DSGE Model-Based Forecasting of Non-modelled Variables

Frank Schorfheide
University of Pennsylvania, CEPR, NBER, and Visiting Scholar, FRB Philadelphia

Keith Sill*
FRB Philadelphia

Maxym Kryshko
University of Pennsylvania

March 25, 2009

*Correspondence: Frank Schorfheide, Maxym Kryshko: Department of Economics, 3718 Locust Walk, University of Pennsylvania, Philadelphia, PA 19104-6297. Email: schorf@ssc.upenn.edu, mkryshko@sas.upenn.edu. Keith Sill: Research Department, Federal Reserve Bank of Philadelphia, Ten Independence Mall, Philadelphia, PA 19106-1574.

We thank the editor, two anonymous referees, and seminar participants at the Board of Governors, the FRB Philadelphia, the University of Richmond, and Texas A&M University for helpful comments. This research was conducted while Schorfheide was visiting the FRB Philadelphia, for whose hospitality he is thankful. Schorfheide gratefully acknowledges financial support from the Alfred P. Sloan Foundation and the National Science Foundation (Grant SES 0617803). The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

Abstract

This paper develops and illustrates a simple method to generate a DSGE model-based forecast for variables that do not explicitly appear in the model (non-core variables). We use auxiliary regressions that resemble measurement equations in a dynamic factor model to link the non-core variables to the state variables of the DSGE model. Predictions for the non-core variables are obtained by applying their measurement equations to DSGE model- generated forecasts of the state variables. Using a medium-scale New Keynesian DSGE model, we apply our approach to generate and evaluate recursive forecasts for PCE inflation, core PCE inflation, the unemployment rate, and housing starts along with predictions for the seven variables that have been used to estimate the DSGE model.

JEL CLASSIFICATION: C11, C32, C53, E27, E47

KEY WORDS: Bayesian Analysis, DSGE Models, Forecasting

1 Introduction

Dynamic stochastic general equilibrium (DSGE) models estimated with Bayesian methods are increasingly used by central banks around the world as tools for projections and policy analysis. Examples of such models are the small open economy model developed by the Sveriges Riksbank (Adolfson, Laseen, Linde, and Villani, 2005 and 2008; Adolfson, Andersson, Linde, Villani, and Vredin, 2007), the New Area-Wide Model developed at the European Central Bank (Coenen, McAdam, and Straub, 2008) and the Federal Reserve Board's new Estimated, Dynamic, Optimization-based model (Edge, Kiley, and Laforte, 2008). These models extend specifications studied by Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2003) to open economy and multisector settings. A common feature is that decision rules of economic agents are derived from assumptions about preferences and technologies by solving intertemporal optimization problems.

Compared to previous generations of macroeconometric models, the DSGE paradigm delivers empirical models with a strong degree of theoretical coherence. The costs associated with this theoretical coherence are two-fold. First, tight cross-equation restrictions potentially introduce misspecification problems that manifest themselves through inferior fit compared to less-restrictive time series models (Del Negro, Schorfheide, Smets, and Wouters, 2007, henceforth DSSW). Second, it is more cumbersome than in a traditional system-of-equations approach to incorporate variables other than a core set of macroeconomic aggregates such as real gross domestic product (GDP), consumption, investment, wages, hours, inflation, and interest rates. Nonetheless, in practical work at central banks it might be important to also generate forecasts for economic variables that do not explicitly appear in medium-scale DSGE models. Our paper focuses on the second problem.

There are in principle two options for generating forecasts for additional variables. First, one could enlarge the structural model to incorporate these variables explicitly. The advantage of a larger model is its ability to deliver a coherent narrative that can accompany the forecasts. The disadvantages are that identification problems are often exacerbated in large-scale models, the numerical analysis, e.g., estimation procedures that utilize numerical optimization or posterior simulation routines, becomes more tenuous, and the maintenance of the model requires more staff resources. The second option is to develop a hybrid empirical model that augments a medium-scale core DSGE model with auxiliary equations that create a link between explicitly modelled variables and non-modelled variables. For brevity we will refer to the latter as non-core variables. One could interpret these auxiliary equations as log-linear approximations of agents' decision rules in a larger DSGE model. This paper explores the second approach.

Recently, Boivin and Giannoni (2006, henceforth BG) integrated a medium-scale DSGE model into a dynamic factor model for a large cross section of macroeconomic indicators, thereby linking non-core variables to a DSGE model. We will refer to this hybrid model as DSGE-DFM. The authors jointly estimated the DSGE model parameters as well as the factor loadings for the non-core variables. Compared to the estimation of a "non-structural" dynamic factor model, the BG approach leads to factor estimates that have a clear economic interpretation. The joint estimation is conceptually very appealing, in part because it exploits information that is contained in the non-core variables when making inference about the state of the economy.[1] The downside of the joint estimation is its computational complexity, which makes it currently impractical for real time forecasting applications.

Our paper proposes a simpler two-step estimation approach for an empirical model that consists of a medium-scale DSGE model for a set of core macroeconomic variables and a collection of measurement equations or auxiliary regressions that link the state variables of the DSGE model with the non-core variables of interest to the analyst. In the first step we estimate the DSGE model using the core variables as measurements. Based on the DSGE model parameter estimates, we apply the Kalman filter to obtain estimates of the latent state variables given the most recent information set. We then use the filtered state variables as regressors to estimate simple linear measurement equations with serially correlated idiosyncratic errors.

The advantage of our procedure is three-fold. First, since the DSGE model estimation is fairly tedious and delicate, in real time applications the DSGE model could be re-estimated infrequently, for instance, once a year. Second, the estimation of the measurement equations is quick and can be easily repeated in real time as new information arrives or interest in additional non-core variables arises. The estimated auxiliary regressions can then be used to generate forecasts of the non-core variables. Third, our empirical model links the non-core variables to the fundamental shocks that are the believed drivers of business cycle fluctuations. In particular, the model allows monetary policy shocks and other structural shocks to propagate through to non-core variables. This allows us to study the effect of unanticipated changes in monetary policy on a broad set of economic variables.[2]

[1] Formally we mean by "state of the economy" information about the latent state variables that appear in the DSGE model.

[2] The The goal of our analysis is distinctly different from recent work by Giannone, Monti, and Reichlin

The remainder of the paper is organized as follows. The DSGE model used for the empirical analysis is described in Section 2. We are using a variant of the Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2003) model, which is described in detail in DSSW. Our econometric framework is presented in Section 3. Section 4 summarizes the results of our empirical analysis. We estimate the DSGE model recursively based on U.S. quarterly data starting with a sample from 1984:I to 2000:IV and generate estimates of the latent states as well as pseudo-out-of-sample forecasts for a set of core variables, that is comprised of the growth rates of output, consumption, investment, nominal wages, the GDP deflator, as well as the levels of interest rates and hours worked. We then estimate measurement equations for four additional variables: personal consumption expenditures (PCE) inflation, core PCE inflation, the unemployment rate, and housing starts. We provide pseudo-out-of-sample forecast error statistics for both the core and non-core variables using our empirical model and compare them to simple AR(1) forecasts. Finally, we study the propagation of monetary policy shocks to auxiliary variables as well as features of the joint predictive distribution. Section 5 concludes and discusses future research. Details of the Bayesian computations are relegated to the Appendix.

2 The DSGE Model

We use a medium-scale New Keynesian model with price and wage rigidities, capital accumulation, investment adjustment costs, variable capital utilization, and habit formation. The model is based on the work of Smets and Wouters (2003) and Christiano, Eichenbaum, and Evans (2005). The specific version is taken from DSSW. For brevity we only present the log-linearized equilibrium conditions and refer the reader to the above-referenced papers for the derivation of these conditions from assumptions on preferences and technologies.

The economy is populated by a continuum of firms that combine capital and labor to produce differentiated intermediate goods. These firms have access to the same Cobb-Douglas production function with capital elasticity α and total factor productivity A_t. Total factor productivity is assumed to be non-stationary. We denote its growth rate by $a_t = \ln(A_t/A_{t-1})$, which is assumed to have mean γ. Output, consumption, investment, capital, and the real wage can be detrended by A_t. In terms of the detrended variables the model

(2008), and Monti (2008) who develop state-space models that allow the analyst to use high frequency data or professional forecasts to update or improve the DSGE-model based forecasts of the core variables.

has a well-defined steady state. All variables that appear subsequently are expressed as log-deviations from this steady state.

The intermediate goods producers hire labor and rent capital in competitive markets and face identical real wages, w_t, and rental rates for capital, r_t^k. Cost minimization implies that all firms produce with the same capital-labor ratio

$$k_t - L_t = w_t - r_t^k \tag{1}$$

and have marginal costs

$$mc_t = (1-\alpha)w_t + \alpha r_t^k. \tag{2}$$

The intermediate goods producers sell their output to perfectly competitive final good producers, which aggregate the inputs according to a CES function. Profit maximization of the final good producers implies that

$$\widehat{y}_t(j) - \widehat{y}_t = -\left(1 + \frac{1}{\lambda_f e^{\widetilde{\lambda}_{f,t}}}\right)(p_t(j) - p_t). \tag{3}$$

Here $\widehat{y}_t(j) - \widehat{y}_t$ and $p_t(j) - p_t$ are quantity and price for good j relative to quantity and price of the final good. The price p_t of the final good is determined from a zero-profit condition for the final good producers.

We assume that the price elasticity of the intermediate goods is time-varying. Since this price elasticity affects the mark-up that intermediate goods producers can charge over marginal costs, we refer to $\widetilde{\lambda}_{f,t}$ as mark-up shock. Following Calvo (1983), we assume that in every period a fraction of the intermediate goods producers ζ_p is unable to re-optimize their prices. These firms adjust their prices mechanically according to steady state inflation π_*. All other firms choose prices to maximize the expected discounted sum of future profits, which leads to the following equilibrium relationship, known as the New Keynesian Phillips curve:

$$\pi_t = \beta I\!E_t[\pi_{t+1}] + \frac{(1-\zeta_p\beta)(1-\zeta_p)}{\zeta_p} mc_t + \frac{1}{\zeta_p}\lambda_{f,t}, \tag{4}$$

where π_t is inflation and β is the discount rate.[3] Our assumption on the behavior of firms that are unable to re-optimize their prices implies the absence of price dispersion in the steady state. As a consequence, we obtain a log-linearized aggregate production function of the form

$$\widehat{y}_t = (1-\alpha)L_t + \alpha k_t. \tag{5}$$

[3] We used the following re-parameterization: $\lambda_{f,t} = [(1-\zeta_p\beta)(1-\zeta_p)\lambda_f/(1+\lambda_f)]\widetilde{\lambda}_{f,t}$, where λ_f is the steady state of $\widetilde{\lambda}_{f,t}$.

Equations (2), (1), and (5) imply that the labor share lsh_t equals marginal costs in terms of log-deviations: $lsh_t = mc_t$.

There is a continuum of households with identical preferences, which are separable in consumption, leisure, and real money balances. Households' preferences display (internal) habit formation in consumption captured by the parameter h. Period t utility is a function of $\ln(C_t - hC_{t-1})$. Households supply monopolistically differentiated labor services. These services are aggregated according to a CES function that leads to a demand elasticity $1 + 1/\lambda_w$. The composite labor services are then supplied to the intermediate goods producers at real wage w_t. To introduce nominal wage rigidity, we assume that in each period a fraction ζ_w of households is unable to re-optimize their wages. These households adjust their nominal wage by steady state wage growth $e^{(\pi_* + \gamma)}$. All other households re-optimize their wages. First-order conditions imply that

$$\begin{aligned} \tilde{w}_t &= \zeta_w \bera \mathbb{E}_t\Big[\tilde{w}_{t+1} + \Delta w_{t+1} + \pi_{t+1} + a_{t+1}\Big] \\ &\quad + \frac{1 - \zeta_w \beta}{1 + \nu_l(1+\lambda_w)/\lambda_w}\left(\nu_l L_t - w_t - \xi_t + \tilde{b}_t + \frac{1}{1-\zeta_w\beta}\phi_t\right), \end{aligned} \tag{6}$$

where $\tilde{w}_t$ is the optimal real wage relative to the real wage for aggregate labor services, w_t, and ν_l would be the inverse Frisch labor supply elasticity in a model without wage rigidity ($\zeta_w = 0$) and differentiated labor. Moreover, $\tilde{b}_t$ is a shock to the household's discount factor[4] and ϕ_t is a preference shock that affects the household's intratemporal substitution between consumption and leisure. The real wage paid by intermediate goods producers evolves according to

$$w_t = w_{t-1} - \pi_t - a_t + \frac{1-\zeta_w}{\zeta_w}\tilde{w}_t. \tag{7}$$

Households are able to insure the idiosyncratic wage adjustment shocks with state contingent claims. As a consequence they all share the same marginal utility of consumption ξ_t, which is given by the expression:

$$\begin{aligned} (e^\gamma - h\beta)(e^\gamma - h)\xi_t &= -(e^{2\gamma} + \beta h^2)c_t + \beta h e^\gamma \mathbb{E}_t[c_{t+1} + a_{t+1}] + he^\gamma(c_{t-1} - a_t) \\ &\quad + e^\gamma(e^\gamma - h)\tilde{b}_t - \beta h(e^\gamma - h)\mathbb{E}_t[\tilde{b}_{t+1}], \end{aligned} \tag{8}$$

where c_t is consumption. In addition to state-contingent claims, households accumulate three types of assets: one-period nominal bonds that yield the return R_t, capital $\bar{k}_t$, and real money balances. Since preferences for real money balances are assumed to be additively separable and monetary policy is conducted through a nominal interest rate feedback rule,

[4]For the estimation we re-parameterize the shock as follows: $b_t = e^\gamma(e^\gamma - h)/(e^{2\gamma} + \beta h^2)\tilde{b}_t$.

money is block exogenous and we will not use the households' money demand equation in our empirical analysis.

