Capturing GDP Nowcast Uncertainty in Real Time

Paul Labonne*

Abstract

Nowcasting methods rely on timely series related to economic growth for producing and updating estimates of GDP growth before publication of official figures. But the statistical uncertainty attached to these forecasts, which is critical to their interpretation, is only improved marginally when new data on related series become available. That is particularly problematic in times of high economic uncertainty. As a solution this paper proposes to model common factors in scale and shape parameters alongside the mixed-frequency dynamic factor model typically used for location parameters in nowcasting frameworks. Scale and shape parameters control the time-varying dispersion and asymmetry round point forecasts which are necessary to capture the increase in variance and negative skewness found in times of recessions. It is shown how cross-sectional dependencies in scale and shape parameters may be modelled in mixed-frequency settings, with a particularly convenient approximation for scale parameters in Gaussian models. The benefit of this methodology is explored using vintages of U.S. economic growth data with a focus on the economic depression resulting from the coronavirus pandemic. The results show that modelling common factors in scale and shape parameters improves nowcasting performance towards the end of the nowcasting window in recessionary episodes.

Keywords: Nowcasting uncertainty, score driven models, dynamic common factors, volatility, skewness
JEL codes: C32, C53, E66

1 Introduction

This paper proposes a novel means for deriving timely measures of GDP nowcast uncertainty using forecasting errors observed in series related to economic growth. Nowcasts can thus be communicated with an improved appreciation of the statistical uncertainty attached to them. The approach relies on filtering and dynamic factor techniques for exploiting dependencies in a panel of time series released on different frequencies and asynchronously. While Gaussian dynamic factor models provide a reliable and popular method for nowcasting, non-Gaussian techniques are used increasingly to model the distribution of conditional GDP growth in a more flexible way. This paper sets out new techniques for updating nowcasting uncertainty in both Gaussian and non-Gaussian frameworks.

*King's College London and Economic Statistics Centre of Excellence, email: paul.labonne@kcl.ac.uk. I thank Martin Weale for his supervision on this project.

As the leading measure of economic growth GDP is a key variable affecting policy makers and economic agents' decisions. But its comprehensiveness comes at the cost of timeliness; it is published with a significant delay. Nowcasting methods provide a means for forecasting current-quarter GDP starting from the beginning of the quarter until its official release by using data related to economic growth but released earlier and more frequently. By modelling the relationship between these related series and GDP, the 'targeted' series, it is thus possible to update the current-quarter GDP forecast, or nowcast, each time new observations on related series become available.

For an efficient use of nowcasts the uncertainty attached to them should be conveyed as well. But while nowcasting methods make use of data related to economic growth to improve point forecasts, the signal on forecasting uncertainty these related series carry is exploited only partially. The model capturing cross-sectional dependencies in the data is typically applied only to conditional means (or locations), whereas forecasting uncertainty is derived from scale and shape parameters. While modelling dependencies in conditional means may be enough for forecasting uncertainty to be propagated across series when related series have dynamic variance parameters, this channel is incomplete and takes place with a potential delay. The sharp downturn brought by the coronavirus pandemic and the increase in forecasting uncertainty linked to it makes this issue particularly salient.

To make full use of the information that related series carry on GDP nowcasting uncertainty this paper explores the benefit of using dynamic factor models for scale and shape parameters, in addition to the one typically used for locations. Hence new observations in the related series lead to timely updates in the point forecasts through the common factor in the locations and, importantly, to the uncertainty attached to them as well through the common factors in the scales and shapes.

Intuitively, if large prediction errors occur in a series related to GDP, and these are associated with large prediction errors in GDP, then the dispersion attached to the GDP nowcast should be adjusted accordingly. This is possible directly through the common factor in the scales. On the other hand, modelling dependencies in the shape parameters is useful to capture the asymmetry in prediction errors typically observed at the onset of recessions : While there is an increase in the dispersion of prediction errors, these are likely to be skewed towards negative values. The onset of the Covid-19 pandemic is a good illustration of this fact. It was clear in April 2020 that the first-quarter GDP growth in the US would be negative or very close to zero, but while the magnitude of the drop was uncertain, very large positive figures were clearly improbable.

The motivation for modelling non-Gaussian features, and a shape parameter in particular, is based on evidence that macroeconomic data in general are not normally distributed (see for instance Haldane (2012) and Jensen et al. (2020)). The features of GDP growth's distribution when conditioning on macroeconomic and financial variables are also likely to vary over the economic cycle. By conditioning on financial variables, Adrian et al. (2019) and Delle-Monache et al. (2020) notably

find increased negative skewness accompanying a rise in volatility in the predictive distribution of US GDP growth in times of recessions. In normal times, however, GDP growth is close to being conditionally normally distributed. These temporary deviations from normality have important implications regarding the communication of a model's predictions. A forecaster should take into account the time-varying shape of the uncertainty attached to a model's forecast, with a particular attention to downside risk.

To model openly non-Gaussian features in the conditional distribution of economic data this paper adopts the score driven methodology of Harvey (2013) and Creal et al. (2013). Score driven models provide a flexibility close to non-Gaussian state space models while remaining easy to estimate and implement. They have been applied successfully to economic forecasting problems notably by Delle-Monache and Petrella (2017), Creal et al. (2014) and Gorgi et al. (2019). More recently, Delle-Monache et al. (2020) use score driven methods to study macroeconomic downside risk by estimating time-varying location, scale and shape parameters. They use a large number of regressors related to financial conditions in a univariate setting. The framework presented here can be seen as a mixed-frequency multivariate alternative to their model.

Some nowcasting models allow variance parameters to change over time, notably through stochastic equations or Markov-switching processes, but their dynamics do not exploit cross-sectional information. Models with stochastic volatility (as in Antolin-Diaz et al. (2017, 2020) for instance) can nevertheless give an improved measure of nowcasting uncertainty through the common factor in conditional means. A large prediction error in a related series increases the expected volatility of the next prediction error in this series. The dynamic in the common factor comes from these prediction errors; therefore, the larger their expected volatility, the larger the expected volatility of the common factor. Since the common factor is the main component of GDP nowcasts, their dispersion increases as well. Modelling dependencies in the scale parameters provide a complementary and direct approach for updating GDP nowcasting uncertainty.

Separately, models with volatility (or scale) common factors are not new but so far have not been used in a mixed-frequency setting (see Carriero et al. (2016, 2018, 2020) and Huber (2016)). Gorgi et al. (2019) model a common scale component in a score driven bivariate model, but they do so by constraining the scale to be identical in both series and do not discuss the issue of temporal aggregation. The model proposed here contributes to this literature on stochastic volatility by addressing openly the temporal aggregation problem arising from modelling common volatility components in mixed-frequency and possibly non-Gaussian frameworks. In doing so it provides an approach specifically fitted for nowcasting which may be applied in a wide range of dynamic factor models.

Modelling common components in series aggregated and sampled at different frequencies raises temporal aggregation issues somewhat different for location, scale and shape parameters. While Mariano and Murasawa (2003) popularised a precise

approximation for conditional means which is applicable to location parameters, similar solutions for scale and shape parameters have not been discussed in the literature.

This paper explores two different strategies to tackle the temporal aggregation problem of scale and shape parameters. First, it is shown how scale parameters may be aggregated temporally in Gaussian models using a convenient approximation, thus providing an approach for modelling volatility common factors in a wide range of mixed-frequency dynamic factor models. The second approach consists of aggregating monthly series into rolling-quarterly figures. The mismatch in the frequency of aggregation is thus alleviated, and scale and shape common factors can be modelled in a non-Gaussian setting.

The mixed-frequency dynamic factor model presented in this paper exploits four US time series : GDP, industrial production, the index of working hours, and the Weekly Economic Index. The first three series are derived from real-time vintages provided by the Federal Reserve Bank of Philadelphia. The Weekly Economic Index (Lewis et al. (2020)) is a recent indicator released by the New York Fed in the aim of capturing and communicating the effect of the Covid-19 pandemic on the US economic activity in a timely way. GDP is quarterly while the other series are monthly.

Different specifications of the model are applied in a real-time setting to study the effect of modelling common factors in scale and shape parameters on nowcasting performance, with a particular interest in density forecasts and recessionary episodes. The results show that models with common factors in their scale and shape parameters perform better than models with only location common factors towards the end of the nowcasting window, while providing competitive performances in other times. Scale and shape common factors prove to be particularly important to capture the forecasting uncertainty generated by the coronavirus pandemic. Separately, modelling fat tails in series related to economic growth complicates the identification of turning points in activity.

2 A Mixed-Measurement Score Driven Approach for Nowcasting with Location, Scale and Shape Common Factors

A dynamic factor approach is used for exploiting dependencies across series in location, scale and shape parameters. A popular method for estimating dynamic factor models consists of writing the model in state space form and using the Kalman filter to evaluate the model's log likelihood. The Kalman filter provides an efficient approach for estimating nowcasting models for two reasons. First, it handles easily the missing values that arise from modelling jointly series aggregated at different frequencies and released asynchronously. Secondly, predictions can be decomposed into latent components, which notably can be used to capture secular changes in addition to common factors (see Harvey (1989) and Durbin and Koopman (2012)). A comprehensive presentation of dynamic factor and state space methods

for nowcasting is given by Banbura et al. (2013).

A limitation of the Kalman filter, however, is that it relies on the data being conditionally normally distributed. Non-Gaussian features can be introduced using importance sampling methods but these can be computationally intensive. Alternatively, Creal et al. (2013) and Harvey (2013) derive a new class of filters relying on the conditional score where the conditional distribution of the observations can arise from a wide range of families. An attractive feature of score driven models is that they provide a general framework for introducing time-variation and latent states in any parameter of the predictive distribution. But unlike state space models, score driven models are fully deterministic conditional on past information. Apart from this shortcoming they retain most of the flexibility of state space models, including the ability to model unobserved components. This section extends the score driven dynamic factor model of Creal et al. (2014) with dynamic factor structures for scale and shape parameters. Hence new observations on related series lead to timely updates in both the point forecasts and the uncertainty attached to them.

2.1 A Mixed-Measurement Approach

As in Creal et al. (2014) and Gorgi et al. (2019) each element of the observation vector $y_t = (y_{1,t}, ..., y_{m,t})'$, $t = 1, ..., N$ can have a distinct conditional (or predictive) density such that

$$y_{i,t} \sim p_i(y_{i,t}|Y_{t-1}), \quad i = 1, ..., m, \tag{1}$$

where m is the number of time series. The set $Y_{t-1} = \{f_t, X_{t-1}, \Theta\}$ includes the information available at time $t-1$. All variables are modelled at a monthly frequency. The GDP series, which is quarterly, shows a quarterly figure in the last month of each quarter with missing values in other months. The vector f_t includes the dynamic states related to location, scale and shape parameters which are fully deterministic conditional on past information. Θ is a set of fixed parameters such as autoregressive coefficients and factor loadings. Importantly the observations are assumed to be cross-sectionally independent conditional on past information. Hence the log likelihood of the observations at time t can be written as

$$\log p(y_t|Y_{t-1}) = \sum_{i=1}^{m} \delta_{i,t} \log p_i(y_{i,t}|Y_{t-1}), \tag{2}$$

$t = 1, ..., N$, where $\delta_{i,t}$ is zero if observation $y_{i,t}$ is missing and one otherwise.

