# Pandemic Priors

Danilo Cascaldi-Garcia*
Federal Reserve Board
danilo.cascaldi-garcia@frb.gov

March, 2023

**Abstract**

The onset of the COVID-19 and the great lockdown caused macroeconomic variables to display complex patterns that hardly follow any historical behavior. In the context of Bayesian VARs, an off-the-shelf exercise demonstrates how a very low number of extreme pandemic observations distorts the estimated persistence of the variables, affecting forecasts and giving a myopic view of the economic effects after a structural shock. I propose an easy and straightforward solution to deal with these extreme episodes, as an extension of the Minnesota Prior by allowing for time dummies. The method is flexible enough to let the econometrician optimally define the level of shrinkage, or arbitrarily choose how much signal to take from these extreme observations, nesting the boundary cases of an uninformative prior that soaks all the variance and a traditional Minnesota Prior. The Pandemic Priors succeed in recovering historical relationships and the proper identification and propagation of structural shocks.

**Keywords:** Bayesian VAR, Minnesota Prior, COVID-19, structural shocks.

**JEL codes:** C11, C32, E32, E44

*I thank Gianni Amisano, Pablo Cuba-Borda, Ana Beatriz Galvão, Jim Hamilton, Ed Herbst, Matteo Iacoviello, Francesca Loria, Matteo Luciani, Hyunseung Oh, Filippos Petroulakis, Michele Piffer, Albert Queralto, Ole Rummel, Tatevik Sekhposyan, Rodrigo Sekkel, Dalibor Stevanovic, Victor Valcarcel, Francesco Zanetti, Ren Zhang, and the participants of the Midwest Macro (Fall) conference and of the SEACEN forecast training for comments and suggestions. Dedicated to Mario Gutiérrez-Roig. The views expressed in this paper are solely the responsibility of the author and should not be interpreted as reflecting the view of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

# 1 Introduction

The onset of the COVID-19 pandemic and the subsequent great lockdown affected our lives and our jobs in an unprecedented way. Macroeconomic variables, which are quantitative mirrors of these effects, displayed complex patterns that hardly follow any historical behavior. Figure 1 exemplifies this unique situation. The variations in the U.S. industrial production and unemployment rate from March 2020 to August 2020 were the largest by far since at least 1976. From an empirical perspective, this episode poses a challenge on how to deal with such unusual behavior and still be able to retain historical relationships, produce reliable forecasts, and provide correct interpretations of economic shocks. I propose an easy and straightforward solution to this challenge, by allowing for irregular relationships of macroeconomic variables in extreme episodes, but conceding that there is uncertainty about these estimations.

Figure 1 Industrial production and unemployment rate variation over time

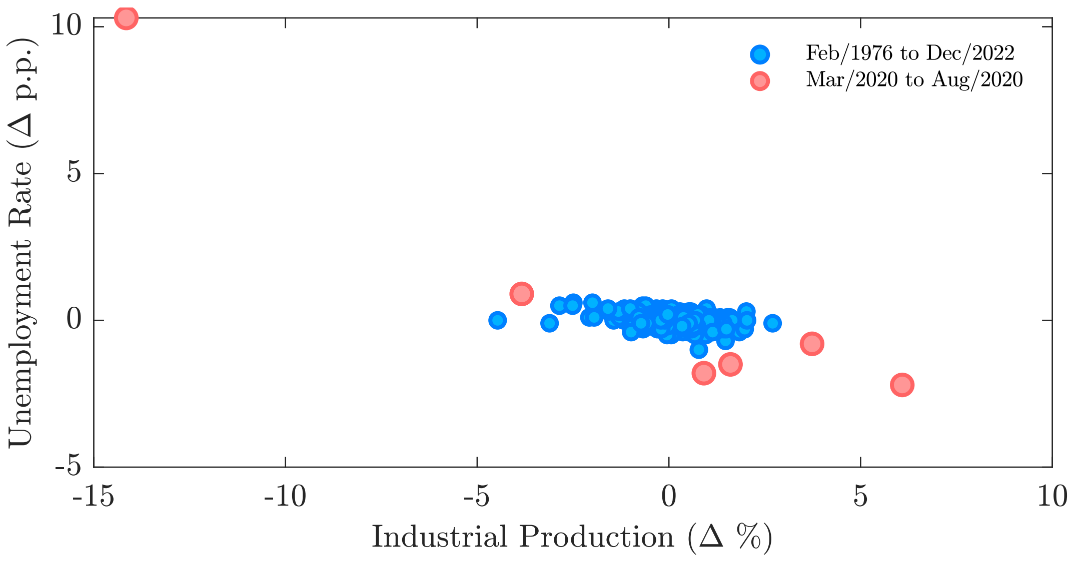


Note: Scatter plot of historical monthly changes of industrial production and unemployment rate. Blue dots correspond to the entire sample (February 1976 to December 2022), and red dots to the most extreme periods of the COVID-19 pandemic March to August 2020).

Bayesian vector autoregressions (VAR) are at the core of the macroeconomic empirical literature and are widely used by researchers, market participants, and policymakers for forecasting and the understanding of economic shocks. The seminal work of Litterman (1986) introducing the Minnesota Prior and future implementation developments

(Bańbura, Giannone, and Reichlin, 2010, Del Negro and Schorfheide, 2011, Carriero, Clark, and Marcellino, 2015, among others) allowed for computationally feasible estimations of large information sets that overcome the curse of dimensionality. I propose an extension of such procedure to allow for time dummies, namely Pandemic Priors, which are able to correctly adjust the historical relationship among the variables for the extreme values observed in specific periods.[1] Importantly, the econometrician can choose how much signal to take from the pandemic period through a single parameter, nesting the boundary cases of time dummies soaking all the pandemic variance and of full signal as in a conventional Minnesota prior. While this choice can be arbitrary and to the taste of the econometrician, I propose a method to define an optimal choice for how much signal to take from these extreme values, and a test to verify if indeed downplaying these observations through shrinkage is advisable.

With an off-the-shelf empirical example, I show that a very low number of extreme observations during the onset of the pandemic imply distorted autoregressive coefficients, affecting the estimated historical relationship among the variables, forecasts, and giving a myopic view of the economic effects after a structural shock. The Pandemic Priors, in turn, succeed in recovering these historical relationships, as confirmed by a Monte Carlo exercise, and the proper identification and propagation of structural shocks. Importantly, the simplicity of the method allows it to be adapted to any conventional or state-of-the-art structural identification procedure, enabling pre-pandemic conclusions to be extended and replicated going forward.

The procedure is akin to other methods that can be nested by the Pandemic Priors. Schorfheide and Song (2020), for example, highlight that forecasts performed after the extreme period around the onset of the pandemic are on par to the ones produced before the pandemic. As such, the authors suggest excluding the observations from March through June 2020 from the information set. This setup is equivalent to estimating the boundary case of the Pandemic Priors where the econometrician arbitrarily chooses a prior where the time dummies soak all the variance of those observations. In an empirical

[1] MATLAB and Julia implementations of the Pandemic Priors are available at www.danilocascaldigarcia.com.

exercise, I show that the optimal choice for the prior is one that downplays most, but not all, the signal from the pandemic period.

Lenza and Primiceri (2021) propose a method of estimating VARs by modeling a common shift and persistence of the volatility of the shocks during the extreme periods of the pandemic. The method takes the assumption that the volatilities of all shocks were scaled up by exactly the same constant and decay by exactly the same rate, so it is possible to establish priors and estimate these scale parameters. I propose a simpler and more parsimonious approach: allowing direct intercept shifts during the pandemic period, which removes the need to assume common volatility scale shifters and persistence. In fact, under the Pandemic Priors, each variable can potentially present different shifts and persistence during the pandemic period, captured by the individual time dummies. In an empirical exercise, the Pandemic Priors recover similar impulse responses to those found when using the Lenza and Primiceri (2021) method. In other words, the Lenza and Primiceri (2021) method is no different than a setup where little signal is taken from the pandemic period, and the results can be nested by the Pandemic Priors.

The procedure I propose is also an easy linear closed-form alternative to complex setups such as modeling extreme observations as random shifts in the stochastic volatility of the VAR, as in Carriero, Clark, Marcellino, and Mertens (2022), through non-parametric methods, as in Huber, Koop, Onorante, Pfarrhofer, and Schreiner (2023), or to estimating the VAR with $t$-distributed errors, as in Bobeica and Hartwig (2023). The Pandemic Priors approach is also related to Ng (2021), who proposes augmenting the VAR with an exogenous variable constructed as the log-differences of the information set during the pandemic period, and to Antolin-Diaz, Drechsel, and Petrella (2021), who model outliers in the context of dynamic factor models.

The outline of the paper is as follows. I discuss the technical implementation of the Pandemic Priors, how to select the optimal shrinkage from the marginal density standpoint, and a test for the applicability of the Pandemic Priors in section 2. Section 3 shows how the Pandemic Priors successfully recover the coefficients of the data generating process through a Monte Carlo simulation. Section 4 presents the implications of the

Pandemic Priors in an empirical example of estimating a medium-scale Bayesian VAR and identifying excess bond premium shocks. Section 5 discusses how sensitive the identification of shocks is to the shrinkage selection. Section 6 compares the Pandemic Priors with alternative methods available in the literature. Section 7 summarizes the findings of this paper.

## 2 Implementation

The Pandemic Priors build on Bańbura et al. (2010), who implements the traditional Minnesota Prior (Litterman, 1986) through dummy observations, by extending it to allow for time dummies on extreme observations. The method has the advantage of easy implementation, and avoiding the curse of dimensionality by allowing for large vector autoregression models through Bayesian shrinkage.

Following the notation from Bańbura et al. (2010), I take a VAR model with $n$ variables, $T$ periods, and $p$ lags as in:

$$\mathbf{Y}_t = \mathbf{c} + \mathbb{1}_{t=a}\mathbf{d}_a + ... + \mathbb{1}_{t=a+h}\mathbf{d}_{a+h} + \mathbf{A}_1\mathbf{Y}_{t-1} + ... + \mathbf{A}_p\mathbf{Y}_{t-p} + \mathbf{u}_t, \tag{1}$$

where $\mathbf{u}_t$ are innovations with $\mathbb{E}[\mathbf{u}_t\mathbf{u}_t'] = \Psi$, $\mathbf{c}$ is a vector of $n$ intercepts, $\mathbf{d}_a$ through $\mathbf{d}_{a+h}$ are $h$ vectors with $n$ time dummies for a pre-defined number of $h$ periods from $t = a$ through $t = a + h$ (which can be the COVID-19 crisis), and $\mathbb{1}_{t=i}$ is an indicator function that takes value $\mathbb{1}_{t=i} = 1$ for the period set $i = a, ..., a + h$, and 0 otherwise.

As in Litterman (1986) and Bańbura et al. (2010), I impose the prior that the variables are centered around the random walk with a drift, but now extending the concept to the idea that the pandemic is an abnormal period where the relationship between the variables may diverge from history. As such, the prior can be represented as

$$\mathbf{Y}_t = \mathbf{c} + \mathbb{1}_{t=a}\mathbf{d}_a + ... + \mathbb{1}_{t=a+h}\mathbf{d}_{a+h} + \mathbf{Y}_{t-1} + \mathbf{u}_t, \tag{2}$$

which is equivalent to shrinking the coefficient matrix $\mathbf{A}_1$ to the identity and the matrices

$\mathbf{A}_2, ..., \mathbf{A}_p$ to zero. The moments for the prior distribution of the coefficients are set as

$$\mathbb{E}\left[(A_k)_{ij}\right] = \begin{cases} \delta_i, & j = i, k = 1 \\ 0, & \text{otherwise} \end{cases} \qquad \mathbb{V}\left[(A_k)_{ij}\right] = \begin{cases} \frac{\lambda^2}{k^2}, & j = i \\ \upsilon \frac{\lambda^2}{k^2} \frac{\sigma_i^2}{\sigma_j^2}, & \text{otherwise} \end{cases} . \tag{3}$$

The coefficients $A_1, ..., A_p$ are assumed to be independent and normally distributed, the covariance matrix of the residuals to be diagonal defined as $\mathbf{\Sigma} = diag(\sigma_1^2, ..., \sigma_n^2)$, and the prior on the intercept is diffuse.

