.$$

In practice, it is not feasible to obtain a closed form expression for the conditional expectation in the above equation. But it can be computed via Monte Carlo simulation. We first notice that for a given sequence $f^{\tau+h} = \{f_1, \dots, f_{\tau+h}\}$ where the f_t's are generated by the DFVAR model, with the unknown parameter vector being replaced by the corresponding ML estimates, all matrices $\Psi_{\tau-h}(h)$ can be computed. We have

$$\mathbb{E}\big(\Psi_{\tau-h}(h)|\mathcal{F}_{\tau-1}\big) = \int \Psi_{\tau-h}(h)p(f^{\tau+h}|\mathcal{F}_{\tau-1})\mathrm{d}f^{\tau+h},$$

where $p(f^{\tau+h}|\mathcal{F}_{\tau-1})$ is the conditional density of $f^{\tau+h}$ given the observed data before the shock at time τ. Given that the DFVAR model admits a linear state space representation as in equation (7), and given the normality of the random disturbance vectors u_t and η_t, the conditional density $p(f^{\tau+h}|\mathcal{F}_{\tau-1})$ is Gaussian. The Monte Carlo estimator for the IRF is then simply obtained by

$$\widehat{\mathbb{E}}\big(\Psi_{\tau-h}(h)|\mathcal{F}_{\tau-1}\big) = M^{-1}\sum_{i=1}^{M}\Psi^i_{\tau-h}(h), \qquad f_i^{\tau+h} \sim p(f^{\tau+h}|\mathcal{F}_{\tau-1}),$$

where $\{f_1^{\tau+h}, \dots, f_M^{\tau+h}\}$, for some $M \in \mathbb{N}$, is a series of independent draws from $p(f^{\tau+h}|\mathcal{F}_{\tau-1})$, and $\Psi^i_{\tau-h}(h)$ is the matrix $\Psi_{\tau-h}(h)$ that is computed by having $f^{\tau+h} = f_i^{\tau+h}$. Under standard regularity conditions, the Monte Carlo estimator converges to the conditional expectation for an increasing M.

As long as the vector autoregressive process is stable, for which the conditions are provided in Theorem 2.1, the impulse response function $\text{IRF}_\tau(h, k)$ converges to zero as h goes to infinity. Hence, the impact of a shock vanishes over time. This approach of deriving impulse response functions allows us to take into account the uncertainty of the estimated factor and does not require the estimated coefficient matrix Φ_t to have spectral radius smaller than one. In fact, at some given points in time, the time-varying autoregressive coefficient matrix

Φ_t can have spectral radius equal to or greater than one, which can occur in periods of high persistence of shocks. However, this will occur only locally as long as the overall DFVAR process is stationary. Finally, we can use the quantiles of the distribution of $\Psi_{\tau-h}(h)$ to derive confidence intervals for the impulse response function. In practice, this is easily incorporated in the presented Monte Carlo method for estimating the IRF.

4 Monte Carlo Evidence

We have carried out two simulation studies. The detailed descriptions of both studies and their results are presented. In the first study, we evaluate the small sample properties of the ML estimator of the parameter vector ψ to verify the overall reliability of our proposed estimation method. In the second study, we verify how the specification of the conditional variance H_t can effectively capture time-variation in the variance of the error term and how robust the estimation of the dynamic factor f_t is to possible misspecifications of H_t.

4.1 Small sample properties of the maximum likelihood estimator

We investigate the properties of the ML estimator through a simulation experiment. We consider a two-dimensional CH-DFVAR model with one lag dependence, one factor f_t for the time-varying autoregressive coefficient matrix, and with the time-varying H_t specification given by the scalar BEKK in equation (13); we have $N = 2$ and $p = r = 1$. This basic CH-DFVAR model specification can be given by

$$y_t = \Phi_t y_{t-1} + u_t, \qquad \Phi_t = \Phi^c + \Phi^f f_t, \qquad f_t = \phi f_{t-1} + \eta_t, \qquad H_t = \Omega + \beta^2 H_{t-1} + \alpha^2 v_{t-1} v_{t-1}',$$

where Var$(u_t) = H_t$ and Var$(\eta_t) = 1 - \phi^2$, for $t = 2, \ldots, T$. We set the true parameter values as follows

$$\Phi^c = 0.3\mathrm{I}_2 + 0.1\mathrm{J}_2, \quad \Phi^f = 0.2\mathrm{I}_2, \quad \phi = 0.95, \quad \Omega = 0.3\mathrm{I}_2 + 0.2\mathrm{J}_2, \quad \beta^2 = 0.75, \quad \alpha^2 = 0.1,$$

where J_2 is a 2×2 matrix with diagonal elements equal to zero and its off-diagonal element equal to unity. This specification entails a common factor that determines time-variation in the autoregressive coefficient matrix. The total number of parameters that we estimate is 12: 4 in Φ^c, 2 in Φ^f, 3 in Ω and 3 scalar parameters. The ML estimation of the parameters is considered for three sample sizes: $T = 500,\ 1000,\ 2500$. This Monte Carlo experiment consists of 1000 replications of simulating data (based on true parameters) and estimation of the parameters by treating the simulated data as the observation data.

We obtain for each parameter a set of 1000 ML estimates based on the 1000 simulated data sets. The distributions of these ML estimates are reported in Figure 1, for each parameter, and for the three sample sizes. As the sample size increases, we observe that the distributions of these estimates are collapsing towards their corresponding true parameter values. It provides some suggestion that our proposed estimation method provides consistent estimators. Furthermore, the distributions appear to be symmetric and normally shaped for most parameters, even at the smallest sample size. An exception is the estimate of the autoregressive coefficient ϕ for the factor f_t: the distribution exhibit some left skewness. This finding is not highly surprising since small sample bias and left skewness are typical features when estimating autoregressive parameters that are close to one, leading to a persistent process for f_t. Overall, we can conclude that our ML estimation method delivers reliable parameter estimates.

