Nowcasting Euro Area GDP Growth Using Bayesian Quantile Regression

James Mitchell (Federal Reserve Bank of Cleveland)
Aubrey Poon (University of Strathclyde)
Gian Luigi Mazzi (Eurostat)

ESCoE EM2021 Conference, May 2021

Nowcasting GDP via quantile regression

- Eurostat's Flash estimates of quarterly GDP growth for the euro area are currently published 30 days after the end of the quarter. For many years, from 2003-2016, they were published at 45 days
- Ahead of these estimates many higher-frequency indicators become available
- This paper uses an application to explore the utility of quantile regression methods in producing density nowcasts of quarterly GDP growth from higher frequency indicators

Quantile regression (QR)

- QR allows for nonlinearities between the indicators and GDP growth
  - The relative importance of different indicators can vary by quantile of GDP growth
- We emphasise the construction of full predictive densities from these QRs
- We propose a Bayesian QR nowcasting strategy that accommodates
  1. the mixed-frequency, and
  2. "ragged-edge" nature of
  3. the increasingly big datasets that characterise recent nowcasting applications

U-MIDAS QR

- We relate the conditional τ-th quantile of quarterly GDP growth in quarter t, y_t, to x_t^m, a vector of monthly indicators, where $m = 1, 2, 3$ denotes the month in quarter t (plus lags of GDP)

$$Q_{y_t}(\tau|\Omega_t) = \beta_{0,\tau} + \beta_{1,\tau}x_t^1 + \beta_{2,\tau}x_t^2 + \beta_{3,\tau}x_t^3, \tau \in (0,1)$$

- We follow Mazzi et al. (2014 OBES) and consider a large set of 124 monthly indicators, including business confidence data, IP data and financial market data for both the euro aggregates and its largest economies
- We use real-time data for all variables (involved us collecting euro area vintage data)
- We adopt a Bayesian estimator and consider various shrinkage priors to avoid parameter proliferation

Bayesian QR (BQR)

- Bayesian inference proceeds by forming the likelihood using asymmetric Laplace densities
- We use the mixture representation of Kozumi and Kobayashi (2011 JSCS)
- We implement different shrinkage priors on $\mathbf{V}_\beta$, where $\beta_\tau \sim N(0, \mathbf{V}_\beta)$
- All can lead to the inclusion of different indicators in different parts of the GDP growth distribution
- Estimation via standard MCMC (Gibbs sampling) methods

Shrinkage priors

We consider three forms of shrinkage prior

1. Lasso, adaptive Lasso and elastic net. Parameter-specific (or local) shrinkage priors
2. Horseshoe and Dirichlet-Laplace. Shrinks small coefficients (or uninformative indicators) to zero, with fat-tails to avoid over-shrinkage of large coefficients associated with important indicators
3. SSVS. A normal mixture prior, the Stochastic Search Variable Selection prior is a discrete mixture of a peaked prior around zero and a vague proper prior

Density nowcasts I

- Having estimated the BQR in-sample ($t = 1, ..., T$) for a given quantile τ, quantile nowcasts can be computed given x_{T+1}:

$$\widehat{Q}_{y_{T+1}}(\tau|\Omega^j_{T+1})^r = \widehat{\beta}^r_0(\tau) + \widehat{\beta}^r_1(\tau)x^1_{T+1} + \widehat{\beta}^r_2(\tau)x^2_{T+1} + \widehat{\beta}^r_3(\tau)x^3_{T+1}$$

  where $\widehat{\beta}^r_k$ ($k = 0, ..., 3$) denotes the r-th MCMC draw from the posterior parameter distribution and Ω^j_{T+1} denotes the j-th available information set. We consider nowcasts at $j = 1$ (t-15 days) and $j = 2$ (t+15 days)
- Recall that the quarter T+1 values of the indicator variables are published ahead of the quarter T+1 values for y_t and are exploited when nowcasting. The density nowcasts can be evaluated when y_{T+1} is subsequently published
- We consider two methods of constructing a predictive density from the BQR quantile nowcasts