The first order condition with respect to bond holdings delivers the standard Euler equation:

$$\xi_t = \mathbb{E}_t[\xi_{t+1}] + R_t - \mathbb{E}_t[\pi_{t+1}] - \mathbb{E}_t[a_{t+1}]. \tag{9}$$

Capital accumulates according to the following law of motion:

$$\bar{k}_t = (2 - e^{\gamma} - \delta)\big[\bar{k}_{t-1} - a_t\big] + (e^{\gamma} + \delta - 1)[i_t + (1+\beta)S'' e^{2\gamma}\mu_t], \tag{10}$$

where i_t is investment, δ is the depreciation rate of capital, and μ_t can be interpreted as an investment-specific technology shock. Investment in our model is subject to adjustment costs, and S'' denotes the second derivative of the investment adjustment cost function at steady state. Optimal investment satisfies the following first-order condition:

$$i_t = \frac{1}{1+\beta}\big[i_{t-1} - a_t\big] + \frac{\beta}{1+\beta}\mathbb{E}_t[i_{t+1} + a_{t+1}] + \frac{1}{(1+\beta)S'' e^{2\gamma}}(\xi^k_t - \xi_t) + \mu_t, \tag{11}$$

where ξ^k_t is the value of installed capital, evolving according to:

$$\xi^k_t - \xi_t = \beta e^{-\gamma}(1-\delta)\mathbb{E}_t\big[\xi^k_{t+1} - \xi_{t+1}\big] + \mathbb{E}_t\big[(1-(1-\delta)\beta e^{-\gamma})r^k_{t+1} - (R_t - \pi_{t+1})\big]. \tag{12}$$

Capital utilization u_t in our model is variable and r^k_t in all previous equations represents the rental rate of effective capital $k_t = u_t + \bar{k}_{t-1}$. The optimal degree of utilization is determined by

$$u_t = \frac{r^k_*}{a''} r^k_t. \tag{13}$$

Here a'' is the derivative of the per-unit-of-capital cost function $a(u_t)$ evaluated at the steady state utilization rate. The central bank follows a standard feedback rule:

$$R_t = \rho_R R_{t-1} + (1-\rho_R)(\psi_1 \pi_t + \psi_2 \widehat{y}_t) + \sigma_R \epsilon_{R,t}. \tag{14}$$

where $\epsilon_{R,t}$ represents monetary policy shocks. The aggregate resource constraint is given by:

$$\widehat{y}_t = (1+g_*)\left[\frac{c_*}{y_*}c_t + \frac{i_*}{y_*}\left(i_t + \frac{r^k_*}{e^{\gamma} - 1 + \delta}u_t\right)\right] + g_t. \tag{15}$$

Here c_*/y_* and i_*/y_* are the steady state consumption-output and investment-output ratios, respectively, and $g_*/(1+g_*)$ corresponds to the government share of aggregate output. The process g_t can be interpreted as exogenous government spending shock. It is assumed that fiscal policy is passive in the sense that the government uses lump-sum taxes to satisfy its period budget constraint.

There are seven exogenous disturbances in the model and six of them are assumed to follow AR(1) processes:

$$\begin{aligned} a_t &= \rho_a a_{t-1} + (1-\rho_a)\gamma + \sigma_a \epsilon_{a,t} \quad (16) \\ \mu_t &= \rho_\mu \mu_{t-1} + \sigma_\mu \epsilon_{\mu,t} \\ \lambda_{f,t} &= \rho_{\lambda_f} \lambda_{f,t-1} + \sigma_{\lambda_f} \epsilon_{\lambda_f} \\ g_t &= \rho_g g_{t-1} + \sigma_g \epsilon_{g,t} \\ b_t &= \rho_b b_{t-1} + \sigma_b \epsilon_{b,t} \\ \phi_t &= \rho_\phi \phi_{t-1} + \sigma_\phi \epsilon_{\phi,t}. \end{aligned}$$

We assume that innovations of these exogenous processes as well as the monetary policy shock $\epsilon_{R,t}$ are independent standard normal random variates and collect them in the vector ϵ_t. We stack all the DSGE model parameters in the vector θ. The equations presented in this section form a linear rational expectations system that can be solved numerically, for instance with the method described in Sims (2002).

3 Econometric Methodology

Our econometric analysis proceeds in three steps. First, we use Bayesian methods to estimate the linearized DSGE model described in Section 2 on seven core macroeconomic time series. Second, we estimate so-called auxiliary regression equations that link the state-variables associated with the DSGE model to other macroeconomic variables that are of interest to the analyst, but not explicitly included in the structural DSGE model (non-core variables). Finally, we use the estimated DSGE model to forecast its state variables and then map these state forecasts into predictions for the core and non-core variables.

3.1 DSGE Model Estimation

The solution of the linear rational expectations system characterized in Section 2 can be expressed as a vector autoregressive law of motion for a vector of non-redundant state variables s_t:

$$s_t = \Phi_1(\theta) s_{t-1} + \Phi_\epsilon(\theta)\epsilon_t. \quad (17)$$

The coefficients of the matrices Φ_1 and Φ_ϵ are functions of the DSGE model parameters θ and the vector s_t is given by

$$s_t = [c_t, i_t, \bar{k}_t, R_t, w_t, a_t, \phi_t, \mu_t, b_t, g_t, \lambda_{f,t}]'.$$

The variables c_t, i_t, $\bar{k}_t$, R_t, and w_t are endogenous state variables, whereas the remaining elements of s_t are exogenous state variables. To estimate the DSGE model based on a sequence of observations $Y^T = [y_t, \ldots, y_T]$, it is convenient to construct a state-space model by specifying a system of measurement equations that link the observables y_t to the states s_t.

The vector y_t used in our empirical analysis consists of quarter-to-quarter growth rates (measured in percentages) of real GDP, consumption, investment, and nominal wages, as well as a measure of hours worked, GDP deflator inflation, and the federal funds rate. Since some of our observables include growth rates, we augment the set of model states s_t by lagged values of output, consumption, investment, and real wages. More specifically, notice that lagged consumption, investment, and real wages are elements of the vector s_{t-1}. Moreover, according to the DSGE model solution, lagged output, $\widehat{y}_{t-1}$, can be expressed as a linear function of the elements of s_{t-1}. Thus, we can write

$$[\widehat{y}_{t-1}, c_{t-1}, i_{t-1}, w_{t-1}]' = M_s(\theta) s_{t-1}$$

for a suitably chosen matrix $M_s(\theta)$ and define

$$\varsigma_t = [s_t', s_{t-1}' M_s'(\theta)]'. \tag{18}$$

This allows us to express the set of measurement equations as

$$y_t = A_0(\theta) + A_1(\theta)\varsigma_t. \tag{19}$$

The state-space representation of the DSGE model is comprised of (17), (18), and (19).

Under the assumption that the innovations ϵ_t are normally distributed, the likelihood function, denoted by $p(Y^T|\theta)$, for the DSGE model can be evaluated with the Kalman filter. The Kalman filter also generates a sequence of estimates of the state vector ς_t:

$$\varsigma_{t|t}(\theta) = I\!\!E[\varsigma_t | \theta, Y^t], \tag{20}$$

where $Y^t = [y_1, \ldots, y_t]$. Our Bayesian estimation of the DSGE model combines a prior $p(\theta)$ with the likelihood function $p(Y^T|\theta)$ to obtain a joint probability density function for data and parameters. The posterior distribution is given by

$$p(\theta|Y^T) = \frac{p(Y^T|\theta)p(\theta)}{p(Y)}, \quad \text{where} \quad p(Y^T) = \int p(Y^T|\theta)p(\theta)d\theta. \tag{21}$$

We employ Markov-Chain-Monte-Carlo (MCMC) methods described in detail in An and Schorfheide (2007) to implement the Bayesian inference. More specifically, a random-walk

Metropolis algorithm is used to generate draws from the posterior distribution $p(\theta|Y^T)$ and averages of these draws (and suitable transformations) serve as approximations for posterior moments of interest.

3.2 Linking Model States to Non-Core Variables

Due to the general equilibrium structure the variables that are included in state-of-the-art DSGE models are limited to a set of core macroeconomic indicators. However, in practice an analyst might be interested in forecasting a broader set of time series. For instance, the DSGE model described in Section 2 generates predictions for hours worked but does not include unemployment as one of the model variables. We use z_t to denote a particular variable that is not included in the DSGE model but nonetheless is of interest to the forecaster. We will express z_t as a function of the DSGE model state variables s_t. According to (18) one can easily recover s_t from the larger vector ς_t using a selection matrix M with the property $s_t = M\varsigma_t$. As discussed in the previous subsection, the Kalman filter delivers a sequence $\varsigma_{t|t}(\theta)$, $t = 1, \ldots, T$. We use $\hat{\varsigma}_{t|t}$ to denote an estimate of $\varsigma_{t|t}(\theta)$ that is obtained by replacing θ with the posterior mean estimate $\hat{\theta}_T$, define $\hat{s}_{t|t} = M\hat{\varsigma}_{t|t}$, and let[5]

$$z_t = \alpha_0 + \hat{s}'_{t|t}\alpha_1 + \xi_t, \quad \xi_t = \rho\xi_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2). \tag{22}$$

Moreover, ξ_t is a variable-specific noise process. The parameters of this auxiliary regression are collected in the vector $\psi = [\alpha_0, \alpha_1', \rho, \sigma_\eta]'$. As for the estimation of the DSGE model, we use Bayesian methods for the estimation of the auxiliary regression (22).

A few remarks about our setup are in order. First, Equations (17), (18), (19), and (22) can be interpreted as a factor model. The factors are given by the state variables of the DSGE model, the measurement equation associated with the DSGE model describes how our core macroeconomic variables load on the factors, and auxiliary regressions of the form (22) describe how additional (non-core) macroeconomic variables load on the factors. The random variable ξ_t in (22) plays the role of an idiosyncratic error term.

Second, our setup can be viewed as a simplified version of BG's framework. Unlike BG, we do not attempt to estimate the DSGE model and the auxiliary equations simultaneously. While we are thereby ignoring information about s_t contained in the z_t variables, our analysis reduces the computational burden considerably and can be more easily used for real time forecasting. The BG approach is computationally cumbersome. A Markov-Chain Monte

[5] At this point it is important that the state vector does not contain redundant elements. If it did, the auxiliary regression (22) would suffer from perfect collinearity.

Carlo algorithm has to iterate over the conditional distributions of θ, ψ, and the sequence of states $S^T = [s_1, \ldots, s_T]$. Drawing from the posterior of S^T is computationally costly because it requires forward and backward iterations of the Kalman filter. Drawing from the distribution of θ requires a Metropolis-Hastings step and, unlike in a stand-alone estimation of the DSGE model, the proposal density needs to be tailored as a function of ψ. In turn, it is more difficult to ensure that the resulting Markov chain properly mixes and converges to its ergodic distribution at a sufficiently fast rate. Our framework de-couples the estimation of the DSGE model and the analysis of the auxiliary regressions. If needed, additional non-core variables can be easily analyzed without having to re-estimate the DSGE model. We view this as a useful feature in real-time applications.

Third, in addition to ignoring the information in the z_t's about the latent states we take one more short-cut. Rather than using estimates of $s_{t|t}$ that depend on θ, we condition on the posterior mean of θ in our construction of $\hat{s}_{t|t}$. As a consequence our posterior draws of DSGE and auxiliary model parameters are uncorrelated and we potentially understate the posterior uncertainty about ψ. However, in practice we found that there are few gains to using a more elaborate sampling procedure.

We proceed by re-writing (22) in quasi-differenced form as

$$\begin{aligned} z_1 &= \alpha_0 + \hat{s}'_{1|1}\alpha_1 + \xi_1 \\ z_t &= \rho z_{t-1} + \alpha_0(1-\rho) + [\hat{s}'_{t|t} - \hat{s}'_{t-1|t-1}\rho]\alpha_1 + \eta_t, \quad t = 2, \ldots, T. \end{aligned} \tag{23}$$

Instead of linking the distribution of ξ_1 to the parameters ρ and σ^2_η we assume that $\xi_1 \sim \mathcal{N}(0, \tau^2)$ and discuss the choice of τ below. A particular advantage of the Bayesian framework is that we can use the DSGE model to derive a prior distribution for the α's for variables z_t that are conceptually related to variables that appear in the DSGE model. Let $\alpha = [\alpha_0, \alpha'_1]'$. Our prior takes the form

$$\alpha \sim \mathcal{N}(\mu_{\alpha,0}, V_{\alpha,0}), \quad \rho \sim \mathcal{U}(-1,1), \quad \sigma_\eta \sim \mathcal{IG}(\nu, \tau), \tag{24}$$

where $\mathcal{N}(\mu, V)$ denotes a normal distribution with mean μ and covariance matrix V, $\mathcal{U}(a,b)$ is a uniform distribution on the interval (a,b), and $\mathcal{IG}(\nu, s)$ signifies an Inverse Gamma distribution with density $p_{\mathcal{IG}}(\sigma|\nu, s) \propto \sigma^{-(\nu+1)}e^{-\nu s^2/2\sigma^2}$. To avoid a proliferation of hyperparameters we use the same τ to characterize the standard deviation of ξ_1 and the prior for σ_η.

