Take Your Model Bowling: Forecasting with General Equilibrium Models

MARCO DEL NEGRO AND FRANK SCHORFHEIDE

Del Negro is a research economist and assistant policy adviser in the Atlanta Fed's research department. Schorfheide is an assistant professor of economics at the University of Pennsylvania. *They thank Eric Leeper, Ellis Tallman, and Rochelle Edge for helpful suggestions.*

In the past twenty years dynamic stochastic general equilibrium (DSGE) models have taken center stage in academic macroeconomic research. The stated goal of DSGE models in the tradition of Kydland and Prescott (1982) is to explain business cycle features of the data and to be usable for quantitative—as opposed to only qualitative—policy analysis. Yet until recently the data that these models have been measured against are not the GDP or inflation figures that appear in newspapers but so-called filtered data. One commonly used filter, the Hodrick-Prescott (1997) filter, decomposes the data into a cyclical component and a growth component and removes the latter. By doing so the filter removes from the data variations that are due to frequencies other than business cycle frequencies. One rationale behind the filtering is that the model is designed to explain business cycles as opposed to very short-run (say, seasonal) or long-run movements (say, due to demographics) in the data. Hence, it seems logical to assess the model's fit in terms of that part of the data that it can explain.

Whatever the motivation behind using filtered data, filtering has two important consequences. First, it implies that the task of forecasting macroeconomic time series stays outside the realm of DSGE models and is left entirely to econometric models or judgmental forecasters. Practitioners are interested in forecasts of actual, as opposed to filtered, data, so they rely on models, or individuals, that deliver such forecasts. The second consequence of using filtered data is that, to this day, policymakers rarely use general equilibrium models, at least in quantitative analysis. Like practitioners, policymakers base their decisions on forecasts of macroeconomic time series. Policymakers want to know, for instance, the expected path of inflation, unemployment, or real output growth in the next few quarters and by how much a 25 basis point cut in the federal funds rate would change such a path. Since very little is known about the forecasting performance of general equilibrium models, policymakers rarely rely on them for quantitative policy assessment.

Many of the models currently used in forecasting and policy analysis belong to one of two categories. The first includes models in the Cowles Commission tradition.[1] These are large-scale simultaneous equation models that were prominent in macroeconomics before the rational expectations revolution, from the late 1950s to the early 1970s (see Diebold 1998 for a brief history of macroeconomic forecasting). These models have been updated to incorporate rational expectations and are still heavily used for forecasting and policy-making by central banks around the world as well as by commercial forecasters.[2] FRB/US—the workhorse model of policy analysis at the Federal Reserve Board of Governors—is one of them.[3] The second category of models includes vector autoregressions (VARs), which were introduced by Sims (1980) in the early 1980s and popularized in the forecasting literature by Litterman (1986). The BVAR model (a VAR with Bayesian priors) used for forecasting at the Federal Reserve Bank of Atlanta belongs to this second category.[4]

Like VARs and Cowles Commission models, DSGE models also aspire to describe the data. Perhaps the main difference between DSGE models on the one side and VARs and Cowles Commission models on the other side is that DSGE models are explicitly derived from first principles. That is, DSGE models describe the general equilibrium allocations and prices of a model economy in which agents (households, firms, financial intermediaries, etc.) dynamically maximize their objectives (utility, profits, and so on) subject to their budget and resource constraints. The DSGE model parameters describe the preferences of agents (tastes), the production function (technology), and other features of the economy. These parameters are called "deep" parameters—parameters that do not vary with policy.

The stated goal of DSGE models in the tradition of Kydland and Prescott (1982) is to explain business cycle features of the data and to be usable for quantitative policy analysis.

To be sure, economic theory also informs Cowles Commission–style models and VARs. For instance, most equations in Cowles Commission–style models, such as consumption equations, investment equations, and so on, are inspired by economic analysis, if not explicitly derived from it. However, in some cases the parameters of these models characterize behavior instead of tastes and technologies. Yet the agents' behavior is not policy invariant, and therefore not all parameters in such models are deep.[5] The modelers typically adopt a block-by-block approach (in which the blocks are the household sector, the business sector, etc.; see Brayton, Levin, et al. 1997; Brayton, Mauskopf, et al. 1997) to describe the various agents and sectors in the economy and often ignore important links among blocks. In particular, when forming expectations, agents in these models often ignore equilibrium restrictions that must hold in all future states of the world. VARs were introduced by Sims (1980) with the intent to overcome the deficiencies of the Cowles Commission approach and to obtain specifications that are consistent with a dynamic general equilibrium. In fact, a linearized DSGE model can be closely approximated by a VAR with a sufficiently large number of lags. The VAR parameters can, in principle, be constrained to be functions of deep parameters for some DSGE model. Typically, however, VARs are not estimated under such constraints, and therefore the VAR parameter estimates cannot be interpreted in terms of deep parameters.

This article reviews some recent attempts to use general equilibrium models for forecasting and policy analysis. In particular, the article focuses on one specific approach, pioneered by Ingram and Whiteman (1994) and further developed by Del Negro and Schorfheide (forthcoming), that relies on the use of general equilibrium models as priors for Bayesian VARs. To motivate this approach, which we will call DSGE-VAR, we first need to address two questions. First, why should one bother to forecast with general equilibrium models? Second, why should one use general equilibrium models as priors instead of forecasting directly with them? The next two sections address these questions.

Why Forecast with DSGE Models?

There are two good reasons to use DSGE models in forecasting (also see Diebold 1998 for a discussion of forecasting with DSGE models). The first reason has to do with improving the forecasting precision. It is well known that loosely parameterized models, such as VARs, are imprecisely estimated unless a very long time series of data is available, which is rarely the case in macroeconomics. Imprecise estimates in turn result in potentially large forecast errors, especially for long forecast horizons. A solution to this problem of too many parameters is to use Bayesian priors. In Bayesian econometrics a *prior* on a set of parameters is a distribution that summarizes beliefs or knowledge about these parameters prior (whence the name) to observing the data. Priors reduce the sample variability in the parameter estimates by "shrinking" them toward a specific point in the parameter space. For this reason, since the seminal work of Litterman (1986) and Doan, Litterman, and Sims (1984), BVARs have earned a reputation for forecasting accuracy (see Robertson and Tallman 1999 for a review of the comparative forecasting accuracy of BVARs). In many BVARs the priors arise from statistics, namely, from the observation that random walk processes describe quite well the behavior of a number of macroeconomic time series.[6] This observation is the rationale, for instance, for the well-known Minnesota prior. The Minnesota prior shrinks the VAR parameters toward a unit root. Ingram and Whiteman (1994) proposed to use a prior that comes from a general equilibrium model, namely, a standard real business cycle (RBC) model. Ingram and Whiteman show that the performance of their VAR with an RBC prior in terms of forecasting real variables (real output, consumption, and investment growth) is comparable to that of a VAR with a Minnesota prior.

The second reason for forecasting with DSGE models has to do with evaluating the impact of changes in policy. The well-known Lucas (1976) critique implies that only models in which the parameters are deep—that is, models in which the parameters do not vary with policy—are suited to evaluate the impact of policy changes. To understand why this is the case, let us consider a model in which the parameters are not deep; this may be a VAR or a Cowles Commission model. The forecaster who uses such a model to predict the effect of a given policy change faces the following dilemma. On the one hand, she can estimate the parameters of the model only on the basis of past data and experiences. On the other hand, unless the policy change has occurred before, she can gain little guidance from past experience about how the policy change affects the decision rules of agents and hence how it affects the parameters of the model.

For example, suppose that the goal is to predict the effects of the 2003 change in the tax code. A forecaster might use the data available prior to the policy change to estimate a consumption equation that describes the behavior of consumers as a function of a number of variables, including wealth and disposable income. Knowing the amount by which wealth and disposable income will be increased by the tax breaks, the forecaster may use the estimated relationship to forecast consumption in the next few quarters. However, the Lucas argument is that the change in policy may induce agents to change their behavior, which in turn may change the relationship between wealth, disposable income, and consumption. For instance, the tax break may affect the agents' propensity to consume. Hence, the forecast for consumption may well turn out to be wrong.

Now suppose that the model being used to forecast the impact of the policy change is a DSGE model. In a DSGE model the parameters are truly deep, that is, invariant with policy, or at least they are assumed to be so. For instance, there is no reason to think that a change in the tax policy would affect either the extent to which people enjoy leisure (tastes) or the current speed of computers (technology).[7] Therefore the forecaster can estimate these parameters using existing data and does not have to worry that they may change with policy. Once the parameters are available, the forecaster can solve the model and work out the impact of the tax change on consumption. For instance, the forecaster using DSGE models can correctly compute agents' propensity to consume under the new policy. If the specification of the DSGE model is appropriate, the effect of the new policy can be correctly evaluated even though it has not occurred in the past.

