Board of Governors of the Federal Reserve System

International Finance Discussion Papers

ISSN 1073-2500 (Print)
ISSN 2767-4509 (Online)

Number 1352

August 2022

Pandemic Priors

Danilo Cascaldi-Garcia

Please cite this paper as:
Cascaldi-Garcia, Danilo (2022). "Pandemic Priors," International Finance Discussion Papers 1352. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/IFDP.2022.1352.

Pandemic Priors

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

July, 2022

Abstract

The onset of the COVID-19 pandemic 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 bias 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 with dummy observations by allowing for time dummies. The Pandemic Priors succeed in recovering these historical relationships and the proper identification and propagation of structural shocks.

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

JEL codes: C32, E32, E44

*I thank Ana Beatriz Galvão, Matteo Iacoviello, Francesca Loria, Matteo Luciani, Hyunseung Oh, Filippos Petroulakis, Michele Piffer, Albert Queralto, and Francesco Zanetti for comments and suggestions, and Maia Crook for excellent research assistance. 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 April 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

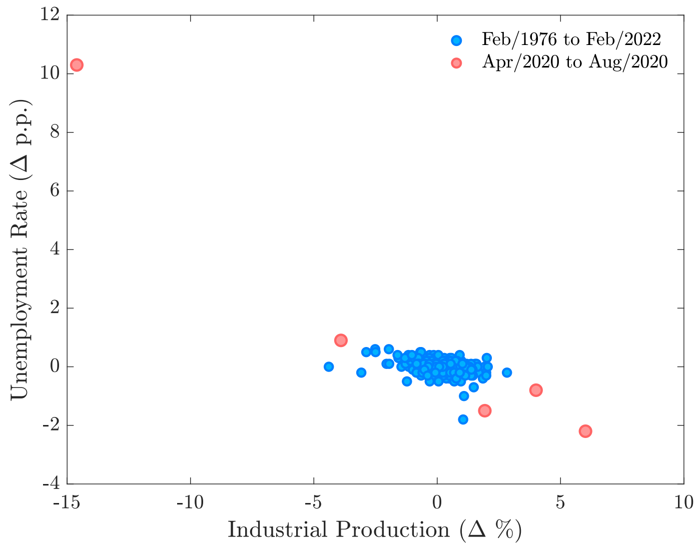

Note: Scatter plot of historical monthly changes of industrial production and unemployment rate. Blue dots correspond to the entire sample (February 1976 to March 2022), and red dots to the most extreme periods of the COVID-19 pandemic (April to August 2020).

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 with uninformative priors, namely Pandemic Priors, which are able to correctly adjust the historical relationship among the variables for the extreme values observed in specific periods.[1]

With an off-the-shelf empirical example, I show that indeed a very low number of extreme observations during the period of March 2020 to August 2020 imply biased 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 Lenza and Primiceri (2021), who 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 volatility 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 COVID-19 period, captured by the individual time dummies. As such, the method proposed here can be directly implemented in closed-form through an extension of the dummy observations procedure described in Bańbura et al. (2010). Still, I show that the Pandemic Priors recover similar impulse responses to those found when using the Lenza and Primiceri (2021) method. The pro-

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

cedure is also an easy linear alternative to discarding extreme observations, as proposed by Schorfheide and Song (2021), or 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). 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 and the Monte Carlo exercise in section 2. Section 3 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 4 shows that the Pandemic Priors recover similar impulse responses to the ones when using the Lenza and Primiceri (2021) estimation procedure. Section 5 summarizes the findings of this paper.

2 Implementation

The Pandemic Priors proposed here builds 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 with Bayesian shrinkage.

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

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

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 a through $a+h$ (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

$$\mathbf{Y}_t = \mathbf{c} + \mathbb{1}_{t=a}\mathbf{d}_a + ... + \mathbb{1}_{t=a+h}\mathbf{d}_{a+h} + \mathbf{Y}_1 + \mathbf{u}_t, \tag{2}$$

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, and the prior on the intercept is diffuse. I take the same diffuse prior stance for the time dummies.

Choices for σ_i, the overall prior tightness λ, the factor $1/k^2$, and the coefficient υ 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, $\mathbf{B}_0$, $\mathbf{\Psi}_0$, S_0, and α_0 are prior expectations, and $\mathbb{E}[\mathbf{\Psi}] = \mathbf{\Sigma}$, or the residual covariance of the Minnesota Prior.

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

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

where $\mathbf{J}_p = \operatorname{diag}(1, 2, ..., p)$, and ϵ 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 ϕ (ordered last in $\mathbf{X}_t$). Following common practice from Litterman (1986), Sims and Zha (1998), and Bańbura et al. (2010), σ_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 ϵ as a very small number makes the prior for the intercept fairly uninformative. I follow the same uninformative approach for ϕ.

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

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

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 τ sets the degree of shrinkage and μ_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 $\mathbf{X}_t^* = [\mathbf{X}_t', \mathbf{X}_d', \mathbf{X}_{sc}']$.

Finally, 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 that happen 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 exercise shows that the (reduced-form) autoregressive coefficients from the data generating process are within the support of the posterior distributions when the Pandemic Priors are applied, while this is not always true for the Minnesota Prior. The larger and persistent the shock is, the more distant the estimated Minnesota Prior coefficient will be from the true value. Also, when facing such unusually large shocks, there is considerably more uncertainty on the autoregressive coefficients with the Minnesota Prior than with the Pandemic Priors. The full experiment is detailed on Appendix B, with posterior distribution comparisons in Figure B.6.

