Bayesian VARs and Prior Calibration in Times of COVID-19

Benny Hartwig*

September 27, 2021

Abstract

This paper investigates whether a heavy-tailed or time-varying volatility error structure is better suited in dealing with the abnormal COVID-19 observations in a Bayesian VAR and discusses pitfalls of using mechanical prior updates. This paper presents evidence that the COVID-19 shock is better captured as a rare event rather than a persistent increase in volatility. Not accounting for heavy-tailed errors may lead to imprecise density forecasts during the pandemic. This paper shows that mechanical updates of prior distributions which depend on scale estimates – such as the commonly used Minnesota prior – may be another source of parameter instability. To mitigate this sensitivity, a COVID-19 robust prior calibration strategy is put forward.

Keywords: COVID-19, Bayesian vector autoregression, t-distributed errors, stochastic volatility, robust prior calibration strategy

JEL: C11, C51, C53.

*Goethe University Frankfurt and European Central Bank, E-mail: benny.hartwig90@gmail.com

I gratefully acknowledge comments and conversations with Michael Binder, Elena Bobeica and Michele Lenza as well as from conference and seminar participants at the 11th RCEA Money-Macro-Finance Conference and the RCC5 Brownbag. The views expressed are those of the authors and do not necessarily reflect those of the European Central Bank.

1 Introduction

The COVID-19 pandemic wrecked havoc in the global economy and triggered unprecedented movements in many key macroeconomic indicators. This unusual shock poses several challenges for the estimation and analysis of macroeconometric models, in general. In particular, Bayesian vector autoregressive models (BVARs) may become unstable and generate implausible forecasts when re-estimated with recent observations, see Lenza and Primiceri (2020); Carriero et al. (2021); Schorfheide and Song (2020) on the U.S. and Bobeica and Hartwig (2021) on the euro area.

The pandemic also raises the question as to how the recent data should be treated in an econometric model. Schorfheide and Song (2020) argue that this depends on the view as to how the pandemic affects fundamental workings of the economy.[1] Though the pandemic may have triggered forces altering macroeconomic interactions in the future, it is still too early to assess this possibility. Therefore, the literature currently focuses on robustifying the standard BVAR model against extreme observations related to the pandemic. One branch of the literature proposes to relax the error structure of the BVAR model.[2] These approaches range from modelling COVID-19 as a common shock via an explicit volatility model (Lenza and Primiceri, 2020), or implicitly by allowing for heavy tails as captured by multivariate t-distributed errors (Bańbura et al., 2020; Bobeica and Hartwig, 2021), or towards allowing for variable-specific outliers or t errors in the volatility process (Carriero et al., 2021).

However, the literature lacks a formal comparison of whether a heavy-tailed or common time-varying volatility error structure is better suited to account for such extreme variation in a BVAR. Modelling the pandemic as a rare event or persistent increase in macroeconomic volatility is an important choice as it has different implications for density forecasts. This paper contributes to the literature by formally assessing both features individually and jointly via multivariate t-distributed errors and common stochastic volatility in a BVAR with generalized error structure of Chan (2020). Because these error structures treat extreme observations as a common shock to volatility, the forecast of these models is compared to that of the Lenza and Primiceri (2020) approach.

[1] If fundamental propagation is viewed to change, the COVID-19 data should be included and previous data should be discounted. In contrast, if fundamentals are believed to remain robust, the recent data should be viewed as outliers and be excluded from the estimation or modelled explicitly.

[2] Alternative approaches to model the extreme the observation is to introduce a dummy variable for the pandemic period (Carriero et al., 2021) or to include other exogenous variables related to the pandemic (Ng, 2021) or to relax the overall model structure by using a non-parameteric and non-linear mixed-frequency VAR approach (Huber et al., 2020).

Modelling COVID-19 as a common shock instead of treating it as a variable-specific shock has some merits to it. First, the influence of extreme observations is downweighed homogeneously instead heterogeneously across equations, implying that the COVID-19 pandemic has a more limited effect on all parameters in the VAR model. Second, inference with a common shock treatment is much simpler and requires considerably less computation time than estimating the more flexible model of Carriero et al. (2021). Moreover, notice that the flexibility of their model comes at the cost of robustness as their modelling choice implies that the time-varying covariance matrix is especially sensitive to the ordering of variables when volatility clusters idiosyncratically, see Hartwig (2020).

A second gap in the literature is that prior calibration in times of COVID-19 has not been discussed yet (except for fixing at pre-COVID-19 period). This paper aims at filling this gap and argues that mechanical updates of standard calibration methods using variable-specific scale estimates, e.g., the commonly used Minnesota prior, are not robust to extreme observations and may be another source of parameter instability. Particularly, these scale estimates are computed as the root mean squared deviation (RMSD) of an AR(p) residual from an ordinary regression. To mitigate this sensitivity, several robust strategies are considered for calibrating the Minnesota prior.

For the empirical application, this paper considers an updated data set of Lenza and Primiceri (2020) with key U.S. macroeconomic indicators for production, labor markets and prices at quarterly frequency until 2020:Q4. The first contribution of this paper is to document that multivariate t-distributed errors are favored over common stochastic volatility during the pandemic as measured by the marginal likelihood. Specifically, the data prefers to interpret the COVID-19 shock as a rare event rather than a persistent increase of macroeconomic volatility. Nevertheless, diagnostics indicate that both features are generally important before and during the pandemic. Moreover, this paper shows that allowing for time-varying volatility but not accounting for heavy-tailed errors may lead to imprecise density forecasts during the pandemic. This complements the findings in Carriero et al. (2021) who model volatility to be variable-specific. Forecasts based on the Lenza and Primiceri (2020) approach yield a similar mean forecast, however, density forecasts are somewhat wider and strongly revised over the last two quarters. Recently, these density forecast became narrower and indicate that the COVID-19 volatility is of transient nature, complementing the evidence on modelling it as a rare event.

The second contribution of this paper is to document another source of parameter instability stemming from a mechanical update of the Minnesota prior due to standard calibration methods. The revised variable-specific scale estimates yield a very different prior distribution as compared to the pre-pandemic calibration. Specifically, leading to a very tight prior distribution for the VAR coefficients of the real variables and a very loose prior for the price variables. This altered prior distribution is associated with a decisive loss of the in-sample fit, becoming even worse than a naïve benchmark prior, and is primarily driven by a strongly heterogeneous response of the variables during the pandemic.

Third, robustifying the prior calibration against extreme observation improves the model fit substantially and yields a comparable prior distribution as of pre-pandemic times. However, the robustified strategies offer key advantages over a manual fine-tuning of the calibration sample as they are easy to implement, do not sacrifice in-sample fit for robustness and process all available information. Based on theoretical considerations and simplicity, the prior calibration with the scaled median absolute deviation and median AR(p) residual is put forward as the COVID-19 robust alternative.

This paper proceeds as follows. Section 2 presents Bayesian VAR models to cope with COVID-19 and discusses prior calibration strategies. Section 3 conducts empirical analysis of these models and priors using U.S. data. Section 4 concludes this paper.

2 Methodology

2.1 Bayesian VARs to cope with the COVID-19 pandemic

This section presents the Bayesian VAR model with a generalized error structure of Chan (2020) and uses it as a framework to assess whether heavy-tailed or heteroskedastic errors should be preferred for modelling the COVID-19 observations.

Let y_t be an $n \times 1$ vector of variables that is observed over the periods $t = 1, ..., T$. Consider the following generic VAR(p) model:

$$y_t = a_0 + A_1 y_{t-1} + ... + A_p y_{t-p} + \epsilon_t$$

where a_0 is an $n \times 1$ vector of intercepts and $A_1, ..., A_p$ are $n \times n$ coefficient matrices. Let $x_t' = (1, y_{t-1}', ..., y_{t-p}')$ be a $1 \times k$ vector of an intercept and lags with $k = 1 + np$ and $A = (a_0, A_1, ..., A_p)'$ is a $k \times n$.

Rewrite the VAR in more compact form

$$y_t' = x_t' A + \epsilon_t$$

and stack the observation over $t = 1, ..., T$, which yields

$$Y = XA + E \tag{1}$$

where Y, X, and E are respectively of dimensions $T \times n$, $T \times k$, and $T \times n$. In a standard VAR, the innovations $\epsilon_1, ..., \epsilon_T$ are assumed to be independent and identically distributed (i.i.d.) as $N(0, \Sigma)$. More compactly, rewrite $E \sim MN(0, \Sigma \otimes I_T)$, where Σ is an $n \times n$ covariance matrix, I_T is the identity matrix of dimension T, MN denotes the matric-variate normal distribution, $\otimes$ is the Kronecker product.

The Kronecker structure allows for more general covariance structures:

$$E \sim MN(0, \Sigma \otimes \Omega) \tag{2}$$

where Ω is a $T \times T$ covariance matrix. Intuitively, the cross-sectional and serial covariance structures of Y are separately modelled by specifying Σ and Ω, respectively. By choosing a suitable serial covariance structure Ω, the model in (1)–(2) includes a variety of flexible error structures such as heavy-tailed and heteroskedastic innovations.

The main difference between these error structures lies in their treatment of large shocks. Under a heavy-tailed distribution large shocks receive more probability mass in general, while with heterosekdastic errors large shocks are captured by allowing volatility to change over time. Intuitively, these error structures may limit the effect of extremely large innovations on the parameter estimates because they downscale their contribution through the (implied) serial covariance Ω, see Section 2.3 for details. This idea of down-weighing observations to stabilize VAR parameter estimates in the context of the COVID-19 pandemic was first proposed by Lenza and Primiceri (2020).

From the class of heavy-tailed distributions, this paper considers the multivariate t-distribution as it can accomodate substantially thicker tails than the normal distribution. Specifically, ϵ_t is assumed to be independently multivariate t-distributed with mean vector 0, scale matrix Σ and degree of freedom parameter ν, or more compactly, $\epsilon_t \sim t(0, \Sigma, \nu)$. As explained by Chan (2020), many non-Gaussian distributions can be used in this framework as they can be written as a scale mixture of Gaussian distributions.[3] The t-distribution can be modelled by letting $\Omega = \text{diag}(\lambda_1, ..., \lambda_T)$, where each λ_t follows independently an inverse-gamma distribution $(\lambda_t|\nu) \sim IG(\nu/2, \nu/2)$, see Geweke (1993). In other words, the scale mixture Gaussian representation of the t-distribution

[3]The outlier adjusted error distribution of Stock and Watson (2016) or, as discussed by Chan (2020), the multivariate Laplace distribution of Eltoft et al. (2006) may also be used in this context.

reveals that the innovations ϵ_t conditional on λ_t exhibit common but independently distributed shifts of the cross-sectional covariance matrix over time. Note Bańbura et al. (2020) and Bobeica and Hartwig (2021) also consider multivariate t-distributed errors to model the abnormal variation due to the COVID-19 pandemic.

As heteroskedastic error structure, the common drift in stochastic volatility model in the spirit of Carriero et al. (2016) is considered with $\epsilon_t \sim N(0, \Sigma \cdot \exp(h_t))$ following Chan (2020). The log variance follows a stationary autoregressive process of order one:

$$h_t = \rho h_{t-1} + \epsilon_t^h \tag{3}$$

with $\epsilon_t^h \sim N(0, \sigma_h^2)$, $|\rho| < 1$ and $\Omega = \text{diag}(\exp(h_1), \exp(h_2), ..., \exp(h_T))$.

Note this covariance structure shares some similarities with the approach of Lenza and Primiceri (2020). Specifically, they assume that the COVID-19 observations scale the cross-sectional covariance matrix persistently by a common factor that decays over time. For the pre-pandemic sample, they assume a constant covariance matrix. In contrast, the model with common stochastic volatility allows volatility to change over the whole sample, not just in response to the pandemic.

The third error structure that is considered is a combination of common drift in volatility and multivariate t innovations, which was also considered by Chan (2020). Specifically, $\epsilon_t \sim N(0, \Sigma \cdot \lambda_t \cdot \exp(h_t))$, where $\lambda_t \sim IG(\nu/2, \nu/2)$ and h_t follows an AR(1) process. This model allows for richer (implied) volatility dynamics as opposed to their single ingredients. It can capture both persistent but also abnormally large changes in volatility and thereby, allows the data to decide whether the COVID-19 shock has transient or persistent effects on volatility.[4] Table 1 summarizes the list of competing Bayesian VAR models with a generalized error structure.

Table 1: List of competing models

Model	Description
BVAR-$\mathcal{N}$	BVAR with i.i.d. Gaussian innovations
BVAR-t	BVAR with i.i.d. t innovations
BVAR-$\mathcal{N}$-CSV	BVAR with Gaussian innovations and common stochastic volatility
BVAR-t-CSV	BVAR with t innovations and common stochastic volatility

Apart from the COVID-19 related macroeconomic disruptions, there is a vast litera-

[4] The idea to combine these error structures to discriminate between high-frequency and low-frequency volatility was first put forward by Jacquier et al. (2004).

ture documenting a departure from classical Gaussian errors in a wide range of macroeconometric models. In the context of VAR models, Christiano (2007) presents evidence that even before the Global Financial Crisis VAR residuals are not normally distributed. Moreover, Chiu et al. (2017), Clark and Ravazzolo (2015) and Chan (2020) document that the data favors stochastic volatility and t-distributed errors over Gaussian errors in a VAR model and that these features improve forecast accuracy. In addition, Cúrdia et al. (2014) and Chib and Ramamurthy (2014) provide evidence that structural shocks in a DSGE models exhibit heavy tails.[5] Antolín-Díaz et al. (2020) show that accounting for both shifts in variances and large shocks improves nowcast accuracy of a dynamic factor model relative to classical Gaussian errors.