Of course, the distinction just drawn between models with and without deep parameters is Manichaean. First, not all parameters in DSGE models are necessarily deep, so some DSGE models may be subject to the same criticism as the other models. Second, not all policy changes result in dramatic changes in agents' behavior. In such cases, models other than DSGE models may well be able to provide reliable forecasts.[8] With these important caveats established, one of the main implications of the Lucas critique is that DSGE models have in principle an important advantage over other models in forecasting the effects of policy changes.

Why Use Priors?

The advantages of DSGE models discussed in the previous section often come at a cost in terms

1. The Cowles Commission (now the Cowles Foundation) was founded by Alfred Cowles in 1932 to promote quantitative research in economics. The Cowles Commission and its fellows played a pivotal role in promoting and developing large-scale econometric models. Hence, its name became associated with the approach (see Fair 1992).
2. Sims (2002) provides a criticism of the way Cowles Commission–style models are currently being estimated and used for policy analysis.
3. See Brayton, Mauskopf, et al. (1997) and Brayton, Levin, et al. (1997) for a description of FRB/US and Reifschneider, Stockton, and Wilcox (1997) for a description of forecasting and policy evaluation at the Federal Reserve Board of Governors.
4. See Zha (1998) for a discussion of the use of VARs in policy analysis.
5. Production functions and utility functions underpin the equations for many sectors of models like FRB/US, however. As a result, one may be able to back out some of the deep parameters from the coefficient estimates.
6. A random walk is a process in which today's best guess about tomorrow's value of a variable is today's value, possibly augmented by a constant.
7. This assumption is not true for all tax policy changes: Some changes may affect spending on research and development and therefore future technology. However, as a first approximation this effect can be ignored for many policy changes.
8. See Sims (1982), Leeper, Sims, and Zha (1996), and Sargent (1984) for a discussion of the relevance of the Lucas critique for VARs.

of the model's fit. A number of papers that study the fit of DSGE models (for example, Altug 1989; Leeper and Sims 1994; Ireland 1997; Schorfheide 2000) find that this fit is far from perfect. As discussed above, economic theory imposes a number of restrictions on the stochastic process followed by the data—the cross-equation restrictions that are the hallmark of rational expectations econometrics. These restrictions imply that DSGE models are scarcely parameterized compared with VARs or Cowles Commission–style models. Hence, DSGE models may match the data in many important dimensions but, being overly simplified model economies, may also fail in several other dimensions, resulting in large forecast errors for some of the variables of interest.

One of the main implications of the Lucas critique is that DSGE models have in principle an important advantage over other models in forecasting the effects of policy changes.

If one wants to forecast with general equilibrium models, using them indirectly as priors may be preferable to the alternative approach of forecasting directly with them. Using general equilibrium models as priors means that the restrictions stemming from economic theory are imposed loosely instead of rigidly. This method implies that the final (posterior) stochastic process used to forecast will respect the restrictions in those dimensions where these restrictions are not rejected by the data but may otherwise depart from them. Of course, if the restrictions are too loosely imposed there is virtually no difference between forecasting with general equilibrium priors and with an unrestricted VAR (a VAR without priors). Therefore, a key input in the process is the "degree of tightness"—which will be denoted hereafter as λ and which ranges from 0 (no prior) to ∞ (rigid restrictions). The remainder of the article will describe how to choose λ optimally.

How Does DSGE-VAR Work?

The discussion in this section offers an intuitive exposition of the procedure in Del Negro and Schorfheide (forthcoming) for using DSGE models as priors in VARs. Assuming that T observations are available for the variables to be forecast (for instance, real output growth, inflation, and the short-term interest rate), the procedure amounts to generating λT observations for real output growth, inflation, and the federal funds rate from the DSGE model; combining these dummy observations with the actual data; and running a VAR on the augmented data set.[9]

The DSGE-VAR procedure assumes that the dynamics for the data to be forecast are reasonably well described by an unrestricted VAR. To the extent that these dynamics are linear and that the VAR has a sufficient number of lags, this is not a very heroic assumption: VARs are parameterized loosely enough to accommodate nearly any linear stochastic process. As discussed above, the problem is precisely that VARs have too many parameters, so the estimates may be imprecise in short samples. Our approach starts from the premise that a DSGE model may provide useful restrictions for the VAR parameters—useful in the sense that the restrictions can improve the model's forecasting performance. We do not want to impose these restrictions dogmatically for the reasons described in the previous section. Rather, we treat the DSGE model as prior information in the estimation. As is well known since the work of Theil and Goldberger (1961), one way to incorporate prior information into the estimation is to augment the sample with dummy observations that reflect the prior (see also Sims and Zha 1998). This is precisely the route we follow: Our dummy observations are simply data generated by the DSGE model.

The next step in the procedure consists of estimating the VAR parameters using both the actual and the dummy observations. To make this step clear, we specify some notation. Y is the $T \times n$ matrix of actual data, where T is the sample size and n is the number of variables. X is the matrix of VAR regressors, which includes the constant as well as the lags of the variables. The VAR, which we assume to be the data-generating process, is given by

$$(1)\quad Y = X\Phi + U,$$

where U is the $T \times n$ matrix of VAR innovations, which are normally distributed with mean 0 and variance Σ_u. The standard OLS estimator for Φ is given by the well-known formula

$$(2)\quad \Phi_{OLS} = (X'X)^{-1}X'Y.$$

Now λT observations are generated for the variables of interest from the DSGE model. As mentioned above, λ is the weight of the prior. So if $T = 100$ and $\lambda = 0.5$, this step generates fifty observations from the DSGE model. Using these dummy observations, the matrices Y^* and X^* are constructed. Finally, the OLS estimates are run again on the augmented dataset that includes both actual and dummy observations, yielding the estimator

$$(3)\quad \Phi_{DSGE\text{-}VAR} = (X'X + X^{*\prime}X^*)^{-1}(X'Y + X^{*\prime}Y^*).$$

This estimator will be used in forecasting.

Now that the formula is determined, a few comments are in order. First, we want to elaborate on the role of λ, the weight of the prior. Notice that the previous formula can be equivalently expressed as

$$(4)\quad \Phi^{\lambda}_{DSGE\text{-}VAR} = \left(\frac{1}{1+\lambda}\frac{X'X}{T}+(1-\frac{1}{1+\lambda})\frac{X^{*\prime}X^*}{\lambda T}\right)^{-1}\left(\frac{1}{1+\lambda}\frac{X'Y}{T}+(1-\frac{1}{1+\lambda})\frac{X^{*\prime}Y^*}{\lambda T}\right)$$

The terms $(X'X)/T$ and $(X'Y)/T$ are the second moments (that is, say, the covariance between real output growth today and interest rates in the previous period) computed from the data. The terms $(X^{*\prime}X^*)/\lambda T$ and $(X^{*\prime}Y^*)/\lambda T$ are the second moments implied by the DSGE model. Our proposed estimator is computed by weighting the second moments from the data with the second moments implied by the DSGE model, with weights that are respectively $1/(1 + \lambda)$ and $1 - 1/(1 + \lambda)$. If $\lambda = 0$, the dummy observations disappear from the formula: Since for $\lambda = 0$ we are using only the second moments from the data, the estimator in this case coincides with the OLS estimator. If $\lambda = \infty$, the weight on the dummy observations becomes 1. Thus, for $\lambda = \infty$ the restrictions coming from the DSGE model are rigidly imposed.

Next, we introduce an important refinement into the procedure. Whenever λT is not too large (say, $\lambda = 1/10$, and $T = 100$), we generate a small number of dummy observations (in the above example, $\lambda T = 10$). Because of sample variability in the Monte Carlo procedure that generates the dummy observation from the DSGE model, whenever λT is small the *sample* second moments, the terms $(X^{*\prime}X^*)/\lambda T$ and $(X^{*\prime}Y^*)/\lambda T$, may provide a poor estimate of the *population* second moments that the DSGE model implies. One way around the problem is to compute the terms $(X^{*\prime}X^*)/\lambda T$ and $(X^{*\prime}Y^*)/\lambda T$ a large number of times and then average across realizations. This way of proceeding has the disadvantage of being computationally expensive because one would have to draw over and over from the DSGE model. We follow an alternative approach and exploit the fact that whenever the DSGE model is linear (or is well approximated by a linear solution) the population second moments, which we call Γ^*_{xx} and Γ^*_{xy}, can be computed analytically. Hence, in the formula for the estimator $\Phi_{DSGE\text{-}VAR}$, we use the population moments Γ^*_{xx} and Γ^*_{xy} in place of $(X^{*\prime}X^*)/\lambda T$ and $(X^{*\prime}Y^*)/\lambda T$.