3 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,[2] namely the excess bond premium (EBP, Gilchrist and Zakrajšek, 2012), (log) of the 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 March 2022. I include 12 lags, with fixed overall prior tightness $\lambda = 0.2$, and $\tau = 10 \times \lambda$.[3] I explicitly model the COVID-19 crisis by applying the Pandemic Priors, with six individual dummies for the period March 2020 through August 2020, coinciding with the onset of the pandemic and the very extreme observations in unemployment rate and industrial production (as illustrated by Figure 1).

3.1 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 2 presents the (reduced-form) posterior distributions from 1,000 draws of the intercept, and the intercept shift (intercept plus time dummy) for the period March 2020 to August 2020.[4] Some variables show quite stable intercepts (EBP and PCE price index), but others show large shifts, with more pronounced examples in April 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 stability 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.

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

[3]Results are robust to different lag selections, and λ and τ specifications, and are available upon request.

[4]Figure A.1 in the Appendix shows the posterior distributions for the time dummies, evidencing the uncertainty around the estimations.

Figure 2 Posterior draws for the intercept and pandemic dummies

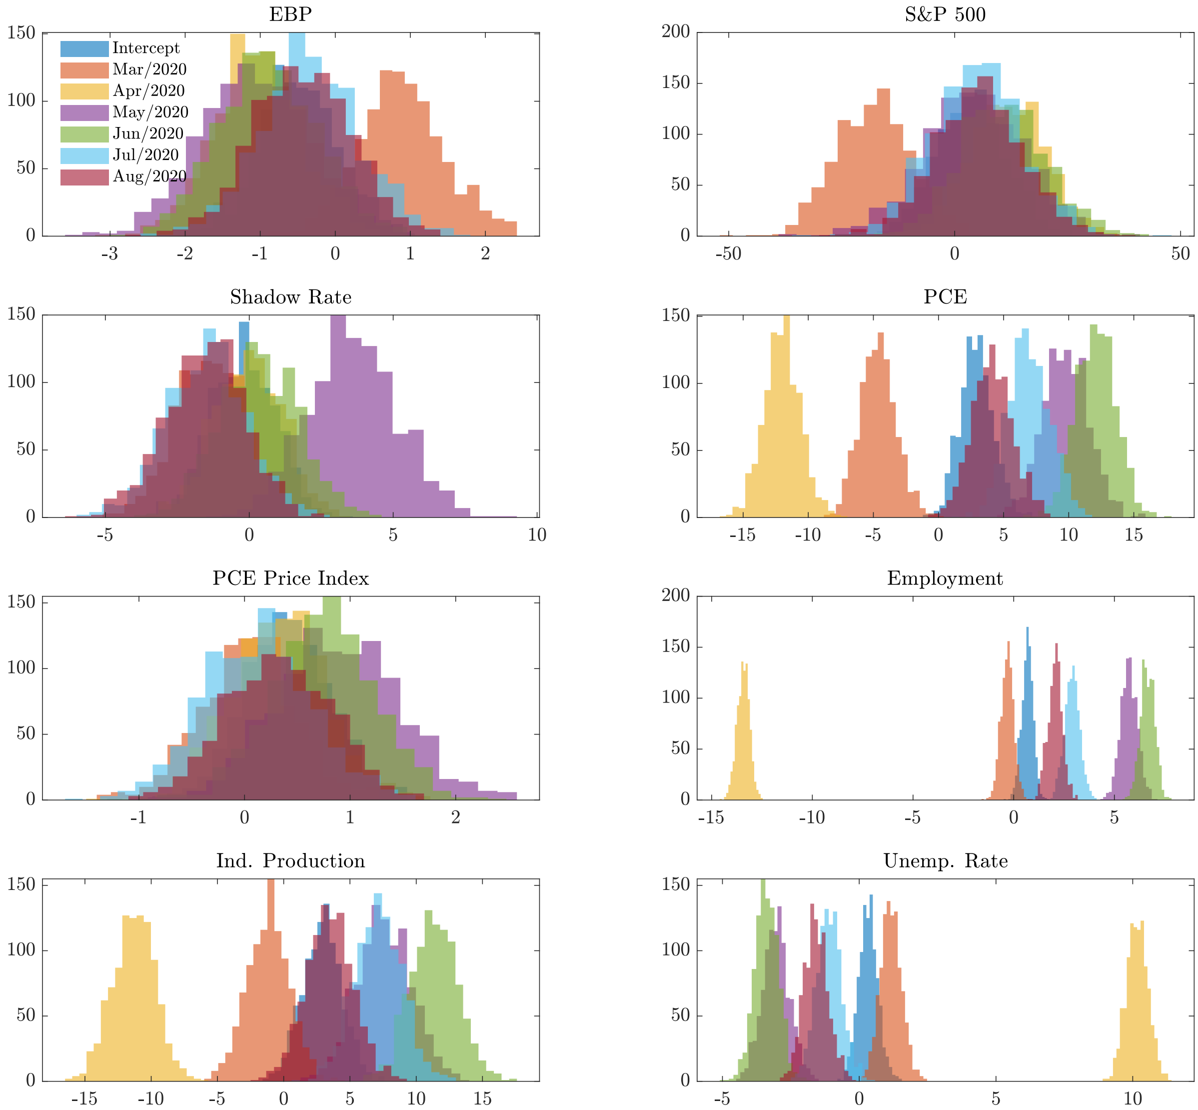

Note: Histograms of the (reduced-form) intercept and the intercept plus the time dummies for the pandemic period (March to August 2020), of each variable in the information set. Distributions constructed after 1,000 draws from the posterior distribution. The VAR is estimated from January 1975 to March 2022.

