Kent Academic Repository

Journal Pre-proof

Bayesian Mixed-Frequency Quantile Vector Autoregression: Eliciting tail risks of Monthly US GDP

Matteo Iacopini, Aubrey Poon, Luca Rossini and Dan Zhu

Please cite this article as: M. Iacopini, A. Poon, L. Rossini et al., Bayesian Mixed-Frequency Quantile Vector Autoregression: Eliciting tail risks of Monthly US GDP, *Journal of Economic Dynamics and Control*, 104757, doi: https://doi.org/10.1016/j.jedc.2023.104757.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Bayesian Mixed-Frequency Quantile Vector Autoregression: Eliciting tail risks of Monthly US GDP *

Matteo Iacopini[†]   Aubrey Poon[‡]   Luca Rossini[§]   Dan Zhu[¶]

October 5, 2023

Abstract

Timely characterizations of risks in economic and financial systems play an essential role in both economic policy and private sector decisions. However, the informational content of low-frequency variables and the results from conditional mean models provide only limited evidence to investigate this problem. We propose a novel mixed-frequency quantile vector autoregression (MF-QVAR) model to address this issue. Inspired by the univariate Bayesian quantile regression literature, the multivariate asymmetric Laplace distribution is exploited under the Bayesian framework to form the likelihood. A data augmentation approach coupled with a precision sampler efficiently estimates the missing low-frequency variables at higher frequencies under the state-space representation.

The proposed methods allow us to analyse conditional quantiles for multiple variables of interest and to derive quantile-related risk measures at high frequency, thus enabling timely policy interventions. The main application of the model is to detect the vulnerability in the US economy and then to nowcast conditional quantiles of the US GDP, which is strictly related to the quantification of Value-at-Risk, the Expected Shortfall and distance among percentiles of real GDP nowcasts.

Keywords: Bayesian inference; mixed-frequency; multivariate quantile regression; nowcasting; VAR.

1 Introduction

Most economic models' primary object of interest is the conditional mean of a given variable or index, as it summarizes the central response to explanatory variables. However, following the financial crises and economic shocks that characterised the last decade, policymakers and researchers have shifted their attention and interest beyond the conditional mean. In particular, the significant effects that exogenous shocks (such as the COVID-19 pandemic), wars (e.g., the Russian invasion of Ukraine), and fluctuations of the business cycles have on the economy highlight the crucial need to investigate the tails and shoulders of the response variable's distribution.

Timely characterisations of risks to the economic outlook play a vital role in both economic policy and private sector decisions, where central bankers and analysts share a demand for timely forecasts of economic activity. To encapsulate and reflect the most recent events, forecasts of macroeconomic or financial variables should blend information collected from a wide array of sources and observed at different intervals or frequencies.

Moreover, research following the global financial crisis has provided substantial empirical evidence that the relationships among macroeconomic and financial time series are characterised by nonlinearities and asymmetries (Kilian and Vigfusson, 2017; Adrian et al., 2019). Thus, investigating the nonlinear effects related to cycles is crucial to policymakers for designing policies targeted at specific phases of the cycles. Researchers in macroeconomics usually base their analysis on linear regression methods, whereas only recently nonlinear methods have been applied to investigate economic policies and financial crises (e.g., Caggiano et al., 2022; Huber and Rossini, 2022). However, using conditional mean regression methods raises several concerns when modelling data with features such as skewness, fat tails, outliers, truncation, censoring, and heteroscedasticity. This relates to the fact that the impact of covariates on the response may significantly vary across the range of the latter, thus highlighting the limitations of methods based on conditional mean only. The issue is exacerbated by nonlinear re-

lationships and non-Gaussian noises, typical features of many economic and financial variables.

In detecting economic and financial crises, appropriate risk measures are needed (Merlo et al., 2021), such as the Value at Risk (VaR) or the Expected Shortfall (ES). The VaR considers the maximum loss an operator can incur over a defined time horizon and for a given confidence level. At the same time, the ES coincides with the conditional expectation of exceedance beyond the VaR. In forecasting, Gneiting and Ranjan (2011) propose a threshold- and quantile-based decomposition of the continuously ranked probability score to assess density forecasting over the whole distribution and specific quantiles or regions of a variable of interest (e.g., the tails).

To address the limitations of standard linear regression models and capture the complete picture of the conditional distribution of a multivariate response variable, we propose a novel mixed-frequency quantile vector autoregressive (MF-QVAR) model. Our approach is based on quantile regression (QR, see Koenker and Bassett, 1978), which offers robust modelling of conditional quantiles and allows for different impacts of covariates on each quantile level. This enables a comprehensive investigation of the entire conditional distribution. To conduct a simultaneous inference under the Bayesian framework on the marginal conditional quantiles of a multivariate response variable, we use the multivariate asymmetric Laplace (MAL) distribution to form the likelihood (see also Petrella and Raponi, 2019).

Our approach allows for a flexible covariance matrix within the multivariate structure, which benefits from cross-sectional information to estimate the marginal conditional quantiles. Considering different quantile levels, our framework permits the investigation of asymmetry in the downside and upside risks, unlike standard models with symmetric second-moment dynamics. This is particularly effective when skewness dynamics accompany the evolution of the distribution. Therefore, the proposed MF-QVAR is more effective in modelling the conditional distribution of a multivariate response variable, especially in the presence of dynamic asymmetries in the distribution.

3

The main feature of the proposed MF-QVAR model is the coupling of multivariate quantile regression with state-space-based mixed-frequency techniques. By combining them, we provide a theoretically coherent and computationally efficient method to forecast and nowcast a desired quantile of low-frequency variables of interest.

Most of the existing approaches for quantile forecasting are actually based on a model for the conditional mean, then rely on the quantiles of the implied predictive distribution (e.g., see Adrian et al., 2019; Carriero et al., 2022). Conversely, we directly model and forecast the conditional quantile of interest, allowing for heterogeneity of the impact of covariates across different quantiles, which is prevented in the above-mentioned approaches. Moreover, by adopting a multivariate framework, the MF-QVAR allows capturing contemporaneous cross-sectional dependence, which is known to characterise economic and financial time series.

We extend the mixed-frequency literature to develop a method for nowcasting quantiles of low-frequency variables using a state-space representation that exploits the information available at a higher frequency. This approach enables policymakers to timely detect early signs of distress and implement corrective measures to counteract the deterioration of the economy. While there is a substantial amount of literature based on the mixed-data sampling (MIDAS) formulation to estimate conditional quantiles of lower frequency variables (e.g., Adams et al., 2021; Carriero et al., 2022), our MF-QVAR framework is the first to provide these conditional quantile estimates at a higher frequency. For example, we provide monthly conditional quantiles of GDP growth, a quarterly variable, instead of the quarterly conditional quantiles provided by the standard MIDAS approach.

Moreover, the proposed approach overcomes the "ragged-edge" problem, which arises due to differences in data release dates that cause the available information set to differ over time within the quarter. This automatic adjustment ensures that the forecaster can use all available information to generate accurate nowcasts and forecasts. Finally, the MF-QVAR model leverages the mixed-frequency structure and flexible multivariate

4

dynamics to blend the high-frequency temporal evolution into estimating conditional quantiles of low-frequency variables in the high-frequency domain. This provides a natural mechanism for incorporating high-frequency information into estimating low-frequency variables, allowing us to capture the high frequency dynamic relationship between the variables more accurately.

To overcome the computational challenge in the mixed-frequency VAR, we follow Chan et al. (2023), who designed a computationally efficient sampler for state-space models with missing observations, such as MF-VARs. They exploit the block-banded structure of the precision matrix of the conditional distribution of the missing observations to adapt the precision-based sampler of Chan and Jeliazkov (2009) to draw the missing low-frequency variables (see also Rue and Held, 2005). The importance of the precision sampler to modern econometric models is proven by its use in a variety of settings, including models with missing observations (Hauber and Schumacher, 2021), and dynamic factor models (Kaufmann and Schumacher, 2019). An earlier method to make inference on unobserved variables in state-space models is the simulation smoother (Durbin and Koopman, 2002). Moreover, we impose linear constraints when sampling the missing observations to ensure that missing high-frequency observations match the observed values of the low-frequency variables. This results in sampling from a linearly constrained Gaussian distribution, which is efficiently performed following the methods in Cong et al. (2017).

To assess the forecasting accuracy of our proposed framework, we conduct two Monte Carlo simulation experiments. In the first experiment, we estimate both the univariate quantile regression and QVAR models using a single observed frequency without any missing data. The purpose is to showcase the effectiveness of jointly estimating conditional quantiles, as opposed to individual estimation, in enhancing forecasting performance. For the second experiment, we consider a mixed-frequency setting where we compare our proposed MF-QVAR model against the conditional mean counterpart of Schorfheide and Song (2015). Overall, we find that our proposed MF-QVAR model is

5

highly accurate in capturing complex dependencies and unobserved variables compared to the standard models employed in the literature.

We apply our novel MF-QVAR model in both an in- and out-of-sample context. For the in-sample analysis, we estimate the MF-QVAR on a sample period from January 1973 to December 2021. Then, we compare our model's in-sample monthly US growth-at-risk estimates against a quarterly measure derived from univariate quantile regression. We find that these two frequency estimates display substantial differences in times of a recession. For instance, our monthly US growth-at-risk estimates appear to detect the vulnerability in the US economy considerably faster than the quarterly quantile regression for the Great Recession of 2007-08. We find that when modelling US growth-at-risks during a recession, the low-frequency estimate can exhibit significantly different dynamics compared to its high-frequency counterpart.

We conduct an out-of-sample analysis using our proposed MF-QVAR model to nowcast monthly US growth-at-risk in real-time for two specific episodes. First, we focus on the Global Financial Crisis (GFC) period, employing data vintage spanning from January 2005 to December 2010. Second, we focus our attention on the period of the COVID-19 pandemic and the Russian invasion of Ukraine, utilising data vintage from January 2016 to March 2022. We generate monthly nowcasts and forecasts based on the three release timings of US real GDP by the US Bureau of Economic Analysis (BEA).

Our analysis reveals four key insights. First, the proposed MF-QVAR appears to produce monthly nowcasts capable of detecting the vulnerability of the US economy at the earlier stage of the GFC compared to its quarterly counterpart. This observation indicates that the MF-QVAR model offers valuable insights into the economic conditions preceding the GFC. Second, the negative growth-at-risk estimates generated by our monthly analysis during the early 2010 period align with the prevailing theme of weakened economic growth experienced by the US economy during that period. In contrast, the quarterly counterpart suggests that the US economy had already rebounded from the recessionary phase that transpired from late 2008 to early 2009.

Third, we observe a downward shift in the monthly nowcasts of US growth-at-risk since the pandemic, with the average dropping from about -3% to -5%. Additionally, we find that the distance between the 10th and 50th percentile of the real GDP nowcasts significantly increases during the pandemic period, indicating that the monthly distribution of US real GDP has become more skewed to the left. Finally, we compare our monthly nowcasts of US growth-at-risk to their corresponding quarterly nowcasts from a quantile regression mixed-data sampling (MIDAS) model. On average, the quarterly nowcasts underestimate US growth-at-risk relative to our monthly nowcasts, which also appear to align with the current post-pandemic economic situation in the US.

