Nowcasting Economic Activity with Fat Tails and Outliers *

Seokki Simon Hong †

First version: 14 September 2023
This version: 15 October 2024

Abstract

The COVID-19 presented macroeconomic models with unique challenges, marked by extreme outliers in economic data. This paper extends dynamic factor models by explicitly incorporating outliers, moving beyond conventional data screening practices. The methodological contribution includes introducing fat tails and outliers multiplicatively into innovation volatility, and two distinct approaches for modeling outliers are presented to address large jumps. Empirical findings demonstrate that outlier-augmented models consistently outperform benchmark models in point and density forecasting, with the most significant improvements observed in nowcasting horizons. Incorporating outliers becomes particularly crucial during major crises, enhancing forecasting accuracy by 44% compared to the benchmark. The uniform-mixture approach is found to be more robust than the student-t models, as it targets extreme variations without disrupting the smoothness of the stochastic volatility process. Overall, this paper enhances macroeconomic modeling by explicitly addressing outliers, improving forecasting accuracy, and providing insights into economic dynamics during and after major crises like the COVID-19 pandemic.

Keywords: Now-casting, Dynamic factor models, Bayesian Methods.
JEL Classification: C11, C32, C38, C53, E37.

*I am grateful to Nikoleta Anesti, Nicholas Carnot, Catherine Doz, Riccardo Degasperi, Francesco del Prato, Thomas Drechsel, Laurent Ferrara, Livia Paranhos, Ivan Petrella, Pierre-Alain Pionnier, Alain Quartier, Giovanni Ricco, and seminar participants at the PSE Workshop on Nowcasting for insightful comments. I acknowledge support from the Measurement in Economics, Nowcasting-Beyond GDP chair at Paris School of Economics - École d'Économie de Paris during my time at PSE. The views expressed in this paper are those of the author and do not necessarily represent the views of the Bank of England or its committees.

†Bank of England, Threadneedle street, London EC2R 8AH, United Kingdom.
Email: simon.hong@bankofengland.co.uk

1 Introduction

The COVID-19 pandemic, with its widespread lockdowns and social distancing measures, presented macroeconomic models with an array of unique challenges. Economic indicators reached unprecedented levels during this period; for instance, in the week ending March 21, 2020, a staggering 3.28 million Americans filed for unemployment claims, surpassing the previous record set during the 1982 recession. The U.S. real GDP plummeted by a staggering 31.4% during the second quarter of 2020, marking an all-time low since the Bureau of Economic Analysis (BEA) began tracking these figures in 1947.[1] Such extreme observations left the New York Federal Reserve's benchmark nowcasting model in disarray, as this constant-parameter model resulted in implausible forecast paths.

Although the economic turbulence witnessed during the COVID-19 pandemic was unparalleled, the presence of outliers within macroeconomic data is a phenomenon with historical antecedents. It is, in fact, a recurring feature within the macroeconomic data landscape, as a manifestation of exceptional events such as labor strikes and natural disasters. Illustratively, Figure 1 portrays selected macroeconomic indicators from the French context spanning the period of 1980 to 2019. A cursory examination of this figure underscores the omnipresence of conspicuous, one-off outliers across four series. It is worth noting that the early contributions in the field of macroeconomic forecasting, including Stock and Watson (2002), recognized this phenomenon and proposed the screening of outliers via replacement with missing values.[2]

This paper extends the dynamic factor models (DFMs) paradigm by introducing explicit modeling of outliers, moving beyond data screening practices. Within the DFM framework, I take the view that outliers primarily represent transient spikes in volatility, not permanent disturbances. The model specification comfortably accommodates conventional models while providing a more intuitive way of dealing with outliers: I introduce and compare two distinct modeling strategies to address large jumps. Importantly, the estimation process is entirely data-driven, allowing the model to estimate outliers based solely on the data without prior constraints.

Subsequently, I take on the role of policymakers and market participants to conduct

[1] The figures are based on preliminary estimates: the FRED database now shows 2.91 million and 28% for initial claims and real GDP growth, respectively, after data revisions.

[2] Stock and Watson (2002) replaced observations exceeding 10 times the interquartile range from the median by missing values.

Figure 1: Existence of one-off outliers in French data

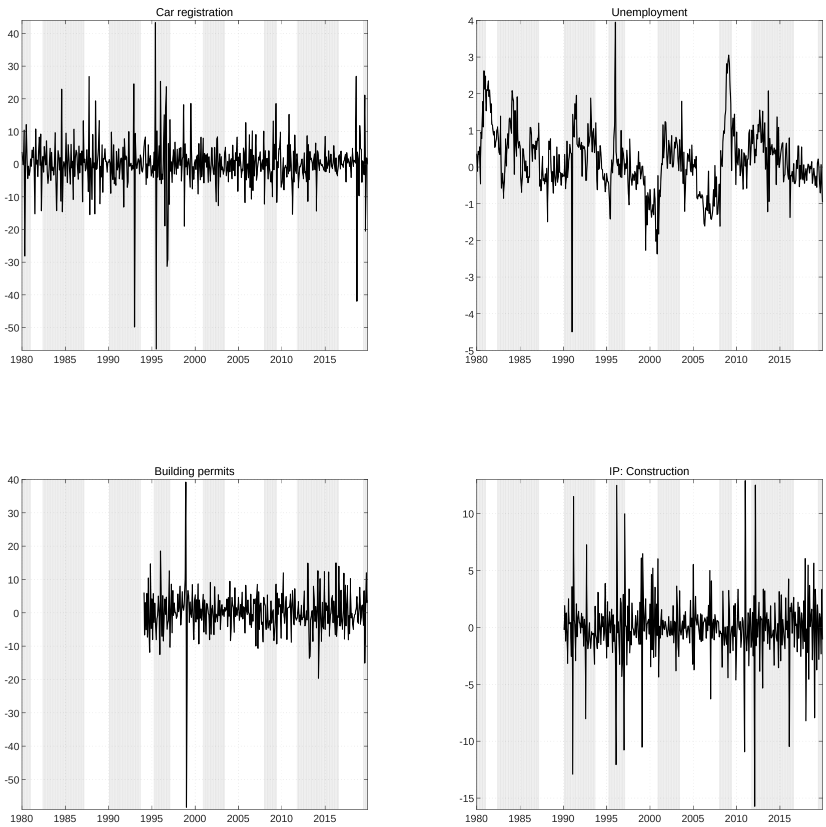

Note: This figure plots four monthly indicators for the French economy: passenger car registrations, unemployment, building permits, and industrial production: construction index. Sample: 1980.1 – 2019.12. The data are transformed in terms of month-on-month % changes. OECD recessions in gray. Source: DBnomics.

pseudo real-time out-of-sample forecasting analyses, comparing performance of the model to the benchmark for both pre- and post-COVID periods.[3] In the context of nowcasting French GDP, I show that the explicit modeling of outliers is particularly beneficial during crises, as shown in Figure 2: while the benchmark model tends to underestimate the severity of recession and overestimate economic dynamics during normal times, the model with outliers adeptly captures economic activity in a more timely manner. As a result, the forecasting accuracy improves by 44% compared to the benchmark even when post-2020 samples are included, and the outlier-augmented models accurately capture the actual GDP growth in the second and third quarters of 2020.

I also explore further refinement of these models that can anticipate the actual economic development during the COVID-19, by accommodating both frequent moderate

[3]Throughout the paper, the benchmark model refers to constant-parameter DFM models of Bańbura and Modugno (2014) and Bańbura et al. (2010), among others.

Figure 2: Nowcasting the French GDP, 2005 – 2019

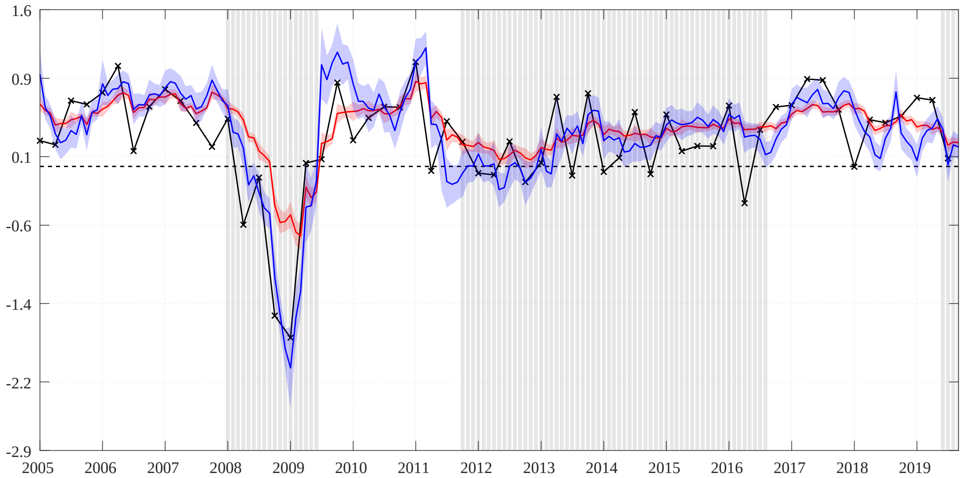

Note: This plot shows the realizations (in black) and pseudo real-time nowcasts of French GDP from 2005 to 2019 constructed from the model with outliers(in blue) and the benchmark DFM model (in red). Shaded areas denote the 68% bands. OECD recessions in gray.

shocks and rare large shocks, which characterise turbulent times.

The methodological contribution of this paper consists of two key components. Firstly, I introduce fat tails and outliers into the DFM framework by incorporating them multiplicatively into innovation volatility. The standard stochastic volatility model represents innovations with two separate error processes: conditional volatility following a log-normal auto-regressive model and a standard normal error. Inspired by the finance literature, particularly Jacquier et al. (2004), I further decompose the latter into an outlier state and noise, enhancing the model's resistance to outliers. This approach introduces pronounced jumps in the data rather than escalating stochastic volatility itself, resulting in a more stable volatility trajectory. While it is possible to model outliers additively, I demonstrate that the multiplicative approach achieves the same goal from a measurement-error perspective and remains independent of the specification of other components in the model. Consequently, outliers genuinely appear as one-off anomalies within this framework.

Secondly, I explicitly model outliers using two distinct approaches, discussing their different treatments of large jumps and interpretations of priors. Outliers are introduced into the model in two ways. One approach assumes that innovations exhibit fat tails and follow the Student-t distribution, a method more commonly employed in the literature. The other approach utilises a uniform mixture distribution that transitions between regular observations and outlier states, in the spirit of Stock and Watson (2016). I show that the latter strategy is better suited for addressing large jumps akin to the COVID-19 shock, by assigning more probability mass to events which more than double the volatility. The t-distribution, on the other hand, concentrates most probability mass on values close to one by construction. The uniform mixture approach is also more intuitive in terms of implementing and interpreting priors, as one can easily determine the prior probability of landing in outlier states by counting the occurrences of outliers in the pre-sample, even variable by variable. As we explore richer economic dynamics, we face the typical trade-off: estimating more parameters increases estimation uncertainty. Therefore, I adopt a Bayesian approach, which shrinks uninformative parameters toward a more parsimonious specification to address the curse of dimensionality.

Upon applying these models to forecast the French economy for both pre- and post-COVID periods, I arrive at three key findings. First, outlier-augmented models, irrespective of the distributional assumptions, consistently outperform the benchmark model in both point and density forecasting. In the pre-pandemic sample, the full model incorporating uniform outliers demonstrates $5-18\%$ improvements compared to the benchmark model, as indicated by reduced Root Mean Squared Error (RMSE) and Continuous Ranked Probability Score (CRPS) values across all forecasting horizons. The t-distribution approach results in similar gains, except for the longest forecasting horizon. However, the 'screening' method, which replaces pre-defined outliers with missing values, induces a negligible effect on accuracy. Notably, the most significant enhancements are observed in the nowcasting horizons, where models with outliers outperform both the benchmark and other stochastic volatility specifications in point and density nowcasts. This performance extends even to the full sample up to 2022, and the findings remain robust when alternative metrics for assessing prediction accuracy are employed.

Second, the significance of incorporating outliers becomes even more pronounced during major crises such as the Great Recession and the COVID-19 pandemic. For instance, when considering post-2020 data, the model with uniform outliers enhances the Root Mean Squared Error (RMSE) of the benchmark specification by 44%. This represents a substantial increase compared to the 13% improvement observed in the pre-COVID scenario. In times of extreme events, outlier-augmented models distinctly differentiate between short-term and long-lasting spikes in uncertainty. Consequently, they adequately capture the peaks and troughs characterizing the Great Recession and the COVID-19 crisis with greater timeliness and accuracy, while the benchmark model tends to underestimate the severity of the economic downturn and the subsequent rebound. The explicit modeling of outliers enhances the forecasting of economic developments, including the 'V-shaped' recovery in the initial phase of the pandemic.

Finally, when comparing different methods to incorporate outliers, the uniform mixture approach appears as a more robust option compared to the commonly employed t-distribution models. The uniform approach specifically targets extreme variations without unduly disrupting the smoothness of the stochastic volatility (SV) process when it is already behaving as expected. In contrast, t-distributed outliers tend to excessively suppress the SV process, unconditionally decreasing the level of volatilities for all indicators in the dataset. It further flattens already smooth time series, such as the SV estimates of GDP, resulting in nearly constant values over 40 years. These characteristics are less desirable in light of the model assumptions and the established narrative in the existing literature. Despite these drawbacks, the t-distribution method offers some benefits in terms of forecasting, particularly during times of crises. Ultimately, the choice between these approaches should be made based on the specific characteristics of the data and research objectives, such as the frequency of outliers and the need for a stable model.

The structure of the paper is as follows: Section 2 describes the econometric framework including the treatment of outliers. Section 3 introduces the data and provides results from the in-sample analyses using the pre-COVID sample, with a focus on comparing two different approaches to incorporate the outliers. Section 4 conducts pseudo real-time forecasting exercises with pre and post-COVID samples. Section 5 studies a modest extension that accommodates both frequent moderate shocks and rare large shocks, which characterised the COVID-19 period. Section 6 concludes.

Related Literature. This paper aims to achieve methodological advances in macroeconomic forecasting with DFMs. After seminal papers such as Stock and Watson (2002), Evans (2005), and Giannone et al. (2008) formalized the application of DFMs to nowcast the economy, there have been many efforts to incorporate structural changes within the DFM framework.[4] Del Negro and Otrok (2008) allow for time-varying factor loadings and volatilities, and Marcellino et al. (2016) introduce the stochastic volatil-

[4]See Bańbura et al. (2010) for the survey on nowcasting with standard DFMs.

ity in innovations to factors and idiosyncratic components. Antolin-Diaz et al. (2017) and Doz et al. (2020) feature shifts in the trend GDP growth and volatility over time. Antolin-Diaz et al. (2023) is the closest to this paper, where they introduce the outliers as an additive component that is assumed to follow the student-t distribution. While building upon their foundation, this paper also examines the limitations of their methodology: their additive specification inadvertently introduces a quasi-differencing induced bias, due to the dependency of outliers on other components within the model. Moreover, using the t-distribution for modeling outliers tends to dampen the stochastic volatility process in inherently smooth series. In response to these challenges, I propose alternative approaches to mitigate these issues by (i) incorporating outliers in a multiplicative fashion within the innovation volatility component and (ii) exploring two different strategies for modeling outliers, student-t and uniform distributions, plus the combination of them.

More broadly, it also speaks to an empirical literature that explores the role of time-variation, non-linearities, and non-gaussian shocks in macroeconomic models. This includes, but not limited to, Cogley and Sargent (2005), Primiceri (2005), Fernández-Villaverde et al. (2015), Carriero et al. (2022), and Chan (2023) in Bayesian Vector Autoregression (BVAR) models. This paper also takes a Bayesian perspective that enables the inclusion of more parameters and provides probabilistic predictions, particularly to take account of time-varying components and fat-tailed outliers.

The focus on the role of fat tails and outliers also relates to the studies addressing new challenges to empirical models after 2020, such as Lenza and Primiceri (2020), Marcellino et al. (2021), Schorfheide and Song (2021), and Cascaldi-Garcia (2022). Diebold (2020) studies the real-time performance of the business conditions index by Aruoba et al. (2009) during the early periods of the pandemic, and Lewis et al. (2022) provide a weekly economic index that tracks the rapid economic developments during the first 10 months of the pandemic. Ng (2021) interprets the variations around the outbreak as outliers and attempts to clean the data by using COVID indicators as controls. This paper also takes a similar view and explores the 2020 episode as a useful case study.

2 Econometric Framework

In this section, I introduce the dynamic factor model (DFM) which incorporates fat tails and outliers. Building on the models of Antolin-Diaz et al. (2017) and Marcellino et al. (2016), that introduce the shifts in long-run growth and stochastic volatility (SV) in innovations, I describe how to model outliers from two distinct approaches: the student-t and uniform mixture distribution. While the outliers are augmented by the SV of innovations in a multiplicative way in both cases, I discuss how they differ in dealing with large jumps and the interpretation of priors. In order to let the data speak, I impose the standard Minnesota-type of priors based on the stylized facts of the data. Finally, I briefly comment on the problem of mixed frequency and provide a sketch of the estimation algorithm.

2.1 The dynamic factor model

Let y_t denote n $\times$ 1 vector of macroeconomic indicators. The main idea of DFM is that a small number of latent common factors, f_t capture the majority of the joint dynamics across variables. Hence, it is assumed that the dimension of f_t is k $\times$ 1 and k $<<$ n. The factors are loaded via coefficients Λ(L), which represent the response of each indicator to a common shock. u_t is an idiosyncratic component that captures variable-specific movements or measurement errors. Eq (1) is what we usually call the "measurement equation" in the standard DFM literature.

$$\Delta y_t = c_t + \Lambda(L) f_t + u_t \tag{1}$$

In addition to the standard formulation, we allow the long-run growth rate of real GDP, and possibly other series, to drift gradually over time: the time-varying intercept c_t catches shifts in the long-run mean of Δy_t.[5] This specification is based on accumulating evidence that the long-run growth rate of GDP in advanced economies is lower than it has been over the past decades (Fernald, 2014; Gordon, 2014, for example). Antolin-Diaz et al. (2017) observe that standard DFM forecasts quickly revert to the unconditional mean of GDP, so taking account of the variation in long-run GDP growth substantially

[5] Macroeconomic indicators are denoted as Δy_t, as I apply appropriate transformation according to Table 1. Note that some indicators enter the estimation in levels, particularly the surveys, even though I take the first differences in most of the series.

improves point and density GDP forecasts even at very short horizons. Specifically, the time-varying intercept c_t is specified in the following way:

$$c_t = \begin{bmatrix} B & 0 \\ 0 & c \end{bmatrix} \begin{bmatrix} a_t \\ 1 \end{bmatrix}, \quad B = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

where c_t is composed of two components: the time-varying long-run growth rate a_t and a vector of constants c. Note that the selection vector B only loads onto the first two variables, GDP and consumption. It implies that a_t captures the time-variation in long-run real GDP growth, which is shared by real consumption growth. This specification is based on the permanent income hypothesis: consumers aim to smooth their consumption throughout the lifetime, so they react more to permanent changes in their income than transitory ones. Hence, rather than simply allowing a time-varying intercept only in GDP by setting B = 1, taking GDP and consumption together may help separate growth into long-run and cyclical fluctuations.[6]

Next, the latent factors and idiosyncratic components evolve with the following law of motion:

$$(1 - \phi(L))f_t = \sigma_{\epsilon_t}\epsilon_t \tag{2}$$

$$(1 - \rho_i(L))u_{it} = \sigma_{\eta_{it}} o_{it}\eta_{it}, \quad i = 1, ..., n \tag{3}$$

where $\phi(L)$ and $\rho_i(L)$ are lag polynomials of order p and q, respectively. Following the standard assumption from the literature, the residuals of the latent factor and idiosyncratic components are orthogonal and uncorrelated with each other, i.e. $\epsilon_t \sim iidN(0, 1)$ and $\eta_{it} \sim iidN(0, 1)$. Importantly, each indicator displays outliers o_{it}, which are incorporated in the stochastic volatility (SV) of innovations in a multiplicative way. A more detailed discussion on the specification of o_{it} follows in the next subsection.

Finally, following Primiceri (2005) and others, the time-varying parameters of the model follow driftless random walks:

$$a_t = a_{t-1} + v_{a,t}, \quad v_{a,t} \sim N(0, \omega_a^2) \tag{4}$$

$$log\sigma_{\epsilon_t} = log\sigma_{\epsilon_{t-1}} + v_{\epsilon,t}, \quad v_{\epsilon,t} \sim N(0, \omega_\epsilon^2) \tag{5}$$

[6]While the low-frequency component of growth could be shared by other series, one also faces the risk of misspecification. So I leave any slow-moving components in other indicators to be absorbed by the idiosyncratic components. More discussion on this can be found in Antolin-Diaz et al. (2017).

$$log\sigma_{\eta_{i,t}} = log\sigma_{\eta_{i,t-1}} + v_{\eta_i,t}, \quad v_{\eta_i,t} \sim N(0, \omega^2_{\eta,i}), \quad i = 1, ..., n \tag{6}$$

where σ_{ϵ_t} and $\sigma_{\eta_{i,t}}$ capture the SV of innovations to factors and idiosyncratic components. Modeling the trend as a random walk may sound unrealistic since it means that the movement of real GDP growth could be unbounded. However, the drift takes place only for a finite time, and the variance of this process is small, so they alleviate the problem.[7]

Note that this specification easily nests the DFMs previously proposed in the literature. By shutting down the outliers and loadings on lagged factors, i.e. $o_{it} = 1$ and $\Lambda(L) = \Lambda$, the model shrinks towards Antolin-Diaz et al. (2017). If we further limit the time-variation in the intercepts, $c_t = c$, we recover the model of Marcellino et al. (2016), which features the DFM with SV. Finally, this model coincides with a more commonly used DFM model of Bańbura and Modugno (2014) by imposing dogmatic priors $\omega^2_a = \omega^2_\epsilon = \omega^2_{\eta,i} = 0$, that turn off the SV features.