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 with dummy observations, but 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

Figure 3 Posterior draws for the autoregressive coefficients

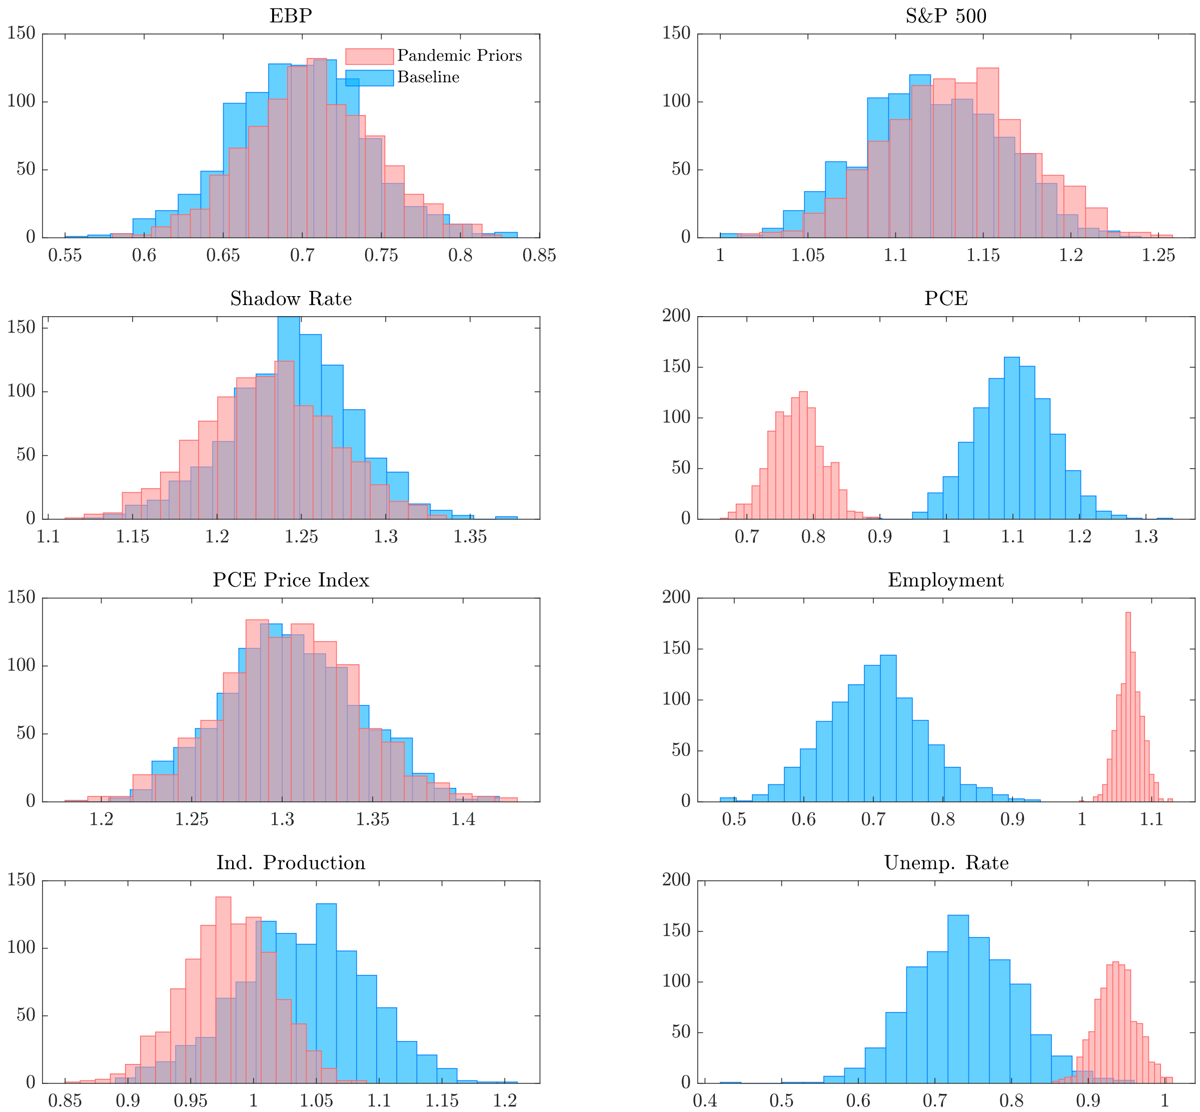

Note: Histograms of the (reduced-form) autoregressive coefficient of the baseline (blue bars) and the Pandemic Priors (pink bars) estimations, for each variable in the information set. Distributions constructed after 1,000 draws from the posterior distribution. The VAR is estimated from January 1975 to March 2022.

second exercise applies the Pandemic Priors. Figure 3 presents the posterior distributions from 1,000 draws of the first lag (reduced-form) autoregressive coefficient of each variable in the information set, for the baseline and the Pandemic Priors setups.[5]

The distributions of posterior draws are substantially different between the baseline and the Pandemic Priors, and there is heterogeneity across variables. While the estimated coefficients are essentially unchanged for EBP, S&P 500, shadow rate, and PCE price

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

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

3.2 ..., 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 method and by the Pandemic Priors, as of March 2022, presented in Figure 4.[6]

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, 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 substantially different unconditional forecasts, implying different economic interpretations. While the baseline model indicates that employment would increase for a couple of months and remain stable over the remainder of the forecast horizon and the unemployment rate would almost immediately increase, the Pandemic Priors provide a picture where employment still rises for about 10 months before stabilizing, and the unemployment rate would decrease for about eight months before increasing.

