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

USING DSGE MODELS

Jonathan Huntley (E-mail: jonathan.huntley@cbo.gov)
Congressional Budget Office
Washington, D.C.

Eric Miller (E-mail: eamiller@artsci.wustl.edu)
Formerly of Congressional Budget Office
Washington, D.C.

November 2009
2009-04

Abstract

This paper is intended to be pedagogical rather than a presentation of original research. We describe a simple dynamic, stochastic general equilibrium (DSGE) model with capital utilization, capital adjustment costs, and a simple Cobb-Douglas technology to illustrate how DSGE models can be used to explain the past and to forecast the future. We identify one method to directly estimate latent variables and parameters in a DSGE model. We then construct estimates of the latent variables and shocks, the latter of which drive observed variations in economic activity. Those latent variables form the foundation of our economic analyses of past events, and the estimated parameters allow us to construct an economic forecast.

1 Introduction

In the past decade, advances in economic modeling—including the application of direct estimation of these models—have resulted in models that can generate results that often closely match real-world dynamics. The latest generation of models can be used to analyze historical economic events, current economic conditions, and hypothetical changes in policy.

In this paper, we discuss one method of directly estimating a class of dynamic stochastic general equilibrium (DSGE) macroeconomic models. The purpose of the exercise is not to develop new research, but rather to demonstrate the capacity for the approach to be used in the service of the analysis produced by the Congressional Budget Office (CBO). The particular method is of interest because DSGE modeling is fairly straightforward and flexible, and it does not require extensive computing power. DSGE models bring together advances in computational economics and time series econometrics to produce estimates of structural parameters, latent variables, and economic shocks that generate economic fluctuations. The analytical process has been applied to real business cycle (RBC) models by Alejandro, Primiceri, and Tambalotti (2007) and by Justiniano and Primiceri (2008).

The paper begins by outlining a model, taking the first-order necessary conditions (FONCs) and resource constraints that describe the solution to the model and linearizing them. Using a method described by Sims (2002) for solving rational expectations models, we reorganize the linear system of equations. What results is a simple, step-ahead transition matrix that approximately describes the expected evolution of the economy, conditioned on some preexisting state.

The solution formulated by Sims (2002) fits the form of the state equation of the Kalman filter. Using the data in the observation equation of the Kalman filter allows us to generate estimates of the model's latent variables and a log-likelihood based on that particular parametrization. After applying prior distributions on those parameters, some of which could be degenerate distributions for those parameters that we want to calibrate directly, we maximize the resulting posterior likelihood over the space of structural parameters that govern the behavior of the agents in the model. The relative simplicity of the Kalman filter allows us to compute the estimate fairly quickly. For the simple example we use, the solution is found within a few minutes.

The approach we describe here offers several advantages. First, the system is internally consistent. Historical fluctuations can be decomposed into shocks whose interpretations are precise and mutually exclusive. Second, we can use the model for forecasting. Several attempts have been made, and a few, such as that of Edge, Kiley, and Laforte (2009), have claimed success using DSGE modeling although much work remains in assessing the models' abilities to forecast economic conditions. Third, DSGE models allow analysts to run experiments to test the effects of policy changes. The models permit a good deal of flexibility in setting up experiments and creating scenarios.

Gathering worthwhile information from DSGE models is predicated on having a specification that represents a reasonable approximation of the economy. Research is adding to the list of reasonable features that characterize DSGE models, many of which rely on microfoundations for support. Although the DSGE model we present is simple, there are variations that include working capital for firms, rigid prices and wages, and other frictions that help explain observable economic dynamics.

In the second section of this paper, we describe an algorithm for estimating a generic model. We review the steps of linearizing, transforming, and estimating the model. In the third section, we present a simple model to which we apply the process. The fourth section details estimates of the model's parameters. The fifth section gives a simple forecast that can be generated from the DSGE model and describes other policy-related applications. The sixth section is the conclusion. The appendix consists of several figures and a table.

2 The Algorithm

2.1 The Model

The solution to the model is characterized by a number of FONCs, market-clearing conditions, laws of motion, and other relationships. Equation 1 defines the vector of relationships:

$$E_t[f(X_t; \Xi)] = 0 \tag{1}$$

The term Ξ represents a vector of the structural parameters, and the term E_t is the expectations operator at time t. The steady-state solution to this problem is $\widetilde{X}$, such that $X_t = \widetilde{X}$ solves Equation 1 at every time t.

2.2 Linearizing the Model

We log-linearize around the steady state $\widetilde{X}$ and define the vector of functions such that

$$0 \approx E_t[f(\widetilde{X}; \Xi)] + E_t[\frac{\partial f(\widetilde{X}; \Xi)}{\partial \ln X} \widehat{x_t}] \tag{2}$$

In this definition, $\widehat{x_t}$ is the vector of log-deviations from steady state. By explicit construction, $E_t[f(\widetilde{X}; \Xi)] = 0$, that term drops out, and we are left with

$$0 \approx E_t[\frac{\partial f(\widetilde{X}; \Xi)}{\partial \ln X} \widehat{x_t}] \tag{3}$$

The elements of the matrix $\frac{\partial f(\widetilde{X};\Xi)}{\partial \ln X}$ are functions of the model's steady-state values and its structural parameters.