We choose the prior mean $\mu_{\alpha,0}$ based on the DSGE model implied factor loadings for a model variable, say $z^\dagger_t$, that is conceptually similar to z_t. For concreteness, suppose that z_t

corresponds to PCE inflation. Since there is only one-type of final good, our DSGE model does not distinguish between, say, the GDP deflator and a price index of consumption expenditures. Hence, a natural candidate for $z_t^\dagger$ is final good inflation. Let $I\!E^D_\theta[\cdot]$ denote an expectation taken under the probability distribution generated by the DSGE model, conditional on the parameter vector θ. We construct $\mu_{\alpha,0}$ by a population regression of the form

$$\mu_{\alpha,0} = \left(I\!E^D_\theta[\tilde{s}_t \tilde{s}_t'] \right)^{-1} I\!E^D_\theta[\tilde{s}_t z_t^\dagger], \tag{25}$$

where $\tilde{s}_t = [1, s_t']'$ and θ is in practice replaced by its posterior mean $\hat{\theta}_T$. If $z_t^\dagger$ is among the observables, then this procedure recovers the corresponding rows of $A_0(\theta)$ and $A_1(\theta)$ in the measurement equation (19). Details on the choice of $z_t^\dagger$ are provided in the empirical section. Our prior covariance matrix is diagonal with the following elements

$$diag(V_{\alpha,0}) = \left[\lambda_0, \frac{\lambda_1}{\omega_1}, \dots, \frac{\lambda_1}{\omega_J} \right]. \tag{26}$$

Here λ_0 and λ_1 are hyperparameters that determine the degree of shrinkage for the intercept α_0 and the loadings α_1 of the state variables. We scale the diagonal elements of $V_{\alpha,0}$ by ω_j^{-1}, $j = 1, \dots, J$, where ω_j denotes the DSGE model's implied variance of the j'th element of $\hat{s}_{t|t}$ (evaluated at the posterior mean of θ).[6] Draws from the posterior distribution can be easily obtained with a Gibbs sampler described in Appendix A.

3.3 Forecasting

Suppose that the forecast origin coincides with the end of the estimation sample, denoted by T. Forecasts from the DSGE model are generated by sampling from the posterior predictive distribution of y_{T+h}. For each posterior draw $\theta^{(i)}$ we start from $\hat{\varsigma}_{T|T}(\theta^{(i)})$ and draw a random sequence $\{\epsilon^{(i)}_{T+1}, \dots, \epsilon^{(i)}_{T+h}\}$. We then iterate the state transition equation forward to construct

$$\begin{aligned} s^{(i)}_{T+h|T} &= \Phi_1(\theta^{(i)}) s^{(i)}_{T+h-1|T} + \Phi_\epsilon(\theta^{(i)}) \epsilon^{(i)}_{T+h}, \quad h = 1, \dots, H \\ \varsigma^{(i)}_{T+h|T} &= [s^{(i)\prime}_{T+h|T}, s^{(i)\prime}_{T+h-1|T} M_s'(\theta^{(i)})]'. \end{aligned} \tag{27}$$

Finally, we use the measurement equation to compute

$$y^{(i)}_{T+h|T} = A_0(\theta^{(i)}) + A_1(\theta^{(i)}) \varsigma^{(i)}_{T+h|T}. \tag{28}$$

[6] Instead of assuming that the elements of α are independent, one could use the inverse of the covariance matrix of $\hat{s}_{t|t}$ to construct a non-diagonal prior covariance matrix for α. To the extent that some of the elements of s_t are highly correlated such a prior will be highly non-informative in the corresponding directions of the α parameter space. We found this feature unattractive and decided to proceed with a diagonal $V_{\alpha,0}$.

The posterior mean forecast $\hat{y}_{T+h|T}$ is obtained by averaging the $y^{(i)}_{T+h|T}$'s.

A draw from the posterior predictive distribution of a non-core variable z_{T+h} is obtained as follows. Using the sequence $s^{(i)}_{T+1|T}, \ldots, s^{(i)}_{T+H|T}$ constructed in (27), we iterate the quasi-differenced version (23) of the auxiliary regression forward:

$$z^{(i)}_{T+h|T} = \rho^{(i)} z^{(i)}_{T+h-1} + \alpha^{(i)}_0 (1 - \rho^{(i)}) + [s^{(i)\prime}_{T+h|T} - s^{(i)\prime}_{T+h-1|T} \rho^{(i)}] \alpha^{(i)}_1 + \eta^{(i)}_{T+h},$$

where the superscript i for the parameters of (22) refers to the i'th draw from the posterior distribution of ψ and $\eta^{(i)}_{T+h}$ is a draw from a $\mathcal{N}(0, \sigma^{2(i)}_\eta)$. The point forecast $\hat{z}_{T+h|T}$ is obtained by averaging the $z^{(i)}_{T+h|T}$'s. While our draws from the posterior distribution of θ and ψ are independent, we maintain much of the correlation in the joint predictive distribution of y_{T+h} and z_{T+h}, because the i'th draw is computed from the same realization of the state vector, $s^{(i)}_{T+h|T}$.

4 Empirical Application

We use post-1983 U.S. data to recursively estimate the DSGE model and the auxiliary regression equations and to generate pseudo-out-of-sample forecasts. We begin with a description of our data set and the prior distribution for the DSGE model parameters. Second, we discuss the estimates of the DSGE model parameters and its forecast performance for the core variables. Third, we estimate the auxiliary regressions and examine their forecasts of PCE inflation, core PCE inflation, the unemployment rate, and housing starts. Finally, we explore multivariate aspects of the predictive distribution generated by our model. We report conditional forecast error statistics and illustrate the joint predictive distribution as well as the propagation of a monetary policy shock to the core and non-core variables.

4.1 Data and Priors

Seven series are included in the vector of core variables y_t that is used for the estimation of the DSGE model: the growth rates of output, consumption, investment, and nominal wages, as well as the levels of hours worked, inflation, and the nominal interest rate. These series are obtained from Haver Analytics (Haver mnemonics are in italics). Real output is computed by dividing the nominal series (*GDP*) by population 16 years and older (*LN16N*) as well as the chained-price GDP deflator (*JGDP*). Consumption is defined as nominal personal consumption expenditures (*C*) less consumption of durables (*CD*). We divide by

LN16N and deflate using *JGDP*. Investment is defined as *CD* plus nominal gross private domestic investment (*I*). It is similarly converted to real per-capita terms. We compute quarter-to-quarter growth rates as log difference of real per capita variables and multiply the growth rates by 100 to convert them into percentages.

Our measure of hours worked is computed by taking non-farm business sector hours of all persons (*LXNFH*), dividing it by *LN16N*, and then scaling to get mean quarterly average hours to about 257. We then take the log of the series multiplied by 100 so that all figures can be interpreted as percentage deviations from the mean. Nominal wages are computed by dividing total compensation of employees (*YCOMP*) by the product of *LN16N* and our measure of average hours. Inflation rates are defined as log differences of the GDP deflator and converted into percentages. The nominal interest rate corresponds to the average effective federal funds rate (*FFED*) over the quarter and is annualized.

Observations for the non-core variables are also obtained from Haver Analytics. We consider PCE-inflation, core PCE inflation, the unemployment rate, and housing starts as candidates for z_t in this paper. We extract quarterly data on the chain price index for personal consumption expenditures (*JC*) and personal consumption expenditures less food and energy (*JCXF*). Inflation rates are calculated as 100 times the log difference of the series. The unemployment rate measure is the civilian unemployment rate for ages 16 years and older (*LR*). Finally, housing starts are defined as millions of new privately owned housing units started (*HST*). We use quarterly averages of seasonally adjusted monthly data, converted to an annual rate.

Our choice of prior distribution for the DSGE model parameters follows DSSW and the specification of what is called a “standard” prior in Del Negro and Schorfheide (2008). The prior is summarized in the first four columns of Table 1. To make this paper self-contained we briefly review some of the details of the prior elicitation. Priors for parameters that affect the steady state relationships, e.g., the capital share α in the Cobb-Douglas production function or the capital depreciation rate are chosen to be commensurable with pre-sample (1955 to 1983) averages in U.S. data. Priors for the parameters of the exogenous shock processes are chosen such that the implied variance and persistence of the endogenous model variables is broadly consistent with the corresponding pre-sample moments. Our prior for the Calvo parameters that control the degree of nominal rigidity are fairly agnostic and span values that imply fairly flexible as well as fairly rigid prices and wages. Our prior for the central bank’s responses to inflation and output movements is roughly centered at Taylor’s (1993) values. The prior for the interest rate smoothing parameter ρ_R is almost uniform on the unit interval.

The 90% interval for the prior distribution on ν_l implies that the Frisch labor supply elasticity lies between 0.3 and 1.3, reflecting the micro-level estimates at the lower end, and the estimates of Kimball and Shapiro (2003) and Chang and Kim (2006) at the upper end. The density for the adjustment cost parameter S'' spans values that Christiano, Eichenbaum, and Evans (2005) find when matching DSGE and vector autoregression (VAR) impulse response functions. The density for the habit persistence parameter h is centered at 0.7, which is the value used by Boldrin, Christiano, and Fisher (2001). These authors find that $h = 0.7$ enhances the ability of a standard DSGE model to account for key asset market statistics. The density for a'' implies that in response to a 1% increase in the return to capital, utilization rates rise by 0.1 to 0.3%.

4.2 DSGE Model Estimation and Forecasting of Core Variables

The first step of our empirical analysis is to estimate the DSGE model. While we estimate the model recursively, starting with the sample 1984:I to 2000:IV and ending with the sample 1984:I to 2007:III, we will focus our discussion of the parameter estimates on the final estimation sample. Summary statistics for the posterior distribution (means and 90% probability intervals) are provided in Table 1. For long horizon forecasts, the most important parameters are γ, π_*, and β. Our estimate of the average technology growth rate implies that output, consumption, and investment grow at an annualized rate of 1.6%. According to our estimates of π_* and β the target inflation rate is 2.9% and the long-run nominal interest rate is 5.5%. The cross-equation restrictions of our model generate a nominal wage growth of about 4.5%.

Our policy rule estimates imply a strong response of the central bank to inflation $\hat{\psi}_1 = 3.05$ and a tempered reaction to deviations of output from its long-run growth path $\hat{\psi}_2 = 0.06$. As discussed in Del Negro and Schorfheide (2008), estimates of wage and price stickiness based on aggregate price and wage inflation data tend to be somewhat fragile. We obtain $\hat{\zeta}_p = 0.66$ and $\hat{\zeta}_w = 0.25$, which means that wages are nearly flexible and the price stickiness is moderate. According to the estimated Calvo parameter, firms re-optimize their prices every three quarters.

The technology growth shocks have very little serial correlation and the estimated innovation standard deviation is about 0.6%. These estimates are consistent with direct calculations based on Solow residuals. At an annualized rate, the monetary policy shock has a standard deviation of 56 basis points. Both the government spending shock g_t and the labor supply shock ϕ_t have estimated autocorrelations near unity. The labor supply shock captures much of the persistence in the hours series.

We proceed by plotting estimates of the exogenous shocks in Figure 1. These shocks are included in the vector $s_t = M\varsigma_t$ that is used as regressor in the auxiliary model (22). Formally, we depict filtered latent variables, $\hat{s}_{j,t|t}$, conditional on the posterior mean $\hat{\theta}_T$ for the period 1984:I to 2007:III. In line with the parameter estimates reported in Table 1, the filtered technology growth process appears essentially *iid*. The processes g_t and ϕ_t exhibit long-lived deviations from zero and in part capture low frequency movements of exogenous demand components and hours worked, respectively. μ_t is the investment-specific technology shock. Its low frequency movements capture trend differentials in output, consumption, and investment.

At this point a comparison between our estimates of the latent shock processes and the estimates reported by BG is instructive. By construction, our filtered state variables $\hat{s}_{t|t}$ are moving averages of the observables y_t. In contrast, BG's estimates of the latent states are functions not just of y_t (in our notation), but also of all the other observables included in their measurement equations, namely numerous measures of inflation as well as 25 principle components constructed from about 70 macroeconomic time series. Due to differences in model specification and data definitions, it is difficult to compare our estimates of the latent states and those reported by BG directly. However, BG overlay smoothed states obtained from the direct estimation of their DSGE model with estimates obtained from their DSGE-DFM. The main difference between the estimated DSGE and DSGE-DFM states is that some of the latter, namely productivity, preferences, and government spending, are a lot smoother. The likely reason is that the DSGE-DFM measurement equations for the seven core variables contain autoregressive measurement errors, which absorb some of the low frequency movements in these series.

Table 2 summarizes pseudo-out-of-sample root-mean-squared error (RMSE) statistics for the seven core variables that are used to estimate the DSGE model: the growth rates of output, consumption, investment, and nominal wages, as well as log hours worked, GDP deflator inflation, and the federal funds rate. We report RMSEs for horizons $h = 1$, $h = 2$, $h = 4$, and $h = 12$ and compare the DSGE model forecasts to those from an AR(1) model, which is recursively estimated by OLS.[7] h-step ahead growth (inflation) rate forecasts refer

[7] The h-step forecast is generated by iterating one-step ahead predictions forward, ignoring parameter uncertainty: $\hat{y}_{i,T+h|T} = \hat{\beta}_{0,OLS} + \hat{\beta}_{1,OLS}\hat{y}_{i,T+h-1|T}$, where the OLS estimators are obtained from the

to percentage changes between period $T+h-1$ and $T+h$. Boldface entries indicate that the DSGE model attains a RMSE that is lower than that of the AR(1) model. We used the Harvey, Leybourne, and Newbold (1998) version of the Diebold-Mariano (1995) test for equal forecast accuracy of the DSGE and the AR(1) model, employing a quadratic loss function. Due to the fairly short forecast period, most of the loss differentials are insignificant.

The RMSE for one-quarter-ahead forecasts of output and consumption obtained from the estimated DSGE model is only slightly larger than the RMSE associated with the AR(1) forecasts. The DSGE model generates a lower RMSE for investment and hours worked forecasts. RMSEs for inflation rates rates are essentially identical across the two models. The AR(1) model performs better than the DSGE model in forecasting nominal wage growth and interest rates. The accuracy of long-run forecasts is sensitive to mean growth estimates, which are restricted to be equal for output, consumption, and investment. Moreover, the DSGE model implies that nominal wage growth equals output plus inflation growth in the long-run.