Choices for $\sigma_i$, the overall prior tightness $\lambda$, the factor $1/k^2$, and the coefficient $\upsilon$ are standard following good practices described in Bańbura et al. (2010), and flexible enough to accommodate beliefs about persistence, shrinkage toward the prior, variance decrease over lags, and the importance of own lags. By taking $\upsilon = 1$, it is possible to impose a normal inverse Wishart as in the Minnesota Prior under the form

$$\text{vec}(\mathbf{B})|\mathbf{\Psi} \sim \mathcal{N}\left(\text{vec}(\mathbf{B_0}), \mathbf{\Psi} \otimes \mathbf{\Omega}_0\right) \quad \textit{and} \quad \mathbf{\Psi} \sim i\mathcal{W}\left(S_0, \alpha_0\right) \tag{4}$$

where $\mathbf{B}$ is the matrix that collects the reduced-form coefficients of the $\mathbf{Y}_t = \mathbf{X}_t\mathbf{B} + \mathbf{U}_t$ vector autoregressive system with $\mathbf{X}_t = (\mathbf{Y}_{t-1}', ..., \mathbf{Y}_{t-p}', \mathbf{1}, \mathbb{1}_{t=a}, ..., \mathbb{1}_{t=a+h})'$, $\mathbf{B}_0$, $\mathbf{\Omega}_0$, $S_0$, and $\alpha_0$ are prior expectations and variance that make $\mathbf{B}$ matching the priors defined in equation 3, and $\mathbb{E}[\mathbf{\Psi}] = \mathbf{\Sigma}$, or the residual covariance.

In practice, these priors can be easily implemented through a series of dummy observations. The simplicity of the procedure makes it computationally efficient, allowing for the estimation of VARs with a large number of variables. I extend the procedure to allow for priors for the $h$ time dummies described in equation 1. Formally, the left-hand and right-hand side dummy observations ($\mathbf{Y}_d$ and $\mathbf{X}_d$, respectively) are defined as

$$\mathbf{Y}_d = \begin{pmatrix} diag(\delta_1\sigma_1, ..., \delta_n\sigma_n)/\lambda \\ \mathbf{0}_{n(p-1)\times n} \\ ... \\ diag(\sigma_1, ..., \sigma_n) \\ ... \\ \mathbf{0}_{n(p-1)\times n} \end{pmatrix} \quad \mathbf{X}_d = \begin{pmatrix} \mathbf{J}_p \otimes diag(\sigma_1, ..., \sigma_n)/\lambda & \mathbf{0}_{np\times 1} & \mathbf{0}_{np\times h} \\ ... & ... & ... \\ \mathbf{0}_{n\times np} & \mathbf{0}_{n\times 1} & \mathbf{0}_{n\times h} \\ ... & ... & ... \\ \mathbf{0}_{1\times np} & \epsilon & \mathbf{0}_{1\times h} \\ \mathbf{0}_{h\times np} & \mathbf{0}_{h\times 1} & diag(\phi_1, ..., \phi_h) \end{pmatrix} \tag{5}$$

where $\mathbf{J}_p = diag(1, 2, ..., p)$, and $\epsilon$ imposes an uninformative prior for the intercept. In comparison to the Bańbura et al. (2010) implementation, the innovation here is on the last column of $\mathbf{X}_d$, which imposes priors also for the $h$ time dummies through $\phi_1, ..., \phi_h = \phi$ (ordered last in $\mathbf{X}_t$). Following common practice from Litterman (1986), Sims and Zha (1998), and Bańbura et al. (2010), $\sigma_j$ can be calibrated from the variance of residuals of univariate autoregressive models with $p$ lags for each variable in the information set, and setting $\epsilon$ as a very small number makes the prior for the intercept fairly uninformative.

The parameter $\phi$ plays an important role for the Pandemic Priors, governing the prior associated with the time dummies. When $\phi \to 0$, the prior for the time dummies is fairly uninformative, and the time dummies soak the variance of the pandemic period. In other words, little signal is taken from those observations, teasing out the (estimated) reduced-form coefficients from that information. When $\phi \to \infty$, full signal is taken from the pandemic period, and that information is treated as any other observation. In other words, the setup converges to a conventional Minnesota Prior. It follows that the Pandemic Priors nest the boundary cases of no-to-full signal from the pandemic observations through the parameter $\phi$.[2] While the selection of $\phi$ can be arbitrary and up to the econometrician's taste of how much information to take from the extreme values from the pandemic period, I propose a method of selecting an optimal level of $\phi$ by finding the one that maximizes the marginal density of the model, or the optimal shrinkage level

[2] Carriero et al. (2022), in a robustness exercise, evaluate the forecasts stemming from uninformative time dummies for the pandemic period, which can also be nested by the Pandemic Priors.

of the pandemic observations, detailed in Section 2.1. Section 5 discusses the implications of different $\phi$ values for the identification of shocks.

Combining the original left-hand side data collected on $\mathbf{Y}_t$ with the dummy observations $\mathbf{Y}_d$ as in $\mathbf{Y}_t^* = [\mathbf{Y}_t', \mathbf{Y}_d']$, and the original right-hand side data collected on $\mathbf{X}_t$ with the dummy observations $\mathbf{X}_d$ as in $\mathbf{X}_t^* = [\mathbf{X}_t', \mathbf{X}_d']$, and adding the improper prior $\mathbf{\Psi} \sim |\mathbf{\Psi}|^{-(n+3)/2}$, leads to the posterior

$$\text{vec}(\mathbf{B})|\mathbf{\Psi}, \mathbf{Y}_t \sim \mathcal{N}\left(\text{vec}(\tilde{\mathbf{B}}), \mathbf{\Psi} \otimes \left(\mathbf{X}_t^{*\prime}\mathbf{X}_t^*\right)^{-1}\right) \quad \textit{and} \quad \mathbf{\Psi}|\mathbf{Y}_t \sim i\mathcal{W}\left(\tilde{\mathbf{\Sigma}}, T_d + 2 + T - m\right), \tag{6}$$

where $T$ is the sample size, $T_d$ is the length of dummy observations, $m = np + 1 + h$, $\tilde{\mathbf{B}} = \left(\mathbf{X}_t^{*\prime}\mathbf{X}_t^*\right)^{-1}\left(\mathbf{X}_t^{*\prime}\mathbf{Y}_t^*\right)$, and $\tilde{\mathbf{\Sigma}} = \left(\mathbf{Y}_t^* - \mathbf{X}_t^*\tilde{\mathbf{B}}\right)'\left(\mathbf{Y}_t^* - \mathbf{X}_t^*\tilde{\mathbf{B}}\right)$, or the reduced-form coefficients and estimated residual variance of the OLS estimation of $\mathbf{Y}_t^*$ on $\mathbf{X}_t^*$.

If the objective of the econometrician is increased forecast performance, it is possible to also adapt the dummy observations that impose a no-cointegration prior by constraining the sum of the coefficients described in Bańbura et al. (2010) to take into account the time dummies proposed here. In this case, it suffices to add an extra set of dummy observations, as in

$$\mathbf{Y}_{sc} = \text{diag}(\delta_1\mu_1, ..., \delta_n\mu_n)/\tau \qquad \mathbf{X}_{sc} = \begin{pmatrix} \mathbf{1}_{1\times p} \otimes \text{diag}(\delta_1\mu_1, ..., \delta_n\mu_n)/\tau & \mathbf{0}_{n\times 1} & \mathbf{0}_{n\times h} \end{pmatrix}, \tag{7}$$

where $\tau$ sets the degree of shrinkage and $\mu_j$ represents the average level of each $j$ variable in the information set. The data can then be combined as $\mathbf{Y}_t^* = [\mathbf{Y}_t', \mathbf{Y}_d', \mathbf{Y}_{sc}']$ and $\mathbf{X}_t^* = [\mathbf{X}_t', \mathbf{X}_d', \mathbf{X}_{sc}']$.

## 2.1 Optimal selection for $\phi$

The selection of $\phi$ can be arbitrary and defines how much signal to take from the extreme observations in the system. I propose here a method to select an optimal level of $\phi$, as the one that maximizes the marginal density of the model. The procedure is an adaptation of the optimal overall prior tightness described in Carriero, Kapetanios, and Marcellino

(2012) and Carriero et al. (2015), and can be defined as

$$\phi^* = \arg\max_{\phi} \ln p_{\phi}(Y), \tag{8}$$

where $p_{\phi}(Y)$ is the marginal density, or marginal likelihood, obtained by integrating the set $\Theta$ of coefficients of the model, defined as

$$p_{\phi}(Y) = p(Y|\phi) = \int p(Y|\Theta, \phi) p(\Theta|\phi) d\Theta. \tag{9}$$

Under the normal inverse Wishart prior, $p_{\phi}(Y)$ can be calculated in closed-form[3] as

$$\begin{aligned} p_{\phi}(Y) = \pi^{-\frac{Tn}{2}} \times \left| \left( \mathbf{I} + \mathbf{X}_t \mathbf{\Omega}_0(\phi) \mathbf{X}_t' \right)^{(-1)} \right|^{\frac{n}{2}} \times |S_0|^{\frac{\alpha_0}{2}} \times \left( \frac{\Gamma_n \frac{\alpha_0+T}{2}}{\Gamma_n \frac{\alpha_0}{2}} \right) \times ... \\ ... \times \left| S_0 + (\mathbf{Y}_t - \mathbf{X}_t \mathbf{B}_0)' \left( \mathbf{I} + \mathbf{X}_t \mathbf{\Omega}_0(\phi) \mathbf{X}_t' \right)^{(-1)} (\mathbf{Y}_t - \mathbf{X}_t \mathbf{B}_0) \right|^{-\frac{\alpha_0+T}{2}}, \end{aligned} \tag{10}$$

where $\alpha_0 = n+2$, $\Gamma_n$ is the $n$-variate gamma function, and the prior variance expectation $\mathbf{\Omega}_0(\phi)$ is now a function of $\phi$. From equation 8, the optimal $\phi^*$ is the one that maximizes the marginal density over a discrete grid of values for $\phi$.

## 2.2 A test for the applicability of the Pandemic Priors

Evaluating the marginal density of the model through different levels of $\phi$ is also an agnostic way of checking if, indeed, the observations for the selected period should be treated differently or not, and if so, how much should the econometrician take signal from them. If the optimal $\phi^* \to \infty$, data favors a model with a conventional Minnesota Prior, the observations from the selected period will be treated as usual, and the time dummies from the Pandemic Priors will be ineffective. If, however, the optimal $\phi^* \to 0$, data favors a model in which the observations from the selected period are downplayed up to some degree governed by $\phi^*$, and the time dummies from the Pandemic Priors become active.

In this section, I propose a test for the applicability of the Pandemic Priors, by

[3] See Bauwens, Lubrano, and Richard (2000) for details.

evaluating the marginal density of the two boundary cases: a Minnesota Prior model ($\phi^* \to \infty$) and an uninformative Pandemic Priors model ($\phi^* \to 0$). The idea is to calculate, over the entire sample $T$, the ratio $R_{t,w}$ between the marginal density (equation 10) of the boundary cases for every possible sub-sample of time periods with a defined length $w$, as

$$R_{t,w} = \frac{\ln p_{\phi\to\infty}(Y)_{t,w}}{\ln p_{\phi\to 0}(Y)_{t,w}}. \tag{11}$$

The $R_{t,w}$ ratio test can be read as follows. At any point in time $t$, if the model favors treating the observations from $t$ to $t+w-1$ as extreme values that should be downplayed by some degree, the marginal density associated to the uninformative Pandemic Priors ($\ln p_{\phi\to 0}(Y)_{t,w}$) will be higher than the marginal density associated to the Minnesota Prior ($\ln p_{\phi\to\infty}(Y)_{t,w}$). $R_{t,w}$ will then be lower than 1 and the application of the Pandemic Priors is advisable for the time period $t$ to $t+w-1$. If, in turn, $R_{t,w}$ is higher than 1, then the model favors a conventional Minnesota Prior, and the application of the Pandemic Priors over the time period $t$ to $t+w-1$ will be ineffective. Sections 3 and 4 present the application of the $R_{t,w}$ ratio test described in equation 11 for simulated and real data, respectively.

# 3 Monte Carlo simulation

The Pandemic Priors are able to recover posterior distributions that encompass the true coefficients from simulated data. I evaluate the method through a Monte Carlo simulation with four variables, for 600 periods, and emulating large and simultaneous shocks to each of them happening at $t = 501$, but with different size (5 to 20 standard deviations) and persistence (0.3 to 0.9), mimicking the behavior of economic variables at the onset of the COVID-19 pandemic. The full experiment is detailed on Appendix B.

