

Nonlinearities in macroeconomic tail risk through the lens of big data quantile regressions

Jan Prüser¹ | Florian Huber²

¹Department of Statistics, TU Dortmund University, Dortmund, Germany

²Department of Economics, University of Salzburg, Salzburg, Austria

Correspondence

Jan Prüser, Department of Statistics, TU Dortmund, CDI Building, Room 122, 44221 Dortmund, Germany.
Email: prueser@statistik.tu-dortmund.de

Funding information

German Research Foundation (DFG), Grant/Award Number: 468814087; Austrian Science, Grant/Award Number: ZK-35

Summary

Modeling and predicting extreme movements in GDP is notoriously difficult, and the selection of appropriate covariates and/or possible forms of nonlinearities are key in obtaining precise forecasts. In this paper, our focus is on using large datasets in quantile regression models to forecast the conditional distribution of US GDP growth. To capture possible nonlinearities, we include several nonlinear specifications. The resulting models will be huge dimensional, and we thus rely on a set of shrinkage priors. Since Markov chain Monte Carlo estimation becomes slow in these dimensions, we rely on fast variational Bayes approximations to the posterior distribution of the coefficients and the latent states. We find that our proposed set of models produces precise forecasts. These gains are especially pronounced in the tails. Using Gaussian processes to approximate the nonlinear component of the model further improves the good performance, in particular in the right tail.

KEYWORDS

global–local priors, growth at risk, large datasets, nonlinear models, quantile regression

1 | INTRODUCTION

Modeling and predicting the conditional distribution of output growth have attracted considerable academic attention in recent years. Starting at least with the influential paper by Adrian et al. (2019), focus has shifted towards analyzing whether there exist asymmetries between a predictor (in their case financial conditions) and output growth across different quantiles of the empirical distribution. Several other papers (Adams et al., 2021; Adrian et al., 2018; Delle Monache et al., 2020; Ferrara et al., 2019; Figueres & Jarociński, 2020; González-Rivera et al., 2019; Mitchell et al., 2022; Plagborg-Møller et al., 2020; Reichlin et al., 2020) have started to focus on modeling full predictive distributions using different approaches and information sets. However, most of these contributions have been confined to models which exploit small datasets and, at least conditional on the quantile analyzed, assume linear relations between GDP growth and the predictors.

Times of economic stress such as the global financial crisis (GFC) or the Covid-19 pandemic have highlighted that exploiting information contained in many time series and allowing for nonlinearities improves predictive performance in turbulent periods (see, e.g., Huber et al., 2023). Since economic dynamics change in volatile economic regimes, models that control for structural breaks allow for different effects of economic shocks over time or imply nonlinear relations between GDP growth and its predictors that often excel in forecasting applications (see Adrian et al., 2021; Carriero

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2023 The Authors. *Journal of Applied Econometrics* published by John Wiley & Sons Ltd.

et al., 2016; Clark et al., 2023; D'Agostino et al., 2013; Huber et al., 2023; Pfarrhofer, 2022). Moreover, another important empirical regularity is that the set of predictors might change over time. This is because variables which are seemingly unimportant in normal periods (such as financial conditions) play an important role in recessions and yield important information on future behavior of output growth.

This discussion highlights that the effect of predictors on output growth depends on the quantile under consideration and thus appears to be state dependent and modeling the transition might call for nonlinear econometric models. The key challenge, however, is to identify the different determinants of GDP growth across quantiles while taking possible nonlinearities into account. In this paper, we aim to solve these issues by proposing a Bayesian quantile regression (QR), which can be applied to huge information sets and is capable of capturing nonlinearities of unknown form. Our model is a standard QR model that consists of two parts. The first assumes a linear relationship between the covariates and quantile-specific GDP growth, whereas the second component assumes an unknown and possibly highly nonlinear relationship between the two. The precise form of nonlinearities is captured through three specifications. One is parametric and based on including polynomials up to a certain order, whereas the remaining two are nonparametric. Among these nonparametric specifications we include B-splines (see Shin et al., 2020) and Gaussian processes (GPs) (see Williams & Rasmussen, 2006). Both have been shown to work well when it comes to functional estimation and forecasting.

The combination of a linear and nonlinear specification implies that the dimension of the parameter space increases substantially. Since all these models can be cast in terms of a linear regression conditional on appropriately transformed covariates, we can use regularization techniques to decide on whether more flexibility is necessary and which variables should enter the model. We achieve this through several popular shrinkage priors that have excellent empirical properties in large dimensions and are relatively easy to implement. These shrinkage priors enable us to select promising subsets of predictors and the degree of nonlinearities for each quantile separately.

Posterior inference using Markov chain Monte Carlo (MCMC) techniques in these dimensions proves to be an issue because we have to estimate a large-scale regression model for all quantiles of interest. This procedure needs to be repeated a large number of times if we wish to carry out an out-of-sample forecasting exercise. To reduce the computational burden enormously, we estimate the QRs using variational Bayes (VB).¹ This estimation strategy approximates the exact full conditional posterior distributions with simpler approximating distributions. These approximating densities are obtained by minimizing the Kullback–Leibler (KL) distance between some known density q and the exact posterior distribution p . Hence, integration in huge dimensions is replaced by a simpler optimization routine. Our approach is fast and allows for computing all results of our forecasting exercise without the use of high performance computing environments. As compared with, for example, Kohns and Szendrei (2021) and Mitchell et al. (2022), who estimate large QRs under flexible Bayesian shrinkage priors, our techniques are highly scalable and can be applied to models even larger than the ones we consider in this paper.

We apply our techniques to the large dimensional FRED-QD dataset (McCracken & Ng, 2016) and focus on single and multistep-ahead forecasting of US GDP growth over a hold-out period ranging from 1991Q2 to 2021Q3. The different nonlinear models we consider are high dimensional and feature up to over 1000 coefficients per equation.

The empirical results can be summarized as follows. We show that forecasts from models estimated with VB are very similar to the ones obtained from estimating the different models using MCMC. When we consider the forecasts from the large-scale VB QRs, we find that using huge information sets and nonlinear models in combination with priors that introduce substantial shrinkage pays off for tail forecasts. In both tails, forecast improvements relative to a large-scale QR estimated using MCMC techniques and a Minnesota prior are sizable. When we focus on the center of the distribution, the differences become smaller. Once we allow for nonlinearities, we find some improvements in predictive accuracy over the full hold-out period. These improvements are mostly related to the right tail, while in the left tail accuracy gains are smaller. Comparing the different nonlinear specifications reveals that GPs offer the largest improvements vis-à-vis the linear QR. Once we condition on the more volatile second half of the sample (starting in 2006Q2), accuracy premia obtained from nonlinear models increase substantially. In particular, nonlinear models yield appreciable improvements in tail forecasting accuracy during the GFC and the pandemic. This indicates that a successful tail forecasting model should be able to extract important information from huge datasets, while controlling for possibly nonlinear relations during turbulent periods. When we focus on the key properties of the proposed priors, we observe that priors that imply a dense model (characterized by many small coefficients) yield good tail forecasts.

¹For an introduction, see Blei et al. (2017) and an algorithm for QRs is provided in Bufrei (2019).

The paper is structured as follows. The next section introduces the general QR and the scale-location mixture representation to cast the model in terms of a standard generalized additive regression with auxiliary latent variables. We then focus on the different priors used, provide additional details on the nonlinear components of the models, briefly discuss VB, outline how we estimate the posterior distributions of the parameters and latent quantities, and illustrate the computational properties of our approach. Section 3 discusses our empirical findings while Section 4 discusses other applications and extensions of the models we develop in this paper. The final section summarizes and concludes the paper. An appendix in the supporting information includes additional technical details, empirical results, and more precise information on the used dataset.

2 | BAYESIAN ANALYSIS OF GENERAL QRS

2.1 | The likelihood function

In this paper, our goal is to model the dependence between the q^{th} quantile of GDP growth y_t and a panel of K predictors in $\{\mathbf{x}_t\}_{t=1}^T$ with K being huge. The covariates include a wide range of macroeconomic and financial indicators. Possible nonlinearities between y_t and \mathbf{x}_t are captured through a function $g_q(\mathbf{x}_t)$, with $g_q : \mathbb{R}^K \rightarrow \mathbb{R}$. The fact that K is large and the inclusion of nonlinear functions of \mathbf{x}_t implies that the number of parameters is large relative to the number of observations T .

Our workhorse model is the QR developed in Koenker and Bassett (1978). As opposed to the standard QR, our model decomposes the q^{th} quantile function $Q_{q(y_t)}$ into a linear and nonlinear part and a nonstandard error distribution:

$$y_t = \mathbf{x}_t' \boldsymbol{\beta}_q + g_q(\mathbf{x}_t) + \varepsilon_t, \tag{1}$$

where $\boldsymbol{\beta}_q$ is a K -dimensional vector of quantile-specific regression coefficients and ε_t is a shock term with density f_q such that the q^{th} quantile equals zero:

$$\int_{-\infty}^0 f_q(\varepsilon_t) d\varepsilon_t = q.$$

Conditional on the quantile, this model resembles a generalized additive model (GAM); see Hastie and Tibshirani (1987). This specification differs from much of the literature (e.g., Adrian et al., 2019; Carriero et al., 2022; Kohns & Szendrei, 2021; Mitchell et al., 2022) which sets $g_q(\mathbf{x}_t) = 0$ for all t .

We approximate $g_q(\mathbf{x}_t)$ using nonlinear transformations of \mathbf{x}_t :

$$g_q(\mathbf{x}_t) \approx \sum_{m=1}^M \gamma_{qm} z_m(\mathbf{x}_t) = \mathbf{z}_t' \boldsymbol{\gamma}_q, \tag{2}$$

with $\boldsymbol{\gamma}_q = (\gamma_{q1}, \dots, \gamma_{qM})'$, $\mathbf{z}_t = (z_1(\mathbf{x}_t), \dots, z_M(\mathbf{x}_t))'$, and $z_m(\mathbf{x}_t)$ denotes a basis function that depends on \mathbf{x}_t with γ_{qm} denoting the corresponding basis coefficient. This basis function depends on the specific approximation model used to infer the nonlinear effects and our additive representation nests models commonly used in the machine learning literature (such as GPs, splines, and neural networks but also more traditional specifications such as time-varying parameter models). We will discuss the precise specification of z_m (and thus \mathbf{z}_t) in more detail in Subsection 2.3. Here it suffices to note that depending on the specification, M could be very large. For instance, in the GP case, $M = T$ and thus the number of regression coefficients would be $K + T$.

If f_q remains unspecified, estimation of $\boldsymbol{\beta}_q$ and $\boldsymbol{\gamma}_q$ is achieved by solving the following optimization problem:

$$\arg \min_{\{\boldsymbol{\beta}_q, \boldsymbol{\gamma}_q\}} = \sum_{t=1}^T \rho_q(y_t - \mathbf{x}_t' \boldsymbol{\beta}_q - \mathbf{z}_t' \boldsymbol{\gamma}_q),$$

with $\rho_q(l) = l[q - \mathbb{I}(l < 0)]$ denoting the loss function. This optimization problem is straightforward to solve but if $K + M$ is large, regularization is necessary. This motivates a Bayesian approach to estimation and inference.

