WORKING PAPER 1/2022 (ECONOMICS AND STATISTICS)

Modelling Okun's Law – Does non-Gaussianity Matter?

Tamas Kiss, Hoang Nguyen and Pär Österholm

ISSN 1403-0586

Örebro University School of Business
SE-701 82 Örebro, SWEDEN

Modelling Okun's Law – Does non-Gaussianity Matter?*

Tamás Kiss[#]
Division of Economics, School of Business, Örebro University

Hoang Nguyen[∀]
Division of Statistics, School of Business, Örebro University

Pär Österholm[◊]
Division of Economics, School of Business, Örebro University
National Institute of Economic Research

Abstract

In this paper, we analyse Okun's law – a relation between the change in the unemployment rate and GDP growth – using data from Australia, the euro area, the United Kingdom and the United States. More specifically, we assess the relevance of non-Gaussianity when modelling the relation. This is done in a Bayesian VAR framework with stochastic volatility where we allow the different models' error distributions to have heavier-than-Gaussian tails and skewness. Our results indicate that accounting for heavy tails yields improvements over a Gaussian specification in some cases, whereas skewness appears less fruitful. In terms of dynamic effects, a shock to GDP growth has robustly negative effects on the change in the unemployment rate in all four economies.

JEL Classification: C11; C32; C52; E32

Keywords: Bayesian VAR; Heavy tails; GDP growth; Unemployment

---

[*] The authors gratefully acknowledge financial support from Jan Wallanders och Tom Hedelius stiftelse (grants number Bv18-0018, P18-0201 and W19-0021).

[#] Corresponding author. Örebro University, School of Business, 701 82 Örebro, Sweden
e-mail: tamas.kiss@oru.se

[∀] e-mail: hoang.nguyen@oru.se

[◊] e-mail: par.osterholm@oru.se

1. Introduction

Okun's law is a key macroeconomic relation which has become a popular tool for analysis and forecasting since its introduction almost sixty years ago (Okun, 1962). Typically relating the change in the unemployment rate to GDP growth,[1] a fairly large literature has analysed various aspects of it, such as its stability over time, its forecasting properties or its validity in different countries; see, for example, Knotek (2007), IMF (2010), Meyer and Tasci (2012), Owyang and Sekhposyan (2012), Rülke (2012), Zanin and Marra (2012), Huang and Yeh (2013), Valadkhani (2015), Economou och Psarianos (2016), Ball *et al*. (2017), Grant (2018), An *et al*. (2019), Ball *et al*. (2019) and Karlsson and Österholm (2020) for some fairly recent contributions. Conclusions regarding the properties of the relation differ somewhat depending on the country and period studied, but Ball *et al*. (2017, p. 1439) nevertheless suggest that Okun's law "… *is strong and stable by the standards of macroeconomics*".

In this paper, we extend the literature on Okun's law by investigating the importance of non-Gaussianity when modelling the relationship between the change in unemployment rate and GDP growth. We consider two aspects of non-Gaussianity. The first of these is *heavy tails* (or "fat tails") – an issue that takes its starting point in the observation that many economic variables seem to experience large swings more frequently than what one would expect if the shocks hitting the economy are drawn from a Gaussian distribution; see, for example, Fagiolo *et al*. (2008), Ascari *et al*. (2015), Cross and Poon (2016), Liu (2019) and Kiss and Österholm (2020). The second aspect is that the unconditional distribution of many variables appears to be characterised by *skewness*. Particular interest has often been paid to GDP growth with respect to this issue; see, for example, Neftci (1984), Acemoglu and Scott (1997) and Bekaert and Popov (2019). The topic of non-Gaussianity appears to have gained interest over time. This is perhaps not surprising in light of recent historical events; we have, for instance, seen both the Global Financial Crisis and the crisis associated with the corona pandemic in less than fifteen years.

[1] Another way to specify the relation is to connect the unemployment rate (or unemployment gap) to the output gap.

Heavy tails and/or skewness in the data can be caused by the disturbances of the model having these properties.[2] Using data from Australia, the euro area, the United Kingdom and the Unites States, we assess the relevance of such non-Gaussianity. This is done by estimating bivariate Bayesian VAR models with stochastic volatility under three different assumptions regarding the error distributions: *i)* Gaussian, *ii)* Student's *t* and *iii)* generalized hyperbolic skew Student's *t*, also known as skew-*t*; see McNeil *et al*. (2015). Our econometric setting – which has been recently developed by Karlsson *et al*. (2021) – allows us to conduct formal model comparison based on the marginal likelihoods of the estimated models. We can accordingly make statements regarding how well the different models fit the data based on formal statistical criteria. By conducting this analysis, we contribute to the literature in two distinct ways. First, we make a general contribution concerning the importance of non-Gaussianity when it comes to macroeconomic modelling. Second, we provide international empirical evidence concerning Okun's law in a state-of-the-art econometric setting.

Our main results are the following: We find that the unconditional distributions of both variables for all four economies exhibit non-Gaussianity. Our main analysis is conducted using quarterly data up until 2019Q4, that is, we do not include data from the corona pandemic. The estimated models using these data suggest that allowing for error terms with heavy tails yields substantial improvements over a Gaussian specification for Australia and the euro area. Also for the United States is a *t*-distribution the preferred specification, but its benefits relative to a Gaussian distribution are minor judging by the marginal likelihoods of the estimated models. For the United Kingdom, the specification with Gaussian error terms is the preferred specification. In no case is the model with skew-*t* error terms supported by the data and we conclude that modelling skewness appears less fruitful in this context. As a sensitivity analysis, we also estimate our models with data up until 2021Q2 to see how the large swings associated with the corona crisis affect our results. Results from this exercise indicate – not surprisingly – that support for non-Gaussianity strengthens when these observations are added. Regarding the dynamic relationship between the variables, we find – regardless

[2] Alternatively, time variation in the second moment also can result in heavier than normal tails in the unconditional distribution of the variables, even if error terms are Gaussian. We account for this effect by estimating models with stochastic volatility.

of which period we study – that Okun's law prevails: A shock to GDP growth has robustly negative effects on the change in the unemployment rate in all four economies.

The rest of this paper is organised as follows: In Section 2, we describe the data we use in our analysis. The econometric framework is described in Section 3. We present our results in Section 4; in addition to our main results, we also provide sensitivity analysis where we include the period of the corona pandemic. Finally, Section 5 concludes.

2. Data

The samples we use for the four economies vary with respect to their starting point due to availability of data, but all have the same end date. In our main analysis, the samples are 1978Q3-2019Q4 for Australia, 1995Q2-2019Q4 for the euro area, 1971Q3-2019Q4 for the United Kingdom and 1948Q2-2019Q4 for the United States. We do not include data from 2020 and later since the corona pandemic induced movements in the variables which were so large that they maybe should be considered outliers. This is illustrated in Figure 1 which shows time series of GDP growth (g_t) and the change in the unemployment rate (Δu_t).[3] As can be seen, particularly the swings in GDP growth associated with the corona pandemic were of a magnitude which had never been seen before in the samples considered here. This is also the case for the change in the unemployment rate in the United States. However, for Australia, the euro area and the United Kingdom, the change in the unemployment rate was obviously large but not extreme by historical standards. We assess the importance of excluding the corona-related observations in a sensitivity analysis in Section 4.2.

In order to assess potential non-Gaussianity of the data, we present some key descriptive statistics in Table 1. We also show histograms which illustrate the unconditional distributions of the variables in Figure A1 in the Appendix. The unconditional distribution of the variables is in all cases associated with excess kurtosis. Regarding skewness, this seems fairly modest for GDP growth; it is negative in three out of four economies but for both Australia and the United States, it is quite close to zero. Turning to the skewness of the

[3] GDP growth is given as the percentage change in seasonally adjusted real GDP from the previous quarter. The change in the seasonally adjusted harmonized unemployment rate is given in percentage points.

change in the unemployment rate, this is found to be positive and more substantial in all four economies. The Jarque-Bera test strongly rejects normality in all cases. This provides an initial indication that a departure from a Gaussian distribution might prove useful when modelling the Okun's law relationship empirically.

Figure 1. Data.