The remainder of this article is organized as follows: Section 2 presents a novel mixed-frequency quantile VAR model under the acronym of MF-QVAR and the intertemporal constraints. In Section 3, the Bayesian approach for inference along with the posterior algorithm is described. Section 4 investigates the performance of our method using simulated data. Section 5 shows the results of the in- and out-of-sample analysis on US real GDP growth-at-risk. Finally, Section 6 draws the conclusions.

2   Mixed-Frequency Quantile VAR

2.1   Notation

Let $\mathbb{S}^k = \{X \in \mathbb{R}^{k \times k} : X = X'\}$ denote the space of symmetric matrices of size $k \times k$ and $\mathbb{S}^k_{++} = \{X \in \mathbb{S}^k : \mathbf{a}'X\mathbf{a} > 0, \ \forall \, \mathbf{a} \in \mathbb{R}^k\}$ be the space of symmetric, positive definite matrices of size $k \times k$. For a matrix $A \in \mathbb{S}^k_{++}$, $A^{1/2}$ represents the Cholesky factor of A. Let $\mathrm{MAL}_n(\boldsymbol{\mu}, \boldsymbol{\delta}, \Sigma)$ indicate a multivariate asymmetric Laplace distribution with location $\boldsymbol{\mu} \in \mathbb{R}^n$, skewness parameter $\boldsymbol{\delta} \in \mathbb{R}^n$, and scale matrix $\Sigma \in \mathbb{S}^n_{++}$. Let $\mathbf{y}^o_t \in \mathbb{R}^{n_o}$ be a vector of variables observed at high-frequency and let $\mathbf{y}^u_t \in \mathbb{R}^{n_u}$ be a vector of variables that are unobserved or only partially observed. Finally, let I_n be the identity matrix of size n, and denote with $\mathbf{1}_n$ and $\mathbf{0}_n$ an n-dimensional vector with all entries

7

equal to one and zero, respectively. The symbol $\otimes$ denotes the Kronecker product.

2.2 Model

We consider a n-dimensional VAR(p) model with p lags for $\mathbf{y}_t = (\mathbf{y}_t^{o'}, \mathbf{y}_y^{u'})'$, with $n = n_o + n_u$, that is

$$\mathbf{y}_t = \mathbf{b}_0 + \sum_{j=1}^{p} B_j \mathbf{y}_{t-j} + \boldsymbol{\epsilon}_t, \qquad \boldsymbol{\epsilon}_t \sim \mathrm{MAL}_n(\mathbf{0}_n, D\boldsymbol{\theta}_{\tau,1}, D\boldsymbol{\theta}_{\tau,2}\Psi\boldsymbol{\theta}_{\tau,2}'D'), \qquad (1)$$

for $t = p+1, \ldots, T$, where $\mathbf{b}_0$ is a n-dimensional vector of intercepts, $B_1, \ldots, B_p$ are $(n \times n)$ autoregressive coefficient matrices, Ψ is a $(n \times n)$ correlation matrix, $D = \mathrm{diag}\left(\Sigma_{11}^{1/2}, \ldots, \Sigma_{nn}^{1/2}\right)$, for $\Sigma_{ii}^{1/2} \in \mathbb{R}$, and

$$\boldsymbol{\theta}_{\tau,1} = \left(\frac{1-2\tau_1}{\tau_1(1-\tau_1)}, \ldots, \frac{1-2\tau_n}{\tau_n(1-\tau_n)}\right)', \qquad \boldsymbol{\theta}_{\tau,2} = \mathrm{diag}\left(\sqrt{\frac{2}{\tau_1(1-\tau_1)}}, \ldots, \sqrt{\frac{2}{\tau_n(1-\tau_n)}}\right),$$

where $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_n)$ is the quantile representation, such that $\tau_i \in (0,1)$ for $i = 1, \ldots, n$. The multivariate asymmetric Laplace distribution, $\mathrm{MAL}_n(\boldsymbol{\mu}, D\boldsymbol{\theta}_{\tau,1}, D\boldsymbol{\theta}_{\tau,2}\Psi\boldsymbol{\theta}_{\tau,2}'D')$, has density function

$$f_Y(\mathbf{y}|\boldsymbol{\mu}, D\boldsymbol{\theta}_{\tau,1}, D\boldsymbol{\theta}_{\tau,2}\Psi\boldsymbol{\theta}_{\tau,2}'D') = \frac{2\exp\left\{(\mathbf{y}-\boldsymbol{\mu})'D^{-1}(\boldsymbol{\theta}_{\tau,2}\Psi\boldsymbol{\theta}_{\tau,2}')^{-1}\boldsymbol{\theta}_{\tau,1}\right\}}{(2\pi)^{n/2}\left|D\boldsymbol{\theta}_{\tau,2}\Psi\boldsymbol{\theta}_{\tau,2}'D\right|^{1/2}}$$
$$\times \left(\frac{\tilde{m}}{2+\tilde{d}}\right)^{\nu} K_{\nu}\left(\sqrt{(2+\tilde{d})\tilde{m}}\right),$$

where $\tilde{m} = (\mathbf{y}-\boldsymbol{\mu})'(D\boldsymbol{\theta}_{\tau,2}\Psi\boldsymbol{\theta}_{\tau,2}'D')^{-1}(\mathbf{y}-\boldsymbol{\mu})$, $\tilde{d} = \boldsymbol{\theta}_{\tau,1}'\boldsymbol{\theta}_{\tau,2}\Psi\boldsymbol{\theta}_{\tau,2}'\boldsymbol{\theta}_{\tau,1}$ and $K_{\nu}(\cdot)$ denotes the modified Bessel function of the third kind with index parameter $\nu = (2-n)/2$. The MAL distribution is closely related to multivariate quantile regression models, as stated in Proposition 1 from Petrella and Raponi (2019), which we report using our notation.

Proposition 1 (Petrella and Raponi (2019)). *Let* $\mathbf{y} \sim \mathrm{MAL}_n(\boldsymbol{\mu}, D\boldsymbol{\theta}_{\tau,1}, D\boldsymbol{\theta}_{\tau,2}\Psi\boldsymbol{\theta}_{\tau,2}'D')$ *and let* $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_n)'$ *be a fixed n-dimensional vector, such that* $\tau_i \in (0,1)$ *for* $i =$

$1, \ldots, n$. *Then* $\mathbb{P}(y_i \leq \mu_i) = \tau_i$ *if and only if*

$$\theta_{\tau,1,i} = \frac{1 - 2\tau_i}{\tau_i(1 - \tau_i)}, \qquad \theta_{\tau,2,i}^2 = \frac{2}{\tau_i(1 - \tau_i)}.$$

Moreover, $y_i \sim \mathcal{AL}(\mu_i, \Sigma_{ii}^{1/2}, \tau_i)$ *follows a univariate asymmetric Laplace distribution (see the Supplement).*

Modern models for investigating macroeconomic and financial time series are inherently multivariate, which calls for developing suitable multivariate quantile regression models. Chavleishvili and Manganelli (2021) provide a structural quantile VAR (QVAR) model to capture nonlinear relationships among macroeconomic variables and propose a quantile impulse response function to perform stress tests. Both the proposed MF-QVAR and the QVAR of Chavleishvili and Manganelli (2021) allow for persistence in all the quantiles. In fact, both specifications include lags of all the endogenous variables in the conditional quantile for every response. To clarify this point, we remark that maximum likelihood and Bayesian inference based on the asymmetric Laplace distribution (multivariate or univariate) with the constraints defined in Petrella and Raponi (2019) yield the same coefficient estimate as the minimisation of the check-loss function. The former approach is adopted by the MF-QVAR model, whereas the latter is followed by Chavleishvili and Manganelli (2021). In short, the location parameter of the constrained asymmetric Laplace distribution corresponds to the conditional quantile of interest.

Most of the models commonly used to investigate macroeconomic data focus on modelling the conditional mean of some indicators of interest. In a recent contribution to this literature, Antolin-Diaz et al. (2021) propose a dynamic factor model with stochastic volatility and outliers for studying the growth rates of macroeconomic variables. Similar to our empirical application, they deal with mixed-frequency data combining quarterly and monthly variables. The key difference between Antolin-Diaz et al. (2021) and the proposed MF-QVAR model is that the latter is a multivariate quantile regres-

sion framework, whereas the former is a conditional mean model. Quantile regression is advantageous over mean models since it allows to investigate the (possibly) heterogeneous impact of covariates on the entire conditional distribution of the response instead of restricting only to the effects on the mean.

Montes-Rojas (2019) develop a reduced form QVAR model to provide reliable forecasts and define a different quantile impulse response function to explore dynamic heterogeneity of the response variables to exogenous shocks. Recently, Adams et al. (2021) used quantile regressions to characterise upside and downside risks around the survey of professional forecasters' median consensus forecasts for each indicator.

The above-mentioned studies adopt a frequentist perspective. In contrast, Bernardi et al. (2015) developed a Bayesian inference for univariate quantile regression models to measure tail risk interdependence using Tobias and Brunnermeier (2016)'s Conditional VaR (CoVaR) indicator, defined as a quantile of a conditional distribution calculated at a given quantile of its conditioning distribution. Despite the increasing interest by policymakers in understanding and forecasting the whole distribution of economic and financial indicators, the literature on quantile VAR models is scant. We aim to fill this gap by proposing a novel fast Bayesian approach to inference for quantile VAR models.