From a Bayesian perspective, carrying out posterior inference requires the specification of a likelihood and suitable priors. Following Yu and Moyeed (2001), we assume that the shocks ε_t follow an asymmetric Laplace distribution (ALD) with density:

$$f_q(\varepsilon_t) = q(1 - q) \exp(-\rho_q(\varepsilon_t)).$$

The key thing to notice is that the q^{th} quantile equals zero and the parameter q controls the skewness of the distribution. Kozumi and Kobayashi (2011) show that one can introduce auxiliary latent quantities to render the model with ALD distributed shocks conditionally Gaussian. This is achieved by exploiting a scale-location mixture representation (West, 1987):

$$\begin{aligned}\varepsilon_{qt} &= \theta_q v_{qt} + \tau_q \sqrt{\sigma_q v_{qt}} u_t, \\ \theta_q &= \frac{1-2q}{q(1-q)}, \quad \tau_q^2 = \frac{2}{q(1-q)}, \quad v_{qt} \sim \mathcal{E}\left(\frac{1}{\sigma_q}\right), \quad u_t \sim \mathcal{N}(0, 1),\end{aligned}$$

where $\mathcal{E}\left(\frac{1}{\sigma_q}\right)$ denotes the exponential distribution and σ_q is a scaling parameter. Hence, conditional on knowing $v_q = (v_{q1}, \dots, v_{qT})'$, θ_q , τ_q , σ_q and appropriately selecting g_q , the model is a linear regression model with response $\hat{y}_t = y_t - \theta_q v_{qt}$ and Gaussian shocks that are conditionally heteroskedastic. This conditional likelihood will form the basis of our estimation strategy.

To complete the model specification, we assume that $\frac{1}{\sigma_q} \sim \mathcal{G}(c_0, d_0)$, where c_0 is the shape and d_0 is the rate parameter of the Gamma distribution. Both to zero in order to obtain a flat prior. The choice of the prior distribution on β_q and γ_q is essential for our high-dimensional QRs. We discuss different suitable choices in the next section.

2.2 | Priors for the quantile regression coefficients

For the large datasets we consider in this paper, $M + K \gg T$ and thus suitable shrinkage priors are necessary to obtain precise inference. Kohns and Szendrei (2021) and Mitchell et al. (2022) use flexible shrinkage priors in large-scale QRs and show that these work well for tail forecasting. We build on their findings by considering a range of different priors on β_q and γ_q . All these priors belong to the class of so-called global-local shrinkage priors (Polson & Scott, 2010) and have the following general form:

$$\begin{aligned}\beta_q | \psi_{q1}^\beta, \dots, \psi_{qK}^\beta, \lambda_q^\beta &\sim \prod_{j=1}^K \mathcal{N}\left(0, \psi_{qj}^\beta \lambda_q^\beta\right), \quad \psi_{qj}^\beta \sim u, \quad \lambda_q^\beta \sim \pi, \\ \gamma_q | \psi_{q1}^\gamma, \dots, \psi_{qM}^\gamma, \lambda_q^\gamma &\sim \prod_{j=1}^M \mathcal{N}\left(0, \psi_{qj}^\gamma \lambda_q^\gamma\right), \quad \psi_{qj}^\gamma \sim u, \quad \lambda_q^\gamma \sim \pi,\end{aligned}$$

with λ_q^s ($s \in \{\beta, \gamma\}$) denoting a quantile-specific global shrinkage parameter and ψ_{qj}^s are local scaling parameters that allow for nonzero coefficients in the presence of strong global shrinkage (i.e., with λ_q^s close to zero). The functions u and π refer to mixing densities which, if suitably chosen, translate into different shrinkage priors. In this paper, all the priors we consider can be cast into this form but differ in the way the mixing densities u and π are chosen. Since these priors are well known, we briefly discuss them in the main text and relegate additional technical details to the Online Appendix.

We focus on six shrinkage priors that have been shown to work well in a wide variety of forecasting applications (see, e.g., Chan, 2021; Cross et al., 2020; Huber & Feldkircher, 2019; Prüser, 2022).

The first prior we consider is the Ridge prior. The Ridge prior is a special case of a global–local prior with local parameters set equal to 1 and a global shrinkage parameter which follows an inverse Gamma distribution. Formally, this implies setting $\psi_{qj}^s = 1$ for all q, j and $\lambda_q^s \sim \mathcal{G}^{-1}(e_0, e_1)$. The hyperparameters e_0 and e_1 control the tightness of the prior. We set these equal to $e_0 = e_1 = 0.01$. This prior shrinks all coefficients uniformly towards zero and provides little flexibility to allow for idiosyncratic (i.e., variable-specific) deviations from the overall shrinkage pattern.

This issue is solved by estimating the local shrinkage parameters. The Horseshoe (HS, see, Carvalho et al., 2010), our second prior, does this. This prior sets u and π to a half-Cauchy distribution: $\sqrt{\psi_{qj}^s} \sim C^+(0, 1)$ and $\sqrt{\lambda_q^s} \sim C^+(0, 1)$. The HS possesses excellent posterior contraction properties (see, e.g., Armagan et al., 2013; Ghosh et al., 2016; van der Pas et al., 2014). Moreover, it does not rely on any additional tuning parameters.

Another popular global–local shrinkage prior is the Normal-Gamma (NG) prior of Griffin and Brown (2010). This prior assumes that u and π are Gamma densities. More formally, $\psi_{qj}^s \sim \mathcal{G}(\vartheta, \lambda_q^s \vartheta / 2)$ and $\lambda_q^s \sim \mathcal{G}(c_0, d_0)$, with ϑ being a hyperparameter that controls the tail behavior of the prior, and c_0 and d_0 are hyperparameters that determine the overall degree of shrinkage. We set $c_0 = d_0 = 0.01$ and $\vartheta = 0.1$. This choice not only implies heavy global shrinkage on the coefficients but also implies fat tails of the marginal prior of the coefficients after integrating out the local scaling parameters. The Bayesian LASSO is obtained as a special case of the NG prior with $\vartheta = 1$.

Next, we consider the Dirichlet–Laplace prior (Bhattacharya et al., 2015) that assumes the local scaling parameter ψ_{qj}^s is a product of a Dirichlet-distributed random variate $\phi_{qj}^s \sim \text{Dir}(\alpha, \dots, \alpha)$ and a parameter $\tilde{\psi}_{qj}^s \sim \mathcal{E}(1/2)$ that follows an exponential distribution. Hence, the Dirichlet–Laplace prior sets $\psi_{qj}^s = (\phi_{qj}^s)^2 \tilde{\psi}_{qj}^s$. On the global scaling parameters, we use a Gamma distribution $\sqrt{\lambda_q^\beta} \sim \mathcal{G}(K\alpha, 1/2)$ and $\sqrt{\lambda_q^\gamma} \sim \mathcal{G}(M\alpha, 1/2)$. We set $\alpha = \frac{1}{K}$ for the linear term and $\alpha = \frac{1}{M}$ for the nonlinear term.

Finally, we also consider the Minnesota prior as implemented in Carriero et al. (2022). This prior assumes that the local scaling parameters are $\psi_{qj}^s = \frac{\sigma_y^2}{\sigma_i^2}$, where σ_y^2 is the OLS variance of a standard linear regression model and σ_i^2 is the error variance of a linear regression with the i^{th} variable used as the dependent variable. This term serves to control for scaling differences. The global parameter $\lambda_q^s = \lambda_1 \lambda_2$ is often fixed and depends on two terms. The first term (λ_1) controls the overall shrinkage of the prior, whereas the second term (λ_2) controls the rate of shrinkage on variables other than lagged GDP growth. λ_2 is then often set to values smaller than 1 if i marks a variable other than lagged GDP growth or it equals 1 if it is related to the first lag of GDP growth. Carriero et al. (2022) suggest using $\lambda_1 = 0.04$ and $\lambda_2 = 0.25$ and find that this choice leads to accurate tail forecasts. However, given that our models are substantially larger, we also include a version of the prior that estimates λ_1 and λ_2 by placing Gamma priors on both of them. We will refer to it as the Minnesota-Gamma (Minn-Gamma) prior.²

2.3 | Capturing nonlinearities in high-dimensional QRs

In extreme periods such as the GFC or the Covid-19 pandemic, nonlinearities in macroeconomic data become prevalent. We control for this by having a nonlinear part in our QR. As stated in Equation (2), we capture possible nonlinearities in \mathbf{x}_t through nonlinear transformations $z_m(\mathbf{x}_t)$.

The first and simplest nonlinear specification maps \mathbf{x}_t into the space of polynomials. Bai and Ng (2008) capture nonlinearities in macro data through polynomials relying on factor-based predictive regressions. We follow this approach and define the corresponding basis function as follows:

$$\mathbf{z}_t = \left((\mathbf{x}_t^2)', (\mathbf{x}_t^3)', \dots, (\mathbf{x}_t^N)' \right)'$$

Deciding on the order of the polynomial N is a model selection issue and suitable shrinkage priors can be adopted. In our empirical work, we focus on the cubic case. This specification will overweight large movements in \mathbf{x}_t and should thus be suitable for quickly capturing sharp downturns in the business cycle. In this case, the number of basis coefficients is $M = 3K$. The resulting nonlinear model is called *Polynomial-QR*.

Adding cubic terms allows us to capture nonlinearities in a relatively restricted manner. Since the precise form of nonlinearities is typically unknown, the remaining two specifications we consider are nonparametric and only require relatively mild prior assumptions on the form of nonlinear interactions. The first of these two is the B-spline (see, e.g., De Boor, 2001, for a review). B-splines have a proven track record in machine learning and computer science (Shin et al., 2020).

For the B-spline, we assume that each element in \mathbf{x}_t exerts a (possibly) nonlinear effect on y_t that might differ across covariates. This implies that $g_q(\mathbf{x}_t)$ equals:

$$g_q(\mathbf{x}_t) \approx \sum_{k=1}^K \Phi_k(\mathbf{x}_{\bullet, j}) \gamma_{q, k}$$

Here, we let Φ_k denote a $T \times r$ matrix of B-spline basis functions that depend on the j^{th} covariate in $\mathbf{X} = (\mathbf{x}'_1, \dots, \mathbf{x}'_T)'$, $\mathbf{x}_{\bullet, j}$, and r is the number of knots. In this case, the number of nonlinear coefficients is $M = rK$. In our empirical work, we place the knots at the following quantiles of $\mathbf{x}_{\bullet, j}$: $\{0, 0.05, 0.1, 0.25, 0.50, 0.75, 0.90, 0.95, 1\}$, implying that $r = 9$ and thus $M = 9K$. We will henceforth call this model *Spline-QR*.

²It is worth stressing that this hierarchical version of the Minnesota prior is closely related to the Ridge prior above if the covariates in \mathbf{x}_t are standardized to have unit variance.

The last specification we consider is the GP regression. GP regression is a nonparametric estimation method that places a GP prior on the function $g_q(\mathbf{x}_t)$:

$$g_q(\mathbf{x}_t) \sim \mathcal{GP}(\mu_q(\mathbf{x}_t), \mathcal{K}(\mathbf{x}_t, \mathbf{x}_t)).$$

The mean function $\mu_q(\mathbf{x}_t)$ is, without loss of generality, set equal to zero and $\mathcal{K}(\mathbf{x}_t, \mathbf{x}_t)$ is a kernel function that encodes the relationship between \mathbf{x}_t and \mathbf{x}_t for $t, t = 1, \dots, T$. It is worth noting that our additive specification implies that if the mean function is set equal to zero, the model is centered on a standard QR.

Since \mathbf{x}_t is observed in discrete time steps, the GP prior implies a Gaussian prior on $\mathbf{g}_q = (g_q(\mathbf{x}_1), \dots, g_q(\mathbf{x}_T))'$:

$$\mathbf{g}_q \sim \mathcal{N}(\mathbf{0}_T, \mathbf{K}(\mathbf{w})),$$

where $\mathbf{K}(\mathbf{w})$ is a $T \times T$ -dimensional matrix with (t, t) th element $\mathcal{K}(\mathbf{x}_t, \mathbf{x}_t)$. $\mathbf{w} = (w_1, w_2)'$ is a set of hyperparameters that determine the properties of the kernel (and thus the estimated function).

