Working Paper Series
Congressional Budget Office
Washington, D.C.

Conditional Forecasting With a Bayesian Vector Autoregression

Byoung Hark Yoo
Congressional Budget Office
hark.yoo@cbo.gov

Working Paper 2023-08

November 2023

To enhance the transparency of the work of the Congressional Budget Office and to encourage external review of that work, CBO's working paper series includes papers that provide technical descriptions of official CBO analyses as well as papers that represent independent research by CBO analysts. Papers in that series are available at http://go.usa.gov/xUzd7.

For valuable comments, I thank Dongho Song of Johns Hopkins University and Domenico Giannone of the University of Washington and Amazon. I also benefited greatly from comments by Sebastien Gay, Michael Falkenheim, Michael McGrane, Jeffrey Kling, Jaeger Nelson, Matthew Wilson, Mark Lasky, and other CBO staff members. Gabe Waggoner edited the paper.

www.cbo.gov/publication/59629

Abstract

This paper describes how the Congressional Budget Office uses a Bayesian vector autoregression (BVAR) method to generate alternative economic projections to the agency's baseline. The BVAR includes a wide range of key economic variables that are needed to approximate budget outcomes. Its estimation methods avoid overfitting, a situation in which a model fits historical data well while having a poor ability to project future values.

Given targets of future values of some variables such as inflation, the BVAR generates economic projections consistent with the targets and historical dynamics of the variables in the model. The BVAR is a flexible framework that can incorporate new variables and impose conditions for alternative economic projections. It also has forecasting performance comparable to that of CBO's baseline forecasting method.

Keywords: Bayesian vector autoregression, conditional forecasting, scenario construction

JEL Classification: C32, C53

Contents

Introduction

The Congressional Budget Office develops alternative economic projections (or economic scenarios) that differ from the agency's baseline projection used in the annual *Budget and Economic Outlook* (for the latest publication, see CBO 2023c). Those scenarios often are created to analyze the economic effects of legislative proposals and their budgetary consequences, and the legislation itself is the cause of any deviations in economic variables from the baseline. CBO's analysis generally uses a suite of structural models to estimate the short- and long-term economic effects of legislation channeled through demand and supply (see CBO 2014).

In other cases, at the request of the Congress, CBO develops economic scenarios based on alternative values or paths of one or more variables such as inflation or interest rates, without specifying what might have caused the variables to reach those levels. For example, the agency may be asked to project the federal budget in an economic scenario in which interest rates remain higher for a few years than they are in CBO's baseline projection. This paper describes how CBO projects economic variables in such cases and uses them to develop a projection of revenue, spending, and the deficit. The Bayesian vector autoregression (BVAR) offers a way by which historical relationships among economic variables are used to generate a forecast for a range of variables given projections for several of the key variables.

CBO's simplified budget model to approximate budget outcomes requires a projection of a large set of economic variables. That model consists of an incomes model and a budgetary feedback model (Frentz and colleagues 2020; CBO 2023a). CBO uses a BVAR to project that large set of economic variables given targets for the future values of some of them. The process is called conditional forecasting, for which CBO uses the procedure described in Crump and colleagues (2021). Once a scenario is generated in the BVAR, CBO runs the simplified budgetary model with the scenario to calculate budget outcomes in comparison with the CBO baseline projection (see Figure 1). In addition to generating scenarios, a BVAR can be used to quantify the forecast's uncertainty.

Despite the large set of variables and many parameters to be estimated, the BVAR overcomes "overfitting"—a situation in which a model fits historical data well while having a poor ability to project the future—by imposing a structure called Bayesian shrinkage. As a result, a projection of a variable is more likely to be influenced by recent data than by older data. The BVAR can flexibly include new variables and impose new conditions and is reported to have forecasting power comparable to that of other macroeconomic models in the literature. The BVAR is suitable for modeling scenarios that lack a clear structural cause.

After some background review of BVARs, this paper shows how to construct a scenario by using a BVAR. The paper also describes how CBO estimates a BVAR and uses one to generate a scenario. The results of an out-of-sample test of the BVAR are then reported.

Background on BVARs

How BVARs Are Developed and Used

Sims (1972, 1980a, 1980b) developed vector autoregressions (VARs) as an alternative to large-scale structural macroeconomics models. Litterman (1980) introduced BVARs to overcome the overfitting issue that occurs when the model includes many variables. Since then, many researchers have further developed BVAR techniques and investigated their performance. For a comprehensive introduction to forecasting with BVARs, see Karlsson (2013).

BVARs use a technique called Bayesian shrinkage to improve the performance of a conventional VAR with many variables. The number of parameters in a VAR is proportional to the square of its number of variables. As the number of parameters in a model increases, its tendency toward overfitting increases. A shrinkage method avoids that problem by imposing tight prior distributions around zero for parameters that are hard to estimate accurately.

Many researchers have examined the forecasting power of BVARs. Bańbura and colleagues (2010) used 131 monthly variables from 1959 to 2003 and showed that large VARs with shrinkage forecast better than VARs with a few (3 or 7) selected variables. Koop (2013) used 168 quarterly variables from 1959 to 2008 and showed that large VARs with shrinkage tended to forecast better than factor methods in which the behavior of a large set of variables is driven by a limited number of unobserved components. Bańbura and colleagues (2015) introduced an algorithm to compute the distribution of conditional forecasts and generated scenarios using 26 quarterly variables. Giannone, Lenza, and Primiceri (2015) suggested a “hierarchical” modeling approach, in which the parameter measuring the level of shrinkage is optimized to improve forecast accuracy. They showed that the hierarchical BVARs forecast better than conventional VARs of the same number of variables. Using U.K. data, Domit, Monti, and Sokol (2019) showed that BVARs forecast growth of real gross domestic product (GDP) better than the microfounded macroeconomic model used by the Bank of England, whereas BVARs and the microfounded macroeconomic model showed similar performance to forecast consumer price index (CPI) inflation. Angelini and colleagues (2019) used quarterly data of the four biggest countries of the euro area and showed that the multicountry BVAR model produced forecasts comparable to the Eurosystem official forecasts. When Crump and colleagues (2021) combined the estimation method by Giannone, Lenza, and Primiceri (2015) and the conditional forecast method by Bańbura and colleagues (2015), BVARs produced out-of-sample predictions comparable to those of the Survey of Professional Forecasters and the Federal Reserve Board of Governors Greenbook.

Applications of BVARs

BVARs are widely used for macroeconomic analysis as an alternative to structural macroeconomic models. Although early versions of BVARs faced challenges in policy analysis or structural breaks, more recent innovations have made progress in overcoming those limitations. See Stock and Watson (2001) for an overview of VARs.

Model flexibility. BVARs can be easily modified to include new variables because they don't rely on structural assumptions or specifications, such as the national income accounting identity. In addition, BVARs can easily incorporate conditions to generate a scenario without specifying the exact causes of the conditions. For example, a BVAR can generate a scenario in which inflation is 4 percent in the next year without specifying whether a demand factor or a supply factor would lead to such level of inflation. Thus, BVARs are useful especially when an analysis requires new variables or new conditions and needs to be implemented within a short period.

Forecasting. Drawing on historical relationships among macroeconomic and financial variables, BVARs have shown forecasting performance comparable to that of other macroeconomic models for a horizon of a few years. The forecasting power is one reason that organizations such as the Federal Reserve often use BVARs to provide guidance and validation for microfounded macroeconomic models, such as dynamic stochastic general equilibrium models.