2.2 A Score Driven Dynamic for the Unobserved Components

Following Creal et al. (2013) and Harvey (2013) the dynamic in the vector of time-varying parameters f_t comes from the conditional score such that

$$f_{t+1} = Bf_t + As_t, \tag{3}$$

where s_t denotes the scaled first derivative of the log density with respect to the vector f_t:

$$s_t = S_t \Delta_t, \quad \Delta_t = \frac{\partial \log p(y_t|Y_{t-1})}{\partial f_t}. \tag{4}$$

The matrix A is a diagonal matrix of gains which are estimated via maximum likelihood along the unknown elements of the matrix B and the initial vector f_1. The matrix S_t is a scaling matrix set to the Moore–Penrose inverse of the expected information matrix[1] given by

$$S_t = \Big[\text{diag}(\mathcal{I}_t)\Big]^{-1}, \quad \mathcal{I}_t = \text{E}\big[\Delta_t \Delta_t' | Y_{t-1}\big]. \tag{5}$$

Following Delle-Monache et al. (2020) and Lucas and Zhang (2016) only the diagonal elements of the information matrix are used to improve estimation.

Since the observations are conditionally cross-sectionally independent the score and the expected information matrix can be expressed conveniently as

$$\Delta_t = \sum_{i=1}^{m} \delta_{i,t} \Delta_{i,t}, \tag{6}$$

where $\Delta_{i,t} = \partial \log p(y_{i,t}|Y_{t-1})/\partial f_t$, while the information matrix becomes

$$\text{E}\big[\Delta_t \Delta_t' | Y_{t-1}\big] = \sum_{i=1}^{m} \delta_{i,t} \text{E}\big[\Delta_{i,t} \Delta_{i,t}' | Y_{t-1}\big]. \tag{7}$$

In score driven models time-varying parameters are fully deterministic conditional on past information. Therefore, if all series are modelled contemporaneously the GDP nowcast cannot be updated once related series become observable for the last month of the quarter. To overcome this problem the related series are lead of one period, as in Gorgi et al. (2019).

The next subsection details the dynamics of the components in f_t.

2.3 Dynamic Factor Models for Location, Scale and Shape Parameters

Each location, scale and shape parameter is decomposed into a stochastic trend representing idiosyncratic secular changes and a common component capturing cyclical variations common to all series which take the following form

$$\lambda_{i,t+1}^j = \lambda_{i,t}^j + A_{\lambda i}^j s_{\lambda i,t}^j, \quad i = 1, ..., m, \tag{8}$$

$$\Phi_j(L) \pi_{t+1}^j = A_\pi^j s_{\pi,t}^j, \tag{9}$$

for all periods $t = 1, ..., T$. Location parameters are indicated by $j = \mu$, scale parameters by $j = \sigma$ and shape parameters by $j = \alpha$.

[1] The information matrix needs to be inverted at each step of the recursion because its components are dynamic which considerably slows down estimation. To overcome this technical issue the score driven recursion (hence the log likelihood evaluation as well) is written in C++ (I am grateful to Caterina Schiavoni for her help with this) while the optimisation and the remainder of the analysis is carried in R.

The common component in the location, π_t^{μ}, is an AR(2) process as is commonly specified for the business cycle; it is constrained to be stationary using the parameter transformation of Osborn (1976). The common factor in the scales captures common volatility shocks, like the one generated by the Covid-19 pandemic, and is specified as a stationary AR(1). The common factor in the shapes captures simultaneous shifts in the asymmetry of prediction errors, like those typically arising at the beginning of recessions, and is also modelled as a stationary AR(1). Adding a second lag in the common factors of scales and shapes has little effect on estimation.

The stochastic trend in each parameter, $\lambda_{i,t}^{j}$, is useful to capture long-term deviations. Antolin-Diaz et al. (2017) notably demonstrate that random-walk specifications for time-varying parameters are robust to discrete breaks. They also stress the importance of capturing secular changes in economic growth, a finding also discussed by Doz et al. (2020).

The next subsection discusses how trends and common components are related to the location, scale and shape parameters.

2.4 A Mixed-Frequency Approach for Location, Scale and Shape Parameters

Each series follows the predictive model :

$$y_{i,t} = \mu_{i,t} + \sigma_{i,t}\epsilon_{i,t}, \tag{10}$$

for $i = 1, ..., m$, $t = 1, ..., N$, and where $v_{i,t} = \sigma_{i,t}\epsilon_{i,t}$ is a prediction error (and $\epsilon_{i,t}$ its standardised counterpart) following an arbitrary distribution with time-varying location, scale and shape parameters ($\mu_{i,t}$, $\sigma_{i,t}$ and $\alpha_{i,t}$). These are decomposed into latent states whose dynamics come from the scaled score as outlined in the previous section. These latent states include a component common to all variables. However, the data used for estimation are of different nature; while GDP is a quarterly variable, the related series are monthly variables. It is important to account explicitly for this mismatch in the frequency of aggregation when relating the latent states to the parameters.

To address this temporal aggregation mismatch it is useful to start from the accounting relationship between monthly and quarterly variables, specifically a quarterly variable in levels $Y_{i,t}$ must be equal to the three-month sum of its monthly sub-components $\tilde{Y}_{i,t-i}$, $i = 0, 1, 2$:

$$Y_{i,t} = \tilde{Y}_{i,t} + \tilde{Y}_{i,t-1} + \tilde{Y}_{i,t-2}. \tag{11}$$

To account for heteroskedasticity and multiplicative components, however, the data are generally taken in logarithms, such that the variables modelled are $Y_{i,t}^{*} = \log Y_{i,t}$, and $\tilde{Y}_{i,t}^{*} = \log \tilde{Y}_{i,t}$. Since the sum of the logarithms is not equal to the logarithm of the sum, the accounting constraint takes a nonlinear form :

$$Y_{i,t}^{*} = \log\Big[\exp\tilde{Y}_{i,t}^{*} + \exp\tilde{Y}_{i,t-1}^{*} + \exp\tilde{Y}_{i,t-2}^{*}\Big]. \tag{12}$$

This type of non-linearity becomes difficult to handle once the first differences in logs are modelled. It is possible, however, to work from a linear approximation. Salazar et al. (1997) and Mitchell et al. (2005) show that

$$\sum_{i=0}^{2} \mathrm{h}(z_{t-i}) \approx 3 \ \mathrm{h}\left(\frac{\sum_{i=0}^{2} z_{t-i}}{3}\right), \tag{13}$$

where $h(.)$ is a non-linear transformation and z_t a smooth variable. The approximation (13) is a second order approximation because the first order errors sum to zero. When monthly values in a quarter are close to the monthly average over the quarter, which is usually the case with seasonally adjusted figures, the approximation error introduced is negligible. Using this approximation equation (12) can be written as

$$Y^*_{i,t} = \log 3 + \frac{1}{3}\tilde{Y}^*_{i,t} + \frac{1}{3}\tilde{Y}^*_{i,t-1} + \frac{1}{3}\tilde{Y}^*_{i,t-2}. \tag{14}$$

Using approximation (14) it is possible to write the quarter-on-quarter log difference $y_{i,t} = Y^*_{i,t} - Y^*_{i,t-3}$ as a function of the monthly log differences $\tilde{y}_{i,t-i} = \tilde{Y}^*_{i,t-i} - \tilde{Y}^*_{i,t-h-1}$, $h = 0, ..., 4$ as

$$y_{t,i} = \frac{1}{3}\tilde{y}_{i,t} + \frac{2}{3}\tilde{y}_{i,t-1} + \tilde{y}_{i,t-2} + \frac{2}{3}\tilde{y}_{i,t-3} + \frac{1}{3}\tilde{y}_{i,t-4}. \tag{15}$$

Equation (15) is also discussed by Mariano and Murasawa (2003) who have popularised its use in mixed-frequency dynamic factor models.

Location Parameters

There is a linear relationship between each variable and its location parameter which makes it possible to split the location of a quarterly variable into monthly components and relate them to the location with (15). Importantly this step does not require any assumption about the distribution of the unobserved monthly variable.

Monthly locations are modelled as

$$\tilde{\mu}_{i,t} = \lambda^{\mu}_{i,t} + \Lambda^{\mu}_{i}(L)\pi^{\mu}_{t}, \tag{16}$$

where the factor loading has a lag structure as in Doz et al. (2020), albeit limited to one period only, specifically $\Lambda^{\mu}(L) = \Lambda^{\mu}_{i,0} + \Lambda^{\mu}_{i,1}L$. The lagged factor loading is used only for series largely derived from employment data, that is the index of working hours and the Weekly Economic Index. The factor loading is constrained to be one for GDP.

Location parameters in quarterly series are related to the monthly model using approximation (15) as

$$\mu_{i,t} = \frac{1}{3}\tilde{\mu}_{i,t} + \frac{2}{3}\tilde{\mu}_{i,t-1} + \tilde{\mu}_{i,t-2} + \frac{2}{3}\tilde{\mu}_{i,t-3} + \frac{1}{3}\tilde{\mu}_{i,t-4}, \tag{17}$$

while for monthly series it is simply

$$\mu_{i,t} = \tilde{\mu}_{i,t}. \tag{18}$$

Scale and Shape Parameters

Diverging from the normal distribution is useful to capture the features of the data in a flexible way, but it also creates challenges in a mixed-frequency framework. Notably, while location parameters can be disaggregated temporally without making any statistical assumptions on the monthly sub-components, that is not possible for scale and shape parameters. This is problematic because the sum of two non-normally distributed variables with known distributions has a distribution which is difficult to evaluate and in many cases unknown. Linking the monthly model to the quarterly model is therefore difficult.

When working in a non-Gaussian framework it is preferable to model the dependencies across series in scale and shape parameters using quarterly models. To do so related series exhibiting common factors in their scale and shape parameters are aggregated into rolling quarterly observations (quarter-on-quarter deviations observed monthly). This introduces serial correlation in these series which is addressed by modelling monthly location components, as outlined above. The temporal aggregation strategies for modelling jointly quarterly and monthly observations and modelling rolling quarterly observations (quarterly observations observed monthly) are essentially the same; in both cases it is necessary to go back to the underlying monthly model. Labonne and Weale (2020) show that when rolling quarterly series are not subject to important measurement errors it is possible to interpolate the monthly path very precisely with (14). Hence there should be little loss of information when using rolling quarterly observations instead of monthly observations.