The GP regression is fully specified if we determine the kernel function \mathcal{K} . In this paper, we use the Gaussian (or squared exponential) kernel:

$$\mathcal{K}(\mathbf{x}_t, \mathbf{x}_t) = w_1 \times \exp\left(-\frac{w_2}{2} \|\mathbf{x}_t - \mathbf{x}_t\|^2\right).$$

The hyperparameters \mathbf{w} are set according to the median heuristic proposed in Arin et al. (2017).

What we discuss above is the function-space view of the GP regression. An alternative way of expressing the GP is the so-called weight-space view. The weight-space view is obtained by integrating out \mathbf{g}_q , yielding the following regression representation:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta}_q + \mathbf{Z}\boldsymbol{\gamma}_q + \boldsymbol{\varepsilon},$$

with \mathbf{y} denoting the stacked dependent variables, \mathbf{Z} is the lower Cholesky factor of \mathbf{K} , and $\boldsymbol{\gamma}_q \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_T)$. Notice that $\mathbf{g}_q = \mathbf{Z}\boldsymbol{\gamma}_q$. Hence, the Cholesky factor of the kernel matrix provides the basis functions, and the parameters can be readily estimated. In this case, the number of nonlinear coefficients is $M = T$. Since we use a shrinkage prior on $\boldsymbol{\gamma}_q$, the corresponding implied kernel is given by $\mathbf{Z}\mathbf{B}_q^y\mathbf{Z}'$. The $M \times M$ matrix \mathbf{B}_q^y is a prior covariance matrix with $\mathbf{B}_q^y = \lambda_q^y \times \text{diag}(\psi_{q1}^y, \dots, \psi_{qM}^y)$. Approximating g_q using GPs leads to the GP-QR specification.

This completes our choice of nonlinear techniques used in the big data QR. Alternative choices (such as allowing for time-varying parameters, neural networks or Bayesian additive regression trees) can be straightforwardly introduced in this general framework.

2.4 | A brief introduction to VB

The high dimensionality of the state space calls for alternative techniques to carry out posterior inference. We opt for using variational approximations to the joint posterior density. In this section, we provide a discussion on how VB works in general. For an excellent in-depth introduction, see Blei et al. (2017). In machine learning, variational techniques have been commonly used to estimate complex models such as deep neural networks (see, e.g., Polson & Sokolov, 2017). In econometrics, recent papers use VB in huge dimensional multivariate time series models such as VARs (Chan & Yu, 2022; Gefang et al., 2023) or state space models to speed up estimation (Koop & Korobilis, 2023b). In a recent paper, Korobilis and Schröder (2022), propose a QR factor model and estimate it using VB techniques.

To simplify the exposition, we fix the prior variances. The appendix provides information on how we estimate the prior variances (and associated hyperparameters) using VB. Let $\boldsymbol{\xi}_q = (\boldsymbol{\beta}_q, \boldsymbol{\gamma}_q, \sigma_q, \mathbf{v}_q)$ denote a generic vector which stores all unknowns of the model, with $\mathbf{v}_q = (v_{q1}, \dots, v_{qT})$ denoting the latent components.

Our aim is to approximate the joint posterior distribution $p(\boldsymbol{\xi}_q | \mathbf{y})$ using an analytically tractable approximating distribution $q(\boldsymbol{\xi}_q)$. This variational approximation is found by minimizing the KL distance between p and q . One can show that minimization of the KL distance is equivalent to maximizing the evidence lower bound (ELBO) defined as follows:

$$\text{ELBO} = \mathbb{E}_{q(\boldsymbol{\xi}_q)}(\log p(\boldsymbol{\xi}_q, \mathbf{y})) - \mathbb{E}_{q(\boldsymbol{\xi}_q)}(\log q(\boldsymbol{\xi}_q)), \quad (3)$$

with $\mathbb{E}_{q(\xi_q)}$ denoting the expectation with respect to $q(\xi_q)$. This implies that finding the approximating density q replaces the integration problem (which is typically solved through MCMC sampling) with an optimization problem (which is fast and thus scales well into high dimensions).

A common and analytically tractable choice of approximating densities assumes that $q(\xi_q)$ is factorized as follows:

$$q(\xi_q) = \prod_{s=1}^S q_s(\xi_{qs}),$$

where ξ_{qs} denotes a partition of ξ_q . A particular example (which we use in this paper) would specify $\xi_{q1} = (\beta'_q, \gamma'_q)'$, $\xi_{q2} = \sigma_q$ and $\xi_{q3} = \nu_q$.

This class is called the mean field variational approximation and assumes that the different blocks ξ_{qs} are uncorrelated.³ Notice that all our priors on ξ_q can be written as follows:

$$p(\xi_q) = \prod_{s=1}^S p(\xi_{qs}),$$

and using the fact that

$$\mathbb{E}_{q(\xi_q)}(\log p(\xi_q, \mathbf{y})) = \mathbb{E}_{q(\xi_q)}(\log p(\mathbf{y}|\xi_q)) + \sum_{s=1}^S \mathbb{E}_{q(\xi_q)}(\log p(\xi_{qs})),$$

the ELBO can be stated as follows:

$$\text{ELBO} = \mathbb{E}_{q(\xi_q)}(\log p(\mathbf{y}|\xi_q)) + \sum_{s=1}^S \mathbb{E}_{q(\xi_q)}(\log p(\xi_{qs})) - \sum_{s=1}^S \mathbb{E}_{q(\xi_q)}(\log q(\xi_{qs})).$$

Wand et al. (2011) prove that under the variational family, the optimal approximating densities are closely related to the full conditional posterior distributions:

$$q_s^*(\xi_q) = \exp \left[\mathbb{E}_{q(\xi_q)}(\log p(\xi_{qs}|\mathbf{y}, \xi_{q,-s})) \right],$$

where $\xi_{q,-s}$ is the vector ξ_q with the s^{th} component excluded. Hence, if $p(\xi_{qs}|\mathbf{y}, \xi_{q,-s})$ is known (which is the case for the QR regression based on the auxiliary representation discussed in the previous subsection), the elements in ξ_{qs} can be updated iteratively (by conditioning on the expected values of $\xi_{q,-s}$) until the squared difference of the ELBO or of all elements of ξ_{qs} is smaller than some small ϵ between two subsequent iterations.

2.5 | Approximate Bayesian inference in general QRs

In this section, we briefly state the three approximating densities ($q_s^*(\xi)$) used to estimate the parameters and latent quantities in the QR regression. We provide derivations for the three approximating densities of the three parameter groups: $\tilde{\beta}_q = (\beta'_q, \gamma'_q)'$, σ_q and ν_q in the supporting information Appendix.

We start by discussing the approximating densities for the regression and basis coefficients. A Gaussian distribution approximates the posterior of $\tilde{\beta}_q$:

$$p(\tilde{\beta}_q|\bullet) \approx \mathcal{N} \left(\mathbb{E}(\tilde{\beta}_q), \hat{\Sigma}_{\kappa_q} \right),$$

³Frazier et al. (2022) state that mean field VB approximations might perform poorly in models with a large number of latent variables. However, they also note that the resulting model forecasts could still perform well in practice.

with variance and mean given by, respectively,

$$\hat{\Sigma}_{\tilde{\beta}_q} = \left[\sum_{t=1}^T \frac{\mathbf{f}_t \mathbf{f}_t'}{\tau_q^2} \mathbb{E} \left(\frac{1}{v_{qt}} \right) \mathbb{E} \left(\frac{1}{\sigma_q} \right) + \mathbf{B}_{0q}^{-1} \right]^{-1},$$

$$\mathbb{E}(\tilde{\beta}_q) = \hat{\Sigma}_{\tilde{\beta}_q} \left[\mathbb{E} \left(\frac{1}{\sigma_q} \right) \sum_{t=1}^T \mathbb{E} \left(\frac{1}{v_{qt}} \right) \frac{\mathbf{f}_t \left(y_t - \theta_q \left[\mathbb{E} \left(\frac{1}{v_{qt}} \right) \right]^{-1} \right)}{\tau_q^2} \right].$$

$\mathbf{f}_t = (\mathbf{x}_t', \mathbf{z}_t')'$ and $\mathbf{B}_{0q}^{-1} = \text{diag}(\mathbf{B}_q^\beta, \mathbf{B}_q^\gamma)^{-1}$ is a prior precision matrix with $\mathbf{B}_q^\beta = \lambda_q^\beta \times \text{diag}(\psi_{q1}^\beta, \dots, \psi_{qK}^\beta)$ and $\mathbf{B}_q^\gamma = \lambda_q^\gamma \times \text{diag}(\psi_{q1}^\gamma, \dots, \psi_{qK}^\gamma)$. The approximating densities used to estimate the prior hyperparameters are provided in Section 1 of the supporting information Appendix.

The latent variable v_{qt} follows a generalized inverse Gaussian (GIG) distribution: $\text{GIG}(r, A, B)^4$ with

$$p(v_{qt}|\bullet) \approx \text{GIG} \left(\frac{1}{2}, \underbrace{2\mathbb{E} \left(\frac{1}{\sigma_q} \right) + \frac{\theta_q^2}{\tau_q^2} \mathbb{E} \left(\frac{1}{\sigma_q} \right)}_{A_q}, \underbrace{\frac{\mathbb{E} \left(\frac{1}{\sigma_q} \right)}{\tau_q^2} \left[(y_t - \mathbf{f}_t' \mathbb{E}(\tilde{\beta}_q))^2 + \mathbf{f}_t' \hat{\Sigma}_{\tilde{\beta}_q} \mathbf{f}_t \right]}_{B_q} \right).$$

The moments of v_{qt} are given by

$$\mathbb{E}(v_{qt}^j) = \left(\frac{\sqrt{B_q}}{\sqrt{A_q}} \right)^j \frac{K_{1/2+j}(\sqrt{A_q B_q})}{K_{1/2}(\sqrt{A_q B_q})},$$

where K_x denotes the modified Bessel function of the second kind.

Finally, we approximate

$$p \left(\frac{1}{\sigma_q} | \bullet \right) \approx \mathcal{G}(c_{q1}, d_{q1})$$

with

$$c_{q1} = c_0 + 1.5T,$$

$$d_{q1} = d_0 + \sum_{t=1}^T \mathbb{E}(v_{qt}) + \frac{1}{2\tau_q^2} \sum_{t=1}^T \left(\mathbb{E} \left(\frac{1}{v_{qt}} \right) (y_t - \mathbf{f}_t' \mathbb{E}(\tilde{\beta}_q))^2 + 2\theta_q (\mathbf{f}_t' \mathbb{E}(\tilde{\beta}_q) - y_t) + \mathbb{E}(v_{qt})\theta_q^2 + \mathbb{E} \left(\frac{1}{v_{qt}} \right) \mathbf{f}_t' \hat{\Sigma}_{\tilde{\beta}_q} \mathbf{f}_t \right),$$

and $\mathbb{E} \left(\frac{1}{\sigma_q} \right) = \frac{c_{q1}}{d_{q1}}$.

2.6 | Comparing computation times between VB and MCMC

The steps for updating ξ_q , in combination with the updating steps for the priors detailed in the supporting information Appendix, form the basis of our VB algorithm. As stated in Section 1, the key advantage of using VB instead of more precise MCMC-based techniques is computational efficiency. Before we turn to our empirical work, we illustrate this point using synthetic data.