Density nowcasts II

1. Collect $r = 1, ..., R$ MCMC draws of the quantile nowcast $\widehat{Q}_{y_{T+1}}(\tau|\Omega^j_{T+1})^r$ across $\tau \in [0.05, 0.10, ..., 0.90, 0.95]$ and then construct the full posterior density nowcast from this stacked vector - use a Gaussian kernel to smooth
2. Follow Adrian et al. (AER 2019) and fit a skew-t density to 5 (we also try 19) conditional mean quantile forecasts

Out-of-sample results I

- We compare the accuracy of the density nowcasts of euro area GDP growth from BQR (with the 6 priors) against

1. The density forecast combination strategy of Mazzi et al. (2014 OBES)
    - Takes linear combinations of the density nowcasts generated from 124 linear Gaussian regression models relating quarterly GDP growth separately to each monthly or quarterly indicator
2. Factor (principal component) augmented classical quantile and linear regressions
    - Likely effective in *dense* data environments
3. An AR(1) Gaussian density

Evaluation

- Evaluate the point and density estimates against Eurostat's first official estimate using
  1. RMSE
  2. Average logarithmic score
  3. Average cumulative ranked probability score
  4. Two variants of the quantile weighted probability score (Gneiting and Ranjan, 2007 JBES) that emphasise left tail events and both tails
- The evaluation sample starts in 2003q2, as this is when Eurostat first published its Flash estimate for GDP growth and ends in 2019q4
- More on the pandemic to come...

Nowcast accuracy

- In summary, we find
  1. BQR with global-local priors delivers the most accurate nowcasts at both horizons according to all five evaluation metrics
  2. The choice of prior matters
  3. BQR outperforms the density forecast combinations and the factor-augmented linear and quantile regressions - suggestive of a *sparse* rather than *dense* dataset and of gains to shrinkage out-of-sample
  4. Using a skewed-*t* does not work as well as using the MCMC draws
  5. Accuracy is better at t+15 days than at t-15 days

Temporal evolution of moments from BQR (HS prior)

