Monitoring multicountry macroeconomic risk

Dimitris Korobilis and Maximilian Schröder

Monitoring multicountry macroeconomic risk

Dimitris Korobilis[1] and Maximilian Schröder[2]

[1]University of Glasgow, and Rimini Center for Economic Analysis
[2]BI Norwegian Business School, and Norges Bank

August 3, 2023

Abstract

We propose a multicountry quantile factor augmeneted vector autoregression (QFAVAR) to model heterogeneities both across countries and across characteristics of the distributions of macroeconomic time series. The presence of quantile factors allows for summarizing these two heterogeneities in a parsimonious way. We develop two algorithms for posterior inference that feature varying level of trade-off between estimation precision and computational speed. Using monthly data for the euro area, we establish the good empirical properties of the QFAVAR as a tool for assessing the effects of global shocks on country-level macroeconomic risks. In particular, QFAVAR short-run tail forecasts are more accurate compared to a FAVAR with symmetric Gaussian errors, as well as univariate quantile autoregressions that ignore comovements among quantiles of macroeconomic variables. We also illustrate how quantile impulse response functions and quantile connectedness measures, resulting from the new model, can be used to implement joint risk scenario analysis.

Keywords: quantile VAR; MCMC; variational Bayes; dynamic factor model

JEL Classification: C11, C32, E31, E32, E37, E66

1   Introduction

The so-called Great Recession of 2008-2009 marked the beginning of an era in which global shocks are more pervasive, are able to generate domino effects, and can increase macroeconomic risks to unprecedented levels. Recent evidence is provided by Adrian et al. (2019) who find that the expected distribution of GDP growth skews left during recessions. This skewness can change over time (Jensen et al., 2020) and is positively related to macroeconomic volatility (Bekaert and Popov, 2019). Similar distributional asymmetries occur for inflation (Korobilis, 2017; López-Salido and Loria, 2019), as for numerous other economic and financial time series. Although some global risks – such as global inflation and oil crises – are eerily familiar to economists, the current unsustainable levels of debt, low growth, and the climate emergency create new challenges for macroeconomic policy-makers. In the euro area, which shares common monetary, regulatory, and other policies, measuring and monitoring the heterogeneous performance of countries facing modern global risks is a particularly challenging quantitative exercise. As an example, following the first response to the coronavirus (Covid-19) health crisis, euro-area GDP was 4.9% below its prepandemic level in the first quarter of 2021. Nevertheless, country-level performance was remarkably heterogeneous, with Ireland reporting growth of 13.2% and Spain experiencing a contraction of -9.3% (Muggenthaler et al., 2021).

In this paper we develop a novel quantile regression methodology to capture, in a parsimonious and interpretable way, the heterogeneous responses of euro-area and country-specific macroeconomic aggregates to global risks. Our approach builds on combining the strengths of vector autoregressions (VARs) for structural inference, with the flexibility of modeling individual percentiles of the data distribution of macro data using quantile regression methods. The main challenge with modeling multivariate quantiles is that different percentiles of different variables might be correlated, which is not typically an issue in univariate quantile models (where each quantile level is estimated independently).[1] With just two macroeconomic variables (y_1, y_2) and three quantile levels (q_1, q_2, q_3), there are numerous ways any two quantiles of these variables could be correlated. As a result, despite the fact that VAR modeling of the mean of y_1, y_2 is a bivariate system, a quantile VAR would require specifying all six variables $y_{1(q_1)}, y_{1(q_2)}, y_{1(q_3)}, y_{2(q_1)}, y_{2(q_2)}, y_{2(q_3)}$ as endogenous. When the interest is in modeling quantiles of macroeconomic variables for many countries, VAR inference on quantiles can quickly become high-dimensional, over-parameterized and, potentially, computationally

---

[1]From another perspective, Adrian et al. (2021) argue that during crises the joint distribution of economic and financial conditions becomes multimodal. Such multimodalities support the argument that percentiles of joint distributions are not symmetrically correlated.

cumbersome.

We solve this modeling conundrum by specifying a quantile factor-augmented vector autoregression (QFAVAR) that extends the popular factor-augmented VAR approach outlined in Bernanke et al. (2005) and Stock and Watson (2005) to the quantile regression setting. In a setting with several macro variables for several countries, we extract variable-specific quantile factors for three percentiles of interest (10th, 50th, and 90th percentiles). As a result, the factors not only capture common dynamics across the cross-section of the data (which is a key feature of multi-country factor studies such as Kose et al. 2003), but also they are a parsimonious way of modeling cross-quantile dependence among the original variables. The benefit is that, similar to the recent methodologies in Chen et al. (2021) and Korobilis and Schröder (2022), the macroeconomic variables load onto quantile factors independently for each quantile level, allowing computational convenience and numerical stability. However, the QFAVAR allows all quantile factors to be dynamically correlated for all percentile levels by means of a joint VAR-state evolution.

Methodologically, the new QFAVAR adds novel features to various interconnected literatures in macroeconometrics. We first build on established literature that uses dynamic factor models (DFMs) to characterize comovements and heterogeneities among different countries. Kose et al. (2003) use multicountry dynamic factor models to measure the degree of synchronization of business cycles; Ciccarelli and Mojon (2010) and Mumtaz and Surico (2012) use common factors to measure global inflation. A common limitation of the traditional DFM/FAVAR approach is the reliance on the normality assumption of common and idiosyncratic disturbances, which does not sufficiently capture any asymmetries in higher moments of the empirical distribution of macroeconomic data.[2] Next, our proposal to estimate quantile factors and combine these with VAR dynamics, adds a new tool to a recent literature that is otherwise limited to specifying static quantile factors. Key papers include Ando and Bai (2020), Chen et al. (2021), Clark et al. (2021), Korobilis and Schröder (2022), and Ma et al. (2021). Finally, the QFAVAR adds to another emerging literature that combines quantile regressions and VAR methods in order to identify the asymmetric effects of various macroeconomic shocks. See for example Castelnuovo and Mori (2022), Forni et al. (2021) and Loria et al. (2019). Although existing methodologies for estimating quantile VARs (QVARs) can be empirically useful and relevant, they may rely on simplifications that are context-specific.[3]

---

[2]Papers such as Korobilis (2013) and Koop and Korobilis (2014) specify flexible FAVARs with time-varying parameters and stochastic volatility; Gorodnichenko and Ng (2017) explicitly estimate mean and volatility factors. Nevertheless, such factor model approaches are restricted to modeling flexibly only the first two moments of the data distribution.

[3]For example, the Bayesian approaches in Mumtaz and Surico (2015) and Schüler (2020) might become

In contrast, the proposed QFAVAR approach is general enough to be used in a multitude of other empirical problems where asymmetric shocks are evident, for example, issues involving climate risks, macroeconomic uncertainty, or financial shocks.

Our first contribution is to establish the workings of such a novel specification and to show how it allows for parsimonious VAR inference for quantiles without sacrificing flexibility and generality. The second contribution is to derive and test numerically likelihood-based estimators for inference in the QFAVAR. In particular, we adopt Bayesian methods and modern priors from the statistics and machine learning literature that lead to tuning-free penalized estimation in high dimensions, and we derive two algorithms for posterior inference. Our benchmark estimation is based on a Markov chain Monte Carlo (MCMC) algorithm that is a generalization of the Gibbs samplers proposed in Bernanke et al. (2005) and Kose et al. (2003). The second algorithm is based on variational Bayes (VB) inference that extends the static quantile factor estimator proposed in Korobilis and Schröder (2022); it provides a fast and convenient approximation to the joint parameter posterior. By leveraging machine learning methods, the second algorithm is particularly useful for applications where computational time is especially important, as are frequent in policy work.

Our third contribution is empirical, as we apply the new tool to the problem of assessing the effects of global risks to a series of euro-area macroeconomic variables. We collect five macroeconomic variables for nine countries, namely inflation, industrial production, the long-term interest rate, an index of economic sentiment, and an index country-level financial stress. We extract quantile factors from each macro variable, by aggregating over all nine countries, such that each quantile factor approximates the distribution of the respective euro area aggregate. For example, the quantile factor extracted from the nine country-level industrial production (IP) series captures the quantiles of aggregate IP for the euro area. From a measurement perspective these quantile factors are superior in capturing cross-sectional and distributional heterogeneities, compared to fitting univariate quantile regressions directly to the euro area aggregates, but allow insights into country-level heterogeneity at the same time.[4] By definition the QFAVAR allows us to augment the

---

computationally cumbersome using the dimensions we consider in this paper. The contribution by Ando et al. (2022) is restricted in that the covariance matrix is based on observed data, such that QVAR estimation is implemented using univariate quantile regressions. Chavleishvili and Manganelli (2020) present a bivariate quantile VAR that does not allow different quantiles of the two series to interact. Finally, the setting in White et al. (2015), while useful for value-at-risk applications in finance, has several limitations for structural macro inference.

[4]A univariate quantile regression fitted to euro-area IP will not capture correctly the distribution of IP, especially in the tails. This is due to the lack of observations in the tails as well as the fact that IP is measured with error. By combining the information in the IP series of multiple countries, we expect the quantile factor to suffer less from these issues and provide superior estimates of the distributional characteristics of aggregate IP.

quantile factors with (observed) global factors, and we choose to include global measures of inflation, supply chain pressures, financial conditions, and economic policy uncertainty.

The empirical facts can be summarized as follows. Quantile factor estimates of 10th, 50th and 90th percentiles are characterized by evident heterogeneity, implying asymmetries in the distribution of unobserved factors. The QFAVAR is significantly better to a FAVAR with symmetric Gaussian errors in forecasting the left and right tails of the distribution of inflation and IP in the euro-area countries. This observation is true, in particular, for the short, one-month-ahead horizon of both variables. Additionally, we find that the QFAVAR is superior to univariate quantile autoregressive models with or without exogenous global predictors, and the quantile dynamic factor model (that is, a special case of the QFAVAR without any global variables). These numerical results clearly show not only the benefits of adopting a multivariate approach to quantile regression forecasting, but also that there are benefits from augmenting the multi-country model with relevant global predictors. We illustrate, by means of quantile impulse response functions, quantile forecast error variance decompositions, and quantile connectedness graphs, that the QFAVAR captures a large variety of heterogeneities across different quantiles of macroeconomic variables of interest. In doing so, we also add to an open debate. Adrian et al. (2019) argue that financial conditions are an important predictor of downside risks to GDP. Plagborg-Møller et al. (2020) question the validity of this argument by providing detailed empirical evidence that no financial time series is persistently informative in univariate quantile regressions. In contrast, we find that (global) financial and economic conditions can be very informative for country-level macro risks, once the potential linkages are modelled as a multivariate system.

The next section describes the exact specification of the QFAVAR and how estimation and inference can be tackled using Bayesian methods. As the QFAVAR is a new model in the literature, Section 3 undertakes a battery of numerical exercises that establish its usefulness for monitoring multi-country macroeconomic risks. Section 4 concludes the paper.

2 Econometric Methodology

Our starting point is the factor augmented vector autoregressive model from Bernanke et al. (2005) and Stock and Watson (2005), adopted for a panel of several macroeconomic time series for several countries. This modeling approach involves extracting a smaller vector of latent factors from the large panel of macroeconomic data. The latent factors evolve jointly with observed variables as a lower-dimensional vector autoregression (VAR). This setting is established in macroeconomics, and the reader should consult Stock and Watson (2016) for

a thorough review. Here we explain a conceptually straightforward extension of the FAVAR to the quantile setting, show how this extension results in an inherently high-dimensional model, outline how Bayesian inference can help tackle estimation challenges, and, finally, we show how the proposed quantile specification can be deployed for traditional structural VAR analysis.

2.1 A multi-country quantile FAVAR (QFAVAR)

Let y_{ijt} denote macroeconomic/financial indicator i for country j observed at time t, for $i = 1, ..., m$, $j = 1, ..., n$ and $t = 1, ..., T$. We characterize the distribution of the $mn \times 1$ vector $\boldsymbol{y}_t = [y_{11t}, ..., y_{1nt}, ..., y_{m1t}, ..., y_{mnt}]'$ by grouping its elements into unobserved, indicator-specific factors $f_{t,(q)}^i$ for each quantile level $q = q_1, ..., q_R$, where $q_r \in (0, 1)$ and $q_{r-1} < q_r$. We also assume global-level factors summarized in the $k \times 1$ vector of observed variables $\boldsymbol{g}_t$. The quantile factor model strategy begins by specifying the qth conditional quantile of y_{ijt} as a linear function of the global indicators and the indicator-specific factor, which is of the form

$$Q_q\left(y_{ijt}|\boldsymbol{g}_t\right) = c_{ij(q)} + \boldsymbol{\gamma}_{ij(q)}\boldsymbol{g}_t + \lambda_{ij(q)}f_{t,(q)}^i, \tag{1}$$

where $c_{ij(q)}$ is a scalar intercept, $\boldsymbol{\gamma}_{ij(q)}$ is a $1 \times k$ vector of loadings on the observed global factors $\boldsymbol{g}_t$, and $\lambda_{ij(q)}$ is the scalar loading (weight) of the scalar, unobserved, indicator-specific quantile factor $f_{t,(q)}^i$.[5] Following the probabilistic approach in Korobilis and Schröder (2022) this quantile factor model can be represented as a parametric regression of the form

$$y_{ijt} = c_{ij(q)} + \gamma_{ij(q)}\boldsymbol{g}_t + \lambda_{ij(q)}f_{t,(q)}^i + u_{ijt(q)}, \tag{2}$$

where $u_{ijt(q)} \sim AL\left(0, \sigma_{ij(q)}^2, q\right)$ is an asymmetric Laplace disturbance term; that is, it has the functional form

$$f\left(u_{ijt(q)}\right) = \frac{q(1-q)}{\sigma_{ij(q)}^2}\left\{e^{\left[(1-q)\frac{u_{ijt(q)}}{\sigma_{ij(q)}^2}\right]}\mathbb{I}(u_{ijt(q)} \leq 0) + e^{\left[(-q)\frac{u_{ijt(q)}}{\sigma_{ij(q)}^2}\right]}\mathbb{I}(u_{ijt(q)} > 0)\right\}, \tag{3}$$

with $\mathbb{I}$ denoting the indicator function. Similar to a Bayesian linear regression, where the Gaussian residual is centred around zero, the asymmetric Laplace residual has the q-th quantile equal to zero.

An important modeling feature of our approach is that the indicator-specific quantile

---

[5]The fact that we extract indicator-specific factors, and we don't just allow all variables to load on all factors, helps with identification of the factor model. We only impose normalization restrictions, where for the factor corresponding to quantile q_r, we normalize the loading of the r^{th} series to be one.

factors and the global factors are contemporaneously and dynamically correlated with each other via a vector autoregressive (VAR) model with p lags. We define the $mr \times 1$ vector $\boldsymbol{F}_t = \left[f_{t(q_1)}^1, ..., f_{t(q_1)}^m, f_{t(q_2)}^1, ..., f_{t(q_2)}^m, ..., f_{t(q_r)}^1, ..., f_{t(q_r)}^m \right]'$ that summarizes all unobserved factors at all quantile levels. The VAR(p) that characterizes the joint dynamics of the quantile factors and the global factors is of the form

$$\begin{bmatrix} \boldsymbol{F}_t \\ \boldsymbol{g}_t \end{bmatrix} = \boldsymbol{v} + \boldsymbol{\Phi}_1 \begin{bmatrix} \boldsymbol{F}_{t-1} \\ \boldsymbol{g}_{t-1} \end{bmatrix} + ... + \boldsymbol{\Phi}_p \begin{bmatrix} \boldsymbol{F}_{t-p} \\ \boldsymbol{g}_{t-p} \end{bmatrix} + \boldsymbol{\varepsilon}_t, \tag{4}$$

where $\boldsymbol{v}$ is an $l \times 1$ vector of intercept terms, $\boldsymbol{\Phi}_c$ are $l \times l$ matrices of autoregressive coefficients for lagged term $c = 1, ..., p$, and $\boldsymbol{\varepsilon}_t \sim N(\boldsymbol{0}, \boldsymbol{\Omega})$ is the disturbance term with $\boldsymbol{\Omega}$ an $l \times l$ full, symmetric and positive definite covariance matrix. Here, $l = mr + k$ is the joint dimension of the vectors $\boldsymbol{F}_t$ and $\boldsymbol{g}_t$. As a result, in contrast to Chen et al. (2021) and Korobilis and Schröder (2022) who estimate factors independently for each quantile, equation (4) allows for complex patterns of dynamic correlations among the quantiles to affect the estimation outcomes of the quantile factors. At the same time, as we show in subsection 2.3, this latter equation is important in order to perform structural VAR inference using the QFAVAR. Finally, the model maintains its computational simplicity, as the disturbances $u_{ijt(q)}$ are independent from each other for all i, j, q. Therefore, equation (2) is a collection of univariate quantile regressions that can be estimated independently from one another. In contrast, the VAR of the state equation (4) can become quite large when either considering many indicator-specific or global factors, or many quantile levels. However, it is trivial to draw from an established literature on large Bayesian VARs in order to alleviate high-dimensionality concerns. In particular, we adopt the efficient algorithm of Carriero et al. (2022). We next discuss likelihood-based inference in the QFAVAR in detail.

2.2 Likelihood, priors, and posterior inference

The QFAVAR consists of seemingly disjointed equations (2) and (4). In order to write them as a joint system and facilitate likelihood-based inference, we simplify our notation by dropping intercepts and assuming one lag in the VAR part of the model. Under these simplifications – and as shown in detail in appendix A – we can combine equations (2) and (4) into the following linear state-space system

$$\begin{bmatrix} \boldsymbol{Y}_t \\ \boldsymbol{g}_t \end{bmatrix} = \begin{bmatrix} \boldsymbol{\Lambda} & \boldsymbol{\Gamma} \\ \boldsymbol{0} & \boldsymbol{I} \end{bmatrix} \begin{bmatrix} \boldsymbol{F}_t \\ \boldsymbol{g}_t \end{bmatrix} + \begin{bmatrix} \boldsymbol{u}_t \\ \boldsymbol{0} \end{bmatrix}, \tag{5}$$

$$\begin{bmatrix} \boldsymbol{F}_t \\ \boldsymbol{g}_t \end{bmatrix} = \boldsymbol{\Phi} \begin{bmatrix} \boldsymbol{F}_{t-1} \\ \boldsymbol{g}_{t-1} \end{bmatrix} + \boldsymbol{\varepsilon}_t. \tag{6}$$

In this matrix notation $\boldsymbol{Y}_t$ is an $nmr \times 1$ vector with the vector $\boldsymbol{y}_t$ repeated r times; $\boldsymbol{\Lambda}$ is an $nmr \times mr$ block diagonal matrix with the quantile-specific loadings $\boldsymbol{\lambda}_{(q)}$ on its diagonal; $\boldsymbol{\lambda}_{(q)}$ in turn is an $nm \times m$ block-diagonal matrix with the factor-specific loading $\boldsymbol{\lambda}_{i(q)} = \left[\lambda_{i1(q)}, ..., \lambda_{im(q)}\right]'$ on its diagonal; and $\boldsymbol{u}_t$ is an $nmr \times 1$ vector of disturbances with q-th element $\boldsymbol{u}_{t(q)} = \left[u_{11t(q)}, ..., u_{mnt(q)}\right]'$. The above equations define a state-space model that characterizes the joint likelihood of the unobserved state variable $[\boldsymbol{F}_t', \boldsymbol{g}_t']'$ and other latent parameters. Because this can become a high-dimensional system with many parameters, we follow Bernanke et al. (2005) in part and adopt Bayesian inference as our preferred likelihood-based approach.

The first reason for addressing estimation using Bayesian inference is the vast availability of suitable prior distributions that provide automatic regularization to the joint likelihood, especially when considering estimation of extreme quantiles where not many observations are available. Following Feldkircher et al. (2022), Korobilis (2022) and others, we specify the Horseshoe prior for sparse signals of Carvalho et al. (2010) for the elements of the matrices $\boldsymbol{\Lambda}$ and $\boldsymbol{\Gamma}$, as well as the elements of the VAR coefficients $\boldsymbol{\Phi}$. For a generic b-dimensional vector of parameters $\boldsymbol{\theta}$ (where $\boldsymbol{\theta}$ represents column vectors obtained from vectorizing the parameter matrices $\boldsymbol{\Lambda}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}$, respectively) the Horseshoe prior takes the form

$$\boldsymbol{\theta}|\xi, \boldsymbol{\eta} \sim \prod_{i=1}^{b} N\left(0, \xi\eta_i\right), \tag{7}$$

$$\xi \sim C^+\left(0, 1\right), \tag{8}$$

$$\eta_i \sim C^+\left(0, 1\right). \tag{9}$$

Shrinkage estimators regularize an equivalent unrestricted estimator by means of a scalar factor κ that determines how much the unrestricted estimator is shrunk towards zero.[6] In the case of the prior above, κ is $Beta(0.5, 0.5)$ distributed, i.e. has a horseshoe shape. This shape means that with smaller parameter spaces the posterior under a horseshoe prior will tend to be unrestricted, but as the parameter space increases relative to the number of observations, an increasingly larger amount of shrinkage towards zero is favored. These properties make it ideal for our large dimensional parameter space. Both theoretically and empirically the horseshoe performs well and it is a default automatic choice for many researchers; see the detailed review of this literature in Korobilis and Shimizu (2022). For the scalar asymmetric Laplace scale coefficients, $\sigma_{ij(q)}$, we specify conjugate inverse gamma priors. Finally, the VAR covariance matrix $\boldsymbol{\Omega}$ has an inverse Wishart prior.

Other than regularized estimation via prior distributions, the second reason for choosing a Bayesian approach to inference is computational convenience. The state-space

---

[6]In this case, the unrestricted estimator results from placing a normal prior on $\boldsymbol{\theta}$ with infinite variance.

model of equations (5)-(6) is linear but non-Gaussian because $\boldsymbol{u}_t \sim \prod_{ijq} AL(0, \sigma_{ij(q)}, q)$. Additionally, the presence of the quantile common component $\boldsymbol{\Lambda F}_t$, which is a product of two high-dimensional unobserved quantities, complicates state-space estimation further. However, estimation via the Gibbs sampler simplifies inference, because conditional distributions in the QFAVAR have a simple form. As discussed, the asymmetric Laplace distribution can be written as a mixture of Gaussian distributions, converting the linear state-space model into conditionally normal form. The parameters $\boldsymbol{\Lambda}$ and $\boldsymbol{F}_t$ can be sampled one at a time, conditional on one other.[7] In practice, our estimation strategy combines established ergodic samplers, and can be outlined in the following steps:

1. Sample $[\boldsymbol{F}_t', \boldsymbol{g}_t']'$ conditional on values of all other system matrices from the state-space model of equations (5)-(6). This step can be implemented using numerous approaches proposed in the literature, most notably the simulation smoother of Carter and Kohn (1994).

2. Sample $\lambda_{ij(q)}, \gamma_{ij(q)}, \sigma_{ij(q)}$ (and c_{ij}, if an intercept is present) conditional on $f_{t(q)}^i$, for each i, j, q, using equation (2). This is simply a univariate quantile regression, and simple conditional posteriors are provided in Khare and Hobert (2012).

3. Sample $\boldsymbol{\Phi}_1, ..., \boldsymbol{\Phi}_p, \boldsymbol{\Omega}$ (and $\boldsymbol{v}$, if an intercept is present), conditional on all quantile factors $\boldsymbol{F}_t$ using equation (4). This is a simple Bayesian VAR, and posterior conditionals are also quite standard, see Koop and Korobilis (2010).

Of course, in steps 2 and 3 one needs to account for the use of the Horseshoe hierarchical prior, but this is also trivial to incorporate using the hierarchical formulation of this prior proposed in Makalic and Schmidt (2016). The outcome is a Gibbs sampler that is not much different to the sampler in Bernanke et al. (2005) for the FAVAR model and, thus, has a simple and user-friendly structure. All our results using the Gibbs sampler are based on 1,000,000 iterations after discarding an initial chain of 100,000 draws. Out of the one million iterations we store every 100th draw, as consecutive draws tend to be heavily autocorrelated.[8] Therefore, all posterior inference (posterior means, variances etc) is built on a remainder of 10,000 samples from the posterior distribution of parameters and factors.

Finally, due to the fact that the dimension of the latent states $\boldsymbol{F}_t$ (quantile factors) can be quite large, Gibbs sampling estimation tends to become computationally cumbersome in certain applications of interest to policy-makers (for example, forecasting

---

[7]While this approach induces high correlation of consecutive Monte Carlo samples of these parameters, this issue can be easily alleviated by storing only every n-th posterior sample, for an appropriate choice of n.