To illustrate the computational merits of employing VB-based approximations, Figure 1 shows the estimation times (in logarithmic scale) for different values of $M + K$ using our VB-based QR (for a specific quantile) and the QR estimated through the Gibbs sampler. The MCMC algorithm is repeated 10,000 times. The figure shows that the computational burden increases lightly in the number of covariates for VB. When we focus on MCMC estimation, the computational

⁴We use the following parametrization of the GIG distribution: $\log(\text{GIG}(x)) \propto (0.5 - 1) \log(x) - \left(Ax + \frac{1}{2} \frac{B}{x} \right)$.




FIGURE 1 Comparison of computation times against the number of covariates $M + K$. Computation times are on the logarithmic scale.

requirements increase sharply in the number of covariates. Especially in our empirical work, where $K + M$ is often above 1000, VB proves to be a fast alternative to MCMC-based quantile regressions. MCMC-based estimation (without leveraging parallel computing facilities) becomes excessively slow (or even infeasible) for forecasting applications (where models have to be estimated sequentially across forecast origins) and if the researcher wishes to estimate a large number of quantiles. In our case, all computations in our empirical work can be carried out on a standard desktop computer. This implies that the models we focus on in this paper can be used for producing forecasts in a timely manner.

3 | FORECASTING OUTPUT GROWTH USING HUGE DIMENSIONAL QRS

In this section, we present our forecasting results. The next subsection provides information on the dataset and the forecasting setup. We then proceed by discussing the results from QRs that exclude the nonlinear part in Subsection 3.2. The question whether nonlinearities are important is investigated in Subsection 3.3, and Subsection 3.4 deals with how forecast accuracy changes over time. Subsection 3.5 discusses the determinants of the different tail forecasts and differences in the shrinkage properties across priors.

3.1 | Data overview and forecasting setup

We use the quarterly version of the McCracken and Ng (2016) dataset. The dataset covers information about the real economy (output, labor, consumption, orders, and inventories), money, prices, and financial markets (interest rates, exchange rates, stock market indexes). All series are transformed to be approximately stationary. The set of variables included in \mathbf{x}_t and their transformation codes are described in Table S1 of the supporting information Appendix. All models we consider also include the first lag of GDP growth.⁵ Forecasts are carried out using direct forecasting by appropriately lagging the elements in \mathbf{x}_t .

Our sample runs from 1971Q1 to 2021Q3, and we use the period 1991Q2 to 2021Q3 as our hold-out period. The forecasting design is recursive. This implies that we estimate all our models on an initial training sample with data until 1991Q1 and produce one-quarter- and four-quarters-ahead predictive distributions for 1991Q2 and 1992Q1, respectively. After obtaining these, we add the next observation (1991Q2) and recompute the models to obtain the corresponding predictive densities for 1991Q3 and 1992Q2. This procedure is repeated until we reach the end of the hold-out period.

As a measure of overall forecasting accuracy, we focus on the continuous ranked probability score (CRPS). The CRPS is a measure of density forecasting accuracy and generalizes the mean absolute error (MAE) to take into account how well a given model predicts higher-order moments of a target variable.

⁵We find that including more lags of GDP growth only has small effects on the empirical results.

The CRPS measures overall density fit. Considering overall CRPSs possibly masks relevant idiosyncrasies of model performance across quantiles. If a decision maker is interested in downside risks to GDP growth, she might value a model more that does well at the critical 5 or 10 percentiles as opposed to the remaining regions of the predictive distribution. To shed light on asymmetries across different predictive quantiles, we focus on the quantile score (QS):

$$QS_{qt} = (y_t - Q_{qt}) \left(q - \mathbf{1}_{\{y_t \leq Q_{qt}\}} \right),$$

where Q_{qt} is the forecast of the q^{th} quantile of y_t and $\mathbf{1}_{\{y_t \leq Q_{qt}\}}$ denotes the indicator function that equals one if y_t is below the forecast for the q^{th} quantile.

The QS can also be used to construct quantile-weighted (qw) CRPS scores (Gneiting & Ranjan, 2011). These qw-CRPSs can be specified to put more weight on certain regions of the predictive distribution. In general, the qw-CRPS is computed as follows:

$$qw\text{-CRPS} = \frac{2}{J-1} \sum_{j=1}^{J-1} \omega(\zeta_j) QS_{\mathfrak{s}_j, t},$$

with $\zeta_j = j/J$, $J-1 = 19$ denoting the number of quantiles we use to set up the qw-CRPS and \mathfrak{s}_j selects the j^{th} element from the set of quantiles we consider. This set ranges from 0.05 to 0.95 with a step size of 0.05, and thus, $\mathfrak{s}_1 = 0.05$, $\mathfrak{s}_2 = 0.10$, ..., $\mathfrak{s}_{19} = 0.95$.

We use two weighting functions $\omega(\zeta_j)$ that focus on different regions of the predictive density. These schemes are motivated in Gneiting and Ranjan (2011). The first (CRPS-l) puts more weight on the left tail (i.e., downside risks) and is specified as $\omega(\zeta_j) = (1 - \zeta_j)^2$, while the second (CRPS-t) puts more weight on both tails as opposed to the center of the distribution: $\omega(\zeta_j) = (2\zeta_j - 1)^2$. Notice that if we use equal weights, we obtain a discrete approximation to the CRPS.

3.2 | Results based on linear QRs

We start discussing the QRs that set $g(\mathbf{x}_t) = 0$ for all t . Since we advocate using VB-based estimation techniques our goal is to show, within feasible dimensions, that our approximation-based approach produces forecast densities of output growth which are competitive to the ones obtained by estimating QRs using MCMC. As a benchmark model, we rely on the QR with a Minnesota prior with fixed hyperparameters estimated using MCMC techniques. We use this as a benchmark since Carriero et al. (2022) show that it displays an excellent overall performance in terms of tail forecasting and is simple to implement.

Then, in addition, we also assess whether the additional information in our dataset translates into predictive gains. This is achieved by also including the model proposed in Adrian et al. (2019). To enable comparison to the original model, we estimate it in the same way as in the original paper and call the model henceforth ABG.

Table 1 shows average (over time) qw-CRPSs relative to the QR-Minnesota model. Values smaller than one suggest that a given model outperforms the benchmark, whereas values exceeding unity indicate that the model produces less precise density forecasts. The upper part of the table shows the results for VB-based approximations, while the lower part of the table displays the results for all models estimated through MCMC sampling.

We first address the question of whether VB-based approximations result in less precise forecasts compared with MCMC-based models. The answer to this question is a clear "not at all." The results indicate that the differences in predictive performance between MCMC and VB-based models are minimal, sometimes even approaching zero. While there are instances where VB-based estimates perform slightly worse, such as for the HS, there are also cases where VB-based models surprisingly yield slightly more accurate tail forecasts. Overall, when comparing the upper and lower parts of the table, we find no significant disparities in forecast performance when employing approximations. Furthermore, conducting our extensive forecasting exercise using VB-based models is considerably faster than obtaining predictive densities through MCMC estimation. Even the occasional small losses incurred are negligible considering the ability to generate accurate forecasts of GDP growth rapidly.

When we zoom into the performance of the VB-based models, we find a great deal of heterogeneity with respect to different priors. Popular GL priors such as the HS, the NG, or the DL lead to forecasts that are often slightly worse than the ones obtained from the benchmark Minnesota QR. However, priors such as the Ridge or the LASSO (which is particularly known for overshrinking significant signals, see, e.g., Griffin & Brown, 2010) yield forecasts that are better than the benchmark forecasts for both horizons and across the different variants of the CRPS. The Minn-Gamma prior yields forecasts which are almost identical to the one of the Ridge. This is not surprising given the fact that the prior introduces strong shrinkage in large dimensions, and most coefficients are forced to zero by estimating λ_1 and λ_2 to be close to zero.

TABLE 1 CRPS for linear models.

Model	One-quarter-ahead				Four-quarters-ahead			
	CRPS	CRPS-t	CRPS-l	CRPS-r	CRPS	CRPS-t	CRPS-l	CRPS-r
VB								
HS	1.06	1.11	1.12	1.03	1.06	1.02	1.07	1.02
Ridge	0.88	0.88	0.90	0.86	0.90	0.89	0.90	0.90
NG	1.01	1.04	1.04	1.02	1.07	1.02	1.08	1.03
LASSO	0.92	0.93	0.98	0.86	0.91	0.89	0.91	0.90
DL	1.12	1.18	1.16	1.11	1.16	1.13	1.21	1.09
Minn-Gamma	0.88	0.88	0.90	0.86	0.90	0.89	0.90	0.90
MCMC								
HS	1.01	1.03	1.02	1.02	1.05	1.06	1.04	1.06
Ridge	0.89	0.89	0.91	0.86	0.87	0.86	0.87	0.87
NG	1.03	1.04	1.01	1.06	1.01	0.96	1.03	1.00
LASSO	0.93	0.94	0.95	0.91	0.97	0.97	0.99	0.94
DL	1.17	1.15	1.06	1.26	1.12	1.11	1.22	1.03
Minn-Gamma	0.89	0.89	0.91	0.86	0.93	0.92	0.95	0.89
ABG	1.00	1.04	1.07	0.95	1.03	1.03	1.03	1.03

Note: We highlight in light gray (dark gray) rejection of equal forecasting accuracy against the benchmark model at significance level 10% (5%) using the test in Diebold and Mariano (1995) with adjustments proposed by Harvey et al. (1997). Results are shown relative to the Minnesota prior with fixed hyperparameter and are based on the full sample.

Abbreviation: CRPS, continuous ranked probability score.

It is also worth stressing that the Minnesota prior with fixed hyperparameters introduces too little shrinkage and leads to overfitting. This points towards the necessity to either specify the global shrinkage parameters to depend on K or estimate them from the data. Finally, when we consider the simple ABG model, we corroborate recent results in Carriero et al. (2022), who show that large QRs with shrinkage improve upon the ABG model (with improvements depending on the prior). This is especially pronounced in the case of the Ridge prior. In this case, the accuracy gains vis-à-vis the ABG model reach around 17%, and in most cases, accuracy differences are statistically significant according to the Diebold and Mariano (1995) test.

Turning to the different forecast horizons reveals that specifications that do well in terms of short-term forecasting also produce precise longer-term predictions. For the LASSO-based model, four-quarters-ahead accuracy gains are slightly more pronounced, whereas for the Ridge, we do not find discernible differences across both forecast horizons.

Next, we drill deeper into the quantile-specific forecasting performance by considering Qs for q ranging from $q \in \{0.05, 0.1, 0.25, 0.5, 0.75, 0.95, 0.99\}$. These, for one-step-ahead forecasts, are shown in Figure 2, and Figure 3 provides the four-steps-ahead results. Before starting our discussion, it is worth stressing that many of these differences to the benchmark model are statistically significant with respect to the DM test. The corresponding results are provided in the supporting information Appendix (see Figures S15 and S16). Moreover, for the sake of brevity, we leave out the results for Minn-Gamma from the graphs that follow. The reason is that they are indistinguishable from the ones of the Ridge.

Similar to the findings based on the CRPSs, there is a great deal of heterogeneity across priors. Both the LASSO and in particular, the Ridge prior improves upon the benchmark for all quantiles by relatively large margins. These gains appear to be more pronounced in the tails, reaching up to 15%–20% in terms of the Qs. When focusing on the center of the distribution (i.e., the median forecast), the gains are a bit smaller. In general, the other priors perform considerably worse. The only exception turns out to be the NG prior, which displays an excellent performance in the left tail, while being still outperformed by the LASSO and the Ridge prior.

This brief discussion gives rise to a simple recommendation for practitioners. If interest is placed on producing precise tail forecasts (irrespective of the forecast horizon), it pays off to use large QRs coupled with either a LASSO, Minn-Gamma, or Ridge-type prior. Since the Ridge prior is the simplest (i.e., it only features a single hyperparameter) and the empirical performance is very similar or better to the LASSO and (almost) identical to the Minn-Gamma, our focus from now on will be on comparing the Ridge-based QR with a range of nonlinear specifications.