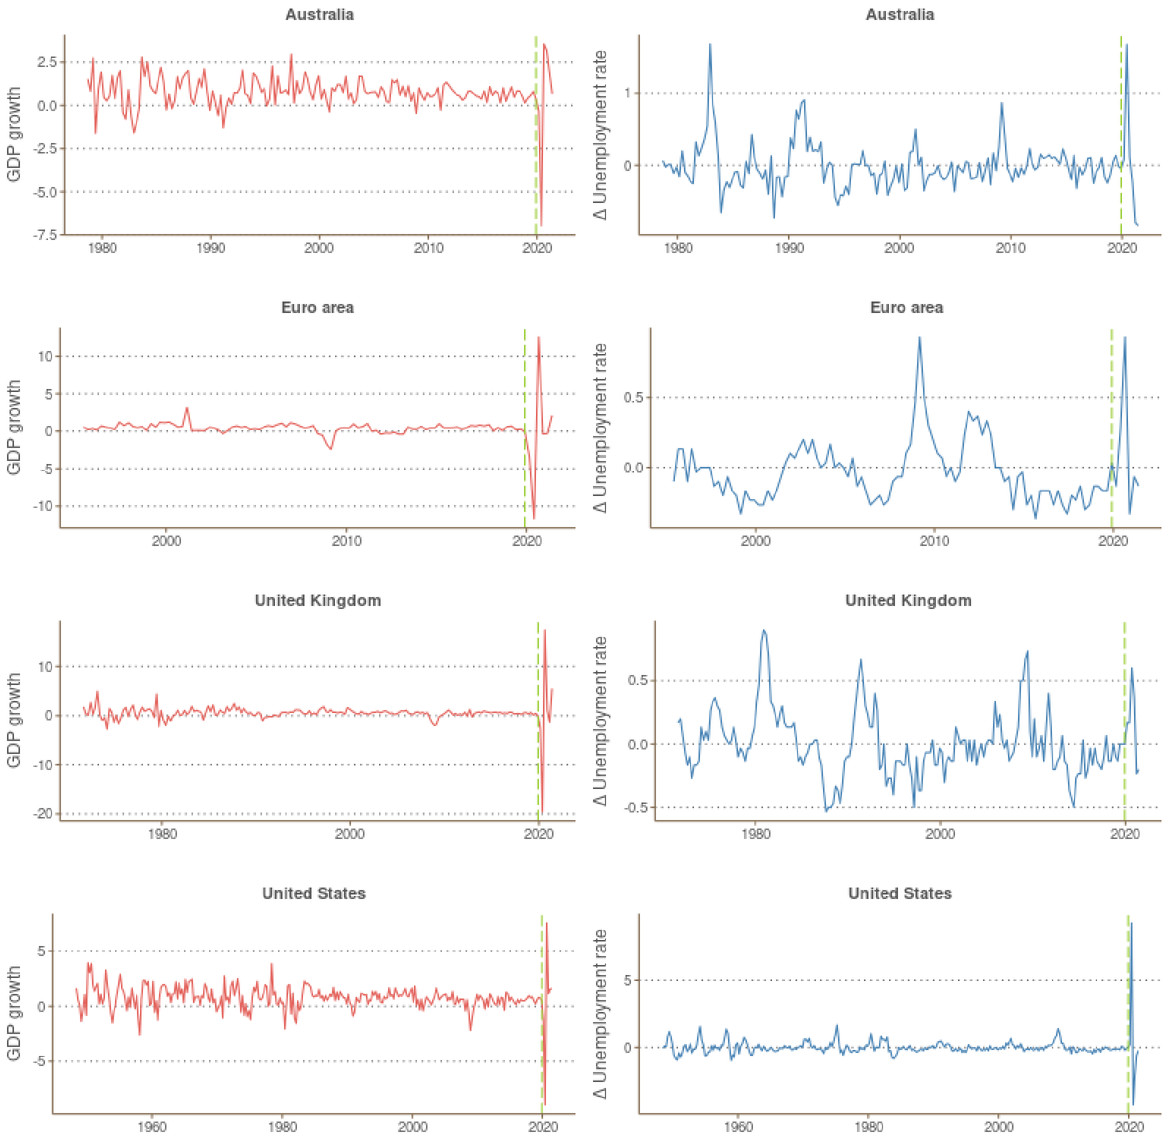

Note: Percent on vertical axis for GDP growth. Percentage points on vertical axis for the change in the unemployment rate. Vertical green dashed line indicates the end of the sample for our main analysis.

Table 1. Descriptive statistics and Jarque-Bera test statistics.

		Mean	Standard deviation	Skewness	Excess kurtosis	Jarque-Bera	Start	End
Australia	g_t	0.773	0.751	-0.053	1.151	10.060	1978Q3	2019Q4
	Δu_t	-0.007	0.300	1.645	6.219	353.491	1978Q3	2019Q4
Euro area	g_t	0.423	0.602	-0.563	8.673	333.042	1995Q2	2019Q4
	Δu_t	-0.034	0.213	1.343	3.128	74.351	1995Q2	2019Q4
United Kingdom	g_t	0.539	0.893	0.265	5.611	264.928	1971Q3	2019Q4
	Δu_t	-0.001	0.256	0.881	1.545	45.976	1971Q3	2019Q4
United States	g_t	0.781	0.936	-0.017	1.602	31.968	1948Q2	2019Q4
	Δu_t	-0.000	0.381	1.256	3.311	210.846	1948Q2	2019Q4

Note: g_t is GDP growth. Δu_t is the change in the unemployment rate. The critical value at the five percent level of the Jarque-Bera test is 5.99.

3. Econometric framework

We rely on bivariate Bayesian VAR(1) models with stochastic volatility for our analysis of Okun's law.[4,5] In that sense, our analysis is closely related to Karlsson and Österholm's (2020) analysis on US data. Unlike Karlsson and Österholm though, we do not allow for time-variation in parameters and, importantly, we have flexible error term distributions that allow for heavy tails and skewness. Denoting $\boldsymbol{y}_t = (g_t, \Delta u_t)'$ for $t = 1, \ldots, T$, we have

$$\boldsymbol{y}_t = \boldsymbol{c} + \boldsymbol{B}\boldsymbol{y}_{t-1} + \boldsymbol{e}_t, \qquad (1)$$

where $\boldsymbol{c}$ is a vector of intercepts and $\boldsymbol{B}$ includes the regression coefficients of the VAR. The error term $\boldsymbol{e}_t$ follows a multivariate skew-*t* distribution with the following stochastic representation,

$$\boldsymbol{e}_t = \boldsymbol{w}_t\boldsymbol{\gamma} + {\boldsymbol{w}_t}^{1/2}\boldsymbol{A}^{-1}{\boldsymbol{H}_t}^{1/2}\boldsymbol{\varepsilon}_t, \qquad (2)$$

where the lower triangular matrix $\boldsymbol{A}$ with unit diagonal contains the structural parameters of the VAR model and $\boldsymbol{w}_t$ is a scalar independent mixing variable drawn from an inverse-gamma

---

[4] Lag length was determined by employing the Schwarz (1978) information criterion to VARs with homoscedastic and Gaussian disturbances, estimated with maximum likelihood. For all four economies, a lag length of one was found optimal.

[5] We use models with stochastic volatility since heteroskedasticity has been shown to be a relevant feature when modelling macroeconomic time series. In a VAR setting, important early contributions include Cogley and Sargent (2005) and Primiceri (2005). Recently Karlsson and Österholm (2020) pointed out that models with constant shock volatility had substantially lower marginal likelihood than models with stochastic volatility when modelling Okun's law in the United States.

distribution with identical scale and shape parameters equal to $\frac{\nu}{2}$, where ν is the degree of freedom; $\boldsymbol{\gamma}$ is the vector of skewness parameters and $\boldsymbol{\varepsilon}_t \sim N(\mathbf{0}, \boldsymbol{I})$. The matrix $\boldsymbol{H}_t = diag(h_{1t}, h_{2t})$ contains the stochastic volatilities of the variables, whose time series evolution is described as

$$\log(h_{it}) = \log(h_{it-1}) + \sigma_i \eta_{it}, \qquad (3)$$

for $i = 1,2$ with $\sigma_i > 0$ and $\eta_{it} \sim N(0,1)$. Finally, $\boldsymbol{w}_t$, $\boldsymbol{H}_t$ and $\boldsymbol{\varepsilon}_t$ are mutually independent. The distribution in (2) allows for both leptokurtic and skewed distributions even after filtering out stochastic volatility. While the mixing variable $\boldsymbol{w}_t$ captures the high frequency shock in mean and variance, the stochastic volatility accounts for the low frequency shocks.

Our proposed specification nests several important models as special cases. Setting $\boldsymbol{\gamma} = 0$, we get the Bayesian VAR model with stochastic volatility and Student's *t*-distributed error terms proposed by Ni and Sun (2005), which has been a workhorse used in empirical modelling of heavy-tailed error terms in the Bayesian VAR context; see, for example, Cross and Poon (2016), Chiu *et al.* (2017), Chan (2020) and Carriero *et al.* (2021). The Gaussian distribution is also nested in this specification ($\boldsymbol{\gamma} = 0, \nu \to \infty$). We accordingly consider three BVAR models with stochastic volatility for each of the economies: the benchmark Gaussian, the Student's *t* and the skew-*t*.

