The time-varying evolution of inflation risks

Dimitris Korobilis
University of Glasgow

Bettina Landau
European Central Bank

Alberto Musso
European Central Bank

Anthoulla Phella
European Central Bank

June 12, 2021

Abstract

This paper develops a Bayesian quantile regression model with time-varying parameters (TVPs) for forecasting inflation risks. The proposed parametric methodology bridges the empirically established benefits of TVP regressions for forecasting inflation, with the ability of quantile regression to model flexibly the whole distribution of inflation. In order to make our approach accessible and empirically relevant for forecasting, we derive an efficient Gibbs sampler by transforming the state-space form of the TVP quantile regression into an equivalent high-dimensional regression form. An application of this methodology points to good forecasting performance of quantile regressions with time-varying parameters augmented with specific credit and money-based indicators for the prediction of the conditional distribution of inflation in the euro area, both in the short and longer run, and specifically for tail risks.

Keywords: Quantile regression; MCMC; time-varying parameters; Bayesian shrinkage; Horseshoe; euro area; inflation tail risks.

JEL Classification: C11, C22, C52, C53, C55, E31, E37, E51

The opinions in this paper are those of the authors and do not necessarily reflect the views of the European Central Bank or the Eurosystem.

We would like to thank Roberto Casarin, Frank Diebold, Kostas Kalogeropoulos, Luca Rossini, Frank Schorfheide, participants at the 11^{th} ECB Workshop on Forecasting Techniques, and participants in seminars at University of Pennsylvania, LSE joint Statistics/Economics seminars, and Ca'Foscari University Venice, for helpful discussions and comments.

Corresponding author: Dimitris Korobilis, Adam Smith Business School, University of Glasgow, G12 8QQ, Glasgow, United Kingdom. email: dikorobilis@googlemail.com

1 Introduction

The objective of this paper is to uncover which models allow for an enhanced prediction of inflation risks in the euro area, that is, the risk of extreme realisations of inflation that correspond to the tails of its distribution. For this purpose, we explore jointly two main modelling directions. First, we assess the role of financial indicators in forecasting the distribution of core inflation, concentrating on its left (negative or low inflation) and right (high inflation) tails. To address this question, we adopt a time-series quantile regression setting that has proven to be particularly useful in macroeconomic forecasting. Using such setting, Adrian et al. (2019) show that the risks of low real economic activity growth are particularly sensitive to deteriorating financial conditions. Similarly, Korobilis (2017) and Lopez-Salido and Loria (2020) use quantile regressions to explore the factors that affect different quantiles of inflation distribution. Within this econometric framework, our second research question is methodological and pertains to understanding the role of time-varying parameters in a quantile regression. Time-varying parameters (TVPs) have a long tradition in macroeconomics (see for example Cooley and Prescott, 1976) and there is a large econometric literature that also attempts to use TVP regressions to identify good predictors of the mean of inflation at different points in time (Koop and Korobilis, 2012).[1] Additionally, it has been argued recently that regressions that feature time-varying variances can forecast output risks as well as constant parameter quantile regression models (Brownlees and Souza, 2021; Carriero et al., 2020). However, little is known about whether one can further improve forecasts by combining the benefits of time-varying parameters with the flexibility of a quantile regression setting.

Taking all the considerations above into account, our proposal is to use a time-varying parameter quantile regression (TVP-QR) model for forecasting the full distribution of inflation. At the conceptual level, specification of a TVP-QR model is not novel. However, serious inference challenges are in order with the implementation of this model in a time-series forecasting context. Kim (2007), Cai and Xu (2009), and Wu and Zhou (2017) use nonparametric methods, such as splines and local polynomials, to estimate TVP-QR models. However, nonparametric estimators are not straightforward to interpret and they are hard to apply to models with more predictors/indicators than time-series observations (as is often the case with euro area macroeconomic data) or if the interpretation of coefficients is key for policy purposes. In contrast to previous contributions, our proposed framework is Bayesian, meaning that error and parameter distributions are all flexible parametric rather than nonparametric. We approximate the quantile regression (QR) problem with an asymmetric Laplace error distribution (Kozumi and Kobayashi, 2011). The evolution of time-varying

[1] Summarising the results of a comprehensive comparison of different models, data and transformations, Faust and Wright (2013) argue that a basic principle in forecasting inflation is to allow for its local mean to be smoothly varying over time, and an obvious way of doing this is via time-varying parameters (Stock and Watson, 2007a).

parameters follows a random walk specification which is traditionally tackled with standard Markov chain Monte Carlo (MCMC) algorithms for state-space models; see for example the Bayesian quantile state-space model of Gonçalves et al. (2020). However, in a QR setting we need to estimate separate regressions for each quantile level, making MCMC estimation cumbersome and costly for the purpose of recursive forecasting.[2] Our main methodological contribution is to propose an MCMC algorithm that makes estimation and forecasting with the TVP-QR model feasible. We borrow ideas from Korobilis (2021) and write the TVP-QR model as an equivalent high-dimensional (quantile) regression. The resulting approximation-free algorithm ends up being a minor reparameterisation of the efficient algorithm of Chan and Jeliazkov (2009), but it is computationally much faster, thereby allowing estimation of the TVP-QR model over a fine grid of quantiles.

We establish the benefits of our approach using both synthetic and real data. When generating synthetic data from regressions with time-varying parameters and flexible error distributions we find that our framework recovers the true parameters with higher accuracy compared to non-quantile TVP regressions that rely on a Gaussian disturbance term. Additionally, we show using these simulated data examples that our algorithm is able to track complex patterns of time-variation in parameters. We achieve this by adopting the "horseshoe prior for sparse signals" of Carvalho et al. (2010), which in our setting allows for shrinkage of the time-varying parameters towards few structural breaks or time-invariance, without any dependence on prior tuning hyperparameters. That way, we fully address concerns in Amir-Ahmadi et al. (2020) about the impact that prior hyperparameter choice has on estimation of TVPs, making our new estimation algorithm both fast and easy to use by less experienced users; forecasts produced by the TVP-QR methodology are fully replicable (up to a typically small asymptotic bias resulting from Monte Carlo estimates of expectations).

We apply this flexible framework to the problem of forecasting core consumer price inflation in the euro area using various financial indicators, both in the short run (four-quarters ahead) and the medium run (twelve-quarters ahead). Using quarterly observations from 1990 to 2019, we find that a number of specific financial volume indicators such as loans to the private sector, loans to households and narrow money (M1) often provide the largest inflation tail risk forecasting gains, especially in the context of quantile regressions with time-varying parameters or quantile regression-based augmented Phillips curve models with time-varying parameters. A comparison of several different models and financial indicators, including both financial prices and financial volumes, allows to conclude that such forecast gains derive both from the predictive informational content of such specific financial volume indicators and from the benefits of modelling time-varying parameters in the context of quantile regressions.

[2] Gonçalves et al. (2020) propose a straightforward MCMC algorithm for general quantile state-space models, but they acknowledge that this algorithm is very slow when iterating over quantiles, so they end up proposing a faster, approximate algorithm. Similarly, Lim et al. (2020) estimate a TVP-QR model but estimation is based on approximate variational Bayes methods.

The article is organised as follows. The next section describes the Bayesian TVP-QR model, its prior distributions, and our efficient MCMC estimation approach. In Section 3 we conduct a small Monte Carlo experiment to establish that the TVP-QR model can recover estimates of the true TVPs much more accurately than the default TVP regression with Gaussian errors and stochastic volatility. In Section 4 we show the quantitative results from a large-scale forecasting exercise involving versions of our proposed model using different indicators, as well as various competing methodologies. Section 5 concludes the paper.

2 Bayesian time-varying parameter quantile regression

Let π_t be the scalar observation of inflation in time periods $t = 1, ..., T$, and $\boldsymbol{x}_t$ a p-dimensional vector of predetermined variables that includes intercept, lags of inflation and exogenous predictors. We want to model the full distribution of π_t by specifying the following model for each of its quantiles $\tau = \{0.05, 0.10, ..., 0.90, 0.95\}$

$$\pi_t = \mathcal{Q}_\tau \left(\pi_t | \boldsymbol{x}_t \right) + \varepsilon_t, \tag{1}$$

where $\mathcal{Q}_\tau$ denotes the conditional quantile function of the τ-th quantile of π_t. Several linear and nonlinear quantile functions have been proposed, especially in microeconometric applications. In a time-series context, we are interested in the following function

$$\begin{aligned} \mathcal{Q}_\tau \left(\pi_t | \boldsymbol{x}_t \right) &= \boldsymbol{x}_t \boldsymbol{\beta}_t(\tau), && (2) \\ \boldsymbol{\beta}_t(\tau) &= \boldsymbol{\beta}_{t-1}(\tau) + \boldsymbol{v}_t, && (3) \end{aligned}$$

where $\boldsymbol{v}_t \sim N_p \left(\boldsymbol{0}, \boldsymbol{V}(\tau) \right)$ is a state error with covariance matrix $\boldsymbol{V}(\tau)$. Under this specification parameters evolve as random walks. When $\boldsymbol{V}(\tau)$ is small the evolution is smooth[3], while for larger values of $V(\tau)$ this specification can capture abrupt jumps. Therefore, the full time-varying parameter quantile regression (TVP-QR) specification in its most general form comprises Equations (1) to (3).

2.1 A reparameterised TVP-QR model

Our first building block for estimating this model is the treatment of the error. In the constant parameter case, $\beta_t(\tau) = \beta(\tau)$, Koenker and Bassett (1978) show that univariate conditional quantiles can be obtained as the solution to the following optimisation problem

$$\widehat{\beta}(\tau) = \min_{\beta(\tau)} \mathbb{E} \sum_{t=1}^{T} \rho_\tau(\varepsilon_t), \tag{4}$$

[3]Notice that the solution $\boldsymbol{V}(\tau) = \boldsymbol{0}$ gives the constant parameter model as a special case of the TVP model.

where $\rho_\tau(u) = (\tau - \mathbb{I}(u < 0))u$ is a loss function. The minimiser of Equation 4 is equivalent to maximising an asymmetric Laplace likelihood (Yu and Moyeed, 2001), that is, the case where ε_t has density given by

$$p(\varepsilon_t; \tau, \sigma) \sim \frac{\tau(1-\tau)}{\sigma(\tau)^2} \left[e^{(1-\tau)\frac{\varepsilon_t}{\sigma(\tau)^2}} \mathbb{I}(\varepsilon_t \leq 0) + e^{-\tau\frac{\varepsilon_t}{\sigma(\tau)^2}} \mathbb{I}(\varepsilon_t > 0) \right], \tag{5}$$

where $\sigma(\tau)$ is a scale parameter.[4] Following Kozumi and Kobayashi (2011), the asymmetric Laplace distribution can be written as a Gaussian-Exponential scale mixture of the form

$$(\varepsilon_t | u_t, z_t) \sim \theta(\tau) z_t + \sqrt{\sigma(\tau)^2 \kappa(\tau)^2 z_t(\tau)} u_t, \tag{6}$$

where $z_t(\tau) \sim Exp(\sigma^2(\tau))$ and $u_t \sim N(0, 1)$, while $\theta(\tau), \kappa(\tau)^2$ are parameters defined as $\theta(\tau) = \frac{1-2\tau}{\tau(1-\tau)}$, $\kappa(\tau)^2 = \frac{2}{\tau(1-\tau)}$. If we marginalise Equation 6 over z_t we obtain Equation 5; see more details in our online Appendix and Khare and Hobert (2012). In a parametric setting, it is trivial to adopt the parametric distribution in Equation 6 in the context of the quantile regression in Equation 1, and this is what we do in this paper. The benefits of this approach are immediately visible: since the likelihood is conditionally (on z_t) Gaussian, the conditional posteriors will be identical to the ones in the simple regression model.[5]

The second building block is the way we treat time variation. We extend ideas in Korobilis (2021) and we rewrite the model in Equations (2) - (3) as a high-dimensional regression with more covariates than observations. In particular, it is easy to show that if we stack all T observations, these equations can be rewritten as

$$\mathcal{Q}_\tau(\boldsymbol{\pi} | \mathcal{X}) = \mathcal{X} \boldsymbol{\beta}^\delta(\tau), \tag{11}$$

$$\boldsymbol{\beta}^\Delta(\tau) = \boldsymbol{v}, \tag{12}$$

[4] In the time-varying parameter mean regression framework it is natural to assume the presence of stochastic volatility, σ_t^2. Combined with the typical assumption of Normality in the errors, time-varying variances are able to produce more flexible, heavy-tailed unconditional distributions of inflation. The assumption of stochastic volatility in the quantile regression model, $\sigma_t(\tau)$, is computationally trivial to incorporate. However, we have found that, unlike longer financial data (Gerlach et al., 2011), it consistently produces inferior fit and out-of-sample forecasts for euro area inflation data, given it relatively short sample size. Even when not allowing the variance parameter to fluctuate over time, in the context of quantile regression it takes a different value in different quantiles, meaning that we are able to capture very complex shapes of distributions.

