

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Sokol, Andrej

Working Paper Fan charts 2.0: Flexible forecast distributions with expert judgement

ECB Working Paper, No. 2624

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Sokol, Andrej (2021) : Fan charts 2.0: Flexible forecast distributions with expert judgement, ECB Working Paper, No. 2624, ISBN 978-92-899-4911-8, European Central Bank (ECB), Frankfurt a. M., https://doi.org/10.2866/966212

This Version is available at: https://hdl.handle.net/10419/249897

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU



Working Paper Series

Andrej Sokol Fan charts 2.0: flexible forecast distributions with expert judgement



Disclaimer: This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Abstract

I propose a new model, conditional quantile regression (CQR), that generates density forecasts consistent with a specific view of the future evolution of some variables. This addresses a shortcoming of existing quantile regression-based models, for example the at-risk framework popularised by Adrian et al. (2019), when used in settings, such as most forecasting processes within central banks and similar institutions, that require forecasts to be conditional on a set of technical assumptions. Through an application to house price inflation in the euro area, I show that CQR provides a viable alternative to existing approaches to conditional density forecasting, notably Bayesian VARs, with considerable advantages in terms of flexibility and additional insights that do not come at the cost of forecasting performance.

JEL Classification: C22, C53, E37, R31.

Keywords: at-risk, quantile regression, house prices, conditional forecasting, density forecast evaluation.

Non-Technical Summary

Risk assessment, that is, a quantitative focus on the entire distribution of possible future outcomes, is an integral part of the forecasting processes and communications strategies of central banks and many other institutions. However, until very recently, its analytical underpinnings have not kept abreast with advances in the academic literature and the technological possibility frontier to the same extent as the production of point forecasts, where various uses of, for example, DSGE, VAR or factor models have become common over the last couple of decades. Instead, risk assessments are typically still carried out by either weighing the likelihood of alternative point forecasts (scenarios), or by agreeing upon a central projection first and only quantifying risks around it at a subsequent stage.

Repeated 'tail events,' such as the Great Financial Crisis or the more recent COVID-19 episode, have put these approaches under increasing strain, and provided a window of opportunity to consider some alternatives. Among these, the 'at-risk' framework has emerged as a viable and popular contender, in academia, but even more so within policy institutions. The at-risk approach is based on quantile regression, a framework which seeks to model the relationship between different parts of the *distribution* of a variable of interest and a set of explanatory variables. Because of its ability to capture changes over time in all moments of that distribution, in a forecast setting it can not only provide a model-based central projection, but also estimates of its uncertainty (variance), balance of risks (skewness) and 'fat-tailedness' (kurtosis), which can all be of independent interest for both analytical and communications purposes.

But while analyses based on quantile regression have become increasingly popular, their implementation to date suffered from a serious disadvantage in many policy settings, namely the inability to incorporate conditioning information (technical assumptions) over the forecast horizon. This limits a model's role to that of an unconditional benchmark, whereas alternative forecasts consistent with a particular view of the economic outlook are often what is required by decision makers.

The main contribution of the paper is to propose a new framework, *conditional* quantile regression (CQR), that allows to incorporate technical assumptions into an otherwise standard quantile regression setup, and to show that it provides a viable alternative to existing approaches to conditional density forecasting, notably Bayesian VARs. I present an application to euro area house price inflation, showcasing some model-based outputs and the novel insights for risk assessment purposes that can be obtained by conditioning forecasts on alternative scenarios for the technical assumptions. I also show that these additional insights do not come at the cost of forecasting performance.

1 Introduction

Risk assessment, that is, a quantitative focus on the entire distribution of possible future outcomes, is an integral part of the forecasting processes and communication strategies of central banks and many other institutions. It has a long tradition, dating back at least to the Bank of England's pioneering use of fan charts (Britton et al., 1998) in its communication of macroeconomic forecasts through the *Inflation Report*.¹ However, until very recently, its analytical underpinnings have not kept abreast with advances in the academic literature and computing power to the same extent as the production of point forecasts, where various uses of, for example, DSGE, VAR or factor models have become common over the last couple of decades.²

Instead, risk assessments are typically still carried out using one of two approaches. Either different 'risks' are viewed as separate point forecasts (scenarios) distinct from the central projection³, explicitly or implicitly weighted by the subjective probabilities attributed to their occurrence and communicated as such. Or a central projection is agreed upon first, and the risks around it are quantified in a second stage and then summarised by means of predictive distributions superimposed to the central forecast. The former approach underpins, for example, the ECB's communication of its staff projections during the COVID-19 crisis (see for example Box 4 in European Central Bank, 2021b). The latter is, by and large, the process underlying the Bank of England's fan charts, where the variance and skewness of predictive distributions centered around a modal forecast are adjusted judgementally by the Monetary Policy Committee to convey their view of the balance of risks.

Repeated 'tail events,' such as the Great Financial Crisis or the COVID-19 episode, have put these approaches under increasing strain,⁴ and provided a window of opportunity to consider some alternatives. Among these, the 'at-risk' framework, popularised by Adrian et al. (2019), has emerged as a viable and successful contender, in academia, but even more so within policy institutions (Gaglianone and Lima, 2012; Korobilis, 2017; Gelos et al., 2019; Eguren-Martin and Sokol, 2020; Ferrara et al., 2020; Figueres and Jarociński, 2020; Hasenzagl and Ricco, 2020; López-Salido and Loria, 2020; Tagliabracci, 2020; Valckx et al., 2020; Adams et al., 2021; Eguren-Martin et al., 2021, among several others).

The at-risk approach is based on quantile regression (Koenker and Bassett, 1978), which allows to model the *entire distribution* of a variable of interest. Because of its ability to capture changes over time in all moments of that distribution, not just the first (mean), in a forecast setting it can not only provide a model-based central projection, but also estimates of its uncertainty (variance), balance of risks (skewness) and 'fat-tailedness' (kurtosis), which can all be of independent interest for both analytical and communications purposes.

¹From November 2019, the Inflation Report became the Monetary Policy Report.

²See for example, Burgess et al. (2013); Coenen et al. (2018); Bok et al. (2018); Cimadomo et al. (2020) for some examples within central banks.

³Usually understood as a modal forecast.

⁴Witness the marked changes to the uncertainty and skewness parameters of the GDP fan chart in the August 2020 Monetary Policy Report (Bank of England, 2020).

But while analyses based on quantile regression are increasingly popular,⁵ their implementation to date suffered from a serious disadvantage: the inability to incorporate conditioning information. Forecasts produced in most institutions typically require that the projections of a variable be consistent with a set of *technical assumptions*⁶, that is, paths for some of its determinants that are assumed to be known with certainty over the forecast horizon.⁷ Technical assumptions are a form of expert judgement imposed upon forecasts for a variety of reasons, ranging from the mixed-frequency nature of some forecasts (see for example Bok et al., 2018; Cimadomo et al., 2020), to the long-standing practice of providing projections conditional on specific paths for fiscal and/or monetary policy instruments, the exchange rate or energy prices (Burgess et al., 2013; Bańbura et al., 2015; Angelini et al., 2019; Boneva et al., 2019; Domit et al., 2019; Chalmovianský et al., 2020; Antolín-Díaz et al., 2021, among others).⁸

There are well-established methods to incorporate technical assumptions in DSGE and VAR models (see Del Negro and Schorfheide, 2013; Antolín-Díaz et al., 2021), which have found many applications also in conditional density forecasting.⁹ However, they suffer from at least one important limitation, namely that their predictive densities are constrained by the distributional assumptions made for the error terms, which are typically Gaussian or otherwise symmetric, and therefore also imply symmetric forecast densities. This limits their usefulness for modelling higher moments, notably the balance of upside and downside risks.¹⁰ Instead, the greater flexibility of quantile regression can be fully leveraged in this respect, but hasn't been to date.