[6] Figure A.2 in the Appendix reports the unconditional forecasts over a longer horizon.

Figure 4 Unconditional forecasts as of March 2022

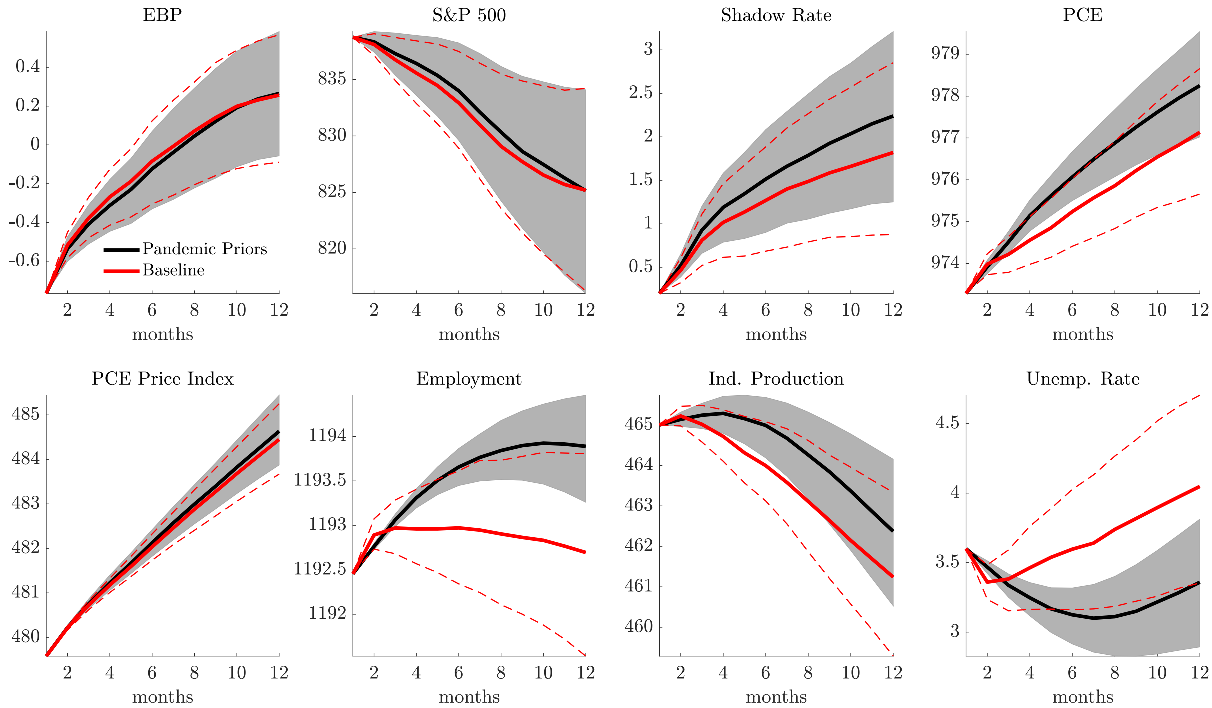

Note: Solid lines are estimated unconditional forecasts and correspond to the posterior median estimates (black with Pandemic Priors, and red as the baseline). The VAR is estimated from January 1975 to March 2022. The gray shaded area and the dashed red lines represent the one standard deviation coverage bands of the forecasts obtained with 1,000 draws from the posterior distribution.

3.3 ..., and for the identification of structural shocks

The extreme observations also impose a tilted view of the economic effects 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 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 5 presents the 12 months ahead impulse response functions of the EBP shock, with solid black lines for the (posterior mean) responses using the Pandemic Priors and solid red lines for the baseline.[7]

The economic effects of an EBP shock using the baseline and the Pandemic Priors

[7]Figure A.3 in the Appendix reports the impulse response functions over a longer horizon.