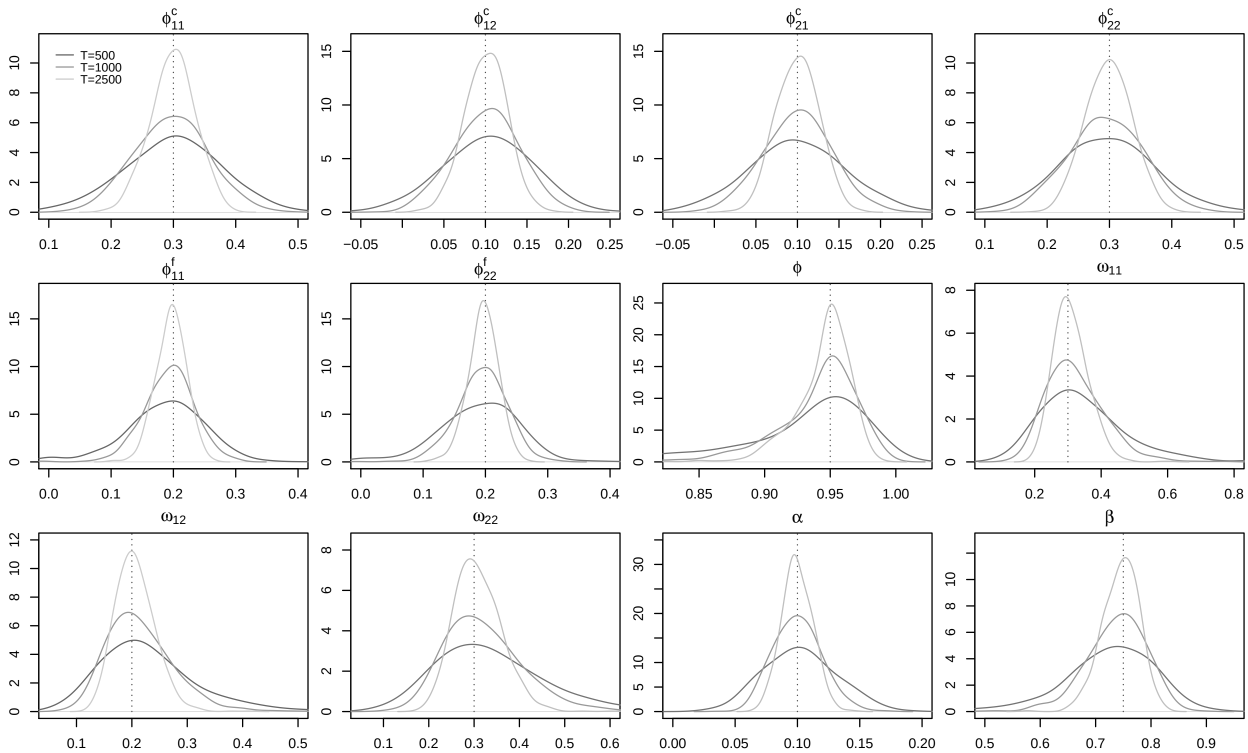

Figure 1: Kernel densities for the ML estimates of the twelve parameters, obtained from 1000 Monte Carlo replications, and for three different sample sizes $T = 500$, 1000, 2500. The dashed vertical line in each plot indicates the true parameter value.

4.2 Accuracy of filtering the conditional heteroskedasticity

We next focus on the accuracy of filtering the conditional heteroskedasticity in the CH-DFVAR model. We evaluate the performance of the scalar BEKK specification for H_t, as in equation (13), in treating different forms of time-variation in the variances of the disturbance vector u_t. We further assess the robustness of the signal extraction of f_t under the misspecification of the conditional covariance matrix H_t.

To focus on these aims of this second Monte Carlo experiment, we slightly simplify the earlier bivariate CH-DFVAR model and consider the true model as given by

$$y_t = \Phi_t y_{t-1} + u_t, \qquad \Phi_t = \Phi^f f_t, \qquad f_t = \phi f_{t-1} + \eta_t, \qquad H_t = h_t H,$$

where $\text{Var}(u_t) = H_t$ and $\text{Var}(\eta_t) = 1 - \phi^2$, for $t = 2, \ldots, T$, and where $\{h_t\}$ is a variance scaling sequence that determines the time-variation in H_t. We set the true parameter values as follows

$$\Phi^f = 0.25\text{I}_2, \qquad \phi = 0.95, \qquad H = \text{I}_2 + 0.5\text{J}_2.$$

We consider the following three specifications for the sequence $\{h_t\}$:

(a) Sine function: $h_t = 0.5\sin(\pi t/250) + 1$,

(b) Step function: $h_t = I(\sin(\pi t/250) > 0) + 0.5$,

(c) Constant: $h_t = 1$,

where $I(\cdot)$ denotes the indicator function. We adopt this model, with the three variants for h_t, to simulate the data. For a simulated data set, we estimate the parameters for the DFVAR model as given above, but with the specification for H_t replaced by the scalar BEKK equation (13). Hence, this model is misspecified for estimation when the model for simulation takes h_t as the sine or the step function. The model for parameter estimation is correctly specified when $h_t = 1$ since the CH-DFVAR model is nested with the DFVAR model, with a constant covariance matrix H. In this case, the nesting conditions are $\alpha^2 = \beta^2 = 0$ in (13).

Figure 2 presents the summary statistics of the extracted scalar sequences for the conditional variances $\{H_t\}$ and for the unobserved dynamic factors $\{f_t\}$, for the three different specifications of h_t as described above. These extracted sequences are obtained from the 1000 Monte Carlo replications and for a sample size of $T = 1000$. In this experiment, the dynamic factor sequence $\{f_t\}$ is simulated only once and its full path is