Policy changes. Structural BVARs (BVARs with identified causal relationships) can incorporate structural relationships among macroeconomic variables and can be used to analyze the effects of unexpected policy changes, such as an unexpected increase in the federal funds rate. Structural BVARs also can be used to examine the effects of expected policy changes by imposing corresponding structural relationships. However, imposing structural relationships for all the variables in a BVAR can be infeasible when it contains many variables. Thus, BVARs may not be well suited for the case when a policy change is introduced by a legislative proposal that states a clear causal link among variables. For an example of scenario analysis with structural VARs, see Antolín-Díaz, Petrella, and Ramírez (2021).

Structural breaks. Standard BVARs incorporate the assumption that the structure of the economy and the dynamics of economic variables were stable during the estimation period and will remain so during the projection period. Moderate changes in policies or dynamics over those periods might have little effect on their performance. However, some significant shifts in economic dynamics or policies have been reported (see Sims and Zha 2006). Challenges with structural breaks are not unique to BVARs: All macroeconomic models face challenges in identifying historical structural breaks and projecting economic variables in the presence of possible economic changes. Structural models can address structural breaks by revising parameters or including new features in the equations when the cause of a structural break is well understood, such as when the break results from a shift in monetary policy stance. When the cause is not clearly understood or is improperly specified, structural models could generate misleading results. BVARs have been developed to allow parameters or volatilities to change over time to incorporate structural breaks even without specifying the causes or timing of possible breaks. Those models have shown that incorporating time-varying parameters can improve forecasting performance (see D'Agostino, Gambetti, and Giannone 2013). The time-varying features have not been introduced much for large-scale BVARs because of computational constraints, although Koop and Korobilis (2013) used time-varying parameter modeling in BVARs.

Long-run dynamics. Some BVARs use variables in levels, in which case long-term forecasts may not be anchored to long-term values—for example, in 10 years, the horizon in the simplified budget model. Extending the forecasting horizon with no long-term anchor often leads to projections that perpetually fluctuate in some scenarios. One simple way to avoid such situations is to anchor the long-run values of variables with a condition, for example, by setting real GDP equal to that of the unconditional forecast in the 10th year or by setting inflation for personal consumption expenditures to 2 percent for an extended period. Anchoring long-run values also can be implemented statistically by shifting the distribution around the long-run value (called entropic tilting; see Crump and colleagues 2021). Controlling long-run values or dynamics also can be implemented by imposing long-run priors based on economic theories. Giannone, Lenza, and Primiceri (2019) showed that the long-run dynamics can be controlled by using priors.

Example of Constructing Scenarios

CBO recently used a BVAR to construct an economic scenario with alternative paths for interest rates and then analyzed its budgetary implications (Swagel 2022). Since CBO published its baseline projection in July 2021, the expectations of interest rates have changed as the Federal Reserve started to raise them. The economic scenarios were designed to incorporate that change of expectations.

First, the agency identified the top six forecasters who projected highest interest rates (3-month and 10-year Treasury rates) among 38 forecasters in the March 2022 *Blue Chip Economic Indicators*. Then CBO constructed the set of conditions for the higher-interest-rates scenario by taking the average of the projected values by the top six forecasters projected for 2022 and 2023 for each 8 variables in the report. For projections for 2024 to 2031, CBO used either the average of the top 10 projections or all the projections according to availability. Likewise, the agency identified the bottom six forecasters and constructed the set of conditions for the lower-interest-rates scenario (see Figure 2A).

Second, CBO calculated the differences between the projections of each scenario and CBO's July 2021 baseline projections for 8 variables in the condition set. Then the agency projected the other 18 variables in the BVAR over 10 years where the sums of the differences and the unconditional forecasts of the BVAR of the 8 variables are fixed as conditions for each scenario (see Figure 2B).

Third, the agency recentered the scenarios by adding the differences between the scenarios and the BVAR unconditional forecast to the agency's baseline projection. Doing so is necessary because the BVAR unconditional forecast is different from CBO's baseline projection in general.

Then the agency used those updated scenarios as inputs for the simplified budget model to approximate budget outcomes of each scenario in comparison with the CBO baseline projection.

The scenario with higher interest rates projected more employment and output than the one with lower interest rates. That finding is consistent with historically dominant dynamics in which economic expansion periods are usually accompanied by higher interest rates. Whereas higher interest rates resulted in more interest payment on the federal debt, higher output led to more revenue, offsetting the negative effects. CBO (2023b) shows another example of constructing scenarios.

How CBO Estimates the Model and Generates Scenarios

CBO's goal in designing the BVAR is to forecast the key economic variables required by the simplified budget model to approximate budget outcomes such as government spending, revenue, and deficits. The agency uses Bayesian methods to estimate the model for that large system of variables without overfitting. Conditional forecasting techniques are then used to generate scenarios given future target values for some variables in the BVAR.

Data

The data set contains historical values of the 26 variables included in the BVAR from the first quarter of 1959 to the fourth quarter of 2021 (see Table 1). The data set contains all the inputs for the simplified budget model, such as nominal gross national product and the consumer price index for medical expenditure. The outputs of the simplified budget model include government spending, revenue, and deficits; the BVAR does not include those variables. The BVAR measures variables in levels, not in difference, to use the full information of the variables, such as cointegration relationships, following the practice of Crump and colleagues (2021). To exclude the unusual dynamics during the COVID-19 pandemic, CBO uses data up to the fourth quarter of 2019 for estimation. That approach is supported by Schorfheide and Song (2022), whereas other papers such as Lenza and Primiceri (2022) and Carriero and colleagues (2022) suggest different approaches to handle data from during the pandemic.

Model and Bayesian Estimation

A general specification of a VAR is as follows:

$$y_t = c + B_1 y_{t-1} + \cdots + B_p y_{t-p} + \epsilon_t \ (t = 1, \cdots, T), \tag{1}$$

where y_t is an $n \times 1$ vector of variables at time t, c is an $n \times 1$ vector of constants, p is the number of lags included in the VAR, and B_s $(s = 1, \cdots, p)$ is an $n \times n$ matrix of coefficients. The error terms are reduced-form errors, not structural errors, because no relationship is explicitly specified among the contemporary variables in the VAR. The $n \times 1$ vector of error terms, ϵ_t, is assumed to be normally distributed with a mean vector of 0 $(n \times 1)$ and a covariance matrix Σ [or $\epsilon_t \sim N(0, \Sigma)$]. T is the end period of the historical data, and the set of parameters to be estimated is $\theta \equiv \{c, B_1, \cdots, B_p, \Sigma\}$. The set θ becomes very large for large *n* because the number of the parameters of a coefficient matrix (B_s) is n^2. For example, the model would have about 1,800 parameters if it had 20 variables and *p* was 4. The traditional VAR techniques would then be vulnerable to overfitting and would tend to generate poor out-of-sample forecasting performance.

To avoid overfitting, CBO uses Bayesian methods to estimate a large-scale VAR with a prior (called the Minnesota prior in the literature) where each variable follows an independent random walk with possible drift. That specification yields an approximation of the dynamics of many economic variables. First, CBO sets the priors of the parameters (θ) as follows:

$$\Sigma \sim IW(\Psi, d) \text{ and} \tag{2}$$

$$\beta|\Sigma \sim N(b, \Sigma \otimes \Omega), \tag{3}$$

where *IW* stands for the inverse Wishart distribution, specified by an $n \times n$ diagonal matrix (Ψ) and a scalar (d). The diagonal elements of Ψ will be determined as explained later, and d (the degree of freedom) is set as $n + 2$ so that the expected value of Σ becomes Ψ or $E[\Sigma] = \Psi$, where E stands for the mathematical expectation. β is the vector of the constants and the coefficients (or $\beta = vec\left(\left[c, B_1, \cdots, B_p\right]'\right)$, where *vec* transforms a matrix to a vector) and is normally distributed. Its mean and covariance matrix are *b* and $\Sigma \otimes \Omega$, respectively, where ⊗ stands for the Kronecker product and Ω is a function of Ψ, as specified later.