Though stochastic volatility is often found to be a more important feature than t-distributed errors for macroeconomic time series, Cúrdia et al. (2014), Chiu et al. (2017), Chan (2020) and Antolín-Díaz et al. (2020) document that inference and forecast accuracy may hinge on allowing for heavy-tailed errors. For instance, ignoring fat tails may yield misleading inference about the historical evolution of volatility, especially during the Global Financial Crisis. Particularly, models that allow for both stochastic volatility and heavy-tailed errors interpret the Global Financial Crisis as a rare event rather than a persistent increase of macroeconomic volatility.

The Bayesian VAR models that combine stochastic volatility with t-distributed errors of Clark and Ravazzolo (2015) and Chiu et al. (2017) or with outlier-adjusted errors of Carriero et al. (2021) differ from above models by their assumption that time-varying volatility and heavy-tailedness is variable-specific rather than common across all variables. Though an individual error distribution may be more realistic and less restrictive, their approach of implementing this flexibility comes at the cost of robustness. Specifically, individual volatility and heavy-tailedness for each variable is modelled by making use of a triangular factorisation of the time-varying covariance matrix in the spirit of Cogley and Sargent (2005). However, a drawback of this approach is that the estimated time-varying covariance matrix depends on the ordering of variables and may lead to systematically different estimates when volatility evolves idiosyncratically, giving rise to ambiguous empirical conclusions, see Hartwig (2020).

[5] Cúrdia et al. (2014) considers both stochastic volatility and t-distributed errors, while Chib and Ramamurthy (2014) focuses on t-distributed errors only.

2.2 Prior calibration

Besides the error structure in a Bayesian VAR, another source of parameter instability may be due to an unusually strong and potentially unintended re-calibration of the prior distribution during the pandemic. Specifically, many prior distributions are calibrated using variable-specific scale estimates which are based on all currently available information in the sample and as such may be sensitive to extreme observations. The widely used Minnesota prior is just one example of such a prior distribution.[6]

To analyze this sensitivity, this section uses a weakly informative but sample independent prior as a naïve benchmark and proposes various approaches to robustify the calibration of the Minnesota prior in the event of extreme observations.

For all considered Bayesian VAR models, this paper assumes a normal-inverse-Wishart prior for the VAR coefficients A and the cross-sectional covariance matrix Σ, i.e. $\Sigma \sim IW(S_0, \nu_0)$, and $(A|\Sigma) \sim MN(A_0, \Sigma \otimes V_a)$. As a naïve benchmark prior, a weakly informative but sample independent calibration is used for the normal-inverse-Wishart prior. The hyperparameters of the inverse Wishart prior are set to $\nu_0 = 3 + n$ and $S_0 = I_n$. The prior mean A_0 is centered at zero and the covariance matrix, V_a, is assumed to be diagonal with i-th diagonal element $v_{a,ii}$ set as:

$$v_{A,ii} = \begin{cases} 1 & \text{for a coefficient associated to a lag of variable } r \\ \kappa_2^2 & \text{for an intercept} \end{cases} \tag{4}$$

where $\kappa_2 = 10$ controls the overall strength of shrinkage for the intercepts.

To maintain the Kronecker structure of the prior, this paper adopts a Minnesota prior without cross-variable shrinkage, see Carriero et al. (2016) and Chan (2020).[7] The Minnesota prior is implemented as normal-inverse-Wishart prior using the same notation as above. For the inverse Wishart prior, the same sample independent calibration is used in order to reduce the effects of a sample dependent calibration to the prior covariance matrix V_a for the VAR coefficients only.[8] V_a is assumed to be diagonal with i-th diagonal

[6] Further examples of commonly used and scale-dependent priors are the sum of coefficient (co-integration) and sum of initial conditions prior of Sims and Zha (1998), the steady-state prior of Villani (2009) and the long-run relations prior of Giannone et al. (2019).

[7] For more details on this specific Minnesota prior, see Section 3.2.1 in Karlsson (2013).

[8] Alternatively, the scale S_0 may be calibrated using sample dependent information as $S_0 = (\nu_0 - n - 1) \cdot \text{diag}(\hat{s}^2_{1,m}, ..., \hat{s}^2_{n,m})$, see Section 3.2.1. Karlsson (2013).

element $v_{a,ii}$ set as:

$$v_{A,ii} = \begin{cases} \frac{\kappa_{1,m}^2}{l^2 \hat{s}_{r,m}^2} & \text{for a coefficient associated to lag } l \text{ of variable } r \\ \kappa_2^2 & \text{for an intercept} \end{cases} \tag{5}$$

where $\hat{s}_{r,m}$ is an estimate of scale of variable r of calibration strategy m, specified below, and $\kappa_{1,m}$ controls the degree of shrinkage for the VAR coefficients. For comparability, $\kappa_{1,m}$ is calibrated for each method m to yield the maximum marginal likelihood of the standard BVAR-$\mathcal{N}$ model at pre-pandemic times, see Section 3.2.

In practice, $\hat{s}_{r,m}$ is calibrated to be the variable-specific sample volatility of the residual from an AR(p) model with Gaussian and homoskedastic errors, see Litterman (1986). Particularly, the root mean squared deviation (RMSD) is used as the sample specific volatility estimate for variable r:

$$RMSD(x_r) = \frac{1}{T}\sqrt{\sum_{t=1}^{T} (x_{r,t} - \text{mean}(x_r))^2} \tag{6}$$

where x_r is the AR(p) residual of variable r in the standard specification.[9]

There are two main reasons why the Minnesota prior is based on variable-specific volatility estimates. First, the prior is not scale invariant, i.e. it matters whether data enters in decimals or percentage points. Second, an estimate of the variable-specific volatility is necessary to express the degree of shrinkage relative to the variability of the underlying stochastic process which is similar in spirit to a ridge regression.

However, a caveat of the classical prior calibration strategy is its sensitivity towards extreme observations. Particularly, some variable-specific volatility estimates may be strongly inflated by the COVID-19 related extreme observations, leading to a substantially tighter prior for some variables and a much looser prior for other variables as compared to the pre-pandemic calibration. This revision may be unattractive for several reasons. First, it is questionable whether a few extreme observations should receive so much weight in revising the prior distribution for the VAR dynamics. Second, it is a-priori not clear if and how quickly macroeconomic transmission channels change in response to the pandemic. Therefore, it might be reasonable to assume that fundamental transmission channels of the economy are a-priori unchanged.

The unique character of the COVID-19 crisis allows the researcher to easily specify

[9] For ease of exposition, the degree of freedom correction is omitted here.

such a prior by cutting the calibration sample at pre-pandemic times, as, for instance, done in Schorfheide and Song (2020) and Lenza and Primiceri (2020). However, this approach has the drawback that it discards all future observations and thus neglects potentially useful from the new normal once macroeconomic adjustment normalize. This raises the question as to how new observations should treated when specifying a prior distribution for common empirical application. Moreover, ad-hoc fine-tuning of the prior calibration sample may be unattractive going forward as the "appropriate" time to include new observations may vary across application and country as well as by the judgement of the researcher. This might adversely affect the comparability across empirical applications.

For this reason, this paper searches for alternative calibration strategies that are robust against extreme observations, can process all information and yield a prior distribution that is comparable to the standard method in the absence of extreme observations. To meet the latter requirement, the search is guided by looking for robust estimators that yield asymptotically the same point estimate as a mean regression when the data is normally distributed. There are two elements that may lead to a revision of the variable-specific volatility estimate: 1) the volatility estimator and 2) the input series x_r used for the volatility estimator. To investigate the importance of theses sources, this paper considers three robust volatility estimators and combines them with three alternative input series to calibrate the Minnesota prior.

The scaled median absolute deviation (MAD) is considered as a first alternative to the RMSD because it is robust to extreme observations, simple to implement and yields asymptotically the same point estimate of volatility when the data is normally distributed, see Rousseeuw and Croux (1993):

$$MAD(x_r) = b \cdot \text{median}(|(x_r - \text{median}(x_r)|) \tag{7}$$

where $b = 1.4826$ is a scaling parameter.

However, a drawback of the MAD is that it takes a symmetric view point on dispersion and has a rather low Gaussian efficiency, see Rousseeuw and Croux (1993). As an alternative, the authors propose the S_n and Q_n statistic which are more efficient and general than the MAD. Specifically, they remain valid even when the input series x_r follows an asymmetric or a heavy-tailed distribution.

For these reasons, the S_n statistic is considered as second alternative which measures dispersion as the typical distance between observations as opposed to the typical distance from the central value like the MAD:[10]

$$S_n(x_r) = c \cdot \text{median}\{\text{median}|x_{r,t} - x_{r,s}|\}, \quad s = 1, ..., T \tag{8}$$

where $c = 1.1926$ is a scaling parameter.

As a third alternative, the Q_n statistic is considered which is the *k-th* order statistic of the $\binom{T}{2}$ interpoint distances:

$$Q_n(x_r) = d \cdot \{|x_{r,t} - x_{r,s}|; t < s\}_{(k)} \tag{9}$$

where $d = 2.219$ is a scaling parameter and $k = \binom{(T/2)+1}{2}$. Note this estimator has the highest Gaussian efficiency among the MAD and the S_n statistic.

The second ingredient for robustifying the calibration is the use of alternative input series. Here, the first difference of the time series (FD), the residual of a standard Gaussian AR(p) model (oReg) and the residual of a median AR(p) model (qReg) for variable r are considered as input series x_r. The median AR(p) model is estimated using the quantile regression framework of Koenker and Bassett (1978).

The advantage of using a first-differenced time series over an AR(p) residual as an input is that the former remains unchanged, while the latter may change due to a revision of model parameters. This invariance property makes the first-differenced time series an ideal input to assess the sensitivity of alternative volatility estimators due to extreme observations related to the pandemic. However, a drawback is that the associated volatility estimate is generally larger but not proportionally larger across the n equations because of different lag properties of the individual variables. Consequently, the prior distribution based on this input series might be somewhat different and the amount of shrinkage may be suboptimal for a VAR model.[11]

Relative to the first-differenced time series, the residual of a median AR(p) model tends to be more similar to the residual of a standard AR(p) model as it take the lag structure into account. For Gaussian data, both estimation methodologies produce asymptotically the same coefficient estimate as the (conditional) median and mean of a normally distributed variable is identical. However, the median regression has the advantage of being more robust and efficient than the mean regression when the error distribution is heavy-tailed, see Koenker and Bassett (1978). Thus, the residual of the median regression is expected to be less sensitive to extreme observations related to the pandemic than that

[10]The lack of location dependence makes the S_n statistic valid for asymmetric distributions.

[11]Note $\kappa_{1,m}$ is common across the n equations and cannot act as a degree of freedom.

of the mean regression.

Table 2 provides an overview of all considered prior specifications. The Minnesota prior is calibrated based on twelve alternative strategies by making use of four alternative volatility estimators: RMSD, MAD, S_n and Q_n; and three input series: first difference, AR(p) residual of a mean and median regression.

Table 2: List of prior specifications

Prior		Description
Weak		Weakly informative prior (naïve benchmark)
RMSD	(FD)	Minnesota prior (first difference and standard scale, RMSD)
MAD	(FD)	Minnesota prior (first difference and robust scale, MAD)
S_n	(FD)	Minnesota prior (first difference and robust scale, S_n)
Q_n	(FD)	Minnesota prior (first difference and robust scale, Q_n)
RMSD	(oReg)	Minnesota prior (standard AR(p) and scale, RMSD)
MAD	(oReg)	Minnesota prior (standard AR(p) but robust scale, MAD)
S_n	(oReg)	Minnesota prior (standard AR(p) but robust scale, S_n)
Q_n	(oReg)	Minnesota prior (standard AR(p) but robust scale, Q_n)
RMSD	(qReg)	Minnesota prior (robust-q AR(p) and scale, RMSD)
MAD	(qReg)	Minnesota prior (robust-q AR(p) and scale, MAD)
S_n	(qReg)	Minnesota prior (robust-q AR(p) and scale, S_n)
Q_n	(qReg)	Minnesota prior (robust-q AR(p) and scale, Q_n)

To complete the prior specification, this paper follows Chan (2020) and assumes a uniform prior on $(2, 100)$ for the degree of freedom parameter, i.e. $\nu \sim U(2, 100)$. This prior implies that degrees of freedom may become sufficiently large to approximate normal distributed errors. For stochastic volatility, independent priors for σ_h^2 and ρ: $\sigma_h^2 \sim IG(\nu_{h0}, S_{h0})$ and $\rho \sim N(\rho_0, V_\rho)\mathbf{1}(|\rho| < 1)$ are assumed, where $\nu_{h0} = 5$, $S_{h0} = 0.04$, $\rho_0 = 0.9$, and $V_\rho = 0.2^2$. These values imply that the prior mean of σ_h^2 is 0.1 and ρ is centered at 0.9. Further, an independent prior for the parameter blocks (A, Σ) and Ω is assumed, i.e., $p(A, \Sigma, \Omega) = p(A, \Sigma)p(\Omega)$.

2.3 Bayesian estimation

The Bayesian VAR model with covariance structure (2) can be easily estimated under a normal-inverse-Wishart prior for (A, Σ) and an independent prior for the parameter blocks (A, Σ) and Ω, see Chan (2020).

