Hauber, Philipp

Working Paper

Real-time nowcasting with sparse factor models

Real-time nowcasting with sparse factor models

Philipp Hauber

Julius-Maximilians-Universität Würzburg, Kiel Insitute for the World Economy

Abstract

Factor models feature prominently in the macroeconomic nowcasting literature, yet no clear consensus has emerged regarding the question of how many and which variables to select in such applications. Examples of both large-scale models, estimated with data sets consisting of over 100 time series as well as small-scale models based on only a few, pre-selected variables can be found in the literature. To adress the issue of variable selection in factor models, in this paper we employ sparse priors on the loadings matrix. These priors concentrate more mass at zero than those conventionally used in the literature while retaining fat tails to capture signals. As a result, variable selection and estimation can be performed simultaneously in a Bayesian framework. Using large data sets consisting of over 100 variables, we evaluate the performance of sparse factor models in real-time for US and German GDP point and density nowcasts. We find that sparse priors lead to relatively small gains in nowcast accuracy compared to a benchmark Normal prior. Moreover, different types of sparse priors discussed in the literature yield very similar results. Our findings are compatible with the hypothesis that large macroeconomic data sets typically used in now- or forecasting applications are not sparse but dense.

Keywords: factor models, sparsity, nowcasting, variable selection

JEL classification: C11, C53, C55, E37.

Contact details: Julius-Maximilians-Universität Würzburg, Chair of Monetary Economics and International Financial Markets, Sanderring 2, 97070 Würzburg. Email: philipp.hauber@uni-wuerzburg.de. I would like to thank Christian Schumacher for many helpful comments and discussions as well as sharing MATLAB code for the point mass Normal mixture and horseshoe priors with me. I am also grateful for Max Schröder who provided research assistance as well as Albrecht Mengel and Simone Knief for giving me access to the University of Kiel's high-performance computating facilities.

1 Introduction

Professional forecasters and policy-makers require timely assessements of the current state of the macroeconomy. However, short-term forecasting or *nowcasting* of the gross domestic product (GDP) - arguably the best single indicator of macroeconomic developments and central in guiding economic policy decisions - faces numerous challenges. The GDP is subject to a considerable publication delay, with a first estimate usually provided roughly four weeks after the end of the respective quarter. At the same time, the information from a large number of potentially informative indicators is available that can be exploited in a more timely manner to improve estimates of current macroeconomic conditions. Often, however, such series are published in an asynchronous manner giving rise to an unbalanced panel or "ragged edge". For example, survey-based sentiment indicators are typically released much sooner after (or even within) the reference period than "hard" indicators such as industrial production or retail sales. Moreover, these time series are usually sampled at higher frequencies than the target variable GDP, e.g. monthly or daily. Nowcasts that reflect on the latest available information need to be able to handle this data flow in real-time, mixing information from different frequencies.

Factor models feature prominently in the nowcasting literature as they can address all the aforementioned issues in a unified modelling framework. Furthermore, they have a proven track record in terms of nowcasting GDP (Giannone et al., 2008; Schumacher and Breitung, 2008; Camacho and Perez-Quiros, 2010; Banbura and Rünstler, 2011; Kuzin et al., 2011; Aastveit et al., 2014, 2018; Marcellino et al., 2016). Notwithstanding this extensive literature, there is no firm consensus about how many and which variables to select when nowcasting GDP. From an asymptotic point of view a larger cross-section should lead to more precise estimates of the underlying factors and, hence, more accurate nowcasts. Boivin and Ng (2006), however, find both in simulations and a real-time forecasting excercise that factors extracted from smaller data sets perform as well or even better than those extracted from a much larger panel. In line with these findings, Bai and Ng (2008) advocate penalized regressions to identify a subset of variables that is closely related to the target series prior to factor extraction ("targeted predictors"). Schumacher (2010) demonstrates that such targeted predictors can improve nowcasts of German GDP when considering a large panel of international data. Similarly, Caggiano et al. (2011) provide evidence for a number of countries that pre-selection of variables can substantially improve forecast performance. In contrast, Alvarez et al. (2016) find no clear benefit from using aggregate headline series representing different economic categories as opposed to a large disaggregated data set in a forecasting excercise for the United States. Rünstler (2016) adresses the issue of variable selection by focusing on the prediction weights inherent in a factor model. Selecting those variables with the largest marginal predictive gains for GDP growth, he finds only moderate gains in forecasting accuracy for short horizons.

As a result of this on-going debate, examples of both approaches can be found in the literature: "large-scale"" factor models comprised of around 70 to 80 time series or more

(Giannone et al., 2008; Banbura and Rünstler, 2011; Kuzin et al., 2011; Aastveit et al., 2014, 2018) or "small-scale " models with up to 10 or 20 variables, pre-selected by expert judgement, statistical procedures or a combination thereof (Camacho and Perez-Quiros, 2010; Marcellino et al., 2016; Bok et al., 2018).

In this paper, the issue of variable selection is adressed by exploring the role sparsity plays in factor models. We investigate to what extent the Bayesian estimation of factor models with sparse priors on the loadings matrix can serve as an alternative to pre-selection of variables. Compared to Normal priors commonly used in the literatur, sparse priors place considerably more mass near zero while still allowing for fat tails to capture signals. Based on the work by George and McCulloch (1993, 1997) on variable selection priors in a multiple regression framework, West (2003) proposes mixtures of Normal distributions and a point mass at zero to induce sparsity in factor models. Employing these priors in a macroeconomic application, Kaufmann and Schumacher (2017) identifiy relevant variables in large panels of international GDP growth rates and disaggregate US CPI data. Global-local shrinkage priors have been proposed as a continuous approximation to discrete variable selection priors (Polson and Scott, 2010) and have been widely used in macroeconomic forecasting applications with Bayesian vector autoregressions (BVAR), e.g. Kastner and Huber (2020), Huber and Feldkircher (2019), Cross et al. (2020) or Chan (2021) in combination with conventional Minnesota-type shrinkage priors.

To assess the performance of sparse factor models, we conduct a real-time evaluation of point and density nowcasts for US and German GDP. Our contribution to the literature can be viewed in two ways: from the methodological side, we perform variable selection and estimation simultaneously in a real-time factor model setting. As such, it follows the recent contributions of Kristensen (2017) and Thorsrud (2020). The former estimates factors nonparametrically via sparse principal components and finds gains in forecasting performance while the latter induces time-varying sparsity in a Bayesian factor model using a latent threshold mechanism (Nakajima and West, 2013). A further novel aspect of our paper is a comparsion of different sparse priors that have been proposed in the literature, with a particular focus on their performance in nowcasting applications. On a conceptual level, we contribute to the current debate about the degree of sparsity in macroeconomic data and its implications for now- and forecasting (Giannone et al., 2018).

The real-time evaluation of nowcasts for US and German GDP suggests that while the factor models outperform autoregressive benchmarks, sparse priors lead to relatively small gains in nowcast accuracy compared to our benchmark Normal prior. Moreover, different types of sparse priors discussed in the literature yield very similar results. Our findings are thus compatible with the hypothesis that large macroeconomic data sets typically used in now- or forecasting applications are not sparse but dense. The remainder of this paper is structured as follows: Section 2 discusses the different sparse priors and lays out their conditional posterior distributions. The results from the real-time nowcast evaluation for US and German GDP along with a number of robustness checks are presented in Section 4. Section 5 concludes.

Notation: Let $z_{1:T} = [z_1, \ldots, z_T]'$ refer to the $T \times N$ matrix where z_t is an $N \times 1$ vector, and $z_{i,1:T} = [z_{i1}, \ldots, z_{iT}]'$ to the vector of length T corresponding to the i-th column of $z_{1:T}$. Finally, let 0_k (I_k) denote a zero column vector (identity matrix) of dimension k. Furthermore, let A_{ij} refer to the element in the i-th row and j-th column of matrix A. Conversely, $A_{i\cdot}$ and $A_{\cdot j}$ denote the i-th row and j-th column, respectively. Additionally, we denote by $\mathcal{N}\left(\mu, \sigma^2\right)$ the Normal distribution with mean μ and variance σ^2. Furthermore, let $\mathcal{N}\left(x; \mu, \sigma^2\right)$ denote the value of the probability density function (pdf) of $\mathcal{N}\left(\mu, \sigma^2\right)$ evaluated at x. We denote by $\mathcal{U}(a, b)$ the continuous uniform distribution on the interval a to b and by $\mathcal{G}(u, U)$ the Gamma distribution with shape and rate parameters u and U, respectively. Its pdf is given by

$$\mathcal{G}(x; u, U) = \frac{U^u}{\Gamma(u)} x^{u-1} \exp(-Ux).$$

Similarly, let $\mathcal{G}^{-1}(u, U)$ denote the Inverse Gamma distribution with pdf

$$\mathcal{G}^{-1}(x; u, U) = \frac{U^u}{\Gamma(u)} x^{-u-1} \exp\left(-\frac{U}{x}\right).$$

$\mathcal{B}(a, b)$ denotes the Beta distribution with pdf

$$\mathcal{B}(x; a, b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} x^{a-1} (1-x)^{b-1}.$$

Also, let $\mathcal{C}^+(\mu, \gamma)$ denote the half-Cauchy distribution with location and scale parameters, μ and γ, respectively and pdf equal to

$$\mathcal{C}^+(x; \mu, \gamma) = 2\left(\pi\gamma\left[1 + ((x-\mu)\gamma)^2\right]\right)^{-1}.$$

2 Sparse factor models

2.1 Factor model

Let x_t denote an $N \times 1$ vector of time series observed at time $t = 1, \ldots, T$. The idea of a factor model is that x_t can be expressed as the sum of two orthogonal elements: a common and an idiosyncratic component, i.e.

$$x_t = \lambda f_t + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \Sigma_\varepsilon) \tag{1}$$

$$\Psi(L) f_t = \eta_t, \quad \eta_t \sim \mathcal{N}\left(0, \Sigma_\eta\right) \tag{2}$$

where λ is the $N \times R$ loadings matrix linking the $R \times 1$ vector of factors, f_t to the observables in x_t. The product of loadings and factors constitutes the common component of the model while ϵ_t captures variable-specific, idiosyncratic developments. The dynamics of the factors are given by stationary vector autoregressions of order P. Σ_ε is a diagonal matrix, implying that the idiosyncratic components are independent of each other and we denote by σ_i^2 the

element corresponding to the i-th variable. Furthermore, the innovations to the common factors and the idiosyncratic components are uncorrelated, i.e. $\mathbb{E}\left[\eta_t \varepsilon_t'\right] = 0$. Lastly, to identify the scale of the factors we set $\Sigma_\eta = I_R$.

2.2 Sparse priors

From a Bayesian perspective, there are two main sparse prior alternatives: discrete mixtures and pure shrinkage priors. The former prior combines a point mass at zero or a continuous distribution with small variance and an absolutely continuous prior providing tail mass, whereas the latter is solely based on a continuous prior providing shrinkage towards zero (Carvalho et al., 2009).

A standard example for a discrete-mixtures prior is the variable selection prior for multiple regression proposed by George and McCulloch (1993) and George and McCulloch (1997). This prior was adopted and modified for factor models by West (2003) and has the following form:

$$\lambda_{ir}|\rho_j, \tau_r \sim (1-\rho_r)\delta_0(\lambda_{ir}) + \rho_r \mathcal{N}(0, \tau_r), \tag{3}$$

$$\rho_r|r_0 s_0 \sim \mathcal{B}(r_0 s_0, r_0(1-s_0)) \tag{4}$$

$$\tau_r|a_0, b_0 \sim \mathcal{G}^{-1}(a_0, b_0). \tag{5}$$

The unit point mass at zero is denoted as $\delta_0(\cdot)$, and nonzero loadings on factor j are drawn from a Normal prior with variance τ_r. Following the applications in Kaufmann and Schumacher (2017) and Kaufmann and Schumacher (2019), the probability of non-zero loading ρ_r is factor-specific as well as the variance τ_r. ρ_r is the probability of non-zero loading and follows a Beta distribution, and scale τ_r an inverse Gamma. An extension to element-wise ρ_{ir} is discussed in Carvalho et al. (2008). However, simulation results in Kaufmann and Schumacher (2017) do not show huge differences in practice. From here on, we call the prior (3) the point-mass normal mixture prior, or PMNM prior in brief.

Global-local shrinkage priors are typically normal distributions, with a global and a local variance component (Polson and Scott, 2010). The idea is that global shrinkage handles the noise, whereas local variances act to detect the signals. Thus, the two components solve the trade-off between shrinking the noise toward zero leaving the large signals unshrunk.

