

Strathclyde

Discussion Papers
in Economics



Investigating Growth at Risk Using a Multi- country Non-parametric Quantile Factor Model*

Todd E. Clark^a, Florian Huber^b, Gary Koop^c, Massimiliano Marcellino^d, and Michael
Pfarrhofer^b

No. 23 – 7

^A Federal Reserve Bank of Cleveland, ^b University of Salzburg, ^c University of Strathclyde, ^d
Bocconi University, IGER and CEPR

Investigating Growth at Risk Using a Multi-country Non-parametric Quantile Factor Model*

Todd E. CLARK^a, Florian HUBER^b, Gary KOOP^c, Massimiliano MARCELLINO^d, and Michael PFARRHOFER^b

^a*Federal Reserve Bank of Cleveland*

^b*University of Salzburg*

^c*University of Strathclyde*

^d*Bocconi University, IGIER and CEPR*

October 8, 2021

We develop a Bayesian non-parametric quantile panel regression model. Within each quantile, the response function is a convex combination of a linear model and a non-linear function, which we approximate using Bayesian Additive Regression Trees (BART). Cross-sectional information at the p^{th} quantile is captured through a conditionally heteroscedastic latent factor. The non-parametric feature of our model enhances flexibility, while the panel feature, by exploiting cross-country information, increases the number of observations in the tails. We develop Bayesian Markov chain Monte Carlo (MCMC) methods for estimation and forecasting with our quantile factor BART model (QF-BART), and apply them to study growth at risk dynamics in a panel of 11 advanced economies.

JEL: C11, C32, C53

KEYWORDS: non-parametric regression, regression trees, forecasting

*The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System. Marcellino thanks MIUR – PRIN Bando 2017 – prot. 2017TA7TYC for financial support; Huber and Pfarrhofer gratefully acknowledge financial support from the Austrian Science Fund (FWF, grant no. ZK 35).

1 Introduction

Empirical macroeconomics has seen an upsurge of interest in modeling the tails of predictive distributions. A recent influential paper is Adrian, Boyarchenko, and Giannone (ABG, 2019), hereafter ABG, which investigated the impact of financial conditions on the conditional distribution of GDP growth and found it to be important in the lower quantiles. Both prior and subsequent to ABG, a large literature has emerged using quantile regression methods to forecast tail risks to economic growth (see, among many others, Adrian, et al. (2018); Cook and Doh (2019); De Nicolò and Lucchetta (2017); Ferrara, Mogliani, and Sahuc (2019); Giglio, Kelly, and Pruitt (2016); González-Rivera, Maldonado, and Ruiz (2019); Delle Monache, De Polis, and Petrella (2020); Plagborg-Møller, et al. (2020); Reichlin, Ricco, and Hasenzagl (2020); Figueres and Jarociński (2020); and Mitchell, Poon, and Mazzi (forthcoming)). Other studies consider tail risks to other macroeconomic variables such as unemployment or inflation (e.g. Galbraith and van Norden (2019), Kiley (2018), Ghysels, Iania, and Striaukas (2018), Manzan (2015), Gaglianone and Lima (2012), Korobilis (2017), Manzan and Zerom (2013, 2015), Korobilis, et al. (2021) and Pfarrhofer (2021)).

The existing literature, with few exceptions, uses quantile models for a single variable of interest. These models are specified conditional on a specific quantile and assume a linear relationship between the predictors and the quantile function of some outcome variable.¹ For macroeconomic data this assumption might be warranted in normal times but in turbulent times it could be that regression relationships change or turn non-linear. Moreover, often several variables rather than a single one are of interest, and in these cases a joint model would be preferable. These observations motivate the model we develop in the present paper.

In contrast to much of the existing literature we propose a non-parametric model which involves multiple equations and allows for assessing whether the quantile response function is linear or unknown and possibly highly non-linear. In particular, the model we propose is a multi-country, non-parametric quantile regression, which we then use to investigate growth at risk in a panel of 11 advanced economies.

The justification for adopting non-parametric methods is provided by Huber, et al. (2020) and Clark, et al. (2021), which found Bayesian non-parametric vector autoregressions (VARs) to be able to successfully model the tails of predictive densities of macroeconomic variables in a flexible and accurate manner. These papers found that Bayesian Additive Regression Trees

¹Examples of exceptions include Korobilis, et al. (2021) and Pfarrhofer (2021), which assume time variation in the quantile regression coefficients. However, even these papers are single-equation and assume particular parametric forms for the time-variation. A multiple-equation exception is Adrian, et al. (2018), which exploits information in the term structure for an empirical application involving linear panel quantile regression.

(BART) are an effective non-parametric method that is particularly useful in crisis times (e.g., the Financial Crisis or the Covid-19 pandemic) when growth at risk issues are of particular importance. However, in normal periods, the predictive gains from using BART are more muted (and sometimes negative). In the present paper we extend the BART regression methods used in these papers to the quantile BART case. Since the predictive gains of BART vary over the business cycle, we assume that within each quantile, the response function is a convex combination of a linear model and some unknown non-linear function, which we approximate using BART. Studies such as [Taddy and Kottas \(2010\)](#) have developed other Bayesian approaches to nonparametric model-based quantile regression.

The justification for use of a multi-country model is that a panel dimension can often improve forecasts with respect to single country models; see, among many others, [Bai, et al. \(2020\)](#) and [Feldkircher, et al. \(2021\)](#). Moreover, and specifically for the quantile case, macroeconomic data sets are short, leading to a small number of observations in the tails of the distribution. We develop a model for the p^{th} quantile that includes a factor that summarizes the available cross-country information at that quantile. In addition, as indicated below, our Bayesian model specification has features that allow information from other countries to inform estimates for a given country. Exploiting this cross-country information through a pooling prior improves predictive accuracy by parsimoniously including international information to inform coefficients associated with domestic quantities.

We then develop Bayesian Markov Chain Monte Carlo (MCMC) methods for estimation and forecasting with our quantile factor BART model (QF-BART). These methods are scalable to large panels with a potentially large number of exogenous regressors.

In terms of empirical results, our proposed models commonly improve on the benchmark single country linear quantile model in recursive growth forecast comparisons, more so in the tails than near the center of the distribution. Importantly, estimating the weight assigned to the BART component of the model as compared to the linear component is helpful to forecast accuracy. The estimated combination weight is smaller (i.e., the model is more linear) for the 25 and 75 quantiles than in the tails, i.e., for the 5 and 95 quantiles. Moreover, some form of international information definitely pays off (either via the new pooling prior, or by outright including non-domestic series). The effects of the common (volatility) factor are also relevant, as it seems to explain a large fraction of the forecast error variance in most countries, in particular in the tails. A shock to this factor, which can be interpreted as an uncertainty shock, has different (stronger negative) effects in the left tail than the right tail of the growth distribution. In the left tail, the BART piece with an estimated weight tends to slightly mitigate the effects of the

shock. Finally, a financial shock in the US spills over to other countries. There is asymmetry in the responses in the sense that a positive shock affects the growth quantiles, whereas a negative shock's effects are not as sharp. Moreover, the effects and asymmetry are more pronounced in the 2010-19 period than earlier.

The remainder of the paper is structured as follows. The next section defines and motivates our QF-BART model including its prior and discusses MCMC estimation. The third section contains our empirical work while the fourth section concludes the paper.

2 A multi-country non-parametric regression model

We model the joint distribution of (for simplicity, demeaned) output growth for a panel of N countries. These are stored in an N -dimensional vector $\mathbf{y}_t = (y'_{1t}, \dots, y'_{Nt})'$ with y_{it} denoting time t output growth in country i . Domestic real activity might depend on the lags of \mathbf{y}_t as well as other, exogenous factors. These pre-determined quantities are included in a K -dimensional vector \mathbf{x}_{it} . We adopt a notational convention where \mathbf{x}_{it} is structured such that J domestic quantities for country i are always ordered first, followed by all $K - J$ non-domestic variables. We assume that y_{it} follows a quantile regression model which, for the p^{th} quantile, is given by:

$$y_{it} = \omega_{ip}g_{ip}(\mathbf{x}_{it}) + (1 - \omega_{ip})\boldsymbol{\beta}'_{ip}\mathbf{x}_{it} + \lambda_{ip}f_{pt} + \epsilon_{ip,t}, \quad \epsilon_{ip,t} \sim \text{ALD}_p(\sigma_{ip}), \quad i = 1, \dots, N, \quad (1)$$

with $g_{ip} : \mathbb{R}^K \rightarrow \mathbb{R}$ denoting unknown country-specific functions and $\boldsymbol{\beta}_{ip}$ is a $K \times 1$ -dimensional vector of regression coefficients. ω_{ip} is a quantile and country-specific parameter that controls how much weight is placed on the non-linear part of the model. The case $\omega_{ip} = 1$ would correspond to a fully non-linear model whereas $\omega_{ip} = 0$ would be a (conditionally) linear quantile regression specification. Contemporaneous relations across the elements in \mathbf{y}_t are introduced through a static factor model with λ_{ip} denoting the country-specific factor loading and f_{pt} the corresponding international factor. Finally, $\epsilon_{ip,t}$ follows an asymmetric Laplace distribution (ALD) scaled by a parameter σ_{ip} with its p^{th} quantile being equal to zero.²

We assume that the latent factor is uncorrelated over time and arises from a Gaussian distribution:

$$f_{pt} \sim \mathcal{N}(0, e^{h_{pt}}),$$

²The density of the $\text{ALD}_p(\sigma_p)$ is given by $p(1-p)/\sigma_p \exp(-\rho_p(\epsilon_t)/\sigma_p)$, where $\rho_p(x) = x(p - \mathbb{I}(x < 0))$ is the check/loss function and $\mathbb{I}(\bullet)$ the usual indicator function. For details on the correspondence between the Bayesian and classical approach to inference in quantile regression, see [Yu and Moyeed \(2001\)](#).

with h_{pt} being a (logarithmic) variance that evolves according to an AR(1) process:

$$h_{pt} = \mu_p + \rho_p(h_{pt-1} - \mu_p) + \zeta_p u_t, \quad u_t \sim \mathcal{N}(0, 1).$$

Here we let μ_p denote the unconditional mean, ρ_p the autoregressive parameter, and ζ_p^2 the error variance of the log-volatility process.

This model possesses several features which should improve not only its predictive capabilities but also allow for additional inferential opportunities. First, the presence of the quantile-specific weights allows for data-based selection of the degree of non-linearities across quantiles. Recent literature indicates that, in the lower tails of the distribution of output growth, macroeconomic relations change and might be subject to substantial non-linearities. While such non-linearities may be important in the extremes of a distribution, linear models might describe the behavior well in tranquil periods of the business cycle (e.g., in the center of the distribution). Our model allows for this by setting the corresponding weights ω_{ip} appropriately. Second, our model allows for lagged relations across countries. The key point to notice is that these dynamic interdependencies can differ across quantiles. For instance, it could be that in the presence of a global adverse economic shock, cross-country dependencies are more important than in tranquil times. Third, the presence of a common static factor that exhibits conditional heteroscedasticity can capture contemporaneous relations across the elements in \mathbf{y}_t for a specific quantile. Moreover, since the factor is conditionally heteroscedastic, it can also control for sudden shifts in the conditional variance of the dependent variables. Inclusion of this stochastic volatility factor allows us to control for unobserved heterogeneity, a feature which might be extremely important during periods such as the recent Covid-19 pandemic (see, e.g., the discussion in [Carriero, et al. \(2021\)](#)).

The model in [Eq. \(1\)](#) is quite general and nests several commonly used alternatives in the literature. For instance, setting $\omega_{ip} = 0, \lambda_{ip} = 0$ for all p , and setting \mathbf{x}_{it} such that it includes only the first lag of GDP plus a (lagged) measure of financial conditions yields a model very closely related to the one proposed in [Adrian, Boyarchenko, and Giannone \(2019\)](#). Notice that this model essentially rules out cross-country relations. Setting $\omega_{ip} = 1$ for all p yields a non-parametric quantile regression model which, depending on λ_{ip} and the choice of \mathbf{x}_{it} , allows for cross-country relations in a flexible manner. Setting $\omega_{ip} = 0$ returns the linear quantile regression model but with a common volatility factor.

2.1 Approximating the unknown functions using BART

We treat the function g_{ip} as unknown and approximate it using BART (Chipman, George, and McCulloch, 2010). Though other alternatives are possible, BART has been successfully employed in economics for forecasting financial time series in Huber and Rossini (2020), nowcasting GDP in selected European economies in Huber, et al. (2020), and tail forecasting of output, inflation, and unemployment in Clark, et al. (2021). BART is a sum-of-trees model that approximates g_{ip} by summing over many individual trees that all take a simple form and act as “weak learners.” The BART approximation for g_{ip} is given by:

$$g_{ip} \approx \hat{g}_{ip} = \sum_{s=1}^S v(\mathbf{x}_{it} | \mathcal{T}_{ip}^s, \boldsymbol{\mu}_{ip}^s),$$

with v denoting a tree function that is determined by a tree structure \mathcal{T}_{ip}^s and a vector of terminal node parameters $\boldsymbol{\mu}_{ip}^s$. This terminal node parameter vector has dimension b_{ip}^s .