Separately, while locations can take any real value, scale parameters are positive and shape parameters can take values only in (0;1) in the distribution used below. To incorporate these constraints during estimation, trends and common factors are related to scale and shape parameters as

$$\sigma_{i,t} = \exp(\lambda^{\sigma}_{i,t} + \Lambda^{\sigma}_{i}\phi^{\sigma}_{t}), \qquad \alpha_{i,t} = 1/(1 + \exp(\lambda^{\alpha}_{i,t} + \Lambda^{\alpha}_{i}\phi^{\alpha}_{t})), \tag{19}$$

with the factor loadings constrained to one for the GDP series ($i = 1$).

The Special Case of a Normal Distribution

In a Gaussian setting it is possible to link monthly variances to quarterly variances using an approximation close to (15) as shown below.

If $y_{i,t}$ is conditionally normally distributed, its conditional variance is equal to the variance of the prediction error such that

$$Var(y_{i,t}|Y_{t-1}) = E[(y_{i,t} - E(y_{i,t}|Y_{t-1}))^2] = E(v^2_{i,t}) = Var(v_{i,t}).$$

Assuming that series i is observable quarterly, the quarterly prediction error can be split into its monthly components using approximation (15) as

$$v_{i,t} = \frac{1}{3}\tilde{v}_{i,t} + \frac{2}{3}\tilde{v}_{i,t-1} + \tilde{v}_{i,t-2} + \frac{2}{3}\tilde{v}_{i,t-3} + \frac{1}{3}\tilde{v}_{i,t-4},$$

where $\tilde{v}_{i,t}$ is a monthly prediction error. Using this approximation the variance of the prediction error can be decomposed into monthly variances as

$$Var(v_{i,t}) = \frac{1}{9}Var(\tilde{v}_{i,t}) + \frac{4}{9}Var(\tilde{v}_{i,t-1}) + Var(\tilde{v}_{i,t-2}) + \frac{4}{9}Var(\tilde{v}_{i,t-3}) + \frac{1}{9}Var(\tilde{v}_{i,t-4}), \tag{20}$$

since, assuming that the monthly errors $\tilde{v}_{i,t}$ are i.i.d, the covariance terms are zero. This step is possible only if the prediction errors are normally distributed. The rolling quarterly scale parameter $\sigma_{i,t}$ can thus be written as a function of the monthly scale parameter $\tilde{\sigma}_{i,t}$ as

$$\sigma_{i,t} = \sqrt{\frac{1}{9}\tilde{\sigma}^2_{i,t} + \frac{4}{9}\tilde{\sigma}^2_{i,t-1} + \tilde{\sigma}^2_{i,t-2} + \frac{4}{9}\tilde{\sigma}^2_{i,t-3} + \frac{1}{9}\tilde{\sigma}^2_{i,t-4}}. \tag{21}$$

It is thus possible to specify a monthly model for all scale parameters where the monthly components are related to the quarterly scale parameter of GDP using (21) with $\tilde{\sigma}_{1,t} = \exp(\lambda^{\sigma}_{1,t} + \Lambda^{\sigma}_{1}\phi^{\sigma}_{t})$.

2.5 A Generalised Asymmetric Student-t Model

Each series' conditional density comes from the family of asymmetric student-t (AST) density of Zhu and Galbraith (2010). The distinctive features of the AST are its shape parameter, or skewness parameter, which controls the asymmetry round the central part of the distribution, and its tail parameters which control tail behaviours independently on each side.

The log AST density of observation t of series i takes the form

$$\begin{aligned} \log p_i(y_{i,t}|Y_{t-1}) = &- \ln\sigma_{i,t} - \frac{\nu_{1,i}+1}{2}\ln\Big[1 + \frac{1}{\nu_{1,i}}\Big(\frac{y_{i,t}-\mu_{i,t}}{2\alpha_{i,t}\sigma_{i,t}K(\nu_{1,i})}\Big)^2\Big]\ 1(y_{i,t} \le \mu_{i,t}) \\ &- \frac{\nu_{2,i}+1}{2}\ln\Big[1 + \frac{1}{\nu_{2,i}}\Big(\frac{y_{i,t}-\mu_{i,t}}{2(1-\alpha_{i,t})\sigma_{i,t}K(\nu_{2,i})}\Big)^2\Big]\ 1(y_{i,t} > \mu_{i,t}), \end{aligned} \tag{22}$$

where $\sigma_{i,t}$ is the scale parameter, $\alpha_{i,t}$ is the shape parameter which can take values in $(0,1)$, and $\nu_{1,i}$ and $\nu_{2,i}$ are respectively the left and right tail parameters which take positive values. $K(\nu) = \Gamma((\nu+1)/2)/(\sqrt{\nu\pi}\Gamma(\nu/2))$ ($\Gamma(.)$ is the Gamma function) and $1(x)$ is an indicator variable equal to one if statement x is true and zero otherwise. The distribution is skewed towards positive values if $\alpha_{i,t} < 0.5$ and towards negative values if $\alpha_{i,t} > 0.5$. When the tail parameters are constrained to be very large and the skewness parameter to 0.5 the AST is equivalent to a (scaled) normal distribution.

Figure 1 illustrates the effects that scale, shape and tail parameters have on the density function. The solid red line shows the AST density when $\sigma = 1$ while other parameters are constrained to be Gaussian ($\alpha = 0.5$, $\nu_1 = \nu_2 = \infty$). The dotted blue line shows the AST density when either the scale, tail or shape parameters are varying. The density at the location changes when the scale parameter is

altered but is independent to variations in the tail and shape parameters. This is a particular feature of this version of the AST distribution (given in equation (5) of Zhu and Galbraith (2010)) where the random variable is scaled with $B_{i,t} = \alpha_{i,t}K(\nu_{i,1}) + (1 - \alpha_{i,t})K(\nu_{i,2})$.

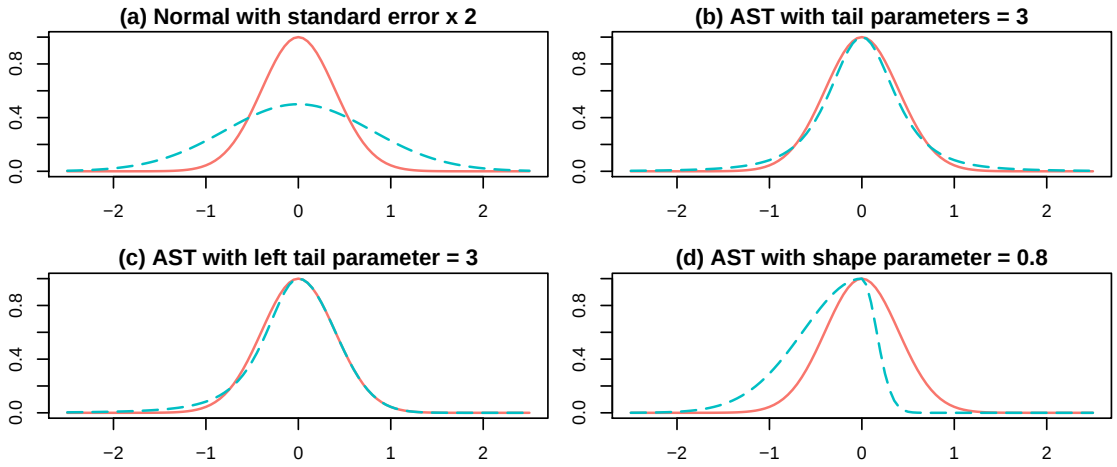

Figure 1: Illustration of the density functions from AST distributions with different set of parameters. With shape parameter = 0.5 and tail parameters = ∞ the AST distribution reduces to a (scaled) normal distribution (red solid line).

Graph (a) shows the effect of increasing the scale parameter, specifically to $\sigma = 2$. The dispersion increases symmetrically on both sides and can be interpreted as a general (or symmetric) increase in forecasting uncertainty when the model is applied to conditional GDP growth.

Graph (b) shows the effect of lowering both tail parameters to three. The probability in the tails increases which can be used to account for extreme events. Graph (c), on the other hand, illustrates asymmetric tails; specifically an heavy tail on the left side and Gaussian tail on the right side ($\nu_1 = 3$ while $\nu_2 = \infty$). This increases the probability of extreme negative events while leaving the probability of large positive extreme events unaffected. This asymmetric tail behaviour can be particularly useful for conditional GDP if, as Jensen et al. (2020) argue, the business cycle has become more asymmetric in the last thirty years. They find that contractions have become more violent while recoveries remain smooth.

Finally, graph (d) shows the effect of negative skewness ($\alpha = 0.8$) on the density function. The central part of the density becomes skewed heavily toward negative values. This statistical behaviour should be especially useful to capture the onset of recessions when coupled with an increase in dispersion : while general economic uncertainty increases, the likelihood of a positive outcome decreases. The real-time analysis of US GDP in the first quarter of 2020 shown in section 7 validates this intuition.

Figure 2 illustrates the response function of the scaled score for location, scale and shape parameters with the prediction error as input.[2] These are useful for

[2]The components of the score and information matrix are shown in appendix A

illustrating two important properties of the scale score. First, introducing a skewness parameter yields a discontinuity round zero. For instance, a distribution skewed towards positive values downweights the effect of positive prediction errors and amplify the effect of positive ones. Secondly and most importantly, the scale score downweights the effect of large prediction errors when tail parameters are small (low rates of decays in the tails). The Gaussian and skewed models, on the other hand, yield linear or increasing responses when prediction errors increase.

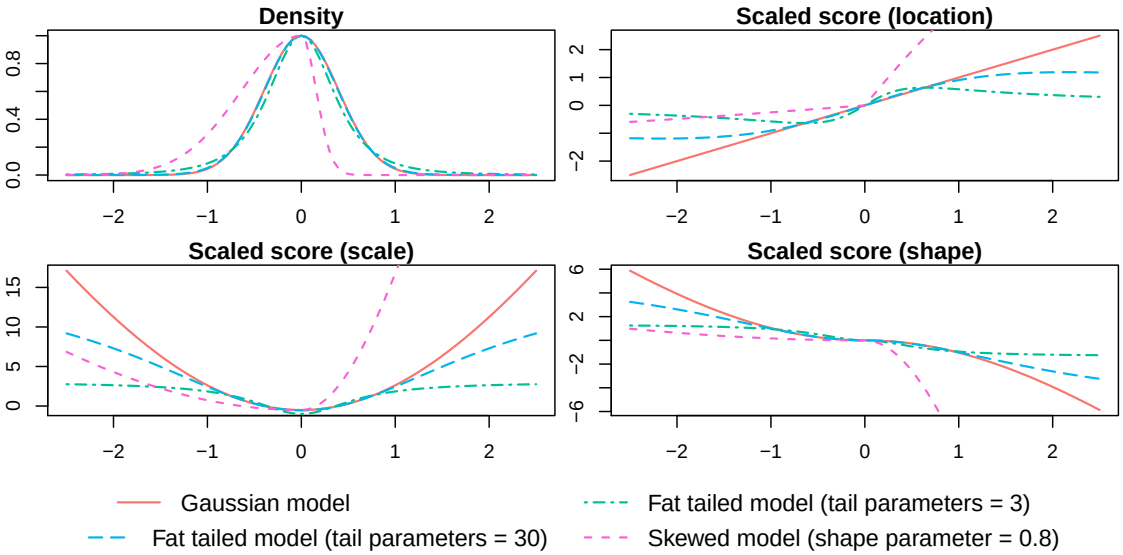

Figure 2: Plots of the density and score functions for an asymmetric student-t (AST) distribution with $\sigma = 1$. X-axes show prediction errors. With shape parameter $\alpha = 0.5$ and tail parameters $\nu_1 = \nu_2 = \infty$ the AST is equivalent to a (scaled) normal distribution.