Then CBO sets *b* and $\Sigma \otimes \Omega$ such that the first and second moments of β should be as follows:

$$E\left[(B_s)_{ij}\middle|\Sigma\right] = \begin{cases} 1 \text{ if } i = j \text{ and } s = 1 \\ 0 \text{ otherwise} \end{cases} \text{ and} \tag{4}$$

$$cov\left[(B_s)_{ij}, (B_t)_{kl}\middle|\Sigma\right] = \begin{cases} \frac{\lambda^2}{s^2}\frac{\Sigma_{ik}}{\psi_j} \text{ if } l = j \text{ and } t = s \\ 0 \text{ otherwise} \end{cases}, \tag{5}$$

where $(B_s)_{ij}$ is the *i*th row and *j*th column element of $B_s (s = 1, \cdots, p)$, Σ_{ij} is the *i*th row and *j*th column element of Σ, ψ_j is the *j*th diagonal element of Ψ, and λ is a parameter to adjust the size of the second moments of β. Ψ and λ are the parameters determining the priors in Equations (4) and (5), and such parameters are called hyperparameters.

The priors are used to address the overfitting issue in which a VAR with many variables has many parameters to be estimated and tends to have poor out-of-sample forecasting accuracy. Equation (4) indicates that the expectation of the VAR coefficient of a variable corresponding to the first lag of the variable [or $(B_1)_{ii}$ for $i = 1, \cdots, n$] is 1 and the other coefficients are 0. That specification corresponds to the idea that each variable follows an independent random walk process, potentially with drift. Meanwhile, Equation (5) specifies the variances and covariances of the coefficients. It shows that the variance of a coefficient decreases by a factor of the square of the lag difference (s^2) when $i = j = k = l$ (that specification is called shrinkage), whereas λ determines the tightness of the variances. For example, the variance of $(B_1)_{ii}$ $\left(\text{or } \frac{\lambda^2}{1^2}\frac{\Sigma_{ii}}{\psi_i}\right)$ is 4 times larger than that of $(B_2)_{ii}$ $\left(\text{or } \frac{\lambda^2}{2^2}\frac{\Sigma_{ii}}{\psi_i}\right)$ for $i = 1, \cdots, n$. That setting ensures that the coefficients at lag 1 have looser prior distribution, whereas those at longer lag have tighter prior distributions. λ determines the overall tightness of the prior distributions.

CBO estimates the hyperparameter λ by using the Bayesian method and calibrates the other hyperparameters by using other methods to ensure stable implementation. ψ_j (for $j = 1, \cdots, n$) is set to be the estimated variance of the error term after an autoregressive model is estimated for each variable. The priors of the constants (c) are set as independent normal distributions with mean 0 and an arbitrary large variance. The variance is arbitrarily chosen to be very large so that the prior should not be binding in the estimation. The mean and the variance of the normal distribution for the constants also are hyperparameters, and CBO fixes the variance as 10,000,000.

The posterior distribution of the parameters (θ) and the hyperparameter (λ) is calculated with Bayes's rule, as follows:

$$p(\lambda, \theta|y) \propto p(\lambda, \theta)p(y|\lambda, \theta) \ \propto p(\lambda)p(\theta|\lambda)p(y|\theta), \tag{6}$$

where y is the vector of the historical values of all the variables or $y = (y_1, \dots, y_T)$, $p(\theta|\lambda)$ is the prior distribution described above, and $p(y|\theta)$ is the likelihood function of the BVAR. Note that $p(y|\theta) = p(y|\lambda, \theta)$. $p(\lambda)$ is the prior distribution of λ, and the gamma distribution is chosen for the distribution to make sure that λ is positive. CBO uses the hierarchical Bayesian approach suggested by Giannone, Lenza, and Primiceri (2015) to generate the draws of λ and θ ($\lambda^{(g)}$ and $\theta^{(g)}$ for $g = 1, \dots, G$) from the posterior distribution. See the appendix for details of the procedure.

Conditional Forecasting

CBO projects the variables in the BVAR with no conditions, an approach called unconditional forecasting. The process is to compute the predictive density function defined as

$$p(y_{T+1}, \dots, y_{T+H}|y) = \int p(y_{T+1}, \dots, y_{T+H}|y, \theta)p(\theta|y)d\theta, \tag{7}$$

where H is the forecasting horizon. The first term in the integral on the right-hand side is determined by Equation (1), and the second term is the marginal density function of the posterior distribution derived in Equation (6). See the appendix for details of the calculation.

CBO also generates scenarios by projecting the variables in the BVAR under some conditions on the future values of some variables in some periods. That process is called conditional forecasting and is done in the form of the conditional predictive density function defined as

$$p(y_{T+1},\ldots,y_{T+H}|y,C)=\int p(y_{T+1},\ldots,y_{T+H}|y,\theta,C)p(\theta|y)d\theta, \tag{8}$$

where C is a set of given conditions for a scenario. See the appendix for details of the calculation. The conditions can be imposed on any variables among those in the BVAR and for any time.

In practice, CBO writes the model in state space form, from which the transition equation is derived from the VAR specification in Equation (1) as

$$x_t = F\,x_{t-1} + u_t, \tag{9}$$

where x_t (called the state variables) is defined as

$$x_t = \begin{pmatrix} y_t \\ y_{t-1} \\ \vdots \\ y_{t-p+1} \\ c \end{pmatrix}, \tag{10}$$

and F is defined as

$$F = \begin{pmatrix} B_1 & B_2 & \cdots & B_p & I_n \\ I_n & 0_n & \cdots & 0_n & 0_n \\ \vdots & \ddots & \ddots & \vdots & \vdots \\ 0_n & \cdots & I_n & 0_n & 0_n \\ 0_n & \cdots & 0_n & 0_n & I_n \end{pmatrix}. \tag{11}$$

The error term, u_t, is defined as

$$u_t = \begin{pmatrix} \epsilon_t \\ 0_n \\ \vdots \\ 0_n \end{pmatrix}. \tag{12}$$

Then the measurement equation of the state space model is defined as

$$y_t^* = G_t x_t (t = T+1,\cdots,T+H) \tag{13}$$

where y_t^* is an $m_t \times 1$ vector ($0 \le m_t \le n$), which contains conditioned future values for some variables in y_t at time t. m_t is the number of the conditions at time t and can vary over the forecasting period ($t = T+1,\cdots,T+H$). If no condition exists at time t, y_t^* should be empty. G_t is an $m_t \times (n \cdot p)$ matrix that consists of 0s and 1s, so that it identifies the variables conditioned at time t. For example, if a scenario has one condition stating that the value of the ith variable in the model is $\tilde{y}_{i,h}$ at time $T + h$, then $y^*_{T+h} = [\tilde{y}_{i,h}]$ and G_t becomes a $1 \times (n \cdot p)$ matrix where the element of ith column is 1 and the others are 0s.

The flexibility of conditional forecasting comes from the fact that one can adjust the measurement equation to set any number of conditions for any variable in the model for any time as a condition. Technically, conditional forecasting is equivalent to estimating the unobservable/missing state variables given the observable/nonmissing variables.