First, I evaluate the $R_{t,w}$ ratio test over the 600 observations, with a defined length of $w = 24$ periods, following procedure described in Section 2.2. I assign $\phi = 0.001$ as the uninformative Pandemic Priors, where the time dummies soak all the variance of the $w$ observations, and $\phi = 5$ for the Minnesota Prior, where the time dummies bring no signal

Figure 2 Marginal density ratio (Minnesota Prior / uninformative Pandemic Priors)

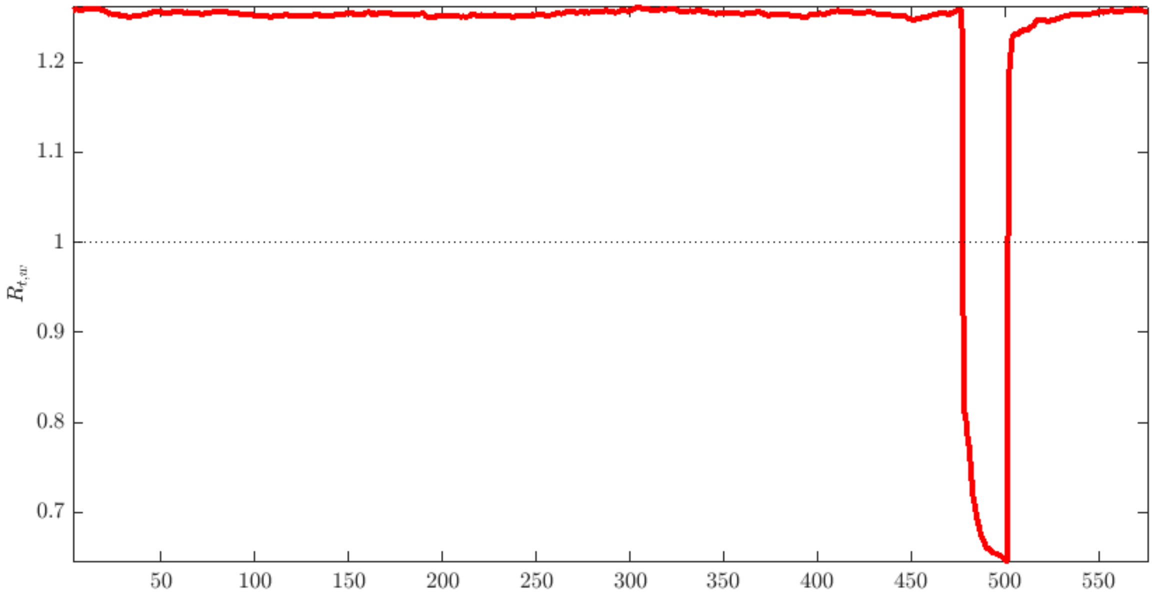


Note: $R_{t,w}$ ratio test of the marginal density of a Minnesota Prior, where $\phi = 5$, and the marginal density of uninformative Pandemic Priors, where $\phi = 0.001$, for simulated data. The test is conducted at each point in time $t$ with time dummies for the period $t$ to $t+w-1$, where $w$ is defined as 24 periods.

to the model. Figure 2 presents the evolution of the $R_{t,w}$ ratio test with the simulated data. The ratio stays consistently above 1 up until observation 477, indicating that the observations of any time window of $w = 24$ periods between $t = 1$ and $t = 477 + w - 1$, or $t = 500$, should not be downplayed as extreme values. The ratio, however, starts to drop and becomes substantially below 1 from observations 478 to 501, indicating that the model favors a system where the $w = 24$ observations between $t$ and $t + w - 1$ are downplayed by some degree. Indeed, applying $w = 24$ time dummies with the Pandemic Priors starting at any point from $t = 478$ to $t = 501$ would cover the simultaneous simulated shocks happening at $t^* = 501$. Finally, the $R_{t,w}$ ratio test reaches its lowest point exactly at $t = 501$.

Following, I estimate two Bayesian VARs in levels with these four series: with the standard Minnesota Prior (baseline), where I do not take into account the large shock observed at $t^* = 501$, and with the Pandemic Priors, treating the first 24 periods when the shock happens with the time dummies ($t = 501, ..., 524$).[4] I set $\phi = 0.075$ as the

[4]The results are robust to extending the period covered by the time dummies to the whole 60 observations with unusually high levels, but 24 periods are sufficient to recover the original data generating process. I estimate both the baseline Minnesota Prior and the Pandemic Priors with fairly loose overall prior tightness ($\lambda = 5$), but the results are also robust to tighter priors.

Figure 3 Posterior draws for the autoregressive coefficients

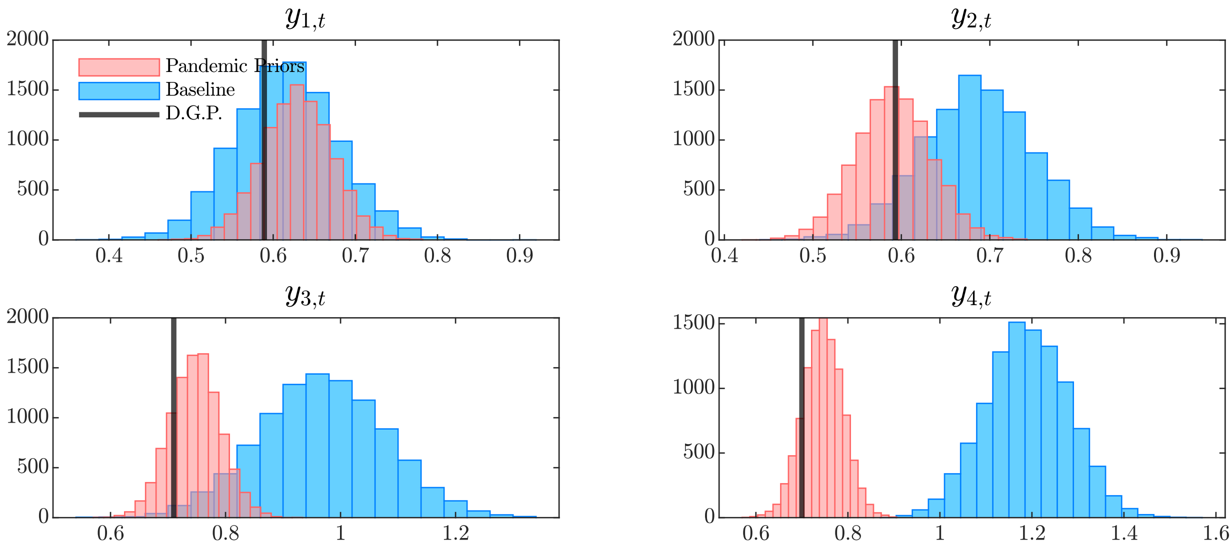


Note: Histograms of the (reduced-form) autoregressive coefficient of the baseline Minnesota Prior (blue bars) and the Pandemic Priors (pink bars) estimations, for each variable in the information set, compared with the data generating process (D.G.P.). $\phi = 0.075$ as the optimal $\phi^*$ found over a discrete grid search. Distributions constructed after 10,000 draws from the posterior distribution. The VAR is estimated for 600 simulated periods.

optimal $\phi^*$ found over a discrete grid search[5] following procedure described in Section 2.1. In Figure 3, I compare the baseline (blue bars) and the Pandemic Priors (pink bars) posterior distributions of the estimated reduced-form autoregressive coefficients with the true coefficients known from the data generating process. Three results stand out from this exercise. First, the larger and more persistent the shock is, the further away the estimated baseline coefficient will be from the true value. For $y_{4,t}$, for example, the true coefficient is not even on the support estimated by the baseline Minnesota Prior. The Pandemic Priors, in turn, successfully incorporate the true coefficient within its posterior support. For $y_{1,t}$, where the shock was considerably smaller and less persistent, both methods manage to have the true coefficient within their supports.

The second result is that, when facing such unusually large shocks, there is considerably more uncertainty on the autoregressive coefficients with the baseline Minnesota Prior than with the Pandemic Priors. The distributions of the baseline posterior coefficients are wider than the Pandemic Priors, indicating that the baseline method had a harder time finding coefficients that fit well with the data.

[5] Grid values of $[0.001, 0.01, 0.025, 0.050, 0.075, 0.10, 0.15, 0.20, 0.25, 0.30, 0.35, 0.40, 0.45, 0.50, 0.75, 1, 2, 5]$.

Third, the choice of an optimal $\phi^*$ matters for achieving a posterior distribution in which the true parameters lie within, particularly for series that experience larger jumps. In this simulation, the optimal $\phi^*$ was found at 0.075, relatively far from the purely uninformative boundary level of 0.001 in the grid search. It follows that the model took some, albeit small, signal from the extreme values of the simulated shocked period. In fact, when arbitrarily setting $\phi$ to an uninformative prior at 0.001, the true parameter of the simulated series $y_{4,t}$ is almost outside of the support of the Pandemic Priors, as shown in Figure A.1 in the Appendix.

# 4 An empirical example

In this section I present an empirical example of how a few pandemic observations can markedly change the estimated relationship among macroeconomic variables and the interpretation of structural shocks. I estimate a monthly Bayesian VAR in levels, where the information set includes eight endogenous variables,[6] namely the excess bond premium (EBP, Gilchrist and Zakrajšek, 2012), (log) S&P 500 index, Federal Funds shadow rate (Wu and Xia, 2016), (log) personal consumption expenditures (PCE), (log) PCE price index, (log) employment, (log) industrial production, and unemployment rate. The estimation sample runs from January 1975 through December 2022. I include 12 lags, set $\phi = 0.05$ as the optimal $\phi^*$ found over a discrete grid search, overall prior tightness $\lambda = 0.2$, and $\tau = 10 \times \lambda$.[7]

The $R_{t,w}$ ratio test favors treating the onset of the COVID-19 pandemic as extreme values that should be downplayed by some degree. Figure 4 presents the evolution of the $R_{t,w}$ ratio test with the empirical exercise, setting the window $w$ to six months. Three results stand out from the test. First, there is no evidence of the need of treating any time window prior to the COVID-19 pandemic onset as an extreme value. The ratio stays consistently above 1 up until 2020. Second, there is overwhelming evidence that the onset of the COVID-19 pandemic is an erratic historical episode, and the model favors a system

[6] Table A.1 in the Appendix presents the full description of the dataset.

[7] Results are robust to different lag selections, and $\lambda$ and $\tau$ specifications, and are available upon request.

Figure 4 Marginal density ratio (Minnesota Prior / uninformative Pandemic Priors)

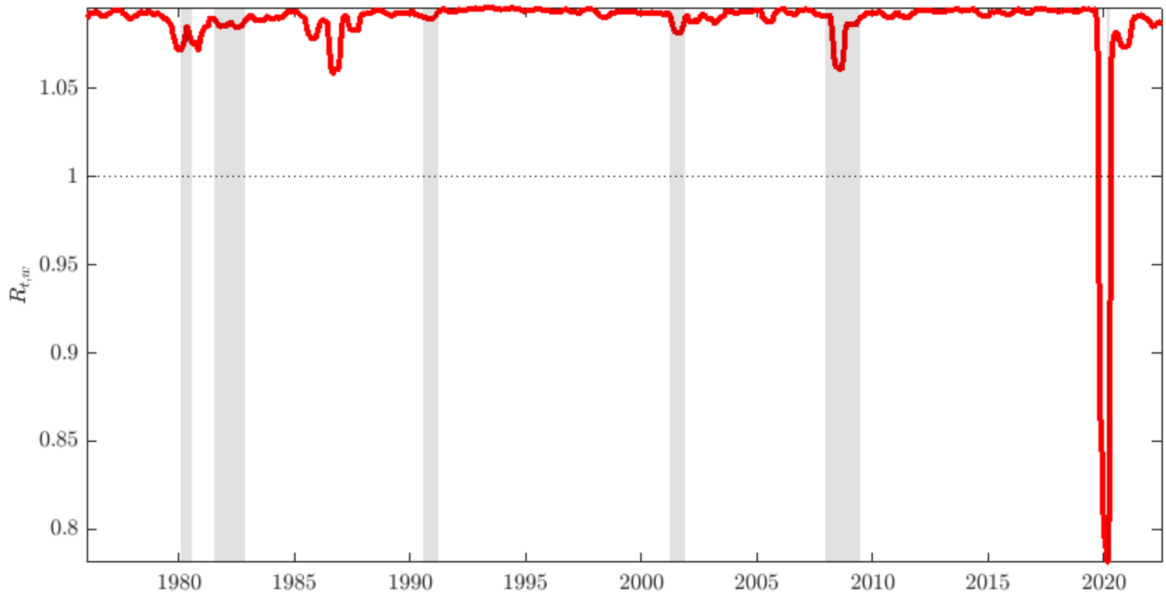


Note: $R_{t,w}$ ratio test of the marginal density of a Minnesota Prior, where $\phi = 5$, and the marginal density of uninformative Pandemic Priors, where $\phi = 0.001$, for empirical data. The test is conducted at each point in time $t$ with time dummies for the period $t$ to $t + w - 1$, where $w$ is defined as 6 months. Shaded areas are recessions identified by the NBER.

where these observations are downplayed by some degree. The ratio drops below 1 from November 2019 to April 2020, with its lowest value in March 2020, indicating that a model where that month and its following five months are treated with time dummies from the Pandemic Priors is preferable over a conventional Minnesota Prior. Finally, the ratio seems to drop near recessions (as identified by the NBER), especially the ones in 1980, 2008, and 2020, indicating that recessions are periods where data does not seem to follow usual historical patterns.