Following the literature on multivariate time series, we define the n_β-dimensional vector $\boldsymbol{\beta} = (\mathbf{b}_0', \mathrm{vec}(B_1)', \ldots, \mathrm{vec}(B_p)')'$, with $n_\beta = n(1 + np)$, and the $n \times n_\beta$-dimensional matrix $X_t = (\mathbf{1}_n, \mathbf{x}_{t,1}', \ldots, \mathbf{x}_{t,p}')$, with $\mathbf{x}_{t,j} = (\mathbf{y}_{t-j} \otimes I_n)$ for each $j = 1, \ldots, p$. Moreover, let us reparametrize the innovation scale by introducing the positive definite matrix $\Sigma = D\Psi D' \in \mathbb{S}_{++}^n$, and relabelling with $D = D(\Sigma) = \mathrm{diag}\left(\Sigma_{11}^{1/2}, \ldots, \Sigma_{nn}^{1/2}\right)$. Owing to the properties of the multivariate asymmetric Laplace distribution, eq. (1) admits a representation as a location-scale mixture of Gaussian distributions (Petrella and Raponi, 2019; Kotz et al., 2001), as follows:

$$\mathbf{y}_t = \mathbf{b}_0 + \sum_{j=1}^{p} B_j \mathbf{y}_{t-j} + D\boldsymbol{\theta}_{\tau,1} w_t + \sqrt{w_t} D\boldsymbol{\theta}_{\tau,2} \Psi^{1/2} \tilde{\mathbf{z}}_t, \qquad \tilde{\mathbf{z}}_t \sim \mathcal{N}_n(\mathbf{0}_n, I_n), \qquad (2)$$

$$= X_t \boldsymbol{\beta} + D(\Sigma) \boldsymbol{\theta}_{\tau,1} w_t + \mathbf{z}_t, \qquad \mathbf{z}_t \sim \mathcal{N}_n(\mathbf{0}_n, w_t \boldsymbol{\theta}_{\tau,2} \Sigma \boldsymbol{\theta}_{\tau,2}), \tag{3}$$

where w_t is an auxiliary variable satisfying[1] $w_t \overset{i.i.d.}{\sim} \mathcal{E}xp(1)$, and define $\mathbf{w} = (w_{p+1}, \ldots, w_T)'$ the vector of latent variables.[2]

Let $\tau = \{\tau_1, \ldots, \tau_n\}$ be quantile levels. Then we denote the vector of quantiles of the marginal distributions, that is, $\mathbb{P}\left(\mathbf{y}_{i,t} \leq Q_{i,t}^{\tau_i} \mid \mathcal{F}_t\right) = \tau_i$, where $\mathcal{F}_t$ is the natural filtration of the whole state space in the highest frequency domain, as follows:

$$Q_t^\tau = \mathbf{b}_0 + \sum_{j=1}^{p} B_j \mathbf{y}_{t-j}.$$

Traditional multivariate time series models, such as the VAR, typically yield distorted results in the presence of outliers, which is caused by the fact that they are modelling the conditional mean of the response vector. Conversely, quantile regression methods, including the proposed MF-QVAR, model the conditional quantiles of the response variable. Therefore, extreme observations are treated as related to the tail behaviour of the response but do not contaminate the estimation of all the other quantiles.

Our MF-QVAR model offers a unique approach compared to the standard MIDAS method. Specifically, we investigate the conditional quantiles of *all* variables (both high- and low-frequency), denoted $\mathbf{y}_t$, in the highest frequency domain, indexed by t. In contrast, the MIDAS approach focuses on conditional quantiles of the low-frequency variables in their respective frequency domain, where t is low-frequency and $\mathbf{y}_t$ only contains low-frequency variables. Therefore, our proposed approach allows us to better capture the dynamics and relationships among variables in the high-frequency domain, which can be especially important for understanding rapid financial market fluctuations and other economic events. By employing the multivariate asymmetric Laplace dis-

---

[1]We use the rate parametrization, such that if $x \sim \mathcal{E}xp(1)$, then $kx \sim \mathcal{E}xp(1/k)$, for $k > 0$.

[2]The variable w_t shares a similar resemblance to the common stochastic volatility (CSV) specification proposed within a large BVAR framework by Chan (2020). As depicted in Figure S.15 of the supplementary material, the posterior estimations of w_t exhibit strikingly similar dynamics to the CSV derived from the large BVAR framework proposed by Chan (2020).

tribution and a Bayesian framework, we can estimate the conditional quantiles of the high-frequency variables and employ them to nowcast the low-frequency ones.

2.3 Mixed-frequency and Inter-Temporal Constraints

Mixed-frequency VAR (MF-VAR) models in macroeconomics and forecasting have become increasingly popular for producing high-frequency nowcasts of low-frequency variables. Specifically, MF-VARs are often used to joint model quarterly macroeconomic variables (low-frequency), such as gross domestic product (GDP), and monthly financial variables (high-frequency), such as surveys, to produce monthly nowcasts of GDP (e.g., see Schorfheide and Song, 2015).

We contribute to this literature by introducing mixed-frequency components in quantile VAR models that allow us to nowcast the conditional mean and the distribution of the low-frequency variables of interest. This is of paramount importance to timely understand the status of the economic system by explaining and nowcasting fundamental indicators, such as GDP and systemic risk indices. Moreover, it allows the researchers to consider any ragged-edge issues arising from the data release calendar.

A stacked or a state-space approach can be used to handle the mixed-frequency variables. The first class includes the MIDAS models, initially proposed by Ghysels et al. (2005) and Ghysels (2016), which consist of a linear model for the lowest observed frequency that includes the high-frequency covariates using particular functional forms for the coefficients. This framework has recently been extended to account for multiple regimes, dynamic panels, and high-dimensional settings (e.g., see Casarin et al., 2018; Mogliani and Simoni, 2021). Conversely, the state-space approach proposed by Schorfheide and Song (2015) treats the high-frequency observations of the low-frequency variables as missing values and estimates them via Kalman filtering and smoothing algorithms. This method has been applied to investigate consumption growth and long-run risks (Schorfheide et al., 2018), regional output growth (Koop et al., 2020), and high-

dimensional macroeconomic systems (Berger et al., 2023).

In particular, we adopt the state-space approach and model all variables at the highest observed frequency to obtain the interpolated estimates of the low-frequency variables at a higher frequency. The main drawback of this approach is the significant computational burden, mainly due to the estimation of high-dimensional latent state vectors (i.e., missing observations of the low-frequency variables) via filtering and smoothing techniques. This cost becomes prohibitive as the dimension of the VAR gets large, thus representing a major obstacle to using state-space methods on datasets with medium-high dimensions.

To investigate the joint distribution of the unobserved variables, conditional on the observations, let us denote with $\mathbf{Y} = (\mathbf{y}'_{p+1}, \ldots, \mathbf{y}'_T)'$ and $\boldsymbol{\zeta} = (\mathbf{z}'_{p+1}, \ldots, \mathbf{z}'_T)'$ the $(T - p)n$-dimensional vectors obtained by stacking all observations and all innovations over time, respectively. It is now possible to rewrite eq. (2) in matrix form as:

$$\mathbf{B}\mathbf{Y} = \mathbf{b} + \boldsymbol{\zeta}, \qquad \boldsymbol{\zeta} \sim \mathcal{N}(\mathbf{0}_{(T-p)n}, \boldsymbol{\Sigma}), \tag{4}$$

where

$$\mathbf{b} = (\mathbf{1}_{T-p} \otimes \mathbf{b}_0) + \mathbf{w} \otimes \left(\frac{(1 - 2\tau_1)\Sigma_{11}^{1/2}}{\tau_1(1 - \tau_1)}, \ldots, \frac{(1 - 2\tau_n)\Sigma_{nn}^{1/2}}{\tau_n(1 - \tau_n)} \right)', \qquad \boldsymbol{\Sigma} = \operatorname{diag}(\mathbf{w}) \otimes \boldsymbol{\theta}_{\tau,2}\Sigma\boldsymbol{\theta}_{\tau,2},$$

and

$$\mathbf{B} = \begin{bmatrix} -B_1 & -B_2 & \ldots & -B_p & I_n & \mathbf{0}_n & \ldots & \mathbf{0}_n & \mathbf{0}_n & \mathbf{0}_n \\ \mathbf{0}_n & -B_1 & \ldots & -B_{p-1} & B_p & I_n & \ldots & \mathbf{0}_n & \mathbf{0}_n & \mathbf{0}_n \\ \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots \\ \mathbf{0}_n & \mathbf{0}_n & \ldots & \mathbf{0}_n & \mathbf{0}_n & \mathbf{0}_n & \ldots & -B_p & I_n & \mathbf{0}_n \\ \mathbf{0}_n & \mathbf{0}_n & \ldots & \mathbf{0}_n & \mathbf{0}_n & \mathbf{0}_n & \mathbf{0}_n & -B_{p-1} & -B_p & I_n \end{bmatrix}$$

as a banded matrix of dimensionality $(Tn \times (T - p)n)$. Notice that $\mathbf{b}$, $\boldsymbol{\Sigma}$, and $\mathbf{B}$ depend

on the model parameters; we drop the notation for exposition purposes.

As stated in Schorfheide and Song (2015), we formulate the mixed-frequency VAR in a state-space structure. To this aim, let us first denote with $\mathbf{y}^o = (\mathbf{y}_1^{o'}, \ldots, \mathbf{y}_T^{o'})'$ and $\mathbf{y}^u = (\mathbf{y}_1^{u'}, \ldots, \mathbf{y}_T^{u'})'$ the Tn_o- and Tn_u-dimensional stacked vectors of observed and unobserved response variables. This allows to represent the vector $\mathbf{y}$ as a linear combination of $\mathbf{y}^o$ and $\mathbf{y}^u$, as follows:

$$\mathbf{y} = M_u \mathbf{y}^u + M_o \mathbf{y}^o, \tag{5}$$

where M_o and M_u are $(Tn \times Tn_o)$ and $(Tn \times Tn_u)$ selection matrices with full column rank. Then, by substituting eq. (5) into eq. (4) along the lines of Chan et al. (2023), one obtains the joint distribution of the missing observations, conditional on the observed data and model parameters as the Gaussian distribution:

$$\mathbf{y}^u | \mathbf{y}^o, \boldsymbol{\beta}, \Sigma, \mathbf{w} \sim \mathcal{N}(\boldsymbol{\mu}_y, K_y^{-1}), \tag{6}$$

where $K_y = M_u' \mathbf{B}' \boldsymbol{\Sigma}^{-1} \mathbf{B} M_u$ and $\boldsymbol{\mu}_y = K_y^{-1} \left(M_u' \mathbf{B}' \boldsymbol{\Sigma}^{-1} \left(\mathbf{b} - \mathbf{B} M_o \mathbf{y}^o \right) \right)$.

In practice, the unobserved data points in $\mathbf{y}^u$ are constrained to match the value of the low-frequency variables at those points in time when the latter are observed. One of the most commonly used restrictions for log-differenced variables is the log-linear approximation of Mariano and Murasawa (2003). This approach assumes that the observed quarterly value of the i-th variable at month t, denoted by $\tilde{y}_{i,t}^u$, is obtained as a linear combination of the missing monthly values at the current and previous four months, denoted by $y_{i,t}^u, \ldots, y_{i,t-4}^u$, as follows:

$$\tilde{y}_{i,t}^u = \frac{1}{3} y_{i,t}^u + \frac{2}{3} y_{i,t-1}^u + y_{i,t-2}^u + \frac{2}{3} y_{i,t-3}^u + \frac{1}{3} y_{i,t-4}^u. \tag{7}$$

This is a log-linear approximation to an arithmetic average of the quarterly variable, where note that $\tilde{y}_{i,t}^u$ is only observed for every third month. Stacking the inter-temporal

14

constraints over time, one gets

$$\tilde{\mathbf{y}}^u = M_a \mathbf{y}^u, \tag{8}$$

where M_a is a $(k \times Tn_u)$ matrix containing the k linear restrictions, and $\tilde{\mathbf{y}}^u$ is a vector containing the observed values of the low-frequency variables. To account for the inter-temporal constraints when sampling the unobserved variables, $\mathbf{y}^u$, it is sufficient to draw from the Gaussian distribution in eq. (6) subject to the restrictions in eq. (8). This is efficiently done following the methods described in Algorithm 2 of Cong et al. (2017), which postulates first to draw a vector from the unconstrained distribution, $\mathbf{u} \sim \mathcal{N}(\boldsymbol{\mu}_y, K_y^{-1})$, then compute

$$\mathbf{y}^u = \mathbf{u} + K_y^{-1} M_a' (M_a K_y^{-1} M_a')^{-1} (\tilde{\mathbf{y}}^u - M_a \mathbf{u}). \tag{9}$$

From a computational perspective, we follow the efficient implementation in Algorithm 1 of Chan et al. (2023).