(a) Sine

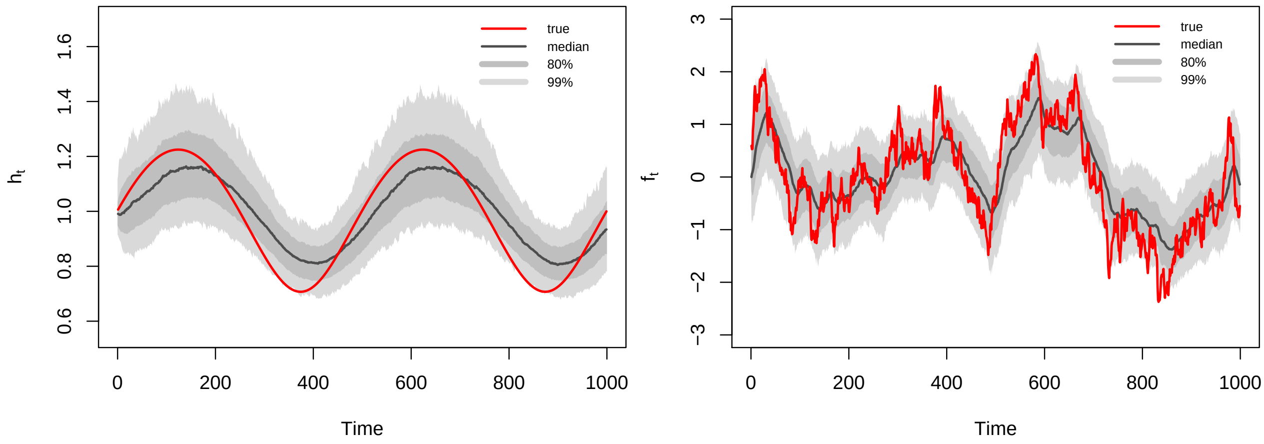

(b) Steps

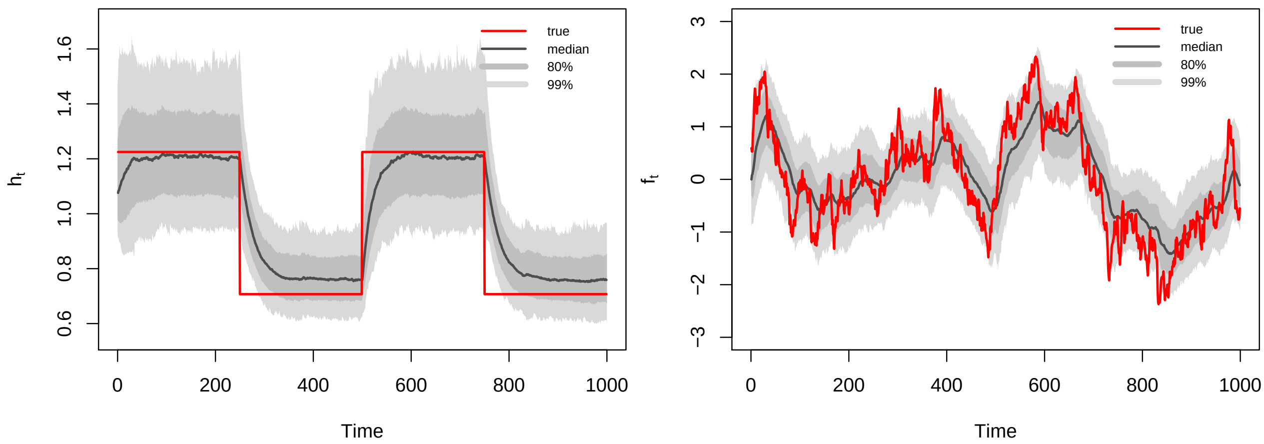

(c) Constant

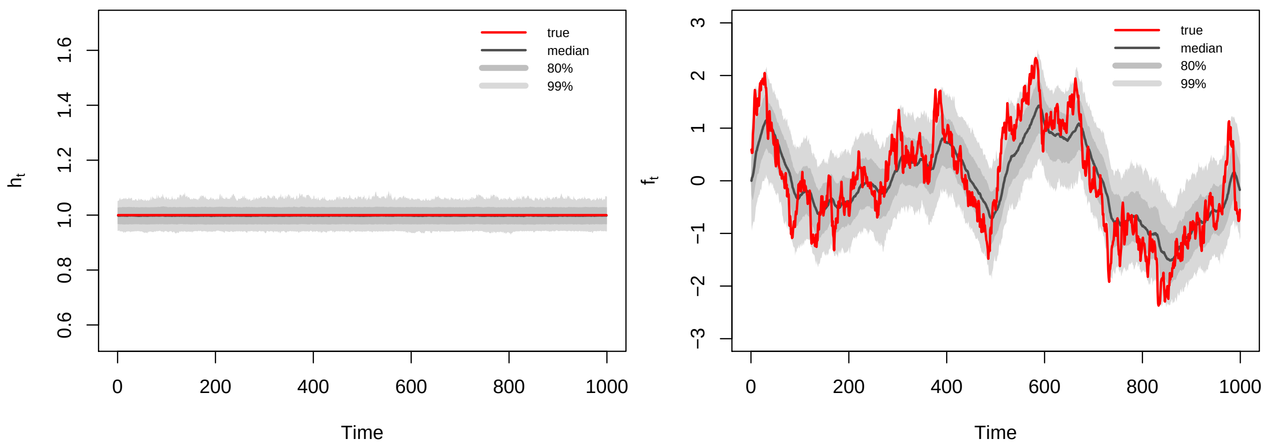

Figure 2: Left plots: the red line displays the true scaling variance h_t for the three different specifications (sine, step and constant functions). The solid black line is the median and the grey areas represent the distribution of the estimates of H_t obtained from the scalar BEKK model that is embedded in the Kalman filter; see Section 2.4. Right plots: the red line displays the true unobserved factor f_t. The factor is the same in all configurations (a), (b) and (c). The solid black line is the median and the grey areas represent the estimated factors for f_t obtained from the Kalman filter, that is $\mathrm{E}(f_t|\mathcal{F}_{t-1}, \hat{\psi})$.

kept fixed across all 1000 Monte Carlo replications. In this way, we can graphically depict both the variance scaling $\{h_t\}$ and the true dynamic factor $\{f_t\}$. In these plots, we also present some summary statistics (median, 80% and 99% percentiles) from the distribution of the estimates of h_t and f_t. We learn from Figure 2 that the BEKK specification of the conditional variance is effective in capturing the time-variation in h_t. The plots (a) and (b) in Figure 2 show this rather convincingly. Furthermore, we do not find any relevant differences in the estimates of the factor f_t for the three configurations of time-variation in h_t (sine, step and constant functions). The case with a constant variance $h_t = 1$ represents the situation where the model is correctly specified. From this Monte Carlo experiment we can conclude that the estimation of the unobserved factor f_t is robust to possible misspecification of H_t. Further confirmation of this finding is presented in Table 1 where we report the precision of the estimates in this Monte Carlo experiment. More specifically, we report the root mean squared errors (RMSE) for the estimate of the factor f_t and the variance scaling h_t, for the different specifications of h_t, for different sample sizes $T = 500, 1000, 2500$. The RMSE is computed as the average over $t = 2, \ldots, T$, and over all 1000 Monte Carlo replications. We observe that the accuracy of the estimates remains similar across the different specifications of h_t. Also, the accuracy of the estimates increases with the sample size. Overall we can conclude that our methods are feasible and are robust to model misspecification.

Table 1: The average root mean squared error (RMSE) for the estimates of the factor f_t and the scaling variance h_t, for the different sample sizes $T = 500, 1000, 2500$. The estimates of f_t are obtained from the Kalman filter, that is $\mathrm{E}(f_t|\mathcal{F}_{t-1})$, while the estimates of h_t are obtained from the scalar BEKK specification that is embedded in the Kalman filter; see Section 2.4.

	$T = 500$		$T = 1000$		$T = 2500$	
	f_t	h_t	f_t	h_t	f_t	h_t
Sine	0.820	0.225	0.800	0.214	0.787	0.213
Steps	0.828	0.299	0.803	0.306	0.790	0.303
Constant	0.820	0.059	0.797	0.039	0.784	0.024

5 Empirical Study: Dynamic Macro-Financial Linkages and Spillovers

Time-varying parameter vector autoregression (VAR) models are frequently used to analyze whether changes in financial market conditions affect the transmission mechanisms of shocks to the real economy; for instance, Hubrich and Tetlow (2015), Prieto et al. (2016) and Galvão and Owyang (2018) provide empirical evidence of the existence of such macro-financial linkages. In these studies, parameter estimation and signal extraction is based on Bayesian Markov chain Monte Carlo (MCMC) methods such as Gibbs sampling; see the seminal work of Primiceri (2005). To show and illustrate that a frequentist approach is also feasible in analyzing macro-financial linkages and spillovers, we provide an empirical illustration for a three-dimensional vector autoregressive model of the U.S. economy. We consider the two macroeconomic variables of industrial production (IP) growth and headline CPI inflation, together with the single financial variable of spread between "BAA" rated corporate bond rates and the ten-year treasury bill rate. This selection of variables is inspired by the study of Galvão and Owyang (2018) where a two-dimensional smooth transition VAR model is considered for IP growth rates and inflation, using a financial index as the transition variable. The index is extracted from a large set of financial variables, but it is shown that the corporate bond spread is a good proxy for this index. Hence, we consider corporate bond spread to represent the financial shocks in our three-dimensional VAR model. From Table 2 of Galvão and Owyang (2018) we learn that the posterior inclusion probabilities of the bond spread is highest, with 98%. The three variables are obtained from the FRED data base (https://fred.stlouisfed.org/) and they all have a monthly frequency. The three time series span from January 1970 until January 2019; each time series consists of 589 observations. Figure 3 presents the time series graphs of the three variables.