For the results in this Section, I explicitly model the COVID-19 crisis by applying the Pandemic Priors with six individual time dummies for the period of March 2020 through August 2020, as this window coincides with the onset of the pandemic, represents the very extreme observations in unemployment rate and industrial production (as illustrated by Figure 1), and is the lowest point evidenced by the $R_{t,w}$ ratio test.

### 4.1 Pandemic Priors matter for estimation, ...

The estimated time dummies build on the assumption that we should potentially observe intercept shifts for the macroeconomic variables in the selected periods. The Pandemic Priors imply that, while we observe the outcome of each variable, there is uncertainty about this shift. Indeed, there is substantial heterogeneity across variables about the size of the intercept shift, the timing of such shifts, and persistence. Figure 5 presents the (reduced-form) posterior distributions from 10,000 draws of the intercept, and the intercept shift (intercept plus time dummy) for the period March 2020 to August 2020.[8] Some variables show quite stable intercepts (e.g., shadow rate and PCE price index), but others show large shifts, with more pronounced examples in March 2020 for PCE, industrial production, employment, and unemployment rate. While the coefficient for the S&P 500 shows a large shift in March 2020 that reverts to the intercept in other periods, the employment variables show a substantial persistence of abnormal intercept shifts over the period March 2020 through August 2020. The Pandemic Priors succeed on capturing these heterogeneous shifts and persistence.

While these six extreme months of the pandemic period correspond to only about 1% of the total sample, not treating them as outliers has direct implications on the (reduced-form) coefficients of the Bayesian VAR. I evaluate this effect by comparing two exercises. First, as a baseline, I estimate the Bayesian VAR without any pandemic time dummy, as in the off-the-shelf Minnesota Prior procedure from Bańbura et al. (2010). The assumption of such a method is that the historical relationship among the endogenous variables have not changed during the pandemic. The second exercise applies the Pandemic Priors. Figure 6 presents the posterior distributions from 10,000 draws of the first lag (reduced-form) autoregressive coefficient of each variable in the information set, for the baseline and the Pandemic Priors setups.[9]

The distributions of posterior draws are substantially different between the baseline and the Pandemic Priors, and there is heterogeneity across variables. While the estimated

[8]Figure A.2 in the Appendix shows the posterior distributions for the time dummies, evidencing the uncertainty around the estimations.

[9]The posterior distribution is truncated to stable coefficient sets, discarding non-stationary draws. Results are also robust to the non-truncated posterior distribution.

Figure 5 Posterior draws for the intercept and pandemic time dummies

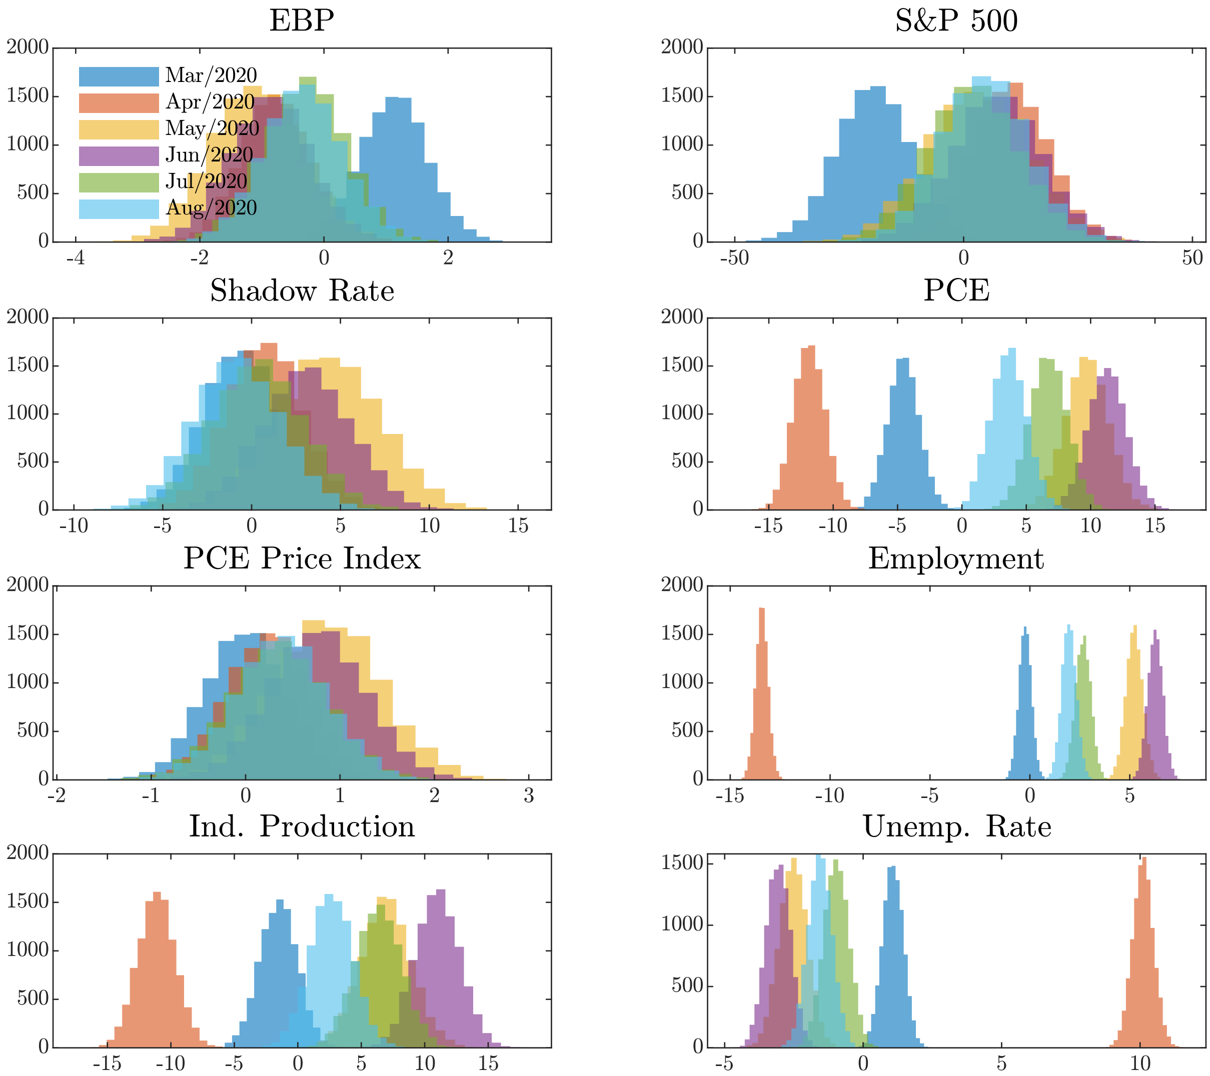


Note: Histograms of the (reduced-form) intercept and the intercept plus the time dummies for the pandemic period (March 2020 to August 2020), of each variable in the information set. Distributions constructed after 10,000 draws from the posterior distribution. The VAR is estimated from January 1975 to December 2022.

coefficients are essentially unchanged for EBP, S&P 500, shadow rate, and PCE price index, the variables with more extreme pandemic observations are also the ones with more disparate coefficients. Not treating the pandemic observations with the time dummies would imply a higher autoregressive coefficient for PCE, and the distributions almost do not overlap. The industrial production coefficient distribution is shifted to the left with the Pandemic Priors. The employment and unemployment rate variables, which have more extreme outliers, present opposite effect, with lower autoregressive coefficients in the baseline setup. Also, there is substantially more parameter uncertainty for the

Figure 6 Posterior draws for the autoregressive coefficients

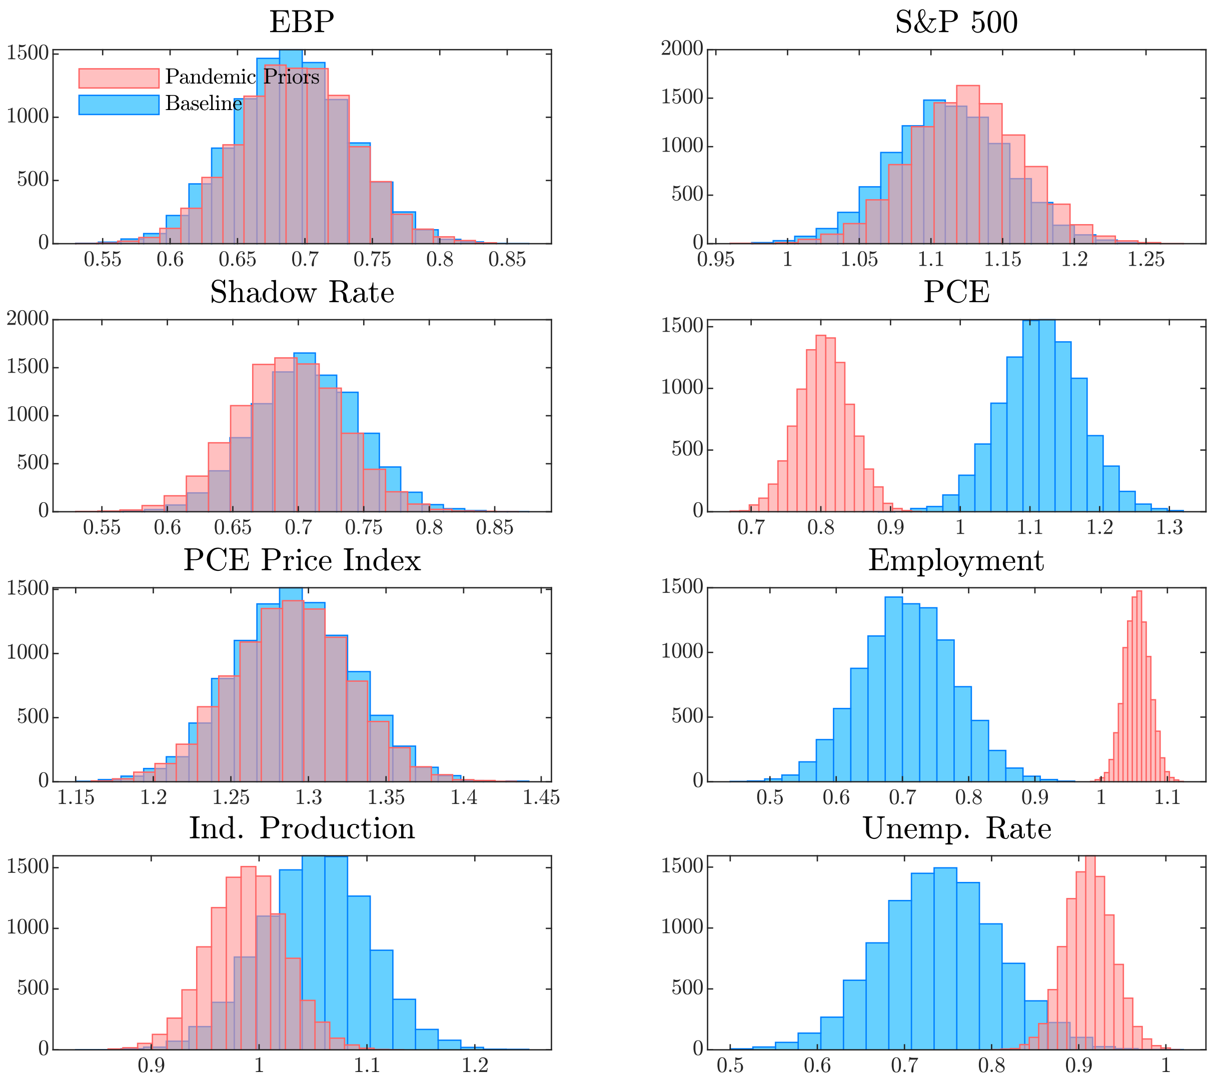


Note: Histograms of the (reduced-form) autoregressive coefficient of the baseline Minnesota Prior (blue bars) and the Pandemic Priors (pink bars) estimations, for each variable in the information set. Distributions constructed after 10,000 draws from the posterior distribution. The VAR is estimated from January 1975 to December 2022.

employment and unemployment rate when using the baseline compared to the Pandemic Priors.