We emphasise that while this is conceptually similar to the standard inter-temporal constraints of mean regression models, it is practically different. Similar to the mean regression models presented in Chan et al. (2023) under the state-space formulation, the variables themselves are subject to constraints. Specifically, the monthly GDP values are linearly combined to obtain the quarterly GDP observations, which are often used for performance evaluation. However, it is important to note that a linear combination of monthly quantiles does not necessarily equal the quarterly quantile:

$$Q^\tau_{\tilde{y}_{i,t}} \neq \frac{1}{3} Q^\tau_{y_{i,t}} + \frac{2}{3} Q^\tau_{y_{i,t-1}} + Q^\tau_{y_{i,t-2}} + \frac{2}{3} Q^\tau_{y_{i,t-3}} + \frac{1}{3} Q^\tau_{y_{i,t-4}},$$

where we made a slight abuse of notation by denoting with $Q^\tau_{y_{i,t}}$ the τth quantile of $y_{i,t}$, the ith response variable at time t. The primary objective of this article is to offer prompt estimates of monthly quantiles for quarterly variables, that is, the quantile estimates for the right-hand side equation above. If the aim is to acquire estimates of

quarterly GDP quantiles, a quantile-MIDAS approach can be used in place of the one proposed in this article.

Remark 1 (Aggregation). In this article, our primary objective is to estimate the high-frequency quantiles of a lower-frequency variable, that is $Q_{i,t}^\tau = \inf\{x \in \mathcal{X} : \mathbb{P}(Y_{i,t} \leq x|\mathcal{F}_t) = \tau\}$, where $Y_{i,t}$ is not observed, which makes a direct evaluation of nowcasts or forecasts from the proposed MF-QVAR model impossible. We illustrate this problem via

$$\tilde{Q}_{i,n_t}^\tau = \inf\left\{x \in \mathcal{X} : \mathbb{P}\left(Y_{i,n_t} = \sum_{j=1}^k w_j Y_{i,t-j+1} \leq x\Big|\mathcal{F}_t\right) = \tau\right\}, \tag{10}$$

where $\{w_j\}_{j=1}^k$ is a sequence of weights stemming from the time constraints that maps the unobserved high-frequency observations $Y_{i,t}$ into the low-frequency observations, Y_{i,n_t}, and n_t denotes the lower frequency count at the time measured in high frequency. Under a standard conditional mean framework, such as the MF-VAR model employed by Schorfheide and Song (2015), nowcasts or forecasts can be evaluated by aggregating the (unobserved) missing high-frequency estimates against the observed low-frequency value. However, in our proposed quantile framework, a direct aggregation across the estimated high-frequency quantiles of a lower-frequency variable cannot be made since:

$$\tilde{Q}_{i,n_t}^\tau \neq \sum_{\ell=1}^k w_\ell Q_{i,t-\ell+1}^\tau, \tag{11}$$

that is, the linear combination of the quantiles is not equal to the quantile of the random variable constructed by the same temporal aggregation as in Eq. (7).

Remark 2 (MIDAS QR vs MF-QVAR). In estimating a conditional quantile of quarterly real GDP, the quantile MIDAS approach is concerned with the conditional distribution:

$$P(y_{2,t}^q|y_{1,t-1}, y_{1,t-2}, y_{2,t-3}^q), \tag{12}$$

where $y_{1,t}$ is the vector of (high-frequency) observed monthly variables, and $y_{2,t}^q$ denotes

16

the (low-frequency) quarterly real GDP that is only observed for $t \in \mathcal{T}_q \subset \{1, \ldots, T\}$.

Instead, we consider a state-space formulation and assume an unobserved monthly evolution of the GDP, such that each quarterly observation is a linear combination of the (unobserved) monthly real GDP:

$$y_{2,t}^q = \sum_{j=1}^{k} w_j y_{2,t-j+1}. \tag{13}$$

In this case, differently from the MIDAS approach, we work with the conditional distribution:

$$P(y_{2,t}|y_{1,t-1}, y_{2,t-1}), \tag{14}$$

where $y_{2,t}$ is never observed and satisfies the set of constraints set in Eq. (13).

It is important to acknowledge that Eq. (12) and (14) represent two distinct distributions. In fact, Eq. (12) is the distribution of an observed variable, which allows direct evaluation of fitness and forecast performances in practical scenarios. Conversely, the distribution in Eq. (14) is related to an unobserved variable, making it challenging to assess. Therefore, considering the practical and theoretical challenges outlined above, we are unable to directly evaluate the nowcasts and forecasts generated by our proposed framework in the out-of-sample forecasting exercise.

3 Bayesian inference

In this section, we provide the details of the estimation of our proposed MF-QVAR model. Initially, we exploit the location-scale mixture representation of the multivariate asymmetric Laplace in eq. (1) and introduce a set of auxiliary variables $w_t \overset{i.i.d.}{\sim} \mathcal{E}xp(1)$, thus the complete-data likelihood is given by:

$$L((\mathbf{y}_1^o, \mathbf{y}_1^u), \ldots, (\mathbf{y}_T^o, \mathbf{y}_T^u), \mathbf{w}|\boldsymbol{\beta}, \Sigma) = \prod_{t=p+1}^{T} P(\mathbf{y}_t|\boldsymbol{\beta}, \Sigma, w_t)P(w_t|\boldsymbol{\beta}, \Sigma)$$

$$= \prod_{t=p+1}^{T} \exp \left\{ -\frac{1}{2} \Big[\big(\mathbf{y}_t - X_t \boldsymbol{\beta} - D(\Sigma)\boldsymbol{\theta}_{\tau,1}\big)' \big(w_t \boldsymbol{\theta}_{\tau,2} \Sigma \boldsymbol{\theta}_{\tau,2}\big)^{-1} \big(\mathbf{y}_t - X_t \boldsymbol{\beta} - D(\Sigma)\boldsymbol{\theta}_{\tau,1}\big) \Big] \right\}$$

$$\times (2\pi)^{-\frac{n}{2}} \left| w_t \boldsymbol{\theta}_{\tau,2} \Sigma \boldsymbol{\theta}_{\tau,2} \right|^{-\frac{n}{2}} \exp\{-w_t\}. \tag{15}$$

Before describing the posterior distributions along with the algorithm used, we define the prior specifications for our parameters. Starting with the coefficient vector $\boldsymbol{\beta} = (\mathbf{b}_0', \mathrm{vec}(B_1, \ldots, B_p)')'$, we assume a conjugate multivariate Gaussian prior distribution

$$\boldsymbol{\beta} \sim \mathcal{N}_{n_\beta}(\underline{\boldsymbol{\mu}}_\beta, \underline{\Omega}_\beta).$$

For this vector of coefficients, one may consider using shrinkage priors such as the global-local shrinkage prior (Polson and Scott, 2010), and the Minnesota prior (Kadiyala and Karlsson, 1997). In this article, we focus on a simple case by setting the prior mean of the coefficient associated with each equation at the frequentist univariate regression estimate, $\underline{\boldsymbol{\mu}}_\beta = \widehat{\boldsymbol{\beta}}$, and a prior variance $\underline{\Omega}_\beta = 100 \cdot I_{n_\beta}$, which results in a relatively flat prior distribution. We leave for further research on the use of more complex shrinkage priors.

The other parameter of interest is the scale matrix, $\Sigma \in \mathbb{S}_{++}^n$, and in this scenario, we assume an inverse Wishart prior distribution

$$\Sigma \sim \mathcal{IW}_n(\underline{\nu}_0, \ \underline{\Phi}_0),$$

where $\underline{\nu}_0 > n - 1$ is the degrees of freedom parameter and $\underline{\Phi}_0 \in \mathbb{S}_{++}^n$ is a scale matrix, such that if $\underline{\nu}_0 > n + 1$ then $\mathbb{E}[\Sigma] = \underline{\Phi}_0 / (\underline{\nu}_0 - n - 1)$. This is equivalent to assuming the Wishart prior distribution for $\Sigma^{-1} \sim \mathcal{W}_n(\underline{\nu}_0, \ \underline{\Phi}_0^{-1})$, where $\mathbb{E}[\Sigma^{-1}] = \underline{\Phi}_0^{-1} \underline{\nu}_0$.

Based on these prior specifications and the likelihood function in eq. (15), we can provide the full conditional distributions for each parameter and latent variable of the model. We remark that our parametrization differs from Tian et al. (2016), as we work with the positive definite matrix Σ, instead of the correlation matrix Ψ and the diagonal matrix D separately. Moreover, we work with multivariate Bayesian analysis of quantile

regression models, while Yu and Moyeed (2001) proposed a Bayesian approach for the univariate framework.

As the joint posterior distribution is not tractable, we rely on data augmentation to obtain closed-form full conditional distributions and to design a Markov chain Monte Carlo (MCMC) algorithm for approximating the posterior distribution. Specifically, we design an efficient Gibbs sampler based on the precision sampler of Chan and Jeliazkov (2009) and Chan et al. (2023), which cycles over the following steps:

1. draw $\mathbf{y}^u$ given $\mathbf{y}^o, \boldsymbol{\beta}, \mathbf{w}$ and Σ from eq. (6) subject to the restrictions in eq. (8) using Algorithm 1 of Chan et al. (2023);

2. draw $\boldsymbol{\beta}$ given $\mathbf{y}^o, \mathbf{y}^u, \mathbf{w}$ and Σ from the Gaussian distribution $\mathcal{N}_{n_\beta}(\overline{\boldsymbol{\mu}}_b, \overline{\Omega}_b)$, with $\tilde{\mathbf{e}}_t = \mathbf{y}_t - D\boldsymbol{\theta}_{\tau,1} w_t$ and parameters

$$\overline{\Omega}_b = \left(\underline{\Omega}_b^{-1} + \sum_{t=p+1}^{T} X_t'(w_t \boldsymbol{\theta}_{\tau,2} \Sigma \boldsymbol{\theta}_{\tau,2})^{-1} X_t\right)^{-1}, \quad \overline{\boldsymbol{\mu}}_b = \overline{\Omega}_b^{-1}\left(\underline{\Omega}_b^{-1}\underline{\boldsymbol{\mu}}_b + \sum_{t=p+1}^{T} \tilde{\mathbf{e}}_t'(w_t \boldsymbol{\theta}_{\tau,2} \Sigma \boldsymbol{\theta}_{\tau,2})^{-1} X_t\right).$$

