WORKING PAPERS
RESEARCH DEPARTMENT

WORKING PAPER NO. 08-17
DSGE MODEL-BASED FORECASTING OF
NON-MODELLED VARIABLES

Frank Schorfheide
University of Pennsylvania and
Visiting Scholar,
Federal Reserve Bank of Philadelphia

Keith Sill
Federal Reserve Bank of Philadelphia

Maxym Kryshko
University of Pennsylvania

September 8, 2008

DSGE Model-Based Forecasting of

Non-modelled Variables

Frank Schorfheide

University of Pennsylvania

CEPR,NBER, and visiting scholar,

Federal Reserve Bank of Philadelphia

Keith Sill[*]

Federal Reserve Bank

of Philadelphia

Maxym Kryshko

University of Pennsylvania

September 8, 2008

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, and the unemployment rate 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). Second, it is more difficult 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 the DSGE model. Our paper focuses on the second problem. For brevity, we will refer to these non-modelled series as non-core variables.

Recently, Boivin and Giannoni (2006) 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. 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 Boivin and Giannoni 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

---

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

of the joint estimation is its computational complexity, which makes it currently impractical for real time forecasting at central banks.

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. 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. 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. This estimation is quick and can be easily repeated in real time as new information arrives or interest in additional non-core variables arises. An attractive feature of our empirical model for policy makers is that we are linking the non-core variables to the fundamental shocks that are believed to drive business cycle fluctuations. In particular, we are creating a link between monetary policy shocks and non-core variables, which allows us to study the effect of unanticipated changes in monetary policy on a broad set of economic variables.

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 Del Negro, Schorfheide, Smets, and Wouters (DSSW, 2007). 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, generate estimates of the latent states as well as pseudo-out-of-sample forecasts for a set of core variables, which is comprised of the growth rates of output, consumption, investment, real wages, the GDP deflator, as well as the levels of interest rates and hours worked. We then estimate measurement equations for three additional variables: personal consumption expenditures (PCE) inflation, core PCE inflation, and the unemployment rate. 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

This section briefly describes the DSGE model to which we apply our methods of constructing prior distributions. 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 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:

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \frac{(1 - \zeta_p\beta)(1 - \zeta_p)}{\zeta_p} mc_t + \frac{1}{\zeta_p}\lambda_{f,t}, \qquad (4)$$

where π_t is inflation and β is the discount rate.[2] 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

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

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

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

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

---

[2]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 $\widetilde{\lambda}_{f,t}$.

factor[3] 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:

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

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, 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:

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

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:

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

where ξ_t^k is the value of installed capital, evolving according to:

$$\xi_t^k - \xi_t = \beta e^{-\gamma}(1 - \delta)\mathbb{E}_t\left[\xi_{t+1}^k - \xi_{t+1}\right] + \mathbb{E}_t\left[(1 - (1 - \delta)\beta e^{-\gamma})r_{t+1}^k - (R_t - \pi_{t+1})\right]. \tag{12}$$

---

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

Capital utilization u_t in our model is variable and r_t^k 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

$$u_t = \frac{r_*^k}{a''} r_t^k.$$ (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:

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

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

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

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

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 macroeconomic 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 state variables ς_t:

$$\varsigma_t = \Phi_1(\theta)\varsigma_{t-1} + \Phi_\epsilon(\theta)\epsilon_t. \tag{17}$$

The coefficients of the matrices Φ_1 and Φ_ϵ are functions of the DSGE model parameters θ. For the model described in Section 2 the non-redundant state variables are given by c_t, i_t, $\bar{k}_t$, R_t, w_t, and the six serially correlated exogenous disturbances.

To estimate the DSGE model based on a sequence of observations $Y^T = [y_t, \ldots, y_T]$ using a likelihood-based method it is convenient to construct a state-space model. In addition to the state transition equation (17) one needs to specify a system of measurement equations that link the observables y_t to the states ς_t. The vector y_t used in our empirical analysis includes quarter-to-quarter growth rates (measured in percentages) of real GDP, consumption, investment, and nominal wages, as well as a measure of hours worked, the GDP deflator, and the federal funds rate. Since some of our observables include growth rates, we augment the set of model states by lagged values of output, consumption, investment, and real wages and augment the matrices Φ_1 and Φ_2 in (17) accordingly. Thus,

$$\varsigma_t = [c_t, i_t, \bar{k}_t, R_t, w_t, a_t, \mu_t, \lambda_{f,t}, g_t, b_t, \phi_t, y_{t-1}, c_{t-1}, i_{t-1}, w_{t-1}].'$$

We express the set of measurement equations generically as

$$y_t = A_0(\theta) + A_1(\theta)\varsigma_t. \tag{18}$$

The state-space representation of the DSGE model is given by Equations (17) and (18).

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:

$$\varsigma_{t|t}(\theta) = I\!\!E[\varsigma_t|\theta, Y^t], \tag{19}$$

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

$$p(\theta|Y^T) = \frac{p(Y^T|\theta)p(\theta)}{p(Y)}, \quad \text{where} \quad p(Y^T) = \int p(Y^T|\theta)p(\theta)d\theta. \tag{20}$$

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 ς_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 $\hat{\theta}_T$ and let[4]

$$z_t = \alpha_0 + \hat{s}'_{t|t}\alpha_1 + \xi_t, \quad \xi_t = \rho\xi_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2), \tag{21}$$

where $\hat{s}_{t|t} = M\hat{\varsigma}_{t|t}$ and M is a matrix composed of zeros and ones that potentially selects a J-dimensional subset of the model state variables. Moreover, ξ_t is a variable-specific noise process.