The tree structure consists of multiple decision rules that ask whether a covariate exceeds a threshold and, according to these simple binary rules, produces (disjoint) partitions of the input space. These take the form $x_{ij,t} > c$ or $x_{ij,t} \leq c$, with $x_{ij,t}$ denoting the j^{th} element of x_{it} and c being a splitting/threshold value. Sequences of these decision rules lead to a terminal node coupled with a corresponding terminal node parameter in $\boldsymbol{\mu}_{ip}^s$.

When S is large, the BART approximation is prone to overfitting if no further regularization is introduced. Chipman, George, and McCulloch (2010) use regularization priors to force the trees v to be simple. We achieve this through shrinkage priors on the tree structure and the terminal node parameters. Following Chipman, George, and McCulloch (1998), the prior on \mathcal{T}_{ip}^s is obtained by constructing a tree-generating stochastic process. The prior $p(\mathcal{T}_{ip}^s)$ comprises of three aspects. First, tree complexity ultimately depends on the depth of intermediate nodes d . If d is large, the tree is complex and thus might overfit the data. To force the individual trees to be simple, we assume that a given node at depth d is non-terminal with probability proportional to:

$$\alpha (1 + d)^{-\zeta},$$

where α is between 0 and 1 and $\zeta > 0$. Notice that this probability decreases in d : growing more complicated trees becomes unlikely if d is already large. The amount of shrinkage is controlled by α and ζ . These hyperparameters are often set to $\alpha = 0.95$ and $\zeta = 2$, implying that trees with two or three terminal nodes receive over 80% of total prior probability. Chipman, George,

and McCulloch (2010) found that, for over 40 data sets, this choice performs well, and extensive cross-validation for α and ζ only improves predictive accuracy by small margins. The second and third aspect are concerned with how decision rules are constructed. To this end, we use discrete uniformly distributed priors to select the variables showing up in the decision rule as well as a uniform prior over the splitting/threshold values.

The second source of shrinkage is a Gaussian shrinkage prior on $\mu_{ij,p}^s$, the j^{th} element of $\boldsymbol{\mu}_{ip}^s$. Chipman, George, and McCulloch (2010) recommend scaling the prior using the range of the data. More specifically, let $y_{i,\min}$ and $y_{i,\max}$ denote the minimum and maximum of the observed data in country i . The corresponding Gaussian prior is then given by:

$$\mu_{ij,p}^s \sim \mathcal{N}(0, v_{ip}^2), \quad v_{ip} = \frac{y_{i,\max} - y_{i,\min}}{2\gamma\sqrt{S}},$$

with $\gamma > 0$ being a prior scaling parameter, typically set equal to 2. The prior implies that if the number of trees S is large, the prior variance decreases and the amount explained by a single tree is decreased. This is consistent with the notion that each tree acts as a “weak learner,” explaining only a small share of variation in the response variable, but the ensemble model provides sufficient flexibility to capture even complicated conditional mean relations. Another feature, noted by Huber, et al. (2020), is that the prior variance increases in the range of the data. Hence, if outliers arise, the prior becomes increasingly loose and allows for more flexibility in terms of capturing observations far outside the range of past data.

The priors on the tree structures and the terminal node parameters constitute the main ingredients of BART. Since our model also features a linear part and additional coefficients, we also need to specify priors on $\boldsymbol{\beta}_{ip}$ and ω_{ip} . We discuss them in the next sub-section.

2.2 Priors on the remaining coefficients of the model

On the coefficients $\boldsymbol{\beta}_{ip}$ we use a horseshoe-type prior on each element $\beta_{ij,p}$:

$$\beta_{ij,p} \sim \mathcal{N}\left(\frac{\beta_{jp}}{\varphi}, \varphi^2 \psi_{ij,p}^2\right), \quad \varphi \sim \mathcal{C}^+(0, 1), \quad \psi_{ij,p} \sim \mathcal{C}^+(0, 1),$$

where \mathcal{C}^+ denotes a half-Cauchy distribution, $\psi_{ij,p}$ is a coefficient and quantile-specific scaling parameter, and φ is a global shrinkage parameter that is common to all coefficients. Notice that the presence of φ introduces dependencies across coefficients (including across countries) and across quantiles. The key advantage is that the presence of the local shrinkage parameters $\psi_{ij,p}$ allows the detection of signals (i.e., non-zero or heterogeneous $\beta_{ij,p}$ over the cross section) even

if φ is close to zero.

The prior mean $\underline{\beta}_{jp}$ pools information over the cross section. In our hierarchical specification, it is estimated from the data using a Gaussian prior for the domestic variables, and deterministically set to zero for non-domestic quantities:

$$\underline{\beta}_{jp} \sim \mathcal{N}(0, \tilde{\varphi}_j) \quad \text{for } j = 1, \dots, J, \quad \underline{\beta}_{jp} = 0 \quad \text{for } j = J + 1, \dots, K.$$

The parameter $\tilde{\varphi}_j$ is the prior variance of the common mean, which we set to a weakly informative value of 10 for the empirical application. We refer to this prior as the pooled horseshoe (HSP), while setting the common mean to a zero vector of size K yields the conventional horseshoe (HS) that we consider as an alternative.

For the factor loadings λ_{ip} , we use a set of independent Gaussian priors for all i, p :

$$\lambda_{ip} \sim \mathcal{N}(0, 1).$$

Note that λ_{ip} is a scalar and, hence, we use this relatively non-informative prior rather than a prior such as the HS which is used to avoid over-parameterization as might occur with high dimensional parameters.

On the weights ω_{ip} we consider a Uniform prior:

$$\omega_{ip} \sim \mathcal{U}(0, 1). \tag{2}$$

This prior introduces no particular prior information on the amount of non-linearities. If we wish to be informative on ω_{ip} we can also use a Beta prior and specify the hyperparameters appropriately.

The remaining coefficients of the model relate to the error term. [Kozumi and Kobayashi \(2011\)](#) write the ALD using a scale-location mixture of Gaussians:

$$\epsilon_{ip,t} = \theta_p \nu_{ip,t} + \tau_p \sqrt{\sigma_{ip} \nu_{ip,t}} e_{ip,t},$$

with $\theta_p = \frac{1-2p}{p(1-p)}$, $\tau_p^2 = \frac{2}{p(1-p)}$, $\nu_{ip,t} = \sigma_{ip} z_{ip,t}$, $e_{ip,t} \sim \mathcal{N}(0, 1)$, and $z_{ip,t} \sim \text{Exp}(1)$. On the scale parameter σ_{ip} we use an inverse Gamma prior:

$$\sigma_{ip} \sim \mathcal{G}^{-1} \left(\frac{a_\sigma}{2}, \frac{b_\sigma}{2} \right),$$

with the relatively non-informative choices of $a_\sigma = 1$ and $b_\sigma = 1$.

This completes the prior setup. In the next sub-section we briefly discuss the Markov chain Monte Carlo (MCMC) algorithm used to carry out estimation and inference.

2.3 Full conditional posterior simulation

We use Markov Chain Monte Carlo (MCMC) techniques to obtain a draw from the joint posterior of the latent quantities and coefficients of the model. Specifically, the following steps of the algorithm are carried out for each equation (i.e., country) i and quantile p :

- **Sampling from $p(\mathcal{T}_{ip}^s|\bullet)$ and $p(\boldsymbol{\mu}_{ip}^s|\bullet)$.** The full conditional posterior of the tree structures takes no well-known form. Chipman, George, and McCulloch (2010) propose a Bayesian backfitting strategy to set up a Metropolis Hastings (MH) algorithm to sample the trees individually, conditionally on the other $S - 1$ trees. This step is carried out marginally of $\boldsymbol{\mu}_{ip}^s$. The terminal node parameters can then, under our conjugate prior, be simulated from a set of independent Gaussian distributions that take a well-known form.
- **Sampling $p(\boldsymbol{\beta}_{ip}|\bullet)$.** The regression coefficients are, conditional on the remaining parameters and latent states, simulated from a multivariate Gaussian posterior distribution with known moments:

$$\boldsymbol{\beta}_{ip}|\bullet \sim \mathcal{N}(\bar{\boldsymbol{\beta}}_{ip}, \bar{\mathbf{V}}_{ip}), \quad \bar{\mathbf{V}}_{ip} = \left(\tilde{\mathbf{X}}_{ip}' \tilde{\mathbf{X}}_{ip} + \mathbf{V}_{ip}^{-1} \right)^{-1}, \quad \bar{\boldsymbol{\beta}}_{ip} = \bar{\mathbf{V}}_{ip} \left(\mathbf{V}_{ip}^{-1} \underline{\boldsymbol{\beta}}_p + \tilde{\mathbf{X}}_{ip}' \tilde{\mathbf{y}}_{ip} \right).$$

$\tilde{\mathbf{y}}_{ip}$ is a T -dimensional response vector with t^{th} element given by $\tilde{y}_{ip,t} = (y_{it} - \omega_{ip} \hat{g}_{ip}(\mathbf{x}_{it}) - \lambda_{ip} f_{pt} - \theta_p \nu_{ip,t}) / (\tau_p \sqrt{\sigma_{ip} \nu_{ip,t}})$, $\tilde{\mathbf{X}}_{ip}$ is a $T \times K$ matrix with typical row $\tilde{\mathbf{x}}_{ip,t} = ((1 - \omega_{ip}) \mathbf{x}_{it}) / (\tau_p \sqrt{\sigma_{ip} \nu_{ip,t}})$, and \mathbf{V}_{ip} is a prior variance matrix with main diagonal element $v_{ij,p} = \varphi^2 \psi_{ij,p}^2$. The prior mean $\underline{\boldsymbol{\beta}}_p = (\underline{\beta}_{1p}, \dots, \underline{\beta}_{Jp}, \mathbf{0}'_{K-J})'$ collects the estimated common means $\underline{\beta}_{jp}$ in the corresponding position of the J domestic variables with the remaining elements being zero for the HSP prior, while $\underline{\boldsymbol{\beta}}_p = \mathbf{0}_K$ for the conventional HS prior.

- **Sampling from $p(\underline{\beta}_{jp}|\bullet)$.** The posterior distribution for the non-zero elements for $j = 1, \dots, J$ of the prior mean for HSP is $\underline{\beta}_{jp} \sim \mathcal{N}(\bar{b}_{jp}, \bar{v}_{jp})$, with moments

$$\bar{v}_{jp} = \left(\left(\sum_{i=1}^N \frac{1}{v_{ij,p}} \right) + \frac{1}{\varphi_j} \right)^{-1}, \quad \bar{b}_{jp} = \bar{v}_{j,p} \left(\sum_{i=1}^N \frac{\beta_{ij,p}}{v_{ij,p}} \right).$$

- **Sampling from $p(\omega_{ip}|\bullet)$:** The full conditional posterior of ω_{ip} takes no well-known form. Since the support of ω_{ip} is bounded and the target density univariate, we adopt a slice sampler (see, e.g., Neal (2003)) that is straightforward to implement and mixes well.

- **Sampling from $p(\lambda_{ip}|\bullet)$.** The factor loadings are obtained by simulating from univariate Gaussian conditional posterior distributions:

$$\lambda_{ip}|\bullet \sim \mathcal{N}(\bar{\lambda}_{ip}, \bar{l}_{ip}), \quad \bar{l}_{ip} = (\hat{\mathbf{f}}_p' \hat{\mathbf{f}}_p + 1)^{-1}, \quad \bar{\lambda}_{ip} = \bar{l}_{ip} \hat{\mathbf{f}}_p' \hat{\mathbf{y}}_{ip}.$$

$\hat{\mathbf{y}}_{ip}$ denotes the $T \times 1$ response vector with typical t^{th} element given by $\hat{y}_{ip,t} = (y_{it} - \omega_{ip} \hat{g}_{ip}(\mathbf{x}_{it}) - (1 - \omega_{ip}) \boldsymbol{\beta}'_{ip} \mathbf{x}_{it}) / (\tau_p \sqrt{\sigma_{ip} \nu_{ip,t}})$, and $\hat{\mathbf{f}}_p$ has typical element $\hat{f}_{pt} = f_{pt} / (\tau_p \sqrt{\sigma_{ip} \nu_{ip,t}})$.

- **Sampling from $p(\sigma_{ip}|\bullet)$.** Kozumi and Kobayashi (2011) show that the conditional posterior of the scaling parameter σ_{ip} follows an inverse Gamma distribution:

$$\sigma_{ip}|\bullet \sim \mathcal{G}^{-1} \left(\frac{\tilde{a}_{ip}}{2}, \frac{\tilde{b}_{ip}}{2} \right),$$

with $\tilde{a}_{ip} = a_\sigma + 3T$, $\tilde{b}_{ip} = b_\sigma + 2 \sum_{t=1}^T \nu_{ip,t} + \sum_{t=1}^T (w_{it} - \theta_p \nu_{ip,t})^2 / (\tau_p^2 \nu_{ip,t})$, and $w_{it} = y_{it} - \omega_{ip} \hat{g}_{ip}(\mathbf{x}_{it}) - (1 - \omega_{ip}) \boldsymbol{\beta}'_{ip} \mathbf{x}_{it} - \lambda_{ip} f_{pt}$.