3. draw w_t, for each $t = p+1, \ldots, T$, given $\mathbf{y}^o, \mathbf{y}^u, \boldsymbol{\beta}$ and Σ from the Generalised inverse Gaussian distribution $\text{GiG}(\overline{p}_w, \overline{a}_w, \overline{b}_{w,t})$, with $\mathbf{u}_t = \mathbf{y}_t - X_t \boldsymbol{\beta}$, and

$$\overline{p}_w = 1 - \frac{n}{2}, \qquad \overline{a}_w = 2 + \boldsymbol{\theta}_{\tau,1}' D(\boldsymbol{\theta}_{\tau,2} \Sigma \boldsymbol{\theta}_{\tau,2})^{-1} D\boldsymbol{\theta}_{\tau,1}, \qquad \overline{b}_{w,t} = \mathbf{u}_t'(\boldsymbol{\theta}_{\tau,2} \Sigma \boldsymbol{\theta}_{\tau,2})^{-1}\mathbf{u}_t.$$

4. draw Σ given $\mathbf{y}^o, \mathbf{y}^u, \boldsymbol{\beta}$ and $\mathbf{w}$ via the slice sampler of Neal (2003). Defining $\mathbf{e}_t(\Sigma) = \mathbf{y}_t - X_t \boldsymbol{\beta} - D(\Sigma)\boldsymbol{\theta}_{\tau,1} w_t$, the target density function is proportional to:

$$\propto |\Sigma|^{-\frac{\nu_0+n+1}{2}} \exp\left\{-\frac{1}{2}\text{tr}\left(\underline{\Phi}_0 \Sigma^{-1}\right)\right\} |\Sigma|^{-\frac{T}{2}} \exp\left\{-\frac{1}{2}\sum_{t=1}^{T} \mathbf{e}_t'(\Sigma)\left(w_t^{-1}\boldsymbol{\theta}_{\tau,2}^{-1}\Sigma^{-1}\boldsymbol{\theta}_{\tau,2}^{-1}\right)\mathbf{e}_t(\Sigma)\right\}.$$

The Supplement provides a detailed description of the MCMC algorithm along with the derivation of the full conditional distributions.

4 Simulation Study

To illustrate the estimation and forecasting accuracy of our proposed framework, we undertake a simulation study across two data-generating processes (DGP) based on the following bivariate process:

$$
\begin{aligned}
Y_{1,t} &= X_t\boldsymbol{\beta}_1 + \exp(X_t\boldsymbol{\gamma}_1)Z_{1,t}, \\
Y_{2,t} &= X_t\boldsymbol{\beta}_2 + \exp(X_t\boldsymbol{\gamma}_2)Z_{2,t}, \qquad t = 1,\ldots,T,
\end{aligned}
\tag{16}
$$

where $X_t = [Y_{1,t-1}, Y_{2,t-1}]$ and $Z_{i,t} \sim \mathcal{N}(0,1)$, for $i = 1, 2$. Given the marginal DGPs in Eq. (16), we model the dependence between the series using:

1. a Gaussian copula, with correlation coefficient $\rho = 0.5$ and variance $\sigma^2 = 1$;
2. a Student-t copula with correlation coefficient, degree of freedom parameter and variance equal to $\rho = 0.5$, $\nu = 5$ and $\sigma^2 = 1$, respectively.

It is worth noting that the Gaussian copula assumes a symmetric dependence with constant strength across the entire range, whereas the t-copula allows for heavy-tailed and asymmetric dependence, accommodating a wider range of dependence patterns. For each model $i = 1, 2$ and variable $j = 1, \ldots, n$, with $n = 2$ in this case, the conditional τ-quantiles are constructed as:

$$
Q_{j,t}^{\tau,(i)} = X_t\boldsymbol{\beta}_j + \exp(X_t\boldsymbol{\gamma}_j)F_i^{-1}(\tau),
\tag{17}
$$

where F_i is the CDF under model i. Finally, we evaluate the performance of each model in terms of the estimated conditional quantiles against their corresponding true values

using the in-sample and out-of-sample mean absolute error (MAE) metrics:[3]

$$\text{MAE}_{in}(i) = \frac{1}{Tn} \sum_{j=1}^{n} \sum_{t=1}^{T} \left| Q_{j,t}^{\tau,(i)} - \widehat{Q}_{j,t}^{\tau,(i)} \right|, \quad \text{MAE}_{out}(i) = \frac{1}{n} \sum_{j=1}^{n} \left| Q_{j,T+1}^{\tau,(i)} - \widehat{Q}_{j,T+1}^{\tau,(i)} \right|, \quad (18)$$

where $Q_{j,t}^{\tau,(i)}$ is the true DGP simulated conditional quantiles by Eq. (17) and $\widehat{Q}_{j,t}^{\tau,(i)}$ is the estimated conditional quantile of the jth variable from model i. We set $T = 600$ and consider $\tau = \{0.1, 0.5, 0.9\}$.[4]

In the first experiment, we estimate univariate quantile regression and QVAR models based on a single observed frequency, that is, without missing observations. The goal of this experiment is to demonstrate the effectiveness of jointly estimating conditional quantiles, as opposed to individual estimation. Instead, the second experiment considers a mixed-frequency setting, where we estimate the conditional mean MF-VAR model of Schorfheide and Song (2015) and our proposed MF-QVAR model. As the primary focus of this article is on estimating the conditional quantiles of the unobserved low-frequency variables, we aim to compare the accuracy of our proposed framework against their corresponding mean counterparts. To estimate all the models, we employ the Bayesian approach and run 10 parallel MCMC chains.

In the first simulation experiment, the MAEs in Table 1 show a consistent pattern across both DGPs. The joint estimation of conditional quantiles consistently outperforms the individual estimation approach, highlighting the effectiveness of considering quantiles jointly rather than in isolation. This is because by jointly estimating the con-

---

[3]In this particular simulated experiment, our primary focus lies in assessing the precision of each model in generating a specific conditional quantile estimate, a central objective of this article. The evaluation and development of density forecasts utilizing our proposed methodology are beyond the scope of this article. If we were to proceed with constructing density forecasts using our proposed model, it would require estimation across a broad spectrum of conditional quantiles. Subsequently, we would have to apply either the skew-t approach as delineated by Adrian et al. (2019) or the non-parametric approach advocated by Mitchell et al. (2023) for constructing the density forecast across all these quantiles. However, it's noteworthy that these two approaches have been primarily employed in the context of univariate quantile regression. Extending these methodologies to a multivariate framework presents a formidable and intricate challenge, and its applicability is not guaranteed. Consequently, we defer this endeavor to a future avenue of research.

[4]The selected time periods and quantiles are chosen to reflect the empirical application undertaken in Section 5.

ditional quantiles, we can leverage the interdependencies and relationships that exist among several variables, capturing the complex dependence structure more accurately. When quantiles are estimated independently, the performance of the upper and lower tails deteriorates dramatically due to limited data. However, the proposed method by jointly estimating the quantiles demonstrates robustness across different correlation structures, effectively capturing potentially complex tail dependencies between the variable, leading to improved estimation accuracy and forecasting performance.

		10% Quantile		50% Quantile		90% Quantile	
	Model	Gaussian	Student-t	Gaussian	Student-t	Gaussian	Student-t
In-sample	Univariate BQR	1.26	1.48	0.13	0.14	1.25	1.46
	QVAR	**0.11**	**0.17**	**0.06**	**0.07**	**0.12**	**0.26**
Out-of-sample	Univariate BQR	1.25	1.48	0.09	0.11	1.24	1.49
	QVAR	**0.12**	**0.19**	**0.03**	**0.07**	**0.06**	**0.23**

Table 1: MAE simulation results for the observed frequency case. Notes: Univariate BQR denotes a univariate Bayesian quantile regression; QVAR denotes a quantile VAR with no missing values. Bold values denote the best model.

In the mixed-frequency case, we focus on estimating the high-frequency quantiles of the low-frequency variable, that is $Q_{2,t}^\tau$. The results in Table 2 provide valuable insights into several key aspects. First, we find evidence that missing values in the data can have an adverse impact on the estimation across all the modelling approaches, due to the induced information loss. Second, the Gaussian MF-QVAR model may struggle to effectively capture tail behaviour even when the Gaussian assumption is correctly specified and the DGP exhibits heterogeneous error variance. This points out the limitations of relying solely on the Gaussian assumption and suggests the need for alternative modelling approaches that can better accommodate the complexities of the data. Lastly, the proposed MF-QVAR demonstrates the capability to capture such complex dependencies, even when a variable is unobserved and only its temporal aggregation is available. This highlights the robustness and effectiveness of the proposed methodology in capturing and incorporating information from unobserved variables, leading to improved estimation accuracy and forecasting performance.

		10% Quantile		50% Quantile		90% Quantile	
	Model	Gaussian	Student-t	Gaussian	Student-t	Gaussian	Student-t
In-sample	MF-VAR	1.27	1.49	0.59	0.68	1.53	2.19
	MF-QVAR	**0.57**	**0.58**	**0.47**	**0.64**	**0.55**	**0.77**
Out-of-sample	MF-VAR	2.51	2.69	0.18	0.35	1.38	1.10
	MF-QVAR	**0.91**	**1.51**	**0.14**	**0.22**	**0.47**	**0.38**

Table 2: MAE simulation results for the mixed-frequency case. Notes: MF-VAR denotes the mixed-frequency state-space VAR of Schorfheide and Song (2015) and Chan et al. (2023); MF-QVAR denotes our proposed mixed-frequency state-space quantile VAR model. Bold values denote the best model.

Overall, the simulation results provide valuable insights into the challenges posed by missing values, the limitations of the Gaussian MF-QVAR model in capturing tail behaviour, and the strengths of the proposed approach in handling complex dependencies and unobserved variables. Furthermore, the Supplement presents the MAE simulation results for a nine-variables DGP encompassing both Gaussian and Student-t copulas. Once again, these outcomes reinforce the aforementioned conclusions drawn from the bivariate cases.

5 Empirical Application

In this empirical application, we illustrate our proposed MF-QVAR model via an in- and out-of-sample analysis. Specifically, we plot the monthly US growth-at-risk estimates from our proposed model against a quarterly frequency counterpart model for the in-sample analysis. Next, we undertake a real-time exercise to nowcast monthly US growth-at-risk for the out-of-sample analysis.

5.1 In-Sample Analysis: Monthly US GDP Quantile Estimates

We estimate our proposed MF-QVAR model on the final vintage date of March 2022, and the sample spans from January 1973 and December 2021. We estimate an MF-QVAR model consisting of the quarterly US real GDP and eight monthly variables,

which were gathered from the St. Louis ALFRED database.[5] Seven of the monthly variables included in our model are broadly similar to the monthly variables chosen in Schorfheide and Song (2015), whereas the last one is the NFCI. This choice is motivated by the recent study of Adrian et al. (2019), which showed that a tightening of the NFCI could lead to a large increase in the growth-at-risk for US real GDP.