Bayesian estimation requires specifying prior distributions for the parameters. We use a diffuse normal prior (with zero mean and variances 10) for the elements of the lower triangular matrix **A**. We impose a Minnesota prior for the regression coefficients ($\boldsymbol{c}$ and $\boldsymbol{B}$) with overall shrinkage $l_1 = 0.2$ and cross-variable shrinkage $l_2 = 0.5$ (Koop and Korobilis, 2010). The priors for the rest of the parameters are given by $\nu \sim \mathcal{G}(2,0.1)$, $\gamma_i \sim N(0,1)$ for $i = 1,2$, and $\sigma_i^2 \sim \mathcal{G}(0.5,0.5)$, where $\mathcal{G}(a, b)$ is a gamma distribution with shape and rate parameters a and b (Kastner and Frühwirth-Schnatter, 2014).

As the error term is written in terms of a variance-mean mixture distribution, it is straightforward to make inference on the model parameters based on the Gibbs sampler of the VAR model with Gaussian stochastic volatility. For example, conditional on $\boldsymbol{w}_t$ for $t = 1, \dots, T$, the conditional posteriors of parameters ($\boldsymbol{c}, \boldsymbol{B}, \boldsymbol{\gamma}, \boldsymbol{A}, \boldsymbol{\sigma}^2$) are conjugate with the prior distribution (Clark, 2011). We sample the mixing variable $\boldsymbol{w}_t$ based on the generalized inverse Gaussian distribution and sample the degrees of freedom ν based on a random walk Metropolis Hastings (Karlsson *et al*., 2021).

In order to compare different specifications of the VAR model with stochastic volatility, we calculate the marginal likelihood based on the cross-entropy method of Chan and Eisenstat (2018). The marginal likelihood provides us with a measure of how well the model and the priors agree with the data, where the model with the highest marginal likelihood is the one preferred by the data. The marginal likelihood requires a high dimensional integration over the fixed parameters $\boldsymbol{\theta} = (\boldsymbol{c}, \boldsymbol{B}, \boldsymbol{\gamma}, \boldsymbol{A}, \boldsymbol{\sigma}^2, \boldsymbol{\nu})$ and the latent states $\boldsymbol{\varphi} = (\boldsymbol{h}_{1:T})$,

$$p(\boldsymbol{y}_{1:T}) = \int p(\boldsymbol{y}_{1:T} | \boldsymbol{\theta})\, p(\boldsymbol{\theta}) d\boldsymbol{\theta} \approx \sum \frac{p(\boldsymbol{y}_{1:T} | \boldsymbol{\theta})\, p(\boldsymbol{\theta})}{f(\boldsymbol{\theta})},$$

$$p(\boldsymbol{y}_{1:T} | \boldsymbol{\theta}) = \int p(\boldsymbol{y}_{1:T} | \boldsymbol{\varphi}, \boldsymbol{\theta})\, p(\boldsymbol{\varphi} | \boldsymbol{\theta})\, d\boldsymbol{\varphi} \approx \sum \frac{p(\boldsymbol{y}_{1:T} | \boldsymbol{\varphi}, \boldsymbol{\theta})\, p(\boldsymbol{\varphi} | \boldsymbol{\theta})}{g(\boldsymbol{\varphi})}.$$

Following Chan and Eisenstat (2018), we use a two-stage importance sampling to calculate the marginal likelihood. In the first stage, we use the cross-entropy method to learn the proposal distribution of the fixed parameters $f(\boldsymbol{\theta})$ based on the posterior samples. Then, we obtain N = 20,000 proposal samples from $f(\boldsymbol{\theta})$ and calculate the integrated likelihood $p(\boldsymbol{y}_{1:T} | \boldsymbol{\theta})$ for each sample of $\boldsymbol{\theta}$ based on an inner importance sampling loop. The proposal distribution of the latent states $g(\boldsymbol{\varphi})$ is based on a sparse matrix representation. For further details concerning posterior inference and marginal likelihood calculations, see Karlsson *et al*. (2021)

4. Results

We initially present results based on our main sample, that is, where the last observation of each sample is 2019Q4. In Section 4.2, we present sensitivity analysis related to the highly volatile period associated with the corona pandemic.

4.1 Main results

The log marginal likelihoods of the estimated models are presented in Table 2 together with posterior means of the degrees of freedom and skewness parameters.

Table 2. Log marginal likelihoods and estimated key parameters.

	Gaussian	Student's t	Skew-t
Australia	-170.39	-168.04	-171.01
ν	-	14.61	24.92
γ_g	-	-	0.18
$\gamma_{\Delta u}$	-	-	0.15
Euro area	-5.17	-0.40	-2.71
ν	-	11.12	22.09
γ_g	-	-	0.35
$\gamma_{\Delta u}$	-	-	0.04
United Kingdom	-85.90	-86.20	-89.43
ν	-	29.03	38.05
γ_g	-	-	0.170
$\gamma_{\Delta u}$	-	-	-0.02
United States	-317.58	-317.31	-319.08
ν	-	26.67	32.76
γ_g	-	-	-0.07
$\gamma_{\Delta u}$	-	-	0.13

Note: Last observation of each sample is 2019Q4. Log marginal likelihoods are calculated using the cross-entropy methods by Chan and Eisenstat (2018). ν is the degrees of freedom. γ_g and $\gamma_{\Delta u}$ are skewness parameters for GDP growth and the change in the unemployment rate respectively.

The log marginal likelihoods suggest that it is beneficial to take heavy tails into account for Australia, the euro area and the United States. For Australia, the support for the Student's t-distribution is "*positive*" against both other models when we use the scale of two times the difference in the log marginal likelihood and the terminology of Kass and Raftery (1995, p. 777). For the euro area, the support for the Student's t-distribution is "*positive*" against the skew-t distribution and "*strong*" against the Gaussian. For the Unites States, the support for the Student's t-distribution is "*not worth more than a bare mention*" when compared to the Gaussian and "*positive*" against the skew-t. Turning to the United Kingdom, we find that the Gaussian model is preferred. The support for it is "*not worth more than a bare mention*" though when compared to the t-distribution and "positive" relative to the skew-t.

These results are also reflected in the estimates of the parameters which describe the shapes of the distributions. For the models with a Student's t-distribution, the estimated degrees of freedom is relatively low for Australia (14.61) and the Euro-area (11.12) signalling modestly heavy tails for the distribution of the error terms. For the United Kingdom and the United States, the estimated degrees of freedom are substantially higher (29.03 and 26.67, respectively). With such high degrees of freedom, the distribution of the error terms is empirically indistinguishable from the Gaussian, which is also reflected in the results for the log marginal likelihoods. In case of the skew-t models, the degrees of freedom parameters are higher for all four economies, hence distributions become less heavy-tailed. It suggests that the allowing for asymmetry helps capture some of the larger movements in the variables. Looking at the asymmetry parameter $\boldsymbol{\gamma}$, we see a positive skewness in most variables; the only exceptions are the change in the unemployment rate for the United Kingdom and GDP growth for the United States, where the estimated skewness of the error terms is negative. Overall though, the estimates of the asymmetry parameter are small in magnitude. Hence, the evidence in favour of allowing for skewness is weak.

We conclude that it seems reasonable to rely on a Student's t-distribution when modelling Australia and the euro area. For the United Kingdom and the United States, both the Gaussian and Student's t-distribution seem like acceptable choices.

Having focused on the question of error distributions so far, we next attract our attention to a key aspect of Okun's law in this framework, namely how the change in the unemployment rate responds to an unexpected increase in GDP growth. These impulse-response functions are presented in Figure 2. For consistency – and comparability – we have used the model based on Student's t-distributed errors for all four economies when conducting this analysis (even though it was not the best model for the United Kingdom).

As can be seen, the response is negative contemporaneously, and remains negative (or zero) over the entire ten-quarter horizon, in all four economies. We note though that the effect of the shock appears somewhat longer lasting in the euro area and the United Kingdom. In light of higher GDP growth than expected, we would accordingly revise our forecast of the unemployment rate downwards. This is in line with our expectations given previous research on Okun's law.

Figure 2 does not give any indication regarding the uncertainty associated with the impulse-response functions. In order to illustrate this, we show the impulse-response functions for all economies at 2019Q4 together with the 90 percent credible interval in Figure 3. At short horizons, the interval does in no case cover the zero line and we conclude that there is indeed a negative effect on the change in the unemployment rate from a shock to GDP growth.

Figure 2. Impulse-response functions. Response of change in the unemployment rate to a shock to GDP growth.