2.3 Solving Linear Rational Expectations Models

We define y_t: a vector of variables such that the union of the elements in y_t and y_{t-1} is a superset of those that comprise $\widehat{x_t}$. For example, assume a very simple model with consumption and capital. Assume that the FONCs reduce to a system of $n = 2$ equations populated by c_t, k_t, and k_{t+1}. Therefore

$$\widehat{x_t} = \begin{pmatrix} \widehat{c_t} \\ \widehat{k_t} \\ \widehat{k_{t+1}} \end{pmatrix} \tag{4}$$

A suitable candidate for y_t is

$$y_t = \begin{pmatrix} \widehat{c_{t+1}} \\ \widehat{k_{t+1}} \end{pmatrix} \tag{5}$$

Then we use the method detailed by Sims (2002) (from which the subsection title is drawn). We rearrange Equation 3, using the newly defined y_t, to get Equation 6:

$$\Gamma_0 y_t = \Gamma_1 y_{t-1} + C + \Psi z_t + \Pi \eta_t \tag{6}$$

In this formulation, the $n \times n$ matrices Γ_0 and Γ_1 are constructed from the elements of $\frac{\partial f(\widetilde{X};\Xi)}{\partial \ln X}$. The term C is any remaining constant term, typically 0. The random shock, z_t, impacts the system by the impact term Ψ. η_t represents an expectational error. For example, a model with habit persistence will have c_{t+1}. At time t, c_{t+1} is not known with certainty, so we must include an expectational error.

Typically, Γ_0 is a singular matrix, so we apply the Sims (2002) algorithm to get Equation 7:

$$y_t = Gy_{t-1} + C_1 + \Delta z_t \tag{7}$$

in which G is an $n \times n$ matrix and Δ is the impact of an independent and identically distributed (i.i.d.) shock z_t. (Equation 7 is more complicated for the case in which the shock is serially correlated). In examples in which we formulate y_t to be a vector of deviations from steady-state values, C_1 is necessarily equal to 0. In cases in which y_t may include another term, such as a scale factor, C_1 may take some nonzero value.

2.4 Using the Kalman Filter

The Kalman filter is characterized by two equations, a state equation and an observation equation:

$$y_{t+1} = Gy_t + C_1 + v_{t+1} \tag{8}$$

$$\overline{y_t} = A'o_t + H'y_t + w_t \tag{9}$$

In these equations, o_t is an exogenous or predetermined term—in this case, it will typically be a constant. The values that comprise $\overline{y_t}$ are the actual measured components of y_t for which there are data (real national accounts data, prices, wages). The error terms are i.i.d. and have the property that $E[v_t v_t'] = Q$. w_t is assumed to be 0, which means that there is no measurement error. Equation 8 is the same as Equation 7. The Q in $E[v_t v_t'] = Q$ is going to be a function of Δ and any structural parameters Ξ that define the relevant moments of z_t.

G and y_t in Equations 8 and 9 were deliberately chosen because they represent the same variables in Equation 7. By carefully constructing A', o_t, and H', we explicitly relate the data to the states in the Kalman filter, which are the model variables' log-deviations from their steady states.

There also can be elements of y_t, such as preference shocks or total factor productivity (TFP), that are not directly observable; the Kalman filter provides estimates for those unobservable values. Finally, the Kalman filter produces a log-likelihood that is a function of the data $\overline{y_t}$ conditioned on the parameter set Ξ.

2.5 Estimating the Model

We are interested in maximizing $P(\Xi, \overline{y_t})$. We begin by rewriting the probability using Bayes' theorem:

$$P(\Xi, \overline{y_t}) = P(\overline{y_t}|\Xi)P(\Xi). \tag{10}$$

The Kalman filter produces $P(\overline{y_t}|\Xi)$; it gives us the likelihood conditioned on the set of parameters used to define the problem. $P(\Xi)$ represents our prior beliefs about the values of the structural parameters—important because some have important restrictions.

We maximize the likelihood $P(\Xi, \overline{y_t})$ over the parameter space Ξ by computing those two components on the right side of Equation 10.

3 The Model

The model we present is a simple version of a model described by Christiano, Eichenbaum, and Evans (2005). We remove money from the economy and, by extension, price and wage stickiness, but we include most other important details. We borrow some elements from Alejandro, Primiceri, and Tambalotti (2007), whose work is similarly based on this class of models. A representative household provides labor and capital to the market, taking a wage and rental rate as given. The household chooses investment, consumption, and a utilization rate subject to the typical budget constraint. In addition, the household pays a capital adjustment cost for investments that deviate from its investment in the last period and a capital utilization cost in periods in which capital is more highly utilized.