5.1 Selection of model specification

In our empirical study for the three time series variables, we consider the modeling framework provided by the equations (2), (3) and (4) for the DFVAR model and equation (13) for the multivariate GARCH model. This general model requires various choices, in particular, for the number of lags p in equation (2), the number of factors r and the composition of the typically sparse matrices $\Phi_1^f, \ldots, \Phi_r^f$ in equation (3). After some initial experimentation where parameters are estimated and results are empirically validated, we have decided to focus on the vector autoregressive model with $p = 2$ in equation (2). Other empirical aspects of the model are investigated in detail below. Given that $N = 3$ and $p = 2$, the number of parameters in Φ^c of equation (3) is $N^2 p = 18$, in φ of equation (4) is r, and in Ω, α, β of equation (13) is $2 + N(N+1)/2 = 8$. Hence, the total number of parameters is $26 + r$ plus the number of parameters in matrices $\Phi_1^f, \ldots, \Phi_r^f$ of equation (3). In the modeling process, we are considering the different choices for the number of factors r and the sparse designs of the matrices $\Phi_1^f, \ldots, \Phi_r^f$. Since we have adopted a frequentist approach in parameter estimation that is based on the straightforward numerical maximisation of the log-likelihood function (11), see Section 2.5 for more details, we base our various model decisions on standard information criteria. Decisions on the

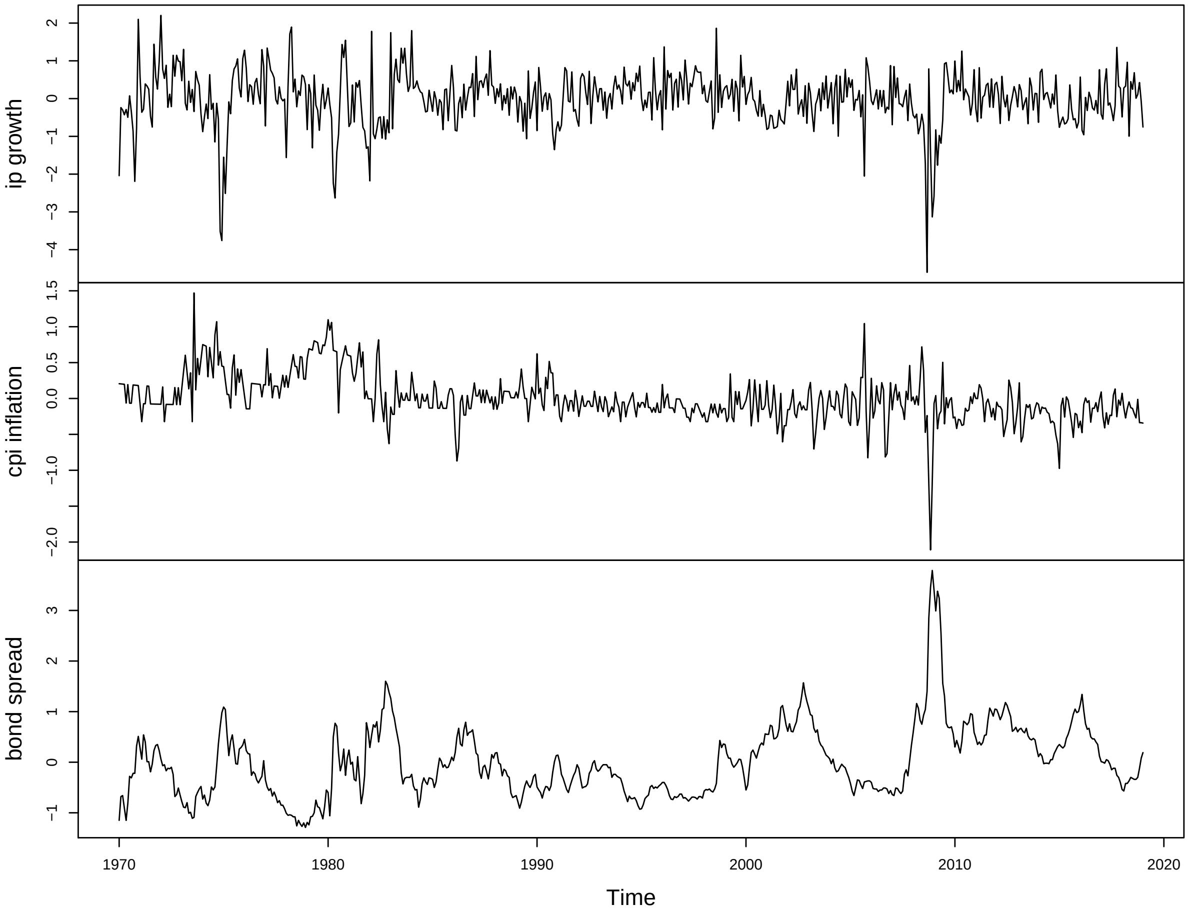

Figure 3: The three time series variables are from the U.S. economy: IP growth (the growth rates of industrial production), inflation (headline CPI inflation) and bond spread (the spread between BAA-rated corporate bond rate and 10-year Treasury rate). The time series are monthly and the sample ranges from January 1970 until January 2019 (589 observations). The data are obtained from the FRED data base (https://fred.stlouisfed.org/).

lag length p, the number of factors r, the compositions of loading matrices $\Phi_1^f, \dots, \Phi_r^f$ are all done using the Bayesian Information Criterion (BIC). We could also have adopted the Akaike Information Criterion (AIC) but empirically it is found that the BIC is somewhat more conservative as it opts for more parsimonious models.