[8]This is not surprising for factor models, where loadings are sampled conditional on the factors and vice-versa.

macroeconomic risks). In order to address this issue, we also propose an approximate two-step algorithm that replaces the factors with fast variational Bayes estimates from the Korobilis and Schröder (2022) probabilistic quantile factor model. This use of machine learning tools avoids the need for sampling using more demanding state-space methods. Conditional on these plug-in estimates of the quantile factors, we can estimate all other parameters ($\mathbf{\Lambda}, \mathbf{\Phi}$ etc) with a variety of methods depending on one's needs. For example, one could use MCMC as in steps 2) and 3) above, or variational Bayes, or even ordinary least squares. In our implementation, for the sake of consistency, we also use variational Bayes to obtain estimates of the parameter matrices. This two-step approach ignores the dynamics of the factors and the dependence on the global variables $\boldsymbol{g}_t$ in the QFAVAR likelihood when producing estimates of $\boldsymbol{F}_t$. However, it is justified on the basis of computational simplicity, and our empirical results suggest that there is no significant information loss from this two-step procedure relative to the one-step Gibbs sampler.[9] A detailed overview and derivation of our proposed two-step procedure is provided in the online supplement.

2.3   VAR inference in the QFAVAR

The QFAVAR implies a joint VAR for the r quantiles of country-level macroeconomic variables $\boldsymbol{y}_t$ and the mean (expectation) of the global variables $\boldsymbol{g}_t$. Here we follow Bernanke et al. (2005), and Stock and Watson (2005) and we assume that the idiosyncratic disturbances $\boldsymbol{u}_t$ in equation (5) are not relevant for structural inference. Indeed, as these disturbances in our model are cross-sectionally uncorrelated (they have a diagonal covariance matrix), they can be treated as nuisance shocks (e.g. due to measurement error or country-level effects). Therefore, $\boldsymbol{\varepsilon}_t$ are the true structural shocks in the system. Given this convention, one can explicitly reduce the two-equation QFAVAR into a one-equation VARMA model on $[\boldsymbol{y}_t', \boldsymbol{g}_t']'$ and then show that the moving average part vanishes asymptotically (thus, simplifying into a VAR model). However, for simplicity we follow Bernanke et al. (2005) and Stock and Watson (2005) and pursue a two-step alternative: first, we implement all VAR exercises of interest (conditional and unconditional forecasting, impulse response analysis, historical decompositions, etc.)

---

[9] Such two-step procedures are very popular in regular FAVARs. For example, both Bernanke et al. (2005) and Stock and Watson (2005) suggest using principal components in a first step and then estimating all system parameters using least squares. This two-step procedure provides a popular and asymptotically consistent estimator (Stock and Watson, 2016). Proving a similar asymptotic result for our simple two-step variational Bayes procedure would be useful, but is beyond the scope of this paper. However, the probabilistic quantile factor analysis estimator in Korobilis and Schröder (2022) is numerically quite similar to the quantile factor analysis estimator in Chen et al. (2021), which in turn is a generalization of the principal components estimator to a quantile setting.

using the VAR in equation (6), and then we project these quantities into the original variables $\boldsymbol{y}_t$ using the projection matrices $\boldsymbol{\Lambda}$ and $\boldsymbol{\Gamma}$ in equation (5). For example, using standard formulas (Lütkepohl, 2005) the VAR in equation (6) can provide h-step-ahead forecasts of $\left[\boldsymbol{F}'_{t+h|t}, \boldsymbol{g}'_{t+h|t}\right]'$, for some $h > 0$, and these can be projected into forecasts for $\boldsymbol{y}_t$ using the formula $\boldsymbol{y}_{t+h|t} = [\boldsymbol{\Lambda}, \boldsymbol{\Gamma}] \left[\boldsymbol{F}'_{t+h|t}, \boldsymbol{g}'_{t+h|t}\right]'$. Similar arguments can be made about impulse response functions and other quantities of interest.

3   Empirical evaluation of the QFAVAR

3.1   Euro-area macroeconomic indicators and global data

We use five macroeconomic variables from nine euro-area countries observed over the sample 1996M1-2022M12. The countries, series, and sources are shown in the top panel of Table 1. Three series come from the Statistical Data Warehouse (SDW) of the European Central Bank. Industrial production is from the OECD data website, and the Economic Sentiment Index is a composite index maintained by the DG ECFIN (and accessible from the website of Eurostat, the main statistical agency in Europe). All series were accessed in March 2023. Consumer prices are not seasonally adjusted, so we convert these to year-on-year growth rates using the transformation $100(\log P_{t+12} - \log P_t)$; we use the same annual growth tranformation for the ESI in order to create a smoother series. IP (which is seasonally adjusted) is converted to month-on-month growth rates using the tranformation $100(\log P_{t+1} - \log P_t)$. Finally, LTIR and CLIFS are left observed in levels. We neither undertake any additional seasonal adjustment, nor do we alter the series in any other way (e.g., outlier adjustment).

Table 1: Euro-area and global indicators

Countries	Macroeconomic Indicators	Acronym	Source
Austria (AT)	Harmonized Index of Consumer Prices – Overall index	HICP	SDW[1]
Belgium (BE)	Industrial Production Index – Total index	IP	OECD[2]
Germany (DE)	Long Term Interest Rate	LTIR	SDW[1]
Spain (ES)	Economic Sentiment Index	ESI	Eurostat[3]
Finland (FI)	Country-Level Index of Financial Stress	CLIFS	SDW[1]
France (FR)			
Italy (IT)			
Netherlands (NL)			
Portugal (PT)			
	Global indicators	Acronym	Source
	Global inflation (OECD countries average)	GINF	OECD[4]
	Global Supply Chain Pressure Index	GSCPI	NY Fed[5]
	Chicago Fed National Financial Conditions Index	FCI	Chicago Fed[6]
	Global Economic Policy Uncertainty	GEPU	EPU webpage[7]

The global series include Global Inflation (GINF), the Global Supply Chain Pressure Index (GSCPI), the Financial Conditions Index (FCI), and Global Economic Policy Uncertainty (GEPU). The FCI is the national index for the United States, produced by the Chicago Federal Reserve Bank; strictly speaking it is not a global average of multiple countries (as is the case with the other three variables). However, this FCI is a factor of over 100 financial time series, covering developments in stock, foreign exchange, bond, and other key U.S. markets at the forefront of global financial activity. Therefore, for euro-area countries in particular, the U.S. FCI can be a proxy for global financial shocks. Global inflation is the month-on-month growth rate of the price index series. All other variables are in levels. Sources and detailed definitions are in the bottom panel of Table 1.

In the benchmark QFAVAR specification we extract one factor per group of country-level macroeconomic indicators in Table 1, we focuses on quantiles $q = 0.1, 0.5, 0.9$, and we set a maximum of $p = 6$ lags in the state equation of the model.[10] Note, however, that one can easily obtain special and restricted cases of our model. When $\boldsymbol{\Gamma} = 0$ all global shocks are only transmitted to country-level data via the quantile factors, making country-level responses symmetric (because each macroeconomic indicator loads on a single factor per quantile with weight $\lambda_{ij(q)}$). If the global shocks are also restricted to be absent from the state equation, the QFAVAR collapses to a quantile dynamic factor model for variables $\boldsymbol{y}_t$. Because the Bayesian QFAVAR has a conditionally normal form and the mean of a

---

[10] Recall that the horseshoe prior can flexibly restrict certain lagged parameters in a flexible way.

normal distribution is identical to its median, the QFAVAR can collapse into a FAVAR by considering estimation only for $q = 0.5$. In practice, whenever we estimate the FAVAR as a benchmark model we do so using algorithms similar to the QFAVAR. When forecasting in subsubsection 3.3.1 we estimate the QFAVAR and the FAVAR with variational Bayes, and when doing structural analysis in subsection 3.4 we estimate both models with MCMC. Finally, as explained in subsection 2.2, the QFAVAR priors are automatic and adjust to varying needs for shrinkage (depending on number of countries, indicators, lags, quantiles, etc.) and we use the same priors whenever estimating the FAVAR as a benchmark for comparison.

3.2   In-sample quantile factor estimates and model fit

It comes as no surprise that a large part of macroeconomic dynamics in the euroarea is characterized by common drivers among countries. Nevertheless, our aim is to showcase that considerable dynamics in the quantiles of macroeconomic variables remain hidden when considering in-mean factors only. In addition, these dynamics are not only interesting by themselves, but also imply strong and economically meaningful heterogeneity in the transmission of economic shocks. Figures 1 and 2 depict the estimated QFAVAR factors, using the one-step MCMC and two-step VB estimators, respectively. In both cases, the quantile factor estimates plotted are defined as posterior means of the 10th, 50th, and 90th percentile factors. For the sake of clarity the mean (FAVAR) factors are not plotted in this figure, as these are fairly indistinguishable from the median (50th percentile) factors. Therefore, median factors are a reference point when comparing those to tail factors.

A first inspection of the two figures reveals that there are no marked differences between the MCMC and VB estimates of the quantile factors, taking into account the fact that the two estimators are based on different modeling assumptions.[11] The VB estimates of the 10th percentile factors sometimes cross with the median or even the 90th percentile (in the case of IP) factors, but such crossing or overlapping of quantile factor estimates is not present when using MCMC. Because latent factors combine distributional information from multiple series, some crossing might be present when using algorithms that assume complete independence between quantiles (for example Chen et al., 2021; Korobilis and Schröder, 2022). Therefore, the one-step estimation performs better, possibly because it takes the correlation of the factors across quantiles into account. At the same time, the numerical results suggest that the use of one-step MCMC is essential in structural exercises where interpretation of the factors is important. When interest is in forecasting, the factors only

---

[11]One-step MCMC estimation recovers the "true" dynamic factors implied by the QFAVAR specification that are also contemporaneously and dynamically correlated with the globals, g_t. Two-step VB only estimates static factors as in Korobilis and Schröder (2022) with no reference to g_t.

serve the role of being reduced-form manifestations of the original data (even if they are not interpretable), in which case VB estimates can be perfectly acceptable; see discussion of this issue in next subsection.
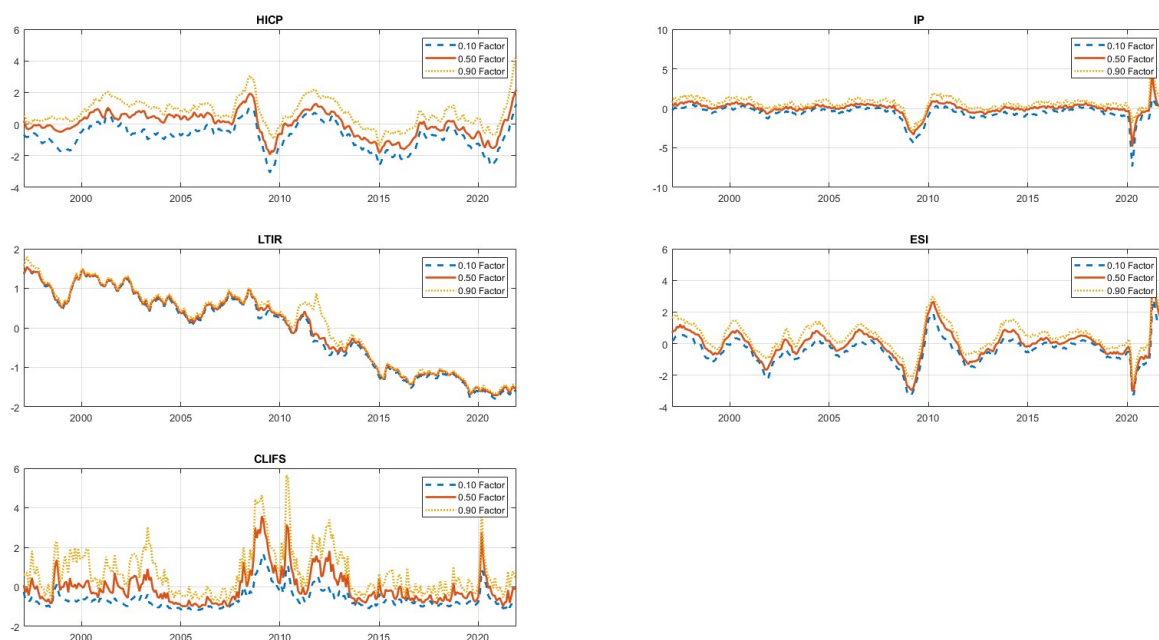

Figure 1: Markov chain Monte Carlo estimates (posterior means) of the five euroarea quantile factors (10th, 50th, and 90th percentiles).
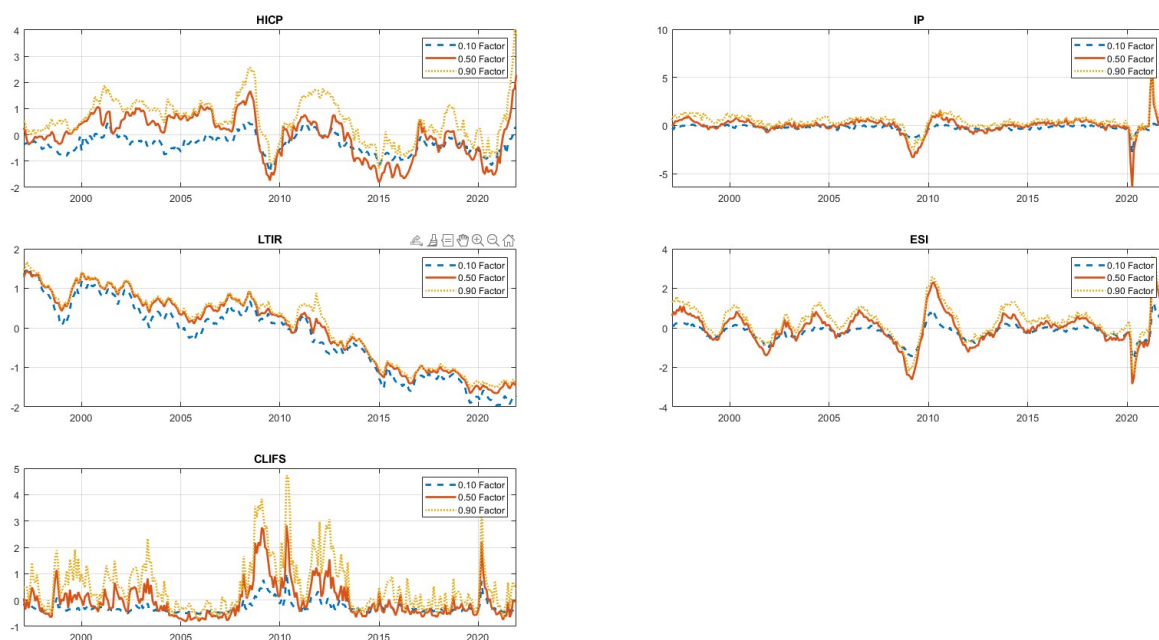

Figure 2: Variational Bayes estimates (posterior means) of the five euro-area quantile factors (10th, 50th, and 90th percentiles).

An interesting question is whether the quantile factors capture more information in the data than the mean factors extracted with the simpler FAVAR method. To explore this question in depth, we provide an in-sample exercise in this section and a thorough out-of-sample exercise in the next section. To formally investigate the additional informational content of the quantile factors we follow Despois and Doz (forthcoming) and computing factor commonalities, defined as the R^2 of the multivariate regression of y_{ij} onto the estimated factors

$$R^2_{y_{ij}, \hat{F}} = \frac{||\hat{F}(\hat{F}'\hat{F})^{-1}\hat{F}'y_{ij}||^2}{||y_{ij}||^2}. \tag{10}$$

In the above, $\hat{F}$ is the posterior mean estimate of the mean factors from the FAVAR, or the posterior mean of the rth quantile factor from the QFAVAR, where both models are estimated using MCMC. Table 2 presents the results for the mean factors in the first column. In the second and third column we then add the tail factors and the full set of QFAVAR factors to the mean factor. For inflation, adding the tail factors consistently increases the R^2, where we observe the largest gains for Finland, the Netherlands, and Portugal. Including the median factors increases the commonality further; performance improves the most for Italy and the least for Portugal. For industrial production, the general pattern is preserved. The R^2 for Belgium, Finland, and the Netherlands increases the most upon including the tail factors. Further small gains emerge from also including the median factors; however, the gains are overall less sizeable compared to inflation.

Table 2: Commonality of factor estimates

	F	$F + F_{0.1} + F_{0.9}$	$F + F_{0.1} + F_{0.5} + F_{0.9}$
HICP.AT	0.809	0.831	0.868
HICP.BE	0.853	0.860	0.881
HICP.DE	0.839	0.851	0.878
HICP.ES	0.931	0.942	0.961
HICP.FI	0.688	0.816	0.821
HICP.FR	0.939	0.945	0.952
HICP.IT	0.908	0.908	0.948
HICP.NL	0.635	0.776	0.804
HICP.PT	0.296	0.456	0.456
IP.AU	0.906	0.916	0.916
IP.BE	0.764	0.809	0.822
IP.DE	0.912	0.916	0.926
IP.ES	0.951	0.954	0.955
IP.FI	0.694	0.828	0.868
IP.FR	0.967	0.971	0.972
IP.IT	0.957	0.962	0.962
IP.NL	0.655	0.751	0.763
IP.PT	0.780	0.788	0.790

Notes: This table contains the factor commonality estimates for inflation and industrial production for the set of the nine countries. The first columns shows the results for the mean FAVAR factor; the second column shows the results for the mean FAVAR and both QFAVAR tail factors; the final column shows the results for the mean FAVAR and all QFAVAR factors.

3.3 Out-of-sample evaluation of tail risks

3.3.1 Is the QFAVAR better than a FAVAR in capturing tail risks?

In this subsection we shed more light onto the empirical fit of the QFAVAR, compared to the more established FAVAR model. Although the comparison of in-sample fit of the factors in the previous subsection (via R^2 statistics) is an informative exercise, we draw a more complete picture by comparing the out-of-sample performance of the QFAVAR versus the FAVAR. Given then importance of accurate projections and forecasts in policy work, this exercise is particularly insightful. We set up a straightforward recursive pseudo-out-of-sample (poos) exercise where we begin the estimation with 50% of the total sample, forecast one to 24 months ahead, add one more observation at the end of the sample, estimate the models, and forecast again up to 24 horizons ahead. We compute forecasts iteratively using the VAR in state equation (6), and subsequently we project them into forecasts for the conditional quantiles of the original variables $\boldsymbol{y}_t$ using the loadings matrices $\boldsymbol{\Lambda}$ and $\boldsymbol{\Gamma}$ in equation (5).

We follow Manzan (2015) and evaluate the tail forecasting performance using the

following quantile score function

$$QS_{t|t-h}(q)^M = [y_{ijt} - Q_q(y_{ijt})] \left[\mathbb{I}\left(y_{ijt} \leq Q_q(y_{ijt})\right) - q \right]. \tag{11}$$

This score function is evaluated for the two competing models $M = \{QFAVAR, FAVAR\}$; the two extreme quantile levels $q = 0.1, 0.9$; four forecast horizons $h = 1, 6, 12, 24$; two variables of interest $i = HICP, IP$; and all nine euro-area countries j. The QS function is a piecewise linear asymmetric loss function; therefore lower values signify better performance. Because we are interested in comparing the performance of only two different classes of models, we again follow Manzan (2015) and calculate the t-statistic testing the equality of the average of the models' QS functions over the whole out-of-sample period. This is defined as the spread between the QS values of the QFAVAR versus the QS values of the FAVAR. Values of this statistic for different variables and forecast horizons are shown in Table 3. The visibly higher proportion of negative values implies that, overall, the QFAVAR experiences less forecast performance loss in both the left and right tails of the distribution of HICP and IP. For HICP, performance gains are statistically significant at the 5% level for $h = 1$ for the bottom tail and for $h = 1$ and $h = 6$ for the top tail of the distribution. For a subset of countries, additional significant gains also emerge for longer forecast horizons. For industrial production the picture is generally comparable with the exception of some countries; however, significant performance gains also arise for the bottom tail of the distribution for the longer horizons $h = 12$ and $h = 24$. Generally, the QFAVAR provides clear performance gains over the symmetric, Gaussian disturbances of the FAVAR, highlighting the importance of flexible asymmetric modeling of conditional quantiles within a multivariate setting.

Table 3: T-statistics of equal quantile predictive accuracy

| | $t-stat_{10}$ | | | | $t-stat_{90}$ | | | |
	$h=1$	$h=6$	$h=12$	$h=24$	$h=1$	$h=6$	$h=12$	$h=24$
HICP.AT	**-3.04**	-0.57	**-2.34**	-0.13	**-5.98**	**-3.71**	**-2.52**	0.12
HICP.BE	**-2.94**	-0.61	-0.02	-0.81	**-5.74**	**-3.01**	-1.00	0.12
HICP.DE	**-2.40**	-1.48	**-2.39**	-0.25	**-5.91**	**-3.75**	-1.63	-0.10
HICP.ES	**-3.30**	-1.53	-1.68	-0.01	**-5.14**	**-2.10**	-0.63	0.73
HICP.FI	**-2.15**	-1.54	-1.57	**-2.42**	**-4.75**	**-3.29**	**-2.95**	-0.71
HICP.FR	**-3.37**	-1.20	-0.11	1.25	**-6.87**	**-2.58**	-1.44	-0.25
HICP.IT	**-3.65**	-1.68	-0.67	0.35	**-5.62**	**-3.23**	-1.86	-0.40
HICP.NL	**-2.68**	**-2.56**	-1.63	1.31	**-2.44**	**-2.18**	-1.94	0.11
HICP.PT	**-4.06**	0.02	-1.07	-1.29	**-4.37**	**-2.34**	-0.35	0.76
IP.AT	**-2.11**	-1.98	-1.92	**-4.38**	**-7.24**	-1.57	-0.63	-0.08
IP.BE	**-2.19**	**-2.77**	**-2.04**	**-3.57**	**-3.67**	**-2.36**	-0.73	-1.07
IP.DE	-1.62	-1.31	**-2.40**	**-3.77**	-1.88	-1.28	-0.04	**-2.44**
IP.ES	**-2.06**	-1.90	**-3.67**	**-3.34**	**-3.66**	-0.77	1.80	1.08
IP.FI	**-2.79**	**-2.48**	**-3.82**	**-5.70**	**-2.90**	-1.71	1.35	-1.87
IP.FR	-0.93	-1.34	-1.90	**-2.78**	**-4.59**	-0.05	1.98	0.55
IP.IT	-0.92	-1.82	**-2.31**	**-2.74**	-1.27	-0.18	1.25	-0.47
IP.NL	**-2.24**	**-3.02**	**-2.04**	**-5.22**	**-2.54**	-0.72	1.04	0.22
IP.PT	-0.74	-1.11	**-2.25**	**-2.02**	**-5.45**	-1.88	-1.69	-1.88

Notes. Entries in this table are t-statistic values for the null hypothesis of equal accuracy of the quantile forecasts from the QFAVAR relative to the FAVAR benchmark. Columns 2-5 show the statistics for 10th percentile forecasts over forecast horizons $h = 1, 6, 12, 24$ months, and columns 6-9 show the statistics for the 90th percentile forecasts over the same forecast horizons. Because all values of the statistic are negative, this shows that the QFAVAR experiences less forecast loss than the benchmark FAVAR. Values lower than the critical value of -2 indicate that QFAVAR tail forecasts are significantly better than the FAVAR forecasts at the 5% level.

Admittedly, evaluating only a certain area (left or right tail) of the distribution of inflation and output is of paramount importance to policy-makers who are interested in assessing worst-case scenarios for these two variables of interest. However, many times the full forecast distribution of the variables of interest is required, most notably when preparing fan charts similar to the ones maintained by the Bank of England (Britton et al., 1998). In such cases, one has to construct the full forecast distribution by interpolating forecasts of a wide range of quantiles. This procedure would not be accurate with the benchmark specification of the QFAVAR, where for parsimony we focus only on the 10th, 50th, and 90th percentiles. For that reason, the model can be trivially extended to consider joint estimation of a larger number of quantiles. Following Chen et al. (2021) we re-estimate the QFAVAR model this time for $q = 0.05, 0.1, 0.25, 0.5, 0.75, 0.9,$ and 0.95. This implies $R = 7$ and the dimension of the state now increases to $Rm + k = 7 \times 5 + 4 = 39$ elements for each time period t. Nevertheless, computational load doesn't increase noticeably when

adopting the two-step variational Bayes algorithm.[12]

Figures 3 and 4 illustrate the one-step-ahead predictive distributions of country-level inflation and industrial production, constructed from the QFAVAR (blue solid line) versus the FAVAR (red dashed line) estimated using the full sample up to December 2021. The date is chosen as it marks the beginning of an era of rapidly increasing food and energy prices, as well as shortages of raw materials and other frictions in production and global trade resulting from the Covid-19 pandemic. The yellow vertical line indicates the realized value in January 2022. We construct these densities by first predicting the one-step-ahead value of each variable for each of the $R = 7$ quantiles and then fitting a Gaussian kernel smoother to these quantiles.[13] Based on these two figures it is clear that there are significant differences between the two modeling approaches. Compared to the symmetric FAVAR distributions, the QFAVAR features asymmetric and, in many cases, bimodal distributions. In addition, the QFAVAR distributions place more probability mass on the right tail of inflation and are left-skewed in case of industrial production. The FAVAR distributions look, by and large, like symmetric weighted averages of the more complex QFAVAR distributions, hiding the true extent of uncertainty around forecasts produced in December 2021. The last few months of 2021 signalled a period of inflationary pressures, and policy-makers faced huge uncertainty deciding whether these pressures would prove transitory or more pervasive. With the advent of the Russian invasion of Ukraine in early 2022, it was clear that food and energy inflation would persist and prospects for positive output growth could be dim due to subsequent interest rate increases and the build-up of financial risk. Given ex-post knowledge of these extreme realized macro risks, it is not surprising that forecast distributions of inflation and industrial production constructed from the QFAVAR look more plausible than the symmetric FAVAR distributions.