3.1 The Household

A representative household's problem is

$$\max_{\{c_t,h_t,u_t,k_{t+1},i_t\}} E_0 \sum_{t=0}^{\infty} \beta^t \xi_t^b \Big(log(c_t - \theta c_{t-1}) - \xi_t^h \phi \frac{h_t^{1+\sigma}}{1+\sigma} \Big) \tag{11}$$

In this model, c is consumption, h is labor effort, u is capital utilization, and i is investment. Consumption and investment, c and i, are scaled variables; actual consumption and investment are multiplied by a subsequently defined scale factor X_t. The household's utility function embodies habit formation; the household utility is a function of consumption's deviation from a fraction θ of the household's consumption in the last period. Furthermore, the disutility of labor is separable from consumption. The disutility of labor varies each period, modified by ξ_t^h, subject to the following law of motion:

$$\xi_t^h = (\xi_{t-1}^h)^{\rho_h} \epsilon_t^h \tag{12}$$

The term $log(\epsilon_t^h)$ is an i.i.d. term distributed normally with zero mean and variance σ_h^2. Similarly, there is variation in the household time value of money β, captured by ξ_t^b, which is subject to

$$\xi_t^b = (\xi_{t-1}^b)^{\rho_b} \epsilon_t^b \tag{13}$$

where $log(\epsilon_t^b)$ term is an i.i.d. term that is distributed normally with zero mean and variance σ_b^2.

The household's budget is described by two equations, the resource constraint

$$c_t + i_t = r_t u_t k_t + w_t h_t \tag{14}$$

and the law of motion for capital

$$\gamma k_{t+1} = (1-\delta)k_t + a(u_t)k_t + s(\frac{i_t}{i_{t-1}})i_t \tag{15}$$

in which $a(.)$ is the cost for capital utilization, $s(.)$ represents the capital adjustment costs, and γ is the rate of population growth.

The functional form of the capital utilization costs is not explicitly formulated, but it is defined by the following features: Capital utilization costs $a(.)$ are defined by a curvature parameter χ. In steady state, u is normalized to 1, $a(1) = 0$, and $\frac{a''(1)}{a'(1)} = \chi$. The capital adjustment cost function $s(.)$ is subject to three conditions. First, $s(\omega) = 1$, where ω is the rate of steady state growth, so that no costs are absorbed by the household along the steady-state-growth path. Second, we impose the condition that $s''(\omega) = \nu > 0$, so that the cost function is along the steady-state-growth path. Third, we set $s'(\omega) = 0$.

3.2 The Firm

The representative firm has access to a simple Cobb-Douglas technology. The firm rents capital at rate r and uses it at utilization rate u. Labor is hired at wage rate w and profits are maximized by solving the following problem:

$$\max_{\{h_t,k_t\}} z_t X_t^{\alpha}(k_t u_t)^{1-\alpha} h_t^{\alpha} - w_t h_t - r_t k_t u_t \tag{16}$$

The z term is a stationary $AR(1)$ component of TFP. X is the trend component. The trend component of TFP grows at a constant rate:

$$X_t = \omega X_{t-1} \tag{17}$$

The law of motion for the z term is represented by

$$z_t = (z_{t-1})^{\rho_z} \epsilon_t^z \tag{18}$$

The term $log(\epsilon_t^z)$ term is an i.i.d. term that is distributed normally with zero mean and variance σ_z^2.

4 Estimates

To estimate the model, we use as many data series as there are shocks to the system. We use data on inflation-adjusted consumption and investment and average hours in manufacturing, all seasonally adjusted. We present the values of those series in Figure 1. The investment and consumption series are displayed in logs. The shaded vertical bars extend from the peak to the trough of each recession.

We assume that there are six latent variables: capital utilization, capital stock, the inflation-adjusted interest rate, the adjustment to the rate of time preference, the adjustment to the disutility of labor, and TFP. The last three of those series are subject to exogenous shocks. Although there are data on some series, such as utilization, we are effectively using the model to derive alternative or competing estimates for those variables. It is unlikely that our estimates are legitimate competition for the published estimates of the series because the model is simplified for the purpose of illustration.

For the structural parameters, we estimate a subset and calibrate the rest. The calibrated subset includes the time value of money, $\beta = 0.99$, and the population growth factor, $\gamma = 1.004$. The remaining subset includes the parameters that define capital adjustment and utilization costs, all elasticities, persistence of the three shocks, the depreciation rate, the habit formation parameter, and the labor share of income.

For each parameter, we establish a set of loose, or unrestrictive, priors, the means for which are based roughly on calibrations that are used in the previously cited studies, such as Justiniano and Primiceri (2008). We use beta distributions for the bounded parameters. For some of the other parameters, particularly a subset with a bound only on one side, we use an inverse gamma prior. For the other parameters, we employ a normal distribution. The formulation of the priors and their moments is described in Table 1. We maximize the posterior likelihood to obtain point estimates of the parameters and employ the Metropolis-Hastings algorithm to simulate the posterior densities for all of the estimates. In Table 1, we report estimates of key moments of the posterior distribution.