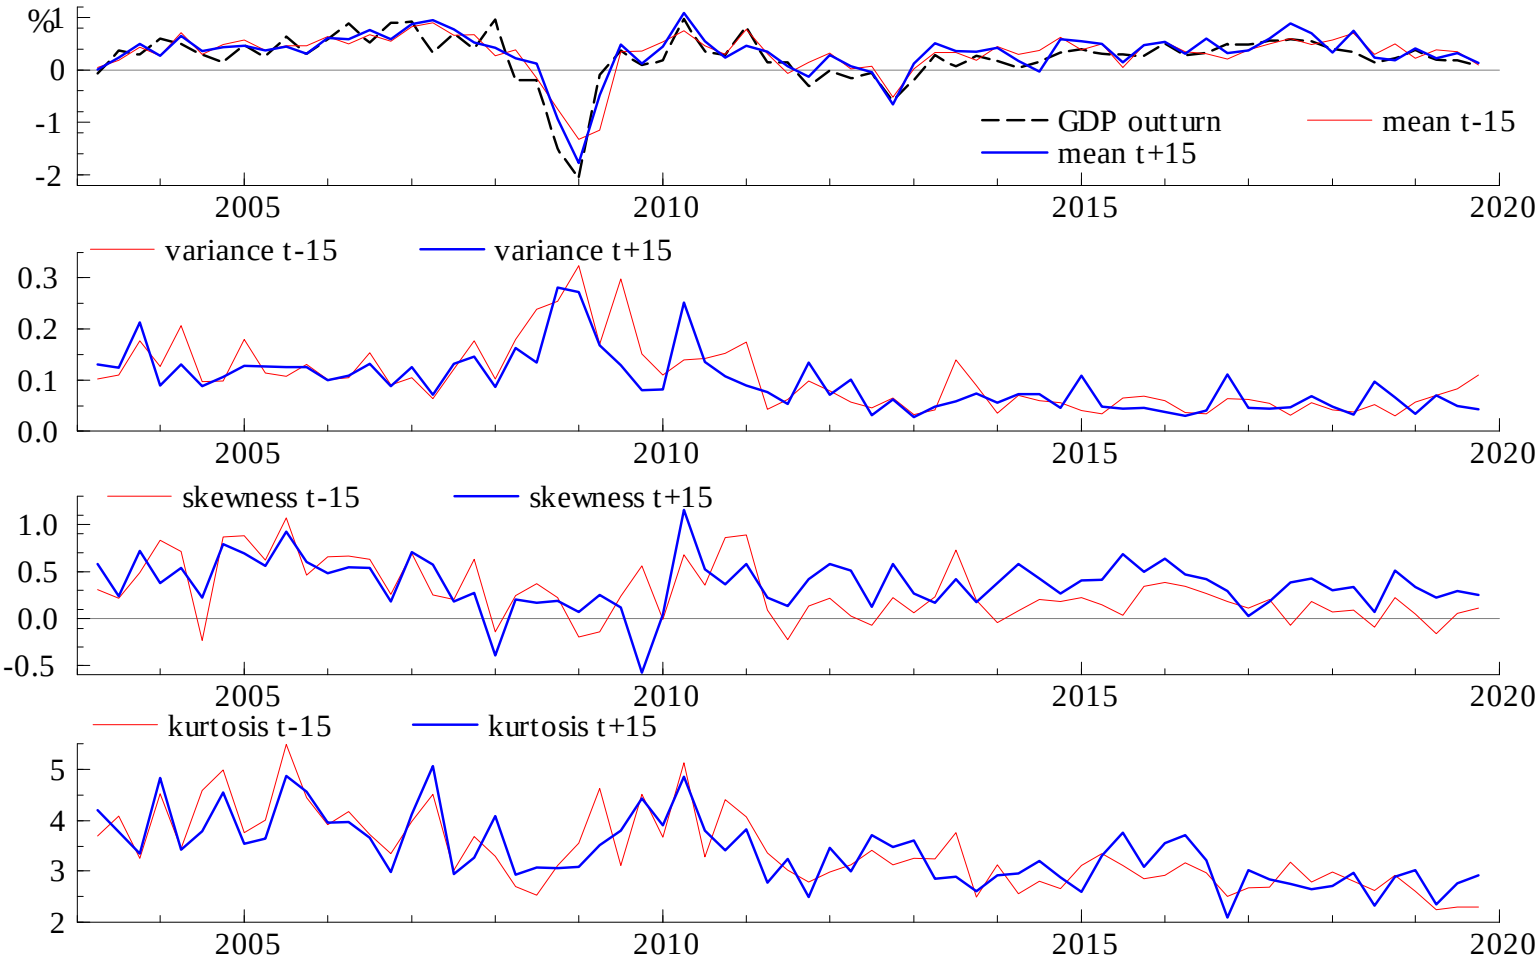

Growth-at-Risk: Prob of negative growth

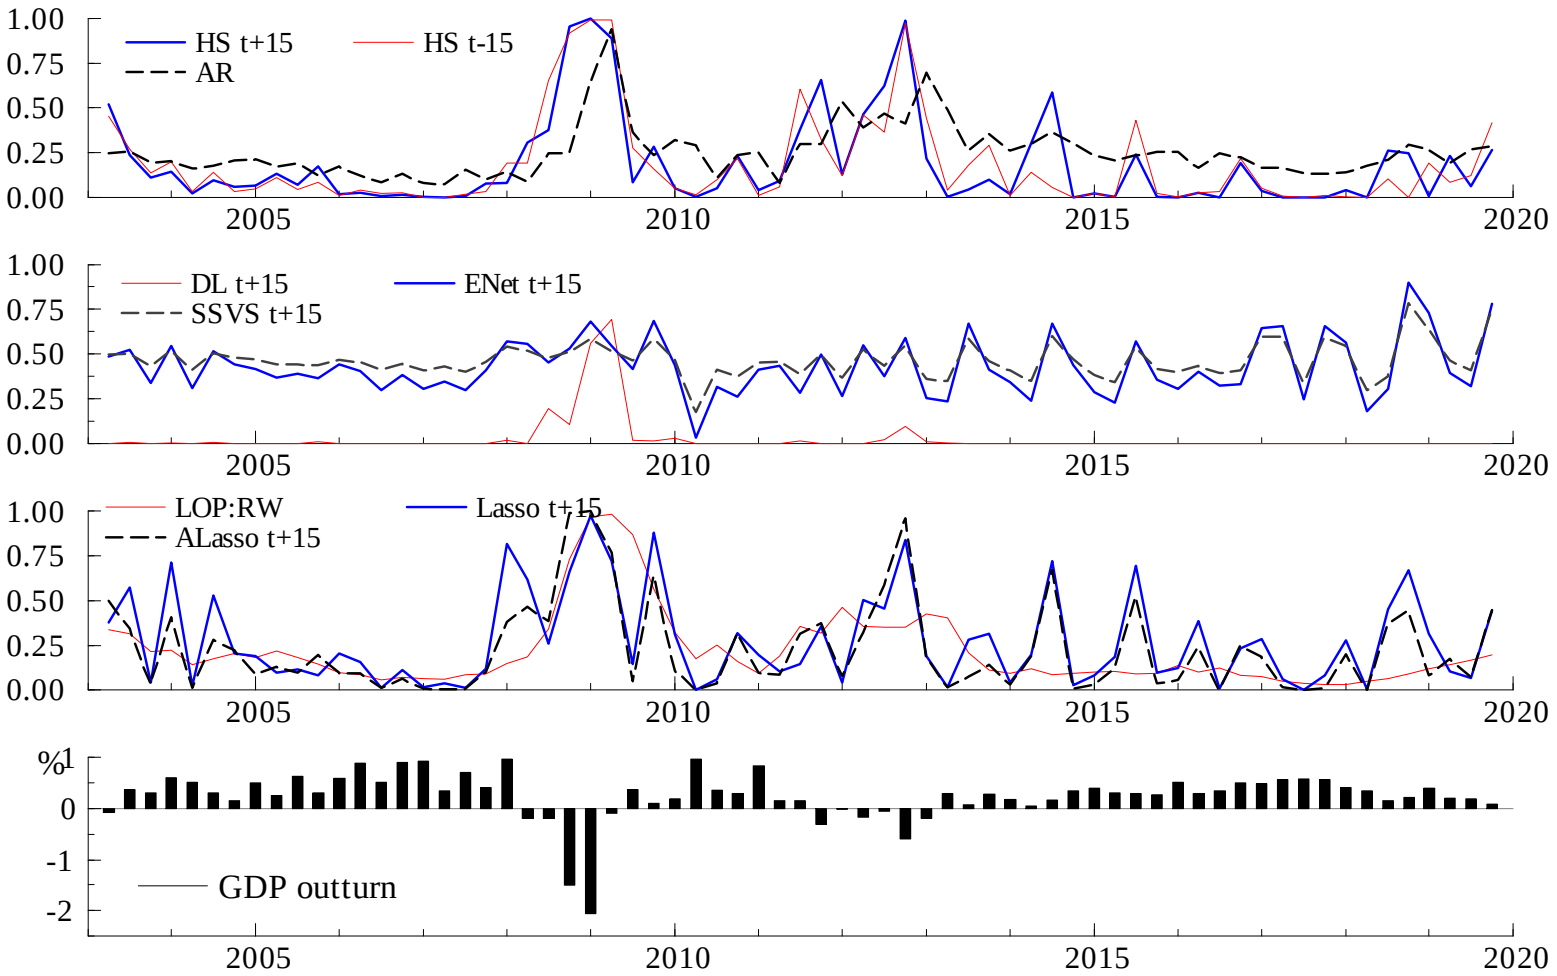

COVID-19: Prob of negative growth

