

Bobeica, Elena; Hartwig, Benny

Working Paper

The COVID-19 shock and challenges for time series models

ECB Working Paper, No. 2558

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Bobeica, Elena; Hartwig, Benny (2021) : The COVID-19 shock and challenges for time series models, ECB Working Paper, No. 2558, ISBN 978-92-899-4558-5, European Central Bank (ECB), Frankfurt a. M., <https://doi.org/10.2866/957422>

This Version is available at:

<https://hdl.handle.net/10419/237697>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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



EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Elena Bobeica, Benny Hartwig The COVID-19 shock and challenges
for time series models

No 2558 / May 2021

Abstract

We document the impact of COVID-19 on frequently employed time series models, with a focus on euro area inflation. We show that for both single equation models (Phillips curves) and Vector Autoregressions (VARs) estimated parameters change notably with the pandemic. In a VAR, allowing the errors to have a distribution with fatter tails than the Gaussian one equips the model to better deal with the COVID-19 shock. A standard Gaussian VAR can still be used for producing conditional forecasts when relevant off-model information is used. We illustrate this by conditioning on official projections for a set of variables, but also by tilting to expectations from the Survey of Professional Forecasters. For Phillips curves, averaging across many conditional forecasts in a thick modelling framework offers some hedge against parameter instability.

JEL: C53, E31, E37

Keywords: COVID-19, Forecasting, Student's t errors, tilting, inflation, VAR

Non-technical summary

We analyse the impact of COVID-19 observations on two widely employed time series models, with a focus on euro area inflation. We investigate single equation models (Phillips curves) and Vector Autoregressions (VARs). We show that in both cases the parameter estimates change notably after adding the COVID-19 observations to the estimation sample, which also affects the forecast path. We discuss solutions for both estimation and forecasting.

In a VAR, we show that allowing the errors to have a distribution with fatter tails than the Gaussian one, more precisely to follow a Student's t -distribution, tackles the parameter instability problem. Intuitively, the fatter tails allow the abnormal observations to be soaked up by residuals rather than distorting coefficients. This solution is simple and appears to characterize the data even before COVID-19. This is a tractable and simple alternative to [Lenza and Primiceri \(2020\)](#) or [Carriero et al. \(2021\)](#) and relates to [Bańbura et al. \(2020a\)](#). Alternatively, we found that a large Gaussian VAR with a higher degree of prior shrinkage also mitigates the problem of changing parameters after adding the COVID-19 observations, but the model with fat-tailed errors outperforms.

These two options also yield unconditional forecasts which are not adversely affected by the change in coefficients induced by the COVID-19 observations. Yet, in abnormal times, unrecorded in the available estimation sample, it is hard to trust forecasts based solely on exploring historical regularities. We explore two ex-ante reliable information sources about the future: official projections and surveys of professional forecasters. We find that within a small-scale standard Gaussian VAR, the conditional forecast paths can be well-behaved (namely insensitive to the change in parameters when adding the COVID-19 observations) when relevant off-model information is employed. Both hard conditioning (using official projections for the explanatory variables) and soft conditioning (using information from the Survey of Professional Forecasters) help. In particular, the drawback of the former approach is that the number of the conditioning variables matters, namely the larger, the more stable the inflation forecasts. Instead, when using survey information, it is enough to tilt the forecast distribution to the SPF 5 years ahead to ensure that inflation forecasts are not affected by the parameter instability.

Single equation Phillips curve models exhibit a set of peculiarities compared to VARs. Also in this case, estimated parameters change notably when adding the COVID-19 observations to the sample. Nevertheless, allowing the errors to have fatter tails than normally distributed ones does not appear to be a panacea this time around. The parameter instability is also transmitted to conditional forecasts. As a possible solution, our findings suggest that in this framework averaging across many models within a 'thick modelling' of the Phillips curve offers some hedge against the instability coming from the change in parameter estimates due to the COVID-19 observations.

1 Introduction

‘In reality, however, the distribution of shocks hitting the economy is more complex ... might exhibit excess kurtosis, commonly referred to as “tail risk” in which the probability of relatively large disturbances is higher than would be implied by a Gaussian distribution.’ (Mishkin, 2011)

As the pandemic unfolded worldwide, macroeconomists have struggled to make sense of their models after adding the new COVID-19 observations. The variation exhibited by some macroeconomic series has been so large (e.g. real activity and labor market indicators), that this sufficed to distort the estimated coefficients since March 2020. This was shown for the U.S. (Schorfheide and Song (2020), Lenza and Primiceri (2020), Carriero et al. (2021)) and we show that this is also true for the euro area.

In this paper we consider two popular time series models and discuss the results by focusing on euro area inflation. We look at single equation models, more specifically at the Phillips curve within a so-called ‘thick modelling’¹ framework (regularly employed in several central banks and in the European Central Bank in particular)² and Vector Autoregressions (VARs), also frequently employed to understand the drivers of inflation. We distinguish between these two types of models because the solutions that work for VARs seem to be less effective for Phillips curves.

We show that the incoming data after March 2020 heavily impacts the parameters of conventionally estimated models and this is something to be aware of when forecasting or conducting empirical analyses covering this period, irrespective of the scope. Here we focus on the analysis of inflation, but the critique is generally valid.

We propose a tractable and elegant solution to the problem of parameters of a VAR model changing when adding the COVID-19 observations, namely to relax the assumption that the errors are normally distributed and assume instead that they follow a Student’s t distribution. Intuitively, large shocks are more likely to be accommodated under t -distributed errors as compared to Gaussian errors, as the former have fatter tails. This allows for abnormal observations to be soaked up by residuals rather than distorting the parameter estimates.

Our paper is related to the strand of literature emerging after the Great Recession trying to accommodate tail events in macro models. For DSGEs, Cúrdia et al. (2014), Chib and Ramamurthy (2014), Ascari et al. (2015) argued that models with a t -distributed shock structure are strongly favoured by the data over standard Gaussian ones. For VARs, Chiu et al. (2017) and Chan (2020) show that allowing for fat-tailed errors improves in-sample fit and forecasting properties.³ We show that also for the euro area fat tails appear to be a better characterisation of the macroeconomic data than Gaussian ones and moreover, this extension suffices to deal with the impact of the pandemic (which is a more rare and extreme event than recessions).

We also contribute to the rising literature on how standard models can be adapted to withstand the impact of the COVID-19 observations. Lenza and Primiceri (2020) and Carriero et al. (2021) propose to downweigh the impact of abnormal observations through assumptions on the associated variance of the residuals. Our solution is similar in spirit. We believe it is still too soon to ascertain a break in the

¹See Granger and Jeon (2004), who advocate for the use of a battery of specifications instead of one single equation.

²As detailed in Ciccarelli and Osbat (2017), ECB (2017), Bobeica and Sokol (2019), Eser et al. (2020).

³Even before the Great Recession some economists warned that the unconditional distribution of macro variables is not Gaussian (see Christiano (2007) or Fagiolo et al. (2008)).

economic transmission mechanisms and thus parameters should not display dramatic changes with respect to pre-COVID-19 times, especially if a normalization in economic developments is to be expected.⁴ And indeed, we find strong evidence in favour of fat-tailed errors even before the pandemic. Our approach is more comparable to [Carriero et al. \(2021\)](#). They opt for combining stochastic volatility with temporary volatility outliers in a VAR; their choice is motivated by the method being more geared towards generating sizable outliers, which can be more suitable for certain purposes, such as mitigating increases in measures of macroeconomic and financial uncertainty, as in [Carriero et al. \(2020\)](#).

As an alternative to fat-tailed errors, we found that adding more variables to the Gaussian BVAR together with prior shrinkage results in more stable results than in a smaller-scale BVAR. This suggests that in large dimensional models the impact of variables with abnormal dynamics is mitigated; also, an appropriately informative prior helps in disciplining the results.⁵ Still, the model with t -distributed errors outperforms in delivering more stable parameters.

We also investigate the impact of the COVID-19 observations on forecasting. Stabilizing the parameters after adding the COVID-19 observations also stabilizes the unconditional forecast paths. Yet, in abnormal times, unrecorded in the available estimation sample, it is hard to trust forecasts based solely on exploring historical regularities. We argue that these times are best suited to add information from outside the model to inform purely model-based results. [Lenza and Primiceri \(2020\)](#), [Primiceri and Tambalotti \(2020\)](#) and [Bańbura et al. \(2020a\)](#) also incorporate some kind of off-model information to produce more plausible forecasts after the pandemic.⁶ Interestingly, [Bańbura et al. \(2020a\)](#) investigate forecasting within a BVAR with t -distributed errors and argue that adding off-model information on projected macroeconomic uncertainty improves the accuracy of density forecasts during the pandemic. We explore two ex-ante reliable information sources about the future: official projections and surveys of professional forecasters. We first analyse ‘hard conditional’ forecasts where we impose the path of a set of explanatory variables as projected within the E(S)CB Macroeconomic Projection Exercise and second, we analyse ‘soft conditional’ forecasts where we tilt unconditional forecasts to expectations from the ECB Survey of Professional Forecasters (SPF). Tilting towards SPF expectations can be a valid way to provide the model some degree of informed judgement, as professional forecasters do not solely rely on models (which might have been affected by COVID-19 observations), but use a substantial degree of judgement when forming their beliefs about the future. Also, there is evidence that such tilting improves inflation forecasting in certain challenging times, such as in the post-Great Recession period, as a way of indirectly accommodating structural changes (see [Tallman and Zaman \(2020\)](#), [Ganics and Odendahl \(2021\)](#)).

We find that within a small-scale standard Gaussian VAR, the conditional forecast paths are well-behaved (namely insensitive to the change in parameters when adding the COVID-19 observations) when

⁴At the same time, we do not exclude the possibility that the COVID-19 observations might affect macroeconomic relationships in the future (as reflected by model parameters).

⁵A Minnesota prior as in [Sims and Zha \(1998\)](#) coupled with dummy observations as in [Bańbura et al. \(2010\)](#) and [Doan et al. \(1983\)](#) seems to be better equipped to yield more robust results in the face of tail events than a weakly informative prior for instance (and undoubtedly a completely flat prior which is equivalent to OLS estimates). The degree of shrinkage matters in a mechanical way, with a larger degree of prior shrinkage yielding more stable parameters. While more shrinkage was found to be suitable for larger VARs, it can be rather restrictive in a smaller model.

⁶[Lenza and Primiceri \(2020\)](#) produce BVAR conditional forecasts imposing the Consensus unemployment projection from the Blue Chip Forecasts. [Primiceri and Tambalotti \(2020\)](#) explain that forecasting with reduced-form time series models during the pandemic requires unusually strong assumptions and illustrate how this might still work by postulating some possible scenarios about the future evolution of the epidemic.

relevant off-model information is used to create the conditional forecasts. Both hard conditioning (using projections for the explanatory variables) and soft conditioning (using information from the Survey of Professional Forecasters) help. In particular, the drawback of the former approach is that the number of the conditioning variables matters, i.e. the larger the number of paths the researcher imposes, the more stable the inflation forecasts. Instead, when using survey information, it is enough to tilt the forecast distribution to the SPF 5 years ahead to ensure that inflation forecasts are not affected by the COVID-19 observations.

Single equation Phillips curve models exhibit a set of peculiarities compared to VARs. Also in this case, estimated parameters change notably when adding the COVID-19 observations to the sample. Allowing for the errors to have fatter tails and follow a Student's t -distribution instead of a normal one does not appear to be a panacea this time around. This relates to the fact that in the multiple equation system of a VAR the fatness of tails pertaining to the error distribution (and hence the extent to which the errors can soak up abnormal observations in order not to distort coefficients) is informed by the simultaneous developments in all variables. In single equation models it is the developments in the dependent variable that matters more in informing the properties of the residuals; inflation was relatively well behaved during the pandemic, so the residuals of a Phillips curve do not have such fat tails as the residuals of the inflation equation in a VAR model including also GDP for instance.

Phillips curve conditional forecast paths are affected by the change in parameters due to COVID-19. As a possible solution, our findings suggest that in this framework averaging across many models within a 'thick modelling' of the Phillips curve offers some hedge against the instability coming from the change in parameter estimates due to the COVID-19 observations. This stability brought about by aggregating information from many models has also been previously emphasized for instance by [Koop and Korobilis \(2012\)](#), [Rossi \(2020\)](#), [Bańbura and Bobeica \(2020\)](#).

In the rest of this paper, Section 2 investigates the impact of COVID-19 on Phillips curve models, Section 3 turns to VAR models, Section 4 explains why t errors are more effective in a multivariate setup than in the single equation setup, and Section 5 concludes.

2 The impact of COVID-19 on Phillips curve models

The Phillips curve is a key macroeconomic relationship used to gauge past and future inflation dynamics. In central banking, it is one of the pillar frameworks in thinking about inflation, with reduced-form models being often employed to grasp the link between inflation and real activity (see [Yellen \(2015\)](#), [Constancio \(2015\)](#)). Rather than solely relying on a single model, it has become standard in policy circles to adopt a 'thick modelling' approach, whereby several proxies for economic slack or inflation expectations are employed in order to deal with uncertainty surrounding the 'true' inflation drivers (see [Bundesbank \(2016\)](#), [ECB \(2017\)](#), [Bobeica and Sokol \(2019\)](#), [Eser et al. \(2020\)](#)).

We employ a reduced-form New Keynesian Phillips curve model within the 'thick modelling' approach as in [Ciccarelli and Oshat \(2017\)](#). More precisely, we estimate the following specification:

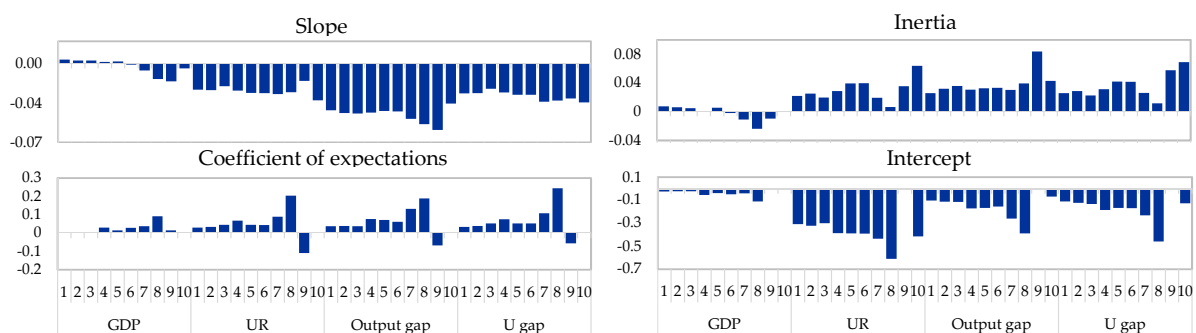
$$\pi_t = \alpha + \rho\pi_{t-1} + \beta[y_{t-1}] + \gamma[exp_t] + \mu imp_{t-2} + \epsilon_t \quad (1)$$

where π_t is the annualized quarterly growth rate of HICP excluding energy and food (in seasonally adjusted terms), y_t is an indicator of the cyclical developments in the economy (henceforth slack measure),

exp_t is a measure of expectations and imp_t are import prices from outside the euro area. $[\cdot]$ indicates that several proxies are being employed, namely 4 measures of slack (real GDP, output gap, unemployment gap and unemployment rate) and 9 measures of survey inflation expectations (Consensus 1,2,...,6 quarters ahead and SPF 1,2,5 years ahead), as well as no expectation term in purely ‘backward’ looking specifications. The estimation sample starts in 1995 for the specifications including Consensus expectations and in 1999 for those including SPF measures

Figure 1 shows the impact on the Phillips curve coefficients when adding three COVID-19 observations, namely 2020Q1, Q2 and Q3, and estimating (1) using an ordinary least squares (OLS) regression. The revision in the parameters brought about by these three observations is quite remarkable and varies with the slack measure. The inertia and coefficient of inflation expectations tend to go up, while the intercept goes down. The change in parameters is not statistically significant, but this does have a bearing on the inflation forecast path. The statistical significance of the inflation drivers is negatively affected, with most slack and expectation measures becoming insignificant, as the error variance is distorted.

Figure 1: Change in Phillips curve parameters. Sample ending in 2020Q3 vs. in 2019Q4



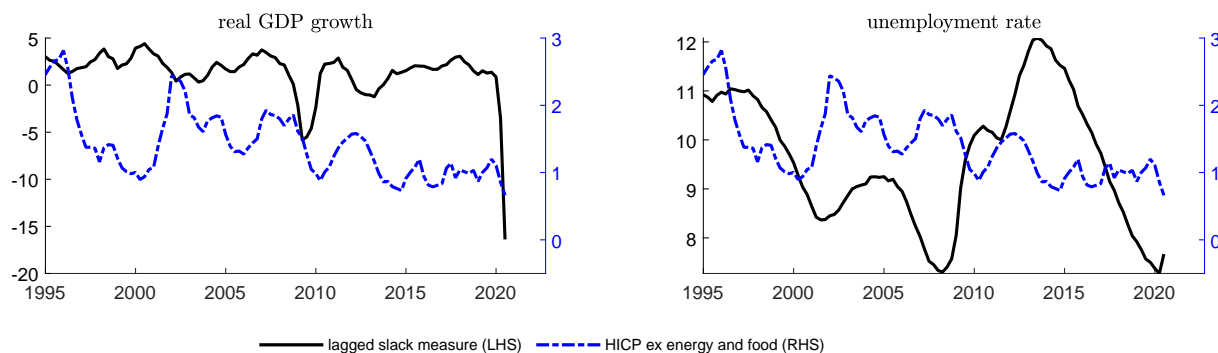
Note: The bars show the change in coefficients for each specification per slack measure. 1,2,...,10 indicate specifications with various measures of expectations, in the following order: 1 - 6 are Consensus 1,2,...,6 quarters ahead, 7 - 9 are SPF 1,2,5 years ahead, 10 - no expectation term. No intercept was included in the specification with SPF 5 years ahead. For the unemployment rate and gap the change in slope refers to coefficients with reversed sign.