- **Sampling from $p(\nu_{ip,t}|\bullet)$.** For each t , we simulate $\nu_{ip,t}$ from a generalized inverse Gaussian (GIG) posterior distribution:

$$\nu_{ip,t}|\bullet \sim \text{GIG} \left(\frac{1}{2}, \tilde{c}_{ip}, \tilde{d}_{ip} \right),$$

where $\tilde{c}_{ip} = w_{it} / (\tau_p^2 \sigma_{ip})$ and $\tilde{d}_{ip} = 2 / \sigma_{ip} + \theta_p^2 / (\tau_p^2 \sigma_{ip})$.³

- **Sampling from $p(\psi_{ij,p}^2|\bullet)$.** The local, coefficient-specific scaling parameters are simulated using the scheme outlined in Makalic and Schmidt (2015). Conditional on auxiliary shrinkage parameters ξ and $\eta_{ij,p}$, the posterior of $\psi_{ij,p}^2$ is inverse Gamma distributed:

$$\psi_{ij,p}^2 \sim \mathcal{G}^{-1} \left(1, \frac{1}{\eta_{ij,p}} + \frac{(\beta_{ij,p} - \beta_{jp})^2}{2\varphi^2} \right), \quad \eta_{ij,p}|\bullet \sim \mathcal{G}^{-1} \left(1, 1 + \frac{1}{\xi} \right).$$

These steps relate to the quantities we have to simulate for each country (or equation) and quantile. Next we turn to the quantities that we simulate per quantile and thus pool over countries.

- **Sampling from $p(f_{pt}|\bullet)$.** For each t , we simulate f_{pt} from a sequence of independent

³The generalized inverse Gaussian (GIG) distribution is parameterized such that a random variable $X \sim \text{GIG}(\lambda, \xi, \psi)$ has probability density function $f(x; \lambda, \xi, \psi) = x^{\lambda-1} \exp(-(\xi/x + \psi x)/2)$.

Gaussian posterior distributions as follows:

$$f_{pt}|\bullet \sim \mathcal{N}(\mathbf{A}_{pt}\tilde{\mathbf{y}}_{pt}, e^{h_{pt}} - \mathbf{A}_{pt}\boldsymbol{\Xi}_{pt}\mathbf{A}'_{pt}), \quad \boldsymbol{\Xi}_{pt} = e^{h_{pt}}(\boldsymbol{\lambda}_p\boldsymbol{\lambda}'_p) + \boldsymbol{\Psi}_{pt}, \quad \mathbf{A}_{pt} = e^{h_{pt}}\boldsymbol{\lambda}'_p\boldsymbol{\Xi}_{pt},$$

whereby $\tilde{\mathbf{y}}_{pt}$ is a $N \times 1$ vector with $\tilde{y}_{ip,t} = y_{it} - \omega_{ip}\hat{g}_{ip}(\mathbf{x}_{it}) + (1 - \omega_{ip})\boldsymbol{\beta}'_{ip}\mathbf{x}_{it} - \theta_p\nu_{ip,t}$, $\boldsymbol{\lambda}_p = (\lambda_{1p}, \dots, \lambda_{Np})'$, and $\boldsymbol{\Psi}_{pt} = \text{diag}(\tau_p^2\sigma_{1p}\nu_{1p,t}, \dots, \tau_p^2\sigma_{Np}\nu_{Np,t})$.

- **Sampling from $p(\mathbf{h}_p|\bullet)$ and $p(\mu_p, \rho_p, \varsigma_p|\bullet)$.** We sample the full history of log-volatilities $\mathbf{h}_p = (h_{p1}, \dots, h_{pT})'$ and the parameters of the state evolution equation using the efficient sampler proposed in [Kastner and Frühwirth-Schnatter \(2014\)](#). This algorithm samples the log-volatilities, conditional on everything else, all without a loop.

The final step refers to the global shrinkage parameter of the horseshoe prior. This step pools information across all equations and quantiles. Since we rely on auxiliary random variables to obtain a well-known full conditional posterior distribution, we first simulate from $p(\xi|\bullet)$ and then from $p(\varphi|\bullet)$.

- **Sampling from $p(\varphi^2|\bullet)$ and $p(\xi|\bullet)$.** The conditional posteriors of the global shrinkage parameter φ and the auxiliary global parameter ξ are, respectively, inverse Gamma distributed:

$$\varphi^2|\bullet \sim \mathcal{G}^{-1}\left(\frac{NK+1}{2}, \frac{1}{\xi} + \frac{1}{2}\sum_p\sum_{i=1}^N\sum_{j=1}^K\frac{(\beta_{ij,p} - \beta_{jp})^2}{\psi_{ij,p}^2}\right), \quad \xi|\bullet \sim \mathcal{G}^{-1}\left(1, 1 + \frac{1}{\varphi^2}\right).$$

This completes our MCMC algorithm. In all our empirical work we repeat the different steps 30,000 times and discard the first 15,000 draws as burn-in. One key advantage of the present algorithm is that it is scalable to larger data sets (i.e., including more countries, additional endogenous variables, or more covariates) because, conditional on the factors and φ , the different posterior quantities are independent across equations and quantiles.

3 Empirical results

In this section we first investigate whether our modeling approach improves upon a set of simpler, nested alternatives by means of a forecasting horse race. We then focus on international growth at risk dynamics in two ways. First, we analyze how GDP growth reacts to changes in the common factor. Afterwards, we focus on how a shock to US financial conditions spills over to the other economies in our sample.

3.1 Data overview, competing models and forecasting design

Our sample runs from 1975Q1 to 2020Q4. We use annualized quarterly growth rates of GDP data from the Main Economic Indicators (MEI) database, maintained by the OECD, and the composite indicator of systemic stress (CISS) by the European Central Bank (ECB). For data availability reasons we include Austria (AT), Denmark (DK), Finland (FI), France (FR), Germany (DE), Italy (IT), Netherlands (NL), Spain (ES), Sweden (SE), United Kingdom (UK), and the United States (US).

We estimate the models for $p \in \{0.05, 0.10, 0.25, 0.50, 0.75, 0.90, 0.95\}$. For each model we consider two different choices for the covariates. The first, which we label **CISS**, includes the CISS and a single lag of y_{it} in \mathbf{x}_{it} , implying that $K = 2$. The second includes cross-country information in \mathbf{x}_{it} by including the first lag of GDP growth and the CISS of all countries; hence, $K = 2N$. The latter is referred to as **CISS-CC** to indicate that the information set includes cross-country data.

Since our model is quite flexible and nests several competing models, we also include a range of restricted variants of the general model outlined in Section 2. First, we obtain the ABG model by using the CISS covariates and setting $\mathbf{\Lambda} = \mathbf{0}$ and $\omega_{ip} = 0$. We use frequentist methods to carry out estimation so as to be the same as ABG, while we estimate all other models using Bayesian methods with either the HS or HSP prior (so, for example, we will consider both CISS-CC-HS and CISS-CC-HSP specifications). ABG will serve as our benchmark model to which we compare all other specifications. We then add features to this benchmark. We begin by remaining linear ($\omega_{ip} = 0$) but adding the international factor to the ABG model by letting $\mathbf{\Lambda} \neq \mathbf{0}$ in order to investigate whether it plays an empirically important role. All subsequent models also let $\mathbf{\Lambda} \neq \mathbf{0}$. We next investigate non-linearities by setting $\omega_{ip} = 1$ and thus obtain a multi-country quantile BART model with a common international factor. Finally our most flexible model allows for ω_{ip} to be estimated from the data. An overview of all model specifications is provided in Table 1.

We compute pseudo out-of-sample forecasts based on a holdout from 1990Q1 to 2020Q4 (so the initial training sample comprises 60 quarters). We compute Quantile Scores (QS, for quantiles 0.1, 0.25, 0.5, 0.75, 0.9) and quantile-weighted cumulative ranked probability scores (qw-CRPS, see Gneiting and Ranjan (2011)) with five weighting schemes (“none” refers to no weighting, i.e., conventional CRPS; both tails “tails;” left tail, “left;” right tail, “right;” and “center”). We compute direct forecasts for $h \in \{1, 4\}$.

Table 1: Model overview.

Data	Prior	Weights	Factor
CISS (domestic)	HS (shrinkage to zero)	$\omega = 0$ (parametric)	$\Lambda = 0$ (independence)
CISS-CC (cross-country)	HSP (pooling cross-section)	$\omega = 1$ (nonparametric) $\omega \in (0, 1)$ (estimated)	$\Lambda \neq 0$ (dependence)

Notes: “Data” refers to the information set for individual country models. “Prior” indicates the prior on the parametric part of the model; we consider the conventional horseshoe prior (HS) shrinking towards zero and the pooled horseshoe (HSP) prior pushing the model towards cross-sectional homogeneity. “Weights” refers to the specification of the conditional quantile function: parametric, nonparametric, or whether we estimate weights on the parametric and nonparametric part. “Factor” indicates whether an international factor modeling the cross-sectional covariance structure within quantiles is present. We consider all possible combinations of these specification choices.

	ALL	AT	DE	DK	ES	FI	FR	IT	NL	SE	UK	US	
CISS-HSP	-0.81 0.97 0.95	0.99 0.97 0.94	1.06 1.02 0.98	0.99 0.99 0.97	1.06 0.98 0.93	1.02 1.01 1.01	0.57 0.88 0.88	1.02 0.91 0.93	1.01 1.03 0.96	0.95 0.94 0.99	1.01 0.94 0.94	1.00 0.97 0.97	none
CISS-HS	-0.86 0.96 0.96	0.99 0.97 0.95	1.05 1.02 0.98	0.99 0.99 0.98	0.78 0.93 0.95	1.01 1.01 1.00	0.70 0.87 0.89	0.99 0.93 0.94	1.02 1.03 0.97	0.94 0.94 1.00	1.03 0.94 0.94	1.00 0.97 0.97	
CISS-CC-HSP	-0.80 0.94 0.85	1.08 0.97 1.10	1.11 0.97 1.12	1.00 0.95 1.07	0.73 0.94 1.10	1.04 0.98 1.12	0.59 0.83 0.68	0.70 0.91 0.78	1.08 0.99 1.11	1.05 0.93 1.11	0.75 0.97 0.85	1.10 0.98 0.93	
CISS-CC-HS	-0.80 0.94 0.82	1.09 0.98 1.05	0.90 0.99 1.09	1.01 0.95 1.07	0.74 0.91 1.10	1.05 0.98 1.10	0.56 0.85 0.99	0.72 0.89 0.73	1.07 0.97 1.09	1.06 0.94 1.11	0.73 0.95 1.18	1.12 0.99 1.17	
CISS-HSP	-0.76 0.93 0.94	0.99 0.95 0.93	1.05 1.01 0.97	1.00 0.99 0.97	1.01 0.94 0.89	1.01 1.01 1.02	0.59 0.82 0.85	1.01 0.88 0.93	0.97 0.99 0.96	0.93 0.92 0.98	0.92 0.87 0.92	0.96 0.94 0.97	tails
CISS-HS	-0.82 0.93 0.95	1.00 0.96 0.93	1.04 1.01 0.97	0.99 0.99 0.98	0.72 0.87 0.92	1.01 1.01 1.00	0.63 0.82 0.86	0.97 0.93 0.94	0.98 0.99 0.97	0.92 0.91 0.99	0.96 0.87 0.94	0.97 0.94 0.96	
CISS-CC-HSP	-0.76 0.92 0.87	1.11 0.98 1.18	1.13 0.95 1.18	1.01 0.95 1.15	0.69 0.89 1.11	1.06 0.99 1.20	0.53 0.76 0.68	0.67 0.89 0.82	1.10 0.99 1.19	1.07 0.91 1.15	0.69 0.91 0.90	1.09 0.96 0.98	
CISS-CC-HS	-0.77 0.91 0.85	1.11 0.96 1.13	0.91 0.98 1.16	1.01 0.95 1.15	0.69 0.84 1.10	1.09 0.99 1.19	0.52 0.80 0.98	0.71 0.86 0.77	1.07 0.97 1.17	1.07 0.92 1.16	0.69 0.89 1.21	1.15 0.99 1.22	
CISS-HSP	-0.89 0.96 0.97	0.97 0.94 0.93	1.10 1.07 1.00	1.00 0.99 0.98	0.96 0.90 0.94	1.01 1.01 1.02	0.70 0.84 0.90	0.93 0.90 0.95	1.05 1.06 0.99	0.96 0.95 1.02	1.01 0.98 0.94	1.05 1.02 0.98	left
CISS-HS	-0.90 0.96 0.98	0.94 0.94 0.94	1.08 1.07 1.00	1.00 1.00 1.00	0.80 0.89 0.96	1.01 1.01 1.01	0.73 0.83 0.91	0.91 0.90 0.95	1.06 1.05 1.01	0.96 0.95 1.04	1.00 0.97 0.94	1.03 1.02 0.98	
CISS-CC-HSP	-0.85 0.93 0.96	0.98 0.97 1.08	1.05 0.98 1.12	0.97 0.96 1.05	0.79 0.87 1.07	1.03 0.99 1.08	0.68 0.80 0.83	0.74 0.90 0.91	0.99 0.95 1.08	0.98 0.91 1.06	0.84 0.95 1.01	1.09 1.05 1.06	
CISS-CC-HS	-0.87 0.94 0.93	0.98 0.95 1.01	0.98 1.01 1.07	0.98 0.95 1.06	0.79 0.87 1.05	1.01 0.99 1.06	0.70 0.82 0.96	0.80 0.89 0.85	0.98 0.95 1.03	0.99 0.92 1.07	0.83 0.96 1.15	1.17 1.06 1.14	
CISS-HSP	-0.71 0.95 0.93	1.03 0.99 0.95	1.02 0.98 0.95	0.98 0.98 0.95	1.14 1.04 0.89	1.02 1.01 0.99	0.59 0.89 0.84	1.11 0.91 0.91	0.94 0.98 0.93	0.92 0.92 0.94	0.96 0.88 0.93	0.92 0.91 0.96	right
CISS-HS	-0.80 0.95 0.94	1.04 0.99 0.95	1.01 0.97 0.95	0.98 0.98 0.96	0.72 0.94 0.92	1.02 1.02 0.99	0.63 0.86 0.86	1.06 0.96 0.93	0.96 0.98 0.93	0.92 0.92 0.95	1.02 0.87 0.94	0.96 0.91 0.96	
CISS-CC-HSP	-0.72 0.94 0.74	1.22 0.98 1.16	1.17 0.95 1.15	1.03 0.94 1.12	0.65 0.98 1.14	1.07 0.97 1.20	0.46 0.82 0.53	0.64 0.91 0.67	1.19 1.03 1.18	1.13 0.94 1.18	0.62 0.96 0.72	1.11 0.90 0.82	
CISS-CC-HS	-0.71 0.93 0.73	1.22 1.00 1.15	0.81 0.97 1.15	1.04 0.95 1.12	0.65 0.91 1.15	1.12 0.97 1.18	0.40 0.85 1.02	0.63 0.88 0.63	1.16 1.00 1.20	1.15 0.95 1.18	0.62 0.92 1.22	1.09 0.92 1.21	
CISS-HSP	-0.84 0.98 0.96	1.00 0.98 0.94	1.06 1.03 0.98	0.99 0.99 0.97	1.09 1.00 0.95	1.02 1.02 1.00	0.61 0.92 0.90	1.03 0.93 0.93	1.03 1.04 0.96	0.96 0.95 0.99	1.05 0.98 0.95	1.01 0.99 0.98	center
CISS-HS	-0.88 0.98 0.97	0.98 0.98 0.96	1.05 1.03 0.98	0.99 0.99 0.98	0.81 0.97 0.97	1.02 1.02 1.00	0.74 0.90 0.91	1.00 0.93 0.95	1.04 1.04 0.98	0.96 0.95 1.01	1.07 0.97 0.95	1.01 0.99 0.97	
CISS-CC-HSP	-0.81 0.95 0.83	1.07 0.97 1.06	1.10 0.98 1.10	1.00 0.95 1.04	0.76 0.95 1.09	1.04 0.97 1.08	0.62 0.86 0.68	0.71 0.91 0.77	1.07 0.98 1.07	1.04 0.94 1.09	0.78 0.99 0.83	1.11 0.98 0.90	
CISS-CC-HS	-0.81 0.96 0.81	1.08 0.98 1.02	0.89 0.99 1.06	1.01 0.95 1.04	0.76 0.94 1.10	1.04 0.98 1.06	0.59 0.87 0.99	0.72 0.90 0.71	1.06 0.97 1.05	1.06 0.95 1.09	0.75 0.99 1.17	1.11 0.99 1.14	