CBO uses the smoothing method suggested by Carter and Kohn (1994) to generate the state variables (x_t for $t = T + 1, \cdots, T + H$) conditioned on $C = Y^*_{T+H} = (y^*_{T+1}, \cdots, y^*_{T+H})$ from the following equation.

$$p(x_{T+1}, \cdots, x_{T+H} | Y^*_{T+H}, \theta) = p(x_{T+H} | Y^*_{T+H}, \theta) \prod_{t=T+1}^{T+H-1} p(x_t | Y^*_t, x_{t+1}, \theta) \tag{14}$$

CBO first generates $x^{(g)}_{T+H}$ from $p(x_{T+H} | Y^*_{T+H}, \theta^{(g)})$ and then generates $x^{(g)}_t$ from $p(x_t \, | Y^*_t, x^{(g)}_{t+1}, \theta^{(g)})$ backward for $t = T + H - 1, \cdots, T + 1$ for each draw of $\theta^{(g)}$. All the conditional distributions are the normal distributions because the state space model is linear (Carter and Kohn 1994). The set of the first n elements of $x^{(g)}_{T+h}$ for $h = 1, \cdots, H$ and $g = 1, \cdots, G$ consists of the conditional predictive density function as in Equation (8). Unconditional forecasting is done through the same process as conditional forecasting, but with no conditions.

Assessing the BVAR's Performance

The cited literature has shown that BVARs have comparable forecasting power to that of other macroeconomic models. This section describes generating out-of-sample forecasts through the BVAR and compares their accuracy with that of CBO's baseline forecasts. For the one-year forecasting horizon, the BVAR and CBO's baseline forecasts have comparable accuracy. For forecasting performance of different BVAR specifications, see Carriero, Clark, and Marcellino (2015).

Setting for Out-of-Sample Forecasting

The BVAR is estimated recursively on a quarterly sample from the first quarter of 1959 to the fourth quarter of the year before the forecasting year. For example, to forecast 2019 the BVAR uses a sample from the first quarter of 1959 to the fourth quarter of 2018. The forecasting horizon is 1 year, and each four-step-ahead forecast is implemented annually. For example, after the BVAR forecasts the first through fourth quarter of 2018 with a sample up to the fourth quarter of 2017, the model projects the first through fourth quarter of 2019 with a sample up to the fourth quarter of 2018, and so on.

The BVAR's out-of-sample forecasting procedure is designed to be compatible with CBO's forecasting evaluation database, in which forecasts are finalized mostly in January of the forecasting year (CBO 2019). From the timing, it is assumed that CBO uses the information of the previous year to forecast the current year. For example, CBO forecasts 2019 by using data up to the fourth quarter of 2018. That assumption on the information set is consistent with the BVAR out-of-sample forecasting procedure.

The forecasting year starts from 2000 to 2019, and CBO evaluates the BVAR's forecasting performance of four variables: real GDP growth rate, CPI inflation rate, unemployment rate, and 10-year Treasury note rate. The real GDP growth rate and CPI inflation rate in a forecasting year are the annual growth rates of averages of projected quarterly real GDP or CPI, respectively. The unemployment rate and 10-year Treasury note rate in a forecasting year are averages of projected quarterly values. That specification is the same as that in the CBO forecasting evaluation database.

For that exercise, CBO uses fewer variables in the BVAR than in the previous section to ensure a stable forecasting evaluation procedure by excluding variables with redundant or little additional information. The BVAR contains 12 major economic and financial variables. The federal funds rate and 3-month Treasury bill rate are not included because they reached the zero lower bound for an extended period. The number of lags of the BVAR is fixed at five, and the same variables are included in the BVAR for each forecasting period. To incorporate the effect of inflation targeting, the price index for personal consumption expenditures is fixed as 6 percent higher in 3 years, which is equivalent to 2 percent annual inflation for 3 years on average.

Results

CBO compares the root mean squared errors (RMSEs) of the forecasts by the BVAR and CBO forecasts to measure forecasting performance. For a benchmark, the agency also calculates RMSEs of a random walk model, of which forecasts are the last observable values. RMSEs of the BVAR are usually smaller than those of the random walk model and greater than those of CBO forecasts. For the CPI inflation rate and 10-year Treasury note rate, the BVAR and CBO turn out to have similar forecasting power, whereas the BVAR shows lower forecasting power for the real GDP growth rate and the unemployment rate (see Table 2).

The agency examines the forecast values of the BVAR and CBO during the forecasting period, along with actual values of the four variables (see Figure 3). The BVAR shows lower forecasting accuracy than CBO on the real GDP growth rate and unemployment rate during the financial crisis, whereas it shows similar accuracy for more recent periods. Because the BVAR projects only on the basis of the past values of the variables in it, it may not immediately track a sudden change in economic trends. The BVAR's CPI inflation forecasts are more volatile than those of CBO, whereas RMSEs are similar between them. The BVAR's forecasts for the 10-year Treasury note rate have been lower than those of CBO for most of the forecasting period.

To assess the BVAR's performance for conditional forecasting, CBO implements the same procedure as above to project variables other than real GDP growth rates, taking real GDP growth rates forecast by CBO as conditions (see the "BVAR Conditioned" column in Table 2). For example, the BVAR forecasts the other variables for 2019 by using CBO's projection on real GDP growth rate in 2019 as a condition. Incorporating the information increases the BVAR's forecasting performance considerably for the unemployment rate, especially for the financial crisis (see Figure 4).

Figures

Figure 1. [Return to Text]

How a Scenario Is Constructed and Used to Calculate Budget Outcomes

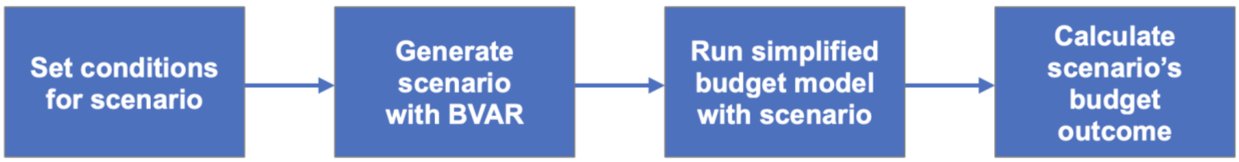

Data source: Congressional Budget Office.

BVAR = Bayesian vector autoregression.