In Table 3 we are comparing the pseudo-out-of-sample RMSEs obtained with our estimated DSGE model to those reported in three other studies, namely (i) DSSW, (ii) Edge, Kiley, and Laforte (EKL, 2008), and (iii) Smets and Wouters (2007). Since all studies differ with respect to the forecast period, we report sample standard deviations over the respective forecast periods, computed from our data set. Unlike the other three studies, EKL use real time data and report mean absolute errors instead of RMSEs. Overall, the RMSEs reported in DSSW are slightly worse than those in the other three studies. This might be due to the fact that DSSW use a rolling window of 120 observations to estimate their DSGE model and start forecasting in the mid 1980s, whereas the other papers let the estimation sample increase and start forecasting in the 1990s. Only EKL are able to attain an RMSE for output growth that is lower than the sample standard deviation. The RMSEs for the inflation forecasts range from 0.22 to 0.27 and are very similar across studies. They are only slightly larger than the sample standard deviations. Finally, the interest rate RMSEs are substantially lower than the sample standard deviations, because the forecasts are able to exploit the high persistence of the interest rate series.

regression $y_{i,t} = \beta_0 + \beta_1 y_{i,t-1} + u_{i,t}$.

4.3 Forecasting Non-Core Variables with Auxiliary Regressions

We now turn to the estimation of the auxiliary regressions for PCE inflation, core PCE inflation, the unemployment rate, and housing starts. The following elements are included in the vector s_t that appears as regressor in (22):

$$s_t = M\varsigma_t = [c_t, i_t, \bar{k}_t, R_t, w_t, a_t, \phi_t, \mu_t, b_t, g_t, \lambda_{f,t}]'.$$

To construct a prior mean for α_1, we link each z_t with a conceptually related DSGE model variable $z_t^\dagger$ and use (25). More specifically, we link the two measures of PCE inflation to the final good inflation π_t, the unemployment rate to a scaled version of log hours worked L_t, and housing starts to scaled percentage deviations i_t of investment from its trend path, see Table 4. Our DSGE model has only a single final good, which is domestically produced and used for consumption and investment. Hence, using identical measurement equations for inflation in consumption expenditures and GDP seems reasonable. Linking the unemployment rate with hours worked can be justified by the observation that most of the variation of hours worked over the business cycle is due to changes in employment rather than variation along the intensive margin. Finally, housing starts can be viewed as a measure of investment, namely investment in residential structures. Since the housing starts series has no apparent trend, we link it to investment deviations from trend.

The four panels of Figure 2 depict the sample paths of the non-core variables z_t and the related DSGE model variables $z_t^\dagger$. The GDP deflator and hours worked are directly observable, while the investment series i_t is latent and obtained from $\hat{s}_{t|t}$. The inflation measures are highly correlated. PCE inflation is more volatile and core PCE inflation is less volatile than GDP deflator inflation. In the bottom left panel we re-scale and re-center log hours such that it is commensurable with the unemployment rate. These two series are also highly correlated. The bottom right panel shows that the DSGE model implied investment series is somewhat smoother than the housing starts series. However, except for the period from 2000 to 2002 the low frequency movements of the two series are at least qualitatively similar.

To proceed with the Bayesian estimation of (23) we have to specify the hyperparameters. In our framework τ can be interpreted as the prior standard deviation of the idiosyncratic error ξ_1. We set τ equal to 0.12 (PCE inflation), 0.11 (core PCE inflation), 0.40 (unemployment rate), and 0.10 (housing starts). These values imply that the prior variance of ξ_1 is about 15% to 20% of the sample variance of z_1. We set the degrees of freedom parameter ν of the inverted gamma prior for σ_η equal to 2, restrict $\lambda_0 = \lambda_1 = \lambda$, and consider three values: 1.00, 0.10, and 1E-5. The value 1E-5 corresponds to a dogmatic prior under which posterior estimate and prior mean essentially coincide. As we increase λ, we allow the factor loading coefficients α to differ from the prior mean.

The estimates of the auxiliary regressions are summarized in Table 5. Rather than providing numerical values for the entire α vector, we focus on the persistence and the standard deviation of the innovation to the idiosyncratic component. By construction, $\hat{s}'_{t|t}\mu_{\alpha_1,0}$, where $\mu_{\alpha_1,0}$ is the prior mean of α_1, reproduces the time paths of the GDP deflator inflation, log hours worked, and investment deviations from trend, respectively. Thus, for 1E-5 the idiosyncratic error term ξ_t essentially picks up the discrepancies between non-core variables and the related DSGE model variables depicted in Figure 2. For the two inflation series the estimate of σ_η falls as we increase the hyperparameter. The larger λ the more of the variation in the variable is explained by $\hat{s}'_{t|t}\hat{\alpha}_1$, where $\hat{\alpha}_1$ is the posterior mean of α_1. For instance, the variability of core PCE inflation captured by the factors is 5 times as large as the variability due to the idiosyncratic disturbance ξ_t if the λ is equal to one. This factor drops to 1.4 if the prior is tightened. For PCE inflation the idiosyncratic disturbance is virtually serially uncorrelated, whereas for core PCE inflation the serial correlation ranges from 0.2 ($\lambda = 1$) to 0.5 ($\lambda =$ 1E-5).

For unemployment, setting $\lambda =$ 1E-5 implies that the prior and posterior means of the factor loadings α are essentially identical. Unemployment loads on c_t, i_t, $\bar{k}_t$, μ_t, and g_t. The intuition is that output in our model can be obtained from consumption, investment, and government spending (see Equation (15)) and hours worked can be determined from the production function as a function of output and capital (see Equation (5)). If the hyperparameter is raised to 0.1 or 1.0 then unemployment also loads on the interest rate, wages, and the shocks a_t and b_t. However, in general we find it difficult to interpret the estimates of particular elements of α_1, because the some of the variables contained in the vector s_t are endogenous equilibrium objects that in turn respond to the exogenous state variables. Hence, we will focus on the estimate of $\hat{s}'_{t|t}\alpha_1$ and the response of z_t to structural shocks below. The most striking feature of the unemployment estimates is the high persistence of ξ_t, with ρ_ξ estimates around 0.98.

For housing starts, the measurement error process is slightly less persistent than for unemployment, but the signal-to-noise ratio is generally low, which is not surprising in view of the fairly large discrepancy between housing starts and i_t shown in the bottom right panel of Figure 2. Unlike for the other three non-core series, the lowest signal-to-noise ratio for housing starts is obtained for $\lambda = 1$. An increase of λ from 1E-5 to 1 decreases the variability of $\hat{s}'_{t|t}\hat{\alpha}_1$ by more than the variability of the measurement error process, as is evident from the bottom right panel of Figure 3.

Figure 3 displays the time path of $\hat{\alpha}_0 + \hat{s}'_{t|t}\hat{\alpha}_1$ for different choices of the hyperparameter. Consider the two inflation series. For $\lambda = 1\text{E-}5$ the factor predicted path for the two inflation rates is essentially identical and reproduces the GDP deflator inflation. As the λ is increased to one they more closely follow the two PCE inflation measures, which is consistent with the estimates of ρ and σ_η reported in Table 5. The predicted paths for the unemployment rate behave markedly different. If we set $\lambda = 1$, then the predicted path resembles the actual path fairly closely, with the exception of the end of the sample. Hence, the implied ξ_t series stays close to zero until about 2002 and then drops to about -2% between 2002 and 2006. As we decrease λ to 1E-5, the predicted path shifts downward. The estimate of ξ_1 is roughly 2% and ξ_t follows approximately a random walk process subsequently that captures the gap between the path predicted with the factors and the actual unemployment series.

The last column of Table 5 contains log marginal likelihood values $\ln p_\lambda(Z^T)$ for the four auxiliary regression models as a function of the hyperparameter λ. These values can be used for a data-driven hyperparameter choice that trades off in-sample fit against complexity of the regression model.[8] According to the marginal likelihoods, the preferred choice for λ is 0.1 for core PCE inflation and the unemployment rate and 1E-5 for PCE inflation and housing starts. The log marginal data density can also be interpreted as a one-step-ahead predictive score:

$$\ln p_\lambda(Z^T) = \sum_{t=0}^{T-1} \int p(z_{t+1}|\psi, Z^t) p_\lambda(\psi|Z^t) d\psi. \tag{29}$$

Thus, we would expect the λ rankings obtained from one-step-ahead pseudo-out-of-sample forecast error statistics to be comparable to the rankings obtained from the marginal likelihoods.

Forecast error statistics for the non-modelled variables are provided in Table 6. We compare RMSEs of the forecasts generated with our auxiliary models to two alternative models. First, as in Section 4.2 we consider an AR(1) model for z_t that is estimated by OLS and from which we generate h-step forecasts by iterating one-step ahead predictions forward. Second, we consider multi-step least squares regressions of the form

$$z_t = \beta_0 + y'_{t-h}\beta_1 + z_{t-h}\beta_2 + u_t, \tag{30}$$

[8] A detailed discussion of hyperparameter selection based on marginal likelihoods can be found, for instance, in DSSW.

estimated for horizons $h = 1$, $h = 2$, $h = 4$, and $h = 12$. Recall that the filtered states $\hat{s}_{t|t}$ are essentially moving averages of y_t and its lags. Hence, both (23) and (30) generate predictions of z_{t+h} as a function of z_t as well as y_t and its lags. However, the restrictions imposed on the parameters of the implied prediction functions are very different. While our least squares estimation of (30) leaves the coefficient vector β_1 essentially unrestricted and excludes additional lags of y_t, the auxiliary regression model (23) tilts the estimates of α_1 toward loadings derived from the DSGE model and additional lags of y_t implicitly enter the prediction through the filtered state vector.

Over short horizons, our auxiliary regression models attain a lower RMSE than the AR(1) benchmark for PCE inflation, the unemployment rate, and housing starts. The improvements of the unemployment forecasts are significant. For one-step-ahead forecasts, the preferred choice of λ is 1E-5. For PCE inflation and housing starts the value of λ that yields the highest marginal likelihood also generates the lowest RMSE. For the unemployment rate the marginal likelihoods for λ equals 0.1 and 1E-5 are very similar and so are the RMSE statistics. The only discrepancy between RMSE and marginal likelihood ranking arises for core PCE inflation. We conjecture that the different rankings could be in part due to the persistent deviations of core PCE inflation from $\hat{s}'_{t|t}\hat{\alpha}_1$ at the beginning of the sample, as evident from the top right panel of Figure 3. According to (29) predictive accuracy at the beginning of the sample affects the marginal likelihood, but it does not enter our RMSE statistics, which are computed from 2001 onward. Over a longer horizon, core PCE and unemployment forecasts from our auxiliary regressions dominate the AR(1) forecasts, whereas the PCE inflation and housing starts forecasts are slightly less precise. Except for short to medium term core PCE inflation forecasts, our auxiliary regressions with λ = 1E-5 are slightly better than the forecasts obtained from the simple predictive regression (30).

4.4 Multivariate Considerations

So far the analysis has focused on univariate measures of forecast accuracy. A conservative interpretation of our findings and those reported elsewhere, e.g., Adolfson *et al.* (2005, 2007) and Edge, Kiley, and Laforte (2008), is that by and large the univariate forecast performance of DSGE models is not worse than that of competitive benchmark models, such as simple AR(1) specifications or more sophisticated Bayesian VARs. The key advantage of DSGE models and the reason that central banks are considering them for projections and policy analysis, is that these models use modern macroeconomic theory to explain and predict comovements of aggregate time series over the business cycle. Historical observations can be decomposed into the contributions of the underlying exogenous disturbances, such as technology, preference, government spending, or monetary policy shocks. Future paths of the endogenous variables can be constructed conditional on particular realizations of the monetary policy shocks that reflect potential future nominal interest rate paths. While it is difficult to quantify some of these desirable attributes of DSGE model forecasts and trade them off against forecast accuracy in a RMSE sense, we will focus on three multivariate aspects. First, we conduct posterior predictive checks for the correlation between core and non-core variables captured by our framework. Second, we present impulse response functions to a monetary policy shock and document how the shock transmits to the non-core variables through our auxiliary regression equations. Third, we examine some features of the predictive density that our empirical model generates for the core and non-core variables.

Posterior predictive checks for correlations between non-core and core variables are summarized in Table 7 for λ = 1E-5, which is the value of λ that leads to the lowest one-step-ahead forecast RMSE. Using the posterior draws for DSGE and auxiliary model parameters we simulate a trajectory of 100 z_t and y_t observations and compute sample correlations of interest. The posterior predictive distribution of these sample correlations is then summarized by 90% credible intervals. Moreover, we report sample correlations computed from U.S. data. The empirical model captures the correlations between non-core and core variables well, if the actual sample correlations do not lie too far in the tails of the corresponding posterior predictive distribution. With the exception of the correlations between output growth and the unemployment rate all of the correlations computed from U.S. data lie inside of the corresponding 90% credible sets.

An important aspect of monetary policy making is assessing the effect of changes in the federal funds rate. In the DSGE model we represent these changes – unanticipated deviations from the policy rule – as monetary policy shocks. An attractive feature of our framework is that it generates a link between the structural shocks that drive the DSGE model and other non-modeled variables through the auxiliary regressions. We can compute impulse response functions of z_t to a monetary policy shock as follows:

$$\frac{\partial z_{t+h}}{\partial \epsilon_{R,t}} = \frac{\partial s'_{t+h}}{\partial \epsilon_{R,t}} \alpha_1,$$

where $\partial s'_{t+h}/\partial \epsilon_{R,t}$ is obtained from the DSGE model.