On the importance of high tail parameters for identifying turning points

Downweighting the effect of large prediction errors is usually a desirable feature of score driven models because it leads to a more robust estimation. However, economic crises, and the Covid-19 period in particular, are examples of cases where outliers most likely have important and long-lasting effects on means and variances of time series. If the related series' distributions are allowed to have fat tails, the effect of large swings in the economic activity on the time-varying parameters are likely to be downweighted. Consequently the model might not capture turning points in the economic activity and the concurrent increase in the dispersion of prediction errors in a timely way.

To investigate thoroughly the effects of fat tails on estimation and forecasting performance in recessionary episodes the empirical analysis compares models with unconstrained tail parameters with models constrained to have Gaussian tail parameters in the related series.

3 A Weighted Maximum Likelihood Estimator

The estimation strategy relies on the weighted maximum likelihood method of Blasques et al. (2016), notably applied in a score driven framework by Gorgi et al. (2019). It accounts explicitly for the fact that, while related series are used for estimation, the primary objective here is to forecast GDP growth, not the related series; these are used solely to improve GDP nowcasts.

To see this more clearly it is helpful to look at the model's log likelihood at period t, given by (2), which is the sum of each observation's log likelihood at time t. In a misspecified setting the parameters that maximise the total log likelihood are not necessarily those that maximise the log likelihood of GDP. This issue is even more prominent in a mixed-frequency framework where GDP is observed once every three months whereas the related series are observed monthly. Indeed when an observation is missing the series it relates to has no impact on the model's log likelihood for this period. The log likelihood associated to GDP has consequently less weight on the total log likelihood than those of the related series.

Using the weighted maximum likelihood approach the vector of unknown parameters Θ which includes the distributions' parameters, the initial values of the time-varying parameters, the autoregressive coefficients and the gains is estimated as

$$\hat{\Theta} = \arg\max_{\Theta} \sum_{t=1}^{N} \log p^{w}(y_t|Y_{t-1}), \tag{23}$$

where $\log p^{w}(y_t|Y_{t-1})$ is the weighted log likelihood defined as

$$\log p^{w}(y_t|Y_{t-1}) = \delta_{i,t} \log p_1(y_{1,t}|Y_{t-1}) + W \sum_{i=2}^{m} \delta_{i,t} \log p_i(y_{i,t}|Y_{t-1}). \tag{24}$$

The weight W is applied to the related series' log likelihood and diminish their contribution to the total log likelihood. It cannot be estimated alongside the other parameters and Blasques et al. (2016) suggest selecting it via cross-validation techniques. Alternatively, Gorgi et al. (2019) set the weight to zero. This complicates the identification of the scale parameters in the related series; a problem they overcome by modelling a unique scale parameter for all series. This would be problematic here since the modelling approach proposed relies on a flexible specification for scale parameters. Furthermore, in their out-of-sample nowcasting exercise Gorgi et al. (2019) do not find that the score driven approach yields a clear benefit when the contribution of the indicator series is null.

Here the weighting factor is set to one third to compensate for the implicit downweighting of the GDP series coming from its lower frequency of observation compared to the related series. The indicator series' log likelihood contributions are not downweighted further because that might deteriorate excessively the model's capacity to predict related series in real time. This could yield to prediction errors loosing their economic meaning and with it their ability to indicate period of economic depressions and uncertainty, which is the focus of this paper.

4 Real-Time Data on US Economic Growth

This study is centred on quarterly GDP which is the leading measure of economic growth and the data are taken from the United States. The estimate of GDP analysed is the *Advance Estimate*, the timeliness estimate of GDP in the US.

Three monthly series related to economic growth are used to improve GDP nowcasts: Industrial production, the index of working hours and the Weekly Economic Index (Lewis et al. (2020)). Industrial production is a component of GDP but is more frequent and timely. The index of working hours is a very timely indicator of employment and notably reflected the effect of the Covid-19 restrictions from the onset of the pandemic. The Weekly Economic Indicator has been created specifically to capture the effect of the pandemic in US in a timely way.

GDP and industrial production are regularly subject to benchmark changes which affect the entire series, such as changes in indexation. To deal with these changes all series are taken in first differences in logs. The sample includes data from January 1973 up to June 2020 for GDP, industrial production and the index of weekly hours whereas the Weekly Economic Index is used from June 2010 up to June 2020. The series are illustrated in figure 3 using the vintage available at the end of July 2020.

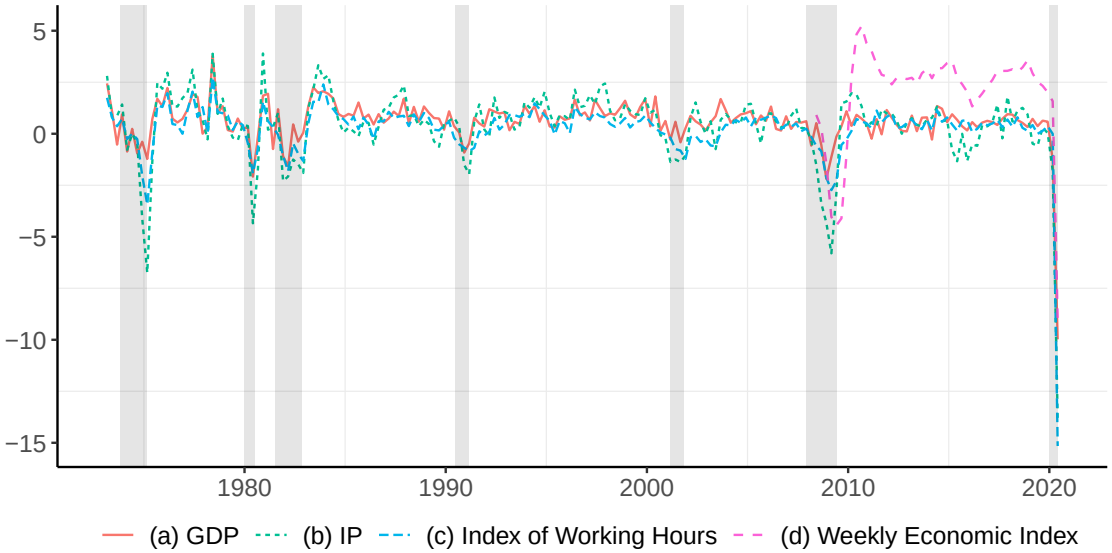

Figure 3: Quarterly data (calendar quarters) from January 1973 up to June 2020. Industrial production, the index of working hours and the Weekly Economic Index are aggregated into quarterly figures for comparison with GDP. Data from the most recent vintage. Shaded areas indicate the US recessions classified by the NBER.

GDP, industrial production and the index of working hours are regularly revised over time because the data used for their compilation accrue gradually. Therefore, to investigate the forecasting performance of the model in real time it is necessary to estimate the model recursively using successive vintages of the data. Such vintages are produced by the Federal Reserve Bank of Philadelphia. The Weekly Economic

Indicator, on the other hand, is revised only marginally.

The most timely series are the index of working hours and the Weekly Economic Index which are released at the beginning of each month and relate to the previous month. Industrial production is released in the middle of each month and relate to the previous month while GDP is released at the end of the month following the quarter it relates to. Hence, two series (the index of working hours, the Weekly Economic Index) are approximately released simultaneously while industrial production and GDP follow different publication schedules. This yields seven rounds of revisions in between GDP releases, that is seven nowcasting steps.

The first nowcasting step is at the end of the first month in the quarter, when the previous-quarter GDP figure is released. The second step is at the beginning of second month in the quarter when monthly figures for the index of working hours and the weekly economic index are released. The next five steps are each approximately two-weeks apart, when figures for either industrial production or the index of working hours and the Weekly Economic Index are released.

5 Model Specifications

The primary objective of this paper is to investigate the potential gains of modelling common factors in scale and shape parameters on dynamic factor models' nowcasting performance. For this it is useful to compare a range of model specifications.

The most flexible specification is model (a) which is an asymmetric student-t dynamic factor model (DFM) with common factors in location, scale and shape parameters. While GDP follows an unrestricted AST distribution, IP and IWH are constrained to have a unique tail parameter and effectively follow skewed Student-t distributions. Model (b) is derived from model (a) but tail parameters in the related series are constrained to be Gaussian; the related series consequently have skewed normal distributions. This means that more attention is placed on sharp movements in the related series.

Separately, model (c) is Gaussian DFM with stochastic volatilities (SV) and a common factor in scale parameters. Since model (c) is Gaussian, the common scale component can be modelled at a monthly frequency and related to the quarterly scale in GDP using (21). Comparing model (c) to the non-Gaussian models (a) and (b) is useful to understand if the gains from modelling non-Gaussian features outweigh the loss in precision resulting from the necessary aggregation of the related series featuring scale and shape common factors.

These three models, which are designed to make full use of the cross-sectional information in the data, are compared to benchmark models which do not feature scale and shape common factors. Model (d) is an asymmetric Student-t model comparable to model (b) but without scale and shape common factors. This specification is closer to the non-Gaussian models of Delle-Monache et al. (2020) (their model is a univariate model with many regressors to improve estimation of time-varying parameters whereas model (d) is a dynamic factor model).

Specification (e) is a Gaussian DFM with stochastic volatilities which is a proxy for current dynamic factor models. It is derived from model (c) but does not feature a common scale component. The relative performance of model (c) compared to model (e) provides an indication of the gains stemming from augmenting dynamic factor models with a scale common factor.

Finally, model (f) extends model (e) to Student-t distributions. Comparing model (e) and (f) is useful for understanding the effect that low tail parameters may have on the ability of the model to capture turning points.

All specifications except (a) and (b) . The Weekly Economic Index, on the other hand, has a normal distribution with constant scale parameter because of the short sample of available observations. Consequently WEI remains monthly which is a common feature in all models. A table summarising the different model specifications is shown in appendix B.

6 In-Sample Analysis

First the models are analysed using the vintage available at the end of July 2020, the latest available vintage in this study.

Modelling common scale and shape components significantly improves in-sample fit

The first two columns of Table 1 show the log likelihood of each model and the log likelihood of the GDP series only, while the third and fourth columns show the AIC and BIC information criteria. The total log likelihood and information criteria of the first two models ((a) and (b)) cannot be compared to the other models because they are estimated with different data. For these two models industrial production and the index of working hours are rolling quarterly whereas they are monthly in the other models.

Table 1: Comparison of the model specifications using the full-sample results.