Figure 1 plots the monthly posterior estimates of US GDP across three selected quantiles: 10, 50 and 90 per cent. All of these quantile estimates display broadly similar dynamics over time and exhibit concurrent fall in their estimates during most NBER recession dates. In addition, these quantile estimates are very volatile over time, and this is unsurprising since we are modelling monthly percentage changes here.
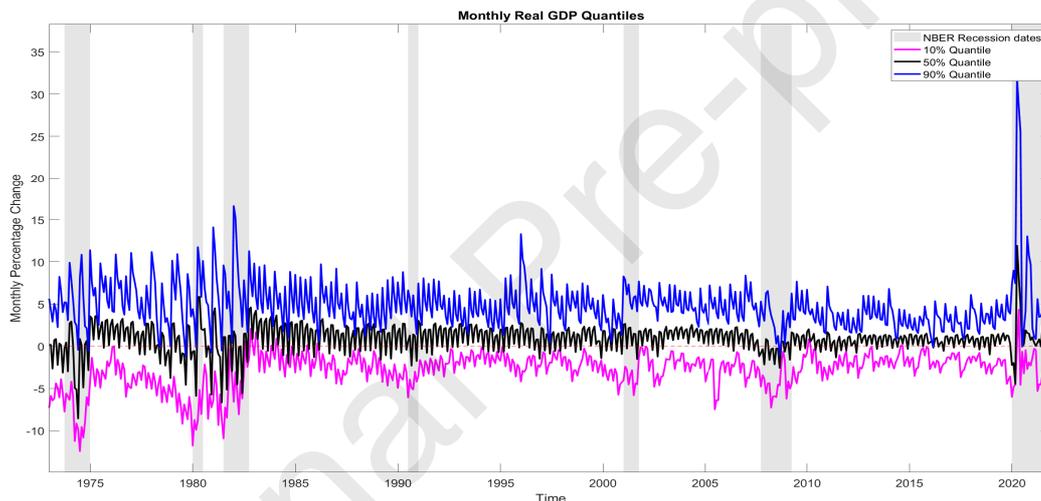
Figure 1: Posterior mean of the monthly US GDP for the 10 (purple), 50 (black) and 90 (blue) per cent quantiles. The shaded grey areas denote the NBER recession dates.

Table 3 showcases the posterior mean of selected coefficients in the MF-QVAR model applied to the GDP equation. Specifically, it elucidates the impacts of various economic indicators, including CPI inflation, the Fed funds rate, the unemployment rate, and NFCI.[6] All the coefficient estimates presented in Table 3 are statistically significant, as evidenced by the fact that the corresponding 68% credible intervals do not include zero. Our analysis reveals an observable negative association between real GDP and both

---

[5]See the Supplement for further details of each data variable and their respective transformations.
[6]See the Supplement for the complete set of posterior estimates of the MF-QVAR model.

CPI inflation and NFCI, indicating clear interdependencies. Importantly, the influence of CPI inflation is more pronounced within the lower range of GDP, while NFCI exerts a greater impact within the median quantile.

An intriguing finding emerges with regard to the coefficient of the Fed funds rate. Its sign changes from positive in the 10% quantile to negative in the 90% quantile, suggesting that a contractionary monetary policy significantly and adversely affects GDP. In contrast, the coefficient of the unemployment rate exhibits an opposite behaviour to that of the Fed funds rate. It demonstrates a negative effect on GDP within the lower and median quantiles, consistent with Okun's law, which states an inverse relationship between the unemployment rate and GDP. However, surprisingly, the unemployment rate positively impacts the upper quantile of GDP, thereby invalidating Okun's law in this context. Consequently, our findings provide evidence of asymmetric behaviour within the conditional distribution and underscore the notion that relationships established based on conditional means may not hold true in the distribution's tails.

Coefficient on the GDP equation	CPI Inflation	Federal Funds Rate	Unemployment Rate	NFCI
10% quantile	-3.09 (-3.19, -2.99)	0.35 (0.26, 0.45)	-0.87 (-0.97, -0.77)	-0.61 (-0.70, -0.51)
50% quantile	-0.31 (-0.40, -0.21)	0.47 (0.37, 0.56)	-1.41 (-1.50, -1.31)	-2.61 (-2.71, -2.51)
90% quantile	-1.26 (-1.36, -1.16)	-1.09 (-1.20, -0.99)	0.95 (0.86, 1.05)	-1.69 (-1.80, -1.58)

Table 3: Posterior mean estimates of selected VAR coefficients in the GDP equation of MF-QVAR, with the associated 68% credible intervals in brackets.

In Carriero et al. (2022), they similarly nowcasts US growth-at-risk using weekly information. However, they implement a MIDAS quantile regression framework where they regress US quarterly real GDP against weekly economic indicators. Consequently, they are still only able to nowcasts US growth-at-risk at the quarterly frequency. In contrast, our proposed MF-QVAR allows us to specifically nowcast the US growth-at-risk at the monthly frequency.
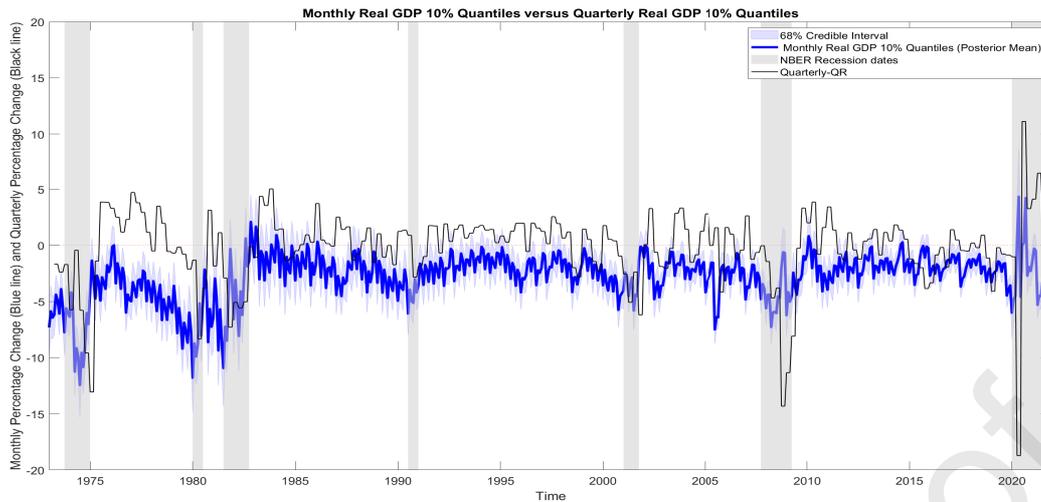
Figure 2: Monthly and quarterly US GDP growth-at-risk estimates. Posterior mean of monthly US GDP 10% quantile (growth-at-risk) estimates (blue line), and the corresponding 68% credible interval (shaded blue area). Quarterly US GDP 10% quantile estimates from a standard frequentist QR (black line). The grey shades denote the NBER recession dates.

Intuitively, we would expect the higher frequency monthly US growth-at-risk estimates to behave very differently from their quarterly counterpart. To illustrate these differences, we plot in Fig. 2 the monthly and quarterly in-sample US growth-at-risk estimates from our MF-QVAR and a quarterly frequency quantile regression, respectively. In regards to the quarterly quantile regression specification (denoted as the Quarterly-QR in Fig. 2), we regress the US real GDP against the lagged values of the other eight variables described in the dataset, all at the quarterly frequency. Essentially, the Quarterly-QR can be interpreted as the univariate version of the US GDP equation of our MF-QVAR. By inspecting Fig. 2, we notice that both the monthly and quarterly US growth-at-risk estimates display similar dynamics over time. However, these two frequencies exhibit substantial differences in their level estimates during the recessions. For instance, focusing on the Great Recession of 2007-08, our monthly US growth-at-risk estimates appear to be detecting the vulnerability of the US economy considerably faster than the quarterly quantile regression. In addition, the Quarterly-QR exhibits a significant negative growth-at-risk compared to our MF-QVAR estimate in both the Great Recession and the COVID-19 pandemic. Therefore, this result suggests that when modelling US growth-at-risk, the low-frequency estimate can display significantly

26

different dynamics compared to its high-frequency counterpart during a recession.[7]

5.2 Out-of-Sample Analysis: Nowcasting Monthly US GDP Growth-at-risk

We illustrate the utility of our proposed MF-QVAR model by undertaking a real-time nowcasting application for the growth-at-risk of US real GDP. To the best of our knowledge, this is the first study that explicitly nowcasts a monthly growth-at-risk estimate for US real GDP, whereas all the previous studies only modelled it at a quarterly frequency. A key advantage of our MF-QVAR model is that it *naturally* allows the forecaster to consider any ragged-edge issues arising from the data release calendar.

We conduct our real-time nowcasting analysis in two distinct episodes. Firstly, we concentrate on the GFC period, employing data vintage spanning from January 2005 to December 2010. Secondly, we shift our attention to the period encompassing the COVID-19 pandemic and the Russian invasion of Ukraine, utilising data vintage from January 2016 to March 2022. All vintage data for the monthly and quarterly series used in the out-of-sample analysis was collected from the St. Louis ALFRED database.[8]

We focus on generating the nowcasts and forecasts under three release timings of US real GDP by the US BEA. This approach is similar to the works of Giannone et al. (2008) and Bańbura et al. (2013). However, for the sake of simplicity and consistency with Schorfheide and Song (2015), we conduct a real-time exercise where we only generate nowcasts and forecasts for US growth-at-risk at the end of each month. In contrast, Giannone et al. (2008) and Bańbura et al. (2013) produce nowcasts of real GDP multiple times within each month. The objective of our real-time exercise is to analyse the variation in the nowcasts and forecasts of US growth-at-risk from month to month rather than focusing on their behaviour within each individual month.

---

[7]We have also estimated our MF-QVAR model using Minnesota prior specified in Carriero et al. (2022). The in-sample results similar to our benchmark are presented in the Supplement.

[8]The data vintage for NFCI is only available from May 2011 onwards in the US ALFRED database. Therefore, for the GFC period, we have to exclude NFCI from the nowcasting exercise.

Our nowcasting approach is demonstrated in Fig. 3, which represents the data release calendar. Let us begin by considering the end of January 2015, denoted as event **(a)** in Fig. 3. At this time, the December 2014 data for the monthly indicators are released, and the BEA provides the first estimate of Real GDP for the fourth quarter of 2014. In this particular scenario, both the monthly variables and quarterly US GDP exhibit a balanced data structure without a ragged edge at the end of the sample. Consequently, we generate a growth-at-risk forecast for the upcoming one to three months (January, February, and March 2015, as illustrated in our example). We categorised this as a *Forecast* in our out-of-sample nowcasting exercise. As we progress to the end of February 2015 (event **(b)** in Fig. 3), the January 2015 data for the monthly indicators and the BEA's second estimate of Real GDP for the fourth quarter of 2014 are released. At this point, there is a ragged edge at the end of the sample due to the unbalanced structure of both the monthly variables and quarterly US GDP. The growth-at-risk nowcast for the following one to three months, generated at this period, is designated as *Nowcast* $T + 1$. Finally, we reach the conclusion of March 2015 (event **(c)** in Fig. 3), denoting the release of the monthly indicators for February 2015 and the BEA's third estimate of Real GDP for the fourth quarter of 2014. Similarly, a ragged edge is observed at the end of the sample period. During this period, the growth-at-risk nowcast is denoted as *Nowcast* $T + 2$.