Most parameters do not wander too far from their prior means. The income share to labor drifts away to 0.57. The capital adjustment cost parameter is estimated at 0.153, which is even smaller than the low calibrated value used by Christiano, Eichenbaum, and Evans (2005). Habit formation comes in at 0.315, suggesting that habit is still an important part of the model. The elasticity of labor parameter, at 5.49, is much higher than the typical values of 1 or 2. All shocks show a good deal of persistence, ranging from 0.9 to 0.95 for each of the three autoregressive processes.

Figure 2 shows three latent variables: multipliers to the disutility of work, TFP, and the discount rate. The values are presented as log deviations from the steady state. The top panel represents estimated changes to the time value of money over the course of the sample. The middle panel represents the same values but for the multiplier that determines the disutility of labor. Finally, deviations in TFP are presented in the bottom panel.

Declines in the discount rate are somewhat correlated with recessionary episodes. More strongly correlated are increases in the disutility of labor and declines in TFP, the latter of which is particularly strong. In this simple model, the disutility of labor moves around substantially, especially during recessions. Attributing declines in economic activity to a change in preferences for labor effort is neither realistic nor satisfying. It is possible to estimate how much of the total output variation is attributable to shocks to each variable, although we do not perform that particular experiment. Because the system is linear, the effects of the orthogonal shocks can be studied separately and then summed to derive the overall response. Using such an approach allows us to decompose the effects of each shock and determine whether relatively large changes in the parameter for disutility of labor have any meaningful economic consequences.

Fortunately, studies of larger models, such as that by Alejandro, Primiceri, and Tambalotti (2007), do not attribute much of the economic variation to this particular shock. In our simple model, we see large changes in that variable because there are few other outlets to explain variations in economic activity.

In addition to providing estimates of unobserved variables, DSGE models can be used to construct latent variables that inform policy. Those variables end up being functions of variables explicitly modeled, estimated structural parameters, and observable data. One simple example that we include is our estimate of the output gap. In this paper, the output gap is a fairly straightforward concept. There exists in the model a unique steady-state-growth path. An economy not subjected to any exogenous shocks and already in a state along that path will continue in perpetuity. That growth path makes for an excellent, if possibly unrealistic, baseline against which to judge the relative strength of the economy. The output gap exists in this model simply as the observed deviations from the estimated path of steady-state growth, which in this case is simply trend growth. We present those series, plotted against the observed values and in log terms as appropriate, in Figure 3.

Figure 4 is a comparison our calculations based on the model with CBO's estimates of the output gap over the past several decades[1]. Figure 4 provides a competing estimate and evaluation based on a structural model for a commonly used and referenced economic indicator. There are many similarities between the two series, again because the model simply measures deviations from an estimated growth trend.

More complicated models could have different views of the output gap: Edge, Kiley, and Laforte (2008) and Justiniano and Primiceri (2008) construct several variations. Simply introducing a unit root shock to TFP confuses the definition of output gap. In this case, there is no simple trend—the trend jumps in response toward unit root shocks to TFP. Instead of using such a trend that shifts over time, we could define baseline economic output as the results of a counterfactual scenario in which only the unit root TFP shock is applied, irrespective of any other perturbations. Either approach would yield a plausible estimate of the output gap. In addition to the output gap, the models we examine can yield alternative estimates of the natural rate of unemployment, the size of the capital stock, or any number of other variables of relevance to public policy.

Interpreting the past is only one use for DSGE models. In the next section, we explore the use of models for economic forecasting.

5 Forecasts

In addition to investigating past economic conditions, the DSGE models can be used for economic forecasting. We generate a forecast by applying Equation 7. At any time $t-1$, we take vector y_t and multiply it repeatedly by the transition matrix G to generate predictions for the T periods $t+1, t+2, ..., t+T$. The resulting prediction assumes that there are no further realized shocks.

[1]An explanation of how the output gap is estimated is available in Congressional Budget (2001).

In Figure 5, we present our forecast compared with CBO's January 2009 forecast. Our model predicts a fairly sharp recovery starting at the end of the 2009, asymptotically approaching the steady-state-growth rate. It predicts a higher economic output through about 2013 than does the CBO's baseline projection, which does not predict recovery in earnest until 2011.

Although we present only a straightforward projection in Figure 5, the method allows us to construct marginal distributions of the possible forecasts as well. By calibrating the start of the experiment to current economic conditions, we can run a simulation to generate a wide range of outcomes and to create quantile graphs.

Policy changes also can sometimes be included in the models. A richer model that included a fiscal sector would allow us to include proposed changes in policy and then to forecast the implications of those changes. Although the Aiyagari (1994) classes of models, except that of Krussel and Smith (1998), generally do not incorporate aggregate uncertainty because of computational limitations, the DSGE class of models is much more flexible. Not only can such models accommodate changes in structural parameters that represent regime movements such as policy changes, they can make forecasts without regard to the current state of the economy.