The main contribution of the paper is therefore to propose a new *conditional* quantile regression (henceforth CQR) model that allows to incorporate technical assumptions into an otherwise standard Bayesian quantile regression setup (Yu and Moyeed, 2001; Kozumi and Kobayashi, 2011; Khare and Hobert, 2012). The CQR model is a straightforward extension to quantile regression of the conditional direct multi-step (DMS) framework studied by McCracken and McGillicuddy (2019), who compare the performance of (conditional) point forecasts generated by iteration and direct projection. The approach works as follows: at the estimation stage, *leads* of the technical assumptions up to the forecast horizon of interest are included among the regressors, so that at the forecasting stage, conditional projections can be made once paths for the technical assumptions are provided. As has become standard since Adrian et al. (2019), the resulting conditional quantiles are then fitted with a flexible parametric distribution, the *skew-t* (Azzalini and Capitanio, 2003). This yields predictive densities that are not only conditional on the available data at the time of the forecast (such as those obtained with current imple-

⁵Witness, for example, their recurrent use in the IMF's *Global Financial Stability Report*.

⁶I follow the terminology used within the ECB's staff macroeconomic projections process. 'Conditioning paths' and other expressions are also commonly used in other policy institutions.

⁷I abstract here from the important issue of the potential uncertainty surrounding technical assumptions (see Andersson et al., 2010).

⁸For the interesting problem of how such conditional forecast should be properly evaluated, see Faust and Wright (2008) or Clark and McCracken (2014).

⁹See for example Angelini et al. (2019); Boneva et al. (2019); Domit et al. (2019); Chalmovianský et al. (2020) for some contributions that specifically focus on the density forecasting performance of standard models.

¹⁰For example, the methods discussed in Antolín-Díaz et al. (2021) are only applicable to VAR models with Gaussian or Student-t errors, which are both symmetric.



Figure 1: Alternative scenarios for the predictive distributions of house price inflation.

(a) Available data only.

(b) Stress.

Note: The fan charts show realised data for house price inflation up to 2020Q2, and then density forecasts for 2020Q3 to 2022Q3 under different technical assumptions for the unemployment and long term real interest rates, with the colour coding matching that of Figure 5: no technical assumptions over the projection (black) and a stress scenario involving a gradual worsening in the technical assumptions (red). The fan chart bands cover 95% of the support around the median: the darkest shade corresponds to the median, while lighter shades represent percentiles increasingly removed from it. The forecasts are based on a CQR model estimated with data up to 2020Q2, and fed projections for the technical assumptions over the forecast horizon.

mentations of quantile regression), but also on the specific view of the future embodied in the technical assumptions fed to the model.¹¹

To showcase the usefulness of the CQR framework, I present an application to euro area house price inflation. House prices are an important variable from both a macroeconomic and a financial stability perspective: they are a driver of households' spending and borrowing decisions, with considerable ramifications for aggregate activity and inflation,¹² as well as a key asset price determining the health of investors' and financial intermediaries' balance sheets. As such, they have a long history of modelling efforts, including, more recently, within the quantile regression framework (Gerdesmeier et al., 2015; Valckx et al., 2020). Building on the former, I specify a parsimonious CQR model where house price inflation forecasts depend both on the histories of house price inflation, household income growth, unemployment and real interest rates up to the time when the forecast is made, and on technical assumptions for the unemployment and real interest rates over the projection period.

¹¹In related work, Chavleishvili and Manganelli (2019) propose a quantile VAR that can in principle also be used to generate density forecasts conditional on a specific path for one or more technical assumptions. However, to compute the conditional forecast distribution, the quantile VAR requires an iterative procedure that can become computationally very costly for larger systems.

¹²See for example the recent focus on owner-occupied housing (OOH) within the ECB's strategy review (European Central Bank, 2021a).

Figure 1 summarises the outputs of the CQR model by means of fan charts,¹³ and illustrates the impact that technical assumptions can have on the conditional predictive densities of house price inflation. The models underlying both panels are estimated on the same data, but while the left fan chart is only conditioned on realised data up to the time when both forecasts are made – as a standard quantile regression model would do – the right fan chart shows the effects of conditioning the projection on a 'stress' scenario for the technical assumptions.¹⁴ As further discussed in Section 5, the stress scenario entails large changes to all moments of the predictive distributions: the red fan chart displays a weaker and downward-sloping median path, is generally wider, and less skewed. All three effects can be of interest to policy makers, but typically only the first and second are captured by existing approaches.

I also show that such usefulness for the purpose of scenario analysis doesn't come at the cost of out-of-sample forecasting performance. I benchmark the CQR's density forecasts to those of a Bayesian VAR (BVAR) estimated with state-of-the-art techniques (Giannone et al., 2015) and conditioned on the same technical assumptions, and find that overall, the two models perform similarly out of sample. Coupled with the additional insights that can be gleaned from the CQR approach, these results should encourage its wider adoption as a risk assessment tool.

Paper structure The rest of the paper is organized as follows. In Section 2 I introduce the conditional quantile regression (CQR) model, then in Section 3 I present an application to house price inflation. In Section 4 I benchmark its forecasting performance to that of a BVAR estimated and conditioned on the same data, while in Section 5 I provide an example of scenario analysis. Section 6 concludes, with additional material reported in Appendices A and B.

2 The Conditional Quantile Regression Model

McCracken and McGillicuddy (2019) are the first to study the (point) forecast performance of conditional direct multi-step (DMS) models. Their conditional DMS setup is a simple linear regression where a standard direct forecast model, which relates a dependent variable to a set of regressors known h periods in advance, is augmented with a set of *technical assumptions*, that is, variables for which a forecast needs to be provided when the projection is made. They focus on the performance of models estimated with OLS, that is, under standard assumptions about the error term. However, their framework is in fact more general, and can be easily adapted to quantile regression, potentially yielding DMS forecasts of the full conditional predictive distribution of a variable of interest – what I call a *conditional* quantile regression (CQR) model. As shown in Sections 4 and 5, this is a viable and appealing alternative to the conditional density forecasting approaches currently in use.

¹³That is, sequences of predictive densities over the forecast horizon. Note that unlike the original Bank of England fan charts (Britton et al., 1998), where the darkest band covers the mode, the fan charts presented here are centered around the median of the underlying distributions. The interpretation of the bands is therefore slightly different in the two cases, but still conveys the overall shape of the distributions.

¹⁴The 'stress' scenario assumes that the unemployment rate is 2pp higher, and the real interest rate 1pp higher, by the end of the projection, see Section 5.

Slightly adapting their notation, for each forecast horizon h one can estimate a regression of the following form

$$y_t = \alpha + \sum_{j=0}^{p-1} \beta'_j \boldsymbol{x}^b_{t-h-j} + \sum_{1 \le i \le h} \gamma'_i \boldsymbol{x}^f_{t-h+i} + \varepsilon_t$$
(1)

where the β coefficients multiply variables known at the time of the projection (t-h), while the γ coefficients multiply leads (relative to t-h) of the technical assumptions, for which a forecast needs then to be provided to project the model. As in McCracken and McGillicuddy (2019), I use realised values of the technical assumptions to estimate the model, but the alternative, namely to use vintages of actual technical assumptions made in the past, is interesting in its own right, especially for institutions with a long history of conditional forecasting, such as most central banks. The key difference between the two alternatives is that by using realised values, the γ coefficients capture the conditional relationship between y_t and $\mathbf{x}_{t-h+1}^{f'} \dots \mathbf{x}_t^{f'}$ in the data, while in the other case they capture the conditional relationship between y_t and the history of technical assumptions made in each period and for each forecast horizon up to h.¹⁵

Equation 1 nests the simpler variant

$$y_t = \alpha + \sum_{j=0}^{p-1} \beta'_j \boldsymbol{x}^b_{t-h-j} + \boldsymbol{\gamma}'_h \boldsymbol{x}^f_t + \varepsilon_t$$
(2)

which is more parsimonious and performs slightly better in my application.¹⁶ For the latter specification, the h-step-ahead conditional forecast is then given by

$$\hat{y}_{t+h|t} = \hat{\alpha} + \sum_{j=0}^{p-1} \hat{\beta}'_{j} \boldsymbol{x}^{b}_{t-j} + \hat{\gamma}'_{h} \boldsymbol{x}^{f}_{t+h|t}$$
(3)

where $\boldsymbol{x}_{t+h|t}^{f}$ is a forecast, made at time t, for the technical assumptions.