Up to this point we have not mentioned the values taken by the deep parameters (preferences, etc.) of the DSGE model. We denote with θ the vector of deep parameters. Clearly, the population moments Γ^*_{xx} and Γ^*_{xy}, and hence our estimator $\Phi_{DSGE\text{-}VAR}$, will depend on the choice of θ. To make this dependence explicit, we rewrite the estimator as

$$(5)\quad \Phi^{\lambda}(\theta)_{DSGE\text{-}VAR} = \left(\frac{1}{1+\lambda}\frac{X'X}{T}+(1-\frac{1}{1+\lambda})\Gamma^*_{xx}(\theta)\right)^{-1}\left(\frac{1}{1+\lambda}\frac{X'Y}{T}+(1-\frac{1}{1+\lambda})\Gamma^*_{xy}(\theta)\right).$$

In the macro literature that follows Kydland and Prescott (1982), a popular approach for choosing θ is calibration (see Kydland and Prescott 1996). Calibration amounts to selecting the values of θ on the basis of information other than that contained in the data we want the model to explain (or forecast). This information may come from microeconomic studies as well as from long-run empirical relationships, such as the labor share of national income or the consumption-output ratio.

We choose to depart from calibration and estimate θ; that is, we let the value of θ be determined by the data we want to fit. We do so on the grounds that if the calibration exercise is poorly performed, or if there is little outside information to pin down some of the elements of θ, the forecasting performance of DSGE-VAR may be severely affected. Still, in order to take advantage of useful outside information (micro studies and so on), we incorporate prior information into the estimation of θ.

How do we learn about θ from the data in our procedure? From equation (1), the data depend on the VAR parameters, not on θ. The answer is that we learn about θ indirectly, via the estimator $\Phi_{DSGE\text{-}VAR}$. As emphasized in equation (5), as long as λ is greater than zero, the choice of θ affects $\Phi_{DSGE\text{-}VAR}$. From the data we learn which $\Phi_{DSGE\text{-}VAR}$ has the best fit. But since for each choice of $\Phi_{DSGE\text{-}VAR}$ there corresponds a choice for θ, we can go back and learn from the data about θ. Note that whenever $\lambda = \infty$,

9. See Del Negro and Schorfheide (forthcoming) for an econometrically detailed description of the approach as well as an appendix on how the procedure works in practice.

that is, whenever the restrictions coming from the DSGE model are imposed rigidly, our estimator for θ coincides with Smith's (1993) SQML (simulated quasi-maximum likelihood) estimator.

How Much Should the DSGE Prior Matter?

The discussion in the previous section emphasized that the choice of λ is crucial in the estimation. Our procedure does not require the forecaster to have strong a priori views on the choice of λ—that is, to pick λ ex ante. Rather, as the forecaster learns about θ from the data, she can also learn about λ. This section shows how λ can be estimated endogenously.

DSGE-VAR addresses regime shifts, trying to strike a balance between the forecasting accuracy of BVARs and the compliance to the Lucas critique of DSGE models.

The intuition about how to choose λ is the same as the one given in the previous section on the estimation of θ. Again, the data do not depend on λ but only on the VAR parameters. However, as formula (5) shows, the estimator $\Phi_{DSGE\text{-}VAR}$ crucially depends on the choice of λ. To make this explicit, let us write $\Phi^{\lambda}_{DSGE\text{-}VAR}$ (for simplicity, in this section we abstract from the choice of θ, which we can think of as fixed). If λ tends to infinity, the resulting estimator $\Phi^{\infty}_{DSGE\text{-}VAR}$ will conform to the restrictions imposed by the DSGE model. Otherwise, it will not. To the extent that the restrictions coming from the DSGE model lead to an estimator that fits the data well, the procedure points toward choosing a high value for λ.

In the above discussion, the definition of "fit" must be clarified. Fit does not simply correspond to large values of the likelihood function or small values of the in-sample sum-of-squared residuals. It is clear that the unrestricted estimator ($\lambda = 0$) always beats the restricted estimator ($\lambda > 0$) in terms of in-sample fit: A constrained optimum cannot fare any better than the unconstrained optimum. What we have in mind is the fit of the model, taking into account the model complexity. Consider the problem of choosing the lag length for a regular VAR. A popular criterion for lag-length selection is the Schwarz criterion. It penalizes the maximized likelihood function by a measure of model complexity, which is a function of the number of parameters to be estimated. The penalty term avoids the problem of the data being overfitted. The choice of λ works similarly, except that complexity cannot be determined by a simple parameter count (the number of VAR parameters is the same for all values of λ). We measure complexity as the degree of uncertainty associated with the parameter estimates. For instance, for $\lambda = 0$ the resulting estimator $\Phi^{0}_{DSGE\text{-}VAR}$ coincides with Φ_{OLS}, which is fairly imprecise. In this case we are using a large penalty. The higher λ, the more the estimator $\Phi^{\lambda}_{DSGE\text{-}VAR}$ is pulled toward the restrictions imposed by the DSGE model and the lower its variance. Hence, for large λ, the penalty that is used to adjust the measure of in-sample fit is small. Overall, if the DSGE model restrictions are very much at odds with the data, one would prefer the uncertainty and choose a low λ. If, however, the model is good, in the sense that the restrictions it imposes are not grossly at odds with the data, then one may welcome the reduction in uncertainty and choose a high value for λ.

In Bayesian terminology, our measure of fit coincides with the marginal likelihood, which is the integral of the likelihood function over all possible parameter values, weighted by the prior density. The marginal likelihood can be approximated by a penalized likelihood function as described above. We use the (exact) marginal data density to find the optimal value of λ (see Del Negro and Schorfheide, forthcoming, section 3.3.1).

The DSGE Model Used to Generate the Artificial Data

The methodology behind DSGE-VAR is general, so it does not depend on the specific DSGE model that is chosen. Of course, the better the DSGE model, in the sense that it captures the important features of the economy, the higher the weight λ it should receive in the composition of the augmented sample. We apply our procedure to a fairly standard and simple neo-Keynesian DSGE model. This section very briefly describes the model (see Del Negro and Schorfheide, forthcoming, section 2, for further details).

When written in log-linearized form (that is, all the variables are expressed in percentage deviations from their stochastic steady state), the model boils down to the following three equations:

1. an IS curve relating real output (x_t) to the level of the real interest rate, computed as the nominal rate minus expected inflation ($R_t - E_t\pi_{t+1}$), as well as to technology shocks (z_t), government

spending shocks (g_t), and expectations of future real activity ($E_t x_{t+1}$):

$$(6)\quad x_t = E_t x_{t+1} - \tau^{-1}(R_t - E_t \pi_{t+1}) + (1-\rho_g)\, g_t + \rho_z \tau^{-1} z_t,$$

where τ, ρ_g, and ρ_z measure the agents' relative risk aversion and the degree of persistence of government and technology shocks, respectively;

2. a Phillips curve relating current inflation (π_t) to expectations of future inflation ($E_t \pi_{t+1}$), output, and government spending:

$$(7)\quad \pi_t = \beta E_t \pi_{t+1} + \kappa(x_t - g_t),$$

where κ measures the slope of the Phillips curve and is a function of deep parameters of the model; and

3. a Taylor rule, by which the monetary authority reacts to deviations of inflation from target and of output from potential output when setting the interest rate, R_t:

$$(8)\quad R_t = \rho_R R_{t-1} + (1-\rho_R)(\psi_1 \pi_t + \psi_2 x_t) + \varepsilon_{Rt},$$

where ρ_R is the degree of persistence of monetary shocks and where the coefficients ψ_1 and ψ_2 represent the sensitivity of interest rates to output and inflation.

Forecasting with DSGE-VAR

The introduction to this article stressed the importance of forecasting, both for practitioners and policymakers. This section investigates the forecasting performance of DSGE-VAR in terms of three of the variables that most interest monetary policymakers: real output growth, inflation, and the federal funds rate. The reader must bear in mind that the results presented in this section are particular to the specific DSGE model described in the previous section. More elaborate models may generate different—and possibly better—results. Nonetheless, it is important to assess how DSGE-VAR fares when applied to a very simple model, if only for comparison.