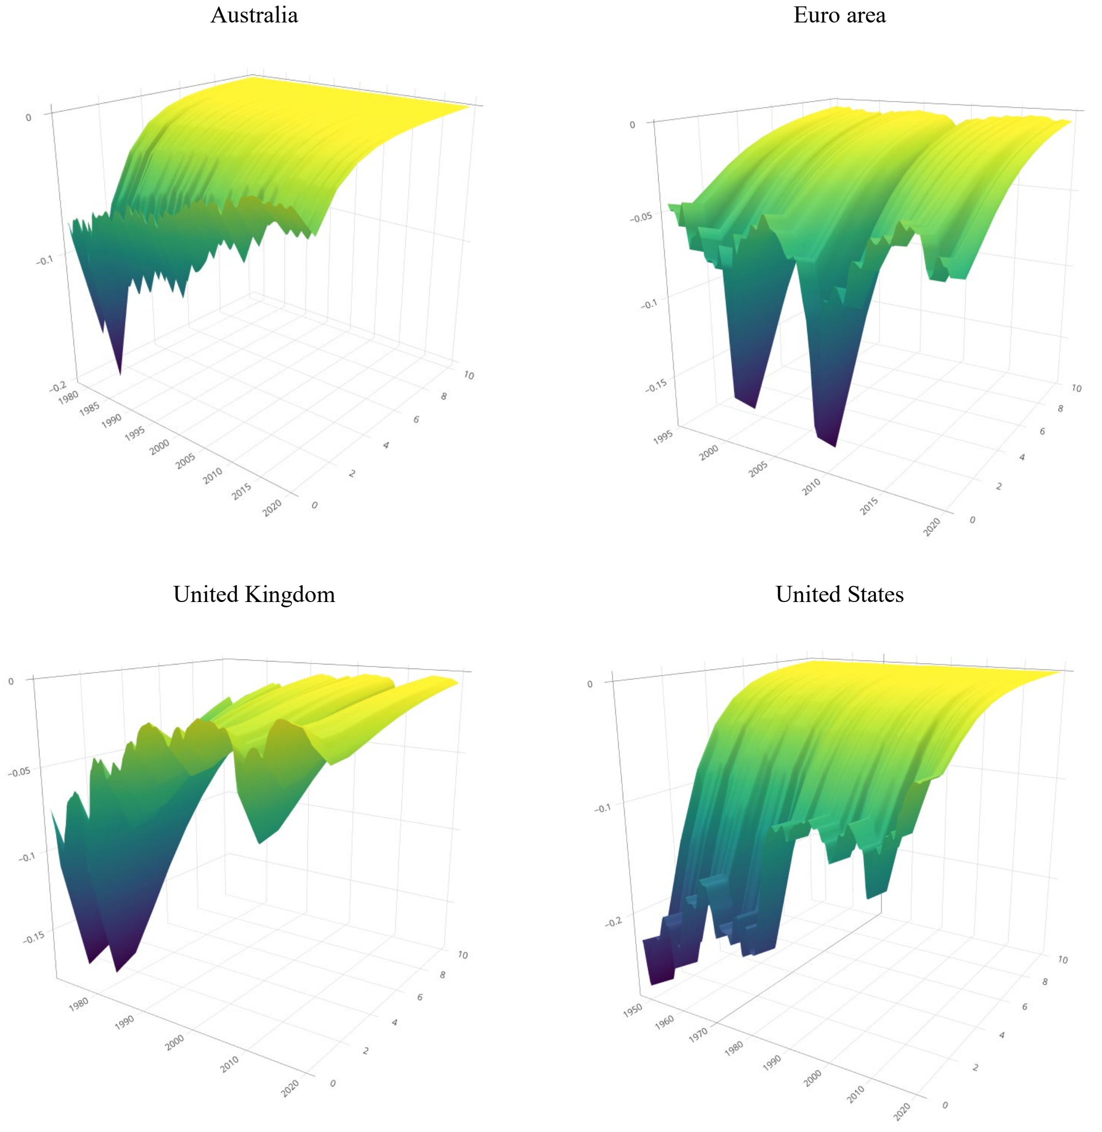

Note: The impulse response functions are based on the model with Student's *t*-distributed errors. Percentage points on the vertical axis. Horizon is given in quarters on the horizontal axes.

Figure 3. Impulse-response functions at 2019Q4. Response of change in the unemployment rate to a shock to GDP growth.

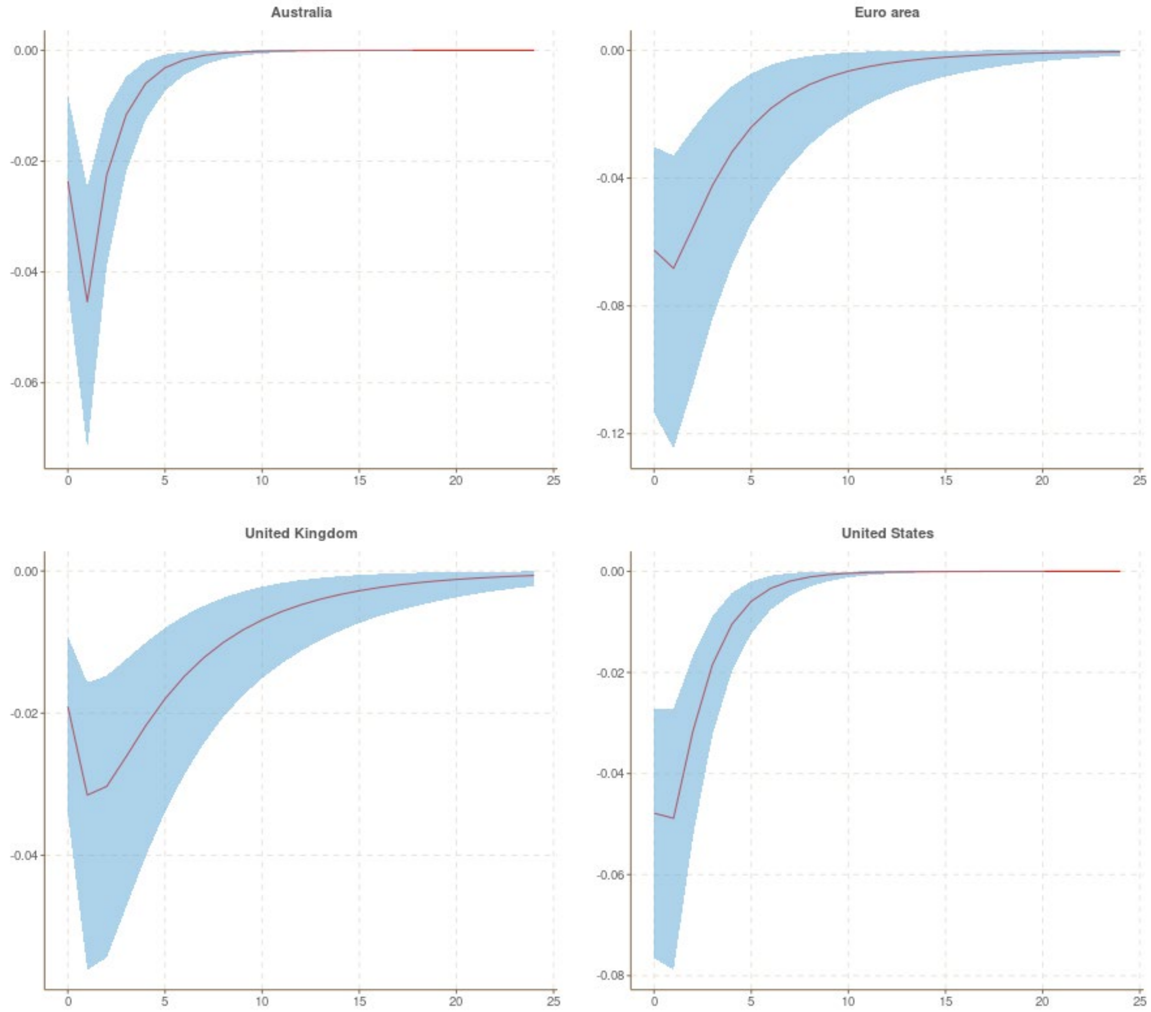

Note: The impulse response functions are based on the model with Student's *t*-distributed errors. Percentage points on the vertical axis. Horizon is given in quarters on the horizontal axes. Coloured band gives 90 percent credible interval.

Returning to Figure 2, it is striking how the magnitude of the impulse response changes considerably over the sample period. Since the regression parameters of the model are constant, these changes are solely attributed to stochastic volatility. Since this is another important feature of the employed modelling framework, we present the posterior mean of the time-varying variance of the shocks to both variables in Figure 4. In order to illustrate the difference between a Gaussian and a Student's *t*-distribution, we provide the estimated variances under both assumptions.

Figure 4. Variance estimates from VAR models.