---

[4]Alternatively, we could define $\hat{\varsigma}_{t|t}$ as $\int \varsigma_{t|t}(\theta)p(\theta|Y^T)d\theta$.

Equations (17), (18), and (21) 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 the auxiliary regression (21) describes how additional macroeconomic variables load on the factors. Our framework can be viewed as a simplified version of Boivin and Giannoni's (2006) DSGE-based factor model. The random variable ξ_t in (21) plays the role of an idiosyncratic error term. Unlike Boivin and Giannoni (2006), 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.

As in the estimation of the DSGE model, we also use Bayesian methods for the estimation of the auxiliary regression for z_t. We re-write (21) in quasi-differenced form as

$$
\begin{aligned}
z_1 &= \alpha_0 + \hat{s}'_{1|1}\alpha_1 + \xi_1 \tag{22} \\
z_t &= \rho z_{t-1} + \alpha_0(1-\rho) + [\hat{s}'_{t|t} - \hat{s}'_{t-1|t-1}\rho]\alpha_1 + \eta_t, \quad t = 2,\ldots,T.
\end{aligned}
$$

Instead of linking the distribution of ξ_1 to the parameters ρ and σ_η^2 we assume that $\xi_1 \sim \mathcal{N}(0,\tau^2)$ and discuss the choice of τ later on. 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

$$
\alpha \sim \mathcal{N}(\mu_{\alpha,0}, V_{\alpha,0}), \quad \rho \sim \mathcal{U}(-1,1), \quad p(\sigma_\eta^2) \propto (\sigma_\eta^2)^{-1}. \tag{23}
$$

We construct the prior mean $\mu_{\alpha,0}$ based on the DSGE model implied factor loadings for a model variable, say $z_t^\dagger$, 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 $\mathbb{E}_\theta^D[\cdot]$ 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

$$
\mu_{\alpha,0} = \left(\mathbb{E}_\theta^D[\tilde{s}_t\tilde{s}'_t] \right)^{-1} \mathbb{E}_\theta^D[\tilde{s}_t z_t^\dagger], \tag{24}
$$

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 essentially recovers[5] the corresponding rows of $A_0(\theta)$

---

[5]Depending on the procedure used to solve the DSGE model, some elements of ς_t might be redundant

and $A_1(\theta)$ in the measurement equation (18). We provide details on the choice of $z_t^\dagger$ in the empirical section. Our prior covariance matrix is diagonal with the following elements

$$diag(V_{\alpha,0}) = \left[\lambda_0, \frac{\lambda_1}{\omega_1}, \ldots, \frac{\lambda_1}{\omega_J}\right]. \tag{25}$$

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 the inverse of ω_j, which denotes the DSGE model's implied variance of the j'th element of $\hat{s}_{t|t}$ (evaluated at the posterior mean of θ). 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[6] $\hat{\varsigma}_{T|T}(\theta^{(i)})$ and draw a random sequence $\{\epsilon_{T+1}^{(i)}, \ldots, \epsilon_{T+h}^{(i)}\}$. We then iterate the state transition equation forward to construct

$$\varsigma_{T+h|T}^{(i)} = \Phi_1(\theta^{(i)})\varsigma_{T+h-1|T}^{(i)} + \Phi_\epsilon(\theta^{(i)})\epsilon_{T+h}^{(i)}, \quad h = 1, \ldots, H. \tag{26}$$

Finally, we use the measurement equation to compute

$$y_{T+h|T}^{(i)} = A_0(\theta^{(i)}) + A_1(\theta^{(i)})\varsigma_{T+h|T}^{(i)}. \tag{27}$$

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

A draw from the posterior predictive distribution of a non-core variable z_{T+h} is obtained as follows. Using the sequence $\varsigma_{T+1|T}^{(i)}, \ldots, \varsigma_{T+H|T}^{(i)}$ constructed in (26), we iterate the quasi-differenced version (22) of the auxiliary regression forward:

$$z_{T+h|T}^{(i)} = \rho^{(i)} z_{T+h-1}^{(i)} + \alpha_0^{(i)}(1 - \rho^{(i)}) + [\varsigma_{T+h|T}^{(i)'}M' - \varsigma_{T+h-1|T}^{(i)'}M'\rho^{(i)}]\alpha_1^{(i)} + \eta_{T+h}^{(i)},$$

where the superscript i for the parameters of (21) refers to the i'th draw from the posterior distribution of $(\alpha, \rho, \sigma_\eta)$ and $\eta_{T+h}^{(i)}$ is a draw from a $\mathcal{N}(0, \sigma_\eta^{2(i)})$. The point forecast $\hat{z}_{T+h|T}$ is obtained by averaging the $z_{T+h|T}^{(i)}$'s. While our draws from the posterior distribution of the DSGE model and auxiliary regression parameters are independent, we maintain some 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, $\varsigma_{T+h|T}^{(i)}$.

---

and linearly dependent, while s_t is assumed to be set of non-redundant states. To the extent that there exists a redundancy in ς_t the matrix $A_1(\theta)$ is not unique.

[6] Alternatively, we could generate a draw $\varsigma_{T|T}^{(i)}$ from $p(\varsigma_T|Y^T, \theta^{(i)})$.