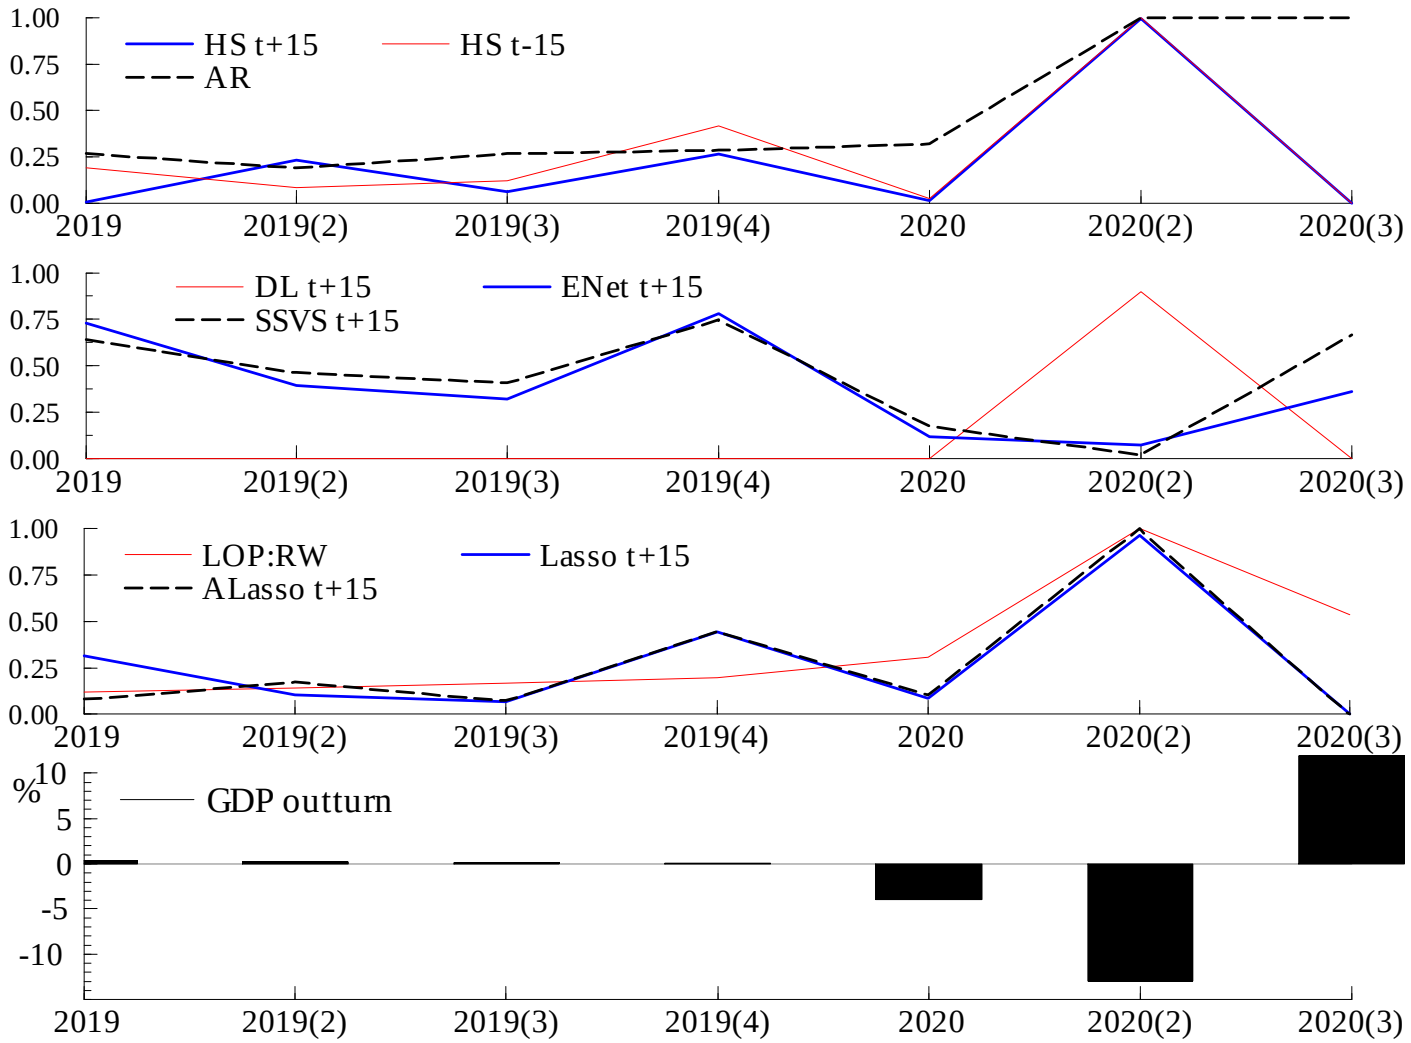

Conclusion

- Bayesian QRs with global-local priors produce accurate density nowcasts of euro area GDP growth, including over the GFC and COVID-19 recessions
- Future work should consider the use of weekly, even daily indicators