The literature on global-local shrinkage priors mainly focusses on the application for multiple regression problems, not on factor models. An exception is the global-local prior by Bhattacharya and Dunson (2011):

$$\lambda_{ir}|\phi_{ir}, \kappa_r \sim \mathcal{N}(0, \phi_{ir}^{-1}\kappa_r^{-1}), \quad \kappa_r = \prod_{l=1}^{r} \delta_l, \tag{6}$$

$$\delta_1|a_1 \sim \mathcal{G}(a_1, 1) \tag{7}$$

$$\delta_l|a_2 \sim \mathcal{G}(a_2, 1) \quad \text{for} \quad l = 2, \ldots, r, \tag{8}$$

$$\phi_{ir} \sim \mathcal{G}(w/2, w/2). \tag{9}$$

In this prior, global shrinkage is governed factor-specific through κ_r. By the multiplicative gamma structure with $\delta_1, \ldots, \delta_r$, the κ_rs are stochastically increasing under the restriction $a_2 > 1$, which favours more shrinkage as the column index of the loading matrix increases. Local shrinkage is governed by loading-specific ϕ_{ir}, which follows a Gamma distribution. From here on, we call this prior the multiplicative-gamma prior, or MG prior in brief.

The horseshoe prior proposed by Carvalho et al. (2009) and Carvalho et al. (2010) is based on a standard half-Cauchy distribution for the local and the global scales. It has the hierarchical representation

$$\lambda_{ir}|\zeta_{ir}^2, \upsilon^2 \sim N(0, \zeta_{ir}^2 \upsilon^2) \tag{10}$$

$$\zeta_{ir} \sim \mathcal{C}^+(1), \qquad \upsilon \sim \mathcal{C}^+(1). \tag{11}$$

With respect to its shrinkage properties, the Cauchy tails in ζ_{ir} allow strong signals to remain large a posteriori. At the same time, its infinitely tall spike at the origin provides severe shrinkage for the zero elements of λ_{ir}. Note that the global variance parameter υ is applied to the whole set of loadings, not to columns of the loading matrix, and also follows a half-Cauchy distribution. Note that the horseshoe prior is free of user-chosen hyper parameters (Carvalho et al., 2010). An extension to accommodate more sparsity is the horseshoe+ prior by Bhadra et al. (2017). It is defined as (10) and (11) before, but with an augmented prior for ζ_{ir} according to $\zeta_{ir}|\chi_{ir} \sim \mathcal{C}^+(0, \chi_{ir})$ and $\chi_{ir} \sim \mathcal{C}^+(0, 1)$. From here on, we abbreviate the two priors as HS for horseshoe and HS+ for horseshoe+.

A natural benchmark to the sparse priors discussed above is a Normal prior on the elements of the loadings matrix, i.e. $\lambda_{ir} \sim \mathcal{N}(0, V_{ir})$. Uninformative priors - common in the nowcasting literature with Bayesian factor models - are imposed by setting V_{ir} to a large value. A hierarchical alternative is to consider $V_{ir} = \tau_r$ with $\tau_r \sim \mathcal{G}^{-1}(g_0, G_0)$ $\forall$ $r = 1, \ldots, R$ and $\forall$ $i = 1, \ldots, N$. We will refer to the benchmark prior as Normal-Inverse Gamma (NIG).

Discussion of the alternative priors

In the literature, discrete mixture priors play a prominent role due to their theoretical properties. In particular, point-mass mixture priors are highly appealing by allowing for separate control of the level of sparsity and the size of the signal coefficients (Bhattacharya et al., 2015). In a multivariate normal mean model context, Castillo and van der Vaart (2012) show that the point-mass mixture prior with an appropriate beta prior on the inclusion probability and suitable tail conditions on the normal component leads to a optimal rate of posterior contraction. The posterior of the location parameter concentrates most of its mass on a ball around its true value. When inferring about sparsity on the true covariance matrix implied by a factor model, Pati et al. (2014) show that the point-mass mixture prior on the factor loadings leads to a consistent estimation of the covariance matrix when $N > T$.

In Polson and Scott (2010), the point-mass mixture prior serves as a benchmark to compare alternative global-local shrinkage priors. Following Polson and Scott (2010), the global-local shrinkage priors implying a weaker concept of sparsity than the point-mass

mixture prior, since all of the entries are assumed to be nonzero, yet most of them small compared to a few large signals. In addition, as argued in Polson and Scott (2010) and Bhattacharya et al. (2015), the global-local shrinkage priors can offer computational savings over point-mass priors.

In practice, the question is which sparsity prior should be chosen. There is a huge and growing literature on sparse priors with a number of proposals regarding new default priors and different fields of applications. However, simulation results and empirical evidence seems to be not fully conclusive. In the factor model context followed in this paper, we focus on the comparison between the Normal Inverse-Gamma prior, which serves as a benchmark in the factor models literature, on the one hand, and a range of prominent sparse priors on the other hand to consider the variety of priors available in the recent literature.

In Figure 1, we compare the alternative priors by histograms obtained from sampling from the priors. By looking at the central bin on the left panel of the Figure 1, we can see how the sparse priors concentrate mass near zero compared to the normal prior. The PMNM and HS+ concentrate considerably more mass than the MG near zero and the Normal prior. The tail behaviour is shown in the right panel. The HS+ and HS have heavier tails than the MG, the PMNM and the Normal prior. By construction, the PMNM converges to zero quicker the Normal prior due to the mass at zero and the same specifications for the variance in the Normal distribution.

2.3 Estimation and conditional posteriors

In order to estimate the model in (1) we need to derive the joint posterior distribution of the parameters of the model and the factors. By assumption, the prior on the loadings and the remaining parameters of the model, denoted by $\Theta = \left\{\Psi, \Sigma_\eta, \Sigma_\varepsilon\right\}$, are independent of each other. The likelihood is given by

$$L\left(x_{1:T}|f_{1:T}, \lambda, \Theta\right) = \prod_{t=1}^{T} p(x_t|f_t, \lambda, \Theta) \tag{12}$$

$$p(x_t|f_t, \lambda, \Theta) = \frac{1}{(2\pi)^{N/2}\left|\Sigma_\varepsilon\right|^{1/2}} \exp\left\{-\frac{1}{2}\sum_{i=1}^{N}\frac{\left[(x_{it} - \lambda f_t)\right]^2}{\sigma_i^2}\right\}. \tag{13}$$

Combining the likelihood with the prior on the parameters (and implicitly the factors) yields the joint posterior

$$p(\lambda, \Theta, f_{1:T}|x_{1:T}, f_{1:T}) \propto L\left(x_{1:T}|f_{1:T}, \lambda, \Theta\right)\, p(f_{1:T}|\lambda, \Theta)\, p(\lambda)\, p(\Theta). \tag{14}$$

Draws from it can be obtained using Gibbs Sampling techniques, i.e. sequentially sampling the parameters conditional on the factors and then conditioning on the draw of the parameters to sample the latent factors. Specifically, at each iteration m the Gibbs Sampler cycles through the following steps or blocks:

B1: draw $f_{1:T}^{(k)}|x_{1:T}, \Theta^{(k-1)}, \lambda^{(k-1)}$

Figure 1: Alternative prior distributions

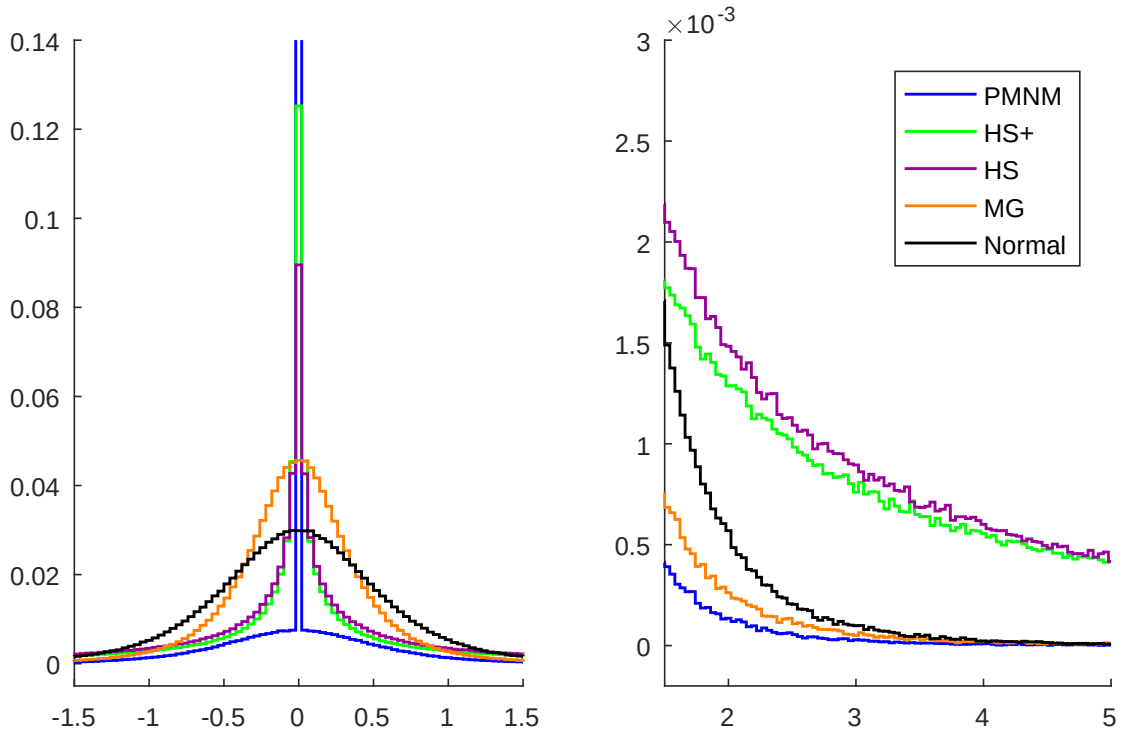

Note: The Figure shows histograms of draws from the prior distributions for 2×10^6 draws. The width of each bin is 0.04, the median bin is centered around zero. The frequency for each bin is shown on the vertical axis. PMNM is the point mass normal mixture prior with hyper parameters $r_0 = 3$, $s_0 = 0.25$, $a_0 = 2$, $b_0 = 0.5$. HS and HS+ denote the horseshoe and horseshoe+ priors, respectively. MG is the multiplicative gamma prior for the first factor with hyper parameters $a_1 = 10$, $a_2 = 2$, and $\upsilon = 3$. The Normal prior has mean zero and the variance follows an inverse Gamma distribution with $a_0 = 2$ and $b_0 = 0.5$ as in the PMNM prior. Details on the priors are provided in Section 2.2 of the main text.

B2: draw $\Psi^{(k)}|f_{1:T}^{(k)}$ and $\Sigma_{\epsilon}^{(k)}|x_{1:T}, f_{1:T}^{(k)}, \lambda^{(k-1)}$

B3: draw $\lambda^{(k)}|x_{1:T}, f_{1:T}^{(k)}, \Sigma_{\epsilon}^{(k)}$

B4: update the hyperparameters of the prior on λ

B1 and **B2** are standard and draws from the respective posterior distributions easily obtained. We leave the details of these blocks to the Appendix and focus here on the conditional posterior distributions of the loadings matrix λ and the hyperparameters governing the different priors (**B3** and **B4**).

We start with the Normal-Inverse Gamma, Multiplicative Gamma and horseshoe(+) priors where standard Bayesian linear regression results can be employed, before turning the point mass-Normal mixture which requires a different treatment. Under a (conditionally) Normal prior, the posterior of $\lambda_{i.}$ is given by (e.g. Kroese and Chan, 2013)

$$\begin{aligned} \lambda_{i.}|\cdot &\sim \mathcal{N}(m_i, M_i) && (15)\\ M_i &= \left(D^p + \frac{1}{\sigma_i^2}\sum_{t=J+1}^{T} f_t' f_t\right)^{-1} && (16)\\ m_i &= M_i\left(\frac{1}{\sigma_i^2}\sum_{t=J+1}^{T} f_t' x_{i,t}\right) && (17) \end{aligned}$$

where D^p is the prior precision matrix of $\lambda_{i.}$ and depends on the hyperparameters of the respective priors $p = \{NIG, MG, HS(+)\}$.

For example, under the Normal-Inverse Gamma prior $D^{NIG} = \text{diag}(\tau_1^{-1}, \ldots, \tau_R^{-1})$. In turn, conditional on a draw of the loadings we can update the hyperparameters by sampling from their Inverse-Gamma posterior distribution, i.e.

$$\tau_r|\lambda \sim \mathcal{G}^{-1}\left(g_0 + 0.5N, G_0 + 0.5\sum_{i=1}^{N}\lambda_{ir}^2\right) \quad (18)$$

for all $r = 1, \ldots, R$.