Different specifications of the CH-DFVAR model, with $N = 3$ and $p = 2$, are considered and for a selection of these we report their features and goodness-of-fit results in Table 2. From the perspective of selecting the model with the lowest BIC, specification (1) is preferred. In this specification, we have the conditional heteroskedasticity included by the scalar BEKK updating in equation (13), and the time-variation of the autoregressive coefficients facilitated by five dynamic factors ($r = 5$). The interpretability and identification of each factor is ensured by a set of restrictions imposed on the factor loading matrices Φ_i^f, $i = 1, \dots, 5$. We treat each Φ_i^f as a zero matrix on the outset. In Table 2, the non-zero elements for Φ_i^f are listed for the first N columns, the same non-zero elements are also imposed to the second set of N columns which represents the autoregressive coefficients associated with lag 2. In case of our preferred specification (1), the interpretation is as follows. Dynamic factor $f_{t,1}$ captures time-variation in the spillover intensity of financial shocks to the real economy, that is the impact of lagged bond spreads on IP growth and inflation. The factors $f_{t,2}, f_{t,3}, f_{t,4}$ account for the respective time-varying persistence in each of the three variables in y_t: IP growth, inflation and bond spread. The last factor $f_{t,5}$ captures the common changes in the spillovers from IP growth to inflation, and vice versa. For other specifications, some other variants of restrictions on the factor loading matrices Φ_i^f, $i = 1, \ldots, 5$ are explored, in particular those for specifications (2) to (5). The results indicate that specification (4) is preferred in terms of AIC: it introduces different financial spillover factors for IP growth and inflation. However, specification (4) requires an additional parameter compared to specification (1) and this this may not be justified in view of a parsimonious model, according to BIC. The specifications (6) and (7) are provided in Table 2 to show the strong support for the time-variation of the autoregressive coefficients and the conditional heteroscedasticity in the CH-DFVAR model. Although the standard VAR(2) model, with or without the scalar BEKK, has a lower number of parameters, the fit is much worse.

5.2 Time-varying autoregressive parameters

The estimated coefficients, together with their asymptotic standard errors, of the parameters in the CH-DFVAR specification (1) are obtained using the maximum likelihood method as discussed in Section 2 and are presented in Table 3. The estimates of Φ^c, the constant long-run part of the autoregressive coefficient matrices in equation (3) for lags 1 and 2, are for more than half of all coefficients in Φ^c significantly different from zero, at the 5% significance level. For example, we find that the bond spread significantly affects both macro variables at both lags. It is well established that interpreting reduced-form coefficients is of limited economic use. We therefore present an orthogonalized impulse response function analysis below. The specification of the time-varying autoregressive coefficients and their estimates are also presented in Table 3. For each factor $f_{t,i}$, for $i = 1, \ldots, 5$, we have at least one loading in Φ_i^f that is estimated as being significantly different from zero. Overall, the coefficients of the factor loading matrices Φ_i^f corresponding to lag 2 of y_t, appear to be more significantly exposed to time-variation. Since lag 2 has a particular impact on cyclical dynamics, we may conclude that the time-varying parameter features may be implied from business cycle features in IP growth,

Table 2: Different CH-DFVAR Model Specifications and Fit

The CH-DFVAR model is for y_t = (IP growth, inflation, bond spread)$'$ and it is given by the equations (3), (4) and (13). We consider seven different specifications of the model. We indicate whether some features are included in the model specification, we report the design of the factor loading matrices Φ_i^f in equation (3) by indicating their non-zero entries, and we provide the maximized value of the log-likehood function together with the corresponding information criteria AIC and BIC. The latter criterion selects the first specification as preferred. The specifications (6) and (7) are performing most poorly amongst these seven specifications.

Specification	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			Features of Model				
# lags	2	2	2	2	2	2	2
Φ^c	✓	✓	✓	✓	✓	✓	✓
Φ^f	✓	✓	✓	✓	✓		
Scalar BEKK	✓	✓	✓	✓	✓	✓	
# factors	5	6	3	6	5	0	0
# parameters	45	46	39	46	45	26	24
		Nonzero Entries in Factor Loading Matrices Φ_i^f for Lags 1 and 2					
Φ_1^f	(1,3) , (2,3)	(1,3) , (2,3)	(1,3) , (2,3)	(1,3)	(1,3)		
Φ_2^f	(1,1)	(1,1)	(2,2)	(2,3)	(2,3)		
Φ_3^f	(2,2)	(2,2)	(1,2) , (2,1)	(1,1)	(1,1) , (3,3)		
Φ_4^f	(3,3)	(3,3)		(2,2)	(2,2)		
Φ_5^f	(1,2) , (2,1)	(1,2)		(3,3)	(1,2) , (2,1)		
Φ_6^f		(2,1)		(1,2) , (2,1)			
			Model selection criteria				
log likelihood	-199.21	-200.46	-227.33	-197.41	-201.82	-275.27	-407.3369
AIC	488.42	492.91	532.67	486.82	493.64	606.54	862.6738
BIC	**685.30**	694.17	703.29	688.08	690.52	729.04	967.6744

inflation and bond spread. The estimates of the autoregressive coefficients $\phi_1, \ldots, \phi_5$ for the dynamic factors $f_{t,1}, \ldots, f_{t,5}$, respectively, imply persistent time-variation for the first and third factors; these are associated with financial-macro spillovers and inflation persistence, respectively. The other three factors do not appear to be persistent. These finding are confirmed in Figure 4 where we present the estimated factors from the Kalman filter smoother. We clearly observe the high persistence in the financial-macro spillover factor and in the financial persistence factor while the remaining three factors are more noisy. The financial-macro spillover factor $f_{t,1}$ shows a temporary but strong increase during the financial crisis. The inflation persistence factor $f_{t,3}$ is high during the 1970s, a period of high inflation rates, and it is low during the early 2000s, a period of low inflation rates. The spillover macro shock factor $f_{t,5}$ exhibits stronger variation in the first part of our sample, up to approximately 1983. This feature may be indicative of the overall decline in macroeconomic volatility from the mid-1980s, which is known as the "Great Moderation"; see Blanchard and Simon (2001) for a discussion.

Table 3: Parameter Estimation Results for Final CH-DFVAR Model

The CH-DFVAR model is for y_t = (IP growth, inflation, bond spread)$'$ and it is given by the equations (3), (4) and (13). We consider the preferred model specification (1) from Table 2, with $N = 3$, $p = 2$ and $r = 5$, for which we present the design of the factor loading matrices Φ_i^f, for $i = 1, \ldots, 5$, and a short description of its interpretation. All parameter estimates are provided with their asymptotic standard errors in parantheses below.

Factor Loading Matrices Φ_i^f for $i = 1, \ldots, 5$, with nonzero elements indicated by *

Factor 1 Φ_1^f	Factor 2 Φ_2^f	Factor 3 Φ_3^f	Factor 4 Φ_4^f	Factor 5 Φ_5^f
financial-real shock spillover	persistence IP growth	persistence inflation rate	persistence bond spread	spillover macro shocks
$\begin{bmatrix} 0 & 0 & * & 0 & 0 & * \\ 0 & 0 & * & 0 & 0 & * \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$	$\begin{bmatrix} * & 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & * & 0 & 0 & * & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & * & 0 & 0 & * \end{bmatrix}$	$\begin{bmatrix} 0 & * & 0 & 0 & * & 0 \\ * & 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$