2.2 Treatment of Outliers

In this paper, I construct an outlier-adjusted stochastic volatility (SV) model within the DFM framework. While a typical SV-DFM model assumes changes in volatility to be highly persistent, as in Eq (6), many macroeconomic time series display large, one-off outliers due to policy changes, strikes, or natural disasters, even in the pre-COVID period.[8] By definition, these extreme observations are more reflective of transitory, rather than permanent, spikes in volatility.

Hence, following Jacquier et al. (2004), I model the outliers in a multiplicative fashion as in the Eq (3). If we set $o_{it} = 1$, it corresponds to the standard SV model where the innovations consist of two components with separate error processes: the conditional volatility, $\sigma_{\eta_{i,t}}$ follows a log-normal auto-regressive model, and the standard normal error. To allow for fat tails or outliers, I further separate the latter into o_{it} and the standard noise $v_{\eta_i,t}$. In the basic SV model, a spike in innovation means the volatility is high, but introducing fat-tails provides an additional source of flexibility: it takes a spike in data by introducing a large o_{it} before increasing $\sigma_{\eta_{i,t}}$. Hence, the intuition is that the basic model results in a more variable sequence of the SV, while this model is able to resist

[7]An alternative strategy would be the model with discrete breaks. However, this is less likely to be robust to misspecification than the reverse case, if the true data generating process is a random walk.

[8]Stock and Watson (2002) already noted the existence of such extreme observations.

outliers and results in more steady dynamics of $\sigma_{\eta_{i,t}}$.

I introduce outliers to the model in two ways. One way is to assume that innovations display fat tails and follow the Student-t distribution, i.e. $\nu_i/o_{i,t} \sim \chi^2_{\nu_i}$, as in Jacquier et al. (2004). The degree of freedom ν_i is jointly estimated with other parameters of the model, and the latent variable o_{it} can be obtained by a scale mixture. While this specification has been adopted more extensively in the literature, Student-t is not the only way to model outliers. The other way is to assume the following uniform mixture distribution for outliers, in the spirit of Stock and Watson (2016):

$$o_{it} = \begin{cases} 1 & \text{with probability} \quad 1 - p_i \\ U(2, 10) & \text{with probability} \quad p_i \end{cases} \tag{7}$$

where the probability of ending up in outlier states, p_i, follows the beta distribution with parameters α, representing the number of occurrences of outliers, and β, that of the normal times. This specification enables a highly persistent volatility state to infrequently and temporarily jump to the outlier states above 1.[9]

This modeling strategy is more advantageous than the student-t approach in at least two aspects. First, it is more geared towards modeling large jumps, like the COVID. Since the t-distribution assigns most probability to values close to one, most of the outliers are of moderate size. It also assigns some mass to values below one; this is not the perfect tool to deal with the COVID period, where the drop in real activity was at least 5 times deeper than any other recession since 1960. In the alternative approach, on the other hand, the outlier states cannot take values below one. There is equal probability for outlier states between 2 and 10, putting relatively more mass on the events that increase volatility by more than twofold: so it is more adequate for modeling large jumps, like COVID.

Moreover, this approach is more intuitive in terms of implementing and integrating priors. In the case of the t-distribution, it is not straightforward to justify the variable-specific prior on the degree of freedom ν_i.[10] However, since the parameter p_i follows the beta distribution in the latter approach, it is easier to implement and justify the priors:

[9] I discretize the distribution of o_{it} using a grid, [1:1:10].

[10] For instance, Antolin-Diaz et al. (2023) impose $\nu_i = 1$ for monthly and 30 for quarterly variables, just to reflect the fact that outliers are observed less in case of the lower frequency. Such priors are not based on exact derivation, so it is hard to extend them to take account of variable-specific rates.

for instance, one can form the prior by counting the number of occurrence of outliers in the pre-sample, variable by variable. I put a prior on p_i based on the stylized fact observed by Stock and Watson (2016) that the outlier occurs once every four years in 10 years of the pre-sample.

Note that despite such differences, both approaches share the same latent state representation: residuals are written as the product of noise and outlier state. Antolin-Diaz et al. (2023) propose an alternative way, where the outliers are present in the additive, rather than multiplicative fashion:

$$\Delta y_t - o_t = c_t + \Lambda(L) f_t + u_t, \quad o_{it} \sim t_{vi}(0, \sigma^2_{oi}) \tag{8}$$

and construct the outlier-adjusted data via the Kalman filter. I show that the multiplicative approach achieves the same goal, while being free from quasi-differencing induced bias that is present in the additive specification. First, by rearranging, we can rewrite the above equation as:

$$\Delta y^*_t = c_t + \Lambda(L) f_t + u_t, \quad u_t \sim N(0,1)$$

$$\Delta y_t = \Delta y^*_t + o_t, \quad o_t \sim t_v(0, \sigma^2_o)$$

and this corresponds to the standard regression with measurement errors: since the residuals are the sum of idiosyncratic components u_t and outlier states o_t, one can interpret this as if the true data were sometimes observed with large errors. Then, identifying the measurement error and subtracting from the data is not the only way to remove it. Instead, we can formulate it in another way. The residual variance is sometimes larger due to the outliers, and the aim is to downweight them when estimating parameters: this corresponds to the multiplicative approach.

Meanwhile, the additive specification may lead to biased estimates. After quasi-differencing the Eq (8), we can express the variance of residuals as $\mathrm{V}[(1-\rho_1 L-\rho_2 L^2)(u_{it}+o_{it})]$[11]. In order to simplify the problem, I add two additional assumptions:

Assumption 1. *The idiosyncratic components u_{it} and outlier states o_{it} are cross-sectionally orthogonal and uncorrelated to each other.*

[11]In this case, I model the idiosyncratic components as the AR(2) process.

FIGURE 3: STOCHASTIC VOLATILITY: ADDITIVE V. MULTIPLICATIVE

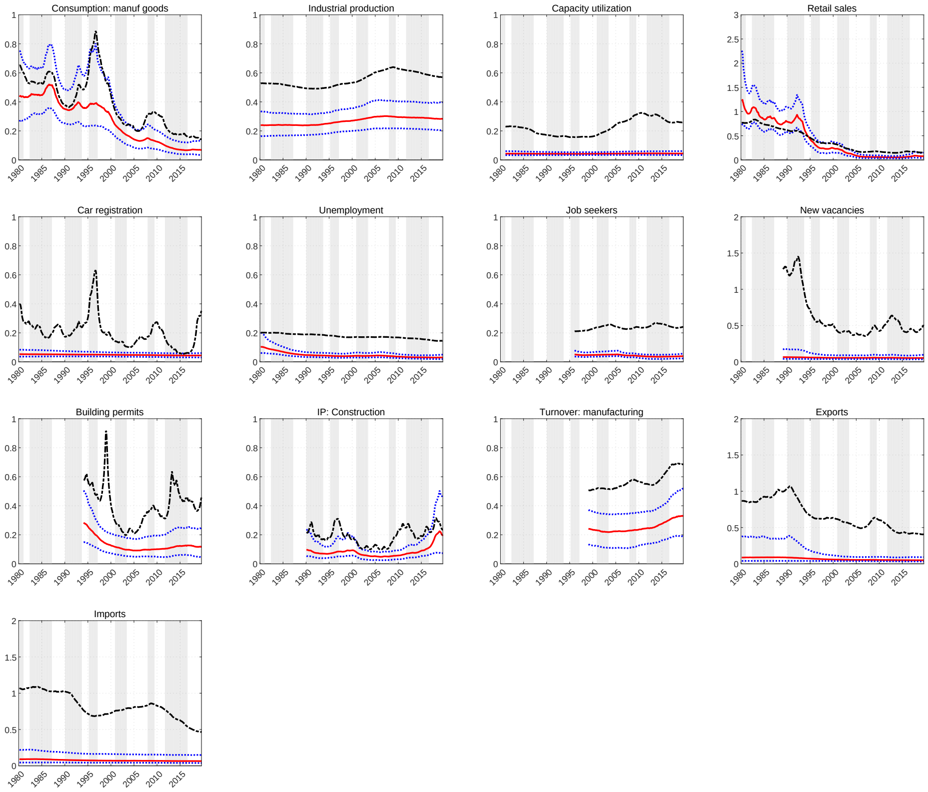

Note: This plot shows the posterior estimate of stochastic volatility in the idiosyncratic components for selected indicators. The black line represents the posterior median from the model with multiplicative outliers. Red and blue lines are the posterior median and 68% bands from the additive approach, respectively. Sample: 1980.1 – 2019.9. OECD recessions in gray.

Assumption 2. *The outlier states o_{it} are serially uncorrelated.*

which are mostly innocuous by the definition of "one-off" outliers. Then, the variance of residuals corresponds to:

$$V[(1-\rho_1 L-\rho_2 L^2)(u_{it}+o_{it})] = \sigma^2_{\eta_{it}} + (1+\rho_1^2+\rho_2^2)V(o_{it}), \quad V(o_{it}) = \sigma^2_{oi}(\frac{\nu_i}{\nu_i-2})$$

I provide a more detailed derivation in Appendix A.3. Here the role of outlier clearly depends on ρ, the autoregressive coefficients of the idiosyncratic component u_{it}, not the outlier states o_{it}. This is not an intended feature of the model. Moreover, in the likely case of $\rho > 0$, it amplifies the variance of the outlier component: resulting in an underestimation of the SV components in innovations.

To provide additional support for the derivation above, I study whether the analytical results also hold empirically. Figure 3 shows the estimates of idiosyncratic volatility for nine indicators from our dataset. The black line displays the results from the model with multiplicative t-outliers, while the red and blue lines indicate the posterior median and credible intervals from the additive t-outliers approach. Estimates from the latter fall significantly below the black line in most cases, confirming the analytical prediction that the additive approach results in an underestimation of the SV components. This underestimation arises as a by-product of the quasi-differencing technique, which is unique to the additive modeling method. In contrast, the multiplicative approach proposed in this paper avoids such unintentional biases.

2.3 Priors and model settings

Based on previous empirical findings, the order of the lag polynomial $\Lambda(L)$ is set to 1, i.e. it loads onto the contemporaneous and the first lag of the factor. Camacho and Perez-Quiros (2010) reported that the survey data were better aligned with a distributed lag of GDP; Antolin-Diaz et al. (2023) found that adding heterogeneous dynamics 'rebalance' the panel by allowing more weight to the "hard" variables, which are otherwise under-weighted compared to the "soft" indicators. The number of lags in the autoregressive coefficients of factors and idiosyncratic components, $\phi(L)$ and $\rho_i(L)$, is set to 2, to allow for medium-term frequency fluctuations as in Stock and Watson (1989).

I adopt a Bayesian approach with informative "Minnesota" style priors (Litterman, 1986) for the coefficients in $\Lambda(L)$, $\phi(L)$ and $\rho_i(L)$.[12] I set the prior mean of $\phi(L)$ to 0.9 for the first lag and zero in other lags, to incorporate a belief that the common, contemporaneous latent factor captures a highly persistent but stationary process. The prior mean of $\Lambda(L)$ is set to the estimate of the standard deviation of each variable for the first lag and zero otherwise, to reflect the belief that the factor corresponds to the cross-sectional average of standardized variables. The prior for $\rho_i(L)$ centers on zero for all lags, the parsimonious model with no serial correlation. This prior aims to shrink the idiosyncratic components towards an iid measurement error. I impose the variance on the priors to be $\frac{\gamma}{l^2}$ for all of the coefficients: the hyperparameter that governs tightness of the prior, γ is set to 0.2, the standard choice from the VAR literature, and l is the lag order. The idea is to shrink distant lags more strongly than the contemporaneous ones.

To express a preference for the more parsimonious model, I impose priors that shrink

[12] These are the most commonly adopted macroeconomic priors for VARs and formalize the view that an independent random-walk model for each variable in the system is a reasonable ground for beliefs about their time series behavior (Sims and Zha, 1998).

variances of the time-varying parameters, $\omega_a^2, \omega_\epsilon^2, \omega_{\eta,i}^2$, close to zero. Specifically, I set an inverse gamma (IG) prior with one degree of freedom for these parameters – just the minimum to obtain proper prior distributions. Regarding the scale parameter of the IG distribution, I set $\omega_a^2, \omega_\epsilon^2$ to 0.001 and $\omega_{\eta,i}^2$ to 0.0001. Following the approach of Primiceri (2005), I incorporate a prior about the time-variation: for example, the value of 0.001 corresponds to the conservative belief that the posterior mean of the long-run growth rate fluctuates with a standard deviation of around 0.4 percentage points in annualized terms over a period of ten years. In the end, the prior specification is conservative enough to shrink the model towards a commonly used DFM model of Bańbura and Modugno (2014) without the time-varying means and the SV, while still being loose enough to let the data speak.

2.4 Mixed frequency and missing data

Macroeconomic indicators are measured at different frequencies. For instance, the real GDP and national accounts variables are usually observed every quarter, while survey data and monthly indicators of real activity are released every month. To efficiently integrate information from the data observed at different frequencies, the standard way is to follow Mariano and Murasawa (2003) which links the *observed* growth rates of quarterly variables y_t^Q to *unobserved* monthly growth y_t^M via the first-order Taylor expansion:

$$y_t^Q = \frac{1}{3}y_t^M + \frac{2}{3}y_{t-1}^M + y_{t-2}^M + \frac{2}{3}y_{t-3}^M + \frac{1}{3}y_{t-4}^M \tag{9}$$

The mixed frequency issue collapses into the problem of missing data, where the model is specified at monthly frequency and we only observe every third observation of y_t^Q. The standard approach is to estimate the latent factors, parameters, and missing data points jointly using the state space representation of the DFM via the Kalman filter. Note that substituting the Eq (1) into (9) yields:

$$y_t^Q = \frac{1}{3}\lambda_y' f_t + \frac{2}{3}\lambda_y' f_{t-1} + \lambda_y' f_{t-2} + \frac{2}{3}\lambda_y' f_{t-3} + \frac{1}{3}\lambda_y' f_{t-4} \tag{10}$$

$$+ \frac{1}{3}u_{yt} + \frac{2}{3}u_{y,t-1} + u_{y,t-2} + \frac{2}{3}u_{y,t-3} + \frac{1}{3}u_{y,t-4} \tag{11}$$

This drastically increases the size of the state vector, which includes 4 lags of f_t and u_t. While the dimension of the state vector is only k $\times$ (p+1) if all indicators are at the monthly frequency (so no approximation is needed), the use of this approximation requires the dimension to be max(p,5)$\times n_Q \times$ k. For instance, in the case of p = 2, k = 1, and n_Q = 4, the size of the state space rises from 3 to 20, more than a six-fold increase. Such an expansion leads to unnecessarily complicated state-space representation and significant rise in computation costs.

To remedy this issue, I specify the model entirely at the monthly frequency by using the interpolated monthly values for quarterly indicators. Specifically, instead of applying the Mariano and Murasawa (2003) approximation to the quarterly variables y_t^Q, we can also implement this for the unobserved quarterly idiosyncratic components, u_t^Q. Then Eq (1) and (3) yield the following state-space representation:

$$u_{i,t}^Q = Hx_{i,t} + \xi_{i,t}, \quad \xi_{i,t} \sim N(0, R) \tag{12}$$

$$x_{i,t} = Ax_{i,t-1} + e_{i,t}, \quad e_{i,t} \sim N(0, Q) \tag{13}$$

with $x_{i,t} = [u_{i,t}^M \; u_{i,t-1}^M \; u_{i,t-2}^M \; u_{i,t-3}^M \; u_{i,t-4}^M]'$, H = $[\frac{1}{3} \; \frac{2}{3} \; 1 \; \frac{2}{3} \; \frac{1}{3}]$, A = $[\rho_1 \; \rho_2 \; 0_{1\times3}; \; I_4 \; 0_{4\times1}]$, Q = $[\sigma_{\eta_{it}} o_{it}; \; 0_4]$, and R is a small number. We estimate this system using the Kalman filter at the beginning of the algorithm to obtain the latent variable $\hat{u}_{i,t}^M$. Then, together with $\hat{\lambda}_y$ and $\hat{f}_t$ obtained from the previous iteration, we obtain an *unobserved* monthly-interpolated quarterly value $y_t^M = \hat{\lambda}_y \hat{f}_t + \hat{u}_{i,t}^M$. A more detailed explanation follows in Appendix A.2.

2.5 Estimation algorithm

I estimate the eq (1) – (7) with the Bayesian approach using a hierarchical Gibbs sampler (Moench et al., 2013), to obtain the posterior distribution of parameters and factors. So whenever possible, I parallelized the steps to reduce the computational costs. Since the idiosyncratic components in the model feature autocorrelation, the state space is rewritten in terms of quasi-differences. The sampling method for the SVs follows Kim et al. (1998), which provide us with faster speed than the exact Metropolis-Hastings algorithm of Jacquier et al. (2004). For the identification of factors, I adopt the steps proposed by Bai and Wang (2015). The full details of the algorithm can be found in

Appendix A.2, and below I provide a sketch.

Algorithm. Let $\theta = \{\lambda, \Phi, \rho, \omega_a, \omega_\epsilon, \omega_\eta, \nu\}$ be the underlying parameters of the model, and Φ, ρ represent the autoregressive coefficients for the factor and idiosyncratic components. The latent states to be estimated are $\{a_t, f_t, \sigma_{\epsilon,t}, \sigma_{\eta i,t}, o_{it}\}_{t=1}^T$. The superscript j denotes a current draw. The algorithm consists of the following steps:

1. *Construct monthly-interpolated values for quarterly variables.* For the quarterly variables i = $1...n_Q$, compute $\Delta y_{it}^Q - c_{it} - \lambda_i(L) f_t$, given $a_t^{j-1}, f_t^{j-1}, \lambda^{j-1}$ from the previous iteration. Then, using the state space in Eq (12) and (13) and conditional on $\rho^{j-1}, \{\sigma_{\eta i,t}^{j-1}\}_{t=1}^T$, draw $\hat{u}_{i,t}^M$ by employing the Kalman filter and smoother. Then we obtain the monthly-interpolated values, $\Delta y_{it}^{M,Q} = c_{it} + \lambda_i(L) f_t + \hat{u}_{i,t}^M$.

2. *Draw the latent factors and autoregressive coefficients.* Conditional on Δy_{it}^M and the parameters θ^{j-1},

   (a) Draw the factor and the trend, $p(\{a_t^j, f_t^j\}_{t=1}^T | \theta^{j-1}, \{\sigma_{\epsilon,t}^{j-1}, \sigma_{\eta i,t}^{j-1}\}_{t=1}^T, y)$ using the Kalman filter and smoother.The state space is rewritten in terms of quasi-differences at this step, conditional on ρ^{j-1} and $\{\sigma_{\epsilon,t}^{j-1}, \sigma_{\eta i,t}^{j-1}\}_{t=1}^T$.

   (b) Draw the variance of the time-varying GDP growth component from the Inverse-Gamma (IG) posterior, $p(\omega_a^{2,j} | \{a_t^j\}_{t=1}^T)$.

   (c) Given $\{f_t^j\}_{t=1}^T$ from the previous step, draw the autoregressive parameters of the factor VAR, Φ^j from the Normal posterior $p(\Phi^j | \{f_t^j, \{\sigma_{\epsilon,t}^{j-1}\}_{t=1}^T)$. Then draw the SV component of innovation to the factors, $p(\{\sigma_{\epsilon,t}^j\}_{t=1}^T | \Phi, \{f_t\}_{t=1}^T)$, using a mixture of normals following Kim et al. (1998). Draw $\omega_\epsilon^{2,j}$ conditional on $\{\sigma_{\epsilon,t}^j\}_{t=1}^T$ and the IG prior.

3. *Draw the factor loadings and serial correlation coefficients of idiosyncratic components.* For each variable i = 1...n and conditional on $\{f_t^j\}_{t=1}^T$,

   (a) Draw the loadings λ^j from $p(\lambda^j | \rho^{j-1}, \{f_t^j, \sigma_{\eta i,t}^{j-1}\}_{t=1}^T, y)$. They can be estimated via GLS, conditional on ρ_i^{j-1} and $\sigma_{\eta i,t}^{j-1}$. Following Bai and Wang (2015), I restrict the loading of GDP on f_t to be unity for the identification.

   (b) Draw serial correlation coefficients ρ^j from $p(\rho | \lambda^j, \{f_t^j, \sigma_{\eta i,t}^{j-1}\}_{t=1}^T, y)$ from the Normal posterior, based on the $\{u_{i,t}^j\}_{t=1}^T = y_t - \lambda^j f_t^j$ and the autoregression with heteroskedasticity.

4. *Draw the outlier-adjusted SV of innovations.* For each variable i = 1...n, I obtain the residuals $u^*_{i,t} = (1-\rho^j(L))u^j_{i,t}$. Then,

   (a) (in the case of t-distribution) draw p$(\nu^j|u^*_{i,t}, \sigma^{j-1}_{\eta i,t})$ following Jacquier et al. (2004), and also draw p$(o^j_{i,t}|\nu^j, u^*_{i,t}, \sigma^{j-1}_{\eta i,t})$, using $\nu^j/o^j_{i,t} \sim \chi^2_{v+1}$.

   (b) Obtain mixture states conditional on $(u^*_{i,t}, o^j_{i,t}, \sigma^{j-1}_{\eta i,t})$ in logs, following Kim et al. (1998). Construct the state space with $log\sigma_{\eta_{it}}$ as the latent variable, and apply the Kalman filter and smoother to draw new $\{\sigma^j_{\eta i,t}\}^T_{t=1}$.

   (c) (in the case of uniform mixture distribution) conditional on the $\{\sigma^{j-1}_{\eta i,t}, o^{j-1}_{i,t}\}^T_{t=1}$, obtain the mixture states and draw p$(o^j_{i,t}|u^*_{i,t}, \sigma^j_{\eta i,t})$. Update p^j_i after calculating the number of outliers based on the cdf of the mixture normals.

   (d) Based on the new estimates of the SV, draw p$(\omega^{2,j}_{\eta i}|\{\sigma^j_{\eta i,t}\}^T_{t=1})$ from the IG posterior distribution.