In Figure 4 we plot impulse responses of the four non-core variables (bottom panels) and the four related DSGE model related variables (top panels: output, inflation, investment, and hours) to a one-standard deviation monetary policy shock. The one standard deviation increase to the monetary policy shock translates into a 40 basis point increase in the funds rate, measured at an annual rate. The estimated DSGE model predicts that output and hours worked drop by 10 basis points in the first quarter and returns to its trend path after seven quarters. Investment is more volatile and drops by about 19 basis points. Quarter-to-quarter inflation falls by 10 basis points and returns to its steady state within two years. Regardless of the choice of hyperparameter, the PCE inflation responses closely resemble the GDP deflator inflation responses both qualitatively and quantitatively. The core PCE inflation, unemployment, and housing starts responses are more sensitive to the choice of hyperparameter. If λ is equal to 1E-5 and we force the factor loadings to match those of hours worked, the unemployment rises by about 3.5 basis points one period after impact. As we relax the hyperparameter, which worsens the RMSE of the unemployment forecast, the initial effect of the monetary policy shock on unemployment is dampened. Likewise, the core PCE response drops from 10 basis points to about 4 basis points. The annualized number of housing starts drops by about 6,000 units for λ = 1E-5 and by 22,000 units if $\lambda = 1$. Unlike for core PCE inflation, housing starts respond more strongly to a monetary policy shock if the restrictions on the factor loadings are relaxed.

Our empirical model generates a joint density forecast for the core and non-core variables, which reflects uncertainty about both parameters and future realizations of shocks. A number of different methods exist to evaluate multivariate predictive densities. To assess whether the probability density forecasts are well calibrated, that is, are consistent with empirical frequencies, one can construct the multivariate analog of a probability integral transform of the actual observations and test whether these transforms are uniformly distributed and serially uncorrelated. A formalization of this idea is provided in Diebold, Hahn, and Tay (1999).

We will subsequently focus on log predictive scores (Good, 1952). To fix ideas, consider the following simple example. Let $x_t = [x_{1,t}, x_{2,t}]'$ be a 2×1 vector and consider the following two forecast models

$$\mathcal{M}_1: \quad x_t \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right), \quad \mathcal{M}_2: \quad x_t \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix} \right).$$

Under a quadratic loss function the two models deliver identical univariate forecasts for each linear combination of the elements of x_t. Nonetheless, the predictive distributions are distinguishable. Let Σ_i be the covariance matrix of the predictive distribution associated with model $\mathcal{M}_i$. The log predictive score is defined as the log predictive density evaluated at a sequence of realizations of x_t, $t = 1, \ldots, T$:

$$LPSC(\mathcal{M}_i) = -\frac{T}{2} ln(2\pi) - \frac{T}{2} \ln|\Sigma_i| - \frac{1}{2} \sum_t^T x_t' \Sigma_i^{-1} x_t.$$

Roughly speaking, if the actual x_t was deemed unlikely by $\mathcal{M}_i$ and falls in a low density region (e.g., the tails) of the predictive distribution, then the score is low. Let Σ_{11}, Σ_{12}, and Σ_{22} denote partitions of Σ that conform with the partitions of x. If we factorize the joint predictive density of x_t into a marginal and a conditional density, we can rewrite the predictive score as

$$\begin{aligned} LPSC(\mathcal{M}_i) &= -\frac{T}{2} ln(2\pi) - \frac{T}{2} \ln|\Sigma_{i,11}| - \frac{1}{2\Sigma_{i,11}} \sum_{t=1}^T x_{1,t}^2 \\ &\quad -\frac{T}{2} \ln|\Sigma_{i,22|11}| - \frac{1}{2\Sigma_{i,22|11}} \sum_{t=1}^T \left(x_{2,t} - \Sigma_{i,21}\Sigma_{i,11}^{-1} x_{1,t} \right)^2, \end{aligned} \tag{31}$$

where

$$\Sigma_{i,22|11} = \Sigma_{22} - \Sigma_{i,21}\Sigma_{i,11}^{-1}\Sigma_{i,12}.$$

We can express the difference between log predictive scores for models $\mathcal{M}_1$ and $\mathcal{M}_2$ as

$$LPSC(\mathcal{M}_1) - LPSC(\mathcal{M}_2) = \frac{T}{2} \ln|1-\rho^2| - \frac{1}{2} \sum_{t=1}^T x_{2,t}^2 + \frac{1}{2(1-\rho^2)} \sum_{t=1}^T (x_{2,t} - \rho x_{1,t})^2.$$

Here the contribution of the marginal distribution of $x_{1,t}$ to the predictive scores cancels out, because it is the same for $\mathcal{M}_1$ and $\mathcal{M}_2$. It is straightforward to verify that for large T the predictive score will be negative if in fact the x_t's are generated from $\mathcal{M}_2$. In fact, the log score differential has similar properties as a log likelihood ratio and is widely used in the prequential theory discussed in Dawid (1992). Moreover, notice that $\frac{1}{T}\sum_{t=1}^T (x_{2,t} - \rho x_{1,t})^2$ can be interpreted as the mean-squared-error of a forecast of $x_{2,t}$ conditional on the realization of $x_{1,t}$. If $x_{1,t}$ and $x_{2,t}$ have non-zero correlation, the conditioning improves the accuracy of the $x_{2,t}$ forecast. We will exploit this insight below.

Figure 5 depicts bivariate scatter plots generated from the joint predictive distribution of core and non-core variables. The predictive distribution captures both parameter uncertainty as well as shock uncertainty. We focus on one-step-ahead predictions for 2001:IV and 2006:III. We use filled circles to indicate the actual values (small, light blue), the unconditional mean predictions (medium, yellow), and the conditional means of ouput growth, PCE inflation, and unemployment given the actual realization of the nominal interest rate. We approximate the predictive distributions by student t distributions with mean μ, variance

Σ, and ν degrees of freedom.[9] We replace μ and Σ by the sample means and covariance matrices computed from the draws from the predictive distributions. Regardless of the degrees of freedom ν the conditional mean of x_2 given the realization of x_1 is given by:

$$\hat{x}_{2|1} = \mu_2 + \Sigma_{21}\Sigma_{11}^{-1}(x_1 - \mu_1). \tag{32}$$

In Figure 5 the nominal interest rate plays the role of the conditioning variable x_1.

First, consider the predictive distribution for output growth and interest rates in 2001:IV. The predictive distribution is centered at an interest rate of 4% and output growth of about 0%. The actual interest rate turned out to be 2% and output grew at about 20 basis points over the quarter. Since the predictive distribution exhibits a negative correlation between interest rates and output growth, conditioning on the actual realization of the interest rate leads to an upward revision of the output growth forecast to about 30 basis points. In 2006:III the actual interest rate exceeds the mean of the predictive distribution, and hence conditioning reduces the output growth forecast.

PCE inflation ($\lambda = $ 1E-5) and the interest rate are strongly positively correlated and the conditioning leads to a downward revision of the inflation forecast in 2001:IV and an upward revision in 2006:III. Our estimation procedure is set up in a way that leaves the coefficients of the auxiliary regression uncorrelated with the DSGE model parameters. Hence, all the correlation in the predictive distribution is generated by shock uncertainty and the fact that the auxiliary regression links the non-core variable to the DSGE model states. Finally, we turn to the joint predictive distribution of unemployment ($\lambda = $ 1E-5) and interest rates. Since the idiosyncratic shock ξ_t plays an important role for the unemployment dynamics according to our estimates and it is assumed to be independent of the DSGE model shocks, the predictive distribution exhibits very little correlation. In this case, conditioning hardly affects the unemployment forecast.

Figure 5 focuses on two particular time periods. More generally, if the family of t-distributions provides a good approximation to the predictive distribution, and our model captures the comovements between interest rates and the other variables, then we should be able to reduce the RMSE of the output, unemployment, and inflation forecasts by conditioning on the interest rate. Tables 8 and 9 provide RMSE ratios of conditional and unconditional forecasts. To put these numbers into perspective we also report the ratio of the conditional versus the unconditional variance computed from a t distribution with $\nu = 5$

[9] Under this parameterization, the density of a m-variate t distribution is proportional to $[1 + (\nu - 2)(x - \mu)'\Sigma^{-1}(x - \mu)]^{-(\nu+m)/2}$.

degrees of freedom and a normal distribution ($\nu = \infty$). Using the subscript j to index the pseudo-out-of-sample forecasts, we define the average theoretical RMSE ratio:

$$\mathcal{R}(\nu) = \frac{1}{J}\sum_{j=1}^{J}\sqrt{\frac{\frac{\nu-2}{\nu}\left(1+\frac{1}{\nu-2}(x_{1,j}-\mu_{1,j})'\Sigma_{11,j}^{-1}(x_{1,j}-\mu_{1,j})\right)(\Sigma_{22,j}-\Sigma_{21,j}\Sigma_{11,j}^{-1}\Sigma_{12,j})}{\Sigma_{22,j}}}. \tag{33}$$

The results obtained when conditioning on the interest rate, reported in Table 8, are somewhat disappointing. Although, except for housing starts, the bivariate correlations between the interest rate and the other variables are non-zero and would imply a potential RMSE reduction between 1% and 12%, the RMSE obtained from the conditional forecasts exceeds that from the unconditional forecasts.[10] If we condition on the realization of the GDP deflator inflation (Table 9), then the results improve and we observe a RMSE reduction at least for output growth and PCE inflation, although not as large as predicted by $\mathcal{R}(\nu)$.

These last results have to be interpreted carefully. It is important to keep in mind that we are examining particular dimensions of the joint predictive density generated by our model. While in the past, researchers have reported log predictive scores and predictive likelihood ratios for DSGE model predictions, these summary statistics make it difficult to disentangle in which dimensions the predictive distributions are well calibrated. We decided to focus on bivariate distributions, trying to assess whether the DSGE model and the auxiliary regressions capture the comovements of, say, interest rates with output growth, inflation, and unemployment. Our results were mixed: bivariate distributions that involved the interest rate were not well calibrated in view of the actual realizations; bivariate distributions that involved the GDP deflator were somewhat more successful capturing the uncertainty about future pairwise realizations. An examination of the sequences of predictive densities and realizations – a few of them were displayed in Figure 5 – suggested to us that the high RMSEs of the conditional forecasts were often caused by a small number of outliers, that is, actual observations that fall far into the tails of the predictive distribution. This suggests that more elaborate distributions for the structural DSGE model shocks might provide a remedy.

[10] 2001:IV and 2006:III are not representative, since conditioning in these periods leads to a reduction of the forecast error.

5 Conclusion

This paper has developed a framework to generate DSGE model-based forecasts for economic variables that are not explicitly modelled but that are of interest to the forecaster. Our framework can be viewed as a simplified version of the DSGE model based factor model proposed by BG. We first estimate the DSGE model on a set of core variables, extract the latent state variables, and then estimate auxiliary regressions that relate non-modelled variables to the model-implied state variables. We compare the forecast performance of our model with that of a collection of AR(1) models based on pseudo-out-of-sample RMSEs. While our approach does not lead to a dramatic reduction in the forecast errors, the forecasts are by and large competitive with those of the statistical benchmark model. We also examined bivariate predictive distributions generated from our empirical model. Our framework inherits the two key advantages of DSGE model based forecasting: it delivers an interpretation of the predicted trajectories in light of modern macroeconomic theory and it enables the forecaster to conduct a coherent policy analysis.

References

Adolfson, Malin, Stefan Laseen, Jesper Linde, and Mattias Villani (2005): "Forecasting Performance of an Open Economy Dynamic Stochastic General Equilibrium Model," *Sveriges Riksbank Working Paper Series*, **190**.

Adolfson, Malin, Michael K. Andersson, Jesper Linde, Mattias Villani, and Anders Vredin (2007): "Modern Forecasting Models in Action: Improving Macroeconomic Analyses at Central Banks," *Sveriges Riksbank Research Paper Series*, **23**.

Adolfson, Malin, Stefan Laseen, Jesper Linde, and Mattias Villani (2008): "Evaluating an Estimated New Keynesian Small Open Economy Model," *Journal of Economic Dynamics & Control*, **32**, 2690-2721.

An, Sungbae and Frank Schorfheide (2007): "Bayesian Analysis of DSGE Models," *Econometric Reviews*, **26**, 113-172.

Barndorff-Nielsen, O. and G. Schou (1973): "On the Parameterization of Autoregressive Models by Partial Autocorrelations," *Journal of Multivariate Analysis*, **3**, 408-419.

Boivin, Jean and Marc Giannoni (2006): "DSGE Models in a Data-Rich Environment," *Manuscript*, HEC Montreal and Columbia University.

Boldrin, Michele, Lawrence Christiano, and Jonas Fisher (2001): "Habit Persistence, Asset Returns, and the Business Cycle," *American Economic Review*, **91**, 149-166.

Calvo, Guillermo (1983): "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, **12**, 383-398.

Chang, Yongsung and Sun-Bin Kim (2006): "From Individual to Aggregate Labor Supply: A Quantitative Analysis based on a Heterogeneous Agent Macroeconomy," *International Economic Review*, **47**, 1-27.

Chib, Siddartha (1995): "Marginal Likelihood from the Gibbs Output," *Journal of the American Statistical Association*, **90**, 1313-1321.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, **113**, 1-45.

Coenen, Günter, Peter McAdam, and Roland Straub (2008): "Tax Reform and Labour-Market Performance in the Euro Area: A Simulation-Based Analysis Using the New Area-Wide Model," *Journal of Economic Dynamics & Control*, **32**, 2543-2583.

Dawid, A.P. (1992): "Prequential Analysis, Stochastic Complexity and Bayesian Inference." In: J.M. Bernado, J.O. Berger, A.P. Dawid and A.F.M. Smith (Eds) *Bayesian Statistics 4*, 109-125.

Del Negro, Marco and Frank Schorfheide (2008): "Forming Priors for DSGE Models (and How it Affects the Assessment of Nominal Rigidities)," *Journal of Monetary Economics*, **55**, 1191-1208.

Del Negro, Marco, Frank Schorfheide, Frank Smets, and Rafael Wouters (2007): "On the Fit of New Keynesian Models," *Journal of Business and Economic Statistics*, **25**, 123-162.