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

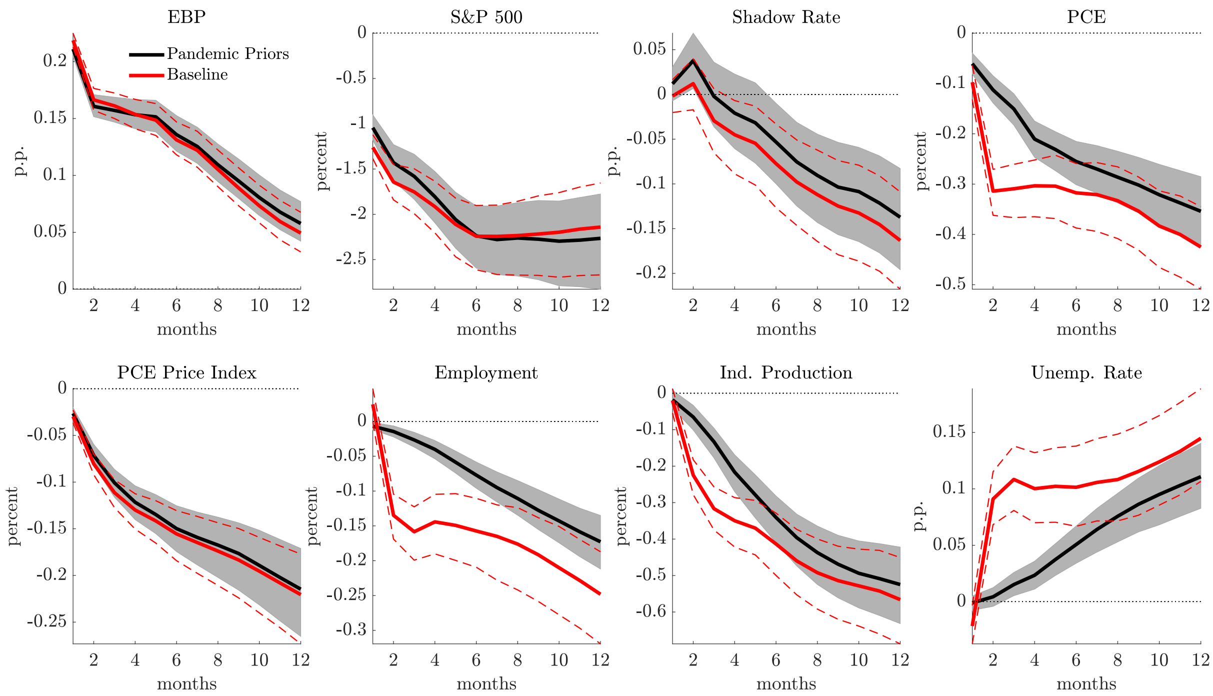

Note: Solid lines are estimated impulse responses to a standard deviation EBP shock and correspond to the posterior median estimates (black with Pandemic Priors, and red as the baseline). The VAR is estimated from January 1975 to March 2022. The gray shaded area and the dashed red lines represent the one standard deviation coverage bands of the EBP shock obtained with 1,000 draws from the posterior distribution.

estimation differ both in size and propagation, and it is heterogeneous over the variables. While the expected effects on the S&P 500, shadow rate, 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, 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 abnormal short-term responses when the pandemic observations are not properly treated.

4 Comparison to Lenza and Primiceri (2021)

The Pandemic Priors are able to recover similar results as with the method proposed by Lenza and Primiceri (2021), but in a simpler and more flexible way. The authors' 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[8]), 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 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 2 and 3, not all variables seem to be reactive to potentially extreme values during the March 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 common scalar shifters and a common decay parameter for the variance of all shocks is a good step to avoid extreme values contaminating 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 dummies for each variable and at each unusual time period, allow for heterogeneous shifts (both in timing and size) and rate of decay over the information set. Still, as shown in Figure 6, when replicating the exercise of an EBP shock with

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

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

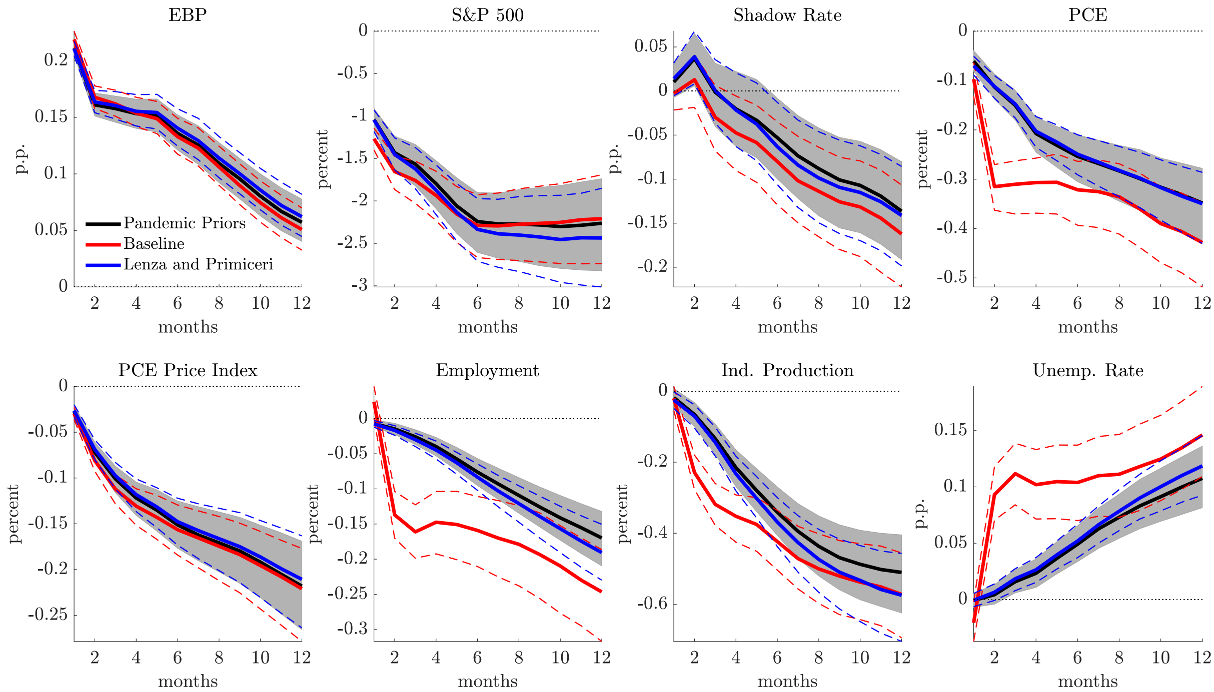

Note: Solid lines are estimated impulse responses to a standard deviation EBP shock and correspond to the posterior median estimates (black with Pandemic Priors, blue with the Lenza and Primiceri (2021) method, and red as the baseline). The VAR is estimated from January 1975 to March 2022. The gray shaded area and the dashed red lines represent the one standard deviation coverage bands of the EBP shock obtained with 1,000 draws from the posterior distribution.

the Pandemic Priors (black lines) and the Lenza and Primiceri (2021) procedure[9] (blue lines), we find results that are almost indistinguishable between the two methods and quite different from the baseline (red lines) with no controls for the pandemic.

