Estimation and Evaluation of DSGE Models: Progress and Challenges *

Frank Schorfheide

University of Pennsylvania, CEPR, and NBER

December 27, 2010

*Correspondence: F. Schorfheide: Department of Economics, 3718 Locust Walk, University of Pennsylvania, Philadelphia, PA 19104. Email: schorf@ssc.upenn.edu. This paper is based on a lecture delivered at the 2010 Econometric Society World Congress in Shanghai. It draws extensively from joint work with Borağan Aruoba and Marco Del Negro. I would like to thank Harald Uhlig and participants of the World Congress session on Macroeconometrics for helpful comments. Financial support from the National Science Foundation under grant SES-0617803 is gratefully acknowledged. Gauss Programs to replicate the empirical analysis as well as a Technical Appendix are available at *http://www.ssc.upenn.edu/˜schorf*.

1 Introduction

Estimated dynamic stochastic equilibrium (DSGE) models are now widely used for empirical research in macroeconomics as well as quantitative policy analysis and prediction at central banks around the world. This chapter summarizes recent advances in the econometric analysis of DSGE models, discusses current challenges, and highlights avenues for future research. To illustrate the advances of the past decade, a prototypical empirical analysis based on a small-scale New Keynesian DSGE model is presented. In this application, Bayesian inference is used to measure the welfare effect of changing the central bank's target inflation rate. The Bayesian inference is implemented through Markov-Chain Monte Carlo (MCMC) methods, which deliver draws from the posterior distribution of DSGE model parameters. These draws can be converted into impulse response functions, welfare effects of policy changes, or other quantities of interest and it is straightforward to obtain point estimates, e.g. posterior means or medians, or credible sets that reflect the posterior uncertainty.

Despite the advances of the DSGE model literature, many important challenges remain. This chapter considers five of them in detail. First, while reported credible (or confidence) sets for DSGE model parameters are often narrow, from a meta perspective estimates of many important parameters tend to be fragile across empirical studies. Second, macroeconomic fluctuations in DSGE models are generated by exogenous disturbances. The estimated shock processes are often highly persistent and their path closely mirrors the path of one of the observables. This raises concerns as to whether these shocks capture aggregate uncertainty or misspecification. Third, many time series exhibit low frequency behavior that is difficult if not impossible to reconcile with the model being estimated. This low frequency misspecification contaminates the estimation of shocks and thereby inference about the sources of business cycle fluctuations. Fourth, in view of more densely parameterized empirical models such as vector autoregressions (VARs), DSGE models often appear to be misspecified in the sense that VARs are favored by statistical criteria that trade off goodness of in-sample fit against model dimensionality. Fifth, the prediction of the effects of rare policy changes often relies exclusively on extrapolation by theory which makes it difficult to provide measures of uncertainty.

This chapter is organized as follows. The above-mentioned advances and challenges are illustrated in Section 2. The remaining sections discuss recent research that addresses some of these problems, including work on identification of DSGE models (Section 3), the generalization of exogenous shock processes (Section 4), methods to construct hybrid models that correct DSGE model misspecification (Section 5), and methods to conduct policy analysis (Section 6). Finally, Section 7 concludes. Details of the empirical illustrations and examples that are presented in this chapter are relegated to a Technical Appendix that is available electronically.

2 A Prototypical Application

A central element of New Keynesian DSGE models is that firms face a cost of adjusting nominal prices. In turn, firms tend to economize on price adjustments if inflation is non-zero. This leads to a distortion of relative prices and an inefficient use of intermediate inputs and ultimately to output and welfare loss. At the same time, non-zero nominal interest rates constitute a tax on money holdings and depress transactions that require the use of money or highly liquid, non-interest bearing funds. These two mechanisms create a trade-off for policy makers. The New Keynesian friction is eliminated by targeting a zero inflation rate which equates nominal and real interest rate. The monetary friction, on the other hand, is eliminated if the nominal interest rate is zero. A DSGE model can be used to estimate the relative strength of the two frictions and to determine a long-run inflation rate that trades-off the opposing mechanisms. The following illustration is based on recent work by Aruoba and Schorfheide (2010), henceforth AS. Section 2.1 provides a description of the model economy. Section 2.2 discusses estimation results that highlight some of the recent advances in the econometric analysis of DSGE models. Finally, Section 2.3 points toward a number of problems and challenges that need to be addressed in future research.

2.1 A Small-Scale DSGE Model

The model economy consists of households, final good producers, intermediate goods producers, a central bank and a fiscal authority. It is a simplified version of the widely cited Smets and Wouters (2003, 2007) model because it abstracts from habit formation in consumption and wage rigidity. I will subsequently provide a brief description of the agents' decision problems, the aggregate resource constraint, and the exogenous shock processes.

Households. The economy is populated by a continuum of identical households. These households take as given the aggregate price level P_t, the gross nominal interest rate R_t on one-period bonds, the wage W_t, the rental rate of capital, R_t^k, and the set of aggregate shocks $\mathcal{S}_t$, along with their laws of motion. The households maximize

$$\mathbb{E}_\tau \left[\sum_{t=\tau}^{\infty} \beta^{(t-\tau)} \left\{ U(C_t) - AH_t + \frac{\chi_t}{1-\nu} \left(\frac{M_t}{P_t} \frac{A}{Z_*^{1/1-\alpha}} \right)^{1-\nu} \right\} \right] \tag{1}$$

subject to the constraints:

$$\begin{aligned} P_t C_t + P_t I_t + B_{t+1} + M_{t+1} &\leq P_t W_t H_t + P_t R_t^k K_t + \Pi_t + R_{t-1} b_t + M_t - T_t + \Omega_t && (2) \\ K_{t+1} &= (1-\delta) K_t + \left[1 - S\left(\frac{I_t}{I_{t-1}} \right) \right] I_t. && (3) \end{aligned}$$

$U(C_t)$ is the instantaneous utility from consuming C_t units of the final good, A is the disutility associated with one unit of labor, H_t is hours worked, and M_t denotes the households' money holdings at the beginning of period t. The assumption of quasi-linear preferences can be motivated by the indivisible labor setup of Rogerson (1988) and is used for convenience in many of the New Keynesian models discussed in Woodford (2003). Money balances enter the utility function to capture the benefits of transaction services. The shock χ_t captures time-varying preferences for money and the parameter ν controls the interest-rate elasticity of money demand.[1]

[1] AS develop an estimable search-based monetary DSGE model, in which money is essential to facilitate bilateral exchanges in a decentralized market. The reduced form specification considered in this chapter serves as a reference model in AS to assess the fit of the search-based DSGE model. The factor $A/Z_*^{1/(1-\alpha)}$ in the utility function can be viewed as a re-parameterization of the steady state level of χ_t that keeps steady state velocity constant as one changes the preference parameter A and the steady state level of technology Z_* (introduced below).

Equation (2) represents the households' budget constraint. Final goods are purchased at the price P_t and used for consumption and investment I_t. The household receives labor income, rental income from lending capital K_t to firms, interest income from bond holdings B_t, and dividends Π_t from intermediate goods producers. T_t is a nominal lump-sum tax and Ω_t is the household's net cash-in-flow from trading state-contingent securities. Equation (3) determines the capital accumulation. The adjustment cost function $S(.)$ satisfies the properties $S(1) = 0$, $S'(1) = 0$ and $S''(1) > 0$. I adopt the timing convention that K_{t+1} (and also M_{t+1}) denote capital and money holdings at the end of period t and do not depend on period $t+1$ shocks.

Final Good Production. The final good Y_t is a composite made of a continuum of intermediate goods $Y_t(i)$:

$$Y_t = \left[\int_0^1 Y_t(i)^{\frac{1}{1+\lambda}} di \right]^{1+\lambda} \tag{4}$$

with elasticity of substitution $(1+\lambda)/\lambda$, where $\lambda \in [0, \infty)$. The final good producers buy the intermediate goods on the market, package them into Y_t units of the composite good, and resell them to consumers. These firms maximize profits in a perfectly competitive environment taking $P_t(i)$ as given, which yields the demand for good i

$$Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\frac{1+\lambda}{\lambda}} Y_t. \tag{5}$$

Combining this demand function with the zero profit condition one obtains the following expression for the price of the composite good

$$P_t = \left[\int_0^1 P_t(i)^{-\frac{1}{\lambda}} di \right]^{-\lambda}. \tag{6}$$

Aggregate inflation is defined as $\pi_t = P_t / P_{t-1}$.

Intermediate Goods Production. Intermediate goods producers, indexed by i, face the demand function (5) and use a Cobb-Douglas technology with fixed costs $\mathcal{F}$ and stochastic total factor productivity Z_t:

$$Y_t(i) = \max \left\{ Z_t K_t(i)^{\alpha} H_t(i)^{1-\alpha} - \mathcal{F}, 0 \right\}. \tag{7}$$

Following Calvo (1983), it is assumed that firms are only able with probability $1-\zeta$ to re-optimize their price in the current period. A random fraction ι of the firms that are not allowed to choose $P_t(i)$ optimally update their price $P_{t-1}(i)$ according to last period's inflation rate π_{t-1}, whereas the remaining $1-\iota$ firms keep their price constant. For a firm that is allowed to re-optimize its price, the problem is to choose a price level $P_t^o(i)$ that maximizes the expected present discounted value of profits in all future states in which the firm is unable to re-optimize its price. This firm uses the time t value of a dollar in period $t+s$ for the consumers, to discount future profits. The solution of this problem leads to a dynamic relationship between inflation and marginal costs, the so-called New Keynesian Phillips Curve (NKPC).

Government Spending. In period t, the government collects a nominal lump-sum tax T_t, spends G_t on final good purchases, issues one-period nominal bonds B_{t+1} that pay R_t gross interest tomorrow and supplies the money to maintain the interest rate rule. It satisfies the following budget constraint every period

$$P_t G_t + R_{t-1} B_t + M_t = T_t + B_{t+1} + M_{t+1}. \tag{8}$$

Government spending G_t is assumed to evolve exogenously.

Aggregate Resource Constraint. Adding the households' budget constraints, the government budget constraint and the profits of intermediate goods producers yields the aggregate resource constraint

$$C_t + I_t + G_t = Y_t. \tag{9}$$

The quantity of final goods is related to the total inputs used by the intermediate goods firms according to

$$Y_t = \frac{1}{D_t}\left[Z_t K_t^{\alpha} H_t^{1-\alpha} - \mathcal{F}\right], \quad D_t = \int \left(\frac{P_t(i)}{P_t}\right)^{-\frac{1+\lambda}{\lambda}} di, \tag{10}$$

where D_t measures the extent of price dispersion. Unless $P_t(i) = P_t$ for all firms, D_t is greater than unity, which in turn implies the economy produces inside its production-possibility frontier. D_t captures the output loss due to the New Keynesian friction.