In this section we specifically address two questions. First, is the DSGE prior useful in terms of forecasting? In other words, does the presence of the DSGE prior increase the forecasting performance relative to that of an unrestricted VAR? This question amounts to asking whether the restrictions that the DSGE model imposes on the VAR do good or harm when forecasting. Table 1 addresses this question and shows that by and large the DSGE prior is useful in terms of forecasting over a VAR with no priors. The table shows the percentage improvement in forecast accuracy relative to an unrestricted VAR for horizons from one to sixteen quarters ahead.[10] The forecast accuracy is measured as the root mean squared error of the forecast using a rolling sample from 1975Q3 to 2003Q3, a period that includes a number of recessions. At each point in the rolling sample, we estimate the model using eighty observations (say, at the first date in the rolling sample, we use data from 1955Q4 to 1975Q3; at the second date we use data from 1956Q1 to 1975Q4, etc.).[11] The importance of the DSGE prior, λ, is chosen optimally as described earlier. Of course, in principle the optimal λ depends on the sample—that is, it might change as we move from the beginning to the end of the rolling sample. In practice, as expected, the optimal λ was fairly constant over the rolling sample, around 0.5. A value for λ of 0.5 means that we used half as many artificial observations from the DSGE model as the number of actual observations.

The numbers in Table 1 are positive whenever the accuracy of DSGE-VAR is greater than that of the unrestricted VAR. One can readily see that the

TABLE 1

Percentage Gain in Root Mean Squared Error: DSGE-VAR versus VAR

Horizon (quarters)	Real GDP growth	Inflation	Federal funds rate
1	17.4	8.4	7.3
2	17.0	7.2	5.0
4	15.1	8.8	5.0
6	14.1	10.5	6.6
8	12.4	11.5	8.4
10	14.4	12.3	8.2
12	15.1	12.6	6.4
14	16.2	13.0	6.1
16	19.1	13.2	5.8

Note: The rolling sample is 1975Q3 to 1997Q3 (ninety periods). At each date in the sample, eighty observations are used to estimate the VAR. The forecasts are computed based on the optimal value of λ chosen ex ante.

10. For real output growth and inflation, the quantities being forecast are cumulative. In other words, for a sixteen-quarter horizon we are trying to forecast the average real output growth in the next four years as opposed to the rate of growth of the economy exactly sixteen quarters from now. Results obtained using the noncumulative forecasts deliver the same conclusion, however.

11. Other details of the exercise, such as the prior used for θ, are described in Del Negro and Schorfheide (forthcoming).

TABLE 2

Percentage Gain (Loss) in Root Mean Squared Error: DSGE-VAR versus BVAR with Minnesota Prior

Horizon (quarters)	Real GDP growth	Inflation	Federal funds rate
1	1.1	1.7	–7.6
2	7.0	1.3	–4.9
4	5.8	4.8	–1.9
6	3.5	7.2	–0.7
8	4.2	7.8	–0.2
10	8.0	8.4	–0.6
12	12.5	9.0	0.7
14	17.2	9.6	1.1
16	21.6	10.1	2.4

Note: The rolling sample is 1975Q3 to 1997Q3 (ninety periods). At each date in the sample, eighty observations are used to estimate the VAR. The forecasts are computed based on the optimal values of λ and ι (the weight of the prior in the BVAR with Minnesota priors) chosen ex ante.

DSGE prior increases the forecasting performance relative to that of an unrestricted VAR. All numbers are positive and most of them are large, indicating a substantial improvement in forecast accuracy.

Since unrestricted VARs are often overparameterized, they are seldom used in practice for forecasting because of the imprecision with which they are estimated. The results in Table 1 are interesting because they show that the restrictions coming from the DSGE model can alleviate this problem. However, from these results one still does not know whether DSGE-VAR can be relied upon as a forecasting tool. Hence, the second question we ask in this section is, How does the accuracy of the forecasts from DSGE-VAR compare with that of benchmark forecasting models?

Table 2 addresses this question. The benchmark chosen here is a VAR with a Minnesota prior, a standard one in the forecasting literature. The Minnesota prior shrinks the parameter estimates of the VAR toward a unit root in levels (or logarithmic levels).[12] Unlike Table 1, Table 2 has both positive and negative numbers, indicating that the VAR with a Minnesota prior is a tougher competitor than the unrestricted VAR. For federal funds rate forecasts, the VAR with Minnesota prior has the upper hand. However, for both inflation and output growth, DSGE-VAR generally outperforms the BVAR with a Minnesota prior in terms of forecasting accuracy, and the gain generally increases with the forecast horizon.

This section has shown that the DSGE-VAR forecasts can be regarded as competitive relative to a standard benchmark, particularly for inflation but also for real output growth. Of course, it would be interesting to know how DSGE-VAR fares relative to other benchmarks, such as FRB/US or commercial models. At this stage, however, the comparison might be premature, as these models are based on dozens of variables while in the current application the DSGE-VAR includes only three. In future research we plan to apply the DSGE-VAR procedure to a more sophisticated DSGE model, such as, for instance, the one in Christiano, Eichenbaum, and Evans (2001). The resulting application would then include enough variables to make the comparison with FRB/US or commercial models meaningful.

Policy Experiments with DSGE-VAR

For DSGE-VAR to be a useful tool for policy analysis, being competitive in terms of forecasting is not enough. DSGE-VAR needs to be able to address policy questions such as the following: (1) What would be the impact on real output growth and inflation of a 50 basis point cut in the federal funds rate? (2) What would be the impact on the volatility of real output growth and inflation, and ultimately on people's welfare, of changing the policy rule followed by the Federal Reserve?[13]

Models that can address the first type of questions are called "identified" in the literature. They are so named because they are able to identify the impact (impulse-response) of monetary policy shocks, as distinguished from other disturbances in the economy, and therefore assess the consequence of a shock that moves the federal funds rate down by 50 basis points. DSGE models are clearly identified. To see what happens after a 50 basis point shock to the variables of interest, one simply feeds a monetary policy shock that generates a 50 point drop in the federal funds rate into the model. Cowles Commission–style models are also identified to the extent that they contain an equation describing monetary policy. As far as VARs are concerned, the papers by Bernanke (1986) and Sims (1986) show how to obtain such identification. Sims and Zha (1998) extend this framework to BVARs, that is, to VARs with priors. The next section discusses identification in the context of DSGE-VAR.

The second question is different in nature from the first one. The monetary policy shock of the first question can be seen as a one-time disturbance that would not affect the view that market participants have of the Fed. The shift in the policy rule of the second question is likely to affect the view of market participants and their expectations. Because of the Lucas critique, the set of models that can successfully address the second question grows thinner relative to those that address the first question. This is not to say that Cowles Commission–style models and VARs cannot successfully address *any* policy-shift type of question.[14] However, there are some regime shifts that these models may not be able to address. DSGE-VAR addresses regime shifts, trying to strike a balance between the forecasting accuracy of BVARs and the compliance to the Lucas critique of DSGE models. An example of the resulting procedure is discussed later in the article.

Identification

To understand how identification works in the DSGE-VAR procedure, one may find it helpful to review the identification problem in standard VARs (see Hamilton 1994, chap. 11). The problem is as follows: One can easily estimate the variance-covariance matrix of the VAR innovations U in equation (1), which we called Σ_u. The problem is that these innovations do not have an economic interpretation: They are not shocks to monetary policy, technology, or government spending, etc. One would like to have a mapping—call it Ω—between these economically interpretable shocks, which we call E, and the shocks that we cannot interpret, U:

(9) $U = \Omega E.$

With Ω in hand, it is straightforward to compute impulse responses to, say, monetary policy shocks, which are one of the elements of E. Using equation (9), one can feed monetary policy shocks into U and then use equation (1) to feed the U shock into the variables of interest, the Ys. The identification problem is that Ω cannot in general be recovered from the data.[15] Identified VARs address this problem by imposing restrictions (zero restrictions, sign restrictions, etc.) on the matrix Ω. The approach taken here is to learn about Ω from the DSGE model at hand, consistent with the rest of the procedure. Once we learn about Ω, the impulse responses are obtained from equation (1) using $\Phi^{\lambda}(\theta)_{DSGE\text{-}VAR}$ as the parameter estimate.