Posterior draws can be obtained by sequentially sampling from 1) $p(A, \Sigma|Y, \Omega)$ and 2) $p(\Omega|Y, A, \Sigma)$. Depending on the covariance structure Ω, additional blocks might be needed to sample some extra hierarchical parameters. In particular, the BVAR with t-distributed errors requires additional sampling of the degree of freedom parameter ν, while (ρ, σ_h^2) need to be sampled for the BVAR with common stochastic volatility. Following Chan (2020), these parameters are fitted using univariate time series models.

To understand how (extreme) observations are treated in this model, it is useful to investigate the conditional posterior for $p(A, \Sigma|Y, \Omega)$. Because of a natural conjugate prior $p(A, \Sigma)$, the conditional posterior for $p(A, \Sigma|Y, \Omega)$ is still normal-inverse-Wishart distributed:

$$p(A, \Sigma|Y, \Omega) \sim MNIW(\hat{A}, \Sigma \otimes K_A^{-1}, \hat{S}, v_0 + T)$$

where

$$K_A = V_A^{-1} + X'\Omega^{-1}X, \quad \hat{A} = K_A^{-1}(V_A^{-1}A_0 + X'\Omega^{-1}Y),$$
$$\hat{S} = S_0 + A_0'V_A^{-1}A_0 + Y'\Omega^{-1}Y - \hat{A}'K_A^{-1}\hat{A}.$$

Hence, $(A, \Sigma|Y, \Omega)$ can be sampled in two steps. First, sample Σ marginally from $(\Sigma|Y, \Omega) \sim IW(\hat{S}, v_0 + T)$. Then, given the Σ drawn sample A from

$$(A|Y, \Sigma, \Omega) \sim MN(\hat{A}, \Sigma \otimes K_A^{-1}).$$

Note under a noninformative prior distribution for the VAR coefficients, i.e. when the inverse of V_A goes to zero, the conditional posterior mean of A converges to the generalized least squared estimator for some conditionally known Ω (weighting matrix).[12] Thus, the data (Y, X) is not equally informative for the parameters (A, Σ), but their contribution is weighted by Ω. Under t-distributed errors and common stochastic volatility, Ω is a diagonal matrix with generic elements $\omega_t^2, t = 1, ..., T$. In period t, the observations (y_t, x_t) are weighted by the (implicit) volatility ω_t^{-1}. Specifically, when ω_t is large, then (y_t, x_t) is less informative for (A, Σ). Intuitively, the common volatility estimate ω_t will be large, when for some variable r or for a set of variables the forecast errors $\epsilon_{t,r}$ are unusually large.

[12]When $\Omega = I_T$, then the posterior mean of A converges to the least squares estimator.

3 Empirical Application

3.1 Data

To investigate how these more flexible Bayesian VARs cope with the COVID-19 observations and to study how the calibration of the Minnesota prior is affected during the pandemic, an updated data set of Lenza and Primiceri (2020) with key U.S. macroeconomic variables at quarterly frequency until 2020:Q4 is considered.[13,14]

This data set includes six variables: (i) employment, (ii) unemployment rate, (iii) consumption, (iv) industrial production, (v) consumer price index, and (vi) PCE core price index. All variables except the unemployment rate enter the VAR models in log-levels times 100 and with a lag length of 4. To compare the models before and during the pandemic, the models are estimated over an expanding window with a common start point from 1988:Q4 to 2019:Q4 until 2020:Q4.

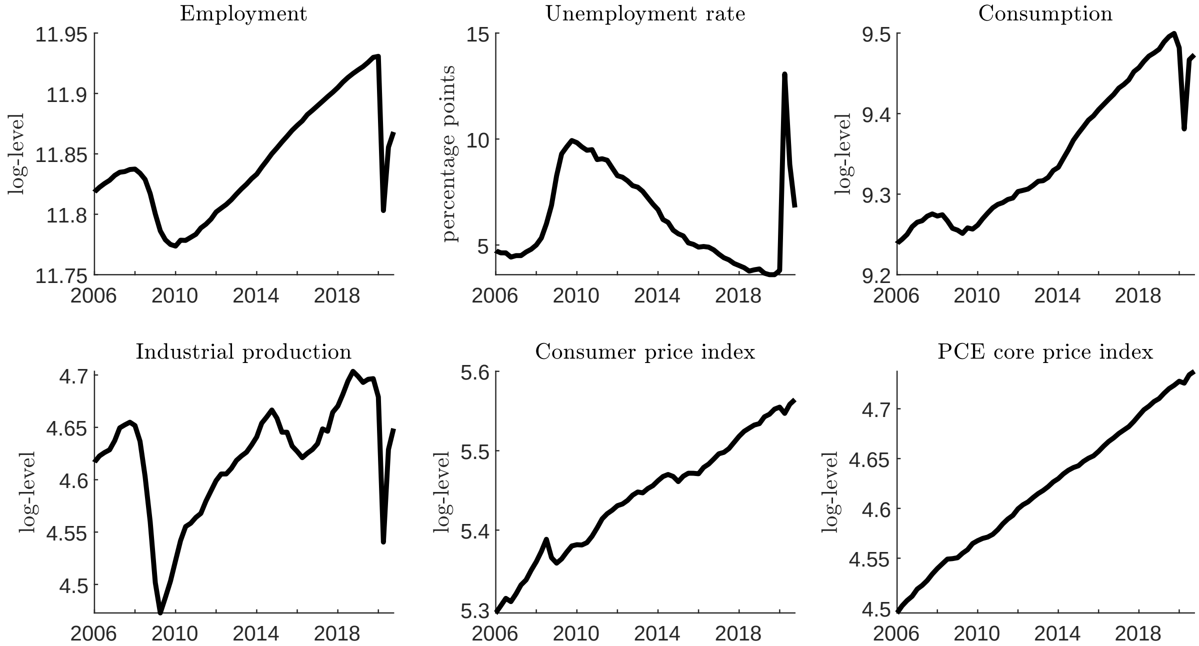

Figure 1: Time-series plots for key U.S. macroeconomic variables

Figure 1 shows an excerpt of the data set for the sample from 2006:Q1 until 2020:Q4. From March 2020 onwards, COVID-19 started to spread through the U.S. and led to a widespread shutdown of the U.S. economy, leading to unparalleled declines in labor

[13]The quarterly frequency is adopted to provide additional evidence on model (in)stability.
[14]This paper considers the data set of the June 2020 version of Lenza and Primiceri (2020).

market indicators like employment but also consumption and production in the second quarter of 2020. Surprisingly, the overall price level did not react particularly strong to the economic turmoil. With lockdown measures easing over the summer, the U.S. economy rebounded somewhat during the third quarter but lost momentum due to a second wave of infections in the fourth quarter.

3.2 (Re)-calibration of the Minnesota prior

Before discussing estimation results, it is useful to consider how the Minnesota prior is affected by alternative calibration strategies and by the pandemic. Table 3 reports the estimated inputs used to calibrate the Minnesota prior at pre-pandemic times (2019:Q4). Across alternative calibration strategies, a few facts stand out: First, Panel (a) shows that estimates of variable-specific volatility are generally larger but not proportionally larger when the first difference is used as an input series, e.g., compare volatility estimates for employment or PCE core. This difference emerges as these first-differenced time series exhibit some variable-specific serial correlation.

Second, estimates of volatility are similar when the AR(p) residual of the mean or the median regression is considered. This is because both regression techniques produces comparable coefficient estimates in the pre-pandemic sample.

Third, the RMSD yields the highest volatility estimate, while the MAD produces the lowest estimate among the robust volatility estimators. The S_n and Q_n statistic are overall (if any) just marginally larger than the MAD except for some cases. Moreover, a larger difference between RMSD and MAD suggests that the input series for variable r exhibits some non-normal behavior. This difference is generally larger when the first difference is used as input series. Therefore, the S_n and Q_n statistic should be preferred over the MAD in these cases because they are more efficient and may be more informative as they can account for skewness in the distribution.

Due to these characteristics, the optimal tuning parameter $\kappa_{1,m}$ differs across calibration strategies and is generally larger for the RMSD than for the robust metrics, see Panel (b).[15] However, by using (5), the bulk of this variation can be explained by adjusting $\kappa_{1,m}$ for an average level shift of the volatility estimates, see Panel (c):

$$\kappa_{1,j} = \kappa_{1,i} \left(\frac{1}{n} \sum_{r=1}^{n} \left(\frac{\hat{\sigma}_{r,j}}{\hat{\sigma}_{r,i}} \right) \right) \tag{10}$$

[15] The marginal likelihood maximizing $\kappa_{1,m}$ was identified via a grid search on $(0.01 : 0.01 : 10)$.

where $\hat{\sigma}_{r,\cdot}$ is the variable r specific volatility estimate of method j and i respectively.

Table 3: Calibration of the Minnesota prior before the pandemic

	RMSD (FD)	MAD (FD)	S_n (FD)	Q_n (FD)	RMSD (oReg)	MAD (oReg)	S_n (oReg)	Q_n (oReg)	RMSD (qReg)	MAD (qReg)	S_n (qReg)	Q_n (qReg)
EMP	0.51	0.28	0.28	0.28	0.19	0.14	0.15	0.16	0.20	0.13	0.13	0.15
UR	0.28	0.16	0.19	0.22	0.19	0.17	0.18	0.18	0.19	0.16	0.17	0.18
CON	0.83	0.42	0.44	0.46	0.39	0.32	0.32	0.35	0.39	0.33	0.32	0.36
IP	1.30	0.70	0.80	0.85	0.85	0.75	0.76	0.81	0.87	0.69	0.71	0.78
CPI	0.77	0.30	0.33	0.35	0.44	0.32	0.33	0.35	0.45	0.33	0.33	0.34
PCE	0.52	0.17	0.18	0.17	0.13	0.13	0.14	0.13	0.13	0.13	0.13	0.13

(a) Estimated volatility in 2019:Q4

$\kappa_{1,m}$	0.82	0.39	0.43	0.45	0.39	0.33	0.33	0.35	0.39	0.31	0.32	0.34

(b) Optimized shrinkage based on BVAR-$\mathcal{N}$ in 2019:Q4

$\hat{\kappa}_{1,m}$	0.88	0.40	0.43	0.45	0.39	0.33	0.34	0.35	0.40	0.32	0.32	0.35
SHIFT	2.26	1.03	1.11	1.15	1.00	0.84	0.88	0.91	1.02	0.82	0.83	0.89
STD	0.88	0.28	0.25	0.22	0.00	0.09	0.11	0.07	0.02	0.10	0.10	0.08

(c) Approximated shrinkage and summary statistics of the relative volatility

Note: Panel (a) shows variable specific volatility estimates until 2019:Q4 for employment (EMP), unemployment rate (UR), consumption (CON), industrial production (IP), consumer price index (CPI) and PCE core prices (PCE) based on different volatility estimators and different input series. Panel (b) reports optimal degree of shrinkage $\kappa_{1,m}$ of the BVAR-$\mathcal{N}$ estimated until 2019:Q4. Panel (c) shows approximated $\hat{\kappa}_{1,m}$ by accounting for an average volatility shift and using optimal value of RMSD(oReg). SHIFT and STD are the mean and standard deviation of variable-specific volatility estimates relative to those of RMSD(oReg).

In this application, these analog $\kappa_{1,m}$ values are fairly close to their optimal values, especially for the robust metrics. This is an encouraging empirical fact as it suggests that established values for $\kappa_{1,m}$ of various empirical applications can be easily mapped into to an alternative calibration strategy by using equation (10). Of course, the goodness of this approximation depends on the degree of proportionality among these alternative volatility estimates. To quantify the degree of proportionality, the standard deviation of the relative volatility estimates may be used, see STD in Panel (c). The statistic shows that strategies using the AR(p) residual are strongly proportional to each other, i.e. they yield a very similar prior distribution, while those using the first-differenced time series are somewhat less proportional.

Table 4 shows how volatility estimates are affected by the pandemic. Panel (a) reports the ratio of variable-specific volatility estimates from 2020:Q4 to 2019:Q4, while Panel (b) shows summary statistics for this volatility ratio.

Table 4: Re-calibration of Minnesota priors during the pandemic

	RMSD (FD)	MAD (FD)	S_n (FD)	Q_n (FD)	RMSD (oReg)	MAD (oReg)	S_n (oReg)	Q_n (oReg)	RMSD (qReg)	MAD (qReg)	S_n (qReg)	Q_n (qReg)
EMP	2.63	1.06	1.03	1.06	6.42	2.13	1.95	1.98	7.57	1.15	1.09	1.08
UR	3.49	1.18	1.03	1.00	4.72	1.64	1.58	1.61	5.40	0.93	0.96	0.99
CON	1.75	1.00	1.04	1.03	3.15	1.49	1.62	1.50	3.52	1.09	1.08	1.05
IP	1.51	1.10	1.05	1.06	2.20	1.04	1.14	1.14	2.29	1.07	1.09	1.06
CPI	1.00	1.05	1.01	1.02	1.03	1.01	1.05	1.04	1.02	0.91	1.02	1.02
PCE	1.00	0.99	1.01	1.02	1.16	1.06	1.04	1.04	1.15	1.11	1.04	1.02

(a) Ratio of volatility: 2020:Q4 to 2019:Q4

	RMSD (FD)	MAD (FD)	S_n (FD)	Q_n (FD)	RMSD (oReg)	MAD (oReg)	S_n (oReg)	Q_n (oReg)	RMSD (qReg)	MAD (qReg)	S_n (qReg)	Q_n (qReg)
MEAN	1.90	1.06	1.03	1.03	3.11	1.40	1.40	1.39	3.49	1.04	1.05	1.04
DEV	1.27	0.09	0.03	0.04	2.87	0.57	0.52	0.52	3.43	0.10	0.07	0.05

(b) Descriptive statistics

Note: Panel (a) presents the ratio of volatility estimates from 2020:Q4 relative to 2019:Q4. Panel (b) reports descriptive statistics of the ratio of volatilities: MEAN is the average and DEV is the root of mean squared deviation from one.