Diebold, Francis and Roberto Mariano (1995): "Comparing Predictive Accuracy," *Journal of Business & Economic Statistics*, **13**, 253-263.

Diebold, Francis, Jinyong Hahn, and Anthony Tay (1999): "Multivariate Density Forecast Evaluation and Calibration in Financial Risk Management: High Frequency Returns on Foreign Exchange," *Review of Economic and Statistics*, **81**, 661-673.

Edge, Rochelle, Michael Kiley, and Jean-Phillipe Laforte (2008): "A Comparison of Forecast Performance Between Federal Reserve Staff Forecasts, Simple Reduced-Form Models, and a DSGE Model," *Manuscript*, Federal Reserve Board of Governors.

Giannone, Domenico, Francesca Monti, and Lucrezia Reichlin (2008): "Incorporating Conjunctural Analysis in Structural Macroeconomic Models," *Manuscript*, ECARES and London Business School.

Good, I.J. (1952): "Rational Decisions" *Journal of the Royal Statistical Society B*, **14**, 107-114.

Harvey, David, Stephen Leybourne, and Paul Newbold (1998): "Tests for Forecast Encompassing," *Journal of Business & Economic Statistics*, **16**, 254-259.

Kimball, Miles and Matthew Shapiro (2003): "Labor Supply: Are the Income and Substitution Effects Both Large or Both Small?" *Manuscript*, University of Michigan.

Monti, Francesca (2008): "Forecast with Judgement and Models," *National Bank of Belgium Working Paper*, **153**, December.

Sims, Christopher (2002): "Solving Rational Expectations Models," *Computational Economics*, **20**, 1-20.

Smets, Frank and Rafael Wouters (2003): "An Estimated Stochastic Dynamic General Equilibrium Model for the Euro Area," *Journal of the European Economic Association*, **1**, 1123-1175.

Smets, Frank and Rafael Wouters (2004): "Forecasting with a Bayesian DSGE Model: An Application to the Euro Area," *Journal of Common Market Studies*, **42**, 841-67.

Smets, Frank and Rafael Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, **97**, 586-606.

Taylor, John B. (1993): "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, **39**, 195-214.

Table 1: Prior and Posterior of DSGE Model Parameters (Part 1)

	Prior			Posterior	
Name	Density	Para (1)	Para (2)	Mean	90% Intv.
Household					
h	$\mathcal{B}$	0.70	0.05	0.65	[0.58 , 0.72]
a''	$\mathcal{G}$	0.20	0.10	0.30	[0.13 , 0.47]
ν_l	$\mathcal{G}$	2.00	0.75	2.29	[1.33 , 3.28]
ζ_w	$\mathcal{B}$	0.60	0.20	0.25	[0.15 , 0.35]
$400(1/\beta - 1)$	$\mathcal{G}$	2.00	1.00	1.034	[0.45 , 1.60]
Firms					
α	$\mathcal{B}$	0.33	0.10	0.20	[0.15 , 0.24]
ζ_p	$\mathcal{B}$	0.60	0.20	0.66	[0.53 , 0.84]
S''	$\mathcal{G}$	4.00	1.50	2.29	[0.84 , 3.91]
λ_f	$\mathcal{G}$	0.15	0.10	0.14	[0.01 , 0.26]
Monetary Policy					
$400\pi_*$	$\mathcal{N}$	3.00	1.50	2.94	[2.08 , 3.78]
ψ_1	$\mathcal{G}$	1.50	0.40	3.05	[2.43 , 3.68]
ψ_2	$\mathcal{G}$	0.20	0.10	0.06	[0.03 , 0.10]
ρ_R	$\mathcal{B}$	0.50	0.20	0.86	[0.83 , 0.89]

Table 1: Prior and Posterior of DSGE Model Parameters (Part 2)

		Prior		Posterior	
Name	Density	Para (1)	Para (2)	Mean	90% Intv.
			Shocks		
400γ	$\mathcal{G}$	2.00	1.00	1.57	[1.13 , 2.02]
g_*	$\mathcal{G}$	0.30	0.10	0.29	[0.13 , 0.43]
ρ_a	$\mathcal{B}$	0.20	0.10	0.19	[0.10 , 0.29]
ρ_μ	$\mathcal{B}$	0.80	0.05	0.80	[0.74 , 0.87]
ρ_{λ_f}	$\mathcal{B}$	0.60	0.20	0.67	[0.30 , 0.94]
ρ_g	$\mathcal{B}$	0.80	0.05	0.96	[0.95 , 0.98]
ρ_b	$\mathcal{B}$	0.60	0.20	0.85	[0.78 , 0.93]
ρ_ϕ	$\mathcal{B}$	0.60	0.20	0.98	[0.96 , 0.99]
σ_a	$\mathcal{IG}$	0.75	2.00	0.62	[0.54 , 0.69]
σ_μ	$\mathcal{IG}$	0.75	2.00	0.53	[0.38 , 0.68]
σ_{λ_f}	$\mathcal{IG}$	0.75	2.00	0.18	[0.15 , 0.21]
σ_g	$\mathcal{IG}$	0.75	2.00	0.33	[0.29 , 0.37]
σ_b	$\mathcal{IG}$	0.75	2.00	0.36	[0.28 , 0.45]
σ_ϕ	$\mathcal{IG}$	4.00	2.00	2.90	[1.99 , 3.80]
σ_R	$\mathcal{IG}$	0.20	2.00	0.14	[0.12 , 0.16]

Notes: Para (1) and Para (2) list the means and the standard deviations for the Beta ($\mathcal{B}$), Gamma ($\mathcal{G}$), and Normal ($\mathcal{N}$) distributions; the upper and lower bound of the support for the Uniform ($\mathcal{U}$) distribution; s and ν for the Inverse Gamma ($\mathcal{IG}$) distribution, where $p_{\mathcal{IG}}(\sigma|\nu,s) \propto \sigma^{-(\nu+1)}e^{-\nu s^2/2\sigma^2}$. The joint prior distribution is obtained as a product of the marginal distributions tabulated in the table and truncating this product at the boundary of the determinacy region. Posterior summary statistics are computed based on the output of the posterior sampler. The following parameters are fixed: $\delta = 0.025$, $\lambda_w = 0.3$. Estimation sample: 1984:I to 2007:III.

Table 2: RMSE Comparison: DSGE Model versus AR(1)

Series	Model	$h=1$	$h=2$	$h=4$	$h=12$
Output Growth (Q %)	DSGE	0.51	0.50	**0.41**	**0.36**
	AR(1)	0.50	0.49	0.44	0.37
Consumption Growth (Q %)	DSGE	0.39	0.38	0.39	0.39
	AR(1)	0.37	0.37	0.34	0.31
Investment Growth (Q %)	DSGE	**1.44**	**1.56**	**1.47****	**1.52**
	AR(1)	1.56	1.67	1.60	1.60
Nominal Wage Growth (Q %)	DSGE	0.67	0.70	0.66	**0.56**
	AR(1)	0.59	0.59	0.59	0.56
100 × Log Hours	DSGE	**0.52****	**0.88****	**1.44****	**2.07****
	AR(1)	0.66	1.20	2.08	3.40
Inflation (Q %)	DSGE	**0.22**	**0.23**	**0.19****	0.24
	AR(1)	0.22	0.23	0.22	0.23
Interest Rates (A %)	DSGE	0.71	1.34	2.13	**2.25**
	AR(1)	0.54**	1.00**	1.73	2.93

Notes: We report RMSEs for DSGE and AR(1) models. Numbers in boldface indicate a lower RMSE of the DSGE model. $*$ ($**$) denotes 10% (5%) significance of the two-sided modified Diebold-Mariano test of equal predictive accuracy under quadratic loss. The RMSEs are computed based on recursive estimates starting with the sample 1984:I to 2000:IV and ending with the samples 1984:I to 2007:III (h=1), 1984:I to 2007:II (h=2), 1984:I to 2006:III (h=4), 1984:I to 2004:III (h=12), respectively. h-step ahead growth (inflation) rate forecasts refer to percentage changes between period $T+h-1$ and $T+h$.

Table 3: One-Step-Ahead Forecast Performance of DSGE Models

Study	Forecast Period	Output Growth (Q %)	Inflation (Q %)	Interest Rate (A %)
Schorfheide, Sill, Kryshko	2001:I to 2007:IV	0.51	0.22	0.71
		(0.47)	(0.22)	(1.68)
Del Negro *et al.* (2007)	1985:IV to 2000:I	0.73	0.27	0.87
		(0.52)	(0.25)	(1.72)
Edge, Kiley, Laforte (2008)	1996:III to 2005:II	0.38	0.22	0.59
		(0.57)	(0.20)	(1.96)
Smets, Wouters (2007)	1990:I to 2004:IV	0.57	0.24	0.43
		(0.57)	(0.22)	(1.97)

Notes: Schorfheide, Sill, Kryshko: RMSEs, DSGE model is estimated recursively with data starting in 1984:I. Del Negro *et al.* (2007, Table 2): RMSEs, VAR approximation of DSGE model estimated based on rolling samples of 120 observations. Edge, Kiley, and Laforte (2008, Table 4): Mean absolute errors, DSGE model is estimated recursively with data starting in 1984:II. Smets and Wouters (2007, Table 3): RMSEs, DSGE model is estimated recursively, starting with data from 1966:I. Numbers in parentheses are sample standard deviations for forecast period, computed from the Schorfheide, Sill, Kryshko data set. Q % is the quarter-to-quarter percentage change, and A % is an annualized rate.

Table 4: NON-MODELLED AND RELATED DSGE MODEL VARIABLES

Non-Modelled Variable	DSGE Model Variable	Transformation
PCE Inflation	Final Good Inflation π_t	None
Core PCE Inflation	Final Good Inflation π_t	None
Unemployment Rate	Hours Worked L_t	$-0.31L_t$
Housing Starts	Investment i_t	$0.033i_t$

Notes: Here L_t and i_t are the DSGE model variables that appear in the model description in Section 2.

Table 5: Auxiliary Regression Estimates

		ρ		σ_η		Signal/Noise	
Series	λ	Mean	90% Intv	Mean	90% Intv	$\frac{v\hat{a}r(\hat{s}'_{t\|t}\hat{\alpha}_1)}{v\hat{a}r(\hat{\xi}_t)}$	$\ln p_\lambda(Z^T)$
PCE Inflation	1.00	0.05	[-0.14, 0.26]	0.03	[0.02, 0.03]	3.15	-0.03
	0.10	0.05	[-0.16, 0.25]	0.03	[0.02, 0.04]	2.62	4.82
	1E-5	0.07	[-0.11, 0.24]	0.04	[0.03, 0.05]	1.47	12.27
Core PCE Inflation	1.00	0.23	[0.03, 0.45]	0.01	[0.01, 0.02]	4.99	29.53
	0.10	0.21	[-0.02, 0.41]	0.01	[0.01, 0.02]	4.88	39.12
	1E-5	0.53	[0.38, 0.68]	0.03	[0.02, 0.04]	1.35	22.58
Unemployment Rate	1.00	0.98	[0.96, 1.00]	0.019	[0.01, 0.02]	3.45	17.71
	0.10	0.97	[0.95, 1.00]	0.019	[0.01, 0.02]	3.67	23.68
	1E-5	0.98	[0.97, 1.00]	0.025	[0.02, 0.03]	1.91	22.78
Housing Starts	1.00	0.89	[0.76, 1.00]	0.007	[0.00, 0.01]	0.74	68.21
	0.10	0.88	[0.74, 1.00]	0.007	[0.01, 0.01]	0.95	80.81
	1E-5	0.96	[0.92, 1.00]	0.009	[0.01, 0.01]	0.88	82.64

Notes: The posterior summary statistics are computed based on the output of the Gibbs sampler. The sample variance ratios are computed using the posterior mean estimate of α_1. Estimation sample: 1984:I to 2007:III.

Table 6: RMSEs for Auxiliary Regressions [PARTIALLY UPDATED]

Series	Model	λ	$h=1$	$h=2$	$h=4$	$h=12$
PCE Inflation (Q %)	Aux	1.00	**0.34**	0.37	0.34	**0.32**
	Aux	0.10	**0.33**	**0.35**	**0.32**	0.35
	Aux	1E-5	**0.32**	**0.34**	**0.30**	0.33
	Regr.		**0.33**	**0.35**	**0.32**	0.49
	AR(1)		0.36	0.35	0.33	0.32
Core PCE Inflation (Q %)	Aux	1.00	0.18	0.19	**0.16**	**0.12**
	Aux	0.10	0.18	0.18	**0.15**	**0.11**
	Aux	1E-5	0.16	0.20	**0.18**	**0.15**
	Regr.		**0.14**	**0.14**	**0.17**	0.35
	AR(1)		0.16	0.16	0.18	**0.17**
Unemployment Rate (%)	Aux	1.00	**0.16****	**0.27**	**0.43**	**1.02**
	Aux	0.10	**0.15****	**0.24**	**0.39**	**0.97**
	Aux	1E-5	**0.15****	**0.23***	**0.37**	**0.74**
	Regr.		**0.20**	0.37	0.72	1.39
	AR(1)		0.21	0.37	0.63	1.01
Housing Starts (4 Million / Q)	Aux	1.00	0.11	0.18	0.31	0.50
	Aux	0.10	0.11	0.17	0.29	0.48
	Aux	1E-5	**0.10**	**0.16**	0.27	0.45
	Regr.		**0.10**	**0.16**	**0.26**	**0.43**
	AR(1)		0.10	0.16	0.27	0.43

Notes: We report RMSEs for DSGE, AR(1), and regression models. Numbers in boldface indicate that DSGE model attains lower RMSE than AR(1) model. $*$ ($**$) denotes 10% (5%) significance of the two-sided modified Diebold-Mariano test of equal predictive accuracy under quadratic loss. The RMSEs are computed based on recursive estimates starting with the sample 1984:I to 2000:IV and ending with the samples 1984:I to 2007:III (h=1), 1984:I to 2007:II (h=2), 1984:I to 2006:III (h=4), 1984:I to 2004:III (h=12), respectively. h-step ahead growth (inflation) rate forecasts refer to percentage changes between period $T+h-1$ and $T+h$.

