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*)

11th ECB Conference on Forecasting Techniques
Macroeconomic forecasting in abnormal times

Tuesday 15 June 2021

Research questions

Setting: We are interested in assessing "inflation risks", i.e. risk of extreme realizations of inflation

Objective 1: Are credit and money indicators useful in predicting the distribution of HICP inflation? → particular attention to tail risks

Literature 1: **Quantile Regressions:**
Financial indicators are useful in the prediction of the distribution of future output growth (Adrian et al., 2019, AER), and inflation (Lopez-Salido and Loria, 2020, CEPR), especially for downside risks; real variables important for long-run forecasting of US inflation (Korobilis, 2016, IJF)
Traditional regressions:
Money and credit are helpful predictors of long-run inflation but only recently (Falagiarda and Sousa, 2017); Predictors of inflation are short-lived ("pockets of predictability"): Stock and Watson (1999, JME); Koop and Korobilis (2012, IER)

Research questions

Setting: We are interested in assessing "inflation risks", i.e. risk of extreme realizations of inflation

Objective 2: Are quantile regressions appropriate for assessing tail risks of macroeconomic variables?

Literature 2: Stochastic volatility (SV) very important for forecasting the distribution of inflation (Stock and Watson, 2007, JMCB)
Recent literature emphasizes that forecasting macro risks with SV/GARCH is as good as quantile regressions → Brownlees and Souza (forthcoming, JME); Carriero et al. (2020a,b, Clev. Fed WP)

Main features of our analysis

- Propose a quantile regression with time-varying parameters (TVPs)
- Previous literature: Kim (2007, Annals) uses regression splines; Cai and Xu (2008, JASA) fit local polynomials; Wu and Zhou (2017, JBES) also nonparametric.
- However, we desire a simple, fast and interpretable method, that can be used on a daily basis and is future-proof
- We specify a parametric Bayesian TVP-QR, using the well-known state-space form (e.g. Cooley and Prescott, 1976, ECMTA)
- Goncalvez et al. (2020, Bayes.Anal.) and Lim et al. (2020, Stat. Sin.) develop approximate inference methods, because MCMC is cumbersome
- **Major methodological contributions of our paper:**
  → Develop a very fast Gibbs sampler algorithm for TVP-QR models
  → Develop automatic shrinkage methods to deal with overparametrization

Main (preliminary) findings

Data: Quarterly Euro-Area data from 1990q1 to 2019q4 for ann. qly core HICP growth (LHS) and for 19 financial variables (RHS)

We find that:

- ☑ Various credit and money aggregates provide marginal value added for specific horizons and tail risks
- ☑ Quantile regressions with TVPs are clearly better (for density and tail forecasts) than TVP-SV regressions and linear quantile regressions
- ☑ Quantile regressions with both TVPs and SVs are impractical – asking too much from the data to estimate SV for each quantile level.

Mean vs Quantile regression model

We model dependence between y_t (scalar) and x_t $(1 \times p)$ using the following equation

$$y_t = f(y_t|x_t) + \varepsilon_t, \tag{1}$$

1. $f(y_t|x_t) = \mathbb{E}(y_t|x_t) = x_t\beta$ gives the linear **mean regression (MR)** model with solution

$$\widehat{\beta} = \min_{\beta} \mathbb{E} \sum_{t=1}^{T} (y_t - x_t\beta)^2, \tag{2}$$

2. $f(y_t|x_t) = \mathbb{Q}_\tau(y_t|x_t)(= x_t\beta(\tau))$ gives the linear **quantile regression (QR)** model, $\tau = \tau_1, \tau_2, ..., \tau_n$, with solution

$$\widehat{\beta}(\tau) = \min_{\beta(\tau)} \mathbb{E} \sum_{t=1}^{T} \rho_\tau(y_t - x_t\beta(\tau)), \tag{3}$$

   where $\rho_\tau(u) = (\tau - \mathbb{I}(u < 0))u$ is a loss function.