5. Repeat the steps until convergence has been reached.

After estimating 7000 draws from the above Gibbs-sampling algorithm, I discard the first 2000 as burn-in draws. The rest 5000 draws of the model parameters and latent variables are used for inference.

3 Implications of outliers: Application to France

In this section, I explore the implications arising from explicit modeling of outliers and conduct a comparative analysis of two distinct approaches by applying these models to the French economy. I begin by describing the datasets closely monitored by market participants. Next, I demonstrate how outlier-augmented models improve the in-sample fit in comparison to the incumbent Dynamic Factor Models (DFMs). Finally, I present the in-sample estimates of various model features, including idiosyncratic volatilities and long-term growth, to facilitate a comparison between the t-distribution and uniform mixture models for incorporating outliers.

3.1 Data

Table 1 describes the dataset which consists of 27 variables in total: 11 soft and 16 hard indicators. Among the latter, three are measured at a quarterly frequency, and the rest

TABLE 1: Indicators Used, France

Type	Variable	Relevance	Delay	Start	Transformation
Quarterly	Real GDP	89	2	Jan-1980	% QoQ
	Real Investment (Gross Fixed Capital Formulation)		2	Jan-1980	% QoQ
	Total volume of hours worked (employees)		2	Jan-1980	% QoQ
Survey (11)	BdF survey: change in output, manufacturing industry		0	Jan-1980	Level
	BdF survey: expected production		0	Jan-1980	Level
	BdF survey: new orders		0	May-1981	Level
	BdF survey: sentiment indicator for manufacturing	51	0	May-1981	Level
	Composite business climate indicator	11	0	Jan-1980	Level
	Manufacturing: order books and demand		0	Jan-1980	Level
	Manufacturing: general outlook	(97)	0	Jan-1980	Level
	Manufacturing: probable trend		0	Jan-1980	Level
	Services: expected activity		0	Jan-1991	Level
	Services: expected demand	(77)	0	Jan-1991	Level
	Consumer confidence	80	0	Jan-1980	Diff
Consumption	Household consumption: manufactured goods	17	1	Jan-1980	% MoM
Output (7)	Industrial production	60	2	Jan-1980	% MoM
	Capacity utilization		1	Jan-1981	Diff
	Retail sales	55	2	Jan-1980	% MoM
	Car registration	90	1	Jan-1980	% MoM
	Building permits		1	Jan-1994	% MoM
	Industrial production: construction	60	2	Jan-1990	% MoM
	Turnover Index: manufacturing	11	2	Jan-1999	% MoM
Labor (3)	Registered unemployment level for France	37	1	Jan-1980	% MoM
	Active job seekers (A,B,C)		1	Jan-1996	% MoM
	New vacancies		1	Jan-1989	% MoM
Trade (2)	Exports: value goods for France	51	2	Jan-1980	% MoM
	Imports: value goods for France	54	2	Jan-1980	% MoM

Note: This table provides the list of variables in the dataset. It also provides the following additional details: the Bloomberg relevance index, the beginning of each series, and the transformation applied to the data. All series are publicly available from the FRED and DBnomics website. % QoQ: quarter-on-quarter changes, % MoM: month-on-month changes, Diff: first differences.

are monthly indicators. The choice of variables is based on the evidence from the literature rather than a data-driven approach. The three quarterly indicators – output, investment, and total hours worked – are known to be strongly procyclical and key indicators of the business cycle. Investment and hours could be informative for capturing the cyclical movements, as they are sensitive to business cycle fluctuations (Stock and Watson, 1999). As the indicator of consumption, I use the monthly household consumption of manufactured goods, which exclude the durable goods spendings. While it lacks the services consumption compared to the quarterly measure in the national accounts, the monthly data have shorter publication lags that may provide a more timely assessment of not only the economic activity but also the trend, as the permanent income hypothesis suggests that taking the GDP and consumption together may help separate the long-run and component from the growth.

Inspired by Cascaldi-Garcia et al. (2021), I select monthly indicators based on the Bloomberg relevance index, among the large number of candidate series available. This index indicates the percentage of Bloomberg users who set an automatic alert for the release of specific data. Specifically, I take the variables whose relevance index is above 50% and the ones classified as the "Market Moving Indicators". The second column in Table 1 reports this index for each indicator in the dataset.[13] It is well known that market participants closely monitor macroeconomic data releases to extract signals on the current state of the economy, since it heavily affects the performance of their asset portfolios.[14] Hence, I believe this is a reasonable starting point to select the monthly indicators.

While in principle I could construct a larger dataset by including broader categories of data and more disaggregated series, Boivin and Ng (2006) found that more data are not always better for factor analysis: a strong correlation in the idiosyncratic components across the disaggregated series within the same category may worsen both the in-sample fit and out-of-sample forecasting performance. Alvarez et al. (2012) reaches a similar conclusion by finding high persistence in either the common factor or the idiosyncratic errors for the large-scale DFMs, concluding that a small-scale dynamic factor model that uses one representative indicator of each category yields better forecasting results.

The choice of monthly indicators in Table 1 reflects the above findings from the literature. The market participants closely monitor variables across many different sectors, such as the labor market (e.g., the unemployment rate), the industrial sector (e.g., the industrial turnover), the construction sector (e.g., the index of production in construction), private consumption (e.g., retail sales and car registrations), and the external sector (e.g., exports and imports of goods). At the same time, most of the hard indicators are the headline indicators for each category: so the economists focus less on the disaggregated data to track the state of the economy. In the end, we end up with a medium-sized panel with representative indicators for each sector.[15] Importantly, I include monthly surveys that represent the perceptions of economic agents about current and future economic prospects from the two major sources: Banque de France and the National Institute of Statistics (INSEE). These "soft" data have a high signal-to-noise ratio and provide timely information to track the economy.

[13]Since the index does not vary dramatically over the sample, I take the relevance index reported in December 2019, which is the last observation before the COVID period.

[14]In fact, Altavilla et al. (2017) reports that asset prices strongly react to data releases when the outcome differs from the expectations of market participants.

[15]Beyond this criterion, I include two extra indicators, i.e. the capacity utilization and building permits, which may lead the existing headline indicators in the industrial and construction sectors.

The dataset spans the period from January 1980 to December 2022, including the recent COVID-19 episode. Since I focus on the ability of this model in a traditional now-casting setting, I first utilize only the pre-COVID sample that ends in December 2019 and then employ the full series up to 2022, in the next section. Moreover, I exclude prices and monetary and financial indicators that are often available daily or weekly. Bańbura et al. (2013) have shown that these types of data do not improve the performance of a now-casting model due to their noisy nature. I also abstain from "alternative" data, such as web searches and text-based measures, at this point and leave these avenues for future research.[16]

3.2 In-sample fit

I start by presenting the in-sample fit across models, with the sample ending in 2019. Here I compare six models: (1) the model with t-distributed outliers, (2) the model with outliers following uniform mixture distribution, (3) the model without outliers but the time-varying long-run growth and stochastic volatilities (SV) are still present (Antolin-Diaz et al., 2017), (4) the benchmark DFM model of (Bańbura and Modugno, 2014), and (5–6) the latter two models with the outliers replaced by missing values ('SV miss' and 'Basic miss').[17] Since they correspond to a simple approach to treat outliers, we employ simpler models by moving right across the columns, until reaching the benchmark.

Here I evaluate the models in two aspects: point and density forecasts. I use two commonly used criteria, the root mean squared error (RMSE) and mean absolute error (MAE) for the former. A lower RMSE or MAE implies that the model generates more accurate point forecasts. Since I adopt the Bayesian approach, density forecasts are also easily obtainable from the DFM models. While there are several measures for density forecast evaluation, I take the one of the most popular metric, the average log score: it assigns a higher value for the model that provides the highest probability to the realizations. I also use another metric, the continuous rank probability score (CRPS), which provides more robustness to outliers and the values that are close but not equal to the realization. In the end, a total of four metrics are used, as shown in the rows of Table 2.

[16] Evidences on the ability of alternative data to improve nowcasting the economy is still mixed: for instance, Aaronson et al. (2022) claim that Google Trends leads to superior real-time forecasts of initial UI claims compared to other models. However, Larson and Sinclair (2022) find the opposite result.

[17] Here I define the outliers as the observations five-interquartile range away from the median.

Table 2: In-sample fit, pre-COVID

Measure/model	Full (t-dist)	Full (uniform)	SV miss	SV	Basic miss	Basic DFM
(a) point forecasting						
RMSE	0.254***	0.277***	0.277***	0.277***	0.362	0.362
MAE	0.211***	0.228***	0.228***	0.229***	0.281	0.282
(b) density forecasting						
Log score	−3.688***	−4.716***	−4.984***	−5.029***	-28.020	-28.103
CRPS	0.172***	0.191***	0.192***	0.192***	0.257*	0.258

Note: This table provides the point-forecasting performance of different models: the DFM model with outliers following the student-t and uniform mixture distribution, respectively; the model with stochastic trend and volatility (Antolin-Diaz et al., 2017); and the basic DFM model (Bańbura and Modugno, 2014). 'SV miss' and 'Basic miss' are the latter two models with the outliers, which are defined as the observations five-interquartile range away from the median, replaced by missing values. Sample: 1980.1 – 2019.9. The asterisks are related to the p-value of the null hypothesis that the basic model performs as well as others against the alternative that the other model performs better, based on the test statistic of Diebold and Mariano (1995). *** is significant at the 1% level, ** at the 5% level, and * at the 10% level.

Results from Table 2 suggest that the model with outliers provide better in-sample fit than the basic DFM. All of them reject the Diebold and Mariano (1995) test of the null hypothesis that the performance of the basic and alternative model is comparable, at the 1% significance level. The student-t and uniform approaches result in roughly 30% and 23% reduction in the RMSE, compared to the benchmark model. Moreover, a simple approach to replace outliers by missing values turns out to be not so effective, except some improvement in terms of density forecasting: performance of the 'basic miss' model is not better than the benchmark except for the CRPS, and when the SV process is explicitly modelled, the model with and without replacement display identical performance for all metrics except the log score. Finally, the t-distribution model provides the best in-sample fit, with the uniform approach being the second best. While the latter is just as good as the model with outliers but the SV present in terms of point forecasting, it still provides a better performance in density forecasting.

3.3 Implications of modeling choices: in-sample estimates

3.3.1 Stochastic Volatility

I begin by displaying the estimates of the SV of innovations to idiosyncratic components, $\sigma_{\eta_{i,t}}$, in Figure 4. Given that the two outlier modeling approaches, and even the standard DFM-SV model without outliers, diverge primarily in this aspect, it serves as an appropriate initial point to gain further insights into the role of outliers and the consequences of modeling choices. Here I select real GDP growth and three monthly indicators that are closely related to the business cycle: industrial production, construction, and unemployment (levels) for France. The full results from all indicators, including quarterly and soft indicators, can be found in the Appendix B.2.

I comment on two observations. First, the t-distributed outliers (top panel) tend to excessively depress the SV process. The differences between the two approaches are clearly visible in Figure 4: the t-distributed outliers lower the level of idiosyncratic volatilities for the all of four indicators, which is not an intended feature of the model. As shown in the Appendix B.2, this phenomenon is not specific to selected variables but all indicators in the dataset, including the quarterly and monthly soft variables. While this approach results in a more persistent SV process, which is consistent with the driftless random walk assumption in Eq (6), the t-distributed outliers flatten out even the smooth series: for instance, the SV process of the industrial production (in black) is already free from strong cyclical variations, but it is unnecessarily suppressed even further. The SV of GDP is almost constant over 40 years, which is still consistent with the assumption in the model but not in line with the established narrative.

On the other hand, those from the uniform mixture distribution (bottom panel) effectively targets only extreme variations. This approach leaves the SV process untouched which it already lacks cyclical variations, resulting in similar dynamics to the baseline model. When the process exhibits strong cyclical variations, however, extreme spikes are removed, as shown in the case of the IP-construction and unemployment. Considering these aspects, the uniform approach seems more capable than the widely used t-distribution models of delivering more model-consistent estimates. The Appendix B.2 shows that this also applies to other monthly and quarterly indicators.

To shed more light on the source of such differences, I present the posterior probabilities of outlier states estimated by both the t-distribution and uniform mixture approaches in Figure 5. These outlier states, drawn from posterior estimates each period, are categorized into three regions: below 2, between 2 and 5, and above 5. These regions correspond to cases of no (or small) outliers, moderate outliers, and large outliers, respectively.[18] The

[18]I have chosen a cutoff value of two to facilitate the comparison between the two modeling choices. Note that while the support for outliers is a positive real line in the student-t case, the uniform approach discretizes it between 1 and 20.

Figure 4: Stochastic volatility: with v. without outliers

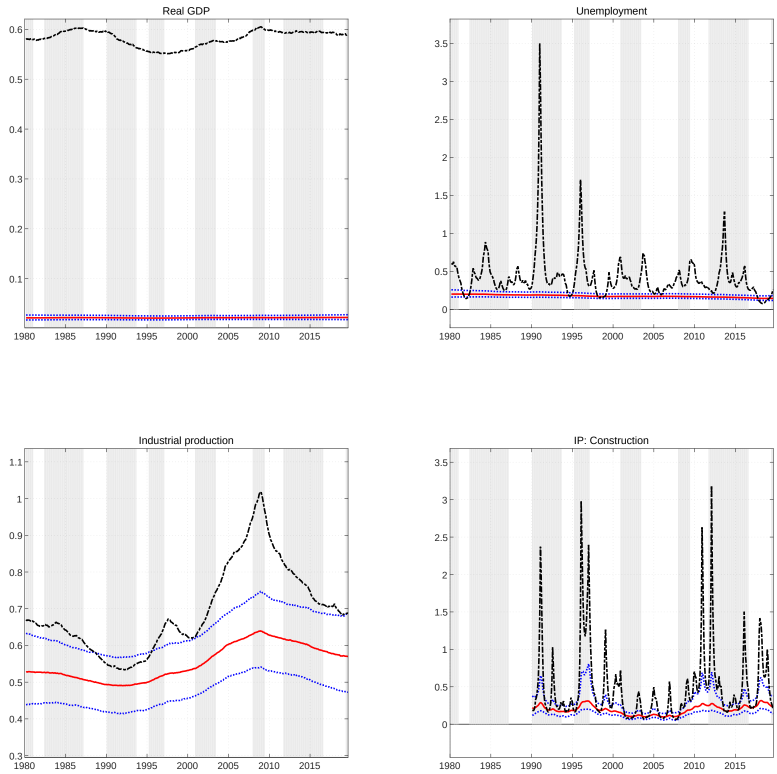

(a) student-t outliers v. no outlier

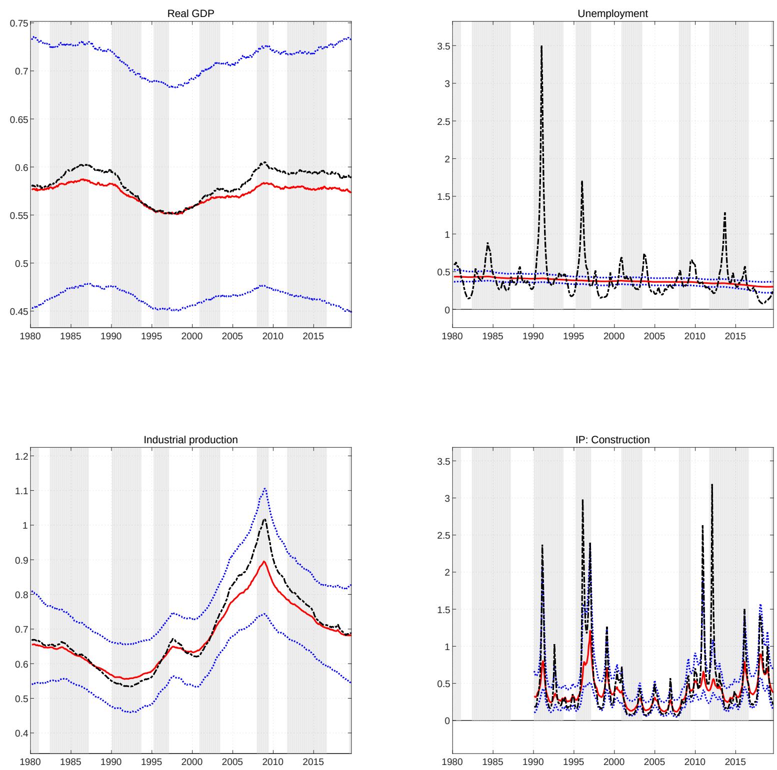

(b) uniform mixture outliers v. no outlier

Note: This plot shows the posterior estimate of stochastic volatility in innovations to the idiosyncratic components. The black line represents the posterior median from the model without outliers. Red and blue lines are the posterior median and 68% bands from the model with outliers following the t-distribution and uniform mixture distribution, respectively. Sample: 1980.1 – 2019.9. OECD recessions in gray.

Figure 5: Posterior probabilities of outlier states

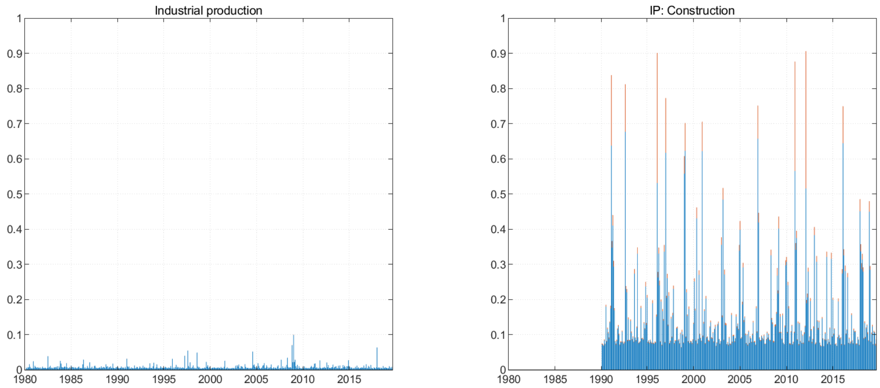

(a) student-t outliers

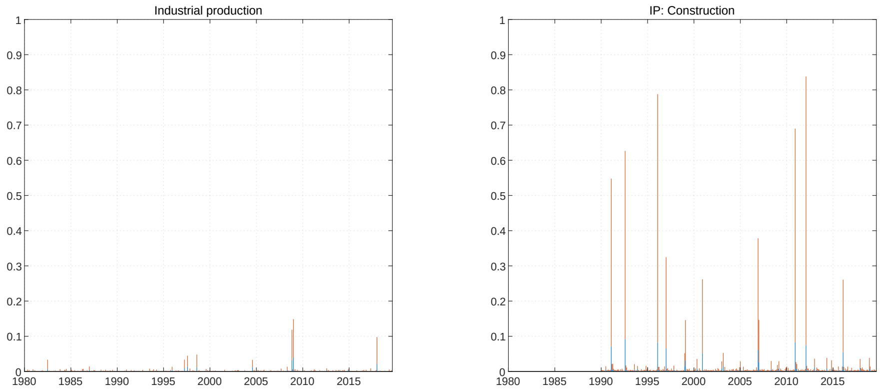

(b) uniform mixture outliers

Note: The stacked bars represent posterior probabilities for realizations of outlier states that are larger than two. The blue bars correspond to the probability that the outliers are within the range between two and five, and the orange bars denote the probability of outlier states that are larger than five to take place. Sample: 1980.1 – 2019.9.

figure directly illustrates the probability of outliers falling within the latter two regions, with the complement representing values below 2 which represent scenarios with no or minor outliers. For this analysis, I have selected two monthly indicators, namely industrial production and construction, which display relatively smooth and volatile processes, respectively. Additional results for other monthly and quarterly indicators, including posterior estimates of outlier states, can be found in Appendix B.3.

Figure 5 clearly illustrates a difference between the two approaches. First, the t-distribution approach tends to produce more moderately sized outliers, and they occur with greater regularity. The model consistently generates moderate-sized outliers in every period, even for the industrial production index, which already lacks high-frequency variations in its stochastic volatility process. In the case of the IP-construction index, the model frequently identifies moderate-to-high outliers across time, often more than 60 out of 100 draws. However, when outliers do occur, they are mostly of moderate size. On the contrary, the uniform approach identifies outlier states less frequently, but when they do occur, they are predominantly large in size. It is noticeable that outliers are drawn more frequently for the IP-construction index, which displays frequent spikes and strong cyclical variations. Industrial production, with its slower-moving SV process, features infrequent outliers over the sample, occurring less than 20 times over the entire period of 477 months. Consequently, while the the t-distribution assumption lowers the level of idiosyncratic volatilities across all indicators, the uniform approach leaves the SV process mostly untouched and target only extreme spikes, for most of the times.

These differences between the two approaches clearly stem from the model specifications detailed in Section 2.2. The key disparity lies in the assumed densities for outlier states. While the density of outlier states peaks around the value of 1 for both methods, the t-distribution concentrates most of the mass on values close to 1. In the case of the uniform approach, there is an equal probability of outlier states between 2 and 10. As a result, the former allocates relatively more mass to values above 1 in total, although these values are centered around 1. The latter, on the other hand, places relatively more mass on the far-right tail. From an empirical standpoint, it translates to the t-distribution approach characterizing outliers as more moderately sized but occurring more frequently, whereas the uniform approach tends to identify large outlier states. In the context of macroeconomic data, outliers are typically perceived as infrequent occurrences with a substantial impact when they do occur. Therefore, the latter approach seems more aligned with our notion.

3.3.2 Long-run growth

In the previous subsection, we noted that the t-distribution approach tends to generate moderate-size outliers too frequently due to its concentration of mass around values close to 1. This behavior may lead to unwanted characteristics in volatility estimation. However, is this the only side effect? To address this question, we now explore additional model estimates.