In the case of the MG prior, the prior precision consists of an additional term reflecting the variable-specific, local shrinkage component. Thus, $D^{MG} = \text{diag}(\kappa_1\phi_{i1}, \ldots, \kappa_R\phi_{iR})$ and updating the hyperparameters requires draws from the conditional posterior distributions given by

$$\begin{aligned} \phi_{ir}|\lambda, \kappa &\sim \mathcal{G}\left(\frac{w+1}{2}, \frac{w+\kappa_r\lambda_{ir}^2}{2}\right) \forall\, i = 1, \ldots, N,\ r = 1, \ldots, R && (19)\\ \delta_1|\lambda, \kappa, \phi &\sim \mathcal{G}\left(a_1 + \frac{NR}{2}, 1 + \frac{1}{2}\sum_{l=1}^{R}\kappa_l^{(1)}\sum_{i=1}^{N}\phi_{il}\lambda_{il}^2\right) && (20)\\ \delta_r|\lambda, \kappa, \phi &\sim \mathcal{G}\left(a_2 + \frac{N}{2}(R-r+1), 1 + \frac{1}{2}\sum_{l=1}^{R}\kappa_l^{(r)}\sum_{i=1}^{N}\phi_{il}\lambda_{il}^2\right) \forall\, r > 1, && (21) \end{aligned}$$

where $\kappa_l^{(r)} = \sum_{k=1, k\neq r}^{l}\delta_k \;\forall\; r = 1, \ldots, R$.

Recall that for both the horseshoe and the horseshoe+ prior, $D^{HS(+)} = \upsilon^2 \text{diag}(\zeta_{i1}^2, \ldots, \zeta_{iR}^2)$. To update the hyperparameters, we follow Makalic and Schmidt (2016), who exploit a scale-mixture representation of the half-Cauchy distribution: if $x^2|a \sim \mathcal{G}^{-1}(1/2, 1/a), \quad a \sim \mathcal{G}^{-1}(1/2, 1/A^2)$ then $x \sim \mathcal{C}^+(0, A)$ (Wand et al., 2011). Therefore, by introducing the auxiliary variables $\beta^\zeta, \beta^\upsilon$ we can rewrite the prior in (11) as

$$\begin{aligned} \zeta_{ir}^2|\beta_{ir}^\zeta &\sim \mathcal{G}^{-1}(1/2, 1/\beta_{ir}^\zeta) \\ \upsilon^2|\beta^\upsilon &\sim \mathcal{G}^{-1}(1/2, 1/\beta^\upsilon) \\ \beta_{11}^\zeta, \ldots, \beta_{NR}^\zeta, \beta^\upsilon &\sim \mathcal{G}^{-1}(1/2, 1). \end{aligned}$$

This leads to convenient posterior distributions for the hyperparameters of the form:

$$\zeta_{ir}^2|\lambda, \beta_{ir}^\zeta \sim \mathcal{G}^{-1}\left(1\,,\, \frac{1}{\beta_{ir}^\zeta} + \frac{\lambda_{ir}^2}{2\upsilon^2}\right) \tag{22}$$

$$\upsilon^2|\lambda, \beta^\upsilon \sim \mathcal{G}^{-1}\left(\frac{NR+1}{2}\,,\, \frac{1}{\beta^\upsilon} + \frac{1}{2}\sum_{r=1}^{R}\sum_{i=1}^{N}\frac{\lambda_{ir}^2}{\zeta_{ir}^2}\right). \tag{23}$$

Similarly, the conditional posteriors of the auxiliary variables β^υ and β^ζ are also Inverse-Gamma and given by

$$\beta_{ir}^\zeta|\zeta_{ir} \sim \mathcal{G}^{-1}\left(1\,,\, 1 + \frac{1}{\zeta_{ir}^2}\right) \forall\, i = 1, \ldots, N,\, r = 1, \ldots, R \tag{24}$$

$$\beta^\upsilon|\kappa \sim \mathcal{G}^{-1}\left(1\,,\, 1 + \frac{1}{\kappa^2}\right). \tag{25}$$

For the HS+ prior, the formula for each β_{ir}^ζ needs to be adjusted to reflect the additional half-Cauchy prior on χ_{ir}, which in turn governs the prior on ζ_{ir}. To this end, we introduce another set of $\mathcal{G}^{-1}(1/2, 1)$ auxiliary variables, β^χ. In this case, the conditional posterior distribution for ζ_{ir}^2 remains unchanged as we are still only conditioning on β_{ir}^ζ. However, the rate parameter of the latter's conditional Inverse-Gamma posterior now also depends on χ_{ir}, i.e.

$$\beta_{ir}^\zeta|\zeta_{ir}, \chi_{ir} \sim \mathcal{G}^{-1}\left(1\,,\, \frac{1}{\chi_{ir}^2} + \frac{1}{\zeta_{ir}^2}\right) \forall\, i = 1, \ldots, N,\, r = 1, \ldots, R \tag{26}$$

while χ_{ir}^2 itself and the corresponding auxiliary variables can be updated at each iteration

of the Gibbs Sampler as follows:

$$\chi_{ir}^2|\beta_{ir}^{\chi} \sim \mathcal{G}^{-1}\left(1, \frac{1}{\beta_{ir}^{\chi}} + \frac{1}{\beta_{ir}^{\zeta}}\right) \tag{27}$$

$$\beta_{ir}^{\chi}|\chi_{ir}^2 \sim \mathcal{G}^{-1}\left(1, 1 + \frac{1}{\chi_{ir}^2}\right), \tag{28}$$

again for all $i = 1, \ldots, N$, $r = 1, \ldots, R$.

The conditional posterior under the PMNM is more involved due to the mixture in (33). We start by defining the transformed observation $x_{it}^* = x_{i,t} - \sum_{l=1,l\neq r}^{R} \lambda_{il} f_{l,t} = \lambda_{ir} f_{r,t} + \epsilon_{it}$ which isolates the impacts of the factors other than r and the corresponding observation density $p\left(x_{it}^*|\cdot\right) = \mathcal{N}\left(\lambda_{ir} f_{rt}, \sigma_i^2\right)$. Then, combining the marginal prior on λ_{ir} with the likelihood yields

$$p(\lambda_{ir}|\cdot) = \prod_{t=J+1:T}^{T} p(x_{it}^*|\cdot)\{(1-\rho_r)\delta_0(\lambda_{ir} + \rho_r \mathcal{N}(0, \tau_r)\} \tag{29}$$

$$= P(\lambda_{ir} = 0|\cdot)\,\delta_0(\lambda ir) + P(\lambda_{ir} \neq 0|\cdot)\mathcal{N}(m_{ir}, M_{ir}) \tag{30}$$

with

$$M_{ir} = \left(\tau^{-1} + \frac{1}{\sigma_i^2}\sum_{t=J+1}^{T} f_{r,t}' f_{r,t}\right)^{-1} \tag{31}$$

$$m_{ir} = M_{ir}\left(\frac{1}{\sigma_i^2}\sum_{t=J+1}^{T} f_{r,t}' x_{i,t}^*)\right). \tag{32}$$

In order to sample from this distribution, we need to evaluate the posterior odds ratio of a non-zero loading, i.e.

$$PO_{ir} = \frac{P(\lambda_{ir} \neq 0|\cdot)}{P(\lambda_{ir} = 0|\cdot)} = \frac{p(\lambda_{ir})|_{\lambda_{ir}=0}}{p(\lambda_{ir}|\cdot)|_{\lambda_{ir}=0}} \frac{\rho_r}{1-\rho_r} = \frac{\mathcal{N}(0; 0, \tau_r)}{\mathcal{N}(0; m_{ir}, M_{ir})} \frac{\rho_r}{1-\rho_r} \tag{33}$$

A draw of λ_{ir} is then obtained by sampling from $\mathcal{N}(m_{ir}, M_{ir})$ and keeping the draw if $u \leq PO_{ir}/(1 + PO_{ir})$, where u is a draw from $\mathcal{U}(0, 1)$. Otherwise, we set $\lambda_{ir} = 0$.

3 Empirical application: nowcasting GDP

In this section, we use the sparse factor models outlined above to nowcast the quarterly gross domestic product (GDP) in the United States (US) and Germany. We first give an overview of the monthly indicators that are included in the dataset along with the GDP (Section 3.1). Next, in Section 3.2 we highlight how the factor model can be estimated while accouting for mixed frequency data. The real-time evaluation set-up and results are

discussed in Section 3.3 and 3.4, respectively. Lastly, we provide robustness checks to our results for different model specifications such as the transformation of survey variables or the estimation window in Section 3.5.

3.1 Data

To nowcast the US GDP, we use the monthly real-time data set provided by McCracken and Ng (2016) of more than 100 time series covering a broad spectrum of macroeconomic activity and mostly ranging back until January 1959. Vintages for this dataset are available from November 1999 onwards. For the German economy we construct a dataset based on vintages from the Deutsche Bundesbank's Real-Time Database, covering series such as production, orders, turnover, prices and the labor market. These vintages are available on a broad basis from the end of 2005 onwards. To these, we add financial market indicators also sourced from the Deutsche Bundesbank. For both countries, we augment the monthly datasets with survey-based sentiment indicators which have proven useful in nowcasting applications due to their timely release.

The estimation samples for the United States and Germany start in January 1985 and 1992, respectively, and are recursively expanded. While the two datasets are similar in terms of economic categories, there are some differences regarding the importance of individual groups. For example, in the German data there is more detailed coverage of production, orders and turnover of the industrial sector, while FRED-MD contains more disaggregated labor market series. A detailed description of the variables we use and the transformations applied to the series prior to estimation can be found in Appendix 7.

3.2 Bayesian estimation of mixed-frequency factor models

In order to combine the monthly indicators with our quarterly target variable - quarter-on-quarter GDP growth - we need to adjust the model in (1) to account for mixed frequencies. This is done by formulating the model at the highest frequency, i.e. monthly, and treating quarterly variables as monthly time series with occasionally missing observations. Thus, we essentially convert the mixed-frequency problem into a missing-value problem which can easily be handled by state space methods. Appendix 5 documents the mixed-frequency factor model and the resulting state space representation. Draws from the predictive density of GDP growth can be obtained from a Gibbs Sampler which alternately draws from the posterior distribution of the factors conditional on the parameters and then updates the parameters given the sampled factors (see Appendix 6 for details). We iteratively draw from the conditional posterior distributions 15000 times, discarding the first 5000 draws as a burn-in. The rest of the Markov chain is thinned by storing every 10-th iteration, yielding a total of $G = 1000$ draws for posterior inference.[1]

[1] By and large, our Gibbs Sampler shows no signs of poor mixing or non-convergence. See Appendix 9 for detailed MCMC diagnostics.

In terms of model specification, we estimate the models for all R = 1 : 10; the number of lags in the factor VAR is set to one, i.e. P = 1. Furthermore, for each prior we combine the ten individual densities into an equally-weighted pool (denoted as "pool" below).

To benchmark the performance of the different priors, we estimate a simple univariate Bayesian autoregressive model of order 1, estimated with loose priors on the autocovariance coefficient - $\mathcal{N}(0,3)$ - and the residual variance - $\mathcal{G}^{-1}(1,0.01)$. This model is subsequently denoted as B-AR(1). Additionally, we also consider the factor model with a Normal prior on the loadings that features a relatively high and fixed variance, e.g. Amisano and Geweke (2017) and Marcellino et al. (2016). In the notation of Section 2, the prior precision D^p is given by $\frac{1}{c}I_R$ where c is some large number, independent of any hyperparameters. We set $c = 10$ and label this "diffuse" prior in what follows as Nd.

The hyperparameters of the remaing priors are as follows: The variance of the loadings in the Normal-Inverse Gamma (NIG) is parametrized as $\mathcal{G}^{-1}(2,1)$. For the Multiplicative-Gamma (MG) prior we choose $a_1 = 5$, $a_2 = 2$ and $w = 3$. In the case of the point mass-Normal mixture (PMNM) we set $r_0 = 5$, $s_0 = 0.5$, while the prior on τ_r is given by $\mathcal{G}^{-1}(2,1)$. As discussed above, the horseshoe prior (HS+) requires no hyperparameters.

3.3 Evaluation set-up

The evaluation period ranges from the first quarter of 2000 (2000Q1) to the fourth quarter of 2018 (2018Q4) in the US. Reflecting the shorter estimation sample for Germany (and lack of available real-time vintages earlier on), we begin to evaluate nowcasts starting in 2006Q1. In both cases, we focus on nowcasts made at the end of the first month of a given quarter e.g. January for nowcasts of GDP growth in Q1. That is to say, the forecast horizon in months is h = 2.

The accuracy of the different models/priors is assessed in terms of point and density nowcasts. For the former, we compute the root mean squared forecast error (RMSFE) of model/prior $m = \{\text{B-AR(1)}, \text{Nd}, \text{NIG}, \text{MG}, \text{PMNM}, \text{HS+}\}$ defined as

$$\text{RMSFE}_m = \left(\frac{1}{S} \sum_{s=1}^{S} (y_{m,T+s}^{f,Q} | \Omega_{v_s} - y_{T+s}^{Q})^2 \right)^{\frac{1}{2}}$$

for a sequence of S nowcasts, where $y_{m,T+s}^{f,Q}|\Omega_{v_s}$ is the mean of the predictive density of GDP growth at time $T+s$, conditional on the information set Ω_{v_s} available in real-time at date v_s when the nowcast for period $T+s$ was made. The corresponding realization of GDP growth which we take to be the first release is denoted by y_{T+s}^{Q} .[2]

Density nowcasts are evaluated by two scoring rules commonly used in the forecasting literature: the log score (logS) and the continuous ranked probability score (CRPS). The former is simply the predictive density of model m for period s evaluated at the realization

[2]Given that the mean absolute revision between first and second release are quite small for both countries, our results are robust to using the latter as the realization. Results are available upon request.

and then averaged over the evaluation period, i.e.

$$\text{logS}_m = \frac{1}{S}\sum_{s=1}^{S}\left(-log\, p_{m,T+s}(y^Q_{T+s})\right)$$

where $p_{m,T+s}(y^Q)$ is estimated from $y^{Q(1)}_{m,T+s}|\Omega_{v_s}, \ldots, y^{Q(G)}_{m,T+s}|\Omega_{v_s}$ using the theta kernel density estimator described in Botev et al. (2010). Furthermore, note that we flip the orientation of the log score to bring its interpretation in line with the RMSFE and CRPS: After premultiplying with -1, lower (higher) values of logS indicate a higher (lower) predictive accuracy.