Classical vs Bayesian quantile regression model

In the linear quantile regression model

$$y_t = x_t\beta(\tau) + \varepsilon_t, \tag{4}$$

- $$\widehat{\beta}(\tau) = \min_{\beta(\tau)} \mathbb{E} \sum_{t=1}^{T} (\tau - \mathbb{I}(\varepsilon_t < 0))\varepsilon_t, \tag{5}$$

  can be solved by first writing it as a linear programming problem and subsequently using simplex methods.
- Yu and Moyeed (2001, Stas & Prob. Let.) show that the above problem is equivalent to maximizing an *asymmetric Laplace (AL)* likelihood for ε_t, that is the density

$$f(\varepsilon;\tau) = \prod_{t=1}^{T} \tau(1-\tau)\left[e^{(1-\tau)\varepsilon_t}\mathbb{I}(\varepsilon_t \leq 0) + e^{-\tau\varepsilon_t}\mathbb{I}(\varepsilon_t > 0)\right]. \tag{6}$$

Bayesian quantile regression model

Bayesian quantile regression is via the following parametric spec

$$y_t = x_t\beta(\tau) + \sigma(\tau)\varepsilon_t, \quad \varepsilon_t \sim AL(\tau), \tag{7}$$

where $\beta(\tau)$ and $\sigma(\tau)$ are parameters for each quantile level.
We can write the asymmetric Laplace as a scale mixture of Normals ("double exponential" distribution, cf Tibsirani, 1996, JRSSB, Section 5)

$$y_t = x_t\beta(\tau) + \theta(\tau)z_t + \sigma(\tau)\kappa(\tau)\sqrt{z_t(\tau)}u_t, \qquad u_t \sim N(0,1), \tag{8}$$

where $\theta(\tau) = \frac{1-2\tau}{\tau(1-\tau)}$ and $\kappa(\tau)^2 = \frac{2}{\tau(1-\tau)}$ and $z_t(\tau) \sim Exp(\sigma(\tau)^2)$. Likelihood of y is

$$\prod_{t=1}^{T} \frac{1}{\sqrt{2\pi z_t(\tau)\sigma(\tau)^2\kappa(\tau)^2}} \exp\left\{-\frac{(y_t - x_t\beta(\tau) - \theta(\tau)z_t(\tau))^2}{2z_t(\tau)\sigma(\tau)^2\kappa(\tau)^2}\right\} \exp\left\{-\frac{z_t(\tau)}{\sigma(\tau)^2}\right\}. \tag{9}$$

Integration of (9) w.r.t. z_t gives the asymmetric Laplace density.

Bayesian quantile regression model: Estimation (sampling)

♣ Likelihood is conditionally Normal → easy to derive conditional posteriors

Indeed, for priors of the form

$$\beta(\tau) \sim N(0, V(\tau)), \tag{10}$$

$$\sigma(\tau) \sim IG(\rho_1, \rho_2), \tag{11}$$

$$z_t(\tau) \sim exp(\sigma(\tau)^2), \tag{12}$$

Conditional posteriors are of the form

$$\beta(\tau)|\bullet \sim N\left(\left(x'Ux + V(\tau)^{-1}\right)^{-1} \times \left(x'U\left[y - \theta(\tau)z(\tau)\right]\right), \left(x'Ux + V(\tau)^{-1}\right)^{-1}\right), \tag{13}$$

$$\sigma(\tau)^2|\bullet \sim iGamma\left(\rho_1 + \frac{3T}{2}, \rho_2 + \sum_{t=1}^{T} \frac{\left(y_t - x_t\beta(\tau) - \theta(\tau)z_t(\tau)\right)^2}{2z_t(\tau)\kappa(\tau)^2} + \sum_{t=1}^{T} z_t(\tau)\right), \tag{14}$$

$$z_t(\tau)|\bullet \sim IG\left(\frac{\sqrt{\theta(\tau)^2 + 2\kappa(\tau)^2}}{|y_t - x_t\beta(\tau)|}, \frac{\theta(\tau)^2 + 2\kappa(\tau)^2}{\sigma(\tau)^2\kappa(\tau)^2}\right), \forall t \in \{1, ..., T\} \tag{15}$$

U is a $T \times T$ diagonal covariance matrix with t-th element $(z_t(\tau)\sigma(\tau)^2\kappa(\tau)^2)^{-1}$

Bayesian quantile regression model: Estimation (sampling)

- We can devise a Gibbs sampler to sample sequentially from conditionally posteriors for each τ
- For $\tau = 0.05, 0.10, ..., 0.90, 0.95$ we need to iterate through these posteriors 19 times per each iteration of the Gibbs sampler
- $\sigma(\tau)$,$z_t(\tau)$ can be updated in one step for all t, τ; $\beta(\tau)$ faster to update sequentially for each τ (e.g. parallelize this step)
- Khare and Hobert (2012, JMA) show that this Gibbs sampler converges at a geometric rate
- This is also true in the “large p, small T” case