Figure 6 plots posterior estimates of the time-varying component in long-run real GDP growth, a_t, from the model. Here I show the estimates from the two different approaches

Figure 6: Posterior estimate of trend GDP

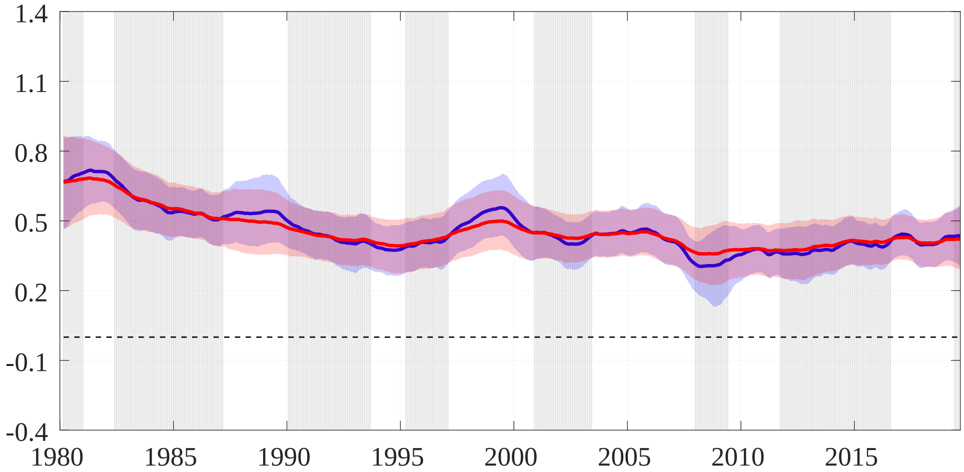

(a) student-t outliers v. no outlier

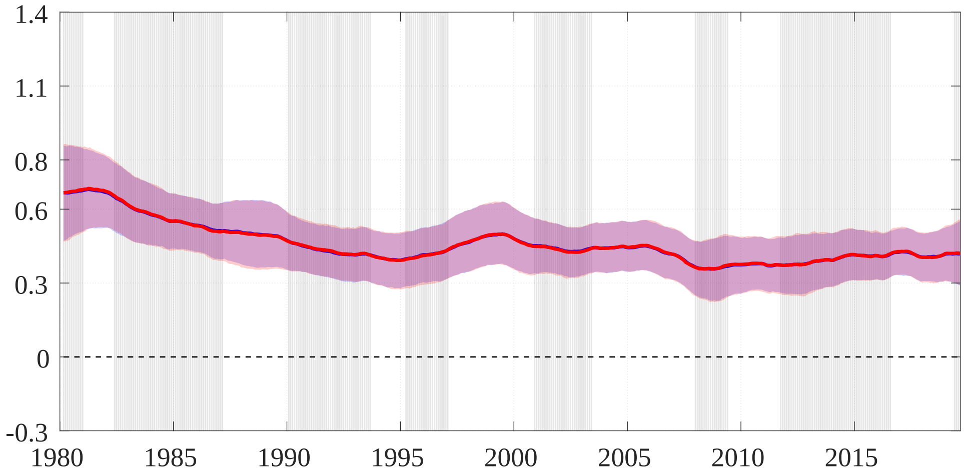

(b) uniform mixture outliers v. no outlier

Note: This plot shows the posterior estimate of the long-run growth from the two models, the one with (in blue) and the one without (in red) outliers. The solid lines and shaded areas denote the posterior median and 90% bands, respectively. Sample: 1980.1 – 2019.9. OECD recessions in gray.

to model outliers (the student-t and uniform mixture) and compare each of them to the model without outliers. Our prior is that they are likely to be identical: the trend is modeled as driftless random walks as in Eq(4), and the idiosyncratic components lack stochastic volatility by the definition of long-run shift. Since no additional components related to the trend were imposed in the model with outliers, we anticipate that the estimates are unlikely to vary. However, it is interesting to note that the way outliers are modeled results in slight differences in the trend estimates. In the case of t-distributed outliers (top panel), trend estimates are less smooth, displaying more high-frequency variations. There is also a spurious IT boom in the late 1990s peaking around 2000, which is true for the U.S.economy but not France. These are less desirable features given the model assumptions and the established narrative in the literature.

On the other hand, when outliers are drawn from the uniform distribution (the bottom panel), the posterior estimates of the long-run growth from the two models, the one with (in blue) and the one without (in red) outliers, are almost identical. The estimated trend captures low-frequency, slow-moving components of growth: it shows the gradual slowdown until 1995, slight rise from 1995 to 2000, and another prolonged decline until the Great Recession. I also plot the posterior estimates of the trend and real GDP growth together in the Appendix B.1. This result is consistent with the narrative of Fernald (2014) that the bulk of the US economy slowdown took place between the turn of the century and the Great recession, as the IT boom faded. Note, however, the size of the IT boom seems much more modest in the case of France, in line with observations from Gordon (2004) that the level of productivity in Europe has been falling behind the US since 1995, due to regulatory barriers in ICT-using industries like wholesale and retail trade and in securities trading.

In Appendix B.1, I also provide posterior estimates of the volatility of the common factor, denoted as f_t, from both models. Despite our assumption that the volatilities of innovations in the factor and idiosyncratic components are independent, the estimates from both models share common dynamics but show slight differences in levels. Similar to what we observed with long-term growth, such differences are more pronounced in the t-distribution model. In summary, the uniform approach appears to be a more robust choice in terms of estimating parameters that may offer structural insights.

4 Out-of-sample forecasting exercises

To replicate the situation of policymakers and market participants, I conduct out-of-sample forecasting analyses with the data available at each point in time. Specifically, using data going back to 1980, I produce nowcasts of French GDP growth from 2005 and evaluate the point forecasting performance of the models relative to the benchmark. Since the historical data vintages are not readily available for France, I run the exercise in pseudo-real time, i.e. based on the recent vintages but taking account of the publication delays. I adopt an expanding out-of-sample window: after January 2005, the model is re-estimated when the new data are released. Since repeating the estimation for hundreds of periods takes huge computational costs, I rely on modern cloud computing facilities for this exercise.[19] Then, I run similar exercises with an extended sample up to 2022, which includes the COVID-19 episode.

4.1 Out-of-sample forecasting: pre-COVID

Figure 7 shows the pseudo real-time nowcasts of French GDP from 2005 to 2019, the pre-COVID period. On average, the nowcasts from the model with outliers (in blue) track the economic development over the last 15 years, including the Great Recession and the Euro Area debt crisis. The model with outliers perform better than the benchmark (in red), the basic DFM model, in two aspects. First, they capture the peaks and troughs during the Great recession more timely and accurately, while the benchmark model underestimates the severity of the recession and the rebound afterwards. Second, while the estimates from the benchmark model are overly persistent after 2012 with an upward bias, those from the model with outliers display more cyclical variations. Between two approaches to model the outliers, the student-t outliers are slightly better in capturing the peaks and troughs during the Great recession, while the uniform approach relatively underestimates the fall in 2009Q1 and the subsequent recoveries. The former, however, also has a drawback: it creates spurious cycles at the end of the sample. Overall, explicitly modeling outliers provide a more accurate picture of the current state of the economy, particularly during the extreme events.

To assess the forecasting ability of models with outliers, I conduct a more formal forecast evaluation in Table 3. As in the Section 3.2, here I compare six models: (1) the model with t-distributed outliers, (2) the model with outliers following uniform mixture distribution, (3) the model without outliers but the time-varying long-run growth and SV are still present (Antolin-Diaz et al., 2017), (4) the benchmark DFM model of (Bańbura and Modugno, 2014), and (5–6) the latter two models with the outliers replaced by missing values ('SV miss' and 'Basic miss'). Models become simpler as we move right across the columns, reaching closer to the benchmark model.

[19] It takes around 40 minutes to complete the estimation for one period on a computer with an Intel i5 processor and 8GB of RAM. I thank Stephane Brice and the IT team at PSE for excellent technical support. One can also employ a more popular alternative channel, e.g. the Amazon Elastic Compute Cloud (EC2) or Microsoft Azure.

Figure 7: Pseudo real-time nowcasts of French GDP, 2005 – 2019

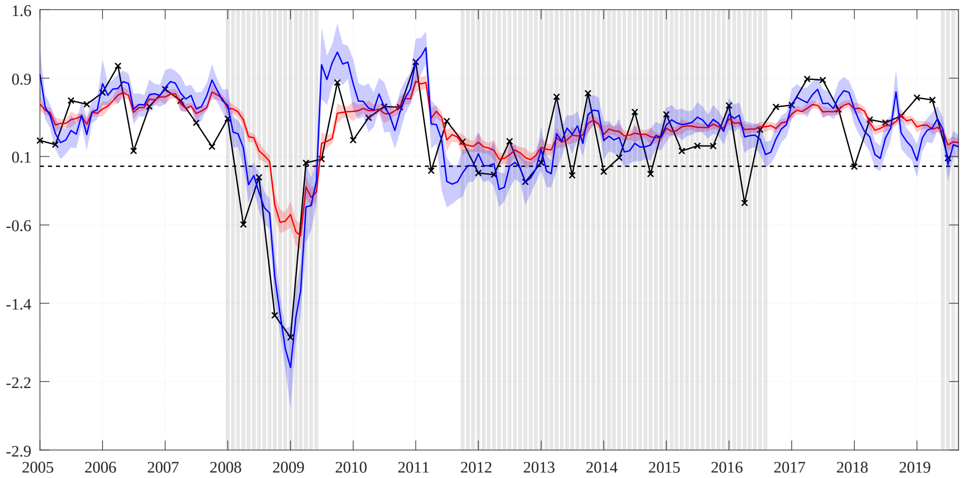

(a) student-t outliers v. benchmark DFM

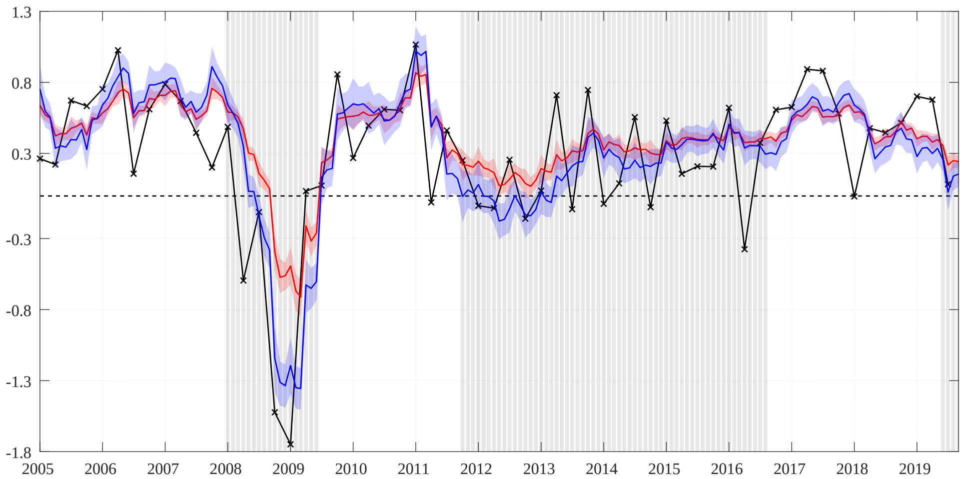

(b) uniform mixture outliers v. benchmark DFM

Note: This plot shows the realizations and pseudo real-time nowcasts of French GDP from 2005 to 2019 constructed from the two models, the one with outliers(in blue) and the benchmark DFM model (in red). The black line represents the actual realizations of GDP growth, and the blue and red solid lines and shaded areas denote the posterior median and 68% bands, respectively. OECD recessions in gray.

I evaluate the models in terms of point and density forecasts, and among many measures available, I use two commonly used metrics for each: the root mean squared error (RMSE) and the continuous ranked probability score (CRPS).[20] Here I present the values in relative terms compared to the benchmark basic DFM model, except the benchmark itself in the last column. The rows in Table 3 correspond to forecasting horizons: I set the last month of the reference quarter as month 0 and start producing forecasts from the first month of the last quarter. In other words, the first three rows are the one-quarter

[20]Appendix C.1 reports the results based on all four measures, including the mean absolute forecast error (MAFE) and log score, in absolute terms.

Table 3: Out-of-sample performance, pre-COVID

Horizon/model	Full (t-dist)	Full (uniform)	SV miss	SV	Basic miss	Basic DFM
(a) point forecasting: relative RMSE						RMSE
-5 month	0.986	0.951	0.949	0.952	0.998	0.484
-4 month	0.961	0.957	0.950	0.953	0.999	0.459
-3 month	0.959	0.928	0.935	0.929	1.001	0.422
-2 month	0.945	0.879	0.877	0.880	0.996	0.374
-1 month	0.881	0.877	0.877	0.876	0.993	0.347
0 month (end of reference Q)	0.900	0.873	0.873	0.874	0.995	0.340
1 month	0.899	0.865	0.872	0.866	0.987	0.338
(b) density forecasting: relative CRPS						CRPS
-5 month	0.987	0.946	0.945	0.950	0.994	0.268
-4 month	0.962	0.955	0.952	0.952	1.002	0.259
-3 month	0.948	0.922	0.931	0.925	1.002	0.249
-2 month	0.883	0.850	0.850	0.853	1.000	0.235
-1 month	0.802	0.829	0.830	0.829	1.002	0.233
0 month (end of reference Q)	0.818	0.826	0.829	0.832	1.001	0.234
1 month	0.822	0.817	0.827	0.822	0.988	0.233

Note: This table provides the forecasting performance of different models: the DFM model with outliers following the student-t and the uniform mixture distribution, respectively; the model with stochastic trend and volatility (Antolin-Diaz et al., 2017), compared to the basic DFM model (Bańbura and Modugno, 2014). 'SV miss' and 'Basic miss' are the latter two models with the outliers, which are defined as the observations five-interquartile range away from the median, replaced by missing values. Using data going back to 1980, I estimate the model at each point in time for 2005.1 – 2019.9 in an expanding window.

ahead predictions, the next three rows are "nowcasts" of the current quarter, and the last one is the forecasts on GDP of the last quarter, or "backcasts". Since the first release of GDP takes place 40 days after the end of the reference period in France, the backcast horizon covers only one month.

I present three takeaways from the results in Table 3. First, the outlier-augmented models, regardless of the distributional assumptions, outperform the benchmark model in terms of point and density forecasting. The full model with uniform outliers display lower RMSE and CRPS than the benchmark model in all horizons, and the same applies to the t-distribution approach. While the 'screening' method, which replaces the pre-defined outliers by missing values, yields negligible impact, the SV model without outliers also dominate the basic DFM. Such results highlight the importance of modeling stochastic volatility in idiosyncratic components: as the sample includes stable expansions (Great Moderation) and sudden depressions (Great Recession), the constant volatility assumption is not enough to track the state of the economy.

Second, improvements are particularly remarkable at the nowcasting horizons. For instance, at the end of the reference quarter, the SV model with uniform outliers improves on the RMSE of the benchmark specification by 13%, and this is the best model in terms of point forecasts. The student-t approach yields a 20% reduction in CRPS relative to the basic DFM, also outperforming all other modeling choices: the model with outliers outperform the benchmark and other SV specifications in point and density nowcasts. The last column and Appendix C.1 show that all models produce more accurate point forecasts as more information become available: the RMSE steadily decreases as the forecasting horizon approaches the end of the quarter, across all columns except for the t-distribution model.

Finally, the explicit modeling of outliers appears to offer marginally more benefit in terms of density than point forecasts. The relative reduction in CRPS becomes noticeably larger than that of RMSE after the two month ahead horizon. Employing an alternative measure, the log score, tells us a similar story as shown in the Appendix C.1.[21] The outlier-augmented models also yield consistently better density forecasts than other SV specifications in all horizons. Interestingly, the uniform approach outperforms the student-t method for point forecasting, but the situation reverses when it comes to density forecasting. Overall, the results highlight a strong forecasting performance of Bayesian DFM models with outliers compared to the benchmark in terms of point and density forecasts for the French GDP.

4.2 Out-of-sample forecasting: post-COVID

I move on to conduct similar pseudo real-time exercises with an extended sample up to September 2022, which includes the COVID episode. Figure 8 shows the pseudo-real time nowcasts of French GDP from the third quarter of 2019 to the end of 2022 with the ex-post realizations (in black) and the major events, e.g. the first nationwide lockdown and subsequent waves. It is clearly visible that the model with outliers (in blue) better tracks the actual dynamics during the year 2020. Nowcasts from the benchmark model (in red) decline only in June, and the degree of the drop and recovery is modest in both cases. The model with student-t or uniform outliers expects a large fall since April 2020,

[21] I abstain from conducting the Diebold and Mariano (1995) test for the out-of-sample analyses in this paper, since Diebold (2015) emphasizes that this test is for comparing forecasts, not models, and the optimal model comparison is based on full-sample residuals, not out-of-sample forecast errors. For this reason, I present the test results only in the Appendix, just for comparative predictive performance.

Figure 8: Pseudo real-time nowcasts of French GDP, 2020

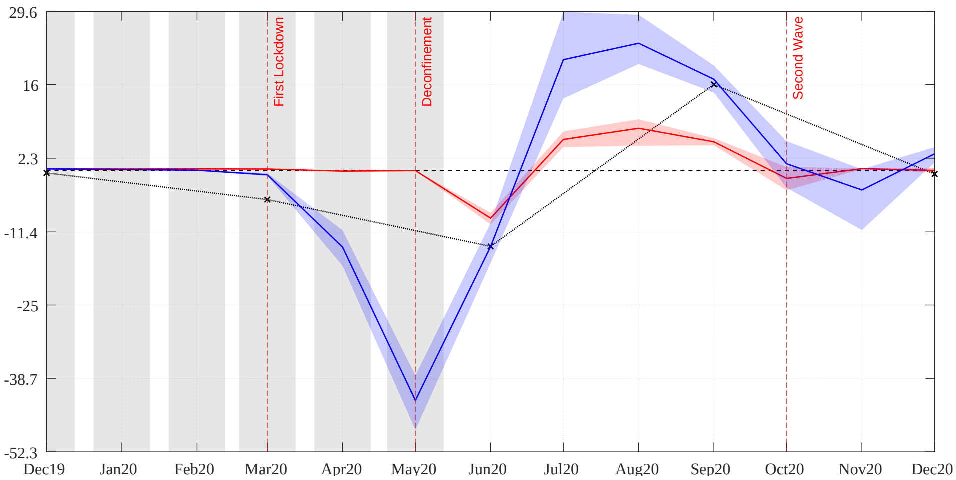

(a) student-t outliers v. benchmark DFM

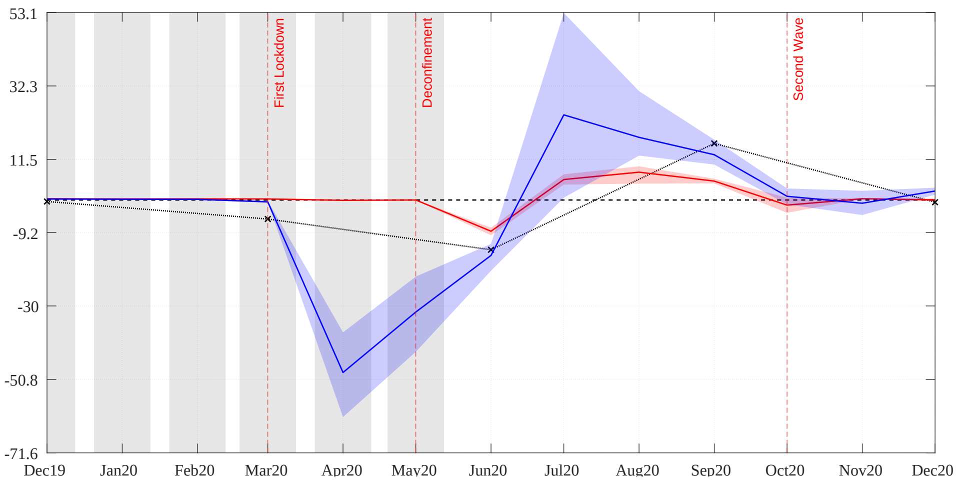

(b) uniform outliers v. benchmark DFM

Note: This plot shows the realizations and pseudo real-time nowcasts of French GDP for 2019.9 – 2020.12 constructed from the two models, the one with outliers(in blue) and the benchmark DFM model (in red). The black line represents the actual realizations of GDP growth, and the blue and red solid lines and shaded areas denote the posterior median and 68% bands, respectively. OECD recessions in gray.

which is a sensible timing given the publication delay. They assign troughs in April and May 2020, respectively, and correctly capture the realization at the second quarter of 2020. The timing of rebound is also consistent with the actual monthly data during this period. Moreover, while the benchmark model reverts to the trend too early, the outlier-augmented models adjust the timing of slowdown to have the actual growth at the third quarter in the ballpark.

In conjunction with the preceding section, I assess the precision of point and density forecasts across various models in comparison to the benchmark in Table 4. Just as in

Table 4: Out-of-sample performance, full

Horizon/model	Full (t-dist)	Full (uniform)	SV miss	SV	Basic miss	Basic DFM
(a) point forecasting: relative RMSE						RMSE
-5 month	1.206	1.899	1.916	1.328	1.040	2.766
-4 month	2.452	1.595	2.342	2.320	1.684	2.722
-3 month	0.682	0.733	1.633	0.956	1.014	2.827
-2 month	0.483	2.027	4.375	0.635	1.199	2.214
-1 month	1.731	1.103	2.162	1.632	1.025	2.133
0 month (end of reference Q)	0.558	0.562	2.324	0.572	2.002	1.626
1 month	0.423	0.570	1.689	0.785	1.451	1.930
(b) density forecasting: relative CRPS						CRPS
-5 month	1.092	1.220	1.252	1.181	1.005	0.835
-4 month	1.629	1.196	1.610	1.482	1.187	0.804
-3 month	0.822	0.792	1.197	0.932	1.014	0.785
-2 month	0.590	1.248	2.327	0.726	1.145	0.673
-1 month	1.187	0.804	1.553	1.301	1.015	0.648
0 month (end of reference Q)	0.662	0.672	1.570	0.691	1.503	0.557
1 month	0.530	0.616	1.236	0.828	1.228	0.643

Note: This table provides the forecasting performance of different models: the DFM model with outliers following the student-t and uniform mixture distribution, respectively; the model with stochastic trend and volatility (Antolin-Diaz et al., 2017), compare to the basic DFM model (Bańbura and Modugno, 2014). 'SV miss' and 'Basic miss' are the latter two models with the outliers, which are defined as the observations five-interquartile range away from the median, replaced by missing values. Using data going back to 1980, I estimate the model at each point in time for 2005.1 – 2022.9 in an expanding window.

the pre-COVID case, improvements are particularly noticeable in the nowcasting horizon for both point and density forecasting. At the end of the reference quarter, it is evident that the SV model with t-distributed outliers outperforms the benchmark and other SV specifications in point and density nowcasts.