5 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 dum-

[9]Estimated using the prior selection procedure proposed by Giannone, Lenza, and Primiceri (2015).

mies for these extreme events that capture such unusual behavior, but accept that there is uncertainty about its potential outcome by imposing a fairly uninformative 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 are potentially (and temporarily) unusual, such as the zero lower bound.

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: Secular trends, large shocks and new data. Tech. rep.

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

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. G., and Mertens, E. (2022). Addressing COVID-19 outliers in BVARs with stochastic volatility. *Review of Economics and Statistics (forthcoming)*.

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.

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. (2021). Real-time forecasting with a (standard) mixed-frequency VAR during a pandemic. Tech. rep., National Bureau of Economic Research.

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)
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 March 2022 period, retrieved on May 2022.

Figure A.1 Posterior draws for the pandemic dummies

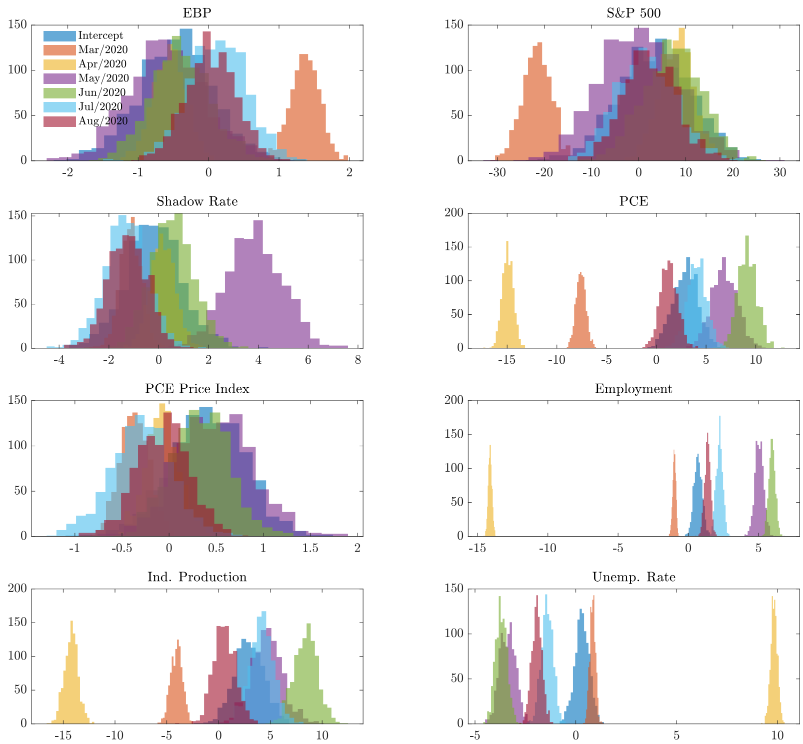

Note: Histograms of the (reduced-form) time dummies for the pandemic period (March to August 2020), of each variable in the information set. Distributions constructed after 1,000 draws from the posterior distribution. The VAR is estimated from January 1975 to March 2022.

Figure A.2 Unconditional forecasts as of March 2022

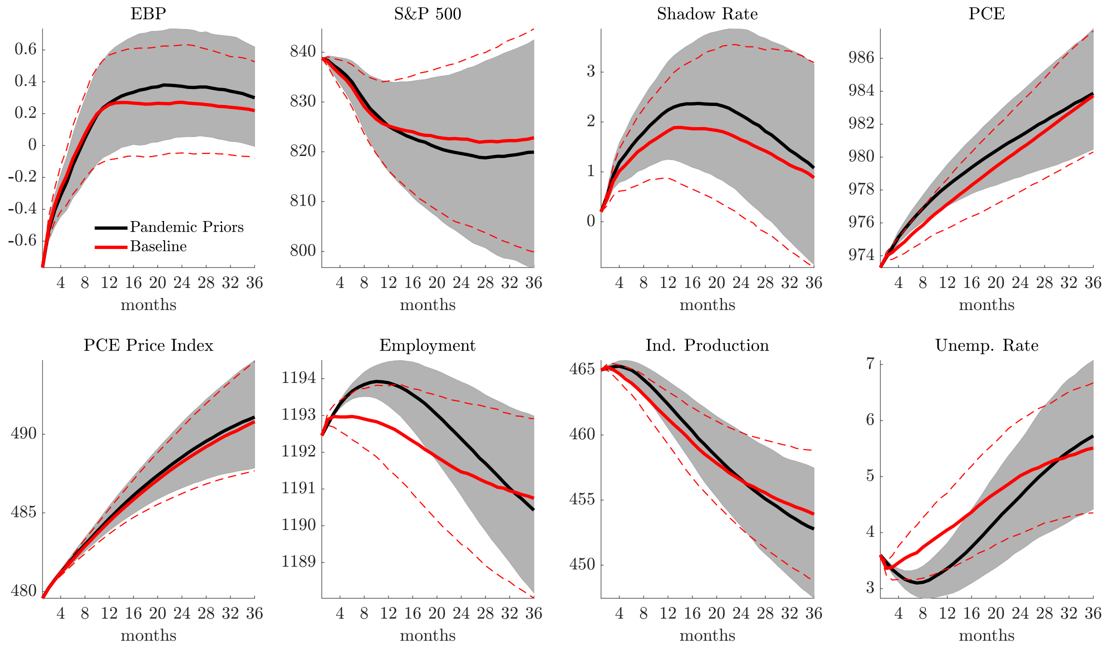

Note: Solid lines are estimated unconditional forecasts and correspond to the posterior median estimates (black with Pandemic Priors, and red as the baseline). The VAR is estimated from January 1975 to March 2022. The gray shaded area and the dashed red lines represent the one standard deviation coverage bands of the forecasts obtained with 1,000 draws from the posterior distribution.