## 4.2 ..., for forecasts, ...

With distinct autoregressive and lagged coefficients between the baseline and the Pandemic Priors, the persistence of each variable is affected, generating direct implications for forecasting. I evaluate the effect on the forecasts by estimating the unconditional 12-month ahead path for each variable implied by the baseline Minnesota Prior method

Figure 7 Unconditional forecasts as of December 2022

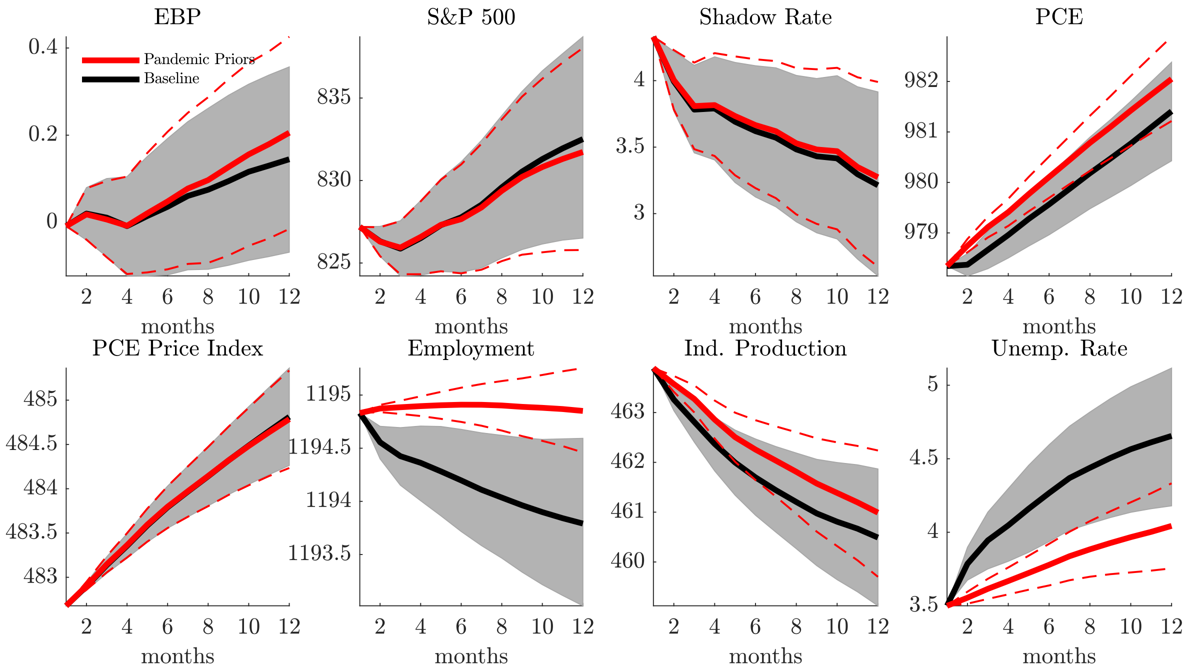


Note: Solid lines are estimated unconditional forecasts and correspond to the posterior median estimates (red with Pandemic Priors, and black as the baseline). The VAR is estimated from January 1975 to December 2022. The gray shaded area and the dashed red lines represent the one standard deviation coverage bands of the forecasts obtained with 10,000 draws from the posterior distribution.

and by the Pandemic Priors, as of December 2022, presented in Figure 7.[10] Red solid lines are the (posterior median) forecasts using the Pandemic Priors and solid black lines using the baseline Minnesota Prior.

As expected, variables where the autoregressive coefficients are essentially unchanged between the baseline and the Pandemic Priors, such as the EBP, the S&P 500, the shadow rate, and the PCE price index, present very similar unconditional forecasts no matter which model is estimated. However, variables that are markedly affected by extreme values during the pandemic, such as employment and unemployment rate, present different unconditional forecasts, implying different economic interpretations. While the baseline model indicates that employment would systematically decrease throughout 2023, especially in the first trimester, the Pandemic Priors provide a picture where employment is essentially unchanged over the forecast horizon. The unemployment rate is forecasted to quickly increase over 2023 under the baseline setup, reaching about 4.5% at the end of the year, while the Pandemic Priors show much slower but steady increase in the unem-

[10] Figure A.3 in the Appendix reports the unconditional forecasts over a longer horizon.

ployment rate to about 4%. The forecast uncertainty for these variables is substantially smaller under the Pandemic Priors than under the baseline. Finally, PCE and industrial production forecasts are also distinct between the methods, with a faster growth of PCE and slower decrease of industrial production with the Pandemic Priors when compared to the foreseen by the baseline.

### 4.3 ..., and for the identification of structural shocks

The extreme observations also impose a tilted view of the economic effects stemming from structural shocks. I evaluate this stance by identifying an excess bond premium shock, in the spirit of Gilchrist and Zakrajšek (2012) and Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek (2016), with the baseline Minnesota Prior estimation and with the Pandemic Priors. For simplicity, I identify the excess bond premium shock recursively, as the first shock in the Bayesian VAR where EBP is ordered first. Of note, the Pandemic Priors are flexible enough to accommodate any other conventional or state-of-the-art identification procedures, such as Proxy VARs, sign restrictions, or maximization of the variance decomposition. Figure 8 presents the 12 months ahead impulse response functions of the EBP shock, with solid red lines for the (posterior median) responses using the Pandemic Priors and solid black lines for the baseline Minnesota Prior.[11]

The economic effects of an EBP shock using the baseline and the Pandemic Priors estimation differ both in size and propagation, and it is heterogeneous over the variables. While the effects on the S&P 500 and PCE price index are almost indistinguishable if one uses the Pandemic Priors or not, there are crucial differences for the other variables. For example, simply ignoring the particular behavior of these six observations would steer an economist to expect a large and sharp short-term effect reduction in employment in response to the increased risk. While in the baseline model employment drops by about 0.15 percent after only two months of the shock, the Pandemic Priors imply a smoother and more delayed employment deterioration, reaching negative 0.15 percent only about nine months after the shock. Similar interpretation applies to the unemployment rate,

[11] Figure A.5 in the Appendix reports the impulse response functions over a longer horizon.

Figure 8 Impulse responses to a 1 s.d. EBP shock

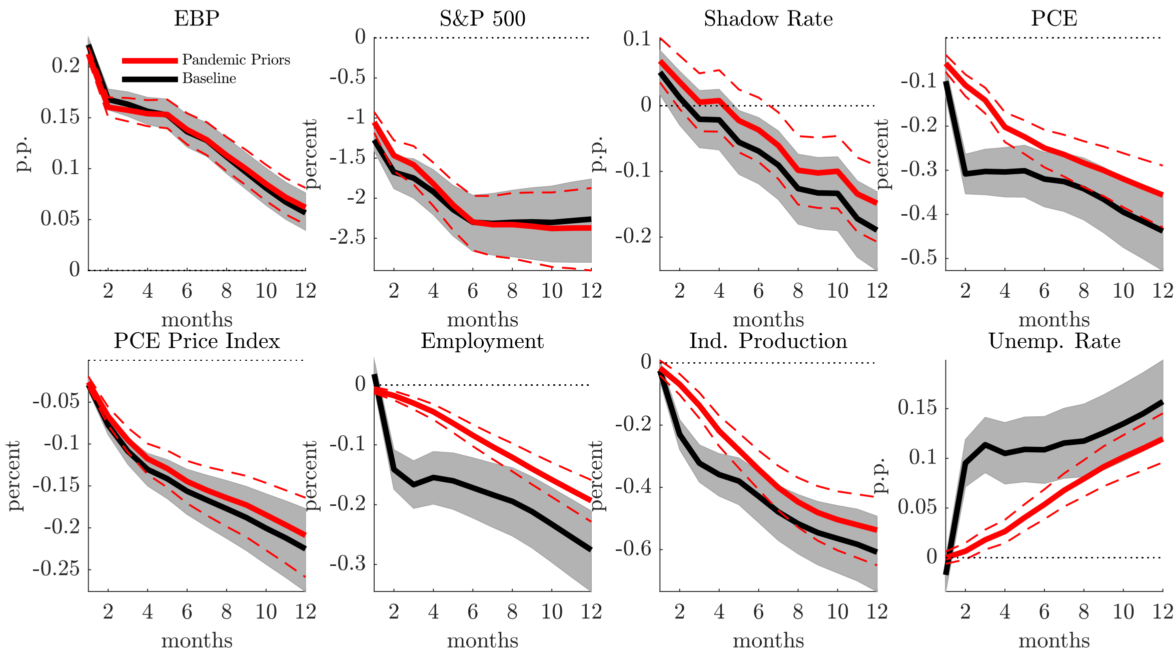


Note: Solid lines are estimated impulse responses to a standard deviation EBP shock and correspond to the posterior median estimates (red with Pandemic Priors, and black as the baseline). The VAR is estimated from January 1975 to December 2022. The gray shaded area and the dashed red lines represent the one standard deviation coverage bands of the EBP shock obtained with 10,000 draws from the posterior distribution.

which sharply increases by 0.1 percentage point in the baseline, but only smoothly reaches that level with the Pandemic Priors. PCE and industrial production also show abnormally sharp short-term deterioration when the pandemic observations are not properly treated.

## 5 Sensitivity to different levels of $\phi$

The parameter $\phi$ governs how informative the data from the pandemic period are for the overall estimation. With $\phi \rightarrow 0$, the prior is fairly uninformative and time dummies soak all the variance of the pandemic period, while with $\phi \rightarrow \infty$ the time dummies are treated as usual observations and we have the traditional Minnesota Prior. Here I revisit the identification of the EBP shock presented in the previous section, but evaluating the sensitivity to different levels of $\phi$.

Figure 9 presents the the posterior median of impulse responses of an EBP shock. Black solid lines are the impulse responses using a Minnesota Prior, and colored lines are

Figure 9 Impulse responses to a 1 s.d. EBP shock under different $\phi$ levels

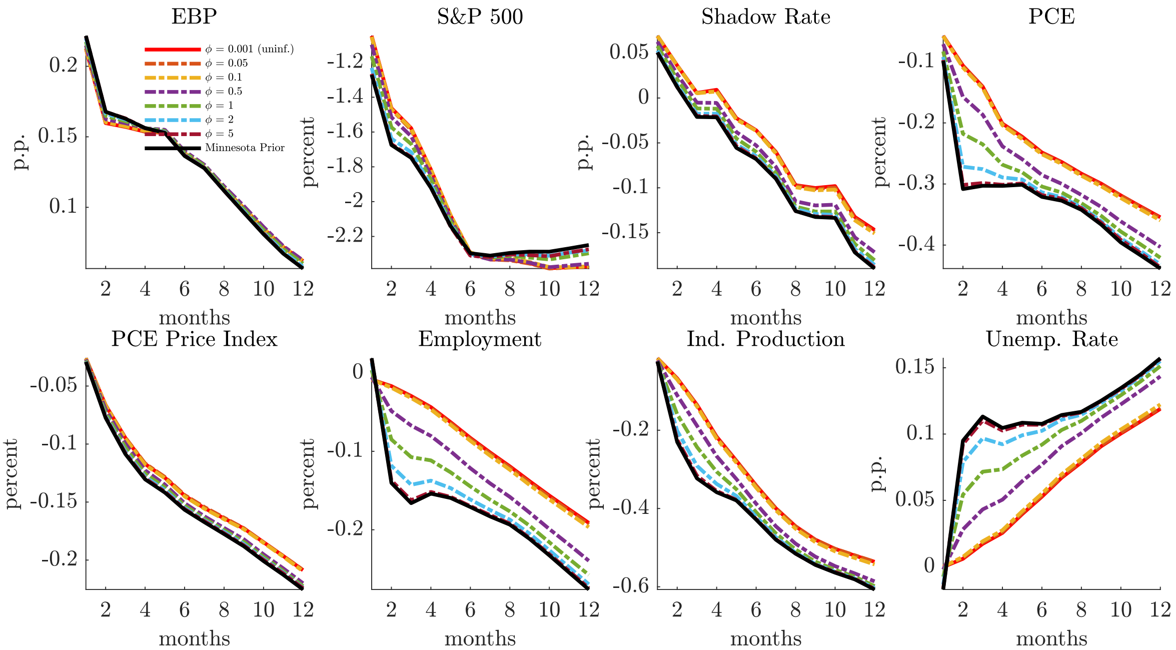


Note: Lines are estimated impulse responses to a standard deviation EBP shock and correspond to the posterior median estimates (black with Minnesota Prior, and colored with Pandemic Priors under different $\phi$). The VAR is estimated from January 1975 to December 2022. Posterior median of the EBP shock obtained with 10,000 draws from the posterior distribution.

the Pandemic Priors with $\phi$ ranging from $\phi = 0.001$, or uninformative, to $\phi = 5$, close to the Minnesota Prior.[12] As before, the posterior median of the impulse responses of variables that did not observe large swings during the pandemic period are quite similar no matter which $\phi$ level is used. This is the case for EBP and the PCE price index, for example. However, employment and unemployment rate show substantially different paths for the impulse responses conditional on how much information the estimation takes from the pandemic observations, bounded by the uninformative Pandemic Priors ($\phi = 0.001$) and the Minnesota Prior ($\phi = 5$).