The CRPS is calculated as (see Krüger et al., 2016):

$$\text{CRPS}_m = \frac{1}{S}\sum_{s=1}^{S}\left(\frac{1}{G}\sum_{i=1}^{G}\left|y^{Q(i)}_{m,T+s}|\Omega_{v_s} - y^Q_{T+s}\right| - \sum_{i=1}^{G}\sum_{j=1}^{G}\left|y^{Q(i)}_{m,T+s}|\Omega_{v_s} - y^{Q(j)}_{m,T+s}|\Omega_{v_s}\right|\right)$$

3.4 Results

The results of the nowcast evaluation for US and Germany GDP are presented in Table 1. For the sake of readability, we only show the results for a selected number of factors as well as the pooled nowcasts.[3] Overall, the factor models in both countries perform well with root mean squared forecast errors relative to the benchmark B-AR(1) as low as 0.7 and 0.5 for US and German GDP, respectively. Density nowcasts yield similar relative gains when evaluated in terms of the CRPS while they are generally even larger under the log score. For the United States, the predictive accuracy generally increases in the number of factors R - both in terms of point and density forecasts. For Germany, the models with $R = \{5, 8\}$ also perform much better than the model with only one factor but for $R = 2$ the performance is similar to the bigger models. Noteworthy is that in both countries the equal-weight pool performs very well, often achieving relative gains almost as large as those of the best individual specification.

Turning to the question of whether sparsity-inducing priors like the MG, PMNM or HS+ generate more accurate nowcasts than convential priors like the NIG or Nd, we find mixed evidence across the two countries. For the United States, the sparse priors do not perform better as the NIG prior. Indeed, for up to $R = 5$ there are virtually no differences between the five priors we consider across both point and density nowcasts. For $R = 8$ we find that a) the PMNM, HS and the NIG do perform slightly better than the Normal-diffuse prior in terms of all three forecast accuracy measures and b) the MG prior's performance deteriorates markedly. This weaker performance carries over to the equal-weight pool where the MG prior's poor performance stands out in terms of the log score and CRPS. To some extent these results hold for the case of Germany as well. There are also no differences between the sparse prior and the NIG. They do, however, outperform the Normal-diffuse prior by a larger amount when considering point and density nowcasts. In contrast to the findings

[3]The results for the remaining specifications are available upon request.

Table 1: Nowcasting results for US (top) and German (bottom) GDP

		full sample			post-crisis sample		
		RMSFE	logS	CRPS	RMSFE	logS	CRPS
		h=0	h=0	h=0	h=0	h=0	h=0
B-AR		0.45	0.64	0.25	0.32	0.50	0.19
R=1	Nd	0.85	0.79	0.89	1.11	0.98	1.07
	NIG	0.85	0.79	0.89	1.10	0.96	1.06
	MG	0.85	0.80	0.89	1.08	0.97	1.05
	PMNM	0.86	0.79	0.89	1.10	0.97	1.06
	HS+	0.86	0.82	0.90	1.10	1.02	1.07
R=2	Nd	0.90	0.82	0.92	1.10	0.95	1.06
	NIG	0.91	0.84	0.93	1.10	0.93	1.05
	MG	0.86	0.81	0.89	1.08	0.95	1.04
	PMNM	0.90	0.82	0.92	1.13	0.96	1.07
	HS+	0.85	0.78	0.89	1.11	0.91	1.05
R=5	Nd	0.82	0.72	0.85	1.03	0.80	0.99
	NIG	0.81	0.69	0.83	0.94	0.72	0.92
	MG	0.78	0.70	0.82	0.99	0.81	0.96
	PMNM	0.79	0.68	0.82	0.99	0.76	0.95
	HS+	0.79	0.69	0.83	1.01	0.78	0.97
R=8	Nd	0.76	0.65	0.79	1.02	0.79	0.97
	NIG	0.77	0.66	0.80	1.03	0.78	0.98
	MG	0.88	0.78	0.88	0.99	0.85	0.97
	PMNM	0.73	0.62	0.77	0.95	0.70	0.91
	HS+	0.74	0.60	0.77	0.93	0.67	0.90
pool	Nd	0.76	0.68	0.80	0.96	0.80	0.95
	NIG	0.77	0.68	0.80	0.92	0.77	0.91
	MG	0.80	0.72	0.83	0.98	0.83	0.96
	PMNM	0.75	0.68	0.79	0.92	0.79	0.92
	HS+	0.76	0.66	0.79	0.94	0.74	0.92

		full sample			post-crisis sample		
		RMSFE	logS	CRPS	RMSFE	logS	CRPS
		h=0	h=0	h=0	h=0	h=0	h=0
B-AR		0.86	1.73	0.41	0.49	0.89	0.29
R=1	Nd	0.86	0.90	0.89	0.86	0.95	0.93
	NIG	0.86	0.89	0.89	0.85	0.94	0.93
	MG	0.88	0.81	0.91	0.90	0.97	0.97
	PMNM	0.89	0.97	0.91	0.88	0.96	0.95
	HS+	0.89	1.00	0.92	0.90	0.97	0.97
R=2	Nd	0.57	0.40	0.65	0.79	0.62	0.78
	NIG	0.52	0.36	0.61	0.76	0.59	0.75
	MG	0.52	0.37	0.61	0.75	0.60	0.75
	PMNM	0.54	0.37	0.62	0.78	0.58	0.77
	HS+	0.54	0.38	0.63	0.77	0.62	0.77
R=5	Nd	0.65	0.46	0.74	0.90	0.75	0.89
	NIG	0.55	0.38	0.63	0.73	0.59	0.74
	MG	0.64	0.41	0.69	0.71	0.61	0.73
	PMNM	0.55	0.37	0.62	0.72	0.58	0.73
	HS+	0.53	0.37	0.61	0.68	0.57	0.70
R=8	Nd	0.56	0.42	0.66	0.85	0.73	0.85
	NIG	0.51	0.37	0.59	0.71	0.62	0.73
	MG	0.56	0.39	0.64	0.74	0.63	0.76
	PMNM	0.49	0.35	0.57	0.68	0.57	0.71
	HS+	0.49	0.33	0.57	0.67	0.54	0.69
pool	Nd	0.51	0.42	0.62	0.73	0.74	0.79
	NIG	0.47	0.37	0.57	0.63	0.64	0.71
	MG	0.54	0.37	0.61	0.67	0.63	0.72
	PMNM	0.48	0.36	0.56	0.62	0.61	0.69
	HS+	0.47	0.35	0.55	0.59	0.59	0.67

RMSFE is the root mean squared forecast error, logS and CRPS are the average log score and continuous ranked probability score. All entries for the factor models are relative to the B-AR benchmark (see text for details) and negatively orientated so that a value in the table below 1 corresponds to a better performance than the benchmark. The forecast horizon h is in months. For the US (Germany) the full sample period is 2000Q1 (2006Q1)-2018Q4. In both cases the post-crisis sample starts in 2010Q1 and ends in 2018Q4.

for US GDP, we do see that the nowcasts from the sparse model perform better than those from the NIG when evaluated in terms of the RMSFE and the log score or CRPS. Similar to the US, however, between the PMNM and HS+ priors there are virtually no differences in predictive accuracy while the MG's poorer performance stands out for larger R.

Lastly, we find that when excluding the Global Financial Crisis (GFC) from the evaluation the gains in predictive accuracy relative to the B-AR(1) are less pronounced for most factor models - for the United States some of the models do not outperform the autoregressive benchmark in terms of point nowcast accuracy. Our main findings with regard to the question whether sparsity matters for nowcasting, however, are unaltered. Focusing only on the post-crisis sample the sparse priors yield highly similar results in terms of all three performance measures considered. In the US, we still find that there are also no large differences between the sparse priors and the NIG or Nd in terms of predictive accuracy, while for German GDP the PMNM and HS+ priors continue to produce much better nowcasts than the Nd and somewhat better nowcasts than the NIG for $R = 8$ and the equal-weight pool.

3.5 Robustness analysis

The results discussed above were obtained under specific modelling/specifications choices. In our application, these concern the transformation applied to the survey indicators or the choice of estimation window. We find, however, that alternative specifications in this regard - such as not first-differencing the survey indicators or estimating the models with a rolling window - do not have a material impact on the results neither in quantitative nor qualitative terms. We discuss these choices and briefly comment on the results below. Details and tables similar to Table 1 for the different robustness checks can be found in Appendix 8.

First, in our baseline specification we included survey indicators in first differences. While this brings their time series behavior more in line with the rest of the dataset, an alternative is to keep the series in levels, as these indicators are by construction stationary. Examples of both approaches can be found in the literature. In our application we find that the nowcast performance overall deteriorates uniformly across models when the surveys are not first-differenced. However, the main findings are unaltered: the differences between the sparse priors and the NIG are very small. In the case of Germany the sparse priors and the NIG perform much better than the Normal-diffuse prior which for large R fails to beat the autoregressive benchmark when the surveys enter the models in levels. **Second**, while a recursively expanded estimation window leads to lower estimation uncertainty, a rolling estimation window might guard against structural instabilities in the forecasting models and therefore generate better nowcasts. However, there is no indication that these instabilities play a large role in our application as the nowcast performance only improves marginally in some cases for the United States, in particular for the smaller models, i.e. $R \leq 2$; in addition the MG prior performs much better under a rolling window. But the performance of the equal-weight pool, for example, is virtually identical. For Germany, with the exception of $R = 1$ the accuracy of point and density nowcasts is somewhat higher when the models'

parameters are estimated with a recursively expanding window. For both countries, the choice of estimation window does not have an impact on the performance of the sparse priors relative to the Normal alternatives.

4 Conclusion

In this paper, we have explored the role that sparsity plays in factor models. In a real-time nowcasting evaluation we find that estimating the model with sparse priors on the loadings matrix in a Bayesian framework does not lead to large gains in nowcast accuracy. Furthermore, we found very similar results for different sparse priors that have been proposed in the literature. This suggests that the practice in parts of the literature of considering large cross sections when nowcasting GDP is justified. Our findings are compatible with the hypothesis that large macroeconomic data sets typically used in now- or forecasting applications are not sparse but dense. However, we caution against generalising our findings too far, as sparsity has been shown to play a role in other macroeconomic settings (Kaufmann and Schumacher, 2017). Moreover, the recent Covid-19 pandemic has accelated a trend in macroeconomic forecasting applications of considering new, unconvential data originating from newspapers, social media, mobile phones or internet search queries. These data sources typically provide vast numbers of time series ("big data"). We leave it to future research to address the issue of variable selection and the role that sparsity plays in such an environment.

Appendix

5 Mixed-frequency factor model

In the following, we outline the mixed-frequency factor model used in the nowcasting application. For the reader's convenience we begin by restating the original model

$$\begin{aligned} x_t &= \lambda f_t + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \Sigma_\varepsilon) \\ \Psi(L) f_t &= \eta_t, \quad \eta_t \sim \mathcal{N}\left(0, \Sigma_\eta\right) \end{aligned} \tag{34}$$

where we now make it explicit that x_t is a vector of N_M stationary monthly variables. To combine the monthly time series with quarterly GDP growth, y_t^Q, we formulate the model at the highest frequency, e.g. monthly. However, we only observe y_t^Q every third period, e.g. at the last month of each quarter. Let T be the number of observations for which at least one monthly variable is available. Assuming for the sake of expostion that the sample starts in the first month a quarter and ends in the third, say January and December, we thus have $y_{1:T}^Q = \left[\text{NaN}, \text{NaN}, y_3^Q, \text{NaN}, \text{NaN}, y_6^Q, \ldots, \text{NaN}, \text{NaN}, y_T^Q\right]$.