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, and the unemployment rate. 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

We include seven series into 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 real wages, as well as the levels of hours worked, inflation, and the nominal interest rate. We obtain these series 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*), and deflating using 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. The labor share is 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.

We consider PCE-inflation, core PCE inflation, and the unemployment rate as candidates for z_t in this paper. Quarterly data on the chain price index for personal consumption

expenditures (JC) and personal consumption expenditures less food and energy ($JCXF$) were obtained from Haver Analytics. Inflation rates are calculated as 100 times the log difference of the series. The unemployment rate measure, also from Haver Analytics, is the civilian unemployment rate for ages 16 years and older (LR).

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) and 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 1.42%. According to our estimates of π_* and β the target inflation rate is 2.5% and the long-run nominal interest rate is 3.5%. The cross-equation restrictions of our model generate a nominal wage growth of about 4%.

Our policy rule estimates imply a strong response of the central bank to inflation $\hat{\psi}_1 = 2.99$ and a tempered reaction to deviations of output from its long-run growth path $\hat{\psi}_2 = 0.04$. 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 $\hat{\zeta}_p = 0.68$ 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 virtually no serial correlation and the estimated innovation standard deviation is about 0.8%. 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 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 that is used as regressor in the auxiliary model (21). Formally, we depict the 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.

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 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, consumption, and nominal wage growth obtained from the estimated DSGE model is slightly larger than the RMSE associated with the AR(1) forecasts. The DSGE model generates a lower RMSE for the investment forecasts. RMSEs for log hours and interest rates are essentially identical for the two models. Over a three-year forecast horizon, the DSGE model attains lower RMSEs than the AR(1) model for the interest rate, nominal wage growth, and hours worked. The AR(1) model does slightly better in forecasting the remaining four variables. The accuracy of long-run forecasts is sensitive to mean growth estimates, which are restricted to be equal for output, consumption, investment, and real wage growth in the DSGE model.

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.

---

[7]The h-step forecast is generated by iterating one-step ahead predictions forward, ignoring parameter uncertainty: $\hat{y}_{i,T+h|T} = \hat{\alpha}_{1,OLS} + \hat{\alpha}_{2,OLS}\hat{y}_{i,T+h-1|T}$.

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, and unemployment. The following elements are included in the vector s_t that appears as regressor in (21):

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

To construct a prior mean for α_1, we are linking each z_t variable with a conceptually related DSGE model variable $z_t^{\dagger}$ and use (24). More specifically, we link the two measures of PCE inflation to the final good inflation π_t and the unemployment rate to a scaled version of log hours worked, 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. The three panels of Figure 2 depict the sample paths of the non-core variables z_t and the elements of the vector of core variables y_t that are used as empirical measures of $z_t^{\dagger}$ in the DSGE model estimation: the GDP deflator and hours worked. The inflation measures are highly correlated. PCE inflation is more volatile and core PCE inflation is less volatile than GDP deflator inflation. In the third 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.

To proceed with the Bayesian estimation of (22) 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), and 0.40 (unemployment rate). These values imply that the prior variance of ξ_1 is about 15% to 20% of the sample variance of z_1. We let $\lambda_0 = \lambda_1$ 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.[8] 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

---

[8]In principle we could use marginal likelihood values to implement a data driven choice of the λ's. Instead, we decided to report the properties of our auxiliary regression model, including the pseudo-out-of-sample forecasting performance, for a variety of values.

construction, $\hat{s}'_{t|t}\mu_{\alpha_1,0}$ ($\mu_{\alpha_1,0}$ is the prior mean of α_1) reproduces the time paths of the GDP deflator inflation and log hours worked, respectively. Thus, for 1E-5 the idiosyncratic error term ξ_t essentially picks up the discrepancies between non-core variables and the related core variables depicted in Figure 2. For the two inflation series the estimate of σ_η increases as we lower the hyperparameter. The larger λ the better the in-sample fit of the auxiliary regression and the more of the variation in the variable is explained by $\hat{s}'_{t|t}\hat{\alpha}_1$. For instance, the variability of core PCE inflation captured by the factors is 5.24 times as large as the variability due to the idiosyncratic disturbance ξ_t if the λ's are equal to one. This factor drops to 1.36 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 0.22 (λ's are 1.00) to 0.54 (λ's are 1.E-5). The most striking feature of the unemployment estimates is the high persistence of ξ_t, with ρ_ξ estimates ranging from 0.96 to 0.98.

Figure 3 displays the time path of $\hat{\alpha}_{0,T} + \hat{s}'_{t|t}\hat{\alpha}_{1,T}$ for different choices of the hyperparameter, where $\hat{\alpha}_{i,T}$ is the posterior mean estimate of α_i. Consider the two inflation series. For $\lambda_0 = \lambda_1 = 1E-5$ the factor predicted path for the two inflation rates is essentially identical and reproduces the GDP deflator inflation. As the λ's are 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 the λ's to one, 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 the λ's 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.

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 those from an AR(1) model. For the inflation measures, decreasing the λ's improves the forecasts and we perform slightly better than an AR(1) model. For the unemployment rate, our auxiliary model yields to lower RMSEs than the AR(1) at all horizons and for all choices of the λ's. The best performance is attained for $\lambda_0 = \lambda_1 = 0.1$.

4.4 Multivariate Considerations

So far the analysis 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 two multivariate aspects. First, 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. Second, we examine some features of the predictive density that our empirical model generates for the core and non-core variables.