Table 2: Relative quantile weighted cumulative ranked probability scores (CRPS) for $h = 1$ and models with $\Lambda \neq \mathbf{0}$. The results are benchmarked to CISS with $\Lambda = \mathbf{0}$ and $\omega = \mathbf{0}$. Lower ratios (shaded in green) indicate better performance (and vice versa, shaded in red).

3.2 Tail forecasting results

Table 2 reports the forecast comparison of the various models based on the relative qw-CRPS. Each cell in the heatmap shows the qw-CRPS relative to the ABG benchmark model. Numbers smaller than one indicate outperformance (green colored) vis-à-vis the ABG model whereas numbers exceeding one suggest a weaker performance (red colored) than the benchmark.

Four main comments can be made. First, and focusing on aggregate results across countries, our proposed models commonly improve on the benchmark ABG model. The gains are about 20 percent when looking at the standard CRPS, decrease to about 10 to 15 percent in the left tail, and increase to about 30 percent in the right tail (based on additional results reported in the Appendix). Overall, these results suggest that at each quantile, and particularly in the right tail, it pays off to allow for non-linearities and for cross-country relations.

Second, while there are small differences between setting $\omega_{ip} = 0$ (linear quantile) or $\omega_{ip} = 1$ (BART quantile), there is often some benefit to estimating the weight ω_{ip} , in turn allowing for both linear and BART pieces in the model. The key advantage of estimating ω_{ip} is that it combines the best of both worlds and thus translates into a model that is strongly

non-linear and non-parametric in the tails and close to a linear quantile regression model in the center of the distribution. Such a behavior is beneficial if loss functions which evaluate the full predictive distribution are used.

Third, the HSP prior, that includes pooling, is typically better than HS, but the differences shrink or are eliminated once cross-country information is included in the model. It is noteworthy that once we use a pooling prior the predictive benefit of adding cross-country information directly diminishes sharply. This points towards the fact that, through pooling, our approach successfully picks up cross-sectional information in a very parsimonious manner.

Finally, there is some heterogeneity across the countries under analysis. In particular, for Spain, France, Italy, and the UK the results are broadly in line with those mentioned above. In contrast, for Austria and Sweden estimating the weight ω yields little gains (setting $\omega_{ip} = 1$ is often best), and for the other countries it is overall difficult to beat the benchmark.

3.3 Estimated weights over time

In the previous sub-section we have shown that our proposed framework yields forecast distributions which are often more precise than the ones obtained from the ABG benchmark and simpler nested alternatives. One key advantage of the model is that it allows for different weights ω_{ip} across countries and quantiles and this improves forecasts when the full predictive density is evaluated. In this sub-section, we investigate whether our intuition that non-linearities are relevant in the tails while linear models are adequate in the center of the distribution is supported by our model.

Figure 1 reports the estimated weight ω_{ip} over our hold-out period. Darkblue cells indicate a weight close to one while gray shaded cells imply a weight close to zero. We focus on two models, the CISS and CISS-CC models coupled with the pooled Horseshoe prior (HSP).⁴

It turns out that for most countries, our conjecture is confirmed. That is, we observe weights which approach unity if we move out in the tails (i.e., the model becomes more non-linear). When we focus on the center of the distribution, the combination weights approach zero (i.e., the model is linear). Comparing the right and left tails reveals that the estimated weight is often larger for the 5 percent quantiles than the 95 percent quantiles. This indicates that non-linearities are important when our focus is on modeling sharp upswings in GDP growth but become even more important when interest centers on capturing downturns in GDP growth that are extreme (i.e., below the 5 percent quantile).

When we compare the model which does not utilize cross-country information (CISS) to

⁴The results for the remaining specifications look similar and may be found in the Appendix.

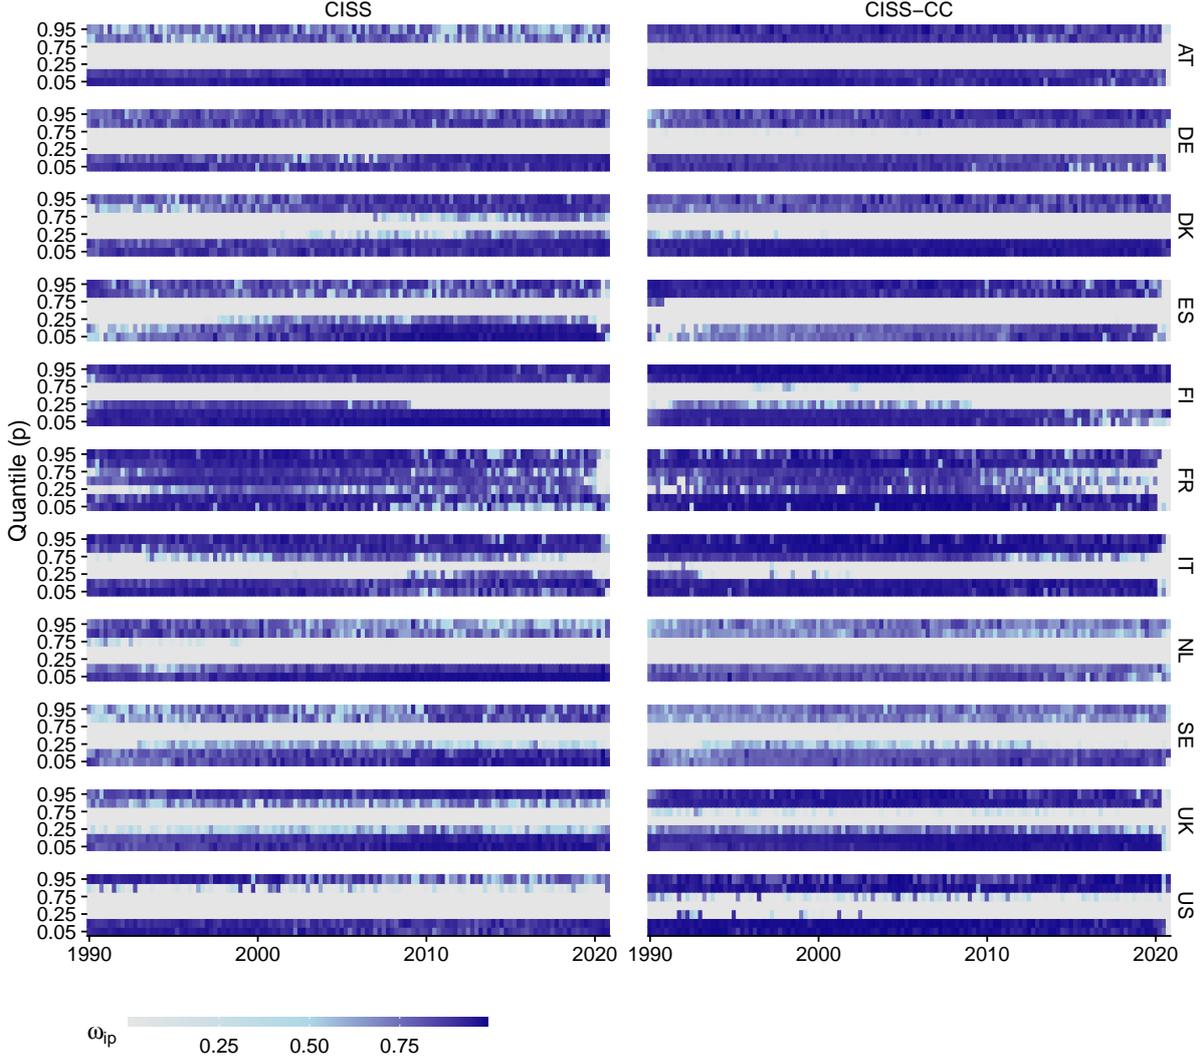


Figure 1: Non-parametric weights ω_{ip} for $\Lambda \neq 0$ with the HSP prior plotted over time with respect to the training/holdout samples ($h = 1$).

the one which explicitly includes cross-sectional data (CISS-CC), we find only modest differences in combination weights. These differences are mainly related to somewhat smaller weights on the BART specification in the upper tail of the distribution (for some selected countries such as AT, DK, the UK, and the US), but the main finding that, in the center of the distribution, our model assigns no weight to the non-linear model still holds.

Zooming into the different results reveals that most countries share the general dynamics described in the previous paragraphs (i.e., ω_{ip} close to one in the tails and $\omega_{ip} \approx 0$ in the center of the distribution). One exception to this broad-based finding is France, which displays larger combination weights across all quantiles. In addition, there exists temporal heterogeneity. For instance, in several countries we observe that, after the global financial crisis (and sometimes slightly earlier), combination weights decrease markedly in the upper tails of the distribution.

Table 3: Time averages of variance decompositions, ω sampled, CC model and HSP prior.

Quantile (p)	ALL	AT	DE	DK	ES	FI	FR	IT	NL	SE	UK	US
0.05	0.83	0.85	0.84	0.79	0.83	0.77	0.86	0.84	0.84	0.86	0.87	0.82
0.1	0.73	0.77	0.72	0.63	0.77	0.61	0.80	0.76	0.73	0.76	0.77	0.75
0.25	0.55	0.63	0.59	0.42	0.61	0.38	0.56	0.61	0.56	0.53	0.61	0.58
0.5	0.67	0.83	0.68	0.30	0.95	0.21	0.98	0.90	0.56	0.40	0.89	0.71
0.75	0.70	0.79	0.76	0.50	0.77	0.48	0.80	0.81	0.69	0.59	0.78	0.69
0.9	0.79	0.83	0.85	0.68	0.81	0.65	0.85	0.84	0.79	0.75	0.82	0.77
0.95	0.89	0.91	0.91	0.82	0.90	0.82	0.92	0.91	0.89	0.87	0.92	0.87

3.4 The role of the common volatility factor

In this sub-section, we investigate the common volatility factor across quantiles. In a first step, we assess the relevance of the common factor volatility specification by considering time averages of variance decompositions. These are computed, by taking the Gaussian representation of the ALD (the distribution of $\epsilon_{ip,t}$ in Equation (1)), as follows:

$$\text{VD}_{ip,t} = \frac{\lambda_{ip}^2 e^{h_{pt}}}{\lambda_{ip}^2 e^{h_{pt}} + \text{Var}(\epsilon_{ip,t})},$$

with $\text{Var}(\epsilon_{ip,t})$ denoting the variance of $\epsilon_{ip,t}$. This decomposition provides information on the share of variation in the shocks (conditional on the quantile) that is explained through the common factor (similar to [Stock and Watson \(2005\)](#)).