3.3.2 Are global variables relevant for forecasting tail risk?

With multiple global shocks affecting the euro-area, it is unclear how these transmit to country-level risks of inflation and industrial production (Panetta, 2023). Understanding the underlying transmission channels of such global predictors is an important issue for policy-makers. A good starting point towards this direction is the related debate on whether financial predictors affect the left quantile of GDP. Despite the fact that (Adrian et al., 2019) find that financial conditions help forecast the left tail of U.S. GDP, Plagborg-Møller

---

[12]This is not true for the one-step MCMC algorithm, as increasing the dimension of the state variable would increase the computational demands of the filtering problem substantially. Additionally, the one-step estimator requires writing the VAR(p) evolution in the state equation into a VAR(1) form, which requires repeatedly sampling a state vector that is even larger; that is, a vector with $p(Rm + k) = 234$ elements for $p = 6$.

[13]This is done using the MATLAB R2022b function `ksdensity` with default settings.
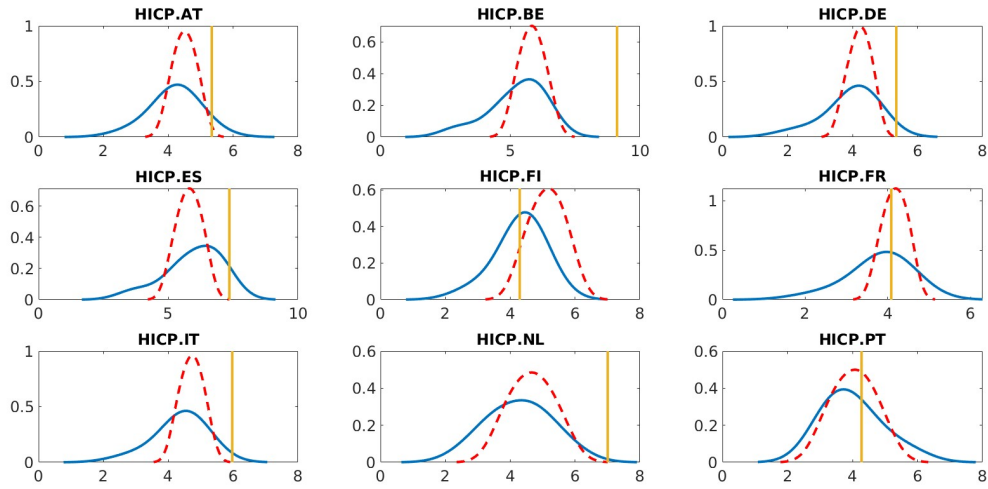
Figure 3: One-step-ahead forecast distributions of inflation from QFAVAR (blue line) and the FAVAR (red line), both estimated using variational Bayes using a finer grid of quantiles ($q = 0.05, 0.1, 0.25, 0.5, 0.75, 0.9,$ and 0.95) compared to the benchmark specifications. The yellow vertical line indicates the realized value in January 2022.
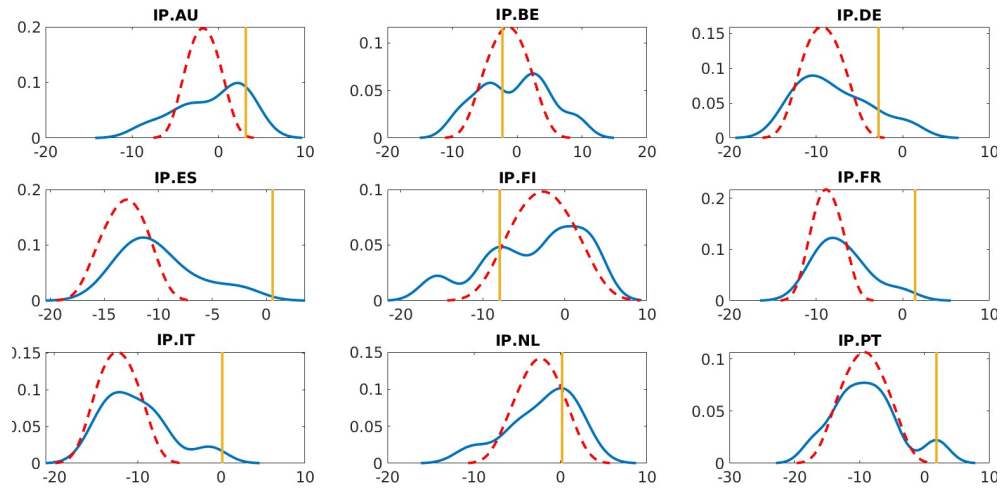

Figure 4: One-step-ahead forecast distributions of industrial production from QFAVAR (blue line) and the FAVAR (red line), both estimated using variational Bayes using a finer grid of quantiles ($q = 0.05, 0.1, 0.25, 0.5, 0.75, 0.9,$ and 0.95) compared to the benchmark specifications. The yellow vertical line indicates the realized value in January 2022.

et al. (2020) argue that no single financial predictor is persistently important for forecasting recessions. Their argument is based on the finding that the relationship between GDP and financial conditions is inherently unstable over time. This finding suggests that policy makers should not mechanically include financial conditions as predictors of tail events.

A major limitation of Adrian et al. (2019) and Plagborg-Møller et al. (2020) is that they rely on univariate quantile regressions; therefore, predictors enter in the form of right-hand-side exogenous variables for each quantile level. This is not the case with the proposed QFAVAR specification, where all estimated conditional quantiles of euro-area macroeconomic variables and the global predictors are endogenous and can interact both contemporaneously and dynamically. In order to explore the effects of global predictors on country-level macro risks, we estimate and forecast with four models: i) our benchmark QFAVAR, ii) a QDFM, which is the QFAVAR with quantile factors only and no global variables $\boldsymbol{g}_t$, iii) a quantile autoregressive (QAR) model, and iv) a quantile autoregressive model with global variables $\boldsymbol{g}_t$ as predictors (QAR-X). To eliminate the impact of the shrinkage priors in driving forecast performance, we estimate all four models with the parsimonious choice of $p = 1$ lag and an intercept, deviating from the remainder of the pager where we use $p = 6$ and an intercept.[14] Finally, also for the sake of consistency we estimate the two univariate quantile models using the same variational Bayes algorithm, despite the fact that MCMC is not cumbersome in this case (see for example the fast algorithm in Korobilis et al., 2021).

Figure 5 plots the cumulative quantile score values for inflation and industrial production, respectively, computed using equation (11) in the previous subsection. Compared to average quantile scores, cumulative sums reflect the evolution of forecast performance of different models over the out-of-sample period. Given that the $QS_{t|t-h}(q)^M$ statistic is a loss function, the best model is the one with the lower cumulative values. The figure has four rows of nine panels, where the first two rows correspond to the QS performance at the 10th and 90th percentile of inflation at horizon $h = 1$ and the columns correspond to the individual euro-area countries. The third and fourth row present the corresponding results for industrial production. Interesting patterns emerge from this graph. At first glance, it is obvious that for the short forecast horizon $h = 1$, multivariate quantile models are superior to univariate quantile models in producing accurate left- and right-tail forecasts of inflation. In particular, the QFAVAR dominates all other models, and in many cases the largest reductions in loss are early in the out-of-sample period, that is, during the 2011 eurozone debt crisis.[15] This is not a

---

[14]In many macroeconomic forecasting problems, longer lag lengths fit very well in-sample, but when it comes to forecasting out-of-sample the choices $p = 1$ or $p = 2$ are hard to beat.

[15]See for example, the 10th percentile QS of HICP.IT, and the 10th percentile QS of IP.ES.

surprising result, as it is well-known that the debt crisis created strong comovements and global economic unrest. The multivariate QFAVAR is able to fit this period better than the univariate competitors and the QDFM.
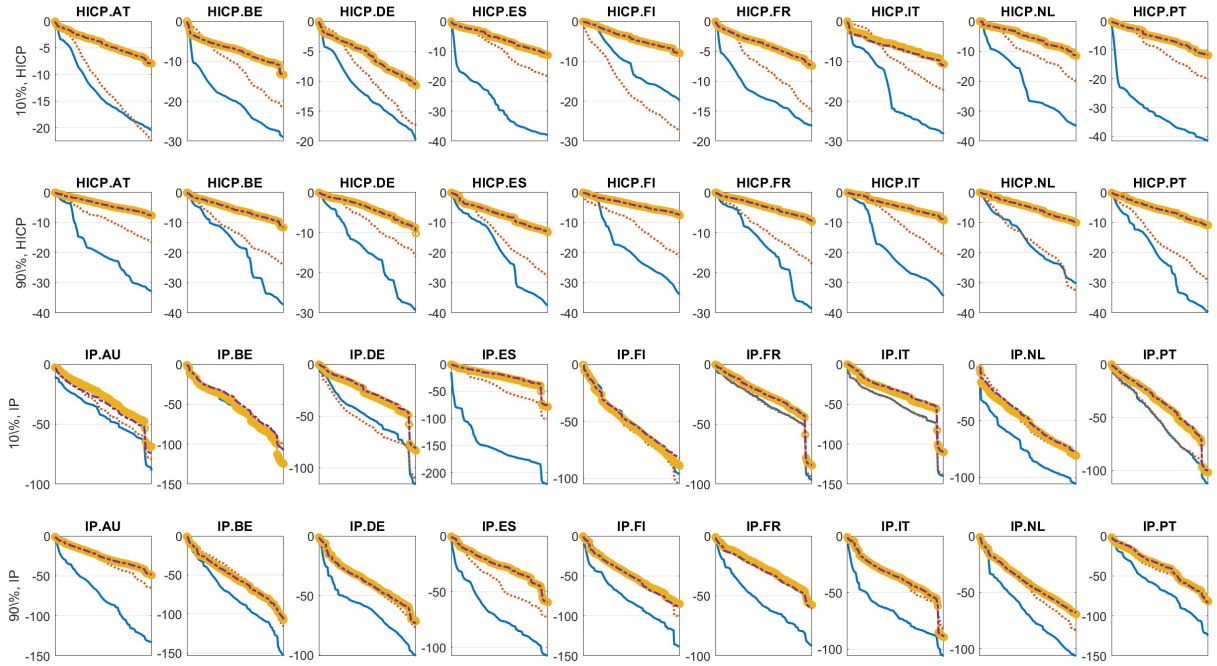

Figure 5: Cumulative quantile score (QS) loss for forecast horizon $h = 1$. We compare four models: QAR (yellow circled line), QAR-X (purple dashed line), QDFM (red dotted line), and QFAVAR (blue solid line). The out-of-sample evaluation period shown on the x-axis is 2011Jan to 2022Dec-h. The first two rows show quantile scores for the 10th and 90th percentiles of inflation, and the third and fourth rows show quantile scores for the 10th and 90th percentiles of industrial production.

Looking more closely at the results, we can infer several stylized facts regarding short-term quantile forecasting. First, global predictors seem to provide irrelevant information when considering univariate forecasting models, but this is not true when considering multivariate models. This stylized fact shows that the conclusions of Plagborg-Møller et al. (2020) might only apply to univariate quantile models. Second, among the multivariate models, the QFAVAR is clearly superior to the QDFM that does not consider global predictors. Third, inflation comovements in multivariate data seem to be important, because both the QFAVAR and QDFM markedly improve the performance of the univariate models. Finally, for industrial production, the QDFM provides identical forecast performance to the univariate models in most cases. The large improvements observed for the QFAVAR therefore likely come from the inclusion of global predictors as endogenous variables.

Moving to longer-term forecasts, this clear picture is lost as there is no clear winner. The only pattern that emerges is that the quantile AR(1) without global predictors is always the worst-performing model. Overall, at longer horizons both the QFAVAR and QDFM and the quantile AR with exogenous predictors performs very well. Figures supporting these results are available among the additional empirical results in the online supplement.

3.4 Evaluating global spillovers using the QFAVAR

3.4.1 Quantile factor responses

We compute generalized impulse response functions[16] of the FAVAR and QFAVAR factors to changes in global inflation, the global supply chain pressure index, financial conditions, and global economic policy uncertainty. While factors from these models are unobserved, in the case of multi-country analysis the indicator-specific factors we estimate by averaging the same series for all nine countries, can be thought of as proxying aggregate euroarea indicators. Therefore, responses of mean (in the FAVAR) and quantile (in the QFAVAR) factors of HICP, IP, LTIR, ESI and CLIFS to global shocks are particularly interesting to policy-makers interested in characterizing aggregate-level dynamics instead of country-level heterogeneities. The FAVAR and QFAVAR impulse response functions (IRFs) are presented in figures 6 and 7, respectively. Note that quantile IRFs can be interpreted in two ways. On one hand, the factor responses reflect features of the same distribution. Consequently, responses of quantile factors suggest changes in skewness and kurtosis, as well as upside and downside risks. If the quantile responses overlap completely, this is indicative of a level shift in the entire distribution. On the other hand, one can interpret the responses as distinct quantities, through the lens of scenario analysis. In this case, depending on the interpretation of the underlying variables, the 10th (90th) percentile factor can be interpreted as an adverse (favorable) benchmark scenario and, hence, contribute important additional information beyond the median IRFs. As a result, both interpretations provide useful supplements to existing models in the policy-making process.

---

[16]There are numerous ways of turning a reduced-form VAR model into a structural econometric model by means of imposing sensible and plausible identification restrictions. These identification schemes range from recursive and long-run restrictions, to the currently popular methodologies of sign restrictions and identification via instrumental variables. However, the QFAVAR has the particularity that it models different quantiles of the same variables jointly. A global shock will not necessarily have the same effect on all quantiles of a given variable. Even if economic theory, intuition or common sense are available to help choose identification restrictions in VARs, these restrictions will hold on average (that is, in the median) and there are no guarantees that structural relationships remain the same during "exceptional times" (that is, at the tails of a distribution). Therefore, the QFAVAR requires careful consideration of plausible restrictions that would help identify structural shocks. We leave this exercise for future research, and in this paper we focus exclusively on generalized impulse response functions that require minimal assumptions about signs, magnitudes or other important features of shocks.
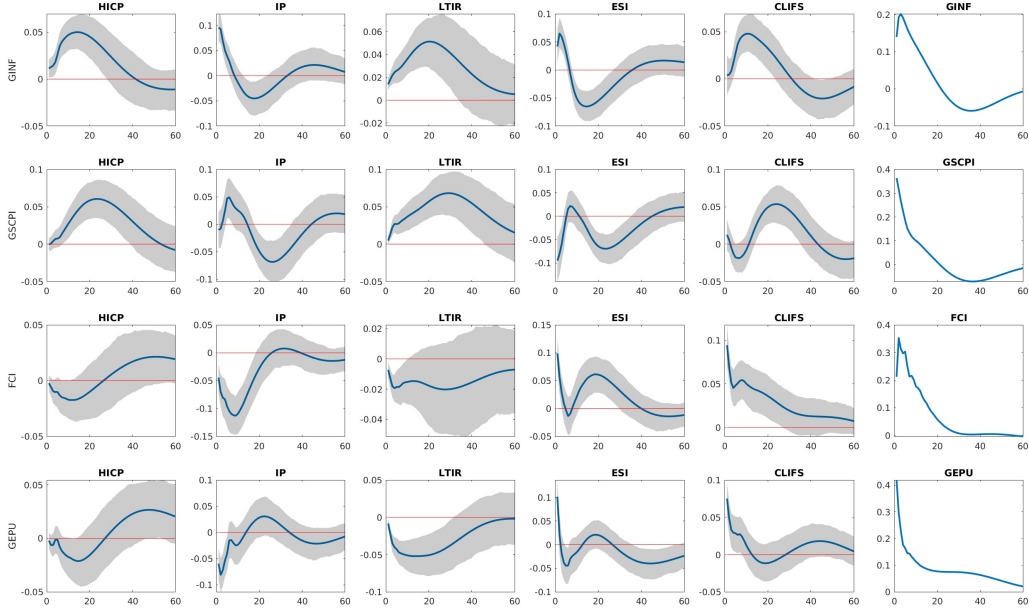
Figure 6: Impulse response functions (posterior median and 68% bands) of mean factors to shocks in the global variables, based on the benchmark FAVAR specification. Each row represents a global shock (GINF, GSCPI, FCI, GEPU); the own-response is shown in the last column. The first five columns show the responses of the five factors extracted from country-level euro-area data.

In the FAVAR responses, shown in Figure 6, an increase in global inflation coincides with a positive impact response for all five factors. Given that we compute generalized IRFs, global inflation might largely be driven by demand-side forces. The profile of the corresponding quantile factor responses is empirically interesting. During the first 20 months following a global inflation shock, the 90th percentile factor of inflation reacts more strongly than the median and the 10th percentile factor. During the first five months, this pattern repeats for industrial production. This fact suggests that the distributions of inflation and industrial production become more leptokurtic as well as positively skewed following the shock. Taking the scenario perspective instead, the 90th percentile factor responses represent a benchmark scenario with larger global spillovers and a stronger response of inflation compared to what is the case in the FAVAR. For the global supply chain pressure indicator, inflation does not respond on impact. During subsequent months, inflation then responds positively, peaking after about two years. In the QFAVAR, however, we see that the 90th percentile inflation factor does respond positively on impact, indicating upside risk to the inflation outlook. In the following months, the median and 10th percentile factor follow up, shifting the entire inflation distribution toward higher realizations. For industrial production, the FAVAR first indicates a contraction on impact. The median factor mirrors this behaviour. The left-tail

23

factor responds negatively and the right-tail factor responds positively on impact, translating into an increase in uncertainty about the industrial production outlook and a more leptokurtic distribution. In subsequent periods, the 90th percentile factor contracts more strongly than the 10th percentile factor, inducing a negative skew in the distribution, accompanied by marked output contraction over the medium turn. Alternatively, one could construct a scenario combining the response of the 90th percentile of inflation and median industrial production factor instead. In this scenario, the shock would lead to an output response similar to what is the case in the FAVAR, but with higher inflation, worsening the policy trade-off. The responses of economic sentiment, financial stress, and especially the long-term interest rate all indicate higher uncertainty about the outlook than under the FAVAR.
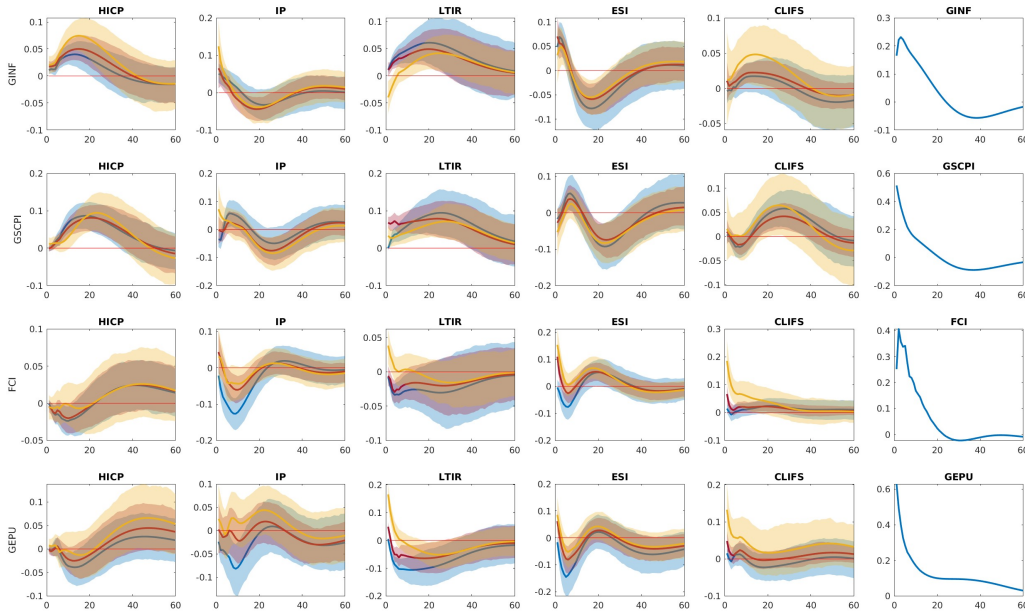

Figure 7: Impulse response functions (posterior median and 68% bands) of 10th, 50th, and 90th percentile factors to global shocks from the benchmark QFAVAR specification. Each row represents a global shock (GINF, GSCPI, FCI, GEPU); the own-response is shown in the last column. The first five columns show the responses of each of the three percentiles of the five factors extracted from country-level euro-area data. Responses of the 10th percentile factors are in blue, reponses of the 50th percentile factors are in red, and responses of the 90th percentile factors are in yellow.

In case of tightening global financial conditions, the left tail and median inflation factors respond similarly to the mean factors during the first months after the tightening. The right tail factor has a slightly more muted response, inducing positive skew. Overall, the distribution shifts toward lower realizations. What is more striking is the response of industrial production. Compared to the mean model, the median and 90th percentile

factor respond less, but the 10th percentile factor responds with a more extreme output contraction. The distribution hence becomes strongly negatively skewed and suggests non-negligible downside risks to the production outlook. Policy-makers might hence pay particular attention to this scenario when monitoring global financial conditions. Combining the 10th percentile inflation and industrial production factors, allows for the construction of a severe scenario in which both prices and output react more negatively than in the FAVAR. The responses of the other variables also show richer dynamics than under the FAVAR. Although the sentiment outlook becomes overall more uncertain, the interest rate and financial stress exhibit considerable positive skew.

Finally, an increase in global economic policy uncertainty is another interesting case. In the FAVAR, inflation and especially industrial production respond negatively during the first 30 and 10 months, respectively. For inflation, this pattern is broadly mirrored by the quantile factors. The 90th and 10th percentile factors respond less and more strongly than the mean factors, respectively. Overall, the inflation distribution thus becomes more leptokurtic. In the case of industrial production, however, the median responds little, the 90th percentile factor responds positively, and the 10th percentile factor responds strongly negatively, suggesting negative skew and pervasive uncertainty about the outlook. Taking the scenario view, one could construct a severe (benign) scenario by combining the 10th (90th) percentile responses of the inflation and industrial production factor. The severe scenario would then correspond to a case with a more pronounced decline in inflation as well as output. Although the ESI shows signs of downside risks to the sentiment outlook, an increase in global economic policy uncertainty correlates with pronounced upside risks for the long-term interest rate and financial stress.

To provide deeper insights into the importance of the four global variables as drivers of the inflation and output responses, figure 8 displays generalized forecast error variance decompositions of the different quantile levels of the inflation and output factors. For the sake of clarity we only show the contribution of the global variables, which is why the contribution shares do not add up to 100%. A few interesting patterns stand out. Overall, changes in global variables explain the most variation in the 10th percentile factors of inflation and IP, and they contribute less to the forecast error of the median and 90th percentile factors. Generally, they are more important for inflation than industrial production, with the global variables accounting for roughly 40% of the forecast error variance of the 10th percentile inflation factor and close to 30% of the 10th percentile industrial production factor. For inflation, global inflation is the most important global variable on impact. This pattern also holds true for the other quantile levels. In the median and longer run, the global supply chain pressure index becomes dominant and contributes by far the largest share. At the 10th percentile, it accounts for almost 20%.

Interestingly, the timing of when the supply chain pressure index becomes dominant differs across quantile levels. At the 10th, 50th, and 90th percentile levels it contributes the largest share after roughly 5, 10, and 15 months, respectively. Global economic policy uncertainty is the third most important global variable, and global financial conditions are the least important global variable driving the inflation quantile factors. This picture is different for industrial production. At the 10% level, the FCI is the most important global variable for most of the forecast horizon, contributing about 10% to the total variation. The other global factors have roughly equal share, with global inflation contributing the least. These findings again show that financial conditions emerge as an important driver of output at risk - once modelled in a multivariate system - adding to the open debate. At the median and 90% level the global supply chain pressure indicator is dominant, followed by global economic policy uncertainty. Together they account for more than half of the total contribution by global variables. The FCI and global inflation contribute relatively equal shares. Overall, these findings point towards non-negligible asymmetries in the inflation and industrial production process, with downside risks being particularly responsive to global developments.
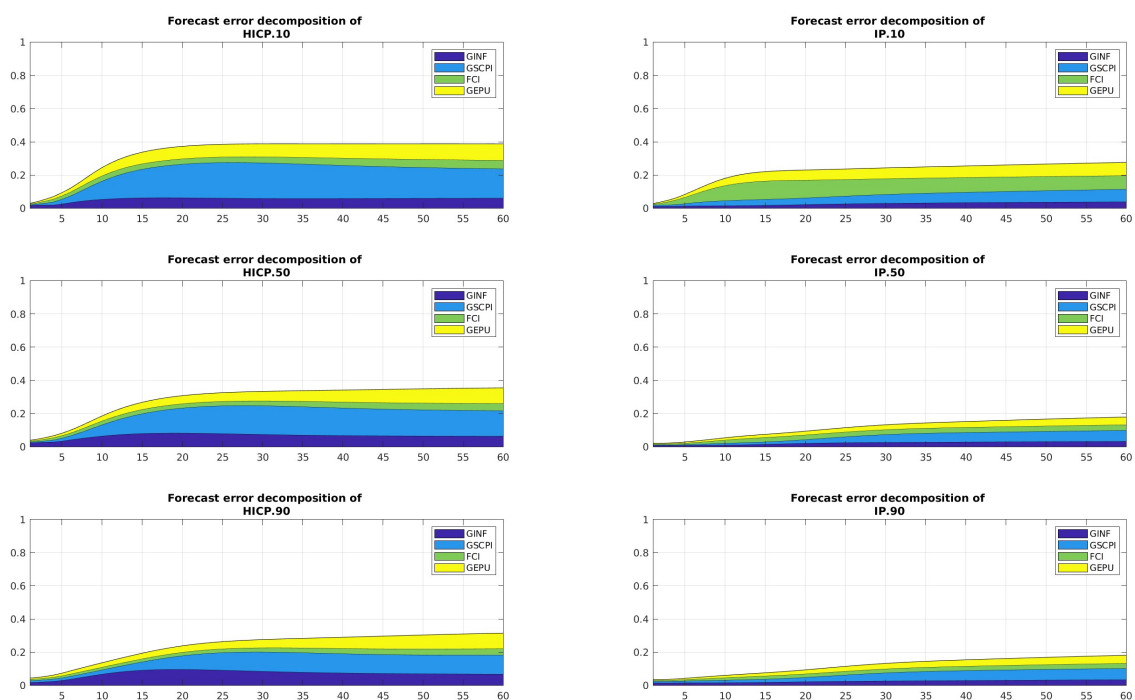

Figure 8: Forecast error variance decompositions of the 10th, 50th, and 90th percentile factors for inflation and industrial production.