Interestingly, the results indicate that the advantages of outlier-augmented models in nowcasting become even more pronounced when the post-COVID sample is included. For instance, the SV model with uniform outliers improves the RMSE of the benchmark specification by 44%, which is more than a threefold increase compared to the 13% improvement observed in the pre-COVID case. The difference in CRPS between the student-t approach and the basic DFM amounts to 0.34, whereas it was only 0.17 in the pre-COVID samples. The improvements over other SV specifications have also expanded. During the pandemic, outlier-augmented models clearly distinguish between short-term and long-lasting increases in uncertainty, an advantage over other models. In contrast to the pre-COVID case, the t-distribution method outperforms the uniform approach for both point and density forecasting. However, the uniform approach still demonstrates comparable capability, dominating models without outliers by a wide margin. The ability of outlier-augmented models to capture extreme events appears particularly advantageous during uncertain times.

In contrast to the pre-COVID case presented in Table 3, models with outliers no longer possess superiority over the benchmark in longer forecasting horizons. This outcome is not surprising, given that the emergence of COVID-19 represented a wholly unexpected event from an economic perspective. Still, when considering the log score as a measure of density forecasting accuracy, as illustrated in Appendix C.1, models incorporating outliers, regardless of their distributional assumptions, consistently generate more precise density forecasts compared to the standard model across all forecasting horizons. In summary, the findings suggest that explicitly modeling outliers significantly enhances forecasting performance, particularly in nowcasting horizons and in terms of density forecasting for both pre and post-COVID datasets.

5 Extension: a hybrid model

So far, we have demonstrated that the outlier-augmented model consistently outperforms conventional models in terms of forecasting accuracy, particularly during periods of heightened uncertainty. While I have also compared the performance across two different outlier models, it may be too restrictive to assume that outliers follow a single distribution. Therefore, I propose a modest extension that accommodates multiple types of outliers simultaneously.

We have introduced outliers multiplicatively within a stochastic volatility (SV) framework, considering two cases where outliers follow either a t-distribution or a uniform mixture distribution. As shown in Section 3, the t-distribution generates frequent small-to-medium outliers, while the uniform mixture model captures rare but large outliers. The distinct outcomes arise from the distributional properties: the heavy tails of the t-distribution capture more frequent moderate deviations from the mean, leading the model to identify a higher number of outliers. Although the tails are heavy, they are not infinite, which limits the most extreme values and results mostly in small-to-medium outliers. In contrast, the uniform mixture model assumes outliers are rare but large, resulting in more conservative forecasts with fewer extreme values.

This does not imply the superiority of a specific model– rather, each has its drawbacks. The uniform mixture approach relies heavily on the prior specification of the outlier probability p_i. When miscalibrated, it may fail to detect true outliers or misclassify regular observations. Moreover, this model may overlook the more nuanced variability that the continuous distribution captures. On the other hand, while the t-distribution is more flexible, it also introduces greater unpredictability by over-identifying moderate outliers, potentially obscuring true rare events.

Since no single model is clearly dominant, I propose an approach that accommodates both types of outliers. Previously, I demonstrated that the t-distribution tends to generate small to medium-sized outliers more frequently, while the uniform mixture captures rare but large outliers. In reality, both types of outliers—small, frequent deviations and large, rare shocks—may coexist. Therefore, instead of relying on a single distributional assumption, I propose a 'hybrid' model, such that:

$$(1 - \rho_i(L))u_{it} = \sigma_{\eta_{it}} o_{it}^{med} o_{it}^{big} \eta_{it}, \quad \eta_{it} \sim N(0, 1) \tag{14}$$

where each of $o_{it}^{med} o_{it}^{big}$ follows the t-distribution, i.e. $\nu_i/o_{i,t} \sim \chi^2_{\nu_i}$, and uniform mixture distribution described in the Eq(7), respectively. To assess whether this hybrid model improves forecasting performance, I repeat the out-of-sample exercises from Section 4 and compare the results to competing models.

Table 5 compares the relative accuracy of point and density forecasts across six models, on top of the hybrid model presented in the first column. As shown in the top panel, there is little to no improvement in forecasting accuracy for both point and density forecasts, for the pre-COVID sample. In terms of point forecasts, there is no comparative gain over other models across any horizon, and in fact, the hybrid model performs the worst among the first five models, excluding the basic DFMs. Specifically, at the 2- and 5-month horizons, its performance is even worse than the benchmark DFM.

The density forecasting results tell a slightly different story. While the hybrid model still underperforms the more advanced models at the one-quarter horizon, it shows modest improvement for the one-month-ahead forecasts, where it outperforms all other models. However, these gains are minimal. The conclusion from the pre-COVID sample is that, because the data predominantly consists of stable periods, the added estimation uncertainty from incorporating a second type of outlier outweighs the potential benefits of

TABLE 5: Out-of-sample performance, with the hybrid model

Horizon/model	t+uniform	Full (t-dist)	Full (uniform)	SV miss	SV	Basic miss	Basic DFM
(a) point forecasting: relative RMSE							RMSE
-5 month	1.015	0.986	0.951	0.949	0.952	0.998	0.484
-4 month	0.973	0.961	0.957	0.950	0.953	0.999	0.459
-3 month	0.958	0.959	0.928	0.935	0.929	1.001	0.422
-2 month	1.000	0.945	0.879	0.877	0.880	0.996	0.374
-1 month	0.895	0.881	0.877	0.877	0.876	0.993	0.347
0 month (end of reference Q)	0.905	0.900	0.873	0.873	0.874	0.995	0.340
1 month	0.889	0.899	0.865	0.872	0.866	0.987	0.338
(b) density forecasting: relative CRPS						CRPS	
-5 month	1.009	0.987	0.946	0.945	0.950	0.994	0.268
-4 month	0.971	0.962	0.955	0.952	0.952	1.002	0.259
-3 month	0.955	0.948	0.922	0.931	0.925	1.002	0.249
-2 month	0.922	0.883	0.850	0.850	0.853	1.000	0.235
-1 month	0.800	0.802	0.829	0.830	0.829	1.002	0.233
0 month (end of reference Q)	0.816	0.818	0.826	0.829	0.832	1.001	0.234
1 month	0.806	0.822	0.817	0.827	0.822	0.988	0.233

(a) Pre-COVID

Horizon/model	t+uniform	Full (t-dist)	Full (uniform)	SV miss	SV	Basic miss	Basic DFM
(a) point forecasting: relative RMSE							RMSE
-5 month	1.046	1.206	1.899	1.916	1.328	1.040	2.766
-4 month	1.251	2.452	1.595	2.342	2.320	1.684	2.722
-3 month	0.564	0.682	0.733	1.633	0.956	1.014	2.827
-2 month	0.421	0.483	2.027	4.375	0.635	1.199	2.214
-1 month	0.967	1.731	1.103	2.162	1.632	1.025	2.133
0 month (end of reference Q)	0.435	0.558	0.562	2.324	0.572	2.002	1.626
1 month	0.404	0.423	0.570	1.689	0.785	1.451	1.930
(b) density forecasting: relative CRPS						CRPS	
-5 month	1.002	1.092	1.220	1.252	1.181	1.005	0.835
-4 month	1.051	1.629	1.196	1.610	1.482	1.187	0.804
-3 month	0.768	0.822	0.792	1.197	0.932	1.014	0.785
-2 month	0.582	0.590	1.248	2.327	0.726	1.145	0.673
-1 month	0.758	1.187	0.804	1.553	1.301	1.015	0.648
0 month (end of reference Q)	0.574	0.662	0.672	1.570	0.691	1.503	0.557
1 month	0.514	0.530	0.616	1.236	0.828	1.228	0.643

(b) Full sample

Note: This table provides the forecasting performance of different models: the DFM model with both student-t and uniform mixture outliers; the model with outliers following the student-t and uniform mixture distribution, respectively; the model with stochastic trend and volatility (Antolin-Diaz et al., 2017), compare to the basic DFM model (Bańbura and Modugno, 2014). 'SV miss' and 'Basic miss' are the latter two models with the outliers, which are defined as the observations five-interquartile range away from the median, replaced by missing values. Using data going back to 1980, I estimate the model at each point in time for 2005.1 – 2019.9 for the top and 2022.9 for the bottom panel, respectively, in an expanding window.

correcting model misspecification.

The situation reverses when the COVID-19 period is included, as illustrated in the bottom panel of Table 5. The hybrid model now outperforms all others at nearly all forecasting horizons, except for the two longest ones. The improvements are notable: at the nowcasting horizon, the hybrid model delivers more than a 56% and 43% gain in point and density forecasting accuracy, respectively, compared to the benchmark model.

It achieves more than a 12% improvement in point forecasts over the t-distribution model, which had previously shown the best performance.

Indeed, such a reversal is anticipated rather than unexpected. The COVID period can be characterized by the increased economic volatility where both small, frequent shocks and large, rare shocks occurring simultaneously. The hybrid model, which accommodates both t-distributed (moderate) and uniform-distributed (large) outliers, is particularly well-suited to capture the complexities of this environment. By incorporating both types of outliers, the model is able to more accurately represent the distribution of shocks, leading to improved forecasting performance. In contrast, models relying solely on t-distributed or uniform-distributed outliers are less capable of capturing the full range of dynamics during the COVID-19. The reduction in misspecification outweighs the added estimation uncertainty introduced by the hybrid approach, in turbulent times.

How does such a better forecasting performance translates into nowcasting? Figure 9 shows the pseudo-real time nowcasts of French GDP from the third quarter of 2019 to the same quarter of 2022 with the ex-post realizations (in black). When compared to the previous outlier-augmented models, the nowcasts from the hybrid model (in blue) are better in two aspects. First, it produces better point nowcasts, especially in the longer horizon. While the hybrid model captures the actual real GDP in the second, third, and fourth quarter of 2020, nowcasts from the uniform mixture model (in red) in the bottom panel of Figure 9 misses the actual value in the third quarter. The t-distribution model (in red) in the top panel also gets the actual real GDP in the second and third quarter of 2020 correctly, it slightly misses the actual value in 2020 Q4. Moreover, the nowcasts from the hybrid model also works pretty well for one and two month ahead nowcasts, which is a longer horizon. There is less swings in the middle, for instance one-month ahead nowcasts made in May and August 2020, for the second and third quarter real gdp growth, respectively.

How does this improved forecasting performance translate into the nowcasts for 2020? Figure 9 presents the pseudo-real-time nowcasts of French GDP from the third quarter of 2019 to the end of 2022, alongside the ex-post realisations (in black). Compared to the previous outlier-augmented models, the nowcasts from the hybrid model (in blue) show improvements in two aspects. First, the hybrid model produces more accurate point nowcasts, particularly over longer horizons. For instance, the hybrid model accurately

Figure 9: Hybrid v. outlier models: nowcasts of French GDP, 2020

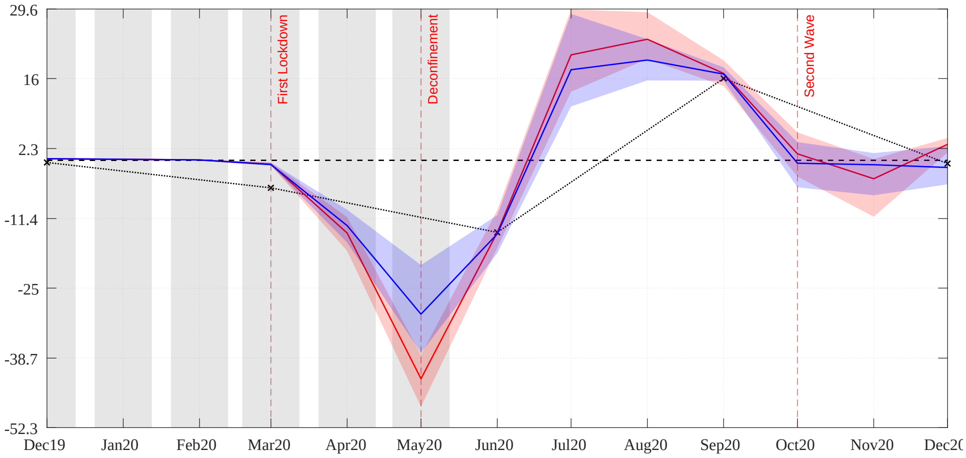

(a) Hybrid model v. student-t outliers

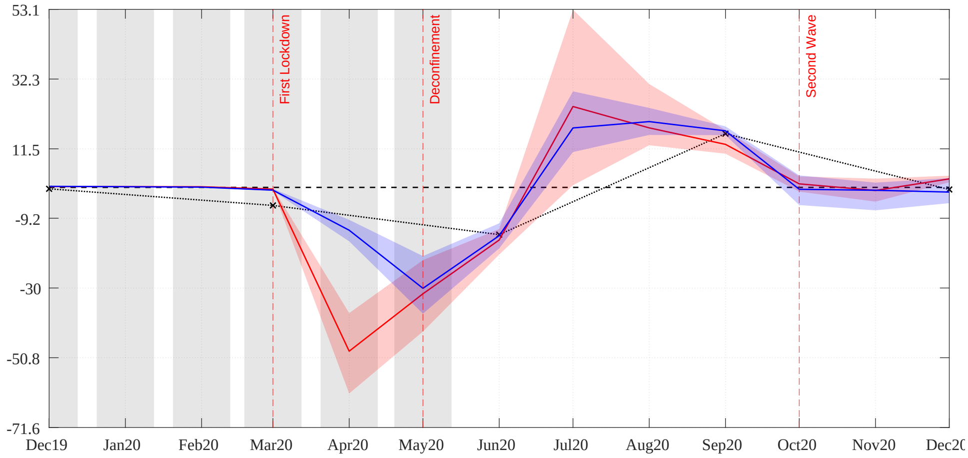

(b) Hybrid model v. uniform outliers

Note: This plot shows the realizations and pseudo real-time nowcasts of French GDP for 2019.9 – 2020.12 constructed from the two models, the hybrid model(in blue) and the model with t-distributed or uniform outliers (in red). The black line represents the actual realizations of GDP growth, and the blue and red solid lines and shaded areas denote the posterior median and 68% bands, respectively. OECD recessions in gray.

captures real GDP in the second, third, and fourth quarters of 2020, whereas the nowcasts from the uniform mixture model (in red), shown in the bottom panel, miss the actual value for the third quarter. Although the t-distribution model (in red in the top panel) correctly captures real GDP in the 2020 Q2 and Q3, it slightly underestimates the actual value in the fourth quarter. Additionally, the hybrid model performs well for nowcasts made at longer horizons, one and two months ahead. There are fewer fluctuations in the middle of the year: one-month-ahead nowcasts for 2020 Q2 and Q3, made in May and

August respectively, illustrate more stability.

Second, the hybrid model demonstrates more accurate density forecasts. It effectively captures the period of heightened uncertainty during the early phase of the pandemic and the summer resurgence, when the number of COVID-19 cases increased again. Since this uncertainty is related to COVID-19, it manifests in the early phase of the pandemic but dissipates quickly by the summer 2020. A similar pattern is observed in winter: while concerns were high, the economic impact was turned out to be less severe. Thus, a moderate rise in uncertainty in October followed by stabilisation reflects the reality.

In contrast, the t-distribution model (in red, top panel), which only partially accounts for large outliers, tends to produce overly persistent volatility when large movements take place in the data. This is reflected in more prolonged swings in summer and winter 2020, leading to larger upward and downward deviations in nowcasts made mid-quarter. Moreover, it fails to capture the significant rise in uncertainty during the 2020 Q2, which is less realistic given that policymakers and markets at the time were focused on increasing uncertainty in the economic outlook. The uniform mixture model, in the bottom panel, better captures the evolution of economic uncertainty compared to the t-distribution model and more in line with the hybrid model. Still, it generally overestimates the degree of uncertainty, as indicated by overly wide confidence intervals. Overall, the hybrid model provides a closer representation of actual developments in both point and density forecasting during this uncertain times.

In sum, the significant improvement in performance of the hybrid model during the COVID-19 stems from its ability to capture both frequent moderate shocks and rare large shocks, which characterized that time. The added complexity may turn into a disadvantage during more stable periods, possibly causing overfitting and particularly weaker performance in point forecasts.

6 Conclusion

This paper aims to advance the dynamic factor model framework in the presence of one-off outliers in the data. While it is of particular interest during unprecedented times like the COVID-19 pandemic, the occurrence of outliers is a common feature in macroeconomic data. Moving beyond conventional data screening practices, I propose an explicit modeling of outliers without unintended biases present in comparable approaches. The methodological contributions of this paper are twofold: firstly, it introduces the incorporation of fat tails and outliers multiplicatively into innovation volatility. Secondly, it investigates two distinct approaches including the student-t and uniform mixture distributed outliers.

The results of applying these models to forecast the French economy before and after the pandemic have yielded several key findings. First, the outlier-augmented models consistently outperform the benchmark model, regardless of distributional assumptions, in point and density forecasting. These enhancements were most pronounced in nowcasting horizons, where models with outliers demonstrated superior performance, capturing economic activity more accurately and in a more timely manner.

Secondly, the significance of incorporating outliers becomes more apparent during major crises, such as the Great Recession and the COVID-19. Outlier-augmented models distinctly differentiated between short-term and long-lasting spikes in uncertainty, effectively capturing economic peaks and troughs with greater accuracy and timeliness. Finally, when comparing two distinct methods to incorporate outliers, the uniform approach is possibly a more robust choice over the commonly employed t-distribution models. The former effectively targets only extreme variations, while the latter tends to overly suppress already smooth processes. Despite such drawbacks, the student-t method also offers some benefits in terms of forecasting, leaving the choice to be made based on the data research objectives.

References

Aaronson, Daniel, Scott A. Brave, R. Andrew Butters, Michael Fogarty, Daniel W. Sacks, and Boyoung Seo, "Forecasting unemployment insurance claims in realtime with Google Trends," *International Journal of Forecasting*, 2022, *38* (2), 567–581.

Altavilla, Carlo, Domenico Giannone, and Michele Modugno, "Low frequency effects of macroeconomic news on government bond yields," *Journal of Monetary Economics*, 2017, *92* (C), 31–46.

Alvarez, Rocio, Maximo Camacho, and Gabriel Perez-Quiros, "Finite sample performance of small versus large scale dynamic factor models," Working Papers 1204, Banco de España February 2012.

Antolin-Diaz, Juan, Thomas Drechsel, and Ivan Petrella, "Tracking the Slowdown in Long-Run GDP Growth," *The Review of Economics and Statistics*, May 2017, *99* (2), 343–356.

_ , _ , and _ , "Advances in Nowcasting Economic Activity: The Role of Heterogeneous Dynamics and Fat Tails," CEPR Discussion Papers 17800, C.E.P.R. Discussion Papers January 2023.

Aruoba, S. Boragan, Francis X. Diebold, and Chiara Scotti, "Real-Time Measurement of Business Conditions," *Journal of Business & Economic Statistics*, 2009, *27* (4), 417–427.

Bai, Jushan and Peng Wang, "Identification and Bayesian Estimation of Dynamic Factor Models," *Journal of Business & Economic Statistics*, April 2015, *33* (2), 221–240.

Bańbura, Marta and Michele Modugno, "Maximum Likelihood Estimation Of Factor Models On Datasets With Arbitrary Pattern Of Missing Data," *Journal of Applied Econometrics*, January 2014, *29* (1), 133–160.

_ , Domenico Giannone, and Lucrezia Reichlin, "Nowcasting," Working Paper Series 1275, European Central Bank December 2010.

_ , _ , Michele Modugno, and Lucrezia Reichlin, "Now-casting and the real-time data flow," Working Paper Series 1564, European Central Bank July 2013.

Boivin, Jean and Serena Ng, "Are more data always better for factor analysis?," *Journal of Econometrics*, May 2006, *132* (1), 169–194.

Camacho, Maximo and Gabriel Perez-Quiros, "Introducing the euro-sting: Short-term indicator of euro area growth," *Journal of Applied Econometrics*, 2010, *25* (4), 663–694.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino, "Nowcasting tail risk to economic activity at a weekly frequency," *Journal of Applied Econometrics*, August 2022, *37* (5), 843–866.

Cascaldi-Garcia, Danilo, "Pandemic Priors," International Finance Discussion Papers 1352, Board of Governors of the Federal Reserve System (U.S.) August 2022.

_ , Thiago Revil T. Ferreira, Domenico Giannone, and Michele Modugno, "Back to the Present: Learning about the Euro Area through a Now-casting Model," International Finance Discussion Papers 1313, Board of Governors of the Federal Reserve System (U.S.) March 2021.

Chan, Joshua C.C., "Comparing stochastic volatility specifications for large Bayesian VARs," *Journal of Econometrics*, 2023, *235* (2), 1419–1446.

Cogley, Timothy and Thomas J. Sargent, "Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S," *Review of Economic Dynamics*, April 2005, *8* (2), 262–302.

Diebold, Francis X., "Comparing Predictive Accuracy, Twenty Years Later: A Personal Perspective on the Use and Abuse of Diebold-Mariano Tests," *Journal of Business & Economic Statistics*, January 2015, *33* (1), 1–1.

—, "Real-Time Real Economic Activity: Exiting the Great Recession and Entering the Pandemic Recession," NBER Working Papers 27482, National Bureau of Economic Research, Inc July 2020.

Diebold, Francis X and Roberto S Mariano, "Comparing Predictive Accuracy," *Journal of Business & Economic Statistics*, July 1995, *13* (3), 253–263.