[5] It is easy to visualize this; for example in the linear QR case

$$\pi_t = \boldsymbol{x}_t \boldsymbol{\beta}(\tau) + \theta(\tau) z_t + \sqrt{\sigma(\tau)^2 \kappa(\tau)^2 z_t(\tau)} u_t, \tag{7}$$

if we condition on z_t (i.e. we treat it as a known parameter) we can write

$$\pi_t - \theta(\tau) z_t = \boldsymbol{x}_t \boldsymbol{\beta}(\tau) + \sigma(\tau)\kappa(\tau)\sqrt{z_t(\tau)} u_t, \Rightarrow \tag{8}$$

$$\frac{\pi_t - \theta(\tau) z_t}{\kappa(\tau)\sqrt{z_t}} = \left(\frac{\boldsymbol{x}_t}{\kappa(\tau)\sqrt{z_t}} \right) \boldsymbol{\beta}(\tau) + \sigma(\tau) u_t, \tag{9}$$

$$\widetilde{\pi}_t = \widetilde{\boldsymbol{x}}_t \boldsymbol{\beta}(\tau) + \widetilde{\varepsilon}_t, \tag{10}$$

which is a linear, Gaussian regression on the data $\widetilde{\pi}_t$ and $\widetilde{\boldsymbol{x}}_t$, and the error $\widetilde{\varepsilon}_t \sim N(0, \sigma(\tau))$. Therefore, it is fairly trivial to derive conditional posteriors for $\boldsymbol{\beta}(\tau)$ and $\sigma(\tau)$ using this form.

where $\boldsymbol{\pi} = [\pi_1, ..., \pi_T]'$, $\boldsymbol{v} = [\boldsymbol{v}_1', ..., \boldsymbol{v}_T']'$ and

$$\mathcal{X} = \begin{bmatrix} x_1 & 0 & ... & 0 & 0 \\ x_2 & x_2 & ... & 0 & 0 \\ ... & ... & ... & ... & ... \\ x_{T-1} & x_{T-1} & ... & x_{T-1} & 0 \\ x_T & x_T & ... & x_{T-1} & x_T \end{bmatrix}, \quad \text{and} \quad \boldsymbol{\beta}^{\Delta}(\tau) = \begin{bmatrix} \boldsymbol{\beta}_1(\tau) \\ \Delta\boldsymbol{\beta}_2(\tau) \\ ... \\ \Delta\boldsymbol{\beta}_{T-1}(\tau) \\ \Delta\boldsymbol{\beta}_T(\tau) \end{bmatrix}. \tag{13}$$

$$T \times Tp \qquad\qquad\qquad Tp \times 1$$

We provide detailed derivations and discussion of this form in the Online Appendix. In this new formulation, all Tp coefficients of the TVP regression are stacked in a single vector, while at the same time they appear in first differences form. Specifically, in this hierarchical (multilevel) regression specification, Equation 12 can be interpreted as a prior for $\boldsymbol{\beta}_{\Delta}(\tau)$ which means that equation Equation 11 can be treated as a linear regression model and estimated using algorithms for constant parameter models. Interpretation of this formulation is straightforward as we can recover the original vector of TVPs, $\boldsymbol{\beta} = [\boldsymbol{\beta}_1(\tau)', ..., \boldsymbol{\beta}_T(\tau)']'$ as the cumulative sum of the vector of first differences, $\boldsymbol{\beta}_{\Delta}(\tau)$.

2.2 Likelihood, priors and a new Gibbs sampler

Putting together these pieces, the new parameterisation now combines Equation 1 with the distributional assumptions on the error term in Equation 6 with the reparameterised TVP function in Equations (11) and (12). The core of the TVP-QR model now has the following form

$$\boldsymbol{\pi} \;=\; \mathcal{X}\boldsymbol{\beta}^{\Delta}(\tau) + \theta(\tau)\boldsymbol{z}(\tau) + \widetilde{\boldsymbol{S}}\boldsymbol{u}, \tag{14}$$

where $\widetilde{\boldsymbol{S}}$ is a $T \times T$ diagonal matrix with t-th diagonal element $\sqrt{\sigma(\tau)^2\kappa(\tau)^2 z_t(\tau)}$. This model is completed by considering priors on various parameters. By definition, $z(\tau) \sim Exponential(\sigma(\tau)^2)$. The scale parameter can take the standard inverse Gamma prior, that is, $\sigma(\tau) \sim inv - Gamma(\rho_1, \rho_2)$. Finally, from Equation 12 we already discussed that $\boldsymbol{\beta}^{\Delta}(\tau)$ has a standard Normal prior of the form $\boldsymbol{\beta}^{\Delta}(\tau) \sim N(0, V(\tau))$. Multiplication of these priors with the reparameterised likelihood implied by the model in Equation (14), gives the following conditional posteriors

$$\begin{aligned}
\beta^{\Delta}(\tau)|\bullet &\sim N\left(\boldsymbol{Q} \times \left(\mathcal{X}'\boldsymbol{U}^{-1}\widetilde{\boldsymbol{y}}\right), \boldsymbol{Q}\right), & (15)\\
\sigma(\tau)^2|\bullet &\sim inv-Gamma\left(\rho_1 + \frac{3T}{2}, \rho_2 + \sum_{t=1}^{T}\frac{(y_t^{\star})^2}{2z_t(\tau)\kappa(\tau)^2} + \sum_{t=1}^{T} z_t(\tau)\right), & (16)\\
z_t(\tau)|\bullet &\sim IG\left(\frac{\sqrt{\theta(\tau)^2 + 2\kappa(\tau)^2}}{|y_t - \mathcal{X}_t\boldsymbol{\beta}^{\Delta}(\tau)|}, \frac{\theta(\tau)^2 + 2\kappa(\tau)^2}{\sigma(\tau)^2\kappa(\tau)^2}\right), & (17)
\end{aligned}$$

where the notation $|\bullet$ means "conditioning on other parameters and data", $\boldsymbol{Q} = \left(\mathcal{X}'\boldsymbol{U}^{-1}\mathcal{X} + V(\tau)^{-1}\right)^{-1}$, $\boldsymbol{U} = \left(\sigma(\tau)^2\kappa(\tau)^2\right) \times diag\left(z_1(\tau), ..., z_T(\tau)\right)$, $\widetilde{\boldsymbol{y}} = (\boldsymbol{y} - \theta(\tau)\boldsymbol{z}(\tau))$, $y_t^{\star} = \left(y_t - \mathcal{X}_t\boldsymbol{\beta}^{\Delta}(\tau) - \theta(\tau)z_t(\tau)\right)$. This Gibbs sampler is a reparameterised version of the ergodic Gibbs sampler for constant parameter quantile regressions developed by Khare and Hobert (2012). Sampling from the conditional distributions is straightforward and computation can be sped up by sampling for all values of τ simultaneously, instead of sampling iteratively for each τ. The only computational challenge is sampling of the Tp elements in $\boldsymbol{\beta}^{\Delta}(\tau)$, since Tp can be very large. Bhattacharya et al. (2016) provide a very efficient way of sampling from such high-dimensional Normal posteriors, and we refer the reader to this paper and our Online Appendix for more information about implementation.

The reparameterised form of the TVP regression shows that this is a model with more predictors than observations (measurement matrix $\mathcal{X}$ has Tp covariates but only T observations). Therefore, it is evident that prior selection for the high-dimensional vector $\boldsymbol{\beta}^{\Delta}(\tau)$ must be done carefully; see Amir-Ahmadi et al. (2020) for a discussion of these issues in traditional TVP models. Consequently, we adopt the horseshoe prior of Carvalho et al. (2010) that is shown in several instances to have excellent theoretical guarantees, thus making it an established estimator in statistics.[6] The horseshoe prior for $\boldsymbol{\beta}^{\Delta}(\tau)$

$$\begin{aligned}
\boldsymbol{\beta}^{\Delta}(\tau)|\lambda(\tau)^2, \{\psi_i(\tau)^2\}_{i=1}^{Tp} &\sim N(0, V(\tau)), & (18)\\
V_{i,i}(\tau) &= \lambda(\tau)^2\psi_i(\tau)^2, \quad i = 1, ..., Tp, &\\
\lambda(\tau) &\sim Cauchy^{+}(0,1), & (19)\\
\psi_i(\tau) &\sim Cauchy^{+}(0,1), \quad i = 1, ..., Tp, & (20)
\end{aligned}$$

where $Cauchy^{+}$ denotes the half-Cauchy distribution on the positive reals. [7] In line with Amir-Ahmadi et al. (2020), $V(\tau)$ comprises hyperparameters that have their own prior distributions

[6] In linear Gaussian regression settings, the horseshoe prior is minimax in l_2 norm (van der Pas et al., 2014), attains risk equal to the Bayes oracle (Ghosh et al., 2016) and posterior credible intervals under the horseshoe prior also have good frequentist coverage properties in an asymptotic sense. More recently, Bhadra et al. (2020) show that horseshoe regularisation retains its excellent properties in several classes of complex models, including non-linear, non-Gaussian regression, and deep neural networks.

[7] The formulation using the half-Cauchy priors is not ideal as it does not allow straightforward derivation of conditional posteriors. Makalic and Schmidt (2016) note that the half-Cauchy distribution can be written as a

and are, thus, updated by information in the data. The fact that the hyperpriors for $\lambda(\tau)$ and $\psi_i(\tau)$ do not depend on further parameters that require tuning/calibration, makes the horseshoe a fully automatic prior that adapts equally well to low-dimensional as well as high-dimensional problems.

3 Simulation study

In this section we use artificial data in order to examine the performance of our proposed algorithm. We generate data from the following time-varying regression model

$$y_t = x_t\beta_t + \varepsilon_t, \quad (26)$$
$$\beta_t = \mu + 0.99(\beta_{t-1} - \mu) + T^{-\frac{1}{2}}u_t, \quad (27)$$

where $x \sim N(0, I_2)$ is a vector of $p = 2$ synthetic predictors, $\mu \sim U(-2, 2)$ is the long-run mean of β_t[8], and $u_t \sim N(0, I)$. Since in the empirical section we are interested in capturing predictors that are short-lived, we artificially shrink all values of $\beta_{1,t}$ to be zero for $t > T/3$, that is, the first predictor is only relevant for y only for the first third of the sample. The second predictor in the vector x is left unrestricted (i.e. not zero) in all periods.

Regarding the distribution of ε_t, we follow the Monte Carlo design in Yu (2017), and consider eight different choices:

1. Gaussian: $N(0, 1^2)$
2. Skewed : $1/5N(-22/25, 1^2) + 1/5N(-49/125, (3/2)^2) + 3/5N(49/250, (5/9)^2)$
3. Kurtotic: $2/3N(0, 1^2) + 1/3N(0, (1/10)^2)$
4. Outlier : $1/10N(0, 1^2) + 9/10N(0, (1/10)^2)$
5. Bimodal : $1/2N(-1, (2/3)^2) + 1/2N(1, (2/3)^2)$

---

mixture of inverse Gamma distributions. Therefore, the horseshoe prior can be written equivalently as

$$\boldsymbol{\beta}^{\Delta}(\tau)|\lambda(\tau)^2, \{\psi_i(\tau)^2\}_{i=1}^{Tp} \sim N(0, V(\tau)), \quad (21)$$
$$V_{i,i}(\tau) = \lambda(\tau)^2\psi_i(\tau)^2, \quad i = 1, ..., Tp,$$
$$\lambda(\tau)^2|\xi(\tau) \sim inv-Gamma\left(1/2, 1/\xi(\tau)\right), \quad (22)$$
$$\xi(\tau) \sim inv-Gamma\left(1/2, 1\right), \quad (23)$$
$$\psi_i(\tau)^2|\zeta_i(\tau) \sim inv-Gamma\left(1/2, 1/\zeta_i(\tau)\right), \quad (24)$$
$$\zeta_i(\tau) \sim inv-Gamma\left(1/2, 1\right), \quad (25)$$

which is a formulation that allows for straightforward calculation of conditional posteriors; see Makalic and Schmidt (2016) for more details on posterior computation. It is trivial to show that one can simply add the formulas for the conditional posteriors of $\lambda(\tau)^2, \xi, \psi_i(\tau)^2, \zeta_i(\tau)$ to the Gibbs sampler in equations (15)-(17).

[8]Notice that even though β_t will be estimated using the random walk evolution outlined in the previous section, in the DGP we generate its path from a persistent yet stationary process. This choice ensures that we generate time-varying parameters that are not explosive and won't cause numerical problems during estimation.

6. Bimodal, separate modes: $1/2N(-3/2,(1/2)^2)+1/2N(3/2,(1/2)^2)$

7. Skewed bimodal: $3/4N(-43/100,1^2)+1/4N(107/100,(1/3)^2)$

8. Trimodal: $9/20N(-6/5,(3/5)^2)+9/20N(6/5,(3/5)^2)+1/10N(0,(1/4)^2)$

This list covers a wide variety of flexible distributions, even though it is far from exhaustive.

We generate $M=500$ datasets of length $T=200$ from the eight time-varying parameter regression data generating processes (DGPs). For each dataset we fit two models, a "mean" TVP regression with stochastic volatility, and our quantile TVP regression. The former model is simply a special case of the latter, where we allow the variance to be time-varying, and we convert the asymmetric Laplace distribution into a Normal distribution.[9] In both cases of these two estimated models we use the same automatically tuned horseshoe prior, such that we do not influence subjectively posterior estimates of time-varying parameters, and the same efficient sampler for TVPs outlined in the previous section. Therefore, both TVP models are almost identical with the exception of the assumption about the error distribution.