The Phillips curve flattens across the considered slack measures; the absolute change in coefficients appears to be small, but given that the Phillips curve was rather flat even before COVID-19, the slope diminishes by around 40 to 50 percent. This change in slope can be gauged also by looking at Figure 2. Compared to the developments in slack measures (particularly on the product market), the reaction of the inflation rate has been rather muted. One argument calling for caution in taking for granted the change in coefficients brought about by the new COVID-19 observations is that they can also reflect real-time mismeasurement of slack. As new data becomes available and the view regarding slack during the pandemic crystallizes, the Phillips curve parameters might again be different.

The change in parameters generated by the COVID-19 observations has an impact on the inflation forecast path. Figure 3 shows the inflation forecast conditioning on the future path of the explanatory variables. The imposed future paths for slack measures and imported inflation are those projected within the European (System of) Central Banks (E(S)CB) Macroeconomic Projection Exercise of September 2020.⁷

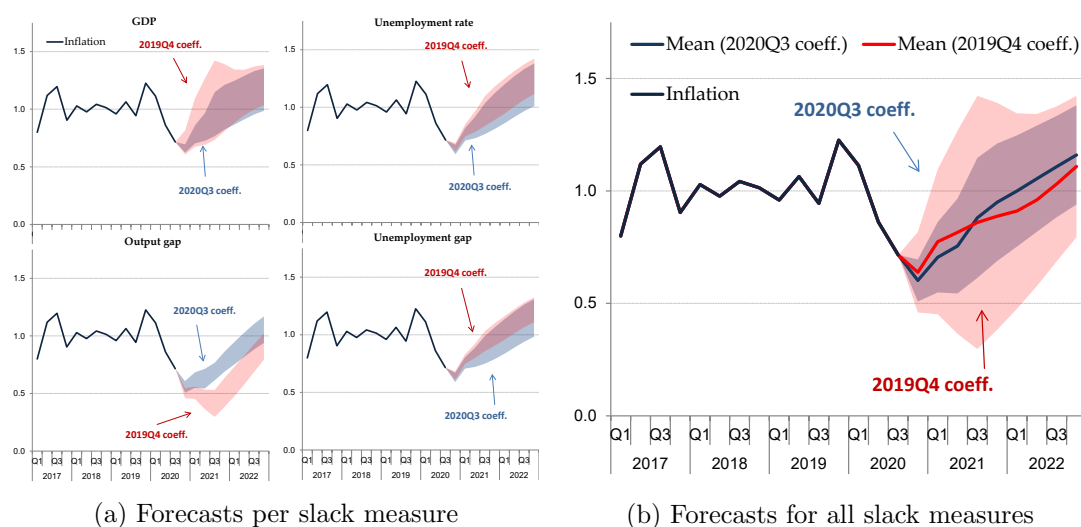
⁷See [ECB \(2016\)](#) for more details on the projection process within the European System of Central Banks.

Figure 2: Indicators of slack and euro area HICP inflation excluding energy and food



Note: The indicators of slack are lagged by one period; real GDP growth and HICP inflation excluding energy and food are expressed as annualized quarterly growth rate of the seasonally adjusted index.

Figure 3: Phillips curve conditional forecasts



Note: The ranges are based on parameter estimates on the sample up to 2019Q4 (red range) and up to 2020Q3 (blue range). Forecast starts in 2020Q4. (a) The ranges show point forecasts covering specifications with various inflation expectation measures. (b) The ranges cover point forecasts from all Phillips curve specifications with various slack and inflation expectation measures.

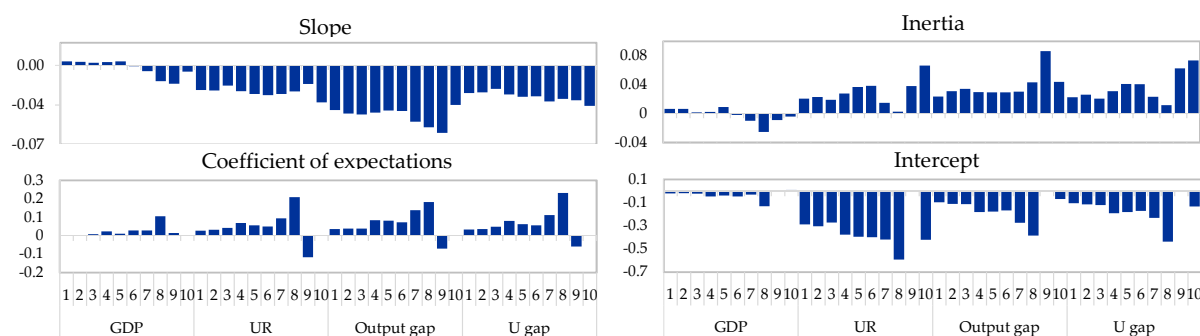
Expectations have been prolonged based on an AR(1) process. In this Figure the path of the explanatory variables, the starting date of the forecast and the forecast horizon stay the same, the only thing that changes is the sample for which the parameter estimates are obtained. The forecast range covers point forecasts across various specifications and it is affected in different ways depending on the slack measure (panel (a)); it is dragged down by the COVID-19 observations in the case of the real GDP and labour market slack measures and up in the case of the output gap.

The ‘thick modelling’ framework offers some hedge with respect to the instability in the forecast path brought about by the change in parameters when estimating over the COVID-19 period. When averaging across conditional forecasts from many models, the mean forecast is almost identical, see

Figure 3, panel (b). This stability brought about by aggregating information from many models has also been previously highlighted for instance by [Koop and Korobilis \(2012\)](#) and [Bańbura and Bobeica \(2020\)](#) in the Phillips curve context or by [Rossi \(2020\)](#) more generally.

Alternative estimation techniques and the parameter instability problem. We investigate whether different modelling assumptions bring about more parameter stability in the Phillips curve setup. We first estimate the Phillips curve models in a Bayesian setting, assuming that the distribution of the errors is heavy tailed by allowing for t -distributed errors instead of Gaussian ones. *A-priori*, one would expect that this should dampen the effect of extreme observations on the parameter estimates. Unfortunately this doesn't solve the parameter instability issue in the Phillips curve models (see Figure 4), but it does in a VAR (see the discussion in Section 3.2.1).

Figure 4: Change in Phillips curve parameters when errors follow a Student's t distribution. Sample ending in 2020Q3 vs. in 2019Q4



Note: The bars show the change in coefficients for each specification per slack measure. 1,2,...,10 indicate specifications with various measures of expectations, in the following order: 1 - 6 are Consensus 1,2,...,6 quarters ahead, 7 - 9 are SPF 1,2,5 years ahead, 10 - no expectation term. No intercept was included in the specification with SPF 5 years ahead. For the unemployment rate and gap the change in slope refers to coefficients with reversed sign.

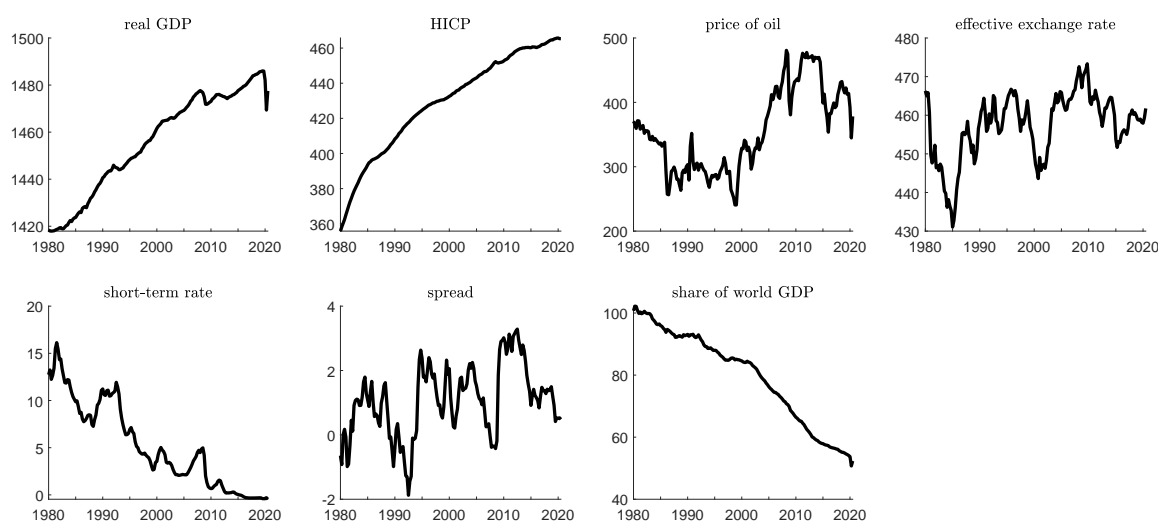
We also investigate whether allowing for smooth time variation in the parameters plus stochastic volatility mitigates the parameter instability problem. We apply 'conventional' time-varying parameter Bayesian models in the spirit of [Primiceri \(2005\)](#) assuming that the Phillips curve coefficients and log-variance of errors follow a random walk. Figure 15 in Appendix C.1 shows the change in the endpoint estimates of the Phillips curve parameters when the models are estimated up to 2020Q3 compared to 2019Q4. The COVID-19 observations shift the entire path of the estimated coefficients altogether and the changes in parameters are even more pronounced than in the fixed parameter case.⁸ Also, when it comes to the impact on the inflation forecast, the shifts are more pronounced than in the case of the fixed coefficients (see Figure 16 in Appendix C.1).

⁸As an alternative modelling technique, we have also checked whether a median regression of the Phillips curve model produces robust parameter estimates. Unfortunately, the estimated parameters exhibit similar sensitivity as in the fat-tailed error model.

3 The impact of COVID-19 on VAR models

In this section we consider the most popular model employed in macroeconomics, namely Sims' Vector Autoregression (VAR). For our analysis, we use the small scale VAR that was employed to understand the drivers of euro area inflation by [Bobeica and Jarociński \(2019\)](#). Figure 5 shows the data over the sample 1980Q1-2020Q3; there is some variation during 2020 in the considered variables, but nothing unprecedented with the exception of the real GDP and the euro area share in world GDP. In contrast, for the U.S. some economic variables, in particular the unemployment rate, have recorded formidable spikes, whereas in the euro area this particular indicator has been more stable. Yet, we show that the variation recorded by the euro area variables suffices to distort the parameter estimates in a standard VAR.

Figure 5: Data set for the small euro area BVAR



Note: Real GDP, HICP, price of oil and nominal effective exchange rate are in log-levels times 100; short-term rate, 10-year government bond spread and share of world GDP are in percentage points.

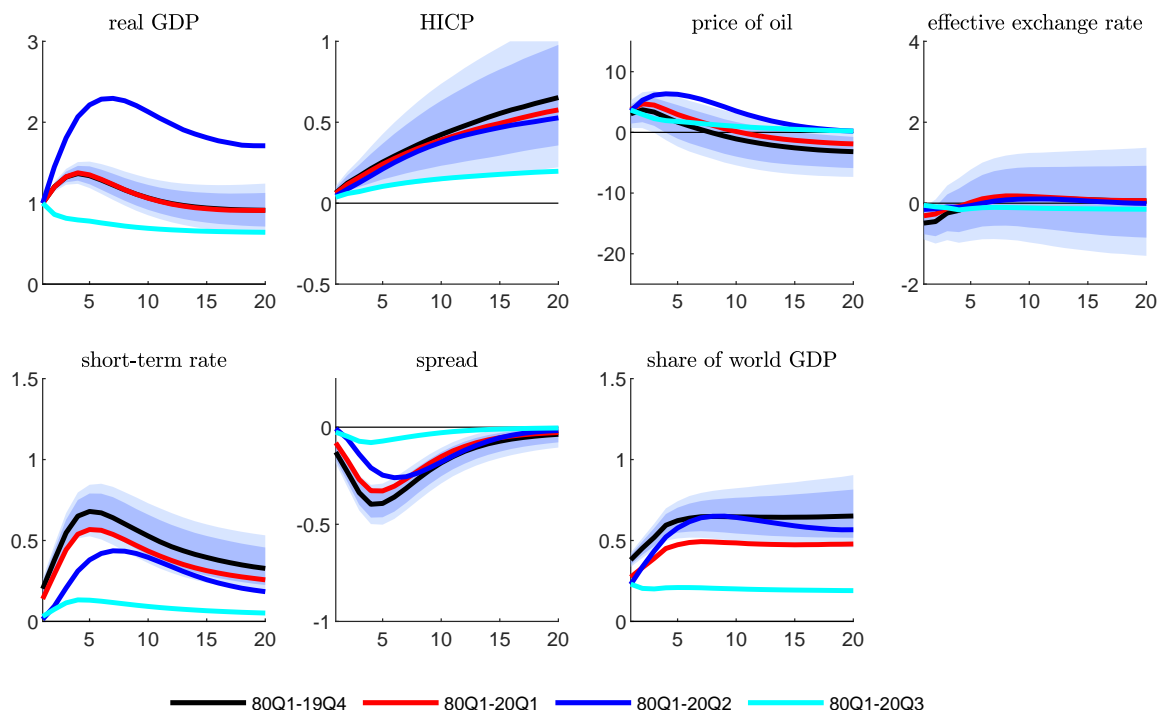
3.1 Impact of COVID-19 observations on estimated parameters

Figure 6 shows estimates of impulse response functions (IRFs) to a one standard deviation shock in real GDP for an expanding estimation sample, which starts in 1980 and ends in 2019Q4, 2020Q1, 2020Q2 and 2020Q3 respectively.⁹ Estimates are obtained using a [Sims and Zha \(1998\)](#) prior.¹⁰ The inclusion of the COVID-19 observation substantially affects the IRF estimates as compared to the pre-COVID-19 times. Specifically, with the inclusion of the 2020Q2 observation, the IRF of real GDP becomes more persistent, while the transmission of the shock to inflation and financial variables is dampened. Estimates become even more different when the sample ends in 2020Q3. The reaction of inflation to the generic GDP shock is flattened. This illustrates that parameter estimates are very sensitive to unusually large innovations in a standard Gaussian VAR, which is a major shortcoming for analysis in times of COVID-19.

⁹We do not attach any structural interpretation to this shock, but merely use it as means to study the dynamic properties of the VAR model.

¹⁰We construct dummy observations along the lines of [Bańbura et al. \(2010\)](#) and [Doan et al. \(1983\)](#), which reflect Minnesota type shrinkage on VAR coefficients coupled with two sum-of-coefficients restrictions. For estimation, we use dummy observations to construct the moments of the implied normal-inverse Wishart prior.

Figure 6: Impulse response functions in a small BVAR with Gaussian errors



Note: Impulse response functions to a one standard deviation shock in real GDP from a BVAR with Gaussian errors using a Sims and Zha prior. Thick lines are median estimates and the dark (light) blue area is the 68% (90%) credible interval for the estimation window until 19Q4.

The prior choice makes a difference, with more informative priors having an upper hand in delivering more stable coefficients once the COVID-19 observations are added. [Lenza and Primiceri \(2020\)](#), [Schorfheide and Song \(2020\)](#) and [Carriero et al. \(2021\)](#) document using monthly U.S. data that estimated VAR models are explosive when the recent COVID-19 observations are included. While the estimated IRFs in Figure 6 are sensitive to the inclusion of the COVID-19 observations, they are not explosive for our chosen prior specification. Figure 17, panel (a), in Appendix C.2 shows that this explosive behaviour is present in the euro area case as well if we adopt a weakly informative prior for the VAR coefficients.¹¹ This suggests that the design of the prior, the chosen amount of shrinkage, as well as the data set at hand dictate whether a Bayesian VAR model becomes explosive once COVID-19 observations are included in the sample. Moreover, we also explore whether a higher degree of shrinkage applied to our baseline [Sims and Zha \(1998\)](#) prior dampens the impact of COVID-19 on the change in the VAR coefficient estimates; we label this prior ‘strong Sims and Zha’.¹² Figure 17, panel (b) shows that the higher degree of shrinkage mechanically disciplines the coefficients in the VAR to some extent, but there is still notable instability in the IRFs once the COVID-19 period is included in the sample, which stems from a change in residual

¹¹The weakly informative prior is calibrated by centering the VAR coefficients around zero and setting the variance of all coefficients to one except for the constant, which we set to 10. The scale of the inverse Wishart distribution is set to the identity matrix.