Furthermore, we assume that y_t, the unobserved monthly analogue of y_t^Q, adheres to the same factor structure as x_t. That is to say, we assume that (unobserved) month-on-month GDP growth can be expressed as

$$y_t = \lambda^y f_t + \varepsilon_t^y, \quad \varepsilon_t^y \sim \mathcal{N}(0, \Sigma_{\varepsilon^y}). \tag{35}$$

The unobserved y_t are linked to the quarterly observations y_t^Q via the following time aggregation rule (Mariano and Murasawa, 2003):

$$y_t^Q \approx \frac{1}{3}y_t + \frac{2}{3}y_{t-1} + \frac{3}{3}y_{t-2} + \frac{2}{3}y_{t-3} + \frac{1}{3}y_{t-4}, \forall \, \mathrm{t} = 3, 6, 9, \ldots, T. \tag{36}$$

By putting together (34), (35) and (36), the model can be cast into state space form with measurement equation

$$\begin{bmatrix} x_{1,t} \\ \vdots \\ x_{N_m,t} \\ y_t^Q \end{bmatrix} = Z \begin{bmatrix} f_t \\ \vdots \\ f_{t-4} \\ \epsilon_{1,t}^y \\ \vdots \\ \epsilon_{1,t-4}^y \end{bmatrix} + \begin{bmatrix} \varepsilon_t \\ 0 \end{bmatrix}, \quad \begin{bmatrix} \varepsilon_t \\ 0 \end{bmatrix} \sim \mathcal{N}(0, H) \tag{37}$$

and transition equation

$$\begin{bmatrix} f_t \\ \vdots \\ f_{t-4} \\ \epsilon_{1,t}^y \\ \vdots \\ \epsilon_{1,t-4}^y \end{bmatrix} = T \begin{bmatrix} f_{t-1} \\ \vdots \\ f_{t-5} \\ \epsilon_{1,t-1}^y \\ \vdots \\ \epsilon_{1,t-5}^y \end{bmatrix} + R \begin{bmatrix} \eta_t \\ \epsilon_t^y \end{bmatrix}, \quad \begin{bmatrix} \eta_t \\ \epsilon_t^y \end{bmatrix} \sim \mathcal{N}(0, Q) \tag{38}$$

where

$$Z = \begin{bmatrix} \lambda & 0_{N_M \times R} & \cdots & \cdots & \cdots & 0_{N_M \times 1} & \cdots & \cdots & \cdots & \cdots \\ \frac{1}{3}\lambda^q & \frac{2}{3}\lambda^q & \frac{3}{3}\lambda^q & \frac{2}{3}\lambda^q & \frac{1}{3}\lambda^q & \frac{1}{3} & \frac{2}{3} & \frac{3}{3} & \frac{2}{3} & \frac{1}{3} \end{bmatrix}, \quad H = \begin{bmatrix} \Sigma_\varepsilon & 0_{N_M \times 1} \\ 0_{1 \times N_M} & 0 \end{bmatrix}$$

and

$$T = \begin{bmatrix} \psi_1 & 0_R & 0_{R \times 1} & \cdots \\ I_{4R} & 0_{4R \times R} & 0_{4R \times 1} & \cdots \\ 0_{1 \times R} & \cdots & 0 & 0_{1 \times 4} \\ 0_{4 \times R} & \cdots & I_4 & 0_{4 \times 1} \end{bmatrix}, \quad Q = \begin{bmatrix} \Sigma_\eta & 0_{Nm \times 1} \\ 0_{1 \times N_m} & \Sigma_{\varepsilon^y} \end{bmatrix}, \quad R = \begin{bmatrix} I_R & 0_{R \times 1} \\ 0_{4R \times R} & 0_{4R \times 1} \\ 0_{1 \times R} & 1 \\ 0_{4 \times R} & 0_{4 \times 1} \end{bmatrix}.$$

When running the Kalman filter or smoother, missing observations can easily be dealt with by either i) removing the missings elements from the vector of observations and ad-

justing the dimensions of Z and H or ii) replacing them with arbitrary values, say 0, and setting the corresponding element in H to a very large number. We found the former to be numerically more stable in our application.

6 Gibbs Sampler

The model is estimated by a Gibbs Sampler which alternately draws from the conditional posterior distribution of the parameters given the factors and the factors given the parameters. Draws from the predictive density of any variable of interest - in our application, quarterly GDP growth - can then be obtained conditional on these values. Let λ denote the $N \times R$ matrix of factor loadings belonging to the monthly variables while λ^y is the loading of unobserved monthly GDP. The remaining parameters of the model are collected in $\Theta = \{\Psi, \Sigma_\varepsilon, \Sigma_{\varepsilon^y}\}$. Furthermore, let φ denote the prior hyperparameters. $x_{1:T}$ is the $N \times T$ matrix of monthly observations - some of which may be missing due different publication delays or because a series has only been collected over part of the estimation sample - where T is the maximum number of periods for which at least one observation is available. Furthermore, $y^Q_{1:T}$ denotes quarterly GDP growth at the monthly frequency, where we follow the convention in the literature and assume that the quarterly observations are available in the last month of each quarter, otherwise, $y^Q_t = \text{NaN}$; the monthly analogue of quarterly GDP growth is denoted as $y_{1:T}$ (see Appendix 5 for more details on the temporal aggregation). The forecast horizon in months is denoted by H, i.e. $T + h$ corresponds to the third month of the quarter that is being nowcasted.

The Gibbs Sampler then cycles through the following steps or blocks to draw from the predictive density $p(y^Q_{T+h}|x_{1:T}, y^Q_{1:T})$:

Step 1: $p(f_{1:T}, y_{1:T}|\lambda, \lambda^y, \Theta, x_{1:T}, y^Q_{1:T})$
Conditional on the observed monthly and quarterly data, $x_{1:T}$ and $y^Q_{1:T}$, and a draw of the parameters, we can sample from the conditional posterior distribution of $f_{1:T}$ using the state space model described in (37) and (38) and the simulation smoother in Durbin and Koopman (2002). As a by-product of the sampled state vector, we also obtain a draw of the monthly analogues of the quarterly time series.

Step 2: $p(x^+_{1:T}|f_{1:T}, \lambda, \Sigma_\varepsilon)$
As some of the elements of $x_{1:T}$ may be missing, we also sample them conditional on the factors. This yields a complete data set $x^+_{1:T}$ which is then used in the subsequent steps. Given the normality of observations and states, the posterior distribution of a generic missing observation $x_{i,t}$ conditional on the factor is also Normal with mean $\lambda_i f_t$. Furthermore, its variance does not depend on the realizations of the conditioning arguments and is simply given by the i-th diagonal element of Σ_ε.

Step 3: $p(\Psi|f_{1:T})$
Conditional on a draw of the factors, we can sample the parameters of the factor VAR. Denote by $\psi^* = \text{vec}([\psi_1, \ldots, \psi_R])$ the vectorized matrix of coefficients. Then, given a Normal prior

$$p(\psi^*) \propto \det\left(\underline{V}_{\psi^*}\right)^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}(\psi^* - \underline{b}_{\psi^*})'\underline{V}_{\psi^*}^{-1}(\psi^* - \underline{b}_{\psi^*})\right\}$$

the (conditional) posterior is Normal with variance and mean given by

$$\bar{V}_{\psi^*} = \left[\underline{V}_{\psi^*}^{-1} + I_R \otimes X'X\right]^{-1}, \quad \bar{b}_{\psi^*} = \bar{V}_{\psi^*}\left\{\underline{V}_{\psi^*}^{-1}\underline{b}_{\psi^*} + I_R \otimes X'X\, b_{OLS}\right\}$$

where $X = [f_{P:T-1}, \ldots, f_{1:T-P}]$, $y = f_{P+1:T}$ and $b_{OLS} = \text{vec}\left((X'X)^{-1}X'y\right)$. We follow the literature on Bayesian VAR and impose a Minnesota type prior on $\psi*$. That is to say, the prior mean is set equal to 0 and the prior variance depends on the lag length. Specifically, for $p = 1 : P$, we have

$$Var(\psi_{r,ij}) = \begin{cases} \frac{\pi_0}{r^2}, & \text{if } i = j \\ \frac{\pi_0 \pi_1}{r^2}, & \text{otherwise.} \end{cases}$$

Common values in the literature are $\pi_0 = 0.2$ and $\pi_1 = 0.1$, thus shrinking coefficients on the lags of other factors stronger towards 0.

Step 4a: $p(\lambda, \lambda^y|f_{1:T}, x^+_{1:T}, y_{1:T}, \Sigma_\varepsilon, \Sigma_{\varepsilon^y}, \varphi)$
Conditional on the factors, data, parameters and hyperparameters we can sample the loadings from their conditional posterior distribution. See Section 2 of the main text for details.

Step 4b: $p(\varphi|\lambda, \lambda^y)$
Conditional on a draw of the loadings, we can update the hyperparameters of the different priors that govern the degree of sparsity. We again refer to the main text for the conditional posterior distributions.

Step 5: $p(\Sigma_\varepsilon, \Sigma_{\varepsilon^y}|\lambda, \lambda^y, f_{1:T}, x^+_{1:T}, y_{1:T})$
Conditional on the idiosyncratic components $\epsilon_{i,1:T}$ we can sample the diagonal elements of the covariance matrix Σ_ε by drawing from

$$\sigma_i^2 \sim \mathcal{G}^{-1}\left(\frac{u+T}{2}, \frac{U + \sum_{t=1}^T \epsilon_{i,t}^2}{2}\right)$$

where u and U are the prior shape and rate. In an analogous manner, we can sample Σ_{ε^y} given $\epsilon^y_{1:T}$. We set $u = 2$ and $U = 1$ so that the prior is centered around one but relatively diffuse.

Step 6: $p(y^Q_{T+h}|x_{1:T}, y^Q_{1:T})$
Draws from the h-step ahead predictive density of y^Q are obtained by iterating forward (34)

Table 2: List of US survey indicators

mnemonic	description	transformation
gacdna	current general activity	diff
gafdna	future general activity	diff
nocdna	current new orders	diff
nofdna	future new orders	diff
shcdna	current shipments	diff
shfdna	future shipments	diff
dtcdna	current delivery time	diff
dtfdna	future delivery time	diff
ivcdna	current inventories	diff
ivfdna	future inventories	diff
uocdna	current unfilled orders	diff
uofdna	future unfilled orders	diff
ppcdna	current prices paid	diff
ppfdna	future prices paid	diff
prcdna	current prices received	diff
prfdna	future prices received	diff
necdna	current employment	diff
nefdna	future employment	diff
awcdna	current workhours	diff
awfdna	future workhours	diff
cefdfna	future capital expenditures	diff

Notes: The table lists the survey indicators from the Federal Reserve Bank of Philadelphia's Manufacturing Business Outlook Survey that are used in nowcasting US GDP growth. For the historical data and further information, see `https://www.philadelphiafed.org/research-and-data/regional-economy/business-outlook-survey/historical-data`.

and using (35) and (36) to compute y^Q_{T+h} (Del Negro and Schorfheide, 2013).

7 Data

This section describes the data sets used in the empirical nowcasting application. The US data set (Section 7.A) is based on FRED-MD. Real-time vintages for Germany (Section 7.B) are compiled from the Deutsche Bundesbank's Real-Time Database. Both datasets are augmented with survey-based sentiment indicators; to guarantee the real-time nature of the evaluation, the raw, unadjusted series are seasonally adjusted in real-time using X11-ARIMA. Furthermore, while the datasets are broadly stable in terms of available series, there are some changes as vintages for some series are added or removed over the period of the evaluation sample. Additionally, for the estimation of the models we require that every variable is available for at least half of the estimation sample. As such, the exact composition of the data sets varies slightly over time. The individual series used in each vintage are available upon request.

7.A United States

Real-time vintages of GDP growth are obtained from ALFRED. Besides the target variable, our data set includes a large number of monthly time series covering various aspects of economic activity. Specifically, we use the large, monthly, real-time data set constructed by McCracken and Ng (2016) and regularly updated by the Federal Reserve Bank of St. Louis.[4] It includes series from categories such as output and income, prices, labor markets, housing and financial markets. Vintages of the "FRED-MD" data set are available from December 1999 onwards, reflecting the information available at the end of the respective month. Regarding transformations prior to estimation, we mainly follow the suggestions in McCracken and Ng (2016). However, given that our sample starts in 1985, some modifications are in order to reflect the shorter span of the time series. For example, in the original FRED-MD dataset some price series and average hourly earnings indicators are included in second (log) differences to achieve stationarity. Over our shorter, "Great Moderation" sample, the means and variances of the log differences are constant so that we do not need to difference these series twice. In contrast, housing starts and permits are included in log levels in the original dataset. Over the shorter sample used in the estimation of the nowcasting models, the log of the series exhibits large and persistent swings around the long-run mean. We therefore consider it more appropriate to difference the series to achvieve stationarity. These modifications bring our dataset more in line with the recent literature on nowcasting US GDP growth, e.g. Aastveit et al. (2018).

As the nowcasting literature has emphasized the importance of "soft", survey-based sentiment indicators, we supplement the McCracken and Ng (2016)-dataset with the Federal Reserve Bank of Philadelphia's Manufacturing Business Outlook Survey. While regional in nature, these indicators are available over a long period of time and provide potentially useful information for the gross domestic product as they are published very timely, usually in the middle of the reference month. We include all 20 series covering e.g. firms' assessement of current and future activity, orders, etc (Table 2). All in all, this yields a monthly real-time data set containing roughly 130 variables, though as mentioned above the exact number varies slightly over the course of the evaluation period.