Monetary Policy. Following authors like Sargent (1999) and Lucas (2000), I assume that low frequency movements of inflation, such as the rise of inflation in the 1970s and the subsequent disinflation episode in the early 1980s, can be attributed to monetary policy changes. Unlike in the learning models considered by Sargent, Zha, and Williams (2006) or Primiceri (2006), in this chapter the DSGE model offers no explanation why monetary policy shifts occur over time and simply assumes a time-varying target inflation rate $\pi_{*,t}$. The central bank supplies money to control the nominal interest rate and reacts to inflation and output growth according to the rule

$$R_t = R_{*,t}^{1-\rho_R} R_{t-1}^{\rho_R} \exp[\sigma_R \epsilon_{R,t}], \quad R_{*,t} = (r_* \pi_{*,t}) \left(\frac{\pi_t}{\pi_{*,t}} \right)^{\psi_1} \left(\frac{Y_t}{\gamma Y_{t-1}} \right)^{\psi_2}, \tag{11}$$

where r_* is the steady state real interest rate, γ is the gross steady state growth rate of the economy, and $\epsilon_{R,t}$ is a monetary policy shock.

Exogenous Shocks. The model economy is subjected to five aggregate disturbances. Z_t is the stochastic total factor productivity process. g_t is a shock that shifts government spending according to

$$G_t = (1 - 1/g_t)\, Y_t. \tag{12}$$

The money demand shock χ_t shifts preferences for real money balances. Finally, the model has two monetary policy shocks: $\epsilon_{R,t}$ is assumed to be serially uncorrelated and captures short-run shifts in monetary policy, whereas the time-varying inflation target $\pi_{*,t}$ captures long-run policy changes. Let $\tilde{Z}_t = \ln(Z_t/Z_*)$, $\tilde{\chi}_t = \ln(\chi_t/\chi_*)$ and $\tilde{g}_t = \ln(g_t/g_*)$, where Z_*, χ_* and g_* are steady state values of the respective exogenous disturbances. It is assumed that these exogenous disturbances evolve according to stationary AR(1) processes $\tilde{Z}_t = \rho_z \tilde{Z}_{t-1} + \sigma_z \epsilon_{z,t}$, $\tilde{\chi}_t = \rho_\chi \tilde{\chi}_{t-1} + \sigma_\chi \epsilon^{\chi}_{\chi,t}$ and $\tilde{g}_t = \rho_g \tilde{g}_{t-1} + \sigma_g \epsilon_{g,t}$. Finally, let $\tilde{\pi}_{*,t} = \ln(\pi_{*,t}/\pi_*)$, where π_* is a constant and $\tilde{\pi}_{*,t}$ evolves as a random walk $\tilde{\pi}_{*,t} = \tilde{\pi}_{*,t-1} + \sigma_\pi \epsilon^{\pi}_{\pi,t}$. The innovations are stacked in the vector $\epsilon_t = [\epsilon_{z,t}, \epsilon_{\chi,t}, \epsilon_{g,t}, \epsilon_{\pi,t}, \epsilon_{R,t}]$ and are assumed to be independently and identically distributed according to a vector of standard normal random variables. The law of motion for the exogenous processes completes the specification of the DSGE model.

State-Space Representation. After log-linearizing the equilibrium conditions of the model the solution of the resulting rational expectations difference equations leads to a state-space representation of the form

$$\begin{aligned} y_t &= \Psi_0(\theta) + \Psi_1(\theta)t + \Psi_s(\theta)s_t \\ s_t &= \Phi_1(\theta)s_{t-1} + \Phi_\epsilon(\theta)\epsilon_t, \end{aligned} \tag{13}$$

where y_t is a vector of observables, such as aggregate output, inflation, and interest rates, s_t contains the unobserved exogenous shock processes as well as the potentially unobserved endogenous state variables of the model economy. The model specification is completed by making a distributional assumption for the vector of innovations ϵ_t and the initial state vector s_0.

2.2 What has the DSGE Model Estimation Literature Delivered?

The goal of the DSGE model estimation literature is to provide quantitative answers to macroeconomic questions as well as probabilistic measures of uncertainty associated with these answers. As an illustration, I will the DSGE model described in Section 2.1 to assess the welfare effects of changes in the target inflation rate. The model is estimated with U.S. data from 1965 to 2005 on linearly detrended log GDP, interest rates, GDP deflator inflation, log inverse M1-velocity, and an empirical measure of the target inflation rate that is constructed from bandpass filtered inflation and long-run inflation expectations. Following the empirical strategy in Aruoba and Schorfheide (2010), the target inflation rate is treated as an observed variable such that it becomes possible to assess the time series fit of the DSGE model and the propagation of unanticipated changes in the target inflation rate through a comparison with a VAR. Except for the use of an observable measure of the target inflation rate, the empirical illustration is representative of the large literature on estimated DSGE models that has emerged recently.

Bayesian Inference. While over the past decades numerous econometric procedures for the analysis of DSGE models have been developed[2], I will focus on Bayesian inference techniques that have first been used in the context of DSGE model estimation in DeJong, Ingram, and Whiteman (2000), Schorfheide (2000), and Otrok (2001) and are by now widely applied in the literature. Let θ denote the collection of parameters of the DSGE model described in Section 2.1. Bayesian inference starts from a prior distribution represented by the density $p(\theta)$. The prior is combined with the conditional density of the data Y given the parameters,

[2] The textbooks by Canova (2007) and DeJong and Dave (2007) provide a detailed overview.

denoted by $p(Y|\theta)$. According to Bayes Theorem the posterior distribution, that is the conditional distribution of parameters given data, is given by

$$p(\theta|Y) = \frac{p(Y|\theta)p(\theta)}{p(Y)}, \quad p(Y) = \int p(Y|\theta)p(\theta)d\theta, \tag{14}$$

where $p(Y)$ is called the marginal likelihood or data density. In DSGE model applications it is typically not possible to derive moments and quantiles of the posterior distribution analytically. Instead, inference is implemented via numerical methods such as MCMC simulation. MCMC algorithms deliver serially correlated sequences $\{\theta_{(s)}\}_{s=1}^{n_{sim}}$ of draws from the density $p(\theta|Y)$. Based on these draws one can approximate the posterior density, its moments and quantiles, and for instance construct credible sets. In addition, the sequence $\{\theta_{(s)}\}_{s=1}^{n_{sim}}$ can be transformed into a sequence $\{f(\theta_{(s)})\}_{s=1}^{n_{sim}}$ to characterize the posterior distribution of $f(\theta)$, where $f(\theta)$ could be a set of steady states or impulse response functions computed from the DSGE model. A more detailed discussion of numerical techniques to implement Bayesian inference for DSGE models can be found, for instance, in An and Schorfheide (2007a) and Del Negro and Schorfheide (2010).

Estimates of Parameters and Transformations Thereof. The output and inflation trade-off faced by a central bank is determined by the NKPC, which for values of the target inflation rate near zero can be approximated as follows:

$$\tilde{\pi}_t = \gamma_b \tilde{\pi}_{t-1} + \gamma_f \mathbb{E}_t[\tilde{\pi}_{t+1}] + \kappa \widetilde{MC}_t, \tag{15}$$

where

$$\gamma_b = \frac{\iota}{1+\beta\iota}, \quad \gamma_f = \frac{\beta}{1+\beta\iota}, \quad \text{and} \quad \kappa = \frac{(1-\zeta)(1-\zeta\beta)}{\zeta(1+\beta\iota)}.$$

$\tilde{x}_t$ denotes percentage deviations from the log-linearization point $\ln(x_t/x_*)$ and MC_t abbreviates the marginal cost of producing an additional unit of the intermediate good. Posterior and prior densities for the coefficient on marginal costs κ and lagged inflation γ_b are depicted in the top panels of Figure 1. The posterior density of κ peaks at about 0.08 and the posterior of γ_b peaks at 0.03, implying that the influence of the lagged inflation term in the NKPC is essentially negligible. The posterior densities reflect the sample information and turn out to be much more concentrated than the prior densities.

The bottom left panel depicts densities for the percentage loss $100|1/D_* - 1|$ in output caused by the inability of a fraction of intermediate goods producers to choose their prices optimally. D_* depends on the steady state mark-up controlled by λ as well as the price setting parameters ζ and ι:

$$D_* = \frac{(1-\zeta)\left[\frac{1}{1-\zeta} - \frac{\zeta}{1-\zeta}\left(\frac{1}{\pi_*}\right)^{-\frac{1-\iota}{\lambda}}\right]^{1+\lambda}}{1-\zeta\pi_*^{\frac{(1+\lambda)(1-\iota)}{\lambda}}}. \tag{16}$$

It can be verified that D_* is bounded below from one. This lower bound is attained if prices are flexible ($\zeta = 0$), all firms that are unable to re-optimize fully index their old prices to inflation ($\iota = 1$), or if the steady state inflation rate is zero (meaning that the gross inflation $\pi_* = 1$). The posterior estimate of the output loss due to the New Keynesian distortion is about 0.6%. Interestingly, it is the combination of modeling assumption about the substitutability of intermediate inputs with information about the correlation between inflation and a measure of aggregate marginal costs that delivers the output loss estimate.[3] The prior density of D_* peaks at about 0.1, because the prior distribution places more weight on large values of γ_b, which imply the New Keynesian friction is reduced by the firms' dynamic indexation.

At last, the bottom right panel shows densities for the interest rate coefficient $1/(\nu(R_* - 1))$ in the log-linearized demand equation for real money balances at the end of period t:

$$\tilde{\mathcal{M}}_{t+1} = -\frac{1}{\nu(R_*-1)}\tilde{R}_t + \frac{\gamma}{\nu}\tilde{X}_t - \frac{1-\nu}{\nu}\mathbb{E}[\tilde{\pi}_{t+1}] + \mathbb{E}[\tilde{\chi}_{t+1}], \tag{17}$$

where R_* is the steady state nominal interest rate. The (partial) interest-rate elasticity of money demand indirectly affects the welfare costs induced by taxing money balances via inflation. In a Bayesian framework the posterior densities plotted in Figure 1 provide a formal characterization of parameter uncertainty. Point and interval estimates can be derived as solutions to decision problems that entail the minimization of posterior expected losses. The most widely used point estimates are the posterior mean and median, and the so-called highest-posterior density interval is the shortest interval among all (including disconnected) intervals that are $1-\alpha$ credible, i.e. have posterior probability $1-\alpha$.

[3] Alternatively, many authors use the frequency of price changes observed in micro-level data sets to determine ζ and hence the magnitude of the aggregate distortion D_*.

Policy Analysis. What are the relative strengths of the monetary and New Keynesian friction and what rate of long-run inflation optimizes the trade-off between these to frictions? The results obtained from the estimated model are depicted in Figure 2. Each line in the left panel of the figure represents the (steady state) welfare loss function for a particular draw of θ from its posterior distribution. The loss is expressed in terms of consumption equivalents relative to a 2.5% (annualized) target inflation rate. Negative values imply welfare gains. The right panel contains (pointwise) posterior means and 90% credible intervals for these losses. The welfare gain is maximized at an inflation rate of near zero, meaning that the New Keynesian friction dominates the policy recommendation.

Summary. The empirical illustration suggests that econometricians have developed a powerful toolkit that enables an elegant econometric analysis of DSGE models. The strengths of the formal econometric analysis are its ability to efficiently extract information about parameters from long-run averages and sample autocovariances of macroeconomic time series and to account for parameter (and model) uncertainty in inference and decision making. Researchers have made widely use of these strengths. The Bayesian approach has the additional advantage that it allows the researcher to coherently combine sample information (contained in the likelihood function) with non-sample information represented by prior distributions. There exist many published papers that to varying degrees follow the template of the empirical analysis presented above, albeit in pursuit of answers to different economic questions. The computations are by now automated in software packages such as DYNARE and accessible to a large community of empirical macroeconomists, which is a reflection of the progress that the literature has made over the past ten years.