¹²Specifically, we divide the prior variances for the coefficients by a factor of four. A higher level of shrinkage is typically recommended for very large VAR with over 100 variables, see [Bańbura et al. \(2010\)](#).

correlations.¹³ Yet, imposing a very dogmatic prior comes with the drawback of silencing the information coming from the data.

3.2 Solutions for estimation within a VAR

3.2.1 Accommodating COVID-19 observations via alternative error structures

Several proposals on how to tackle the problem of changing parameters due to the COVID-19 observations have been put forward in the literature, see [Lenza and Primiceri \(2020\)](#) and [Carriero et al. \(2021\)](#). What they have in common is that macroeconomic relationships which held in the past are assumed to be still valid. That is, they propose to downweigh the abnormal observations by allowing the variance of the residuals to react quickly to the COVID-19 observations.

[Lenza and Primiceri \(2020\)](#) propose an explicit volatility model for the residuals corresponding to the COVID-19 observations. Their solution is tractable and easy to understand, but requires that one specifies exactly when the abnormal observations start and estimate the decay of volatility based on few observations (especially when working with quarterly data).

[Carriero et al. \(2021\)](#) propose a more flexible solution that allows for variable-specific outliers in the volatility process, leaving fewer choices to the modeller. Here volatility is a function of its past values, but it is also allowed to jump with abnormal observations with a certain probability.¹⁴ One possible limitation of this model class is related to the sensitivity to the ordering of variables (see [Primiceri \(2005\)](#), [Cogley and Sargent \(2005\)](#)). Specifically, when volatility changes idiosyncratically, estimates of the time-varying covariance matrix are extremely sensitive to the ordering of variables due to the structure imposed by the triangular factorization of the time-varying covariance matrix in their model (see [Hartwig \(2020\)](#)).

Our proposal does not involve making ad-hoc choices and is thus more in line with the approach of [Carriero et al. \(2021\)](#). We propose to abandon the Gaussian error structure and assume multivariate t -distributed errors instead, see [Ni and Sun \(2005\)](#) and [Chan \(2020\)](#). Intuitively, large shocks are more likely to occur under errors following a Student's t distribution, which have fatter tails than the Gaussian one. This means that the COVID-19 observations do not heavily influence the relationships between variables, as reflected in the parameter estimates (but they are also not prevented from doing that as is the case when including dummies for this period). The idea that the data might favour models with errors with fatter tails is inspired by the literature emerging after the Great Recession (see [Cúrdia et al. \(2014\)](#); [Chiu et al. \(2017\)](#); [Chan \(2020\)](#)), when Gaussian models failed to properly account for the strong variations in some variables; we show that the t -distributed errors work also during the pandemic. This solution shares the spirit of the aforementioned ones in the sense that the COVID-19 observations are downweighed by some form of conditional variance within the estimation framework. Similar to [Lenza and Primiceri \(2020\)](#), we treat the change in shock volatility as being common across all variables instead of considering variable-specific volatility as in [Carriero et al. \(2021\)](#). This is a simple and tractable assumption. As opposed to the latter paper, in our VAR, residuals have constant second-order moments (see Appendix A for more technical details on the VAR with t errors). Student's t -distributed errors – or other fat-tailed error distributions – are better suited during the pandemic as the COVID-19 observations

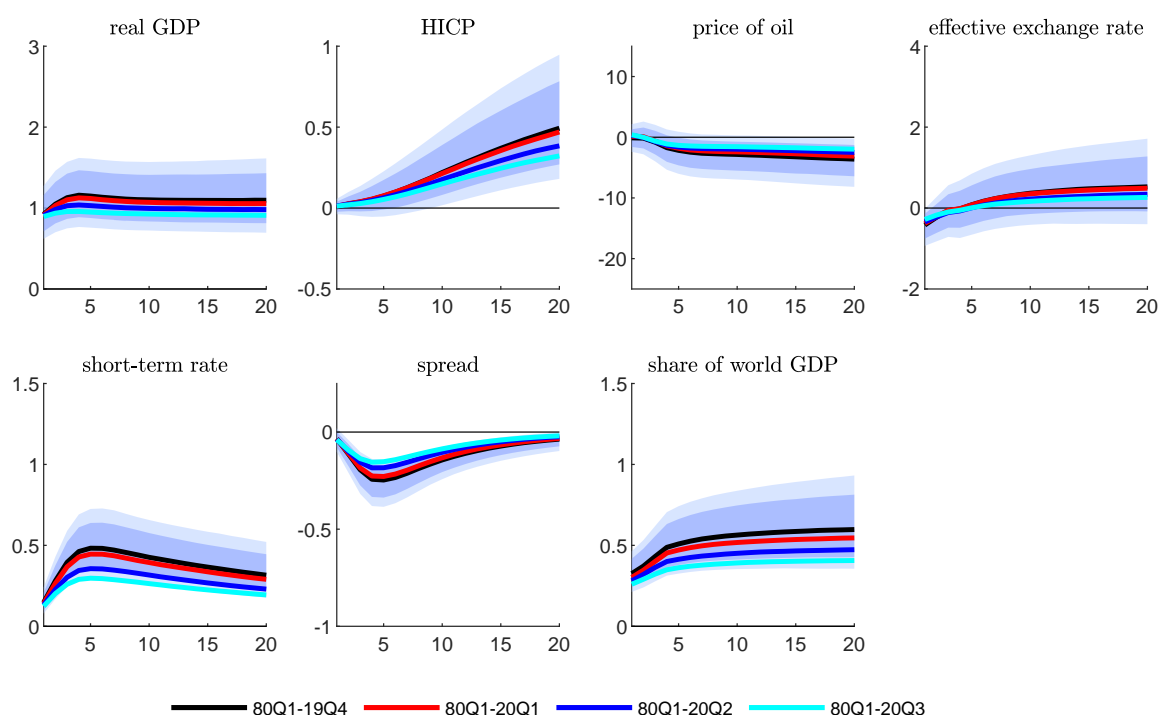
¹³For brevity, we do not report changes in residual correlation for the considered VAR models.

¹⁴The authors also explore a dummy VAR using separate dummies in every equation for each month since March 2020. Abstracting from the impact of the COVID-19 observations on the parameters, they show that this approach does not perform reasonably when it comes to forecasting, yielding implausibly narrow forecast bands.

may inflate the volatility process for an extended period of time in widely employed stochastic volatility models (without outlier correction). This may lead to rather imprecise density forecasts, see [Carriero et al. \(2021\)](#) and [Hartwig \(2021\)](#). In addition, [Carriero et al. \(2020\)](#) and [Hartwig \(2021\)](#) discuss that a simple stochastic volatility process inflates estimated error volatility also before COVID-19 occurred due to the ex-post smoothing of the spike during the pandemic.

Figure 7 shows the counterpart to Figure 6 under the assumption that the errors in the Bayesian VAR are multivariate t -distributed. For the pre-COVID-19 sample, the reaction of inflation to real GDP is somewhat weaker in the VAR with fat-tailed errors compared to the one with Gaussian residuals. When adding the COVID-19 observations the response of inflation to a shock in real GDP is only marginally diminished and in general the IRFs are much more stable across different estimation samples compared to their Gaussian version.¹⁵

Figure 7: Impulse response functions in a small BVAR with fat-tailed errors



Note: Impulse response functions to a one standard deviation shock in real GDP from a BVAR with fat-tailed errors using a Sims and Zha prior. Thick lines are median estimates and the dark (light) blue area is the 68% (90%) credible interval for the estimation window until 19Q4.

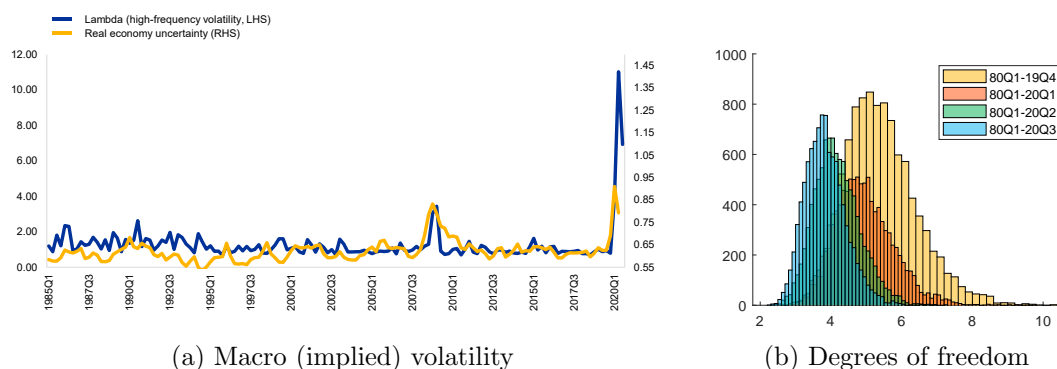
This solution allows for the impact of COVID-19 observations to be soaked up by the heavy tails of the VAR residuals, leaving parameter estimates largely unaffected. One interesting question is why allowing for a fat-tailed error distribution seems to be a solution in a multiple equation model such as the VAR but not in the single equation framework, such as the Phillips curve (see Figure 1). We discuss this issue in more detail in Section 4.

This estimation technique yields, as a by-product, an interesting latent variable, which we label

¹⁵The IRFs of the fat-tailed VAR are stable also under the weakly informative prior and under the strong Sims and Zha prior, see panels (a) and (b) in Figure 18 in Appendix C.3.

λ ,¹⁶ that can be interpreted as the high-frequency volatility of the data set (Chiu et al., 2017)¹⁷ or can be used as an outlier detection tool (Jacquier et al., 2004). Technically, this variable mixes the multivariate normal distribution such that it mimics the fatter tails of the multivariate t -distribution with a particular degree of freedom, see Geweke (1993) and Ni and Sun (2005). Panel (a) of Figure 8 depicts the estimated high-frequency volatility in our data set alongside a measure of macroeconomic uncertainty.¹⁸ λ spiked during the Great Recession and sky-rocketed during the COVID-19 crisis, capturing the fact that there was abnormal variation in the data. The timing of these spikes coincides with the peaks in measured macroeconomic uncertainty.¹⁹ λ capturing the spikes in macroeconomic uncertainty is another argument in favour of this estimation technique: the model with t -distributed errors captures some salient macroeconomic features in the shape of volatility/uncertainty, which may not be accounted for by standard Gaussian models. Therefore, standard Gaussian models may yield distorted estimates of macroeconomic relationships especially after COVID-19.

Figure 8: Macroeconomic uncertainty and tail risk



Note: (a) Posterior mean of the high-frequency volatility λ in the BVAR with t -distributed errors and Sims and Zha prior. (b) Posterior distribution of the degrees of freedom parameter estimated on an expanding sample.

Panel (b) of Figure 8 shows that the posterior distribution of the degree of freedom parameter (ν) becomes sharper and shifts to the left once the COVID-19 observations are included. Specifically, the posterior median of ν is 4.3 in the pre-COVID period and declines to 3.3 in 2020Q3. Therefore, the estimated tail risk in the economy increases somewhat with the COVID-19 observations. Nevertheless, the log probability of large tail events is rather similar for a degree of freedom parameter between 3

¹⁶In Appendix A, λ_t is defined in the variance scale. In the text, we use $\lambda = \sqrt{\lambda_t}$.

¹⁷The term high-frequency volatility derives from the fact that under the t -distribution, these unobserved volatility changes are independently distributed. In contrast, in stochastic volatility models, volatility is not independently distributed, but follows a persistent process and can be regarded as a lower-frequency volatility.

¹⁸The measure of macroeconomic uncertainty is the one of Jurado et al. (2015) for the U.S. and we used it as an imperfect proxy for the euro area as well. The available Economic Policy Uncertainty index for the euro area derived by Baker et al. (2016), focuses too much on policy and does not reflect the macroeconomic uncertainty embedded in our variables.

¹⁹By construction, our measure of high-frequency volatility (λ) is rather short-lived, as under t errors large shocks are modelled as independent events and are not allowed to be serially correlated. This can explain why during the Great Recession the proxy for macroeconomic uncertainty has been elevated for longer than what λ would imply.

and 5, while for larger values of ν , the tails become substantially thinner.²⁰ We interpret these findings as evidence that the data not only exhibits fat tails due to the extreme COVID-19 observations, but exhibited substantial tail risk even before, which speaks in favour of the t -distribution assumption as an appropriate treatment of the data in general.

An alternative way to check if the data asks for fat-tailed errors is to compare the log marginal likelihood of the Bayesian VARs with Gaussian and t -distributed errors, respectively. Table 1 shows these estimates for various prior specifications and on an expanding estimation window.²¹ By comparing the log marginal likelihood figures in panel (a) and (b) we find decisive evidence in favour of fat tails both before and during the COVID-19 pandemic, as well as across all different prior specifications. The optimal degree of parameter shrinkage is not affected by the COVID-19 observations. Specifically, both models favour our standard calibration for the Sims and Zha prior over the weakly informative or the strong Sims and Zha prior with a higher degree of parameter shrinkage.²² This is in contrast to the result reported in [Lenza and Primiceri \(2020\)](#), who document that the Gaussian VAR model likes less shrinkage after COVID-19.

Table 1: Log marginal likelihood for the small scale BVAR model

	Weak	SZ	Strong SZ		Weak	SZ	Strong SZ
19Q4	-1767.66	-1348.28	-1450.39	19Q4	-1700.91	-1267.89	-1387.99
20Q1	-1814.78	-1391.95	-1489.12	20Q1	-1722.95	-1289.83	-1409.72
20Q2	-1907.49	-1499.75	-1599.82	20Q2	-1756.07	-1322.09	-1441.35
20Q3	-1976.63	-1539.87	-1626.97	20Q3	-1784.04	-1350.60	-1467.01
(a) BVAR with Gaussian errors				(b) BVAR with fat-tailed errors			

Note: The bold figure indicates maximum log marginal likelihood for each model in a selected estimation window. Weak stands for a weakly informative prior and SZ stands for Sims and Zha.

3.2.2 Expanding the data set by adding more variables

As an alternative to fat-tailed errors in the VAR we investigate whether employing a larger model ensures more stable parameters during the pandemic. Intuitively, in large dimensional models the impact of variables with abnormal dynamics might be mitigated. A larger set of data might offer a hedge when it comes to estimation and forecasting in the presence of instabilities, see the discussion in [Rossi \(2020\)](#). To explore this possibility, we estimate a BVAR with 23 variables used in [Bobeica and Jarociński \(2019\)](#) to understand inflation dynamics in the euro area.²³

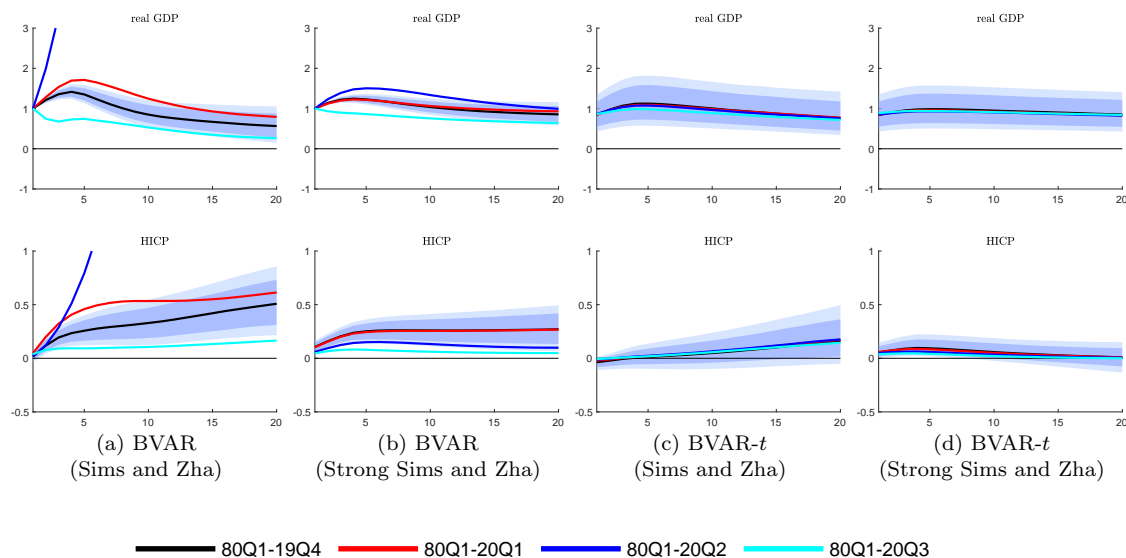
²⁰Compared to the Gaussian case, a ten standard deviation shock receives a log probability of -50 , i.e. the probability of such an event is 7.69×10^{-21} percentage points, which practically means that such an event will never happen. While, under t errors with a degree of freedom parameter between 3 and 5, the probability of such an event is roughly comparable to a four standard deviation shock under normal errors, which is about 0.0134 percentage points.

²¹For the VAR with t errors, we use the algorithm in [Chan \(2020\)](#).

²²We also conducted a more granular grid search on the tightness of the dummy observation prior and found that the t error model generally favors a lower level of tightness as compared to the Gaussian model.