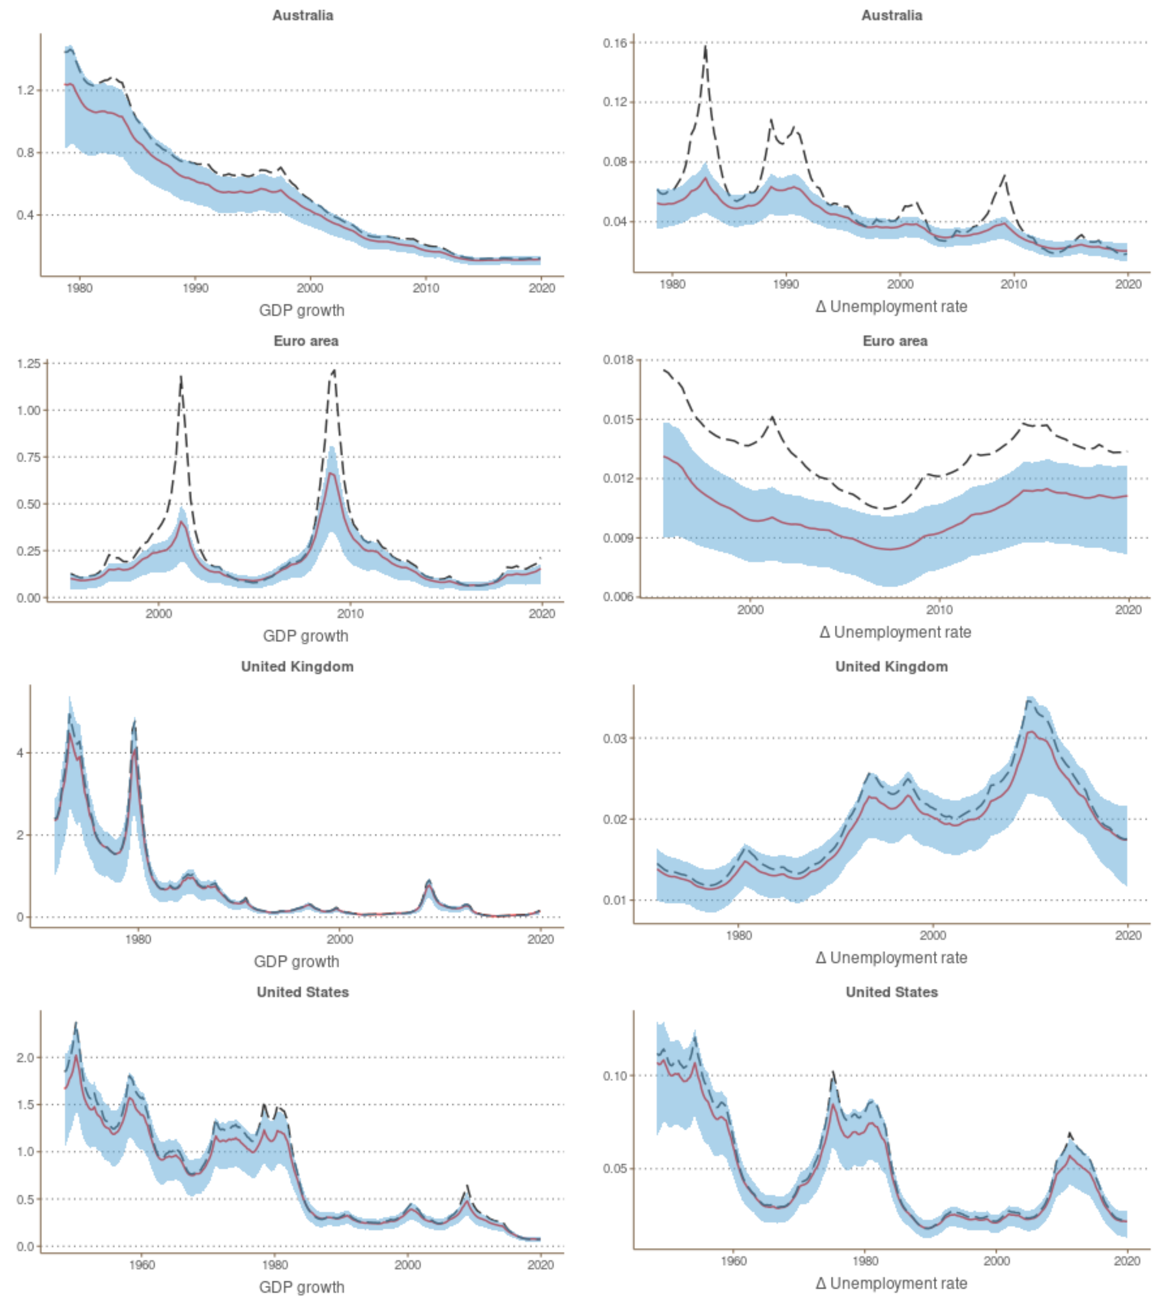

Note: The red solid line gives the volatility under the assumption of a Student's *t*-distribution with the coloured bands showing the 50 percent credible interval. The black dashed line gives the volatility under the assumption of a Gaussian distribution.

Regarding the time-varying variances, the patterns obviously differ across economies and variables. However, some features tend to be common. For example, except for Australian GDP growth, there is an increase in volatility around the Global Financial Crisis in 2008. From a modelling perspective, we also note that there is substantial time variation in the estimates of the variances. This shows that it is relevant to use models which account for heteroskedasticity, such as the models employed here.

The variance estimates based on a Gaussian and a Student's *t*-distribution show overall similar patterns, but substantial differences can also be observed in some cases. For example, under the assumption of a Student's *t*-distribution the variance of the change in the unemployment rate in Australia is clearly smoother, and the overall level of the variance is lower in the euro area. Further, the spikes in output growth volatility in crisis periods are more pronounced in the euro area if we use Gaussian error terms. This latter observation can be attributed to the fact that, in the absence of flexible error distributions, the effect of larger swings in the variables appears through increased volatility. However, in line with what we would expect based on the results presented in Table 1, we can also see that some differences are minor. For the United Kingdom and the United States – where the marginal likelihoods of the two models were quite similar and the estimated degrees of freedom of the Student's *t*-distribution high for both variables – estimated volatilities are similar for both GDP growth and the change in the unemployment rate.

4.2 Sensitivity analysis

Our results so far indicate that for some economies, there might be improvements to be made when it comes to modelling Okun's law if error terms are assumed to be drawn from a Student's *t*-distribution. However, our analysis has been based on a sample which excludes the corona pandemic and the economic crisis and recovery which followed it. As pointed out in Section 2, the swings in the variables associated with this period were very large – so large that they perhaps should be considered outliers. These issues associated with modelling the corona pandemic have recently been discussed in the literature; see for example Bobeica and Hartwig (2021), Carriero *et al.* (2021) and Hartwig (2021). Seeing that these data are something that empirical macroeconomists will have to handle in the future, we next assess the effects that they have in the context of the analysis in this paper.

In Table 3, we first provide descriptive statistics and results from the Jarque-Bera test for normality. The effects of the large movements in the variables around the corona pandemic affect higher order moments of the unconditional distribution. The standard deviation of the variables increases somewhat compared to the baseline results in Table 1. A negative skewness of GDP growth and a positive skewness of the change in the unemployment rate also become salient features of the data in all economies (although skewness is overall modest in these variables). The most striking difference compared to the baseline sample, which excludes the corona observations, is found in the excess kurtosis of the variables which shoots up by including observations from 2020 and 2021. The increase is most striking for GDP growth in all economies and the change in the unemployment rate in the United States, in which cases excess kurtosis becomes six to ten times larger.

Table 3. Descriptive statistics and Jarque-Bera test statistics – sample including corona pandemic.

		Mean	Standard deviation	Skewness	Excess kurtosis	Jarque-Bera	Start	End
Australia	g_t	0.757	0.994	-2.518	20.170	3179.736	1978Q3	2021Q2
	Δu_t	-0.006	0.334	1.724	6.935	442.994	1978Q3	2021Q2
Euro area	g_t	0.388	1.833	-0.014	33.988	5268.788	1995Q2	2021Q2
	Δu_t	-0.027	0.231	1.590	3.899	116.631	1995Q2	2021Q2
United Kingdom	g_t	0.525	2.106	-1.684	59.661	30388.72	1971Q3	2021Q2
	Δu_t	0.003	0.258	0.854	1.359	41.093	1971Q3	2021Q2
United States	g_t	0.770	1.165	-1.310	19.262	4686.325	1948Q2	2021Q2
	Δu_t	0.008	0.717	6.712	96.729	118071.7	1948Q2	2021Q2

Note: g_t is GDP growth. Δu_t is the change in the unemployment rate. The critical value at the five percent level of the Jarque-Bera test is 5.99.