3.4.2 Country-level responses to global shocks

Finally, our proposed QFAVAR allows mapping the quantile factor responses back to the panel of individual country-level variables. To do so, we simply project the quantile factor responses back to the measurement equation using the loading matrices $\mathbf{\Lambda}$ and $\mathbf{\Gamma}$. Note that $\mathbf{\Gamma}$ plays a particular role for the quantile dynamics of the country-level variables. With $\mathbf{\Gamma} = \mathbf{0}$, the global factors have no direct impact on the country-level variables and only affect their quantile responses indirectly through the factors. The entire dynamics in the measurement equation are then driven by the common component. Because we allow the individual variables to load on their group-specific factors only, this also implies that the country-level IRFs are just the factor-level IRFs rescaled by the individual loadings. In this case, heterogeneity between the individual EA countries and across quantile levels is limited to the magnitude of the responses. To allow for more flexible degrees of heterogeneity, in this section we focus on the case where $\mathbf{\Gamma}$ is unrestricted, and we report the full set of responses for the $\mathbf{\Gamma} = \mathbf{0}$ case in the appendix as a robustness exercise. In addition, given the large set of countries and variables under study, we focus on interesting subsets of countries and variables in the main body of this paper and refer the reader to the appendix for the full set of responses.

Figures 9, 10, 11, and 12 present the generalized IRFs for HICP and IP following an increase in the four global variables. Starting with global inflation, we show the IRFs for France, which experienced relatively low inflation compared to the European average, Belgium, which had inflation rates close to the European average, and the Netherlands and Italy, which were among the countries with the highest inflation rates in the EA during the recent high-inflation period. A few interesting features stand out. In terms of dynamics, the inflation responses in France and Belgium look roughly similar; however, the inflation response in France is more leptokurtic and negatively skewed. Overall, the inflation outlook in France is more uncertain, but lower inflation realizations are relatively more likely compared to Belgium. The opposite is true for both, the Netherlands and Italy. Here the inflation distribution is characterized by a marked response of the upper tail and hence upside risk to the inflation outlook following an increase in global inflation. Although the median responses are comparable across both the low- and high-inflation countries, what sets them apart is pronounced heterogeneity in the tails of the distributions. Policy makers in the euro-area might hence want to pay particular attention to cross-country heterogeneity following global inflation shocks. For industrial production, the dynamics are rather similar across countries. In Belgium, the production outlook is more uncertain overall on impact, and in the Netherlands upsides risks are slightly more pronounced.
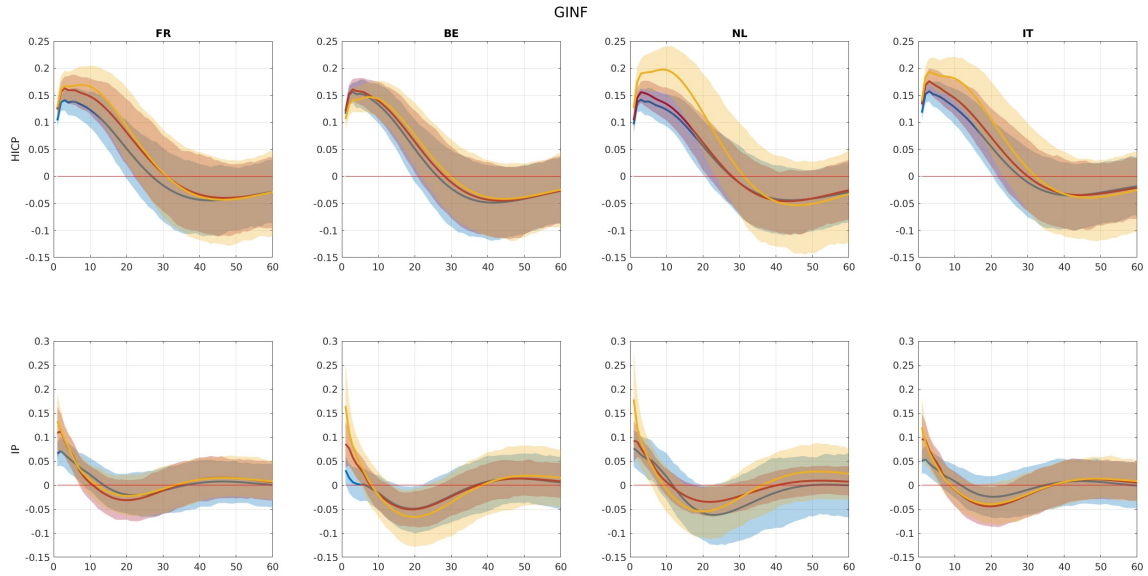
Figure 9: Impulse response functions (median and 68% bands) at the 10th, 50th, and 90th percentiles of country-level inflation (first row) and industrial production (second row). Shock is to the global inflation (GINF). Responses at the 10th percentiles are in blue, responses at the 50th percentiles are in red, and responses at the 90th percentiles are in yellow.

To showcase the effects of an increase in the global supply chain pressure indicator, we select Germany and Italy as examples of a core and periphery country with a large manufacturing share of overall industrial output and we select France and Spain as large core and periphery countries with a comparatively low manufacturing share. On impact, the inflation outlook becomes more uncertain in all countries but Italy and is positively skewed in France and Spain. In addition, although inflation picks up in Germany, France, and Spain on impact, in Italy it only does with delay. A different pictures emerges for industrial production. In all countries, the production outlook becomes uncertain on impact; however, this uncertainty is pervasive in the countries with large manufacturing share (Germany and Italy).
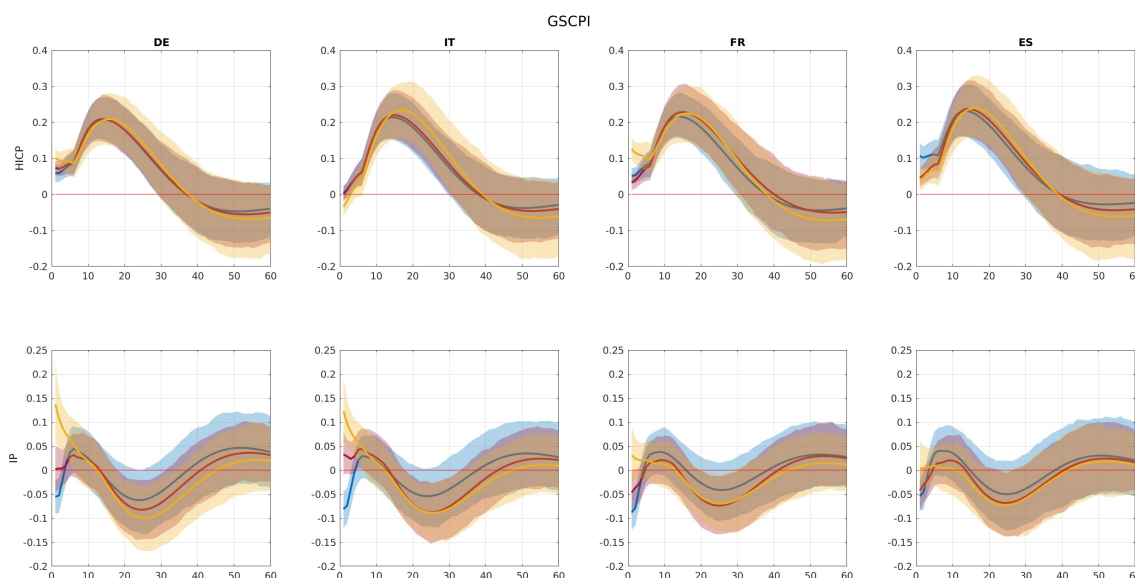
28

Figure 10: Impulse response functions (median and 68% bands) at 10th, 50th, and 90th percentiles of country-level inflation (first row) and industrial production (second row). Shock is to the global supply chain pressure index (GSCPI). Responses at the 10th percentiles are in blue, reponses at the 50th percentiles are in red, and responses at the 90th percentiles are in yellow.

Next, we turn to the IRFs in response to tightening global financial conditions. Here we focus on Germany and the Netherlands as examples of large and small core countries, and we focus on Italy and Portugal as examples of large and small periphery countries. On impact, inflation responds negatively in all countries. In Portugal, inflation responds more negatively overall, but the response in the Netherlands is characterized by greater uncertainty and a particularly marked response at the 10th percentile. Tail risks are hence more pronounced than in the other countries. Equally interesting are the responses of industrial production. On average the 90th percentile responds the least and the 10th percentile the most, indicating an overall increase in uncertainty in all countries. In Germany, Italy, and Portugal, the response of the median and 90th percentile are rather similar. In the Netherlands, the response is more muted and the 90th percentile responds positively on impact. The largest differences emerge for the response at the 10th percentile. The response is similar in Germany and Italy but more pronounced in Portugal and especially the Netherlands. In this stylized study, smaller countries are hence more susceptible to downside risks to the production outlook, following an increase in global financial conditions. More generally, these results are well in line with our findings from above as well as the literature that finds that financial conditions are an important driver of downside risks to output (Adrian et al., 2019).
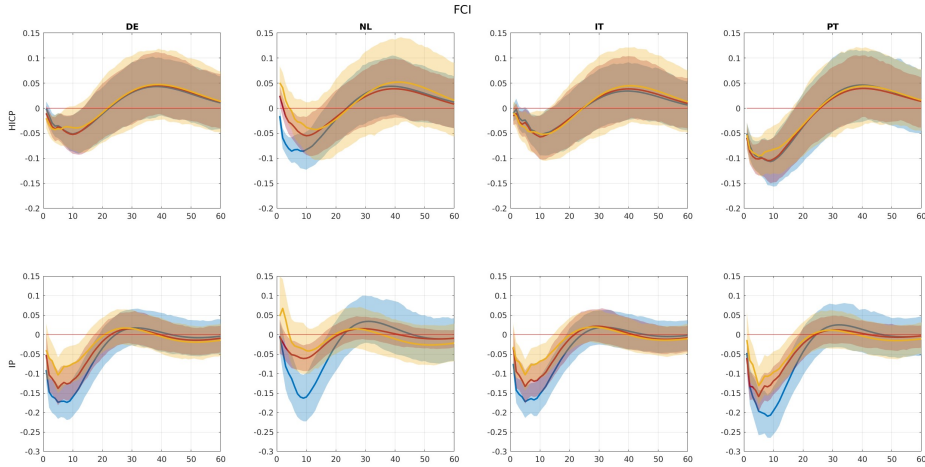
Figure 11: Impulse response functions (median and 68% bands) at 10th, 50th, and 90th percentiles of country-level inflation (first row) and industrial production (second row). Shock is to the financial conditions index (FCI). Responses at the 10th percentiles are in blue, reponses at the 50th percentiles are in red, and responses at the 90th percentiles are in yellow.

Finally, we turn attention to global economic policy uncertainty. Here we show the responses of Germany, France, Spain, and Portugal. The inflation response of Germany and France are roughly similar and symmetric. Both countries only respond little on impact and the dynamics are generally more muted compared to the other two countries. What stands out is that the periphery countries (Portugal and Spain) respond negatively on impact. The responses of industrial production show greater heterogeneity across quantile levels. With the exception of Portugal, the on-impact responses are stronger from the 90th to the 10th percentile. As before, this is indicative of increasing uncertainty about the outlook, a feature hidden in a classical FAVAR.

Our stylized analysis focuses on generalized impulse response functions and hence does not allow statements linked to specific structural shocks. Nonetheless, the country-level IRFs are evidence of strong heterogeneity, not only across quantile levels, but also across individual EA countries. Given the structure of the monetary union and the conduct of a common monetary policy, the QFAVAR offers policy-makers a tool to monitor not only these heterogeneities, but also risks to fragmentation of the broader transmission of economic shocks. Taking the scenario-analysis view based on the proposed QFAVAR might hence be a particularly useful addition to the policy analysis toolbox.

3.4.3 Quantile connectedness

To conclude the empirical section, we extend our FEVD analysis and derive quantile connectedness measures similar to Diebold and Yılmaz (2014). In a first step, we use the
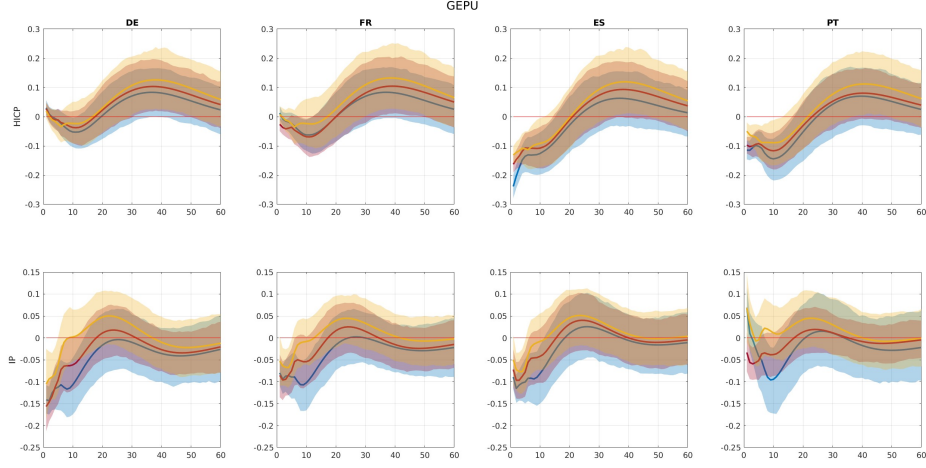
Figure 12: Impulse response functions (median and 68% bands) at 10th, 50th, and 90th percentiles of country-level inflation (first row) and industrial production (second row). Shock is to the global economic policy uncertainty (GEPU). Responses at the 10th percentiles are in blue, reponses at the 50th percentiles are in red, and responses at the 90th percentiles are in yellow.

state space form of our model to map the FEVDs from the state equation to the panel of individual country-level variables. In a second step, we construct a variance decomposition matrix for each forecast horizon,

$$D^h = \begin{bmatrix} d^h_{1,1} & \cdots & d^h_{1,l+k} \\ \vdots & \ddots & \vdots \\ d^h_{mnr,1} & \cdots & d^h_{mnr,l+k} \end{bmatrix}, \tag{12}$$

where element $d^h_{i,j}$ denotes the fraction of variable i^{th} forecast error attributed to state variable j at forecast horizon h. Alternatively, we relate the variance decomposition matrix, D^h, to a weighted network adjacency matrix, A. Compared to Diebold and Yılmaz (2014), however, the matrices D^h and A are not directly equivalent, because our state-space system has $l + k$ shocks and $n \times m \times r$ country-level variables. To represent the variance decomposition matrix as directed network, we stack the observed variables, quantile factors, and global variables, and we construct the augmented variance decomposition matrix,

$$\tilde{D}^h = \begin{bmatrix} \mathbf{0} & D^h \\ \mathbf{0} & D^h_s \end{bmatrix} = A, \tag{13}$$

where D^h_s denotes the variance decomposition matrix for the state equation that is constructed analogously to D^h. In line with Diebold and Yılmaz (2014) we now define

$$C_{i \leftarrow j}^h = \tilde{d}_{i,j}^h \qquad (14)$$

as the pairwise directional connectedness from j to i. D^h hence stores the directional connectedness from the global variables and quantile factors to the country-level variables, and D_s^h contains the pairwise connectedness among the global variables and quantile factors. The **0** entries are an artefact of the QFAVAR's structure and imply that there is no directional connectedness among the country-level variables or from the country-level variables to the global variables and quantile factors.

To illustrate the QFAVAR connectedness, figure 13 displays the directed network graphs implied by $\tilde{D}^h$. Given the large dimensions of our model, in line with the previous sections we focus on interesting sub graphs. The left column of figure 13 contains the connectedness of country-level inflation with the global variables for the three quantile levels of interest, $q_r \in [0.1, 0.5, 0.9]$, for $h = 24$. The right column contains the analogous directed graphs for country-level industrial production. In addition, we exclude edges for which the directional connectedness is smaller than 5%. In general, the thicker the edge, the stronger the directional connectedness. A few features stand out immediately. For inflation, country-level inflation overall is strongly connected to the global variables. The global supply chain pressure index, and to a lesser extent global inflation, have the largest directional connectedness to country-level inflation for all quantile levels. Across quantile levels, the patterns are relatively similar; however, although global economic policy uncertainty is connected to inflation in all countries at the 10% level, it only connects to a subset of countries at the 50% and 90% level. In contrast, for industrial production the pattern is very heterogeneous across quantile levels. At the 10% level the FCI has the strongest connectedness, especially to Spain. Global economic policy uncertainty emerges as the second most connected global variable. At the 50% level, the FCI connects to industrial production in all countries but the Netherlands. Global economic policy uncertainty only connects to industrial production in the Netherlands, Austria, and Germany. At the 90% level the global FCI only connects to Spain, France, Austria, and Portugal, and the global supply chain pressure index now connects to industrial production in Italy, Germany, and Belgium. Overall, though the connectedness networks reflect the observations for the FEVD of the quantile factors in section 3.4.1, they allow us to observe previously hidden heterogeneity across EA countries.
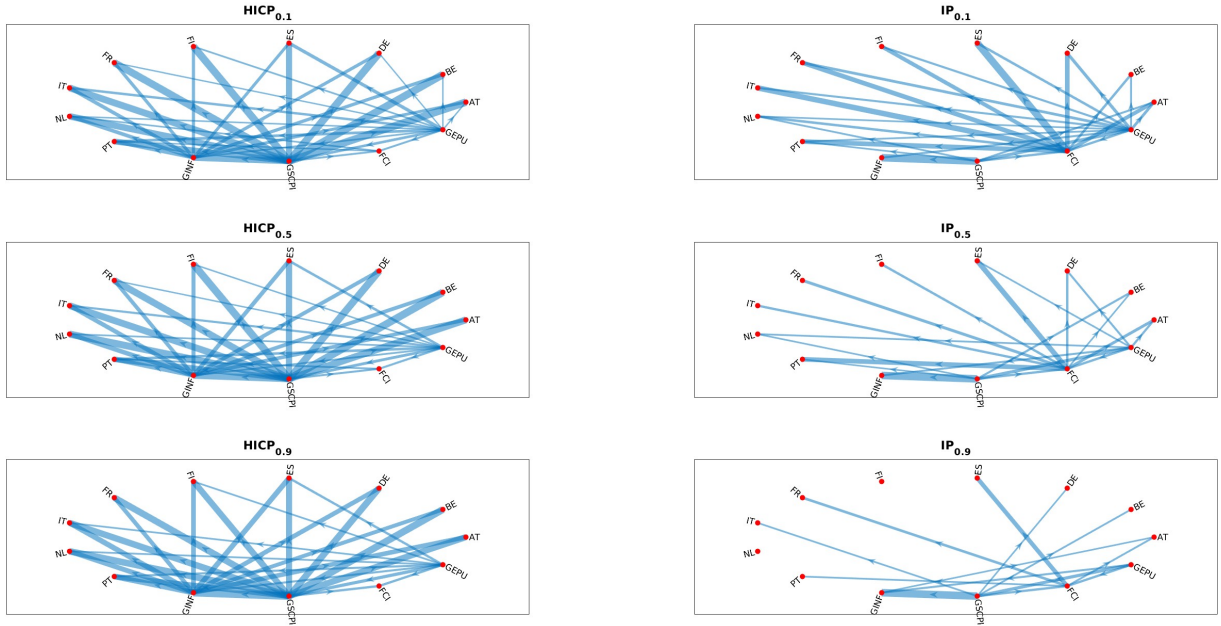
Figure 13: Pairwise Directional Connectedness Networks.

4 Conclusions

We develop a new quantile factor augmented vector autoregressive (QFAVAR) model that is a natural extension of the popular FAVAR for targeting specific quantiles of the distribution of macroeconomic data. The advantage of the factor-based approach over quantile VAR modeling is that it is both flexible and parsimonious. The factors not only summarize cross-sectional correlations across different countries, but also across different quantiles of the data distribution. Using a Bayesian perspective, estimation of the QFAVAR adds only a minor level of complexity relative to a Markov chain Monte Carlo algorithm for the classical FAVAR model (Bernanke et al., 2005). As MCMC can be computationally demanding in high dimensions, we also develop a simple two-step variational Bayes estimator that is appropriate for the computational demands associated with recursive out-of-sample forecasting exercises.

The proposed QFAVAR is fully parametric (likelihood-based), meaning that our proposed specification can easily extend to incorporate other formulations without inducing huge estimation and setup costs. First, we can obtain quantile dynamic factor models and univariate quantile autoregressions as special cases of the QFAVAR simply by restricting certain parameters of the state-space form of the model. Second, we can incorporate interesting features such as time-varying parameters and stochastic volatility in the measurement and/or state equations of the model in order to allow for more

33

flexible inference. Finally, although our approach to impulse response analysis follows a more neutral approach, by focusing on generalized impulse response functions, the QFAVAR can trivially be treated as a formal structural VAR model. We leave these extensions for future research.

A Technical Appendix: Derivation of the linear state-space form

Throughout the following, lower case letters indicate scalars, bold lower case letters indicate vectors, and bold upper case letters denote matrices. Let $y_{ij,t}$ denote a macroeconomic or financial indicator $i = 1, \ldots, m$ for country $j = 1, \ldots, n$ that is observed for time $t = 1, \ldots, T$. Additionally, let $q \in (0, 1)$ denote a given quantile level. The quantile factor regression for variable $y_{ij,t}$ for quantile level q is of the form

$$y_{ijt} = c_{ij(q)} + \lambda_{ij(q)} f_{t(q)}^i + \boldsymbol{\gamma}_{ij(q)} \boldsymbol{g}_t + u_{ij,t(q)}, \quad u_{ij,t(q)} \sim AL(0, \sigma_{ij(q)}, q) \tag{A.1}$$

Stacking all variables over i, j, $i = 1, ..., m$, $j = 1, ..., n$, into the vector $\boldsymbol{y}_t = [y_{11t}, ..., y_{1nt}, ..., y_{m1t}, ..., y_{mnt}]'$ we obtain the model

$$\boldsymbol{y}_t = \boldsymbol{c}_q + \boldsymbol{\lambda}_{(q)} \boldsymbol{f}_{t(q)} + \boldsymbol{\gamma}_{(q)} \boldsymbol{g}_t + \boldsymbol{u}_{t(q)}, \tag{A.2}$$

where $\boldsymbol{c}_{(q)} = \left[c_{11(q)}, ..., c_{1n(q)}, ..., c_{m1(q)}, ..., c_{mn(q)} \right]'$ is an $nm \times 1$ vector, $\boldsymbol{\gamma}_{(q)} = \left[\boldsymbol{\gamma}_{11(q)}', ..., \boldsymbol{\gamma}_{1n(q)}', ..., \boldsymbol{\gamma}_{m1(q)}', ..., \boldsymbol{\gamma}_{mn(q)}' \right]'$ is an $nm \times k$ matrix, $\boldsymbol{u}_{t(q)} = \left[u_{11t(q)}, ..., u_{1nt(q)}, ..., u_{m1t(q)}, ..., u_{mnt(q)} \right]'$ is an $nm \times 1$ vector,

$$\boldsymbol{\lambda}_{(q)} = diag\left(\boldsymbol{\lambda}_{1(q)}, \boldsymbol{\lambda}_{2(q)}, ..., \boldsymbol{\lambda}_{m(q)} \right) = \begin{bmatrix} \begin{bmatrix} \lambda_{11(q)} \\ \vdots \\ \lambda_{1n(q)} \end{bmatrix} & \boldsymbol{0} & \dots & \boldsymbol{0} \\ \boldsymbol{0} & \begin{bmatrix} \lambda_{21(q)} \\ \vdots \\ \lambda_{2n(q)} \end{bmatrix} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \boldsymbol{0} \\ \boldsymbol{0} & \dots & \boldsymbol{0} & \begin{bmatrix} \lambda_{m1(q)} \\ \vdots \\ \lambda_{mn(q)} \end{bmatrix} \end{bmatrix}$$

and $\boldsymbol{f}_{t(q)} = \begin{bmatrix} f^1_{t(q)} \\ f^2_{t(q)} \\ \vdots \\ f^m_{t(q)} \end{bmatrix}$. Equation (A.2) shows the FAVAR measurement equation for each

quantile level q. Stacking across r quantiles $q = q_1, ..., q_r$ we obtain

$$\boldsymbol{Y}_t = \boldsymbol{c} + \boldsymbol{\Lambda}\boldsymbol{F}_t + \boldsymbol{\Gamma}\boldsymbol{g}_t + \boldsymbol{u}_t, \tag{A.3}$$

where $\boldsymbol{c} = \left[\boldsymbol{c}'_{(q_1)}, ..., \boldsymbol{c}'_{(q_r)}\right]'$ is an $(nmr \times 1)$ vector, $\boldsymbol{\Gamma} = \left[\boldsymbol{\gamma}'_{(q_1)}, ..., \boldsymbol{\gamma}'_{(q_r)}\right]'$ is an $(nmr \times k)$ matrix, $\boldsymbol{F}_t = \left[\boldsymbol{f}'_{t(q_1)}, ..., \boldsymbol{f}'_{t(q_r)}\right]'$ is an $(mr \times 1)$ vector and $\boldsymbol{\Lambda} = \left[\boldsymbol{\lambda}'_{(q_1)}, ..., \boldsymbol{\lambda}_{(q_r)'}\right]'$ is an $(nmr \times mr)$ matrix.