²³The VAR includes HICP, real GDP, real private consumption, real private investment, total employment, unemployment rate, capacity utilisation, consumer confidence, PMI, the rest-of-the-world real GDP, nominal effective exchange rate, EUR/USD exchange rate, the oil price in dollars, non-energy commodity prices, U.S. real GDP, U.S. CPI, U.S. short term interest rate, EONIA, two-year government bond spread, 10-year government bond spread, Mortgage bank lending spread, VSTOXX index of financial volatility, economic policy uncertainty. See [Bobeica and Jarociński \(2019\)](#) for details regarding the data set.

Figure 9: Selected impulse response functions in a large BVAR



Note: Impulse response functions to a one standard deviation shock in real GDP from the large BVAR with Gaussian (BVAR) and fat-tailed errors (BVAR- t) using various priors.

Figure 9 shows the impulse response functions for real GDP and HICP inflation to a one standard deviation shock in real GDP from both the Gaussian and the fat-tailed error VAR with more variables using the two variants of the Sims and Zha prior.²⁴ The IRFs of the large Gaussian VAR with a Sims and Zha prior with the same parameterisation used for the small-scale VAR (panel (a)) exhibit substantial instability as soon as the COVID-19 observations are included and become explosive when estimation runs until 2020Q2. When imposing more shrinkage, as recommended for larger models, the VAR coefficients become better-behaved, see panel (b). The IRFs are more stable than the ones presented in Figure 6, but there is still some change in the IRFs when adding the COVID-19 observations.²⁵ In contrast, estimates from a VAR with fat tails are well-behaved across different prior specifications, see panel (c) and (d). Therefore, even in large VARs, fat-tailed errors are a sufficient extension to make the model robust against the extreme variations triggered by the pandemic.

Since the large Gaussian BVAR produces more stable dynamics only with the stronger Sims and Zha prior, a natural question is whether the marginal likelihood would also favour this model. Table 2 shows the log marginal likelihood for these VARs. The models with fat-tailed errors yield higher marginal likelihoods, suggesting that they are preferred by the data. At the same time, the standard Sims and Zha prior which was used for the small-scale BVAR is preferred over the strong variant for all estimation samples. This criterion would thus guide us to choose an explosive over a stable model, which is a clear caveat for selecting a good model based on the marginal likelihood. Therefore, we argue that in times of COVID-19 one has to consider other aspects as well when selecting the most appropriate model.

²⁴Figures 19, 20, 21 in Appendix C.4 show the IRFs for all variables for all considered prior specifications.

²⁵The differences in the IRFs before and after COVID-19 differ across variables: for some there is hardly any change (e.g. employment, the EUR/USD exchange rate), whereas for others the changes are more pronounced (consumption, foreign real GDP, VSTOXX).

Table 2: Log marginal likelihood for the large BVAR

	Weak	SZ	Strong SZ		Weak	SZ	Strong SZ
19Q4	-6434.98	-3686.72	-3855.57	19Q4	-5871.95	-3404.98	-3620.00
20Q1	-6581.73	-3801.36	-3961.20	20Q1	-5946.10	-3462.14	-3675.51
20Q2	-6725.73	-3907.66	-4084.26	20Q2	-6037.22	-3526.69	-3736.27
20Q3	-6842.95	-4012.50	-4187.80	20Q3	-6095.78	-3629.02	-3823.51
(a) BVAR with Gaussian errors				(b) BVAR with fat tailed errors			

Note: The bold figure indicates maximum log marginal likelihood for each model in a selected estimation window. Weak stands for weakly informative and SZ stands for Sims and Zha.

3.3 Impact of COVID-19 observations on unconditional forecasts

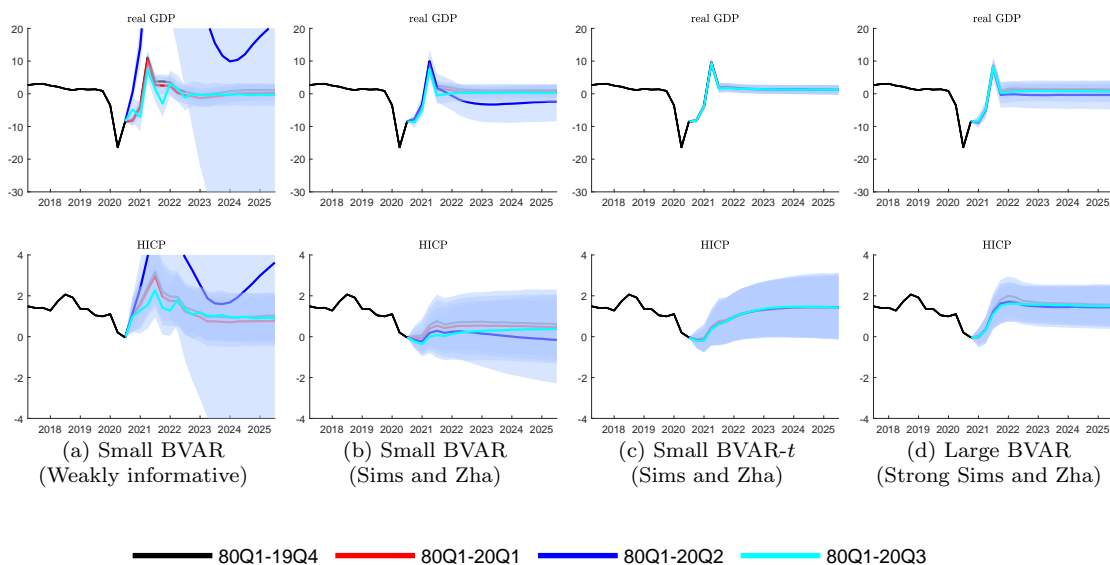
It is difficult to ascertain what works in forecasting after the pandemic as there are not enough realisations to perform a proper forecast evaluation. Instead, we focus on a different criterion to judge the appropriateness of these models, namely the lack of instability in the forecast path due to a change in parameters induced by adding the COVID-19 observations to the estimation sample.

Figure 10 shows unconditional forecasts starting 2020Q4 for the VAR models previously described. In producing these forecasts the data and the starting date of the forecast are the same, the only thing that differs is the sample on which the coefficients have been estimated. In other words, the change in the forecast path solely reflects revisions in VAR parameter estimates. Panels (a) and (b) show that in the case of the small scale BVAR, unconditional forecasts for both a weakly informative and the Sims and Zha prior change visibly once the COVID-19 observations are included in the estimation sample (see blue and cyan line), more so for the weakly informative prior. In contrast, unconditional forecasts are not affected by the change in parameters induced by expanding the sample through 2020 in the BVAR with fat-tailed errors (panel (c)) or in the large scale Gaussian BVAR with a strong Sims and Zha prior (panel (d)).

Interestingly, the projected paths of inflation are very similar for both the small-scale fat-tailed BVAR and the large-scale Gaussian BVAR. They both imply a return to levels below, but close to 2 percent, though the fat-tailed model implies a somewhat slower convergence. The forecasts bands are somewhat wider for the fat-tailed BVAR than for the large scale Gaussian model. Also, the bands under the fat-tailed BVAR become wider (hardly visible in the figures) as soon as the COVID-19 observations are included in the estimation sample. This is due to the fact that the estimated degrees of freedom parameter declines, which increases the unconditional variance of the forecast errors.

Whether such a quick pick-up in inflation is reasonable is subject to discussion. Unconditional forecasts based on models that only look at past historical patterns of the data are hard to be trusted during the pandemic as COVID-19 is an unprecedented macroeconomic shock. It has complex effects on the demand- and supply-side of the economy (see the discussion in [Guerrieri et al. \(2020\)](#)), highly heterogeneous impacts across sectors and agents and unknown behavioural consequences (for a stylised representation of the channels through which the COVID-19 shock affects inflation, see Figure 22 in Appendix C.5). As such, these times are ideally suited to adding information from outside the model to inform purely model-based forecasts, as also argued by [Primiceri and Tambalotti \(2020\)](#).

Figure 10: Unconditional forecasts starting 2020Q4



Note: The forecast starts in 2020Q4, parameters are estimated on an expanding estimation window. Thick lines are median estimates and the (overlapping) light blue areas are the 68% credible interval for each respective estimation window.

3.4 Impact of COVID-19 observations on conditional forecasts

There are two key sources of plausible information about possible future developments of macroeconomic variables: projections published by economic institutions and surveys of professional forecasters. We illustrate ways to incorporate both types of external information and how the conditional forecasts behave in some selected VARs after including the COVID-19 observations.²⁶

3.4.1 ‘Hard conditioning’ a la Waggoner and Zha (1999) using external projections for a set of variables

Figure 11 shows conditional forecasts of inflation and real GDP growth in our VAR when imposing a set of off-model projected paths for a set of variables.²⁷ The external projections are obtained from the E(S)CB Macroeconomic Projection Exercise of September 2020.

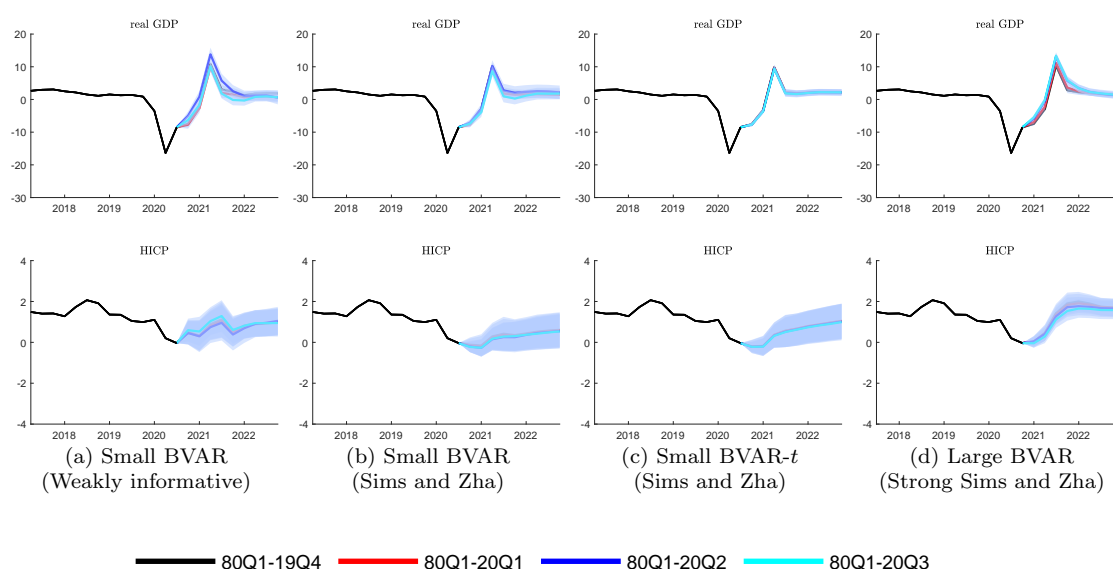
The path of the variables is strictly imposed using the algorithm of Waggoner and Zha (1999) without the updating step. The figure shows that conditional forecasts of the small-scale Gaussian VAR are now robust to the change in parameter estimates over the expanding estimation window for all prior

²⁶ An interesting proposal on how to add useful off-model information when forecasting with a fat-tailed error BVAR after COVID-19 is put forward by Bańbura et al. (2020a). The authors explore the link between the high frequency volatility (λ) and macroeconomic uncertainty and link the future path of λ to externally projected volatility paths in order to increase the accuracy of the density forecasts.

²⁷ For the small VARs, the forecast is conditioned on the price of oil, the effective exchange rate, the short term rate, the spread and the share of euro area GDP in world GDP. For the large VAR, the forecast is conditioned on employment, unemployment rate, rest of the world GDP, nominal effective exchange rate, EUR USD exchange rate, oil price, commodity prices, real GDP foreign, consumer prices foreign, domestic and foreign short-term interest rate and the spread between ten-year government bond and short-term interest rate.

specifications, even for the weakly informative one (panels (a) and (b)). Thus, the problem of unstable forecasts can be solved by conditioning on a sufficient set of variables that exhibit well-behaved future paths and are strongly interlinked with the target variable.²⁸ The conditional forecast of the VAR with fat-tailed errors are stable as well (panel (c)). This model yields a slightly less pessimistic outlook for inflation than the Gaussian one under this scenario and compared to the unconditional version, inflation is now expected to return to somewhat lower levels. The fact that the derived future inflation path is now somewhat lower than indicated by purely unconditional forecasts highlights the effect of adding off-model information also incorporating assessments on the pandemic impact on several economic variables. Conditional forecasts are stable also for the large VAR (panel (d)) and this result is robust when conditioning on various blocks of variables, with the mention that the larger the set of conditioning variables, the more stable the conditional forecasts. Overall, it seems that conditioning on a relevant information set alleviates the problem of in-sample parameter instability.

Figure 11: Conditional forecast based on macroeconomic projections starting 2020Q4



Note: Conditional forecasts imposing the paths for a set of variables obtained from the E(S)CB Macroeconomic Projection Exercise of September 2020; parameters are estimated on an expanding estimation window. Thick lines are median estimates and the (overlapping) light blue areas are the 68% credible interval for each respective estimation window.

3.4.2 ‘Soft conditioning’ a la [Robertson et al. \(2005\)](#) using inflation expectations from the Survey of Professional Forecasters

Using the previously described methodology of producing conditional forecasts, the imposed future paths of the conditioning variables always hold exactly. However, in particularly uncertain times, one may want to consider off-model information with some degree of uncertainty around it. This is known as ‘soft con-

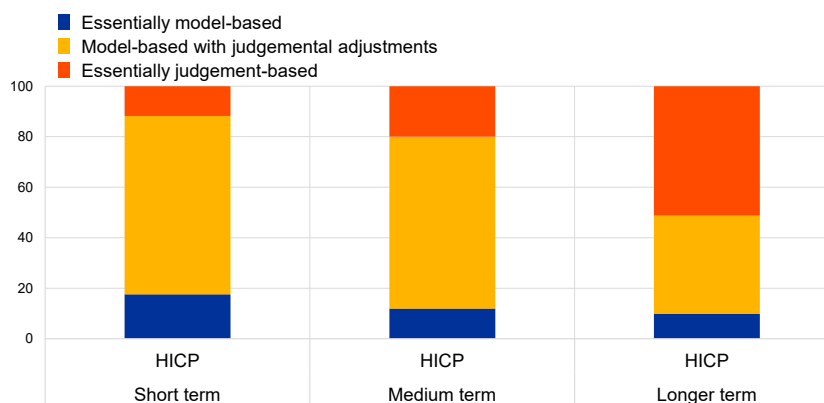
²⁸When we condition for instance only on the oil price or short-term interest rate, the conditional forecast of the Gaussian VAR becomes explosive with the weakly informative prior and the Sims and Zha prior. This is because the correlation of these variables with inflation is too weak to discipline the forecast.

ditioning’, see the relative entropy method by [Robertson et al. \(2005\)](#). The main idea of this approach is to derive a new predictive distribution ‘tilted’ to the mean or median of the distribution of the off-model information (and possibly also to some quantile values) that is as close as possible to the initial unconditional forecast distribution (which might not be the case in the [Waggoner and Zha \(1999\)](#) approach).

In principle, one can apply this procedure to produce forecasts for any variable included in the VAR by ‘tilting’ its unconditional forecast to some off-model information regarding its future path. Here we illustrate this approach by considering a valuable information set, namely the expectations of the Survey of Professional Forecasters (SPF) for the euro area ([Grothe and Meyler \(2018\)](#) show that these expectations have a non-negligible predictive power for euro area inflation developments, as compared to statistical benchmark models). This alternative approach alleviates the issue of having to come up with a set of strong assumptions as in [Primiceri and Tambalotti \(2020\)](#) or imposing the future path of many variables included in the model. It can also constitute a valid cross-check for results obtained when using external projections for some variables, such as those from the E(S)CB Macroeconomic Projection Exercise.

Professional forecasters do not solely rely on models, which have likely been adversely affected by COVID-19 observations, but use judgement when submitting their forecasts (see the discussion in [de Vincent-Humphreys et al. \(2019\)](#)). In a special ECB SPF questionnaire (2018), respondents indicated that a sizeable share of their forecast can be attributed to their judgement, especially when it comes to long-term forecasts, see Figure 12.