Doz, Catherine, Laurent Ferrara, and Pierre-Alain Pionnier, "Business cycle dynamics after the Great Recession: An Extended Markov-Switching Dynamic Factor Model," PSE Working Papers halshs-02443364, HAL January 2020.

Evans, Martin D. D., "Where Are We Now? Real-Time Estimates of the Macroeconomy," *International Journal of Central Banking*, September 2005, *1* (2).

Fernald, John G., "Productivity and Potential Output before, during, and after the Great Recession," in "NBER Macroeconomics Annual 2014, Volume 29" NBER Chapters, National Bureau of Economic Research, Inc, January 2014, pp. 1–51.

Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, Keith Kuester, and Juan Rubio-Ramírez, "Fiscal Volatility Shocks and Economic Activity," *American Economic Review*, November 2015, *105* (11), 3352–3384.

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

Gordon, Robert J., "Why was Europe Left at the Station When America's Productivity Locomotive Departed?," NBER Working Papers 10661, National Bureau of Economic Research, Inc August 2004.

—, "The Demise of U.S. Economic Growth: Restatement, Rebuttal, and Reflections," NBER Working Papers 19895, National Bureau of Economic Research, Inc February 2014.

Jacquier, Eric, Nicholas G. Polson, and P.E.Peter E. Rossi, "Bayesian analysis of stochastic volatility models with fat-tails and correlated errors," *Journal of Econometrics*, September 2004, *122* (1), 185–212.

Kim, Sangjoon, Neil Shephard, and Siddhartha Chib, "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models," *Review of Economic Studies*, 1998, *65* (3), 361–393.

Larson, William D. and Tara M. Sinclair, "Nowcasting unemployment insurance claims in the time of COVID-19," *International Journal of Forecasting*, 2022, *38* (2), 635–647.

Lenza, Michele and Giorgio E. Primiceri, "How to estimate a VAR after March 2020," Working Paper Series 2461, European Central Bank August 2020.

Lewis, Daniel J., Karel Mertens, James H. Stock, and Mihir Trivedi, "Measuring real activity using a weekly economic index," *Journal of Applied Econometrics*, June 2022, *37* (4), 667–687.

Litterman, Robert B, "Forecasting with Bayesian Vector Autoregressions-Five Years of Experience," *Journal of Business & Economic Statistics*, January 1986, *4* (1), 25–38.

Marcellino, Massimiliano, Mario Porqueddu, and Fabrizio Venditti, "Short-Term GDP Forecasting With a Mixed-Frequency Dynamic Factor Model With Stochastic Volatility," *Journal of Business & Economic Statistics*, January 2016, *34* (1), 118–127.

—, Todd Clark, Andrea Carriero, and Elmar Mertens, "Addressing COVID-19 Outliers in BVARs with Stochastic Volatility," CEPR Discussion Papers 15964, C.E.P.R. Discussion Papers March 2021.

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

Moench, Emanuel, Serena Ng, and Simon Potter, "Dynamic Hierarchical Factor Model," *The Review of Economics and Statistics*, December 2013, *95* (5), 1811–1817.

Negro, Marco Del and Christopher Otrok, "Dynamic factor models with time-varying parameters: measuring changes in international business cycles," Technical Report 2008.

Ng, Serena, "Modeling Macroeconomic Variations after Covid-19," NBER Working Papers 29060, National Bureau of Economic Research, Inc July 2021.

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

Schorfheide, Frank and Dongho Song, "Real-Time Forecasting with a (Standard) Mixed-Frequency VAR During a Pandemic," NBER Working Papers 29535, National Bureau of Economic Research, Inc December 2021.

Sims, Christopher A and Tao Zha, "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, November 1998, *39* (4), 949–968.

Stock, James H. and Mark W. Watson, "New Indexes of Coincident and Leading Economic Indicators," in "NBER Macroeconomics Annual 1989, Volume 4" NBER Chapters, National Bureau of Economic Research, Inc, January 1989, pp. 351–409.

_ and _ , "Business cycle fluctuations in us macroeconomic time series," in J. B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, Vol. 1 of *Handbook of Macroeconomics*, Elsevier, 1999, chapter 1, pp. 3–64.

Stock, James H and Mark W Watson, "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business & Economic Statistics*, April 2002, *20* (2), 147–162.

Stock, James H. and Mark W. Watson, "Core Inflation and Trend Inflation," *The Review of Economics and Statistics*, October 2016, *98* (4), 770–784.

A Details on the estimation of the model

A.1 The state space system

Let y_t denote n × 1 vector of macroeconomic indicators. While it includes n_Q quarterly and n_M monthly variables, I employ the interpolated monthly values $y_t^{Q,M}$ for the quarterly indicators as described in Section 2.4. So I specify in the model that the time index t is always monthly for all variables. The number of lags in the autoregressive coefficients of factors and idiosyncratic components is set to 2, i.e. p = q = 2. Since the idiosyncratic components in the model feature autocorrelation, the state space is rewritten in terms of quasi-differences:

$$
\begin{aligned}
\tilde{y}_t &= \tilde{\Lambda} F_t + \tilde{\eta}_t, \quad \tilde{\eta}_t \sim N(0, R_t) \\
F_t &= A F_{t-1} + e_t, \quad e_t \sim N(0, Q_t)
\end{aligned}
$$

where the observables are defined as:

$$
\tilde{y}_t = \begin{bmatrix} y_{1,t}^{Q,M} - \rho_{1,1} y_{1,t-1}^{Q,M} - \rho_{1,2} y_{1,t-2}^{Q,M} \\ \vdots \\ y_{n_Q,t}^{Q,M} - \rho_{n_Q,1} y_{n_Q,t-1}^{Q,M} - \rho_{n_Q,2} y_{n_Q,t-2}^{Q,M} \\ y_{1,t}^{M} - \rho_{n_Q+1,1} y_{1,t-1}^{M} - \rho_{n_Q+1,2} y_{1,t-2}^{M} \\ \vdots \\ y_{n_M,t}^{M} - \rho_{n,1} y_{n_M,t-1}^{M} - \rho_{n,2} y_{n_M,t-2}^{M} \end{bmatrix}
$$

and the state vector $F_t = [a_t \ a_{t-1} \ a_{t-2} \ f_t \ f_{t-1} \ f_{t-2}]'$. As we extract the trend and factor, the factor loadings consist of two blocks, $\tilde{\Lambda} = [\Lambda_a \ \Lambda_f]$, which are respectively defined as the following:

$$
\Lambda_a = \begin{bmatrix} 1 & -\rho_{1,1} & -\rho_{1,2} \\ 1 & -\rho_{2,1} & -\rho_{2,2} \\ b_c & -b_c\rho_{3,1} & -b_c\rho_{3,2} \\ & 0_{(n-3)\times 3} & \end{bmatrix}, \quad \Lambda_f = \begin{bmatrix} 1 & -\rho_{1,1} & -\rho_{1,2} \\ \Lambda_2 & -\Lambda_2\rho_{2,1} & -\Lambda_2\rho_{2,2} \\ \vdots & \vdots & \vdots \\ \Lambda_n & -\Lambda_n\rho_{2,1} & -\Lambda_n\rho_{n,2} \end{bmatrix}
$$

where b_c loads on the consumption. The matrix of autocorrelation coefficients A also consists of two blocks, $[F_1; \ F_2]$, which are defined as:

$$
F_1 = \begin{bmatrix} 1 & 0_{1\times 2} \\ I_2 & 0_{2\times 1} \end{bmatrix}, \quad F_2 = \begin{bmatrix} \phi_1 & \phi_2 & 0 \\ I_2 & 0_{2\times 1} & \end{bmatrix}
$$

Finally, the innovations to the law of motion can be described as $e_t = [v_{a,t}\ 0_{2\times1}\ \epsilon_t\ 0_{2\times1}]$', which follows the normal distribution with mean zero and the covariance matrix $Q_t = \text{diag}(\omega_a^2$, $0_{1\times2}$, $\sigma^2_{\epsilon,t}$,$0_{1\times2}$). The covariance matrix of the measurement equation R_t, is the diagonal matrix with $\omega^2_{\eta_i,t}$, for i = 1...n.

A.2 The estimation algorithm

Let $\theta = \{\lambda, \Phi, \rho, \omega_a, \omega_\epsilon, \omega_\eta, p, \alpha, \beta\}$ be the underlying parameters of the model, and Φ, ρ represent the autoregressive coefficients for the factor and idiosyncratic components.[22] The latent states to be estimated are $\{a_t, f_t, \sigma_{\epsilon,t}, \sigma_{\eta i,t}, o_{it}\}_{t=1}^T$. All variables are standardized, and the superscript j denotes a current draw. The algorithm consists of the following steps.

Initialization

The model parameters are initialized at the arbitrary starting values for the parameters and the stochastic volatilities, θ^0 and $\{\sigma^0_{\epsilon,t}, \sigma^0_{\eta i,t}, o^0_{it}\}_{t=1}^T$. Specifically, I set $\lambda = 1$, $\Phi = \rho = 0$, $\omega_a = \omega_\epsilon = \omega_\eta = 10^{-5}$, p = 0.1. I set α and β to reflect stylized facts observed by Stock and Watson (2016) that the outlier occurs once every four years in the 10 years of the pre-sample. I also select $\sigma_{\epsilon,t} = 0.1$, $\sigma_{\eta i,t} = 1$ and $o_{it} = 1$ for the latent states. Set j = 1.

1. Construct monthly-interpolated values for quarterly variables

For the quarterly variables i = 1...n_Q, compute $\Delta y^Q_{it} - c_{it} - \lambda_i(L) f_t$, given $a_t^{j-1}, f_t^{j-1}, \lambda^{j-1}$ from the previous iteration. Then, using the state space in Eq (12) and (13) and conditional on $\rho^{j-1}, \{\sigma^{j-1}_{\eta i,t}\}_{t=1}^T$, draw $\hat{u}^M_{i,t}$ by employing the Kalman filter and smoother. Following Bai and Wang (2015), I initialize the filter from a normal distribution, i.e. $x_0 \sim N(0, 10^4 I_5)$, and the covariance matrix of the measurement equation R_t is set to 10^{-4}. Then we obtain the monthly-interpolated values, $\Delta y^{M,Q}_{it} = c_{it} + \lambda_i(L) f_t + \hat{u}^M_{i,t}$. When j = 1, I interpolate quarterly variables based on the cubic splines using the MATLAB command *interp1*. In the case of the out-of-sample analysis, the unknown values are set to NaN.

2. Draw the latent factors and autoregressive coefficients

Based on the previous draws of $\phi^{j-1}, \Lambda^{j-1}, \rho^{j-1}, \omega_a^{j-1}, \{\sigma^{j-1}_{\epsilon,t}, \sigma^{j-1}_{\eta i,t}\}_{t=1}^T$, construct the state space in the quasi-differences as shown in Appendix A.1. Also apply the quasi-differencing to the monthly and monthly-interpolated quarterly indicators obtained from the previous step. Draw the factor and the trend, $p(\{a_t^j, f_t^j\}_{t=1}^T | \theta^{j-1}, \{\sigma^{j-1}_{\epsilon,t}, \sigma^{j-1}_{\eta i,t}\}_{t=1}^T, y)$ using the Kalman filter and smoother, which are initialized from a normal distribution, i.e. $x_0 \sim N(0, 10^4 I)$. Then, conditional on the new draws of $\{a_t^j\}_{t=1}^T$, obtain $v^j_{a,t}$ using Eq (4). Draw the variance of the time-varying GDP growth component, $p(\omega_a^{2,j} | \{a_t^j\}_{t=1}^T)$, from the Inverse-Gamma (IG) posterior with the prior of one degree of freedom and the scale of 0.001.

Given $\{f_t^j\}_{t=1}^T$ from the previous step, draw the autoregressive parameters of the factor

[22] These are the parameters of the model where outliers follow the piecewise-uniform distribution. In the case of the student-t distributed outliers, we estimate parameters ν instead of $\{p, \alpha, \beta\}$.

VAR, Φ^j from the Normal-Inverse Wishart posterior $p(\Phi^j|\{f_t^j, \{\sigma_{\epsilon,t}^{j-1}\}_{t=1}^T)$. Specifically, I run the standard Bayesian VAR routine with the Minnesota-style priors described in Section 2.3 and p = 2. Then, using the residuals from the BVAR, run the Kalman filter and smoother and use the mixture of normals following Kim et al. (1998) to draw the SV component of innovation to the factors, $p(\{\sigma_{\epsilon,t}^j\}_{t=1}^T|\Phi, \{f_t\}_{t=1}^T)$. Finally, draw $\omega_\epsilon^{2,j}$ conditional on $\{\sigma_{\epsilon,t}^j\}_{t=1}^T$ and from IG posterior with the prior of the one degree of freedom and the scale of 10^{-4}.

3. Draw the factor loadings and serial correlation coefficients of idiosyncratic components

Draw the loadings λ^j from $p(\lambda^j|\rho^{j-1}, \{f_t^j, \sigma_{\eta i,t}^{j-1}\}_{t=1}^T, y)$. They can be estimated via GLS, conditional on $\rho_i j - 1$ and $\sigma_{\eta i,t}^{j-1}$. To be more specific, for each variable i = 1...n, I first divide the indicator (monthly or monthly-interpolated quarterly) and the factor $\{f_t^j\}_{t=1}^T$ by the volatility of the idiosyncratic component from the previous draw, $\{\sigma_{\eta i,t}^{j-1}\}_{t=1}^T$. Then I also apply quasi-differencing to the resulting indicator and factor. After these steps, I stack all the indicators and factors and run the standard OLS: for the prior on λ^j, I set the mean and variance as the matrix of ones and 0.2, respectively. Finally, I correct the estimated loadings λ^j to reflect the restriction that the loading of GDP on f_t to be unity, following Bai and Wang (2015).

Having obtained λ^j and $\{a_t^j, f_t^j\}_{t=1}^T$, calculate the idiosyncratic components $u_{it}^j = \Delta y_{it} - c_{it} - \lambda_i(L)f_t$. Draw serial correlation coefficients ρ^j from $p(\rho|\lambda^j, \{f_t^j, \sigma_{\eta i,t}^{j-1}\}_{t=1}^T, y)$ in a similar way to the above, i.e. for each variable i = 1...n, first divide u_{it}^j by the volatility of the idiosyncratic component from the previous draw, $\{\sigma_{\eta i,t}^{j-1}\}_{t=1}^T$. Then estimate Eq (3) via OLS. I set the prior mean and the variance of ρ_i^j as the vector of zeros and 0.2 with the lag decay of 2. In the case of explosive roots, I discard the draw and repeat this step.

4. Draw the SV of innovations and the outlier states.

For each variable i = 1...n, I obtain the residuals $u_{i,t}^{*,j} = (1 - \rho^j(L))u_{i,t}^j$. Then, (*in the case of the uniform mixture distribution*) conditional on $\{\sigma_{\eta i,t}^{j-1}, o_{i,t}^{j-1}, u_{i,t}^{*,j}\}_{t=1}^T$ in log terms, obtain the mixture states. Then obtain the new draws of $\{\sigma_{\eta i,t}^j\}_{t=1}^T$ using the Kalman filter and smoother, where we construct the observations by subtracting the mean of the mixture states and $\log(o_{i,t}^{j-1})$ from the $\log(u_{i,t}^{*,j})$. I initialize the filter from $x_0 \sim N(0, 10I)$, with the transition matrix and the loadings set to the Identity. The covariance matrices for the state and measurement equations consist of $\omega_{\eta i}^{2,j-1}$ and the volatility of the mixture states, respectively. Then draw $p(o_{i,t}^j|u_{i,t}^*, \sigma_{\eta i,t}^j)$ based on the cdf of standardized draws for each component of mixture normals and calculate the number of outliers to update α_i, β_i. Finally, draw p_i^j from the beta (α_i, β_i) posterior. The entire step can be parallelized and run in a univariate state-space system.

(*in the case of the student-t distribution*) conditional on $\{\sigma^{j-1}_{\eta i,t}, o^{j-1}_{i,t}, u^{*,j}_{i,t}\}^T_{t=1}$ in log terms, update the posterior distribution of the ν_i, which is proportional to the product of t-distribution ordinates, i.e. $\mathrm{p}(\nu^j_i|u^*_{i,t}, \sigma^{j-1}_{\eta i,t}) = \mathrm{p}(\nu)\Pi^T_{t=1}\frac{\nu^{\nu/2}\Gamma((\nu+1)/2)}{\Gamma(1/2)\Gamma(\nu/2)}(\nu + (u^{*,j}_{i,t}/\sigma^{j-1}_{\eta i,t})^2)^{-(\nu+1)/2}$, following Jacquier et al. (2004). I impose a weakly informative prior, i.e. $\mathrm{p}(\nu) \sim \Gamma(2, 10)$ which is discretized on the support [3:40], to ensure the existence of a conditional variance at the lower bound. Then I draw $\mathrm{p}(o^j_{i,t}|\nu^j, u^*_{i,t}, \sigma^{j-1}_{\eta i,t})$, using $\nu^j/o^j_{i,t} \sim \chi^2_{v+1}$. After this step, I follow the same process as in the previous case of the uniform mixture distribution, to update $\{\sigma^j_{\eta i,t}\}^T_{t=1}$ using the Kalman filter and smoother, where we construct the observations by subtracting the mean of the mixture states and $\log(o^{j-1}_{i,t})$ from the $\log(u^{*,j}_{i,t})$.

Then, in both cases, draw $\mathrm{p}(\omega^{2,j}_{\eta i}|\{\sigma^j_{\eta i,t}\}^T_{t=1})$ from the IG posterior with the prior of the one degree of freedom and the scale of 10^{-4}.

Increase j by 1 and repeat the steps above until j = 7000. After discarding the first 2000 as burn-in draws, use the rest 5000 draws of the model parameters and latent variables for the inference.

A.3 The unwanted dependency

Antolin-Diaz et al. (2023) model the outliers in an additive rather than multiplicative way, with the following measurement equation:

$$\Delta y_t - o_t = c_t + \Lambda(L) f_t + u_t, \quad o_{it} \sim t_{vi}(0, \sigma_{oi}^2)$$

where Δy_t is the transformed data, which can be in levels or differences, and so is o_t. By re-arranging and quasi-differencing $(1 - \rho_i(L))u_{it} = \sigma_{\eta_{it}} o_{it} \eta_{it}$ with the lag order of 2, we obtain:

$$\tilde{y}_t = \tilde{\Lambda} F_t + (1 - \rho_1 L - \rho_2 L^2)(u_t + o_t)$$

where $\tilde{y}_t$ and $\tilde{\Lambda} F_t$ are in terms of the quasi-differences, e.g. $\tilde{y}_t = (1 - \rho_1 L - \rho_2 L^2)\Delta y_t$, as shown in the state-space representation in Appendix A.1. Then, the variance of residuals, the last term on the right-hand side, can be expressed as:

$$\begin{aligned} V[(1 - \rho_1 L - \rho_2 L^2)(u_{it} + o_{it})] &= V(\sigma_{\eta_{it}} \eta_{it} + (1 - \rho_1 L - \rho_2 L^2) o_{it}) \\ &= \sigma_{\eta_{it}}^2 + (1 + \rho_1^2 + \rho_2^2) V(o_{it}) \end{aligned}$$

where the first equality reflects $(1 - \rho_i(L))u_{it} = \sigma_{\eta_{it}} o_{it} \eta_{it}$, and the second line is based on the two extra assumptions (1) – (2): u_{it} and o_{it} are uncorrelated and o_{it} are serially uncorrelated. These are mostly innocuous by the definition of "one-off" outliers. Then, the role of outlier depends on ρ, the autoregressive coefficients of the idiosyncratic component u_{it}, not the outlier states o_{it}. This is not an intended feature of the model, hence the 'unwanted' dependency. In the likely case of $\rho > 0$, it amplifies the variance of the outlier component, resulting in an underestimation of the SV components in innovations. We can calculate further to obtain the variance of outlier components:

$$\begin{aligned} V(o_{it}) &= V(\sqrt{\psi_{it}} z_{it}) \\ &= V[E(\sqrt{\psi_{it}} z_{it} | \psi_{it})] + E[V(\sqrt{\psi_{it}} z_{it} | \psi_{it})] \\ &= V[\sqrt{\psi_{it}} \underbrace{E(z_{it} | \psi_{it})}_{=0}] + E[\psi_{it} \underbrace{V(z_{it} | \psi_{it})}_{=\sigma_{oi}^2}] \\ &= \sigma_{oi}^2 E(\psi_{it}) \end{aligned}$$

where we apply the specification of outliers from Antolin-Diaz et al. (2023), i.e. the latent scale mixture variable ψ_{it} follows $\nu_i / \psi_{it} \sim \chi_{\nu_i}^2$, and the noise $z_{it} \sim N(0, \sigma_{oi}^2)$, so that $o_{it} = \sqrt{\psi_{it}} z_{it} \sim t_{\nu_i}(0, \sigma_{oi}^2)$. Since $\nu_i / \psi_{it} \sim \chi_{\nu_i}^2$ implies that $\psi_{it} / \nu_i \sim IG(\frac{\nu_i}{2}, \frac{1}{2})$, we can complete the above expression by:

$$E(\frac{\psi_{it}}{\nu_i}) = \frac{1}{\nu_i - 2}, \quad E(\psi_{it}) \qquad = \frac{\nu_i}{\nu_i - 2}$$

We can reach the same expression for $V(o_{it})$ by both additive and multiplicative modeling of outliers. However, when the quasi-differencing is applied, the dependeny issue only arises in the former case.