Augmenting the measurement equation (A.3) with an identity for $\boldsymbol{g}_t$ and combining it with the state equation (4) we obtain the state-space form of the QFAVAR

$$\begin{bmatrix} \boldsymbol{Y}_t \\ \boldsymbol{g}_t \end{bmatrix} = \boldsymbol{c} + \begin{bmatrix} \boldsymbol{\Lambda} & \boldsymbol{\Gamma} \\ \boldsymbol{0} & \boldsymbol{I} \end{bmatrix} \begin{bmatrix} \boldsymbol{F}_t \\ \boldsymbol{g}_t \end{bmatrix} + \begin{bmatrix} \boldsymbol{u}_t \\ \boldsymbol{0} \end{bmatrix}, \tag{A.4}$$

$$\begin{bmatrix} \boldsymbol{F}_t \\ \boldsymbol{g}_t \end{bmatrix} = \boldsymbol{v} + \boldsymbol{\Phi}_1 \begin{bmatrix} \boldsymbol{F}_{t-1} \\ \boldsymbol{g}_{t-1} \end{bmatrix} + ... + \boldsymbol{\Phi}_p \begin{bmatrix} \boldsymbol{F}_{t-p} \\ \boldsymbol{g}_{t-p} \end{bmatrix} + \boldsymbol{\varepsilon}_t. \tag{A.5}$$

Each element of the vector $\boldsymbol{u}_t$ is distributed as independent univariate asymmetric Laplace. As explained in the main text we can write the asymmetric Laplace as a Gaussian-Exponential location-scale mixture, in which case we the state-space model above is in conditionally normal form and sampling of the state-vector $[\boldsymbol{F}'_t, \boldsymbol{g}'_t]'$ is feasible using the simulation smoother of Carter and Kohn (1994). Finally, note that sampling of the state-form requires first-order Markov dependence of the state variable. In equation (A.5) above the state vector follows a VAR(p) but we can use standard tools for writing it in VAR(1) companion form, (see Lütkepohl, 2005, Chapter 2). Detailed derivations of MCMC and variational Bayes algorithms for inference, are provided in the online supplement.

References

ADRIAN, T., N. BOYARCHENKO, AND D. GIANNONE (2019): "Vulnerable Growth," *American Economic Review*, 109, 1263–89.

——— (2021): "Multimodality in macrofinancial dynamics," *International Economic Review*, 62, 861–886.

ANDO, T. AND J. BAI (2020): "Quantile Co-Movement in Financial Markets: A Panel Quantile Model With Unobserved Heterogeneity," *Journal of the American Statistical Association*, 115, 266–279.

ANDO, T., M. GREENWOOD-NIMMO, AND Y. SHIN (2022): "Quantile Connectedness: Modeling Tail Behavior in the Topology of Financial Networks," *Management Science*, 68, 2401–2431.

BEKAERT, G. AND A. POPOV (2019): "On the Link Between the Volatility and Skewness of Growth," *IMF Economic Review*, 67, 746–790.

BERNANKE, B. S., J. BOIVIN, AND P. ELIASZ (2005): "Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach*," *The Quarterly Journal of Economics*, 120, 387–422.

BLEI, D. M., A. KUCUKELBIR, AND J. D. MCAULIFFE (2017): "Variational Inference: A Review for Statisticians," *Journal of the American Statistical Association*, 112, 859–877.

BRITTON, E., P. FISHER, AND J. WHITLEY (1998): "The Inflation Report projections: Understanding the fan chart," *Bank of England Quarterly Bulletin*, Q1, 30–37.

CARRIERO, A., J. CHAN, T. E. CLARK, AND M. MARCELLINO (2022): "Corrigendum to "Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors" [J. Econometrics 212 (1) (2019) 137–154]," *Journal of Econometrics*, 227, 506–512.

CARTER, C. K. AND R. KOHN (1994): "On Gibbs Sampling for State Space Models," *Biometrika*, 81, 541–553.

CARVALHO, C. M., N. G. POLSON, AND J. G. SCOTT (2010): "The Horseshoe Estimator for Sparse Signals," *Biometrika*, 97, 465–480.

CASTELNUOVO, E. AND L. MORI (2022): "Uncertainty, Skewness, and the Business Cycle through the MIDAS Lens," CESifo Working Paper Series 10062, CESifo.

CHAN, J. C. (2013): "Moving average stochastic volatility models with application to inflation forecast," *Journal of Econometrics*, 176, 162–172.

CHAVLEISHVILI, S. AND S. MANGANELLI (2020): "Forecasting and Stress Testing with Quantile Vector Autoregression," Working Paper No 2330, European Central Bank.

CHEN, L., J. J. DOLADO, AND J. GONZALO (2021): "Quantile Factor Models," *Econometrica*, 89, 875–910.

CICCARELLI, M. AND B. MOJON (2010): "Global Inflation," *The Review of Economics and Statistics*, 92, 524–535.

CLARK, T. E., F. HUBER, G. KOOP, M. MARCELLINO, AND M. PFARRHOFER (2021): "Investigating Growth at Risk Using a Multi-country Non-parametric Quantile Factor Model," .

DESPOIS, T. AND C. DOZ (forthcoming): "Identifying and interpreting the factors in factor models via sparsity: Different approaches," *Journal of Applied Econometrics*, n/a.

DIEBOLD, F. X. AND K. YILMAZ (2014): "On the network topology of variance decompositions: Measuring the connectedness of financial firms," *Journal of econometrics*, 182, 119–134.

FELDKIRCHER, M., F. HUBER, G. KOOP, AND M. PFARRHOFER (2022): "Approximate Bayesian inference and forecasting in huge-dimensional multicountry VARs," *International Economic Review*, 63, 1625–1658.

FORNI, M., L. GAMBETTI, AND L. SALA (2021): "Downside and Upside Uncertainty Shocks," CEPR Discussion Papers 15881, C.E.P.R. Discussion Papers.

GORODNICHENKO, Y. AND S. NG (2017): "Level and volatility factors in macroeconomic data," *Journal of Monetary Economics*, 91, 52–68, the Swiss National Bank/Study Center Gerzensee Special Issue: "Modern Macroeconomics: Study Center Gerzensee Conference in Honor of Robert G. King" Sponsored by the Swiss National Bank and the Study Center Gerzensee.

JENSEN, H., I. PETRELLA, S. H. RAVN, AND E. SANTORO (2020): "Leverage and Deepening Business-Cycle Skewness," *American Economic Journal: Macroeconomics*, 12, pp. 245–281.

KHARE, K. AND J. P. HOBERT (2012): "Geometric ergodicity of the Gibbs sampler for Bayesian quantile regression," *Journal of Multivariate Analysis*, 112, 108–116.

KOOP, G. AND D. KOROBILIS (2010): "Bayesian Multivariate Time Series Methods for Empirical Macroeconomics," *Foundations and Trends® in Econometrics*, 3, 267–358.

——— (2014): "A new index of financial conditions," *European Economic Review*, 71, 101–116.

KOROBILIS, D. (2013): "Assessing the Transmission of Monetary Policy Using Time-varying Parameter Dynamic Factor Models," *Oxford Bulletin of Economics and Statistics*, 75, 157–179.

——— (2017): "Quantile regression forecasts of inflation under model uncertainty," *International Journal of Forecasting*, 33, 11–20.

——— (2022): "A new algorithm for structural restrictions in Bayesian vector autoregressions," *European Economic Review*, 148, 104241.

KOROBILIS, D., B. LANDAU, A. MUSSO, AND A. PHELLA (2021): "The time-varying evolution of inflation risks," Working Paper Series 2600, European Central Bank.

KOROBILIS, D. AND M. SCHRÖDER (2022): "Probabilistic quantile factor analysis," Tech. Rep. 2212.10301, Arxiv.

KOROBILIS, D. AND K. SHIMIZU (2022): "Bayesian Approaches to Shrinkage and Sparse Estimation," *Foundations and Trends in Econometrics*, 11, 230–354.

KOSE, M. A., C. OTROK, AND C. H. WHITEMAN (2003): "International Business Cycles: World, Region, and Country-Specific Factors," *American Economic Review*, 93, 1216–1239.

LÓPEZ-SALIDO, J. D. AND F. LORIA (2019): "Inflation at Risk," Discussion Papers No. 14074, Center for Economic Policy Research.

LORIA, F., C. MATTHES, AND D. ZHANG (2019): "Assessing Macroeconomic Tail Risk," Finance and Economics Discussion Series 2019-026, Board of Governors of the Federal Reserve System (U.S.).

LÜTKEPOHL, H. (2005): *New Introduction to Multiple Time Series Analysis*, no. 978-3-540-27752-1 in Springer Books, Springer.

MA, S., O. LINTON, AND J. GAO (2021): "Estimation and inference in semiparametric quantile factor models," *Journal of Econometrics*, 222, 295–323, annals Issue:Financial Econometrics in the Age of the Digital Economy.

MAKALIC, E. AND D. F. SCHMIDT (2016): "A Simple Sampler for the Horseshoe Estimator," *IEEE Signal Processing Letters*, 23, 179–182.

MANZAN, S. (2015): "Forecasting the Distribution of Economic Variables in a Data-Rich Environment," *Journal of Business & Economic Statistics*, 33, 144–164.

MUGGENTHALER, P., J. SCHROTH, AND Y. SUN (2021): "The heterogeneous economic impact of the pandemic across euro area countries," ECB Economic Bulletin No. 5, European Central Bank.

MUMTAZ, H. AND P. SURICO (2012): "Evolving International Inflation Dynamics: World and Country-Specific Factors," *Journal of the European Economic Association*, 10, 716–734.

——— (2015): "The transmission mechanism in good and bad times," *International Economic Review*, 56, 1237–1260.

ORMEROD, J. T., C. YOU, AND S. MÜLLER (2017): "A variational Bayes approach to variable selection," *Electronic Journal of Statistics*, 11, 3549 – 3594.

PANETTA, F. (2023): "Everything everywhere all at once: Responding to multiple global shocks," Speech at a panel on "Global shocks, policy spillovers and geo-strategic risks: how to coordinate policies" at The ECB and its Watchers XXIII Conference [Accessed: 2023 03 23].

PLAGBORG-MØLLER, M., L. REICHLIN, G. RICCO, AND T. HASENZAGL (2020): "When is Growth at Risk?" *Brookings Papers on Economic Activity*, 2020, 167–229.

SCHÜLER, Y. S. (2020): "The impact of uncertainty and certainty shocks," Discussion Papers 14/2020, Deutsche Bundesbank.

STOCK, J. AND M. WATSON (2016): "Chapter 8 - Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics," in *Handbook of Macroeconomics*, ed. by J. B. Taylor and H. Uhlig, Elsevier, vol. 2, 415–525.

STOCK, J. H. AND M. W. WATSON (2005): "Implications of Dynamic Factor Models for VAR Analysis," Working Paper 11467, National Bureau of Economic Research.

WHITE, H., T.-H. KIM, AND S. MANGANELLI (2015): "VAR for VaR: Measuring tail dependence using multivariate regression quantiles," *Journal of Econometrics*, 187, 169–188.

YU, K. AND R. A. MOYEED (2001): "Bayesian quantile regression," *Statistics & Probability Letters*, 54, 437–447.

Online supplement to "Monitoring multicountry macroeconomic risk"

Dimitris Korobilis and Maximilian Schröder

A Bayesian estimation of the Quantile FAVAR model

A.1 Derivation of the linear, Gaussian state-space form

Throughout the following, lower case letters indicate scalars, bold lower case letters indicate vectors, and bold upper case letters denote matrices. Let $y_{ij,t}$ denote a macroeconomic or financial indicator $i = 1, \ldots, m$ for country $j = 1, \ldots, n$ that is observed for time $t = 1, \ldots, T$. Additionally, let $q \in (0,1)$ denote a given quantile level. The measurement equation of the QFAVAR for variable $y_{ij,t}$ for quantile level q is then of the form

$$y_{ij,t} = c_{ij(q)} + \beta_{ij(q)} y_{ij,t-1} + \boldsymbol{\lambda}'_{ij(q)} \boldsymbol{f}_{t(q)} + \boldsymbol{\gamma}'_{ij(q)} \boldsymbol{g}_t + u_{ij,t(q)}, \quad u_{ij,t(q)} \sim AL(0, \sigma_{ij(q)}, q) \quad (1)$$

where $\boldsymbol{g}_t$ denotes the set of observed global variables, $\boldsymbol{f}_t(q)$ denotes the quantile factors, and $\boldsymbol{\gamma}_{ij(q)}$ and $\boldsymbol{\lambda}_{ij(q)}$ are conformable loading vectors. $c_ij(q)$ and $\beta_{ij(q)}$ denote a constant and an autoregressive coefficient, respectively. To ease notation, we will suppress $c_{ij(q)}$ and $\beta_{ij(q)}$ in the following expressions without loss of generality. Finally, $AL(0, \sigma_{ij(q)}, q)$ denotes the univariate asymmetric Laplace density with the location parameter set to 0, scale parameter $\sigma_{ij(q)}$, and asymmetry parameter q. The use of the univariate asymmetric Laplace later implies a diagonal covariance matrix in the measurement equation and thus mirrors the standard identifying assumption used in linear Gaussian factor models. This distribution is parametrized as

$$u_{ij,t(q)} \sim \frac{q(1-q)}{\sigma_{ij,t}} \left[e^{(1-q)\frac{u_{ij,t(q)}}{\sigma_{ij(q)}}} \mathbb{I}(u_{ij,t(q)} \leq 0) + e^{-q\frac{u_{ij,t(q)}}{\sigma_{ij(q)}}} \mathbb{I}(u_{ij,t(q)} > 0) \right]. \quad (2)$$

The first step towards tractable Bayesian inference is to rewrite the asymmetric Laplace likelihood as a conditionally Gaussian likelihood, which greatly simplifies inference. Following Yu and Moyeed (2001) the AL distribution can equivalently be expressed as a normal-exponential mixture of the form

$$u_{ij,t(q)}|z_{ij,t(q)} \sim \frac{1}{\sqrt{2\pi z_{ij,t(q)}\sigma_{ij(q)}\kappa_{2(q)}^2}} exp\left\{-\frac{\left(y_{ij,t} - \boldsymbol{\lambda}'_{ij(q)}\boldsymbol{f}_{t(q)} - \boldsymbol{\gamma}'_{ij(q)}\boldsymbol{g}_t - \kappa_{1(q)}z_{ij,t(q)}\right)^2}{2z_{ij,t(q)}\sigma_{ij(q)}\kappa_{2(q)}^2}\right\} exp\left\{-\frac{z_{ij,t(q)}}{\sigma_{ij(q)}}\right\},$$

(3)

or compactly $u_{ij,t(q)}|z_{ij,t(q)} \sim N(\kappa_1(q)z_{ij,t(q)}, \kappa_{2(q)}^2\sigma_{ij(q)}z_{ij,t(q)})$, with $z_{ijt,(q)} \sim Exp(\sigma_{ij(q)})$, where $Exp(\bullet)$ denotes the exponential distribution, $\kappa_{1(q)} = \frac{1-2q}{q(1-q)}$, and $\kappa_{2(q)}^2 = \frac{2}{q(1-q)}$. Using this scale mixture of normals representation, we can rewrite 1 as

$$y_{ij,t} = \boldsymbol{\lambda}'_{ij(q)}\boldsymbol{f}_{t(q)} + \boldsymbol{\gamma}'_{ij(q)}\boldsymbol{g}_t + \kappa_{1(q)}z_{ij,t(q)} + \kappa_{2(q)}\sqrt{\sigma_{ij(q)}z_{ij,t(q)}}\nu_{ij,t}, \quad \nu_{ij,t} \sim N(0,1).$$

(4)

For a given quantile level q, we can now collect the set of variables and stack their respective measurement equations across countries, j, and indicators, i

$$
\begin{bmatrix} y_{11,t} \\ \vdots \\ y_{1n,t} \\ y_{21,t} \\ \vdots \\ y_{2n,t} \\ \vdots \\ y_{mn,t} \end{bmatrix}
=
\begin{bmatrix} \boldsymbol{\lambda}'_{11(q)} & \boldsymbol{\gamma}'_{11(q)} \\ \vdots & \vdots \\ \boldsymbol{\lambda}'_{1n(q)} & \boldsymbol{\gamma}'_{1n(q)} \\ \boldsymbol{\lambda}'_{21(q)} & \boldsymbol{\gamma}'_{21(q)} \\ \vdots & \vdots \\ \boldsymbol{\lambda}'_{2n(q)} & \boldsymbol{\gamma}'_{2n(q)} \\ \vdots & \vdots \\ \boldsymbol{\lambda}'_{mn(q)} & \boldsymbol{\gamma}'_{mn(q)} \end{bmatrix}
\begin{bmatrix} \boldsymbol{f}_{t(q)} \\ \boldsymbol{g}_t \end{bmatrix}
+
\begin{bmatrix} \tilde{z}_{11,t(q)} \\ \vdots \\ \tilde{z}_{1n,t(q)} \\ \tilde{z}_{21,t(q)} \\ \vdots \\ \tilde{z}_{2n,t(q)} \\ \vdots \\ \tilde{z}_{mn,t(q)} \end{bmatrix}
+
\begin{bmatrix} \tilde{\nu}_{11,t(q)} \\ \vdots \\ \tilde{\nu}_{1n,t(q)} \\ \tilde{\nu}_{21,t(q)} \\ \vdots \\ \tilde{\nu}_{2n,t(q)} \\ \vdots \\ \tilde{\nu}_{mn,t(q)} \end{bmatrix}
$$

(5)

where $\tilde{\nu}_{ij,t(q)}$ collects $\kappa_{2(q)}\sqrt{\sigma_{ij(q)}z_{ij,t(q)}}\nu_{ij,t}$ and $\tilde{z}_{ij,t(q)}$ denotes $\kappa_{1(q)}z_{ij,t(q)}$ to simplify notation. Importantly, different quantile levels maintain the same structure. To model multiple quantile levels simultaneously, we now stack the block of equations in 5 for different quantile levels. For example, for three arbitrary quantile levels $q = \{q_1, q_2, q_3\}$, this yields

$$
\begin{bmatrix} \boldsymbol{y}_{1\bullet,t} \\ \vdots \\ \boldsymbol{y}_{m\bullet,t} \\ \hline \boldsymbol{y}_{1\bullet,t} \\ \vdots \\ \boldsymbol{y}_{m\bullet,t} \\ \hline \boldsymbol{y}_{1\bullet,t} \\ \vdots \\ \boldsymbol{y}_{m\bullet,t} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\Lambda}'_{1\bullet,t(q_1)} & \mathbf{0} & \mathbf{0} & \boldsymbol{\gamma}'_{1\bullet,t(q_1)} \\ \vdots & \vdots & \vdots & \vdots \\ \boldsymbol{\Lambda}'_{m\bullet(q_1)} & \mathbf{0} & \mathbf{0} & \boldsymbol{\gamma}'_{m\bullet,t(q_1)} \\ \mathbf{0} & \boldsymbol{\Lambda}'_{1\bullet(q_2)} & \mathbf{0} & \boldsymbol{\gamma}'_{1\bullet,t(q_2)} \\ \vdots & \vdots & \vdots & \\ \mathbf{0} & \boldsymbol{\Lambda}'_{m\bullet(q_2)} & \mathbf{0} & \boldsymbol{\gamma}'_{m\bullet,t(q_2)} \\ \mathbf{0} & \mathbf{0} & \boldsymbol{\Lambda}'_{1\bullet(q_3)} & \boldsymbol{\gamma}'_{1\bullet,t(q_3)} \\ \vdots & \vdots & \vdots & \\ \mathbf{0} & \mathbf{0} & \boldsymbol{\Lambda}'_{m\bullet(q_3)} & \boldsymbol{\gamma}'_{m\bullet,t(q_3)} \end{bmatrix} \begin{bmatrix} \boldsymbol{f}_{t(q_1)} \\ \boldsymbol{f}_{t(q_2)} \\ \boldsymbol{f}_{t(q_3)} \\ \boldsymbol{g}_t \end{bmatrix} + \begin{bmatrix} \tilde{z}_{1\bullet,t(q_1)} \\ \vdots \\ \tilde{z}_{m\bullet,t(q_1)} \\ \tilde{z}_{1\bullet,t(q_2)} \\ \vdots \\ \tilde{z}_{m\bullet,t(q_2)} \\ \tilde{z}_{1\bullet,t(q_3)} \\ \vdots \\ \tilde{z}_{m\bullet,t(q_3)} \end{bmatrix} + \begin{bmatrix} \tilde{\nu}_{1\bullet,t(q_1)} \\ \vdots \\ \tilde{\nu}_{m\bullet,t(q_1)} \\ \tilde{\nu}_{1\bullet,t(q_2)} \\ \vdots \\ \tilde{\nu}_{m\bullet,t(q_2)} \\ \tilde{\nu}_{1\bullet,t(q_3)} \\ \vdots \\ \tilde{\nu}_{m\bullet,t(v3)} \end{bmatrix}
\tag{6}
$$

Note that the LHS of equation 6 repeatedly stacks the observed variables $y_{ij,t}$. A key feature of this formulation is that the conditional quantiles of each variable only load onto the quantile factors that are defined at the same quantile level. This induces sparsity in the loadings matrix and hence keeps the model parsimonious which in turn facilitates inference.

With the measurement equations in place, we now turn to defining the state equation. The state equation generally follows a VAR(p) and describes the joint evolution of the quantile factors $\boldsymbol{f}_{t(q)}$ and the observed global variables $\boldsymbol{g}_t$. Without loss of generality, we restrict attention to a VAR(1) in order to ease notation.

$$
\begin{bmatrix} \boldsymbol{f}_{t(q_1)} \\ \boldsymbol{f}_{t(q_2)} \\ \boldsymbol{f}_{t(q_3)} \\ \boldsymbol{g}_t \end{bmatrix} = \boldsymbol{v} + \boldsymbol{\Phi} \begin{bmatrix} \boldsymbol{f}_{t-1(q_1)} \\ \boldsymbol{f}_{t-1(q_2)} \\ \boldsymbol{f}_{t-1(q_3)} \\ \boldsymbol{g}_{t-1} \end{bmatrix} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim N(\boldsymbol{0}, \Omega)
\tag{7}
$$

where $\boldsymbol{\Phi}$ denotes a conformable coefficient matrix, $\boldsymbol{v}$ is a vector of constants, and $\boldsymbol{\varepsilon}_t$ denotes the vector of reduced form residuals. While the variables in the measurement equation evolve independently conditional on the factors, the state equation allows for the factors at different quantile levels to co-move with one another and with the global variables. Overall, we hence model the co-movements of a potentially high dimensional quantile surface, by compressing them down to a lower dimensional space. This ensures parsimony and computational tractability. Fully cast into state space form, the system is given by

$$
\begin{bmatrix} \boldsymbol{y}_t \\ \boldsymbol{g}_t \end{bmatrix} = \begin{bmatrix} \boldsymbol{\Lambda} & \boldsymbol{\Gamma} \\ 0 & \boldsymbol{I} \end{bmatrix} \begin{bmatrix} \boldsymbol{F}_t \\ \boldsymbol{g}_t \end{bmatrix} + \tilde{\nu}_t
\tag{8}
$$

$$
\begin{bmatrix} \boldsymbol{F}_t \\ \boldsymbol{g}_t \end{bmatrix} = \boldsymbol{\Phi} \begin{bmatrix} \boldsymbol{F}_{t-1} \\ \boldsymbol{g}_{t-1} \end{bmatrix} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim N(\boldsymbol{0}, \boldsymbol{\Omega})
\tag{9}
$$

3

where $\boldsymbol{y}_t = [y'_{11,t(q_1)}, \ldots, y'_{mn,t(q_3)}]'$, $\boldsymbol{F}_t = [\boldsymbol{f}'_{t(q_1)}, \ldots, \boldsymbol{f}'_{t(q_3)}]'$, $\tilde{\nu}_t = [\tilde{\nu}'_{11,t(q_1)}, \ldots, \tilde{\nu}'_{mn,t(q_r)}]'$, and $\boldsymbol{\Lambda}$ and $\boldsymbol{\Gamma}$ are conformable coefficient matrices collecting $\boldsymbol{\lambda}_{ij(q)}$ and $\boldsymbol{\gamma}_{ij(q)}$, respectively, as shown above. Further, $\boldsymbol{\Omega}$ denotes the covariance matrix of the reduced form residuals. The resulting QFAVAR is hence a linear Gaussian state-space model and standard filtering and smoothing algorithms apply.