Some of the non-robust volatility estimates are substantially inflated as compared to their pre-pandemic values. Estimates based on the RMSD and coupled with an AR(p) residual are particularly affected, increasing on average by a factor of three, c.f. MEAN in Panel (b). Moreover, the increase across variables is strongly heterogeneous, with real variables being primarily affected and price variables almost left unchanged. For employment, the volatility estimate by RMSD(oReg) is over 6 times larger as compared to its pre-pandemic estimate. With unchanged $\kappa_{1,m}$, this means that the prior variance for the first lag of the employment coefficient becomes implausibly tight, being roughly $41(\approx 6.42^2)$ times tighter than its pre-pandemic value (the revised prior variance is 0.1022 in 2020:Q4 as compared to 4.213 in 2019:Q4).

Thus, the revised prior distribution differs considerably from its pre-pandemic calibration. Also, re-optimizing $\kappa_{1,m}$ or accounting for an average volatility shift cannot generate a comparable pre-pandemic prior distribution as the change of volatility estimates is too heterogeneous. Therefore, the mechanical revision induced by the standard calibration may be viewed as another source of parameter instability in a Bayesian VAR.

In contrast, estimates based on robust volatility estimators and coupled with the first difference as input series are hardly affected by the pandemic observations. The average increase of the volatility ratio is close to one and the deviation from one is close to zero. This is not surprising as this input series is not revised by model parameters and extreme observations are ignored using these robust statistics.

The robust volatility estimators coupled with an AR(p) residual are somewhat sensitive to these observations. The median regression is less sensitive than the mean regression since it is more robust to extreme observations. In fact, the median regression yields a largely comparable residual as in pre-pandemic times. Notice that there is still a sizeable revision for some variables when the residual of the mean regression is used. For employment, the volatility is roughly doubled, highlighting the importance of a revision in the input series. Therefore, the calibration strategies with robustified volatility estimators and input series FD and qReg yield a moderately changed prior distribution as compared to pre-pandemic times.

3.3 Model and prior comparison

To discriminate between these various Bayesian VAR models and prior specifications, the marginal likelihood is used as a formal Bayesian model selection criterion following Chan (2020).[16] The marginal likelihood under model $M_{k,m}$ is defined as

$$p(y|M_{k,m}) = \int p(y|\theta_k, M_{k,m})p(\theta_k|M_{k,m})d\theta_k$$

where $p(y|\theta_k, M_{k,m})$ is the likelihood function, $p(\theta_k|M_{k,m})$ is the prior distribution, θ_k is a model-specific parameter vector and m is a prior calibration. For the BVAR-$\mathcal{N}$ with a natural conjugate prior, an analytical formula is available and used to compute the marginal likelihood, while for the other BVAR models no closed-form exists and Chib's method (Chib, 1995) is used instead.

Specifically, if the marginal likelihood under model $M_{i,s}$ is larger than under $M_{j,q}$, then the data is more likely under model $M_{i,s}$ as compared to $M_{j,q}$. Given all models are a-priori equally likely, the weight of evidence between two models can be measured by the Bayes factor defined as the ratio of marginal likelihoods, see Chan (2017). Table 5 presents estimated log marginal likelihoods before and during the pandemic based on updated parameter estimates and prior calibrations.[17]

Panel (a) shows that before the pandemic all Minnesota priors are decisively favored over the naïve benchmark prior by an average log Bayes factor of 64, i.e. the BVARs with Minnesota prior are 6.23×10^{27} more likely than with a naïve benchmark prior. This

[16] The marginal likelihood does not always favor the most general model as it trades off model fit against complexity. Intuitively, the higher the dimension of the parameter space, the more spread is the prior with respect to the likelihood. This is known as Occam's razor.

[17] The marginal likelihood cannot be directly compared across different points in time as it is based on different information sets due to different sample sizes.

Table 5: Log marginal likelihood

Prior		BVAR-$\mathcal{N}$		BVAR-t		BVAR-$\mathcal{N}$-CSV		BVAR-t-CSV	
Weak		−353.53	–	−342.51	(0.07)	−311.43	(0.07)	**-310.02**	(0.12)
RMSD	(FD)	−293.20	–	−282.03	(0.04)	−244.81	(0.07)	**-243.43**	(0.15)
MAD	(FD)	−295.94	–	−285.22	(0.04)	−247.18	(0.07)	**-245.90**	(0.10)
S_n	(FD)	−294.93	–	−284.21	(0.04)	−246.85	(0.06)	**-245.54**	(0.13)
Q_n	(FD)	−294.44	–	−283.96	(0.08)	−246.15	(0.07)	**-244.80**	(0.10)
RMSD	(oReg)	−290.90	–	−281.32	(0.06)	−242.42	(0.08)	**-240.89**	(0.10)
MAD	(oReg)	−291.04	–	−280.77	(0.03)	−242.31	(0.07)	**-240.97**	(0.17)
S_n	(oReg)	−290.69	–	−280.44	(0.04)	−240.84	(0.06)	**-239.38**	(0.13)
Q_n	(oReg)	−291.01	–	−281.03	(0.05)	−241.37	(0.06)	**-239.98**	(0.08)
RMSD	(qReg)	−290.92	–	−281.46	(0.05)	−241.27	(0.08)	**-239.86**	(0.11)
MAD	(qReg)	−290.14	–	−280.29	(0.06)	−240.51	(0.08)	**-239.06**	(0.12)
S_n	(qReg)	−290.16	–	−280.12	(0.07)	−240.93	(0.07)	**-239.17**	(0.10)
Q_n	(qReg)	−290.91	–	−280.95	(0.05)	−240.91	(0.08)	**-239.79**	(0.12)

(a) Estimation sample until 2019:Q4

Prior		BVAR-$\mathcal{N}$		BVAR-t		BVAR-$\mathcal{N}$-CSV		BVAR-t-CSV	
Weak		−633.29	–	−434.63	(0.11)	−412.87	(0.30)	**-404.58**	(0.24)
RMSD	(FD)	−601.67	–	−414.38	(0.11)	−349.88	(0.19)	**-343.58**	(0.19)
MAD	(FD)	−582.45	–	−375.78	(0.09)	−345.78	(0.22)	**-339.21**	(0.25)
S_n	(FD)	−583.05	–	−375.21	(0.16)	−348.74	(0.30)	**-339.57**	(0.15)
Q_n	(FD)	−582.56	–	−374.84	(0.12)	−346.98	(0.14)	**-339.23**	(0.18)
RMSD	(oReg)	−666.59	–	−502.21	(0.08)	−381.78	(0.16)	**-376.25**	(0.13)
MAD	(oReg)	−585.30	–	−385.98	(0.12)	−342.62	(0.17)	**-335.22**	(0.16)
S_n	(oReg)	−585.62	–	−385.84	(0.15)	−339.57	(0.16)	**-333.20**	(0.23)
Q_n	(oReg)	−585.75	–	−386.50	(0.13)	−341.57	(0.18)	**-333.71**	(0.17)
RMSD	(qReg)	−690.02	–	−530.33	(0.08)	−393.06	(0.13)	**-388.36**	(0.22)
MAD	(qReg)	−582.41	–	−372.07	(0.14)	−342.54	(0.16)	**-334.12**	(0.17)
S_n	(qReg)	−581.77	–	−371.78	(0.12)	−342.97	(0.21)	**-334.59**	(0.22)
Q_n	(qReg)	−581.92	–	−372.41	(0.11)	−343.46	(0.16)	**-334.86**	(0.20)

(b) Estimation sample until 2020:Q4

Note: This table shows log marginal likelihood of various BVARs and prior specifications estimated until 2019:Q4 and 2020:Q4. Bold figures indicate maximum marginal likelihood for each prior specification. Brackets report numerical standard error.

overwhelming support for shrinking more distant VAR coefficients to zero is not surprising as more distant lags are relatively less informative for current dynamics. Moreover, estimated marginal likelihoods of robust calibration strategies are very similar to those of comparable classical calibration strategies. Specifically, the figures based on AR(p) residuals are almost identical across volatility estimators, while those based on the first difference are slightly smaller than that of the residuals. Thus, robustified calibration strategies do not sacrifice in-sample-fit for robustness.

Across all prior specifications, the BVAR-t-CSV is the best performing model albeit the marginal likelihood is just immaterially larger than that of the BVAR-$\mathcal{N}$-CSV model. Both models have, however, a substantial margin over the BVAR-t model, while the latter clearly outperforms the standard Gaussian BVAR-$\mathcal{N}$ model. Thus, both error extensions are favored over standard Gaussian errors by the data. But on a single ingredient basis, stochastic volatility is a more important feature than a heavy-tailed error distribution. This is line with the findings of Cúrdia et al. (2014), Chiu et al. (2017) and Chan (2020) who also look at the U.S. but consider a different data set and an earlier period.

How does the pandemic affect the estimation of the BVARs? Panel (b) shows that the ranking across BVARs remains unchanged, i.e. the BVAR-t-CSV is still the best model. However, the relative marginal likelihood of the BVAR-t with respect to the BVAR-t-CSV improves. Specifically, considering the naïve benchmark prior, the change of the log Bayes factor in 2019:Q4 to 2020:Q4 increases by 2.44 log points for the BVAR-t.[18] In contrast, this metric falls by 6.88 log points for the BVAR-$\mathcal{N}$-CSV. Therefore, the unprecedentedly large variation of the data can be better captured by multivariate t-distributed errors as compared to common stochastic volatility. This relative gain of a heavy-tailed distribution is due to the fact this type of variation is hard to capture by a persistent stochastic volatility process. Moreover, the change of the log Bayes factor is -228.71 for the BVAR-$\mathcal{N}$, suggesting that the classical Bayesian VAR is poorly equipped to deal with these extreme observations.

Turning to prior sensitivity, the Minnesota calibrations based on the AR(p) residual with the RMSD as volatility estimator are now dominated by the naïve benchmark prior for the BVAR with Gaussian and t-distributed errors. Moreover, log marginal likelihoods based on robust calibration strategies are very similar across specifications. Therefore, for calibrating the Minnesota prior, the S_n or Q_n statistic of Rousseeuw and Croux (1993) do not provide a substantial improvement in terms of the model fit over the scaled MAD. Furthermore, log marginal likelihoods based on robust volatility estimators are also significantly higher than those of analog calibration strategies. For instance, the log Bayes factor in favor of MAD(oReg) against the standard RMSD(oReg) for the BVAR-N-CSV is about 39.2.

To complement this analysis and gauge the impact of the prior revision on the model fit, Figure 2 shows estimated marginal likelihoods obtained over an expanding sample

[18] The change of the log Bayes factor for two models i and j with prior specification m in t to $t+h$ is computed as $\Delta BF_{t,t+h,m} = BF_{t,m} - BF_{t+h,m}$ where $BF_{t,m} = ML_{i,t,m} - ML_{t,j,m}$ and $ML_{t,i,m}$ is the log marginal likelihood on sample t of model i with prior m. For the weakly informative prior, the log Bayes factor between BVAR-t-CSV and BVAR-t is 32.49 in 2019:Q4 and 30.05 in 2020:Q4, hence, the change of the log Bayes factor is 2.44.