Figure 12: Information used by professional forecasters of inflation at various forecast horizons

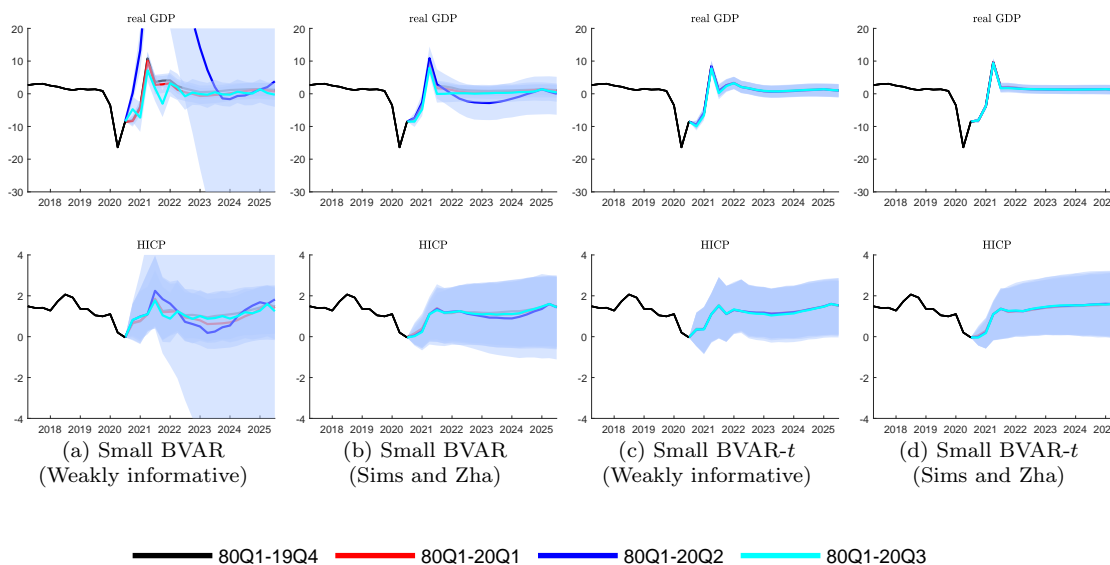


Note: Professional forecasters’ response to question: To what extent are your point forecasts model or judgement based? (percent of responses)

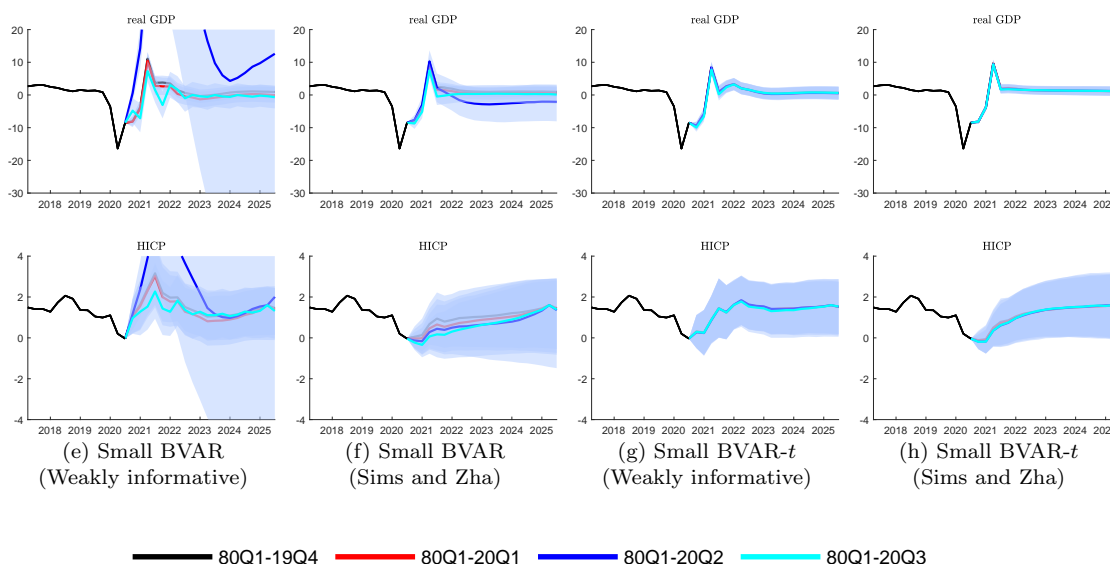
Source: ECB SPF special questionnaire 2018.

Tilting model-based inflation forecasts to SPF expectations has been recently investigated, but for somewhat different reasons. [Ganics and Odendahl \(2021\)](#) focus on the euro area and show that survey forecasts can help mitigate the effects of structural breaks on the forecasting performance in a VAR. They also find that professional forecasters are better at forecasting than a standard Bayesian VAR model around the two recent euro area recessions, as well as the slow recovery thereafter. This is a particularly challenging period for forecasting euro area inflation as reduced-form time series models have a hard time in providing accurate forecasts (as also discussed in [Bańbura and Bobeica \(2020\)](#)). Similarly, [Tallman and Zaman \(2020\)](#) show that such tilting produces superior inflation forecasts after the Great Recession also in the U.S., acting as a way of indirectly accommodating structural changes (and mitigating mis-

Figure 13: Conditional forecasts based on tilting to SPF expectations



(I) Forecast tilted to SPF expectations for inflation (short- and long-run) and GDP growth (long-run)



(II) Forecast tilted to long-run SPF inflation expectations

Note: Forecast tilted to the SPF forecast for inflation and real GDP growth. The forecast starts in 2020Q4 and is using parameters estimated on an expanding estimation window. Thick lines are median estimates and the (overlapping) light blue areas are the 68% credible interval for each respective estimation window.

specification issues in a VAR). Bańbura et al. (2020b) document forecast gains for euro area inflation and GDP when an optimal forecast combination from many models is tilted to the more subjective SPF mean; it is worth noting that when tilting to both first and second moments of SPF there is a general

worsening of the forecasting performance, so the authors recommend tilting only to the mean.

We produce soft conditional forecasts by considering the median inflation expectation for SPF one, two and five-year ahead, but also long-term real GDP growth expectations. The upper panel of Figure 13 shows that tilting the unconditional inflation forecasts to median SPF values for inflation and GDP yields inflation forecasts which are robust to the impact of the COVID-19 observations on the parameter estimates.²⁹ Nevertheless, these conditions are not sufficient to ensure overall stability in the Gaussian VAR with a weakly informative prior. Furthermore, the lower panel of Figure 13 shows forecasts when the unconditional distribution is tilted to five-ahead inflation only. Thus, whereas the long-run expectations are crucial to ensure this stability, the short-run expectations are important to ensure a plausible profile of the future inflation path.

4 Modelling the COVID-19 shock with univariate and multivariate t -distributed errors

Why are the coefficient estimates of a VAR with multivariate fat-tailed errors more resilient against the impact of the COVID-19 observations as compared to the single equation Phillips curve models with univariate heavy-tailed errors?

In the single equation framework fatter tails may primarily capture abnormal variations associated with the dependent variable, but not so much those induced by an exogenous regressor (since its coefficients can always be shrunk to zero to mitigate the impact on the errors). Intuitively, in the multivariate setup all variables are re-scaled by a common volatility factor, which is informed by the innovations in all endogenous variables, while in the single equation setup this volatility factor is informed only by the residuals of the dependent variable. For technical details, see Appendix B.

To illustrate this point, consider the following example with a simplified version of the Phillips curve model in (1) based only on one slack measure, namely the output gap (Ogap), its own lag and no other exogenous regressors:

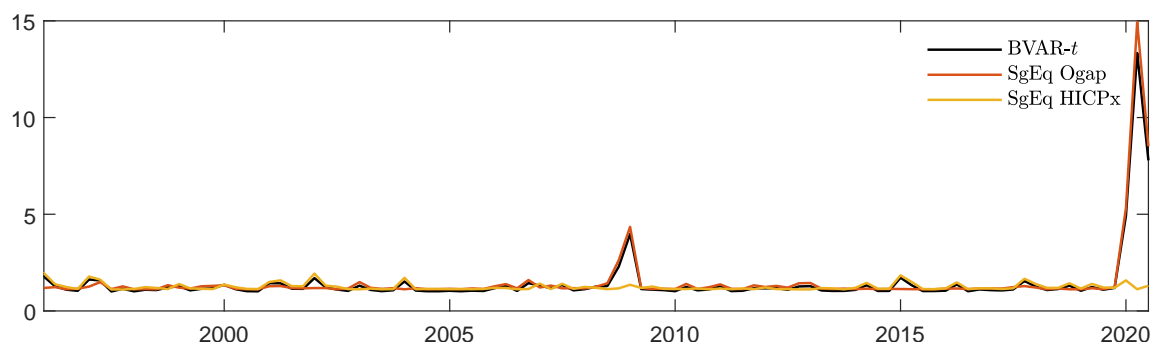
$$\pi_t = \alpha + \rho\pi_{t-1} + \beta\text{Ogap}_{t-1} + \epsilon_t \quad (2)$$

We estimate the coefficients of (2) using system and single equation techniques. For the system estimation, we use: (i) a Bayesian VAR with Gaussian errors, (ii) a Bayesian VAR with multivariate t -distributed errors and (iii) a counterfactual BVAR with exogenous time-varying variance Λ (as detailed below). For the single equation estimation we use a univariate version the three models listed above. As prior distribution we assume a diffuse (non-informative) prior for all parameters. We obtain different estimates of λ from the models with t errors: the BVAR- t (λ_{VAR}) and the single equation regression of the VAR equation for the output gap (λ_{Ogap}) and HICP inflation excluding energy and food (λ_{HICPx}), respectively. Subsequently, the posterior medians of these λ 's are fed into the model with exogenous time-varying variance to study the impact of allowing for fat tails on the coefficient estimates.

Figure 14 shows the implied volatility by these λ 's. Both λ_{VAR} and λ_{Ogap} spike when the COVID-19 pandemic wrecked havoc in the euro area, whereas λ_{HICPx} remains relatively flat during this period. This illustrates that different information sets may crucially affect these λ 's in the multi- and univariate setup with t errors.

²⁹Our codes adjust the tilting function included in the BEAR toolbox, see Dieppe et al. (2016).

Figure 14: Implied volatility under multivariate and univariate t errors



Note: Posterior median of implied volatility in the following models with t errors: (i) BVAR- t using one lag (BVAR- t) and the single equation regression of the VAR model for the output gap (SgEq Ogap) and inflation (SgEq HICPx).

How do these different estimates of λ affect the inference about the coefficients in equation (2)? Table 3 shows the posterior mean of the coefficients using the estimation choices described above for (a) the pre-COVID-19 sample and (b) with the COVID-19 observations included in the estimation sample.

Table 3: Simple Phillips curve coefficients with output gap

Coeff.	Gaussian		Fat-tailed		Gaussian(λ_{VAR})		Gaussian(λ_{Ogap})		Gaussian(λ_{HICPx})	
	System	SgEq	System	SgEq	System	SgEq	System	SgEq	System	SgEq
α	0.603	0.605	0.621	0.657	0.631	0.631	0.582	0.581	0.672	0.671
ρ	0.566	0.564	0.546	0.515	0.538	0.539	0.588	0.589	0.503	0.503
β	0.078	0.078	0.085	0.089	0.086	0.086	0.074	0.074	0.091	0.091
(a) Simple Phillips curve estimated until 2019:Q4										
Coeff.	Gaussian		Fat-tailed		Gaussian(λ_{VAR})		Gaussian(λ_{Ogap})		Gaussian(λ_{HICPx})	
	System	SgEq	System	SgEq	System	SgEq	System	SgEq	System	SgEq
α	0.575	0.574	0.626	0.612	0.635	0.635	0.586	0.587	0.619	0.622
ρ	0.581	0.582	0.540	0.544	0.533	0.534	0.583	0.583	0.538	0.536
β	0.055	0.054	0.084	0.059	0.085	0.084	0.074	0.074	0.059	0.059
(b) Simple Phillips curve estimated until 2020:Q3										

Note: Posterior mean of Phillips curve coefficients based on various modelling assumption. The first top row denotes the class of the error distribution, i.e. Gaussian is normal errors, Fat-tailed is t -errors, Gaussian(\cdot) is normal errors with exogenous time-varying variance and the second row describes whether the equation was estimated as a system (System) or as single equation (SgEq).

Focusing on the pre-COVID-19 sample (panel (a)), the estimated slope of the Phillips curve β is the same for the Gaussian model estimated as a system or as a single equation (column Gaussian). However, for the fat-tailed model (column Fat-tailed), the slope coefficients differ slightly between the two versions. These differences can be explained by the fact that the λ 's differ across the two options (multi- versus univariate t errors), mainly because of differences during the financial crisis, see Figure 14.

As discussed in Section 2, the COVID-19 observations lead to a decrease in the Phillips curve slope by about 50% percent for the output gap, see panel (b) (columns Gaussian and Fat-tailed, SgEq). This strong parameter revision in the univariate fat-tailed model (but also in the Gaussian model) occurs because the COVID-19 observations are not heavily discounted since the λ_{HICPx} remains rather flat.

In contrast, when β is estimated under multivariate t -distributed errors (column Fat-tailed, System), the slope coefficient hardly changes. In addition, when estimating this equation with known λ from the VAR (column Gaussian(λ_{VAR})) or the output gap equation (column Gaussian(λ_{Ogap})), then the slope coefficients also remain largely unchanged, while when plugging in λ_{HICP_x} it changes substantially (column Gaussian(λ_{HICP_x})).

5 Conclusions

The parameters of models frequently employed to analyse inflation by central bankers and other practitioners are affected by the COVID-19 observations and this has a bearing also on the inflation forecast path.

For VAR models, relaxing the assumption of normal errors and allowing them to follow a multivariate t -distribution appears to tackle the parameter instability problem. Within this model tail events are accommodated by the residuals, which diminishes the impact on the parameters. Within a standard Gaussian VAR, the choice of prior matters, with more informative specifications or a higher degree of prior shrinkage (more suitable for larger VARs) acting in favour of stabilising the results. Both a fat-tailed VAR and a large Gaussian VAR with more prior shrinkage also ensure stable unconditional forecasts. Nevertheless, as past data does not record an event similar to COVID-19, adding relevant off-model information to the purely model-based forecast is crucial in producing more trustworthy forecasts.

We find that a small-scale Gaussian VAR can still be used for producing conditional forecasts which are well-behaved (namely insensitive to the change in parameters when adding the COVID-19 observations) when relevant off-model information is employed. Both hard conditioning (using official projections for the explanatory variables) and soft conditioning (using information from the Survey of Professional Forecasters) help.

For single equation models, such as the ‘thick modelling’ Phillips curve framework, allowing for the distribution of the errors to have fatter tails does not ensure parameter stability. Our observations point to the idea that multiple equation systems such as VARs are better equipped to deal with the impact of abnormal observations than single equation models because they are more efficient in processing the information jointly captured in the developments in all considered variables. When it comes to forecasting, averaging information from conditional forecasts from many models within a thick modelling framework appears to offer a hedge with respect to parameter instability due to the COVID-19 observations.

As a lesson for practitioners, we suggest that before deriving conclusions based on existing models, one has to check and be aware of how the COVID-19 observations affect the parameter estimates and the implied forecasts. It can be that as new observations pile up the distortion in traditionally estimated parameters diminishes. Yet, a simple model with Gaussian and homoscedastic errors will not account for the heightened volatility in the COVID-19 times.

This paper shows that allowing for VARs to have fat-tailed errors ensures more stable parameters and forecasts. Yet, a more in-depth investigation on how such a model would have fared historically in forecasting is warranted. Also, as yet, it is not possible to ascertain based on existing data what works well in forecasting during the pandemic.

References

- ANGELINI, E., M. DARRACQ PARIÈS, S. ZIMIC, AND M. DAMJANOVIĆ (2020): “ECB-BASIR: a primer on the macroeconomic implications of the Covid-19 pandemic,” Working Paper 2431, European Central Bank.
- ASCARI, G., G. FAGIOLO, AND A. ROVENTINI (2015): “Fat-Tail Distributions And Business-Cycle Models,” *Macroeconomic Dynamics*, 19, 465–476.
- BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2016): “Measuring Economic Policy Uncertainty,” *The Quarterly Journal of Economics*, 131, 1593–1636.
- BAQAEI, D. AND E. FARHI (2020): “Supply and Demand in Disaggregated Keynesian Economies with an Application to the Covid-19 Crisis,” Working Paper 27152, National Bureau of Economic Research.
- BAÑBURA, M. AND E. BOBEICA (2020): “Does the Phillips curve help to forecast euro area inflation?” Working Paper 2471, European Central Bank.
- BAÑBURA, M., F. BRENNA, M. CICCARELLI, C. MONTES-GALDON, F. MAROTTA, AND J. PAREDES (2020a): “Forecasting with Conditional Uncertainty,” Mimeo.
- BAÑBURA, M., F. BRENNA, J. PAREDES, AND F. RAVAZZOLO (2020b): “How to better combine Bayesian VARs with SPF forecasts: a performance assessment,” Mimeo.
- BAÑBURA, M., D. GIANNONE, AND L. REICHLIN (2010): “Large Bayesian Vector Auto Regressions,” *Journal of Applied Econometrics*, 25, 71–92.
- BOBEICA, E. AND M. JAROCIŃSKI (2019): “Missing Disinflation and Missing Inflation: A VAR Perspective,” *International Journal of Central Banking*, 15, 199–232.
- BOBEICA, E. AND A. SOKOL (2019): “Drivers of underlying inflation in the euro area over time: a Phillips curve perspective,” Economic Bulletin 4, Article 2, European Central Bank.
- BODENSTEIN, M., G. CORSETTI, AND L. GUERRIERI (2020): “Social Distancing and Supply Disruptions in a Pandemic,” Discussion Paper 14629, Centre for Economic Policy Research.
- BUNDESBANK (2016): “The Phillips curve as an instrument for analysing prices and forecasting inflation in Germany,” Monthly Report 4, Article 1, Deutsche Bundesbank.
- CARRIERO, A., T. E. CLARK, M. MARCELLINO, AND E. MERTENS (2020): “VARs in 2020: Dealing with outliers and the lower bound on interest rates,” Working Paper 20-32, Federal Reserve Bank of Cleveland.
- (2021): “Addressing COVID-19 Outliers in BVARs with Stochastic Volatility,” Working Paper 21-02, Federal Reserve Bank of Cleveland.
- CHAN, J. C. AND C. Y. HSIAO (2014): “Estimation of Stochastic Volatility Models with Heavy Tails and Serial Dependence,” in *Bayesian Inference in the Social Sciences*, John Wiley & Sons, Ltd, chap. 6, 155–176.