Bayesian quantile regression model: Shrinkage

- A major benefit of parametric Bayesian inference is the vast availability of model selection and shrinkage priors
- Shrinkage is **imperative** in QR models: only few observations available for each quantile
- We consider the Horseshoe of Carvalho et al. (2010, Biometrika)
- Bayes estimates are consistent a-posteriori, with risk equivalent to the (Bayes) oracle (Armagan et al., 2013, Biometrika; Ghosh et al., 2016, Bayes.Anal.)
- Results are for Normal regression, while our model is conditionally Normal

$$\beta(\tau)_i|\sigma(\tau)^2, \lambda(\tau)^2, \psi_i(\tau)^2 \sim N\left(0, \sigma(\tau)^2\lambda(\tau)^2\psi_i(\tau)^2\right), \tag{16}$$

$$\lambda(\tau) \sim Cauchy^+\left(0, 1\right), \tag{17}$$

$$\psi_i(\tau) \sim Cauchy^+\left(0, 1\right), \tag{18}$$

Time-varying parameter Bayesian quantile regression model

Linear Bayesian quantile regression is

$$y_t = x_t\beta(\tau) + \varepsilon_t, \qquad \varepsilon_t \sim AL(\sigma(\tau)^2), \tag{19}$$

In line with the macro TVP regression literature, the extension is straightforward

$$\begin{aligned} y_t &= x_t\beta_t(\tau) + \varepsilon_t, \qquad \varepsilon_t \sim AL(\sigma(\tau)^2), && (20) \\ \beta_t(\tau) &= \beta_{t-1}(\tau) + v_t, \qquad v_t \sim N(0, V(\tau)), && (21) \end{aligned}$$

- Allowing for stochastic volatility (SV) $\sigma_t(\tau)^2$ is a bad idea for short (quarterly) EA data $\rightarrow$ Gerlach et al. (2011, JBES): estimate a Bayesian quantile SV model for daily stock indices
- TVP-QR is a conditionally Gaussian & linear state-space model

Time-varying parameter Bayesian quantile regression model

We follow ideas in Korobilis (forthcoming, JBES) and Goulet Coulombe (2020, Arxiv) and rewrite the TVP regression by modeling the increments $\Delta\beta_t(\tau)$

$$
\begin{aligned}
y_t &= x_t\beta_t(\tau) + \varepsilon_t && (22)\\
&= x_t\Delta\beta_t(\tau) + x_t\beta_{t-1}(\tau) + \varepsilon_t && (23)\\
&= x_t\Delta\beta_t(\tau) + x_t\Delta\beta_{t-1}(\tau) + x_t\beta_{t-2}(\tau) + \varepsilon_t && (24)\\
&\ldots && (25)\\
&= x_t\Delta\beta_t(\tau) + x_t\Delta\beta_{t-1}(\tau) + \ldots + x_t\Delta\beta_2(\tau) + x_t\beta_1(\tau) + \varepsilon_t && (26)
\end{aligned}
$$

The above holds for the t-th equation of the TVP-QR.

Time-varying parameter Bayesian quantile regression model

The previous equation in matrix form becomes

$$\underset{(T \times 1)}{\begin{bmatrix} y_1 \\ y_2 \\ \dots \\ y_{T-1} \\ y_T \end{bmatrix}} = \underset{(T \times Tp)}{\begin{bmatrix} x_1 & 0 & \dots & 0 & 0 \\ x_2 & x_2 & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots \\ x_{T-1} & x_{T-1} & \dots & x_{T-1} & 0 \\ x_T & x_T & \dots & x_T & x_T \end{bmatrix}} \underset{(Tp \times 1)}{\begin{bmatrix} \beta_1(\tau) \\ \Delta\beta_2(\tau) \\ \dots \\ \Delta\beta_{T-1}(\tau) \\ \Delta\beta_T(\tau) \end{bmatrix}} + \underset{(T \times 1)}{\begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \dots \\ \varepsilon_{T-1} \\ \varepsilon_T \end{bmatrix}} \quad (27)$$

♣ In this form the TVP regression is a high-dimensional model with more predictors than observations