Figure 2. [Return to Text]

Example of Using the BVAR to Construct Scenarios

A. Conditioned Variables for Scenarios

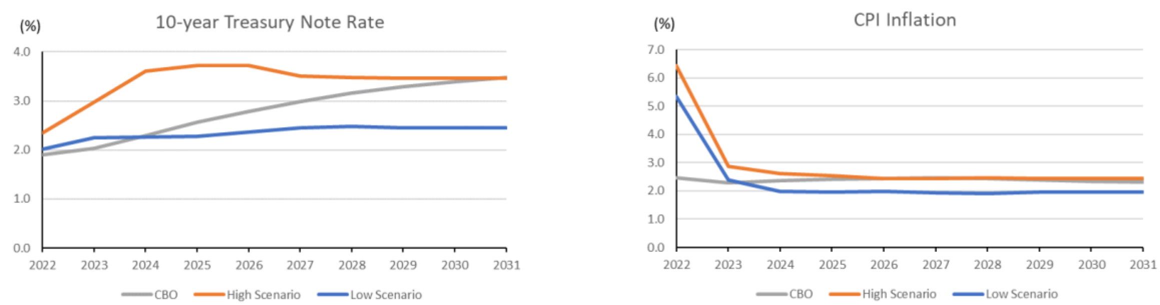

B. Projected Variables by the BVAR

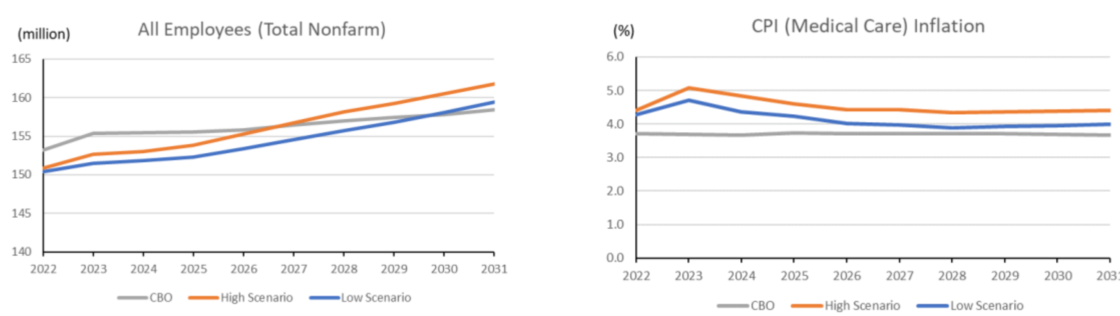

Data source: Congressional Budget Office.

BVAR = Bayesian vector autoregression; CPI = consumer price index.

Graphs in panel A show conditioned values of two variables among 8 variables in the condition set for each scenario: high scenario and low scenario. The high scenario is a scenario based on the projection by top forecasters who projected the highest interest rates among 38 forecasters in the *Blue Chip Economic Indicators* published in March 2022, whereas the low scenario is one by bottom forecasters who projected the lowest interest rates. CBO's July 2021 forecast also is shown.

Graphs in panel B show projected valued of two variables among 18 variables projected by the BVAR for each scenario.

Figure 3. [Return to Text]

Out-of-Sample Forecasts of the BVAR

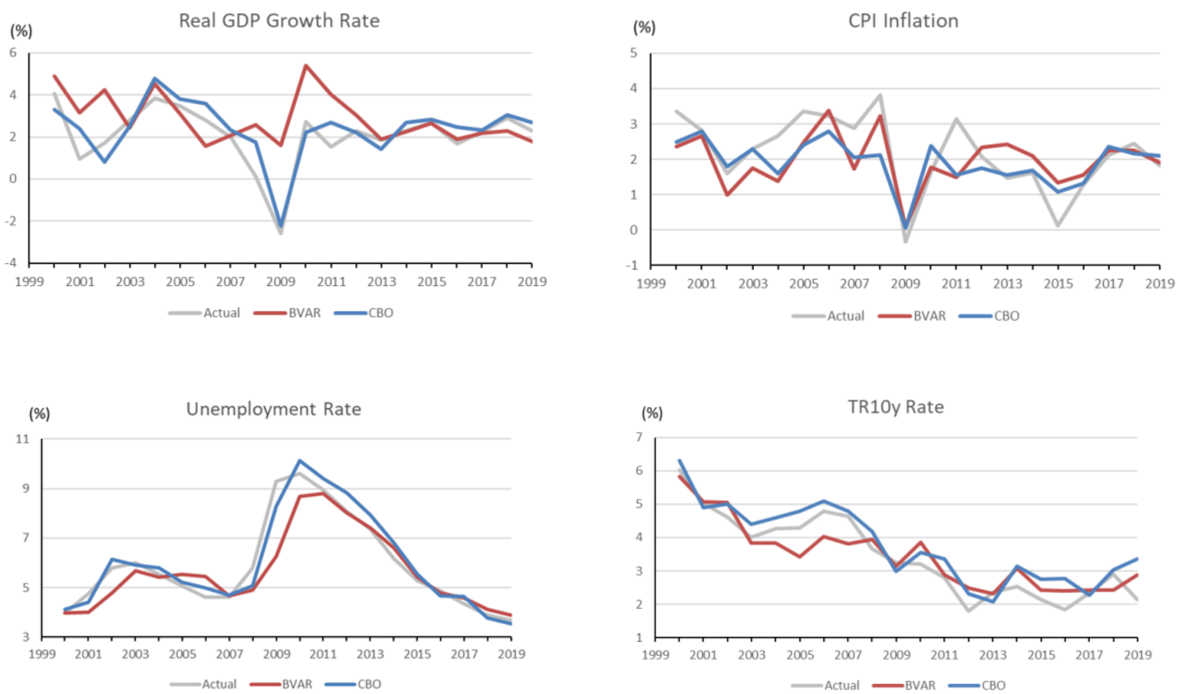

Data source: Congressional Budget Office.

BVAR = Bayesian vector autoregression; CPI = consumer price index; GDP = gross domestic product; TR10y = 10-year Treasury note.

Graphs show the one-year-ahead predicted values by the BVAR and CBO's forecasts for each year along with actual values of the four variables.

Figure 4. [Return to Text]

Out-of-Sample Forecasts of the BVAR Conditioned on CBO's Real GDP Forecasts

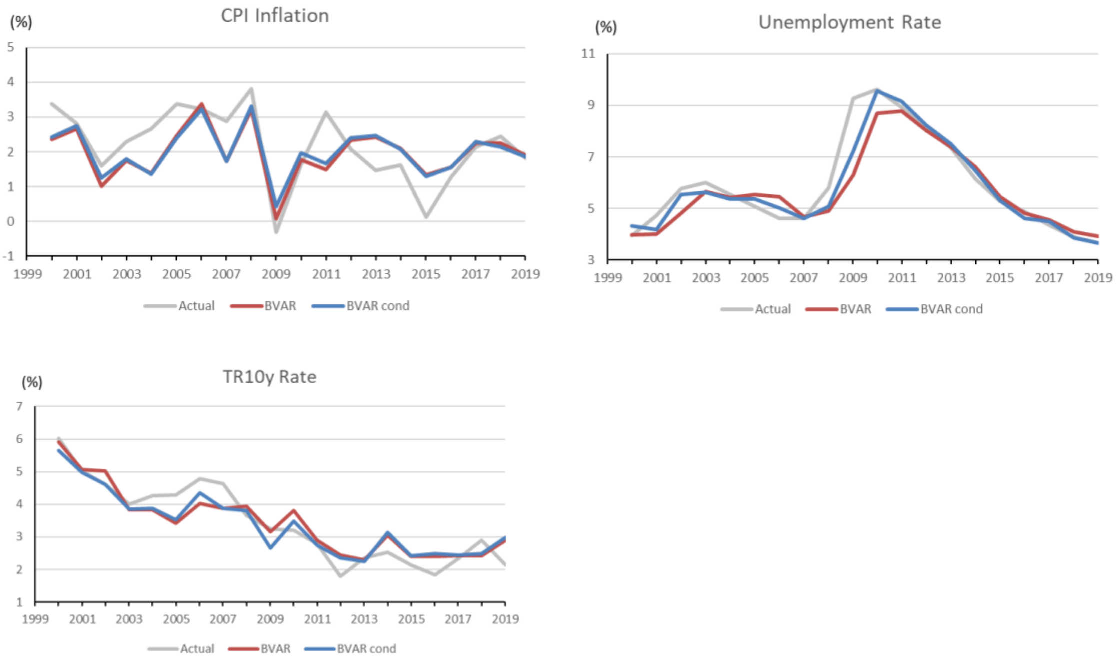

Data source: Congressional Budget Office.