2.3 Challenges

The smooth execution of the empirical analysis in the previous section may give the impression that the literature has solved most of the key conceptual problems associated with the estimation of DSGE models. Unfortunately – for those who are applying the methods – and fortunately – for those who are developing them, this is not the case. Computational constraints put bounds on the degree of realism and complexity of macroeconometric models. In light of the steady progression of computational capabilities, much of the ongoing research focuses on enriching endogenous propagation mechanisms (e.g., by incorporating labor market frictions, financial frictions, informational frictions and learning, heterogeneity impulses), the use of richer exogenous shock processes (e.g., anticipated shocks and shocks with regime-switching or stochastic volatility dynamics), and accounting for model-implied nonlinear dynamics of endogenous variables in the estimation of DSGE models. Rather than scrutinizing the latest advances in enriching DSGE models, I will focus on some methodological and conceptual challenges that have plagued the field for a while. Recent advances in the estimation of nonlinear DSGE models are discussed in the chapter by Fernandez-Villaverde and Rubio-Ramirez (This Volume), with a special emphasis on time-varying volatility dynamics in macroeconomic data.

Challenge 1: Fragile Parameter Estimates. The NKPC (15) appears in many DSGE models. In Schorfheide (2008) I compiled a table of 42 DSGE model-based estimates of κ and γ_b that had been published in academic journals. The large number of estimates is testament to a widespread use of the estimation techniques that have been developed in recent years. The estimates range from essentially zero to about four. A value near zero implies that monetary policy changes have a large effect on output but very little effect on inflation. A value of four, on the other hand, means that prices are essentially flexible and that output does not react to monetary policy changes. This remarkable range is due to differences in model specification, choice of observables and sample period, data definitions, and detrending. Unfortunately, the measures of uncertainty reported in the individual studies give no indication about the fragility of the results from a meta perspective. To illustrate this point, Figure 3 depicts a 90% credible set for γ_b and κ in (15) based on the estimation of the DSGE model described in Section 2.1 as well as the 42 parameter estimates surveyed in Schorfheide (2008). It is apparent that the posterior uncertainty conditional on a specific model and data choice is dwarfed by the variation across model specifications and data sets. The fragility of parameter estimates potentially translates into other objects of interest such as inference about the sources of business cycle fluctuations, forecasts, as well as policy prescriptions. Thus, accounting for model uncertainty as well as for different approaches of relating model variables to observables is of first-order importance.

Challenge 2: Aggregate Uncertainty versus Misspecified Endogenous Propagation. Figure 4 depicts the time series of inverse velocity used for the estimation of the DSGE model. In addition, the figure shows a counterfactual path for velocity that is obtained by setting all exogenous shocks except the money demand shock equal to zero. The sequence of money demand shocks is kept at its estimated value. A visual inspection of Figure 4 suggests that the money demand shock explains most of the historical variation in velocity. This finding has two possible interpretations. On the one hand, it could be the case that velocity fluctuations are overwhelmingly due to changes in money demand. On the other hand, it is conceivable that the endogenous transmission of technology, government spending, and monetary policy shocks into monetary aggregates is misspecified and the exogenous money demand shock absorbs mostly specification error. In the absence of other empirical evidence, formal econometric methods have difficulties distinguishing these two interpretations. The phenomenon that the variation in certain time series is to a large extent explained by shocks that are inserted into intertemporal or intratemporal optimality conditions is fairly widespread and has lead to criticisms of existing DSGE models, e.g. Chari, Kehoe, and McGrattan (2007).

Challenge 3: Trends. The DSGE model of Section 2.1 implies that velocity follows a stationary process with a constant mean. Figure 4 shows that inverse velocity was falling from 1960 to 1982 and then rising subsequently, which suggests that that its path would be better captured by a trend-stationary model with a structural break. The problem of a mismatch between trends in the data and trends in DSGE models is fairly widespread and extends beyond the velocity series. Most DSGE models impose strict balanced growth path restrictions implying, for instance, that consumption-output, investment-output, government spending-output, and real-wage output ratios should exhibit stationary fluctuations around a constant mean. In the data, however, many of these ratios exhibit trends. As a consequence, counterfactual low frequency implications of DSGE models manifest themselves in exogenous shock processes that are estimated to be highly persistent. To the extent that inference about the sources of business cycles and the design of optimal economic policies is sensitive to the persistence of shocks misspecified trends are a reason for concern.

Challenge 4: Statistical Fit. Macroeconometrics is plagued by a trade-off between theoretical coherence and empirical fit. Theoretically coherent DSGE models impose tight restrictions on the autocovariance sequence of a vector time series, which often limit its ability to track macroeconomic time series as well as, say, a less restrictive vector autoregression (VAR). A Bayesian framework allows researchers to assign probabilities to competing model specifications. If $\pi_{0,i}$ are prior probabilities assigned to models M_i, $i = 1, 2$, then the posterior odds of the two models after observing a sample of T observations are given by

$$\frac{\pi_{1,T}}{\pi_{2,T}} = \frac{\pi_{1,0}}{\pi_{2,0}} \frac{p(Y|\mathcal{M}_1)}{p(Y|\mathcal{M}_2)}. \tag{18}$$

The marginal likelihood $p(Y)$, omitting the conditioning on $\mathcal{M}_i$, was defined in (14) and implicitly penalizes the in-sample fit of a model by a measure of complexity. The log marginal likelihoods for the DSGE model and the VAR are -940.22 and -924.14, respectively, and shift the prior odds in favor of the VAR by a factor of e^{16}.

To shed some light on the difference in (penalized) fit of DSGE model and VAR, Figure 5 depicts the impulse responses to an unanticipated change in the target inflation rate. In both DSGE model and VAR the response is identified by the assumption that the target inflation rate evolves exogenously. The target inflation shock raises inflation and nominal interest rates by about 22 basis points in the long run. Output falls, because the higher inflation rate exacerbates both the New Keynesian and the monetary distortion. While the estimated responses of output, inflation, and interest rates are similar, the inverse velocity response is very different and points toward a source of misspecification of the DSGE model: it is unable to capture the rather large long-run elasticity of money demand with respect to interest rate changes.

If the goal of the empirical analysis is to provide an impulse response function to an unanticipated change in the target inflation rate, one might feel more comfortable relying on the VAR prediction because a formal econometric analysis suggests to place more weight on them (though the VAR does not provide a coherent economic explanation for the responses). If, on the other hand, the goal is to determine the welfare effect of the change in the inflation target, then the VAR is of limited use. While the drop in output and money balances might suggest a welfare loss, it is unclear how to trade-off a decrease in consumption against an increase in leisure. At the same time the discrepancy between VAR and DSGE model responses is disconcerting as money balances enter directly the households' utility function.

In order to narrow the gap between the DSGE and VAR impulse responses to a target-inflation rate shock, I reduce the value of ν from 31.7 to 3 to increase the (partial) elasticity of money demand to interest rate changes without re-estimating the remaining parameters. A comparison of impulse responses obtained under the two values of ν is provided in the top panels of Figure 6. With $\nu = 3$ there is overlap of the VAR and DSGE impulse response bands over a horizon of 5 to 20 quarters. While real deficiency of the DSGE model is its inability to deliver a small short-run and a large long-run interest elasticity of money demand, it is possible to adopt a "loss-function-based" estimation approach for ν and choose a value that matches the properties of the DSGE models with the VAR evidence on the long-run effect of target inflation changes. The bottom panels of Figure 6 illustrate how the change in ν affects the policy implications. A higher interest elasticity increases the welfare cost of inflation caused by the monetary distortion and shifts the optimal inflation towards -2%, which yields a zero nominal interest rate.

Challenge 5: Reliability of Policy Predictions. Estimated DSGE models are often used as laboratories for policy experiments. An example of such an experiment is a change in the target inflation rate discussed above. While our sample contains observations from a high inflation episode as well as observations from low inflation episodes, there are no extended periods of zero or negative inflation, which are the inflation rates at which the New Keynesian and the monetary friction create a trade-off for policy makers. More generally, to the extent that there exist no (or very few) observations on the behavior of households and firms under a counterfactual policy, the DSGE model is used to derive the agents' decision rules by solving intertemporal optimization problems assuming that the preferences and production technologies are unaffected by the policy change. In most cases the policy invariance is simply an assumption and there is always concern that the assumption is unreliable. This concern is typically exacerbated by evidence of model misspecification. While it is conceivable that a model with the worse statistical fit delivers the better policy prediction as illustrated by Kocherlakota (2007), bad fit is certainly no guarantor of good policy predictions.

In the remainder of this chapter I will discuss recent progress in overcoming these five challenges. I will begin by reviewing current work on the identification of DSGE model parameters (Section 3). Lack of identification contributes to the fragility of parameter estimates. A second factor contributing to the fragility of estimates is model misspecification. Misspecification is also plays a leading role in the other four challenges. While misspecification can be alleviated through improving the endogenous propagation mechanisms of the DSGE model, I will focus on two other directions of research, namely the generalization of exogenous shock processes in Section 4 and the development of hybrid models that correct DSGE model misspecification in Section 5. Finally, in Section 6 I discuss some simulation experiments that illustrate how even simple forms of heterogeneity and asset market incompleteness can undermine the policy invariance of preference and technology parameters in a representative agent model and lead to an understatement of the uncertainty associated with policy predictions.

3 Identification and Inference

The fragility of estimates discussed in Section 2.3 is in part due to lack of identification of key DSGE model parameters. Identification in DSGE models, even if they are linearized, is much less transparent than identification in linear simultaneous equations models. This lack of transparency is reflected in the fact that the system matrices of the state-space representation (13) are complicated nonlinear functions of the underlying DSGE model parameters θ, which for all but the most rudimentary and unrealistic DSGE models can only be evaluated numerically. While the early literature on DSGE model estimation had paid very little attention to identification, more recently researchers have realized that estimation objective functions are often uninformative with respect to important parameters such as the Phillips curve coefficients in (15) or the parameters in the monetary policy rule (11). Canova and Sala (2009), for instance, document identification problems in popular New Keynesian DSGE models. Section 3.1 provides a simple example that illustrates the identification problems. Section 3.2 presents recently developed conditions for identification of DSGE model parameters and consequences for econometric inference are discussed in Section 3.3.