Model	L.L.	L.L. GDP	AIC	BIC
(a): AST DFM with Scale+Shape CFs	-381.4	-124	854.7	1097.5
(b): AST DFM (Sk Covs) with Scale+Shape CFs	-402.3	-124.6	892.7	1124.9
(c): DFM with SV and Scale CF	-548.9	-135.9	1159.8	1323.4
(d): AST DFM (Sk Covs) with SV and TV shape	-578.1	-136.3	1220.2	1389.1
(e): DFM with SV	-580.1	-138	1212.1	1349.4
(f): St-t DFM with SV	-569	-139.4	1196	1349.1

Note: L.L. refers to the model's total log likelihood. L.L. GDP refers to the log likelihood of the GDP series only. AIC $= -2 \times$ Log Lik. $+ 2 \times p$; BIC $= -2 \times$ Log Lik. $+ \log N \times p$; N is the number of observations and p the number of parameters estimated.

The results from model (a), (b) and (c) show that modelling common factors in scale and shape parameters yields to large improvements in the log likelihood

of GDP, with a particularly large benefit when modelling a shape common factor. Separately, the Gaussian model with scale common factor (model (c)) does better, in all metrics, than the benchmark models which do not feature a scale common factor. This highlights the importance of capturing cross-sectional dependencies in scale parameters.

Modelling fat tails does not necessarily improve the log likelihood of GDP

Modelling covariates with fat tails in the AST model improves the total log likelihood, the log likelihood of GDP (albeit marginally) and the information criteria. But while modelling fat tails in a dynamic factor model with stochastic volatility (specification (f)) improves the total log likelihood and information criteria compared to a model featuring Gaussian series, the likelihood of GDP deteriorates.

The decline in the predictive capacity of GDP when related series are allowed to have fat tails can be explained from the downweighting of large prediction errors in the related by the score functions. These prediction errors provide a signal on unforeseen turning points in the economic activity, and low tail parameters prevent this signal to be propagated to the GDP predictions. This result is consistent with the real-time analysis below which shows that fat-tails in the related series increase GDP nowcasting performance in normal times but reduce it in recessions.

The common scale component provides a consistent historical picture of forecasting uncertainty

Figure 4 shows the estimated common scale components. All three models featuring a common scale component yield broadly similar estimates. First, they reflect an overall decrease in volatility starting in the second half of the eighties which has been documented extensively in the macroeconomics literature since McConnell and Perez-Quiros (2000). Secondly, they feature sharp spikes during the recessions triggered by the 2007 financial crisis and lately by the coronavirus pandemic. While the common scale component at the time of the financial crisis reaches levels not seen since the seventies and early eighties, its level in the first two quarters of 2020 is unprecedented. This reflects the stark forecasting uncertainty generated by the government lockdown policies in response to the pandemic. Overall the common volatility factor carries economic meaning and gives a consistent picture of forecasting uncertainty historically and across different model specifications.

Recessions are generally associated with rising negative skewness

Location, scale and shape parameters derived from all models are compared in Figure (5). Unlike scale and shape parameters, location parameters show little divergence across models except at the time of the Covid-19 pandemic. All models feature a decrease in forecasting uncertainty starting in the mid-eighties. Interestingly the

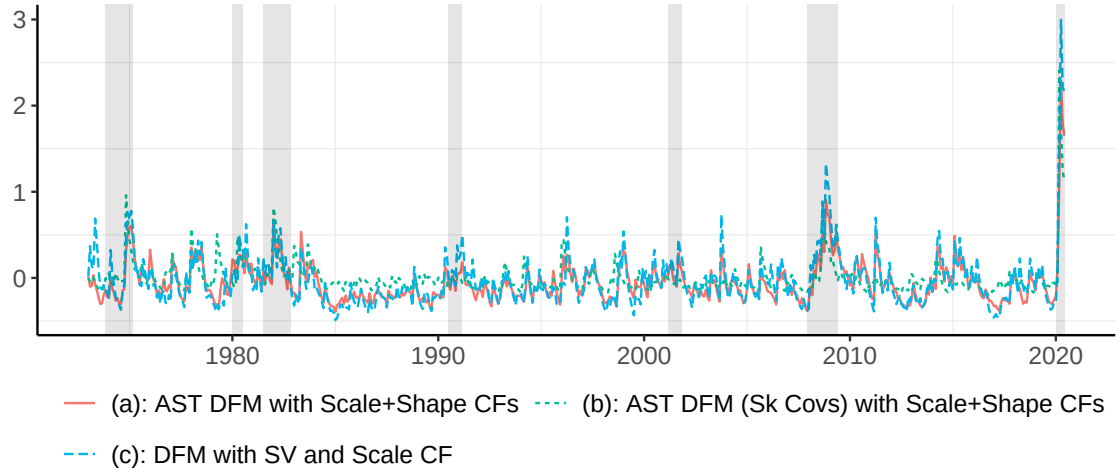

Figure 4: Scale common factors. Rolling quarterly estimate for model (a) and (b) and monthly estimate for model (c). Shaded areas indicate the US recessions classified by the NBER.. DFM = Dynamic Factor Model. SV = Stochastic Volatility; CF = Common Factor.

increase in forecasting uncertainty in recessions is generally associated with a sharp increase in the skewness parameter: the probability of negative prediction errors increases while positive prediction errors become less likely.

The results from the model with dynamic shape components are consistent with the findings of Carriero et al. (2020) and Delle-Monache et al. (2020) who show that, while statistical evidence for skewness in GDP growth is generally weak, this masks an erratic behaviour with periods of positive and negative skewness. While skewness fluctuates round zero until the mid-eighties, it becomes increasingly negative afterwards when overall volatility decreases. The financial crisis marks a change with a period of positive skewness.

Analysing the different model specifications using the entire sample is useful to understand the statistical features of the models but cannot be used to infer their effectiveness in real time. The next section shows the results from a recursive estimation simulating a real-time environment.

7 Real-Time Out-of-Sample Analysis

In this section the models are estimated recursively over a 20-year period from June 2000 up to June 2020 using historical vintages. Thus it is possible to analyse their performances in pseudo[3] real-time.

[3]It is not in exactly in real-time because the analysis was carried in August 2020 when all GDP figures were already released.

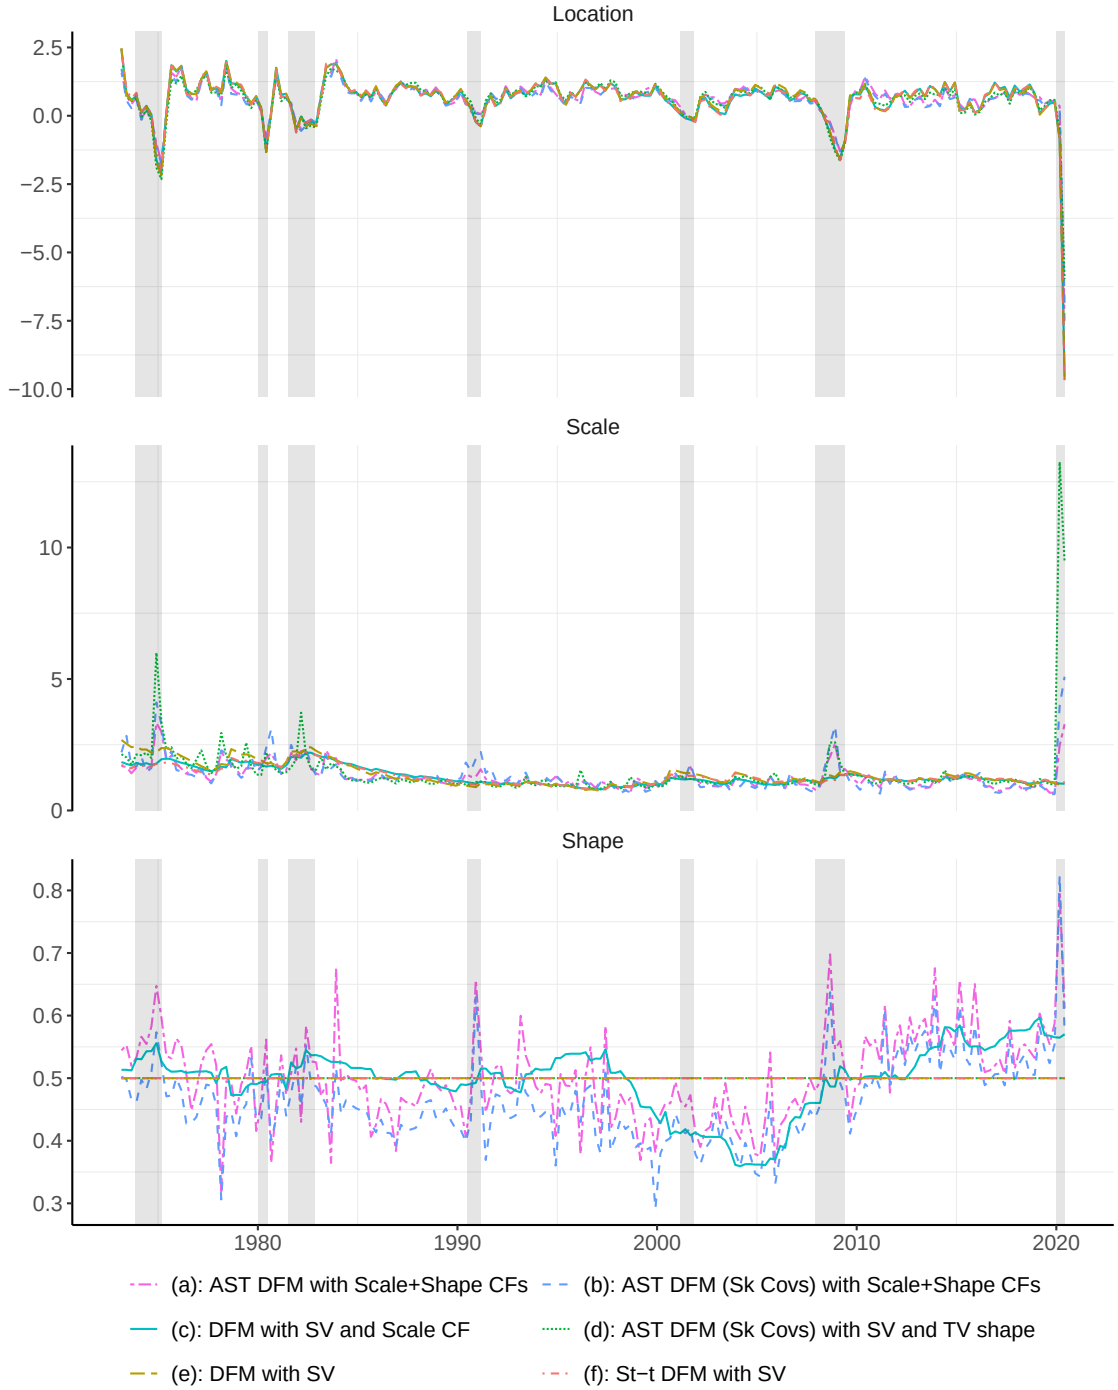

Figure 5: Quarterly figures corresponding to calendar quarters. Shaded areas indicate the US recessions classified by the NBER. SV = Stochastic Volatility; CF = Common Factor.