With this Monte Carlo-based exercise, we aim to find out how well the asymmetric Laplace distribution can capture all of the eight error distributions we assume that generated the data. In particular, we want to find out how large is the estimation error of the TVPs under the two estimated models in each of the eight DGPs, as accuracy of estimates of the TVPs will have an immediate impact on forecasting the synthetic outcome variable y_t (and, as a result, inflation π_t, when we input real data). On that account, we measure estimation accuracy using the following mean squared deviation (MSD) measure:

$$MSD_j = \frac{1}{M}\sum_{m=1}^{M}\left\{\frac{1}{T}\sum_{t=1}^{T}\left[\frac{1}{2}\sum_{i=1}^{2}\left(\beta_{t,i}^{(m)}-\widehat{\beta}_{t,i}^{(m),j}\right)\right]\right\}, \tag{28}$$

where $j=\{mean\ TVP\ regression,\ quantile\ TVP\ regression\}$ and $\widehat{\beta}_{t,i}^{(m),j}$ is the posterior mean of the m-th Monte Carlo iteration, of coefficient $\beta_{t,i}$, in model j. For this loss function, lower values imply lower estimation error.

Results of this exercise are presented in Table 1. Numerical entries in this table are the MSDs for the mean regression (first row) and quantile regression for quantiles $\tau=0.05,0.10,0.25,0.50,0.75,0.90,0.95$ (rows 2-8). There are eight columns in this table that are associated with each of the error distributions assumed in the DGP. We observe that when the data are generate from a TVP regression with Normal errors, then the mean TVP regression estimator is optimal as it is based on the assumption that the disturbance term is Normal. In this case the quantile TVP regression is overparameterised and the asymmetric Laplace distribution assumption does not perform well. However, once we drop the assumption

[9] Using the scale mixtures of Normals representation in (6), we can obtain the Normal distribution by fixing $z_t(\tau)=\frac{1}{\kappa(tau)^2}=\frac{\tau(1-\tau)}{2}$ for all t and setting $\tau=0.5$.

of Normality in the DGP, the mean regression model is performing worse than the quantile regression model in terms of estimation error. In the case of more complex distributions (Skewed bimodal, trimodal) the error produced by the mean regression model can be substantially larger.

Table 1: Mean squared deviations (MSDs) of estimated vs true time-varying parameters, using mean and quantile regressions

	Gaussian	Skewed	Kurtotic	Outlier	Bimodal	Bimodal sep. modes	Bimodal skewed	Trimodal
				MSD REGRESSION				
mean	0.01	0.25	0.04	0.01	0.10	0.21	0.45	1.03
				MSD QUANTILE REGRESSION				
$\tau = 0.05$	0.06	0.05	0.05	0.02	0.09	0.13	0.05	0.09
$\tau = 0.10$	0.05	0.05	0.05	0.02	0.09	0.13	0.04	0.08
$\tau = 0.25$	0.05	0.04	0.04	0.01	0.08	0.12	0.04	0.08
$\tau = 0.50$	0.05	0.04	0.04	0.01	0.06	0.11	0.03	0.06
$\tau = 0.75$	0.04	0.04	0.04	0.01	0.07	0.12	0.03	0.07
$\tau = 0.90$	0.05	0.05	0.05	0.02	0.08	0.13	0.03	0.08
$\tau = 0.95$	0.05	0.05	0.05	0.01	0.09	0.13	0.03	0.08

Notes: The mean regression model is a TVP regression with stochastic volatility assuming Normal measurement error distribution. The quantile regression model allows for time-varying coefficients of predictors and constant intercept and variance in each quantile.

In order to assist our understanding of how severe estimation error is in the TVP regression with stochastic volatility and Normality, Figure 1 plots parameter estimates from the mean and quantile regression models in the case of the true error distribution being trimodal (eighth case). The left-hand side panel shows estimates of the coefficient on the first predictor ($\beta_{t,1}$), and the right-hand side panel estimates on the coefficient of the second predictor ($\beta_{t,2}$). In all plots the black, dotted line shows the (average over 500 iterations) true generated time-varying coefficient. The green line is the average over 500 iterations of the posterior median of the TVPs and the shaded areas are the 68% probability bands. Looking at $\beta_{t,1}$ in the left-hand side panel of the figure, the mean regression model (top graph) produces some error, as the true value of the coefficient (black dotted line) is not close to the posterior median (green line), i.e. it is not always inside the grey shaded area. In a few periods the true value of this coefficient is even outside the 68% bands. This is not the case for the quantile regression estimates for $\tau = 0.05, 0.10, 0.90, 0.95$ (bottom four graphs), where the posterior median is much closer to the true value of this coefficient. The picture of large estimation errors in the TVPs are more pronounced when we look at $\beta_{t,2}$. The "mean" regression estimates completely miss the true path of this coefficient that is used to generate the data. By contrast, the estimates of this time-varying parameter from the quantile regression are much more accurate. This graphical

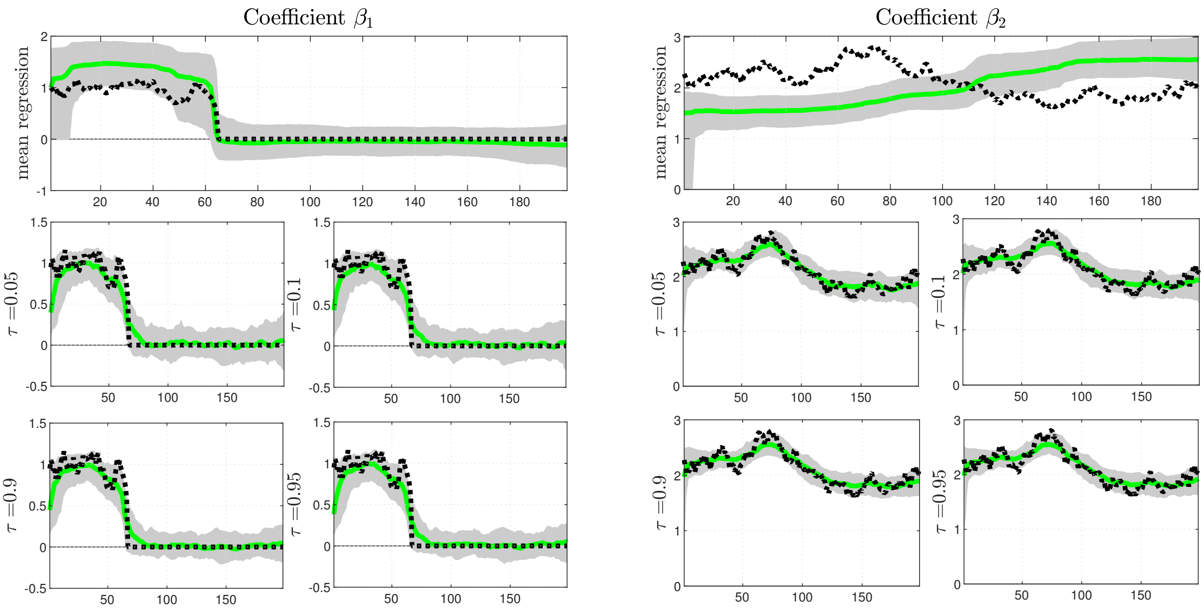

Figure 1: *Posterior estimates of time-varying parameters (TVPs) estimated using mean (upper panels) and quantile (middle and bottom panels) regressions. The five panels on the left pertain to coefficient β_{1t} in our DGP, and the five panels on the right to coefficient β_{2t}. The DGP used to produce this figure is that of a time-varying regression model with a trimodal error distribution. Mean regression is done under the standard assumption of a Normal error distribution, while the quantile regression is estimated using the flexible asymmetric Laplace distribution. Black lines are the true TVPs, which are the same for both the mean and quantile regressions. The green lines are the averages (over 100 Monte Carlo iterations) of the estimated posterior means, and the shaded areas and the 68 percent probability bands.*

illustration gives an example of how the MSDs of the previous table translate into significant estimation errors. Consequently, when our data distribution is non-Gaussian (which is the case with inflation and several other macroeconomic and financial variables), then our proposed TVP-QR methodology will dominate traditional TVP regressions with stochastic volatility.

4 Forecasting inflation risks in the euro area

4.1 Data

In order to assess the practical usefulness of the proposed approach for the projection of the conditional distribution of inflation, we concentrate on euro area HICP excluding energy and food inflation (henceforth referred to as core HICP inflation) developments from the first quarter of 1990 to the fourth quarter of 2019 (see Figures B1 and B2 in Appendix B). Taking core inflation instead of headline (total) HICP inflation as reference allows abstracting from the influence of temporary factors such as oil price and exchange rate swings and focusing on more fundamental forces driving inflation tail risks. For core HICP quarterly inflation (measured by the annualised quarter-on-quarter growth rate of HICP excluding energy and food) over the whole sample size (1990Q1-2019Q4) the 5th and 95th percentiles correspond to quarterly inflation rates at

0.7 percent and 4.2 percent, respectively (see Table B2 in Appendix B). The unconditional distribution of core HICP quarterly inflation is slightly skewed to the left and exhibits a heavy right-hand side tail.

Nineteen financial indicators are considered in the analysis, including both financial volumes and financial prices. Specifically, these include four money volume indicators (M1 and M3, each expressed in annualised quarterly growth rates and as a ratio to GDP), eight credit volume indicators (total credit to the non-financial private sector, bank lending to the non-financial private sector, bank lending to non-financial corporations and bank lending to households, each expressed in annualised quarterly growth rates and as ratio to GDP), four credit spreads (for the 10-year government bond yields, lending rates to non-financial corporations, lending rates to households and investment grade corporate bond yields, all as deviations from the 3-month Euribor rate) and three additional financial indicators (stock prices, house prices and the composite indicator of systemic stress). Table B1 in Appendix B reports the details of these data.

4.2 Models and forecast setup

The forecast performance of several models for the prediction of the conditional distribution of core HICP inflation in the short and medium term (i.e. four-quarters and twelve-quarters ahead, respectively) is assessed. The performance of various categories of models is analysed, including bivariate models conditional on financial indicators and Phillips-curve based models augmented with financial indicators, which are compared to the forecasting ability of various univariate models, as well as standard Phillips-curve based models (i.e., not augmented with any financial indicator). For each model we entertain various specifications, ranging from quantile regressions to mean regressions, with or without time-varying parameters and with or without stochastic volatility.

We consider forecasts from the TVP-QR and numerous competitors, all of which can be written as special cases of the general formulation

$$\pi_{t+h} = c_t(\tau) + \phi_{1t}(\tau)\pi_t + \phi_{2t}(\tau)\pi_{t-1} + \beta_t(\tau)\boldsymbol{x}_t + \varepsilon_{t+h}, \quad \varepsilon_{t+h} \sim ALD(\sigma_t(\tau)), \tag{29}$$

where all the time-varying parameters follow the standard random walk assumption we established in the previous section (but which we omit here for the sake of brevity). When $\beta_t(\tau) = 0$ the exogenous predictors $\boldsymbol{x}_t$ are absent, and we have the class of AR(2) models. When also $\phi_{1t}(\tau) = \phi_{2t}(\tau) = 0$ no lags of inflation are present and the model belongs to the class of time-varying intercept (local level) models. The assumption that ε_{t+h} follows the asymmetric Laplace distribution (ALD), coupled with the additional assumption that $\sigma_t(\tau) = \sigma(\tau)$ (see footnote 4), provides us with our proposed class of TVP quantile regression models. We already argued in footnote 9 that, when using the scale mixture of Normals representation of the ALD, we can obtain traditional TVP regression with stochastic volatility as a special case, and this result also holds for equation (30). Finally, we can obtain constant parameter regressions and quantile regressions simply by fixing time-varying parameters to be constant over time (e.g. by setting the state variance in equation (3) to zero). Therefore, we see that by placing appropriate restrictions in the specification above we can nest a wide class of popular forecasting models for inflation. These range from the TVP-QR, as the most flexible specification we can obtain from that equation, to the simple AR(2) model, being the most parsimonious special case. In particular, we consider forecasts from the following models

1. AR(2) model with constant parameters and variance (AR(2)) – this model is our benchmark upon which we measure the performance of all other models
2. AR(2) model with TVPs and stochastic volatility (TVP-AR-SV)
3. Time-varying intercept only model with stochastic volatility[10] (TVI-SV)
4. Quantile AR(2) with time-varying parameters (TVP-QAR)
5. Quantile regression model with time-varying intercept (TVI-QR)
6. Mean regressions with constant parameters, exogenous predictors, and stochastic volatility (AR-SV-X)
7. Mean regressions with time-varying parameters, exogenous predictors, and stochastic volatility (TVP-AR-SV-X)
8. Quantile AR(2) with constant parameters augmented with exogenous predictors (QAR-X), and
9. Quantile AR(2) with time-varying parameters augmented with exogenous predictors (TVP-QAR-X)

On top of these purely time-series models, we also consider a semi-structural Phillips Curve (PC) formulation of equation (29), similar to López-Salido and Loria (2019). In its most general form, the PC formulation is

$$\pi_{t+h} = (1-\lambda_t(\tau))\pi_t^* + \lambda_t(\tau)\pi_t^{LTE} + \theta_t(\tau)\left(y_t - y_t^*\right) + \gamma_t(\tau)\pi_t^I + \beta_t(\tau)\boldsymbol{x}_t + \varepsilon_{t+h}, \tag{30}$$

where π_t^* is lagged inflation (computed as the average over the previous four quarters), π_{t+h}^{LTE} are the long-term inflation expectations (measured using Consensus 6 to 10 years ahead inflation expectations), $(y_t - y_t^*)$ is the output gap (calculated as the principal component of available estimates), and π_{t+h}^I are relative prices (measured as the spread between import deflator inflation

[10]This model is similar to the unobserved components stochastic volatility (UCSV) model of Stock and Watson (2007b), although it does not assume stochastic volatility in the equation for trend inflation.

and domestic inflation). Accordingly, we consider the following specifications based on the PC restrictions:

1. Mean PC regression with stochastic volatility, no additional predictors (PC-SV)
2. Mean PC regression with time-varying parameters and stochastic volatility, no additional predictors (TVP-PC-SV)
3. Quantile PC regression, no additional predictors (QPC)
4. Quantile PC regression with time-varying parameters, no additional predictors (TVP-QPC)
5. Mean PC regression with stochastic volatility, with additional predictors (PC-SV-X)
6. Mean PC regression with time-varying parameters and stochastic volatility, with additional predictors (TVP-PC-SV-X)
7. Quantile PC regression, with additional predictors (QPC-X)
8. Quantile PC regression with time-varying parameters, with additional predictors (TVP-QPC-X)

All the models above, other than the AR(2) which is based on least squares, are estimated using the same default, automatic horseshoe prior we specified in the previous two sections. Whenever we consider models with exogenous predictors, we only estimate each class of models with one predictor at a time. Even though the horseshoe prior can accommodate ultra-high dimensional models, we are particularly interested in understanding the role of individual variables for forecasting inflation risks. For that reason, we do not consider here forecasting using the full model (all predictors), or principal components from the predictors, or forecast combinations. These are all reliable methods for improving forecast accuracy in any class of models, but they do not allow us to pin down the informational content of each individual predictor. Given these considerations, overall, 160 models are estimated, eight of which are models with no predictors and 152 of which are models with one individual financial indicator as an exogenous predictor.