3.1 A Simple Example

The following stylized example illustrates identification problems that may arise in the context of DSGE models. Suppose that the structural model is nested in the following state-space representation, which resembles (13):

$$y_t = [\ 1 \quad 1\ ] s_t, \quad s_t = \begin{bmatrix} \phi_1 & 0 \\ \phi_3 & \phi_2 \end{bmatrix} s_{t-1} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} \epsilon_t. \tag{19}$$

The first state, $s_{1,t}$, resembles an exogenous shock, such as technology, whereas the transition equation for $s_{2,t}$ mimics that of an endogenous state variable such as the capital stock. Moreover, for the sake of concreteness suppose that the relationship between the reduced-form (state-space) parameter $\phi = [\phi_1, \phi_2, \phi_3]'$ and the structural (DSGE Model) parameter $\theta = [\theta_1, \theta_2]'$ is given by:

$$\phi_1 = \theta_1^2, \quad \phi_2 = (1 - \theta_1^2), \quad \phi_3 - \phi_2 = -\theta_1 \theta_2. \tag{20}$$

In order to understand the identification problems, it is useful to rewrite the state-space model as ARMA(2,1) process

$$(1 - \phi_1 L)(1 - \phi_2 L) y_t = (1 - (\phi_2 - \phi_3) L) \epsilon_t. \tag{21}$$

First, (20) implies that θ_2 becomes non-identifiable as θ_1 approaches zero, because for $\theta_1 = 0$ the law of motion of y_t is invariant to θ_2. Second, it can be easily verified that the following two parameterizations are observationally equivalent:

$$\theta_1^2 = \rho, \quad (1 - \theta_1^2) = \theta_1 \theta_2 \quad \text{versus} \quad \tilde{\theta}_1^2 = 1 - \rho, \quad \tilde{\theta}_1^2 = \tilde{\theta}_1 \tilde{\theta}_2.$$

Under both parameterizations, y_t follows an AR(1) process with autocorrelation parameter ρ, because one factor of the autoregressive polynomial of the ARMA(2,1) process cancels against the moving-average polynomial.

3.2 Conditions for Identifiability

Recent work by Iskrev (2010) and Komunjer and Ng (2009) develops necessary and sufficient conditions for the identifability of DSGE model parameters. The conditions are meant to be comparable to the rank and order conditions that exist for simultaneous equations models and focus on linearized DSGE models with Gaussian innovations that can be cast in the state-space form (13). Iskrev (2010), develops a condition based on the direct relationship between the parameter vector θ and first and second population moments $m_T(\theta)$ of a sequence of observations $Y_{1:T} = [y_1, \ldots, Y_t]'$. A sufficient condition for global identifiability is that $m_T(\tilde{\theta}) = m_T(\theta)$ implies that $\tilde{\theta} = \theta$ for each pair $(\theta, \tilde{\theta})$. If the condition only holds in an open neighborhood of θ, then θ is locally identifiable. Since the state-space model is linear, the identifiability condition is necessary if the structural shocks ϵ_t as well as the initial state s_0 are normally distributed. If $m_T(\theta)$ is continuously differentiable, then θ is locally identifiable if the Jacobian matrix $\partial m_T(\theta)/\partial\theta'$ is of full column rank. Since even linearized DSGE models are nonlinear in the parameters, the rank condition needs to be verified for a large number of empirically relevant parameter values. The simple model in Section 3.1 is not globally identifiable but it is locally identifiable for many values of θ. However, local identification fails, for instance, if $\theta_1 = 0$.

Komunjer and Ng (2009) extend Iskrev's conditions from a finite number of second moments stacked in $m_T(\theta)$ to the infinite-dimensional autocovariance sequence, represented by the spectral density of y_t. To do so, the authors develop rank conditions that ensure that the mapping between θ and the reduced-form parameters of the state-space representation is (locally) one-to-one. The difficulty in developing such conditions arises from the fact that the parameters of the state-space representation themselves suffer from identification problems. Thus, a naively defined reduced form parameter vector

$$\phi = [vec(\Psi_0)',\ vec(\Psi_1)',\ vec(\Psi_s)',\ vec(\Phi_1)',\ vec(\Phi_\epsilon)']',$$

where the Ψ and Φ matrices refer to the coefficient matrices in (13), needs to be reparameterized in terms of an identifiable subvector before rank conditions can be stated.

3.3 Consequences for Inference

From an inferential viewpoint there are two basic reactions to a potential lack of identification. The first perspective is represented in the large literature on weakly or partially identified econometric models: taking data and model as given, the econometrician should use inferential procedures that are robust to a potential lack of identification. The second perspective is reflected in the following quote from Dreze (1974), p. 164: "The econometrician who is concerned with inference about parameters that are not identified may try to overcome this difficulty by collecting richer data, or by resorting to a more restrictive theory."[4] I will subsequently focus on identification-robust inference in Bayesian and frequentist analysis as well as the notion of collecting richer data sets.

Bayesian Inference. Bayesian inference with proper priors does not require identifiability as a regularity condition. As long as the prior distribution is proper (meaning the total probability mass is one), so is the posterior, see for instance Poirier (1998). What matters for inference is the curvature in the likelihood function, as priors do not get updated in directions in which the likelihood function is flat. This leads to a number of practical challenges. First, inference becomes more sensitive to the choice of prior distributions, making a careful, systematic, and well-documented choice of prior distribution important for compelling empirical work.[5] Second, lack of identification may complicate the generation of parameter draws from the posterior distribution.

Figure 7 depicts two likelihood functions for the stylized model of Section 3.1, constructed by simulating 100 artificial observations based on two different sets of "true" θ values. In the top panel the "true" value of θ_1 is fairly close to zero, which makes it difficult to identify θ_2. Accordingly, the likelihood function has a ridge and is fairly flat in the direction of θ_2. The second parameterization highlights the global identification problem. While not directly visible from the contours plotted in the figure, the likelihood function is in fact bimodal. It is typically the lack of global identification and the resulting multi-modal posterior surfaces that cause problems for posterior simulators.[6] While many of the posterior

[4]The debate between Lubik and Schorfheide (2004, 2007) and Beyer and Farmer (2007) illustrates how a restrictive theory can lead to identification and the disagreement between researchers as to whether such restrictions should be imposed in empirical work.

[5]Müller (2010) develops measures of prior sensitivity and informativeness tailored toward DSGE model applications.

[6]Suppose the likelihood function of a DSGE model were completely uninformative with respect to all parameters. In this case one would simply have to generate draws from the prior, which typically can be

simulators that are used in practice, most notably the version of the random-walk Metropolis (RWM) algorithm described in An and Schorfheide (2007a), in principle deliver consistent approximations of posterior moments and quantiles even if the posterior is multi-modal, the practical performance can be poor as documented in An and Schorfheide (2007a).

Recent research on posterior simulators tailored toward DSGE models tries to address the shortcomings of the "default" approaches that are being used in empirical work. An and Schorfheide (2007b) use transition mixtures to deal with a multi-modal posterior distribution. This approach works well if the researcher has knowledge about the location of the modes, obtained, for instance, by finding local maxima of the posterior density with a numerical optimization algorithm. Chib and Ramamurthy (2010) propose to replace the commonly used single block RWM algorithm with a Metropolis-within-Gibbs algorithm that cycles over multiple, randomly selected blocks of parameters. Kohn, Giordani, and Strid (2010) propose an adaptive hybrid Metropolis-Hastings samplers and Herbst (2010) develops a Metropolis-within-Gibbs algorithm that uses information from the Hessian matrix to construct parameter blocks that maximize within-block correlations at each iteration and Newton steps to tailor proposal distributions for the various conditional posteriors.

Frequentist Inference. Standard large sample approximations of sampling distributions of estimators and test statistics require parameter identifiability as regularity conditions. The literature on identification-robust inference procedures relaxes this regularity condition while maintaining that the coverage probability of a confidence interval $CS_T(Y_{1:T})$ constructed from a sequence of observations $Y_{1:T}$ should converge uniformly in the following sense:

$$\lim_{T \longrightarrow \infty} \inf_{\tilde{\phi} \in \mathcal{P}} \inf_{\theta \in \Theta(\tilde{\phi})} P_{\tilde{\phi}}\{\theta \in CS_T(Y_{1:T})\} = 1 - \alpha. \tag{22}$$

Here $\tilde{\phi}$ denotes an *identifiable* reduced form parameter which indexes the probability distribution of the data. $\Theta(\tilde{\phi})$ denotes the set of structural parameters that is consistent with a particular reduced-form parameter $\tilde{\phi}$. This set degenerates to a singleton in a point-identified model. In the context of the example presented in Section 3.1 $\tilde{\phi}$ could be defined as the au-

done by direct sampling or acceptance sampling given the highly informative prior distributions that are used in the literature.

tocovariances of order zero to three. If the autocovariances of order one to three are zero then θ_2 is non-identifiable and $\Theta(\tilde{\phi})$ corresponds to a line in $\mathbb{R}^3$.

The standard approach of constructing confidence sets by taking a point estimate and adding and subtracting multiples of the associated standard error estimate does typically not lead to valid inference in models with identification problems (meaning (22) is violated). Instead, confidence sets are often obtained through point-wise testing procedures. Suppose that inference for the reduced form parameter $\tilde{\phi}$ is regular in the sense that[7] $\sqrt{T}(\hat{\tilde{\phi}} - \tilde{\phi}) \Longrightarrow N(0, \Lambda)$ and that the relationship between $\tilde{\phi}$ and θ can be expressed by a function $\tilde{\phi}^*(\theta)$. To obtain a valid confidence set, choose a grid $\mathcal{T}$ for θ and conduct point-wise tests of the hypothesis $\tilde{\phi} = \tilde{\phi}^*(\theta)$ for all $\theta \in \mathcal{T}$. The confidence set for θ is comprised of those values of θ for which the null hypothesis cannot be rejected. This approach is explored in Guerron-Quintana, Inoue, and Kilian (2010). While the procedure leads to valid inference in the sense of (22), it has several drawbacks. In high-dimensional parameter spaces the procedure requires many point-wise tests. Moreover, the method is conservative in regions of the parameter space in which the parameters are well identified. The development of efficient methods to construct identification-robust confidence intervals for a DSGE model remains an open area of research.

Richer Data Sets. DSGE models are typically estimated with observations on only a subset of all the variables that appear in the model, because due to its stylized structure it is only able to capture the dynamics of some but not all variables in a realistic manner. For instance, the simple structure of the labor market of the model in Section 2 (infinitely elastic labor supply, absence of search frictions) makes it difficult to match the dynamics of hours worked and wages, which is why these observations are omitted from the likelihood function. However, long-run properties of series that are excluded from the likelihood function, e.g. the average labor share, remain informative about some of the DSGE model parameters. This *non-sample* information can and should be used for inference. Non-sample information might

[7] The example in Section 3.1 illustrates that the sampling distribution of estimators of the state-space coefficients may be irregular. To overcome this problem $\tilde{\phi}$ could be defined as the coefficients of the VAR approximation of a DSGE model, which leads to a standard normal sampling distribution provided the process y_t is stationary.

also include evidence from microeconometric panel studies on demand or supply elasticities.

The non-sample information may be able to resolve some identification problems inherent in the likelihood function. In a Bayesian framework, it is most natural to use this information in the specification of a prior distribution, which was the approach taken in the empirical analysis in Section 2. I started with marginal densities for the model parameters θ_i, $i = 1, \dots, k$ and then combined them with a function $f(\theta)$ that incorporates some information from long-run averages of observations that do not enter the construction of the likelihood function:

$$p(\theta) \propto f(\theta) \prod_{i=1}^{k} p_i(\theta_i). \tag{23}$$

where

$$f(\theta) \quad = \quad \exp\left\{ -\frac{1}{2}\left[\frac{(I_*(\theta)/\mathcal{Y}_*(\theta) - 0.16)^2}{0.005^2} + \frac{(lsh(\theta) - 0.60)^2}{0.01^2} \right] \right\}.$$

Here $I_*(\theta)$, $\mathcal{Y}_*(\theta)$, and $lsh(\theta)$ are functions that define the steady state levels of investment, output, and the labor share. 0.16 and 0.60 are long-run averages of the investment-output ratio and the labor share computed from U.S. data. This method of constructing prior distributions is formalized in Del Negro and Schorfheide (2008). The underlying assumption in the application of Bayes Theorem in this case is that sample and non-sample information are approximately independent.