An important aspect of monetary policy making is to assess 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-modelled 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 are plotting impulse responses of three core variables (top panels: output, GDP inflation, interest rates) and the three non-core variables (bottom panels) 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 drops by 10 basis points in the first quarter and returns to its trend path after seven quarters. Quarter-to-quarter inflation also 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 and unemployment responses are more sensitive to the choice of hyperparameter. If the λ's are equal to 1E-5 and we force the factor loadings to match those of hours worked, the unemployment rises by about five basis points immediately after impact. As we relax the hyperparameter, which improves the RMSE of unemployment forecast, the unemployment response becomes more hump-shaped and the core PCE response drops from 10 basis points to about five basis points.

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 $y_t = [y_{1,t}, y_{2,t}]'$ be a 2×1 vector and consider the following two forecast models

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

Under a quadratic loss function the two models deliver identical univariate forecasts for each linear combination of the elements of y. 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 y:

$$LPSC(\mathcal{M}_i) = -\frac{H}{2}ln(2\pi) - \frac{H}{2}\ln|\Sigma_i| - \frac{1}{2}\sum_{h=1}^{H} x'_{T+h}\Sigma_i^{-1}x_{T+h}.$$

Roughly speaking, if the actual x_{T+h} 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 into a marginal and a conditional density, we can rewrite the

predictive score as

$$LPSC(\mathcal{M}_i) = -\frac{H}{2}ln(2\pi) - \frac{H}{2}\ln|\Sigma_{i,11}| - \frac{1}{2\Sigma_{i,11}}\sum_{h=1}^{H}x_{1,T+h}^2 \qquad (28)$$

$$-\frac{H}{2}\ln|\Sigma_{i,22|11}| - \frac{1}{2\Sigma_{i,22|11}}\sum_{h=1}^{H}\left(x_{2,T+h} - \Sigma_{i,21}\Sigma_{i,11}^{-1}x_{1,T+h}\right)^2,$$

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

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

Here the contribution of the marginal distribution of $y_{1,T+h}$ 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 H the predictive score will be negative if in fact the y's are generated from $\mathcal{M}_2$. In fact, the log score differentials has similar properties as a log likelihood ratio and is widely used in the prequential theory discussed in Dawid (1992). Moreover, notice that $\frac{1}{H}\sum_{h=1}^{H}(x_{2,T+h} - \rho x_{1,T+h})^2$ can be interpreted as the mean-squared-error of a forecast of x_2 conditional on the realization of x_1. If x_1 and x_2 have non-zero correlation, the conditioning improves the accuracy of the x_2 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 distribution by a normal distribution with mean μ and variance Σ and compute the prediction of a variable x_2 given the realization of x_1 from the conditional mean formula for a multivariate normal distribution:

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

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 (λ's are 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 (λ's are 0.1) 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 normal distribution is 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 7 and 8 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 normal distribution:

$$\sqrt{(\Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21})/\Sigma_{11}}.$$

Since the covariance matrix of the predictive distribution changes over time, we compute averages of the theoretical RMSE reduction.

The results obtained when conditioning on the interest rate, reported in Table 7, are somewhat disappointing. Although except for the unemployment rate, the bivariate correlations between the interest rate and the other variables are non-zero and would imply a potential RMSE reduction between 2% and 12%, the RMSE obtained from the conditional forecasts exceeds that from the unconditional forecasts.[9] If we condition on the realization

---

[9] 2001:IV and 2006:III are not representative, since conditioning in these periods leads to a reduction of the forecast error.

of the GDP deflator inflation (Table 8), then the results improve and we observe a RMSE reduction. For output growth, hours worked, PCE inflation, and the unemployment rate, the actual RMSE ratios appear to be broadly in line with those predicted from a bivariate normal distribution. The joint distribution of core PCE inflation and the GDP deflator exhibits are strong correlation as is evident from the fairly large RMSE reduction factor. The actual RMSE ratio tends to be much larger than the one predicted by the normal distribution and it is greater than one for multi-step forecasts.

The results reported in this section have to be interpreted carefully. First, 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 more successful capturing the uncertainty about future pairwise realizations. We think that our statistics are a useful addition to the univariate forecast accuracy measures that have been reported for DSGE models. It might be worthwhile to consider non-parametric approximations to the predictive distribution in future work.

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 Boivin and Giannoni (2006). 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.

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.

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)," *NBER Working Paper No. 13741.*

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.

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.

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.

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	0.67	[0.61 , 0.74]
a''	$\mathcal{G}$	0.20	0.10	0.32	[0.15 , 0.48]
ν_l	$\mathcal{G}$	2.00	0.75	2.29	[1.33 , 3.16]
ζ_w	$\mathcal{B}$	0.60	0.20	0.25	[0.15 , 0.35]
$400 * (1/\beta - 1)$	$\mathcal{G}$	2.00	1.00	0.998	[0.44 , 1.54]
Firms					
α	$\mathcal{B}$	0.33	0.10	0.21	[0.16 , 0.27]
ζ_p	$\mathcal{B}$	0.60	0.20	0.68	[0.58 , 0.85]
S''	$\mathcal{G}$	4.00	1.50	2.44	[0.99 , 3.83]
λ_f	$\mathcal{G}$	0.15	0.10	0.19	[0.02 , 0.34]
Monetary Policy					
$400\pi^*$	$\mathcal{N}$	3.00	1.50	2.47	[2.25 , 2.71]
ψ_1	$\mathcal{G}$	1.50	0.40	2.99	[2.29 , 3.69]
ψ_2	$\mathcal{G}$	0.20	0.10	0.04	[0.01 , 0.07]
ρ_R	$\mathcal{B}$	0.50	0.20	0.87	[0.84 , 0.90]