Figure 2: Prior revision and log marginal likelihood

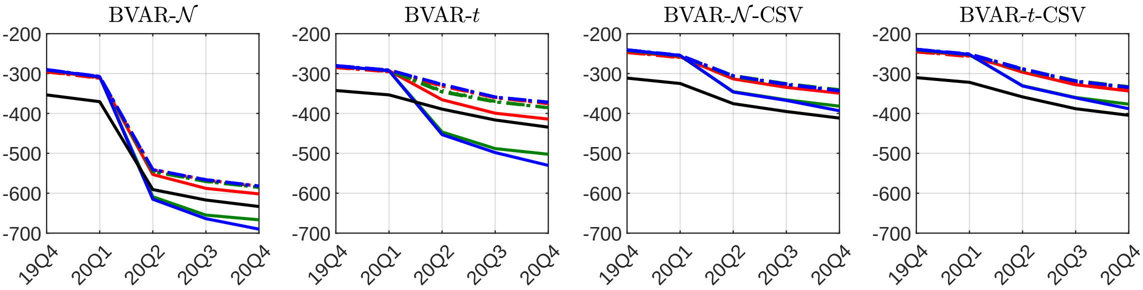

(a) Estimated with a re-calibrated Minnesota prior

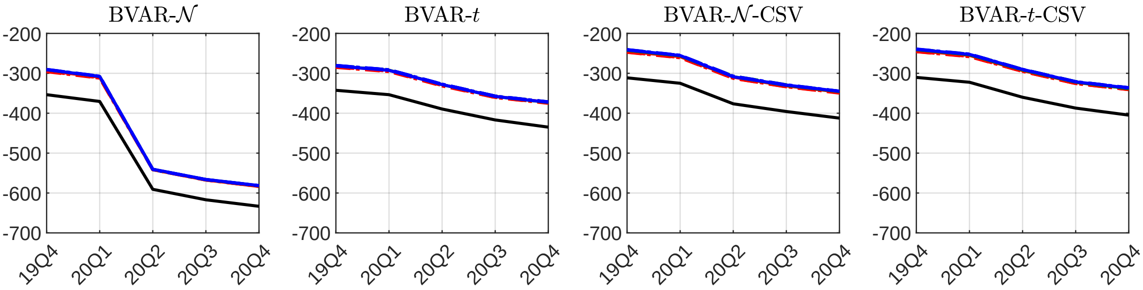

(b) Estimated with a 2019:Q4-fixed Minnesota prior

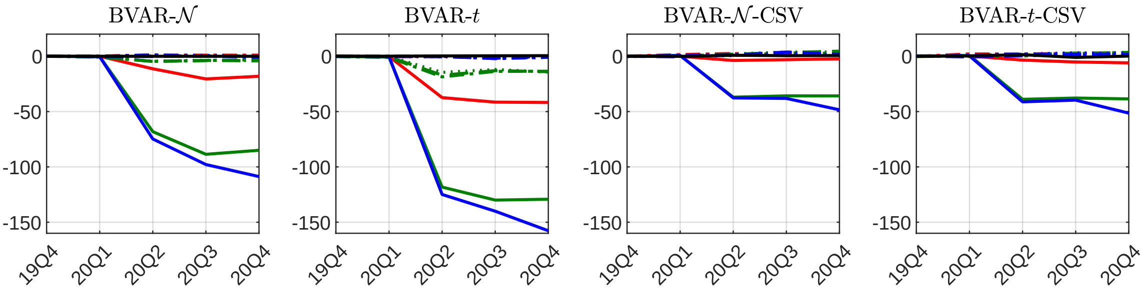

(c) Relative marginal likelihood (a)–(b)

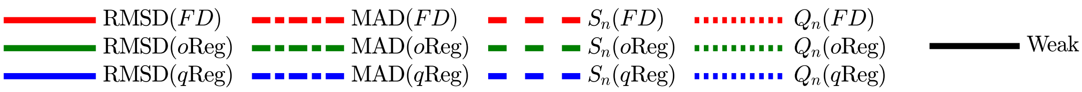

Notes: The figure shows marginal likelihoods of the BVAR models based on alternative prior distributions and estimated over an expanding sample from 2019:Q4 to 2020:Q4.

from 2019:Q4 to 2020:Q4 with (a) a re-calibrated prior, (b) a 2019:Q4-fixed prior and (c) the relative marginal likelihood for these calibrations. These graphs show that re-calibrating the Minnesota prior with standard calibration strategies (thick colored lines) in times of the pandemic has a strongly negative effect on the overall in-sample-fit starting in 2020:Q2, see Panel (c). For instance, the log Bayes factor of the BVAR-$\mathcal{N}$ with a fixed prior against re-calibrated RMSD(oReg) prior is about 80.97 in 2020:Q4. This means that this BVAR model with a fixed prior is 1.46×10^{35} times more likely then the same model with a re-calibrated prior. In addition, Panel (c) shows that these prior updates increase the time-variation of the marginal likelihood (e.g. see blue and green thick line).

The figure also shows that for the robust volatility estimators log marginal likelihoods are hardly sensitive except for BVARs with Gaussian or t errors coupled with the AR(p) residual of the mean regression. Here, the updated residual leads to a substantial decline of the model fit. Thus, these results suggest that the robust calibration strategies based on the AR(p) residual of the quantile regression or the first difference should be preferred over standard calibration methods. In addition, these robust methods should also be favored over fixing the prior distribution at pre-pandemic times because they are easier to work with in practice as they do not require a manual fine-tuning of the calibration sample.

In the remainder of the paper, empirical results are presented for the Minnesota prior calibrated by the MAD(qReg) strategy. This prior calibration is chosen because it robustifies the calibration of RMSD(oReg) along the two dimensions and yields from a theoretical perspective asymptotically the same point estimates when the data is normally distributed. The MAD is favored over the S_n and Q_n statistic as it is simpler to implement and yields a comparable in-sample-fit.

3.4 Macroeconomic tail risk and volatility

Given model diagnostics provide overwhelming evidence against Gaussianity, how does the pandemic affect macroeconomic tail risk and volatility? To dig deeper into this issue, Figure 3 shows (a) the posterior distribution of the degree of freedom parameter for the multivariate t-distribution, (b) the log common stochastic volatility and (c) the implied log common volatility for the respective BVARs.

Recall the lower the degree of freedom parameter ν of the t-distribution, the heavier the tails and the stronger it departs from Gaussianity. Panel (a) shows that before the pandemic the posterior of the degree of freedom parameter peaks around 7 and 9 for the BVAR-t and the BVAR-t-CSV (red area), indicating a substantial tail risk. Nevertheless, the posterior distribution for the BVAR-t-CSV is more wide-spread than that of the BVAR-t. This is because the former model can explain large shocks either by changes in stochastic volatility or by allowing for heavier tails.

With the pandemic observations included in the estimation sample, the posterior distribution becomes much tighter and shifts to the left, peaking around 4 for both models (blue area). For the BVAR-t, this is not surprising as this is the main channel of the model to accommodate more extreme innovations. The much sharper identification of this parameter for the BVAR-t-CSV means that the tail risk of the economy increased substantially besides changes in stochastic volatility.

Stochastic volatility is another important feature of the data. Due to the scale of the shock, Panel (b) shows the log instead of the raw series of common stochastic volatility $h_t/2$. Before the pandemic, there is some marked variation in the log stochastic volatility series, especially during the Global Financial Crisis. Moreover, the BVAR-$\mathcal{N}$-CSV estimate is somewhat different, specifically, more volatile than that of the BVAR-t-CSV (red

Figure 3: Macroeconomic tail risk and volatility

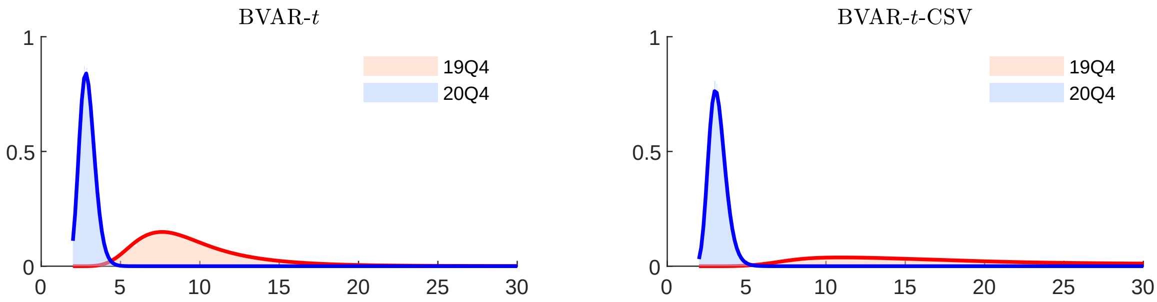

(a) Posterior distribution of the degree of freedom parameter

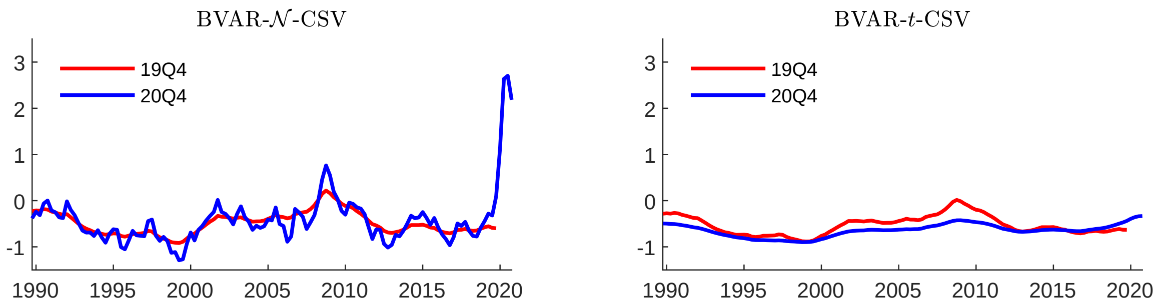

(b) Posterior mean estimates of log common stochastic volatility

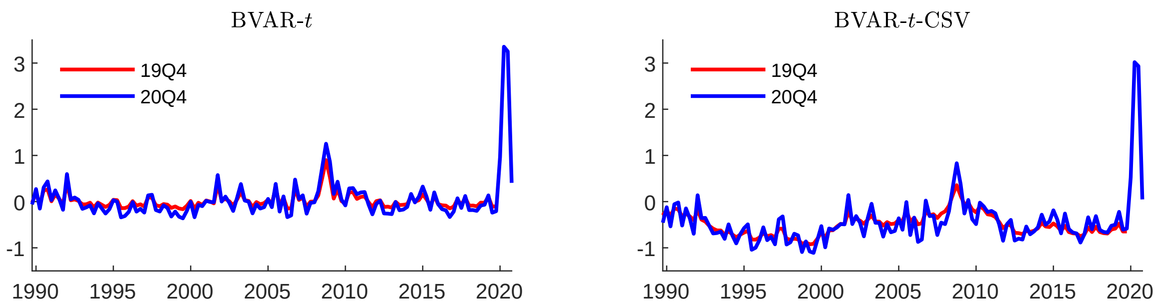

(c) Posterior mean estimates of implied log common volatility

Notes: The figure shows posterior estimates for the degree of heavy-tailedness, volatility and implied volatility of the error for the respective BVAR models estimated using the MAD(qReg) Minnesota prior calibration up to 2019:Q4 and 2020:Q3.

line). This indicates that some unusually large shocks are more of a transient nature and that fat tails are an important feature of the data. In other words, ignoring fat-tailedness of the errors may lead to misleading inference about the historical evolution of stochastic volatility. This general observation has been previously documented by Jacquier et al. (2004), Cúrdia et al. (2014), Chiu et al. (2017) and Chan (2020) for different data sets and sample periods.

Estimated stochastic volatility changes substantially with the COVID-19 observations (blue line). The series based on Gaussian errors is more volatile than the pre-pandemic estimate, while that of the t-error model is even less volatile than before. This pronounced difference emerges as the parameters governing the stochastic volatility process are revised in a different manner. Specifically, the disruption triggered by the pandemic is interpreted as a persistent spike in volatility by the BVAR-$\mathcal{N}$-CSV of roughly 14 standard deviations ($\approx \exp(2.65)$) in 2020:Q2. To fit this increase, the shock size hitting stochastic volatility process had to be upward revised. In contrast, the BVAR-t-CSV interprets this shock as a rare event and upward revises the tail risk instead of the innovation variance of stochastic volatility.

Panel (c) shows how the implied volatility under t-distributed errors is affected. Specifically, the implied log common volatility, $\log(\lambda_t)/2$, of the BVAR-t is more volatile during the pandemic, see left chart (cf. red and blue line) which is due to the increased tail risk in the economy. But also the log common volatility of the BVAR-t-CSV, $h_t/2 + \log(\lambda_t)/2$, has become more volatile, see right chart. This change of the volatility pattern is however primarily driven by the increased tail risk rather than by a changed stochastic volatility process, cf. right chart in Panel (b).

Besides notice that the compound log volatility series of the BVAR-t-CSV is somewhat more volatile but shares many similarities with the log stochastic volatility series of the BVAR-$\mathcal{N}$-CSV. In contrast, the implied volatility series of the BVAR-t exhibits by construction no persistent volatility pattern which is however an important feature of the data before the pandemic. The BVAR-t and BVAR-t-CSV quantify the pandemic to be a shock 29 and 20 standard deviations, respectively.

3.5 Model (in)stability and forecasting

The COVID-19 observations decisively affect macroeconomic tail risk and volatility. However, what is their influence on the common VAR parameters and how are forecasts affected in these models? Figure 4 shows scatter plots of the posterior mean VAR coefficients and residual correlation matrix of 2019:Q4 against 2020:Q4.

Panel (a) shows that for the BVAR-$\mathcal{N}$ some VAR coefficients are substantially affected by the pandemic observations. Specifically, intercepts and coefficients measuring (cross)-dynamics of employment, unemployment rate, consumption and industrial production are particularly strong affected. For instance, the first-order autoregressive coefficient of employment changes from 1.24 in 2019:Q4 to -0.44 in 2020:Q4. In contrast, the change of these parameters is more moderate for the BVARs with a more flexible error structure, remaining roughly on the 45 degree line.

Not only VAR coefficients, but also residual correlations of the BVAR-$\mathcal{N}$ change drastically with the outbreak of the pandemic, see Panel (b). For instance, the pre-pandemic correlations between residuals of the unemployment rate with employment, consumption and industrial production was around -0.45 while it declined to around -0.95 in 2020:Q4. This change is solely driven by the very synchronized co-movement of these variables in response to the COVID-19 shock. Because of its extreme size, residual correlations and variances are likely to remain overshadowed by this shock for an extended period of time. However, by allowing the VAR residuals to have heavy tails or time-varying volatility, the impact of the COVID-19 shock is downweighed by its scale and therefore, only mildly affects correlation estimates.

Figure 5 shows unconditional forecasts starting in 2021:Q1 for the BVAR models and adds the Lenza and Primiceri (2020) BVAR as another benchmark (denoted as BVAR-LP).[19,20] A few facts stand out: First, as noted by Lenza and Primiceri (2020), Carriero et al. (2021) Schorfheide and Song (2020), the forecast based on the standard Gaussian BVAR does not yield a meaningful forecast when parameters are re-estimated with the recent COVID-19 data (blue line).