For the sake of simplicity, we do not delve into the technical details of the identification procedure (see Del Negro and Schorfheide, forthcoming, section 4.3). A comment about the role of λ in the identification procedure is in order, however. As λ increases, DSGE-VAR will tend to coincide with the VAR approximation of the DSGE model. Hence, the impulse responses from DSGE-VAR will become closer and closer to those from the DSGE model. Figure 1 makes this point visually. The figure plots the impulse responses of (cumulative) real output growth, inflation, and the federal funds rate to a monetary policy shock. The impulse responses are based on the sample 1981Q4–2001Q3. The gold lines in the plot are the impulse responses for the DSGE model, the solid black lines are the mean impulse responses for DSGE-VAR, and the dashed lines are 90 percent confidence bands, which measure the uncertainty surrounding the estimates for the impulse responses. One can readily see that as λ increases from 0.5 to 5, the mean impulse responses for DSGE-VAR move closer to the DSGE model's impulse responses, and the bands narrow.

In our procedure λ is computed endogenously and measures the extent to which we can trust the DSGE model used as a prior. We therefore view positively the fact that the identification procedure

In the best of all possible worlds we would have a DSGE model that forecasts well, so we could forget the VAR correction.

12. Since two of the variables, real output and the price level, enter the VAR as growth rates, in the equations corresponding to these variables we shrink the coefficient on the first lag of the "own" variable toward zero. For instance, in the real output equation we shrink the coefficient on the first lag of real output growth toward zero. This restriction corresponds to the unit root in log level. In the equation corresponding to the federal funds rate, the prior on the first lag is one since this variable enters as a level.
13. In asking this question we assume that the Federal Reserve implicitly follows a monetary policy rule, as in Taylor (1993). Whether this is indeed the case is an issue beyond the scope of this paper.
14. See Leeper and Zha (2003) for an interesting analysis of what identified VARs can and cannot address.
15. Note that since $var(U) = \Sigma_u$, it must be the case that $\Omega var(E)\Omega' = \Sigma_u$. Because of this restriction, it is customary to decompose Ω as $\Omega = chol(\Sigma_u)\Omega^*$, where $chol(\Sigma_u)$ is the Cholesky decomposition of Σ_u, Ω^* is an orthonormal matrix, and $var(E)$ is the identity.

FIGURE 1

Impulse Response Functions to Monetary Policy Shocks

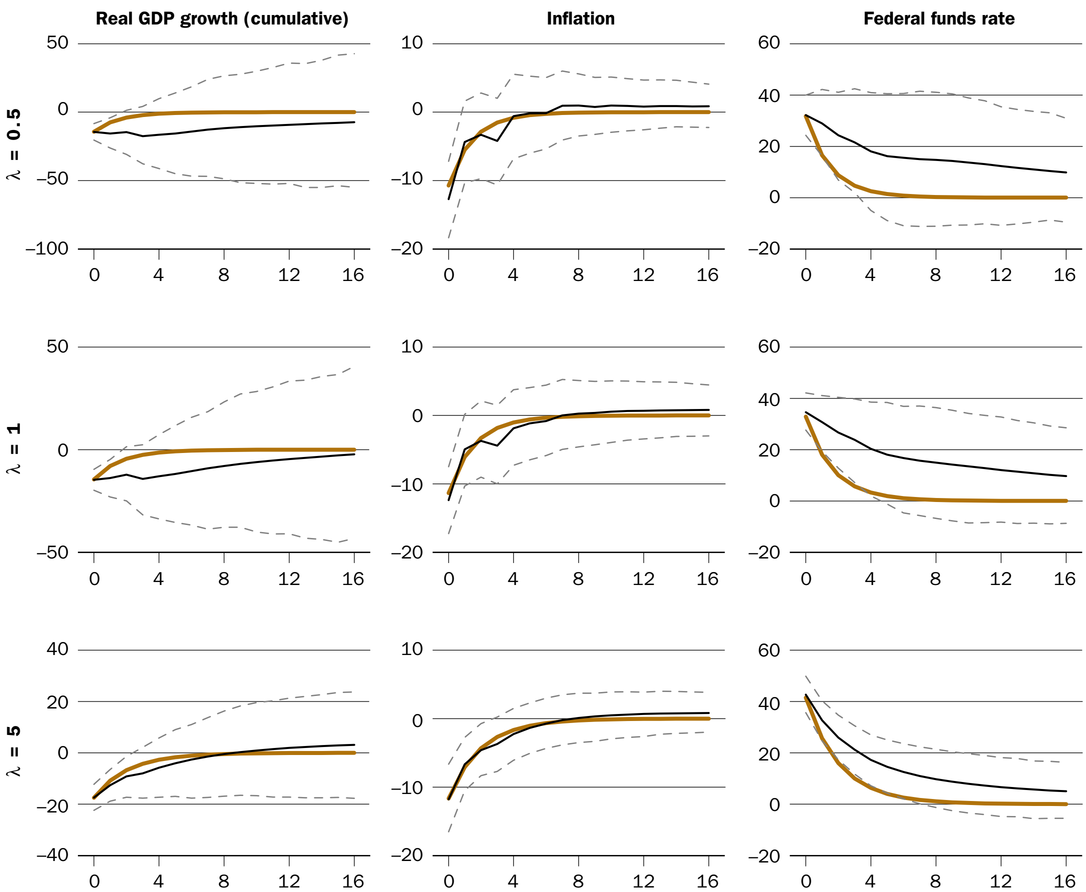

Notes: The solid black lines represent the posterior means of the VAR impulse response functions. The dashed lines are 90 percent confidence bands. The gold lines represent the mean impulse responses from the DSGE model. The impulse responses are based on the sample 1983Q3 to 2003Q2.

hinges on λ: The higher λ, the more we feel confident about the DSGE model at hand and the more reasonable it becomes to use it as a base for identification. Our approach therefore complements the existing literature, where economic theory is often used as an implicit metric to decide whether a given identification procedure works or not.

We now use the identification in DSGE-VAR to address an issue that is relevant to the current policy discussion: What shocks hit the economy during the past four years and, in particular, during the recession? Was the recession the result of monetary policy shocks, as some have claimed, or was it the result of technology or other shocks? Figure 2 plots the time paths of the identified shocks—that is, the *E* variables in equation (9)—as well as the actual paths of the variables entering the VAR: real output growth, inflation, and the federal funds rate.

As described in the model section, the identified shocks are (1) monetary policy shocks, ε_{Rt}; (2) government spending shocks, g_t; and (3) technology shocks, z_t. To describe the findings in Figure 2, it is necessary to discuss the impulse response functions with respect to these shocks, plotted in Figure 3. Impulse response functions simply trace the impact of a one-standard-deviation shock on the variables of interest. A one-standard-deviation shock can be interpreted as the average shock. The impulse responses in Figure 3 are obtained for a value of λ equal to 1, which is the same value under which the identified shocks in Figure 2 are obtained. Although the impulse responses change with λ, as shown in Figure 1, the overall conclusions of this exercise are fairly robust to the choice of λ. As in Figure 1, the gold lines in the plot are the impulse responses for the DSGE model, the solid black lines are the mean

FIGURE 2

Real GDP Growth, Inflation, the Federal Funds Rate, and Identified Shocks from the DSGE-VAR