Figure A.3 Impulse responses to a 1 s.d. EBP shock

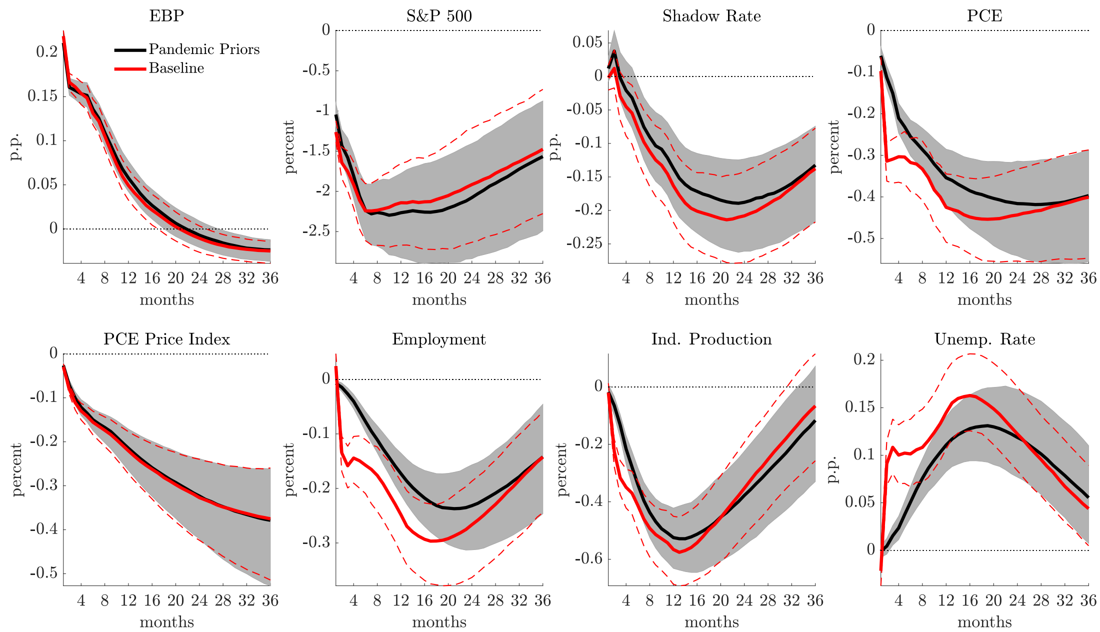

Note: Solid lines are estimated impulse responses to a standard deviation EBP shock and correspond to the posterior median estimates (black with Pandemic Priors, and red as the baseline). The VAR is estimated from January 1975 to March 2022. The gray shaded area and the dashed red lines represent the one standard deviation coverage bands of the EBP shock obtained with 1,000 draws from the posterior distribution.

Figure A.4 Comparison of impulse responses to a 1 s.d. EBP shock

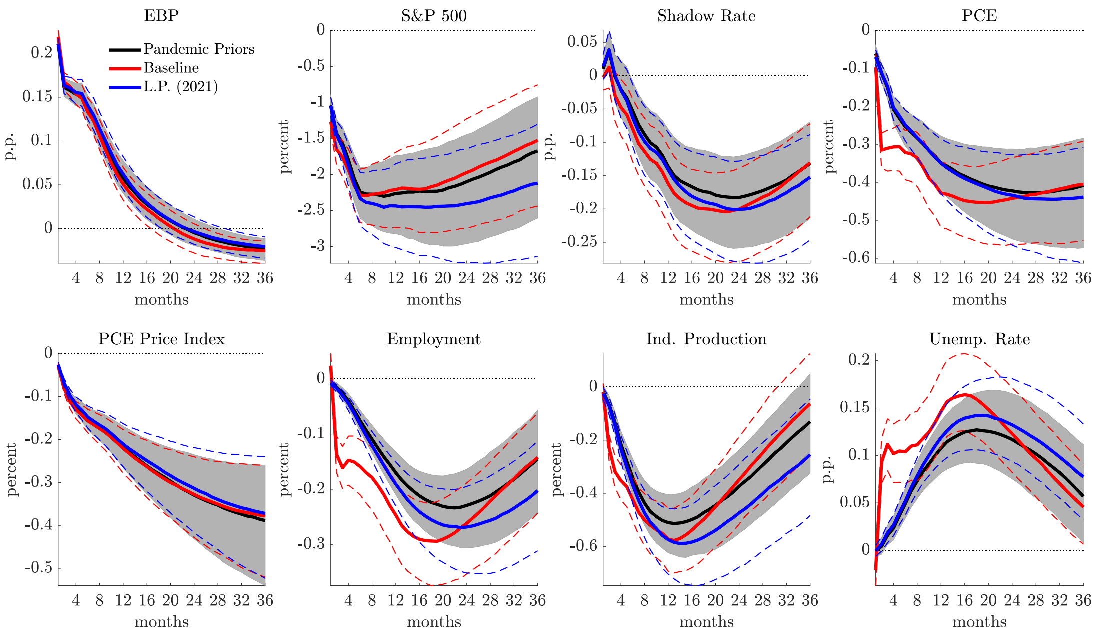

Note: Solid lines are estimated impulse responses to a standard deviation EBP shock and correspond to the posterior median estimates (black with Pandemic Priors, blue with the Lenza and Primiceri (2021) method, and red as the baseline). The VAR is estimated from January 1975 to March 2022. The gray shaded area and the dashed red lines represent the one standard deviation coverage bands of the EBP shock obtained with 1,000 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

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

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

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

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.5 presents the time series of the simulated variables $y_{1,t}$ to $y_{4,t}$.