Second all considered BVARs with a more flexible error structure produce a stable forecast over the pandemic period. Overall, the revision in parameters implies a somewhat dampened response of labor market variables and a stronger response of industrial production as compared to fixing the parameters in 2019:Q4. Nevertheless, the difference between projected paths is not significant. Thus, the COVID-19 shock does not (yet) indicate a significant structural break in VAR coefficients but rather distorts parameter estimates when the extreme variation is not modelled.

Third, density forecasts differ across error structures. In response to the COVID-19

[19] The BVAR-LP is estimated with the MAD(qReg) Minnesota prior and $\kappa_{1,m}$ is not re-optimized using Giannone et al. (2015) method for comparability. For estimation details, see Appendix A

[20] Appendix B reports forecasts as of 2020:Q3. Relative to the findings presented in this section, parameter estimates of the generalized BVARs are less affected by the COVID-19 observations and the BVAR-LP model produces much wider density forecasts than the BVARs with t errors.

Figure 4: Scatter of parameter estimates: 2019:Q4 against 2020:Q4

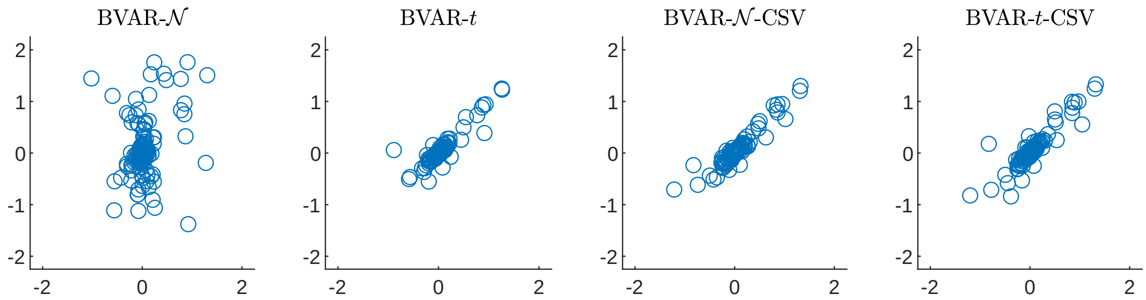

(a) Posterior mean estimates of VAR coefficients

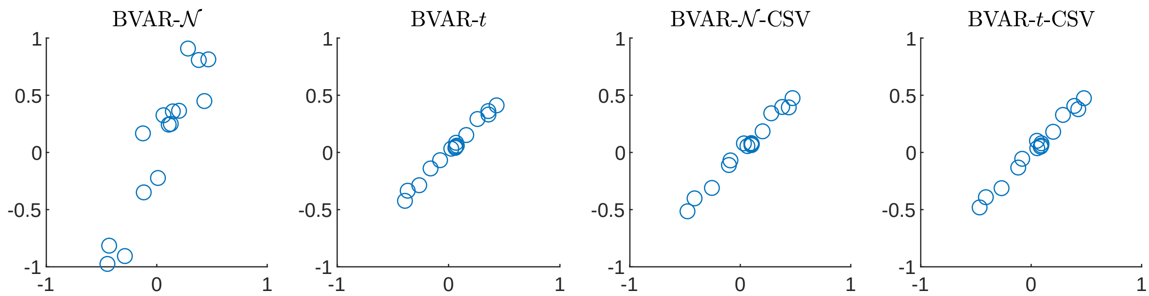

(b) Posterior mean estimates of VAR residual correlations

Notes: The figure shows scatter plots of parameters estimates from 2019:Q4 against 2020:Q4 for various BVARs using the MAD(qReg) Minnesota prior calibration.

shock, BVARs accounting only for time-varying volatility but not for heavy tails exhibit much wider forecasts intervals, see also Carriero et al. (2021). The BVAR-t-CSV indicates that macroeconomic uncertainty did not increase drastically in response to the pandemic shock but rather catapulted the economy on a new trajectory.

However, these unconditional forecasts should be interpreted with a grain of salt as these models have never seen the transmission of such a large and unusual shock before. For instance, the projected recovery of the unemployment rate appears to be too optimistic, reaching pre-pandemic levels already in 2022 and declining even further below a record low rate of about 2.5 percentage points in mid 2023. These dynamics are a consequence of the VAR treating the COVID-19 shock as a usual macroeconomic shock, but this may be inappropriate as this shock constitutes a mixture of supply and demand disturbances (Guerrieri et al., 2020; Baqaee and Farhi, 2020; Baldwin and Di Mauro, 2020; del Rio-Chanona et al., 2020) as well as might have a different degree of persistence as compared to traditional macroeconomic shocks (Primiceri and Tambalotti, 2020).

To produce a more meaningful forecast, off-model information may be used to disci-

Figure 5: Unconditional forecasts with re-estimated parameters

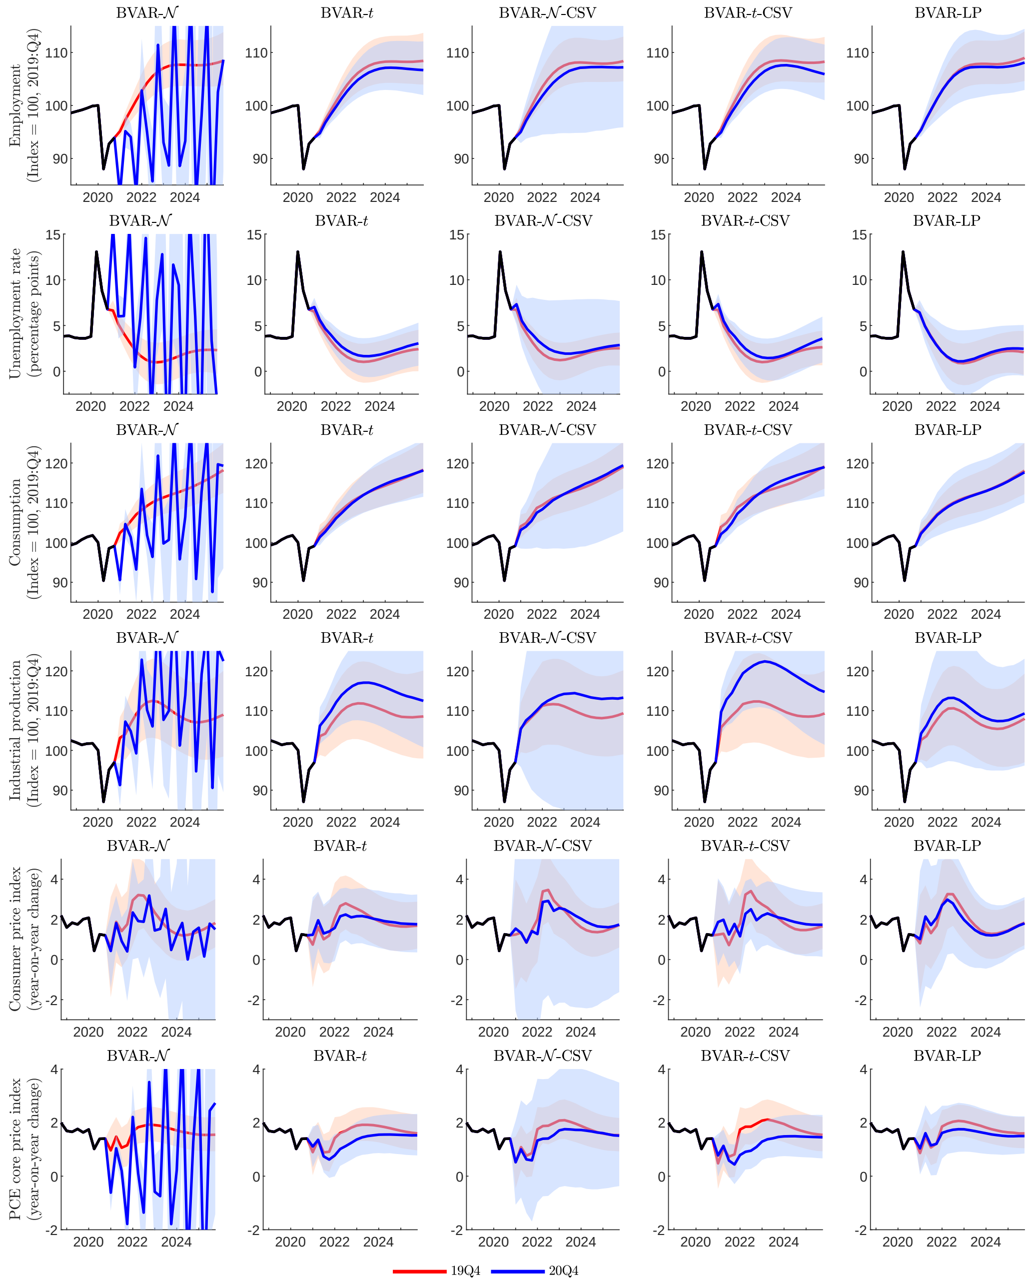

Notes: The figure shows unconditional forecasts starting in 2021:Q1 of various BVAR models estimated until 2019:Q4 and 2020:Q4 respectively. The thick lines are posterior mean estimates and the colored area depicts the 16% to 84% credible interval.

pline the evolution of some variables. This paper follows Lenza and Primiceri (2020) and conditions on the consensus forecast of the survey of professional forecasters (SPF) for the unemployment rate using the January 2021 release.[21,22]

Figure 6 shows the forecast based on this scenario and the accompanying SPF forecast.[23] Focusing on the COVID-19 robust BVARs with re-estimated parameters (blue line), the conditional mean forecast for the real variables is similar to that of the SPF forecast (green line), but substantially more pessimistic than the unconditional forecast in Figure 5. Note the sharp reduction in density forecast intervals of real variables is due to the conditioning scenario and the correlation of forecast errors.

Under the SPF scenario, employment is projected to remain below pre-pandemic levels for an extended period of time, while consumption is forecasted to recover from the COVID-19 shock already in 2022. For industrial production, the VAR forecast exhibits a somewhat unreasonable jump as compared to the SPF forecast. Price dynamics are ambiguous under this scenario. Similar to the SPF forecast, consumer price inflation exhibits some inflationary tendencies in the short to medium term. However, PCE core inflation is projected to remain subdued over the coming years, hovering around 1.3 percentage points. Thus, this exercise illustrates that these BVARs can still be used to produce a meaningful forecast in times of the pandemic.

4 Conclusion

Estimation of many standard macroeconomic models became a challenge with the outbreak of the COVID-19 pandemic and calls for an appropriate treatment of these extreme observations, in general. For a Bayesian VAR in particular, this paper documents that a heavy-tailed error distribution is better suited to deal with these extreme observations than a common time-varying volatility error structure. It shows that the BVAR accounting for both multivariate t-distributed errors and common stochastic volatility prefers to interpret the pandemic as a rare event rather than a persistent increase of macroeconomic volatility. Moreover, density forecasts of BVARs not accounting for heavy-tailed errors

[21] Alternatively, a conditional forecast may be produced by constructing a synthetic COVID-19 shock consisting of both macroeconomic supply and demand disturbances and tilting its propagation to different scenarios, see Primiceri and Tambalotti (2020) or by tilting the unconditional forecast to the long-run forecast of surveys only, see Bobeica and Hartwig (2021).

[22] The 2021:Q1 forecast for unemployment rate is 6.3 percentage points and prolonged using a smooth exponential curve to the long-run forecast of 4 percentage points in 2024. https://www.philadelphiafed.org/-/media/frbp/assets/surveys-and-data/survey-of-professional-forecasters/2021/spfq121.pdf

[23] For employment, consumption, and industrial production the SPF forecast is available until 2022, while for the CPI and PCE core the SPF provides one additional year.