3.3 | Allowing for nonlinearities in large-scale QRs

In the previous subsection, we have shown that using big QRs leads to tail forecasts that are superior to the ones of the benchmark ABG specification. Conditional on the quantile, these models are linear in the parameters. However, recent




FIGURE 2 One-quarter-ahead quantile scores for different values of q , averaged over the hold-out period.

literature (see, e.g., Clark et al., 2023) suggests that nonlinearities become more important in the tails. Hence, we now address this question within our approximate framework. As discussed in Subsection 3.3, once we control for nonlinearities the number of covariates becomes huge. In this case, carrying out MCMC-based inference becomes impractical (i.e., for around 1000 covariates estimating a model for a specific quantile takes over 15 min with MCMC and under a minute with VB). Hence, after having showed that VB-based forecasts are close to the ones obtained by estimating the models through MCMC, we focus on the comparison between linear QRs and nonlinear extensions, both estimated via VB-based methods.

Table 2 shows relative CRPSs for the different nonlinear models. As opposed to Table 1, all results are now benchmarked against the QR with the Ridge prior. This allows us to directly measure the performance gains from introducing nonlinearities relative to setting $g_q(\mathbf{x}_t) = 0$. Notice that the absence of gray shaded cells in the table indicates that the DM test does not point towards significant differences in forecast accuracy between the linear and the different nonlinear QRs.

Despite this, a few interesting insights emerge from the table. First, many numbers in the table are close to unity and differences are not statistically significant from the best performing linear QR.⁶ This indicates that once we include many predictors, additionally controlling for nonlinearities of different forms only yields small positive (and sometimes negative) gains in terms of tail forecasting accuracy. Second, the first finding strongly depends on the approximation techniques chosen. Among all three specifications, using GPs is superior to using either polynomials or B-splines to approximate the unknown function g_q . Second, and focusing on GP-QR specifications, the specific prior chosen matters appreciably. Whereas the results for the conditionally linear models clearly suggest that the LASSO and Ridge priors are producing the most precise density forecasts. The results for the nonlinear models tell a slightly different story. We observe that the Ridge does well again but, for one-quarter-ahead tail forecasts, is outperformed by the HS. The LASSO, in contrast, is the weakest specification. Since the LASSO is known to overshrink significant signals (see, e.g., Griffin & Brown, 2010), it could be that it misses out important information arising from the GP-based basis functions. Third, and finally, if we consider four-quarters-ahead predictions, the QR coupled with a GP and a Ridge prior becomes the single best performing model again.

⁶For the GP-QR specifications with Ridge prior we obtain p-values between 0.1 and 0.2 using the test in Diebold and Mariano (1995) with adjustments proposed by Harvey et al. (1997).




FIGURE 3 Four-quarters-ahead quantile scores for different values of q , averaged over the hold-out period.

To again gain a better understanding on which quantiles of the predictive distribution drive the CRPSs, Figures 4 and 5 (similar to Figures 2 and 3) show the QSs for different quantiles. These are normalized to the linear QR with Ridge prior so that numbers smaller than one indicate that nonlinearities improve predictive accuracy for a given quantile and numbers exceeding one imply that nonlinearities decrease forecasting accuracy.

In general, both figures tell a consistent story: nonlinearities help in the right tail across both forecast horizons, for all three nonlinear specifications, and for most priors considered. The only exception to this pattern is four-quarters-ahead right tail forecasts of GDP growth when B-splines are used. When there are gains, they are often sizable. For instance, in the case of the QR-GP model, we observe accuracy improvements up to 25% relative to the linear QR model.

When we focus on the left tail, accuracy premia often turn negative. In some cases (such as for GP models with Ridge, NG and HS priors), there are accuracy gains for predicting downside risks, but these gains are only rather small (reaching 5% in the case of the QR-GP regression with a Ridge prior).

To sum up this discussion, our results indicate that linear models that exploit large datasets are sufficient if interest lays in capturing the left tail. In the right tail, it pays off to use nonlinear models. In Subsection 3.5, we investigate the properties of left- and right-tail forecasts in order to shed some light on why the model is doing well in terms of predicting strong rebounds in economic activity.

3.4 | Heterogeneity of forecast accuracy over time

Up to this point, our analysis focused on averages over time. In the next step, we will consider how the forecasting performance changes over the hold-out period. The overall results point towards limited gains of nonlinearities when the full hold-out period is considered. As a starting point, we ask whether nonlinear models improve tail forecasts if we focus only on the second half of the sample (2006Q2 to 2021Q3). This period includes the GFC and the substantial downturn in output during the Covid-19 pandemic.

Table 3 is the same as Table 2 but includes only averages of CRPSs over the 2006Q2 to 2021Q3 period. The table reveals much larger gains from using nonlinear models if the volatile part of the sample is used as our verification period. In

TABLE 2 CRPSs for nonlinear models.

Model	One-quarter-ahead				Four-quarters-ahead			
	CRPS	CRPS-t	CRPS-l	CRPS-r	CRPS	CRPS-t	CRPS-l	CRPS-r
Polynomials								
HS	1.02	1.00	1.08	0.94	0.96	0.97	1.00	0.92
Ridge	0.98	0.94	1.01	0.92	0.96	0.97	1.00	0.92
NG	1.03	0.99	1.07	0.96	0.95	0.96	1.00	0.90
LASSO	1.05	1.04	1.08	1.01	1.02	1.00	0.99	1.04
DL	1.07	1.04	1.15	0.98	1.22	1.20	1.29	0.94
Minn-Gamma	0.98	0.94	1.01	0.92	0.96	0.97	1.00	0.92
B-splines								
HS	1.13	1.16	1.18	1.09	1.10	1.14	1.05	1.18
Ridge	1.08	1.08	1.08	1.08	1.09	1.13	1.04	1.17
NG	1.15	1.17	1.20	1.12	1.13	1.17	1.07	1.22
LASSO	1.07	1.06	1.09	1.05	1.02	1.01	0.99	1.05
DL	0.98	1.00	1.01	0.95	1.04	1.08	1.02	1.09
Minn-Gamma	1.08	1.08	1.08	1.08	1.09	1.13	1.04	1.17
Gaussian Processes								
HS	0.96	0.94	0.97	0.94	1.05	1.02	1.10	0.99
Ridge	0.97	0.95	0.98	0.94	0.96	0.95	0.97	0.94
NG	0.98	0.95	0.98	0.97	1.06	1.02	1.08	1.01
LASSO	1.04	1.04	1.07	1.01	1.01	0.99	1.00	1.00
DL	1.02	0.97	1.00	1.02	1.22	1.20	1.29	1.10
Minn-Gamma	0.97	0.95	0.98	0.94	0.96	0.95	0.97	0.94

Notes: We highlight in light gray (dark gray) rejection of equal forecasting accuracy against the benchmark model at significance level 10% (5%) using the test in Diebold and Mariano (1995) with adjustments proposed by Harvey et al. (1997). Results are shown relative to the linear QR with a Ridge prior and are based on the full sample.

particular, we find significant improvements of up to 7% in the tails for the GP-QR-Ridge model. When we focus on QR with polynomials, we find even stronger gains (which reach around 17% for the right tail), but these are never significant. There are also some nonlinear specifications that improve upon the linear benchmark when the overall CRPS is used (such as GP-QR-Ridge).

When we consider higher-order predictions, gains become even larger, corroborating the findings in Clark et al. (2023). For longer horizons, spline-based models coupled with a HS or Ridge prior yield one-year-ahead left tail forecasts which are around 12% more accurate than the ones of the linear QR model with the Ridge prior.

This discussion shows that if we focus on periods characterized by a rather volatile macro environment, nonlinearities appear to be important. In the next step, we investigate whether a split of the hold-out sample in pre-2006 and post-2006 still masks important temporal idiosyncrasies. To focus on the time evolution of the CRPSs, Figures 6 and 7 show the cumulative CRPSs relative to the linear benchmark QR (with a Ridge prior) for one-quarter and four-quarters-ahead forecasts.

We start by focusing on the one-quarter-ahead forecasts. For this specification, the density forecasting performance is heterogenous over time. In the first part of the sample, models using either polynomials or GPs coupled with a DL prior yield CRPSs that are superior to the linear benchmark. However, these accuracy gains vanish during the GFC. When we put more weight on tail forecasting accuracy (and consider GP-QRs), the gains disappear as early as during the 2001 recession that followed the 9/11 terrorist attacks and the burst of the dot-com bubble.

In the pandemic, we observe a sharp increase in predictive accuracy for several priors (most notably the Ridge and NG priors). This pattern is more pronounced for the weighted variants of the CRPSs. Considering the other nonlinear model specifications gives rise to similar insights. Spline-based approximations to g_q generally perform poorly up until the pandemic. During the pandemic, even this specification improves sharply against the linear benchmark specification. This pattern is particularly pronounced for the GP-QRs.

Considering the performance of the models and priors that did well on average (GP-QRs with Ridge and the HS) reveals that most of these gains are actually driven by a superior performance during turbulent times. Specifically, gains arise from a superior performance during the GFC and the Covid-19 pandemic, as evidenced by declining ratios of CRPSs over these periods.

Turning to four-quarters-ahead forecasts provides little new insights. Models using the DL prior do not excel in the first part of the hold-out period and are generally outperformed by the linear QR. However, accuracy improvements during




FIGURE 4 One-quarter-ahead quantile scores for different values of q , averaged over the hold-out period and normalized to the quantile regression (QR) with a Ridge prior.

the GFC and the pandemic are quite pronounced for splines and the GP-QRs. This pattern is even more pronounced than in the case of the one-quarter-ahead forecasts.

To sum up this discussion, our results indicate that forecast performance is heterogenous over time. Different models such as the Polynomial-QR and the GP-QR with a DL prior outperform in the early part of the hold-out period. This performance premium vanishes during the first two recessions observed in the sample. In contrast, other models such as QR-GP with either the NG or the LASSO do not gain much in tranquil periods but excel during recessions.

3.5 | Properties and determinants of the quantile forecasts

The previous subsections have outlined that QRs and QRs with nonlinear components perform well in terms of tail forecasting. In this subsection, our goal is to investigate which variables determine the quantile forecasts and in what respect successful shrinkage priors differ from their less successful counterparts.

The presence of nonlinearities complicates our investigation since it is not clear how to measure the effect of \mathbf{x}_t on a given quantile of y_t in the presence of nonlinearities. Moreover, given that K is large, it is difficult to understand which variables have an important effect on the quantile forecast of y_t . This issue is further intensified since priors such as the Ridge imply that most elements in \mathbf{x}_t have a small effect and thus the importance of a single variable is difficult to quantify. As a simple solution to both issues, we follow Clark et al. (2022) and approximate the nonlinear, quantile-specific model using a linear posterior summary (see Woody et al., 2021). Specifically, we estimate the following regression model:

$$Q_{q,t} = \mathbf{x}'_t \hat{\boldsymbol{\alpha}}_q + \hat{\varepsilon}_t, \quad \hat{\varepsilon}_t \sim \mathcal{N}(0, \sigma_{\hat{\varepsilon}_q}^2).$$

On the linearized coefficients, we use a HS prior and on the error variances an inverse Gamma prior. To achieve interpretability and decouple shrinkage and selection (see Hahn & Carvalho, 2015), we then apply the SAVS estimator




FIGURE 5 Four-quarters-ahead quantile scores for different values of q , averaged over the hold-out period and normalized to the quantile regression (QR) with a Ridge prior.

proposed in Ray and Bhattacharya (2018) to the posterior mean of $\hat{\alpha}_q$.⁷ This will yield a sparse variant of $\hat{\alpha}_q$ that is easy to interpret and can be understood as the best linear approximation to the corresponding quantile forecast arising from the nonlinear model. We label predictors that survive the sparsification process as robust. Notice that this does not imply that the corresponding forecasting model is sparse. It simply tells us that, conditional on a given shrinkage prior, a given predictor plays an important role in approximating a given quantile.