Using data from 1990Q1 to 2019Q4, each model is estimated on the basis of the first half of the sample and used to produce four-quarters ahead and twelve-quarters ahead forecasts for 19 quantiles ($\tau = 0.05, 0.10, ..., 0.90, 0.95$), and thereafter forecasting follows a recursive scheme. We concentrate on the relative quantile scores at the first and last quantiles, that is the 5th and the 95th percentiles, to assess the ability of these models and indicators to predict low and high tail risks to inflation. The analysis will take into consideration the quantile score of each competing model relative to the benchmark AR(2) model, by taking an average of such scores across all the forecasting periods. For each competing model j, the relevant quantile score at each quantile τ is defined as:

$$QS_t^j(\tau) = [\pi_t - \hat{Q}_\tau(\pi_t|\boldsymbol{x}_t)][\mathbb{I}\{\pi_t \leq \hat{Q}_\tau(\pi_t|\boldsymbol{x}_t)\}]. \tag{31}$$

Smaller values of this loss function indicate better performance. We also use the predictive likelihood as a measure of general performance of the whole predictive density from each model we estimate (Korobilis, 2017). The predictive likelihood (PL) is obtained as the h-step ahead predictive density evaluated at the h-step ahead out-of-sample realisation of inflation. For the case of the PL, higher values indicate better performance.

4.3 Best performing categories of models

Starting with an assessment of the forecasting performance by category of model, in Table 2 we report the relative scores for each of the four univariate models, the average relative score of the four standard Phillips curve models (i.e., not augmented with any financial indicator), the average relative score of the four Phillips curve models augmented with a financial indicator, and the average relative score of the four bivariate models estimated with one additional financial indicator at a time, for the two forecast horizons considered. We report these relative scores for the first and last standard tails, alongside the Predictive Likelihood (PL) which provides an evaluation of the forecast of the whole distribution of inflation.

The results suggest that for the four-quarters ahead horizon, most of the model classes considered do not outperform the simple AR(2) benchmark model for the lowest quantile. Indeed, only one univariate model (TVP-QAR), one group of standard PC models on average (TVP-PC-SV) and one group of augmented PC models on average (TVP-PC-X) outperform the AR(2) model at the lowest quantile (Q5) (see second column in Table 2). For the higher quantile the relative performance of these models improves, with various quantile regression univariate and bivariate model classes with time-varying parameters (TVP-QAR, TVI-QR,and TVP-QAR-X) marking the strongest forecast improvements (see third column in Table 2). The relative scores for the PL suggest that, along with one univariate model (TVI-SV), the two quantile regression PC model categories with time-varying parameters (TVP-QPC, and TVP-QPC-X) on average outperform all other model classes (see fourth column in Table 2).

For the twelve-quarters ahead horizon, no model class considered does outperform the simple AR(2) benchmark model for the lowest quantile (see fifth column in Table 2). By contrast, for the highest quantile some model classes outperform the benchmark, especially two univariate models (TVP-AR-SV and TVI-SV) and a standard PC model class (QPC) (see sixth column in Table 2). The two univariate quantile regression models with time-varying parameters (TVP-QAR and TVI-QR) outperform all others when looking at the relative scores for the PL (see last column in Table 2).

Overall, looking at categories of models, it appears that the best performing specifications tend to be represented by quantile regressions with time-varying parameters, either for univariate or multivariate models. The classes of models including financial indicators most often do not appear to be among the best performing ones on average. However, as we will see next, the poor average performance of the latter models is due to the fact that within these classes of models some specific ones perform very well while other ones display poor forecasting properties.

Table 2: *Average relative scores by group of models for euro area core HICP inflation*

	4-quarters ahead			12-quarters ahead		
	QScore5	**Qscore95**	**PL**	**QScore5**	**Qscore95**	**PL**
TVP-AR-SV	1.496	0.896	1.158	2.423	0.752	1.202
TVI-SV	1.654	0.914	1.314	2.412	0.691	1.210
TVP-QAR	0.988	0.839	1.218	1.144	0.852	1.245
TVI-QR	1.025	0.853	1.260	1.235	0.826	1.234
PC-SV	1.519	0.974	1.077	1.823	0.935	0.888
TVP-PC-SV	0.909	0.938	1.245	2.495	1.244	0.804
QPC	2.383	1.543	1.204	3.356	0.782	1.066
TVP-QPC	1.129	0.926	1.308	3.652	0.959	0.928
PC-SV-X	1.560	1.039	1.051	2.222	0.951	0.890
TVP-PC-SV-X	0.973	0.963	1.214	2.362	1.231	0.820
QPC-X	2.411	1.583	1.175	3.413	0.989	1.070
TVP-QPC-X	1.152	0.994	1.285	2.614	0.922	0.976
AR-SV-X	1.166	0.905	1.095	1.412	0.923	1.030
TVP-AR-SV-X	1.388	0.908	1.180	2.181	0.873	1.099
QAR-X	1.440	0.921	1.261	2.161	0.892	1.168
TVP-QAR-X	1.023	0.843	1.243	1.163	0.895	1.194

Notes: The quantile score refers to the average of the $QS_t^j(\tau)$ over all forecasting periods relative to the AR(2) benchmark, as an average across across all competing models with the specific model specification. The table reports the ratios of the quantile score of the models described in each row to that of the AR(2) model. For the four quantiles shown, any value below 1 signals an improvement of the forecast relative to the benchmark model and the lower the value of such ratio the larger the improvement. For the Predictive Likelihood (PL) (fourth and last columns), any value above 1 signals an improvement of the forecast relative to the benchmark model and the higher the value of such ratio the larger the improvement.

4.4 Best performing specific models

When looking at the forecasting performance of the specific models, it appears that a number of credit and money volume indicators are particularly useful in predicting tail risks of core HICP inflation, outperforming several other financial indicators, especially in the context of quantile regression models featuring time-varying parameters.

For example, looking at forecasts for inflation tail risks four-quarters ahead (see upper panel in Table 3) a number of specific models appear to be very useful especially for forecasting the upper quantile (Q95). This is the case, in particular, of quantile regressions with time-varying parameters (TVP-QAR-X) featuring bank loans to the private sector and total credit to the private sector, which lead to forecast gains very close to the best performing model, that is the PC model with time-varying parameters and stochastic volatility (TVP-PC-SV-X) featuring bank loans to firms (see results under QScore95). For the lowest quantile (Q5) the improvement is more limited but non-negligible, especially for PC models with time-varying parameters and stochastic volatility (TVP-PC-SV-X) augmented with private sector loans, as well as for quantile regressions with time-varying parameters (TVP-QAR-X) featuring the M1 to GDP ratio or house prices (see results under QScore5). The overall density seems to be improved especially for quantile regressions with time-varying parameters (TVP-QAR-X) featuring the M1 to GDP ratio or loans to the private sector (see results under PL).

Also twelve-quarters ahead inflation tail risk forecasts can be improved, especially for upper tail risks (Q95), by considering quantile regression-based PC models with time-varying parameters (TVP-QPC-X) featuring loans to households (see lower panel in Table 3). The latter specific model also appears to lead to the strongest improvement for the overall density (see results under PL). By contrast, gains for the lowest quantile (Q5) are more limited but can still be detected especially for quantile regressions with time-varying parameters (TVP-QAR-X) featuring the ratio of private sector loans to GDP or the ratio to private sector total credit to GDP.

Overall, when assessing specific models, financial volume indicators such as loans to the private sector, loans to firms, loans to households, total credit to the private sector and the ratio of M1 to GDP often provide the largest forecasting gains, especially in the context of quantile regressions with time-varying parameters (TVP-QAR-X) or quantile regression-based Phillips curve models with time-varying parameters (TVP-QPC-X).

The good forecasting properties of credit and money indicators for inflation tail risks appear consistent with economic intuition. Credit growth indicators, the coefficients of which tend to be positive within the estimated quantile regression models considered, can be seen as proxies of the degree of tightness of financial constraints for firms and households, which affect the price setting decisions of firms and the spending decisions of households. By contrast, the estimated coefficient for the ratio of M1 to GDP tends to be negative, which could reflect the fact that increases in M1 in excess of GDP are likely to capture increased uncertainty which is typically associated with the postponement of expenditure thereby creating disinflationary pressures.

Table 3: *Best financial indicators and models for the prediction of core HICP inflation tail risks*

Measure	Ranking	Indicator	Specification	Score
		4-QUARTERS AHEAD		
	1st	loans to private sector	TVP-PC-SV-X	0.891
QScore5	2nd	M1/GDP	TVP-QAR-X	0.900
	3rd	house prices	TVP-QAR-X	0.900
	1st	loans to firms	TVP-PC-SV-X	0.761
QScore95	2nd	loans to private sector	TVP-QAR-X	0.767
	3rd	credit to private sector	TVP-QAR-X	0.780
	1st	M1/GDP	TVP-QAR-X	1.474
PL	2nd	loans to private sector	TVP-QAR-X	1.429
	3rd	house prices	QAR-X	1.383
		12-QUARTERS AHEAD		
	1st	private sector loans/GDP	TVP-QAR-X	0.951
QScore5	2nd	private sector credit/GDP	TVP-QAR-X	0.962
	3rd	yield curve	TVP-QAR-X	0.963
	1st	loans to households	TVP-QPC-X	0.635
QScore95	2nd	private sector credit/GDP	QAR-X	0.685
	3rd	loans to households/GDP	QAR-X	0.692
	1st	loans to households	TVP-QPC-X	1.552
PL	2nd	loans to households	QAR-X	1.336
	3rd	loans to households/GDP	QAR-X	1.295

Notes: PC stands for Phillips curve, QAR for quantile autoregression with two lags, AR for mean autoregression with two lags, TVP for time-varying parameters and SV for stochastic volatility. Low inflation tail risks are captured by the 5th percentile (Q5) while high inflation tail risks are captured by the 95th percentile (Q95). The lower blocks for each horizon reports the Predictive Likelihood (PL) which provides an evaluation of the forecast of the whole distribution of inflation. The last column reports the ratios of the predictive quantile score of the models described in each row to that of the AR(2) model. For the quantiles, a value below 1 signals an improvement of the forecast relative to the benchmark and the lower the value the larger is the improvement. For the PL, any value above 1 signals an improvement of the forecast relative to the benchmark model and the higher the value of such ratio the larger the improvement.

4.5 The role of financial indicators

In order to understand what role the inclusion of financial indicators might play, it can be interesting to compare the evolution over time of the upper and lower quantile scores, as a measure of tail risk forecast errors, for different models. By comparing the evolution of the quantile scores of the best models for the upper and lower quantiles (Q5 and Q95) among the quantile regression models with time-varying parameters - which are often among the best and which are the main focus of the paper - as highlighted in Table 3, with those of a similar model without the financial indicator we can derive some indication on the importance of considering such additional indicator when forecasting inflation tail risks. A similar comparison with the quantile scores of the same models with constant parameters, but all else equal, can give an idea on the importance of this specific modelling choice.

Starting with the best TVP-QAR model for the prediction of four-quarters ahead inflation low quantiles (Q5), that is the TVP-QAR model featuring the M1 to GDP ratio, and comparing its quantile score with that of the TVP-QAR model (i.e., without any financial indicator) we can notice that the latter model leads to additional large forecast errors especially during the global financial crisis, that is around 2008 and 2009 (see left-hand side chart of the top panel of Figure 2). Similarly, the model with constant parameters (QAR-X) for M1 to GDP produces additional large forecast errors not only in 2008 and 2009, though smaller than those of the TVP-QAR model, but also in 2013, 2016 and 2017, by contrast to the benchmark TVP-QAR model for M1 to GDP (see right-hand side chart of the top panel of Figure 2).