B Additional figures

B.1 Posterior estimate of trend and factor volatility

Figure B.1: Posterior estimate of trend v. Real GDP growth

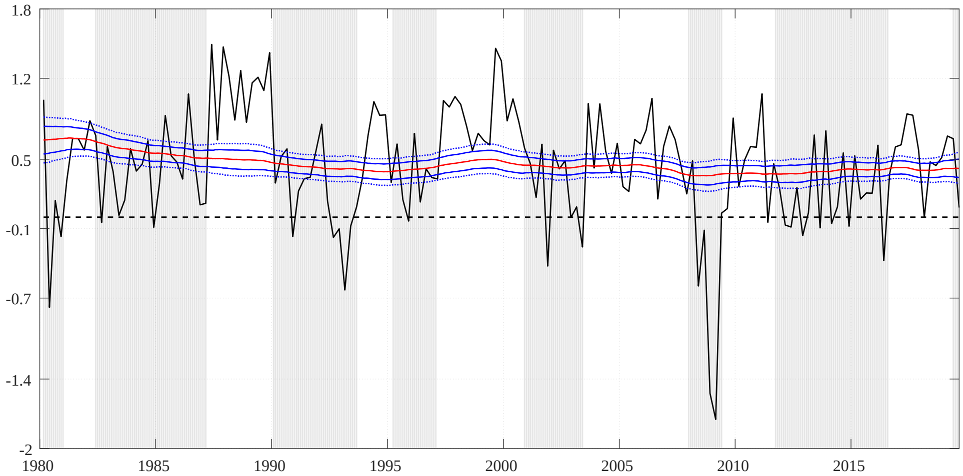

Note: This plot shows the posterior estimate of the long-run growth from the model with outliers following the uniform mixture distribution (in red and blue) and the real GDP growth of France in the black solid line. The solid red and blue lines denote the posterior median and 68% bands, respectively, while the dotted blue line represents 90% bands. Sample: 1980.1 – 2019.9. OECD recessions in gray.

Figure B.2: Volatility of the common activity factor

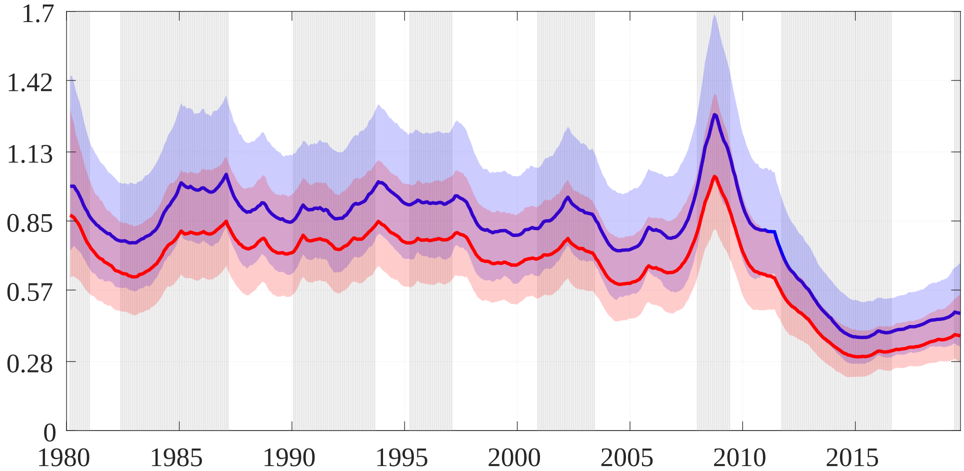

(a) student-t outliers v. no outlier

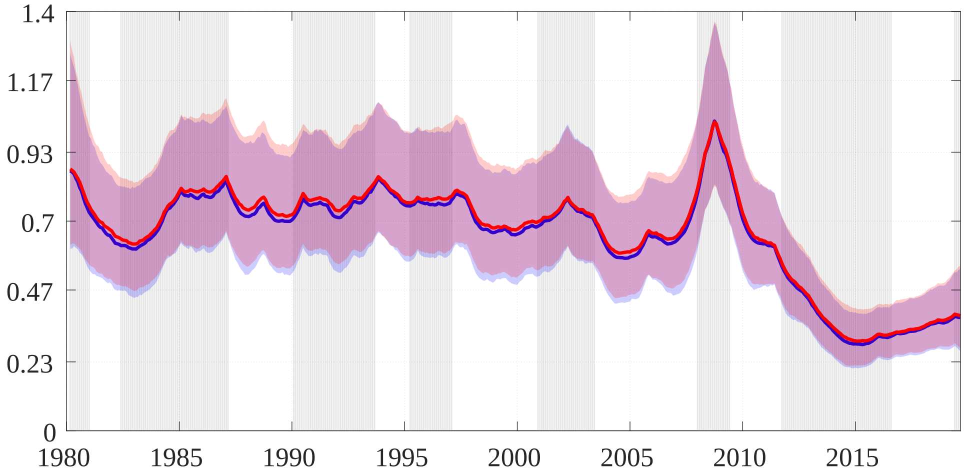

(b) uniform mixture outliers v. no outlier

Note: This plot shows the posterior estimate of volatility of the common activity factor from the two models, the one with (in blue) and the one without (in red) outliers. The solid lines and shaded areas denote the posterior median and 90% bands, respectively. Sample: 1980.1 – 2019.9. OECD recessions in gray.

B.2 Stochastic volatility: with v. without outliers

Figure B.3: SV estimates: student-t v. without outliers, monthly hard

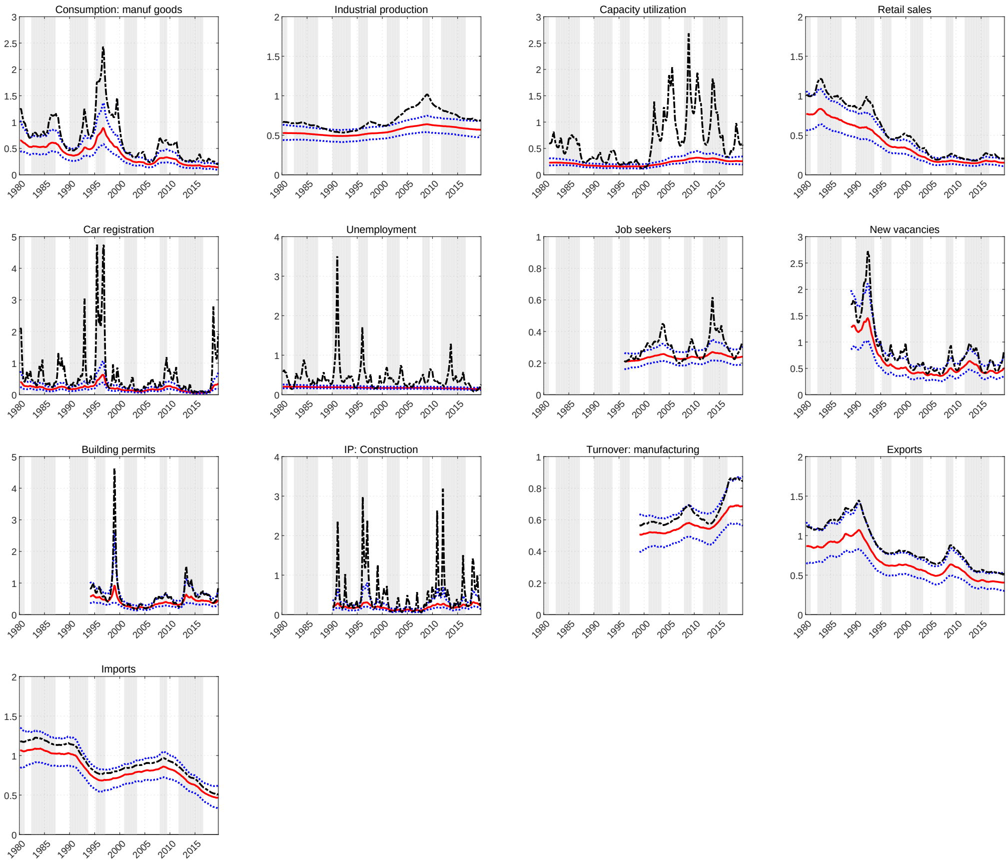

Note: This plot shows the posterior estimate of stochastic volatility in innovations to the idiosyncratic components of monthly hard indicators. The black line represents the posterior median from the model without outliers. Red and blue lines are the posterior median and 68% bands from the model with outliers following the t-distribution, respectively. Sample: 1980.1 – 2019.9. OECD recessions in gray.

Figure B.4: SV estimates: student-t v. without, quarterly and soft

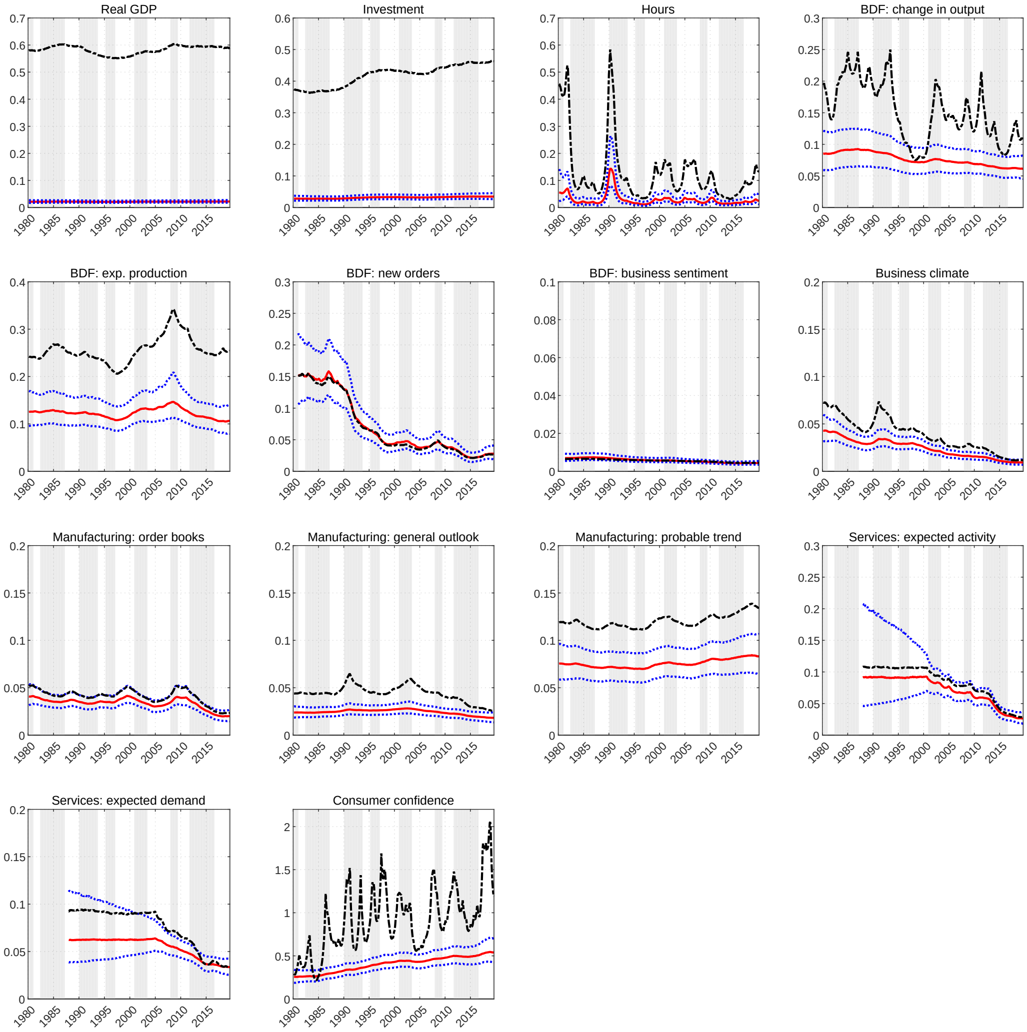

Note: This plot shows the posterior estimate of stochastic volatility in innovations to the idiosyncratic components of quarterly and monthly soft indicators. The black line represents the posterior median from the model without outliers. Red and blue lines are the posterior median and 68% bands from the model with outliers following the t-distribution, respectively. Sample: 1980.1 – 2019.9. OECD recessions in gray.

FIGURE B.5: SV ESTIMATES: UNIFORM MIXTURE V. WITHOUT, MONTHLY HARD

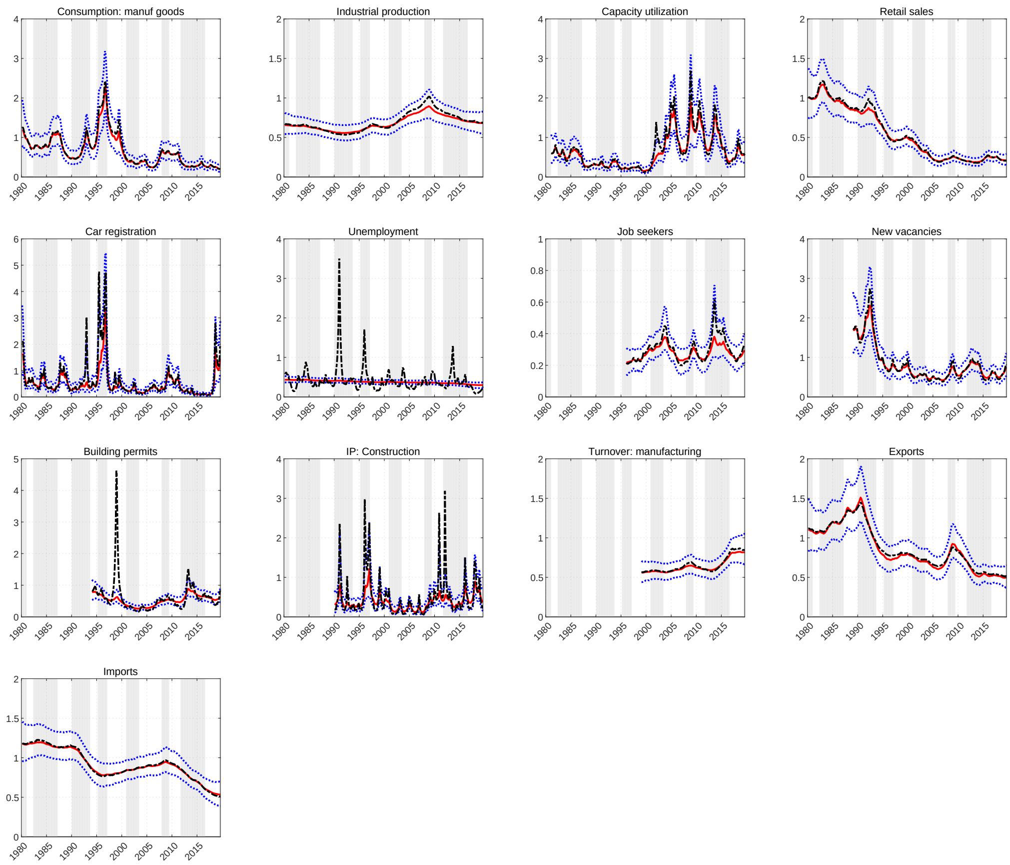

Note: This plot shows the posterior estimate of stochastic volatility in innovations to the idiosyncratic components of monthly hard indicators. The black line represents the posterior median from the model without outliers. Red and blue lines are the posterior median and 68% bands from the model with outliers following the uniform mixture distribution, respectively. Sample: 1980.1 – 2019.9. OECD recessions in gray.

Figure B.6: SV estimates: uniform mixture v. without, quarterly and soft

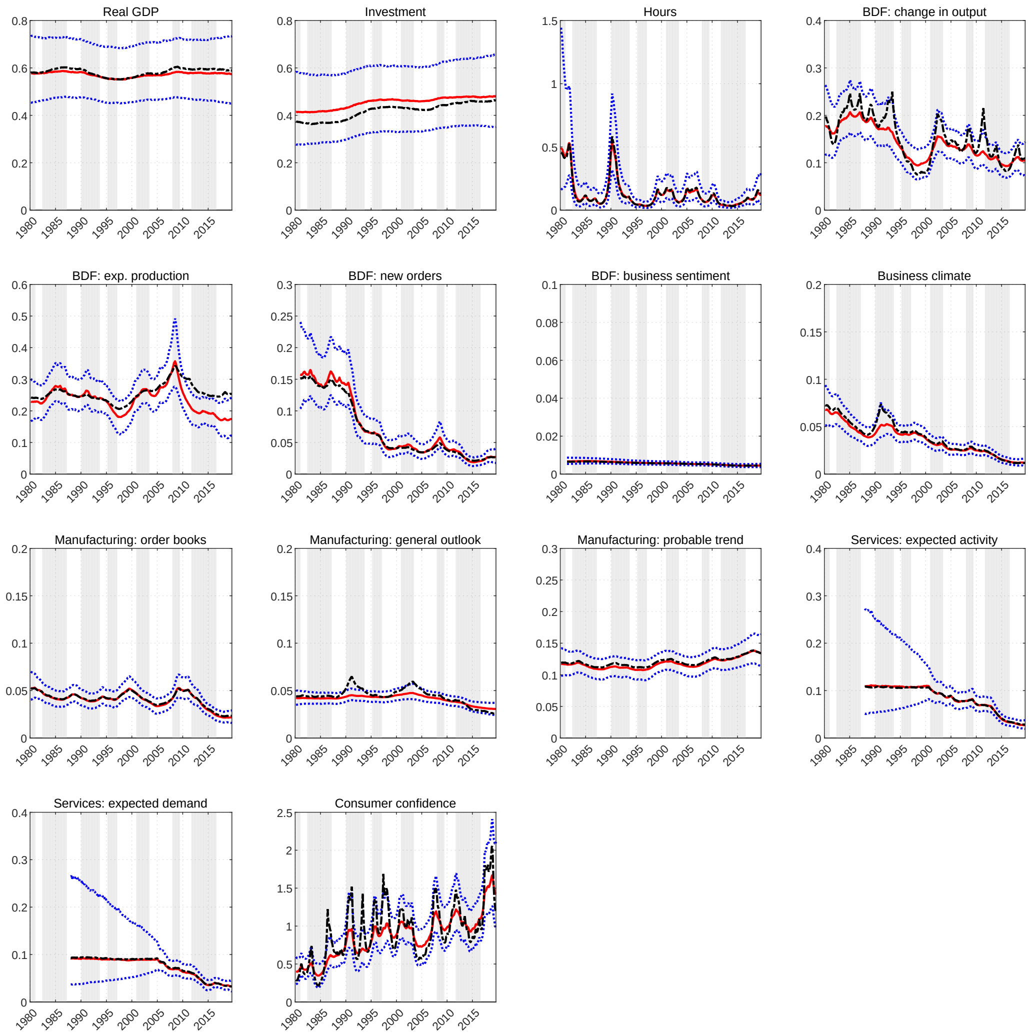

Note: This plot shows the posterior estimate of stochastic volatility in innovations to the idiosyncratic components of quarterly and monthly soft indicators. The black line represents the posterior median from the model without outliers. Red and blue lines are the posterior median and 68% bands from the model with outliers following the uniform mixture distribution, respectively. Sample: 1980.1 – 2019.9. OECD recessions in gray.

B.3 Posterior estimates of outlier states

Figure B.7: Posterior probabilities of outlier states: t-distribution

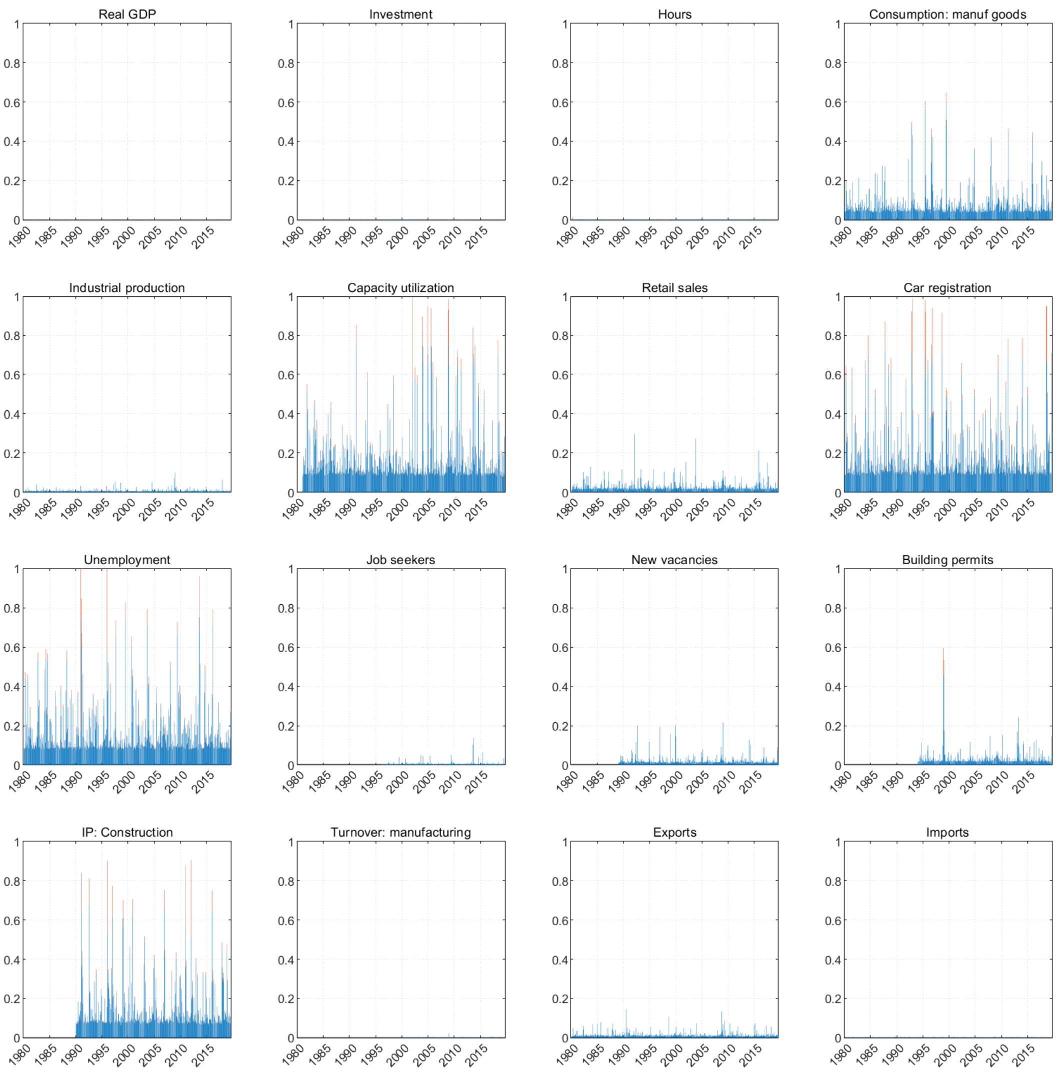

Note: The stacked bars represent posterior probabilities for realizations of outlier states that are larger than two. The blue bars correspond to the probability that the outliers are within the range between two and five, and the orange bars denote the probability of outlier states that are larger than five to take place. Monthly and quarterly hard indicators in the dataset are included. Sample: 1980.1 – 2019.9.