♣ Therefore, the state equation can be thought of as the regularizing prior $\beta^{\Delta}(\tau) \sim N(0, V(\tau)) \rightarrow$ sensitivity to choice $V(\tau)$ (Amir-Ahmadi, 2020, JBES)

♣ Convert this to the Horseshoe specification and allow the data to select $V(\tau)$

Time-varying parameter Bayesian quantile regression model

- Using the previous, linear form, we can sample the TVP regression without relying on Kalman filter
- Great transformation because we avoid sequential sampling that cannot be parallelized
- However, in the linear form we have Tp coefficients $\rightarrow$ How to sample from a $N_{Tp}(A^{-1}a, A^{-1})$ posterior when inversion and Cholesky operators on A require $\mathcal{O}((Tp)^3)$ operations?
- For monthly US data Tp could be more than $100,000$!
- Bhattacharya et al. (2016, Biometrika) provide a simple but clever algorithm that utilizes Woodbury matrix identity not only to invert A, but also sample from the Normal distribution
- Worst case algorithmic complexity of $\mathcal{O}(T^2p)$

Fast Sampling from Normal posterior

Assume the following model (ignoring the variance parameter)

$$\begin{aligned} y &\sim N(X\beta, \sigma), & \texttt{likelihood} \quad (28) \\ \beta &\sim N(0, D), & \texttt{prior} \quad (29) \\ \beta|y &\sim N(V \times X'y, V) & \texttt{posterior} \quad (30) \end{aligned}$$

where $V = (X'X + D^{-1})^{-1}$.

Algorithm *Bhattacharya et al. (2016) fast sampling of* β

---

i. Sample $u \sim N\left(0, D\right)$ and $\delta \sim N\left(0, I_T\right)$
ii. Set $v = Xu$
iii. Solve $\left(XDX' + I_T\right) w = y - v$ to obtain w
iv. Set $\beta = u + DX'w$

Proof is trivial and relies on application of Woodbury identity

Time-varying parameter Bayesian quantile regression model: summary

To summarize the TVP-QR model is now

$$\begin{aligned}
y &= \mathcal{X}\beta^{\Delta}(\tau) + \theta(\tau)z(\tau) + \widetilde{S}u, && (31)\\
\beta^{\Delta}(\tau) &\sim N(0, V(\tau)), && (32)\\
V_{i,i}(\tau) &= \sigma(\tau)^2\lambda(\tau)^2\psi_i(\tau)^2, \quad i = 1, ..., Tp, &&\\
\lambda(\tau) &\sim Cauchy^{+}(0,1), && (33)\\
\psi_i(\tau) &\sim Cauchy^{+}(0,1), && (34)\\
\sigma(\tau) &\sim IG(\rho_1, \rho_2), && (35)
\end{aligned}$$

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

♣ We can use the Gibbs sampler presented earlier for the linear QR model

♣ As long as $\beta^{\Delta}(\tau)$ is sampled efficiently using the Bhattacharya et al. (2016, Biometrika) trick, all other posteriors are scalar and easy to sample from

♣ We sample $\beta^{\Delta}(\tau)$ but trivial to recover β_t (cumsum)

Quantile noncrossing

- The quantile regression model fits each conditional quantile independently
- In practice neighboring quantiles will correlated
- A typical problem is *quantile crossing*, i.e. estimated quantile curves $\widehat{\mathbb{Q}}_{\tau}(y_t|x_t)$ are not monotonic functions of τ
- A typical solution is to post-process the quantiles $\rightarrow$ such procedures might cause some bias
- Here we use the algorithm of Rodrigez and Fan (2017, JCGS) for Bayesian QR
- As long as MCMC draws are independent (use thinning) their noncrossing procedure ensures monotonicity and posterior consistency

Quantile noncrossing

Exactly because adjacent quantiles are correlated, use asymmetric Laplace quantile function to obtain the auxiliary quantile model

$$\mathbb{Q}_{\tau,\tau^\star}(y_t|x_t) = \begin{cases} x_t\beta(\tau^\star) + \frac{\sigma(\tau^\star)}{1-\tau^\star}\log\left(\frac{\tau}{\tau^\star}\right), & if \quad 0 \leq \tau \leq \tau^\star, \\ \\ x_t\beta(\tau^\star) - \frac{\sigma(\tau^\star)}{\tau^\star}\log\left(\frac{1-\tau}{1-\tau^\star}\right), & if \quad \tau^\star \leq \tau \leq 1, \end{cases} \tag{36}$$

where $\mathbb{Q}_{\tau,\tau^\star}(y_t|x_t)$ is the induced quantile, and $\tau, \tau^\star \in \{0.05, 0.10, ..., 0.90, 0.95\}$.