BVAR = Bayesian vector autoregression; CPI = consumer price index; GDP = gross domestic product; TR10y = 10-year Treasury note.

Graphs show the one-year-ahead predictions of the BVAR and the BVAR conditioned on CBO's real GDP prediction for each year along with actual values of the three variables.

Tables

Table 1. [Return to Text]

Variables in the BVAR

Variable	Units	Seasonal Adjustment	Transformation	Input for Budget Model[a]
Real GDP	Billions of chained (2012) dollars	Yes	$\log_e \times 100$	Yes
Real PCE	Billions of chained (2012) dollars	Yes	$\log_e \times 100$	No
Real private nonresidential fixed investment	Billions of chained (2012) dollars	Yes	$\log_e \times 100$	No
Real exports of goods and services	Billions of chained (2012) dollars	Yes	$\log_e \times 100$	No
Real imports of goods and services	Billions of chained (2012) dollars	Yes	$\log_e \times 100$	No
GDP: chain price index	Fixed as 100 in 2012	Yes	$\log_e \times 100$	Yes
PCE: chain price index	Fixed as 100 in 2012	Yes	$\log_e \times 100$	No
PCE less food and energy: chain price index	Fixed as 100 in 2012	Yes	$\log_e \times 100$	No
CPI-U: all items	Fixed as 100 in 1982–1984	Yes	$\log_e \times 100$	Yes
CPI-U: food at home	Fixed as 100 in 1982–1984	Yes	$\log_e \times 100$	Yes
CPI-U: medical care	Fixed as 100 in 1982–1984	Yes	$\log_e \times 100$	Yes
All employees, total nonfarm	Millions	Yes	$\log_e \times 100$	Yes
Civilian labor force: 16 years and over	Millions	Yes	$\log_e \times 100$	Yes
Wage and salaries	Billions of dollars	Yes	$\log_e \times 100$	Yes
Nonfarm business sector: hours of all persons	Fixed as 100 in 2012	Yes	$\log_e \times 100$	Yes
Civilian unemployment rate: 16 years and over	Percent	Yes	$\times 100$[b]	Yes
Federal funds effective rate	Percent (annualized)	No	$\times 100$	Yes
3-month Treasury bill rate: secondary market	Percent (annualized)	No	$\times 100$	Yes
5-year Treasury note yield at constant maturity	Percent (annualized)	No	$\times 100$	No
10-year Treasury note yield at constant maturity	Percent (annualized)	No	$\times 100$	Yes
Moody's seasoned Aaa corporate bond yield	Percent (annualized)	No	$\times 100$	No
Moody's seasoned Baa corporate bond yield	Percent (annualized)	No	$\times 100$	No
GNP	Billions of dollars	Yes	$\log_e \times 100$	Yes
Private nonresidential investment: equipment	Billions of dollars	Yes	$\log_e \times 100$	Yes
Real potential GDP	Billions of chained (2012) dollars	n.a.	$\log_e \times 1{,}000$[c]	Yes
Total factor productivity	Fixed as 100 in 2007	n.a.	$\log_e \times 100$	Yes

Data source: Congressional Budget Office, using data from Bureau of Economic Analysis, Bureau of Labor Statistics, Federal Reserve Board, and Moody's.

BVAR = Bayesian vector autoregression; CPI-U = consumer price index for all urban consumers; GDP = gross domestic product; GNP = gross national product; $\log_e$, natural (base *e*) logarithm; PCE = personal consumption expenditure; n.a. = not applicable.

Real potential GDP and total factor productivity are constructed by CBO.

a. "Budget model" refers to CBO's simplified budget model to approximate budget outcomes.

b. Applied to variables whose units are percentages to make the variations similar to those of other variables.

c. Used for real potential GDP to make the variation similar to those of other variables.

Table 2. [Return to Text p. 13 | p. 14]

Results of Out-of-Sample Forecasting

Variable	Random Walk	BVAR	CBO	BVAR Conditioned
Real GDP growth rate	1.74	1.62	0.73	n.a.
CPI inflation rate	1.33	0.76	0.74	0.75
Unemployment rate	1.04	0.83	0.45	0.55
10-year Treasury note rate	0.56	0.49	0.50	0.46

Data source: Congressional Budget Office.

BVAR = Bayesian vector autoregression; CPI = consumer price index; GDP = gross domestic product; RMSE = root mean squared error; n.a. = not applicable.

Values are RMSEs of predictions by the random walk model, the BVAR, CBO, and the BVAR conditioned on CBO's real GDP forecasts. Lower RMSEs mean lower prediction error.

References Cited

Angelini, Elena, Magdalena Lalik, Michele Lenza, and Joan Paredes. 2019. "Mind the Gap: A Multi-Country BVAR Benchmark for the Eurosystem Projections." *International Journal of Forecasting* 35, no. 4 (October–December): 1658–68. https://doi.org/10.1016/j.ijforecast.2018.12.004. [Return to p. 5]

Antolín-Díaz, Juan, Ivan Petrella, and Juan F. Rubio-Ramírez. 2021. "Structural Scenario Analysis With SVARs." *Journal of Monetary Economics* 117 (January): 798–815. https://doi.org/10.1016/j.jmoneco.2020.06.001. [Return to p. 6]

Bańbura, Marta, Domenico Giannone, and Michele Lenza. 2015. "Conditional Forecasts and Scenario Analysis With Vector Autoregressions for Large Cross-Sections." *International Journal of Forecasting* 31, no. 3 (July–September): 739–56. https://doi.org/10.1016/j.ijforecast.2014.08.013. [Return to p. 5 (i) | p. 5 (ii)]

Bańbura, Marta, Domenico Giannone, and Lucrezia Reichlin. 2010. "Large Bayesian Vector Auto Regressions." *Journal of Applied Econometrics* 25, no. 1 (January–February): 71–92. https://doi.org/10.1002/jae.1137. [Return to p. 5]

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino. 2015. "Bayesian VARs: Specification Choices and Forecast Accuracy." *Journal of Applied Econometrics* 30, no. 1 (January–February): 46–73, https://doi.org/10.1002/jae.2315. [Return to p. 12]

Carriero, Andrea, Todd E. Clark, Massimiliano Marcellino, and Elmar Mertens. 2022. "Addressing COVID-19 Outliers in BVARs With Stochastic Volatility." *Review of Economics and Statistics* (June). https://doi.org/10.1162/rest_a_01213. [Return to p. 8]

Carter, C. K., and R. Kohn. 1994. "On Gibbs Sampling for State Space Models." *Biometrika* 81, no. 3 (September), pp. 541–53. https://doi.org/10.1093/biomet/81.3.541. [Return to p. 12 (i) | p. 12 (ii)]

CBO (Congressional Budget Office). 2014. *How CBO Analyzes the Effects of Changes in Federal Fiscal Policies on the Economy* (November). www.cbo.gov/publication/49494. [Return to p. 4]

———. 2019. *CBO's Economic Forecasting Record: 2019 Update* (October). www.cbo.gov/publication/55505. [Return to p. 13]

———. 2023a. *Assessing the Budgetary Implications of Economic Uncertainty With CBO's Incomes Model and Budgetary Feedback Model* (January). www.cbo.gov/publication/58885. [Return to p. 4]

———. 2023b. *Estimating the Uncertainty of the Economic Forecast Using CBO's Bayesian Vector Autoregression Model* (January). www.cbo.gov/publication/58883. [Return to p. 8]