Equations 1 or 2 can be readily adapted to quantile regression, and thus to estimate quantile- (as well as horizon-) specific CQR models, by assuming that the error terms ε_t have an Asymmetric Laplace distribution, as explained in Section 2.1. For each horizon, individual quantile estimates can then be fitted with a continuous distribution.¹⁷ As in Adrian et al. (2019), I use Azzalini and Capitanio (2003)'s *skew-t* distribution, which allows a very flexible modelling of higher

$$y_t = \alpha + \sum_{j=0}^{p-1} \beta'_j \boldsymbol{x}_{t-h-j}^b + \sum_{1 \le i \le H} \boldsymbol{\gamma}'_i \boldsymbol{x}_{t-h+i}^f + \varepsilon_t$$

¹⁵The latter model can therefore also capture systematic features of the process that leads to the formulation of technical assumptions within an institution, which could potentially affect its forecasting performance. Instead, the baseline setup is entirely data-driven at the estimation stage, and a portion of its forecast errors are readily attributable to the discrepancy between technical assumptions and outturns for the conditioning variables.

¹⁶Another theoretically interesting variant is

where H denotes the maximum forecast horizon, i.e. when the direct forecast for each horizon includes the conditioning information up to the longest forecast horizon.

¹⁷Following a quantile sorting step, as proposed by Chernozhukov et al. (2010).

moments, and thus of one- and two-sided risks.¹⁸ Then, predictive distributions for different horizons can be combined into a fan chart, or studied individually. A point worth noting is that especially with Equation 1, the number of regressors can increase quickly with the forecast horizon and number of variables for which technical assumptions are provided. This makes Bayesian estimation particularly attractive, as shrinkage allows to mitigate the adverse effect of a large number of regressors on parameter uncertainty.

2.1 Bayesian Quantile Regression

Let Y be a response variable and X a set of explanatory variables. Then the linear quantile regression model of Y is given by

$$y_t = \boldsymbol{x}_t' \boldsymbol{\beta} + \sigma \boldsymbol{\epsilon}_t \quad (t = 1, \dots, T) \tag{4}$$

where $\boldsymbol{\beta} = \boldsymbol{\beta}(q)$, the error term $\sigma \epsilon_t = \sigma(q) \epsilon_t(q)$ has a density function $f_q(\cdot)$ such that the q-th quantile equals zero, i.e. $\int_{-\infty}^0 f_q(\sigma \epsilon_t) d\epsilon_t = q$, and $\sigma = \sigma(q) > 0$ is a scale parameter. With a slight abuse of notation, Equation 1 can be cast in the above format by letting $\boldsymbol{\beta} = [\alpha \ \boldsymbol{\beta}_0 \dots \boldsymbol{\beta}_{p-1} \ \boldsymbol{\gamma}_1 \dots \boldsymbol{\gamma}_h]', \boldsymbol{x}_t = [1 \ \boldsymbol{x}_{t-h}^{b'} \dots \boldsymbol{x}_{t-h-p+1}^{b'} \ \boldsymbol{x}_{t-h+1}^{f'} \dots \boldsymbol{x}_t^{f'}]'$ and $\sigma \epsilon_t = \varepsilon_t$.

The classical (frequentist) solution to the problem of finding β obtains from the minimisation of the criterion function

$$\sum_{t=1}^{T} \rho_q \left(y_t - \boldsymbol{x}_t' \boldsymbol{\beta} \right) \tag{5}$$

where $\rho_q(u) = u [q - I (u < 0)]$ is the so-called check function (see for example Koenker and Bassett, 1978). Yu and Moyeed (2001) pointed out that solving (5) is equivalent to finding the maximum likelihood estimator of β in (4) under the assumption that the errors have an Asymmetric Laplace distribution. This in turn opens the way to re-casting (4) as a Normal variance-mean mixture model (Barndorff-Nielsen et al., 1982), as shown by Kozumi and Kobayashi (2011) and further refined by Khare and Hobert (2012). Indeed, assuming that the ϵ_t are iid^{19} Asymmetric Laplace with probability density function

$$f_q(\epsilon) = q(1-q)\exp\left\{\rho_q(\epsilon)\right\}$$
(6)

the scaled error terms in (4) have an equivalent latent variable representation (Kotz et al., 2001),

¹⁸I also obtain very similar results, not reported to conserve space, by fitting a non-parametric distribution to the estimated quantiles, as in Eguren-Martin and Sokol (2020).

¹⁹The *iid* assumption is not entirely innocuous in a time series context, and especially with a DMS setup such as Equation 2, because for h > 1, the error terms are auto-correlated by construction. Various solutions to this problem have been provided in a frequentist setup: Fitzenberger (1998) showed that for quantile regression, the block bootstrap is robust to autocorrelation of unknown form; Manzan and Zerom (2013) proposed an adhoc procedure for 'purging' the error terms of their MA structure prior to running the quantile regressions; and Adrian et al. (2019) rely on artificial data generated by a VAR to compute confidence bands around their quantile regression point estimates. In a Bayesian setting, Sokol (2020) proposes a solution based on a modification of the framework outlined in this Section.

Table 1: Data sources and transformations.

Variable	Symbol	Formula/Transformation	Source
Nominal house price index (SA)	HP_t	Log	ECB and Eurostat
Harmonised index of consumer prices (SA)	P_t	Log	ECB and Eurostat
Nominal long-term interest rate p.a.	i_t		ECB
Unemployment rate	U_t		ECB and Eurostat
Real household disposable income (SA)	Y_t	Log	ECB and Eurostat
Annual HICP inflation	π_t^p	$100 * (P_t - P_{t-4})$	
Annual real house price inflation	π_t^{hp}	$100 * (HP_t - HP_{t-4}) - \pi_t^p$	
Real long-term interest rate change	r_t	$(i_t - i_{t-1}) - (\pi_t^p - \pi_{t-1}^p)$	
Unemployment rate change	u_t	$U_t - U_{t-1}$	
Real household disposable income annual growth	y_t	$100 * (Y_t - Y_{t-4})$	

given by

$$\sigma \epsilon_t \stackrel{a}{=} \theta z_t + \tau \sqrt{\sigma z_t} u_t \tag{7}$$

where $Z = Z(q) \sim Exp(\sigma)$, $U = U(q) \sim \mathcal{N}(0,1)$, Z and U are mutually independent, $\theta = \theta(q) = \frac{1-2p}{q(1-q)}$ and $\tau^2 = \tau^2(q) = \frac{2}{q(1-q)}$. Then Y conditional on Z is Normal and given by

$$y_t | z_t \sim \mathcal{N} \left(\boldsymbol{x}_t' \boldsymbol{\beta} + \theta z_t, \sigma \tau^2 z_t \right),$$

the likelihood of $\{Y_t, Z_t\}_{t=1}^T$ is

$$f^*(\boldsymbol{y}, \boldsymbol{z}; \boldsymbol{\beta}, \boldsymbol{\delta}, \sigma, \boldsymbol{x}) = \prod_{t=1}^T \frac{1}{\sqrt{2\pi\sigma\tau^2 z_t}} \exp\left\{-\frac{1}{2\sigma\tau^2 z_t} \left(y_t - \boldsymbol{x}_t' \boldsymbol{\beta} - \theta z_t\right)^2\right\} \sigma^{-1} \exp\left\{-\frac{z_t}{\sigma}\right\}, \quad (8)$$

and given an appropriate (conjugate) choice of priors, the posterior distributions of all parameters and of the latent variables Z_t are known and amenable to Gibbs sampling, as shown in Appendix A.