For brevity, we focus on one-step-ahead forecasts. Results for four-quarters-ahead are included in the supporting information Appendix.

Figure 8 shows the results of this exercise across nonlinear specifications and priors. Starting with the left tail forecasts and linear models suggests that most quantile forecasts are not related to elements in \mathbf{x}_t in a robust manner. There is only one exception. In the case of the QR with the NG prior, we find that real money growth (M1real) survives the sparsification step and the relationship indicates that declines in money growth imply an increase in tail risks (i.e. a decline in GDP growth in the ten percent quantile).

If we focus on nonlinear models, other variables appear to be correlated with forecasts of tail risks. Among the different priors, we find some variables which show up repeatedly. Among these are all nonfarm employees (PAYEMS), money growth, capacity utilization in manufacturing (CUMFNS), and private fixed investment (both residential and nonresidential). Most of these variables are forward looking in nature and thus consistent with our intuition that economic agents form expectations about the state of the economy in the future and thus change their investment decisions accordingly. Notice that the relationship between private fixed investment is particularly pronounced for GP-QRs under the HS, NG, and the DL prior. Another pattern worth mentioning is that the LASSO-based forecasts are generated from sparse models across both linear and nonlinear specifications.

Once we focus on the center of the distribution, we find that forecasts from linear models are driven by one or two variables. Most prominently, specifications that do well in terms of point forecasts (such as the Ridge and LASSO) yield point forecasts that display a strong relationship with (lagged) money growth. In case we adopt a nonlinear specification,

⁷Huber et al. (2021) and Hauzenberger et al. (2021) apply SAVS to multivariate time series models and show that it works well for forecasting.

TABLE 3 CRPSs for nonlinear models: 2006Q2 to 2021Q3.

Model	One-quarter-ahead				Four-quarters-ahead			
	CRPS	CRPS-t	CRPS-l	CRPS-r	CRPS	CRPS-t	CRPS-l	CRPS-r
Polynomials								
HS	0.97	0.93	0.97	0.86	0.88	0.86	0.94	0.80
Ridge	0.91	0.86	0.96	0.83	0.88	0.87	0.94	0.81
NG	0.98	0.92	1.01	0.89	0.87	0.84	0.94	0.76
LASSO	1.03	1.02	1.01	0.99	0.98	0.95	0.95	0.99
DL	1.01	1.05	1.02	0.98	1.02	0.97	1.05	0.86
Minn-Gamma	0.91	0.86	0.96	0.83	0.88	0.87	0.94	0.81
B-splines								
HS	1.02	1.02	1.08	0.95	0.94	0.93	0.88	0.99
Ridge	0.96	0.92	0.95	0.94	0.93	0.92	0.88	0.99
NG	1.02	1.01	1.07	0.97	0.95	0.94	0.89	1.01
LASSO	1.07	1.05	1.09	1.02	0.99	0.97	0.95	1.01
DL	0.90	0.88	0.93	0.85	0.96	0.93	0.94	0.96
Minn-Gamma	0.96	0.92	0.95	0.94	0.93	0.92	0.88	0.99
Gaussian Processes								
HS	0.94	0.91	0.96	0.91	0.97	0.96	1.02	0.91
Ridge	0.95	0.93	0.97	0.92	0.94	0.92	0.95	0.91
NG	0.97	0.93	0.97	0.95	0.99	0.96	1.01	0.94
LASSO	1.06	1.05	1.09	1.01	1.00	0.98	0.99	0.99
DL	1.02	0.96	0.99	1.02	1.05	1.04	1.03	0.96
Minn-Gamma	0.95	0.93	0.97	0.92	0.94	0.92	0.95	0.91

Note: We highlight in light gray (dark gray) rejection of equal forecasting accuracy against the benchmark model at significance level 10% (5%) using the test in Diebold and Mariano (1995) with adjustments proposed by Harvey et al. (1997). Results are shown relative to the linear QR with a Ridge prior.

Abbreviation: CRPS, continuous ranked probability score.

some differences arise across specifications. For polynomials, median forecasts under all priors except the DL are related to very few predictors, with money growth and short-run unemployment showing up for the NG and LASSO models. The DL prior implies a more dense model. This could be a possible reason for the rather weak performance of this specification. When we turn to spline-based models, we again find a similar pattern. Money growth shows up in the case of the NG and LASSO and short-run unemployment predicts median output growth if we adopt a HS, Ridge, or NG prior. Models that capture nonlinearities through GPs, our best performing nonlinear specifications, give rise to a very consistent pattern across priors. In all cases, lagged money growth appears to be a robust predictor of GDP growth, and it impacts GDP growth forecasts negatively.

Finally, when our focus is on right-tail forecasts, all models become much more dense. Variables that have been showing up in the case of left-tail and point forecasts show up again (most notably money growth and short-term unemployment). Additional variables such as initial unemployment claims or prices remain in the sparse model as well. But there is no clear pattern across models, except for the fact that money growth also remains in the set of robust predictors even if much shrinkage is introduced. In general, we can say that as opposed to median forecasts, right tail predictions are driven by more variables. This pattern also holds (for some nonlinear specifications) for left tail forecasts but to a lesser degree.

The analysis based on linearized coefficients provides information on which variables are predictive for output growth forecasts across quantiles. However, the analysis in Subsections 3.2 and 2.3 suggests that differences in forecast performance are driven by the prior. To understand which properties of a given prior exert a positive effect on predictive accuracy, we now focus on the shrinkage hyperparameters of the different priors. Comparing the amount of shrinkage introduced through the different priors is not straightforward. Here, our measure of choice is based on using the rescaled log determinant of the prior covariance matrices as a measure of overall shrinkage for each respective prior. Since all prior covariance matrices are diagonal, this simply amounts to summing over the log of the diagonal elements of \mathbf{B}_{0q} and then normalizing through by the number of diagonal elements. This constitutes a rough measure of overall shrinkage, and we can compute it for each quarter in the hold-out period. Again, we will focus on shrinkage introduced in one-quarter-ahead predictive regressions. The four-quarters-ahead results are qualitatively similar and included in the supporting information Appendix.

Log-determinants of the prior covariance matrices over the hold-out period are depicted in Figure 9. The figure includes (if applicable) solid lines which refer to the amount of shrinkage introduced on the linear coefficients and dashed lines

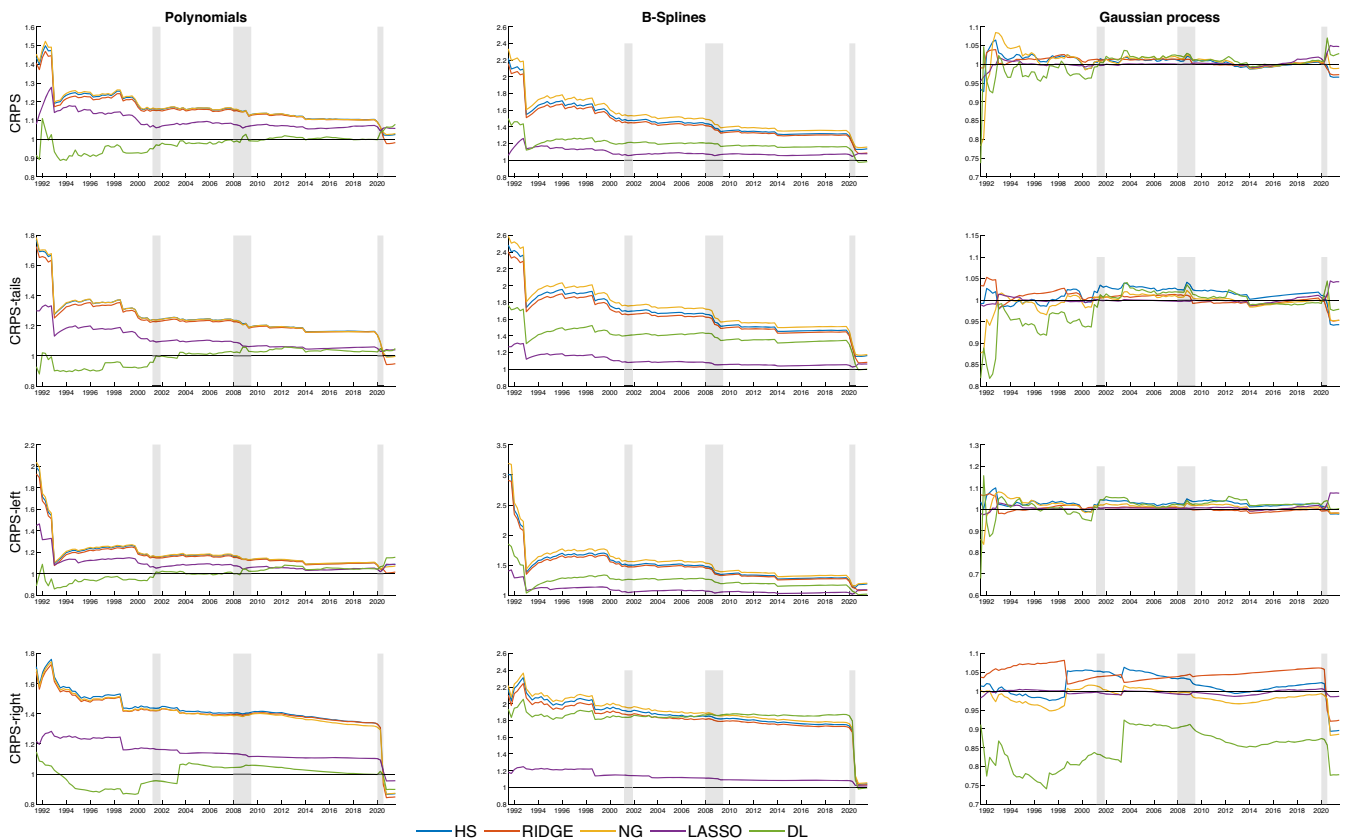


FIGURE 6 Cumulative one-quarter-ahead continuous ranked probability score (CRPS) relative to the linear quantile regression (QR) with the Ridge prior over the hold-out period.

which refer to the log-determinants of the prior covariances that relate to the shrinkage factors on the basis coefficients of the different nonlinear models.

From this figure, a few interesting insights emerge. First, the different priors introduce different degrees of shrinkage. Overall, two priors stand out in terms of the amount of shrinkage they introduce. The first one is the DL. This is rather surprising given the fact that this prior performs worst in the forecasting horse race but also leads to posterior summaries which feature several nonzero coefficients. Our conjecture is that this prior forces the vast majority of coefficients to effectively zero, but several coefficients remain sizable and the corresponding set of variables is still too large causing overfitting issues to arise. The prior that introduces the largest amount of shrinkage is the LASSO. In this case, almost all coefficients are very small. These observations are corroborated by boxplots, included in the supporting information Appendix (see Figures S3 to S6 in the supporting information Appendix), which show the scaling parameters over three subsamples. Our results imply that models which feature a large number of shrunk coefficients provide better forecasts than models featuring many coefficients that are effectively zero and some coefficients that are nonzero and sizeable. This is consistent with findings in Giannone et al. (2021), who provide empirical evidence that macroeconomic data are rather dense as opposed to sparse. Notice that the fact that dense models produce accurate tail forecasts is not inconsistent with our analysis based on linearized posterior summaries. This is because the linearized model under a shrinkage and sparsification approach strikes a balance between achieving a good model fit while keeping the model as simple as possible. Hence, if the covariates in the panel co-move, shrinkage and sparsification techniques will select one of these variables.