Parameter Estimates

	DFVAR Φ^c in equation (3)			DFVAR Φ_i^f in equation (3)		Factors φ_i in equation (4)		Scalar BEKK in equation (13)	
	Lag 1	Lag 2		Lag 1	Lag 2				$\alpha^2, \beta^2, \Omega$
$\Phi^c(1,1)$	0.089	0.098	$\Phi_1^f(1,3)$	-0.2771	0.2726	φ_1	0.8788	α^2	0.0735
	(0.054)	(0.051)		(0.313)	(0.068)		(0.0716)		(0.013)
$\Phi^c(2,1)$	-0.026	-0.006	$\Phi_1^f(2,3)$	-0.0989	0.4638	φ_2	0.0607	β^2	0.8532
	(0.016)	(0.015)		(0.0893)	(0.0771)		(0.1619)		(0.0226)
$\Phi^c(3,1)$	-0.027	-0.024	$\Phi_2^f(1,1)$	0.0907	-0.4232	φ_3	0.9771	$\Omega(1,1)$	0.1214
	(0.010)	(0.010)		(0.327)	(0.0979)		(0.0141)		(0.0187)
$\Phi^c(1,2)$	-0.009	-0.291	$\Phi_3^f(2,2)$	0.1126	0.0513	φ_4	0.0042	$\Omega(2,1)$	-0.0033
	(0.118)	(0.121)		(0.1045)	(0.0262)		(0.1569)		(0.0059)
$\Phi^c(2,2)$	0.426	0.158	$\Phi_4^f(3,3)$	0.3979	0.205	φ_5	0.1756	$\Omega(3,1)$	-0.0032
	(0.047)	(0.108)		(0.076)	(0.1711)		(0.1891)		(0.0038)
$\Phi^c(3,2)$	-0.017	0.005	$\Phi_5^f(2,1)$	0.1448	0.0609			$\Omega(2,2)$	0.0535
	(0.024)	(0.025)		(0.086)	(0.0252)				(0.0064)
$\Phi^c(1,3)$	-0.656	0.504	$\Phi_5^f(1,2))$	-0.0432	-0.7517			$\Omega(3,2)$	-0.0032
	(0.162)	(0.153)		(0.0479)	(0.1677)				(0.0031)
$\Phi^c(2,3)$	-0.160	0.104						$\Omega(3,3)$	0.0324
	(0.056)	(0.055)							(0.0046)
$\Phi^c(3,3)$	1.294	-0.339							
	(0.050)	(0.048)							

The index pair (i, j) refers to the matrix element (i, j) with $i, j = 1$ (IP growth), 2 (inflation), 3 (bond spread).

5.3 Time-varying variances and covariances

In Table 3 we also report estimates for the coefficients of the scalar BEKK parameters in equation (13). The estimates for the diagonal elements of the variance intercept matrix Ω are all strongly significant, as well as those for the persistence parameter β^2 and the parameter capturing the impact of innovations, α^2. We expect macroeconomic and financial volatility to follow persistent processes and we confirm that the overall persistence is $\alpha^2 + \beta^2 = 0.93$. The three estimated covariances in Ω are negative but they are not significantly different from zero. The filtered estimates of the variances and covariances in H_t are provided in Figure 5. The time-varying variance estimates show high volatility levels during the rescession periods in the 1970s and

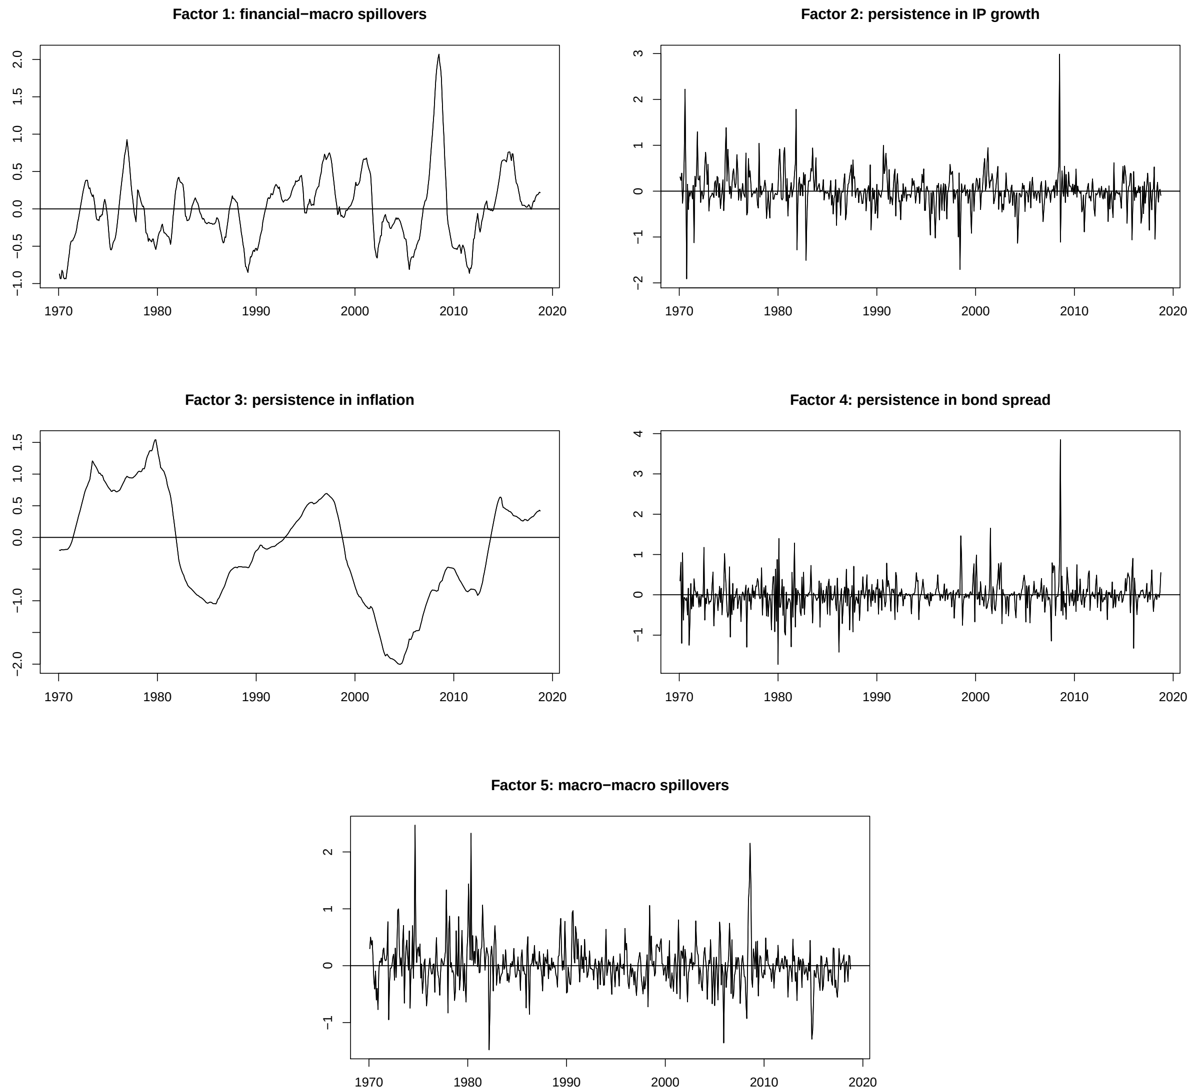

Figure 4: The CH-DFVAR model is for $y_t = (\text{IP growth, inflation, bond spread})'$ and it is given by the equations (3), (4) and (13). We consider the preferred model specification (1) from Table 2 and with the estimated parameters reported in Table 3, where also the features of the model and the interpretation of the factors are provided. We present the estimated factors which are obtained from the Kalman filter smoother.

early 1980s, and in particular, during the financial crisis and its aftermath. We observe strong variations in the filtered estimates of the covariances over time, but they fluctuate around a long-term mean of zero.