It follows that the Pandemic Priors nest any setup for the prior belief the econometrician may have on how much information she wants to be stemmed from the pandemic period. If the econometrician would want to have the time dummies soaking up completely the variance, a fairly uninformative prior can be implemented by making $\phi$ substantially small. If the econometrician would want to take a downplayed signal from those observations, that is also possible with a larger $\phi$. In this setup, $\phi = 5$ is already sufficiently

[12] Figure A.4 in the Appendix reports the impulse response functions over a longer horizon.

large to make it indistinguishable from the Minnesota Prior, and levels between 0.001 and 5 represent a mix of diffuse and rigid priors over the pandemic period.

## 6 Comparison to alternative methods

In this section, I compare the Pandemic Priors estimated with an optimal $\phi^*$ with the alternative methods proposed by Schorfheide and Song (2020) and by Lenza and Primiceri (2021). According to Schorfheide and Song (2020), there are two main options for dealing with the COVID-19 observations: either by increasing the complexity of the VAR by directly modeling the outliers, or by simply excluding the extreme observations of that period. The authors show that VAR forecasts perform well from July 2020 onward, on par with pre-pandemic performance. As such, the authors propose excluding the observations from March through June 2020. In essence, such alternative would be similar to imposing time dummies with fully uninformative priors ($\phi \to 0$) on those months, which is not necessarily the optimal choice from the perspective of the marginal density of the model. The Pandemic Priors, by allowing the econometrician to choose how much signal to draw from the pandemic observations, nest the approach proposed by Schorfheide and Song (2020), on one extreme, the Minnesota Prior on the other, and any combination in between, including the optimal that maximizes the marginal density.

The Lenza and Primiceri (2021) procedure conjectures that the shocks observed at the onset of the pandemic presented substantially larger volatility. If the volatility of all shocks were scaled up by exactly the same constant, with exactly the same persistence thereafter (the commonality assumption),[13] it is possible to establish priors and estimate these parameters. In practice, the procedure estimates common scale parameters for the volatility of all shocks observed in March, April, and May 2020, and assumes that the residual variance decays at a fixed rate after May 2020. As pointed out by the authors, the commonality assumption is an approximation that works well in a period in which all variables experienced record variation. However, several aggregate variables commonly used in monthly VARs to characterize U.S. macroeconomic relationships did

[13] Also present in the stochastic volatility model of Carriero, Clark, and Marcellino (2016).

not show unreasonably large variance shocks during the onset of the COVID-19 pandemic in comparison to historical standards. For example, the S&P 500 index fell by 19% in March 2020, which is on par with the global financial crisis (-20% in October 2008), and with 10 other events with monthly double-digit variations since 1975. Indeed, as in the Bayesian VAR example shown on Figures 5 and 6, not all variables seem to be reactive to potentially extreme values during the March 2020 to August 2020 period. The EBP and the S&P 500, for example, showed relatively stable intercepts and autoregressive coefficients throughout the most acute period of the pandemic.

In summary, assuming a common scalar shifter and a common decay parameter for the variance of all shocks is a good step to avoid that extreme values contaminate the stochastic process of the variables in the VAR, but may not be the most appropriate when there is heterogeneity on the size and persistence of the volatility shift. The Pandemic Priors, by assigning individual time dummies for each variable and at each unusual period, allow for heterogeneous shifts (both in timing and size) and rate of decay over the information set.

I compare the Schorfheide and Song (2020) and the Lenza and Primiceri (2021) procedures with the Pandemic Priors by estimating the empirical exercise of the previous sections and identifying the EBP shock. I employ the optimal $\phi^* = 0.05$ for the Pandemic Prior, meaning that the time dummies will soak most of the variance of the pandemic period, but not all of it, as a pure uninformative prior would suggest. I exclude the March through June 2020 observations for the Schorfheide and Song (2020) procedure, as suggested by the authors. I estimate the Lenza and Primiceri (2021) procedure using the prior selection proposed by Giannone, Lenza, and Primiceri (2015). Figure 10 presents the posterior median impulse responses for the EBP shock, with the Pandemic Priors depicted with red lines, the Schorfheide and Song (2020) procedure with yellow dashed lines, the Lenza and Primiceri (2021) procedure with blue dashed lines, and the baseline Minnesota Prior with black lines.[14] All three methods present very similar results, and quite different from the baseline (black lines), with no controls for the pandemic.

[14]Figure A.6 in the Appendix reports the impulse response functions over a longer horizon.

Figure 10 Comparison of impulse responses to a 1 s.d. EBP shock

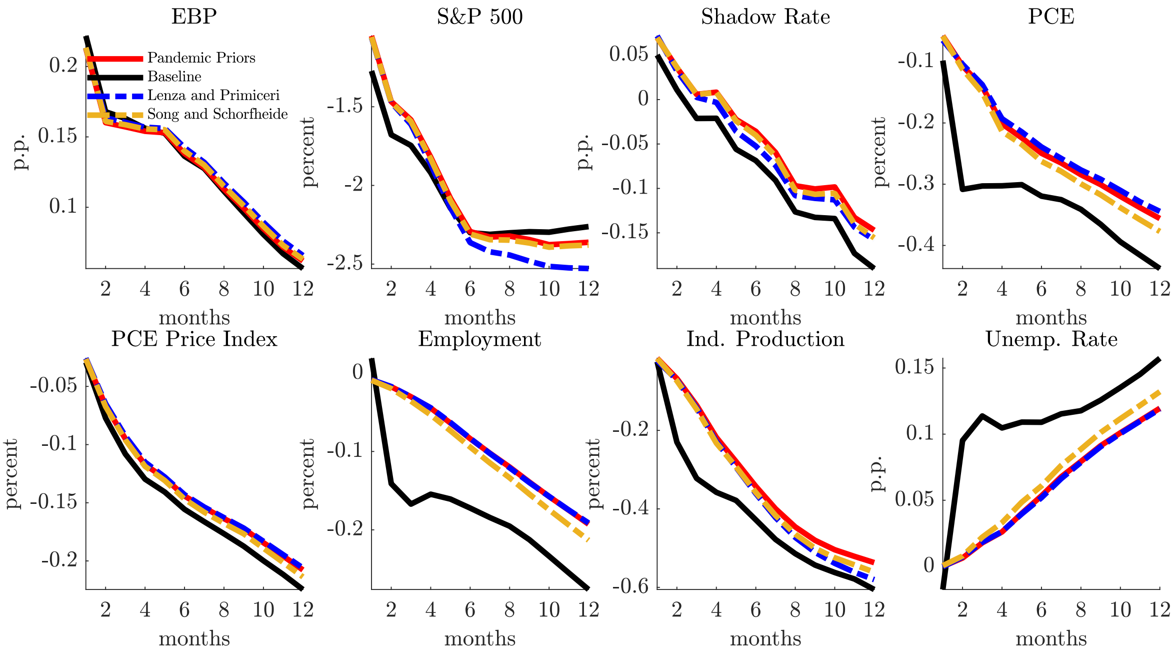


Note: Solid lines are estimated impulse responses to a standard deviation EBP shock and correspond to the posterior median estimates (red with Pandemic Priors, blue with the Lenza and Primiceri (2021) method, yellow with the Schorfheide and Song (2020) method, and black as the baseline). The VAR is estimated from January 1975 to December 2022. Posterior median of the EBP shock obtained with 10,000 draws from the posterior distribution.

Three results can be drawn from these comparisons. First, the similarity of the impulse responses indicates that most of the information distorting the reduced-form autoregressive coefficients must be coming from the earlier months of the onset of the pandemic, or March and April 2020. Those dates are the ones that have a common special treatment among the three methods: excluded in Schorfheide and Song (2020), dummy shifter on the volatility in Lenza and Primiceri (2021), and downplayed signal in the Pandemic Priors. Second, while the impulse responses are similar, they are not the same. Excluding the observations, as proposed by Schorfheide and Song (2020), means making an active choice of taking no signal from the pandemic observations, while the Pandemic Priors show that downplaying most, but not all, of the signal from those observations would be optimal. The similarity of the Lenza and Primiceri (2021) procedure with the other two indicates that the method is essentially no different than a setup where no signal (or downplayed signal) is taken from the pandemic period, and the results can be nested by the Pandemic Priors.

## 7 Conclusion

Extreme observations, such as the ones observed during the most acute periods of the COVID-19 pandemic, blur our interpretation of historical relationships among macroeconomic variables and the economic effects of shocks. In this paper, I show an easy and straightforward way of dealing with such episodes in empirical macroeconomics by proposing Pandemic Priors for Bayesian VAR estimations. The assumption is that macroeconomic variables present an abnormal behavior in extreme episodes such as the pandemic, but resume their historical relationship once conditions normalize. I propose time dummies for these extreme events that capture such unusual behavior, but accept that there is uncertainty about its potential outcome. Importantly, the method is flexible enough to indicate the optimal level of shrinkage of the pandemic period, or to let the econometrician choose how much signal to take from these extreme observations, nesting the boundary cases of an uninformative prior that soaks all the variance of the period and a traditional Minnesota Prior. While the COVID-19 pandemic is a natural candidate for such modeling, the method presented here can also be applied to other periods where the macroeconomic relationship among the variables is potentially (and temporarily) unusual, such as the zero lower bound of interest rates.

The empirical example of estimating and identifying an excess bond premium shock confirms the substantial intercept shifts during the period of March 2020 to August 2020, affecting the estimated historical coefficients, unconditional forecasts, and structural identification. The Pandemic Priors recover historical relationships, as confirmed by a Monte Carlo exercise, and the proper identification and propagation of structural shocks, allowing for estimating Bayesian VARs without having to restrict the sample to pre-pandemic periods, dropping observations, or resorting on more complex methods, such as volatility changes or $t$-distributed shocks. As the Pandemic Priors are flexible enough to accommodate any sort of structural identification, they also allow policymakers to make well-informed decisions about responses to economic shocks going forward.

# References

Antolin-Diaz, J., Drechsel, T., and Petrella, I. (2021). Advances in nowcasting economic activity: The role of heterogeneous dynamics and fat tails. Tech. rep.

Bańbura, M., Giannone, D., and Reichlin, L. (2010). Large bayesian vector auto regressions. *Journal of Applied Econometrics*, *25(1)*, 71–92.

Bauwens, L., Lubrano, M., and Richard, J.-F. (2000). *Bayesian inference in dynamic econometric models*. OUP Oxford.

Bobeica, E., and Hartwig, B. (2023). The COVID-19 shock and challenges for inflation modelling. *International Journal of Forecasting*, *39*(1), 519–539.

Caldara, D., Fuentes-Albero, C., Gilchrist, S., and Zakrajšek, E. (2016). The macroeconomic impact of financial and uncertainty shocks. *European Economic Review*, *88*, 185–207.