———. 2023c. *The Budget and Economic Outlook: 2023 to 2033* (February). www.cbo.gov/publication/58848. [Return to p. 4]

Crump, Richard K., Stefano Eusepi, Domenico Giannone, Eric Qian, and Argia M. Sbordone. 2021. *A Large Bayesian VAR of the United States Economy*, staff report no. 976 (August). Federal Reserve Bank of New York. www.newyorkfed.org/research/staff_reports/sr976.html. [Return to p. 4 | p. 5 | p. 7 | p. 8]

D'Agostino, Antonello, Luca Gambetti, and Domenico Giannone. 2013. "Macroeconomic Forecasting and Structural Change." *Journal of Applied Econometrics* 28, no. 1 (January–February): 82–101. https://doi.org/10.1002/jae.1257. [Return to p. 6]

Domit, Sílvia, Francesca Monti, and Andrej Sokol. 2019. "Forecasting the UK Economy With a Medium-Scale Bayesian VAR." *International Journal of Forecasting* 35, no. 4 (October–December): 1669–78. https://doi.org/10.1016/j.ijforecast.2018.11.004. [Return to p. 5]

Frentz, Nathaniel, Jaeger Nelson, Dan Ready, and John Seliski. 2020. *A Simplified Model of How Macroeconomic Changes Affect the Federal Budget*, Working Paper 2020-01 (Congressional Budget Office, January). www.cbo.gov/publication/55884. [Return to p. 4]

Giannone, Domenico, Michele Lenza, and Giorgio E. Primiceri. 2015. "Prior Selection for Vector Autoregressions." *Review of Economics and Statistics* 97, no. 2 (May): 436–51. https://doi.org/10.1162/REST_a_00483. [Return to p. 5 (i) | p. 5 (ii) | p. 10 | p. 24 | p. 25]

———. 2019. "Priors for the Long Run." *Journal of the American Statistical Association* 114, no. 526: 565–80. https://doi.org/10.1080/01621459.2018.1483826. [Return to p. 7]

Koop, Gary M. 2013. "Forecasting With Medium and Large Bayesian VARs." *Journal of Applied Econometrics* 28, no. 2 (March): 177–203. https://doi.org/10.1002/jae.1270. [Return to p. 5]

Koop, Gary, and Dimitris Korobilis. 2013. "Large Time-Varying Parameter VARs." *Journal of Econometrics* 177, no. 2 (December): 185–98. https://doi.org/10.1016/j.jeconom.2013.04.007. [Return to p. 7]

Karlsson, Sune. 2013. "Forecasting With Bayesian Vector Autoregression." In *Handbook of Economic Forecasting*, edited by Graham Elliott, Clive Granger, and Allan Timmermann (Elsevier), vol. 2, part B: 791–897. https://doi.org/10.1016/B978-0-444-62731-5.00015-4. [Return to p. 5]

Lenza, Michele, and Giorgio E. Primiceri. 2022. “How to Estimate a Vector Autoregression After March 2020.” *Journal of Applied Econometrics* 37, no. 4 (June–July): 688–99. https://doi.org/10.1002/jae.2895. [Return to p. 8]

Litterman, Robert B. 1980. “A Bayesian Procedure for Forecasting With Vector Autoregression” (working paper, Department of Economics, Massachusetts Institute of Technology, Cambridge, MA). [Return to p. 5]

Schorfheide, Frank, and Dongho Song. 2022. “Real-Time Forecasting With a (Standard) Mixed-Frequency BVAR During a Pandemic.” Centre for Economic Policy Research Discussion Paper DP16760. https://ssrn.com/abstract=4026603. [Return to p. 8]

Sims, Christopher A. 1972. “Money, Income, and Causality.” *American Economic Review* 62, no. 4 (September): 540–52. www.jstor.org/stable/1806097. [Return to p. 5]

———. 1980a. “Comparison of Interwar and Postwar Business Cycles: Monetarism Reconsidered.” *American Economic Review* 70, no. 2 (May): 250–7. www.jstor.org/stable/1815476. [Return to p. 5]

———. 1980b. “Macroeconomics and Reality.” *Econometrica* 48, no. 1 (January): 1–48. https://doi.org/10.2307/1912017. [Return to p. 5]

Sims, Christopher A., and Tao Zha. 2006. “Were There Regime Switches in U.S. Monetary Policy?” *American Economic Review* 96, no. 1 (March): 54–81. http://doi.org/10.1257/000282806776157678. [Return to p. 6]

Stock, James. H., and Mark W. Watson. 2001. “Vector Autoregressions.” *Journal of Economic Perspectives* 15, no. 4 (Fall): 101–15. http://doi.org/10.1257/jep.15.4.101. [Return to p. 6]

Swagel, Phillip L. 2022. “Budgetary Implications of Economic Scenarios With Higher and Lower Interest Rates.” *Brookings Papers on Economic Activity* (Spring): 233–49. https://doi.org/10.1353/eca.2022.0013. [Return to p. 7]

Appendix

[Return to Text p. 10 (i) | p. 10 (ii) | p. 11]

Deriving Predictive Density Functions

Equation (7) for the unconditional forecasting in the main text is derived as follows:

$$p(y_{T+1}, \cdots, y_{T+H}|y) = \iint p(y_{T+1}, \cdots, y_{T+H}, \lambda, \theta|y) d\lambda d\theta$$
$$= \int p(y_{T+1}, \cdots, y_{T+H}, \theta|y) d\theta$$
$$= \int p(y_{T+1}, \cdots, y_{T+H}|y, \theta) p(\theta|y) d\theta,$$

where $p(\theta|y)$ is the marginal distribution of the posterior distribution, $p(\lambda, \theta|y)$ in Equation (6). The equation shows that one can approximate the predictive density function by using the draws of θ from the posterior distribution.

Similarly, Equation (8) for the conditional forecasting in the main text is derived as follows:

$$p(y_{T+1}, \cdots, y_{T+H}|y, C) = \iint p(y_{T+1}, \cdots, y_{T+H}, \lambda, \theta|y, C) d\lambda d\theta$$
$$= \int p(y_{T+1}, \cdots, y_{T+H}, \theta|y, C) d\theta$$
$$= \int p(y_{T+1}, \cdots, y_{T+H}|y, \theta) p(\theta|y, C) d\theta,$$

where *C* is a set of conditions for a scenario.

How to Estimate BVARs

The Congressional Budget Office uses the Markov chain Monte Carlo (MCMC) algorithm suggested by Giannone, Lenza, and Primiceri (2015) to generate draws of the hyperparameter (λ) and the parameters (θ) conditional on the data (*y*) from the posterior distribution. CBO first sets the initial values ($\lambda^{(0)}$ and $\theta^{(0)}$) as their modes in the posterior distribution and implements the following procedure:

1. Generate $\lambda^{(g)}$ from $p(\lambda|y)$.
2. Generate $\theta^{(g)}$ from $p\left(\theta|\lambda^{(g)}, y\right)$.
3. Repeat the above steps for $g = 1, \cdots, G$.