Second, in almost all cases, the amount of shrinkage introduced on the nonlinear part of the different models is much larger than the degree of shrinkage on linear coefficients. This holds for most priors, nonlinear methods, and over all time periods. One exception is the spline-QR specification with a DL prior and when the right tail is considered. Interestingly, this specific combination of much stronger shrinkage on the linear part of the model and less shrinkage on the nonlinear part leads to good forecasts in the right tail (see Figure 4).

Third, and finally, there is (with some notable exceptions) relatively little time-variation in the amount of shrinkage over the hold-out period. The only exception are the GP-QRs. In this case, the amount of shrinkage decreases appreciably from 2013 onward.




FIGURE 7 Cumulative four-quarter-ahead continuous ranked probability score (CRPS) relative to linear quantile regression (QR) with a Ridge prior over the hold-out period.

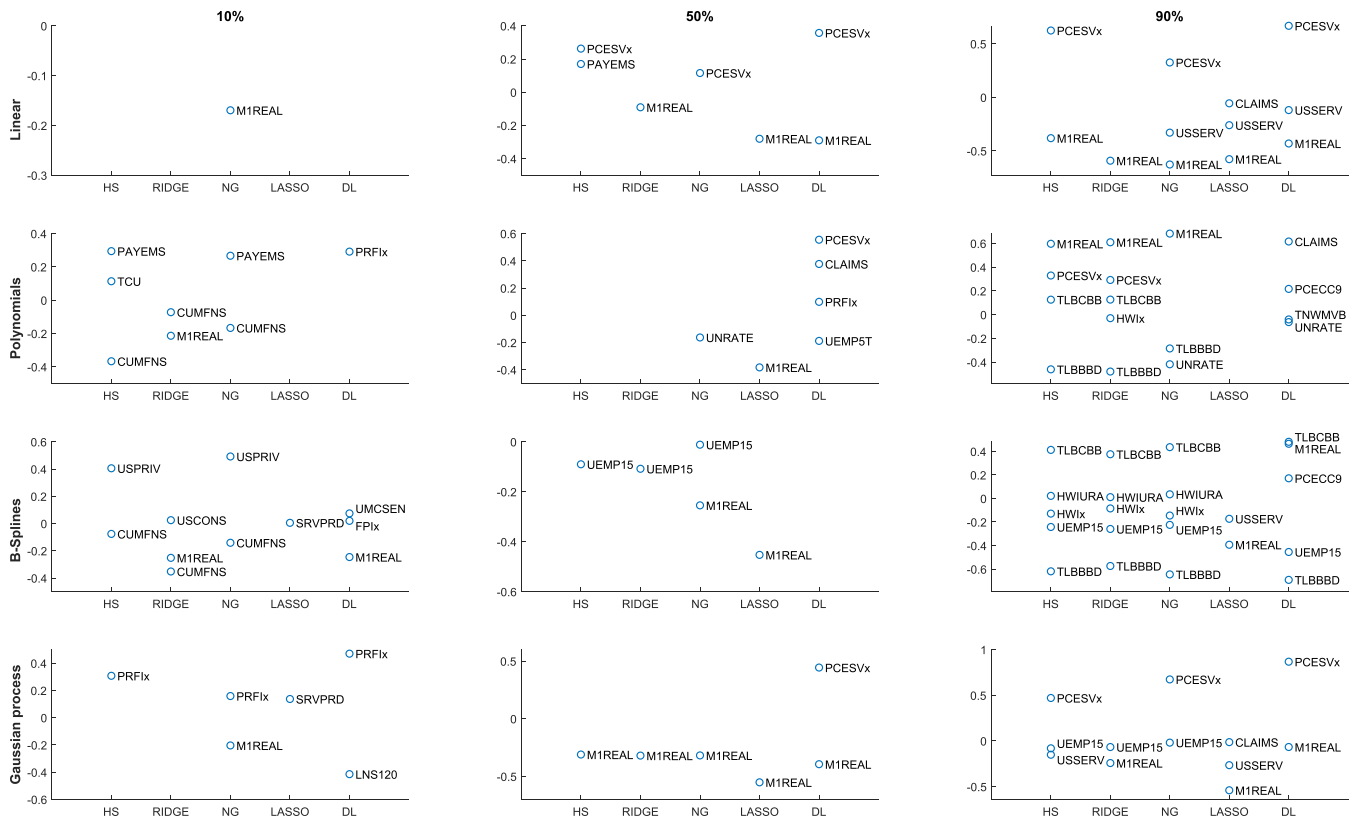


FIGURE 8 One-quarter-ahead linearized posterior summaries across quantiles.

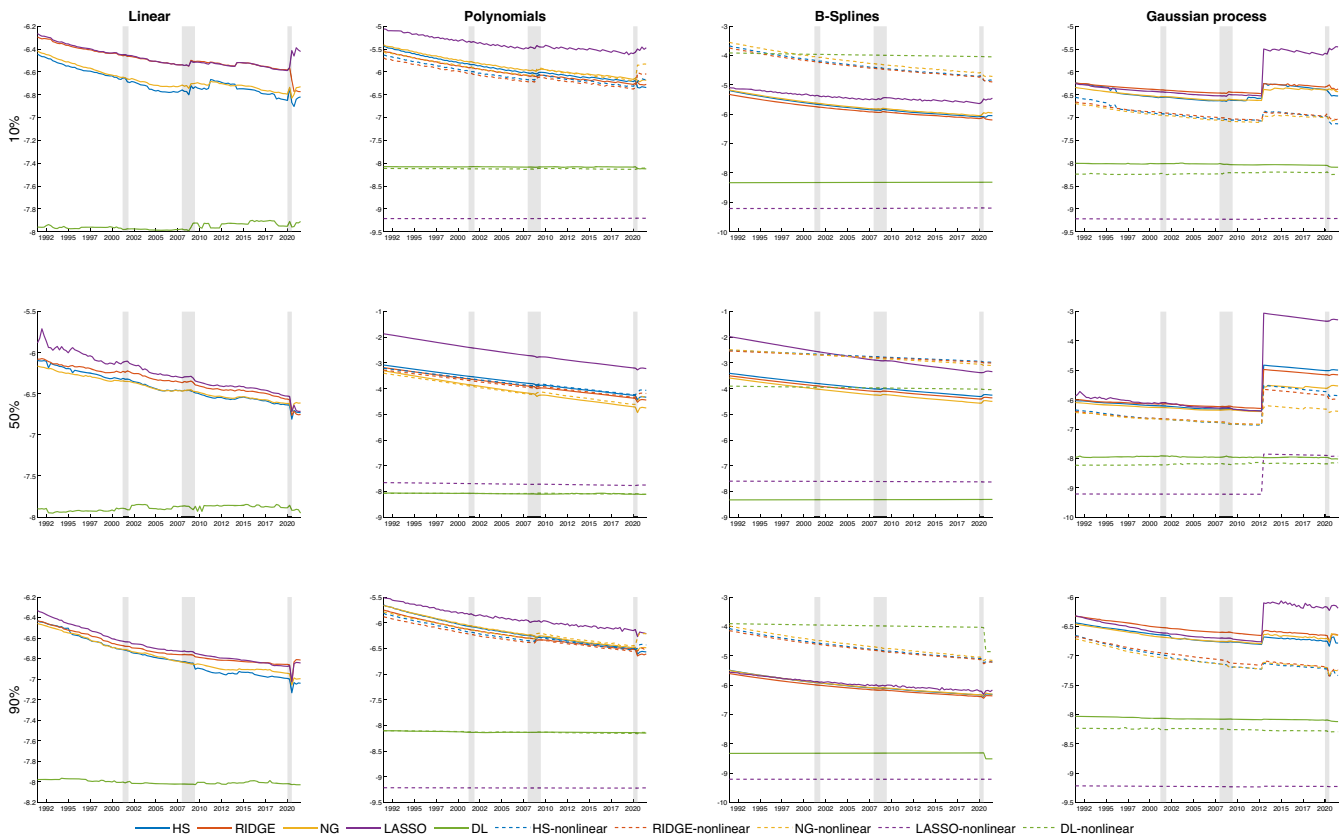


FIGURE 9 Overall shrinkage in one-step-ahead predictive quantile regressions (QRs).

4 | OTHER APPLICATIONS OF OUR VB-BASED QR ESTIMATOR

As repeatedly stated in the paper, our approach is highly scalable and the general framework can be applied to estimate specifications which are huge dimensional. In this section, we sketch how our approach can be used to estimate two popular models in the macro forecasting literature.

When predicting GDP growth, it might pay off to exploit information from higher-frequency sources. This can be achieved through mixed frequency models such as the U-MIDAS (Forni et al., 2015; Ghysels et al., 2007). In this case, $g_q(\mathbf{x}_t)$ can be replaced with a linear function $g_q(\mathbf{s}_t)$, where \mathbf{s}_t contains the quarterly observations of the higher-frequency variables where the different elements in \mathbf{s}_t relate to different periods (days, weeks, months) within a quarter. As an example, suppose that \mathbf{s}_t consists of N different weekly series and a quarter consists of 13 weeks. In this case, $M = 13N$ and if N is large, the number of parameters becomes huge. If the researcher is interested in using high-frequency information without introducing parametric restrictions, our VB-based approach could handle this situation particularly well.

Another prominent example which we mentioned in Subsection 3.3 are time-varying parameter QRs (Korobilis et al., 2021; Pfarrhofer, 2022). In this case, we set $\mathbf{g}_q = \mathbf{Z}\boldsymbol{\gamma}_q$ where \mathbf{Z} is a lower triangular matrix given by the following:

$$\mathbf{Z} = \begin{pmatrix} \mathbf{x}'_1 & 0 & \dots & 0 \\ \mathbf{x}'_2 & \mathbf{x}'_2 & \dots & 0 \\ \vdots & \vdots & \ddots & 0 \\ \mathbf{x}'_T & \mathbf{x}'_T & \dots & \mathbf{x}'_T \end{pmatrix},$$

which is a $T \times TK$ matrix and $\boldsymbol{\gamma}_q$ is a TK -dimensional vector of regression coefficients. In this case, the model is a QR with time-varying parameters that evolve according to a random walk. Since K and T can be large, the dimension of the problem quickly becomes intractable. Again, our approach could handle such a situation well and several recent papers advocate using VB to carry out estimation and inference in TVP regression models (see Koop & Korobilis, 2023a). By contrast, MCMC estimation of such models is possible but requires sophisticated tricks or other approximations to speed up computation (for a recent contribution that uses singular value decompositions, see Hauzenberger et al., 2022).

5 | CONCLUSION

In this paper, we have shown that combining QRs with nonlinear specifications and large datasets leads to precise quantile forecasts of GDP growth. Since the resulting models are high dimensional, we consider several popular shrinkage priors to regularize estimates. MCMC-based estimation of these huge dimensional models is slow. Hence, we speed up computation by using VB approximation methods that approximate the joint posterior distribution using simpler approximating densities.

The empirical results indicate that our methods work remarkably well when the CRPS is taken under consideration. When we put more weight on the tail forecasting performance, we find that most of the overall gains are driven by a strong performance in both the left and right tails, while the performance in the center of the distribution is close to the predictive accuracy of the simple quantile regression proposed in Adrian et al. (2019). These results, however, differ across priors and nonlinear specifications. In principle, it can be said that models featuring simple shrinkage priors, such as the LASSO or Ridge, in combination with GPs to capture nonlinearities of arbitrary form yield the most precise forecasts.