6 Conclusion

We have described the method we use to directly estimate structural economic models. After describing the algorithm for approximating and estimating the models, we apply the algorithm to a very simple RBC model. Our model has a few frictions, such as capital adjustment costs, capital utilization costs, and habit formation, and three stochastic processes that drive the observed fluctuations. Shocks to the time value of money, the disutility of labor, and TFP account for all of the variation observed in the data.

With our DSGE model, we can estimate values of the latent variables (providing insight into history), and we can make economic forecasts. In this paper, we show samples for both exercises. Because the model desribed here is overly simple, there is little reason to interpret the results as anything but a reduced-form representation for other shocks and frictions that could be driving the observed fluctuations in the economy. Regardless of one's faith in the model's specification, however, the exercise offers the opportunity to tell a story that is internally consistent and that is informed by a wide range of economic indicators.

We chose not to include a variety of frictions that could easily be added to the stock model, although many of those features have been useful in a variety of applications. Some are implemented in models in which there is a central bank that controls the money supply. One possible constraint is a cash-in-advance borrowing constraint under which agents must set aside money in advance to smooth the purchase of consumption and investment goods. In such models, there is a technology associated with money that determines the cost in terms of real resources of economic transactions.

Staggered price and wage setting as described by Calvo (1983) can add yet another dimension of realism that allows models to capture the delayed response of inflation and nominal wages to shocks to the monetary supply. In such models, firms are allowed to reset their prices at intervals determined by a stochastic process. Although that is a poor representation of price-setting behavior in high-inflation economies, as shown by Gagnon (2007), it could be a useful reduced-form representation of the process for economies in which inflation is more modest. Similarly, staggered wage setting, such as that detailed by Erceg, Henderson, and Levin (2000), approximates nominal frictions in wage setting.

Alejandro, Primiceri, and Tambalotti (2007) detail sources of variation not covered in this simple model. Their experiment includes shocks to monetary policy objectives, price mark-ups (monopoly-pricing power), wage mark-ups (wage-setting power), the productivity of installing investment goods, and government spending. Any number of those features can be added to the model to meet a variety of analytical demands.

The principle advantage of DSGE models is that they can be adapted to suit a wide variety of applications. Whereas vector autoregression (VAR) models can be used for forecasting and to describe past economic shocks, the VAR approach cannot be adapted easily to policy analysis. A structural model, such as the DSGE model, can be used to test policy experiments. Because all agents in the DSGE model make decisions, the model can capture the endogenous effects of changes in policy in a way that a VAR model cannot. Another advantage of a DSGE model comes from its transparency. Even when the features of the model are reduced-form representations for more complex processes, the DSGE model can provide a more precise insight about the nature of the shocks that drive observed economic variations than alternative modeling strategies might be able to provide.

References

Aiyagari, S.R., "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 1994, pp. 659–684.

Alejandro, J., G. Primiceri, and A. Tambalotti, "Investment Shocks and Business Cycles," *mimeo, Northwestern University*, 2007.

Calvo, G., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 1983, *12* (3), 383–98.

Christiano, L.J., M. Eichenbaum, and C.L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 2005, *113* (1), 1–45.

Congressional Budget Office, "CBO's Method for Estimating Potential Output: An Update," 2001.

_ , "The Budget and Economic Outlook: Fiscal Years 2009 to 2019," January 2009.

Edge, R., M.T. Kiley, and J.P. Laforte, "Natural Rate Measures in an Estimated DSGE Model of the US Economy," *Journal of Economic Dynamics and Control*, 2008, *32* (8), 2512–2535.

_ , _ , and _ , "A Comparison of Forecast Performance Between Federal Reserve Staff Forecasts, Simple Reduced-Form Models, and a DSGE Model," *mimeo, Federal Reserve Board*, 2009.

Erceg, C.J., D.W. Henderson, and A.T. Levin, "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, 2000, *46* (2), 281–313.

Gagnon, E., "Price Setting During Low and High Inflation: Evidence from Mexico," *mimeo, Federal Reserve Board*, 2007.

Justiniano, A. and G.E. Primiceri, "Potential and Natural Output," *mimeo, Northwestern University*, 2008.

Krusell, P. and A.A. Smith Jr., "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 1998, *106* (5), 867–896.

Sims, C.A., "Solving Linear Rational Expectations Models," *Computational Economics*, 2002, *20* (1), 1–20.