Deriving multi-step ahead density nowcasts

The GDP density nowcast in the last step of the nowcasting window is directly given by the one-step ahead prediction error density of GDP. But for earlier steps it is important to account for the uncertainty induced by missing values in the related series. This is done by drawing vectors of observations for each period with

missing observations and using the score driven recursion to retrieve the time-varying parameters in the next period for each draw, which are then used to generate new observations. Specifically, if the target is the GDP nowcast in period t, but related series are missing from $t-2$, then the one-step ahead densities of the related series at $t-2$ are used to draw vectors of prediction errors (a vector is composed of a draw for each related series) centred round the set of location estimates in $t-2$. When observations in related series are only missing partially the draws are replaced by the observed prediction errors for the related series which have been released. The score driven recursion is used on each vector of prediction errors which yields sets of scale, shape and location parameters for period $t-1$. These new sets of scale and shape parameters yield new sets of one-step ahead densities which are used to draw vectors of prediction errors round the vector of locations in $t-1$. Eventually, each vector of prediction errors in $t-1$ yields a GDP nowcast for t through the score driven recursion (related series are modelled with a lead of one period compared to GDP). The density nowcast is given by the empirical density attached to these GDP nowcasts.

Average Log Score

The accuracy of probabilistic predictions, or density nowcasts, is analysed using the average log score at each nowcasting step. The log score of a given nowcast is given by the nowcasted log density at the observed value. The better the probabilistic forecast, the higher the log score.

Forecasters and policy makers are particularly interested in the performance of forecasting methods in times of economic troubles. To address this question the results are analysed over five samples which cover the Covid-19 pandemic, the great financial crisis (GFC), all recessions taken from the NBER classifications since 2000, normal times (all periods except recessions) and finally a sample covering all quarters.

As more data accrue during the nowcasting quarter the average log score should increase monotonically. However, while the average log score is generally higher towards the end of the nowcasting window, it does not exhibit a monotonic improvement. This can be explained by the relatively low numbers of series modelled and the resulting low numbers of release dates during the quarters. Nevertheless this statistic remains useful to compare the models' predictive capacities.

First, scale and shape common factors improve the forecasting performance during recessions with a particularly large benefit at the end of the nowcasting window. Up until the last nowcasting step the Gaussian DFM with scale common factor performs especially well. The common factor in the scales captures sharp increases in forecasting uncertainty in a timely way, and unlike with the AST models featuring both scale and shape common factors, all related series can remain monthly; hence there is no loss in precision at the beginning of the nowcasting window. At the end of the nowcasting window, however, there is a benefit in modelling non-Gaussian

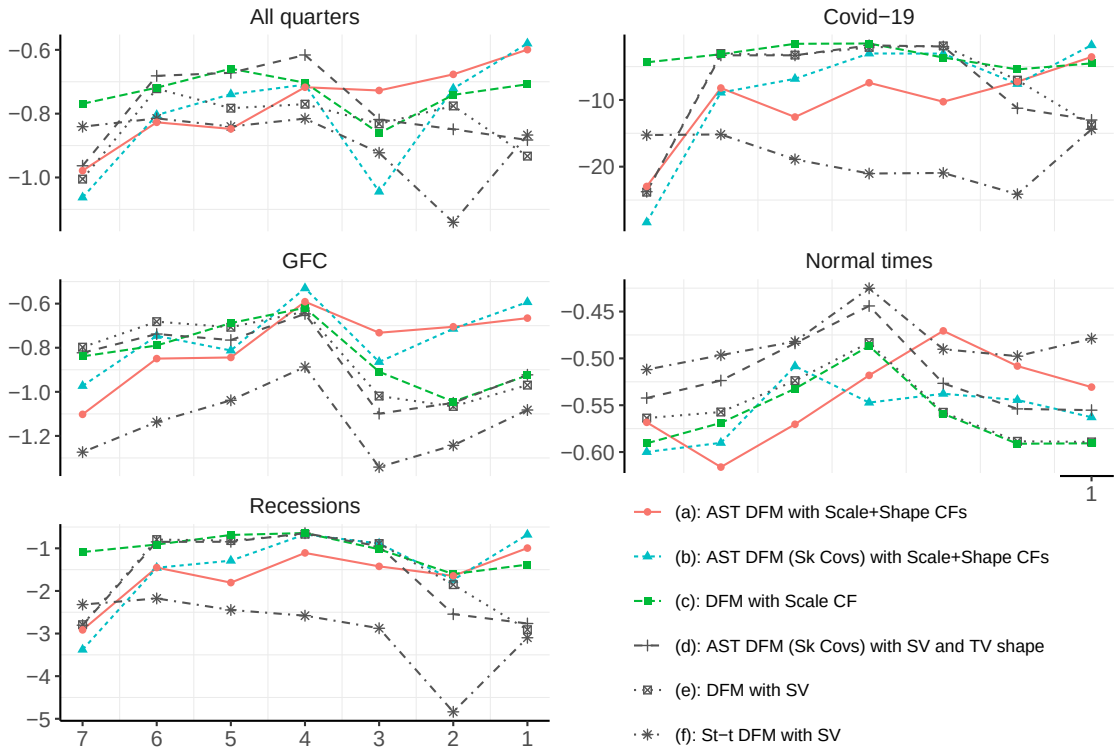

Figure 6: Average log score. Higher values imply better density forecasts. x-axes show the numbers of steps towards the release of the GDP number in the nowcasting window. New data are released at each step. DFM = Dynamic Factor Model. SV = Stochastic Volatility; CF = Common Factor. Covs = Covariates.

features and a common factor in the shape parameters in particular.

Both Gaussian models yield relatively similar performances up until the last nowcasting step, where the benefit of modelling a scale common factor clearly gives an advantage. The uncertainty attached to the GDP nowcast in both models can be updated through the dispersion of the common factor in the locations. But this mechanism takes place with a lag of one period because the dispersion of the location common factor is derived from the expected dispersion of the related series' forecast errors. Therefore, the uncertainty attached to the GDP nowcast is not updated once all related series have been observed in the quarter, unless there is a common factor in the scale parameters.

Separately, in the first half of the nowcasting window the AST models with scale and shape common factors (model (a) and (b)) perform less well than the other models excluding the Student-t DFM model (f). This most likely represents the cost of aggregating the related series featuring scale and shape common factors into rolling quarterly figures. However, AST models with scale and shape common factors do consistently better at the end of the nowcasting window.

The Student-t model performs especially poorly during recessions. This stems from the model's inability to capture turning points in both the location and scale parameters due to the downweighting of prediction errors by the score function when

tail parameters are low, as illustrated in figure 2. This downweighting is subdued in the AST model (a) because the related series feature larger tail parameters when aggregated quarterly. In normal times, however, this downweighting is beneficial because the central part of GDP density nowcasts is not affected by relatively large errors in the related series. In fact the Student-t DFM model does consistently better in normal times.

Figure 7 makes it easier to compare the accuracy of density nowcasts across models when considering both recessions and normal times. It shows the average log score at each nowcasting step for each model. The darker the colour, the higher log score and thus the better the density nowcast. Model (c), the Gaussian DFM with scale common factor, is the only model to perform well during the entire nowcasting window. It is only out-performed by the AST models with scale and shape common factors in the last nowcasting step.

	7	6	5	4	3	2	1
Model (a)	-5.71	-2.39	-3.33	-2.07	-2.72	-2.16	-1.27
Model (b)	-6.87	-2.50	-2.04	-1.10	-1.28	-2.26	-0.84
Model (c)	-1.53	-1.23	-0.83	-0.80	-1.40	-1.88	-1.62
Model (d)	-5.78	-1.22	-1.21	-0.83	-1.08	-3.25	-3.63
Model (e)	-5.80	-1.16	-1.22	-0.93	-1.04	-2.26	-3.83
Model (f)	-4.04	-3.96	-4.74	-5.15	-5.31	-6.37	-3.98

Figure 7: Average log score across all periods at each nowcasting step. The lower the nowcasting step the closer the release. The higher the log score the better the density forecast. And the darker the background colour the higher the log score.

7.1 The Covid-19 Pandemic

To understand why models with a scale common factor or both scale and shape common factors do better as more data accrue during the quarter and the GDP release gets closer, it is useful to analyse a period of stark economic uncertainty and the first two quarters of 2020 provide a good illustration for this.

Figure 8 shows the density nowcasts for the first quarter of 2020 starting from the March release of the industrial production figure relating to February. The effect of Covid-19 started to appear in the related series in March, and it was clear at the time that the economic environment was growing more uncertain. But the models with cross-sectionally independent stochastic volatilities (model (d) and (e) and (f)) show a relatively stable nowcasting uncertainty. Although this is counter-intuitive it is an inevitable feature of these models.

In normal times, the uncertainty attached to a nowcast should decline gradually as more data accrue during the quarter. This is because the uncertainty attached to each coming release is alleviated once the figures are published. In other words, the forecast horizon of the related series decreases after each release, and since

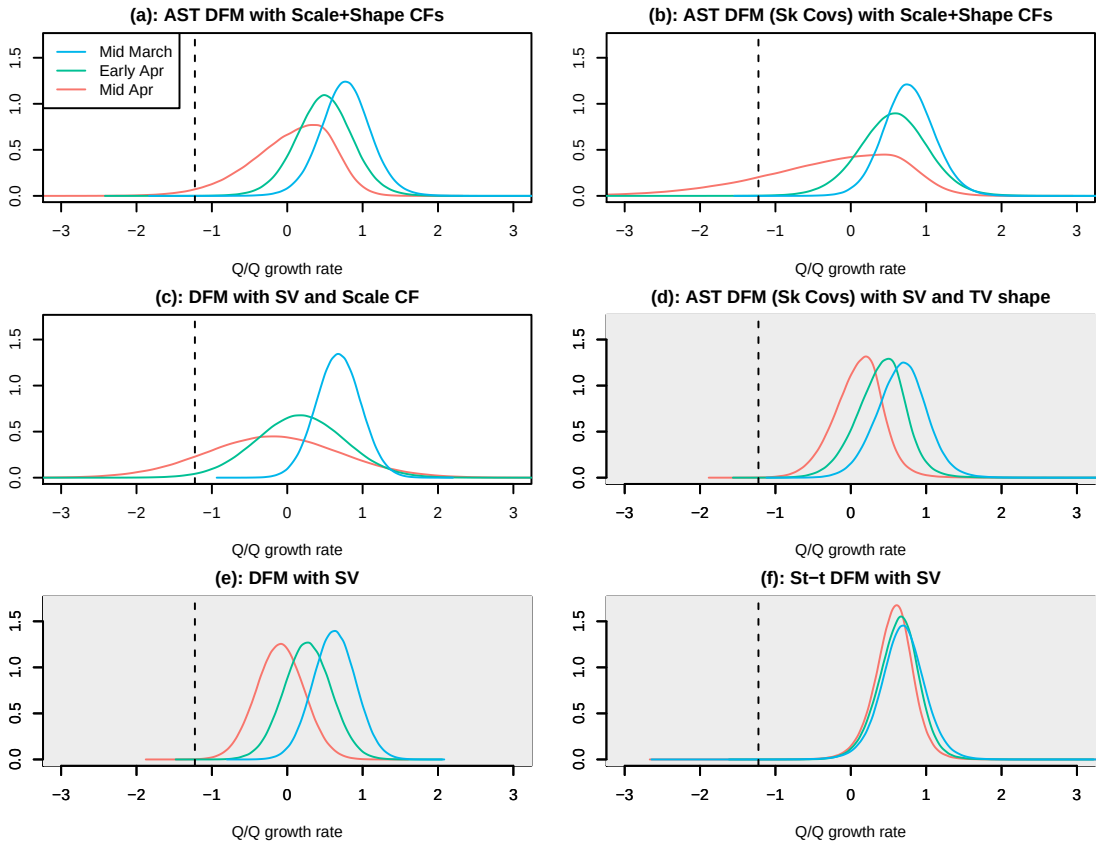

Figure 8: Real-time density nowcasts of US GDP Q1. The dashed line indicates the published figure. DFM = Dynamic Factor Model. SV = Stochastic Volatility; CF = Common Factor. TV = Time Varying.