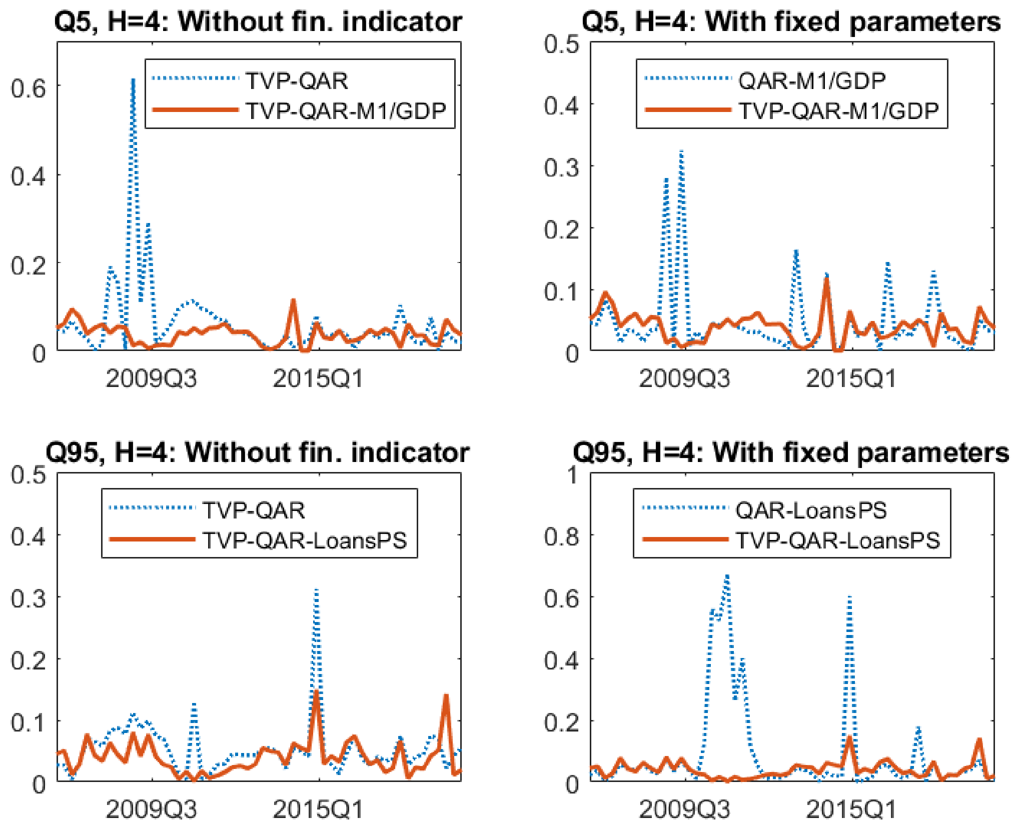

Figure 2: *Quantile score evolution for the best TVP-QAR-X model and alternative models: 4-quarter ahead*

As regards the best TVP-QAR model for the prediction of the highest quantile (Q95) for infation four-quarters ahead, that is the TVP-QAR model featuring loans to the private sector, the same model without this financial indicator leads to higher forecasting errors both in both prolonged periods, such as between 2007 and 2009, and in specific quarters, such as in mid-2011 and late 2014 (see left-hand side chart of the bottom panel of Figure 2). The same model with constant parameters implies much larger forecast errors especially between 2010 and 2011 and in late-2014, compared to the model with time-varying parameters (see right-hand side chart of the bottom panel of Figure 2).

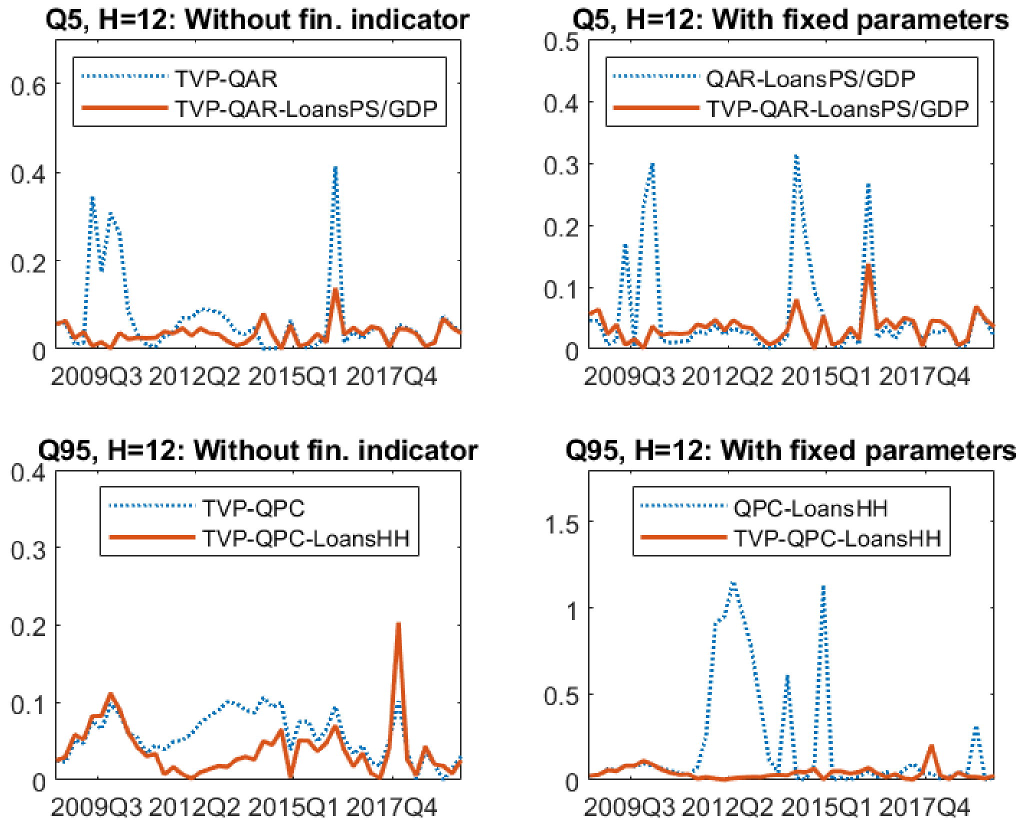

Figure 3: *Quantile score evolution for the best TVP-QAR-X model and alternative models: 12-quarter ahead*

A similar exercise with the best TVP-QAR model for the prediction of the lowest quantile (Q5) of inflation twelve-quarters ahead, that is the TVP-QAR model featuring the ratio of loans to the private sector to GDP, suggests that the same model without this financial indicator leads to higher forecasting errors around 2009 and in mid-2016 (see left-hand side chart of the top panel of Figure 3). The same model with constant parameters implies much larger forecast errors especially around 2009, 2014 and 2016, compared to the model with time-varying parameters (see right-hand side chart of the top panel of Figure 3). As regards the upper quantile (Q95), taking the time-varying parameters quantile regression PC augmented with loans to households as reference, the forecast error of the same model either without financial indicator or with fixed parameters are clearly higher for prolonged periods (bottom panel of Figure 3).

Overall, for the prediction of tail risks of inflation both one-year and three-years ahead, it appears that the inclusion of the financial indicator outperforming other ones within TVP-QAR-X or TVP-QPC-X models explains a significant fraction of the reduction in forecasting errors compared to competing specific models. However, also the key modelling specification characterising such models, that is time-varying parameters, seems to be instrumental in explaining the good forecasting performance of the models highlighted.

5 Conclusion

We develop a methodology for modelling and forecasting inflation risks flexibly via time-varying parameter quantile regressions. A key methodological contribution is represented by a new Gibbs sampler for time-varying parameters that is highly efficient and can be easily adapted to other settings that admit familiar state-space forms and are notoriously computationally intensive, such as dynamic factor models and time-varying parameter VARs. In terms of forecasting accuracy, we show that our feasible TVP-QAR model indeed provides improvements over its constant parameter alternative, as well as traditional TVP regressions with stochastic volatility.

An application of this methodology to the prediction of euro area core inflation tail risks with data spanning the past thirty years points to a very good forecasting performance of quantile regressions with time-varying parameters augmented with specific credit and money-based indicators, both in the short and the medium run.

Follow-up work will concentrate on assessing the ability of the proposed modelling framework to enhance the prediction of the tail risks of other core macroeconomic and financial variables, such as real GDP and asset prices, both in the euro area and in other countries.

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019): "Vulnerable Growth," *American Economic Review*, 109, 1263–89.

Amir-Ahmadi, P., C. Matthes, and M.-C. Wang (2020): "Choosing Prior Hyperparameters: With Applications to Time-Varying Parameter Models," *Journal of Business & Economic Statistics*, 38, 124–136.

Bhadra, A., J. Datta, Y. Li, and N. Polson (2020): "Horseshoe Regularisation for Machine Learning in Complex and Deep Models1," *International Statistical Review*, 88, 302–320.

Bhattacharya, A., A. Chakraborty, and B. K. Mallick (2016): "Fast sampling with Gaussian scale mixture priors in high-dimensional regression," *Biometrika*, 103, 985–991.

Brownlees, C. and A. B. Souza (2021): "Backtesting global Growth-at-Risk," *Journal of Monetary Economics*, 118, 312–330.

Cai, Z. and X. Xu (2009): "Nonparametric Quantile Estimations for Dynamic Smooth Coefficient Models," *Journal of the American Statistical Association*, 104, 371–383.

Carriero, A., T. E. Clark, and M. Marcellino (2020): "Nowcasting Tail Risks to Economic Activity with Many Indicators," Working Paper No. 20-13, Federal Reserve Bank of Cleveland.

Carvalho, C. M., N. G. Polson, and J. G. Scott (2010): "The Horseshoe Estimator for Sparse Signals," *Biometrika*, 97, 465–480.

Chan, J. and I. Jeliazkov (2009): "Efficient simulation and integrated likelihood estimation in state space models," *International Journal of Mathematical Modelling and Numerical Optimisation*, 1, 101–120.

Cooley, T. F. and E. C. Prescott (1976): "Estimation in the Presence of Stochastic Parameter Variation," *Econometrica*, 44, 167–184.

Faust, J. and J. H. Wright (2013): "Forecasting Inflation," in *Handbook of Forecasting*, Elsevier, vol. 2, chap. 1, 2–56.

Gerlach, R. H., C. W. S. Chen, and N. Y. C. Chan (2011): "Bayesian Time-Varying Quantile Forecasting for Value-at-Risk in Financial Markets," *Journal of Business & Economic Statistics*, 29, 481–492.

Ghosh, P., X. Tang, M. Ghosh, and A. Chakrabarti (2016): "Asymptotic Properties of Bayes Risk of a General Class of Shrinkage Priors in Multiple Hypothesis Testing Under Sparsity," *Bayesian Anal.*, 11, 753–796.

Gonçalves, K. C. M., H. S. Migon, and L. S. Bastos (2020): "Dynamic Quantile Linear Models: A Bayesian Approach," *Bayesian Analysis*, 15, 335–362.

Johnson, N. L., S. Kotz, and N. Balakrishnan (1994): *Continuous univariate distributions. Vol. 1.*, New York: Wiley & Sons, 2nd ed.

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

Kim, M.-O. (2007): "Quantile regression with varying coefficients," *Annals of Statistics*, 35, 92–108.

Koenker, R. and G. Bassett (1978): "Regression Quantiles," *Econometrica*, 46, 33–50.

Koop, G. and D. Korobilis (2012): "FORECASTING INFLATION USING DYNAMIC MODEL AVERAGING*," *International Economic Review*, 53, 867–886.

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

——— (2021): "High-Dimensional Macroeconomic Forecasting Using Message Passing Algorithms," *Journal of Business & Economic Statistics*, 39, 493–504.

Kozumi, H. and G. Kobayashi (2011): "Gibbs sampling methods for Bayesian quantile regression," *Journal of Statistical Computation and Simulation*, 81, 1565–1578.

Lim, D., B. Park, D. Nott, X. Wang, and T. Choi (2020): "Sparse signal shrinkage and outlier detection in high-dimensional quantile regression with variational Bayes," *Statistics and Its Interface*, 13, 237–249.

Lopez-Salido, D. and F. Loria (2020): "Inflation at Risk," SSRN Scholarly Paper ID 3600420, Social Science Research Network, Rochester, NY.

López-Salido, J. D. and F. Loria (2019): "Inflation at Risk," Discussion Papers No. 14074, Center for Economic Policy Research.

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

Rodrigues, T. and Y. Fan (2017): "Regression Adjustment for Noncrossing Bayesian Quantile Regression," *Journal of Computational and Graphical Statistics*, 26, 275–284.

Rossi, B. and T. Sekhposyan (2019): "Alternative tests for correct specification of conditional predictive densities," *Journal of Econometrics*, 208, 638–657.

Stock, J. H. and M. W. Watson (2007a): "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking*, 39, 3–33, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1538-4616.2007.00014.x.

——— (2007b): "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking*, 39, 3–33.

van der Pas, S. L., B. J. K. Kleijn, and A. W. van der Vaart (2014): "The horseshoe estimator: Posterior concentration around nearly black vectors," *Electronic Journal of Statistics*, 8, 2585–2618.

Wu, W. and Z. Zhou (2017): "Nonparametric Inference for Time-Varying Coefficient Quantile Regression," *Journal of Business & Economic Statistics*, 35, 98–109.

Yu, K. (2017): "Bayesian Quantile Regression," *https://expertsinuncertainty.net*, accessed Nov. 08, 2020.