Subsequently, we maintain the iterative process throughout the ensuing three months within each quarter, ensuring consistent classification of both nowcasts and forecasts. Our main focus is on nowcasting the growth-at-risk at the 10th percentile, denoted as $\tau = 0.1$. To comprehensively explore the dynamics of the entire distribution of real GDP, we also generate nowcasts for the 50th ($\tau = 0.5$) and 90th ($\tau = 0.9$) percentiles.

5.2.1 Global Financial Crisis Period

We investigate the nowcasting behavior of the monthly growth-at-risk estimates during the GFC period. To obtain a quarterly series to compare with the low-frequency ob-

28
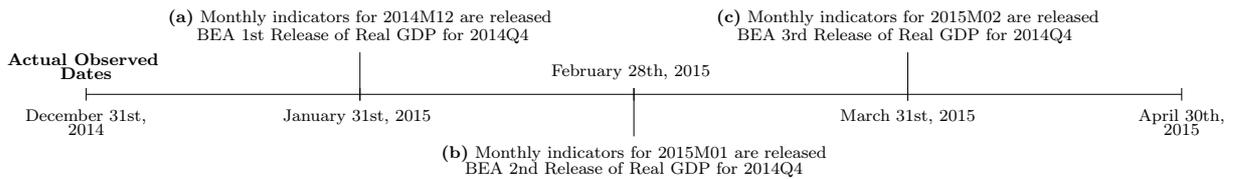
Figure 3: Data Release Calendar. We denote the event as **(a)** Forecast, **(b)** Nowcast $T + 1$ and **(c)** Nowcast $T + 2$.

served GDP, Fig. 4 plots the rolling three-month posterior mean average of the monthly changes of the growth-at-risk estimates for the three types of nowcasts. Moreover, for comparison purposes, in Fig. 4 we also plot the corresponding quarterly growth-at-risk nowcasts from a U-MIDAS quantile regression (QR) model. This U-MIDAS QR model extends the quarterly frequency QR model used in Adrian et al. (2019) and incorporates all the monthly indicators specified in our data table of the Supplement. The resulting quarterly growth-at-risk nowcasts are denoted as the dashed lines in Fig. 4.

The findings depicted in Fig. 4 indicate that the MF-QVAR model demonstrates a notable capability in detecting the vulnerability of the US economy at an earlier stage compared to the U-MIDAS QR model. The end of 2005 was characterised by a slowdown in housing prices and a rising delinquency rates in mortgage loans. During this period, our proposed model generates forecast and nowcast estimates of approximately -8%, whereas the quarterly counterpart model provides an estimate of -1%. A possible explanation for the discrepancy is that the proposed framework might better captures the large negative effects of the above-mentioned events on the economy.

29
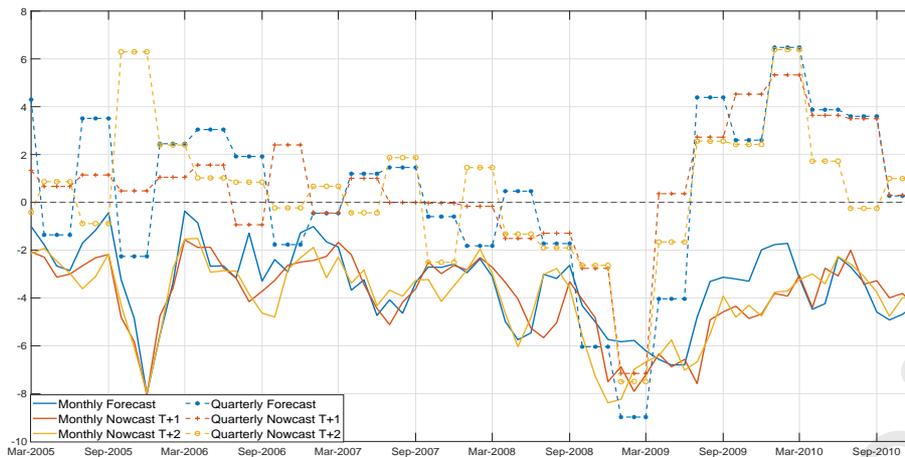
Figure 4: Posterior mean of the rolling three-month average growth-at-risk changes for the 10th percentile ($\tau = 0.1$), from MF-QVAR (monthly, solid), and U-MIDAS QR (quarterly, dashed). Forecasts of T in blue, nowcasts of $T+1$ in red, and nowcasts of $T+2$ in yellow.

During the period of intensified recession in late 2008 and early 2009, both the monthly and quarterly growth-at-risk estimates exhibit a significant level of similarity. However, the monthly growth-at-risk estimates demonstrate a tendency to identify the onset of recessionary phases earlier compared to their quarterly counterparts. Moreover, as the early 2010 period unfolded, the quarterly growth-at-risk estimates indicated positive values, suggesting a recovery of the US economy following the recession of late 2008 to early 2009. Conversely, the monthly growth-at-risk estimates showed negative values, indicating that the US economy remained vulnerable and distant from achieving a state of recovery. Additionally, we plot the posterior estimates of the monthly nowcast for $T+2$ at the 50th ($\tau = 0.5$) and 90th ($\tau = 0.9$) percentiles, as depicted in Fig. 5. This figure illustrates that the monthly nowcasting distribution of real GDP becomes increasingly negatively skewed towards the end of 2005 and during the heightened recessionary period of late 2008 and early 2009.[9] The persistence of negative growth-at-risk estimates during the early 2010 period aligns with the prevailing theme of weakened economic growth experienced by the US economy during that time.

---

[9] We find similar conclusions in the monthly forecast and nowcast of $T+1$, please see the Supplement.

Figure 5: Posterior mean of the rolling three-month average of the Monthly Nowcast $T + 2$ for the 10th ($\tau = 0.1$, blue), 50th ($\tau = 0.5$, red), and 90th percentile ($\tau = 0.9$, yellow).

5.2.2 COVID-19 Pandemic and Russian Invasion of Ukraine Period

In our final out-of-sample study, we focus on generating the forecast and nowcasts for the period, including the COVID-19 pandemic and the Russian invasion of Ukraine. Similarly, Fig. 6 displays the rolling three-month posterior mean average of the monthly growth-at-risk estimates' changes, specifically pertaining to the aforementioned period of interest. It is evident from Fig. 6 that the growth-at-risk nowcasts were, on average, about -3% during the pre-pandemic period, whereas this average has fallen to about -5% since the pandemic. This implies that COVID-19 has caused US monthly real GDP to be more skewed to the left and increased the vulnerability of the US to enter a recession. In contrast, considering the 50th and 90th percentile, the monthly nowcast between spring and summer 2020 evolved in the opposite direction compared to the quarterly predictions (see the Supplement). This further highlights the importance of the proposed nowcasting approach in providing a timely characterization of risks.

Based on the observations derived from Fig. 6, the majority of the quarterly growth-at-risk nowcasts exhibit positive values, with the exception of the initial year marked by

31
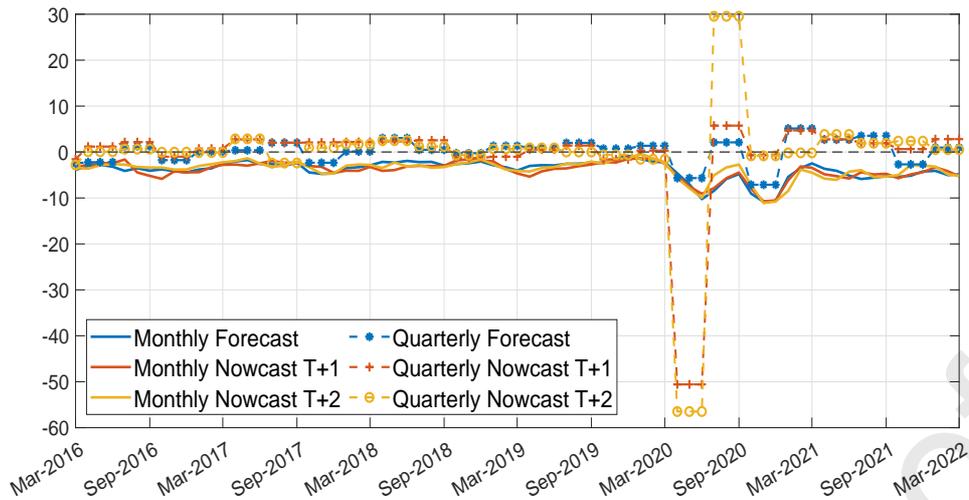
Figure 6: Posterior mean of the rolling three-month average growth-at-risk changes for the 10th percentile ($\tau = 0.1$), from MF-QVAR (monthly, solid), and U-MIDAS QR (quarterly, dashed). Forecasts of T in blue, nowcasts of $T + 1$ in red, and nowcasts of $T + 2$ in yellow.

the onset of the COVID-19 pandemic. In contrast, the monthly growth-at-risk nowcasts from the MF-QVAR model are all negative. These results suggest that the nowcasts from the U-MIDAS QR model may be underestimating the underlying growth-at-risk measure for US Real GDP. For instance, the growth-at-risk nowcasts from the U-MIDAS QR model bounced back to their pre-pandemic level at the end of 2020, which is inconsistent with recent global events. In fact, since the pandemic period, the US has experienced weaker growth and high inflation, and intuitively one would expect the US to be more prone to a recession than an expansion. Conversely, the nowcasting results from our MF-QVAR model are consistent with this idea. These differences between our MF-QVAR and U-MIDAS QR model are consistent with the findings in the in-sample analysis, where we found lower frequency growth-at-risk estimates can display different dynamics to their monthly counterparts during a recession.

To further investigate the skewness of the real GDP nowcasts, we also generated the nowcasts for the 50th ($\tau = 0.5$) and 90th ($\tau = 0.9$) percentiles. Figure 7 plots all the posterior estimates of the percentiles for the monthly nowcast of $T + 2$. The pattern displayed in Fig. 7 indeed confirms our previous finding that the COVID-19 pandemic has caused the real GDP nowcasts to become more negatively skewed. Similar

32

conclusions can be drawn from the monthly forecast and monthly nowcast of $T+1$ (see Supplement). Moreover, in Table 4, we report the posterior monthly nowcasts of all the percentiles for selected pandemic periods. For December 2019, the nowcast estimate for the 10th percentile was -1.44%, dropping to -5.67% in April 2020 of the first wave of the pandemic. In addition, the uncertainty associated with the real GDP nowcasts appears to have widened since the pandemic. For example, the distance between the nowcasts of the 10th and 90th percentile has significantly increased since December 2019, as reported in the fifth column in Table 4. This result suggests that nowcasting real GDP has become more challenging since the pandemic.