Carriero, A., Clark, T. E., and Marcellino, M. (2015). Bayesian VARs: specification choices and forecast accuracy. *Journal of Applied Econometrics*, *30(1)*, 46–73.

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

Carriero, A., Clark, T. E., Marcellino, M., and Mertens, E. (2022). Addressing COVID-19 Outliers in BVARs with Stochastic Volatility. *The Review of Economics and Statistics*, (pp. 1–38).

Carriero, A., Kapetanios, G., and Marcellino, M. (2012). Forecasting government bond yields with large bayesian vector autoregressions. *Journal of Banking & Finance*, *36*(7), 2026–2047.

Del Negro, M., and Schorfheide, F. (2011). Bayesian macroeconometrics. *The Oxford Handbook of Bayesian Econometrics*, *293*, 389.

Giannone, D., Lenza, M., and Primiceri, G. E. (2015). Prior selection for vector autoregressions. *Review of Economics and Statistics*, *97*(2), 436–451.

Gilchrist, S., and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, *102(4)*, 1692–1720.

Huber, F., Koop, G., Onorante, L., Pfarrhofer, M., and Schreiner, J. (2023). Nowcasting in a pandemic using non-parametric mixed frequency VARs. *Journal of Econometrics*, *232*(1), 52–69.

Lenza, M., and Primiceri, G. E. (2021). How to Estimate a VAR after March 2020. *Journal of Applied Econometrics (forthcoming).*

Litterman, R. B. (1986). Forecasting with bayesian vector autoregressions: five years of experience. *Journal of Business & Economic Statistics*, *4(1)*, 25–38.

Ng, S. (2021). Modeling macroeconomic variations after COVID-19. Tech. rep., National Bureau of Economic Research.

Schorfheide, F., and Song, D. (2020). Real-time forecasting with a (standard) mixed-frequency VAR during a pandemic. *International Journal of Central Banking (forthcoming).*

Sims, C. A., and Zha, T. (1998). Bayesian methods for dynamic multivariate models. *International Economic Review*, (pp. 949–968).

Wu, J. C., and Xia, F. D. (2016). Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound. *Journal of Money, Credit and Banking*, *48*(2-3), 253–291.

Zakrajsek, E., Lewis, K. F., and Favara, G. (2016). Updating the recession risk and the excess bond premium. *FEDS Notes*, (2016-10), 06.

# A Appendix: Tables and figures

Table A.1 Description of variables

| | Name | Description | Source |
|---|---|---|---|
| 1 | EBP | Excess bond premium as computed by Gilchrist and Zakrajšek (2012). | Zakrajsek, Lewis, and Favara (2016) |
| 2 | S&P 500 | S&P 500 stock index in log levels. | Nasdaq Data Link |
| 3 | Shadow Rate | Fed funds rate shadow rate. | Wu and Xia (2016) and Fred |
| 4 | Consumption (PCE) | Real consumption in log levels. | Fred |
| 5 | Price index | PCE Price Index in log levels. | Fred |
| 6 | Employment | PCE Total nonfarm payroll in log levels. | Fred |
| 7 | Ind. production | Real industrial output in log levels. | Fred |
| 8 | Unemployment rate | Number of unemployed as a percentage of the labor force. | Fred |

Note: All for the January 1975 to December 2022 period, retrieved on February 2023.

Figure A.1 Posterior draws for the autoregressive coefficients with uninformative $\phi$ prior

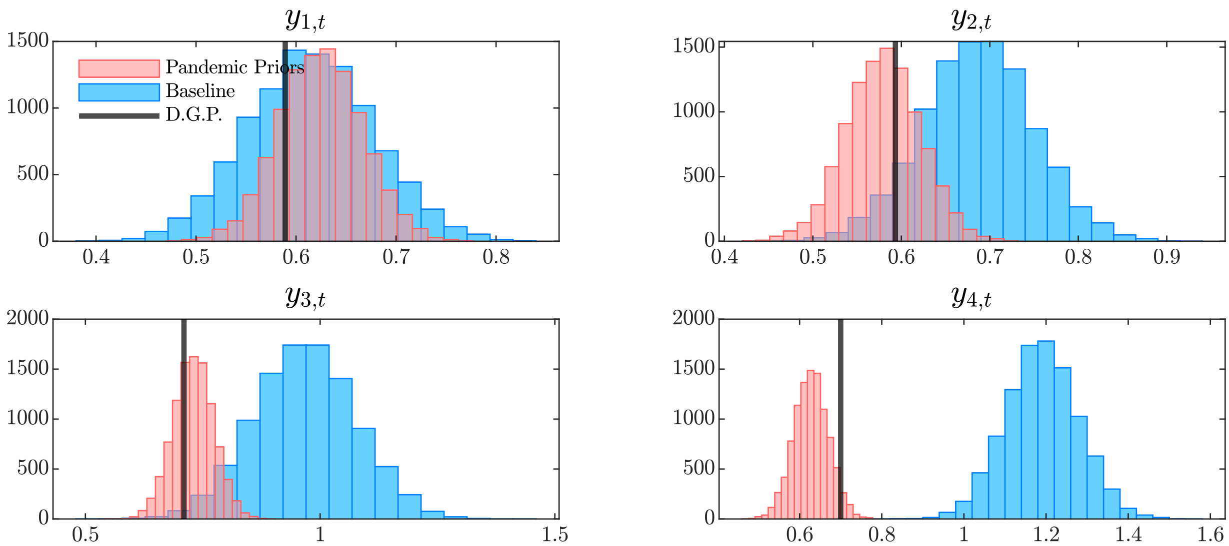


Note: Histograms of the (reduced-form) autoregressive coefficient of the baseline Minnesota Prior (blue bars) and the Pandemic Priors (pink bars) estimations, for each variable in the information set, compared with the data generating process (D.G.P.). $\phi$ set arbitrarily to 0.001 to make it uninformative. Distributions constructed after 10,000 draws from the posterior distribution. The VAR is estimated for 600 simulated periods.

## Figure A.2 Posterior draws for the pandemic time dummies

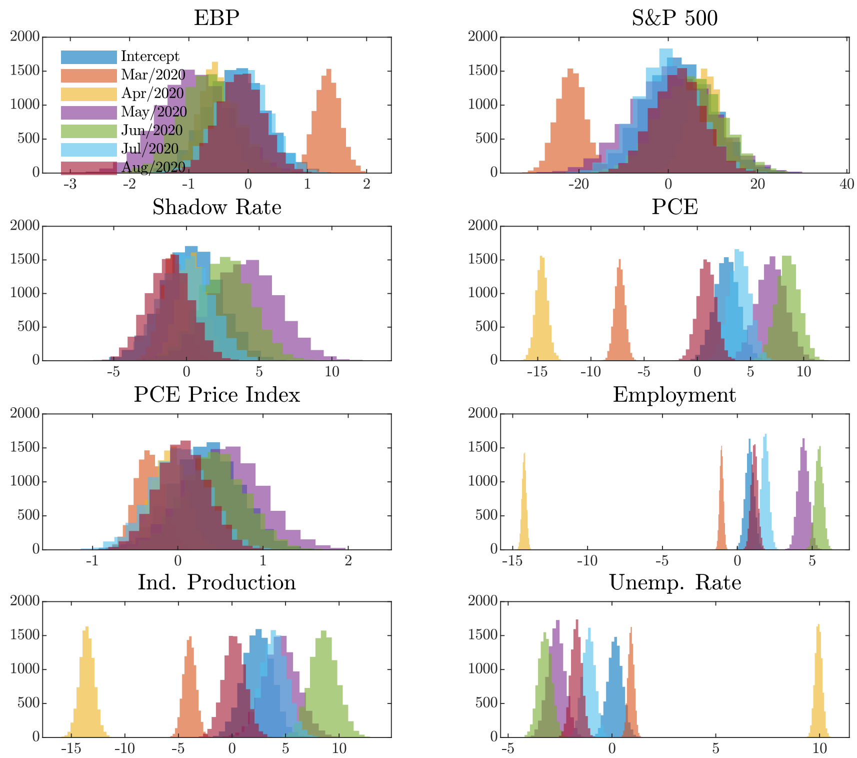


Note: Histograms of the (reduced-form) time dummies for the pandemic period (March 2020 to August 2020), of each variable in the information set. Distributions constructed after 10,000 draws from the posterior distribution. The VAR is estimated from January 1975 to December 2022.

## Figure A.3 Unconditional forecasts as of December 2022

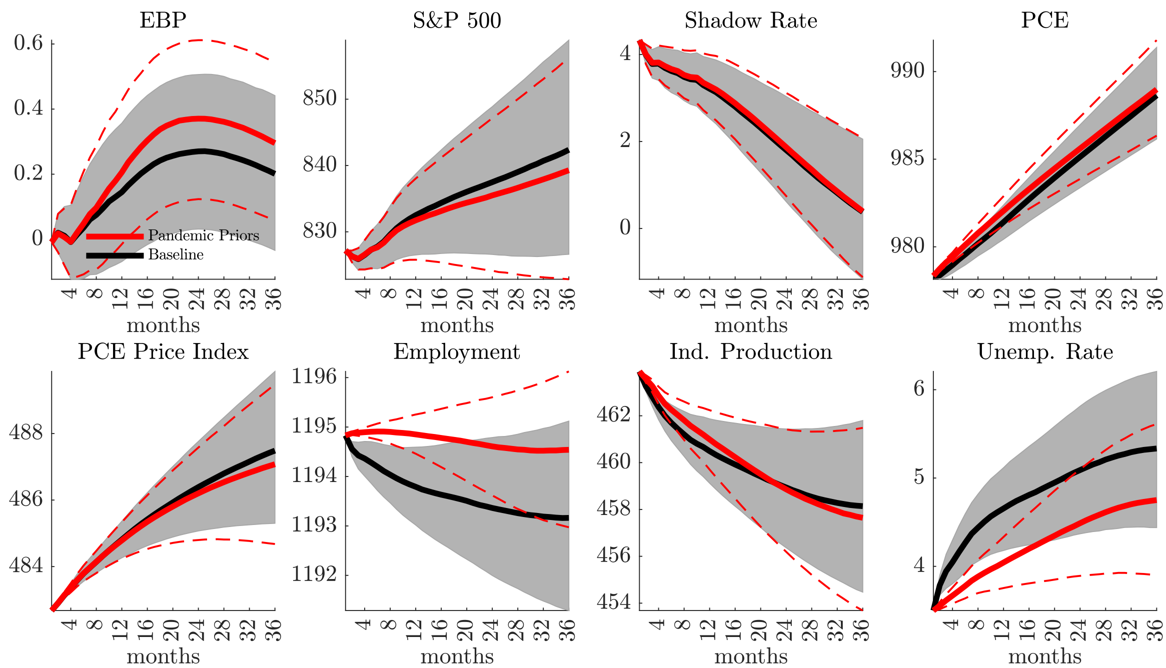


Note: Solid lines are estimated unconditional forecasts and correspond to the posterior median estimates (red with Pandemic Priors, and black as the baseline). The VAR is estimated from January 1975 to December 2022. The gray shaded area and the dashed red lines represent the one standard deviation coverage bands of the forecasts obtained with 10,000 draws from the posterior distribution.