5.4 Impulse response function analysis

To conclude our analysis, we carry out an impulse response function (IRF) analysis based on the CH-DFVAR model with the parameters replaced by their corresponding estimates as reported in Table 3. Similarly to Prieto et al. (2016), we identify the shocks by carrying out a Cholesky decomposition of the covariance matrix, see

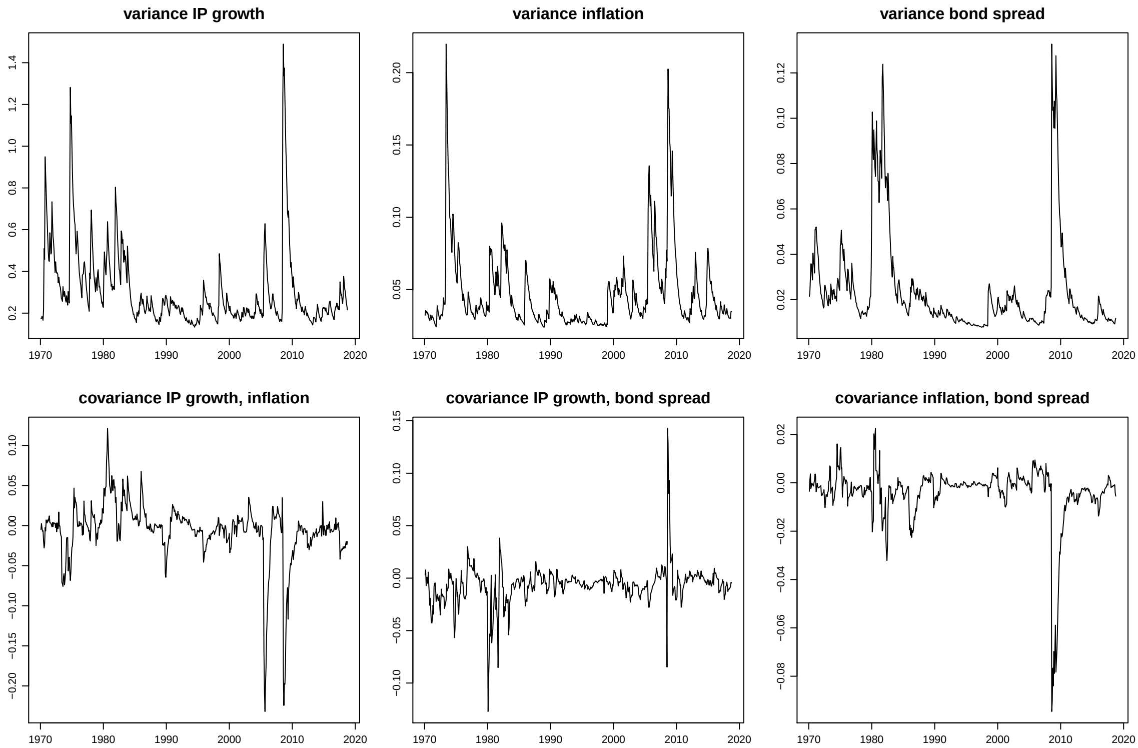

Figure 5: The CH-DFVAR model is for y_t = (IP growth, inflation, bond spread)$'$ and is the same as the one used for Figure 4. We present conditional variances and covariances which are obtained as discussed in Sections 2.4 and 2.5.

the discussion in Section 3. We also adopt their reasoning of ordering the macro variables before the financial ones, because the impact of financial shocks on macroeconomic variables often occurs with a delay. Therefore, the order in the chain is IP growth $\rightarrow$ inflation rate $\rightarrow$ bond spread. Due to the time-variation in both the autoregressive coefficient matrices and covariance matrices, the obtained impulse response functions differ from period to period. We are particularly interested in the potentially adverse impacts of financial shocks on macro variables, which have been found to be more pronounced during crisis periods; see the discussion in Hubrich and Tetlow (2015). Therefore, we average the time-varying responses of IP growth and the inflation rate to financial shocks over the recession months as defined by the National Bureau of Economic Research (NBER); the dates are obtained from the FRED data base https://fred.stlouisfed.org/. In Figure 6 we present the average impulse responses with 68% confidence intervals. We find that financial shocks, which imply a widening of the bond spread, have bigger impacts in times of crises, in terms of magnitude as well as in terms of persistence. While the deflationary effect of positive financial shocks is significant, both in crisis and non-crisis periods, we find a strong significant effect on output growth only for the recession period 2008/09.