3 A Simple CQR Model of House Price Inflation

To illustrate the usefulness of the framework in a policy setting, I present an application to house price inflation. As discussed in the introduction, house prices are closely monitored for both macroeconomic and financial stability purposes, and have therefore also a long history of modelling efforts. As they share some features with both asset prices and durable goods, Gerdesmeier et al. (2015), in their review of the existing literature, distinguish between asset price-based and 'structural' (i.e. relying on macroeconomic fundamentals affecting the demand for and supply of housing) approaches. They also provide an early example of an application of quantile regression to the study of house prices,²⁰ albeit with important methodological differences to this paper, most notably the absence of technical assumptions in the formulation of out-of-sample forecasts.

I follow these authors in specifying a very parsimonious 'structural' model of house price inflation that includes disposable income, real long-term interest rates and the unemployment rate as its

 $^{^{20}}$ Another useful precedent is Valckx et al. (2020) and earlier work carried out at the IMF for the *Global Financial Stability Report*.



Figure 2: House price inflation fan chart for 2018Q3-2020Q3.

Note: The fan chart shows realised data for house price inflation up to 2018Q2, and then density forecasts for 2018Q3 to 2020Q3, with actual data outturns for the forecast period overlaid in red. The fan chart bands cover 95% of the support around the median: the darkest shade of blue corresponds to the median, while lighter shades represent percentiles increasingly removed from it. The forecasts are based on a CQR model estimated with data up to 2018Q2, and fed realised data for the technical assumptions over the projection horizon.

main determinants. However, an important difference is the inclusion of technical assumptions about the future evolution of both real long-term interest rates and the unemployment rate. The exact CQR model specification is, following Equation 2,

$$\pi_t^{hp} = \alpha + \sum_{j=0}^{1} \beta_j' \left[\pi_{t-h-j}^{hp}, \ y_{t-h-j}, \ r_{t-h-j}, \ u_{t-h-j} \right]' + \gamma_h' \left[r_t, \ u_t \right]' + \varepsilon_t$$
(9)

where all coefficients and the error term are both horizon- (h) and quantile- (q) specific, but with the dependencies suppressed to simplify the notation. Equation 9 is essentially a quantile autoregression model (Koenker and Xiao, 2006), with additional exogenous regressors that include both variables known at the time of the forecast and technical assumptions over the projection horizon made at the time of the forecast. The estimation sample is 1981Q1:2020Q2, and information about data sources and transformations is provided in Table 1.

Figures 2 and 3 show a fan chart for horizons up to 9 quarters ahead and the corresponding predictive densities for h = 5 and h = 9, respectively. To obtain the plots, I estimated the model

Figure 3: House price inflation forecasts as of 2018Q2 and outturns.



(a) Conditional forecast for 2019Q3.

(b) Conditional forecast for 2020Q3.

Note: The two panels show density forecasts for 2019Q3 and 2020Q3, respectively, with actual data outturns overlaid in red and the unconditional distribution of house price inflation over the estimation sample shown in dashed black. The forecasts are based on a CQR model estimated with data up to 2018Q2, and fed realised data for the technical assumptions over the projection horizon.

with data available as of 2018Q2, and then used realised interest rate and unemployment data up to 2020Q3 as technical assumptions²¹ to produce conditional forecasts for each horizon. House price inflation outturns over the same forecast horizon are shown in red. The fan chart shows that until the end of 2019, the model, once fed with the actual developments of real interest and unemployment rates, would have assigned a high probability to house price inflation turning out as it did in the data, as evidenced by the red line falling within the darker region of the fan chart. The conditional forecast for 2019Q2 shown in 3a (which is nothing but a 'slice' of the fan chart at h = 5) provides a different illustration of the same fact, with the vertical red line very close to the mode of the predictive distribution. The panel also illustrates another interesting feature, namely how much the conditional predictive distribution, which incorporates knowledge of past developments up to the time of the forecast as well as technical assumptions over the projection, differs from the unconditional distribution of house price inflation over the entire sample, shown by the dashed black line.

From 2020Q1 onwards, with the onset of the COVID-19 pandemic, the real interest and unemployment rate data used as technical assumptions showed a marked deterioration (see the solid black lines in Figure 5). Based on the historical regularities captured by the model estimates, this would have implied both a substantially weaker central projection and, by 2020Q3, also some downside risks to the outlook (Figures 2 and 3b). Instead, house price inflation didn't show any signs of weakness and even picked up to historically high levels. Figure 3b shows that by 2020Q3, it was a complete outlier relative to the conditional predictive distribution for

 $^{^{21}}$ These would of course not been available in real time, so the forecasts shown in Figure 2 are best thought of as a (counterfactual) 'upper bound' to the model's out-of-sample forecasting ability.

the same quarter (blue curve), but also high by historical standards, as shown by its location within the unconditional distribution (dashed black curve). This corroborates the narrative put forward by Battistini et al. (2021), who note that compared to their behaviour during both the Global Financial Crisis and the Sovereign Debt Crisis, house prices remained surprisingly resilient in the face of both macroeconomic and financial stress, and attribute these differences to both the different nature of the COVID-19 shock and of the ensuing policy responses.

Beyond this narrative validation over the recent past, in the next Section I also provide a formal evaluation of the model's predictive ability, benchmarking it against a standard Bayesian VAR (BVAR). The results show that not only does the CQR model provide reasonable and well-calibrated out-of-sample forecasts, but also that its forecasts are comparable, and for longer horizons somewhat better, than those of a BVAR sharing the same conditioning information, which should allay concerns about the use of the model in a live policy setting. In the following Section I then show what that might look like by means of a simple scenario analysis example.

4 Forecast Evaluation

To assess the CQR model's forecasting performance, I carry out a (pseudo-) out-of-sample evaluation. Because the data used do not, as a rule, get revised (except for sporadic corrections and methodological changes), the qualifier is not particularly relevant for the estimation of the model in real time. However, at the time of each forecast, technical assumptions would have needed to be provided, instead of which I use ex-post realised data (as in the example discussed in Section 3). The evaluation sample spans the 1999Q1 to 2020Q3 period, for a total of 79 forecast vintages that can be compared against outturns up to 9 quarters ahead.²²

As a benchmark, I estimate a Bayesian VAR on the same data, and feed it the exact same set of technical assumptions as the baseline model. The BVAR model is given by

$$z_t = A_0 + A_1 z_{t-1} + \dots + A_p z_{t-p} + \eta_t, \tag{10}$$

where, following the notation of Table 1, $z_t = [(HP_t - P_t), Y_t, (i_t - \pi_t^p), U_t]', A_0, \ldots, A_p$ are matrices of coefficients, the lag order is p = 5 and η_t is a Normally-distributed white noise process with full covariance matrix Σ_{η} . All variables are purposefully specified in levels so as not to put the BVAR at an undue disadvantage because of the data transformations required by the quantile regression setup. The BVAR is estimated following Giannone et al. (2015), that is, adopting Minnesota (Litterman, 1979), sum-of-coefficients (Doan et al., 1984) and dummyinitial-observations (Sims, 1993) priors and choosing their optimal degree of tightness based on the model's marginal likelihood. Technical assumptions are implemented with the Kalman

²²The date of the vintage indicates the first quarter of the forecast. So for example, in the 2020Q3 vintage, house price inflation data up to 2020Q2 are available, and the 2020Q3 number needs to be forecasted. Note that due to the comparatively long delay in the release of house prices data, 2020Q3 data for all explanatory variables included in the model are available well in advance, and therefore technical assumptions only need to be provided for 2020Q4 and subsequent quarters. A similar 'ragged edge' situation arises in each vintage.