♣ The above gives a 19×19 $\mathbb{Q}_{\tau,\tau^\star}(y_t|x_t)$ of quantiles for each MCMC draw

♣ Use a GP regression to obtain a weighted average (over draws, and over 19 auxiliary quantiles)

Background

Methodology

Simulation study

Model evaluation using real data

Synthetic data experiment

We generate data from the following time-varying regression model

$$y_t = x_t \beta_t + \varepsilon_t, \tag{37}$$

$$\beta_t = \mu + 0.99(\beta_{t-1} - \mu) + T^{-\frac{1}{2}} u_t \tag{38}$$

where $x \sim N(0, I_2)$ is a vector of two artificial predictors, $\mu \sim U(-2, 2)$ is the long-run mean of $\beta_t = [\beta_{1,t}, \beta_{2,t}]'$, and $u_t \sim N(0, I_2)$.
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.

$\rightarrow$ All that remains is to specify various distributions of ε_t

Flexible error distributions in DGP

We follow the Monte Carlo design in Yu (2017, JASA), and consider eight different choices. These are the following:

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 : $1/2N(-1, (2/3)^2) + 1/2N(1, (2/3)^2)$
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.

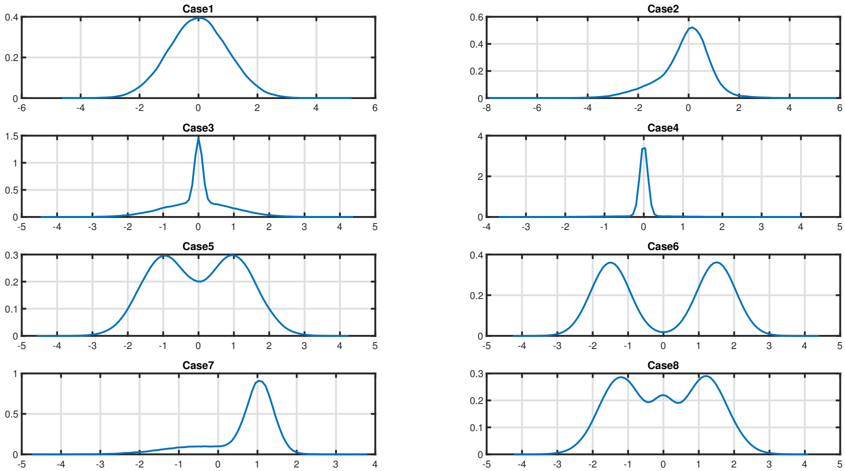

Figure: *Error distributions generated in the Monte Carlo study: 1) Normal, 2) Skewed, 3) Kurtotic, 4) Outlier, 5) Bimodal, 6) Bimodal, separate modes, 7) Skewed bimodal, and 8) Trimodal.*

Monte Carlo evaluation

- We generate 500 datasets of length $T = 200$ from each of the 8 DGPs
- We fit two models, "mean" TVP regression; quantile TVP regression
- In a Bayesian setting one is a special case of the second (Normal vs Normal-Exponential errors)
- The precision of estimation β_t affects how well we forecast (either mean or quantiles) so our loss function is

$$MSD_j = \frac{1}{500} \sum_{r=1}^{500} \left\{ \frac{1}{200} \sum_{t=1}^{200} \left[\frac{1}{2} \sum_{i=1}^{2} \left(\widehat{\beta}_{i,t}^{(j,r)} - \beta_{i,t} \right)^2 \right] \right\}, \tag{39}$$

where $\widehat{\beta}_{i,t}^{(j,r)}$ is j model's r-th Monte Carlo estimate of coefficient $\beta_{i,t}$, $i = 1, 2$, $t = 1, ..., 200$, $j = \{mean \;\; TVP\}, \{quantile \;\; TVP\}$, $r = 1, ..., 500$.

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

	DGP1	DGP2	DGP3	DGP4	DGP5	DGP6	DGP7	DGP8
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	0.03	0.08
$\tau = 0.50$	0.05	0.04	0.04	0.01	0.06	0.11	0.03	0.06
$\tau = 0.75$	0.05	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	0.04	0.08
$\tau = 0.95$	0.05	0.05	0.05	0.01	0.09	0.13	0.05	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.