Figure B.5 Simulated series

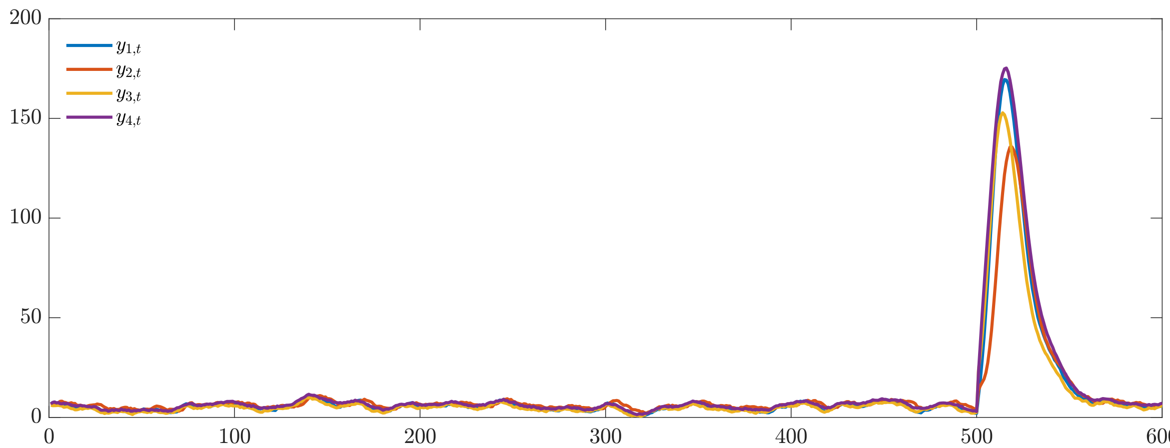

Note: Simulated series with data generating process described in Appendix B, for 600 periods. All series receive a simultaneous shock at $t = 501$, with different magnitudes and persistence.

I estimate two Bayesian VARs 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 ($t = 501, ..., 524$).[10]

In Figure B.6, I compare the baseline (blue bars) and Pandemic Priors (pink bars) posterior distributions of the estimated reduced-form autoregressive coefficients (main

[10]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 B.6 Posterior draws for the autoregressive coefficients

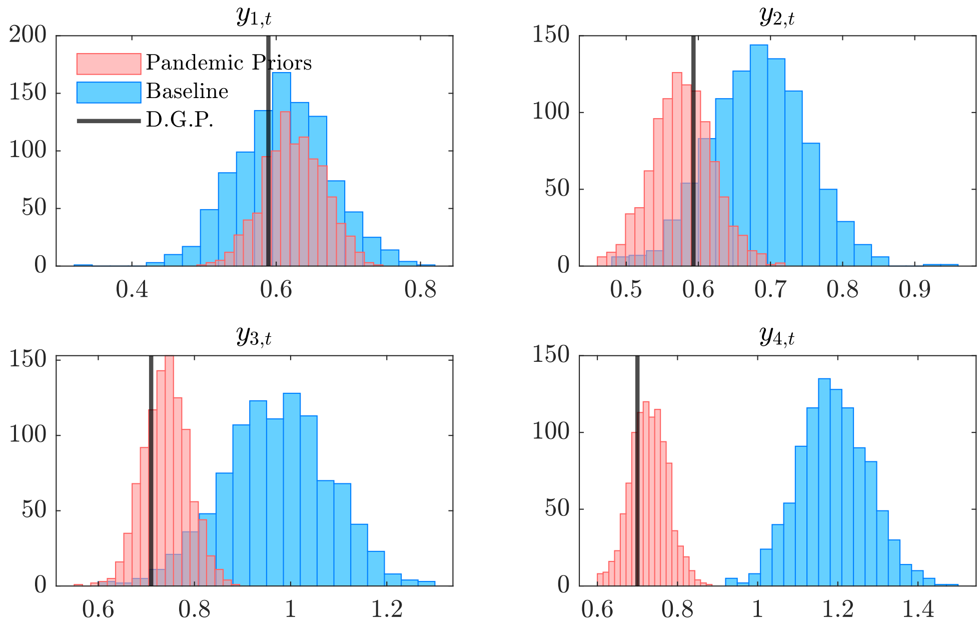

Note: Histograms of the (reduced-form) autoregressive coefficient of the baseline (blue bars) and the Pandemic Priors (pink bars) estimations, for each variable in the information set, compared with the data generating process (D.G.P.). Distributions constructed after 1,000 draws from the posterior distribution. The VAR is estimated for 600 simulated periods.

diagonal of the matrix $\mathbf{B}_1 = \mathbf{A}_0^{-1}\mathbf{A}_1$) with the true coefficient known from the data generating process. Two results stand out from this exercise. First, the larger and persistent the shock is, the more distant the estimated baseline coefficient is 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 incorporates 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 on 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.