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	1.42	[0.94 , 1.89]
g^*	$\mathcal{G}$	0.30	0.10	0.27	[0.12 , 0.41]
ρ_a	$\mathcal{B}$	0.20	0.10	0.09	[0.02 , 0.15]
ρ_μ	$\mathcal{B}$	0.80	0.05	0.79	[0.73 , 0.86]
ρ_{λ_f}	$\mathcal{B}$	0.60	0.20	0.78	[0.43 , 0.98]
ρ_g	$\mathcal{B}$	0.80	0.05	0.96	[0.95 , 0.98]
ρ_b	$\mathcal{B}$	0.60	0.20	0.86	[0.80 , 0.93]
ρ_ϕ	$\mathcal{B}$	0.60	0.20	0.98	[0.96 , 0.99]
σ_a	$\mathcal{IG}$	0.75	2.00	0.78	[0.68 , 0.89]
σ_μ	$\mathcal{IG}$	0.75	2.00	0.51	[0.38 , 0.64]
σ_{λ_f}	$\mathcal{IG}$	0.75	2.00	0.17	[0.14 , 0.19]
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 , 0.44]
σ_ϕ	$\mathcal{IG}$	4.00	2.00	3.13	[2.15 , 4.07]
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	DSGE	0.51	0.51	**0.41**	0.39
	AR(1)	0.50	0.49	0.44	0.37
Consumption Growth	DSGE	0.39	0.39	0.42	0.45
	AR(1)	0.37	0.37	0.34	0.31
Investment Growth	DSGE	**1.46**	**1.55**	**1.31***	**1.53**
	AR(1)	1.56	1.67	1.60	1.60
Nominal Wage Growth	DSGE	0.65	0.67*	0.61	**0.54**
	AR(1)	0.59	0.59	0.59	0.56
100× Log Hours	DSGE	**0.57**	**0.99**	**1.67**	**2.00**
	AR(1)	0.66	1.20	2.08	3.40
Inflation	DSGE	0.25	0.26	0.24	0.28
	AR(1)	0.22	0.23	0.22	0.23
Interest Rates	DSGE	**0.54**	**0.97**	**1.58**	**2.07**
	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	Inflation	Interest Rate
		(Q %)	(Q %)	(A %)
Schorfheide, Sill, Kryshko	2001:I to 2007:IV	0.51	0.25	0.55
		(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, Krysho: 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	$-(100/3)(L_t - \bar{L})$

Table 5: Auxiliary Regression Estimates

Series	(λ_0, λ_1)	ρ Mean	ρ 90% Intv	σ_η Mean	σ_η 90% Intv	Signal/Noise $\frac{v\hat{a}r(\hat{s}'_{t|t}\hat{\alpha}_1)}{v\hat{a}r(\hat{\xi}_t)}$
PCE Inflation	(1.00, 1.00)	0.09	[-0.12, 0.30]	0.03	[0.02, 0.04]	3.03
	(0.10, 0.10)	0.10	[-0.14, 0.30]	0.03	[0.02, 0.04]	2.53
	(1E-5, 1E-5)	0.06	[-0.11, 0.24]	0.04	[0.03, 0.05]	1.48
Core PCE Inflation	(1.00, 1.00)	0.22	[0.02, 0.45]	0.01	[0.01, 0.02]	5.24
	(0.10, 0.10)	0.22	[-0.01, 0.44]	0.01	[0.01, 0.02]	5.11
	(1E-5, 1E-5)	0.54	[0.37, 0.67]	0.03	[0.03, 0.04]	1.36
Unemployment Rate	(1.00, 1.00)	0.98	[0.94, 1.00]	0.02	[0.01, 0.03]	2.88
	(0.10, 0.10)	0.96	[0.93, 1.00]	0.02	[0.01, 0.03]	3.24
	(1E-5, 1E-5)	0.98	[0.97, 1.00]	0.04	[0.03, 0.05]	1.94

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: RMSE Comparisons: Auxiliary Regressions versus AR(1)

Series	Model	(λ_0, λ_1)	$h = 1$	$h = 2$	$h = 4$	$h = 12$
PCE Inflation	Aux	(1.00, 1.00)	**0.36**	0.40	0.39	0.36
	Aux	(0.10, 0.10)	0.37	0.41	0.39	0.40
	Aux	(1E-5, 1E-5)	**0.32**	0.35	0.33	0.36
	AR(1)		0.36	0.35	0.33	0.32
Core PCE Inflation	Aux	(1.00, 1.00)	0.21*	0.22**	0.19	**0.17**
	Aux	(0.10, 0.10)	0.22*	0.22**	0.19	**0.14**
	Aux	(1E-5, 1E-5)	0.16	0.16	**0.15**	**0.11****
	AR(1)		0.16	0.16	0.18	0.17
Unemployment Rate	Aux	(1.00, 1.00)	**0.16****	**0.26**	**0.46**	**0.74***
	Aux	(0.10, 0.10)	**0.15****	**0.25**	**0.45**	**0.64**
	Aux	(1E-5, 1E-5)	**0.18**	**0.29**	**0.47**	**0.74**
	AR(1)		0.21	0.37	0.63	1.00

Notes: See Table 2.