Table 7: Posterior Predictive Check: Cross-Correlations

		Output Growth	Inflation	Interest Rates
PCE Inflation	90% Intv.	[-0.46, 0.01]	[0.50, 0.91]	[0.11, 0.63]
λ = 1E-5	Data	-0.07	0.75	0.42
Core PCE Inflation	90% Intv.	[-0.47, 0.03]	[0.50, 0.91]	[0.07, 0.63]
λ = 1E-5	Data	0.01	0.68	0.61
Unemployment Rate	90% Intv.	[-0.32, 0.09]	[-0.26, 0.36]	[-0.24, 0.63]
λ = 1E-5	Data	0.15	0.17	0.12
Housing Starts	90% Intv	[-0.11, 0.33]	[-0.26, 0.33]	[-0.47, 0.43]
λ = 1E-5	Data	0.23	0.05	-0.22

Notes: We report 90% credible intervals of the posterior predictive distribution for the sample correlations of non-modelled variables with core variables. The data entries refer to sample correlations calculated from U.S. data.

Table 8: RMSE Ratios: Conditional (on Interest Rates) versus Unconditional

Series	$h = 1$	$h = 2$	$h = 4$	$h = 12$
Output Growth (Q %)	1.08	1.18	1.22	1.17
	(0.93, 0.94)	(0.91, 0.99)	(0.93, 1.06)	(0.97, 0.99)
100 $\times$ Log Hours	1.23	1.42	1.57	2.05
	(0.96, 0.98)	(0.91, 0.98)	(0.93, 1.03)	(0.96, 0.97)
Inflation (Q %)	1.14	1.18	1.86	2.02
	(0.80, 0.81)	(0.90, 0.97)	(0.90, 1.01)	(0.93, 0.94)
PCE Inflation (Q %)	0.95	1.01	1.40	1.69
λ = 1E-5	(0.90, 0.90)	(0.90, 0.97)	(0.90, 1.01)	(0.90, 0.91)
Core PCE Inflation (Q %)	0.99	1.05	1.91	3.26
λ = 1E-5	(0.88, 0.87)	(0.89, 0.96)	(0.90, 1.01)	(0.91, 0.92)
Unemployment Rate (%)	1.16	1.43	1.60	1.45
λ = 1E-5	(0.98, 0.98)	(0.97, 1.04)	(0.96, 1.08)	(0.93, 0.94)
Housing Starts (4 Million / Q)	1.01	1.00	0.99	1.00
λ = 1E-5	(1.00, 0.99)	(1.00, 1.08)	(1.00, 1.12)	(1.00, 1.02)

Notes: Using the draws from the posterior predictive distribution of two variables x_1 and x_2 we construct conditional mean forecasts for x_2 given x_1, assuming that the predictive distribution is student-t with $\nu = 5$ or $\nu = \infty$ degrees of freedom. We report RMSE ratios for conditional and unconditional recursive h-step ahead pseudo-out-of-sample forecast as well as the theoretical reductions $\mathcal{R}(\infty)$ and $\mathcal{R}(5)$ in parenthesis (see Equation (33) for a definition).

Table 9: RMSE Ratios: Conditional (on GDP Deflator Inflation) versus Unconditional

Series	$h = 1$	$h = 2$	$h = 4$	$h = 12$
Output Growth (Q %)	0.94	0.91	0.94	1.04
	(0.94, 0.88)	(0.74, 0.69)	(0.76, 0.68)	(0.98, 0.90)
100 $\times$ Log Hours	1.01	1.03	1.06	0.92
	(0.98, 0.92)	(0.74, 0.69)	(0.74, 0.66)	(0.98, 0.90)
PCE Inflation (Q %)	0.71	0.68	0.83	0.83
λ = 1E-5	(0.69, 0.65)	(0.68, 0.63)	(0.67, 0.60)	(0.67, 0.61)
Core PCE Inflation (Q %)	1.07	0.98	1.26	2.11
λ = 1E-5	(0.58, 0.55)	(0.63, 0.59)	(0.67, 0.60)	(0.68, 0.62)
Unemployment Rate (%)	1.06	1.08	1.09	1.10
λ = 1E-5	(0.99, 0.92)	(0.99, 0.92)	(0.99, 0.89)	(0.95, 0.86)
Housing Starts (4 Million / Q)	1.00	1.00	1.00	1.00
λ = 1E-5	(1.00, 0.93)	(1.00, 0.93)	(1.00, 0.90)	(1.00, 0.91)

Notes: See Table 8.

Figure 1: Latent State Variables of the DSGE Model

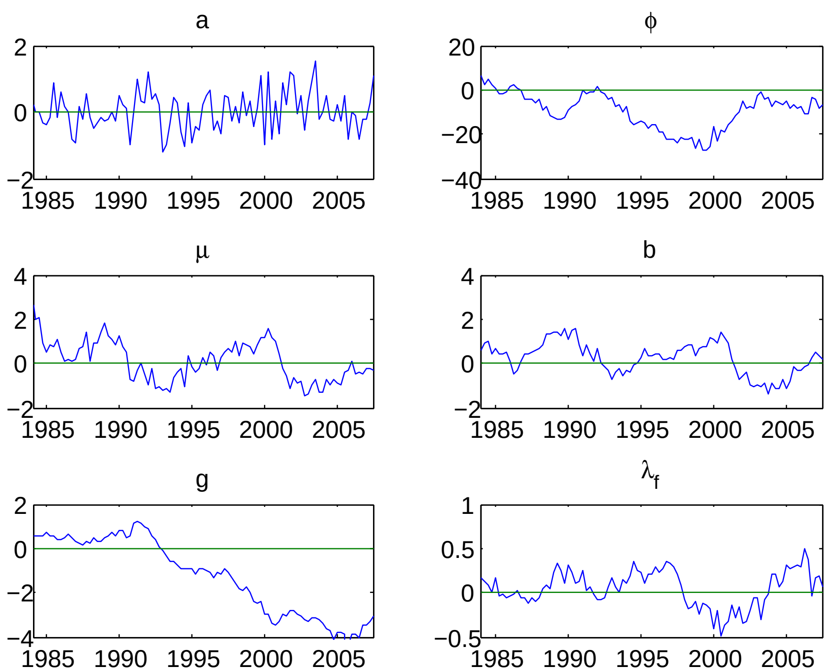

Notes: The six panels of the figure depict time series of elements of $\hat{s}_{t|t}$. Estimation sample: 1984:I to 2007:III.

Figure 2: Non-Core Variables and Related Model Variables

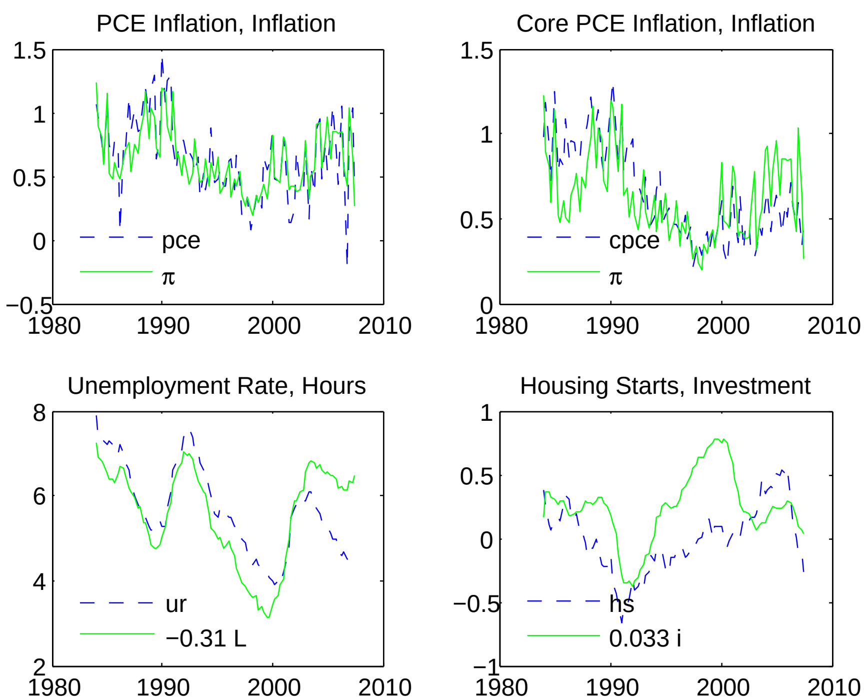

Notes: The top two panels depict quarter-to-quarter inflation rates. In the bottom panels we add constants to the scaled log of hours worked and investment deviations from trend to match the means of the unemployment rate and housing starts over the period 1984:I to 2007:III.

Figure 3: Non-Core Variables and Factors

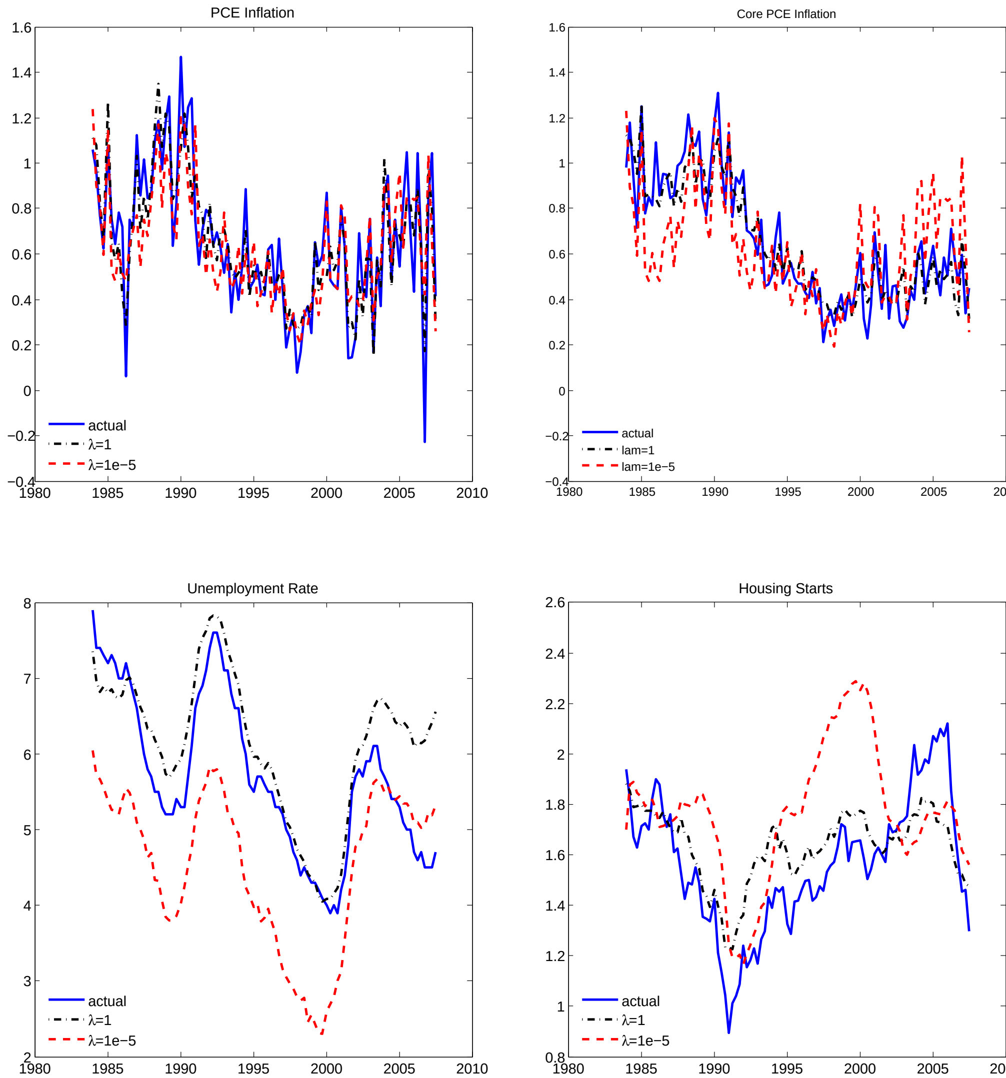

Notes: Figure depicts the actual (blue, solid) path of the non-core variables as well as the factor predictions $\hat{\alpha}_0 + \hat{s}'_{t|t}\hat{\alpha}_{1,T}$ for $\lambda = 1\text{E-}5$ (light blue, dashed) and $\lambda = 1$ (red, dotted).