Table 3 reports time averages of variance decompositions resulting from the CC-HSP model. Interestingly, for most countries the commonality is larger and more substantial in the tails than at the center of the distribution, and a bit larger in the right than in the left tail. These larger contributions in extreme periods can be traced back to the fact that several of the recessions in our hold-out period can be viewed as shocks with a pronounced global dimension (such as the global financial crisis or the Covid-19 pandemic) and the factor is picking this up.

Across countries, we find a considerable degree of homogeneity within country groups. For instance, Finland, Denmark, and Sweden feature commonalities that are very pronounced in the tails but decline once we approach the center of the distribution both from left and right. The US and the UK share a rather similar pattern in terms of commonalities (high shares in the tails and for the median, smaller shares for the quantiles in between).

The heterogeneity across quantiles in the role of the volatility factor is further supported by Figure 2, which reports estimates of the factors (upper panels) and associated log-volatility per quantile (lower panels). In the upper panel, we observe that especially in the tails the factor moves sharply during global events such as the global financial crisis and the Covid-19 pandemic.

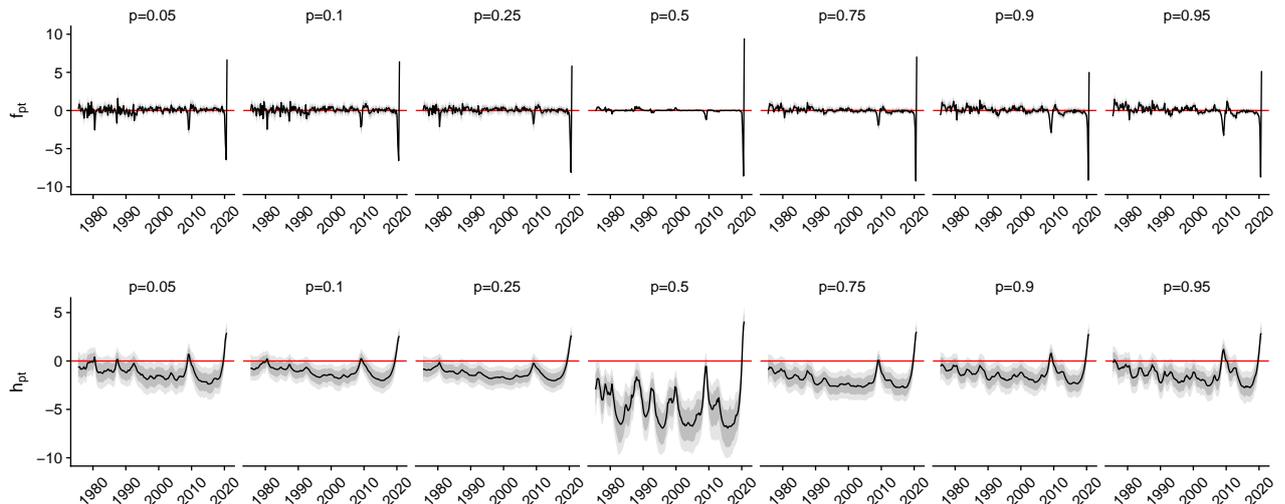


Figure 2: Estimates of the factors and associated log-volatility per quantile, HSP prior.

To a somewhat smaller extent the results also suggest declines in the beginning of the 1990s and the early 2000s. When we focus attention on the 50 percent quantile we find strikingly different results. In the center of the distribution, the factor is small and very close to zero throughout the sample. During the pandemic we find a strong pronounced decrease in 2020:Q2, which was triggered by an unprecedented downturn in real activity globally but also a strong increase in 2020:Q3 (which was accompanied with sharply increasing GDP growth rates throughout all our countries).

Turning to the evolution of the log-volatilities in the lower panel generally yields consistent insights with the findings discussed for the level of the factor. The log-volatility spikes during recessions (i.e., in the early 1990s, 2008-2009, and 2020), and for $p = 0.5$ the level of the log-volatility is much smaller than for the other quantiles but then exceeds the increases in volatility observed for the other quantiles of the distribution. This finding also sheds light on why the amount of variation explained through the factor for most countries is lowest but still sizable in the 50 percent quantile. In most periods, the volatility factor is small (around -5 to -10 on the log-scale) but then during the pandemic it rapidly increases and reaches values of around 5 on the log-scale. This suggests that in tranquil periods, the factor only explains little variation in the shocks but in recessions (or turbulent times) this share increases appreciably and approaches 1.

3.5 Generalized impulse responses to a global business cycle shock

The discussion on the qualitative and quantitative properties of the estimated factor provides evidence that it can be interpreted as a global business cycle shock since, depending on the quantile adopted, it closely tracks events such as global recessions. Following Stock and Watson (2005), we now consider how changes to the factor, labeled factor shocks, impact GDP growth across countries and quantiles.

The posterior quantiles of the generalized impulse response functions (GIRFs) for a common factor shock as estimated with the CISS-CC-HSP model are reported in Figure 3. This figure includes the GIRFs for our model with ω_{ip} estimated (gray shaded areas) and for $\omega_{ip} = 0$ (solid blue lines).

A first interesting finding is that, for all countries, a factor shock has different effects in the left tail than the right tail. In both tails, growth is negatively affected, confirming that higher volatility/uncertainty is detrimental for growth, but the size of the effect (and persistence of the negative effect) is much larger in the left tail. Moreover, notice that for the right tail we also observe an overshoot in real activity in response to an adverse business cycle shock.

Second, when we consider the left tail, the BART piece with an estimated weight tends to mitigate the effects of the shock. This is most likely driven by the fact that, if we rule out non-linearities, there is more to be explained through the factor model and this might translate into factor dynamics which not only pick up business cycle shocks but also soak up information left in the error term potentially arising from ignoring non-linear dynamics between GDP growth and the CISS.

A third striking pattern is the pronounced degree of cross-country heterogeneity in the 5 percent quantile (and, to a somewhat lesser extent, in the 10 percent quantile). When our focus is on the left tail, we observe that France, Italy, the UK, and Spain exhibit sharp declines in GDP growth. Once we consider higher quantiles the GIRFs become much more similar across countries. For instance, we find only modest differences if we focus on $p = 0.5$.

To sum up, we find that the countries in our sample display pronounced reactions to an international business cycle shock. These reactions differ not only across quantiles but also across countries.

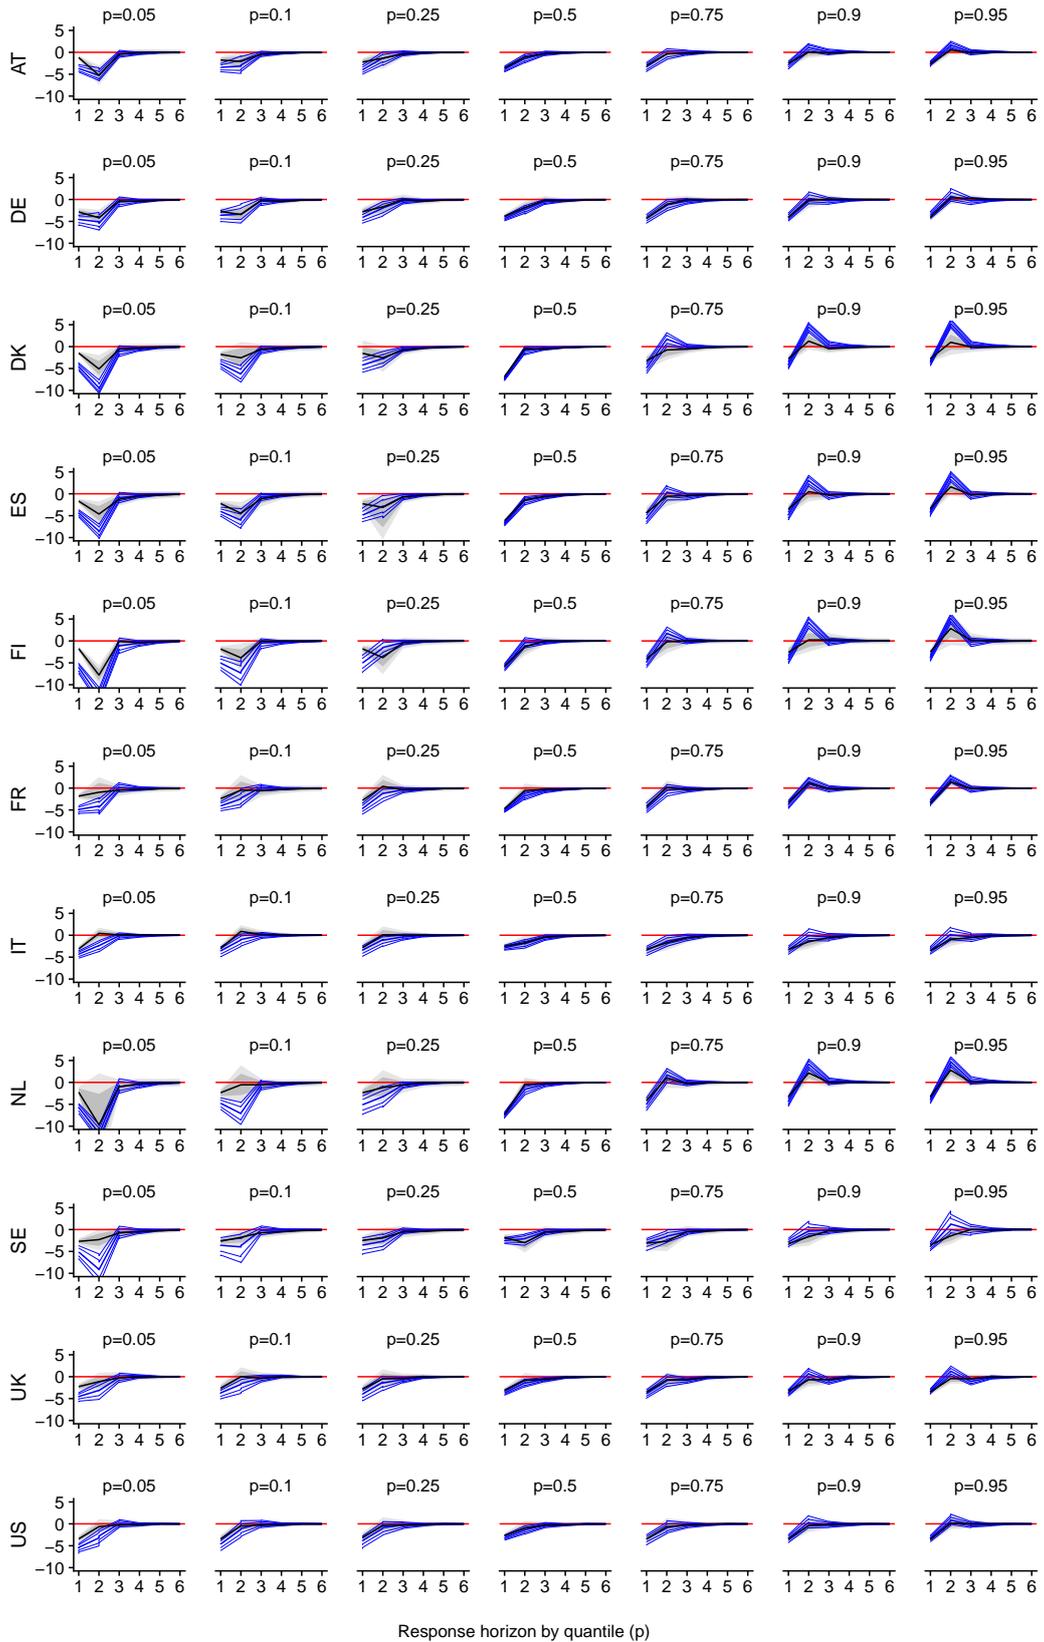


Figure 3: Generalized impulse response functions to a unit shock in f_{pt} across countries and quantiles (average over time). CISS-CC with estimated ω in grey (16th, 84th percentiles), parametric CISS-CC with $\omega = 0$ in blue (16th and 84th percentiles), HSP prior. The black and dark blue lines refer to the posterior median, respectively.

3.6 Generalized impulse responses to a US financial conditions shock

The previous sub-section emphasized that our latent factor can be interpreted as a global business cycle shock. In this sub-section we will instead focus attention on the international effects of a shock to US financial conditions and whether the real effects of such a shock differ from the ones arising from changes in f_{pt} .