As one additional extension, in our application we introduce stochastic volatility to the state equation and hence allow the diagonal elements of Ω_t to be time-varying. Again, standard methods for estimating models with stochastic volatility apply so we leave the details to the appendix on estimation.

A.2 Identification

As is common in standard linear Gaussian factor models, while the common component, $\boldsymbol{\Lambda F}_t$, is identified the loadings, $\boldsymbol{\Lambda}$, and factors, $\boldsymbol{F}_t$, are not identified individually. While forecasting exercises can still proceed even without identification, identification is necessary for structural analysis and the computation of IRFs. In our empirical exercise, we facilitate factor identification with three strategies.

1. First, we group the variables $y_{ij,t}$ by their economic interpretation and extract only one factor per block. The loadings matrix $\boldsymbol{\Lambda}_{(q)}$ collecting all loadings at a given quantile level is hence block-diagonal. Specifically, for $\boldsymbol{f}_{t(q)} = [f^1_{t(q)}, f^2_{t(q)}, \cdots, f^i_{t(q)}, \cdots, f^m_{t(q)}]'$ the corresponding loadings are given by

$$\boldsymbol{\lambda}'_{ij(q)} = [0, 0, \cdots, \lambda_{ij(q)}, \cdots, 0], \tag{10}$$

and

$$\boldsymbol{\Lambda}'_{i\bullet(q)} = \begin{bmatrix} 0 & 0 & \cdots & \lambda_{i1(q)} & \cdots & 0 \\ 0 & 0 & \cdots & \lambda_{i2(q)} & \cdots & 0 \\ \vdots & \vdots & & \vdots & & \vdots \\ 0 & 0 & \cdots & \lambda_{in(q)} & \cdots & 0 \end{bmatrix} \tag{11}$$

For two reasons, this has the additional benefit of giving the factors a specific interpretation. First, for e.g. all industrial production series, we will identify one industrial factor per quantile level. A similar strategy is already suggested in the original FAVAR framework proposed by Bernanke et al. (2005). Second, grouping series that move similarly and are conceptually related, such as industrial production from different EA countries, ensures that the quantile factors can be assigned to the corresponding quantile levels and maintain their economic interpretation. Grouping e.g. industrial production and unemployment

instead would mix across different economic concepts as the resulting e.g. 10^{th} percentile factor would be extracted from low unemployment, but also low industrial production. This greatly complicates interpretation. Following our strategy instead, we interpret the factors as e.g. the quantile industrial production factors.

Note, however, that this choice also has clear implications for the dynamics of the model. In the case of $\mathbf{\Gamma} = \mathbf{0}$, $\boldsymbol{\beta} = \mathbf{0}$, and $\boldsymbol{c} = \mathbf{0}$, i.e. the absence of global variables and constants in the measurement equation, the conditional quantiles of $\boldsymbol{y}_t$, $\boldsymbol{y}_t(q)$, are governed by the common component, $\mathbf{\Lambda}\boldsymbol{f}_t$, only. Given that we allow for only one factor per group of variables, this implies that the conditional quantiles $y_{ijt(q)}$ are just rescaled versions of the factors, $\boldsymbol{f}_t$. The conditional quantiles of a given variables for different countries, $\boldsymbol{y}_{i\bullet t(q)}$ then also feature identical dynamics and only differ in their scale. In order to highlight more meaningful heterogeneities, it is hence important to allow for additional RHS variables in the measurement equation, such as global variables.

2. To fix the scale of the quantile factors within their groups, we further restrict one loading per quantile to unity. This yields e.g.

$$\mathbf{\Lambda}'_{i\bullet(q)} = \begin{bmatrix} 0 & 0 & \cdots & 1 & \cdots & 0 \\ 0 & 0 & \cdots & \lambda_{i2(q)} & \cdots & 0 \\ \vdots & \vdots & & \vdots & & \vdots \\ 0 & 0 & \cdots & \lambda_{in(q)} & \cdots & 0 \end{bmatrix} \tag{12}$$

3. To identify the sign of the quantile factors, we extract quantile factors using the probabilistic quantile factor algorithm proposed by Korobilis and Schröder (2022) from the same set of variables prior to estimation. Each iteration of the algorithm, we then check whether the QFAVAR factors are positively correlated with the corresponding VBQFA factors and invert them if this is not the case.

A.3 Markov Chain Monte Carlo estimation

The Gibbs sampler requires deriving conditional posteriors. As outlined above, we formulated the QFAVAR as a linear Gaussian state space model and hence standard practices apply.

A.3.1 Prior Distributions

In our proposed approach, we specify the following prior distributions

$$\sigma_{ij(q)} \sim G^{-1}(r_0, s_0) \tag{13}$$

$$z_{ij,t(q)} \sim Exp(\sigma_{ij(q)}) \tag{14}$$

To simplify notation, let $\boldsymbol{\Phi}$ collect the coefficient vectors and matrices in the measurement equations, i.e. $\boldsymbol{\Lambda}$ and $\boldsymbol{\Gamma}$. Correspondingly, let $\hat{F}_t = [\boldsymbol{F}'_t, \boldsymbol{g}'_t]$. As outlined in the main body of the text, we impose shrinkage on $\boldsymbol{\Phi}$ by means of the horseshoe prior following Makalic and Schmidt (2016), given by

$$\boldsymbol{\phi}_{ij(q)} \big| \left\{ \bar{\lambda}_{ij(q),k,\phi}, \upsilon_{ij(q),k,\phi} \right\}_{k=1}^{l_\phi}, \tau_{ij(q),\phi}, \xi_{ij(q),\phi} \sim N(\mathbf{0}, \sigma_{ij(q)} \tau_{ij(q),\phi}^2 \bar{\boldsymbol{\Lambda}}_{ij(q),\phi}), \tag{15}$$

$$\bar{\lambda}_{ij(q),k,\phi}^2 \big| \upsilon_{ij(q),k,\phi} \sim G^{-1}\left(\frac{1}{2}, \frac{1}{\upsilon_{ij(q),k,\phi}}\right), \quad \text{for } k = 1, \cdots, l_\phi, \tag{16}$$

$$\upsilon_{ij(q),k,\phi} \sim G^{-1}\left(\frac{1}{2}, 1\right), \quad \text{for } k = 1, \cdots, l_\phi, \tag{17}$$

$$\tau_{ij(q),\phi}^2 \big| \xi_{ij(q),\phi} \sim G^{-1}\left(\frac{1}{2}, \frac{1}{\xi_{ij(q),\phi}}\right), \tag{18}$$

$$\xi_{ij(q),\phi} \sim G^{-1}\left(\frac{1}{2}, 1\right), \tag{19}$$

where $\bar{\Lambda}_{ij(q),\phi} = diag(\bar{\lambda}_{ij(q),1,\phi}^2, \cdots, \bar{\lambda}_{ij(q),l,\phi}^2)$, subscript k denotes the element of $\boldsymbol{\phi}_{ij(q)}$ for $k = 1, \cdots, l_\phi$, where l_ϕ is equal to the number of factors, n_f, number of global variables, n_g, plus the constant. Further, subscript ϕ indicates that the parameters of the shrinkage prior apply to the elements of the measurement equation. Similarly, because the VAR in the state equation 7 is heavily parametrized, we define another set of horseshoe priors to the elements of Ψ.

$$\boldsymbol{\psi}_r \big| \left\{ \bar{\lambda}_{s,\psi}, \upsilon_{s,\psi} \right\}_{s=1}^{l_\psi}, \tau_{r,\psi}, \xi_{r,\psi} \sim N(\mathbf{0}, \omega_r \tau_{r,\psi}^2 \bar{\boldsymbol{\Lambda}}_{r,\psi}), \tag{20}$$

$$\bar{\lambda}_{r,s,\psi}^2 \big| \upsilon_{r,s,\psi} \sim G^{-1}\left(\frac{1}{2}, \frac{1}{\upsilon_{r,s,\psi}}\right), \quad \text{for } s = 1, \cdots, l_\psi, \tag{21}$$

$$\upsilon_{r,s,\psi} \sim G^{-1}\left(\frac{1}{2}, 1\right), \quad \text{for } s = 1, \cdots, l_\psi, \tag{22}$$

$$\tau_{r,\psi}^2 \big| \xi_{r,\psi} \sim G^{-1}\left(\frac{1}{2}, \frac{1}{\xi_{r,\psi}}\right), \tag{23}$$

$$\xi_{r,\psi} \sim G^{-1}\left(\frac{1}{2}, 1\right), \tag{24}$$

where $r = 1, \cdots, R$, with $R = n_f + n_g$ indexes the state equation, $s = 1, \cdots, l_\psi$ indexes the individual parameter in the respective state equation and $l_\psi = (n_f + n_g) \times p$, where p denotes the number of lags, and subscript ψ indicates that the prior hyperparameters belong to the state equation.

To be able to estimate the QFAVAR with and without stochastic volatility in the state equation, we take the following approach to sample the elements of the covariance matrix. Let

$$\Omega_t = A H_t A' \tag{25}$$

6

with $H_t = diag(e^{h_{1,t}}, \cdots, e^{h_{R,t}})$. Further, A is a lower unitriangular coefficient matrix. Sampling the elements of A this way has a drawback. In principle, the triangular structure implies that the order of the states in the state equation matters. In a standard VAR setting, the likelihood quickly dominates the prior for the covariance matrix, such that inference likely remains unaffected. In our setting, it is unclear how much information the data contains about the quantile factors. Depending on the exercise inference might hence suffer. To sample the coefficients in A, we impose a normal prior.

$$\boldsymbol{a}_r \sim N(\boldsymbol{\mu}_{r,a}, \boldsymbol{\Sigma}_{r,a}) \tag{26}$$

with r denoting the columns of A. What remains is to define the prior for $h_{r,t}$. In the case without stochastic volatility, we set

$$\log(h_{r,t}) = \log(h_r) = G^{-1}(r_h, s_h) \tag{27}$$

with stochastic volatility, we introduce the state equation

$$\log(h_{r,t}) = \log(h_{r,t-1}) + v_t, \quad v_t \sim N(0, \sigma_\omega^2) \tag{28}$$

with $h_{r,1} \sim N(0, V_0)$ and prior $\sigma_\omega^2 \sim IG(r_\omega, s_\omega)$.

A.3.2 Estimation Algorithm

Conditional on all parameter matrices being known (e.g. $\boldsymbol{\Lambda}$, $\boldsymbol{S}$, $\boldsymbol{\Phi}$), sampling of the state vector $[\boldsymbol{F}_t', \boldsymbol{g}_t']'$ can be obtained using a simulation smoother such as the one proposed by Carter and Kohn (1994). Conditional on these states, parameters can be obtained by using standard methodologies for linear and quantile regression models. First, we need to obtain the parameters in the measurement equation. As mentioned above, the univariate asymmetric Laplace distributions of the individual measurement equations imply a diagonal covariance matrix. Conditioning on the factors $f_{t(q)}^i$ and treating these as observed, the individual measurement equations for y_{ijt} are independent. They can hence be treated as $m \times n \times q$ individual univariate quantile regressions. We can thus obtain samples from all parameters of equation (1) using standard formulas provided in Khare and Hobert (2012). Next, we obtain samples from the parameters of the state equation using standard formulas for Bayesian vector autoregressions. Indeed, once we condition on $[\boldsymbol{F}_t', \boldsymbol{g}_t']'$, equation (7) becomes a VAR and the conditional posteriors of $\boldsymbol{v}$, $\boldsymbol{\Phi}$, $\boldsymbol{\Omega}$ can be derived in a straightforward way (see Koop and Korobilis, 2010, for a thorough examination of Bayesian inference in VARs).

All sampling steps of the MCMC algorithm are provided in algorithm 1.

Algorithm 1: MCMC algorithm for the estimation of the QFAVAR

begin

  Define $\kappa_{1(q)} = \frac{1-2q}{q(1-q)}$ and $\kappa_{2(q)}^2 = \frac{2}{q(1-q)}$ for each quantile level q.

  [**1**] Sample the parameters in the measurement equation leveraging results from Bayesian quantile regression:

  for $q = 1 : nq$ **do**

   for $i = 1 : m$ **do**

    for $j = 1 : n$ **do**

     [i] Sample $\boldsymbol{\phi}_{ij(q)}$ from the full conditional $(\boldsymbol{\phi}_{ij(q)}|-) \sim N(\boldsymbol{\mu}_{ij(q)}^\phi, \boldsymbol{\Sigma}_{ij(q)}^\phi)$, with

     $$\boldsymbol{\mu}_{ij(q)}^\phi = \boldsymbol{\Sigma}_{ij(q)}^\phi \left\{ \hat{\boldsymbol{F}}' diag(\tilde{\boldsymbol{\nu}}_{ij(q)}^{-1}) \tilde{\boldsymbol{y}}_{ij(q)} \right\}, \boldsymbol{\Sigma}_{ij(q)}^\phi = \left\{ \hat{\boldsymbol{F}}' diag(\tilde{\boldsymbol{\nu}}_{ij(q)})^{-1} \hat{\boldsymbol{F}} + \tilde{\boldsymbol{\Lambda}}_{ij(q)}^{-1} \right\}^{-1},$$

     where $\tilde{\boldsymbol{\nu}}_{ij(q)} = \sigma_{ij(q)} \kappa_{2(q)}^2 \boldsymbol{z}_{ij,\bullet(q)}$, $\tilde{\boldsymbol{y}}_{ij(q)} = \boldsymbol{y}_{ij,\bullet(q)} - \kappa_{1(q)} \boldsymbol{z}_{ij,\bullet(q)}$, and
     $\tilde{\boldsymbol{\Lambda}}_{ij(q)}^{-1} = diag(\bar{\boldsymbol{\lambda}}_{ij(q),\bullet,\phi}^2 \tau_{ij(q),\phi}^2)^{-1}$.

     [ii] Sample the parameters corresponding to the horseshoe prior from

     $$\left(\bar{\lambda}_{ij(q),k,\phi}^2|-\right) \sim IG\left(1, \frac{\phi_{ij,k(q)}^2}{2\tau_{ij(q),\phi}^2} + \frac{1}{v_{ij(q),k,\phi}}\right), \left(v_{ij(q),k,\phi}|-\right) \sim IG\left(1, 1 + 1/(\bar{\lambda}_{ij(q),k,\phi}^2)\right),$$

     $$\left(\tau_{ij(q),\phi}^2|-\right) \sim IG\left(\frac{l_\phi+1}{2}, \frac{1}{\xi_{ij(q),\phi}} + \sum_k^{l_\phi} \frac{\phi_{ij(q)}^2}{2\bar{\lambda}_{ij(q),k,\phi}^2}\right), \left(\xi_{ij(q),\phi}|-\right) \sim IG\left(1, 1 + \frac{1}{\tau_{ij(q),\phi}^2}\right),$$

     for $k = 1, \cdots, l_\phi$.

     [iii] Sample the latent quantile variables from $\left(z_{ij,t(q)}|-\right) \sim GIG\left(\chi_z, \psi_z\right)$, with

     $$\chi_z = \frac{\left\{y_{ij,t(q)} - \hat{\boldsymbol{\phi}}_{ij(q)}' \boldsymbol{F}_t\right\}^2}{\kappa_{1(q)}^2 + 2\kappa_{2(q)}^2}, \psi_z = \frac{\kappa_{1(q)}^2 + 2\kappa_{2(q)}^2}{\sigma_{ij(q)} \kappa_{2(q)}^2}$$

     for all $t = 1, \cdots, T$.

     [iv] Sample the factor regression variances $\left(\sigma_{ij(q)}|-\right) \sim IG\left(r_\sigma, s_\sigma\right)$, where $r_\sigma = r_0 + 3T/2$ and

     $$s_\sigma = s_0 + \sum_{t=1}^T \frac{\left\{y_{ij,t(q)} - \hat{\boldsymbol{\phi}}_{ij(q)}' \boldsymbol{F}_t - \kappa_{1(q)} z_{ij,t(q)}\right\}^2}{2 z_{ij,t(q)} \kappa_{2(q)}^2}$$

    end

   end

  end

  [2] Sample the quantile factors F_t. Define $\tilde{\boldsymbol{y}}_{ij(q)} = \boldsymbol{y}_{ij,\bullet(q)} - \kappa_{1(q)} \boldsymbol{z}_{ij,\bullet(q)}$ and stack the loadings according to 6. Sample F_t using the Kalman filter & smoother.

  [3] Sample the diagonal elements, $h_{r,t}$ of covariance matrix Ω_t in state equation 7.

  if *stochastic volatility* **then**

   set $y_{r,t}^* = log\left\{(\boldsymbol{F}_{r,t} - \boldsymbol{F}_{t-1}\boldsymbol{\Psi}_r')^2\right\}$ and sample $log(h_{r,t})$ using the Chan (2013) filter. And update $\sigma_{t,\omega}^2$

   from $(\sigma_{t,\omega}^2|-) \sim IG\left(r_\omega + 0.5(T - p - 1), s_\omega + 0.5 \cdot \sum_{t=1}^T (h_{r,t} - h_{r,t-1})^2\right)$

   for every $r = 1, \cdots, R$.

  else

   Sample from $(\sigma_{t,\omega}^2 = \sigma_\omega^2|-) \sim IG\left(r_h + 0.5 \cdot (T - p), s_h + 0.5 \cdot \sum_{t=1}^T (\boldsymbol{F}_{r,t} - \boldsymbol{F}_{t-1}\boldsymbol{\Psi}_r')^2\right)$

   for all $r = 1, \cdots, R$.

  end

  [4] Sample the off-diagonal elements of $\boldsymbol{\Omega}_t$, $\boldsymbol{A}$, and the VAR coefficients $\boldsymbol{\Psi}$

  for $r = 1 : R$ **do**

   [i] Sample $\boldsymbol{\psi}_r$ and $\boldsymbol{a}_r$ jointly from the full conditional $(\boldsymbol{\psi}_r, \boldsymbol{a}_r|-) \sim N(\boldsymbol{\mu}_r^\psi, \boldsymbol{\Sigma}_r^\psi)$, with

   $\boldsymbol{\mu}_r^\psi = \boldsymbol{\Sigma}_r^\psi \left\{ \boldsymbol{X}' diag(\sigma_{t,\omega}^2)^{-1} \hat{\boldsymbol{f}}_r \right\}, \boldsymbol{\Sigma}_r^\psi = \left\{ \boldsymbol{X}' diag(\sigma_{t,\omega}^2)^{-1} \boldsymbol{X} + \tilde{\boldsymbol{\Lambda}}_r^{-1} \right\}^{-1}$, where

   $\boldsymbol{X} = [L\boldsymbol{F}', L\boldsymbol{G}', \boldsymbol{E}_{1:r-1}']'$, $\boldsymbol{e}_r = \hat{\boldsymbol{f}}_r - \hat{\boldsymbol{F}}\boldsymbol{\psi}_r' - \boldsymbol{E}_{1:r-1}'\boldsymbol{a}_r'$, and $\tilde{\boldsymbol{\Lambda}}_r^{-1} = diag\left([(\bar{\boldsymbol{\lambda}}_{r,\bullet,\psi}^2,' \tau_{r,\psi}^2)^{-1}, \boldsymbol{\mu}_{r,a}/\boldsymbol{\Sigma}_{r,a}]\right)$.

   [ii] Sample the parameters of the horseshoe prior for ψ_r from $\left(\bar{\lambda}_{r,s,\psi}^2|-\right) \sim IG\left(1, \frac{\psi_{r,s}^2}{2\tau_{r,\psi}^2} + \frac{1}{v_{r,s,\psi}}\right)$,

   $\left(v_{r,s,\psi}|-\right) \sim IG\left(1, 1 + 1/(\bar{\lambda}_{r,s,\psi}^2)\right), \left(\tau_{r,\psi}^2|-\right) \sim IG\left(\frac{l_\psi+1}{2}, \frac{1}{\xi_{r,\psi}} + \sum_k^{l_\psi} \frac{\psi_r^2}{2\bar{\lambda}_{r,s,\psi}^2}\right),$

   $\left(\xi_{r,\psi}|-\right) \sim IG\left(1, 1 + \frac{1}{\tau_{r,\psi}^2}\right),$

   for $s = 1, \cdots, l_\psi$.

  end

end

Note: We use the rate parametrization of the IG distribution throughout. We use L to denote the lag operator.

A.4 Variational Bayes estimation

We also propose a 2-step variational QFAVAR algorithm which is appropriate for forecasting or other computationally cumbersome applications of the model. For a general introduction to variational Bayes (VB) see Blei et al. (2017). Generally, variational inference and MCMC both provide approximations to a given posterior distribution. While MCMC provides approximations through sampling, VB approximates the objective by solving an optimization problem.While in practice both algorithms end up being iterative, VB is computationally less intensive than MCMC due to the fact that it requires significantly less iterations. As a caveat, while MCMC approximates the full posterior distribution and its uncertainty, VB only provides the posterior mean and (a possibly biased estimate of the) posterior variance. Therefore, VB is particularly useful for tasks where precise inference is less of a concern, such as big data applications, large scale forecasting exercises, or real-time monitoring tasks.