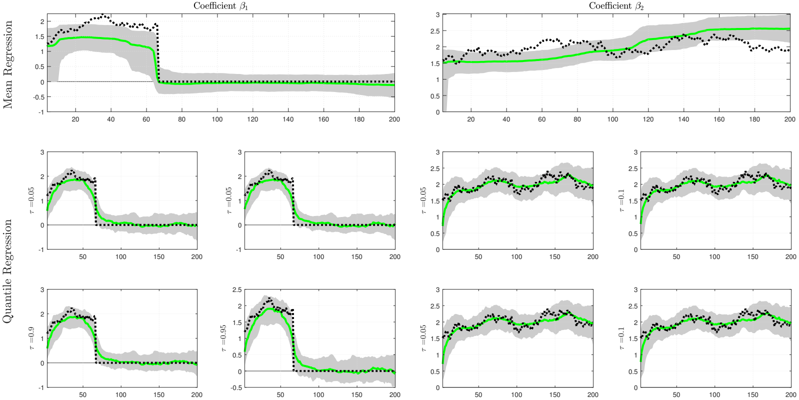

Figure: Posterior estimates of time-varying parameters (TVPs) estimated using mean (upper panels) and quantile (middle and bottom panels) regressions. Black lines are the true TVPs, which are the same for both the mean and quantile regressions. The green lines are the averages (over 100 Monte Carlo iterations) of the estimated posterior means, and the shaded areas and the 68 percent probability bands.

Forecasting EA inflation (HICP)

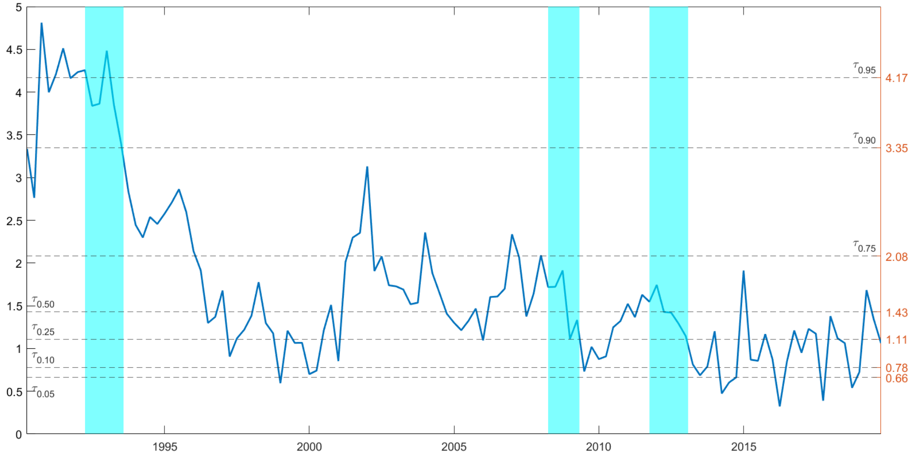

Figure: Core inflation data and quantiles

Data

VARIABLE	FULL DESCRIPTION	UNIT	SOURCE
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
M12GDP	M1 to GDP ratio	percent	ECB
M3	M3 nominal stock	index	ECB
M32GDP	M3 to GDP ratio	percent	ECB
CRNFPS	Credit to the non-fin. priv. sector (NFPS) nom. stock	index	BIS
CRNFPS2GDP	Credit to the NFPS to GDP ratio	percent	BIS
LONFPS	Bank loans to the non-fin. priv. sector (NFPS) nom. stock	index	BIS
LONFPS2GDP	Bank loans to the NFPS to GDP ratio	percent	BIS
LONFC	Bank loans to non-fin. corporations (NFC) nom. stock	index	ECB
LONFC2GDP	Bank loans to NFC to GDP ratio	percent	ECB
LOHH	Bank loans to households (HH) nom. stock	index	ECB
LOHH2GDP	Bank loans to HH to GDP ratio	percent	ECB
CISS	Composite Indicator of Systemic Stress	index	ECB
STP	Dow Jones Euro Stoxx Price Index	index	ECB
HP	Residential property price index	index	ECB
CRSPR	Corporate bond spread (IG-3M Euribor)	percentage points	ECB
YC	Slope of the Yield Curve: 10Y gov. bond yield - 3M Euribor	percentage points	ECB
LRHHSPR	Mortgate lending rate minus 3M Euribor	percentage points	ECB
LRNFCSPR	NFC lending rate minus 3M Euribor	percentage points	ECB