7.B Germany

Nowcasts of German GDP growth are based on a real-time dataset comprised of over 100 monthly variables similar to that employed by Schumacher (2007). Vintages are compiled from the Deutsche Bundesbank Real-Time Database and augmented with financial market and the survey-based ifo indicators. Real-time vintages for a sufficiently large number of variables are available as of November 2005. For others, these become available over the course of the evaluation period. As a result, the size of the dataset increases somewhat from 167 monthly variables in the January 2006 vintage to 172 series as of December 2017. The

[4] All vintages can be downloaded from `https://research.stlouisfed.org/econ/mccracken/fred-databases/`.

Table 3: Real-time data, Germany: production

variable/sector	adj.	trafo	1st obs.	1st vintage
industrial production	p,s,c	3	Jan 1991	6.2.1995
industrial production and construction	p,s,c	3	Jan 1991	6.2.1995
total construction	p,s,c	3	Jan 2010	3.8.2013
main construction industry	p,s,c	3	Jan 1991	6.2.1995
finishing trade	p,s,c	3	Jan 2010	3.8.2013
building construction	p,s,c	3	Jan 1991	7.11.2005
civil engineering	p,s,c	3	Jan 1991	7.11.2005
industry	p,s,c	3	Jan 1991	6.2.1995
intermediate goods	p,s,c	3	Jan 1991	6.2.1995
investment goods	p,s,c	3	Jan 1991	6.2.1995
consumption goods	p,s,c	3	Jan 1991	6.2.1995
durable goods	p,s,c	3	Jan 1991	6.2.1995
non-durable goods	p,s,c	3	Jan 1991	6.2.1995
energy	p,s,c	3	Jan 1991	6.2.1995

Source: Deutsche Bundesbank.

Table 4: Real-time data, Germany: orders

variable/sector	adj.	trafo	1st obs.	1st vintage
industry	p,s,c	3	Jan 1991	6.9.2001
industry (domestic)	p,s,c	3	Jan 1991	6.9.2001
industry (abroad)	p,s,c	3	Jan 1991	6.9.2001
intermediate goods	p,s,c	3	Jan 1991	8.6.1995
intermediate goods (domestic)	p,s,c	3	Jan 1991	8.6.1995
intermediate goods (abroad)	p,s,c	3	Jan 1991	8.6.1995
investment goods	p,s,c	3	Jan 1991	8.6.1995
investment goods (domestic)	p,s,c	3	Jan 1991	8.6.1995
investment goods (abroad)	p,s,c	3	Jan 1991	8.6.1995
consumption goods	p,s,c	3	Jan 1991	8.6.1995
consumption goods (domestic)	p,s,c	3	Jan 1991	8.6.1995
consumption goods (abroad)	p,s,c	3	Jan 1991	8.6.1995
building construction	p,s,c	3	Jan 1991	22.11.2005
civil engineering	p,s,c	3	Jan 1991	22.11.2005
residential construction	p,s,c	3	Jan 1991	22.11.2005
construction industry (private)	p,s,c	3	Jan 1991	22.11.2005
construction industry (public)	p,s,c	3	Jan 1991	22.11.2005

Source: Deutsche Bundesbank.

starting point of our sample is January 1992.

Below we list the time-series comprising the seven groups of our dataset. The second column refers to the type of adjustment that have been applied to the time series (**p**rice-adjusted, **s**easonally adjusted, **c**alendar adjusted) while the third columns lists how each series is transformed to achieve stationarity (3 = difference of logarithm, 2 = difference, 1 = level). Furthermore, we also highlight if a series has a later starting point than January 1991 or if the vintages do not go back as far as November 2005 to indicate that the variable is only included in part of the evaluation period.

Table 5: Real-time data, Germany: turnover

variable/sector	adj.	trafo	1st obs.	1st vintage
industry	s,c	3	Jan 1991	4.11.2005
industry (domestic)	s,c	3	Jan 1991	4.11.2005
industry (abroad)	s,c	3	Jan 1991	4.11.2005
intermediate goods	s,c	3	Jan 1991	4.11.2005
intermediate goods (domestic)	s,c	3	Jan 1991	4.11.2005
intermediate goods (abroad)	s,c	3	Jan 1991	4.11.2005
investment goods	s,c	3	Jan 1991	4.11.2005
investment goods (domestic)	s,c	3	Jan 1991	4.11.2005
investment goods (abroad)	s,c	3	Jan 1991	4.11.2005
consumption goods	s,c	3	Jan 1991	4.11.2005
consumption goods (domestic)	s,c	3	Jan 1991	4.11.2005
consumption goods (abroad)	s,c	3	Jan 1991	4.11.2005
durable goods	s,c	3	Jan 1991	4.11.2005
non-durable goods	s,c	3	Jan 1991	4.11.2005
residential construction	s,c	3	Jan 1991	22.11.2005
construction industry (private)	s,c	3	Jan 1991	20.11.2005
construction (public)	s,c	3	Jan 1991	4.11.2005
retail sales	p,s,c	3	Jan 1994	17.11.2005
retail sales excluding cars	p,s,c	3	Jan 1994	17.11.2005
retail sales: cars	p,s,c	3	Jan 1994	17.11.2005

Source: Deutsche Bundesbank.

Table 6: Real-time data, Germany: prices

variable/sector	adj.	trafo	1st obs.	1st vintage
consumper price index (CPI)	s,c	3	Jan 1991	28.11.1995
CPI excl. energy	s	3	Jan 1991	28.1.1999
CPI excl. energy and food	s	3	Jan 1995	13.4.2017
CPI: food	s	3	Jan 1991	28.1.1999
CPI: other non-durables and durables	s	3	Jan 1995	14.11.2003
CPI: energy	s,c	3	Jan 1991	29.2.2008
CPI: services	s,c	3	Jan 2010	13.4.2017
CPI: services (excluding rents)	s	3	Jan 2000	14.11.2003
CPI: rents	s	3	Jan 1991	11.11.2005
CPI: rents excl. ancillary costs	s	3	Jan 1995	16.4.2008
produer price index (PPI): industrial products	s	3	Jan 1991	18.11.2005
PPI: industrial products excl. energy	s	3	Jan 1994	18.11.2005
PPI: agricultural products	s	3	Jan 1968	7.11.2005
export price index	s	3	Jan 1970	24.11.2005
import price index	s	3	Jan 1970	24.11.2005

Source: Deutsche Bundesbank.

Table 7: Real-time data, Germany: labor market

variable/sector	adj.	trafo	1st obs.	1st vintage
employment	s	3	Jan 1991	8.8.1995
employment: manufacturing and mining	s	3	Jan 1991	16.11.2005
employment: main construction industry	s	3	Jan 1991	22.11.2005
hours worked	s,c	3	Jan 1991	22.11.2005
hours worked: manufacturing and mining	s,c	3	Jan 1991	16.11.2005
hours worked: main construction industry	s,c	3	Jan 1991	22.11.2005
employees subject to social security contributions	s	3	Jan 1991	30.3.2006
gross wages and salaries: manufacturing and mining	s,c	3	Jan 1991	16.11.2005
gross wages and salaries: main construction industry	s,c	3	Jan 1991	22.11.2005

Source: Deutsche Bundesbank.

Table 8: Real-time data, Germany: financial markets

variable/sector	adj.	trafo	1st obs.	1st vintage
yields on debt securities issued by residents	none	2	Jan 1991	-
yields: bank debt securities	none	2	Jan 1991	-
yields: mortgage Pfandbriefe	none	2	Jan 1991	-
yields: public Pfandbriefe securities	none	2	Jan 1991	-
yields: special purpose credit institutions	none	2	Jan 1991	-
yields: other bank debt securities	none	2	Jan 1991	-
yields: corporate debt securities	none	2	Jan 1991	-
yields: public debt securities	none	2	Jan 1991	-
yields: state government securities	none	2	Jan 1991	-
government bond yields: 6 month maturity	none	2	Jan 1991	-
government bond yields: 1 year maturity	none	2	Jan 1991	-
government bond yields: 2 year maturity	none	2	Jan 1991	-
government bond yields: 3 year maturity	none	2	Jan 1991	-
government bond yields: 4 year maturity	none	2	Jan 1991	-
government bond yields: 5 year maturity	none	2	Jan 1991	-
government bond yields: 6 year maturity	none	2	Jan 1991	-
government bond yields: 7 year maturity	none	2	Jan 1991	-
government bond yields: 8 year maturity	none	2	Jan 1991	-
government bond yields: 9 year maturity	none	2	Jan 1991	-
government bond yields: 10 year maturity	none	2	Jan 1991	-
CDAX index	none	3	Jan 1991	-
Nominal effective exchange rate (narrow)	none	3	Jan 1991	-
Nominal effective exchange rate (broad)	none	3	Jan 1991	-

Source: Deutsche Bundesbank.

Table 9: Real-time data, Germany: survey indicators

variable/sector	adj.	trafo	1st obs.	1st vintage
ifo: manufacturing, current situation	s	2	Jan 1991	-
ifo: manufacturing, climate	s	2	Jan 1991	-
ifo: manufacturing, expectations	s	2	Jan 1991	-
ifo: manufacturing, demand	s	2	Jan 1991	-
ifo: manufacturing, prices	s	2	Jan 1991	-
ifo: manufacturing, employment expectations	s	2	Jan 1991	-
ifo: manufacturing, export expectations	s	2	Jan 1991	-
ifo: manufacturing, orders from abroad	s	2	Jan 1991	-
ifo: manufacturing, inventories	s	2	Jan 1991	-
ifo: manufacturing, orders	s	2	Jan 1991	-
ifo: manufacturing, production expecations	s	2	Jan 1991	-
ifo: manufacturing, orders (m/m)	s	2	Jan 1991	-
ifo: manufacturing, production (m/m)	s	2	Jan 1991	-
ifo: manufacturing, price expectation	s	2	Jan 1991	-
ifo: wholesale, current situation	s	2	Jan 1991	-
ifo: wholesale, climate	s	2	Jan 1991	-
ifo: wholesale, expectations	s	2	Jan 1991	-
ifo: wholesale, employment expectations	s	2	Jan 1991	-
ifo: wholesale, inventories	s	2	Jan 1991	-
ifo: wholesale, order expectations	s	2	Jan 1991	-
ifo: wholesale, price expectations	s	2	Jan 1991	-
ifo: wholesale, inventories	s	2	Jan 1991	-
ifo: wholesale, prices (m/m)	s	2	Jan 1991	-
ifo: wholesale, turnover (m/m)	s	2	Jan 1991	-
ifo: retail, current situation	s	2	Jan 1991	-
ifo: retail, climate	s	2	Jan 1991	-
ifo: retail, expectations	s	2	Jan 1991	-
ifo: retail, employment expectations	s	2	Jan 1991	-
ifo: retail, inventories	s	2	Jan 1991	-
ifo: retail, order expectations	s	2	Jan 1991	-
ifo: retail, price expectations	s	2	Jan 1991	-
ifo: retail, inventories	s	2	Jan 1991	-
ifo: retail, prices (m/m)	s	2	Jan 1991	-
ifo: retail, turnover (m/m)	s	2	Jan 1991	-
ifo: construction, current situation	s	2	Jan 1991	-
ifo: construction, climate	s	2	Jan 1991	-
ifo: construction, expectations	s	2	Jan 1991	-
ifo: construction, prices	s	2	Jan 1991	-
ifo: construction, capacity utilisation	s	2	Jan 1991	-

All series are seasonally adjusted in real-time using X-11-ARIMA. *Source*: ifo-Institute.

8 Additional results

This section presents the results of the nowcasting evaluation for several robustness checks discussed in Section 3.5.