Yu, K. and R. A. Moyeed (2001): "Bayesian quantile regression," *Statistics & Probability Letters*, 54, 437 – 447.

Online Appendix to
"The time-varying evolution of inflation risks"

Dimitris Korobilis
University of Glasgow

Bettina Landau
European Central Bank

Alberto Musso
European Central Bank

Anthoulla Phella
European Central Bank

A Bayesian inference in the quantile regression model

A.1 Linear quantile regression setting

Following Yu and Moyeed (2001) the quantile regression model has a parametric representation

$$y_t = x_t \beta(\tau) + \sigma(\tau)\varepsilon_t, \tag{A.1}$$

where $\beta(\tau)$ and $\sigma(\tau)$ are the regression coefficients and the scale parameter, respectively, for each quantile level τ, and ε_t are i.i.d. from a joint Asymmetric Laplace density of the form $\prod_{t=1}^{T} \frac{\tau(1-\tau)}{\sigma(\tau)^2} \left[e^{(1-\tau)\frac{\varepsilon_t}{\sigma(\tau)^2}} \mathbb{I}(\varepsilon_t \leq 0) + e^{-\tau \frac{\varepsilon_t}{\sigma(\tau)^2}} \mathbb{I}(\varepsilon_t > 0) \right]$. We can write the Asymmetric Laplace distribution using the following mixture representation (cf Kozumi and Kobayashi, 2011)

$$y_t = x_t \beta(\tau) + \theta(\tau) z_t + \sigma(\tau)\kappa(\tau)\sqrt{z_t(\tau)} u_t, \qquad u_t \sim N(0,1), \tag{A.2}$$

where $\theta(\tau) = \frac{1-2\tau}{\tau(1-\tau)}$ and $\kappa(\tau)^2 = \frac{2}{\tau(1-\tau)}$ and $z_t(\tau) \sim exp(\sigma(\tau)^2)$. Under this mixture representation the density of y_t is of the form

$$\prod_{t=1}^{T} \frac{1}{\sqrt{2\pi z_t(\tau)\sigma(\tau)^2\kappa(\tau)^2}} \exp\left\{ -\frac{\left(y_t - x_t\beta(\tau) - \theta(\tau) z_t(\tau)\right)^2}{2 z_t(\tau)\sigma(\tau)^2\kappa(\tau)^2} \right\} \exp\left\{ -\frac{z_t(\tau)}{\sigma(\tau)^2} \right\}, \tag{A.3}$$

which conditionally on $z_t(\tau)$ is a Normal density, while marginalizing (i.e. integrating) over the unknown $z_t(\tau)$ gives the desired Asymmetric Laplace density (see the Appendix of Khare and Hobert, 2012, for a proof).

Given priors of the form

$$\beta(\tau) \sim N(0, V(\tau)), \tag{A.4}$$
$$\sigma(\tau) \sim IG(\rho_1, \rho_2), \tag{A.5}$$
$$z_t(\tau) \sim exp(\sigma(\tau)), \tag{A.6}$$

for each $\tau = 0.05, 0.10, ..., 0.90, 0.95$ [1], we obtain conditional posteriors of the form

$$\beta(\tau)|\bullet \sim N\left(\left(x'Ux + V(\tau)^{-1}\right)^{-1} \times \left(x'U\left[y - \theta(\tau)z(\tau)\right]\right), \left(x'Ux + V(\tau)^{-1}\right)^{-1}\right), \tag{A.7}$$

$$\sigma(\tau)^2|\bullet \sim inv - Gamma\left(\rho_1 + \frac{3T}{2}, \rho_2 + \sum_{t=1}^{T}\frac{(y_t - x_t\beta(\tau) - \theta(\tau)z_t(\tau))^2}{2z_t(\tau)\kappa(\tau)^2} + \sum_{t=1}^{T} z_t(\tau)\right), \tag{A.8}$$

$$z_t(\tau)|\bullet \sim IG\left(\frac{\sqrt{\theta(\tau)^2 + 2\kappa(\tau)^2}}{|y_t - x_t\beta(\tau)|}, \frac{\theta(\tau)^2 + 2\kappa(\tau)^2}{\sigma(\tau)^2\kappa(\tau)^2}\right), \tag{A.9}$$

where $x = [x_1', ..., x_T']'$, $x = [y_1, ..., y_T]'$, and U is a $T \times T$ diagonal covariance matrix with t-th element $(z_t(\tau)\sigma(\tau)^2\kappa(\tau)^2)^{-1}$. In equation (A.9) IG denotes the Inverse Gaussian distribution.[2] Kozumi and Kobayashi (2011) derive a posterior for $z_t(\tau)$ as in equation (A.9) that is Generalized Inverse Gaussian (GIG) with different hyperparameters. The expression in (A.9) is derived by noting the property that the $IG(\mu, \lambda)$ distribution is a $GIG(\lambda/\mu^2, \lambda, -1/2)$ distribution, see Johnson et al. (1994).

A.2 Dealing with high-dimensional settings

Before we proceed with the time-varying parameter (TVP) version of the previous sampling algorithm, we discuss how we deal with high-dimensional versions of the Bayesian quantile regression model (the TVP quantile regression is such a model). Our focus is both on fast and efficient computation, as well as automatic shrinkage of the vector of regression coefficients.

Unlike the regular regression model, notice that in the quantile regression shrinkage is imperative even in the case where the number of predictors, p, is much smaller relative to the number of time series observations, T, that is, even when $p \ll T$. This is because we need to estimate a p-dimensional vector $\beta(\tau)$ for each quantile level $\tau = 0.05, 0.10, ..., 0.90, 0.95$. While around the median nearby quantiles can help assist estimate the β's more accurately, estimation in more extreme quantiles will rely on only a small part of the sample. Therefore, even small or moderate values of p can induce large estimation error of unrestricted estimators.

We specify a hiearchical shrinkage prior for $\beta(\tau)$ and, in particular, we follow Bhattacharya et al. (2016) who adopt a horseshoe prior of the form

$$\beta(\tau)_i|\lambda(\tau)^2, \psi_i(\tau)^2 \sim N\left(0, \lambda(\tau)^2\psi_i(\tau)^2\right), \tag{A.10}$$

$$\lambda(\tau) \sim Cauchy^+(0, 1), \tag{A.11}$$

$$\psi_i(\tau) \sim Cauchy^+(0, 1), \tag{A.12}$$

[1] Note that each $\beta(\tau)$ has its own prior variance $V(\tau)$ for each quantile level τ, while $\sigma(\tau)$ depends on the same hyperparmeters ρ_1, ρ_2 for all τ. This is because it is trivial to be noninformative for the variance parameters $\sigma(\tau)$, while this is not the case for $\beta(\tau)$: as we detail in the next subsection we are interested in regularizing this parameter for the sake of estimation accuracy, therefore, $V(\tau)$ will be estimated adaptively (i.e. for each $\tau = 0.05, 0.10, ..., 0.90, 0.95$).

[2] While the inverse of a Gamma variate is distributed inverse Gamma, and vice-versa, the same is not true for the Normal (Gaussian) and Inverse Gaussian distributions.

for $i = 1, ..., p$ and $\tau = 0.05, 0.10, ..., 0.90, 0.95$. The conditional posteriors of λ, ψ_i can be obtained if we consider the formulation of the horseshoe prior adopted in Makalic and Schmidt (2016). These authors write the horseshoe prior using the equivalent hierarchical notation

$$\begin{aligned} \beta(\tau)_i|\lambda(\tau)^2, \psi_i(\tau)^2 &\sim N(0, \lambda(\tau)^2\psi_i(\tau)^2), & \text{(A.13)} \\ \lambda(\tau)^2|\xi(\tau) &\sim inv - Gamma\,(1/2, 1/\xi(\tau))\,, & \text{(A.14)} \\ \xi(\tau) &\sim inv - Gamma\,(1/2, 1)\,, & \text{(A.15)} \\ \psi_i(\tau)^2|\zeta_i(\tau) &\sim inv - Gamma\,(1/2, 1/\zeta_i(\tau))\,, & \text{(A.16)} \\ \zeta_i(\tau) &\sim inv - Gamma\,(1/2, 1)\,. & \text{(A.17)} \end{aligned}$$

Conditional posteriors under this prior formulation are trivial to derive and exact formulas can be found in Makalic and Schmidt (2016).

The next step in our analysis is to sample efficiently the large dimensional vector $\beta(\tau)$ using equation (A.7), especially when $p \gg T$. We follow again Bhattacharya et al. (2016) who propose an efficient way to sample from the Normal distribution using the Woodbury matrix identity. Calculation of the posterior covariance matrix of $\beta(\tau)$ relies on inverting the $p \times p$ matrix $\left(x'Ux + V^{-1}\right)$ which requires $\mathcal{O}(p^3)$ algorithmic operations. The same number of algorithmic operations are needed to obtain the Cholesky decomposition of the posterior covariance, which is essential in order to generate from the desired Normal posterior distribution. In high dimensions, that is when p gets large, both these operations become computationlly cumbersome. Bhattacharya et al. (2016) propose instead the following sampling scheme[3]:

Algorithm for efficient sampling from (A.7)

Step 1	Sample $\eta \sim N(0, V(\tau))$ and $\delta \sim N(0, I_T)$
Step 2	Set $v = \widetilde{x}\eta + \delta$
Step 3	Set $w = (\widetilde{x}V\widetilde{x}' + I_T)^{-1}[y - \theta(\tau)z(\tau) - v]$
Step 4	Set $\beta(\tau) = \eta + V\widetilde{x}'w$

where $\widetilde{x} = xU^{-1/2}$ where $U^{-1/2}$ is a $T \times T$ diagonal matrix with elements $(\sqrt{z_t(\tau)}\sigma(\tau)\kappa(\tau))^{-1}$ on its main diagonal. Instead of generating from a p-variate Normal posterior distribution, the algorithm above involves generating from the p-variate Normal prior distribution, and a T-variate standard Normal. As long as the prior covariance matrix is diagonal, generating from $\eta \sim N(0, V)$ is computationally trivial. Similar arguments hold for $\delta \sim N(0, I_T)$. The remaining transformations in the algorithm result from the Woodbudy identity (see Bhattacharya et al., 2016, for a straightforward proof) and they can be extremely efficient. The worst case complexity

[3] This algorithm only works when the prior covariance matrix is diagonal, which is the case with the Horseshoe prior and the vast majority of Bayesian prior specifications.

of the algorithm is $\mathcal{O}(T^2p)$, which provides huge gains relative to inverting and taking the Cholesky factor of the $p \times p$ matrix $\left(x'Ux + V^{-1}\right)$.

Since our algorithm for estimating the quantile regression model is written in MATLAB, further gains can be achieved by replacing for loops with vector operators. This is very relevant here, because the algorithm requires for each MCMC iteration to also iterate through equations (A.7)-(A.9) for each quantile level τ. MATLAB allows to generate from matrix-variate versions of the required posterior distributions, such that all parameters can be generated at once $\forall$ $\tau \in \{0.05, 0.10, ..., 0.90, 0.95\}$.

A.3 Bayesian inference in the quantile regression model with time-varying parameters

Following the previous sections, the Bayesian time-varying parameter quantile regression model can be written as

$$y_t = x_t\beta(\tau)_t + \varepsilon_t, \qquad \varepsilon_t \sim ALD(\sigma(\tau)^2), \tag{A.18}$$

$$\beta_t(\tau) = \beta_{t-1}(\tau) + v_t, \qquad v_t \sim N(0, V(\tau)) \tag{A.19}$$

subject to the initial condition $\beta_0(\tau) \sim N\left(0, V_0(\tau)\right)$, where x_t is a $1 \times p$ vector of predictors, and $V(\tau)$ is a $p \times p$ covariance matrix. Given that the ALD distribution for ε_t admits a mixture of Normals representation, it is trivial to treat the system above as a linear, conditionally Gaussian state-space model. However, doing so would result in recursive sampler that would be quite inefficient.