4 Sensitivity to Shock Specification

A DSGE model consists of endogenous propagation mechanisms, e.g. investment and capital accumulation, derived from some primitive assumptions about agents' preferences and production technologies, as well as exogenous propagation mechanisms. While most of the modelling efforts in the DSGE model literature are rightly directed toward the specification of the endogenous propagation mechanism, this section focuses on the specification of exogenous shock processes and its consequences for inference based on estimated DSGE models. These shocks themselves are frequently assumed to follow independent AR(1) processes as in Section 2.1. The lag length restriction for the individual shock processes is in many instances arbitrary. The assumption that the exogenous processes are independent of each other is a reflection of a modeling strategy that tries to explain the comovements of macroeconomic aggregates with economic mechanisms rather than through correlated exogenous shocks.

A careful specification of the law of motion for the exogenous shocks can help to overcome model misspecification, in particular if one means by misspecification inferior time series fit (adjusted for model dimensionality) relative to more flexible time series models such as VARs. More specifically, recent empirical work has documented that the fit of a DSGE model can be improved by relaxing the restriction that the exogenous shocks exhibit AR(1) dynamics. Smets and Wouters (2007) use an ARMA mark-up shock to improve model fit and Del Negro and Schorfheide (2009) let their government spending shock follow an higher-order autoregressive process. Curdia and Reis (2010) proposed to introduce correlation among the exogenous processes and replaced independent univariate shock processes by a vector process. At the same time some of the current arbitrariness in the specification of the exogenous shock processes as well as potential generalizations to improve the model fit contribute to the identification problems discussed in Section 3 and thereby to the fragility of parameter estimates.

Generalization of Shock Dynamics and Identification. Consider a DSGE model in which a representative firm has access to a Cobb-Douglas production function of the form $Y_t = Z_t K_t^{\alpha} H_t^{1-\alpha}$ and capital accumulates according to $K_{t+1} = (1-\delta)K_t + I_t$. To the extent that α can be measured from labor share data, δ from NIPA data on capital stock depreciation, and output, hours, and investment are used as observables in the estimation, the latent total factor productivity process Z_t is essentially identified as (Solow) residual in the production function. As discussed in Section 3.3, in a Bayesian estimation the information about α and δ can be incorporated through a prior distribution. The Kalman filter that is used to compute the likelihood function delivers estimates of the latent capital stock K_t as well as Z_t. Given observations on Y_t, H_t, and I_t as well as fairly tight priors on α and δ the only source of uncertainty with regard to the latent variables is the initialization of the capital stock. In turn, it is fairly straightforward to identify the coefficients of a flexible time series model for the exogenous technology process. In practice, AR(1) or AR(2) models are widely used for the TFP process because they are fairly successful in capturing the stochastic properties of the Solow residual.

Alternatively, consider a simplified version of the monetary policy rule (11):

$$\tilde{R}_t = \rho_R \tilde{R}_{t-1} + (1-\rho_R)(\psi_1 \tilde{\pi}_t + \psi_2 \tilde{Y}_t) + \epsilon_{R,t}. \tag{24}$$

Unlike in the production function example the slope coefficients in the monetary policy rule are not tied to steady states of macroeconomic aggregates that could be identified through long-run averages. As a consequence, assumptions about the stochastic properties of the exogenous monetary policy shock $\epsilon_{R,t}$ are closely tied to the identification of the policy rule coefficients. The assumption that $\epsilon_{R,t}$ is an *iid* sequence provides identification in the sense that lagged inflation and output can serve as instrumental variables in the estimation of the policy rule coefficients. This source of identification vanishes if $\epsilon_{R,t}$ is allowed to be serially correlated. Unfortunately, in many instances of DSGE model estimation the identification of key economic mechanisms is determined by somewhat arbitrary and restrictive assumptions about the stochastic properties of exogenous shocks. More general shock processes, on the other hand, are likely to exacerbate the problem of multi-modal estimation objective functions as illustrated in Herbst (2010).

Documenting Sensitivity to Auxiliary Modelling Assumptions. In particular in medium to large-scale DSGE models that are estimated on seven or more observables, the choice of several of the shocks is somewhat arbitrary. While there is little controversy about technology and monetary policy shocks, the inclusion of inter and intratemporal preference shocks, price mark-up shocks, or risk-premium shocks tends to be controversial and typically guided by improving model fit. To the extent that there is modeling uncertainty about the exogenous shock structure and that assumptions about to shock structure affect the identification of key parameters and propagation mechanisms, it is useful to document the sensitivity to modeling assumptions in a systematic manner. In Ríos-Rull, Schorfheide, Fuentes-Albero, Kryshko, and Santaeulalia-Llopis (2009) this is done by Bayesian model averaging across model specifications with different exogenous shock specifications.

5 Hybrid Models

Econometric modelling typically faces a trade-off between theoretical coherence and empirical fit. The DSGE paradigm delivers empirical models with a strong degree of theoretical coherence that often fit worse than more densely parameterized time series models, e.g. VARs, as illustrated in Section 2. There exist essentially two approaches in the literature to construct empirical models that relax DSGE model restrictions. I will refer to these models as additive hybrid models (Section 5.1) and hierarchical hybrid models (Section 5.2), respectively. Hybrid models provide a complete characterization of the law of motion of the data, as opposed to empirical procedures that remove some variation from the data that the DSGE model is unable to capture. At the same time, hybrid models retain important dynamic properties of the DSGE model.

5.1 Additive Hybrid Models

The additive hybrid model augments the state-space model (13) by a latent process z_t that bridges the gap between data and theory:

$$\begin{aligned} y_t &= \Psi_0(\theta) + \Psi_1(\theta)t + \Psi_s(\theta)s_t + \Lambda_0 + \Lambda_1 t + \Lambda_z z_t \\ s_t &= \Phi_1(\theta)s_{t-1} + \Phi_\epsilon(\theta)\epsilon_t, \quad z_t = \Gamma_1 z_{t-1} + \Gamma_\eta \eta_t. \end{aligned} \tag{25}$$

The process z_t is often called *measurement error*, blaming the data collectors rather than the DSGE model builders for the gap between data and theory.[8] Unlike in the generalization of the exogenous shocks of the DSGE model described in Section 4, the agents in the model economy do not account for z_t in their decision making and consequently there is no interaction with the economic states s_t.

Special Cases. Without any restrictions on Λ and Γ the model (25) is not identifiable. The following two restrictions have been widely used in practice. First, a low dimensional vector of structural shocks ϵ_t is combined with a diagonal Γ_1 matrix, e.g. Altug (1989). In this

[8] The use of measurement errors in the estimation of optimization-based macro models dates back at least to Sargent (1989) and Altug (1989) and has been advocated more recently by Ireland (2004).

setup the ϵ_t's generate the comovements between the observables, whereas the elements of z_t pick up idiosyncratic dynamics that are not captured by the structural part of the hybrid model. Second, if one sets Ψ_0, Ψ_1, and Λ_z to zero, then the hybrid model uses the DSGE component to describe the fluctuations of y_t around a deterministic trend path, but it ignores the common trend restrictions of the structural model. This version of the additive hybrid model is typically estimated in two steps, e.g. Smets and Wouters (2003). In the first step deterministic trends are removed from the data and in the second step the DSGE model is estimated based on the linearly detrended observations.

Correcting Low Frequency Misspecification. Section 2.3 illustrated that some of the misspecification of DSGE models rests in their inability to capture certain long-run features of the data. The hybrid model can be used to correct these deficiencies. Canova (2010) proposes the following specification:

$$\begin{aligned} y_t &= \Psi_s(\theta)s_t + \Lambda_0 + z_t \\ s_t &= \Phi_1(\theta)s_{t-1} + \Phi_\epsilon(\theta)\epsilon_t \\ z_t &= z_{t-1} + \bar{z}_{t-1} + \eta_t, \quad \bar{z}_t = \bar{z}_{t-1} + \nu_t. \end{aligned} \tag{26}$$

Depending on the restrictions imposed on the variances of η_t and ν_t the process z_t is integrated of order one or two and can generate a variety of stochastic trend dynamics.

Connecting DSGE Models with Large Data Sets. Macroeconomists have access to large cross sections of aggregate variables that include measures of sectoral economic activities and prices as well as numerous financial variables. Additive hybrid models can also be used to link DSGE models with aggregate variables that are not explicitly modeled. Using these additional variables in the estimation potentially sharpens inference about latent state variables. Moreover, the link enables researchers to construct impulse response functions and predictions for economic variables that are not explicitly modeled.

Let y_t denote the observable variables that are described by the DSGE model and let x_t denote a large vector of non-modeled variables. The joint law of motion of y_t and x_t is given by

$$y_t = \Psi_0(\theta) + \Psi_1(\theta)t + \Psi_s(\theta)s_t + z_{y,t} \tag{27}$$

$$x_t = \Lambda_0 + \Lambda_1 t + \Lambda_s s_t + z_{x,t} \tag{28}$$

$$s_t = \Phi_1(\theta)s_{t-1} + \Phi_\epsilon(\theta)\epsilon_t. \tag{29}$$

Since the structure of this model resembles that of a dynamic factor model (DFM), e.g. Sargent and Sims (1977), Geweke (1977), and Stock and Watson (1989), I refer to the system (27) to (29) as DSGE-DFM. The vector of factors is given by the state-variables associated with the DSGE model. The processes $z_{y,t}$ and $z_{x,t}$ are uncorrelated across series and capture idiosyncratic but potentially serially correlated movements (or measurement errors) in the observables. (28) links the variables x_t to the DSGE model. This linkage generates comovements between the y_t's and the x_t's and allows the computation of impulse responses to the structural shocks ϵ_t. The DSGE-DFM was originally proposed by Boivin and Giannoni (2006). Kryshko (2010) improves some computational aspects of the Bayesian inference for the DSGE-DFM. Moreover, using a DSGE model very similar to the one described in Section 2.1, he documents that the space spanned by factors extracted from the DSGE-DFM is similar to the space spanned by the factors estimated with an unrestricted DFM. This finding gives an economic interpretation to the factors extracted with a reduced-form factor model and lends credibility to the state transitions implied by the DSGE model. Schorfheide, Sill, and Kryshko (2010) study the forecast performance of the DSGE-DFM with respect to some specific variables x_t that are not explicitly modeled in the DSGE model.

5.2 Hierarchical Hybrid Models

Now consider the following modification of the additive hybrid model:

$$y_t = \Lambda_0 + \Lambda_1 t + \Lambda_s s_t, \quad s_t = \Gamma_1 s_{t-1} + \Gamma_\epsilon \epsilon_t, \tag{30}$$

where

$$\Lambda_i = \Psi_i(\theta) + \eta_i^\Psi, \quad i = 0, 1, s, \quad \Gamma_i = \Phi_i(\theta) + \eta_i^\Phi, \quad i = 1, \epsilon. \tag{31}$$

In this setup $\Psi_i(\theta)$ and $\Phi_i(\theta)$ are interpreted as restrictions on the unrestricted state-space matrices Λ_i and Γ_i. The disturbances η_i^Ψ and η_i^Φ can capture deviations from the restriction functions $\Psi_i(\theta)$ and $\Phi_i(\theta)$. The smaller the variance of the η's the closer the empirical model stays to the DSGE model. In a Bayesian framework, the stochastic restrictions (31) correspond to a prior distribution of the unrestricted state-space matrices conditional on the DSGE model parameters θ.