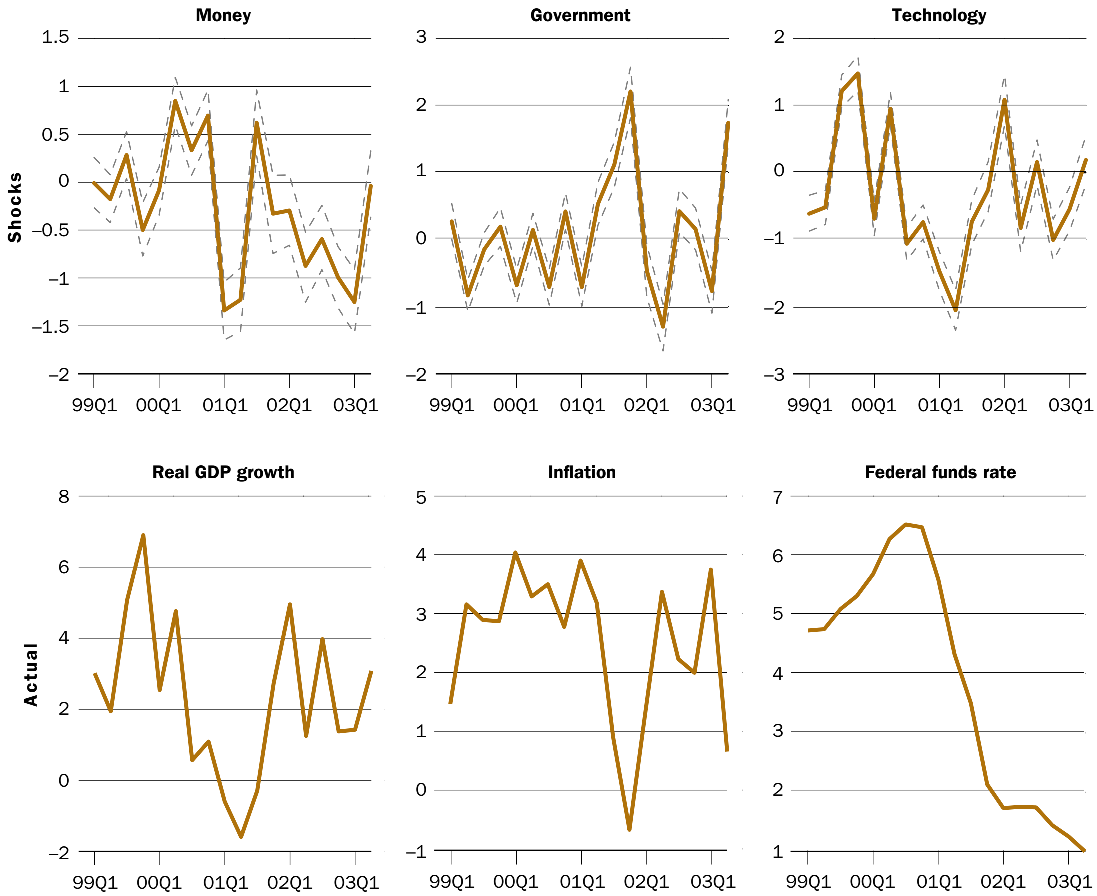

Notes: The solid lines in the three upper plots represent the posterior means of the identified shocks from the DSGE-VAR (1999Q2–2003Q2). The dashed lines are 90 percent confidence bands. The sold lines in the three lower plots represent the actual paths of real GDP growth, inflation, and the federal funds rate. The estimates are based on the sample 1983Q3 to 2003Q2.

impulse responses for DSGE-VAR, and the dashed lines are 90 percent confidence bands.

The impulse responses of output, inflation, and interest rates to a monetary policy shock conform to well-known patterns. A one-standard-deviation shock raises the federal funds rate by 25–30 basis points, decreases real output growth, and lowers inflation (by convention, here a positive monetary policy shock is contractionary). Note that the response of output in DSGE-VAR is more persistent than in the model. The impulse responses to a government shock deserve more explanation because what is called a government shock in the model is not what people generally have in mind. In the model, a positive government shock is essentially equivalent to a shock to the marginal utility of consumption: For given output, an increase in government spending reduces the resources available from consumption and hence increases the marginal utility of consumption.

The increase in the marginal utility of consumption has two effects. On the cost side, it lowers the real wage and hence the marginal cost faced by firms because in equilibrium the real wage is inversely proportional to the marginal utility of consumption: Since agents value their wages more, all else being equal, they need to be paid less. In sticky price models, a decrease in the marginal cost paid by the firm has the effect of lowering inflation. This reasoning explains the negative impact of a government shock on inflation. On the supply side, the other effect of an increase in the marginal utility of consumption is an increase in output: Again, since agents value output more, they have an incentive to produce more. This reasoning explains the positive impact of a government shock on the output growth

F I G U R E 3

Impulse Response Functions to All Shocks

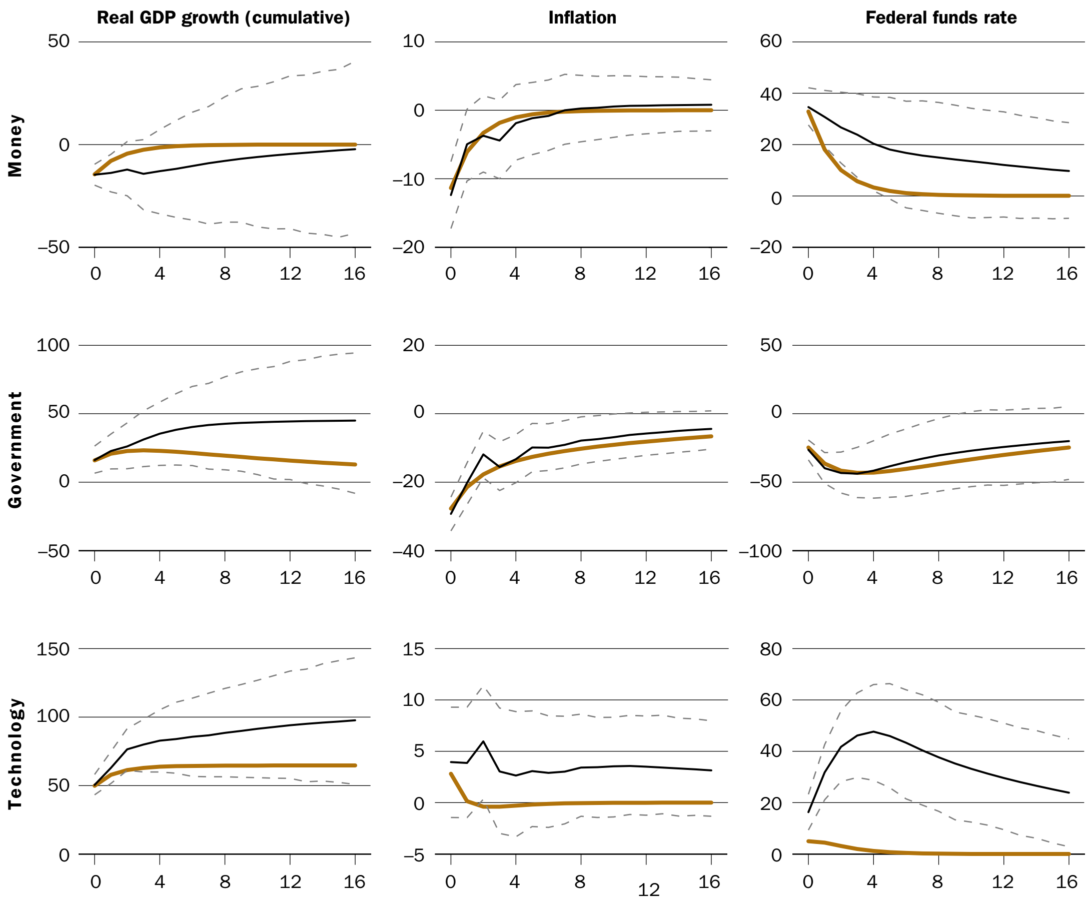

Notes: The solid black lines represent the posterior means of the DSGE-VAR impulse response functions for λ = 1. The dashed lines are 90 percent confidence bands. The gold lines represent the mean impulse responses from the DSGE model. The impulse responses are based on the sample 1983Q3 to 2003Q2.

rate, which is in any case fairly small. The response of the federal funds rate simply mirrors the decline in inflation as it feeds through the Taylor rule. In summary, positive (negative) government spending shocks in the model look very much like positive (negative) oil price shocks, which drive inflation down (up) and output up (down). Finally, technology shocks drive output up. Since the technology shocks in the model are permanent, the increase in output is permanent as well. The impact on inflation is negligible and insignificant.

The shocks plotted in the three upper panels of Figure 2 are measured in terms of standard deviations: A value of 1 (–1) indicates a positive (negative) shock of one standard deviation. In interpreting the plots, one must bear in mind that shocks between –1 and 1 are the norm while shocks outside this range are the exception. The path of monetary policy shocks indicates that by and large such shocks were not responsible for the last recession. It is true that before the recession most monetary policy shocks were positive (contractionary), but they were fairly small. After the start of the recession, most monetary policy shocks were negative, indicating an accommodative monetary policy stance. In particular, according to the model the beginning of 2001 witnessed two large expansionary shocks.

The driving forces of the recession, according to the model, were technology shocks. Figure 2 shows that technology shocks were positive in 1999 but then turned negative, and sizably so, in 2000 and 2001. The only large positive technology shock was associated with the output rebound in the first quarter of 2002. Finally, government spending shocks were negligible up to the third quarter of 2001, when a large positive shock occurred, associated with the sharp decline in inflation. As Figure 2 shows, government spending shocks have the largest impact on inflation. The recent decline in inflation, resulting from the decline in energy prices, is also associated with a positive government spending shock. This result is not surprising because the effect of government spending shocks in the model is similar to the perceived effect of oil shocks in reality, as discussed above. As energy shocks are not part of the model, their effect is likely attributed to government spending shocks. This remark underscores that the analysis just conducted is in many ways heroic because it is done with a very stylized model and using only a few variables. Yet the purpose of the analysis was to illustrate how, in general, DSGE-VAR can be used to uncover the disturbances affecting the economy.