Figure B.8: Posterior probabilities of outlier states: uniform mixture

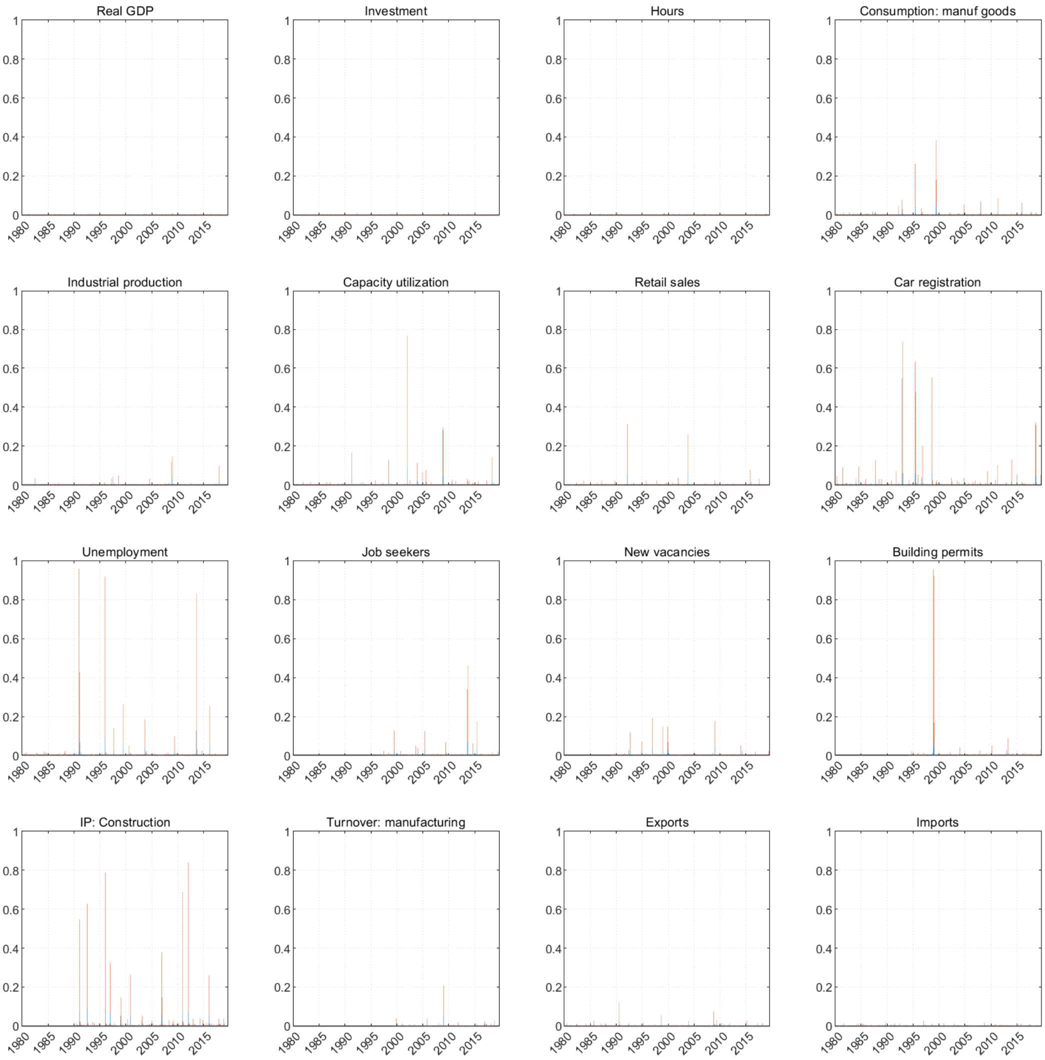

Note: The stacked bars represent posterior probabilities for realizations of outlier states that are larger than two. The blue bars correspond to the probability that the outliers are within the range between two and five, and the orange bars denote the probability of outlier states that are larger than five to take place. Monthly and quarterly hard indicators in the dataset are included. Sample: 1980.1 – 2019.9.

Figure B.9: Posterior median of outlier states: t-distribution

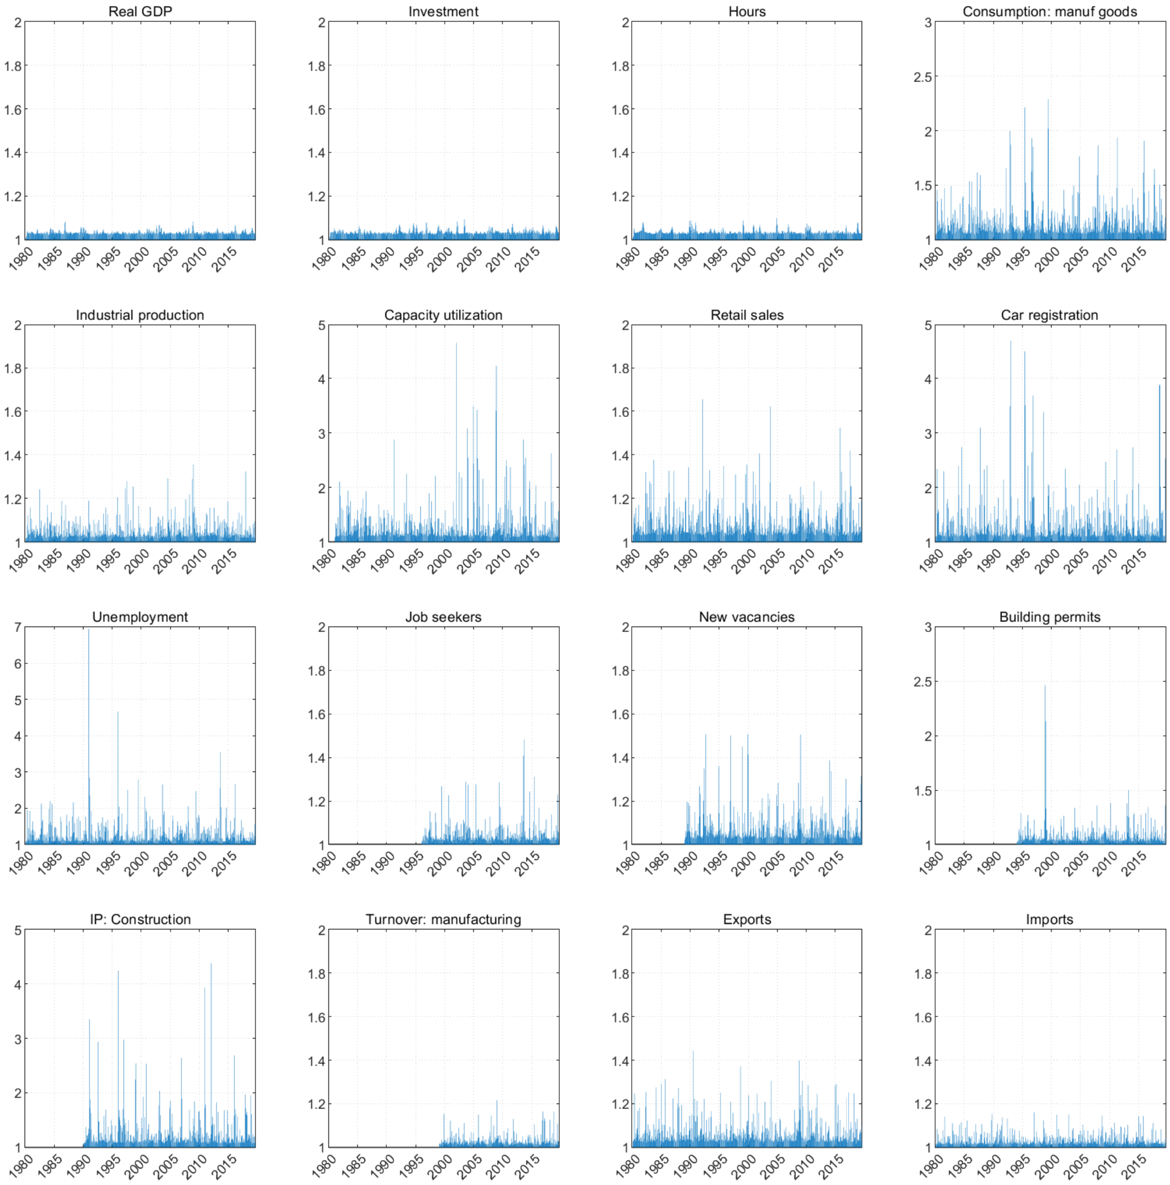

Note: This plot shows the posterior median estimates of outlier states for selected monthly and quarterly hard indicators in the dataset. Sample: 1980.1 – 2019.9.

Figure B.10: Posterior median of outlier states: uniform mixture

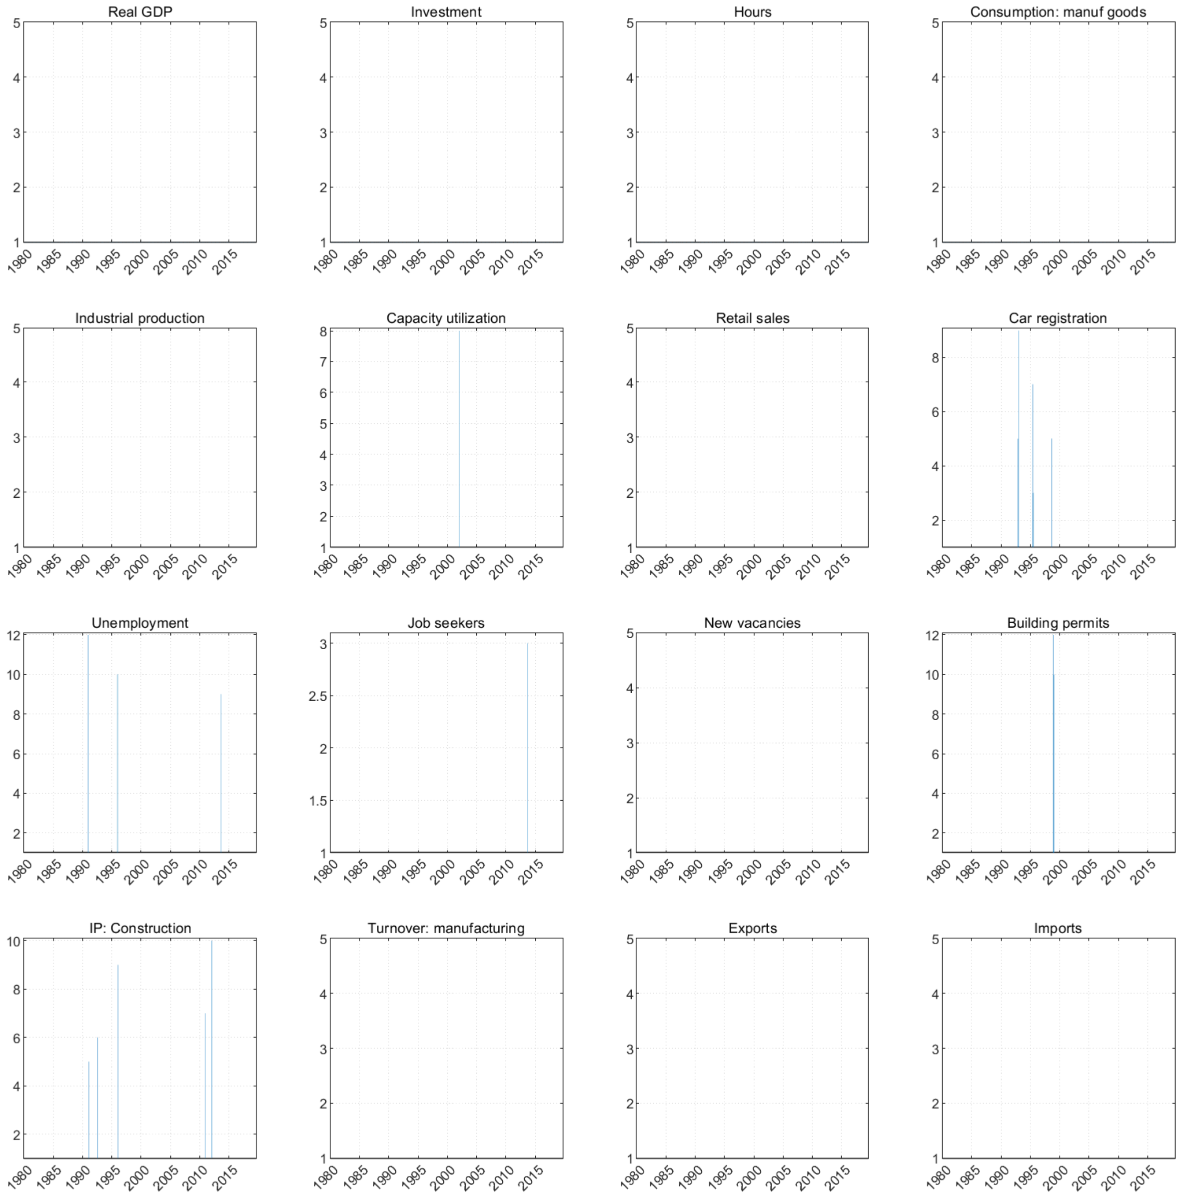

Note: This plot shows the posterior median estimates of outlier states for selected monthly and quarterly hard indicators in the dataset. Sample: 1980.1 – 2019.9.

C Additional tables

C.1 Out-of-sample point forecasting evaluation

Table C.1: Out-of-sample point forecasting evaluation, pre-COVID

Horizon/model	Full (t-dist)	Full (uniform)	SV miss	SV	Basic miss	Basic DFM
(a) point forecasting: RMSE						
-5 month	0.4774	0.4607	0.4598	0.4612	0.4835	0.484
-4 month	0.4407	0.4392	0.4359	0.4374	0.4583	0.459
-3 month	0.4052	0.3920	0.3951	0.3922	0.4230	0.422
-2 month	0.3533	0.3287	0.3278	0.3290	0.3725	0.374
-1 month	0.3054	0.3039	0.3041	0.3037	0.3443	0.347
0 month (end of reference Q)	0.3061	0.2969	0.2970	0.2974	0.3383	0.340
1 month	0.3036	0.2924	0.2948	0.2927	0.3336	0.338
(b) point forecasting: MAFE						
-5 month	0.3651	0.3467	0.3438	0.3458	0.3391	0.3402
-4 month	0.3394	0.3274	0.3270	0.3271	0.3281	0.3302
-3 month	0.3219	0.2982	0.2997	0.2983	0.3054	0.3050
-2 month	0.2803	0.2542	0.2532	0.2531	0.2749	0.2742
-1 month	0.2464	0.2431	0.2423	0.2422	0.2632	0.2633
0 month (end of reference Q)	0.2566	0.2378	0.2389	0.2396	0.2592	0.2596
1 month	0.2491	0.2337	0.2364	0.2347	0.2552	0.2577

Note: This table provides the forecasting performance of different models: the DFM model with outliers following the student-t and the uniform mixture distribution, respectively; the model with stochastic trend and volatility (Antolin-Diaz et al., 2017); and the basic DFM model (Bańbura and Modugno, 2014). 'SV miss' and 'Basic miss' are the latter two models with the outliers, which are defined as the observations five-interquartile range away from the median, replaced by missing values. Using data going back to 1980, I estimate the model at each point in time for 2005.1 – 2019.9 in an expanding window.

TABLE C.2: Out-of-sample density forecasting evaluation, pre-COVID

Horizon/model	Full (t-dist)	Full (uniform)	SV miss	SV	Basic miss	Basic DFM
(a) density forecasting: log score						
-5 month	-0.826	-0.853	-0.871	-0.887	-2.245	-2.176
-4 month	-0.792	-0.983	-0.987	-0.994	-2.695	-2.332
-3 month	-0.804	-1.181	-1.227	-1.233	-3.561	-3.988
-2 month	-0.774	-1.604	-1.694	-1.715	-7.750	-8.319
-1 month	-1.011	-2.452	-2.534	-2.506	-15.319	-15.209
0 month (end of reference Q)	-1.102	-2.907	-2.983	-3.032	-23.034	-21.222
1 month	-1.448	-2.988	-3.090	-3.018	-22.273	-23.052
(b) density forecasting: CRPS						
-5 month	0.2643	0.2534	0.2530	0.2544	0.2660	0.2677
-4 month	0.2486	0.2468	0.2461	0.2462	0.2591	0.2585
-3 month	0.2362	0.2296	0.2319	0.2304	0.2497	0.2492
-2 month	0.2075	0.1998	0.1996	0.2003	0.2348	0.2349
-1 month	0.1870	0.1933	0.1935	0.1934	0.2338	0.2332
0 month (end of reference Q)	0.1914	0.1932	0.1939	0.1946	0.2341	0.2339
1 month	0.1913	0.1901	0.1925	0.1912	0.2300	0.2327

Note: This table provides the forecasting performance of different models: the DFM model with outliers following the student-t and the uniform mixture distribution, respectively; the model with stochastic trend and volatility (Antolin-Diaz et al., 2017); and the basic DFM model (Bańbura and Modugno, 2014). 'SV miss' and 'Basic miss' are the latter two models with the outliers, which are defined as the observations five-interquartile range away from the median, replaced by missing values. Using data going back to 1980, I estimate the model at each point in time for 2005.1 – 2019.9 in an expanding window. The asterisks are related to the p-value of the null hypothesis that the basic DFM model performs as well as others, against the alternative that the other model performs better, based on the test statistic of Diebold and Mariano (1995). *** is significant at the 1% level, ** at the 5% level, and * at the 10% level.

Table C.3: Out-of-sample point forecasting evaluation, full sample

Horizon/model	Full (t-dist)	Full (uniform)	SV miss	SV	Basic miss	Basic DFM
(a) point forecasting: RMSE						
-5 month	3.3347	5.2520	5.2986	3.6733	2.8766	2.766
-4 month	6.6743	4.3425	6.3752	6.3158	4.5847	2.722
-3 month	1.9276	2.0719	4.6172	2.7033	2.8664	2.827
-2 month	1.0688	4.4881	9.6858	1.4062	2.6536	2.214
-1 month	3.6929	2.3537	4.6124	3.4823	2.1873	2.133
0 month (end of reference Q)	0.9078	0.9130	3.7779	0.9293	3.2543	1.626
1 month	0.8157	1.0999	3.2592	1.5153	2.7994	1.930
(b) point forecasting: MAFE						
-5 month	1.0806	1.3226	1.3501	1.1407	0.8873	0.9244
-4 month	1.6074	1.2093	1.4959	1.3530	1.1311	0.9075
-3 month	0.7929	0.8089	1.1182	0.8786	0.8486	0.8480
-2 month	0.5057	1.0224	1.8311	0.5804	0.8212	0.7292
-1 month	0.9206	0.6419	1.1349	0.9448	0.7278	0.6950
0 month (end of reference Q)	0.4567	0.4647	0.9741	0.4729	0.8791	0.6024
1 month	0.4453	0.4854	0.8828	0.6015	0.8268	0.6783

Note: This table provides the forecasting performance of different models: the DFM model with outliers following the student-t and the uniform mixture distribution, respectively; the model with stochastic trend and volatility (Antolin-Diaz et al., 2017); and the basic DFM model (Bańbura and Modugno, 2014). 'SV miss' and 'Basic miss' are the latter two models with the outliers, which are defined as the observations five-interquartile range away from the median, replaced by missing values. Using data going back to 1980, I estimate the model at each point in time for 2005.1 – 2022.9 in an expanding window.

Table C.4: Out-of-sample density forecasting evaluation, full sample

Horizon/model	Full (t-dist)	Full (uniform)	SV miss	SV	Basic miss	Basic DFM
(a) density forecasting: log score						
-5 month	-21.463	-28.105	-30.715	-29.218	-65.136	-244.853
-4 month	-24.854	-37.859	-38.779	-38.305	-100.828	-462.350
-3 month	-5.565	-8.832	-8.653	-8.866	-53.639	-318.273
-2 month	-8.049	-16.548	-17.443	-17.694	-79.132	-559.888
-1 month	-12.247	-23.190	-26.051	-28.649	-190.255	-871.559
0 month (end of reference Q)	-4.454	-6.801	-8.377	-6.449	-134.327	-367.203
1 month	-2.020	-3.642	-20.577	-12.662	-108.009	-419.691
(b) density forecasting: CRPS						
-5 month	0.9115	1.0188	1.0450	0.9860	0.8390	0.835
-4 month	1.3096	0.9614	1.2940	1.1914	0.9540	0.804
-3 month	0.6457	0.6224	0.9400	0.7321	0.7966	0.785
-2 month	0.3972	0.8403	1.5662	0.4886	0.7708	0.673
-1 month	0.7695	0.5212	1.0064	0.8430	0.6578	0.648
0 month (end of reference Q)	0.3693	0.3745	0.8752	0.3850	0.8379	0.557
1 month	0.3408	0.3960	0.7950	0.5328	0.7903	0.643

Note: This table provides the forecasting performance of different models: the DFM model with outliers following the student-t and the uniform mixture distribution, respectively; the model with stochastic trend and volatility (Antolin-Diaz et al., 2017); and the basic DFM model (Bańbura and Modugno, 2014). 'SV miss' and 'Basic miss' are the latter two models with the outliers, which are defined as the observations five-interquartile range away from the median, replaced by missing values. Using data going back to 1980, I estimate the model at each point in time for 2005.1 – 2022.9 in an expanding window.