Figure A.4 Impulse responses to a 1 s.d. EBP shock under different $\phi$ levels

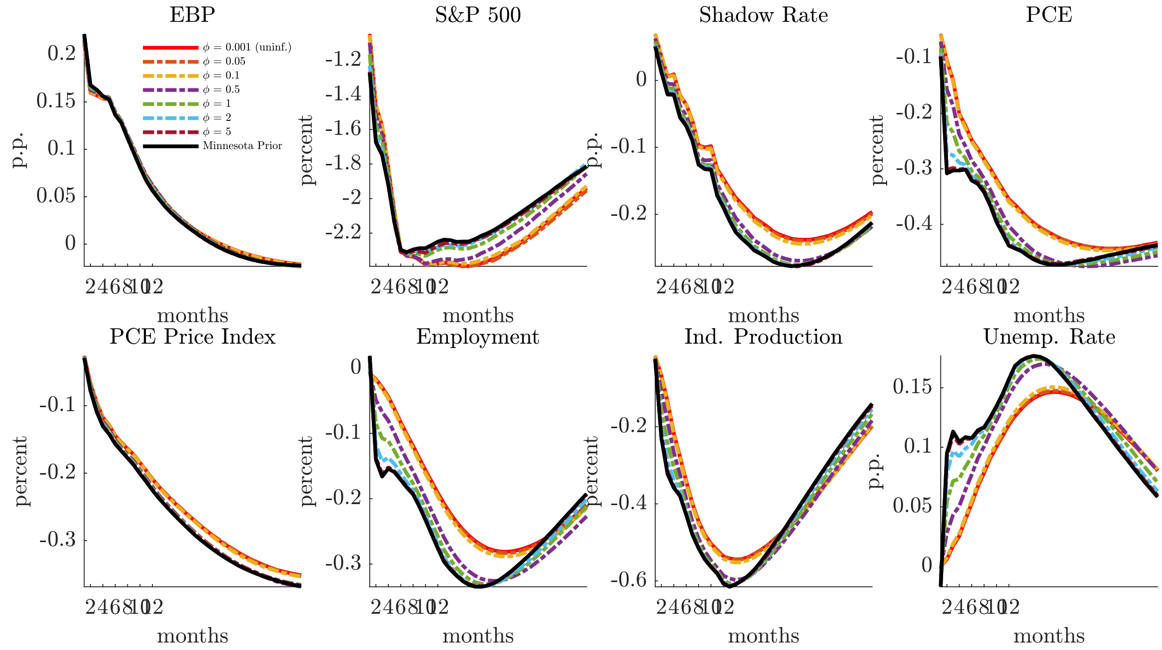


Note: Lines are estimated impulse responses to a standard deviation EBP shock and correspond to the posterior median estimates (black with Minnesota Prior, and colored with Pandemic Priors under different $\phi$). The VAR is estimated from January 1975 to December 2022. Posterior median of the EBP shock obtained with 10,000 draws from the posterior distribution.

Figure A.5 Impulse responses to a 1 s.d. EBP shock

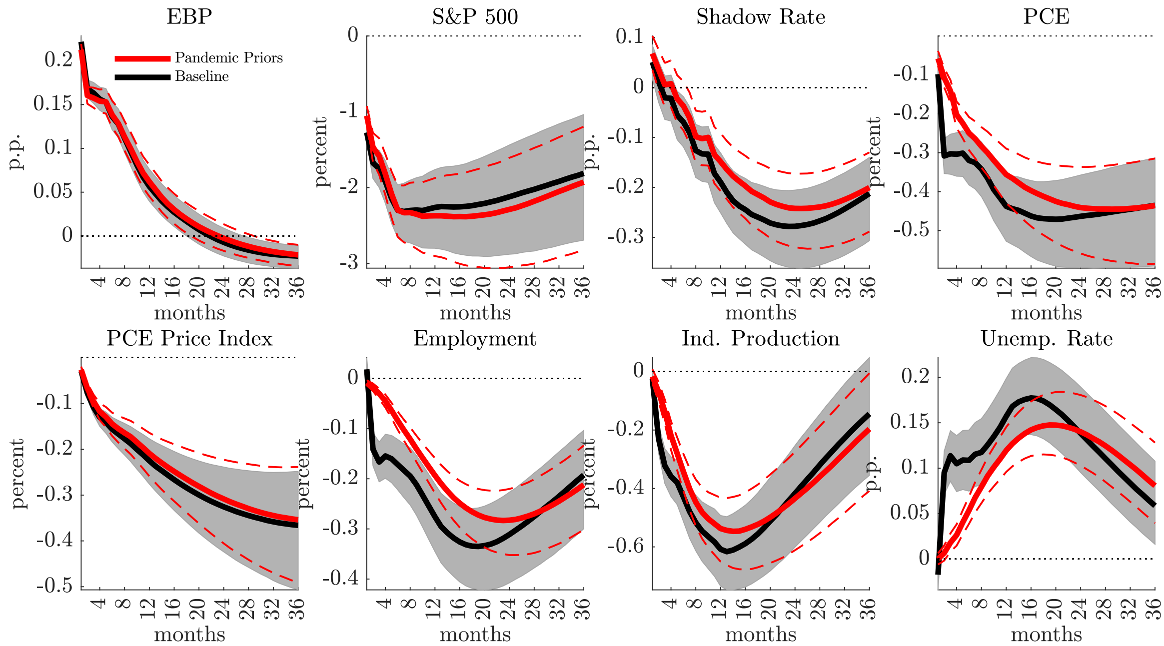


Note: Solid lines are estimated impulse responses to a standard deviation EBP shock and correspond to the posterior median estimates (red with Pandemic Priors, and black as the baseline). The VAR is estimated from January 1975 to December 2022. The gray shaded area and the dashed red lines represent the one standard deviation coverage bands of the EBP shock obtained with 10,000 draws from the posterior distribution.

Figure A.6 Comparison of impulse responses to a 1 s.d. EBP shock

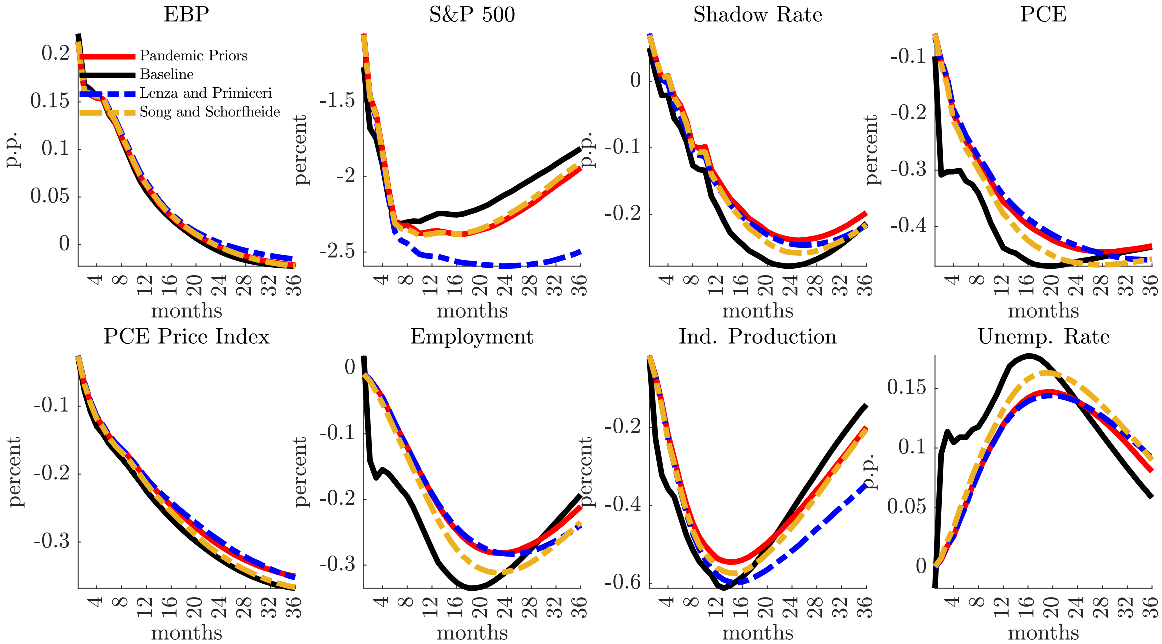


Note: Solid lines are estimated impulse responses to a standard deviation EBP shock and correspond to the posterior median estimates (red with Pandemic Priors, blue with the Lenza and Primiceri (2021) method, yellow with the Schorfheide and Song (2020) method, and black as the baseline). The VAR is estimated from January 1975 to December 2022. Posterior median of the EBP shock obtained with 10,000 draws from the posterior distribution.

# B Appendix: Monte Carlo simulation

I test the ability of the Pandemic Priors to recover the true (reduced-form) coefficients by employing the method on simulated data, with a known data generating process. I produce "abnormal" shocks, affecting all variables simultaneously at a pre-defined time, but with different size and persistence, emulating the environment observed during the COVID-19 pandemic. I simulate a stationary system of four variables and two lags, as in

$$\mathbf{D}_0 \begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \end{bmatrix} = \mathbf{C} + \mathbf{D}_1 \begin{bmatrix} y_{1,t-1} \\ y_{2,t-1} \\ y_{3,t-1} \\ y_{4,t-1} \end{bmatrix} + \mathbf{D}_2 \begin{bmatrix} y_{1,t-2} \\ y_{2,t-2} \\ y_{3,t-2} \\ y_{4,t-2} \end{bmatrix} + \begin{bmatrix} e_{1,t} \\ e_{2,t} \\ e_{3,t} \\ e_{4,t} \end{bmatrix} + \begin{bmatrix} e^*_{1,t} \\ e^*_{2,t} \\ e^*_{3,t} \\ e^*_{4,t} \end{bmatrix}, \tag{B.1}$$

where $e_{i,t}$ are *i.i.d.* innovations with mean 0 and standard deviation 1, and $e^*_{i,t}$ are abnormal shocks that happen simultaneously to all variables at a specific time $t = t^*$, as

$$e^*_{i,t} = \begin{cases} 0, & t < t^* \\ e^*_{i,t^*}, & t = t^* \\ \rho_i e^*_{i,t-1}, & t > t^* \end{cases}. \tag{B.2}$$

I simulate data for 600 periods, with structural coefficients defined as

$$\mathbf{C} = \begin{bmatrix} 0.10 \\ 0.15 \\ 0.05 \\ 0.20 \end{bmatrix}, \quad \mathbf{D}_0 = \begin{bmatrix} 1 & 0.20 & -0.15 & -0.1 \\ 0 & 1 & -0.15 & 0.20 \\ 0 & 0 & 1 & -0.30 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$
$$\mathbf{D}_1 = \begin{bmatrix} 0.65 & -0.10 & 0.10 & 0.05 \\ 0.20 & 0.60 & 0.10 & -0.10 \\ -0.10 & -0.20 & 0.65 & 0.15 \\ -0.05 & -0.15 & 0.20 & 0.80 \end{bmatrix}, \quad \mathbf{D}_2 = \begin{bmatrix} 0.15 & 0 & 0.05 & 0 \\ 0.10 & 0.10 & 0.05 & 0 \\ 0 & -0.01 & 0.10 & 0.05 \\ 0 & -0.05 & 0.10 & 0.10 \end{bmatrix}, \tag{B.3}$$

and abnormal shocks happening at $t^* = 501$ with different size (measured in standard deviations) and persistence for each variable, defined as

$$\begin{bmatrix} \epsilon_{1,t^*} \\ \epsilon_{2,t^*} \\ \epsilon_{3,t^*} \\ \epsilon_{4,t^*} \end{bmatrix} = \begin{bmatrix} 5 \\ 10 \\ 15 \\ 20 \end{bmatrix}, \quad \begin{bmatrix} \rho_1 \\ \rho_2 \\ \rho_3 \\ \rho_4 \end{bmatrix} = \begin{bmatrix} 0.6 \\ 0.7 \\ 0.3 \\ 0.9 \end{bmatrix}. \tag{B.4}$$

Since these shocks are substantially larger than observed in normally distributed series of 600 periods, varying from 5 to 20 standard deviations, the series all jump at $t^* = 501$, and stay at unusually high levels for about 60 periods. Figure B.1 presents the time series of the simulated variables $y_{1,t}$ to $y_{4,t}$.

Figure B.1 Simulated series

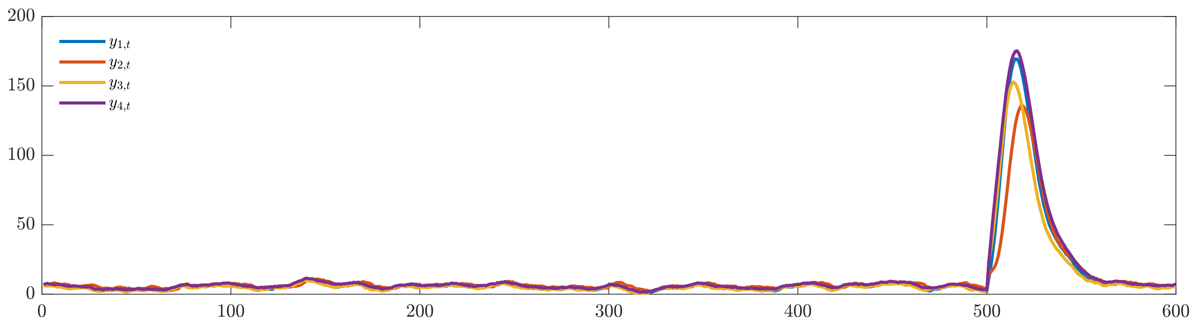


Note: Simulated series with pre-defined data generating process for 600 periods. All series receive a simultaneous shock at $t = 501$, with different magnitudes and persistence.