Regime Shifts

This section describes how DSGE-VAR works under a hypothetical policy experiment. Let us put ourselves in the shoes of Paul Volcker as he took office as chairman of the Federal Reserve Board at the end of the second quarter of 1979. Suppose that he had two options: of being either soft on inflation (labeled policy A) or tough on inflation (labeled policy B). In terms of the Taylor rule in equation (8), policy A corresponds to a low reaction to deviations of inflation from target in the Taylor rule (a low value for ψ_1, say, $\psi_1 = 1.1$) while policy B corresponds to a high value for ψ_1 (say, $\psi_1 = 1.7$).[16] In this hypothetical policy experiment, Chairman Volcker uses macroeconomic stability—measured by the standard deviations of output growth, inflation, and the interest rate in the next twenty years—as the criterion to choose between policies A and B.

To understand how the policy experiment under DSGE-VAR works, it is instructive to see how it would work under a DSGE model. Recall that θ is the vector of deep parameters and that ψ_1 is one of the elements of this vector. Let us assume that the only difference between policy A and B lies in the choice of ψ_1. To perform the policy experiment under a DSGE model, one would estimate the remaining elements of θ using pre-1979Q3 data. Call θ_p, $p = A, B$ the vector of deep parameters corresponding to policies A and B. One would then use the DSGE model to make twenty-year forecasts for the variables of interest and, finally, compute the standard deviation of the forecast paths. To the extent that the dynamics of the DSGE model are reasonably well approximated by a VAR, the forecasts can be obtained from a vector autoregression with coefficients:

$$(10)\quad \Phi^{\lambda=\infty}(\theta_p)_{DSGE\text{-}VAR} = \Gamma^*_{xx}(\theta_p)^{-1}\Gamma^*_{xy}(\theta_p),\quad \text{for } p = A, B.$$

Note that in equation (10) the second moments $\Gamma^*_{xx}(\theta_p)$ and $\Gamma^*_{xy}(\theta_p)$ are computed in full compliance with the Lucas critique. That is, these second moments reflect the fact that agents would behave differently when policy moves from A to B.

To perform the policy experiment under the DSGE-VAR procedure, one would estimate the vector of deep parameters θ using pre-1979Q3 data as described earlier in the article. One would then replace the estimate of ψ_1 with the values 1.1 for policy A and 1.7 for policy B and obtain twenty-year forecasts for the variables of interest using equation (5), which is shown below written in a slightly different way:

$$\Phi^{\lambda}(\theta_p)_{DSGE\text{-}VAR} = \left(\Gamma^*_{xx}(\theta_p) + \frac{1}{\lambda}\frac{X'X}{T}\right)^{-1} \left(\Gamma^*_{xy}(\theta_p) + \frac{1}{\lambda}\frac{X'Y}{T}\right),$$

$p = A, B$, where Y and X represent the pre-1979Q3 data. Note that the second moments computed from the data, $(X'X)/T$ and $(X'Y)/T$, do not depend on policy and do not reflect the change in the agent's behavior resulting from the policy shift. This outcome implies that policy experiments under the DSGE-VAR procedure are in full compliance with the Lucas critique only in the $\lambda = \infty$ case. For λ less than infinity, the backward-looking components $(X'X)/T$ and $(X'Y)/T$ are still present. To the extent that the DSGE model does not fit the data well enough, these terms work as a data-driven "correction" to achieve a good forecasting performance. Clearly, in the best of all possible worlds we would have a DSGE model that forecasts well, so we could set $\lambda = \infty$ and forget the backward-looking correction.

The remainder of the section shows the results of the Volcker policy experiment. Figure 4 plots the distributions of outcomes according to policies A and B. Since the assumed criterion of choice is macroeconomic stability, the outcomes we are inter-

16. The values of ψ_1 used in the two policy regimes are broadly consistent with estimated Taylor-rule inflation coefficients obtained over pre- and post-Volcker sample periods by authors such as Clarida, Gali, and Gertler (2000).

FIGURE 4

Effects of a Policy Regime Shift

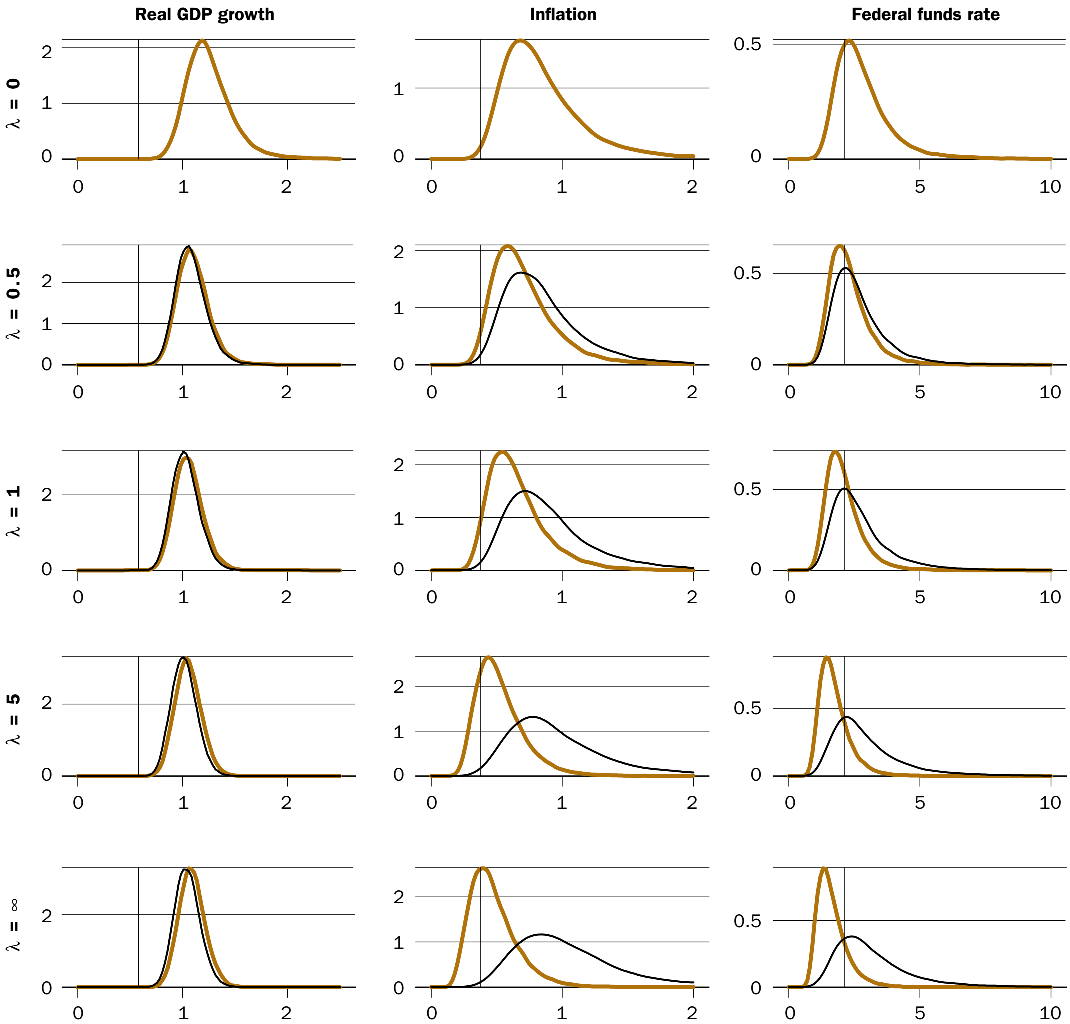

Notes: The vertical lines correspond to the sample standard deviation of the actual data from 1982Q4 to 1999Q2. The solid black and gold lines are posterior predictive distributions of sample standard deviations for the same time period, obtained using data up to 1979Q2. The solid black line corresponds to $\psi_1 = 1.1$; the gold line corresponds to $\psi_1 = 1.7$.