Figure 4: 1-year-ahead conditional forecasts for 2000Q1-2019Q2.

Note: The chart shows successive vintages of 1-year-ahead forecasts for house price inflation, with corresponding data outturns overlaid in red. The bands cover 68% and 95% of the support around the median. The forecasts are based on a CQR model estimated with data available as of the time of each forecast, and fed realised data for the technical assumptions over the projection horizon.

filtering techniques described in Bańbura et al. (2015).

Figure 4 shows that 1-year-ahead conditional forecasts from the CQR model would have mostly tracked well the evolution of the data, including during the two sharp downturns experienced over the 2007-2009 and 2011-2013 periods, corresponding to the fallouts on the housing market of the Great Financial Crisis and European Sovereign Debt Crisis episodes. This remains largely true also for 2-years-ahead forecasts (Figure B.8). Such ability to track realised data is of course at least partly due to the model's counterfactual knowledge of the outturns for the technical assumptions at each forecast horizon, but is nevertheless reassuring. A comparison with the corresponding Figures B.10a and B.10b for the BVAR is also instructive. The confidence bands around the BVAR forecasts are wider, and of course symmetric by construction given the distributional assumption for the error terms in Equation 10. Instead, the somewhat narrower bands in Figures 4 and B.8 display noticeable changes in both variance and skewness.

This illustrates a key advantage of a more flexible approach to modelling predictive distributions: the BVAR is constrained to convey changes to forecast uncertainty over time solely through changes to the first and second moments of its predictive distributions, which are generally wider, and thus potentially less informative, as a result.²³ Instead, the CQR model can in

 $^{^{23}}$ A BVAR with stochastic volatility, such as the one proposed by Carriero et al. (2019), might be better at modeling changes to the second moments of predictive distributions, but would still not address the symmetry point.

	Log score		Quantile score					
			q = 0.05		q = 0.5		q = 0.95	
	CQR	BVAR	CQR	BVAR	CQR	BVAR	CQR	BVAR
h = 1	-1.27	-0.80	0.08	0.05	0.26	0.20	0.08	0.07
h=2	-1.69	-1.39	0.12	0.10	0.38	0.36	0.09	0.12
h = 3	-1.76	-1.79	0.15	0.15	0.51	0.52	0.12	0.18
h = 4	-1.95	-2.09	0.18	0.20	0.65	0.72	0.15	0.24
h = 5	-2.06	-2.16	0.18	0.21	0.77	0.77	0.16	0.25
h = 6	-2.17	-2.23	0.22	0.22	0.83	0.87	0.20	0.26
h = 7	-2.20	-2.32	0.19	0.24	0.85	0.97	0.20	0.27
h = 8	-2.28	-2.39	0.27	0.25	0.90	1.05	0.21	0.29
h = 9	-2.53	-2.49	0.33	0.28	0.99	1.16	0.31	0.32

Table 2: Performance metrics of CQR and BVAR out-of-sample conditional forecasts.

Note: The table reports average logarithmic and quantile scores over the forecast evaluation period 1999Q1 to 2020Q3, for a total of 79 vintages. A larger (less negative) average logarithmic score, and a smaller average quantile score indicate better forecasting performance.

addition also change their higher moments, at times avoiding the need for large changes to their mean or variance. The ability to also convey information about the balance of risks, which are captured by a distribution's skewness, is an important advantage in a policy setting, and was, for example, the motivation behind the choice of the two-part Normal distribution for the Bank of England's official fan charts (see Britton et al., 1998).

The probability integral transforms (PITs) of the two sets of forecasts (Figure B.11) show that the CQR's greater flexibility doesn't come at the cost of a worse calibration. The PIT is defined as the (predictive) cumulative distribution function evaluated at the actual data outturn, formally

$$PIT_{v,t_v+h} = \int_{-\infty}^{\pi_{t_v+h}^{hp}} \hat{p}_{v,h}(w) dw \equiv \hat{P}_{v,h}(\pi_{t_v+h}^{hp}), \tag{11}$$

where v indexes vintages, $\hat{p}_{v,h}(w)$ is the predictive density function estimated in vintage v for forecast horizon h, and $\hat{P}_{v,h}(w)$ the corresponding cumulative distribution function. An ideally-calibrated model should deliver a sequence of predictive distributions whose PITs are distributed uniformly over the unit interval, that is, should lie on the diagonal of each panel of Figure B.11. The figure shows that the two sets of PITs are rather similar, and in fact, for horizons up to 7 quarters ahead, on this metric the CQR's PITs are somewhat better calibrated than the BVAR's.

A comparison of the two models' density forecasting performance also shows that the two models typically score similarly, with a few exception discussed below. Table 2 reports two sets of metrics: the average logarithmic score as a summary measure for the entire predictive distribution, and three average quantile scores measuring local fit in the tails (5^{th} and 95^{th} percentiles) and for the median of the predictive distribution. The average logarithmic score is defined as

$$LS_{h} = \frac{1}{N_{v}} \sum_{v} \ln \hat{p}_{v,h}(y_{t_{v}+h})$$
(12)

where N_v denotes the number of forecast vintages. The log score penalises forecasts that assign a low probability to actual outturns. The average quantile score is defined as

$$QS_{q,h} = \frac{1}{N_v} \sum_{v} \rho_q \left(y_{t_v+h} - \hat{P}_{v,h}^{-1}(q) \right)$$
(13)

The quantile score penalises outturns that are more extreme (i.e. fall further in the corresponding tail) than the predictive quantile $\hat{P}_{v,h}^{-1}(\tau)$.

Except for the first two quarters, the log scores of the CQR model are similar or slightly better than those of the BVAR (first two columns of Table 2), indicating that the two models' overall density forecast ability is comparable. The difference in log scores in the first two quarters is however mostly due to a very small number of outliers for the CQR model, rather than to a more systematic discrepancy, and the two models otherwise perform similarly also in the first two quarters.²⁴ As for the quantile scores, the CQR model generally performs very similarly to the BVAR in the left tail and for the conditional median (columns 4 to 6 in Table 2), and a little bit better for the right tail, indicating that the CQR model is somewhat better at forecasting upside risk to house prices inflation.

The overall picture from the forecast evaluation exercise is therefore that despite its greater flexibility in modelling asymmetric risks, the CQR model is able to deliver out-of-sample density forecasts that are for the most part better or at a par with those from an established benchmark such as the BVAR. In the next Section, I show how this greater flexibility can be exploited for the purposes of scenario analysis of risks.

5 Scenario Analysis for the Balance of Risks

Scenario analyses are ubiquitous in central banks and other policy institutions. As discussed in the Introduction, model-based scenarios typically take the form of forecasts for one or more variables of interest conditional on the path of of some other variables (technical assumptions in the terminology of this paper). In this section I show how the greater flexibility of the CQR model can be leveraged for the purposes of scenario analysis of risks.