Table 1: Estimated Parameters

			Prior		Posterior	
Coefficient	Description	Dist.	Mean	Std. Dev.	Median	Std. Dev.
α	labor share	beta	0.66	0.1	0.57	1.08e-5
δ	depreciation	beta	0.025	0.01	0.025	2.9e-5
σ	elasticity of labor	normal	2.0	0.75	5.49	1.43e-5
ω	trend TFP	normal	1.005	1.0025	1.004	2.63e-8
χ	capital utilization costs	beta	0.025	0.01	0.027	2.9e-5
θ	habit persistence	normal	0.5	0.2	0.315	1.26e-4
ν	capital adjustment costs	inverse gamma[1]	2.0	0.2	0.153	3.8e-5
ρ_b	persistence of ξ_b	beta	0.85	0.1	0.9	2.66e-5
ρ_h	persistence of ξ_h	beta	0.85	0.1	0.95	1.21e-5
ρ_z	persistence of z	beta	0.85	0.1	0.91	1.16e-5
σ_b^2	variance of ϵ_b	inverse gamma[1]	0.002	2.0	0.0000	2.1e-6
σ_h^2	variance of ϵ_h	inverse gamma[1]	0.002	2.0	0.0001	4.6e-6
σ_z^2	variance of ϵ_z	inverse gamma[1]	0.002	2.0	0.0003	1.62e-6
X_0	scale factor	log-normal	2.0	1.0	3.0	1.72e-5

[1] For inverse gamma distribution priors, instead of mean, we report mode and instead of standard deviation, we report degrees of freedom.

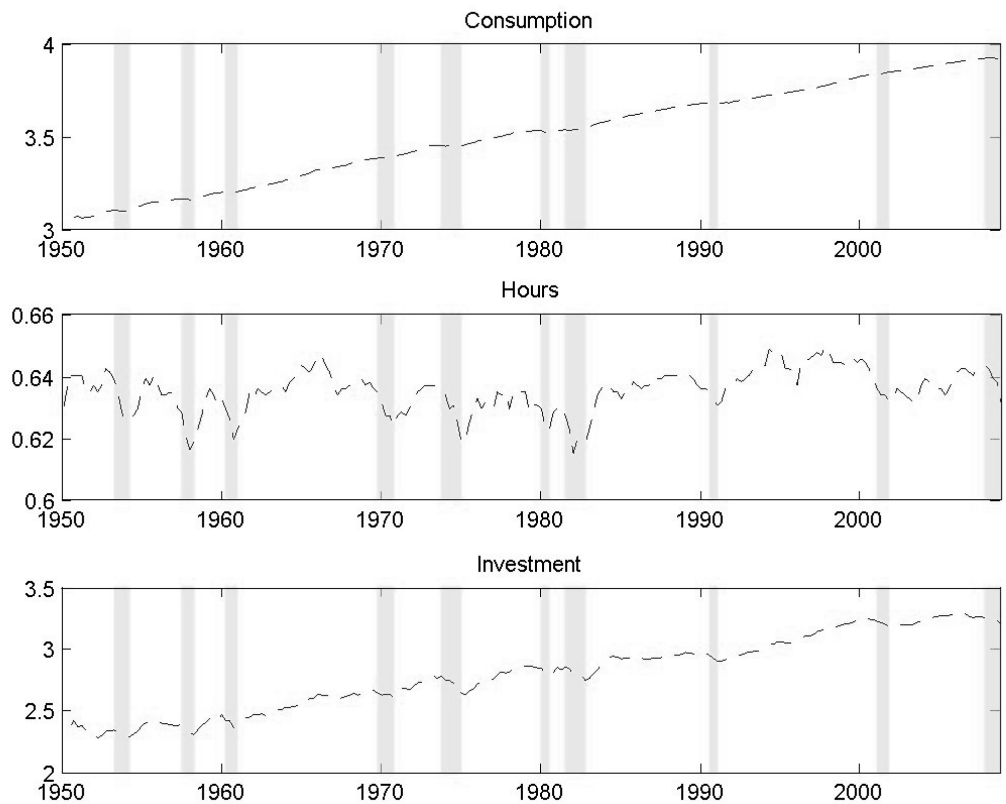

Figure 1: Estimated values of selected variables. Gray bars represent recessions from peak to trough. Consumption and investment in log values.