- CHAN, J. C. C. (2020): “Large Bayesian VARs: A Flexible Kronecker Error Covariance Structure,” *Journal of Business & Economic Statistics*, 38, 68–79.
- 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): “Comment on ‘On the Fit of New Keynesian Models’,” *Journal of Business and Economic Statistics*, 25, 143–151.
- CICCARELLI, M. AND C. OSBAT (2017): “Low inflation in the euro area: Causes and consequences,” Occasional Paper Series 181, European Central Bank.
- COCHRANE, J. (2020): “Coronavirus monetary policy,” in *Economics in the time of COVID-19: A new eBook*, ed. by R. Baldwin and B. W. di Mauro, Centre for Economic Policy Research, chap. 12, 105–108.
- COGLEY, T. AND T. J. SARGENT (2005): “Drifts and volatilities: monetary policies and outcomes in the post WWII US,” *Review of Economic dynamics*, 8, 262–302.
- CONSTANCIO, V. (2015): “Understanding Inflation Dynamics and Monetary Policy,” Speech at the Jackson Hole Economic Policy Symposium.
- CÚRDIA, V., M. DEL NEGRO, AND D. L. GREENWALD (2014): “Rare shocks, great recessions,” *Journal of Applied Econometrics*, 29, 1031–1052.
- DE VINCENT-HUMPHREYS, R., I. DIMITROVA, E. FALCK, L. HENKEL, AND A. MEYLER (2019): “Twenty years of the ECB Survey of Professional Forecasters,” Economic Bulletin 1, Article 1, European Central Bank.
- DIEPPE, A., B. VAN ROYE, AND R. LEGRAND (2016): “The BEAR toolbox,” Working Paper 1934, European Central Bank.
- DOAN, T., R. B. LITTERMAN, AND C. A. SIMS (1983): “Forecasting and Conditional Projection Using Realistic Prior Distributions,” Working Paper 1202, National Bureau of Economic Research.
- ECB (2016): “A guide to the Eurosystem/ECB staff macroeconomic projection exercises,” Technical report, European Central Bank.
- (2017): “Understanding the subdued inflation dynamics,” Annual Report, Box 2, European Central Bank.
- EICHENBAUM, M. S., S. REBELO, AND M. TRABANDT (2020): “Epidemics in the Neoclassical and New Keynesian Models,” Working Paper 27430, National Bureau of Economic Research.
- ESER, F., P. KARADI, P. R. LANE, L. MORETTI, AND C. OSBAT (2020): “The Phillips Curve at the ECB,” Working Paper 2400, European Central Bank.

- FAGIOLO, G., M. NAPOLETANO, AND A. ROVENTINI (2008): “Are output growth-rate distributions fat-tailed? some evidence from OECD countries,” *Journal of Applied Econometrics*, 23, 639–669.
- FORNARO, L. AND M. WOLF (2020): “Covid-19 Coronavirus and Macroeconomic Policy,” Working Paper 1168, Barcelona Graduate School of Economics.
- GANICS, G. AND F. ODENDAHL (2021): “Bayesian VAR Forecasts, Survey Information and Structural Change in the Euro Area,” *International Journal of Forecasting*, 37, 971–999.
- GEWEKE, J. (1993): “Bayesian treatment of the independent student-t linear model,” *Journal of Applied Econometrics*, 8, 19–40.
- GRANGER, C. W. J. AND Y. JEON (2004): “Thick modeling,” *Economic Modelling*, 21, 323–343.
- GROTHE, M. AND A. MEYLER (2018): “Inflation Forecasts: Are Market-Based and Survey-Based Measures Informative?” *International Journal of Financial Research*, 9, 171–188.
- GUERRIERI, V., G. LORENZONI, L. STRAUB, AND I. WERNING (2020): “Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?” Working Paper 26918, National Bureau of Economic Research.
- HARTWIG, B. (2020): “Robust Inference in Time-Varying Structural VAR Models: The DC-Cholesky Multivariate Stochastic Volatility Model,” Discussion Paper 34/2020, Deutsche Bundesbank.
- (2021): “Bayesian VARs and prior calibration in times of COVID-19,” Mimeo.
- 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.
- JAROCIŃSKI, M. (2010): “Conditional forecasts and uncertainty about forecast revisions in vector autoregressions,” *Economics Letters*, 108, 257 – 259.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): “Measuring Uncertainty,” *American Economic Review*, 105, 1177–1216.
- KOOP, G. AND D. KOROBILIS (2012): “Forecasting Inflation Using Dynamic Model Averaging,” *International Economic Review*, 53, 867–886.
- LENZA, M. AND G. PRIMICERI (2020): “How to Estimate a VAR after March 2020,” Working Paper 27771, National Bureau of Economic Research.
- MISHKIN, F. S. (2011): “Monetary Policy Strategy: Lessons from the Crisis,” Working Paper 16755, National Bureau of Economic Research.
- NI, S. AND D. SUN (2005): “Bayesian Estimates for Vector Autoregressive Models,” *Journal of Business & Economic Statistics*, 23, 105–117.
- O’BRIEN, D., C. DUMONCEL, AND E. GONÇALVES (2021): “The role of demand and supply factors in HICP inflation during the COVID-19 pandemic – a disaggregated perspective,” Economic Bulletin 1, Article 1, European Central Bank.

- PRIMICERI, G. AND A. TAMBALOTTI (2020): “Macroeconomic Forecasting in the Time of COVID-19,” Manuscript June 2020, Northwestern University.
- PRIMICERI, G. E. (2005): “Time Varying Structural Vector Autoregressions and Monetary Policy,” *Review of Economic Studies*, 72, 821–852.
- ROBERTSON, J. C., E. W. TALLMAN, AND C. H. WHITEMAN (2005): “Forecasting Using Relative Entropy,” *Journal of Money, Credit and Banking*, 37, 383–401.
- ROSSI, B. (2020): “Forecasting in the Presence of Instabilities: How Do We Know Whether Models Predict Well and How to Improve Them,” Discussion Paper 14472, Center for Economic Policy Research.
- SCHORFHEIDE, F. AND D. SONG (2020): “Real-Time Forecasting with a (Standard) Mixed-Frequency VAR During a Pandemic,” Working Paper 20-26, Federal Reserve Bank of Philadelphia.
- SIMS, C. A. AND T. ZHA (1998): “Bayesian Methods for Dynamic Multivariate Models,” *International Economic Review*, 39, 949–968.
- TALLMAN, E. W. AND S. ZAMAN (2020): “Combining survey long-run forecasts and nowcasts with BVAR forecasts using relative entropy,” *International Journal of Forecasting*, 36, 373–398.
- WAGGONER, D. F. AND T. ZHA (1999): “Conditional Forecasts In Dynamic Multivariate Models,” *The Review of Economics and Statistics*, 81, 639–651.
- YELLEN, J. (2015): “Inflation dynamics and monetary policy, Speech at the Philip Gamble Memorial Lecture, University of Massachusetts, Amherst, Massachusetts,” 24 September 2015, Federal Reserve System.

Appendix A The proposed econometric solution for estimating BVARs after March 2020. Errors with a t -distribution Dummy observation prior

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

$$y_t = a_0 + A_1 y_{t-1} + \dots + A_p y_{t-p} + \epsilon_t, \quad \epsilon_t \sim t(0, \Sigma, \nu) \quad (3)$$

where a_0 is an $n \times 1$ vector of intercepts and A_1, \dots, A_p are $n \times n$ coefficient matrices. Let $x'_t = (1, y'_{t-1}, \dots, y'_{t-p})$ be a $1 \times k$ vector of an intercept and lags with $k = 1 + np$ and $A = (a_0, A_1, \dots, A_p)'$ is a $k \times n$ matrix. The error term ϵ_t follows an independent multivariate t distribution with covariance matrix Σ of dimension $n \times n$ and degree of freedom ν .

The multivariate t distribution of the error term can be represented by a scale mixture of normal distributions, see [Geweke \(1993\)](#) and [Ni and Sun \(2005\)](#).

$$y'_t = x'_t A + \epsilon_t, \quad \epsilon_t | \lambda_t \sim N(0, \lambda_t \Sigma) \quad (4)$$

where λ_t is a latent state variable that is inverse gamma distributed, i.e. $\lambda_t \sim IG(\nu/2, \nu/2)$.

Next, we stack the observation over $t = 1, \dots, T$, which yields

$$Y = XA + E, \quad E \sim MN(0, \Sigma \otimes \Lambda) \quad (5)$$

where Y, X , and E are respectively of dimensions $T \times n$, $T \times k$, and $T \times n$. MN denotes the matrix-variate normal distribution, \otimes is the Kronecker product, $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_T)$ is of dimension $T \times T$.

We assume that the prior for $p(A, \Sigma)$, $p(\Lambda)$ and $p(\nu)$ are independent. For (A, Σ) we adopt a standard normal-inverse-Wishart prior

$$\Sigma \sim IW(S_0, v_0), \quad A | \Sigma \sim MN(A_0, \Sigma \otimes V_A) \quad (6)$$

For Λ , we assume that λ_t is independently inverse gamma distributed

$$\lambda_t \sim IG(\nu/2, \nu/2) \quad (7)$$

and for ν we follow [Chan and Hsiao \(2014\)](#) and assume a uniform prior

$$\nu \sim U(2, \bar{\nu}) \quad (8)$$

where $\bar{\nu} = 50$ is sufficiently large to approximate normal errors.

Because of the prior structure, the posterior draws can be obtained by sequentially sampling from: (1) $p(A, \Sigma | Y, \Lambda)$; (2) $p(\Lambda | Y, A, \Sigma)$ and (3) $p(\nu | Y, A, \Sigma, \Lambda)$. As detailed in [Chan \(2020\)](#), natural conjugacy of the prior $p(A, \Sigma)$ survives even with a general covariance structure in Λ . Therefore, one can derive conditional posterior quantities for $p(A, \Sigma | Y, \Lambda)$ that are still normal-inverse-Wishart distributed. The

conditional posterior are given by

$$(A|Y, \Sigma, \Lambda) \sim MN(\hat{A}, \Sigma \otimes K_A^{-1}) \quad (9)$$

$$(\Sigma|Y, \Lambda) \sim IW(\hat{S}, v_0 + T) \quad (10)$$

where

$$\begin{aligned} K_A &= V_A^{-1} + X' \Lambda^{-1} X \\ \hat{A} &= K_A^{-1} (V_A^{-1} A_0 + X' \Lambda^{-1} Y) \\ \hat{S} &= S_0 + A_0' V_A^{-1} A_0 + Y' \Lambda^{-1} Y - \hat{A}' K_A^{-1} \hat{A}. \end{aligned}$$

The posterior distribution for $p(\Lambda|Y, A, \Sigma)$ can be easily derived as the $\lambda_1, \dots, \lambda_T$, are conditionally independent given the data and (A, Σ) . Therefore, one can show that the posterior of λ_t is given by

$$(\lambda_t|Y, A, \Sigma) \sim IG\left(\frac{\nu+n}{2}, \frac{1}{2}(\nu + (y'_t - x'_t A)\Sigma^{-1}(y'_t - x'_t A)')\right)$$

Next, the posterior distribution for $p(\nu|Y, \Lambda, A, \Sigma)$ conditional on Λ is independent from the data and (A, Σ) . As shown by [Chan and Hsiao \(2014\)](#), the conditional posterior for ν is given by

$$\begin{aligned} (\nu|\Lambda) &\propto p(\Lambda|\nu)p(\nu) \\ &\propto \frac{(\nu/2)^{T\nu/2}}{\Gamma(\nu/2)^T} (\prod_{t=1}^T \lambda_t)^{-\frac{\nu}{2}+1} \exp\left(-\frac{\nu}{2} \sum_{t=1}^T \lambda_t^{-1}\right) \end{aligned}$$

for $2 < \nu < \bar{\nu}$, and 0 otherwise. This density is non-standard and we use the independence-chain Metropolis-Hastings algorithm of [Chan and Hsiao \(2014\)](#) to sample ν .

To summarize, we can sequentially sample the parameters $(A^{(s)}, \Sigma^{(s)}, \Lambda^{(s)}, \nu^{(s)})$ using a Gibbs sampler with Metropolis-Hasting step:

1. Draw $\Sigma^{(s)}$ from an $IW(\hat{S}, T + \nu_0)$ distribution
2. Draw $A^{(s)}$ from an $MN(\hat{A}, \Sigma \otimes K_A^{-1})$
3. Draw $\lambda_t^{(s)}$ from an $IG\left(\frac{\nu+n}{2}, \frac{1}{2}(\nu + (y'_t - x'_t A)\Sigma^{-1}(y'_t - x'_t A)')\right)$ for $t = 1, \dots, T$
4. Draw $\nu^{(s)}$ from a proposal distribution and accept/reject using a Metropolis-Hasting step

Inference in the BVAR model with t errors. To conduct inference in the BVAR model with t errors, it is convenient to work with the normal mixture representation of (3).

$$y_t = a_0 + A_1 y_{t-1} + \dots + A_p y_{t-p} + \lambda_t^{\frac{1}{2}} C u_t, \quad u_t \sim N(0, I), \quad \lambda_t \sim IG\left(\frac{\nu}{2}, \frac{\nu}{2}\right) \quad (11)$$

where $CC' = \Sigma$. Thus, conditional on λ_t , the BVAR with t errors can be represented as a standard Gaussian VAR with independent but not identical distributed errors. Specifically, the error term exhibits idiosyncratic heteroskedasticity which is controlled by λ_t .

To compute impulse response functions, unconditional and conditional forecasts from (11), we follow the exposition of [Jarociński \(2010\)](#) and derive impulse response coefficients as well as derive the projection

matrices H for the variables and R for the residuals of the BVAR model with t errors. We start by simplifying the derivation and define the VAR in terms $z_t = y_t - \mu$ – deviations of y_t from the central level μ defined as:

$$\mu = (I - (A_1 + \dots + A_p))^{-1}a_0$$

The VAR in the companion form is

$$\tilde{z}_t = F\tilde{z}_{t-1} + G\tilde{u}_t \quad (12)$$

where

$$F = \begin{pmatrix} A_1 & \dots & A_p \\ I & & 0 \end{pmatrix}, G = \begin{pmatrix} \lambda_t^{\frac{1}{2}}C & 0 \\ 0 & 0 \end{pmatrix}, \tilde{z}_t = \begin{pmatrix} z_t \\ \vdots \\ z_{t-p+1} \end{pmatrix}, \tilde{u}_t = \begin{pmatrix} u_t \\ \vdots \\ 0 \end{pmatrix}.$$

Take \tilde{z}_{S+h} and, using (12), recursively substitute $\tilde{z}_{S+h-1}, \dots, \tilde{z}_{S+1}$. In the first N rows we obtain z_{S+h} expressed in terms of data up to S and subsequent errors

$$z_{S+h} = F_{(1..N,.)}^h \tilde{z}_S + \lambda_{S+h}^{\frac{1}{2}} C u_{S+h} + \Phi_1 \lambda_{S+h-1}^{\frac{1}{2}} C u_{S+h-1} + \dots + \Phi_{h-1} \lambda_{S+1}^{\frac{1}{2}} C u_{S+1} \quad (13)$$

where Φ_j is the upper left $n \times n$ block of F^j (F to the power j). $\Phi_j C \lambda_t^{\frac{1}{2}}$ is the matrix of orthogonalized impulse response after j periods. $F_{(1..N,.)}^h$ is the matrix composed of the first N rows of F^h . The stacked vector of $z_{S+1} \dots z_T$ can be written as

$$\begin{pmatrix} z_{S+1} \\ \vdots \\ z_T \end{pmatrix} = \begin{pmatrix} F_{(1..N,.)} \\ \vdots \\ F_{(1..N,.)}^{T-S} \end{pmatrix} \tilde{z}_S + \begin{pmatrix} \lambda_{S+1}^{\frac{1}{2}} C & 0 & 0 \\ \Phi_1 \lambda_{S+1}^{\frac{1}{2}} C & \lambda_{S+2}^{\frac{1}{2}} C & 0 \\ \Phi_{T-S+1} \lambda_{T-S+1}^{\frac{1}{2}} C & \dots \Phi_1 \lambda_{T-1}^{\frac{1}{2}} C & \lambda_T^{\frac{1}{2}} C \end{pmatrix} \begin{pmatrix} u_{S+1} \\ \vdots \\ u_T \end{pmatrix} \quad (14)$$

or shortly as:

$$z = H\tilde{z}_S + Ru \quad (15)$$

The above derivation are very similar to those of standard stationary VARs. The only difference is that the forecast errors are now conditionally heteroskedastic on the sequence of $\lambda_{S+1}, \dots, \lambda_{T-S+1}$ which follows directly from the normal mixture representation of the multivariate t distribution.