By inspecting the difference between nowcasts of the 50th-10th and the 90th-50th percentiles, we find evidence of an increased negative skewness during the pandemic. This is motivated by the shift of mass from the right tail (as proxied by $|Q_{0.9} - Q_{0.5}|$) to the left tail (as proxied by $|Q_{0.5} - Q_{0.1}|$), whose difference in the last column of Table 4 can be interpreted as a proxy for the skewness. In particular, the distance between the nowcasts of the 10th and 50th percentile during the pandemic outweighs the increased distance between the 50th and 90th percentile, thus providing evidence supporting our previous finding that the US monthly real GDP has become more negatively skewed during the pandemic (see last column in Table 4).
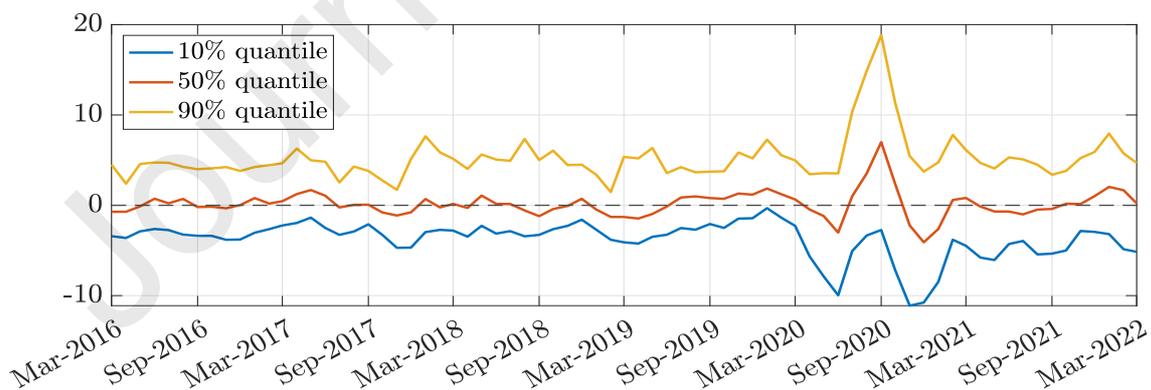
Figure 7: Posterior mean of the rolling three-month average of the Monthly Nowcast $T+2$ for the 10th ($\tau = 0.1$, blue), 50th ($\tau = 0.5$, red), and 90th percentile ($\tau = 0.9$, yellow).

				Percentile			
Dates	$Q_{0.1}$	$Q_{0.5}$	$Q_{0.9}$	$Q_{0.9} - Q_{0.1}$	$Q_{0.5} - Q_{0.1}$	$Q_{0.9} - Q_{0.5}$	$\begin{array}{c} \lvert Q_{0.9} - Q_{0.5}\rvert - \\ \lvert Q_{0.5} - Q_{0.1}\rvert \end{array}$
December 2019	-1.44	1.17	5.20	6.63	2.61	4.03	1.42
April 2020	-5.67	-0.46	3.44	9.11	5.21	3.90	-1.31
September 2020	-2.74	6.99	18.84	21.58	9.73	11.85	2.12
January 2021	-8.48	-2.63	4.75	13.23	5.85	7.38	1.53
December 2021	-2.95	1.01	5.90	8.85	3.96	4.89	0.93

Table 4: Posterior estimates of the 10th, 50th and 90th percentiles for the monthly nowcast $T + 2$ across selected periods during the pandemic period.

6   Conclusions

Motivated by the limitations of popular VAR models for low-frequency economic variables, we introduce a novel mixed-frequency quantile vector autoregression (MF-QVAR) model. The proposed method exploits the informational content of high- and low-frequency variables to produce forecasts and nowcasts of conditional quantiles for indicators of interest. This permits to derive quantile-related risk measures at high frequency, thus enabling timely policy interventions.

The MF-QVAR model admits a state-space representation where the measurement equation follows a multivariate asymmetric Laplace distribution. Bayesian inference is performed by means of an efficient MCMC algorithm that exploits a data augmentation scheme coupled with a precision sampler to estimate the missing low-frequency variables at higher frequencies. To illustrate the forecasting accuracy of the proposed framework, we conduct two Monte Carlo simulation experiments and showcased that our proposed MF-QVAR model is capable of handling complex dependencies and unobserved variables relative to the standard models prevalent in the existing literature.

The proposed method is applied to US macroeconomic data via an in- and out-of-sample context. The in-sample analysis shows that the low-frequency US growth-at-risk

34

estimate can display significant dynamics compared to its corresponding high-frequency counterpart during a recession. For the out-of-sample exercise, we obtain real-time nowcasts for the US growth-at-risk. The results show the ability of MF-QVAR to produce meaningful monthly nowcasts that outperform the quarterly U-MIDAS QR benchmark and reveal interesting patterns during the GFC period and the onset and subsequent phases of the COVID-19 pandemic.

References

Adams, P. A., T. Adrian, N. Boyarchenko, and D. Giannone (2021). Forecasting macroeconomic risks. *International Journal of Forecasting 37*(3), 1173–1191.

Adrian, T., N. Boyarchenko, and D. Giannone (2019). Vulnerable growth. *American Economic Review 109*(4), 1263–89.

Antolin-Diaz, J., T. Drechsel, and I. Petrella (2021). Advances in nowcasting economic activity: Secular trends, large shocks and new data. Technical report, CEPR Discussion Paper No. DP15926.

Bańbura, M., D. Giannone, M. Modugno, and L. Reichlin (2013). Now-casting and the real-time data flow. In *Handbook of Economic Forecasting*, Volume 2, pp. 195–237. Elsevier.

Berger, T., J. Morley, and B. Wong (2023). Nowcasting the output gap. *Journal of Econometrics 232*(1), 18–34.

Bernardi, M., G. Gayraud, and L. Petrella (2015). Bayesian tail risk interdependence using quantile regression. *Bayesian Analysis 10*(3), 553–603.

Caggiano, G., E. Castelnuovo, and G. Nodari (2022). Uncertainty and monetary policy in good and bad times: A replication of the vector autoregressive investigation by bloom (2009). *Journal of Applied Econometrics 37*(1), 210–217.

Carriero, A., T. E. Clark, and M. Marcellino (2022). Nowcasting tail risk to economic activity at a weekly frequency. *Journal of Applied Econometrics 37*(5), 843–866.

Casarin, R., C. Foroni, M. Marcellino, and F. Ravazzolo (2018). Uncertainty through the lenses of a mixed-frequency Bayesian panel Markov-switching model. *The Annals of Applied Statistics 12*(4), 2559–2586.

Chan, J. C. (2020). Large bayesian vars: A flexible kronecker error covariance structure. *Journal of Business & Economic Statistics 38*(1), 68–79.

Chan, J. C. and I. Jeliazkov (2009). Efficient simulation and integrated likelihood estimation in state space models. *International Journal of Mathematical Modelling and Numerical Optimisation 1*(1-2), 101–120.

Chan, J. C., A. Poon, and D. Zhu (2023). High-dimensional conditionally gaussian state space models with missing data. *Journal of Econometrics 236*(1), 105468.

Chavleishvili, S. and S. Manganelli (2021). Forecasting and stress testing with quantile vector autoregression. Technical Report 2330, ECB Working Paper.

Cong, Y., B. Chen, and M. Zhou (2017). Fast simulation of hyperplane-truncated multivariate normal distributions. *Bayesian Analysis 12*(4), 1017–1037.

Durbin, J. and S. J. Koopman (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika 89*(3), 603–616.

Ghysels, E. (2016). Macroeconomics and the reality of mixed frequency data. *Journal of Econometrics 193*(2), 294–314.

Ghysels, E., P. Santa-Clara, and R. Valkanov (2005). There is a risk-return trade-off after all. *Journal of Financial Economics 76*(3), 509–548.

Giannone, D., L. Reichlin, and D. Small (2008). Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics 55*(4), 665–676.

Gneiting, T. and R. Ranjan (2011). Comparing density forecasts using threshold- and quantile-weighted scoring rules. *Journal of Business & Economic Statistics 29*(3), 411–422.

Hauber, P. and C. Schumacher (2021). Precision-based sampling with missing observations: A factor model application. *Deutsche Bundesbank Discussion Paper*.

Huber, F. and L. Rossini (2022). Inference in Bayesian additive vector autoregressive tree models. *The Annals of Applied Statistics 16*(1), 104–123.

Kadiyala, K. R. and S. Karlsson (1997). Numerical methods for estimation and inference in bayesian var-models. *Journal of Applied Econometrics 12*(2), 99–132.

Kaufmann, S. and C. Schumacher (2019). Bayesian estimation of sparse dynamic factor models with order-independent and ex-post mode identification. *Journal of Econometrics 210*(1), 116–134.

Kilian, L. and R. J. Vigfusson (2017). The role of oil price shocks in causing US recessions. *Journal of Money, Credit and Banking 49*(8), 1747–1776.

Koenker, R. and G. Bassett (1978). Regression quantiles. *Econometrica*, 33–50.

Koop, G., S. McIntyre, J. Mitchell, and A. Poon (2020). Regional output growth in the United Kingdom: More timely and higher frequency estimates from 1970. *Journal of Applied Econometrics 35*(2), 176–197.

Kotz, S., T. Kozubowski, and K. Podgórski (2001). *The Laplace distribution and generalizations: a revisit with applications to communications, economics, engineering, and finance*. Number 183. Springer Science & Business Media.

Mariano, R. S. and Y. Murasawa (2003). A new coincident index of business cycles based on monthly and quarterly series. *Journal of Applied Econometrics 18*(4), 427–443.

Merlo, L., L. Petrella, and V. Raponi (2021). Forecasting VaR and ES using a joint quantile regression and its implications in portfolio allocation. *Journal of Banking & Finance 133*, 106248.

Mitchell, J., A. Poon, and D. Zhu (2023). Constructing density forecasts from quantile regressions: Multimodality in macro-financial dynamics. Technical report.

Mogliani, M. and A. Simoni (2021). Bayesian MIDAS penalized regressions: Estimation, selection, and prediction. *Journal of Econometrics 222*(1), 833–860.

Montes-Rojas, G. (2019). Multivariate quantile impulse response funcitons. *Journal of Time Series Analysis 40*, 739–752.

Neal, R. M. (2003, 6). Slice sampling. *The Annals of Statistics 31*(3), 705–767.

Petrella, L. and V. Raponi (2019). Joint estimation of conditional quantiles in multivariate linear regression models with an application to financial distress. *Journal of Multivariate Analysis 173*, 70–84.

Polson, N. G. and J. G. Scott (2010). Shrink globally, act locally: Sparse Bayesian regularization and prediction. *Bayesian Statistics 9*(501-538), 105.

Rue, H. and L. Held (2005). *Gaussian Markov random fields: Theory and applications.* Chapman and Hall/CRC.

Schorfheide, F. and D. Song (2015). Real-time forecasting with a mixed-frequency VAR. *Journal of Business & Economic Statistics 33*(3), 366–380.

Schorfheide, F., D. Song, and A. Yaron (2018). Identifying long-run risks: A Bayesian mixed-frequency approach. *Econometrica 86*(2), 617–654.

Tian, Y., E. Li, and M. Tian (2016). Bayesian joint quantile regression for mixed effects models with censoring and errors in covariates. *Computational Statistics 31*(3), 1031–1057.

Tobias, A. and M. K. Brunnermeier (2016). CoVaR. *American Economic Review 106*(7), 1705.

Yu, K. and R. A. Moyeed (2001). Bayesian quantile regression. *Statistics & Probability Letters 54*(4), 437–447.