Table 7: RMSE Ratios: Conditional (on Interest Rates) versus Unconditional

| Series	(λ_0, λ_1)		$h = 1$	$h = 2$	$h = 4$	$h = 12$
Output Growth		Actual	1.10	1.19	1.28	1.02
		(Theory)	(0.91)	(0.88)	(0.89)	(0.90)
100× Log Hours		Actual	1.14	1.30	1.43	1.65
		(Theory)	(0.98)	(0.96)	(0.95)	(0.96)
Inflation		Actual	1.20	1.29	1.50	1.70
		(Theory)	(0.84)	(0.84)	(0.86)	(0.89)
PCE Inflation	(1E-5, 1E-5)	Actual	1.10	1.15	1.32	1.44
		(Theory)	(0.90)	(0.90)	(0.90)	(0.91)
Core PCE Inflation	(1E-5, 1E-5)	Actual	1.15	1.55	2.04	3.08
		(Theory)	(0.88)	(0.89)	(0.90)	(0.91)
Unemployment Rate	(0.10, 0.10)	Actual	1.04	1.08	1.25	1.12
		(Theory)	(0.99)	(0.99)	(0.98)	(0.98)

Notes: Using the draws from the posterior predictive distribution of two variables x_1 and x_2 we construct a Gaussian approximation with means μ_1 and μ_2 and covariances Σ_{11}, Σ_{12}, and Σ_{22}. The conditional forecast of x_1 given x_2 is $\mu_1 + \Sigma_{12}\Sigma_{22}^{-1}(x_2 - \mu_2)$. and the unconditional forecast is μ_1. The theoretical RMSE ratio is $\sqrt{(\Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21})/\Sigma_{11}}$.

Table 8: RMSE Ratios: Conditional (on GDP Deflator Inflation) versus Unconditional

Series	(λ_0, λ_1)		$h = 1$	$h = 2$	$h = 4$	$h = 12$
Output Growth		Actual	0.93	0.88	0.99	1.18
		(Theory)	(0.92)	(0.93)	(0.93)	(0.94)
$100\times$ Log Hours		Actual	0.95	0.96	0.99	0.92
		(Theory)	(0.99)	(0.99)	(0.99)	(0.98)
PCE Inflation	(1E-5, 1E-5)	Actual	0.79	0.76	0.83	0.75
		(Theory)	(0.69)	(0.65)	(0.64)	(0.67)
Core PCE Inflation	(1E-5, 1E-5)	Actual	0.94	1.48	1.54	2.16
		(Theory)	(0.59)	(0.63)	(0.64)	(0.68)
Unemployment Rate	(0.10, 0.10)	Actual	1.03	1.01	0.99	1.05
		(Theory)	(1.00)	(1.00)	(1.00)	(0.99)

Notes: See Table 7.

Figure 1: LATENT STATE VARIABLES OF THE DSGE MODEL
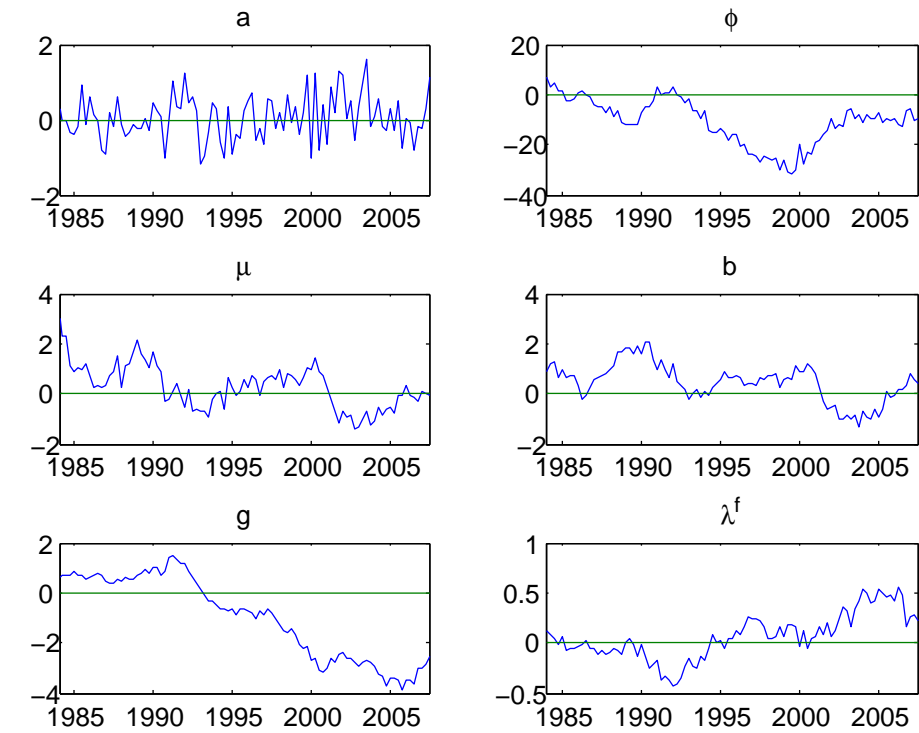
Notes: The six panels of the figure depict time series of elements of $\hat{s}_{t|t}$. Estimation sample: 1984:I to 2007:III.