We first note that the t-th observation y_t can be solved for $\Delta\beta_t(\tau) = \beta_t(\tau) - \beta_{t-1}(\tau)$ as

$$y_t = x_t\beta_t(\tau) + \varepsilon_t \tag{A.20}$$

$$= x_t\Delta\beta_t(\tau) + x_t\beta_{t-1}(\tau) + \varepsilon_t \tag{A.21}$$

$$= x_t\Delta\beta_t(\tau) + x_t\Delta\beta_{t-1}(\tau) + x_t\beta_{t-2}(\tau) + \varepsilon_t \tag{A.22}$$

$$... \tag{A.23}$$

$$= x_t\Delta\beta_t(\tau) + x_t\Delta\beta_{t-1}(\tau) + ... + x_t\Delta\beta_2(\tau) + x_t\beta_1(\tau) + \varepsilon_t \tag{A.24}$$

which shows that the coefficients at time t, $\beta(\tau)_t$ is simply the cumulative sum of changes over the previous time periods. More intuition can be built if we stack for all observations t and rewrite equations (A.18)-(A.19) in the form

$$\begin{bmatrix} y_1 \\ y_2 \\ ... \\ y_{T-1} \\ y_T \end{bmatrix} = \begin{bmatrix} x_1 & 0 & ... & 0 & 0 \\ x_2 & x_2 & ... & 0 & 0 \\ ... & ... & ... & ... & ... \\ x_{T-1} & x_{T-1} & ... & x_{T-1} & 0 \\ x_T & x_T & ... & x_{T-1} & x_T \end{bmatrix} \begin{bmatrix} \beta_1(\tau) \\ \Delta\beta_2(\tau) \\ ... \\ \Delta\beta_{T-1}(\tau) \\ \Delta\beta_T(\tau) \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ ... \\ \varepsilon_{T-1} \\ \varepsilon_T \end{bmatrix}, \quad \text{(A.25)}$$

$$\begin{bmatrix} \beta_1(\tau) \\ \Delta\beta_2(\tau) \\ ... \\ \Delta\beta_{T-1}(\tau) \\ \Delta\beta_T(\tau) \end{bmatrix} = \begin{bmatrix} v_1 \\ v_2 \\ ... \\ v_{T-1} \\ v_T \end{bmatrix} \quad \text{(A.26)}$$

or more compactly

$$y = \mathcal{X}\beta^{\Delta}(\tau) + \varepsilon, \quad \text{(A.27)}$$

$$\beta^{\Delta}(\tau) = v, \quad \text{(A.28)}$$

where $\beta^{\Delta} = [\Delta\beta_1(\tau)', \Delta\beta_2(\tau)', ..., \Delta\beta_T(\tau)]'$ and $\mathcal{X}$ is the block triangular matrix shown analytically above. The first equation is a static linear regression with parameters $\beta^{\Delta}(\tau)$. The main characteristic of this equation is that it is represents a high-dimensional setting, since the lower-triangular matrix $\mathcal{X}$ has dimensions $T \times Tp$, i.e. more covariates than observations. The second equation is an identity and, instead of having the interpretation of a state equation, it can be seen as a standard Normal prior for the difference between $\beta_t(\tau)$ and $\beta_{t-1}(\tau)$. This high-dimensional representation shows clearly why shrinkage in TVP models is imperative, and why choice of the state covariance matrix ($V(\tau)$ here) is of paramount importance; see the discussion in Amir-Ahmadi et al. (2020).

It is notable that this representation of the TVP regression is equivalent to a minor reparametrization of the formulation and algorithm of Chan and Jeliazkov (2009). Using their methods, the time-varying parameter quantile regression would be written as

$$y = \mathsf{X}\beta(\tau) + \varepsilon, \quad \text{(A.29)}$$

$$\beta(\tau) = H^{-1}v, \quad \text{(A.30)}$$

where

$$H = \begin{bmatrix} I_p & 0 & \dots & 0 & 0 \\ -I_p & I_p & \ddots & 0 & 0 \\ \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & -I_p & I_p & 0 \\ 0 & 0 & 0 & -I_p & I_p \end{bmatrix}, \qquad \mathsf{X} = \begin{bmatrix} x_1 & 0 & \dots & 0 & 0 \\ 0 & x_2 & \ddots & 0 & 0 \\ \vdots & \ddots & \ddots & 0 & \vdots \\ 0 & \dots & 0 & x_{T-1} & 0 \\ 0 & \dots & 0 & 0 & x_T \end{bmatrix} \tag{A.31}$$

and $\beta = [\beta_1(\tau)', ..., \beta_T(\tau)']'$. Note that if we compare (A.29) with (A.27) it holds that:

$$\mathsf{X}\beta(\tau) = \mathsf{X}H^{-1}H\beta(\tau) = (\mathsf{X}H^{-1})(H\beta(\tau)) = \mathcal{X}\beta^{\Delta}(\tau). \tag{A.32}$$

Similarly, if we left-multiply both sides of (A.30) with H, we obtain (A.28), showing that our suggested specification is a simple rotation of Chan and Jeliazkov (2009) using matrix H that creates first differences of the coefficients $\beta(\tau)$ and turns the diagonal matrix X into a lower triangular matrix. Despite the similarities, the formulation we propose makes full use of the fast algorithm of Bhattacharya et al. (2016) because the prior distribution of $\beta^{\Delta}(\tau)$ is diagonal (i.e. easy to sample from), while the distribution of $\beta(\tau)$ is tridiagonal; see Chan and Jeliazkov (2009) for its exact form. Therefore, our proposed Gibbs sampler is faster than the one proposed by Chan and Jeliazkov (2009), especially in high-dimensions (large p) and in the case of the quantile regression model where we need to sample TVPs for each quantile level τ.

Replacing ε with its mixture of Normal representation of the previous subsection, and adding the Horseshoe prior on $V(\tau)$ we can write the full Bayesian TVP quantile regression model using the following equations

$$\begin{aligned} y &= \mathcal{X}\beta^{\Delta}(\tau) + \theta(\tau)z(\tau) + \widetilde{(S)}u, && \text{(A.33)} \\ \beta^{\Delta}(\tau) &\sim N(0, V(\tau)), && \text{(A.34)} \\ V_{i,i}(\tau) &= \sigma(\tau)^2\lambda(\tau)^2\psi_i(\tau)^2, \quad i = 1, ..., Tp, && \\ \lambda(\tau) &\sim Cauchy^{+}(0,1), && \text{(A.35)} \\ \psi_i(\tau) &\sim Cauchy^{+}(0,1), && \text{(A.36)} \\ \sigma(\tau) &\sim IG(\rho_1, \rho_2), && \text{(A.37)} \\ z_t(\tau) &\sim exp(\sigma(\tau)), \quad t = 1, ..., T, && \text{(A.38)} \end{aligned}$$

where $\widetilde{S}$ is a $T \times T$ diagonal matrix with diagonal element $\sigma(\tau)\kappa(\tau)\sqrt{z_t(\tau)}$.

A.4 Noncrossing quantiles

All estimation algorithms for quantile regression models, included the one presented above, assume that the quantile curves are fitted independently from each other. If we write the quantile regression model using the generic form

$$y_t = \mathcal{Q}_\tau(y_t|x_t) + \varepsilon_t, \tag{A.39}$$

then each curve $\mathcal{Q}_\tau(y_t|x_t)$ is estimated independently for each $\tau = 0.05, 0.1, ..., 0.90, 0.95$. This independence means that there no mechanism in place in order to guarantee that estimates $\widehat{\mathcal{Q}}_\tau(y_t|x_t)$ satisfy the very definition of a quantile, that is the fact that $\widehat{\mathcal{Q}}_{\tau_1}(y_t|x_t) < \widehat{\mathcal{Q}}_{\tau_2}(y_t|x_t)$ when $\tau_1 < \tau_2$. This condition is also known as "quantile noncrossing". Various algorithms have been proposed in the literature to deal with this issue. Most algorithms propose to post-process the estimated quantile functions using some smoothing procedure/function. Such post-processing can work well, however, inevitably, will introduce some bias in quantile estimates, therefore choice of an appropriate algorithm is essential.

Here we use the recently proposed algorithm of Rodrigues and Fan (2017) for Bayesian quantile regression. This algorithm involves to first use a consistent MCMC-based estimator to obtain quantile regression estimates (such as the one outlined above), and then use a Gaussian process regression to smooth out the quantile estimates. In order to achieve this, Rodrigues and Fan (2017) note that exactly because adjacent quantiles are correlated, one can use the following auxiliary model

$$\mathcal{Q}_{\tau,\tau^\star}(y_t|x_t) = \begin{cases} x_t\beta(\tau^\star) + \frac{\sigma(\tau^\star)}{1-\tau^\star}\log\left(\frac{\tau}{\tau^\star}\right), & if \quad 0 \leq \tau \leq \tau^\star, \\ \\ x_t\beta(\tau^\star) - \frac{\sigma(\tau^\star)}{\tau^\star}\log\left(\frac{1-\tau}{1-\tau^\star}\right), & if \quad \tau^\star \leq \tau \leq 1, \end{cases} \tag{A.40}$$

where $\mathcal{Q}_{\tau,\tau^\star}(y_t|x_t)$ is the induced quantile, and $\tau, \tau^\star \in \{0.05, 0.10, ..., 0.90, 0.95\}$. When $\tau = \tau^\star$ then $\mathcal{Q}_{\tau,\tau^\star}(y_t|x_t) \equiv \mathcal{Q}_\tau(y_t|x_t)$, that is, the induced quantile is equivalent to the estimated quantile based on our model. However, for all other levels of $\tau^\star$ we obtain additional induced quantile values that provide information for the quantile curve at τ. The principle is that the closer $\tau^\star$ is to τ, the more information its quantile curve can provide for estimation of the quantile curve at τ.

Given that there are 19 values in the set $\tau, \tau^\star \in \{0.05, 0.10, ..., 0.90, 0.95\}$, in our application $\mathcal{Q}_{\tau,\tau^\star}(y_t|x_t)$ is a 19×19 matrix. The diagonal elements of this matrix are identical to $\mathcal{Q}_\tau(y_t|x_t)$. Rodrigues and Fan (2017) specify a Gaussian process regression that ends up being equivalent to a weighting scheme where for each τ the quntiles $\mathcal{Q}_{\tau,\tau^\star}(y_t|x_t)$ take increasingly more weight the closer $\tau^\star$ is to τ. That way, the induced quantile at $\tau = \tau^\star$ (i.e. $\mathcal{Q}_\tau(y_t|x_t)$) takes the most weight and very distant quantiles get decreasing weights. It is very trivial to specify and implement this weighting/smoothing scheme, and we refer the reader to Rodrigues and Fan (2017) for more details.[4] Proposition 2 in that paper shows that this smoothed estimate of the quantiles is

[4]Their method introduces two new parameters, σ_κ^2 and b (using their notation). We follow the authors and set $\sigma_\kappa^2 = 100$ and we select the minimum b that provides a non-crossing solution.

consistent, and a Monte Carlo study supports the good properties of this method.

B Data

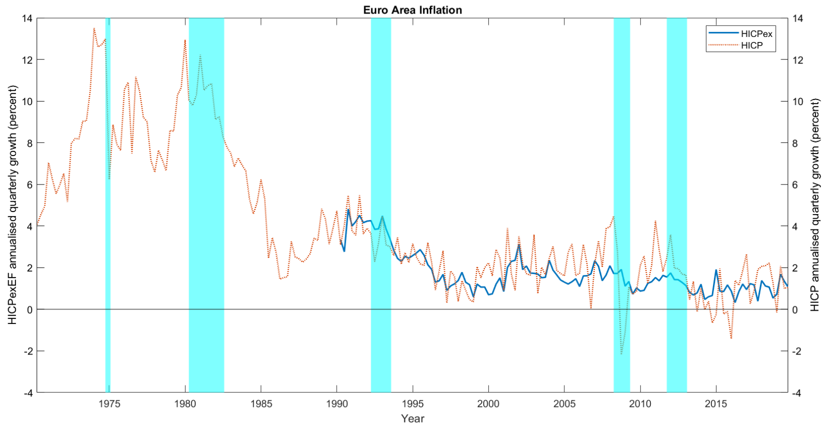

Figure B1: Euro area inflation

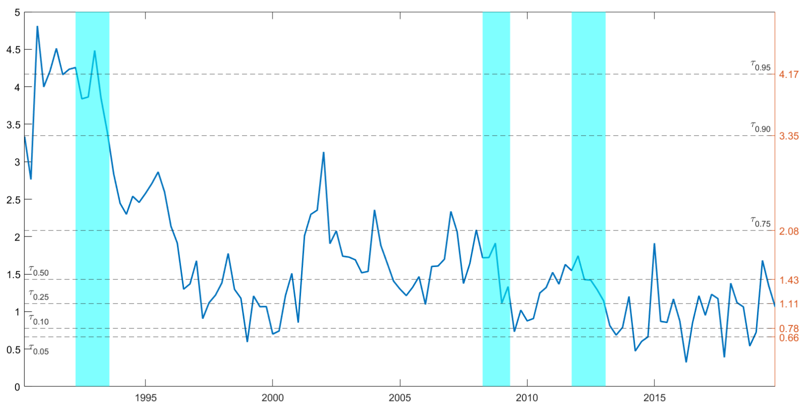

Figure B2: Euro area core inflation

Table B1: List of euro area Indicators

VARIABLE	FULL DESCRIPTION	UNIT	SOURCE	CATEGORY
HICPCORE	HICP - All-items excluding energy and food	index	Eurostat	
LTIE	Consensus Long-Term Inflation Expectations 6-10Y	percent	Consensus	
OG	Output gap (PC of EC, IMF and OECD estimates)	percentage points	EC, IMF, OECD	
IMPP	Relative import prices	index	Eurostat	
M1	M1 nominal stock	index	ECB	money
M12GDP	M1 to GDP ratio	percent	ECB	money
M3	M3 nominal stock	index	ECB	money
M32GDP	M3 to GDP ratio	percent	ECB	money
CRNFPS	Credit to the non-fin. priv. sector (NFPS) nom. stock	index	BIS	credit
CRNFPS2GDP	Credit to the NFPS to GDP ratio	percent	BIS	credit
LONFPS	Bank loans to the non-fin. priv. sector (NFPS) nom. stock	index	BIS	credit
LONFPS2GDP	Bank loans to the NFPS to GDP ratio	percent	BIS	credit
LONFC	Bank loans to non-fin. corporations (NFC) nom. stock	index	ECB	credit
LONFC2GDP	Bank loans to NFC to GDP ratio	percent	ECB	credit
LOHH	Bank loans to households (HH) nom. stock	index	ECB	credit
LOHH2GDP	Bank loans to HH to GDP ratio	percent	ECB	credit
CISS	Composite Indicator of Systemic Stress	index	ECB	other financial variable
STP	Dow Jones Euro Stoxx Price Index	index	ECB	other financial variable
HP	Residential property price index	index	ECB	other financial variable
CRSPR	Corporate bond spread (IG-3M Euribor)	percentage points	ECB	interest rate spread
YC	Slope of the Yield Curve: 10Y gov. bond yield - 3M Euribor	percentage points	ECB	interest rate spread
LRHHSPR	Mortgate lending rate minus 3M Euribor	percentage points	ECB	interest rate spread
LRNFCSPR	NFC lending rate minus 3M Euribor	percentage points	ECB	interest rate spread