Estimation results are also impacted by the strong increase in excess kurtosis. In Table 4, we present the log marginal likelihoods and estimated key parameters from the models when relying on the sample including the observations during the corona pandemic, that is, up until 2021Q2. Considering the log marginal likelihoods first, it can be seen that the Student's *t*-distribution is preferred in three cases, namely for Australia, the United Kingdom and the United States. The strength of the evidence varies though; comparing to the Gaussian model, we find that it is "*very strong*" for Australia and the Unites States but "*not worth more than a bare mention*" for the United Kingdom. For the euro area, the skew-*t* model is now the preferred one and the support in favour of it is "*very strong*" and "*strong*" when compared to the models assuming a Gaussian and a Student's *t*-distribution respectively. As a general tendency, we see that the skew-*t* model is doing much better in this sample; while it is the preferred model only for the euro area, one should recall that for the sample excluding the corona pandemic, it is always the worst model (see Table 1). Not surprisingly, we also see that the support for the Gaussian model declines when using the sample including the corona pandemic.

Table 4. Log marginal likelihoods and estimated key parameters – sample including corona pandemic.

	Gaussian	Student's t	Skew-t
Australia	-213.20	-206.29	-210.59
ν	-	5.35	23.86
$\boldsymbol{\gamma_g}$	-	-	0.72
$\boldsymbol{\gamma_{\Delta u}}$	-	-	0.18
Euro area	-48.79	-46.71	-42.96
ν	-	15.84	29.09
$\boldsymbol{\gamma_g}$	-	-	-0.94
$\boldsymbol{\gamma_{\Delta u}}$	-	-	0.24
United Kingdom	-135.24	-135.17	-136.01
ν	-	23.94	33.85
$\boldsymbol{\gamma_g}$	-	-	0.48
$\boldsymbol{\gamma_{\Delta u}}$	-	-	-0.09
United States	-399.62	-380.73	-387.45
ν	-	3.92	4.31
$\boldsymbol{\gamma_g}$	-	-	-0.02
$\boldsymbol{\gamma_{\Delta u}}$	-	-	0.02

Note: Last observation of each sample is 2019Q4. Log marginal likelihoods are calculated using the cross-entropy methods by Chan and Eisenstat (2018). ν is the degrees of freedom. $\boldsymbol{\gamma_g}$ and $\boldsymbol{\gamma_{\Delta u}}$ are skewness parameters for GDP growth and the change in the unemployment rate respectively.

Turning to the estimated degrees of freedom and skewness parameters, we also see an effect of the pandemic observations. The degrees of freedom radically decrease for Australia and the United States if one considers the model with Student's t-distribution. Interestingly, for the United States it remains low (signalling very heavy tails) even when allowing for skewness. In contrast, allowing for skewness in the case of Australia helps capturing large movements as the degrees of freedom parameter jumps up for the skew-t specification. For the other two economies, the estimated degrees of freedom do not change substantially. The skewness parameters for each country also change somewhat (they tend to increase) but they retain the same sign as in the baseline sample. The only exception is the skewness of GDP growth in the euro area where the sign of the skewness parameter switches from positive to negative.

Another way of illustrating the influence of the observations associated with the corona pandemic is by looking at the posterior distribution of the estimated degrees of freedom. This is shown in Figure 5. First, we can note that the posterior distributions in the euro area and the United Kingdom change only slightly. For the euro area, the sample including corona actually puts somewhat more weight on higher values of degrees of freedom, that is, the tails become less heavy. The changes are more drastic for Australia and the United States though, where the posterior distribution becomes heavily concentrated at low values (which is also reflected in the radically decreased point estimate which is taken to be the posterior mean).

Figure 5. Posterior distributions of the estimated degrees of freedom – samples including and excluding the corona pandemic.

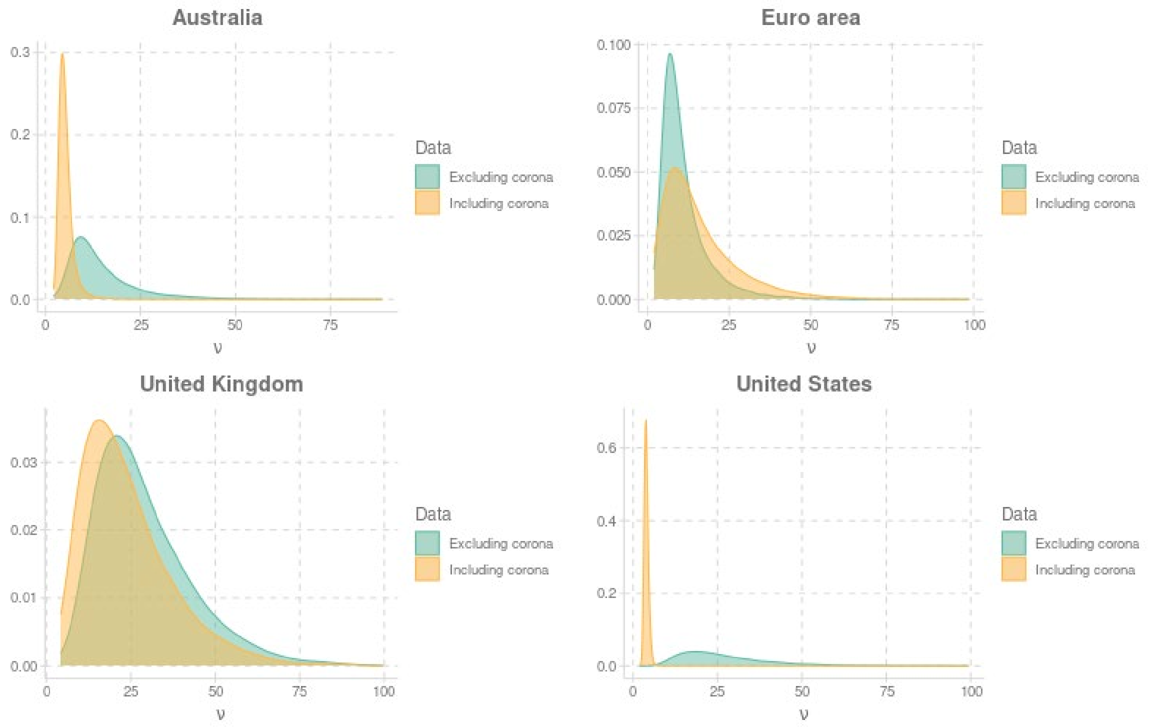

Note: The blue density gives the degrees of freedom based on data which do not include the corona pandemic, that is, the samples end in 2019Q4. The yellow density gives the degrees of freedom based on data which do include the corona pandemic, that is, the samples end in 2021Q2. All densities are based on the model with Student's *t*-distributed error terms.

As we established above, including the corona pandemic has implications for which model is preferred by the data. It also tends to affect the estimated volatilities of the models in a substantial manner, particularly near the end of the sample. Figure 6 shows the estimated log volatilities from the models. We see that in most cases there is a sharp jump in volatility, during 2020 and 2021, often reaching previously unprecedented levels. Note however, that using our baseline sample ending in 2019, the volatility estimates were on a stable low level or slightly on the way down in most cases, signalling tranquil times. This drastically changes when the observations from 2020 and 2021 are included: Not only does the volatility spike during these latter years, but in order to match the high volatility during the crisis associated with the corona pandemic, volatility is also on the rise even a few years before that during the second half of the 2010s. Furthermore, for the United States – where the evidence on heavy tails is also the strongest in the longer sample – we also see that using the Student's *t*-distribution allows the model to capture the large movements around the end of the sample. Using a non-Gaussian specification allows stochastic volatility for both the change in unemployment rate and GDP growth to remain at a modest level (in contrast to the large upward jump in the Gaussian case). A similar effect can be observed for Australia regarding the change in unemployment rate. The volatility estimates for the euro area and the United Kingdom are almost identical regardless of the distributional assumption, which is in line with the fact that the Student's *t*-distribution is less useful for these economies.

Finally, we also look at the impulse-response functions of the change in the unemployment rate to a shock to GDP growth; for comparability with the main specification, we continue to use the model based on a Student's *t*-distribution. We present the impulse-response functions at 2021Q2, that is, the end of the extended sample. These are given in Figure 7. The shape of the impulse responses remains similar to the ones reported in Figure 3. We still find that all the impulse responses start in the negative region and remain significantly negative for several quarters. It can be noted though that the magnitudes are quite different to those in Figure 3. This is of course only to be expected given the much higher volatility in 2021Q2.

Figure 6. Log volatility estimates from VAR models – samples including and excluding the corona pandemic.