Our algorithm evolves as follows. We first extract the quantile factors, F_t, using the VBQFA algorithm proposed in Korobilis and Schröder (2022). These are the same factors that we use to identify the sign of the factors in our proposed MCMC algorithm. Conditional on these factors, we then update the loadings and the remaining parameters in measurement equation 1. Finally, we update the parameters of the VAR in state equation 7 and generate forecasts. Treating the quantile factors estimated in the first stage as observed, the model parameters and latent variables are given by $\boldsymbol{\theta} = (\boldsymbol{\Phi}, \boldsymbol{\sigma}, \boldsymbol{z}, \boldsymbol{\Psi}, \boldsymbol{\Omega}, \boldsymbol{\tau}_\phi^2, \boldsymbol{\xi}_\phi, \bar{\boldsymbol{\lambda}}_\phi^2, \boldsymbol{v}_\phi, \boldsymbol{\tau}_\psi^2, \boldsymbol{\xi}_\psi, \bar{\boldsymbol{\lambda}}_\psi^2, \boldsymbol{v}_\psi)$. For a family of tractable densities $q(\boldsymbol{\theta})$, we aim to find a density $q^\star$ that best approximates the posterior $p(\boldsymbol{\theta}|\boldsymbol{x})$ by minimizing

$$q^\star(\boldsymbol{\theta}|\boldsymbol{x}) = \underset{q \in \mathscr{Q}}{\operatorname{argmin}} \, \mathbb{D}_{KL}\left(q(\boldsymbol{\theta}||p(\boldsymbol{\theta}|\boldsymbol{x})\right), \tag{29}$$

which is equivalent to maximizing

$$ELBO = \mathbb{E}_{q(\boldsymbol{\theta}|\boldsymbol{x})}\left[\log p(\boldsymbol{x}|\boldsymbol{\theta})\right] + \mathbb{E}_{q(\boldsymbol{\theta}|\boldsymbol{x})}\left[\log p(\boldsymbol{\theta})\right] - \mathbb{E}_{q(\boldsymbol{\theta}|\boldsymbol{x})}\left[\log q(\boldsymbol{\theta}|\boldsymbol{x})\right], \tag{30}$$

where KL denotes the Kullback-Leibler divergence. Note that we need to optimize over a family of distribution functions. Finding the solution to the problem hence requires the application of variational calculus. Usually, optimization can be simplified by factorizing the variational posterior into L groups of independent densities. In our proposed two-step estimator, we apply the following factorization

$$q(\boldsymbol{\theta}|\boldsymbol{x}) \equiv q\left(\boldsymbol{\Phi}, \boldsymbol{\sigma}, \boldsymbol{z}, \boldsymbol{\Psi}, \boldsymbol{\Omega}, \boldsymbol{\tau}_\phi^2, \boldsymbol{\xi}_\phi, \bar{\boldsymbol{\lambda}}_\phi^2, \boldsymbol{\upsilon}_\phi, \boldsymbol{\tau}_\psi^2, \boldsymbol{\xi}_\psi, \bar{\boldsymbol{\lambda}}_\psi^2, \boldsymbol{\upsilon}_\psi\right) \tag{31}$$

$$= \prod_{q=1}^{nq} \prod_{j=1}^{m} \prod_{i=1}^{n} \left[q(\boldsymbol{\phi}_{ij(q)}|\boldsymbol{x})q(\sigma_{ij(q)}|\boldsymbol{x})q(\tau_{ij(q),\phi}^2|\boldsymbol{x})q(\xi_{ij(q),\phi}|\boldsymbol{x}) \prod_{k=1}^{l_\phi} q(\bar{\lambda}_{ij(q),k,\phi}^2|\boldsymbol{x})q(\upsilon_{ij(q),k),\phi}|\boldsymbol{x}) \prod_{t=1}^{T} q(z_{ij,t(q)}|\boldsymbol{x}) \right]$$

$$\cdot \prod_{r=1}^{R} \left[q(\boldsymbol{\psi}_r, \boldsymbol{a}_r|\boldsymbol{x})q(\tau_{r,\psi}^2|\boldsymbol{x})q(\xi_{r,\psi}|\boldsymbol{x})q(\omega_r|\boldsymbol{x}) \prod_{s=1}^{l_\psi} q(\bar{\lambda}_{r,s,\psi}^2|\boldsymbol{x})q(\upsilon_{r,s,\psi}|\boldsymbol{x}) \right],$$

which implies partial posterior independence between the regression parameters and residual variances in the individual measurement equations, the VAR parameters and covariance matrix elements in the state equations, as well as the hyperparameters of the horseshoe priors. With the additional assumption that the priors are conditionally independent, we can write the joint prior as

$$p(\boldsymbol{\theta}) = p\left(\boldsymbol{\Phi}, \boldsymbol{\sigma}, \boldsymbol{z}, \boldsymbol{\Psi}, \boldsymbol{\Omega}, \boldsymbol{\tau}_\phi^2, \boldsymbol{\xi}_\phi, \bar{\boldsymbol{\lambda}}_\phi^2, \boldsymbol{\upsilon}_\phi, \boldsymbol{\tau}_\psi^2, \boldsymbol{\xi}_\psi, \bar{\boldsymbol{\lambda}}_\psi^2, \boldsymbol{\upsilon}_\psi\right) \tag{32}$$

$$= \prod_{q=1}^{nq} \prod_{j=1}^{m} \prod_{i=1}^{n} \left[p(\boldsymbol{\phi}_{ij(q)}|\bar{\boldsymbol{\lambda}}_{ij(q),\phi}^2, \tau_{ij(q),\phi}^2, \sigma_{ij(q)})p(\sigma_{ij(q)})p(\tau_{ij(q),\phi}^2|\xi_{ij(q),\phi})p(\xi_{ij(q),\phi}) \right.$$

$$\left. \cdot p(\boldsymbol{z}_{ij(q)}) \prod_{k=1}^{l_\phi} p(\bar{\lambda}_{ij(q),k,\phi}^2|\upsilon_{ij(q),k,\phi})p(\upsilon_{ij(q),k,\phi}) \right]$$

$$\cdot \prod_{r=1}^{R} \left[p(\boldsymbol{\psi}_r|\bar{\boldsymbol{\lambda}}_{r,\psi}^2, \tau_{r,\psi}^2, \omega_r)p(\omega_r)p(\boldsymbol{a}_r)p(\tau_{r,\psi}^2|\xi_{r,\psi})p(\xi_{r,\psi}) \prod_{s=1}^{l_\psi} p(\bar{\lambda}_{r,s,\psi}^2|\upsilon_{r,s,\psi})p(\upsilon_{r,s,\psi}) \right].$$

Given this partitioning, it can be shown that the solution to the optimization problem can be obtained by sequentially iterating over the densities

$$q(\boldsymbol{\theta}_l|\boldsymbol{x}) \propto \exp \mathbb{E}_{q(\boldsymbol{\theta}_{(-l)}|\boldsymbol{x})}\left(\log p(\boldsymbol{\theta}_l|\boldsymbol{\theta}_{(-l)}, \boldsymbol{x})\right) \tag{33}$$

where $\boldsymbol{\theta}_{(-l)}$ denotes all elements of $\boldsymbol{\theta}$, excluding those in the l^{th} group, $l = 1, \cdots, L$. Consequently, the variational posterior can be obtained by calculating the variational expectation of the conditional posterior. Generally, the accuracy of the variational approximation depends on how well the chosen partitioning matches the independence structure of the parameters in the target posterior. In essence, this is where we trade off computational tractability and speed with accuracy. For a general discussion of this issue see e.g. Ormerod et al. (2017). In order to arrive at the final variational densities, we now need to insert 31 and 32 into 33. Given that this expression relies on expectations of logarithms of standard densities it is rather easy to derive, however, we suppress it here due to its very lengthy nature. The expressions for $q(\boldsymbol{\theta}_l|\boldsymbol{x})$ are presented in algorithm 2 together with the full variational algorithm.

Algorithm 2: 2-step Variational Bayes QFAVAR algorithm

Define $\kappa_{1(q)} = \frac{1-2q}{q(1-q)}$ and $\kappa_{2(q)}^2 = \frac{2}{q(1-q)}$ for each quantile level q.

Step 1: Estimate the factors $\hat{F}$ using the variational VBQFA algorithm proposed in Korobilis and Schröder (2022) for each quantile level, q, and for the variable blockings in 4.

Step 2: Conditional on the factors, estimate the remaining parameters of the model.

begin

[**1**] Update the parameters in the measurement equation:

for $q = 1 : nq$ **do**

    for $i = 1 : m$ **do**

        for $j = 1 : n$ **do**

            [i] Update $\phi_{ij(q)}$ from $q(\phi_{ij(q)}|\boldsymbol{x}) = N(\boldsymbol{\mu}_{ij(q)}^{\phi}, \boldsymbol{\Sigma}_{ij(q)}^{\phi})$, with $\boldsymbol{\mu}_{ij(q)}^{\phi} = \boldsymbol{\Sigma}_{ij(q)}^{\phi}\left\{\hat{\boldsymbol{F}}' diag(\tilde{\boldsymbol{\nu}}_{ij(q)}^{-1})\tilde{\boldsymbol{y}}_{ij(q)}\right\}$,

            $\boldsymbol{\Sigma}_{ij(q)}^{\phi} = \left\{\hat{\boldsymbol{F}}' diag(\tilde{\boldsymbol{\nu}}_{ij(q)})^{-1}\hat{\boldsymbol{F}} + \tilde{\boldsymbol{\Lambda}}_{ij(q)}^{-1}\right\}^{-1}$, where $\tilde{\boldsymbol{\nu}}_{ij(q)}^{-1} = \kappa_{2(q)}^{-2}\mathbb{E}\left(\sigma_{ij(q)}^{-1}\right)\mathbb{E}\left(\boldsymbol{z}_{ij,\bullet(q)}^{-1}\right)$,

            $\tilde{\boldsymbol{y}}_{ij(q)} = \boldsymbol{y}_{ij,\bullet(q)} - \kappa_{1(q)}\mathbb{E}\left(\boldsymbol{z}_{ij,\bullet(q)}\right)$, and $\tilde{\boldsymbol{\Lambda}}_{ij(q)}^{-1} = diag\left(\mathbb{E}\left(\bar{\boldsymbol{\lambda}}_{ij(q),\bullet,\phi}^{-2}\right)\mathbb{E}\left(\tau_{ij(q),\phi}^{-2}\right)\right)$.

            Assign: $\mathbb{E}\left(\phi_{ij(q)}\right) = \boldsymbol{\mu}_{ij(q)}^{\phi}$ and $\mathbb{E}\left(\phi_{ij(q)}^2\right) = \boldsymbol{\mu}_{ij(q)}^{2,\phi} + diag\left(\Sigma_{ij(q)}^{\phi}\right)$

            [ii] Update the parameters of the horseshoe prior from

            $q(\bar{\lambda}_{ij(q),k,\phi}^2|\boldsymbol{x}) = IG(a_{\bar{\lambda},\phi}, b_{\bar{\lambda},\phi}) = IG\left(1, \frac{\mathbb{E}(\phi_{ij,k(q)}^2)}{2} + \mathbb{E}\left(\frac{1}{v_{ij(q),k,\phi}}\right)\right)$,

            $q(v_{ij(q),k,\phi}|\boldsymbol{x}) = IG(a_{v,\phi}, b_{v,\phi}) = IG\left(1, \mathbb{E}\left(\frac{1}{\bar{\lambda}_{ij(q),k,\phi}^2}\right) + b_{\phi}^{-2}\mathbb{E}\left(\frac{1}{\tau_{ij,\phi}^2}\right)\right)$,

            $q(\tau_{ij(q),k,\phi}^2|\boldsymbol{x}) = IG(a_{\tau,\phi}, b_{\tau,\phi}) = IG\left(1, b_{\phi}^{-2}\mathbb{E}\left(\frac{1}{v_{ij(q),k,\phi}}\right) + \mathbb{E}\left(\frac{1}{\xi_{ij(q),\phi}}\right)\right)$,

            $q(\xi_{ij(q),\phi}|\boldsymbol{x}) = IG(a_{\xi,\phi}, b_{\xi,\phi}) = IG\left(\frac{l_{\phi}+1}{2}, 1 + \sum_{k=1}^{l_{\phi}}\mathbb{E}\left(\frac{1}{\tau_{ij(q),k,\phi}^2}\right)\right)$,

            for $k = 1, \cdots, l_{\phi}$.

            Assign: $\mathbb{E}\left(\frac{a_{\bar{\lambda},\phi}}{b_{\bar{\lambda},\phi}}\right)$, $\mathbb{E}\left(\frac{a_{v,\phi}}{b_{v,\phi}}\right)$, $\mathbb{E}\left(\frac{a_{\tau,\phi}}{b_{\tau,\phi}}\right)$, and $\mathbb{E}\left(\frac{a_{\xi,\phi}}{b_{\xi,\phi}}\right)$.

            [iii] Update the latent quantile indicators, $\boldsymbol{z}_{ij,t(q)}$, from $q(\boldsymbol{z}_{ij,t(q)}|\boldsymbol{x}) = GIG(1/2, \delta_z, \rho_{t,z})$ with

            $\delta_z = \mathbb{E}\left(\frac{1}{\sigma_{ij(q)}}\right)\frac{\kappa_{1(q)}^2}{\kappa_{2(q)}^2}$, $\rho_{t,z} = \mathbb{E}\left(\frac{1}{\sigma_{ij(q)}}\right)\frac{\left\{y_{ij,t(q)} - \mathbb{E}(\phi_{ij(q)})'\hat{\boldsymbol{f}}_t\right\}^2 + \hat{\boldsymbol{f}}_t'\boldsymbol{\Sigma}_{ij(q)}^{\phi}\hat{\boldsymbol{f}}_t}{\kappa_{2(q)}^2}$,

            for $t = 1, \cdots, T$

            Assign: $\mathbb{E}(z_{ij,t(q)}) = \frac{\sqrt{\rho_{t,z}}K_{3/2}\left(\sqrt{\delta_z\rho_{t,z}}\right)}{\sqrt{\delta_z}K_{1/2}\left(\sqrt{\delta_z\rho_{t,z}}\right)}$, $\mathbb{E}\left(\frac{1}{z_{ij,t(q)}}\right) = \frac{\sqrt{\delta_z}K_{3/2}\left(\sqrt{\delta_z\rho_{t,z}}\right)}{\sqrt{\rho_{t,z}}K_{1/2}\left(\sqrt{\delta_z\rho_{t,z}}\right)} - \frac{1}{\rho_{t,z}}$

            [iv] Update the factor regression variances, $\sigma_{ij(q)}$, from $q(\sigma_{ij(q)}|\boldsymbol{x}) = IG(r_{\sigma}, s_{\sigma})$ with

            $r_{\sigma} = r_0 + 3T$ and

            $s_{\sigma} = s_0 + \sum_{t=1}^{T}\left[\mathbb{E}\left(\frac{1}{z_{ij,t(q)}}\right)\frac{M}{2\kappa_{2(q)}^2} - \kappa_{1(q)}^2\frac{y_{ij,t(q)} - \mathbb{E}(\phi_{ij(q)})'\hat{\boldsymbol{f}}_t}{\kappa_{2(q)}^2} + \left(1 + \frac{\kappa_{1(q)}^2}{2\kappa_{2(q)}^2}\right)\mathbb{E}(z_{ij,t(q)})\right]$ and

            $M = \left\{y_{ij,t(q)} - \mathbb{E}(\phi_{ij(q)})'\hat{\boldsymbol{f}}_t\right\}^2 + \hat{\boldsymbol{f}}_t'\boldsymbol{\Sigma}_{ij(q)}^{\phi}\hat{\boldsymbol{f}}_t$.

            Assign: $\mathbb{E}\left(\frac{1}{\sigma_{ij(q)}}\right) = \frac{r_{\sigma}}{s_{\sigma}}$.

        end

    end

end

[**2**] Update the diagonal elements of Ω from $q(\omega_r|\boldsymbol{x}) = IG(r_{\omega}, s_{\omega})$ with

$r_{\omega} = r_h + T/2$ and $s_{\omega} = s_h + \sum_{t=1}^{T}(\hat{\boldsymbol{f}}_{r,t} - \hat{\boldsymbol{F}}_{t-1}\boldsymbol{\psi}_r')^2)$, for all $r = 1, \cdots, R$.

Assign: $\mathbb{E}\left(\frac{1}{\omega_r}\right) = \frac{r_{\omega}}{s_{\omega}}$.

[3] Sample the off-diagonal elements of $\boldsymbol{\Omega}$, $\boldsymbol{A}$, and the VAR coefficients $\boldsymbol{\Psi}$

for $r = 1 : R$ **do**

    [i] Sample $\boldsymbol{\psi}_r$ and $\boldsymbol{a}_r$ jointly from $q(\boldsymbol{\psi}_r, \boldsymbol{a}_r|\boldsymbol{x}) = N(\boldsymbol{\mu}_r^{\psi}, \boldsymbol{\Sigma}_r^{\psi})$, with $\boldsymbol{\mu}_r^{\psi} = \boldsymbol{\Sigma}_r^{\psi}\left\{\boldsymbol{X}'\mathbb{E}(\sigma_{\omega}^{-1})\hat{\boldsymbol{f}}_r\right\}$,

    $\boldsymbol{\Sigma}_r^{\psi} = \left\{\boldsymbol{X}'\mathbb{E}(\sigma_{\omega}^{-1})\boldsymbol{X} + \tilde{\boldsymbol{\Lambda}}_r^{-1}\right\}^{-1}$, where $\boldsymbol{X} = [L\boldsymbol{F}', L\boldsymbol{G}', \boldsymbol{E}_{1:r-1}']'$, $\boldsymbol{e}_r = \hat{\boldsymbol{f}}_r - \hat{\boldsymbol{F}}\mathbb{E}(\boldsymbol{\psi}_r)' - \boldsymbol{E}_{1:r-1}'\mathbb{E}(\boldsymbol{a}_r)'$, and

    $\tilde{\boldsymbol{\Lambda}}_r^{-1} = diag\left([\mathbb{E}(\bar{\boldsymbol{\lambda}}_{r,\bullet,\psi}^{-2}), \boldsymbol{\mu}_{r,a}/\boldsymbol{\Sigma}_{r,a}]\right)$.

    [ii] Update the parameters of the horseshoe prior from $q(\bar{\lambda}_{r,s,\psi}^2|\boldsymbol{x}) = IG(a_{\bar{\lambda},\psi}, b_{\bar{\lambda},\psi}) = IG\left(1, \frac{\mathbb{E}(\boldsymbol{\psi}_{r,s}^2)}{2} + \frac{1}{\mathbb{E}(v_{r,s,\psi})}\right)$,

    $q(v_{r,s,\psi}|\boldsymbol{x}) = IG(a_{v,\psi}, b_{v,\psi}) = IG\left(1, \mathbb{E}\left(\frac{1}{\bar{\lambda}_{r,s,\psi}^2}\right) + b_{\psi}^{-2}\mathbb{E}\left(\frac{1}{\tau_{r,\psi}^2}\right)\right)$,

    $q(\tau_{r,s,\psi}^2|\boldsymbol{x}) = IG(a_{\tau,\psi}, b_{\tau,\psi}) = IG\left(1, b_{\psi}^{-2}\mathbb{E}\left(\frac{1}{v_{r,s,\psi}}\right) + \mathbb{E}\left(\frac{1}{\xi_{r,\psi}}\right)\right)$,

    $q(\xi_{r,\psi}|\boldsymbol{x}) = IG(a_{\xi,\psi}, b_{\xi,\psi}) = IG\left(\frac{l_{\psi}+1}{2}, 1 + \sum_{s=1}^{l_{\psi}}\mathbb{E}\left(\frac{1}{\tau_{r,s,\psi}^2}\right)\right)$,

    for $s = 1, \cdots, l_{\psi}$.

    Assign: $\mathbb{E}\left(\frac{a_{\bar{\lambda},\psi}}{b_{\bar{\lambda},\psi}}\right)$, $\mathbb{E}\left(\frac{a_{v,\psi}}{b_{v,\psi}}\right)$, $\mathbb{E}\left(\frac{a_{\tau,\psi}}{b_{\tau,\psi}}\right)$, and $\mathbb{E}\left(\frac{a_{\xi,\psi}}{b_{\xi,\psi}}\right)$.

end

end

Note: We use the rate parametrization of the IG distribution throughout. $K_p(\bullet)$ denotes the Bessel function of order p and L denotes the Lag operator.

Finally, note that the horseshoe priors take a slightly different form in our VB algorithm with

$$\boldsymbol{\phi}_{ij(q)} \big| \big\{ \bar{\lambda}_{ij(q),k,\phi}, \upsilon_{ij(q),k,\phi} \big\}_{k=1}^{l_\phi}, \tau_{ij(q),\phi}, \xi_{ij(q),\phi} \sim N(\mathbf{0}, \bar{\boldsymbol{\Lambda}}_{ij(q),\phi}), \tag{34}$$

$$\bar{\lambda}_{ij(q),k,\phi}^2 \big| \upsilon_{ij(q),k,\phi} \sim G^{-1}\left(\frac{1}{2}, \frac{1}{\upsilon_{ij(q),k,\phi}}\right), \quad \text{for } k = 1, \cdots, l_\phi,$$

$$\upsilon_{ij(q),k,\phi} \big| \tau_{ij(q),\phi}^2 \sim G^{-1}\left(\frac{1}{2}, \frac{1}{b_\phi^2 \tau_{ij(q),\phi}^2}\right), \quad \text{for } k = 1, \cdots, l_\phi,$$

$$\tau_{ij(q),k,\phi}^2 \big| \xi_{ij(q),\phi} \sim G^{-1}\left(\frac{1}{2}, \frac{1}{\xi_{ij(q),\phi}}\right),$$

$$\xi_{ij(q),\phi} \sim G^{-1}\left(\frac{1}{2}, 1\right),$$

and

$$\boldsymbol{\psi}_r \big| \big\{ \bar{\lambda}_{r,s,\psi}, \upsilon_{r,s,\psi} \big\}_{s=1}^{l_\psi}, \tau_{r,\psi}, \xi_{r,\psi} \sim N(\mathbf{0}, \bar{\boldsymbol{\Lambda}}_{r,\psi}), \tag{35}$$

$$\bar{\lambda}_{r,s,\psi}^2 \big| \upsilon_{r,s,\psi} \sim G^{-1}\left(\frac{1}{2}, \frac{1}{\upsilon_{r,s,\psi}}\right), \quad \text{for } s = 1, \cdots, l_\psi,$$

$$\upsilon_{r,s,\psi} \big| \tau_{r,\psi}^2 \sim G^{-1}\left(\frac{1}{2}, \frac{1}{b_\psi^2 \tau_{r,\psi}^2}\right), \quad \text{for } s = 1, \cdots, l_\psi,$$

$$\tau_{r,s,\psi}^2 \big| \xi_{r,\psi} \sim G^{-1}\left(\frac{1}{2}, \frac{1}{\xi_{r,\psi}}\right),$$

$$\xi_{r,\psi} \sim G^{-1}\left(\frac{1}{2}, 1\right),$$

where b_ϕ and b_ψ are hyperparameters, which we set to 0.0001 in application.

B   Additional empirical results
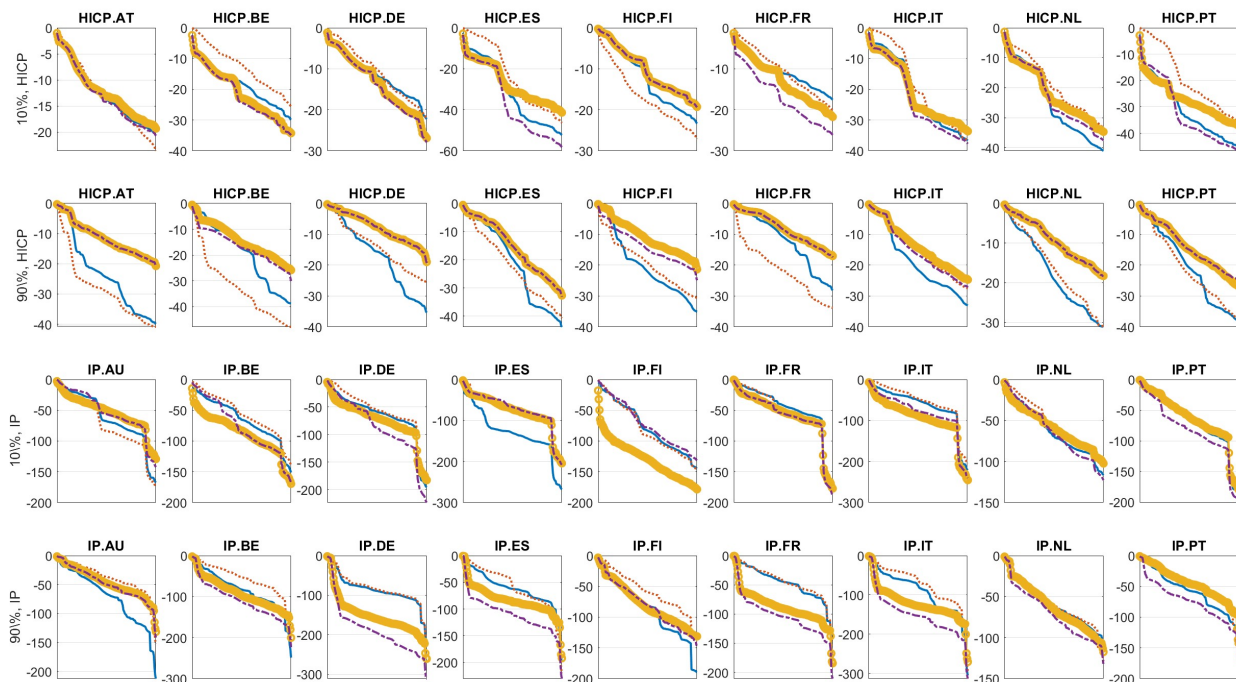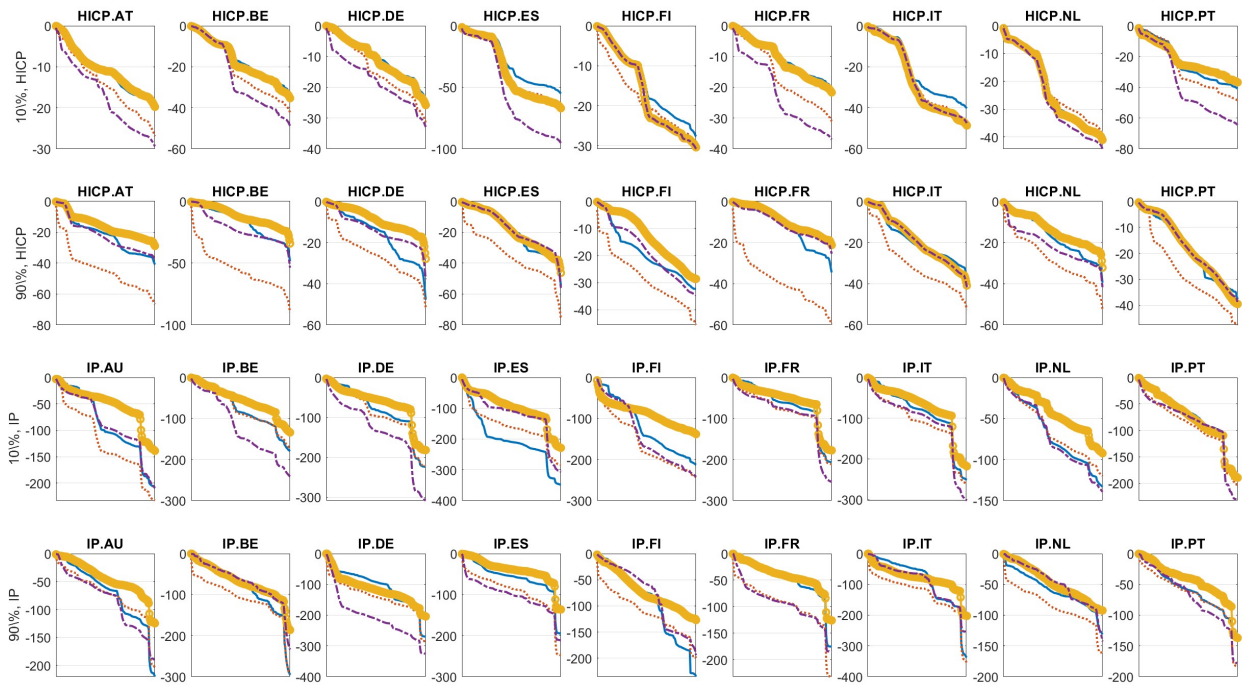
B.1   Forecast evaluation