United States

surveys in first differences, rolling estimation sample, P=1, first release
surveys in levels, recursive estimation sample, P=1, first release
surveys in levels, rolling estimation sample, P=1, first release

Germany

surveys in first differences, rolling estimation sample, P=1, first release
surveys in levels, recursive estimation sample, P=1, first release
surveys in levels, rolling estimation sample, P=1, first release

Table 10: Additional results: United States (first, diff, rolling)

		full sample			post-crisis sample		
		RMSFE	logS	CRPS	RMSFE	logS	CRPS
		h=0	h=0	h=0	h=0	h=0	h=0
B-AR		0.47	0.69	0.26	0.32	0.55	0.20
R=1	Nd	0.81	0.82	0.86	1.04	0.97	1.02
	NIG	0.92	0.89	0.93	1.05	0.94	1.02
	MG	0.84	0.82	0.88	1.09	0.97	1.05
	PMNM	0.86	0.85	0.89	1.06	0.95	1.03
	HS+	0.84	0.83	0.88	1.03	0.99	1.03
R=2	Nd	0.86	0.83	0.89	0.96	0.89	0.96
	NIG	0.89	0.85	0.91	0.99	0.87	0.98
	MG	0.84	0.80	0.87	0.98	0.88	0.97
	PMNM	0.84	0.81	0.88	1.05	0.93	1.02
	HS+	0.84	0.80	0.87	1.02	0.92	1.00
R=5	Nd	0.81	0.80	0.87	1.08	0.87	1.03
	NIG	0.79	0.75	0.84	1.12	0.86	1.05
	MG	0.79	0.72	0.83	1.12	0.85	1.04
	PMNM	0.79	0.74	0.84	1.09	0.87	1.03
	HS+	0.79	0.73	0.84	1.10	0.84	1.03
R=8	Nd	0.84	0.88	0.89	1.04	0.97	1.03
	NIG	0.80	0.75	0.83	1.05	0.82	1.00
	MG	0.75	0.70	0.80	1.02	0.78	0.98
	PMNM	0.72	0.66	0.77	0.87	0.66	0.87
	HS+	0.74	0.66	0.78	0.94	0.69	0.90
pool	Nd	0.80	0.85	0.86	1.03	0.95	1.01
	NIG	0.76	0.74	0.81	0.99	0.81	0.96
	MG	0.75	0.71	0.80	1.00	0.82	0.97
	PMNM	0.74	0.70	0.79	0.95	0.76	0.93
	HS+	0.77	0.70	0.81	1.01	0.77	0.96

RMSFE is the root mean squared forecast error, logS and CRPS are the average log score and continuous ranked probability score. All entries for the factor models are relative to the B-AR benchmark (see text for details) and negatively orientated so that a value in the table below 1 corresponds to a better performance than the benchmark. The forecast horizon h is in months. The full sample period is 2000Q1-2018Q4, the post-crisis sample starts in 2010Q1 and ends in 2018Q4.

Table 11: Additional results: United States (first, level, rec)

		full sample			post-crisis sample		
		RMSFE	logS	CRPS	RMSFE	logS	CRPS
		h=0	h=0	h=0	h=0	h=0	h=0
B-AR		0.45	0.64	0.25	0.32	0.50	0.19
R=1	Nd	1.04	1.13	1.05	0.99	1.07	1.03
	NIG	1.04	1.08	1.04	0.97	1.08	1.01
	MG	1.03	1.10	1.03	0.99	1.09	1.03
	PMNM	1.02	1.09	1.02	0.92	1.04	0.98
	HS+	1.01	1.10	1.01	0.91	1.04	0.98
R=2	Nd	0.93	0.89	0.96	1.20	1.04	1.14
	NIG	0.95	0.93	0.99	1.21	1.06	1.15
	MG	0.94	0.93	0.98	1.19	1.05	1.13
	PMNM	0.94	0.89	0.97	1.13	0.96	1.08
	HS+	0.93	0.88	0.96	1.15	0.98	1.09
R=5	Nd	0.83	0.75	0.86	1.07	0.86	1.01
	NIG	0.87	0.80	0.90	1.19	0.98	1.11
	MG	0.83	0.76	0.86	0.97	0.81	0.95
	PMNM	0.87	0.79	0.90	1.19	1.00	1.13
	HS+	0.86	0.77	0.89	1.19	0.97	1.12
R=8	Nd	0.81	0.72	0.85	1.02	0.80	0.98
	NIG	0.80	0.68	0.83	1.02	0.78	0.98
	MG	0.91	0.85	0.93	1.04	0.91	1.00
	PMNM	0.80	0.69	0.84	1.01	0.76	0.97
	HS+	0.78	0.67	0.82	1.03	0.78	0.98
pool	Nd	0.80	0.75	0.84	1.02	0.88	0.99
	NIG	0.81	0.76	0.85	1.03	0.87	1.00
	MG	0.89	0.84	0.91	1.05	0.95	1.02
	PMNM	0.80	0.75	0.85	1.02	0.86	0.99
	HS+	0.80	0.74	0.84	1.01	0.84	0.98

RMSFE is the root mean squared forecast error, logS and CRPS are the average log score and continuous ranked probability score. All entries for the factor models are relative to the B-AR benchmark (see text for details) and negatively orientated so that a value in the table below 1 corresponds to a better performance than the benchmark. The forecast horizon h is in months. The full sample period is 2000Q1-2018Q4, the post-crisis sample starts in 2010Q1 and ends in 2018Q4.

Table 12: Additional results: United States (first, level, rolling)

		full sample			post-crisis sample		
		RMSFE	logS	CRPS	RMSFE	logS	CRPS
		h=0	h=0	h=0	h=0	h=0	h=0
B-AR		0.47	0.69	0.26	0.32	0.55	0.20
R=1	Nd	0.88	0.97	0.93	0.87	1.05	0.98
	NIG	0.87	0.96	0.92	0.87	1.07	0.99
	MG	0.87	0.96	0.93	0.87	1.07	0.98
	PMNM	1.00	1.06	1.02	0.90	1.10	1.00
	HS+	0.99	1.08	1.01	0.88	1.09	0.99
R=2	Nd	0.84	0.83	0.88	0.97	0.86	0.96
	NIG	0.84	0.83	0.89	0.97	0.87	0.96
	MG	0.84	0.83	0.88	0.97	0.88	0.97
	PMNM	0.91	0.88	0.94	1.07	0.94	1.03
	HS+	0.87	0.83	0.90	1.01	0.87	0.98
R=5	Nd	0.80	0.76	0.85	0.92	0.70	0.90
	NIG	0.74	0.68	0.79	0.91	0.69	0.89
	MG	0.74	0.67	0.79	0.95	0.70	0.92
	PMNM	0.86	0.78	0.89	1.04	0.80	0.99
	HS+	0.80	0.72	0.84	1.00	0.74	0.95
R=8	Nd	0.86	0.92	0.92	1.11	0.94	1.06
	NIG	0.77	0.75	0.82	1.01	0.80	0.97
	MG	0.74	0.71	0.80	1.01	0.77	0.96
	PMNM	0.81	0.75	0.85	1.04	0.81	1.00
	HS+	0.74	0.66	0.79	0.96	0.73	0.93
pool	Nd	0.75	0.84	0.83	0.92	0.88	0.95
	NIG	0.75	0.76	0.81	0.90	0.79	0.91
	MG	0.73	0.74	0.79	0.90	0.80	0.92
	PMNM	0.80	0.78	0.85	0.97	0.86	0.96
	HS+	0.78	0.75	0.82	0.93	0.79	0.92

RMSFE is the root mean squared forecast error, logS and CRPS are the average log score and continuous ranked probability score. All entries for the factor models are relative to the B-AR benchmark (see text for details) and negatively orientated so that a value in the table below 1 corresponds to a better performance than the benchmark. The forecast horizon h is in months. The full sample period is 2000Q1-2018Q4, the post-crisis sample starts in 2010Q1 and ends in 2018Q4.

Table 13: Additional results: Germany (first, diff, rolling)

		full sample			post-crisis sample		
		RMSFE	logS	CRPS	RMSFE	logS	CRPS
		h=0	h=0	h=0	h=0	h=0	h=0
B-AR		0.91	1.35	0.45	0.52	1.04	0.32
R=1	Nd	0.73	0.69	0.78	0.96	0.68	0.88
	NIG	0.72	0.69	0.77	0.92	0.66	0.85
	MG	0.71	0.66	0.76	0.92	0.65	0.85
	PMNM	0.73	0.76	0.78	0.93	0.66	0.86
	HS+	0.72	0.70	0.76	0.91	0.62	0.83
R=2	Nd	0.60	0.58	0.66	0.85	0.60	0.80
	NIG	0.54	0.51	0.60	0.76	0.53	0.72
	MG	0.55	0.51	0.61	0.78	0.54	0.73
	PMNM	0.56	0.52	0.62	0.81	0.55	0.75
	HS+	0.56	0.54	0.63	0.81	0.57	0.75
R=5	Nd	0.77	0.64	0.78	0.97	0.71	0.89
	NIG	0.58	0.50	0.61	0.69	0.50	0.66
	MG	0.60	0.50	0.63	0.69	0.48	0.66
	PMNM	0.59	0.53	0.63	0.73	0.53	0.70
	HS+	0.61	0.53	0.64	0.70	0.50	0.68
R=8	Nd	0.66	0.65	0.72	1.04	0.78	0.94
	NIG	0.55	0.49	0.58	0.74	0.53	0.69
	MG	0.58	0.49	0.60	0.72	0.50	0.68
	PMNM	0.54	0.49	0.59	0.66	0.52	0.66
	HS+	0.62	0.52	0.63	0.67	0.50	0.65
pool	Nd	0.59	0.59	0.65	0.82	0.67	0.79
	NIG	0.54	0.49	0.58	0.68	0.52	0.66
	MG	0.55	0.49	0.59	0.68	0.51	0.66
	PMNM	0.57	0.52	0.61	0.71	0.55	0.69
	HS+	0.57	0.51	0.61	0.67	0.51	0.66

RMSFE is the root mean squared forecast error, logS and CRPS are the average log score and continuous ranked probability score. All entries for the factor models are relative to the B-AR benchmark (see text for details) and negatively orientated so that a value in the table below 1 corresponds to a better performance than the benchmark. The forecast horizon h is in months. The full sample period is 2006Q1-2018Q4, the post-crisis sample starts in 2010Q1 and ends in 2018Q4.

Table 14: Additional results: Germany (first, level, rec)

		full sample			post-crisis sample		
		RMSFE	logS	CRPS	RMSFE	logS	CRPS
		h=0	h=0	h=0	h=0	h=0	h=0
B-AR		0.86	1.73	0.41	0.49	0.89	0.29
R=1	Nd	0.88	0.66	0.98	1.10	1.01	1.09
	NIG	0.88	0.70	0.98	1.10	1.01	1.10
	MG	0.88	0.66	0.98	1.10	1.00	1.08
	PMNM	0.88	0.66	0.97	1.09	1.01	1.08
	HS+	0.88	0.69	0.97	1.06	1.01	1.07
R=2	Nd	0.87	0.65	0.92	1.02	0.99	1.01
	NIG	0.74	0.53	0.79	0.82	0.85	0.88
	MG	0.70	0.51	0.75	0.81	0.84	0.87
	PMNM	0.75	0.53	0.80	0.82	0.85	0.87
	HS+	0.76	0.53	0.80	0.81	0.85	0.87
R=5	Nd	1.01	0.86	1.14	1.18	0.96	1.08
	NIG	0.74	0.53	0.82	0.92	0.74	0.87
	MG	0.68	0.53	0.81	1.05	0.92	1.04
	PMNM	0.74	0.53	0.82	0.82	0.69	0.80
	HS+	0.71	0.50	0.79	0.85	0.70	0.82
R=8	Nd	1.17	0.87	1.17	1.28	1.04	1.21
	NIG	0.68	0.51	0.75	0.92	0.76	0.89
	MG	0.73	0.52	0.80	0.86	0.80	0.88
	PMNM	0.51	0.37	0.60	0.75	0.62	0.76
	HS+	0.51	0.36	0.59	0.72	0.59	0.73
pool	Nd	0.69	0.55	0.82	0.85	0.88	0.90
	NIG	0.55	0.45	0.66	0.73	0.76	0.80
	MG	0.59	0.48	0.70	0.82	0.83	0.87
	PMNM	0.51	0.44	0.63	0.61	0.70	0.72
	HS+	0.50	0.42	0.61	0.65	0.68	0.73

RMSFE is the root mean squared forecast error, logS and CRPS are the average log score and continuous ranked probability score. All entries for the factor models are relative to the B-AR benchmark (see text for details) and negatively orientated so that a value in the table below 1 corresponds to a better performance than the benchmark. The forecast horizon h is in months. The full sample period is 2006Q1-2018Q4, the post-crisis sample starts in 2010Q1 and ends in 2018Q4.

Table 15: Additional results: Germany (first, level, rolling)

		full sample			post-crisis sample		
		RMSFE	logS	CRPS	RMSFE	logS	CRPS
		h=0	h=0	h=0	h=0	h=0	h=0
B-AR		0.91	1.35	0.45	0.52	1.04	0.32
R=1	Nd	0.89	1.07	0.97	1.09	0.98	1.07
	NIG	0.88	1.11	0.97	1.08	0.99	1.07
	MG	0.89	1.12	0.96	1.07	0.96	1.04
	PMNM	0.87	1.20	0.95	1.01	0.96	1.01
	HS+	0.87	1.20	0.94	1.01	0.95	1.01
R=2	Nd	0.76	0.77	0.83	0.99	0.76	0.93
	NIG	0.65	0.68	0.71	0.89	0.76	0.87
	MG	0.65	0.69	0.71	0.88	0.77	0.86
	PMNM	0.73	0.73	0.79	0.94	0.77	0.90
	HS+	0.65	0.68	0.70	0.87	0.72	0.84
R=5	Nd	0.97	1.00	1.05	1.23	0.97	1.17
	NIG	0.62	0.64	0.68	0.77	0.64	0.76
	MG	0.60	0.61	0.66	0.73	0.62	0.73
	PMNM	0.62	0.61	0.67	0.71	0.59	0.71
	HS+	0.58	0.63	0.65	0.70	0.61	0.71
R=8	Nd	0.87	0.85	0.96	1.39	1.08	1.32
	NIG	0.58	0.55	0.64	0.82	0.63	0.78
	MG	0.61	0.56	0.66	0.76	0.58	0.73
	PMNM	0.57	0.55	0.62	0.70	0.55	0.69
	HS+	0.62	0.57	0.66	0.70	0.53	0.68
pool	Nd	0.61	0.77	0.75	0.90	0.92	0.94
	NIG	0.53	0.58	0.61	0.66	0.67	0.72
	MG	0.56	0.60	0.63	0.67	0.65	0.71
	PMNM	0.53	0.58	0.60	0.62	0.63	0.68
	HS+	0.55	0.59	0.62	0.63	0.61	0.68

RMSFE is the root mean squared forecast error, logS and CRPS are the average log score and continuous ranked probability score. All entries for the factor models are relative to the B-AR benchmark (see text for details) and negatively orientated so that a value in the table below 1 corresponds to a better performance than the benchmark. The forecast horizon h is in months. The full sample period is 2006Q1-2018Q4, the post-crisis sample starts in 2010Q1 and ends in 2018Q4.