Note: The red (blue) solid line gives the logarithm of the volatility of the variables using a model with Gaussian (Student's *t*) error terms up to 2021Q2, that is, including the pandemic period. The orange (green) dashed line gives the logarithm of the volatility of the variables using a model with Gaussian (Student's *t*) error terms up to 2019Q4, that is, excluding the pandemic period.

Figure 7. Impulse-response functions at 2021Q2. Response of change in the unemployment rate to a shock to GDP growth – sample including corona pandemic.

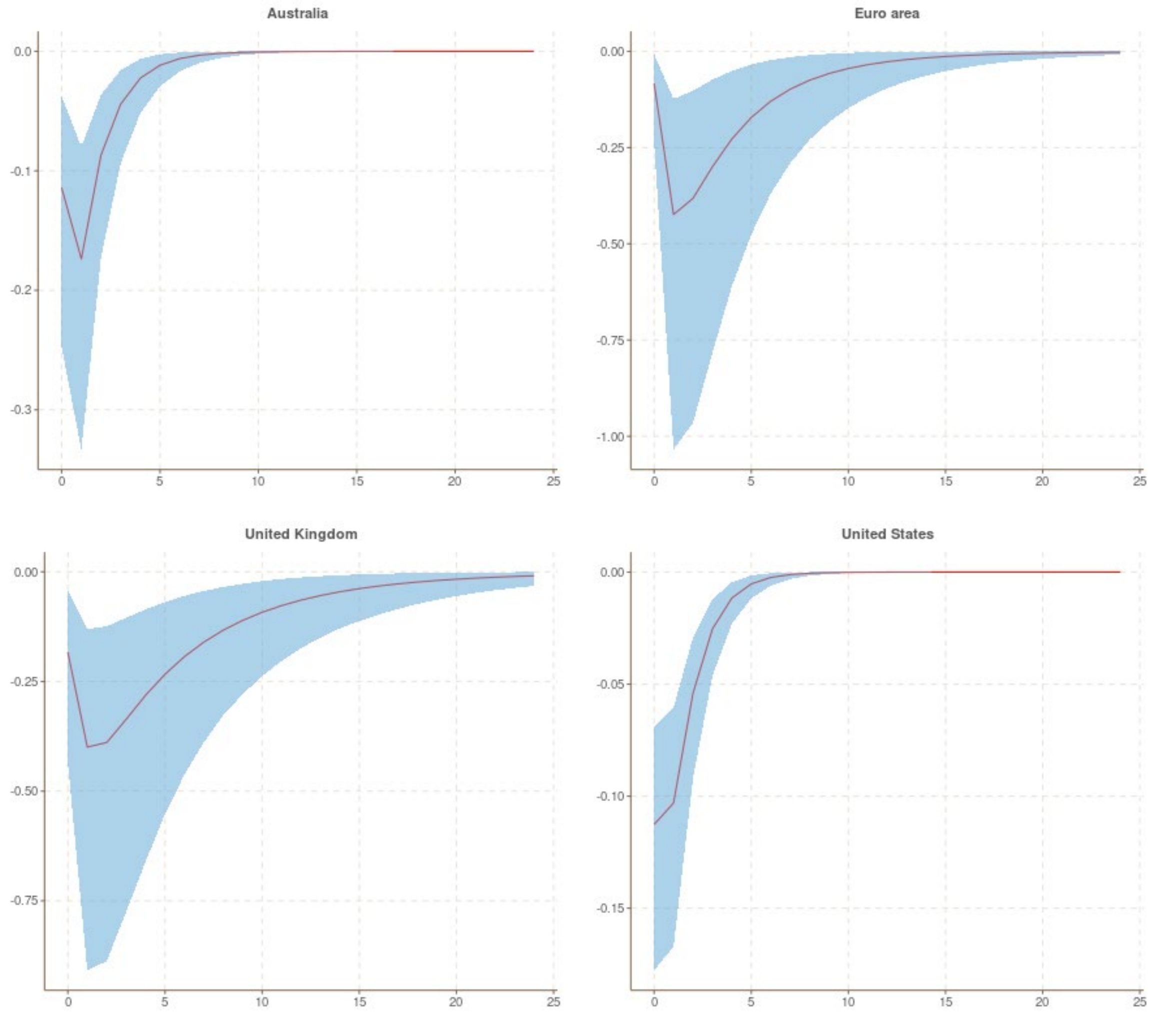

Note: The impulse response functions are based on the model with Student's *t*-distributed errors. Percentage points on the vertical axis. Horizon is given in quarters on the horizontal axes. Coloured band gives 90 percent credible interval.

Summing up this sensitivity analysis, we conclude that including the corona pandemic has non-negligible effects on the results. The large swings in the variables during this time generally result in stronger evidence against Gaussianity. This is supported by the radically decreasing degrees of freedom for the distribution of the error terms for Australia and the United States, and the fact that models with non-Gaussian error terms (either with Student-*t* or Skew-*t* distribution) become the preferred model based on log marginal likelihoods in all four economies. The cases of Australia and the United States also highlight that accounting for heavier tails in the error terms also helps avoiding large jumps in stochastic volatility.

5. Conclusion

In this paper, we have analysed the relevance of taking non-Gaussianity into account when empirically modelling Okun's law in Australia, the euro area, the United Kingdom and the United States. Our results based on Bayesian VAR models with stochastic volatility suggest that heavier-than-Gaussian tails find support in some cases. Taking skewness into account is, however, less beneficial in this context considering our baseline sample. Our results confirm that it is important to account for heavy tails in the distribution of macroeconomic variables, an argument put forward by Fagiolo *et al.* (2008) and Ascari *et al.* (2015) among others.

It should be noted though that our results to some extent depend on whether data from the corona pandemic are included or not. We believe that including them might be problematic since they should probably be treated as outliers (see the discussion in Carriero *et al.*, 2021). If they nevertheless are treated as regular observations, our analysis indicates that the evidence of non-Gaussianity strengthens. In addition, it can be noted that accounting for non-Gaussianity not only improves the model fit in several cases but it also captures the large swings in the variables without causing large swings in the stochastic volatility.

Apart from the modelling perspective, our analysis has also provided updated international empirical evidence concerning Okun's law. We find that the dynamic relationship between the variables in all four economies is such that a shock to GDP growth has robustly negative effects on the change in the unemployment rate. This finding is robust to whether we include the period associated with the corona pandemic or not. It confirms Ball *et al.* (2017) and Ball *et al.* (2019) who argue that Okun's law continues to be a robust relationship in empirical macroeconomics. This should be highly relevant information to the central banks of the economies studied here, suggesting that Okun's law – which has been an important empirical relationship when modelling the economy – continues to be useful regardless of modelling choices and time periods.

References

Acemoglu, D., Ozdaglar, A. and Tahbaz-Salehi, A. (2017), "Microeconomic Origins of Macroeconomic Tail Risks", *American Economic Review* 107, 54-108.

An, Z., Ball, L., Jalles, J. and Loungani, P. (2019), "Do IMF Forecasts Respect Okun's Law? Evidence for Advanced and Developing Economies", *International Journal of Forecasting* 35, 1131-1142.

Ascari, G., Fagiolo, G. and Roventini, A. (2015), "Fat-Tail Distributions and Business-Cycle Models", *Macroeconomic Dynamics* 19, 465-476.

Ball, L., Leigh, D. and Loungani, P. (2017), "Okun's Law: Fit at 50?", *Journal of Money, Credit and Banking* 49, 1413-1441.

Ball, L., Furceri, D., Leigh, D. and Loungani, P. (2019), "Does One Law Fit All? Cross-Country Evidence on Okun's Law", *Open Economies Review* 30, 841-874.

Bekaert, G. and Popov, A. (2019), "On the Link between the Volatility and Skewness of Growth", *IMF Economic Review* 67, 746-790.

Bobeica, E., and Hartwig, B. (2021), "The COVID-19 Shock and Challenges for Time Series Models", ECB Working Papers No. 2558

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

Chan, J. C., and Eisenstat, E. (2018), "Bayesian Model Comparison for Time-varying Parameter VARs with Stochastic Volatility", *Journal of Applied Econometrics* 33, 509-532.

Chan, J. C. (2020), "Large Bayesian VARs: A Flexible Kronecker Error Covariance Structure", *Journal of Business and Economic Statistics* 38, 68-79.

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

Clark, T. E. (2011), "Real-time Density Forecasts from Bayesian Vector Autoregressions with Stochastic Volatility", *Journal of Business and Economic Statistics* 29, 327-341.

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

Cross, J. and Poon, A. (2016), "Forecasting Structural Change and Fat-Tailed Events in Australian Macroeconomic Variables", *Economic Modelling* 58, 34-51.