Figure 4: Impulses Response to a Monetary Policy Shock

(i) Core Variables: Output, GDP Deflator Inflation, Hours, Investment

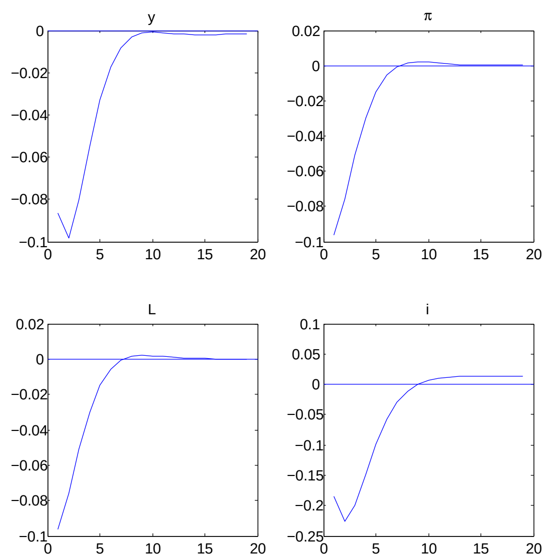

(ii) Non-core Variables: PCE Inflation, core PCE Inflation, Unemployment, Housing Starts

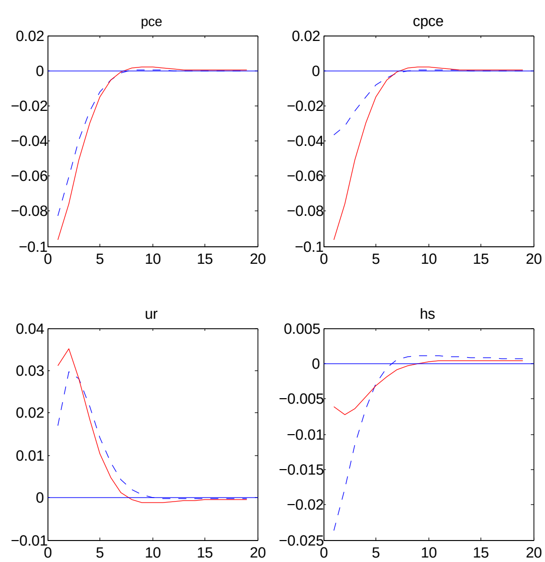

Notes: Non-core variables: we depict log level responses for output, hours, and investment. (ii) Non-core variables: we overlay two responses, corresponding to the auxiliary regressions estimated with $\lambda = 1\text{E-}5$ (red, solid), and $\lambda = 1$ (blue, dashed). Estimation sample: 1984:I to 2007:III.

Figure 5: Bivariate One-Step-Ahead Predictive Distributions

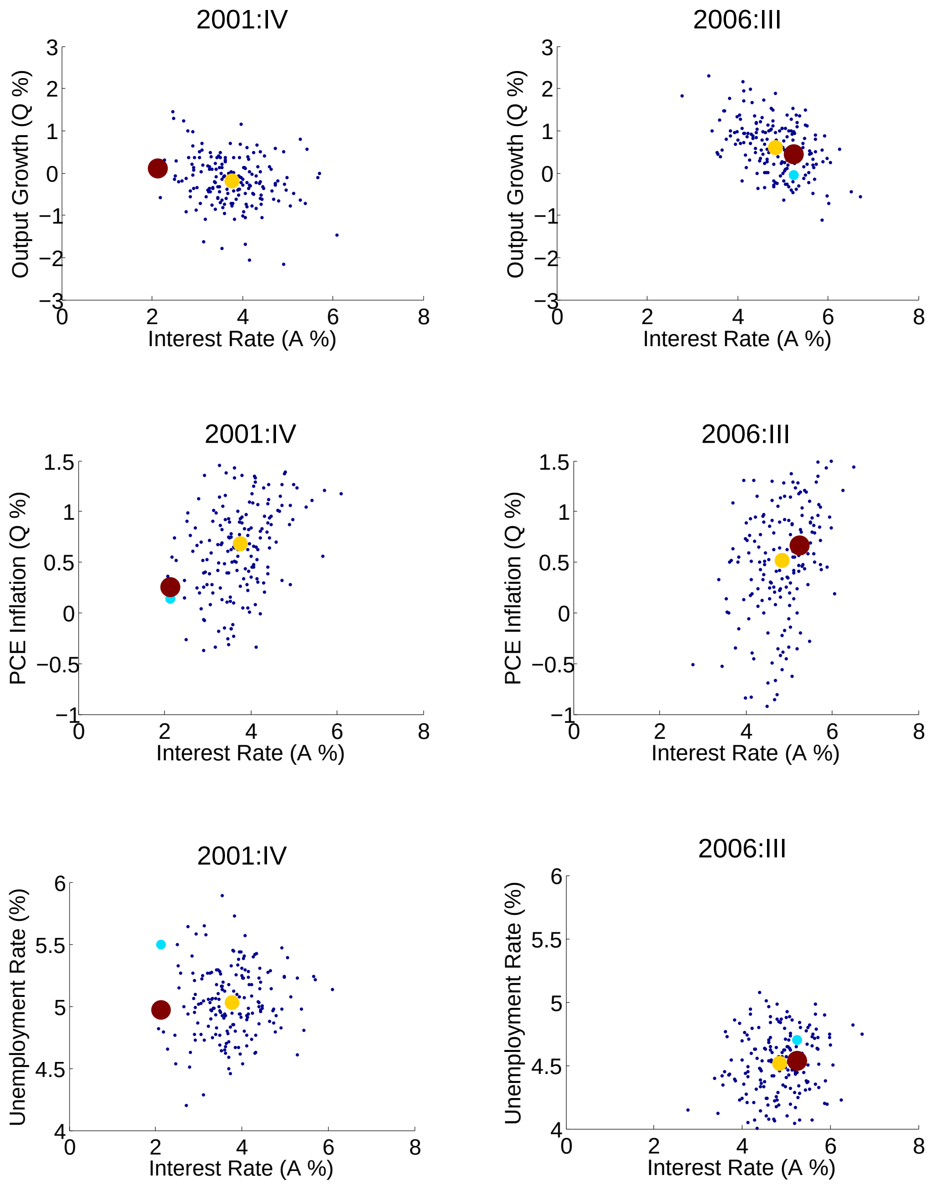

Notes: The panels depict a scatter plot of draws from the one-step-ahead predictive distribution. The three filled circles denote: the actual value (small, light blue), the unconditional mean predictor (medium, yellow), and the conditional mean predictor (large, brown). We set $\lambda = 1\text{E-}5$.

A MCMC Implementation

DSGE model coefficients. The posterior sampler for the DSGE model is described in An and Schorfheide (2007).

Gibbs sampler for the coefficients that appear in the measurement equations. We will in turn derive the conditional distributions for a Gibbs sampler that iterates over the conditional posteriors of α, ρ, and σ_η^2. We will start from the quasi-differenced form (23) of the auxiliary regression. τ, λ_0, and λ_1 are treated as hyperparameters and considered as fixed in the description of the Gibbs sampler. Let L denote the lag operator.

Conditional posterior of α: The posterior density is of the form

$$p(\alpha|\rho,\sigma_\eta^2,Z^T,S^T) \propto p(Z^T|S^T,\alpha,\rho,\sigma_\eta^2)p(\alpha). \tag{A.1}$$

Define

$$\begin{aligned} y_1 &= \frac{\sigma_\eta}{\tau} z_1, \quad x_1' = \frac{\sigma_\eta}{\tau}[1,\ \hat{s}_{1|1}'] \\ y_t &= (1-\rho L)z_t, \quad x_t' = [1-\rho,\ (1-\rho L)\hat{s}_{t|t}']', \quad t=2,\dots,T, \end{aligned}$$

which implies that (23) can be expressed as linear regression

$$y_t = x_t'\alpha + \eta_t. \tag{A.2}$$

If we let Y be a $T \times 1$ matrix with rows y_t and X be a $T \times k$ matrix with rows x_t', then we can rewrite the regression in matrix form

$$Y = X\alpha + E.$$

We deduce

$$\begin{aligned} p(\alpha|\rho,\sigma_\eta^2,Z^T,S^T) &\propto \exp\left\{-\frac{1}{2\sigma_\eta^2}(\alpha-\hat{\alpha})'X'X(\alpha-\hat{\alpha})\right\} \\ &\quad \times \exp\left\{-\frac{1}{2}(\alpha-\mu_{\alpha,0})'V_{\alpha,0}^{-1}(\alpha-\mu_{\alpha,0})\right\}, \end{aligned} \tag{A.3}$$

where

$$\hat{\alpha} = (X'X)^{-1}X'Y.$$

Thus, the conditional posterior of α is $\mathcal{N}(\mu_{\alpha,T}, V_{\alpha,T})$ with

$$\begin{aligned} \mu_{\alpha,T} &= V_{\alpha,T}\left[V_{\alpha,0}^{-1}\mu_{\alpha,0} + \frac{1}{\sigma_\eta^2}X'X\hat{\alpha}\right] \\ V_{\alpha,T} &= \left(V_{\alpha,0}^{-1} + \frac{1}{\sigma_\eta^2}X'X\right)^{-1}. \end{aligned}$$

Conditional posterior of ρ: Given the $\mathcal{U}(-1,1)$ prior for ρ, the posterior density is of the form

$$p(\rho|\alpha,\sigma_\eta^2,Z^T,S^T) \propto p(Z^T|S^T,\alpha,\rho,\sigma_\eta^2)\mathcal{I}\{|\rho|<1\}. \tag{A.4}$$

We now define

$$y_t = z_t - \alpha_0 - \hat{s}'_{t|t}\alpha_1, \quad x_t = z_{t-1} - \alpha_0 - \hat{s}'_{t-1|t-1}\alpha_1.$$

Again, we can express (23) as linear regression model

$$y_t = x_t\rho + \eta_t. \tag{A.5}$$

Using the same arguments as before we deduce that

$$p(\rho|\alpha,\sigma_\eta^2,Z^T,S^T) \propto \mathcal{I}\{|\rho|<1\}\exp\left\{-\frac{1}{2\sigma_\eta^2}(\rho-\hat{\rho})'X'X(\rho-\hat{\rho})\right\} \tag{A.6}$$

with

$$\hat{\rho} = (X'X)^{-1}X'Y.$$

Thus, the conditional posterior is truncated normal: $\mathcal{I}\{|\rho|<1\}\mathcal{N}(\mu_{\rho,T},V_{\rho,T})$ with

$$\mu_{\rho,T} = \hat{\rho}, \quad V_{\rho,T} = \sigma_\eta^2(X'X)^{-1}.$$

Conditional posterior of σ_η^2: The posterior density is of the form

$$p(\sigma_\eta^2|\alpha,\rho,Z^T,S^T) \propto p(Z^T|S^T,\alpha,\rho,\sigma_\eta^2)p(\sigma_\eta^2). \tag{A.7}$$

Solve (23) for η_t:

$$\eta_t = (1-\rho L)z_t - (1-\rho)\alpha_0 - (1-\rho L)\hat{s}'_{t|t}\alpha_1. \tag{A.8}$$

Now, notice that

$$p(\sigma_\eta^2|\alpha,\rho,Z^T,S^T) \propto (\sigma_\eta^2)^{-(T+2)/2}\exp\left\{-\frac{1}{2\sigma_\eta^2}\sum\eta_t^2\right\}. \tag{A.9}$$

This implies that the conditional posterior of σ_η^2 is inverted Gamma with T degrees of freedom and location parameter $s^2 = \sum\eta_t^2$. To sample a σ_η^2 from this distribution generate T random draws $Z_1,\ldots,Z_T$ from a $\mathcal{N}(0,1/s^2)$ and let $\tilde{\sigma}_\eta^2 = \left[\sum_{j=1}^T Z_j^2\right]^{-1}$.

Marginal Data Density: Can be approximated using Chib's (1995) method. Let $\hat{\alpha}$, $\hat{\rho}$, and $\hat{\sigma}_\eta^2$ be the posterior mean estimates computed from the output of the Gibbs sampler. According to Bayes Theorem,

$$p(Y) = \frac{p(Y|\hat{\alpha},\hat{\rho},\hat{\sigma}_\eta^2)p(\hat{\alpha})p(\hat{\rho})p(\hat{\sigma}_\eta^2)}{p(\hat{\alpha}|\hat{\rho},\hat{\sigma}_\eta^2,Y)p(\hat{\rho}|\hat{\sigma}_\eta^2,Y)p(\hat{\sigma}_\eta^2|Y)} \tag{A.10}$$

All but the following two terms are straightforward to evaluate. First, let $\alpha_{(i)}$ and $\rho_{(i)}$ denote the i'th draw from the Gibbs sampler. Then we can use the approximation:

$$\widehat{p}(\hat{\sigma}^2_\eta|Y) = \frac{1}{n}\sum_{i=1}^{n} p(\hat{\sigma}^2_\eta|\alpha_{(i)}, \rho_{(i)}, Y). \tag{A.11}$$

Now consider a "reduced" run of the Gibbs sampler, in which we fix $\sigma^2_\eta = \hat{\sigma}^2_\eta$ and iterate over $p(\alpha|\rho, \hat{\sigma}^2_\eta, Y)$ and $p(\rho|\alpha, \hat{\sigma}^2_\eta, Y)$ using the conditional densities in (A.3) and (A.6). Denote the output of this Gibbs sampler by $\rho_{(s)}$ and $\alpha_{(s)}$. Then,

$$\widehat{p}(\hat{\rho}|\hat{\sigma}^2_\eta, Y) = \frac{1}{m}\sum_{s=1}^{m} p(\hat{\rho}|\alpha_{(s)}, \hat{\sigma}^2_\eta, Y). \tag{A.12}$$

Generalization to AR(p): Let $\rho(L) = 1 - \sum_{j=1}^{p} \rho_p L^p$, where L is the lag operator, then we can express the auxiliary model as:

$$\begin{aligned} z_t &= \alpha_0 + \hat{s}'_{t|t}\alpha_1 + \xi_t, \quad t = 1, \ldots, p \\ \rho(L) z_t &= \rho(1)\alpha_0 + \rho(L)\hat{s}'_{t|t}\alpha_1 + \eta_t, \quad t = p+1, \ldots, T \end{aligned}$$

where $[\xi_1, \ldots, \xi_p]' \sim \mathcal{N}(0, \tau^2\Omega(\rho(L)))$ and $\Omega(\rho(L))$ is the correlation matrix associated with the stationary AR(p) specification of ξ_t. The conditional posteriors of α and σ^2_η are obtained from a straightforward generalization of (A.3) and (A.9). The conditional posterior distribution of $\rho_1, \ldots, \rho_p$ is now non-normal and requires a Metropolis step. A generalization of (A.6) can serve as proposal density. To conveniently enforce stationarity of the autoregressive measurement error process it could be re-parameterized in terms of partial autocorrelations as in Barndorff-Nielson and Schou (1973).