ested in are the standard deviations of the variables of interest. Remember that for each policy option there is not only one possible outcome but a whole distribution of outcomes, reflecting the uncertainty about the parameters of the model as well the shocks that may hit the economy. Figure 4 is organized as a matrix. The columns of the matrix correspond to real output growth, inflation, and the interest rate, respectively. The rows of the matrix correspond to the relative weight of artificial versus actual data in the augmented sample. The first row ($\lambda = 0$) uses only actual data: This amounts to using the unrestricted VAR only. The last row ($\lambda = \infty$) uses the DSGE model only. The rows in between show the results for values of λ that are between 0 and ∞. Each entry of the matrix plots the standard deviation of the corresponding macroeconomic variable according to option A (solid black line) and option B (gold line). For each plot the vertical line shows what actually occurred from 1982Q4 to 1999Q2.[17]

Notice first that when the DSGE model is not used (first row, $\lambda = 0$) the black and gold lines overlap since the effect of the policy change is embodied only in the artificial data (the $\Gamma^*_{xx}(\theta_p)$ and $\Gamma^*_{xy}(\theta_p)$ terms), which have no weight in this case. As the weight of the artificial data (λ) increases, the predictions from policy A and policy B start to diverge. The forecasts suggest that policy B (tough on inflation) delivers lower variability in both inflation and the interest rate than does policy A (soft on inflation): The gold densities are shifted to the left of, and are narrower than, the black densities. Interestingly, policy B delivers not only a lower variability in inflation, as expected, but also a lower variability of interest rates in spite of the fact that the interest rate reacts more, and not less, to inflation under policy B. This effect works through agents' expectations: Since agents expect monetary policy to reign in inflation under policy B, they will expect lower inflation variability. Their expectations will be realized, and in equilibrium the interest rate will not have to move much. In other words, the threat to react to inflation is enough to lower inflation variability, avoiding wide swings in interest rates.

Although this is a one-time experiment and not a test of the forecasting accuracy of the model, it is interesting to consider how accurate the predictions from DSGE-VAR are in this case. Again, for each plot the dotted vertical lines correspond to the sample standard deviation of the actual data from 1982Q4 to 1999Q2. As far as output is concerned, there is no difference across policies. This result is expected because the difference between policy A and policy B regards the response to inflation and not to output. Not surprisingly, both models overpredict the standard deviation of real output growth: Both the parameters of the BVAR and those of the DSGE model are estimated using data up to 1979, that is, a period in which real output volatility was higher than in the 1980s and 1990s. In terms of inflation, policy B is clearly more on target than policy A is, as it should be since the Taylor rule parameters in policy B are broadly consistent with estimated coefficients obtained over the post-Volcker sample period. Policy A overpredicts the variability of inflation. Also, its forecasts are much more uncertain than those from policy B. For interest rate prediction, for high values of λ policy B appears to underpredict the volatility of the federal funds rate. Policy A, on the other hand, tends to overpredict the rate's volatility. As discussed earlier, the current application of DSGE-VAR is not very accurate in forecasting interest rates.

Summary

This article describes the workings of DSGE-VAR, a procedure that aims to combine VARs and DSGE models. The ultimate goal of the procedure is to provide a proper assessment of the impact of different monetary policy rules and at the same time provide a tool that can also be relied upon for forecasting. It may well be that in the not-too-distant future a full-fledged DSGE model will attain both goals. In the meanwhile, DSGE-VAR may provide a viable alternative to the models that are currently in use for forecasting and policy analysis.

17. Following Clarida, Gali, and Gertler (2000), we compute the actual excluding the pre-1983 disinflation period.

REFERENCES

Altug, Sumru. 1989. Time-to-build and aggregate fluctuations: Some new evidence. *International Economic Review* 30, no. 4:889–920.

Bernanke, Ben S. 1986. Alternative explanations of the money-income correlation. *Carnegie-Rochester Series on Public Policy* 25 (Autumn): 49–99.

Brayton, Flint, Andrew Levin, Ralph Tryon, and John C. Williams. 1997. The evolution of macro models at the Federal Reserve Board. *Carnegie-Rochester Conference Series on Public Policy* 47 (December): 43–81.

Brayton, Flint, Eileen Mauskopf, Dave Reifschneider, Peter Tinsley, and John Williams. 1997. The role of expectations in the FRB/US macroeconomic model. *Federal Reserve Bulletin* 83 (April): 227–45.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2001. Nominal rigidities and the dynamic effects of a shock to monetary policy. Federal Reserve Bank of Chicago Working Paper 2001-08.

Clarida, Richard, Jordi Gali, and Mark Gertler. 2000. Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics* 115, no. 1:147–80.

Del Negro, Marco, and Frank Schorfheide. Forthcoming. Priors from general equilibrium models for VARs. *International Economic Review*.

Diebold, Francis X. 1998. The past, present, and future of macroeconomic forecasting. *Journal of Economic Perspectives* 12 (Spring): 175–92.

Doan, Thomas, Robert Litterman, and Christopher A. Sims. 1984. Forecasting and conditional projections using realistic prior distributions. *Econometric Reviews* 3, no. 1:1–100.

Fair, Ray. 1992. The Cowles Commission approach, real business cycle theories, and new Keynesian economics. Yale University, Cowles Foundation Discussion Paper No. 1004.

Hamilton, James D. 1994. *Time series analysis*. Princeton, N.J.: Princeton University Press.

Hodrick, Robert J., and Edward C. Prescott. 1997. Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit, and Banking* 29, no. 1:1–16.

Ingram, Beth F., and Charles H. Whiteman. 1994. Supplanting the "Minnesota" prior: Forecasting macroeconomic time series using real business cycle model priors. *Journal of Monetary Economics* 34 (December): 497–510.

Ireland, Peter N. 1997. A small, structural, quarterly model for monetary policy evaluation. *Carnegie-Rochester Conference Series on Public Policy* 47 (December): 83–108.

Kydland, Finn E., and Edward C. Prescott. 1982. Time to build and aggregate fluctuations. *Econometrica* 50 (November): 1345–70.

———. 1996. The computational experiment: An econometric tool. *Journal of Economic Perspectives* 10 (Winter): 69–85.

Leeper, Eric M., and Christopher A. Sims. 1994. Toward a modern macroeconomic model usable for policy analysis. *NBER macroeconomics annual*. Cambridge and London: MIT Press.

Leeper, Eric M., Christopher A. Sims, and Tao Zha. 1996. What does monetary policy do? *Brookings Papers on Economic Activity* 2:1–78.

Leeper, Eric M., and Tao Zha. 2003. Modest policy interventions. *Journal of Monetary Economics* 50 (November): 1673–1700.

Litterman, Robert B. 1986. Forecasting with Bayesian vector autoregressions—five years of experience. *Journal of Business and Economic Statistics* 4 (January): 25–38.

Lucas, Robert E. 1976. Econometric policy evaluation: A critique. In *The Phillips curve and the labor market*, edited by Karl Brunner and Allan Meltzer. *Carnegie-Rochester Conference Series* 1. Amsterdam: North-Holland.

Reifschneider, David L., David J. Stockton, and David W. Wilcox. 1997. Econometric models and the monetary policy process. *Carnegie-Rochester Conference Series on Public Policy* 47 (December): 1–37.

Robertson, John C., and Ellis W. Tallman. 1999. Vector autoregressions: Forecasting and reality. Federal Reserve Bank of Atlanta *Economic Review* 84 (First Quarter): 4–18.

Sargent, Thomas J. 1984. Autoregressions, expectations, and advice. *American Economic Review* 74, no. 2:408–15.

Schorfheide, Frank. 2000. Loss function-based evaluation of DSGE models. *Journal of Applied Econometrics* 15 (November–December): 645–70.

Sims, Christopher A. 1980. Macroeconomics and reality. *Econometrica* 48 (January): 1–48.

———. 1982. Policy analysis with econometric models. *Brookings Papers on Economic Activity* 1:107–64.

———. 1986. Are forecasting models usable for policy analysis? Federal Reserve Bank of Minneapolis *Quarterly Review* 10 (Winter): 2–16.

———. 2002. The role of models and probabilities in the monetary policy process. *Brooking Papers on Economic Activity* 2:1–40.

Sims, Christopher A., and Tao Zha. 1998. Bayesian methods for dynamic multivariate models. *International Economic Review* 39, no. 4:949–68.

Smith, Anthony A., Jr. 1993. Estimating nonlinear time-series models using simulated vector autoregressions. *Journal of Applied Econometrics* 8:S63–S84.

Taylor, John B. 1993. Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy* 39:195–214.

Theil, Henry, and Arthur S. Goldberger. 1961. On pure and mixed estimation in economics. *International Economic Review* 2:65–78.

Zha, Tao. 1998. A dynamic multivariate model for use in formulating policy. Federal Reserve Bank of Atlanta *Economic Review* 83 (First Quarter): 16–29.