Figure 2: Non-Core Variables and Related Model Variables
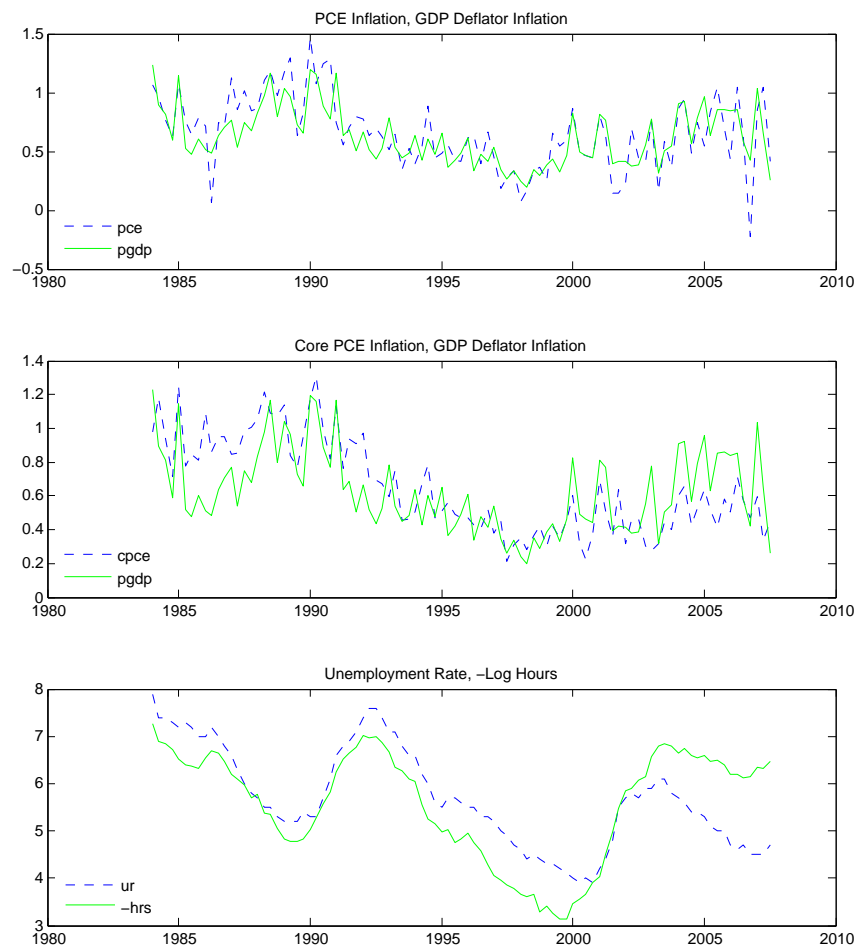
Notes: The top two panels depict quarter-to-quarter inflation rates. In the third panel we re-scale the log of hours worked by a factor of -(100/3) and add a constant to match the mean of unemployment over the period 1984:I to 2007:III.

Figure 3: Non-Core Variables and Factors
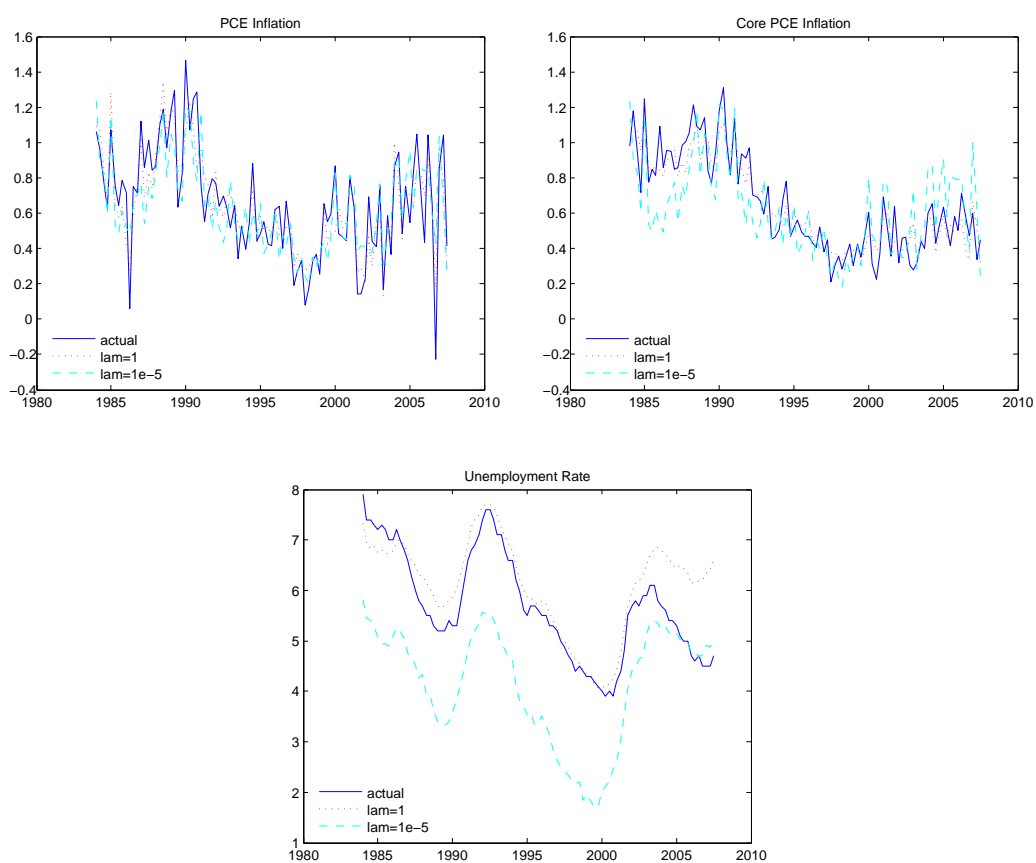
Notes: Figure depicts the actual (blue, solid) path of the non-core variables as well as the factor predictions $\hat{\alpha}_0 + \hat{s}'_{t|t}\hat{\alpha}_{1,T}$ for $\lambda = $ 1E-5 (light blue, dashed) and $\lambda = 1$ (red, dotted).