Appendix B Heavy tails in single and multiple equation models

To understand the effect of univariate and multivariate t -distributed errors on the coefficients of a (multivariate) linear regression model, consider the VAR in (12) and rewrite it in seemingly unrelated regression form using the scale-mixture Gaussian representation of the t distribution:

$$y = (I_n \otimes X)a + e, \quad e \sim N(0, \Sigma \otimes \Lambda) \quad (16)$$

where $y = \text{vec}(y)$, $a = \text{vec}(A)$, $e = \text{vec}(E)$ and I_n is an identity matrix of dimension n and Λ is diagonal covariance matrix. For simplicity it is instructive to focus on the conditional posterior mean of the coefficients under a diffuse (non-informative) prior.³⁰

The conditional posterior mean of a is given by

$$\begin{aligned}\hat{a}_{GLS} &= ((I_n \otimes X)'(\Sigma^{-1} \otimes \Lambda^{-1})(I_n \otimes X))^{-1} (I_n \otimes X)'(\Sigma^{-1} \otimes \Lambda^{-1})y \\ &= ((\Sigma^{-1} \otimes X' \Lambda^{-1} X))^{-1} (\Sigma^{-1} \otimes X' \Lambda^{-1})y \\ &= (I_n \otimes (X' \Lambda^{-1} X))^{-1} X' \Lambda^{-1} y \\ &= \begin{pmatrix} (X' \Lambda^{-1} X)^{-1} X' \Lambda^{-1} y_1 \\ (X' \Lambda^{-1} X)^{-1} X' \Lambda^{-1} y_2 \\ \dots \\ (X' \Lambda^{-1} X)^{-1} X' \Lambda^{-1} y_n \end{pmatrix} \underbrace{=}_{\Lambda=I_T} \begin{pmatrix} (X' X)^{-1} X' y_1 \\ (X' X)^{-1} X' y_2 \\ \dots \\ (X' X)^{-1} X' y_n \end{pmatrix} = \hat{a}_{OLS}\end{aligned}$$

where GLS and OLS denote the generalized and ordinary least squares estimate.

These formulas reveal two interesting properties. First, given Λ , the system estimation (System) is equivalent to single equation estimation (SgEq). Second, given $\Lambda = I_T$ then $\hat{a}_{GLS} = \hat{a}_{OLS}$ (error term is homoskedastic and independently Gaussian distributed). Note that under t -distributed errors $\Lambda \neq I_T$, hence, $\hat{a}_{GLS} \neq \hat{a}_{OLS}$.

Because of the first property, if system estimation does not yield the same estimate as equation by equation estimation, then it must be the case that the implied volatility must be different in the univariate setting as compared to the multivariate setting, i.e. $\Lambda^{SgEq} \neq \Lambda^{System}$.

To see this, suppose a and Σ are fixed, using (16), the error is $\epsilon_t = y_t - (I_n \otimes x'_t)a$. The conditional posterior of λ_t from a regression with univariate t -distributed errors for variable $y_{i,t}$ is given by

$$\lambda_t^{SgEq} \sim IG\left(\frac{\nu+1}{2}, \frac{1}{2}\left(\nu + \frac{1}{\sigma_{ii}^2} \epsilon_{i,t}^2\right)\right)$$

whereas the conditional posterior of λ_t with multivariate t -distributed errors for y_t is given by

$$\lambda_t^{System} \sim IG\left(\frac{\nu+n}{2}, \frac{1}{2}\left(\nu + \epsilon'_t \Sigma^{-1} \epsilon_t\right)\right)$$

Thus, λ_t depends on different information sets (abstracting from distributional differences). Particularly, λ_t^{SgEq} depends only on the forecast error $\epsilon_{i,t}$ of the i -th variable ($s_t^2 = \frac{1}{\sigma_{ii}^2} \epsilon_{i,t}^2$) in the univariate settings, while under the multivariate distribution, λ_t^{System} depends on the vector of forecast errors ϵ_t ($s_t^2 = \epsilon'_t \Sigma^{-1} \epsilon_t$, scalar).

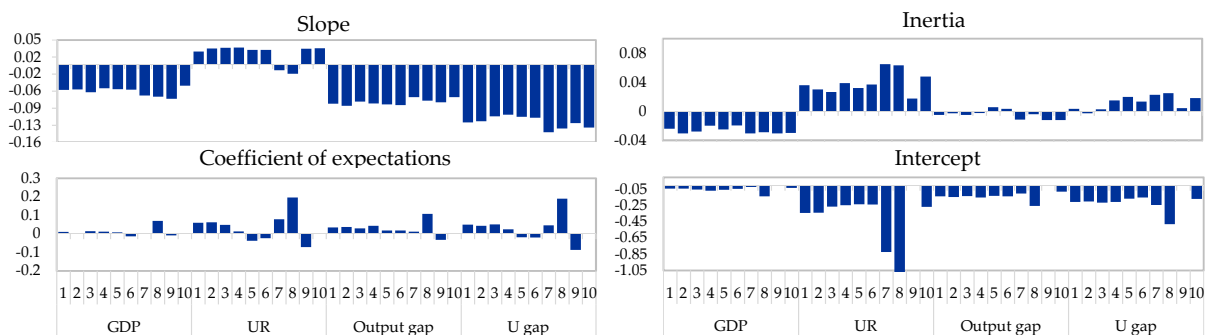
For instance, suppose there is a large shock in the j -th variable but not in the i -th variable. Heuristically, the posterior of λ_t^{System} will be larger than λ_t^{SgEq} because there was no unusually large forecast error in the i -th variable. As a consequence, when the parameters are updated in the next iteration of the Gibbs sampler, the data will be downweighed more strongly in the multivariate setup as opposed to the univariate setup. The estimated coefficients a are then different because $\lambda_t^{System} \neq \lambda_t^{SgEq}$.

³⁰Note that this posterior mean is equivalent to the generalized least squares estimate.

Appendix C Additional Tables and Figures

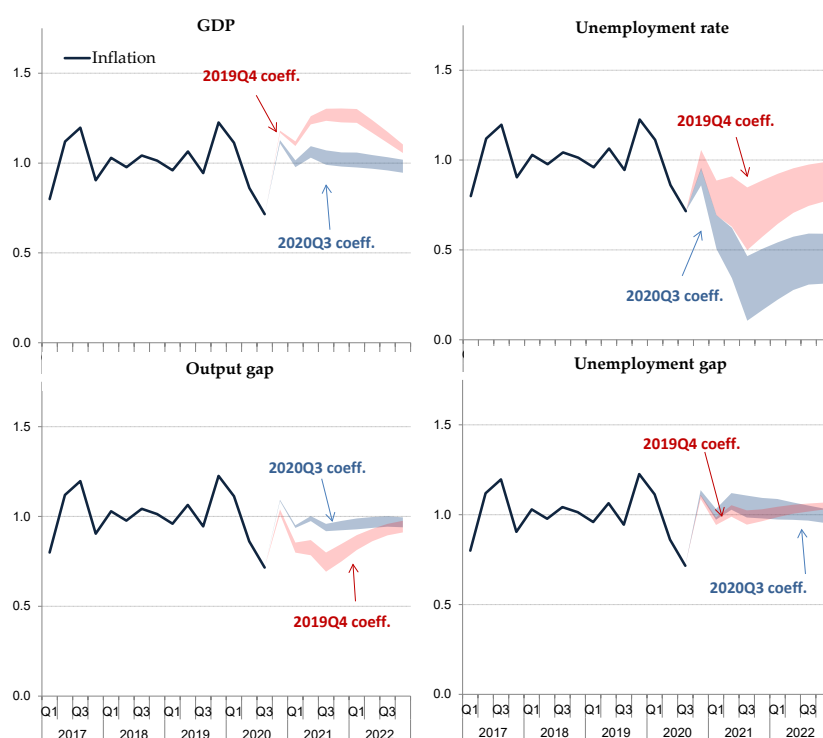
C.1 Phillips curves with time-varying parameters

Figure 15: Change in Phillips curve endpoint parameters under time variation in coefficients and error variance. Sample ending in 2020Q3 vs. in 2019Q4



Note: The bars show the change in coefficients for each specification per slack measure. 1,2,...,10 indicate specifications with various measures of expectations, in the following order: 1 - 6 are Consensus 1,2,...,6 quarters ahead, 7 - 9 are SPF 1,2,5 years ahead, 10 - no expectation term. No intercept was included in the specification with SPF 5 years ahead. For the unemployment rate and gap the change in slope refers to coefficients with reversed sign. The specification allows the coefficients and log variance of the errors to follow a random walk. The choice of priors is very conservative, a priori the coefficients are assumed to barely move over time, while the log-volatility process is expected to capture more variation.

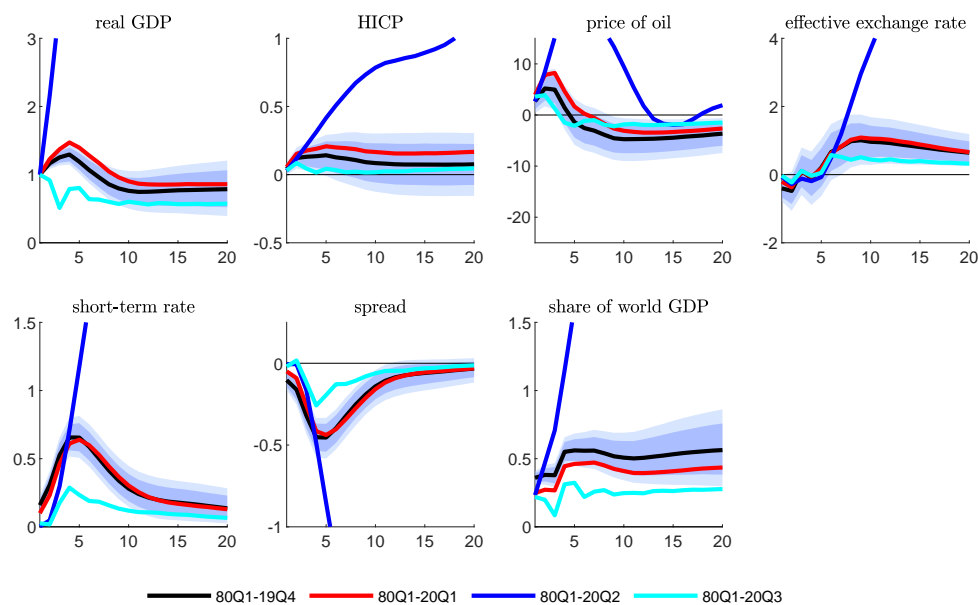
Figure 16: Phillips curve conditional forecasts (time-varying parameter models)



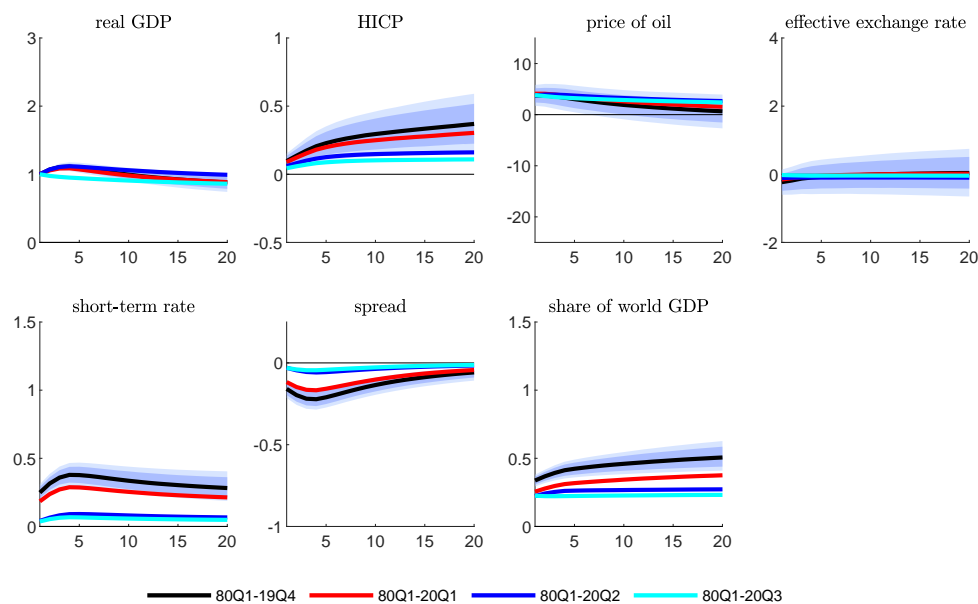
Note: The ranges are based on the endpoint parameter estimates on the sample up to 2019Q4 (red range) and up to 2020Q3 (blue range). Forecast starts in 2020Q4. The ranges show point forecasts covering specifications with various inflation expectation measures.

C.2 BVAR with Gaussian errors

Figure 17: Impulse response functions in a small BVAR (various prior distributions)



(a) Weakly informative prior

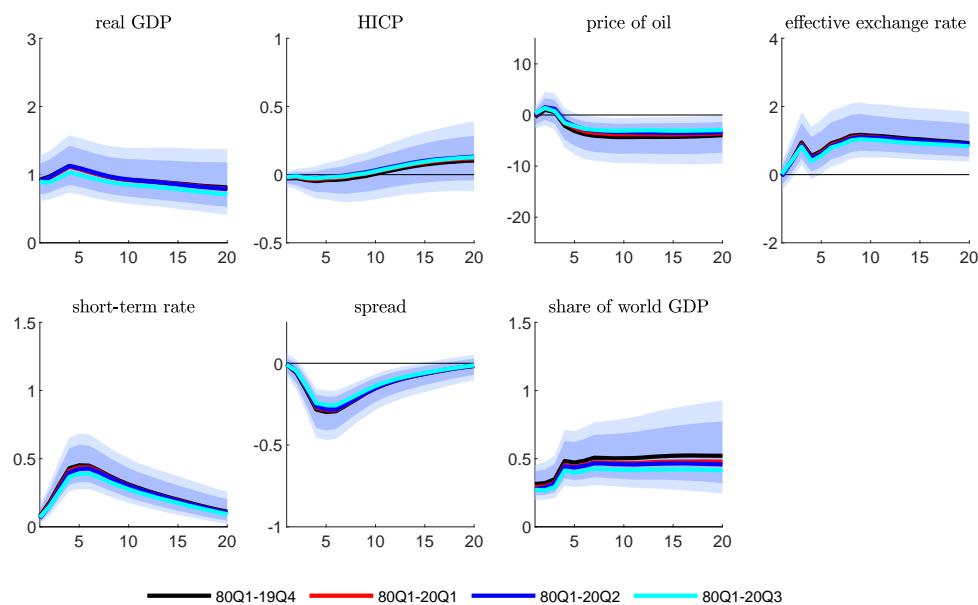


(b) Strong Sims and Zha prior

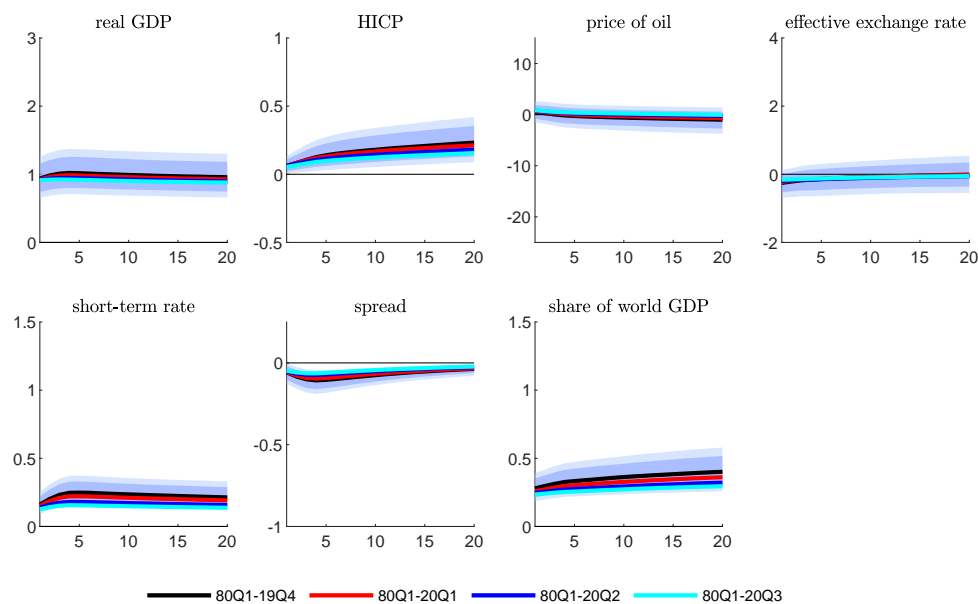
Note: Impulse response functions to a one standard deviation shock in real GDP. Thick lines are median estimates and the dark (light) blue area is the 68% (90%) credible interval for the estimation window until 19Q4.