DSGE-VARs. It turns out that the formal Bayesian analysis of the model comprised of (30) and (31) is computationally challenging and the subject of ongoing research. The analysis is considerably easier to implement if the state-space model in (30) is replaced by a VAR:

$$y_t = B_1 y_{t-1} + \ldots + B_p y_{t-p} + B_c + \Sigma_{tr}\Omega\epsilon_t, \tag{32}$$

where Σ_{tr} is the unique lower triangular Cholesky factor of the one-step-ahead VAR forecast error covariance matrix Σ, Ω is an orthogonal matrix, and $\epsilon_t \sim N(0, I)$. Let $B = [B_1, \ldots, B_p, B_c]'$. Suppose that DSGE model parameters θ and VAR parameters are linked through binding functions denoted by

$$B^*(\theta), \quad \Sigma^*(\theta), \quad \Omega^*(\theta). \tag{33}$$

The prior distribution for the VAR coefficients (B, Σ, Ω) conditional on θ is chosen such that it is centered at the binding functions (33) but allows for deviations through a non-zero covariance matrix, as in (31).[9] This covariance matrix is scaled by a hyperparameter λ. Overall, the setup leads to a hierarchical model of the form

$$p_\lambda(Y, B, \Sigma, \theta) = p(Y|B, \Sigma)p_\lambda(B, \Sigma, \Omega|\theta)p(\theta), \tag{34}$$

where $p(\theta)$ is a prior for the DSGE model parameters and $p(Y|B, \Sigma)$ is the likelihood function associated with (32). Details of the specification of $p_\lambda(B, \Sigma, \Omega|\theta)$ can be found in Del Negro and Schorfheide (2004) or Del Negro and Schorfheide (2010). The resulting empirical model is more flexible than the DSGE model itself while it still inherits many of its dynamic properties for a wide range of hyperparameter settings.

[9]The basic idea of using a DSGE model to formulate a prior distribution for VAR coefficients dates back to Ingram and Whiteman (1994).

Empirical Illustration. The DSGE model from Section 2 is now used to create a hierarchical hybrid model. The analysis differs in three dimensions from the DSGE-VARs in Del Negro and Schorfheide (2004) and Del Negro, Schorfheide, Smets, and Wouters (2007). First, the prior distribution used in the analysis is a combination of the Minnesota prior[10] and the DSGE model prior. For $\lambda = 0$, no information is used from the DSGE model and and a VAR with Minnesota prior is estimated. For $\lambda = \infty$, on the other hand, the DSGE model restrictions are dogmatically imposed. Second, the DSGE model implies that the target inflation rate evolves according to a unit root process, which was not covered by the existing DSGE-VAR setup. Consequently, I generalized the construction of the prior distribution to allow for unit roots in the DSGE model. Third, in order to identify the target inflation rate shock, I use the assumption that $\pi_{*,t}$ is the first element of y_t and simply restrict Ω in (32) to be the identity matrix. Thus, the target inflation rate does not react to the other shocks contemporaneously.

The top left panel of Figure 8 depicts the log marginal data density as a function of the hyperparameter λ, given by

$$\ln p_\lambda(Y) = \ln \int p_\lambda(Y, \Phi, \Sigma, \theta) d(\Phi, \Sigma, \theta). \tag{35}$$

This function peaks approximately at $\lambda = 0.5$. Thus, the DSGE model restrictions improve the fit of the empirical model relative to the fit attained with only the Minnesota prior. However, since the marginal likelihood is much larger at $\lambda = 0.5$ than at $\lambda = \infty$, the plot provides evidence for model misspecification. The remaining panels of Figure 8 depict posterior mean impulse responses to an inflation target shock as a function of λ. First, the responses of the target inflation rate, output, and inflation does not substantially change as one varies λ, suggesting that the DSGE model seems to be well specified in this dimension. Second, the response of real money balances is highly sensitive to the choice of λ. The rather low value of λ favored by the marginal likelihood, implies a real money balance response that is much stronger than the response predicted by the DSGE model.

[10] Details on the version of the Minnesota prior used for the empirical analysis can be found in Del Negro and Schorfheide (2010).

6 Econometric Policy Evaluation

As illustrated in Section 2, estimated DSGE models can serve as laboratory for policy experiments, such as changes in the target level of inflation or changes in tax policies. The key assumption underlying such experiments is that the primitives of the model, in particular the parameters that characterize preferences and technologies are policy-invariant. Chang, Kim, and Schorfheide (2010) conduct a simulation experiment to assess the policy invariance of the parameters in a simple neoclassical stochastic growth model. The data generating process is a heterogeneous agent economy in which individuals face idiosyncratic productivity shocks, idiosyncratic productivity risk is uninsurable, individuals face a borrowing constraint, and labor supply is indivisible. Based on aggregated data from this economy, a representative agent model is estimated. The question of interest is to what extent the effect of labor and capital tax changes can be correctly predicted with the estimated representative agent model, assuming the invariance of the "structural" parameters. According to the simulations, the parameters of the representative agent model are not invariant to the policy changes. Moreover, the bias in the policy predictions is large relative to the size of the predictive predictive intervals obtained from the Bayesian analysis. Interestingly, there is little evidence of misspecification when the representative agent model is estimated based on data from the heterogeneous agent economy. Unlike in applications with actual U.S. data, posterior odds favor the DSGE model over a less restrictive and more densely parameterized VAR.

If the empirical analysis does reveal strong evidence of model misspecification in the sense of a violation of the cross coefficient restrictions imposed by a DSGE model on a state-space representation, e.g. (30), or a VAR approximation, e.g. (32), then there is not only concern as to whether the structural parameters of the DSGE model should be treated as policy invariant, but also whether the discrepancies between the restricted and unrestricted representations are policy invariant. Del Negro and Schorfheide (2009) develop DSGE-VAR based methods to assess the robustness of policy predictions to perturbations in the model misspecification.

7 Conclusion

The literature on the econometric analysis of DSGE models has made substantial progress over the past decade and the econometric analysis of DSGE models has become a fairly standard procedure that is now taught in many Ph.D. programs around the world. Nonetheless, many challenges that need to be tackled in the future remain. The purpose of this chapter was to review several of them and to discuss current research that tries to address them.

References

Altug, S. (1989): "Time-to-Build and Aggregate Fluctuations: Some New Evidence," *International Economic Review*, 30(4), 889–920.

An, S., and F. Schorfheide (2007a): "Bayesian Analysis of DSGE Models," *Econometric Reviews*, 26(2-4), 113–172.

——— (2007b): "Bayesian Analysis of DSGE Models–Rejoinder," *Econometric Reviews*, 26(2-4), 211–219.

Aruoba, S. B., and F. Schorfheide (2010): "Sticky Prices versus Monetary Frictions: An Estimation of Policy Trade-offs," *American Economic Journal: Macroeconomics*, forthcoming.

Beyer, A., and R. E. Farmer (2007): "Testing For Indeterminancy: An Application to U.S. Monetary Policy: Comment," *American Economic Review*, 97(1), 524 – 529.

Boivin, J., and M. P. Giannoni (2006): "DSGE Models in a Data Rich Enviroment," *NBER Working Paper*, 12772.

Canova, F. (2007): *Methods for Applied Macroeconomic Research*. Princeton University Press.

——— (2010): "Bridging Cyclical DSGE Models and the Raw Data," *Manuscript*.

Canova, F., and L. Sala (2009): "Back to Square One: Identification in DSGE Models," *Journal of Monetary Economics*, 56(4), 431–449.

Chang, Y., S.-B. Kim, and F. Schorfheide (2010): "Labor Market Heterogeneity and the Lucas Critique," *Manuscript.*

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2007): "Business Cycle Accounting," *Econometrica*, 75, 781–836.

Chib, S., and S. Ramamurthy (2010): "Tailored Randomized Block MCMC Methods with Application to DSGE Models," *Journal of Econometrics*, 155(1), 19–38.

Curdia, V., and R. Reis (2010): "Correlated Disturbances and U.S. Business Cycles," *Manuscript, Columbia University and FRB New York.*

DeJong, D., and C. Dave (2007): *Structural Macroeconometrics.* Princeton University Press.

DeJong, D. N., B. F. Ingram, and C. H. Whiteman (2000): "A Bayesian Approach to Dynamic Macroeconomics," *Journal of Econometrics*, 98(2), 203 – 223.

Del Negro, M., and F. Schorfheide (2004): "Priors from General Equilibrium Models for VARs," *International Economic Review*, 45(2), 643 – 673.

——— (2008): "Forming Priors for DSGE Models (and How it Affects the Assessment of Nominal Rigidities)," *Journal of Monetary Economics*, 55(7), 1191–1208.

——— (2009): "Monetary Policy with Potentially Misspecified Models," *American Economic Review*, 99(4), 1415–1450.

——— (2010): "Bayesian Macroeconometrics," in *Handbook of Bayesian Econometrics*, ed. by H. K. van Dijk, G. Koop, and J. Geweke. Oxford University Press.

Del Negro, M., F. Schorfheide, F. Smets, and R. Wouters (2007): "On the Fit of New Keynesian Models," *Journal of Business and Economic Statistics*, 25(2), 123–162.

Dreze, J. (1974): "Bayesian Theory of Identification in Simultaneous Equations Models," in *Studies in Bayesian Econometrics and Statistics*, ed. by S. Fienberg, and A. Zellner. North Holland, Amsterdam.

Geweke, J. (1977): "The Dynamic Factor Analysis of Economic Time Series," in *Latent Variables in Socio-Economic Models*, ed. by D. J. Aigner, and A. S. Goldberger, chap. 19. North Holland, Amsterdam.

Guerron-Quintana, P., A. Inoue, and L. Kilian (2010): "Frequentist Inference in Weakly Identified DSGE Models," *Manuscript*.

Herbst, E. (2010): "Gradient and Hessian-based MCMC for DSGE Models," *Manuscript*.

Ingram, B., and C. Whiteman (1994): "Supplanting the Minnesota Prior- Forecasting Macroeconomic Time Series Using Real Business Cycle Model Priors," *Journal of Monetary Economics*, 49(4), 1131–1159.

Ireland, P. N. (2004): "A Method for Taking Models to the Data," *Journal of Economic Dynamics and Control*, 28(6), 1205–1226.

Iskrev, N. (2010): "Local Identification in DSGE Models," *Journal of Monetary Economics*, 57(2), 189–202.

Kocherlakota, N. (2007): "Model Fit and Model Selection," *FRB St. Louis Review*, 89(4), 349–360.

Kohn, R., P. Giordani, and I. Strid (2010): "Adaptive Hybrid Metropolis-Hastings Samplers for DSGE Models," *Riksbank Manuscript*.

Komunjer, I., and S. Ng (2009): "Dynamic Identification of DSGE Models," *Manuscript, Columbia University and UCSD*.

Kryshko, M. (2010): "Data-Rich DSGE and Dynamic Factor Models," *Manuscript, University of Pennsylvania*.

Lubik, T. A., and F. Schorfheide (2004): "Testing for Indeterminancy: An Application to U.S. Monetary Policy," *American Economic Review*, 94(1), 190–217.