Economou, A. och Psarianos, I. N. (2016), "Revisiting Okun's Law in European Union Countries", *Journal of Economic Studies* 43, 275-287.

Fagiolo, G., Napoletano, M. and Roventini, A. (2008), "Are Output Growth-Rate Distributions Fat-Tailed?", *Journal of Applied Econometrics* 23, 639-669.

Grant, A. L. (2018), "The Great Recession and Okun's Law", *Economic Modelling* 69, 291-300.

Hartwig, B. (2021), "Bayesian VARs and prior calibration in times of COVID-19", SSRN Working Paper No. 3792070.

Huang, H.-C. and Yeh, C. C. (2013), "Okun's Law in Panels of Countries and States", *Applied Economics* 45, 191-199

IMF (2010), "Unemployment Dynamics during Recessions and Recoveries: Okun's Law and Beyond", *World Economic Outlook* April 2010.

Karlsson, S., Mazur, S. and Nguyen, H. (2021), "Vector Autoregression Models with Skewness and Heavy Tails", Working Paper 8/2021, School of Business, Örebro University.

Karlsson, S. and Österholm, P. (2020), "A Hybrid Time-Varying Parameter Bayesian VAR Analysis of Okun's Law in the United States", *Economics Letters* 197, 109622.

Kass, R.E. and Raftery, A.E. (1995), "Bayes Factors", *Journal of the American Statistical Association* 90, 773-795.

Kastner, G. and Frühwirth-Schnatter, S. (2014), "Ancillarity-Sufficiency Interweaving Strategy (ASIS) for Boosting MCMC Estimation of Stochastic Volatility Models", *Computational Statistics and Data Analysis* 76, 408-423.

Kiss, T. and Österholm, P. (2020), "Fat Tails in Leading Indicators", *Economics Letters* 193, 109317.

Knotek, E. S. (2007), "How Useful is Okun's Law?", *Federal Reserve Bank of Kansas City Economic Review* 92, 73-103.

Koop, G. and Korobilis, D. (2010), *Bayesian Multivariate Time Series Methods for Empirical Macroeconomics*, Now Publishers Inc.

Liu, X. (2019), "On Tail Fatness of Macroeconomic Dynamics", *Journal of Macroeconomics* 62, 103154.

McNeil, A. J., Frey, R. and Embrechts, P. (2015), *Quantitative Risk Management: Concepts, Techniques and Tools (Revised Edition)*, Princeton University Press.

Meyer, B., and Tasci, M. (2012), "An Unstable Okun's Law, not the Best Rule of Thumb", *Economic Commentary*, 2012-08.

Neftci, S. N. (1984), "Are Economic Time Series Asymmetric over the Business Cycle?", *Journal of Political Economy* 92, 307-328.

Ni, S., and Sun, D. (2005), "Bayesian Estimates for Vector Autoregressive Models", *Journal of Business and Economic Statistics* 23, 105-117.

Okun, A.M. (1962), "Potential GNP: Its Measurement and Significance", In: *Proceedings of the Business and Economics Statistics Section*, Washington DC, American Statistical Association.

Owyang, M. T. and Sekhposyan, T. (2012), "Okun's Law over the Business Cycle: Was the Great Recession All That Different?", *Federal Reserve Bank of St. Louis Review* September/October 2012, 399-418.

Primiceri, G. E. (2005), "Time Varying Structural Vector Autoregressions and Monetary Policy", *Review of Economic Studies*, 72(3), 821-852.

Rülke, J.-C. (2012), "Do Professional Forecasters Apply the Phillips Curve and Okun's Law? Evidence from Six Asian-Pacific Countries", *Japan and the World Economy* 24, 317-324.

Schwarz, G. (1978), "Estimating the Dimension of a Model", *Annals of Statistics* 6, 461-464.

Valadkhani, A. (2015), "Okun's Law in Australia", *Economic Record* 91, 509-522.

Zanin, L. and Marra, G. (2012), "Rolling Regression versus Time-Varying Coefficient Modelling: An Empirical Investigation of the Okun's Law in Some Euro Area Countries", *Bulletin of Economic Research* 64, 91-108.

Appendix

Figure A1. Unconditional distributions of the data.

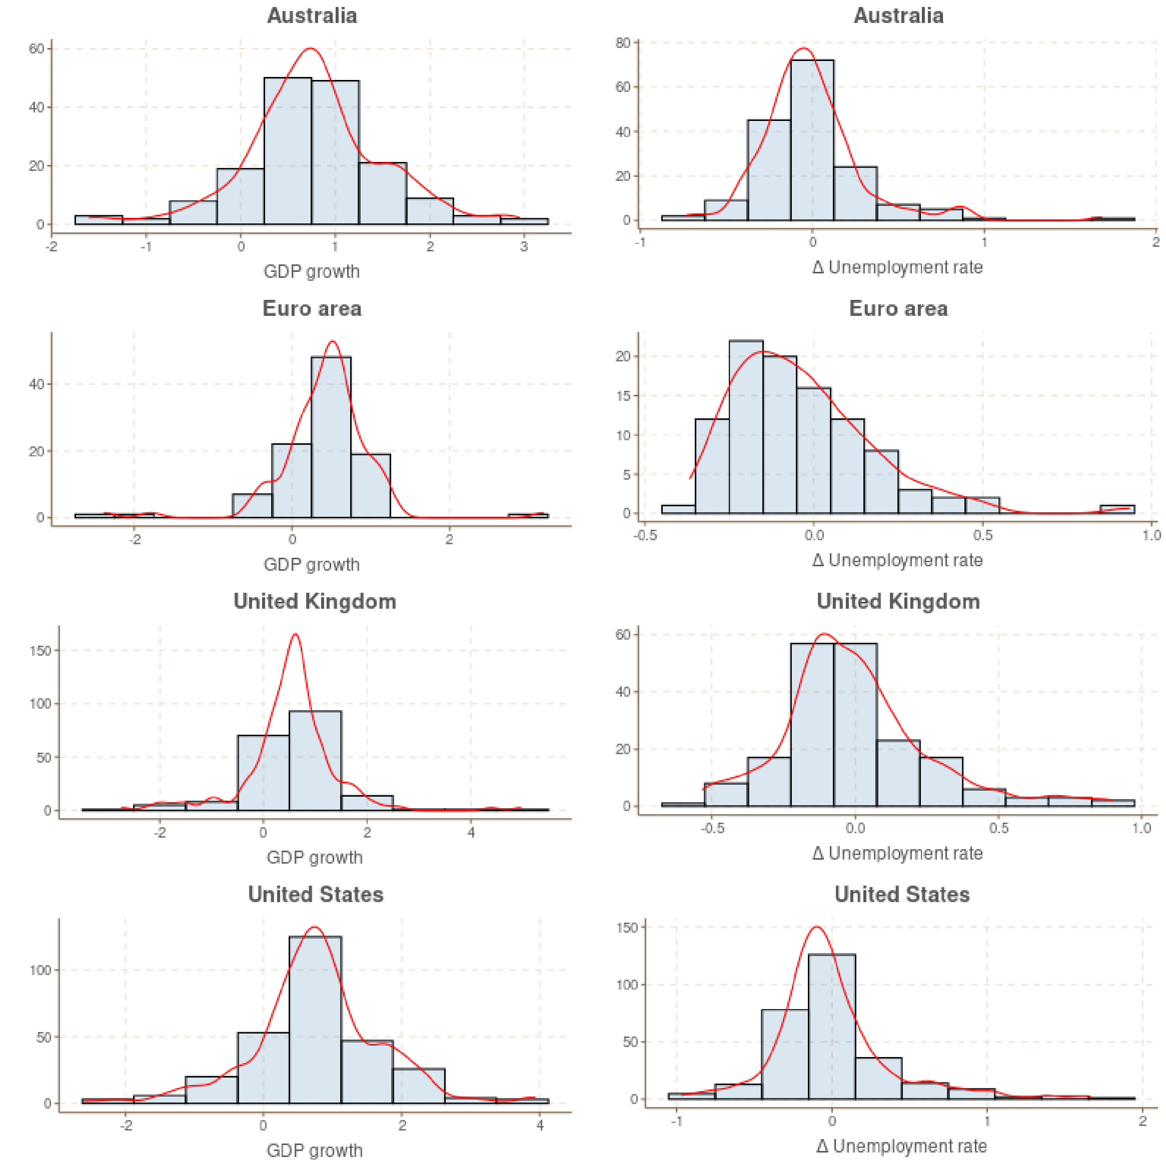

Note: Histograms and smoothed kernel density estimates (red line) of the data series ending in 2019Q4. Frequency on vertical axis. Percent on horizontal axis for GDP growth. Percentage points on horizontal axis for the change in the unemployment rate.