the scale parameters remains relatively stable in normal times (or changes very moderately due to the expanding sample), forecast uncertainty decreases naturally. A shortcoming of models with constant variance parameters is that this reasoning also applies to periods of economic stress. As the forecast horizon decreases, nowcasting uncertainty can only decrease, and this even if the economic environment becomes more uncertainty.

Economic uncertainty is generally signalled by increasing prediction errors in series related to economic growth. By modelling stochastic scale parameters in these related series, larger prediction errors can be translated into greater conditional variances (the one-step ahead forecast error variances) in the related series.

Importantly, changes in the related series' forecast uncertainty also affect the uncertainty attached to the GDP nowcast, albeit with a lag. A greater forecast uncertainty in the related series, or more precisely an increase in the one-step ahead forecast error variances, increases the dispersion of the two-step ahead estimate of the common location component. Since the common location component is shared by all series, the dispersion in the GDP forecast increases as well. Hence, even though the scale parameters in stochastic volatility models are contemporaneously cross-sectionally independent, forecast error variances are related through the location

common factor.

But there is a lag of a one period before forecast errors in the related series affect the forecast error variance of GDP; the uncertainty attached to the one-step ahead GDP forecast is not affected.[4] The progress of the Covid-19 pandemic in the US provides an important illustration of this fact.

Modelling a common factor in the scale parameter alleviates this problem because the one-step ahead forecast error variances are now related directly. In the Gaussian model with scale common factor and the AST models with scale and shape common factors the uncertainty attached to the GDP nowcast for March increases gradually following the release of the related series. This increase in uncertainty is necessary to generate realistic probabilistic forecasts for March.

Moreover, the AST models with a shape common component can capture the asymmetry in the uncertainty attached to the nowcast in addition to a greater dispersion. Negative prediction errors are significantly more likely than positive ones. However, the AST model with fat tails (model (a)) produces a poorer density nowcast compared to the AST model with Gaussian tail parameters in the related series. Low tail parameters downweight the signal from the related series' prediction errors; hence the location, scale and shape parameters do not adjust enough to yield a satisfactory density nowcast.

The analysis of the June nowcast shown in figure 9 is more complicated because of the erratic behaviour of the projections in this quarter coming from the unprecedented downturn in the activity and resulting magnitude of the prediction errors in the related series. However, it is the AST model with scale and shape common factor and Gaussian tail parameters in the related series (model (b)) that yields the best probabilistic forecast at the end of the nowcasting window. Again the Gaussian model with common volatility component and the AST models with scale and shape common factors capture the uncertainty arising late in the nowcasting window which the other models do not identify.

8 Conclusion

This paper has shown how common components in scale and shape parameters may be modelled in Gaussian (scale only) and non-Gaussian mixed-frequency dynamic factor models. The empirical application in pseudo real time using US economic growth data shows that modelling scale and shape common factors yields better density nowcasts towards the end of the nowcasting window in recessionary episodes. While this paper uses a score driven approach for estimation, the gains in real-time performance stemming from modelling a common volatility factor should arise in other modelling strategies such as mixed-frequency vector auto-regressions and state space models which allow for stochastic volatility specifications.

[4]This is not necessarily true for models where unobserved states have their own sources of errors. But modelling the dependencies in conditional variances directly though scale parameters should nevertheless improve nowcast uncertainty in these models.

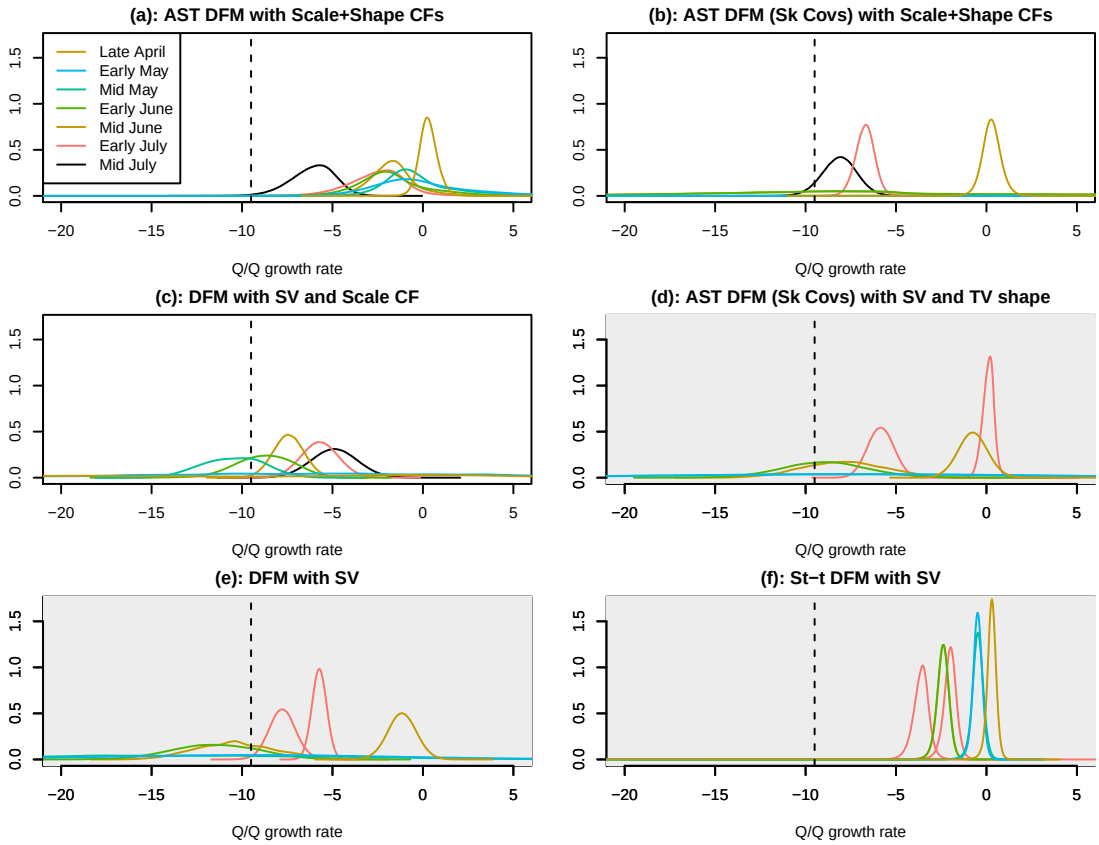

Figure 9: Real-time density nowcasts of US GDP Q1. The dashed line indicates the published figure. DFM = Dynamic Factor Model. SV = Stochastic Volatility; CF = Common Factor. TV = Time Varying.

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019, April). Vulnerable growth. *American Economic Review 109*(4), 1263–1289.

Antolin-Diaz, J., T. Drechsel, and I. Petrella (2017, May). Tracking the slowdown in long-run gdp growth. *The Review of Economics and Statistics 99*(2), 343–356.

Antolin-Diaz, J., T. Drechsel, and I. Petrella (2020). Advances in Nowcasting Economic Activity:Secular Trends, Large Shocks and New Data.

Banbura, M., D. Giannone, M. Modugno, and L. Reichlin (2013). Now-Casting and the Real-Time Data Flow. In G. Elliott, C. Granger, and A. Timmermann (Eds.), *Handbook of Economic Forecasting*, Volume 2 of *Handbook of Economic Forecasting*, Chapter 0, pp. 195–237. Elsevier.

Blasques, F., S. Koopman, M. Mallee, and Z. Zhang (2016). Weighted maximum likelihood for dynamic factor analysis and forecasting with mixed frequency data. *Journal of Econometrics 193*(2), 405–417.

Carriero, A., T. E. Clark, and M. Marcellino (2016). Common drifting volatility in large bayesian vars. *Journal of Business & Economic Statistics 34*(3), 375–390.

Carriero, A., T. E. Clark, and M. Marcellino (2018). Measuring uncertainty and its impact on the economy. *The Review of Economics and Statistics 100*(5), 799–815.

Carriero, A., T. E. Clark, and M. Marcellino (2020). Assessing international commonality in macroeconomic uncertainty and its effects. *Journal of Applied Econometrics 35*(3), 273–293.

Creal, D., S. J. Koopman, and A. Lucas (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics 28*(5), 777–795.

Creal, D., B. Schwaab, S. J. Koopman, and A. Lucas (2014). Observation-driven mixed-measurement dynamic factor models with an application to credit risk. *Review of Economics and Statistics 96*(5), 898–915.

Delle-Monache, D., A. De-Polis, and I. Petrella (2020). Modelling and Forecasting Macroeconomic Downside Risk. *Economic Modelling and Forecasting Group* (34).

Delle-Monache, D. and I. Petrella (2017). Adaptive models and heavy tails with an application to inflation forecasting. *International Journal of Forecasting 33*(2), 482–501.

Doz, C., L. Ferrara, and P.-A. Pionnier (2020, January). Business cycle dynamics after the great recession: An extended markov-switching dynamic factor model. (2020/01).

Durbin, J. and S. J. Koopman (2012). *Time series analysis by state space methods*. Oxford University Press.

Gorgi, P., S. J. Koopman, and M. Li (2019). Forecasting economic time series using score-driven dynamic models with mixed-data sampling. *International Journal of Forecasting 35*(4), 1735–1747.

Haldane, A. (2012). Tails of the unexpected. *Paper delivered at a conference sponsored by the University of Edinburgh Business School, Edinburgh, June 8*.

Harvey, A. C. (1989). *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press.