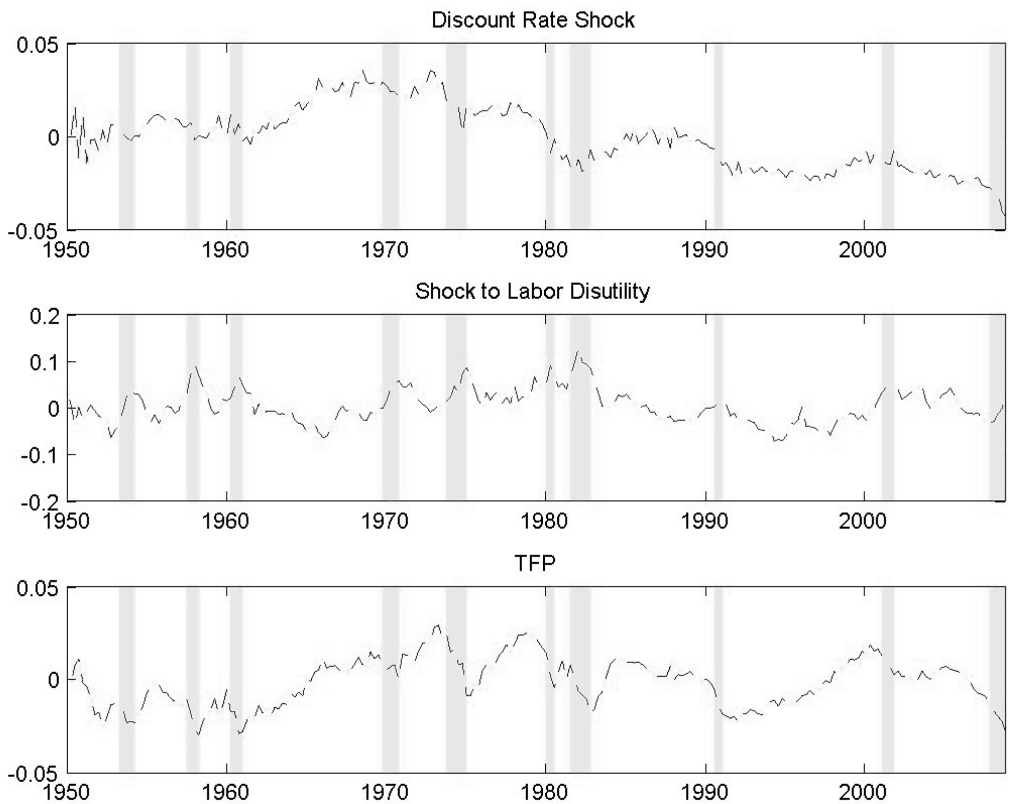

Figure 2: Estimated values of latent variables. Gray bars represent recessions from peak to trough. Variables in log deviations from steady state.

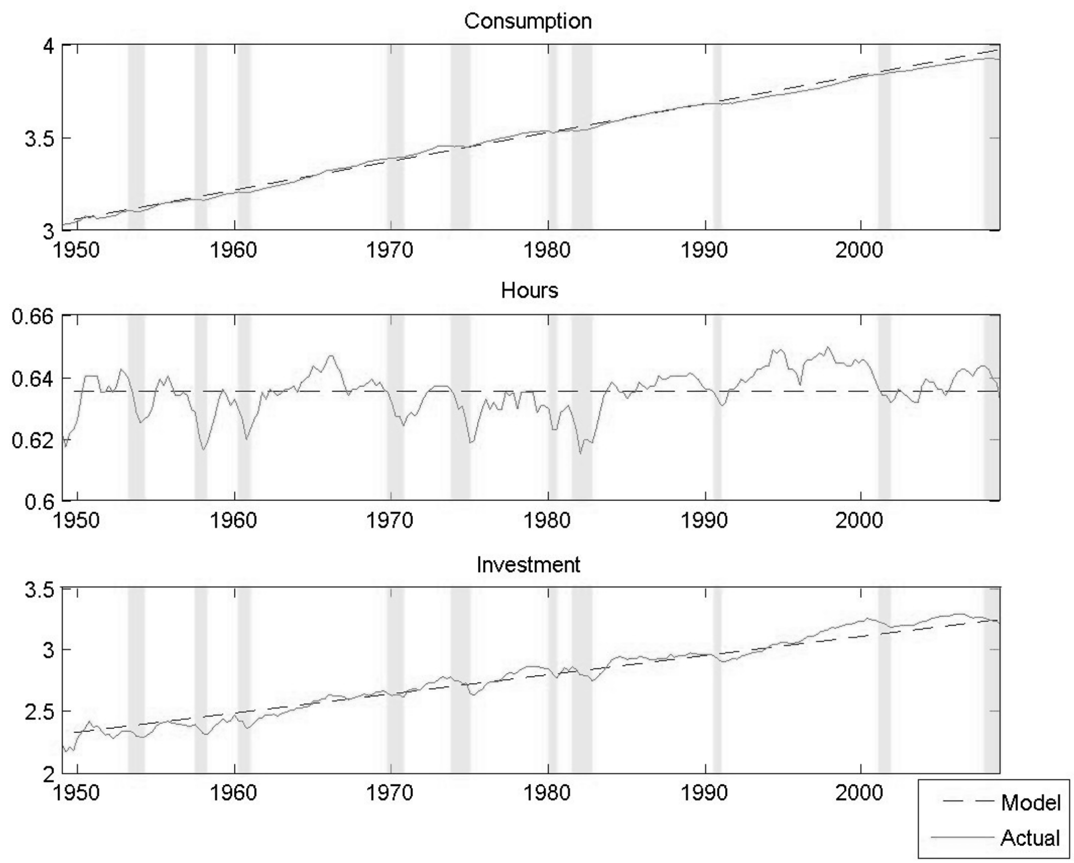

Figure 3: Estimated balanced growth path for observed variables. Variables in log values.