9 MCMC diagnostics

To assess the performance of the Gibbs Sampler in terms of convergence and mixing, we calculate the inefficiency factors of the draws from the predictive density in the nowcasting application in Section 3. The inefficiency factor is defined as (e.g. Chib, 2011):

$$\text{Ineff}(\hat{h}_G) = \frac{Var(\hat{h}_G)}{s^2/G} \tag{39}$$

where G is the length of the chain, s^2 its sample variance and $Var(\hat{h}_G)$ an estimate of the variance of the simulation error taking into account autocorrelation in the Markov chain. For independent draws from the posterior distribution the inefficiency factor is equal to 1; higher values of Ineff($\hat{h}_G$) thus signal autocorrelation in the chain as result of poor mixing or lack of convergence. An estimate of the inefficiency factors is given by the sum of the autocorrelation coefficients ρ of the posterior draws, i.e.

$$\text{Ineff}(\hat{h}_G) = 1 + 2\sum_{l}^{L} \rho_l \tag{40}$$

where L is some suitably chosen upper bound. Alternatively, the inefficiency factor can also be expressed as

$$\text{Ineff}(\hat{h}_G) = \frac{G}{ESS(\hat{h}_G)} \tag{41}$$

where $ESS(\hat{h}_G)$ is the "effective sample size" of the Markov chain. The latter can readily be computed using the R package **coda** (Plummer et al., 2006).

Figure 2 shows boxplots of the inefficiency factors for the predictive densities from the different priors and model specifications. The majority of inefficiency factors of the predictive densities are concentrated at the lower end of the scale, indicating close to independent draws from the posterior distribution. There are a few outliers - defined as any observations that exceeds the median by 1.5 times the interquartile range and highlighted in the plot by a dot - for each specification, with some inefficiency factors as high as 50 or even 100 in one case. But overall the chains appear to mix well. An exception is the Multiplicative Gamma prior in the case of the United States where the number of outliers is much larger. We note, however, that this is only the case for models with $P = 1$, a recursive estimation sample and the survey indicators in levels. Furthermore, judging from the tables presented in Appendix 8 this does not seem to have impacted the nowcasting performance in a substantial way as the results are qualitatively and quantitatively similar to those obtained for the other specifications where the inefficiency factors are smaller.

Figure 2: Boxplots of inefficiency factors for different model specifications

Note: The figure shows the inefficiency factors for the different model specifications. *rec/rolling*: recursive or rolling estimation window, *level/diff*: survey indicators in levels or first differences. The different priors are denoted as follows: HS+ = horseshoe plus, MG = multiplicative Gamma, NIG = Normal Inverse Gamma, Nd = Normal diffuse, PMNM = point mass normal mixture. The number of lags in the factor VAR equals $P = 1$. For details, see the main text.

References

Aastveit, Knut Are, Francesco Ravazzolo, and Herman K. van Dijk, "Combined Density Nowcasting in an Uncertain Economic Environment," *Journal of Business & Economic Statistics*, 2018, *36* (1), 131–145.

_ , Karsten R. Gerdrup, Anne Sofie Jore, and Leif Anders Thorsrud, "Nowcasting GDP in Real Time: A Density Combination Approach," *Journal of Business & Economic Statistics*, January 2014, *32* (1), 48–68.

Alvarez, Rocio, Maximo Camacho, and Gabriel Perez-Quiros, "Aggregate versus disaggregate information in dynamic factor models," *International Journal of Forecasting*, 2016, *32* (3), 680–694.

Amisano, Gianni and John Geweke, "Prediction Using Several Macroeconomic Models," *The Review of Economics and Statistics*, 2017, *99* (5), 912–925.

Bai, Jushan and Serena Ng, "Forecasting economic time series using targeted predictors," *Journal of Econometrics*, 2008, *146* (2), 304–317. Honoring the research contributions of Charles R. Nelson.

Banbura, Marta and Gerhard Rünstler, "A look into the factor model black box: Publication lags and the role of hard and soft data in forecasting GDP," *International Journal of Forecasting*, 2011, *27* (2), 333–346.

Bhadra, Anindya, Jyotishka Datta, Nicholas G. Polson, and Brandon Willard, "The horseshoe+ estimator of ultra-sparse signals," *Bayesian Analysis*, 2017, *12* (4), 1105–1131.

Bhattacharya, Anirban and David B. Dunson, "Sparse Bayesian infinite factor models," *Biometrika*, 2011, *98*, 291–306.

_ , Debdeep Pati, Natesh S. Pillai, and David B. Dunson, "Dirichlet–Laplace priors for optimal shrinkage," *Journal of the American Statistical Association*, 2015, *110* (512), 1479–1490.

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

Bok, Brandyn, Daniele Caratelli, Domenico Giannone, Argia Sbordone, and Andrea Tambalotti, "Macroeconomic Nowcasting and Forecasting with Big Data," *Annual Review of Economics*, 08 2018, *10*.

Botev, Zdravko I., Joseph F. Grotowski, and Dirk P. Kroese, "Kernel density estimation via diffusion," *Annals of Statistics*, 10 2010, *38* (5), 2916–2957.

Caggiano, Giovanni, George Kapetanios, and Vincent Labhard, "Are more data always better for factor analysis? Results for the euro area, the six largest euro area countries and the UK," *Journal of Forecasting*, 2011, *30* (8), 736–752.

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.

Carvalho, Carlos M., Jeffrey Chang, Joseph E. Lucas, Joseph R. Nevins, Quanli Wang, and Mike West, "High-dimensional sparse factor modeling: applications in gene expression genomics," *Journal of the American Statistical Association*, 2008, *103* (484), 1438–1456.

_ , Nicholas G. Polson, and James G. Scott, "Handling sparsity via the horseshoe," *Journal of Machine Learning Research*, 2009, *5*, 73–80.

_ , _ , and _ , "The horseshoe estimator for sparse signals," *Biometrika*, 2010, *97* (2), 465–480.

Castillo, Ismaël and Aad van der Vaart, "Needles and straw in a haystack: Posterior concentration for possibly sparse sequences," *The Annals of Statistics*, 2012, *40* (4), 2069–2101.

Chan, Joshua C.C., "Minnesota-type adaptive hierarchical priors for large Bayesian VARs," *International Journal of Forecasting*, 2021.

Chib, Siddhartha, "Introduction to Simulation and MCMC Methods," in "The Oxford Handbook of Bayesian Econometrics" 2011, pp. 183–217.

Cross, Jamie L., Chenghan Hou, and Aubrey Poon, "Macroeconomic forecasting with large Bayesian VARs: Global-local priors and the illusion of sparsity," *International Journal of Forecasting*, 2020, *36* (3), 899–915.

Del Negro, Marco and Frank Schorfheide, "DSGE Model-Based Forecasting," in Graham Elliott and Allan Timmermann, eds., *Handbook of Economic Forecasting*, Vol. 2 of *Handbook of Economic Forecasting*, Elsevier, 2013, pp. 57–140.

Durbin, James and Siem Jan Koopman, "A simple and efficient simulation smoother for state space time series analysis," *Biometrika*, 2002, *89* (3), 603–616.

George, Edward I. and Robert E. McCulloch, "Variable selection via Gibbs sampling," *Journal of the American Statistical Association*, 1993, *88* (423), 881–889.

_ and _ , "Approaches for Bayesian variable selection," *Statistica Sinica*, 1997, *7*, 339–373.

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

_ , Michele Lenza, and Giorgio E. Primiceri, "Economic predictions with big data: the illusion of sparsity," Staff Reports 847, Federal Reserve Bank of New York April 2018.

Huber, Florian and Martin Feldkircher, "Adaptive Shrinkage in Bayesian Vector Autoregressive Models," *Journal of Business & Economic Statistics*, January 2019, *37* (1), 27–39.

Kastner, Gregor and Florian Huber, "Sparse Bayesian vector autoregressions in huge dimensions," *Journal of Forecasting*, November 2020, *39* (7), 1142–1165.

Kaufmann, Sylvia and Christian Schumacher, "Identifying relevant and irrelevant variables in sparse factor models," *Journal of Applied Econometrics*, 2017, *32* (6), 1123–1144.

_ and _ , "Bayesian estimation of sparse dynamic factor models with order-independent and ex-post mode identification," *Journal of Econometrics*, 2019, *210* (1), 116 – 134.

Kristensen, Johannes Tang, "Diffusion Indexes With Sparse Loadings," *Journal of Business & Economic Statistics*, 2017, *35* (3), 434–451.

Kroese, Dirk P. and Joshua C.C. Chan, *Statistical Modeling and Computation*, Springer Publishing Company, Incorporated, 2013.

Krüger, Fabian, Sebastian Lerch, Thordis L. Thorarinsdottir, and Tilmann Gneiting, "Predictive Inference Based on Markov Chain Monte Carlo Output," *arXiv e-prints*, Aug 2016, p. arXiv:1608.06802.

Kuzin, Vladimir, Massimiliano Marcellino, and Christian Schumacher, "MIDAS vs. mixed-frequency VAR: Nowcasting GDP in the Euro Area," *International Journal of Forecasting*, 2011, *27* (2), 529–542.

Makalic, Enes and Daniel F. Schmidt, "A Simple Sampler for the Horseshoe Estimator," *IEEE Signal Processing Letters*, 2016, *23* (1), 179–182.

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*, 2016, *34* (1), 118–127.

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.

McCracken, Michael W. and Serena Ng, "FRED-MD: A Monthly Database for Macroeconomic Research," *Journal of Business & Economic Statistics*, 2016, *34* (4), 574–589.

Nakajima, Jouchi and Mike West, "Bayesian Analysis of Latent Threshold Dynamic Models," *Journal of Business & Economic Statistics*, 2013, *31* (2), 151–164.

Pati, Debdeep, Anirban Bhattacharya, Natesh S. Pillai, and David B. Dunson, "Posterior contraction in sparse Bayesian factor models for massive covariance matrices," *The Annals of Statistics*, 2014, *42* (3), 1102–1130.

Plummer, Martyn, Nicky Best, Kate Cowles, and Karen Vines, "CODA: Convergence Diagnosis and Output Analysis for MCMC," *R News*, 2006, *6* (1), 7–11.

Polson, Nicholas G. and James G. Scott, "Shrink globally, act locally: Sparse Bayesian regularization and prediction," in J M Bernardo, M J Bayarri, J O Berger, A P Dawid, D Heckerman, A F M Smith, and M West, eds., *Bayesian statistics*, Vol. 9, Oxford University Press, 2010, pp. 501–538.

Rünstler, Gerhard, "On the Design of Data Sets for Forecasting with Dynamic Factor Models," in Eric Hillebrand and Siem Jan Koopman, eds., *Dynamic Factor Models*, Vol. 35 of *Advances in Econometrics*, Emerald Publishing Ltd, 2016, pp. 629–662.

Schumacher, Christian, "Forecasting German GDP using alternative factor models based on large datasets," *Journal of Forecasting*, 2007, *26* (4), 271–302.

_ , "Factor forecasting using international targeted predictors: The case of German GDP," *Economics Letters*, May 2010, *107* (2), 95–98.

_ and Jörg Breitung, "Real-time forecasting of German GDP based on a large factor model with monthly and quarterly data," *International Journal of Forecasting*, 2008, *24* (3), 386–398.

Thorsrud, Leif Anders, "Words are the New Numbers: A Newsy Coincident Index of the Business Cycle," *Journal of Business & Economic Statistics*, 2020, *38* (2), 393–409.

Wand, Matthew P., John T. Ormerod, Simone A. Padoan, and Rudolf Frühwirth, "Mean field variational Bayes for elaborate distributions," *Bayesian Analysis*, 2011, *6* (4), 847–900.

West, Mike, "Bayesian factor regression models in the large p, small n paradigm," in "Bayesian Statistics," Vol. 7 Oxford University Press 2003, pp. 723–732.