Figure 6: Conditional forecasts with re-estimated parameters

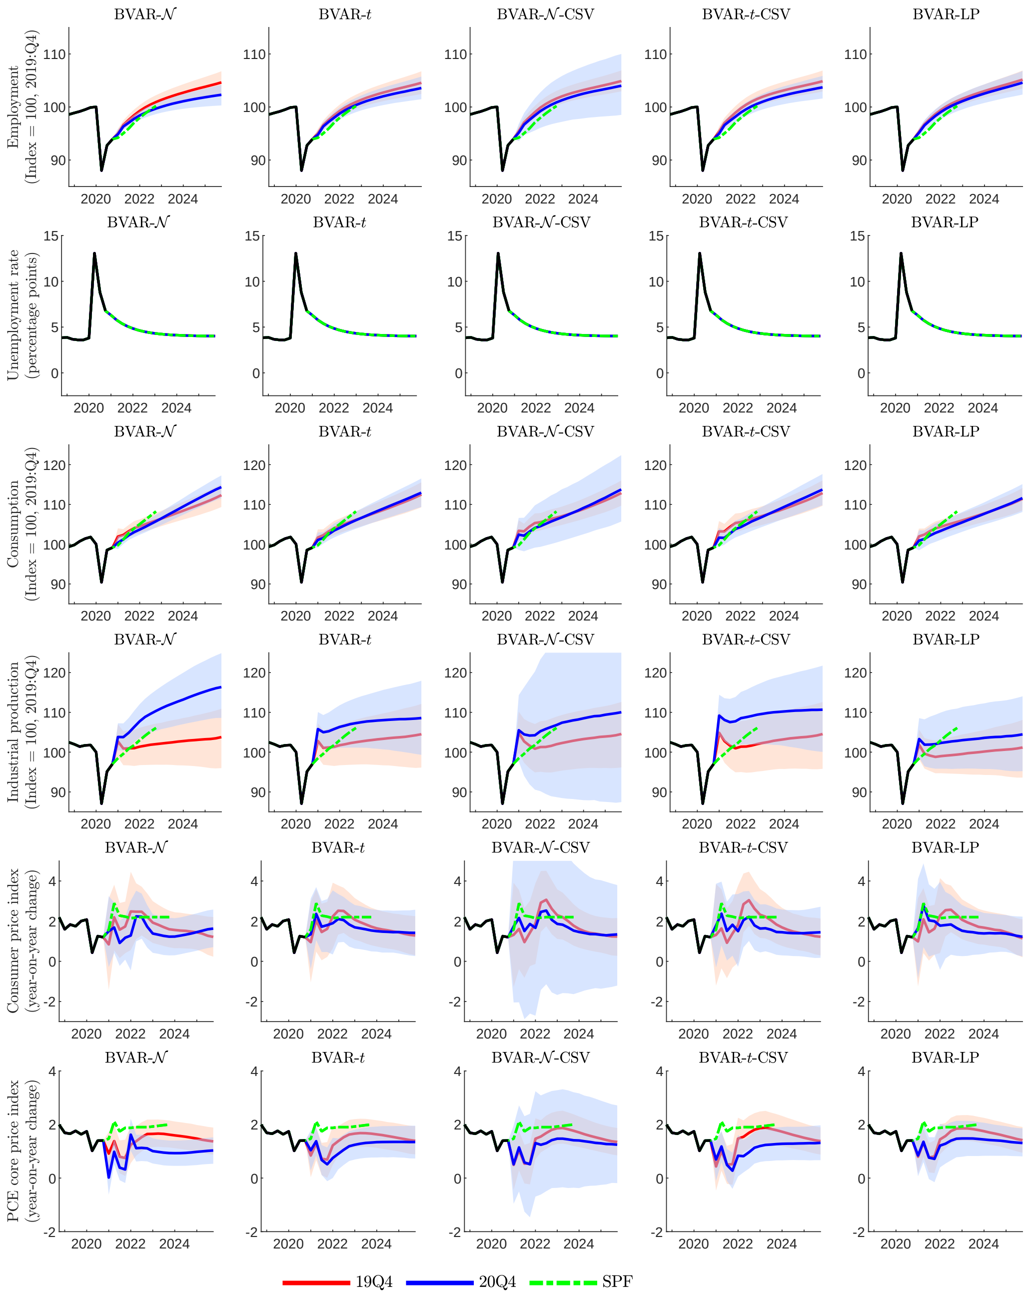

Notes: The figure shows conditional forecasts starting in 2021:Q1 of various BVAR models estimated until 2019:Q4 and 2020:Q4 respectively. The thick lines are posterior mean estimates and the colored area depicts the 16% to 84% credible interval.

but allowing for time-varying volatility yield less precise forecast during the pandemic. However, mean forecasts of the BVARs with t errors and stochastic volatility imply stable variable dynamics as opposed to a standard Gaussian model and produce a similar projection as compared to the BVAR of Lenza and Primiceri (2020).

Choosing between these alternative modelling approaches thus depends on the research question at hand and the view as to how COVID-19 should be reflected in the parameter estimates. The explicit treatment of the pandemic observations by Lenza and Primiceri (2020) approach offers the advantage that the overall volatility process and tail risk is left unaffected. However, a drawback of this approach is that these ad-hoc assumptions may pose the risk of mis-specification and are less flexible than a model based treatment with t-distributed errors and stochastic volatility. Thus, the results of this paper suggest that both of these approaches provide some room for improvement by making explicit modelling less prone to mis-specification and implicit modelling less prone to a revision in hyperparameters which provides a fruitful avenue for future research.

Apart from that, this paper documents another source of parameter instability stemming from a revision of the Minnesota prior distribution due to standard calibration methods. Particularly, the revised prior implies an extremely tight prior distribution for real variables but very loose distribution for price variables which results in a decisive loss of the in-sample fit. To mitigate this prior sensitivity, the prior calibration with the scaled MAD and median AR(p) residual is put forward as the COVID-19 robust alternative. The proposed COVID-19 robust prior calibration strategy is not only limited to the Minnesota prior, but may be also useful for calibrating other prior distributions that depend on variable-specific scale estimates such as the sum of coefficient (co-integration) and sum of initial conditions prior of Sims and Zha (1998), the steady-state VAR prior of Villani (2009) and the long-run relations prior of Giannone et al. (2019).

References

Antolín-Díaz, J., T. Drechsel, and I. Petrella (2020): "Advances in Nowcasting Economic Activity: Secular Trends, Large Shocks and New Data," Available at https://ssrn.com/abstract=3669854.

Baldwin, R. and B. W. Di Mauro (2020): *Economics in the Time of COVID-19: A New eBook*, VOX CEPR Policy Portal.

Bańbura, M., F. Brenna, M. Ciccarelli, C. Montes-Galdón, and J. Paredes (2020): "Forecasting with Conditional Uncertainty," Unpublished.

Baqaee, D. and E. Farhi (2020): "Supply and Demand in Disaggregated Keynesian Economies with an Application to the COVID-19 Crisis," NBER Working Paper 27152.

Bobeica, E. and B. Hartwig (2021): "The COVID-19 Shock and Challenges for Time Series Models," ECB Working Paper 2558.

Carriero, A., T. E. Clark, and M. Marcellino (2016): "Common Drifting Volatility in Large Bayesian VARs," *Journal of Business & Economic Statistics*, 34, 375–390.

Carriero, A., T. E. Clark, M. Marcellino, and E. Mertens (2021): "Addressing COVID-19 Outliers in BVARs with Stochastic Volatility," Federal Reserve Bank of Cleveland Working Paper 21-02.

Chan, J. C. C. (2017): "Notes on Bayesian Macroeconometrics," Available at https://joshuachan.org/papers/BayesMacro.pdf, accessed at 09.03.2021.

——— (2020): "Large Bayesian VARs: A Flexible Kronecker Error Covariance Structure," *Journal of Business & Economic Statistics*, 38, 68–79.

Chib, S. (1995): "Marginal Likelihood from the Gibbs Output," *Journal of the American Statistical Association*, 90, 1313–1321.

Chib, S. and S. Ramamurthy (2014): "DSGE Models with Student-t Errors," *Econometric Reviews*, 33, 152–171.

Chiu, C.-W. J., H. Mumtaz, and G. Pinter (2017): "Forecasting with VAR Models: Fat Tails and Stochastic Volatility," *International Journal of Forecasting*, 33, 1124–1143.

Christiano, L. J. (2007): "On the Fit of New Keynesian Models: Comment," *Journal of Business & Economic Statistics*, 25, 143–151.

Clark, T. E. and F. Ravazzolo (2015): "Macroeconomic Forecasting Performance under Alternative Specifications of Time-Varying Volatility," *Journal of Applied Econometrics*, 30, 551–575.

Cogley, T. and T. J. Sargent (2005): "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S." *Review of Economic Dynamics*, 8, 262–302.

Cúrdia, V., M. Del Negro, and D. L. Greenwald (2014): "Rare Shocks, Great Recessions," *Journal of Applied Econometrics*, 29, 1031–1052.

del Rio-Chanona, R. M., P. Mealy, A. Pichler, F. Lafond, and J. D. Farmer (2020): "Supply and Demand Shocks in the COVID-19 Pandemic: An Industry and Occupation Perspective," *Oxford Review of Economic Policy*, 36, 94–137.

Eltoft, T., T. Kim, and T.-W. Lee (2006): "On the Multivariate Laplace Distribution," *IEEE Signal Processing Letters*, 13, 300–303.

Geweke, J. (1993): "Bayesian Treatment of the Independent Student-t Linear Model," *Journal of Applied Econometrics*, 8, 19–40.

Giannone, D., M. Lenza, and G. E. Primiceri (2015): "Prior Selection for Vector Autoregressions," *The Review of Economics and Statistics*, 97, 436–451.

——— (2019): "Priors for the Long Run," *Journal of the American Statistical Association*, 114, 565–580.

Guerrieri, V., G. Lorenzoni, L. Straub, and I. Werning (2020): "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" NBER Working Paper 26918.

Hartwig, B. (2020): "Robust Inference in Time-Varying Structural VAR Models: The DC-Cholesky Multivariate Stochastic Volatility Model," Deutsche Bundesbank Discussion Paper Series 34/2020.

Huber, F., G. Koop, L. Onorante, M. Pfarrhofer, and J. Schreiner (2020): "Nowcasting in a pandemic using non-parametric mixed frequency VARs," *Journal of Econometrics*.

Jacquier, E., N. G. Polson, and P. E. Rossi (2004): "Bayesian Analysis of Stochastic Volatility Models with Fat-Tails and Correlated Errors," *Journal of Econometrics*, 122, 185–212.

Karlsson, S. (2013): ""Chapter 15 - Forecasting with Bayesian Vector Autoregression"," in *Handbook of Economic Forecasting*, ed. by G. Elliott and A. Timmermann, Elsevier, vol. 2 of *Handbook of Economic Forecasting*, 791 – 897.

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

Lenza, M. and G. E. Primiceri (2020): "How to Estimate a VAR after March 2020," NBER Working Paper 27771.

Litterman, R. B. (1986): "Forecasting with Bayesian Vector Autoregressions: Five Years of Experience," *Journal of Business & Economic Statistics*, 4, 25–38.

Ng, S. (2021): "Modeling Macroeconomic Variations After COVID-19 ," NBER Working Paper 29060.

Primiceri, G. E. and A. Tambalotti (2020): "Macroeconomic Forecasting in the Time of COVID-19," Available at https://faculty.wcas.northwestern.edu/~gep575/PredictionCovid1-5.pdf, accessed at 09.03.2021.

Rousseeuw, P. J. and C. Croux (1993): "Alternatives to the Median Absolute Deviation," *Journal of the American Statistical Association*, 88, 1273–1283.

Schorfheide, F. and D. Song (2020): "Real-Time Forecasting with a (Standard) Mixed-Frequency VAR During a Pandemic," Federal Reserve Bank of Philadelphia Working Paper 20-26.

Sims, C. A. and T. Zha (1998): "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, 949–968.

Stock, J. H. and M. W. Watson (2016): "Core Inflation and Trend Inflation," *The Review of Economics and Statistics*, 98, 770–784.

Villani, M. (2009): "Steady-State Priors for Vector Autoregressions," *Journal of Applied Econometrics*, 24, 630–650.

A Estimation Details of the BVAR-LP

The BVAR of Lenza and Primiceri (2020) is estimated from 1988:Q4 to 2020:Q4 and 2021:Q1 (t^*) is chosen to be the start of the pandemic episode in which volatility is scaled by a common factor s_{t^*}, otherwise s_t is equal to one. Specifically, $s_{t^*} = \bar{s}_0$, $s_{t^*+1} = \bar{s}_1$, $s_{t^*+2} = \bar{s}_2$, and $s_{t^*+j} = 1 + (\bar{s}_2 - 1)\bar{\rho}^{j-2}$. For the VAR coefficients, the Minnesota calibration based on the MAD(qReg) is used and the hyperparameter $\kappa_{1,m}$ is fixed at its optimal pre-pandemic value and not re-optimized using the method of Giannone et al. (2015) to ensure comparability with the other BVAR models.

Figure 7 show estimated hyperparameters that characterize the explicit volatility process during the COVID-19 pandemic. The posteriors of $\bar{s}_0$, $\bar{s}_1$ and $\bar{s}_2$ peak around 2.4, 37 and 28 which implies a substantially higher volatility spikes than indicated by the BVAR-t, BVAR-$\mathcal{N}$-CSV and BVAR-t-CSV in 2020:Q2. The posterior of $\bar{\rho}$ is centered around 0.35, implying an extremely short-lived volatility process in the aftermath of the pandemic.

Figure 7: Volatility parameters of the BVAR-LP

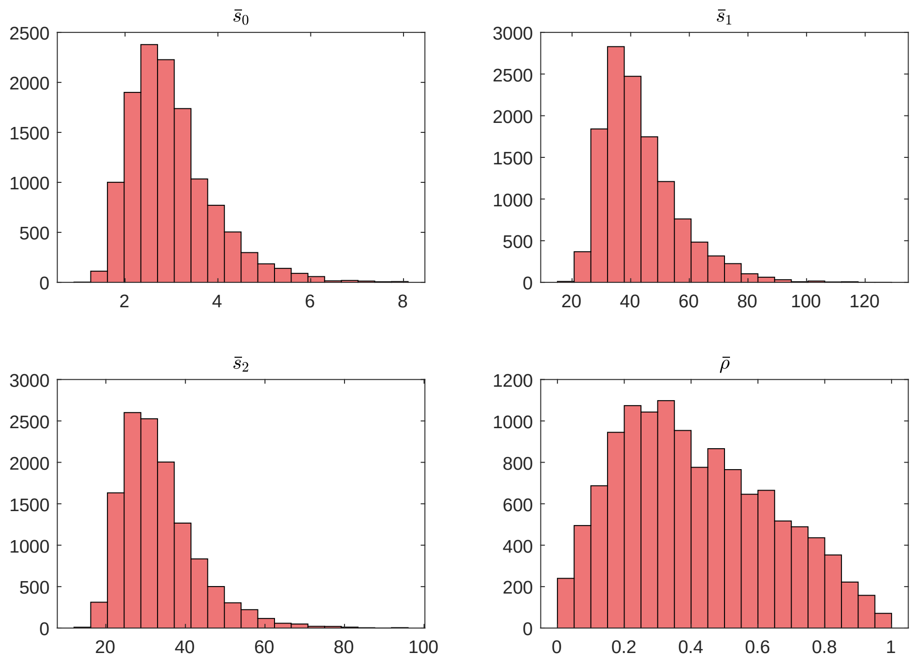

Notes: The figure shows posterior distribution of the volatility scaling factors for 2020:Q1 ($\bar{s}_0$), 2020:Q2 ($\bar{s}_1$), 2020:Q3 ($\bar{s}_2$) and the volatility decay factor $\bar{\rho}$.