Over the projection horizon, policy makers might wish to understand the implications of alternative paths for the technical assumptions. To continue the house price inflation example of Section 3, they might query the impact on house price inflation of alternative paths for the future evolution of the unemployment rate and the long-term real interest rate. Figure 5 shows the time series of the two variables,²⁵ extended over the projection horizon with three alternative paths: a 'no change' scenario (blue), in which both variables remain at their current

²⁴As is well-known, the log score imposes a disproportionately heavy penalty on outturns that fall in regions of very low probability density. For the CQR model, this occurs 3 and 4 times over the forecast evaluation sample for h = 1 and h = 2, respectively.

²⁵Note that as per Table 1, both variables enter the model in first differences, but are shown in levels in Figure 5 for ease of interpretation.





Note: The charts show the time series over the available sample period (up to and including 2020Q3) and three alternative scenarios over the projection horizon (up to 9 quarters ahead) for the long-term real interest rate and the unemployment rate, as defined in Table 1.

value throughout the projection; a 'stress' scenario (red), in which the unemployment rate is 2pp higher, and the real interest rate 1pp higher, by the end of the projection; and a recovery scenario, in which the unemployment rate is 1pp lower, and the real interest rate also 1pp lower, by the end of the projection (green). In this case, the three scenarios are purely arbitrary, but a scenario for the technical assumptions might equally be the output of another forecast which policy makers want their house price inflation forecast to be consistent with. I also consider a fourth scenario, in which no technical assumptions are provided, but only knowledge of the latest available data is exploited, as would be the case with current implementations of quantile regression (black). Figure 6 shows the resulting fan charts for the CQR model estimated with data up to 2020Q2.²⁶ Moreover, Figure 7 shows the corresponding predictive densities 1 and 2 years ahead, benchmarked against the *unconditional* distribution of house price inflation over the estimation sample.

The forecast based on available data only, shown in black, is a purely statistical forecast without any additional inputs (very much like a VAR forecast obtained by pure iteration from current conditions), and is therefore a useful starting point. It shows a *median* forecast settling around 2.5% within a year's time, but within a wide range of possible outcomes and risks mostly skewed to the downside. For forecasts 2 years ahead (Figure 7), this is the most optimistic forecast (in terms of the *mode*), but also the one with the highest variance and displaying most downside skewness. In terms of the mode (the most likely outcome), a no-change forecast, shown in blue, is a bit more optimistic 1 year ahead, but considerably more cautious 2 years out; 1 year ahead, this comes with much less dispersion, but still signalling downside risks, while 2 years ahead the risks are much more balanced, while the variance of the distribution is somewhat smaller than for the available data only scenario.

²⁶And knowledge of unemployment, household income and interest data up to 2020Q3.



Figure 6: Alternative scenarios for the predictive distributions of house price inflation.

(c) Recovery.

(d) Stress.

Note: The fan charts show realised data for house price inflation up to 2020Q2, and then density forecasts for 2020Q3 to 2022Q3 under different technical assumptions for the unemployment and long term real interest rates, with the colour coding matching that of Figure 5: no technical assumptions over the projection (black), no change in technical assumptions over the projection (blue), a recovery scenario involving a gradual improvement in the technical assumptions (green), and a stress scenario involving a gradual worsening in the technical assumptions (red). The fan chart bands cover 95% of the support around the median: the darkest shade corresponds to the median, while lighter shades represent percentiles increasingly removed from it. The forecasts are based on a CQR model estimated with data up to 2020Q2, and fed projections for the technical assumptions over the forecast horizon.

Figure 7: Alternative scenarios for the predictive distributions of house price inflation 1 and 2 years ahead.



Note: The two panels show density forecasts for 2021Q3 and 2022Q3, respectively, under different technical assumptions for the unemployment and long term real interest rates. The colour coding matches that of Figure 5: no technical assumptions over the projection (black), no change in technical assumptions over the projection (blue), a recovery scenario involving a gradual improvement in the technical assumptions (green), and a stress scenario involving a gradual worsening in the technical assumptions (red). The unconditional distribution of house price inflation over the estimation sample is also shown (dashed black). The forecasts are based on a CQR model estimated with data up to 2020Q2, and fed projections for the technical assumptions over the forecast horizon.

The stress and recovery scenarios for the technical assumptions, shown in red and green, respectively, also imply interesting concomitant changes to the predictive distributions. First of all, the modal (and median) forecasts in the stress scenario are unambiguously the lowest ones, while in the recovery scenario they are at the high end,²⁷ as one would expect. Second, the recovery scenario displays consistently lower variance than the stress one, and is in fact the one with the lowest variance across all four scenarios and all horizons. And finally, while both scenarios signal some downside risks 1 year ahead, by the end of the projection the risks are broadly balanced under both sets of paths for the technical assumptions.

The above discussion illustrates the value added that the CQR approach can bring to the assessment of risks. Different projections for the technical assumptions can lead to meaningful changes not only to the first and second moments of predictive distributions, as would be the case, for example, if conditional forecasts based on them had been generated by the BVAR introduced in the previous Section, but also to higher moments. In the proposed framework, these changes can be easily quantified and communicated using well-established tools such as fan charts and plots of the predictive densities, allowing for example a model-based sensitivity analysis of upside and downside risks around a modal forecast, an essential element of most institutional forecasting processes, but one that so far has been largely limited to qualitative assessments or an interpretation of model-based outputs not explicitly designed for this purpose.

 $^{^{27}}$ The modal recovery forecast is below the one based on available data only, but given the absence of skew, its median forecast is the highest one, as can be gleaned from the fan charts in Figure 6.

6 Conclusion

I propose a novel approach, conditional quantile regression (CQR), to incorporate technical assumptions over the projection horizon into an otherwise standard quantile regression framework. The approach addresses a shortcoming of existing quantile regression-based models, for example the by now common at-risk framework, when used in settings, such as most forecasting processes within central banks and similar institutions, that require forecasts to be consistent with a specific view of the future evolution of some variables. Through an application to house price inflation in the euro area, I show that CQR provides a viable alternative to existing approaches to conditional density forecasting, notably VARs, with considerable advantages in terms of flexibility and additional insights that do not come at the cost of forecasting performance.

References

ADAMS, P. A., T. ADRIAN, N. BOYARCHENKO, AND D. GIANNONE (2021): "Forecasting macroeconomic risks," *International Journal of Forecasting*, 37, 1173–1191.

ADRIAN, T., N. BOYARCHENKO, AND D. GIANNONE (2019): "Vulnerable Growth," *American Economic Review*, 109, 1263–1289.

ANDERSSON, M. K., S. PALMQVIST, AND D. F. WAGGONER (2010): "Density-Conditional Forecasts in Dynamic Multivariate Models," Working Paper Series 247, Sveriges Riksbank (Central Bank of Sweden).

ANGELINI, E., M. LALIK, M. LENZA, AND J. PAREDES (2019): "Mind the gap: a multi-country BVAR benchmark for the Eurosystem projections," Working Paper Series 2227, European Central Bank.

ANTOLÍN-DÍAZ, J., I. PETRELLA, AND J. F. RUBIO-RAMÍREZ (2021): "Structural scenario analysis with SVARs," *Journal of Monetary Economics*, 117, 798–815.

AZZALINI, A. AND A. CAPITANIO (2003): "Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution," *Journal of the Royal Statistical Society Series B*, 65, 367–389.

BAŃBURA, M., D. GIANNONE, AND M. LENZA (2015): "Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections," *International Journal of forecasting*, 31, 739–756.

BANK OF ENGLAND (2020): "August 2020 Monetary policy report," Tech. rep.