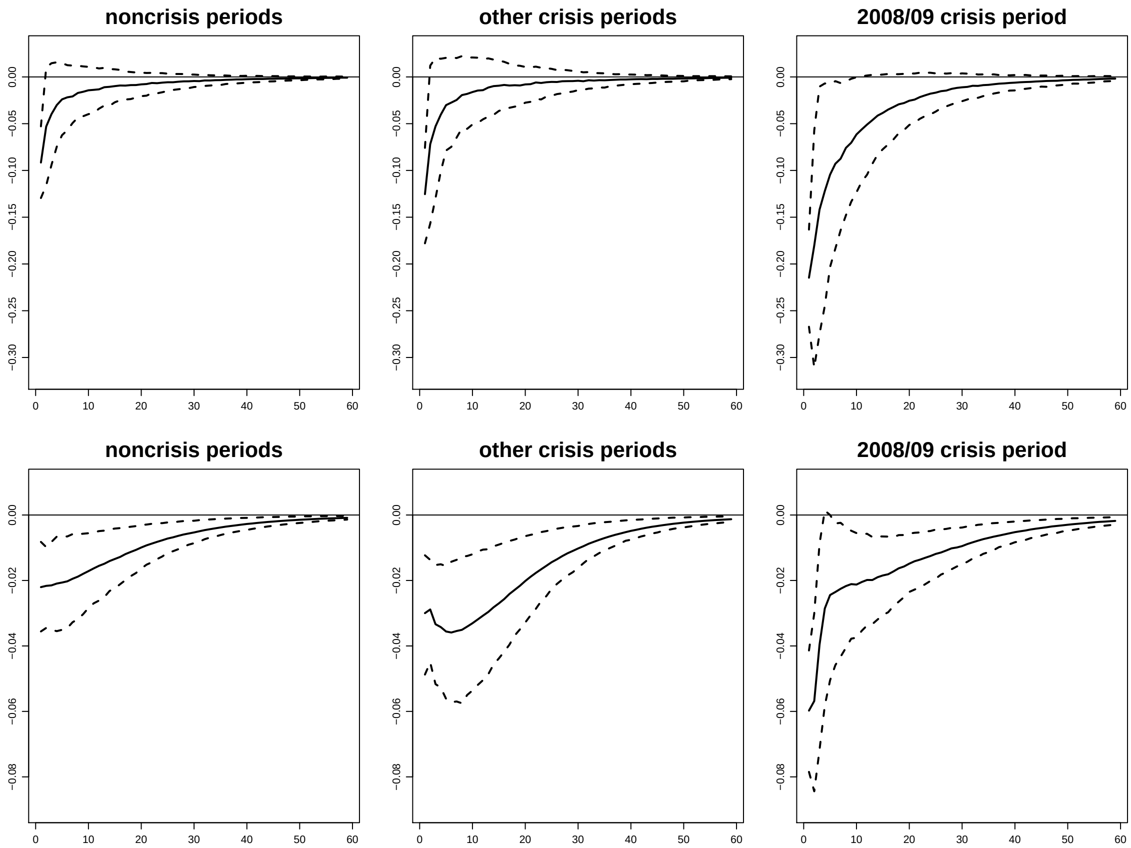

Figure 6: The CH-DFVAR model is for $y_t = (\text{IP growth, inflation, bond spread})'$ and is the same as the one used for Figure 4. We present impulse responses, with 68% confidence intervals, of IP growth (upper panels) and inflation rate (lower panels)) to the bond spread shock. The impulse responses are averages over different sub-periods in our sample: noncrises, crises, and the financial crisis (2008/09) periods.

6 Conclusion

We have considered a vector autoregressive (VAR) model with time-varying autoregressive coefficient matrices and with conditionally heteroskedastic disturbances. The specification for the time-varying VAR matrices relies on dynamic factors and is flexible as it allows a range of different specifications. All elements in the VAR matrices can vary while our framework also allows for more parsimonious formulations of the dynamic evolution of the autoregressive coefficients. The conditions for identification are easily verified. The analysis relies on the maximum likehood estimation method and the estimation of the dynamic factors (signal extraction) is done via the Kalman filter and related methods. The full estimation process is computationally fast and it offers an alternative to the widely used Bayesian methods of analysis. We carry out a Monte Carlo study to confirm the reliability of the estimation method. In the empirical study, we find that financial shocks in the U.S. economy have larger impacts during recession periods, both in terms of magnitude and in terms of persistence.

References

Abbate, A., Eickmeier, S., Lemke, W., and Marcellino, M. (2016). The changing international transmission of financial shocks: evidence from a classical time-varying favar. *Journal of Money, Credit and Banking*, 48(4):573–601.

Blanchard, O. and Simon, J. (2001). The long and large decline in us output volatility. *Brookings papers on economic activity*, 2001(1):135–174.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327.

Bougerol, P. and Picard, N. (1992). Strict stationarity of generalized autoregressive processes. *The Annals of Probability*, pages 1714–1730.

Brunnermeier, M., Eisenbach, T., and Sannikov, Y. (2013). Macroeconomics with financial frictions: A survey. *in: Advances in Economics and Econometrics: Tenth World Congress of the Econometric Society, Vol. 2. Edited by D. Acemoglu, M. Arellano and E. Dekel*, pages 3–94.

Canova, F. (2007). *Methods for Applied Macroeconomic Research*. Princeton University Press.

Canova, F. and Ciccarelli, M. (2004). Forecasting and turning point predictions in a bayesian panel var model. *Journal of Econometrics*, 120:327–359.

Canova, F. and Ciccarelli, M. (2009). Estimating multicountry var models. *International Economic Review*, 50:929–959.

Chauvet, M. and Potter, S. (2013). Forecasting output. *in: Handbook of Economic Forecasting, Volume 2 Part A, Edited by G. Elliot and A. Timmermann*, pages 141–194.

Christiano, L., Eichenbaum, M., and Evans, C. (1996). The effects of monetary policy shocks: Evidence from the flow of funds. *The Review of Economics and Statistics*, 78(1):16–34.

Clements, M. P. and Galvão, A. B. (2013). Forecasting with vector autoregressive models of data vintages: US output growth and inflation. *International Journal of Forecasting*, 29(4):698–714.

Cooley, T. F. and LeRoy, S. F. (1985). Atheoretical Macroeconomics: A Critique. *Journal of Monetary Economics*, 16:283–308.

Durbin, J. and Koopman, S. J. (2012). *Time series analysis by state space methods*. Oxford University Press, 2nd edition.

Eickmeier, S., Lemke, W., and Marcellino, M. (2015). Classical time varying factor-augmented vector autoregressive models - estimation, forecasting and structural analysis. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 178(3):493–533.

Engle, R. F. and Kroner, K. F. (1995). Multivariate simultaneous generalized arch. *Econometric theory*, 11(1):122–150.

Galvão, A. B. and Owyang, M. (2018). Financial stress regimes and the macroeconomy. *Journal of Money, Credit and Banking*, 50(7):1479–1505.

Gertler, M. and Kiyotaki, N. (2011). Financial intermediation and credit policy in business cycle analysis. *in: Handbook of Monetary Economics, Vol. 2A. Edited by B. Friedman and M. Woodford*, pages 547–599.

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

Hubrich, K. and Tetlow, R. J. (2015). Financial stress and economic dynamics: The transmission of crises. *Journal of Monetary Economics*, 70:100–115.

Inoue, A. and Kilian, L. (2013). Inference on impulse response functions in structural VAR models. *Journal of Econometrics*, 177(1):1–13.

Justiniano, A. and Primiceri, G. E. (2008). The time-varying volatility of macroeconomic fluctuations. *American Economic Review*, 98(3):604–41.

Koopman, S. J., Mallee, M. I., and van der Wel, M. (2010). Analyzing the term structure of interest rates using the dynamic nelson-siegel model with time-varying parameters. *Journal of Business and Economic Statistics*, 28:329–343.

Lütkepohl, H. (2005). *New introduction to multiple time series analysis*. Springer Science & Business Media.

Magnus, J. R. and Neudecker, H. (2019). *Matrix Differential Calculus with Applications in Statistics and Econometrics*. John Wiley, New York, 3rd edition.

Pagan, A. (1987). Three Econometric Methodologies: A Critical Appraisal. *Journal of Economic Surveys*, 1(1):3–24.

Prieto, E., Eickmeier, S., and Marcellino, M. (2016). Time variation in macro-financial linkages. *Journal of Applied Econometrics*, 31:1215.123.

Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72:821–852.

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

Stock, J. H. and Watson, M. W. (2001). Vector autoregressions. *Journal of Economic Perspectives*, 15(4):101–115.

Stock, J. H. and Watson, M. W. (2011). Dynamic factor models. *in: The Oxford Handbook of Economic Forecasting. Edited by Michael P. Clements and David F. Hendry*.