B Model (In)stability as of 2020:Q3

Figure 8 shows estimated hyperparameters that characterize the explicit volatility process during the COVID pandemic. The posteriors of $\bar{s}_0$, $\bar{s}_1$ and $\bar{s}_2$ peak around 2.4, 31 and 26 which implies a substantially higher volatility spikes than indicated by the BVAR-t, BVAR-$\mathcal{N}$-CSV and BVAR-t-CSV in 2020:Q2. The posterior of $\bar{\rho}$ is centered around 0.8, implying a rather persistent volatility process during the pandemic. Note, however, the COVID-19 volatility is deterministic and therefore fades quicker off than that of the common stochastic volatility model.

Figure 8: Volatility parameters of the BVAR-LP

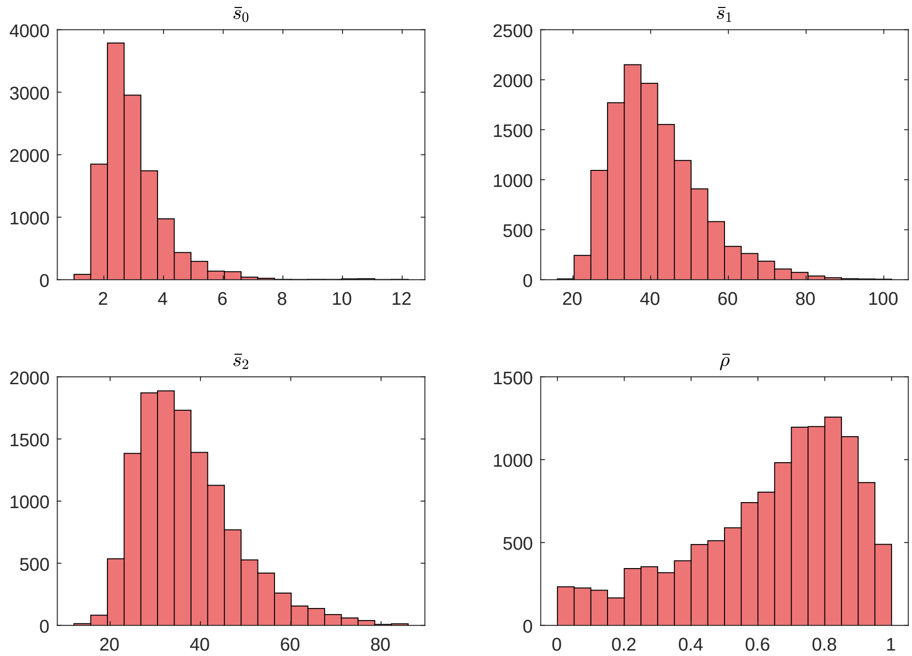

Notes: The figure shows posterior distribution of the volatility scaling factors for 2020:Q1 ($\bar{s}_0$), 2020:Q2 ($\bar{s}_1$), 2020:Q3 ($\bar{s}_2$) and the volatility decay factor $\bar{\rho}$.

Figure 9 shows unconditional forecasts starting in 2020:Q4 onwards and Figure 10 shows conditional forecasts starting in 2020:Q4 based on the unemployment projection from the October 2020 release of the consensus SPF forecast.[24]

[24] The 2020:Q4 forecast for unemployment is 7 percentage points and prolonged using a smooth exponential curve to the long-run forecast of 4.6. https://www.philadelphiafed.org/-/media/frbp/assets/surveys-and-data/survey-of-professional-forecasters/2020/spfq420.pdf

Figure 9: Unconditional forecasts with re-estimated parameters

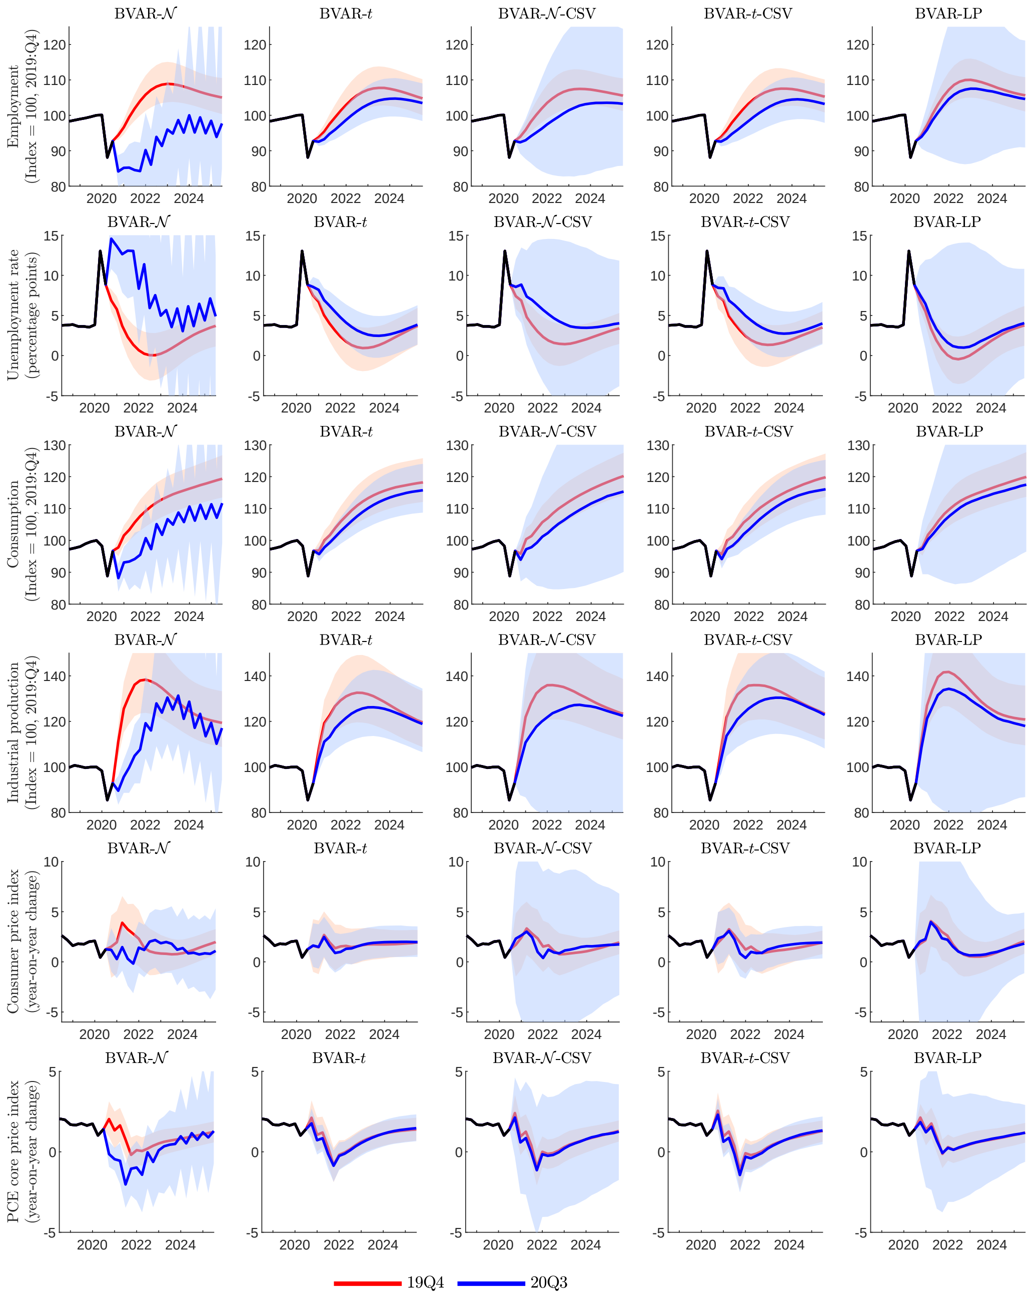

Notes: The figure shows unconditional forecasts starting in 2020:Q4 of various BVAR models estimated until 2019:Q4 and 2020:Q3 respectively. The thick lines are posterior mean estimates and the colored area depicts the 16% to 84% credible interval.

Figure 10: Conditional forecasts with re-estimated parameters

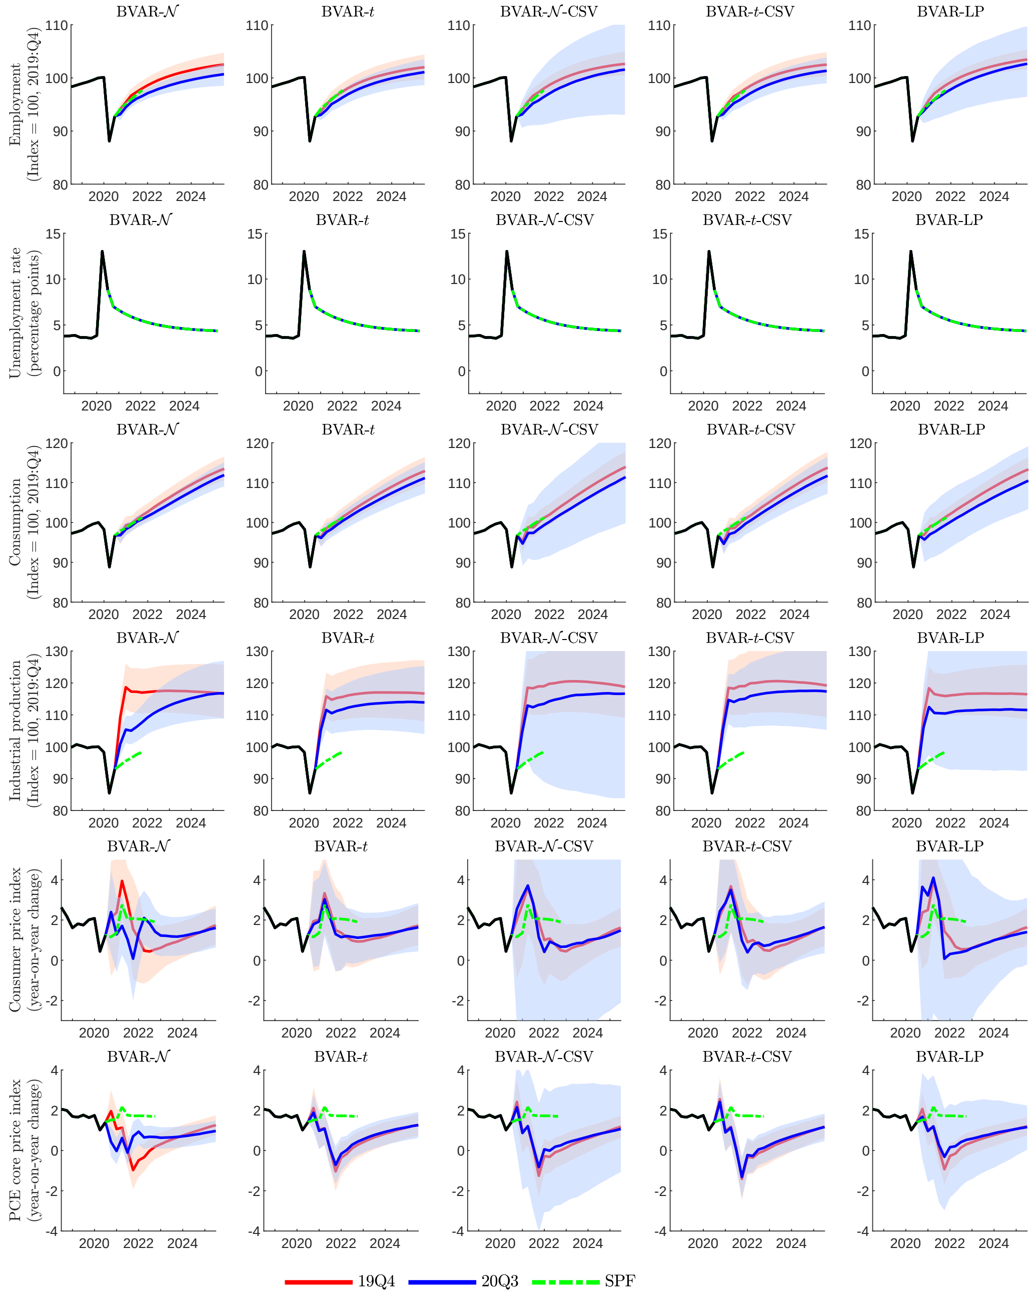

Notes: The figure shows conditional forecasts starting in 2020:Q4 of various BVAR models estimated until 2019:Q4 and 2020:Q3 respectively. The thick lines are posterior mean estimates and the colored area depicts the 16% to 84% credible interval.