Figure 4: Impulses Response to a Monetary Policy Shock

Core Variables: Output, GDP Deflator Inflation, Interest Rates

Non-modelled Variables: PCE Inflation, core PCE Inflation, Unemployment

Notes: For the 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
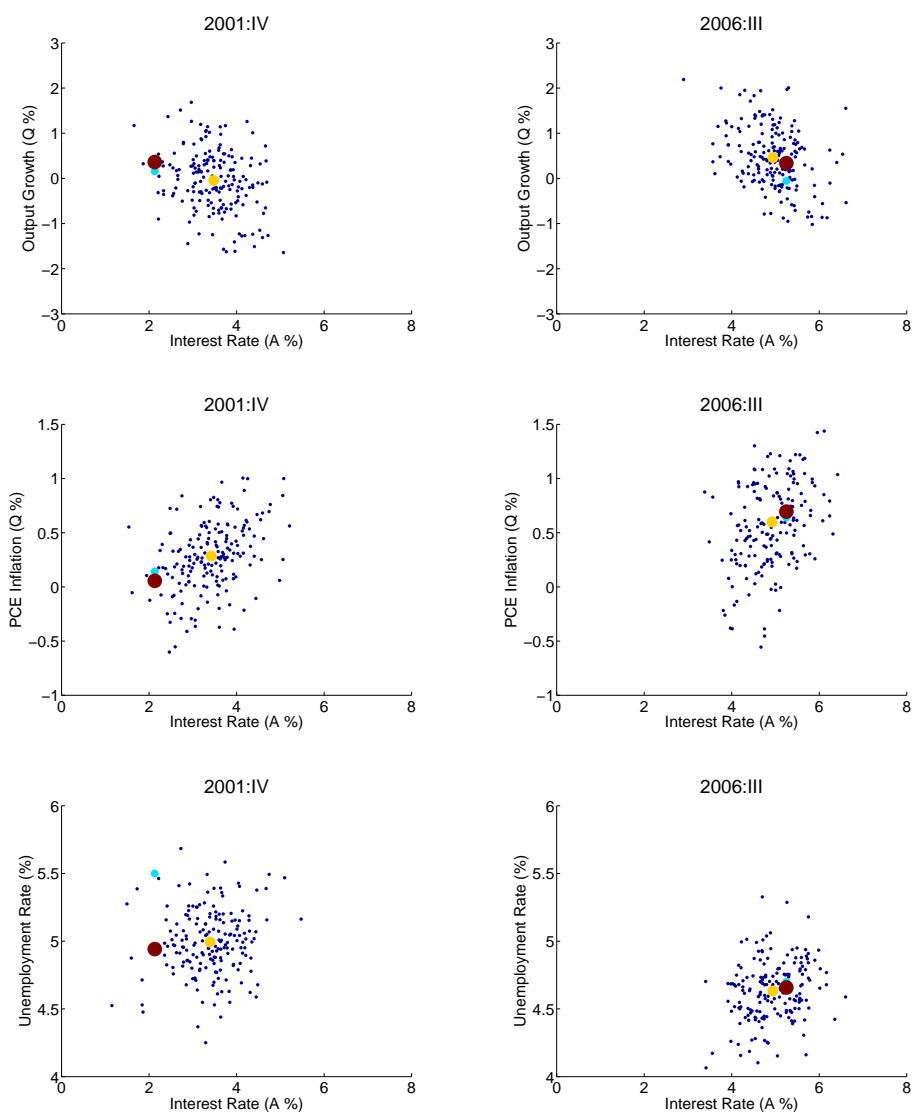
Notes: The panels depict a scatter plot of draws from the one-step-ahead predictive distri-
bution. The three filled circles denote: the actual value (small, light blue), the unconditional
mean predictor (medium, yellow), and the conditional mean predictor (large, brown). For
PCE inflation we use $\lambda_0 = \lambda_1 = 1\text{E-}5$ and for the unemployment rate we let $\lambda_0 = \lambda_1 = 0.1$.

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. We will start from the quasi-differenced form (22) 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.

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

$$p(\alpha|\rho, \sigma_\eta^2, Z^T, S^T) \propto p(Z^T|S^T, \alpha, \rho, \sigma_\eta^2)p(\alpha).$$

Define

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

We can now write (3.3) as linear regression

$$y_t = x_t'\alpha + \eta_t.$$

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

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

From this linear regression model we obtain

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

where

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

We deduce that 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: The posterior density is of the form

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

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 (3.3) as linear regression model

$$y_t = x_t\rho + \eta_t.$$

Using the same arguments as before we deduce that

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

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: The posterior density is of the form

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

Solve (3.3) for η_t:

$$\eta_t = z_t - \left[\rho z_{t-1} + \alpha_0(1 - \rho) + [\hat{s}'_{t|t} - \hat{s}'_{t-1|t-1}\rho]\alpha_1\right].$$

Now, notice that

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

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 $\tilde{\sigma}_\eta^2 = \left[\sum_{j=1}^T Z_j^2\right]^{-1}$.