ACKNOWLEDGMENTS

We would like to thank the editor, Mike McCracken, three anonymous referees, and participants at the International Symposium on Forecasting 2022 at the University of Oxford and at the Statistische Woche in Münster for helpful comments. Jan Prüser gratefully acknowledges the support of the German Research Foundation (DFG, 468814087). Florian Huber acknowledges support of the Austrian Science Fund (FWF, ZK: 35).

OPEN RESEARCH BADGES



This article has been awarded Open Data Badge for making publicly available the digitally-shareable data necessary to reproduce the reported results. Data is available at [10.15456/jae.2023318.1313109186](https://doi.org/10.15456/jae.2023318.1313109186).

REFERENCES

- Adams, P. A., Adrian, T., Boyarchenko, N., & Giannone, D. (2021). Forecasting macroeconomic risks. *International Journal of Forecasting*, 37(3), 1173–1191.
- Adrian, T., Boyarchenko, N., & Giannone, D. (2019). Vulnerable growth. *American Economic Review*, 109(4), 1263–89.
- Adrian, T., Boyarchenko, N., & Giannone, D. (2021). Multimodality in macrofinancial dynamics. *International Economic Review*, 62(2), 861–886.
- Adrian, T., Grinberg, F., Liang, N., & Malik, S. (2018). The term structure of growth-at-risk. (18/180): IMF Working Paper.
- Arin, C., Kakde, D., Sadek, C., Gonzalez, L., & Kong, S. (2017). The mean and median criteria for kernel bandwidth selection for support vector data description. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, IEEE, pp. 882–849.
- Armagan, D., Bajwa, W., Lee, J., & Strawn, N. (2013). Posterior consistency in linear models under shrinkage priors. *Biometrika*, 100(4), 1011–1018.
- Bai, J., & Ng, S. (2008). Forecasting economic time series using targeted predictors. *Journal of Econometrics*, 146(2), 304–317.
- Bhattacharya, A., Pati, D., Pillai, N., & Dunson, D. (2015). Dirichlet-Laplace priors for optimal shrinkage. *Journal of the American Statistical Association*, 110, 1479–1490.
- Blei, D. M., Kucukelbir, A., & McAuliffe, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518), 859–877.
- Bufrei, G. (2019). Variational inference for quantile regression. *Arts & Sciences Electronic Theses and Dissertations*, 1743.
- Carriero, A., Clark, T. E., & Marcellino, M. (2016). Common drifting volatility in large Bayesian VARs. *Journal of Business & Economic Statistics*, 34(3), 375–390.
- Carriero, A., Clark, T. E., & Marcellino, M. G. (2022). Specification choices in quantile regression for empirical macroeconomics.
- Carvalho, C. M., Polson, N. G., & Scott, J. G. (2010). The horseshoe estimator for sparse signals. *Biometrika*, 97(2), 465–480.
- Chan, J. C. (2021). Minnesota-type adaptive hierarchical priors for large Bayesian VARs. *International Journal of Forecasting*, 37(3), 1212–1226.
- Chan, J. C. C., & Yu, X. (2022). Fast and accurate variational inference for large Bayesian VARs with stochastic volatility. *Journal of Economic Dynamics and Control*, 143, 104505.
- Clark, T. E., Huber, F., Koop, G., & Marcellino, M. (2022). Forecasting us inflation using Bayesian nonparametric models. arXiv preprint arXiv:2202.13793.
- Clark, T. E., Huber, F., Koop, G., Marcellino, M., & Pfarrhofer, M. (2023). Tail forecasting with multivariate Bayesian additive regression trees. *International Economic Review*, 64(3), 979–1022.
- Cross, J. L., Hou, C., & Poon, A. (2020). Macroeconomic forecasting with large Bayesian VARs: Global-local priors and the illusion of sparsity. *International Journal of Forecasting*, 36(3), 899–915.

- D'Agostino, A., Gambetti, L., & Giannone, D. (2013). Macroeconomic forecasting and structural change. *Journal of applied econometrics*, 28(1), 82–101.
- De Boor, C. (2001). *A practical guide to splines*: Springer.
- Delle Monache, D., De Polis, A., & Petrella, I. (2020). Modeling and forecasting macroeconomic downside risk. (15109): CEPR Discussion Paper Series.
- Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13(3), 253–265.
- Ferrara, L., Mogliani, M., & Sahuc, J. G. (2019). Real-time high frequency monitoring of growth-at-risk.
- Figueres, J. M., & Jarociński, M. (2020). Vulnerable growth in the euro area: Measuring the financial conditions. *Economics Letters*, 191, 109126.
- Foroni, C., Marcellino, M., & Schumacher, C. (2015). Unrestricted mixed data sampling (MIDAS): MIDAS regressions with unrestricted lag polynomials. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 178(1), 57–82.
- Frazier, D. T., Loaiza-Maya, R., & Martin, G. M. (2022). Variational Bayes in state space models: Inferential and predictive accuracy. (01/22): Department of Econometrics and Business Statistics, Monash University.
- Gefang, D., Koop, G., & Poon, A. (2023). Forecasting using variational Bayesian inference in large vector autoregressions with hierarchical shrinkage. *International Journal of Forecasting*, 39(1), 346–363.
- Ghosh, P., Tang, X., Ghosh, M., & Chakrabarti, A. (2016). Asymptotic properties of Bayes risk of a general class of shrinkage priors in multiple hypothesis testing under sparsity. *Bayesian Analysis*, 11(3), 753–796.
- Ghysels, E., Sinko, A., & Valkanov, R. (2007). MIDAS regressions: Further results and new directions. *Econometric reviews*, 26(1), 53–90.
- Giannone, D., Lenza, M., & Primiceri, G. E. (2021). Economic predictions with big data: The illusion of sparsity. *Econometrica*, 89(5), 2409–2437.
- Gneiting, T., & Ranjan, R. (2011). Comparing density forecasts using threshold-and quantile-weighted scoring rules. *Journal of Business & Economic Statistics*, 29(3), 411–422.
- González-Rivera, G., Maldonado, J., & Ruiz, E. (2019). Growth in stress. *International Journal of Forecasting*, 35(3), 948–966.
- Griffin, J. E., & Brown, P. J. (2010). Inference with normal-gamma prior distributions in regression problems. *Bayesian Analysis*, 5(1), 171–188.
- Hahn, P. R., & Carvalho, C. M. (2015). Decoupling shrinkage and selection in Bayesian linear models: A posterior summary perspective. *Journal of the American Statistical Association*, 110(509), 435–448.
- Harvey, D., Leybourne, S., & Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13(2), 281–291.
- Hastie, T., & Tibshirani, R. (1987). Generalized additive models: Some applications. *Journal of the American Statistical Association*, 82(398), 371–386.
- Hauzenberger, N., Huber, F., Koop, G., & Onorante, L. (2022). Fast and flexible Bayesian inference in time-varying parameter regression models. *Journal of Business & Economic Statistics*, 40(4), 1904–1918.
- Hauzenberger, N., Huber, F., & Onorante, L. (2021). Combining shrinkage and sparsity in conjugate vector autoregressive models. *Journal of Applied Econometrics*, 36(3), 304–327.
- Huber, F., & Feldkircher, M. (2019). Adaptive shrinkage in Bayesian vector autoregressive models. *Journal of Business & Economic Statistics*, 37(1), 27–39.
- Huber, F., Koop, G., & Onorante, L. (2021). Inducing sparsity and shrinkage in time-varying parameter models. *Journal of Business & Economic Statistics*, 39(3), 669–683.
- Huber, F., Koop, G., Onorante, L., Pfarrhofer, M., & Schreiner, J. (2023). Nowcasting in a pandemic using non-parametric mixed frequency VARs. *Journal of Econometrics*, 232, 52–69.
- Koenker, R., & Bassett, G. (1978). Regression quantiles. *Econometrica: journal of the Econometric Society*, 33–50.
- Kohns, D., & Szendrei, T. (2021). Decoupling shrinkage and selection for the Bayesian quantile regression. arXiv preprint arXiv:2107.08498.
- Koop, G., & Korobilis, D. (2023a). Bayesian dynamic variable selection in high dimensions. *International Economic Review*, 64(3), 1047–1074.
- Koop, G., & Korobilis, D. (2023b). Variational Bayes inference in high-dimensional time-varying parameter models.
- Korobilis, D., Landau, B., Musso, A., & Phella, A. (2021). The time-varying evolution of inflation risks. manuscript.
- Korobilis, D., & Schröder, M. (2022). Probabilistic quantile factor analysis. arXiv preprint arXiv:2212.10301.
- Kozumi, H., & Kobayashi, G. (2011). Gibbs sampling methods for Bayesian quantile regression. *Journal of Statistical Computation and Simulation*, 81(11), 1565–1578.
- McCracken, M. W., & Ng, S. (2016). FRED-MD: A monthly database for macroeconomic research. *Journal of Business & Economic Statistics*, 34(4), 574–589.
- Mitchell, J., Poon, A., & Mazzi, G. L. (2022). Nowcasting Euro area GDP growth using Bayesian quantile regression, *Essays in honor of M. Hashem Pesaran: Prediction and macro modeling*: Emerald Publishing Limited, pp. 51–72.
- Pfarrhofer, M. (2022). Modeling tail risks of inflation using unobserved component quantile regressions. *Journal of Economic Dynamics and Control*, 143, 104493.
- Plagborg-Møller, M., Reichlin, L., Ricco, G., & Hasenzagl, T. (2020). When is growth at risk? *Brookings Papers on Economic Activity*, 167–229.
- Polson, N. G., & Scott, J. G. (2010). Shrink globally, act locally: Sparse Bayesian regularization and prediction. *Bayesian Statistics*, 9(501-538), 105.
- Polson, N. G., & Sokolov, V. (2017). Deep learning: A Bayesian perspective. *Bayesian Analysis*, 12(4), 1275–1304.
- Prüser, J. (2022). Data-based priors for vector error correction models. *International Journal of Forecasting*, 39(1), 209–227.
- Ray, P., & Bhattacharya, A. (2018). Signal adaptive variable selector for the horseshoe prior. arXiv preprint arXiv:1810.09004.
- Reichlin, L., Ricco, G., & Hasenzagl, T. (2020). Financial variables as predictors of real growth vulnerability. (05/2020): Deutsche Bundesbank Discussion Paper <http://hdl.handle.net/10419/214829>

- Shin, M., Bhattacharya, A., & Johnson, V. E. (2020). Functional horseshoe priors for subspace shrinkage. *Journal of the American Statistical Association*, 115(532), 1784–1797.
- van der Pas, S., Kleijn, B., & van der Vaart, A. (2014). The horseshoe estimator: Posterior concentration around nearly black vectors. *Electronic Journal of Statistics*, 8(2), 2585–2618.
- Wand, M. P., Ormerod, J. T., Padoan, S. A., & Frühwirth, R. (2011). Mean field variational Bayes for elaborate distributions. *Bayesian Analysis*, 6(4), 847–900.
- West, M. (1987). On scale mixtures of normal distributions. *Biometrika*, 74(3), 646–648.
- Williams, CKI, & Rasmussen, C. E. (2006). *Gaussian processes for machine learning*, Vol. 2: MIT Press.
- Woody, S., Carvalho, C. M., & Murray, J. S. (2021). Model interpretation through lower-dimensional posterior summarization. *Journal of Computational and Graphical Statistics*, 30(1), 144–161.
- Yu, K., & Moyeed, R. A. (2001). Bayesian quantile regression. *Statistics & Probability Letters*, 54(4), 437–447.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of the article.

How to cite this article: Prüser, J., & Huber, F. (2024). Nonlinearities in macroeconomic tail risk through the lens of big data quantile regressions. *Journal of Applied Econometrics*, 39(2), 269–291. <https://doi.org/10.1002/jae.3018>