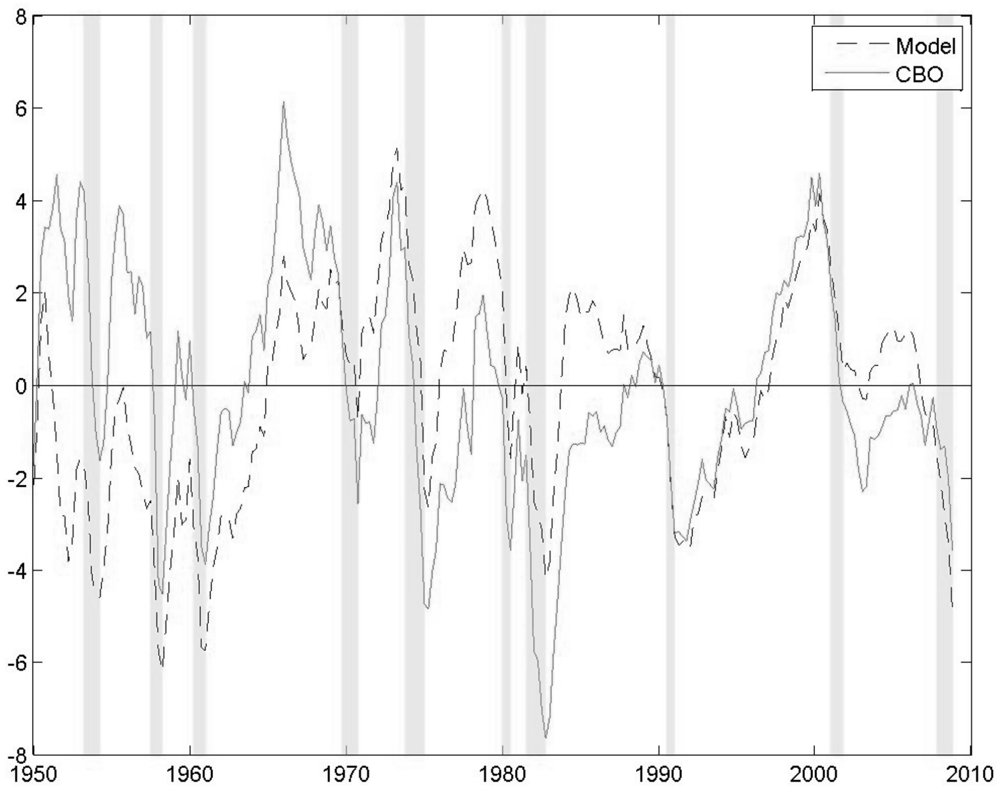

Figure 4: Estimated log difference from steady-state-growth path.

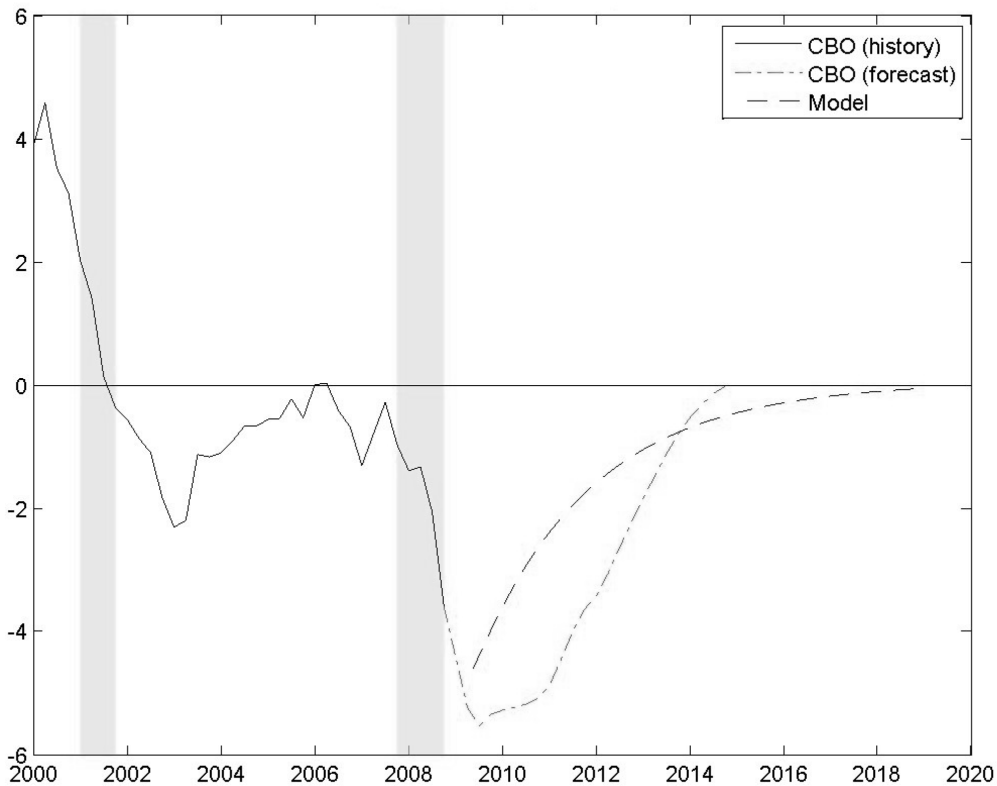

Figure 5: Forecast output gap.