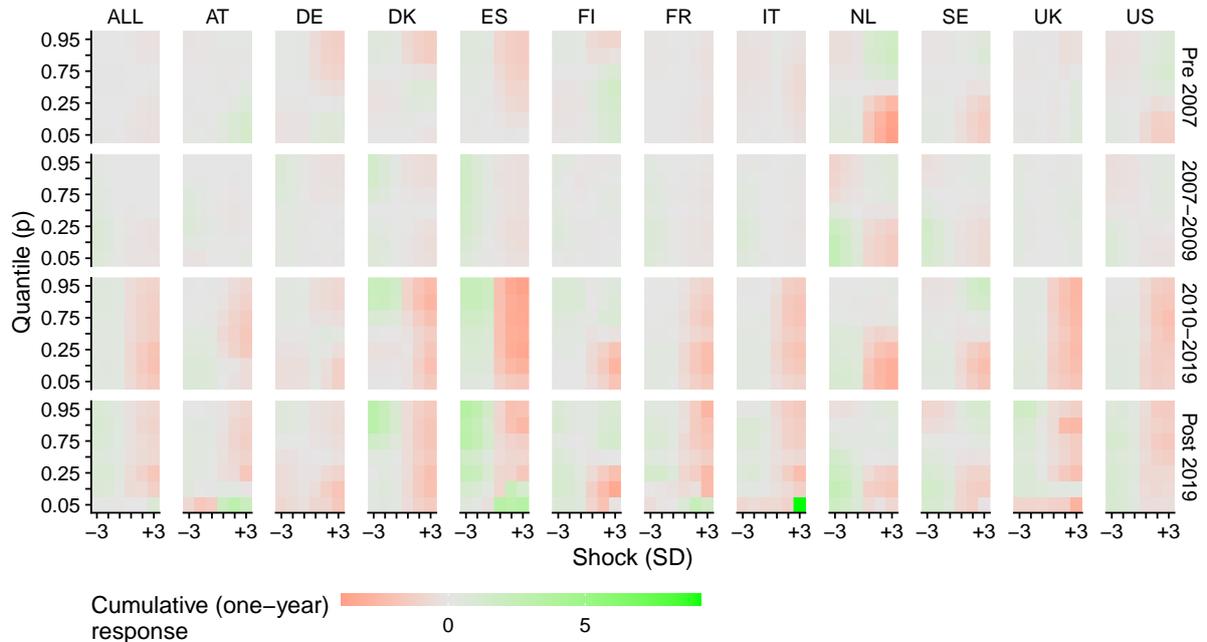


Figure 4: Cumulative (one-year ahead) generalized impulse response functions of GDP by quantile (p) to a $\{-3, -2, -1, 1, 2, 3\}$ in-sample standard deviation US financial conditions shock. Model: CISS-CC, $\Lambda \neq 0$ and estimated ω , HSP prior. Difference between conditional and unconditional forecast across countries (grouped average over time by indicated period).

Figures 4 and 5 report the posterior median of the cumulative (one-year ahead) generalized impulse response functions of GDP by quantile (p) to a $-3, -2, -1, 1, 2$ and 3 in-sample standard deviation US financial conditions shock, based on the CISS-CC-HSP model with estimated weight ω and either with (Figure 4) or without (Figure 5) the common factor in the model’s innovation component. Recall that the CISS is defined so that higher values represent tighter financial conditions; a positive shock may be expected to reduce GDP growth.

As the model is non-linear, the sign and size of the shocks can matter to determine the effects (i.e., contrary to the linear case, the effects are not proportional to the size of the shock, or symmetric). Hence, we have recomputed the model for various sub-samples to analyze how different global business cycle conditions impact the estimates of the GIRFs.

Comparing the two figures shows that the factor volatility has little effect on the results, which is not surprising as it should not affect much the point estimates of the GIRFs (rather

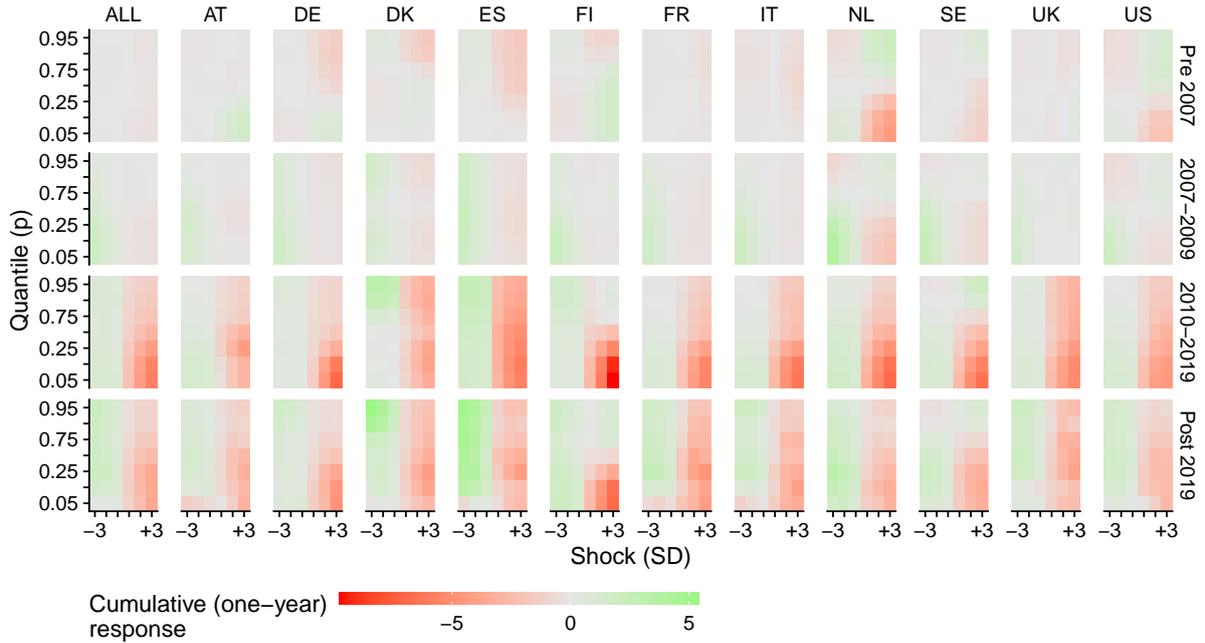


Figure 5: Cumulative (one-year ahead) generalized impulse response functions of GDP by quantile (p) to a $\{-3, -2, -1, 1, 2, 3\}$ in-sample standard deviation US financial conditions shock. Model: CISS-CC, $\Lambda = 0$ and estimated ω , HSP prior. Difference between conditional and unconditional forecast across countries (grouped average over time by indicated period).

their precision). In both cases, a financial shock in the US spills over to other countries. There is asymmetry in the sense that a positive shock affects the growth quantiles, whereas a negative shock's effects are not as sharp. Moreover, the effects and asymmetry are sharper in the 2010-2019 period than earlier.

Analyzing cross-country differences provides additional interesting insights. For some countries (DE, FR, DK, and ES), we find substantial time variation in the GIRFs. Prior to 2010, the corresponding heatmaps feature a great deal of gray colored cells, implying no reactions at all. After the global financial crisis, US-CISS shocks have pronounced effects for these countries that are mostly located in the left tail of the distribution of GDP growth. For other countries (IT, UK, US, and NL), we find less evidence in favor of time-variation in the GIRFs. In these countries, positive (negative) shocks to the US-CISS have negative (positive) effects on GDP growth for $p \in \{0.05, 0.1\}$. Notice, however, that when we consider $p \in \{0.9, 0.95\}$, the effect of a CISS shock seems to reverse sign; a positive shock to the CISS has positive effects on GDP growth if it is already historically high and a negative shock triggers a decline in GDP growth.

Figures 6 and 7 are similar to Figures 4 and 5 but now the weight ω is set to zero, so that the quantile part is linear. Here, too, the common (volatility) factor component does not seem to have an obvious effect on the results. Instead, the effects of financial conditions are now

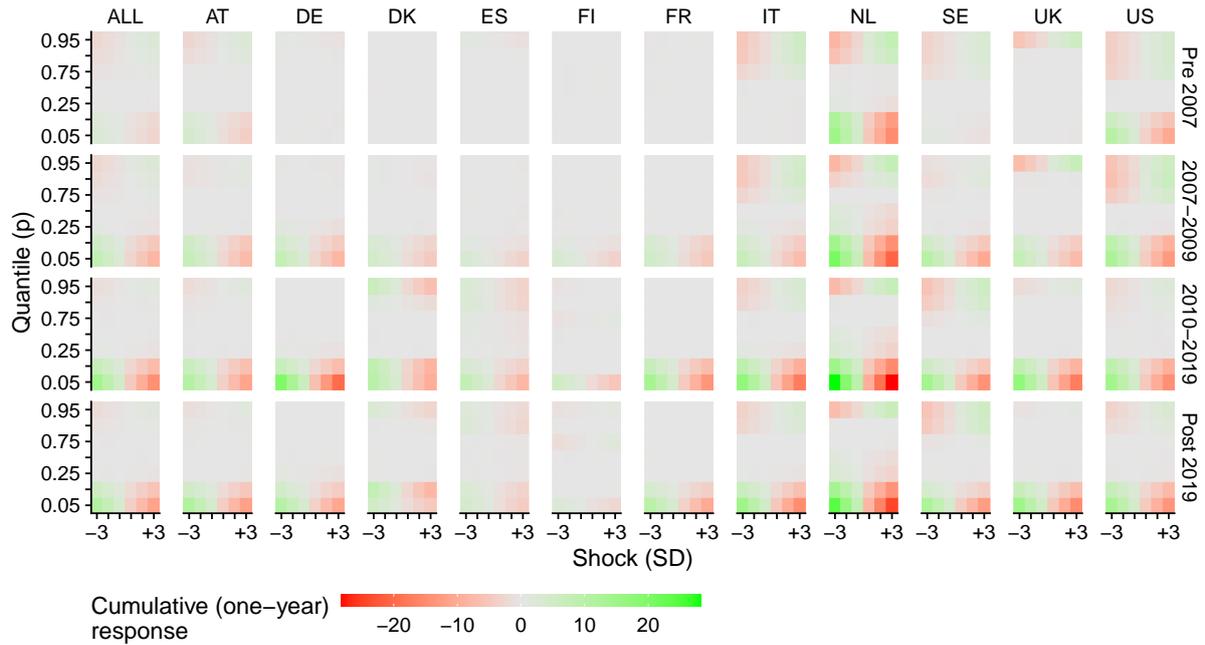


Figure 6: Cumulative (one-year ahead) generalized impulse response functions of GDP by quantile (p) to a $\{-3, -2, -1, 1, 2, 3\}$ in-sample standard deviation US financial conditions shock. Model: CISS-CC, $\Lambda = 0$ and $\omega = \mathbf{0}$, HSP prior. Difference between conditional and unconditional forecast across countries (grouped average over time by indicated period).

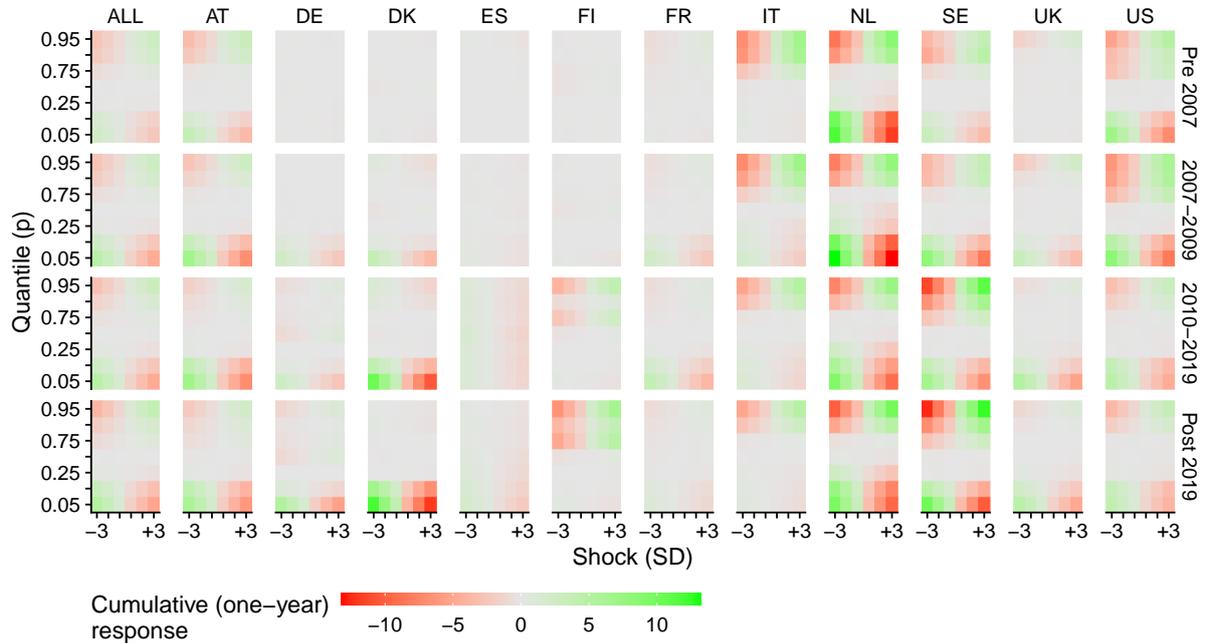


Figure 7: Cumulative (one-year ahead) generalized impulse response functions of GDP by quantile (p) to a $\{-3, -2, -1, 1, 2, 3\}$ in-sample standard deviation US financial conditions shock. Model: CISS-CC, $\Lambda \neq 0$ and $\omega = \mathbf{0}$, HSP prior. Difference between conditional and unconditional forecast across countries (grouped average over time by indicated period).

more linear; positive and negative shocks to financial conditions in the US have similar effects, of opposite sign. In addition, the pattern is clearer in the data since 2007 than before. The effects also look smaller and more stable over time than in the non-linear quantile specification, and in general they are more marked at the lower quantiles. These marked differences in the empirical findings highlight the importance of allowing for non-linearities also in the context of quantile regressions.