BARNDORFF-NIELSEN, O., J. KENT, AND M. SØRENSEN (1982): "Normal Variance-Mean Mixtures and z Distributions," *International Statistical Review / Revue Internationale de Statistique*, 50, 145–159.

BARNDORFF-NIELSEN, O. E. AND A. E. KOUDOU (1998): "Trees with Random Conductivities and the (Reciprocal) Inverse Gaussian Distribution," *Advances in Applied Probability*, 30, 409–424.

BATTISTINI, N., M. FALAGIARDA, J. GAREIS, A. HACKMANN, AND M. ROMA (2021): "The euro area housing market during the COVID-19 pandemic," *Economic Bulletin Articles*, 7/21.

BOK, B., D. CARATELLI, D. GIANNONE, A. M. SBORDONE, AND A. TAMBALOTTI (2018): "Macroeconomic Nowcasting and Forecasting with Big Data," *Annual Review of Economics*, 10, 615–643.

BONEVA, L., N. FAWCETT, R. M. MASOLO, AND M. WALDRON (2019): "Forecasting the UK economy: Alternative forecasting methodologies and the role of off-model information," *International Journal of Forecasting*, 35, 100–120.

BRITTON, E., P. FISHER, AND J. WHITLEY (1998): "The Inflation Report projections: understanding the fan chart," *Bank of England Quarterly Bulletin*, February, 30–37.

BURGESS, S., E. FERNANDEZ-CORUGEDO, C. GROTH, R. HARRISON, F. MONTI, K. THEODORIDIS, AND M. WALDRON (2013): "The Bank of England's forecasting platform: COMPASS, MAPS, EASE and the suite of models," Bank of England working papers 471, Bank of England.

CARRIERO, A., T. E. CLARK, AND M. MARCELLINO (2019): "Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors," *Journal of Econometrics*, 212, 137–154.

CHALMOVIANSKÝ, J., M. PORQUEDDU, AND A. SOKOL (2020): "Weigh(t)ing the basket: aggregate and component-based inflation forecasts for the euro area," Working Paper Series 2501, European Central Bank.

CHAVLEISHVILI, S. AND S. MANGANELLI (2019): "Forecasting and stress testing with quantile vector autoregression," Working Paper Series 2330, European Central Bank.

CHERNOZHUKOV, V., I. FERN-NDEZ-VAL, AND A. GALICHON (2010): "Quantile and Probability Curves Without Crossing," *Econometrica*, 78, 1093–1125.

CIMADOMO, J., D. GIANNONE, M. LENZA, F. MONTI, AND A. SOKOL (2020): "Nowcasting with large Bayesian vector autoregressions," Working Paper Series 2453, European Central Bank.

CLARK, T. E. AND M. W. MCCRACKEN (2014): "Evaluating Conditional Forecasts from Vector Autoregressions," Working Papers 2014-25, Federal Reserve Bank of St. Louis.

COENEN, G., P. KARADI, S. SCHMIDT, AND A. WARNE (2018): "The New Area-Wide Model II: an extended version of the ECB's micro-founded model for forecasting and policy analysis with a financial sector," Working Paper Series 2200, European Central Bank.

DEL NEGRO, M. AND F. SCHORFHEIDE (2013): "DSGE Model-Based Forecasting," in *Handbook of Economic Forecasting*, ed. by G. Elliott, C. Granger, and A. Timmermann, Elsevier, vol. 2 of *Handbook of Economic Forecasting*, chap. 0, 57–140.

DOAN, T., R. LITTERMAN, AND C. SIMS (1984): "Forecasting and conditional projection using realistic prior distributions," *Econometric reviews*, 3, 1–100.

DOMIT, S., F. MONTI, AND A. SOKOL (2019): "Forecasting the UK economy with a medium-scale Bayesian VAR," *International Journal of Forecasting*, 35, 1669–1678.

EGUREN-MARTIN, F., C. O'NEILL, A. SOKOL, AND L. V. D. BERGE (2021): "Capital flowsat-risk: push, pull and the role of policy," Working Paper Series 2538, European Central Bank. EGUREN-MARTIN, F. AND A. SOKOL (2020): "Attention to the tail(s): global financial conditions and exchange rate risks," Working Paper Series 2387, European Central Bank.

EUROPEAN CENTRAL BANK (2021a): "Inflation measurement and its assessment in the ECB's monetary policy strategy review," Occasional Papers 265.

(2021b): "September 2021 ECB staff macroeconomic projections for the euro area," Tech. rep.

FAUST, J. AND J. H. WRIGHT (2008): "Efficient forecast tests for conditional policy forecasts," *Journal of Econometrics*, 146, 293–303.

FERRARA, L., M. MOGLIANI, AND J.-G. SAHUC (2020): "High-frequency monitoring of growth-at-risk," CAMA Working Papers 2020-97, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University.

FIGUERES, J. M. AND M. JAROCIŃSKI (2020): "Vulnerable growth in the euro area: Measuring the financial conditions," *Economics Letters*, 191.

FITZENBERGER, B. (1998): "The moving blocks bootstrap and robust inference for linear least squares and quantile regressions," *Journal of Econometrics*, 82, 235 – 287.

GAGLIANONE, W. P. AND L. R. LIMA (2012): "Constructing Density Forecasts from Quantile Regressions," *Journal of Money, Credit and Banking*, 44, 1589–1607.

GELOS, R. G., L. GORNICKA, R. KOEPKE, R. SAHAY, AND S. SGHERRI (2019): "Capital Flows at Risk: Taming the Ebbs and Flows," IMF Working Papers 19/279, International Monetary Fund.

GERDESMEIER, D., A. LENARČIČ, AND B. ROFFIA (2015): "An alternative method for identifying booms and busts in the Euro area housing market," *Applied Economics*, 47, 499–518.

GIANNONE, D., M. LENZA, AND G. E. PRIMICERI (2015): "Prior Selection for Vector Autoregressions," *The Review of Economics and Statistics*, 97, 436–451.

HASENZAGL, THOMAS, M. P.-M. L. R. AND G. RICCO (2020): "When is growth at risk?" *Brookings Papers on Economic Activity*, Spring, 167–229.

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

KOENKER, R. AND G. BASSETT (1978): "Regression Quantiles," Econometrica, 46, 33-50.

KOENKER, R. AND Z. XIAO (2006): "Quantile Autoregression," Journal of the American Statistical Association, 101, 980–990.

KOROBILIS, D. (2017): "Quantile regression forecasts of inflation under model uncertainty," International Journal of Forecasting, 33, 11 – 20. KOTZ, S., T. KOZUBOWSKI, AND K. PODGORSKI (2001): The Laplace Distribution and Generalizations: A Revisit with Applications to Communications, Economics, Engineering, and Finance, Birkhäuser, Boston, MA.

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

LITTERMAN, R. B. (1979): "Techniques of forecasting using vector autoregressions," Tech. rep.

LÓPEZ-SALIDO, J. D. AND F. LORIA (2020): "Inflation at Risk," Finance and Economics Discussion Series 2020-013, Board of Governors of the Federal Reserve System (U.S.).

MANZAN, S. AND D. ZEROM (2013): "Are macroeconomic variables useful for forecasting the distribution of U.S. inflation?" *International Journal of Forecasting*, 29, 469–478.

MCCRACKEN, M. W. AND J. T. MCGILLICUDDY (2019): "An empirical investigation of direct and iterated multistep conditional forecasts," *Journal of Applied Econometrics*, 34, 181–204.

SIMS, C. A. (1993): "A nine-variable probabilistic macroeconomic forecasting model," in *Business cycles, indicators, and forecasting*, University of Chicago press, 179–212.

SOKOL, A. (2020): "Bayesian Quantile Regression with Correlated Errors," mimeo.