Models

We use two classes of models, a dynamic regression

$$\pi_{t+h} = c_t(\tau) + \phi_{1t}(\tau)\pi_t + \phi_{2t}(\tau)\pi_{t-1} + \beta_t(\tau)\boldsymbol{x}_t + \varepsilon_{t+h}, \quad \varepsilon_{t+h} \sim ALD(\sigma_t(\tau)), \quad (40)$$

and a semi-structural (Phillips curve) model of the form

$$\pi_{t+h} = (1 - \lambda_t(\tau))\pi_t^* + \lambda_t(\tau)\pi_t^{LTE} + \theta_t(\tau)\left(y_t - y_t^*\right) + \gamma_t(\tau)\pi_t^I + \beta_t(\tau)\boldsymbol{x}_t + \varepsilon_{t+h}, \quad (41)$$

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 and domestic inflation).

Models

$$\pi_{t+h} = c_t(\tau) + \phi_{1t}(\tau)\pi_t + \phi_{2t}(\tau)\pi_{t-1} + \beta_t(\tau)\boldsymbol{x}_t + \varepsilon_{t+h}, \tag{42}$$

$$\varepsilon_{t+h} \sim ALD(\sigma_t(\tau)), \tag{43}$$

- If $\sigma_t(\tau) = \sigma(\tau)$ we obtain our proposed class of TVP-QR models
- If $\tau = 0.5$ and $z_t(\tau) = \frac{1}{\kappa(\tau)^2} = \frac{\tau(1-\tau)}{2}$ then $\varepsilon_{t+h} \sim Normal(0, \sigma_t)$
- All TVP can become constant, e.g. if we fix the state variance to be zero
- If x_t is the empty set, we can obtain the class of AR models

♣ Similar arguments can be made about the Phillips curve specification

List of regression-based models

AR(2) is the benchmark. Remaining models are:

1. AR(2) model with TVPs and stochastic volatility (TVP-AR-SV)
2. Time-varying intercept only model with stochastic volatility[1] (TVI-SV)
3. Quantile AR(2) with time-varying parameters (TVP-QAR)
4. Quantile regression model with time-varying intercept (TVI-QR)
5. Mean regressions with constant parameters, exogenous predictors, and stochastic volatility (AR-SV-X)
6. Mean regressions with time-varying parameters, exogenous predictors, and stochastic volatility (TVP-AR-SV-X)
7. Quantile AR(2) with constant parameters augmented with exogenous predictors (QAR-X), and
8. Quantile AR(2) with time-varying parameters augmented with exogenous predictors (TVP-QAR-X)

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

List of PC-based models

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)

Forecast Metrics

- We use two forecast metrics, quantile score (Manzan, 2015, JBES) and predictive likelihood (e.g. numerous papers by John Geweke)
- The (average) quantile score (QS) is the following loss function:

$$QS_h^j(\tau) = \frac{1}{R_h}\sum_{t=1}^{R_h} \pi_{t+h} - \hat{Q}_\tau(\pi_{t+h}|\boldsymbol{x}_t)][\mathbb{I}\{\pi_{t+h} \leq \hat{Q}_\tau(\pi_{t+h}|\boldsymbol{x}_t)\}], \tag{44}$$

  where R_h is the length of the forecast evaluation sample.
  We evaluate this metric at $\tau = 0.05, 0.95$, and we give the names QScore5, QScore95.
- All QScore results are relative to an AR(2) benchmark
- We present results for $h = 4, 12$ quarters ahead forecasts

Top models, $h = 4$

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

Top models, $h = 12$

Measure	Ranking	Indicator	Specification	Score
		12-QUARTERS AHEAD		
QScore5	1st	private sector loans/GDP	TVP-QAR-X	0.951
	2nd	private sector credit/GDP	TVP-QAR-X	0.962
	3rd	yield curve	TVP-QAR-X	0.963
QScore95	1st	loans to households	TVP-QPC-X	0.635
	2nd	private sector credit/GDP	QAR-X	0.685
	3rd	loans to households/GDP	QAR-X	0.692
PL	1st	loans to households	TVP-QPC-X	1.552
	2nd	loans to households	QAR-X	1.336
	3rd	loans to households/GDP	QAR-X	1.295

Thank you!