Table B2: Main properties of the euro area HICP inflation

	Whole Sample 1990Q1-2019Q4	**Pre-Great Recession 1990Q1-2007Q2**	**Post-Great Recession 2007Q3-2019Q4**
		Summary Statistics	
Observations	116	69	47
Mean	1.7	2.1	1.1
Standard deviation	1.0	1.0	0.4
Maximum	0.3	0.6	0.3
Minimum	54.5	4.5	1.9
		Empirical Quantiles	
τ=**0.05**	0.66	0.87	0.50
τ=**0.10**	0.78	1.09	0.66
τ=**0.90**	3.35	4.03	1.72
τ=**0.95**	4.17	4.25	1.84
		Skewness and fat tails (Test Statistics)	
Skewness	4.988***	3.029***	0.145
Kurtosis	2.357***	0.247	1.004
Normality	0.695***	0.784***	0.644***

C Additional Results

Table C1: *Average relative scores by group of models for euro area core HICP inflation: 4-quarters ahead*

	QScore5	**Qscore10**	**QScore90**	**Qscore95**	**PL**
TVP-AR-SV	1.496	1.449	0.956	0.896	1.158
TVI-SV	1.654	1.461	0.922	0.914	1.314
TVP-QAR	0.988	1.243	0.901	0.839	1.218
TVI-QR	1.025	1.279	0.865	0.853	1.260
PC-SV	1.519	1.534	1.045	0.974	1.077
TVP-PC-SV	0.909	1.006	0.920	0.938	1.245
QPC	2.383	1.817	1.266	1.543	1.204
TVP-QPC	1.129	1.225	1.042	0.926	1.308
PC-SV-X	1.560	1.574	1.093	1.039	1.051
TVP-PC-SV-X	0.973	1.044	0.976	0.963	1.214
QPC-X	2.411	1.859	1.304	1.583	1.175
TVP-QPC-X	1.152	1.291	1.060	0.994	1.285
AR-SV-X	1.166	1.230	0.897	0.905	1.095
TVP-AR-SV-X	1.388	1.367	0.945	0.908	1.180
QAR-X	1.440	1.247	0.876	0.921	1.261
TVP-QAR-X	1.023	1.229	0.885	0.843	1.243

Notes: The quantile score refers to the average of the $QS_t^j(\tau)$ over all forecasting periods relative to the AR(2) benchmark, as an average across across all competing models with the specific model specification. The table reports the ratios of the quantile score of the models described in each row to that of the AR(2) model. For the four quantiles shown, any value below 1 signals an improvement of the forecast relative to the benchmark model and the lower the value of such ratio the larger the improvement. For the Predictive Likelihood (PL) (last column), any value above 1 signals an improvement of the forecast relative to the benchmark model and the higher the value of such ratio the larger the improvement.

Table C2: *Average relative scores by group of models for euro area core HICP inflation: 12-quarters ahead*

	QScore5	Qscore10	QScore90	Qscore95	PL
TVP-AR-SV	2.423	1.785	0.774	0.752	1.202
TVI-SV	2.412	1.670	0.786	0.691	1.210
TVP-QAR	1.144	1.368	0.783	0.852	1.245
TVI-QR	1.235	1.383	0.737	0.826	1.234
PC-SV	1.823	1.815	0.912	0.935	0.888
TVP-PC-SV	2.495	2.217	1.246	1.244	0.804
QPC	3.356	2.254	0.834	0.782	1.066
TVP-QPC	3.652	3.323	1.011	0.959	0.928
PC-SV-X	2.222	1.907	0.951	0.951	0.890
TVP-PC-SV-X	2.362	2.029	1.228	1.231	0.820
QPC-X	3.413	2.287	0.965	0.989	1.070
TVP-QPC-X	2.614	2.790	0.985	0.922	0.976
AR-SV-X	1.412	1.278	0.921	0.923	1.030
TVP-AR-SV-X	2.181	1.681	0.886	0.873	1.099
QAR-X	2.161	1.566	0.886	0.892	1.168
TVP-QAR-X	1.163	1.378	0.832	0.895	1.194

Notes: The quantile score refers to the average of the $QS_t^j(\tau)$ over all forecasting periods relative to the AR(2) benchmark, as an average across across all competing models with the specific model specification. The table reports the ratios of the quantile score of the models described in each row to that of the AR(2) model. For the four quantiles shown, any value below 1 signals an improvement of the forecast relative to the benchmark model and the lower the value of such ratio the larger the improvement. For the Predictive Likelihood (PL) (last column), any value above 1 signals an improvement of the forecast relative to the benchmark model and the higher the value of such ratio the larger the improvement.

Table C3: *Best financial indicators and models for the prediction of core HICP inflation tail risks: 4-quarters ahead*

Measure	Ranking	Indicator	Specification	Score
	1st	loans to private sector	TVP-PC-SV-X	0.891
QScore5	2nd	M1/GDP	TVP-QAR-X	0.900
	3rd	house prices	TVP-QAR-X	0.900
	1st	loans to firms	TVP-PC-SV-X	0.913
QScore10	2nd	loans to firms	TVP-QAR-X	0.949
	3rd	loans to private sector	TVP-PC-SV-X	0.969
	1st	loans to private sector	TVP-QAR-X	0.788
QScore90	2nd	private sector credit/GDP	QAR-X	0.799
	3rd	credit spread	QAR-X	0.800
	1st	credit to private sector	AR-SV-X	0.761
QScore95	2nd	loans to private sector	TVP-QAR-X	0.767
	3rd	credit to private sector	TVP-QAR-X	0.780
	1st	M1/GDP	TVP-QAR-X	1.474
PL	2nd	loans to private sector	TVP-QAR-X	1.429
	3rd	house prices	QAR-X	1.383

Notes: PC stands for Phillips curve, QAR for quantile autoregression with two lags, AR for mean autoregression with two lags, TVP for time-varying parameters and SV for stochastic volatility. Low inflation tail risks are captured by the 5th and 10th percentiles (Q5 and Q10, respectively) while high inflation tail risks are captured by the 90th and 95th percentiles (Q90 and Q95, respectively). The lower block reports the Predictive Likelihood (PL) which provides an evaluation of the forecast of the whole distribution of inflation. The last column reports the ratios of the predictive quantile score of the models described in each row to that of the AR(2) model. For the quantiles, a value below 1 signals an improvement of the forecast relative to the benchmark and the lower the value the larger is the improvement. For the PL, any value above 1 signals an improvement of the forecast relative to the benchmark model and the higher the value of such ratio the larger the improvement.

Table C4: *Best financial indicators and models for the prediction of core HICP inflation tail risks: 12-quarters ahead*

Measure	Ranking	Indicator	Specification	Score
	1st	private sector loans/GDP	TVP-QAR-X	0.951
QScore5	2nd	private sector credit/GDP	TVP-QAR-X	0.962
	3rd	yield curve	TVP-QAR-X	0.963
	1st	house prices	AR-SV-X	0.969
QScore10	2nd	M1	AR-SV-X	0.977
	3rd	private sector credit/GDP	AR-SV-X	1.054
	1st	loans to households	TVP-QPC-X	0.675
QScore90	2nd	lending rate to firms	TVP-QAR-X	0.706
	3rd	loans to households/GDP	QR-X	0.715
	1st	loans to households	TVP-QPC-X	0.635
QScore95	2nd	private sector credit/GDP	QAR-X	0.685
	3rd	loans to households/GDP	QAR-X	0.692
	1st	loans to households	TVP-QPC-X	1.552
PL	2nd	loans to households	QAR-X	1.336
	3rd	loans to households/GDP	QAR-X	1.295

Notes: PC stands for Phillips curve, QAR for quantile autoregression with two lags, AR for mean autoregression with two lags, TVP for time-varying parameters and SV for stochastic volatility. Low inflation tail risks are captured by the 5th and 10th percentiles (Q5 and Q10, respectively) while high inflation tail risks are captured by the 90th and 95th percentiles (Q90 and Q95, respectively). The lower block reports the Predictive Likelihood (PL) which provides an evaluation of the forecast of the whole distribution of inflation. The last column reports the ratios of the predictive quantile score of the models described in each row to that of the AR(2) model. For the quantiles, a value below 1 signals an improvement of the forecast relative to the benchmark and the lower the value the larger is the improvement. For the PL, any value above 1 signals an improvement of the forecast relative to the benchmark model and the higher the value of such ratio the larger the improvement.

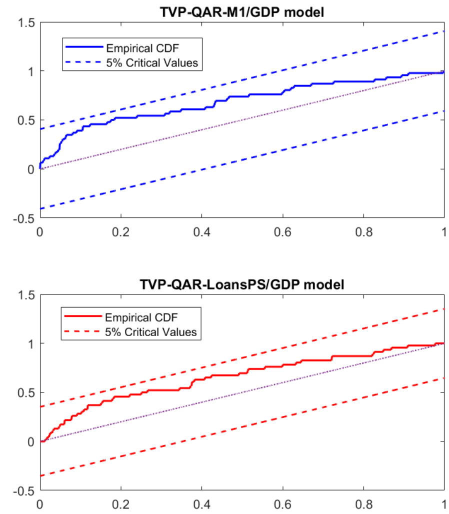

Figure C1: Calibration tests of Rossi and Sekhposyan (2019) for the best models

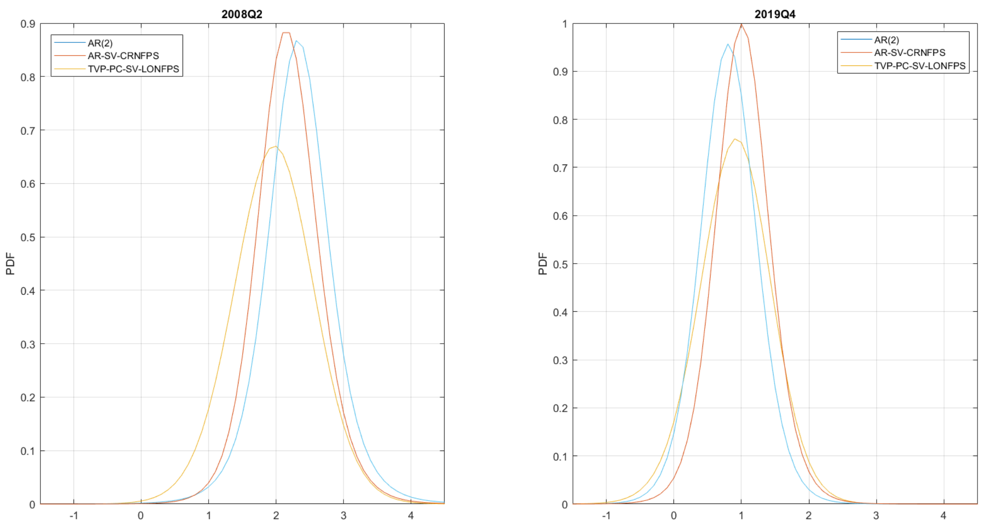

Figure C2: Conditional Distributions of euro area HICP Core inflation quarter-on-quarter growth rate for the best and benchmark models: 4-quarter ahead

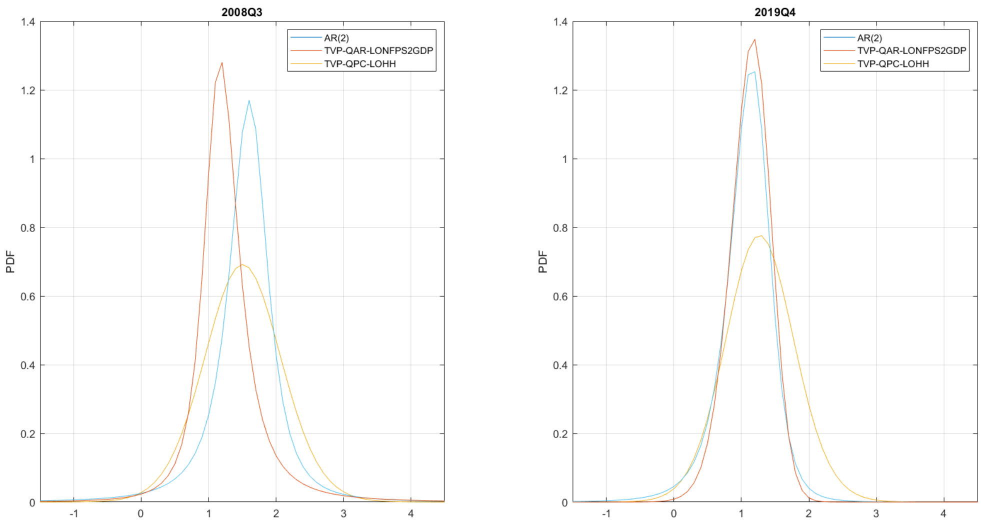

Figure C3: Conditional Distributions of euro area HICP Core inflation quarter-on-quarter growth rate for the best and benchmark models: 12-quarter ahead