4 Conclusions

In this paper we propose a non-parametric quantile panel regression model which assumes that the conditional mean is a convex combination of a linear and an unknown non-linear function. We learn the unknown functions using BART, a successful tool closely related to random forests. To decide on how much weight the BART piece should receive in the p^{th} quantile, we estimate it alongside the remaining model parameters. This non-parametric feature enhances model flexibility, especially in the tails. Using cross-sectional information, in addition, enables us to improve predictive accuracy. This is achieved by proposing a novel pooling prior as well as introducing cross-country information directly. To carry out estimation and inference we design a scalable MCMC algorithm and apply the model to investigate "growth at risk" using an international panel of 11 countries.

In terms of empirical results, our proposed models commonly improve on the benchmark single country linear quantile model in recursive growth forecast comparisons, more so in the tails than near the center of the distribution and in particular when estimating the weight ω , in turn allowing for both linear and BART pieces in the model. The estimated combination weight is smaller (i.e., the model is more linear) for the 25 and 75 percent quantiles than in the tails, i.e., for the 5 and 95 percent quantiles. Moreover, some form of international information definitely pays off (either via the new pooling prior, or by outright including non-domestic series). The effects of the common (volatility) factor are also relevant, as it seems to explain a large fraction of the forecast error variance in most countries, in particular in the tails. A shock to this factor, which can be interpreted as an uncertainty shock, has different (stronger negative) effects in the left tail than the right tail of the growth distribution. In the left tail, the BART piece with an estimated weight tends to mitigate a bit the effects of the shock. Finally, a financial shock in the US spills over to other countries. There is asymmetry in the responses in the sense that a positive shock affects the growth quantiles, whereas a negative shock's effects are not as sharp. Moreover, the effects and asymmetry are sharper in the 2010-20 period than earlier. The

responses are instead much more proportional and symmetric in the linear model, highlighting the importance of allowing for non-linearities in the specification of quantile regressions.

References

- Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone (2019), “Vulnerable growth,” *American Economic Review*, 109, 1263–89, <https://doi.org/10.1257/aer.20161923>.
- Adrian, Tobias, Federico Grinberg, Nellie Liang, and Sheheryar Malik (2018), “The term structure of growth-at-risk,” *IMF Working Paper*, 18/180, <https://doi.org/10.5089/9781484372364.001>.
- Bai, Yu, Andrea Carriero, Todd Clark, and Massimiliano Marcellino (2020), “Macroeconomic forecasting in a multi-country context,” *manuscript*.
- Carriero, Andrea, Todd E. Clark, Massimiliano Marcellino, and Elmar Mertens (2021), “Addressing COVID-19 outliers in BVARs with stochastic volatility,” *Federal Reserve Bank of Cleveland Working Papers*, 21-02R, <https://doi.org/10.26509/frbc-wp-202102r>.
- Chipman, Hugh A., Edward I. George, and Robert E. McCulloch (1998), “Bayesian CART model search,” *Journal of the American Statistical Association*, 93, 935–948, <https://doi.org/10.2307/2669832>.
- (2010), “BART: Bayesian additive regression trees,” *The Annals of Applied Statistics*, 4, 266–298, <https://doi.org/10.1214/09-A0AS285>.
- Clark, Todd E., Florian Huber, Gary Koop, Massimiliano Marcellino, and Michael Pfarrhofer (2021), “Tail forecasting with multivariate bayesian additive regression trees,” *FRB of Cleveland Working Paper*, 21-08, <https://doi.org/10.2139/ssrn.3809866>.
- Clark, Todd E., Florian Huber, Gary Koop, Massimiliano Marcellino, and Michael Pfarrhofer (2021), “Tail forecasting with multivariate bayesian additive regression trees,” *FRB of Cleveland Working Paper*, 21-08.
- Cook, Thomas, and Taeyoung Doh (2019), “Assessing macroeconomic tail risks in a data-rich environment,” *Federal Reserve Bank of Kansas City Research Working Paper*, 19-12, <https://doi.org/10.18651/RWP2019-12>.
- De Nicolò, Gianni, and Marcella Lucchetta (2017), “Forecasting tail risks,” *Journal of Applied Econometrics*, 32, 159–170, <https://doi.org/10.1002/jae.2509>.
- Delle Monache, Davide, Andrea De Polis, and Ivan Petrella (2020), “Modeling and forecasting macroeconomic downside risk,” *CEPR Discussion Paper Series*.
- Feldkircher, Martin, Florian Huber, Gary Koop, and M. Pfarrhofer (2021), “Approximate bayesian inference and forecasting in huge-dimensional multi-country vars,” *arXiv*, 2103.04944, <https://arxiv.org/abs/2103.04944>.
- Ferrara, Laurent, M. Mogliani, and J.G. Sahuc (2019), “Real-time high frequency monitoring of growth-at-risk,” Technical report.
- Figueres, Juan Manuel, and Marek Jarociński (2020), “Vulnerable growth in the euro area: Measuring the financial conditions,” *Economics Letters*, 191, p. 109126.
- Gaglianone, Wagner Piazza, and Luiz Renato Lima (2012), “Constructing density forecasts from quantile regressions,” *Journal of Money, Credit and Banking*, 44, 1589–1607, <https://doi.org/10.1111/j.1538-4616.2012.00545.x>.
- Galbraith, John W., and Simon van Norden (2019), “Asymmetry in unemployment rate forecast errors,” *International Journal of Forecasting*, 35, 1613–1626, <https://doi.org/10.1016/j.ijforecast.2018.11.006>.
- Ghysels, Eric, Leonardo Iania, and Jonas Striaukas (2018), “Quantile-based inflation risk models,” *National Bank of Belgium Research Working Paper*, 349, <http://hdl.handle.net/10419/207729>.
- Giglio, Stefano, Bryan Kelly, and Seth Pruitt (2016), “Systemic risk and the macroeconomy: An empirical evaluation,” *Journal of Financial Economics*, 119, 457–471, <https://doi.org/10.1016/j.jfineco.2016.01.010>.
- Gneiting, Tilmann, and Roopesh Ranjan (2011), “Comparing density forecasts using threshold- and quantile-weighted scoring rules,” *Journal of Business & Economic Statistics*, 29, 411–422, <https://doi.org/10.1198/jbes.2010.08110>.
- González-Rivera, Gloria, Javier Maldonado, and Esther Ruiz (2019), “Growth in stress,” *International Journal of Forecasting*, 35, 948–966, <https://doi.org/10.1016/j.ijforecast.2019.04.006>.
- Huber, Florian, Gary Koop, Luca Onorante, Michael Pfarrhofer, and Josef Schreiner (2020), “Nowcasting in a pandemic using non-parametric mixed frequency VARs,” *Journal of Econometrics*, in-press, <https://doi.org/10.1016/j.jeconom.2020.11.006>.
- Huber, Florian, and Luca Rossini (2020), “Inference in Bayesian additive vector autoregressive tree models,” *arXiv*, 2006.16333, <https://arxiv.org/abs/2006.16333>.
- Kastner, Gregor, and Sylvia Frühwirth-Schnatter (2014), “Ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC estimation of stochastic volatility models,” *Computational Statistics & Data Analysis*, 76, 408–423, <https://doi.org/10.1016/j.csda.2013.01.002>.
- Kiley, Michael T. (2018), “Unemployment risk,” *Board of Governors of the Federal Reserve System Finance and Economics Discussion Series*, 2018-067, <https://doi.org/10.17016/FEDS.2018.067>.

- Korobilis, Dimitris (2017), “Quantile regression forecasts of inflation under model uncertainty,” *International Journal of Forecasting*, 33, 11–20, <https://doi.org/10.1016/j.ijforecast.2016.07.005>.
- Korobilis, Dimitris, Bettina Landau, Alberto Musso, and Anthoulla Phella (2021), “The time-varying evolution of inflation risks,” *manuscript*.
- Kozumi, Hideo, and Genya Kobayashi (2011), “Gibbs sampling methods for bayesian quantile regression,” *Journal of statistical computation and simulation*, 81, 1565–1578.
- Makalic, Enes, and Daniel F. Schmidt (2015), “A simple sampler for the horseshoe estimator,” *IEEE Signal Processing Letters*, 23, 179–182, <https://doi.org/10.1109/LSP.2015.2503725>.
- Manzan, Sebastiano (2015), “Forecasting the distribution of economic variables in a data-rich environment,” *Journal of Business and Economic Statistics*, 33, 144–164, <https://doi.org/10.1080/07350015.2014.937436>.
- Manzan, Sebastiano, and Dawit Zerom (2013), “Are macroeconomic variables useful for forecasting the distribution of US inflation?” *International Journal of Forecasting*, 29, 469–478, <https://doi.org/10.1016/j.ijforecast.2013.01.005>.
- (2015), “Asymmetric quantile persistence and predictability: the case of US inflation,” *Oxford Bulletin of Economics and Statistics*, 77, 297–318, <https://doi.org/10.1111/obes.12065>.
- Mitchell, James, Aubrey Poon, and Gian Luigi Mazzi (forthcoming), “Nowcasting euro area GDP growth using quantile regression,” *Advances in Econometrics*.
- Neal, Radford (2003), “Slice sampling,” *Annals of Statistics*, 31, 705–767.
- Pfarrhofer, Michael (2021), “Tail forecasts of inflation using time-varying parameter quantile regressions,” *arXiv*, 2103.03632, <https://arxiv.org/abs/2103.03632>.
- Plagborg-Møller, Mikkel, Lucrezia Reichlin, Giovanni Ricco, and Thomas Hasenzagl (2020), “When is growth at risk?” *Brookings Papers on Economic Activity*, 167 – 229.
- Reichlin, Lucrezia, Giovanni Ricco, and Thomas Hasenzagl (2020), “Financial variables as predictors of real growth vulnerability,” *Deutsche Bundesbank Discussion Paper*, 05/2020, <http://hdl.handle.net/10419/214829>.
- Stock, James H., and Mark W. Watson (2005), “Understanding changes in international business cycle dynamics,” *Journal of the European Economic Association*, 3, 968–1006, <https://doi.org/10.1162/1542476054729446>.
- Taddy, Matthew, and Athanasios Kottas (2010), “A bayesian nonparametric approach to inference for quantile regression,” *Journal of Business and Economic Statistics*, 28, 357–369, <https://doi.org/10.1198/jbes.2009.07331>.
- Yu, Keming, and Rana A Moyeed (2001), “Bayesian quantile regression,” *Statistics & Probability Letters*, 54, 437–447.