——— (2007): "Testing for Indeterminancy: An Application to U.S. Monetary Policy: Reply," *American Economic Review*, 97(1), 530–533.

Lucas, Robert E., J. (2000): "Inflation and Welfare," *Econometrica*, 68(2), 247–274.

Müller, U. (2010): "Measuring Prior Sensitivity and Prior Informativeness in Large Bayesian Models," *Manuscript*.

Otrok, C. (2001): "On Measuring the Welfare Costs of Business Cycles," *Journal of Monetary Economics*, 45(1), 61–92.

Poirier, D. (1998): "Revising Beliefs in Nonidentified Models," *Econometric Theory*, 14(4), 483–509.

Primiceri, G. (2006): "Why Inflation Rose and Fell: Policymakers Beliefs and US Postwar Stabilization Policy," *Quarterly Journal of Economics*, 121, 867–901.

Ríos-Rull, J.-V., F. Schorfheide, C. Fuentes-Albero, M. Kryshko, and R. Santaeulalia-Llopis (2009): "Methods versus Substance: Measuring the Effects of Technology Shocks," *NBER Working Paper*, 15375.

Rogerson, R. (1988): "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics*, 21(1), 3–16.

Sargent, T. (1999): *Conquest of American Inflation*. New Jersey: Princeton University Press.

Sargent, T. J. (1989): "Two Models of Measurements and the Investment Accelerator," *Journal of Political Economy*, 97(2), 251–287.

Sargent, T. J., and C. A. Sims (1977): "Business Cycle Modeling Without Pretending To Have Too Much A Priori Economic Theory," in *New Methods in Business Cycle Research*. FRB Minneapolis, Minneapolis.

Sargent, T. J., T. Zha, and N. Williams (2006): "Shocks and Government Beliefs: The Rise and Fall of American Inflation," *American Economic Review*, 94(4), 1193–1224.

Schorfheide, F. (2000): “Loss Function-based Evaluation of DSGE Model,” *Journal of Applied Econometrics*, 15(6), 645–670.

——— (2008): “DSGE Model-Based Estimation of the New Keynesian Phillips Curve,” *FRB Richmond Economic Quarterly*, Fall Issue, 397–433.

Schorfheide, F., K. Sill, and M. Kryshko (2010): “DSGE Model-Based Forecasting of Non-Modelled Variables,” *International Journal of Forecasting*, 26(2), 348–373.

Smets, F., and R. Wouters (2003): “An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area,” *Journal of the European Economic Association*, 1(5), 1123–1175.

——— (2007): “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 97(3), 586–606.

Stock, J. H., and M. W. Watson (1989): “New Indices of Coincident and Leading Economic Indicators,” in *NBER Macroeconomics Annual 1989*, ed. by O. J. Blanchard, and S. Fischer, vol. 4, pp. 351–394. MIT Press, Cambridge.

Woodford, M. (2003): *Interest and Prices*. Princeton University Press.

Figure 1: Posterior (and Prior) Densities

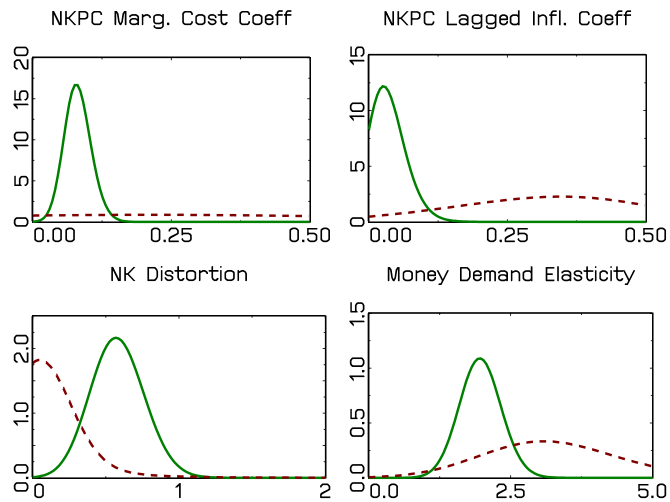

Notes: NKPC marginal cost coefficient is κ in (15), NKPC lagged inflation coefficient is γ_b in (15), NK Distortion is $100|1/D_* - 1|$ in (16), and money demand elasticity is $1/(\nu(1-R_*)$ in (17). Solid lines depict posterior densities and dashed lines represent prior densities.

Figure 2: Welfare Implications of Estimated DSGE Model

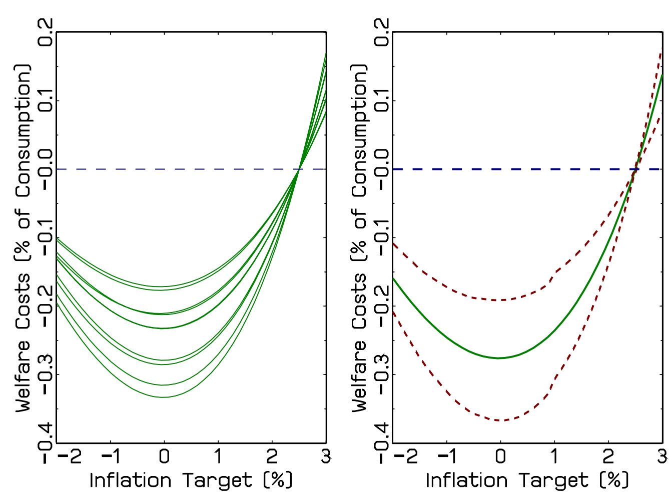

Notes: The left panel depicts several draws from the posterior distribution of steady state welfare costs (in percent of consumption) of deviating from 2.5% inflation as a function of counterfactual target inflation. The right panel depicts pointwise posterior means and credible intervals.

Figure 3: New Keynesian Phillips Curve Estimates

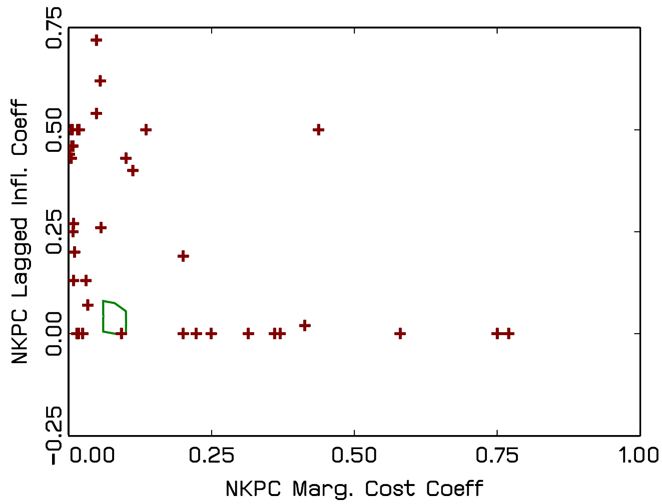

Notes: 90% credible set obtained from estimated DSGE model is denoted by solid contours. Point estimates reported in the papers surveyed in Schorfheide (2008) are indicated by "+".

Figure 4: Inverse Velocity: Actual and Counterfactual Path

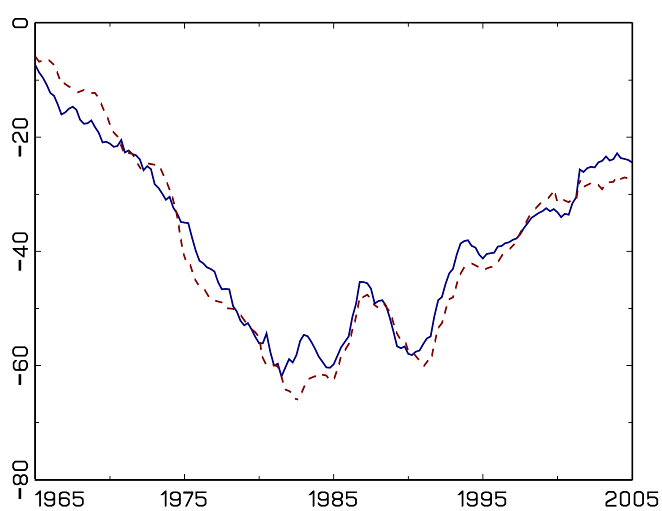

Notes: The solid line depicts actual inverse velocity and the dashed line depicts a counterfactual path that is solely based on money demand shocks.

Figure 5: TARGET INFLATION SHOCK IMPULSE RESPONSES – DSGE VS. VAR

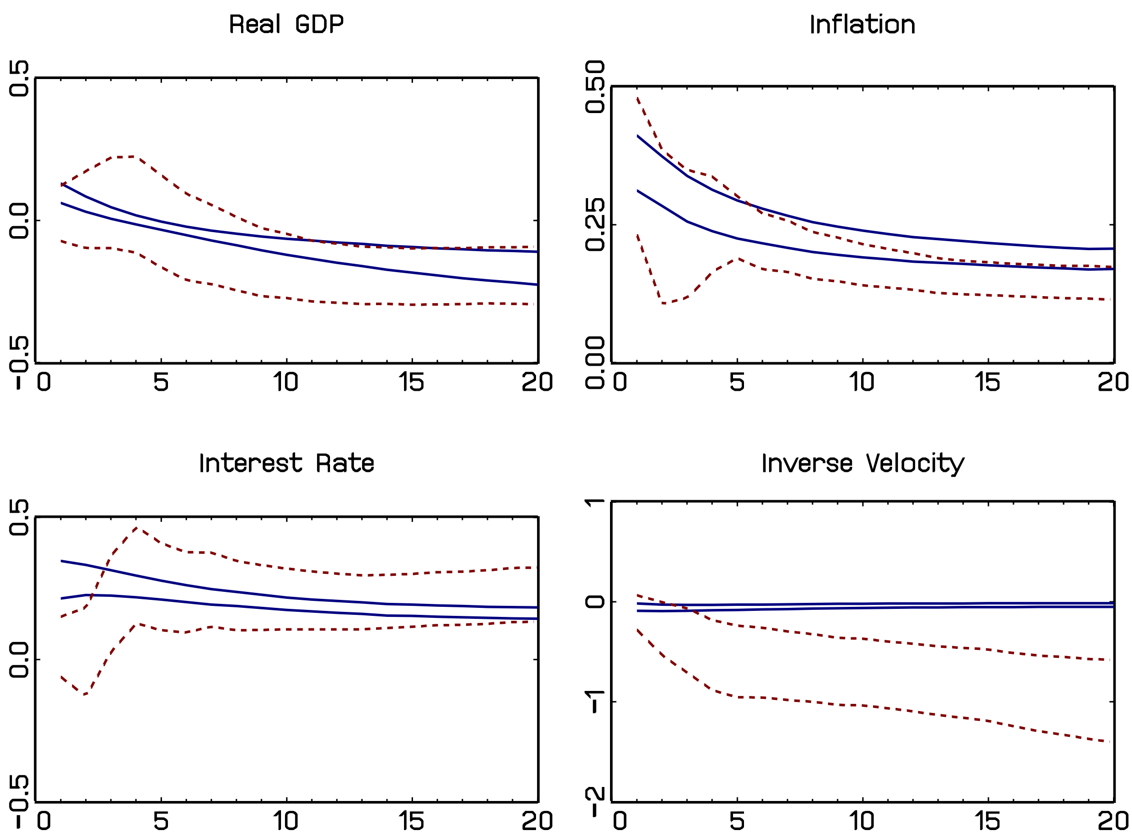

Notes: 90% Credible bands for impulse responses to a change in the target inflation rate for DSGE model (solid) and VAR (dashed).