TAGLIABRACCI, A. (2020): "Asymmetry in the conditional distribution of euro-area inflation," Temi di discussione (Economic working papers) 1270, Bank of Italy, Economic Research and International Relations Area.

VALCKX, N., M. S. SHAHID, M. KATAGIRI, AND A. DEGHI (2020): "Predicting Downside Risks to House Prices and Macro-Financial Stability," IMF Working Papers 2020/011, International Monetary Fund.

YU, K. AND R. A. MOYEED (2001): "Bayesian quantile regression," *Statistics & Probability Letters*, 54, 437–447.

A Gibbs Sampling for Bayesian Quantile Regression

Combining equation 8 with conjugate prior distributions for β and σ , namely $\beta \sim \mathcal{N}(\beta_0, \Sigma_0)$ and $\sigma \sim Inv\Gamma(\alpha_0, \gamma_0)$ implies

$$z_t | \boldsymbol{\beta}, \boldsymbol{\delta}, \sigma, \boldsymbol{x}, \boldsymbol{y} \propto \frac{1}{\sqrt{z_t}} \exp\left\{-\frac{(y_t - \boldsymbol{x}_t' \boldsymbol{\beta})^2}{2\sigma \tau^2 z_t} - \frac{\theta^2 + 2\tau^2}{2\sigma \tau^2} z_t\right\},\tag{A.1}$$

which is the density of a Reciprocal Inverse Gaussian random variable²⁸ with parameters $\lambda_t = \frac{\theta^2 + 2\tau^2}{2\sigma\tau^2}$ and $\mu_t = \frac{\sqrt{\theta^2 + 2\tau^2}}{|y_t - x'_t \beta|}$;

$$\boldsymbol{\beta}|\boldsymbol{\delta},\sigma,\boldsymbol{x},\boldsymbol{y},\boldsymbol{z}\sim\mathcal{N}\left(\hat{\boldsymbol{\beta}},\hat{\boldsymbol{\Sigma}}\right)$$
 (A.2)

with

$$\hat{\mathbf{\Sigma}} = \left(oldsymbol{x}'oldsymbol{W}oldsymbol{x} + \mathbf{\Sigma}_0^{-1}
ight)^{-1} \qquad \hat{oldsymbol{eta}} = \hat{\mathbf{\Sigma}} \left(oldsymbol{x}'oldsymbol{W}oldsymbol{y} - rac{ heta}{\sigma au^2}oldsymbol{x}'oldsymbol{\ell} + \mathbf{\Sigma}_0^{-1}oldsymbol{eta}_0
ight),$$

where \boldsymbol{W} is a $T \times T$ diagonal matrix with diagonal elements given by $(\sigma \tau^2 z_t)^{-1}$, $\boldsymbol{\ell}$ is a T vector of ones, and $\boldsymbol{\Sigma}_0$ is a diagonal matrix with diagonal elements given by $\boldsymbol{\delta}$, whose individual elements have (hyper-)prior distributions $\delta_i^2 \sim InvGamma(\eta_0, \zeta_0)$ and posterior distributions given by

$$\delta_i^2 | \boldsymbol{\beta}, \sigma, \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \sim Inv\Gamma\left(\hat{\eta}, \hat{\zeta}\right),$$
(A.3)

where $\hat{\eta} = \eta_0 + \frac{1}{2}$ and $\hat{\zeta} = \zeta_0 + \frac{\beta_i^2}{2}$; and finally,

$$\sigma | \boldsymbol{\beta}, \boldsymbol{\delta}, \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \sim Inv\Gamma\left(\hat{\alpha}, \hat{\gamma}\right), \tag{A.4}$$

with

$$\hat{\alpha} = \alpha_0 + \frac{3T}{2} \qquad \hat{\gamma} = \gamma_0 + \sum_{t=1}^T \frac{\left(y_t - \boldsymbol{x}_t'\boldsymbol{\beta} - \boldsymbol{\theta}z_t\right)^2}{2\tau^2 z_t} + \boldsymbol{z}'\boldsymbol{\ell}.$$

$$Z \sim Exp\left(\sigma\right) \stackrel{d}{=} \Gamma\left(1, \frac{1}{\sigma}\right) \stackrel{d}{=} \mathcal{GIG}\left(1, \frac{2}{\sigma}, 0\right).$$

Moreover, the following relationships hold between the Generalised-Inverse-Gaussian and the (Reciprocal) Inverse Gaussian distribution (Barndorff-Nielsen and Koudou, 1998):

$$X \sim \mathcal{GIG}\left(-\frac{1}{2}, \frac{\lambda}{\mu^2}, \lambda\right) \stackrel{d}{=} \mathcal{IG}\left(\mu, \lambda\right) \Longleftrightarrow \frac{1}{X} \sim \mathcal{GIG}\left(\frac{1}{2}, \lambda, \frac{\lambda}{\mu^2}\right) \stackrel{d}{=} \mathcal{RIG}\left(\lambda, \mu\right)$$

 $^{^{28}{\}rm This}$ follows directly from the fact that in a Normal variance-mean mixture model, the conjugate prior for Z is Generalised-Inverse-Gaussian, and

B Additional Figures



Figure B.8: 2-year-ahead conditional forecasts for 2001Q1-2020Q2.

Note: The chart shows successive vintages of 2-year-ahead forecasts for house price inflation, with corresponding data outturns overlaid in red. The bands cover 68% and 95% of the support around the median. The forecasts are based on a CQR model estimated with data available as of the time of each forecast, and fed realised data for the technical assumptions over the projection horizon.





Note: The fan chart shows realised data for house price inflation up to 2018Q2, and then density forecasts for 2018Q3 to 2020Q3, with actual data outturns for the forecast period overlaid in red. The fan chart bands cover 95% of the support around the median: the darkest shade of orange corresponds to the median, while lighter shades represent percentiles increasingly removed from it. The forecasts are based on a BVAR model estimated with data up to 2018Q2, and fed realised data for the technical assumptions over the projection horizon.





(a) 1-year-ahead conditional forecasts for 2000Q1- (b) 2-years-ahead conditional forecasts for 2001Q1-2019Q2, BVAR 2020Q2,BVAR

Note: The two panels show successive vintages of 1- and 2-year-ahead forecasts for house price inflation, with corresponding data outturns overlaid in red. The bands cover 68% and 95% of the support around the median. The forecasts are based on a BVAR model estimated with data available as of the time of each forecast, and fed realised data for the technical assumptions over the projection horizon.



Figure B.11: Probability integral transforms (PITs), CQR and BVAR.

Note: The PIT of a forecast is defined as the predictive cumulative distribution function evaluated at the data outturn. Each panel shows the empirical distributions of the two models' forecast vintages. An well-calibrated model should deliver an empirical distribution close to the diagonal, shown in black. For each panel, the mean absolute deviation between each model's PITs and those of an ideal (uniform) distribution are also reported. A smaller value indicates that the model is closer to the ideal, that is, better-calibrated.

Acknowledgements

I am grateful to Domenico Giannone, Marek Jarociński, Mike McCracken, Matteo Mogliani, Thomas Westermann and seminar participants at the European Central Bank for helpful comments.

Andrej Sokol

Centre for Macroeconomics, London, United Kingdom; email: sokolandrej@gmail.com

© European Central Bank, 2021

Postal address60640 Frankfurt am Main, GermanyTelephone+49 69 1344 0Websitewww.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the Social Science Research Network electronic library or from RePEc: Research Papers in Economics. Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website.

² DF ISBN 978-92-899-4911-8	ISSN 1725-2806	doi:10.2866/966212	Q
--	----------------	--------------------	---

QB-AR-21-115-EN-N