C.3 BVAR with fat-tailed errors

Figure 18: Impulse response functions in a small BVAR- t (various prior distributions)



(a) Weakly informative prior

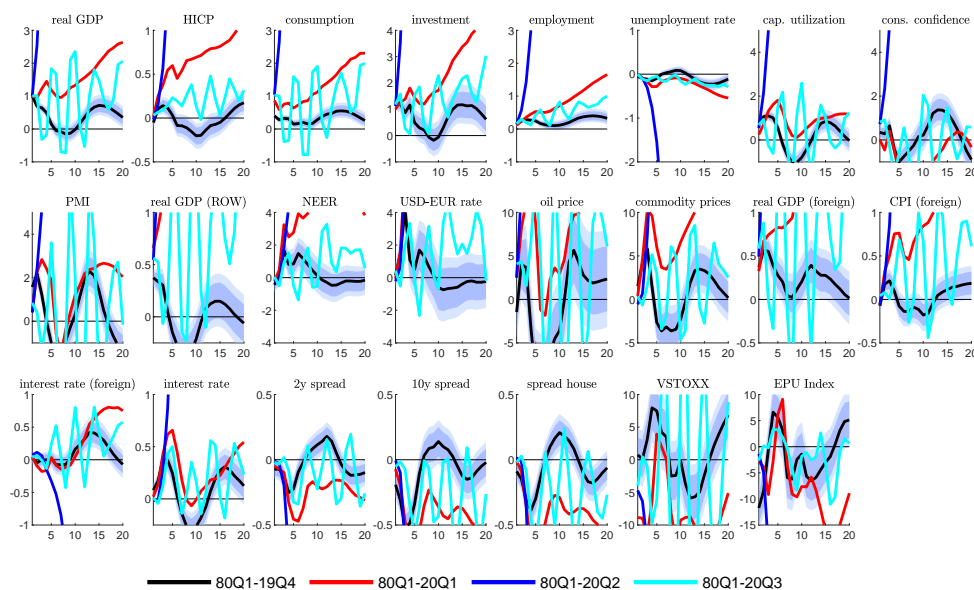


(b) Strong Sims and Zha prior

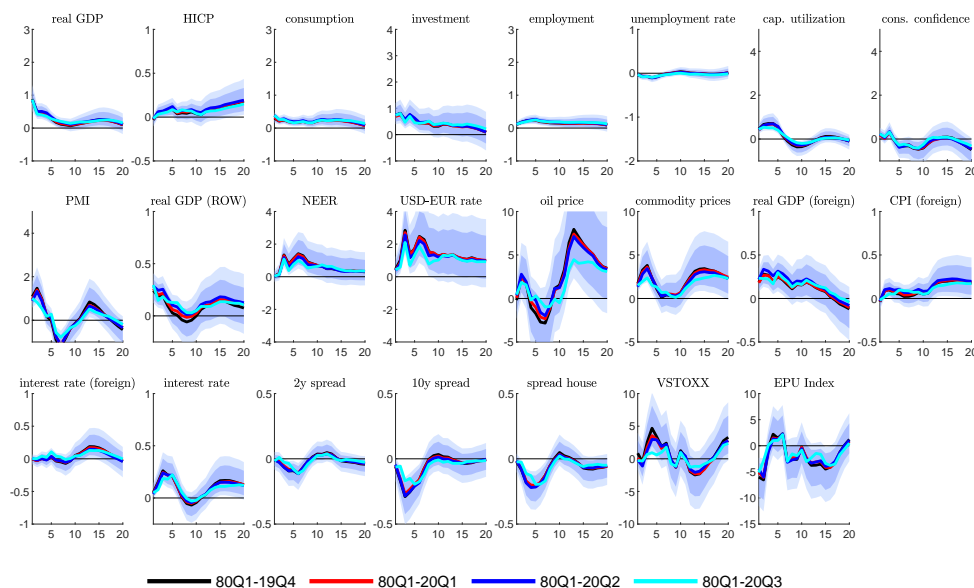
Note: Impulse response functions to a one standard deviation shock in real GDP from a BVAR with multivariate t -distributed errors with (a) weakly informative prior and (b) strong Sims and Zha prior. Thick lines are median estimates and the dark (light) blue area is the 68% (90%) credible interval for the estimation window until 19Q4.

C.4 Additional results for the large BVAR

Figure 19: Impulse response functions in a large BVAR (weakly informative prior)



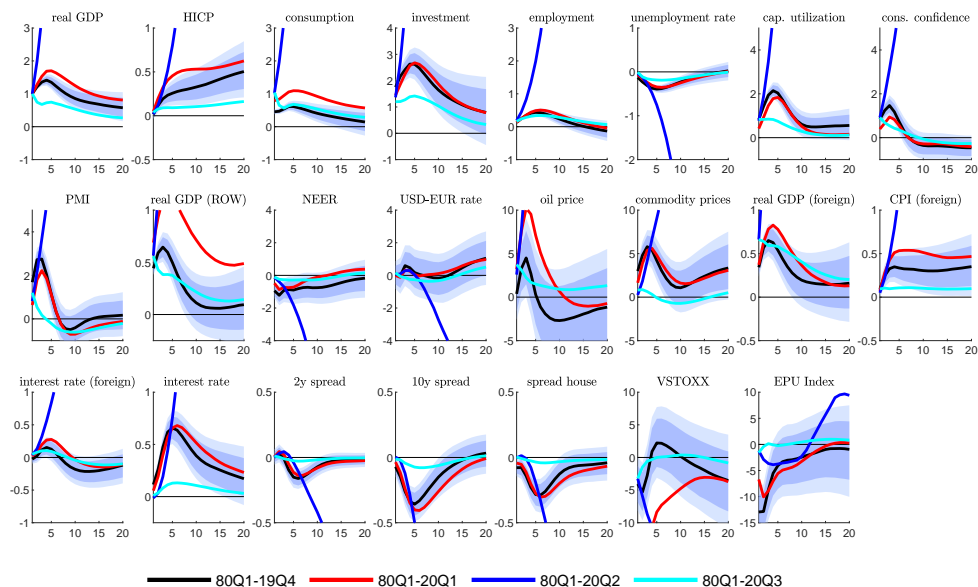
(a) BVAR with Gaussian errors



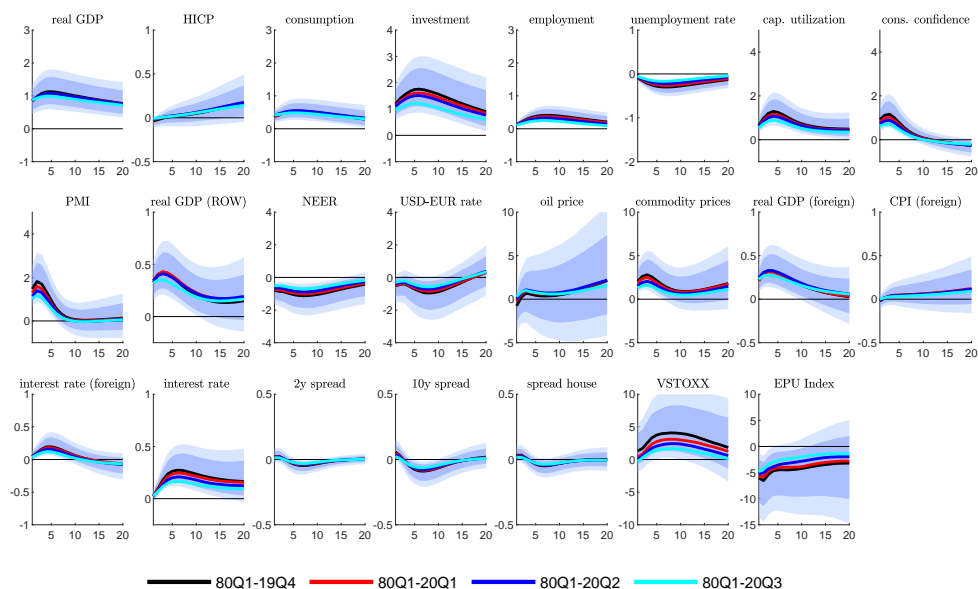
(b) BVAR with fat-tailed errors

Note: Impulse response functions to a one standard deviation shock in real GDP from a large BVAR with (a) Gaussian and (b) multivariate t -distributed errors using a weakly informative prior. Thick lines are median estimates and the dark (light) blue area is the 68% (90%) credible interval for the estimation window until 19Q4.

Figure 20: Impulse response functions in a large BVAR (Sims and Zha prior)



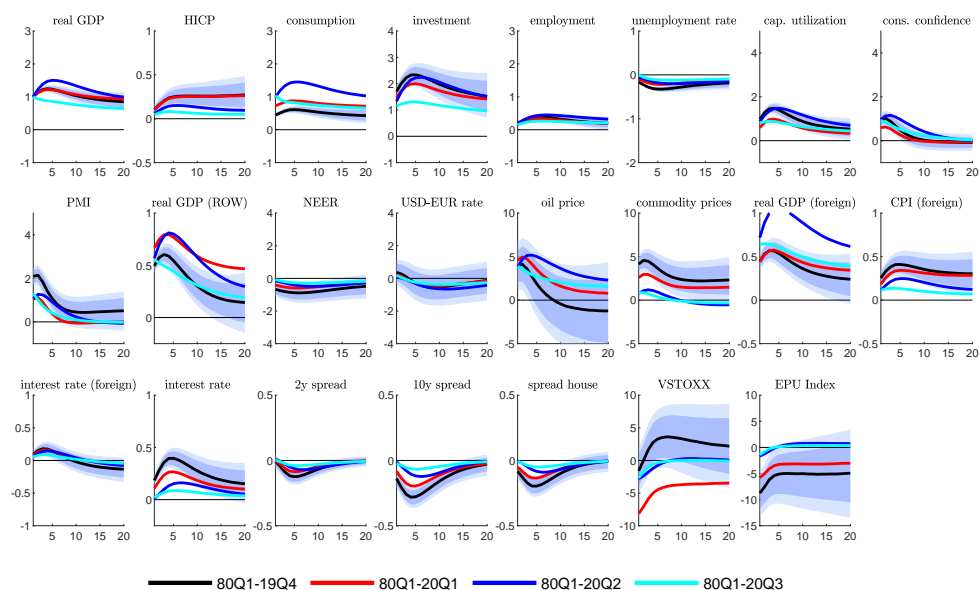
(a) BVAR with Gaussian errors



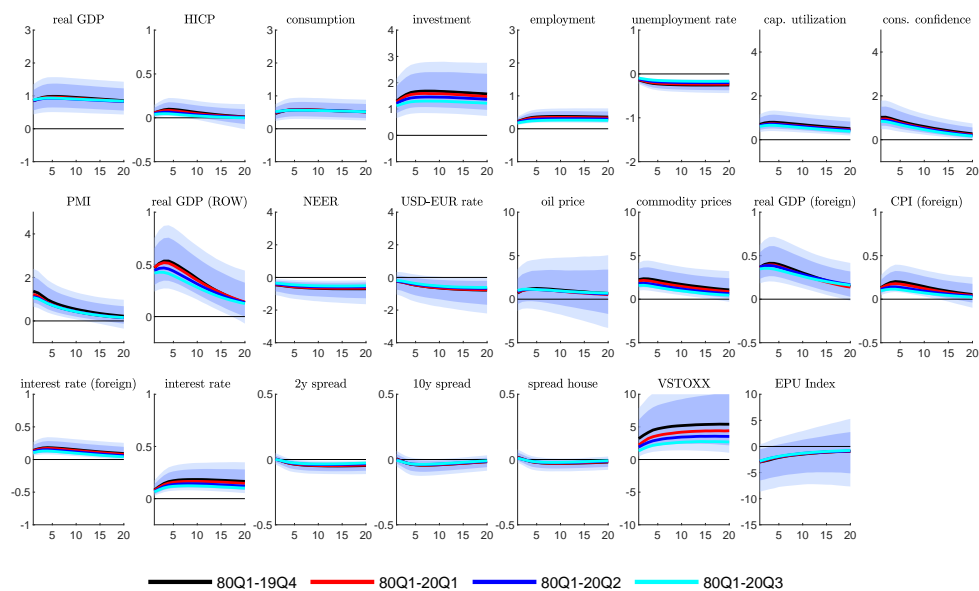
(b) BVAR with fat-tailed errors

Note: Impulse response functions to a one standard deviation shock in real GDP from a large BVAR with (a) Gaussian and (b) multivariate t -distributed errors using a Sims and Zha prior. Thick lines are median estimates and the dark (light) blue area is the 68% (90%) credible interval for the estimation window until 19Q4.

Figure 21: Impulse response functions in a large BVAR (strong Sims and Zha prior)



(a) BVAR with Gaussian errors

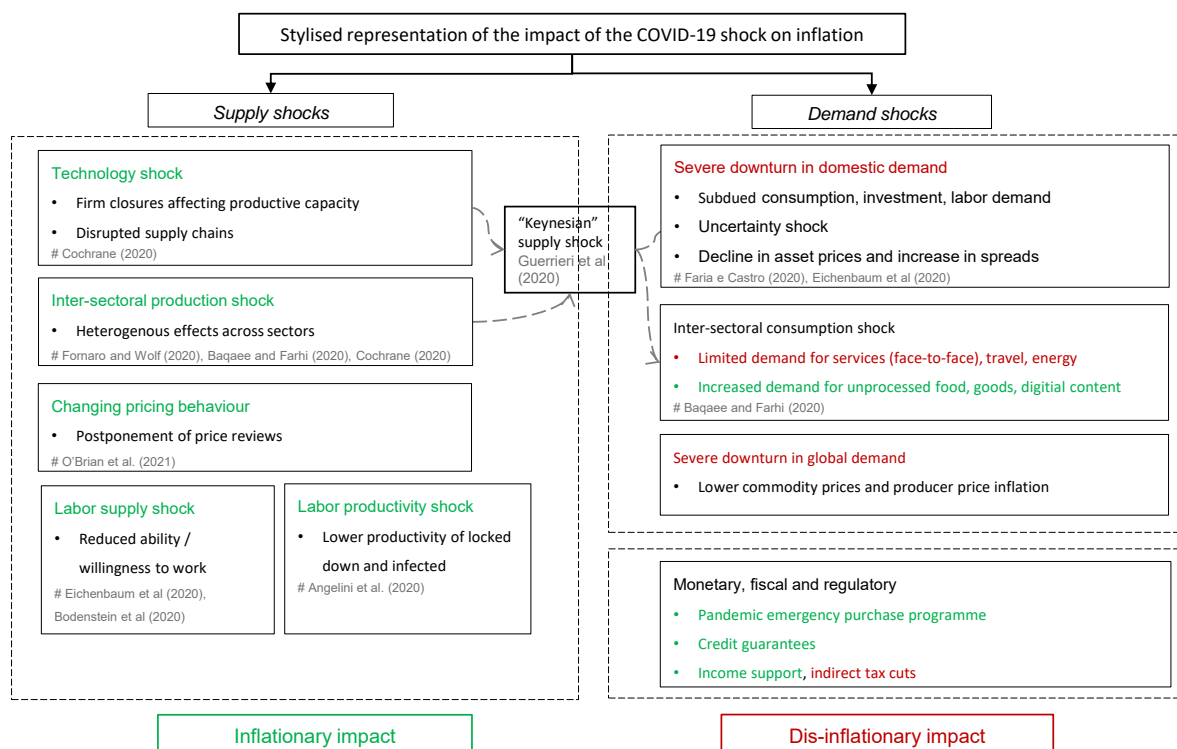


(b) BVAR with fat-tailed errors

Note: Impulse response functions to a one standard deviation shock in real GDP from a large BVAR with (a) Gaussian and (b) multivariate t -distributed errors using a strong Sims and Zha prior. Thick lines are median estimates and the dark (light) blue area is the 68% (90%) credible interval for the estimation window until 19Q4.

C.5 Stylised representation of the impact of the COVID-19 shock on inflation

Figure 22: Impact of the COVID-19 shock on inflation



Note: The color coding points to likely inflationary or disinflationary effects, but actual impacts are less straightforward.

Acknowledgements

We thank Marek Jarociński, Michele Lenza, Chiara Osbat and one anonymous referee for their valuable comments, the participants in the 9th ECB Near Term Inflation Projection Workshop and the participants in the November 2020 DG-ECFIN internal seminar for their feedback. The views expressed in this paper are of the authors only and do not necessarily reflect those of the European Central Bank or the Eurosystem. Data disclaimer: The dataset employed in this paper contains confidential statistical information. Its use for the purpose of the analysis described in the text has been approved by the relevant ECB decision making bodies. All the necessary measures have been taken during the preparation of the analysis to ensure the physical and logical protection of the information.

Elena Bobeica

European Central Bank, Frankfurt am Main, Germany; email: elena.bobeica@ecb.europa.eu

Benny Hartwig

European Central Bank, Frankfurt am Main, Germany; Goethe University Frankfurt, Frankfurt am Main, Germany;
email: benny.hartwig.external@ecb.europa.eu

© European Central Bank, 2021

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

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

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

PDF

ISBN 978-92-899-4558-5

ISSN 1725-2806

doi:10.2866/957422

QB-AR-21-049-EN-N