Harvey, A. C. (2013). *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series (Econometric Society Monographs)*. Cambridge University Press.

Huber, F. (2016). Density forecasting using bayesian global vector autoregressions with stochastic volatility. *International Journal of Forecasting 32*(3), 818–837.

Jensen, H., I. Petrella, S. H. Ravn, and E. Santoro (2020). Leverage and deepening business-cycle skewness. *American Economic Journal: Macroeconomics 12*(1), 245–281.

Labonne, P. and M. Weale (2020). Temporal disaggregation of overlapping noisy quarterly data: estimation of monthly output from uk value-added tax data. *Journal of the Royal Statistical Society: Series A (Statistics in Society) 183*(3), 1211–1230.

Lewis, D. J., K. Mertens, and J. H. Stock (2020). U.S. Economic Activity During the Early Weeks of the SARS-Cov-2 Outbreak. *Federal Reserve Bank of Dallas* (2011).

Lucas, A. and Z. Zhang (2016). Score-driven exponentially weighted moving averages and value-at-risk forecasting. *International Journal of Forecasting 32*(2), 293–302.

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

McConnell, M. M. and G. Perez-Quiros (2000). Output fluctuations in the united states: What has changed since the early 1980's? *American Economic Review 90*(5), 1464–1476.

Mitchell, J., R. J. Smith, M. R. Weale, S. Wright, and E. L. Salazar (2005). An indicator of monthly gdp and an early estimate of quarterly gdp growth. *The Economic Journal 115*(501), 108–129.

Osborn, D. R. (1976). Maximum likelihood estimation of moving average processes. *Annals of Economic and Social Measurement 5*(1), 75–87.

Salazar, E., R. Smith, M. Weale, and S. Wright (1997). A monthly indicator of gdp. *National Institute Economic Review 161*(1), 84–89.

Zhu, D. and J. W. Galbraith (2010). A generalized asymmetric student- distribution with application to financial econometrics. *Journal of Econometrics 157*(2), 297–305.

Appendix A Elements of the score and information matrix

The score with respect to location parameters is

$$\Delta^{\mu}_{i,t} = \frac{\partial \log p(y_{i,t}|Y_{t-1})}{\partial \mu_{i,t}} = \frac{\nu_{i,1}+1}{1+\frac{1}{\nu_{i,1}}\Big(\frac{y_{i,t}-\mu_{i,t}}{2\alpha_{i,t}\sigma_{i,t}K(\nu_{i,1})}\Big)^2}\cdot\frac{y_{i,t}-\mu_{i,t}}{\nu_{i,1}(2\alpha_{i,t}\sigma_{i,t}K(\nu_{i,1}))^2}\ 1(y_{i,t}\leq\mu_{i,t}) + \frac{\nu_{i,2}+1}{1+\frac{1}{\nu_{i,2}}\Big(\frac{y_{i,t}-\mu_{i,t}}{2(1-\alpha_{i,t})\sigma_{i,t}K(\nu_{i,2})}\Big)^2}\cdot\frac{y_{i,t}-\mu_{i,t}}{\nu_{i,2}(2(1-\alpha_{i,t})\sigma_{i,t}K(\nu_{i,2}))^2}\ 1(y_{i,t}>\mu_{i,t}). \tag{25}$$

The elements of the information matrix corresponding to location parameters are given by

$$\mathcal{I}^{\mu}_{i,t} = \mathrm{E}\big[\Delta^{\mu}_{i,t}\Delta^{\mu'}_{i,t}|Y_{t-1}\big] = \frac{1}{\sigma^2_{i,t}}\Big[\frac{\nu_{i,1}+1}{\alpha_{i,t}(\nu_{i,1}+3)K^2(\nu_{i,1})} + \frac{\nu_{i,2}+1}{(1-\alpha_{i,t})(\nu_{i,2}+3)K^2(\nu_{2,i})}\Big]. \tag{26}$$

The score with respect to scale parameters is

$$\begin{aligned}\Delta^{\sigma}_{i,t} = \frac{\partial \log p(y_{i,t}|Y_{t-1})}{\partial \sigma_{i,t}} &= \Big[\frac{(\nu_{i,1}+1)}{1+(\frac{y_{i,t}-\mu_{i,t}}{2\alpha_{i,t}K(\nu_{i,1})\sigma_{i,t}\sqrt{\nu_{i,1}}})^2} \\ &\times(\frac{y_{i,t}-\mu_{i,t}}{2\alpha_{i,t}\sigma_{i,t}K(\nu_{i,1})\sqrt{\nu_{i,1}}})^2-1\Big]/\sigma_{i,t}\ 1(y_{i,t}<\mu_{i,t}), \\ &+\Big[(\nu_{i,2}+1)\frac{1}{1+(\frac{y_{i,t}-\mu_{i,t}}{2(1-\alpha_{i,t})K(\nu_{i,2})\sigma_{i,t}\sqrt{\nu_{i,2}}})^2} \\ &\times\frac{y_{i,t}-\mu_{i,t}}{2(1-\alpha_{i,t})\sigma_{i,t}K(\nu_{i,2})\sqrt{\nu_{i,2}}})^2-1\Big]/\sigma_{i,t}\ 1(y_{i,t}>\mu_{i,t}).\end{aligned} \tag{27}$$

The elements of the information matrix corresponding to scale parameters are given by

$$\mathcal{I}^{\sigma}_{i,t} = \mathrm{E}\big[\Delta^{\sigma}_{i,t}\Delta^{\sigma'}_{i,t}|Y_{t-1}\big] = \frac{2}{\sigma^2_{i,t}}\Big[\frac{\alpha_{i,t}\nu_{i,1}}{\nu_{i,1}+3} + \frac{(1-\alpha_{i,t})\nu_{i,2}}{\nu_{i,2}+3}\Big]. \tag{28}$$

The score with respect to shape parameters is

$$\begin{aligned}\Delta^{\alpha}_{i,t} = \frac{\partial \log p(y_{i,t}|Y_{t-1})}{\partial \alpha_{i,t}} &= \frac{\nu_{i,1}+1}{\nu_{i,1}}\frac{1}{1+(\frac{y_{i,t}-\mu_{i,t}}{2\alpha_{i,t}K(\nu_{i,1})\sigma_{i,t}\sqrt{\nu_{i,1}}})^2} \\ &\times((\frac{y_{i,t}-\mu_{i,t}}{2\sigma_{i,t}K(\nu_{i,1})\sqrt{\nu_{i,1}}})^2\frac{1}{\alpha^3_{i,t}}\ 1(y_{i,t}<\mu_{i,t}), \\ &+\frac{\nu_{i,2}+1}{\nu_{i,2}}\frac{1}{1+(\frac{y_{i,t}-\mu_{i,t}}{2(1-\alpha_{i,t})K(\nu_{i,2})\sigma_{i,t}\sqrt{\nu_{i,2}}})^2} \\ &\times((\frac{y_{i,t}-\mu_{i,t}}{2\sigma_{i,t}K(\nu_{i,2})\sqrt{\nu_{i,2}}})^2\frac{1}{(1-\alpha_{i,t})^3}\ 1(y_{i,t}>\mu_{i,t}).\end{aligned} \tag{29}$$

The elements of the information matrix corresponding to shape parameters are given by

$$\mathcal{I}^{\alpha}_{i,t} = \mathrm{E}\big[\Delta^{\alpha}_{i,t}\Delta^{\alpha'}_{i,t}|Y_{t-1}\big] = 3\Big[\frac{\nu_{i,1}+1}{\alpha_{i,t}(\nu_{i,1}+3)} + \frac{\nu_{i,2}+1}{(1-\alpha_{i,t})(\nu_{i,2}+3)}\Big]. \tag{30}$$

The formulae for the information matrix can be found in Zhu and Galbraith (2010).

Finally, by defining the vector $a_{i,t} = (\mu_{i,t}, \sigma_{i,t}, \alpha_{i,t})'$, the score with respect to the time-varying parameters of series i is:

$$\Delta_{i,t} = \frac{\partial \log p(y_{i,t}|Y_{t-1})}{\partial f_t} = \frac{\partial \log p(y_{i,t}|Y_{t-1})}{\partial a_{i,t}} . \frac{\partial a_{i,t}}{\partial f_t} \tag{31}$$

while the information matrix is

$$\mathrm{E}\big[\Delta_{i,t}\Delta_{i,t}'|Y_{t-1}\big] = \Big(\frac{\partial a_{i,t}}{\partial f_t}\Big)' \begin{pmatrix} \mathcal{I}_{i,t}^{\mu} & 0 & \mathcal{I}_{i,t}^{\mu,\alpha} \\ 0 & \mathcal{I}_{i,t}^{\sigma} & 0 \\ \mathcal{I}_{i,t}^{\mu,\alpha} & 0 & \mathcal{I}_{i,t}^{\alpha} \end{pmatrix} \frac{\partial a_{i,t}}{\partial f_t}. \tag{32}$$

Appendix B Summary of the model specifications

Table 2: Description of the models

Model	Label	Conditional distributions	Specifications for location, scale and shape parameters	Data
(a)	AST DFM with Scale+Shape CFs	Asymmetric Student-t (AST) for GDP; Skewed Student-t for IP, and IWH; Normal for WEI.	Dynamic factor models for location, scale and shape parameters.	Quarterly GDP; rolling quarterly IP and IWH; monthly WEI
(b)	AST DFM (Sk Covs) with Scale+Shape CFs	Asymmetric Student-t (AST) for GDP; Skewed Normal for IP, and IWH; Normal for WEI.	Dynamic factor models for location, scale and shape parameters.	Quarterly GDP; rolling quarterly IP and IWH; monthly WEI
(c)	DFM with SV and Scale CF	Normal for all series	Dynamic factor model for locations (means) and scales (volatilities);	Quarterly GDP; monthly IP, IWH and WEI
(d)	AST DFM (Sk Covs) with SV and TV shape	Asymmetric Student-t (AST) for GDP; Skewed Normal for IP, and IWH; Normal for WEI.	Dynamic factor model for locations; Random walk models for scale and shape parameters.	Quarterly GDP; monthly IP, IWH and WEI
(e)	DFM with SV	Normal for all series	Dynamic factor model for locations (means) and stochastic scales (volatilities).	Quarterly GDP; monthly IP, IWH and WEI
(f)	St-t DFM with SV	Student-t for GDP, IP, and IWH; Normal for WEI.	Dynamic factor model for locations (means) and stochastic scales (volatilities).	Quarterly GDP; monthly IP, IWH and WEI

Note : The skewed normal distribution here is derived by constraining the AST distribution to have Gaussian tail parameters ($\nu_1 = \nu_2 = \infty$). The skewed Student-t is derived by constraining the AST distribution to have a unique tail parameter ($\nu_1 = \nu_2$). The Student-t is derived by constraining the AST distribution to have Gaussian shape and tail parameters ($\alpha = 0.5$ and $\nu_1 = \nu_2 = \infty$). IP = Industrial Production; IWH = Index of working hours; WEI = Weekly Economic Index