The posterior distribution of λ given the data, or $p(\lambda|y)$, is derived via Bayes's rule as:

$$p(\lambda|y) \propto p(\lambda)p(y|\lambda),$$

where $p(y|\lambda)$ is called the marginal likelihood, which is computed as follows:

$$p(y|\lambda) = \int p(y|\lambda)p(\theta|\lambda)d.$$

CBO uses the closed-form solution of the marginal likelihood derived by Giannone, Lenza, and Primiceri (2015). To generate λ, CBO uses a Metropolis–Hastings algorithm as follows:

1. Draw a candidate value, λ^*, from a normal distribution with the mean equal to the previous draw, $\lambda^{(g-1)}$, and the variance equal to $k \cdot W$, where W is the inverse Hessian of the negative of the log-posterior of the hyperparameter at the mode of the posterior distribution. k is a scaling constant to adjust the acceptance rate of the algorithm.
2. Set

$$\lambda^{(g)} = \begin{matrix} \lambda^* \text{ with probability of } \alpha^{(g)} \\ \lambda^{(g-1)} \text{ with probability of } 1 - \alpha^{(g)}, \end{matrix}$$

where

$$\alpha^{(g)} = min\left(1, \frac{p(\lambda^*|y)}{p\left(\lambda^{(g-1)}\middle|y\right)}\right).$$

Then, CBO uses Gibbs sampling to generate θ from $p\left(\theta\middle|\lambda^{(g)}, y\right)$, the density function of the Normal-inverse-Wishart distribution.

Figure A-1 shows the MCMC draws and the density functions of λ from $p(\lambda|y)$ and the corresponding marginal likelihood $[p(y|\lambda)]$. The figure indicates that the MCMC draws converge quickly.

Figure A-2 shows the density functions of some parameters in the equation of real GDP. The top-left panel shows the density function of the coefficient of real GDP with a lag of one period. The other panels show those of the coefficients of lagged other variables such as real consumption, real nonresidential fixed investment, and real exports. As the Minnesota prior implies, the density function of the coefficient of its own lagged variable is around 1, whereas those of the other variables are around 0.

Figure A-1. [Return to Text]

Draws of the Hyperparameter and Marginal Likelihood

A. Draws (left) and density function (right) of the hyperparameter (λ).

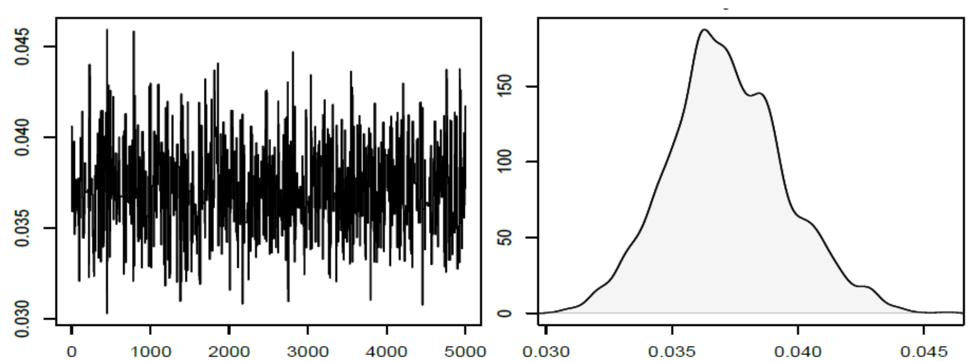

B. Draws (left) and density function (right) of the marginal likelihood.

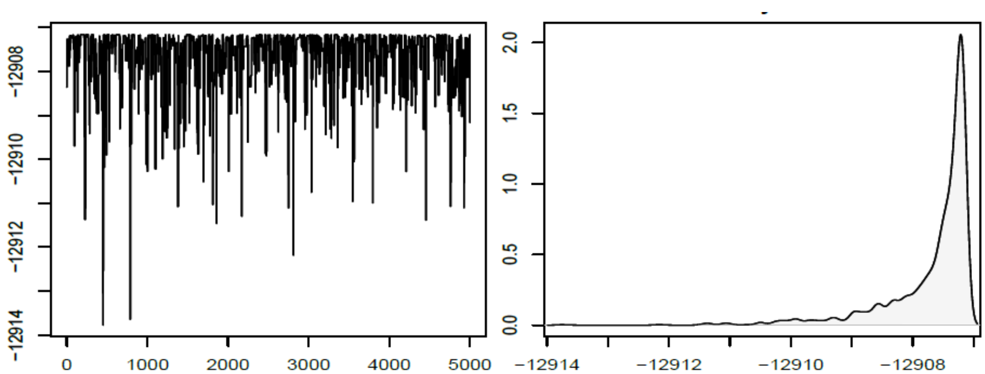

Data source: Congressional Budget Office.

Figure A-2. [Return to Text]

Density Functions of Selected Parameters in the BVAR

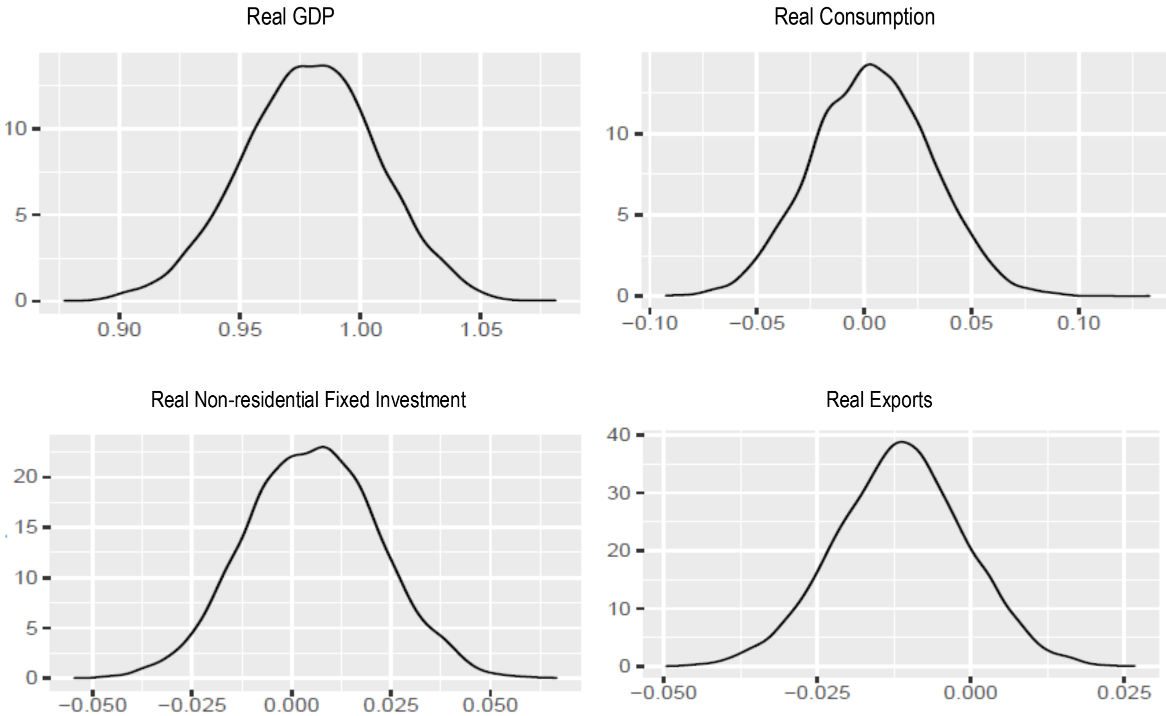

Data source: Congressional Budget Office.

BVAR = Bayesian vector autoregression; GDP = gross domestic product.

Graphs show density functions of some parameters in the equation of real GDP. The top-left panel shows the density function of the coefficient of real GDP with a lag of one period. The other panels show those of the coefficients of lagged real consumption, real nonresidential fixed investment, and real exports.