Figure 6: THE ROLE OF ν: IMPULSE RESPONSES AND WELFARE

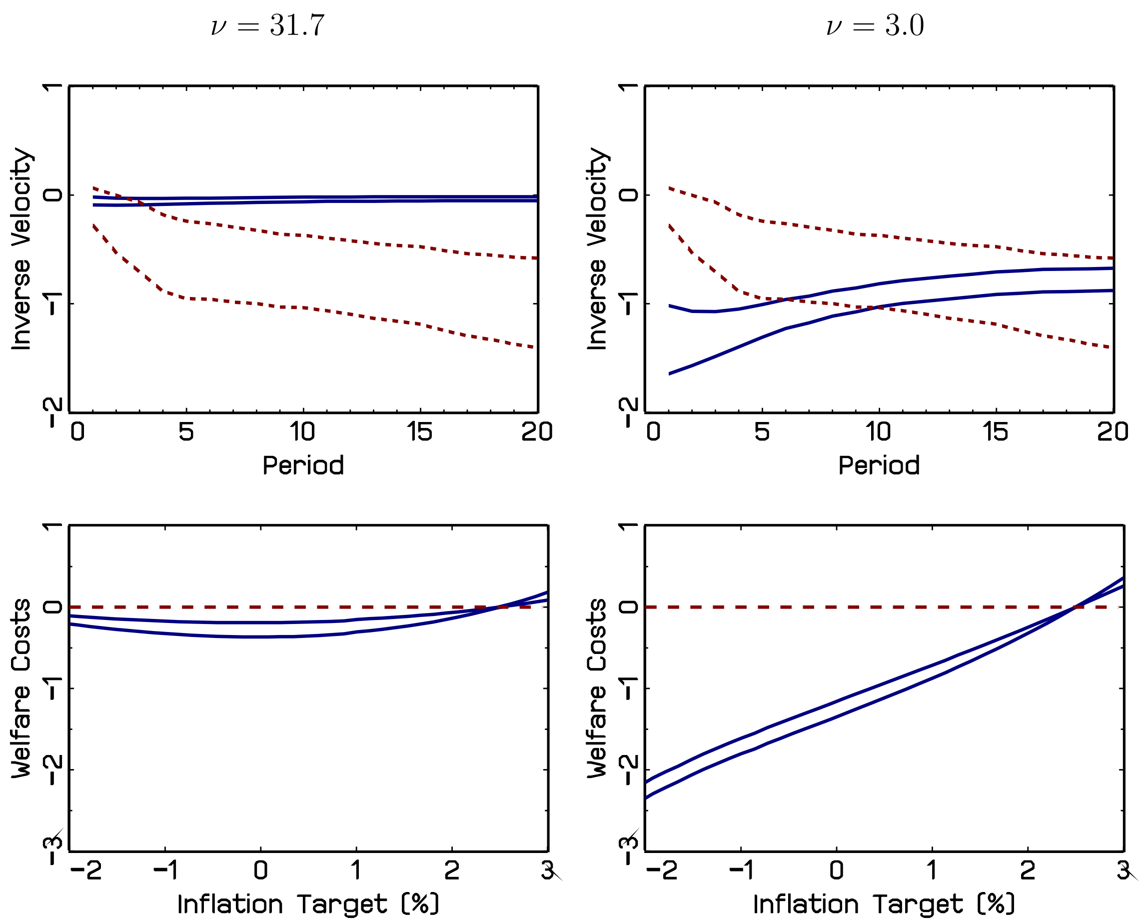

Notes: Top panels: 90% Credible bands for impulse responses to a change in the target inflation rate for DSGE model (solid) and VAR (dashed). Bottom panels: pointwise 90% credible intervals of steady state welfare costs (in percent of consumption) of deviating from 2.5% inflation as a function of counterfactual target inflation. The left-hand-side panels are generated based on the posterior distribution of ν, which has a mean of $\hat{\nu} = 31.7$. The right-hand-side panels are based on fixing $\nu = 3$.

Figure 7: Example – Contours of the Likelihood Function for $T = 100$

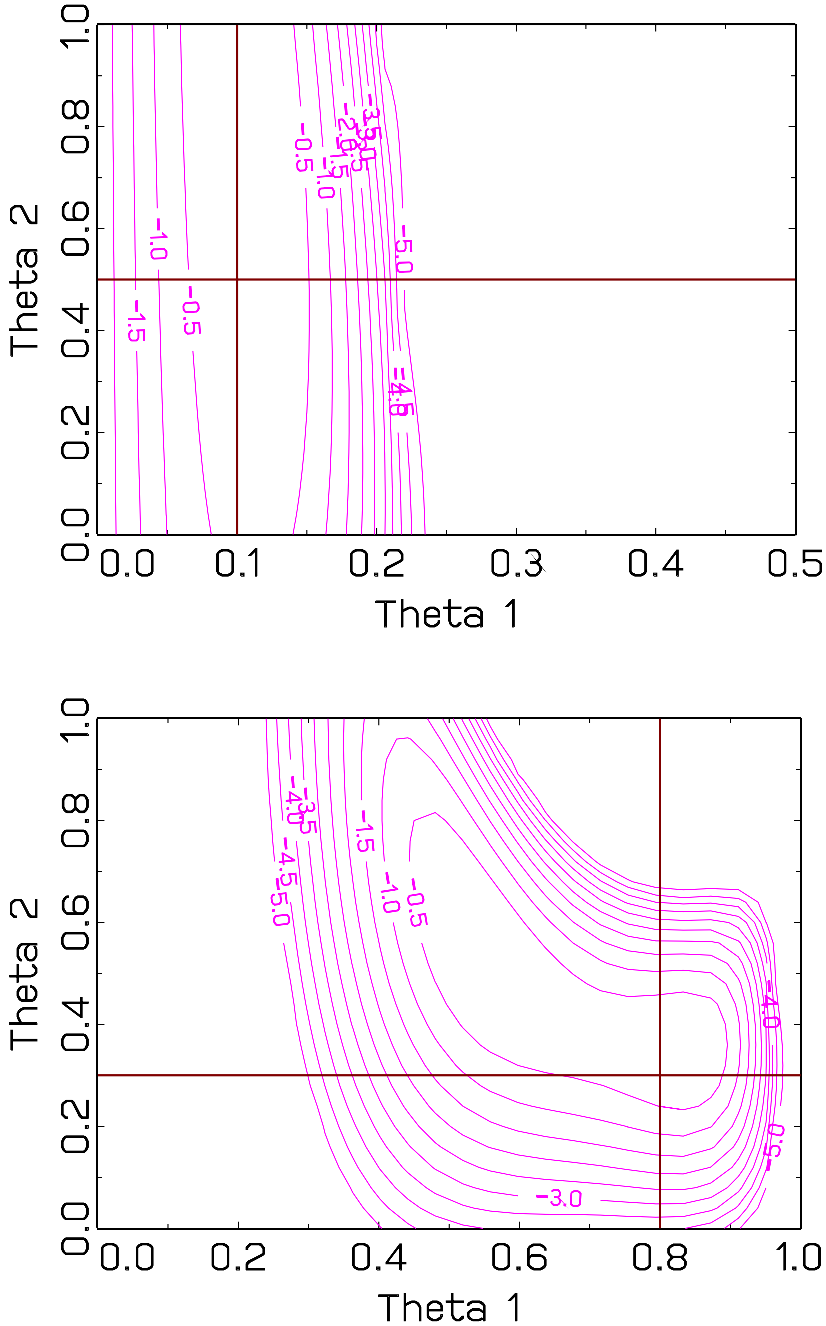

Notes: The intersection of the solid line indicates the parameter value that was used to simulate the observations from which the likelihood function is constructed.

Figure 8: DSGE-VAR Estimation

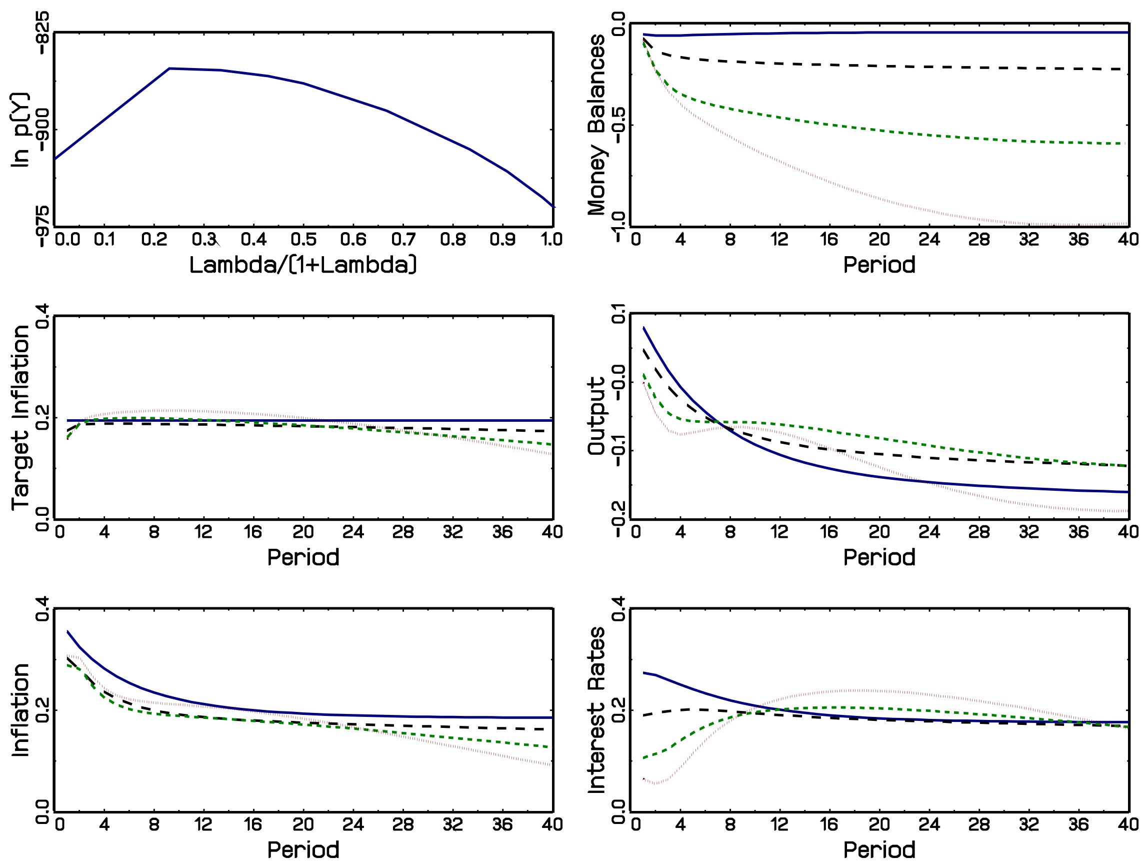

Notes: The top left panel depicts the log marginal data density of the DSGE-VAR as a function of $\lambda/(1+\lambda)$. The remaining panels depict posterior mean impulse responses computed from the DSGE-VAR for various values of λ, ranging from $\lambda = 0$ (solid) to $\lambda = \infty$ (dotted).