Appendix

Further forecast results

	ALL	AT	DE	DK	ES	FI	FR	IT	NL	SE	UK	US	
CISS-HSP	-1.01 1.01 0.98	0.99 1.00 0.97	1.00 1.01 0.99	0.97 0.97 0.98	1.01 1.01 0.98	1.02 1.02 0.98	1.00 1.00 1.00	1.00 0.99 0.97	1.01 1.01 0.96	1.03 1.03 0.96	1.04 1.04 0.98	1.01 1.01 0.97	none
CISS-HS	-1.01 1.01 0.98	0.99 1.00 0.97	1.00 1.00 0.99	0.97 0.97 0.97	1.01 1.01 0.99	1.02 1.02 0.98	1.00 1.00 0.99	0.99 1.00 0.96	1.01 1.01 0.96	1.03 1.03 0.96	1.04 1.04 0.98	1.00 1.01 0.97	
CISS-CC-HSP	-1.01 1.01 0.99	1.00 1.00 0.97	1.01 1.00 1.02	0.98 0.97 0.99	1.02 1.02 0.99	1.00 1.00 0.97	1.03 1.02 1.02	0.99 0.99 0.96	1.07 1.09 0.98	1.02 1.02 0.99	1.01 1.01 0.98	1.03 1.02 1.00	
CISS-CC-HS	-1.01 1.01 0.99	1.00 1.00 0.97	1.01 1.00 1.04	0.98 0.97 0.99	1.02 1.02 1.01	1.00 1.00 0.97	1.03 1.03 1.01	0.99 0.99 0.94	1.08 1.08 1.01	1.02 1.03 1.00	1.01 1.01 0.99	1.03 1.02 1.00	
CISS-HSP	-1.00 1.00 0.98	0.99 0.99 0.97	1.01 1.02 1.01	0.94 0.94 0.96	0.99 0.99 0.98	1.01 1.01 0.99	1.01 1.01 1.00	0.99 0.98 0.97	1.01 1.00 0.96	1.00 1.00 0.96	1.02 1.02 0.97	0.99 0.99 0.98	tails
CISS-HS	-0.99 1.00 0.98	0.99 0.99 0.97	1.02 1.02 1.00	0.94 0.94 0.96	0.99 0.99 0.98	1.01 1.01 0.99	1.01 1.01 1.00	0.98 0.98 0.96	1.00 1.00 0.97	0.99 1.00 0.96	1.02 1.02 0.97	0.99 0.99 0.98	
CISS-CC-HSP	-1.00 1.00 0.99	0.98 0.98 0.96	1.02 1.01 1.05	0.94 0.94 0.97	1.02 1.02 0.99	1.01 1.01 0.99	1.04 1.04 1.02	0.98 0.98 0.96	1.03 1.05 1.00	1.00 1.00 0.99	0.98 0.98 0.97	1.01 1.00 1.00	
CISS-CC-HS	-1.00 1.00 1.00	0.98 0.99 0.96	1.02 1.01 1.06	0.94 0.94 0.98	1.02 1.02 1.00	1.01 1.01 1.00	1.04 1.04 1.02	0.99 0.99 0.95	1.05 1.04 1.01	1.00 1.00 1.00	0.98 0.98 0.99	1.01 1.00 1.01	
CISS-HSP	-1.01 1.01 0.98	1.02 1.03 0.98	1.03 1.03 1.02	0.97 0.97 0.98	1.00 1.00 0.98	0.99 0.99 0.97	1.03 1.03 1.01	1.02 1.02 0.98	1.00 1.00 0.96	1.00 1.00 0.95	1.02 1.02 0.97	1.01 1.00 0.98	left
CISS-HS	-1.01 1.01 0.98	1.02 1.02 0.98	1.03 1.03 1.01	0.96 0.97 0.97	1.00 1.00 0.99	0.99 0.99 0.98	1.03 1.03 1.01	1.02 1.02 0.98	1.00 1.00 0.97	1.00 1.00 0.95	1.02 1.02 0.97	1.00 1.01 0.98	
CISS-CC-HSP	-1.03 1.03 0.99	1.04 1.03 0.98	1.04 1.04 1.03	0.98 0.98 1.00	1.02 1.02 0.99	0.98 0.98 0.96	1.07 1.06 1.04	1.04 1.04 0.98	1.08 1.10 0.98	1.02 1.01 0.98	1.00 1.01 0.98	1.05 1.04 1.02	
CISS-CC-HS	-1.03 1.03 1.00	1.03 1.04 0.99	1.05 1.04 1.04	0.98 0.98 0.99	1.02 1.02 1.00	0.98 0.98 0.96	1.07 1.07 1.04	1.04 1.04 0.97	1.10 1.08 1.00	1.01 1.01 0.99	1.00 1.00 1.00	1.05 1.04 1.04	
CISS-HSP	-1.00 1.00 0.97	0.96 0.96 0.97	0.98 0.98 0.98	0.96 0.95 0.96	1.01 1.01 0.99	1.05 1.05 0.99	0.98 0.98 0.98	0.96 0.96 0.96	1.02 1.02 0.96	1.05 1.05 0.97	1.06 1.06 0.99	1.00 1.00 0.97	right
CISS-HS	-1.00 1.00 0.97	0.96 0.96 0.97	0.98 0.98 0.97	0.95 0.95 0.96	1.01 1.01 0.99	1.05 1.05 0.99	0.98 0.98 0.98	0.96 0.96 0.95	1.02 1.02 0.96	1.05 1.04 0.97	1.06 1.06 0.99	1.00 1.00 0.97	
CISS-CC-HSP	-0.99 0.99 0.98	0.95 0.95 0.95	0.98 0.97 1.02	0.96 0.95 0.98	1.01 1.02 0.99	1.02 1.02 1.00	0.99 0.99 1.00	0.93 0.93 0.94	1.04 1.06 0.99	1.02 1.03 1.01	1.00 1.00 0.98	0.99 0.99 0.98	
CISS-CC-HS	-0.99 0.99 0.99	0.95 0.95 0.94	0.98 0.98 1.04	0.96 0.95 0.98	1.01 1.01 1.01	1.02 1.02 0.99	0.99 0.99 0.98	0.93 0.93 0.92	1.04 1.05 1.02	1.03 1.03 1.02	1.00 1.00 0.98	0.99 0.99 0.98	
CISS-HSP	-1.01 1.01 0.98	1.00 1.00 0.97	0.99 1.00 0.99	0.98 0.98 0.98	1.02 1.02 0.99	1.03 1.03 0.98	1.00 1.00 0.99	1.00 1.00 0.97	1.02 1.01 0.96	1.05 1.04 0.96	1.06 1.06 0.99	1.02 1.01 0.97	center
CISS-HS	-1.01 1.01 0.98	1.00 1.00 0.98	1.00 1.00 0.98	0.98 0.98 0.98	1.02 1.02 0.99	1.03 1.03 0.98	1.00 1.00 0.99	1.00 1.00 0.96	1.02 1.02 0.96	1.04 1.04 0.96	1.06 1.06 0.99	1.01 1.01 0.97	
CISS-CC-HSP	-1.02 1.01 0.99	1.01 1.01 0.97	1.01 1.00 1.01	0.99 0.99 1.00	1.01 1.01 1.00	1.00 0.99 0.96	1.02 1.02 1.00	0.99 0.99 0.97	1.09 1.10 0.97	1.03 1.03 0.99	1.02 1.03 0.99	1.04 1.02 1.01	
CISS-CC-HS	-1.02 1.01 0.99	1.01 1.01 0.97	1.01 1.00 1.02	0.99 0.99 0.99	1.01 1.01 1.01	0.99 0.99 0.95	1.02 1.02 1.01	0.99 0.99 0.94	1.09 1.09 1.00	1.03 1.04 1.01	1.02 1.02 1.00	1.03 1.03 1.02	
	$\frac{1}{3} \parallel \frac{0}{3} \parallel \frac{0}{3}$												

Figure 8: Relative quantile weighted cumulative ranked probability scores (CRPS) for $h = 4$ and models with $\Lambda \neq \mathbf{0}$. The results are benchmarked to ABG with $\Lambda = \mathbf{0}$ and $\omega = \mathbf{0}$. Lower ratios (shaded an green) indicate better performance (and vice versa, shaded in red).

	ALL	AT	DE	DK	ES	FI	FR	IT	NL	SE	UK	US	
CISS-HSP	-0.93 0.92 0.96	0.90 0.87 0.88	1.09 1.07 0.99	1.01 1.00 0.97	0.85 0.81 0.90	1.01 1.01 1.05	0.77 0.76 0.89	0.83 0.85 0.96	1.02 1.03 1.01	0.98 0.94 1.08	0.93 0.92 0.93	1.05 1.02 0.98	QS10
CISS-HS	-0.91 0.92 0.96	0.88 0.88 0.88	1.06 1.07 1.00	1.00 1.00 1.00	0.82 0.81 0.93	1.01 1.01 1.02	0.73 0.75 0.89	0.84 0.87 0.93	1.03 1.03 1.04	0.95 0.94 1.10	0.92 0.92 0.93	1.02 1.02 0.98	
CISS-CC-HSP	-0.86 0.89 1.09	0.93 0.95 1.13	1.07 0.96 1.19	0.94 0.93 1.12	0.77 0.79 1.09	1.05 1.02 1.14	0.69 0.72 0.95	0.74 0.86 1.05	0.94 0.90 1.17	1.00 0.91 1.12	0.84 0.89 1.22	1.08 1.08 1.25	
CISS-CC-HS	-0.89 0.90 1.04	0.91 0.88 1.05	1.11 1.03 1.13	0.93 0.92 1.14	0.76 0.79 1.05	1.02 1.02 1.12	0.71 0.75 0.97	0.86 0.85 0.98	0.92 0.91 1.09	0.99 0.92 1.14	0.84 0.90 1.24	1.30 1.11 1.22	
CISS-HSP	-0.94 0.99 0.98	1.05 1.01 0.98	1.13 1.09 1.04	1.00 0.99 1.01	1.01 0.96 1.00	1.01 1.01 1.02	0.78 0.88 0.91	0.98 0.93 0.96	1.09 1.07 0.98	0.95 0.95 1.00	1.03 1.01 0.93	1.06 1.03 0.98	QS25
CISS-HS	-0.93 0.99 0.99	1.02 1.02 1.00	1.11 1.10 1.02	1.00 1.00 1.02	0.85 0.96 1.02	1.00 1.00 1.02	0.78 0.88 0.93	0.92 0.92 0.98	1.09 1.07 0.99	0.96 0.95 1.01	1.03 1.01 0.95	1.03 1.03 0.98	
CISS-CC-HSP	-0.86 0.96 1.03	1.00 0.99 1.06	1.01 1.00 1.10	0.98 1.00 1.01	0.96 0.92 1.06	1.00 0.98 1.03	0.68 0.85 0.94	0.90 0.93 1.02	0.99 0.97 1.03	0.93 0.88 0.99	0.85 0.98 1.05	1.07 1.05 1.12	
CISS-CC-HS	-0.87 0.96 1.01	1.01 1.03 0.99	1.02 1.00 1.04	1.00 0.99 1.03	0.87 0.93 1.05	0.98 0.98 1.03	0.93 0.86 0.92	0.93 0.92 0.95	0.99 0.98 0.98	0.93 0.88 1.00	0.85 0.99 1.03	1.07 1.04 1.07	
CISS-HSP	-1.06 1.01 0.96	0.96 0.97 0.94	1.07 1.04 0.99	1.00 1.00 0.98	1.11 0.99 0.94	1.04 1.03 0.99	1.19 0.96 0.94	1.05 0.96 0.92	1.06 1.08 0.96	0.97 0.96 0.99	1.16 1.06 0.96	1.07 1.03 0.98	QS50
CISS-HS	-1.05 0.99 0.97	0.95 0.97 0.96	1.06 1.04 0.99	0.99 1.00 0.98	1.12 0.94 0.95	1.03 1.04 1.00	1.17 0.91 0.94	1.04 0.93 0.96	1.06 1.08 0.98	0.97 0.97 1.02	1.16 1.03 0.94	1.04 1.04 0.98	
CISS-CC-HSP	-0.85 0.97 1.02	1.03 0.96 0.98	1.08 1.00 1.05	0.99 0.95 1.00	1.09 0.96 1.01	1.02 0.97 1.04	0.71 0.92 0.99	0.76 0.94 0.98	1.04 0.98 1.00	1.02 0.96 1.06	1.22 1.03 1.06	1.15 1.03 1.09	
CISS-CC-HS	-0.83 0.96 1.01	1.06 0.96 0.95	1.09 1.00 1.00	1.01 0.95 0.98	1.06 0.93 1.03	1.02 0.97 1.00	0.63 0.88 0.97	0.74 0.92 0.93	1.05 0.95 0.99	1.06 0.96 1.04	1.14 1.01 1.14	1.13 1.01 1.11	
CISS-HSP	-1.07 1.00 0.95	1.01 0.98 0.95	1.01 0.97 0.93	0.97 0.96 0.91	1.32 1.19 1.00	1.01 1.00 0.98	1.39 1.07 0.91	1.05 0.94 0.91	0.99 1.02 0.96	0.99 0.98 0.98	1.11 0.99 0.98	0.94 0.92 0.97	QS75
CISS-HS	-1.06 1.00 0.95	0.99 0.98 0.95	0.99 0.97 0.93	0.96 0.96 0.92	1.31 1.20 1.01	1.01 1.00 0.98	1.25 1.06 0.92	1.05 0.93 0.91	0.99 1.02 0.96	0.97 0.98 0.98	1.12 1.00 0.98	0.96 0.92 0.97	
CISS-CC-HSP	-1.14 0.97 1.13	1.15 0.96 1.07	1.17 0.96 1.08	1.01 0.90 1.02	1.29 1.15 1.27	1.06 0.96 1.10	1.12 0.96 1.16	1.04 0.88 1.08	1.20 1.00 1.08	1.17 0.99 1.19	1.25 1.07 1.24	1.10 0.89 1.15	
CISS-CC-HS	-1.12 1.01 1.13	1.14 1.01 1.05	1.14 1.00 1.09	1.02 0.93 1.01	1.34 1.19 1.29	1.08 0.98 1.09	1.02 1.02 1.15	1.05 0.92 1.04	1.15 1.04 1.10	1.19 1.02 1.20	1.21 1.12 1.28	1.08 0.90 1.16	
CISS-HSP	-0.83 0.88 0.89	1.09 1.02 0.96	0.99 0.93 0.93	1.00 0.99 0.99	1.06 0.97 0.80	1.02 1.02 1.00	0.75 0.78 0.77	1.21 0.86 0.90	0.82 0.87 0.87	0.81 0.83 0.86	0.79 0.73 0.89	0.82 0.84 0.95	QS90
CISS-HS	-1.00 0.87 0.90	1.16 1.02 0.95	1.00 0.93 0.94	1.00 0.99 0.99	1.10 0.78 0.83	1.03 1.02 0.99	1.01 0.78 0.79	1.10 0.99 0.94	0.85 0.87 0.87	0.82 0.83 0.87	0.90 0.72 0.91	0.91 0.83 0.94	
CISS-CC-HSP	-1.12 0.90 1.23	1.44 1.02 1.39	1.25 0.90 1.27	1.10 0.98 1.35	1.08 0.89 1.10	1.13 0.97 1.46	0.84 0.71 0.96	1.17 0.93 1.25	1.30 1.10 1.40	1.19 0.89 1.26	1.02 0.85 1.21	1.11 0.84 1.31	
CISS-CC-HS	-1.12 0.85 1.24	1.45 1.02 1.39	1.25 0.92 1.30	1.11 0.97 1.34	1.04 0.73 1.11	1.27 0.95 1.44	0.80 0.75 0.96	1.16 0.81 1.24	1.27 0.99 1.43	1.20 0.89 1.26	1.04 0.75 1.23	1.07 0.88 1.33	
	$\frac{1}{3} \parallel \frac{0}{3} \parallel \frac{0}{3}$												

Figure 9: Relative quantile scores (QSs) for $h = 1$ and models with $\Lambda \neq \mathbf{0}$. The results are benchmarked to ABG with $\Lambda = \mathbf{0}$ and $\omega = \mathbf{0}$. Lower ratios (shaded an green) indicate better performance (and vice versa, shaded in red).