Figure 1: Cumulative quantile score (QS) loss for forecast horizon $h = 6$. Four models are compared, QAR (yellow circled line), QAR-X (purple dashed line), QDFM (red dotted line) and QFAVAR (blue solid line). The out-of-sample evaluation period shown on the x-axis is 2011Jan to 2022Dec-h. First two rows show quantile scores for the 10th and 90th percentiles of inflation and third and fourth rows show quantile scores for the 10th and 90th percentiles of industrial production.
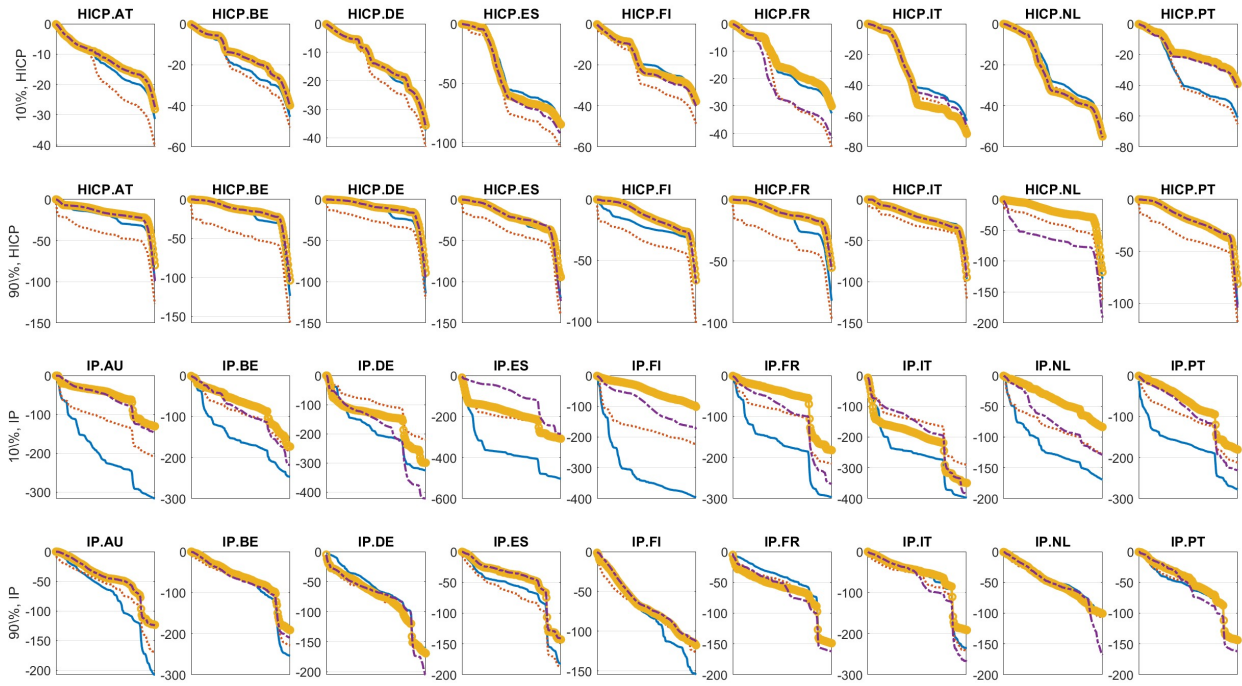
Figure 2: Cumulative quantile score (QS) loss for forecast horizon $h = 12$. Four models are compared, QAR (yellow circled line), QAR-X (purple dashed line), QDFM (red dotted line) and QFAVAR (blue solid line). The out-of-sample evaluation period shown on the x-axis is 2011Jan to 2022Dec-h. First two rows show quantile scores for the 10th and 90th percentiles of inflation and third and fourth rows show quantile scores for the 10th and 90th percentiles of industrial production.

Figure 3: Cumulative quantile score (QS) loss for forecast horizon $h = 24$. Four models are compared, QAR (yellow circled line), QAR-X (purple dashed line), QDFM (red dotted line) and QFAVAR (blue solid line). The out-of-sample evaluation period shown on the x-axis is 2011Jan to 2022Dec-h. First two rows show quantile scores for the 10th and 90th percentiles of inflation and third and fourth rows show quantile scores for the 10th and 90th percentiles of industrial production.

B.2 Detailed quantile IRFs
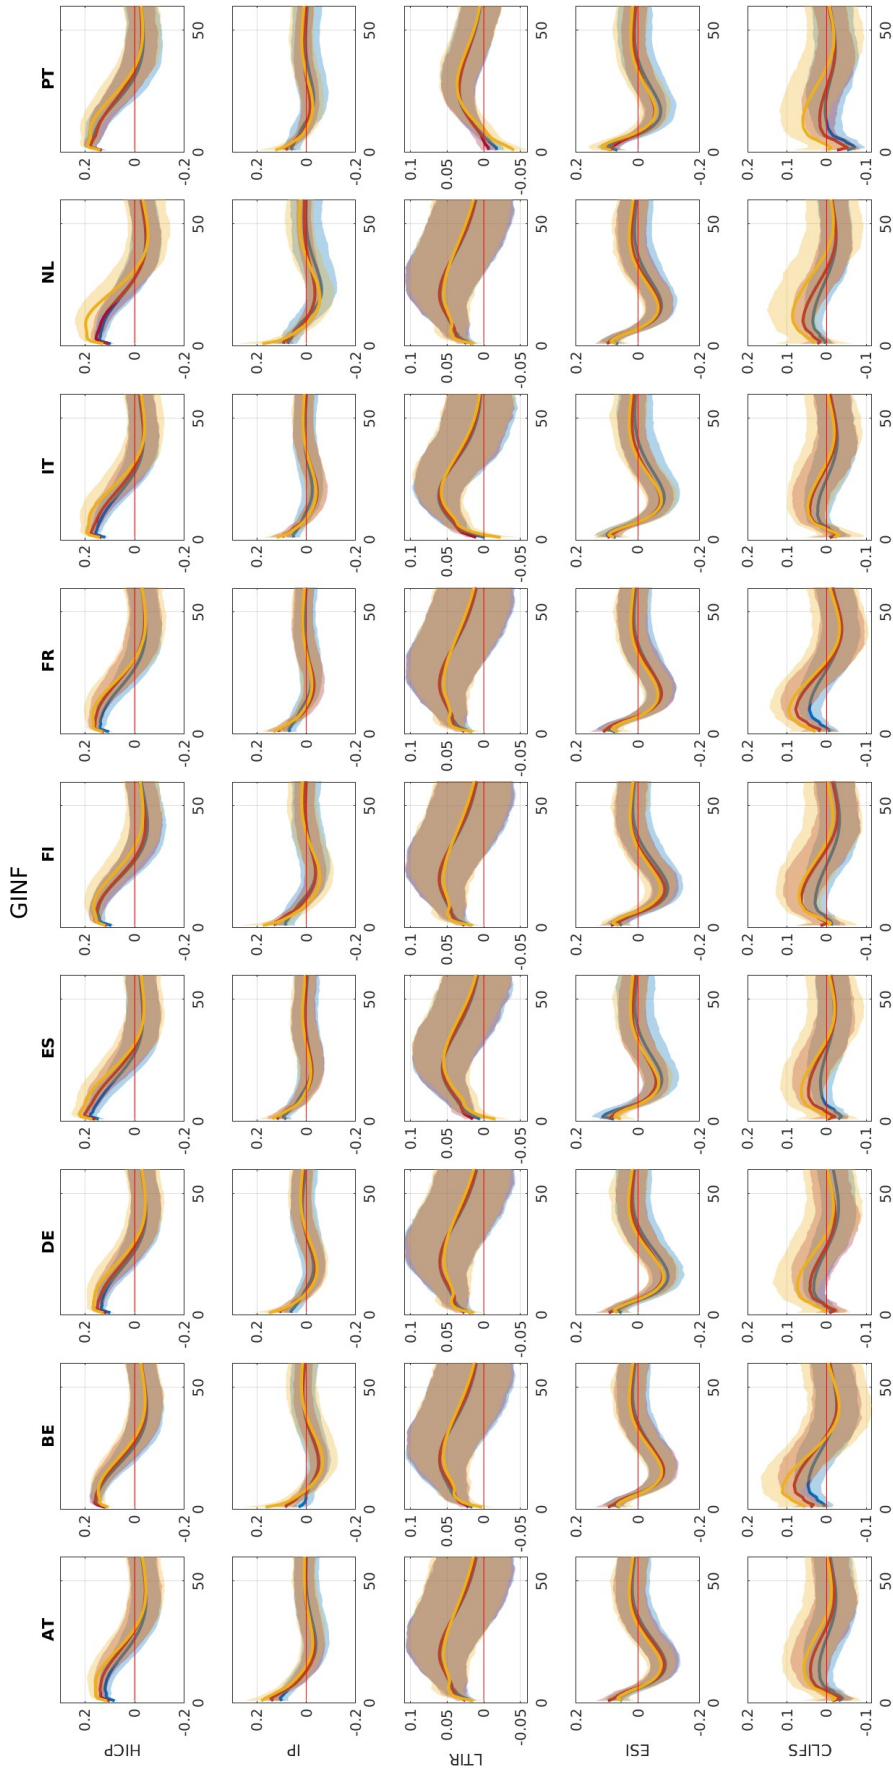
Figure 4: Detailed quantile impulse response functions, shock to global inflation (GINF).
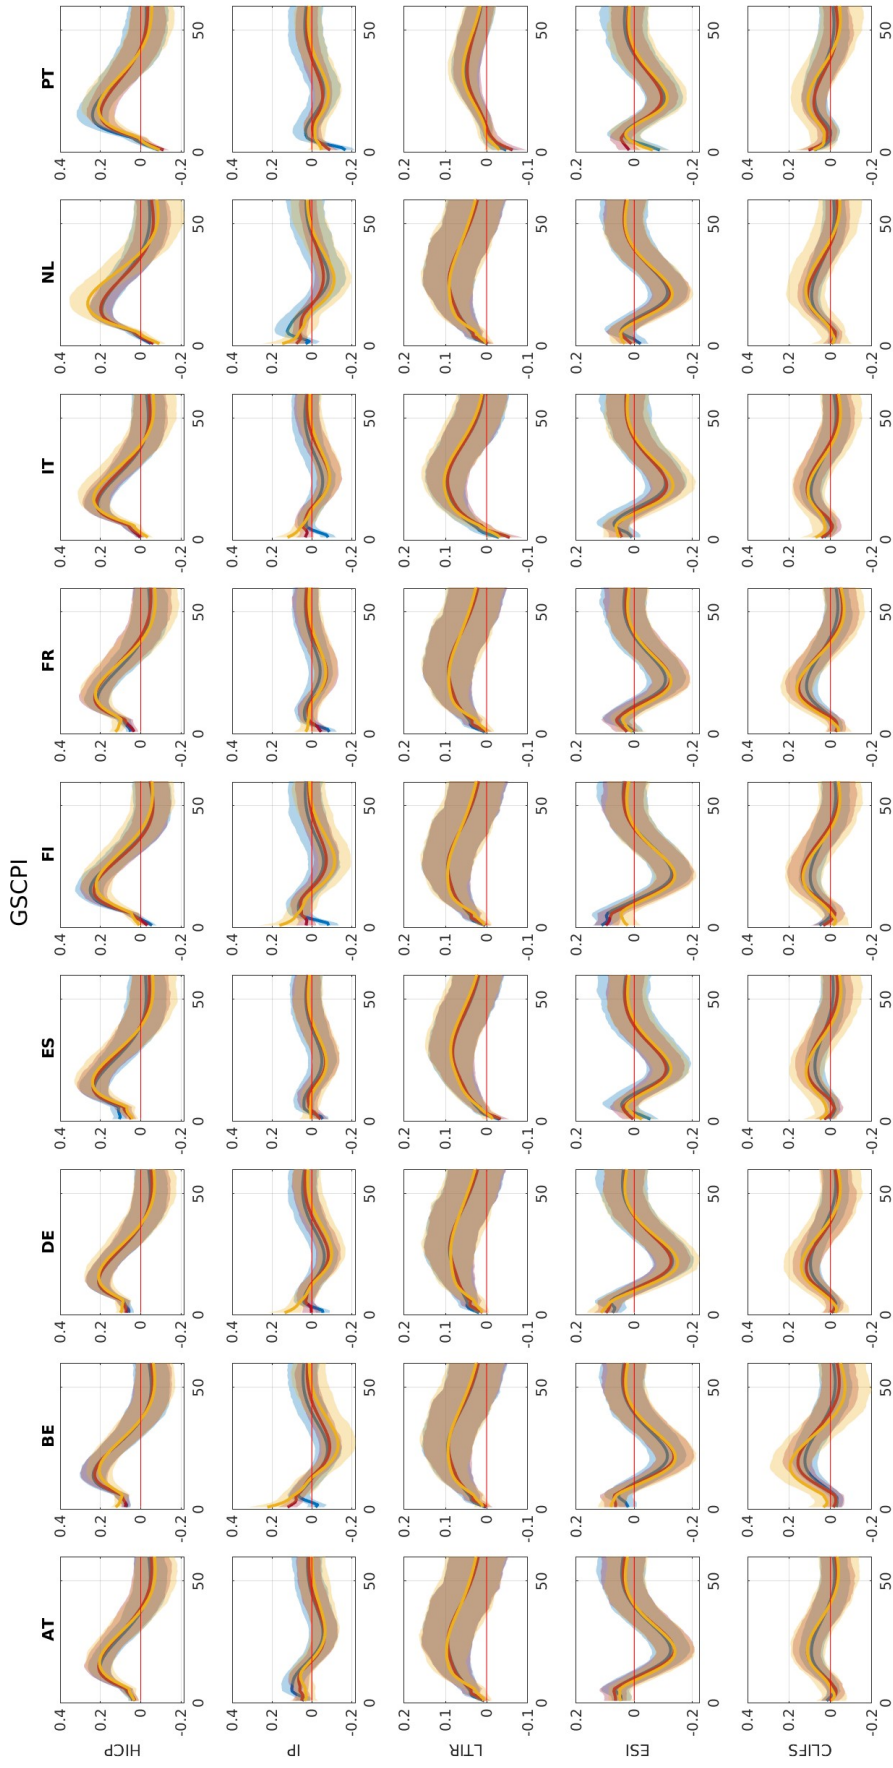
Figure 5: Detailed quantile impulse response functions, shock to global supply chain pressure index (GSCPI).
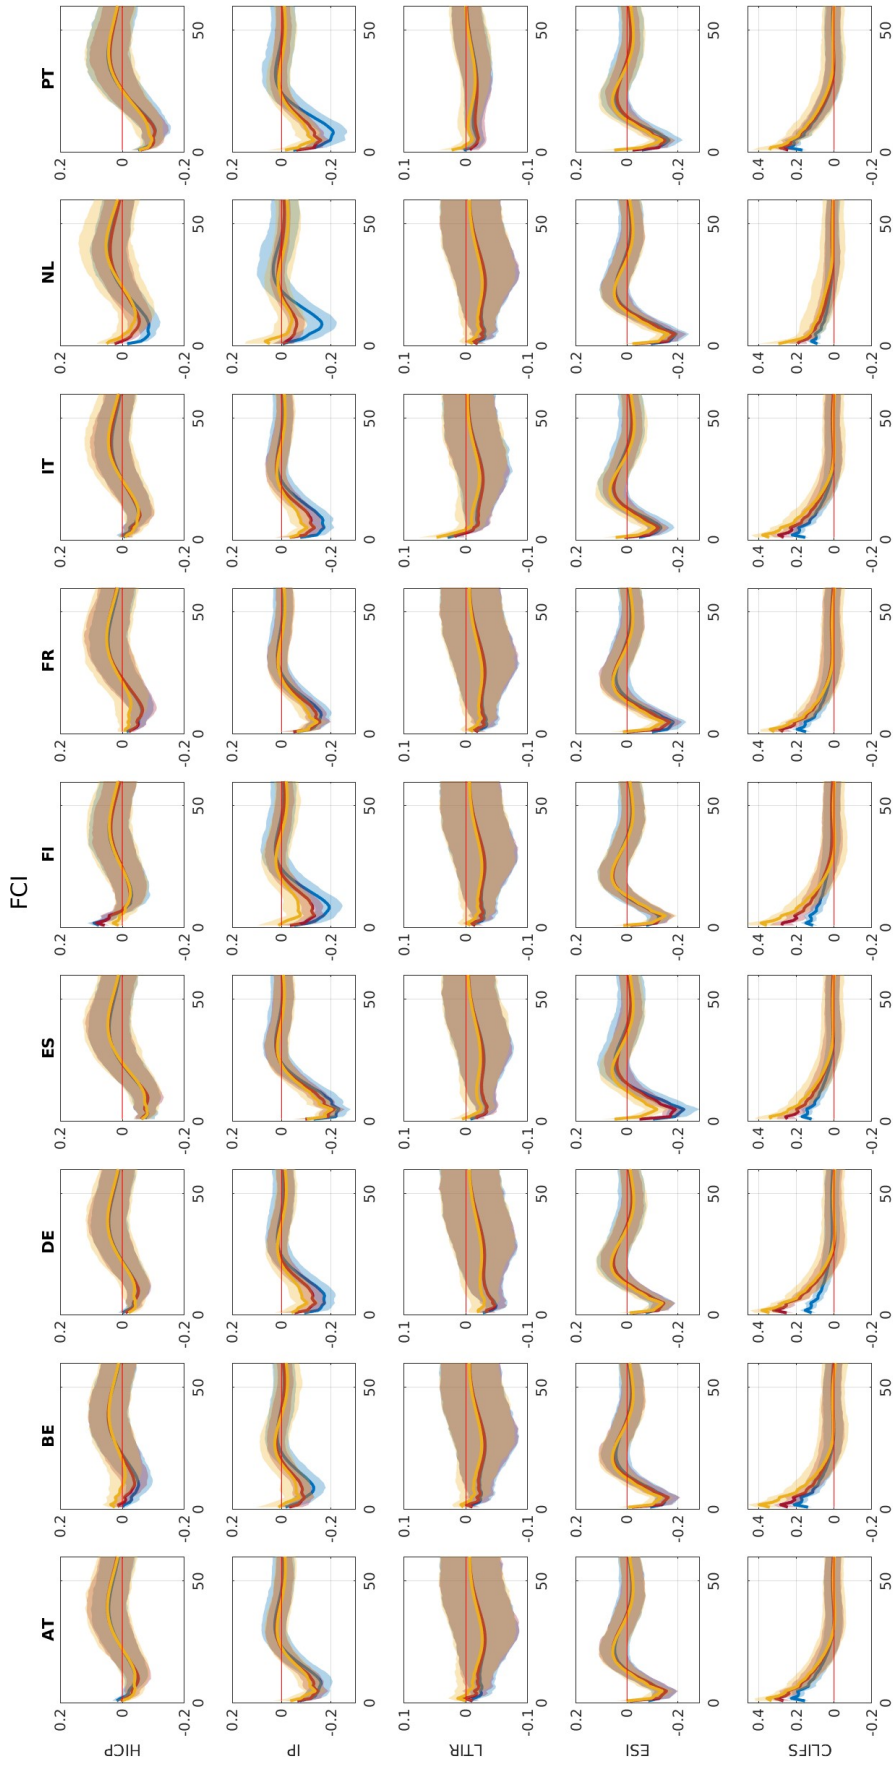
Figure 6: Detailed quantile impulse response functions, shock to global (US) financial conditions (FCI).
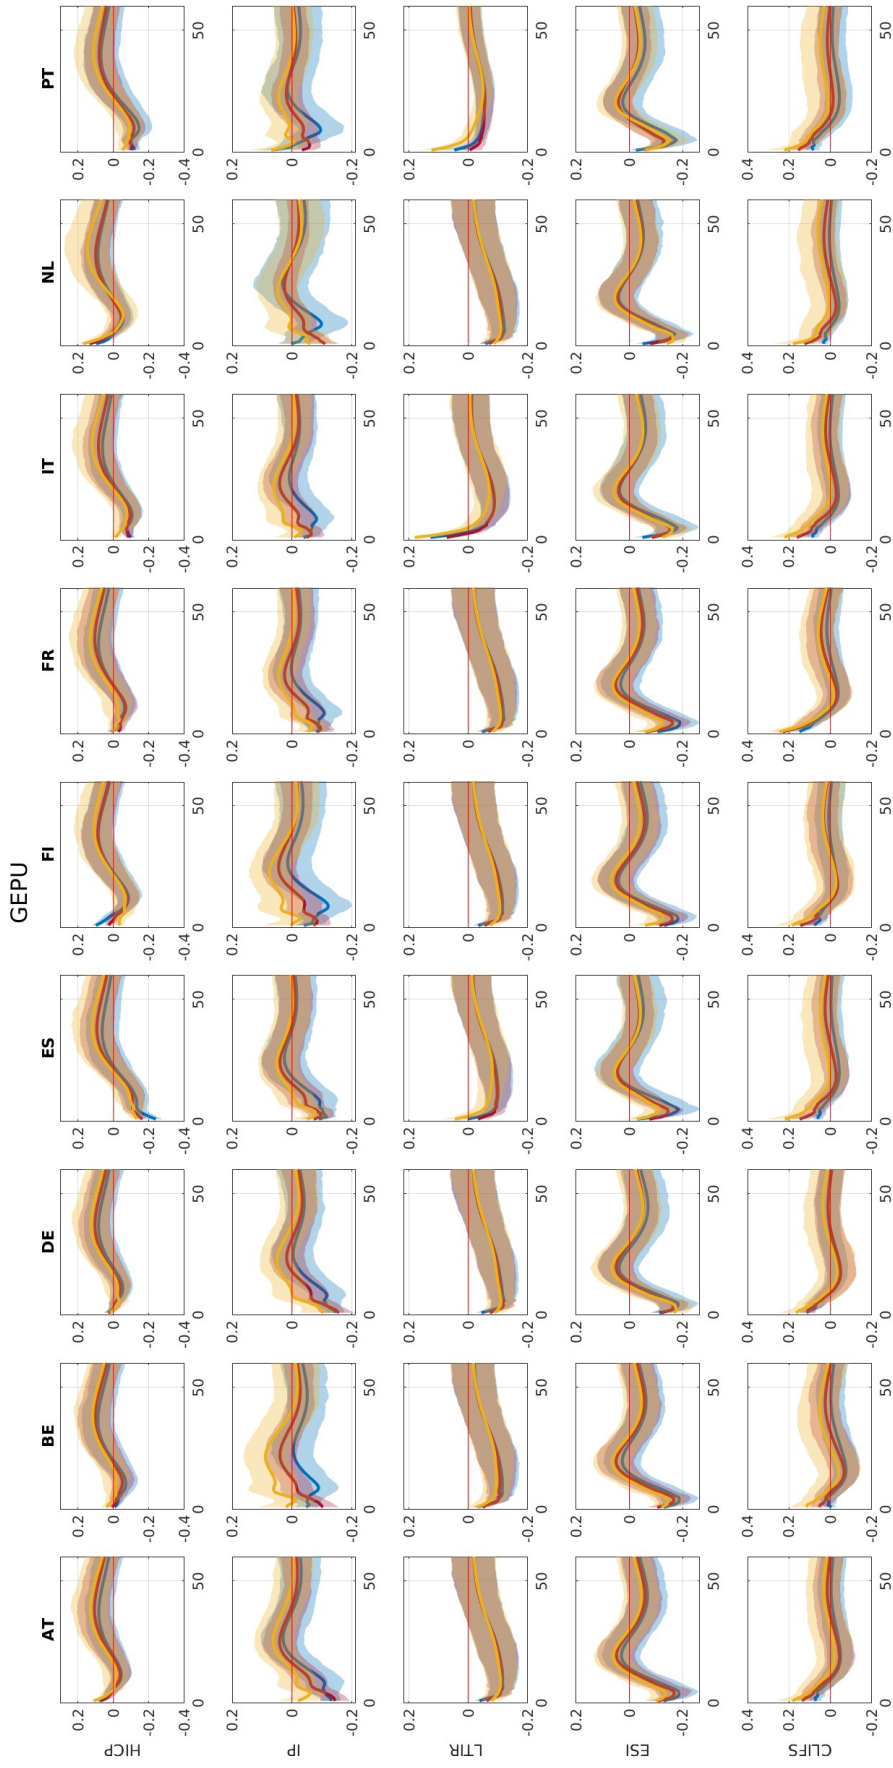
18

Figure 7: Detailed quantile impulse response functions, shock to global economic policy uncertainty (GEPU).

CAMP Working Paper Series
ISSN: 1892-2198

Centre for Applied Macroeconomics and Commodity Prices (CAMP)

will bring together economists working on applied macroeconomic issues, with special emphasis on petroleum economics.

BI Norwegian Business School
Centre for Applied Macro - Petroleum economics (CAMP)
N-0442 Oslo

www.bi.no/camp