

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

The Empirical Implications of the Interest-Rate Lower Bound

**Christopher Gust, Edward Herbst, David Lopez-Salido, and
Matthew E. Smith**

2012-083

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

The Empirical Implications of the Interest-Rate Lower Bound*

Christopher Gust
Federal Reserve Board

Edward Herbst
Federal Reserve Board

David López-Salido
Federal Reserve Board

Matthew E. Smith[†]
Hutchin Hill Capital

January 2017

Abstract

Using Bayesian methods, we estimate a nonlinear DSGE model in which the interest-rate lower bound is occasionally binding. We quantify the size and nature of disturbances that pushed the U.S. economy to the lower bound in late 2008 as well as the contribution of the lower bound constraint to the resulting economic slump. We find that the interest-rate lower bound was a significant constraint on monetary policy that exacerbated the recession and inhibited the recovery, as our mean estimates imply that the zero lower bound (ZLB) accounted for about 30 percent of the sharp contraction in U.S. GDP that occurred in 2009 and an even larger fraction of the slow recovery that followed.

Keywords: zero lower bound, DSGE model, Bayesian estimation

JEL classification: C11, C32, E32, E52

*Email Addresses: christopher.gust@frb.gov; edward.p.herbst@frb.gov; david.lopez-salido@frb.gov; and matt.e.smith@gmail.com, respectively. We thank Marty Eichenbaum, Chris Sims, and three anonymous referees for helpful comments, and Kathryn Holston for excellent research assistance. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

[†]This paper represents the personal work product and views of Matthew E. Smith and have neither been reviewed nor endorsed by Hutchin Hill Capital, LP, its affiliates, any funds advised by them, or any of their officers, directors, employees, agents, partners, principals, owners, or shareholders (collectively, "Hutchin Hill Parties"). No representation or warranty is made by Matthew E. Smith or any Hutchin Hill Party about the completeness or accuracy of any information provided herein, and such parties expressly disclaim all liability for any of the information contained in these materials. These materials have been prepared for educational use only and may not be relied upon for any other purpose (including without limitation any commercial purpose). These materials do not constitute an offer, or the solicitation of an offer, to buy or sell any security or financial instrument or investment advisory services. No Hutchin Hill Party shall be liable or have any obligation with respect to these slides or any reliance thereon.

1 Introduction

As a result of the financial crisis of 2007-2009, the U.S. economy experienced a sharp contraction in activity followed by a markedly slow recovery. At the same time, the federal funds rate fell and remained at its effective lower bound for a protracted period of time. Accordingly, the zero lower bound (ZLB) constraint became a practical concern for monetary policymakers, limiting their ability to offset contractionary disturbances to the economy. In light of these developments, an important question remains: What role did this constraint play in exacerbating the economic slump that occurred in the aftermath of the financial crisis?

In this paper, we address this question by formally estimating and evaluating a dynamic stochastic general equilibrium (DSGE) model. In doing so, the paper offers a methodological contribution by estimating a nonlinear DSGE model in which a lower bound is occasionally binding. We apply this methodology to the interest-rate lower bound, but our approach is more general and could be used in many different contexts, including models with financial constraints. When applied to the interest-rate lower bound, the technique allows us to identify the nature and size of disturbances that pushed the interest rate to the lower bound in the United States in late 2008 and quantify the role of the constraint in exacerbating the resulting economic slump. Moreover, our methodology provides a new perspective on the relative contributions of the endogenous and exogenous sources of business cycle fluctuations.

We estimate a nonlinear version of a medium-scale model widely used in macroeconomics and monetary economics. In particular, we use a model similar to those of Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007), as these models have been successful in providing an empirically plausible account of key macroeconomic variables including output, consumption, investment, inflation, and the nominal interest rate. The model includes several real and nominal frictions including habit persistence in consumption, costly adjustment of investment, and costs to adjusting nominal wages and prices. Because of these features, the model has many state variables, and we demonstrate that it is possible albeit challenging to estimate a model of this size nonlinearly using Bayesian techniques.

Fluctuations in the model are driven by five shocks including those to total factor productivity, government spending, and monetary policy. We also include two other shocks that can be interpreted in a reduced-form manner as financial shocks. The first is the marginal efficiency of investment (MEI) shock discussed in Justiniano, Primiceri, and Tambalotti (2011), which affects the transformation through which goods can be turned into productive capital. As shown in Justiniano, Primiceri, and

Tambalotti (2011), this shock can be viewed as a disturbance to the financial sector’s ability to channel savings into investment. The second is a disturbance to the household’s intertemporal Euler equation governing the purchase of a risk-free bond as in Smets and Wouters (2007) and Christiano, Eichenbaum, and Trabandt (2015). Following Smets and Wouters (2007), we refer to this shock as a ‘risk premium’ shock since it affects the spread between the risk-free rate and the return on risky assets.

In the model, the lower bound constraint is potentially an important factor influencing the stance of monetary policy and hence economic outcomes. By constraining the current interest rate, the lower bound limits the degree of monetary stimulus. However, monetary policy can still be effective in influencing expectations about future output and inflation, thereby affecting the current levels of prices and spending. Similarly, the current decisions of households and firms are influenced by the expectation that the constraint may be binding in future states.

Our solution method takes into account the effect that uncertainty about the ZLB has on the economic decisions of households and firms, as we do not impose certainty-equivalence or perfect foresight to solve the model. With this solution, we estimate the nonlinear model using Bayesian methods via a Markov chain Monte Carlo (MCMC) algorithm. Since we cannot evaluate the likelihood using the Kalman filter, we use the particle filter. Although others have estimated nonlinear models with these techniques, this paper is the first to our knowledge to do so using a nonlinear model with an occasionally binding constraint.

Our empirical approach offers several advantages. First, as noted above, it explicitly takes into account the effect of future uncertainty about whether the ZLB will bind or not on the economic decisions of agents, and this uncertainty is accounted for in the estimation of the model. In contrast, others including Ireland (2011) estimate a linearized DSGE model without the constraint, which neglects to consider how the lower bound systematically changes monetary policy and thus the behavior of economic agents—a key aspect at the heart of our analysis. Moreover, we find evidence that the risk of a binding ZLB constraint affected economic behavior during the 2003-2004 episode in which the federal funds rate was near-but-never-at the constraint.

A second advantage of our approach is that it uses all available data, including the period over which the ZLB binds, to estimate the model. In contrast, an alternative approach is to estimate a linear model over a sample period in which the ZLB does not bind and use those estimates to simulate a nonlinear

version of the model that imposes the constraint.¹ This approach precludes jointly estimating the parameters of the model and the shocks that occurred over the zero lower bound period and then using them to quantify the economic effects of the constraint, a central feature of our analysis.

We also find that there are important differences between the parameter estimates coming from the linearized version of the model and those coming from the nonlinear version, suggesting that inference based on the estimates of the linearized model is flawed. In particular, we find that the linear version of the model attributes an overly large role to the marginal efficiency of investment shock in explaining economic fluctuations. This result arises because the nonlinear version of the model has a stronger internal propagation mechanism than its linear counterpart. This stronger propagation reflects, in part, the nonlinearity associated with the ZLB but more importantly that the linear dynamics understate the sensitivity of investment to economic disturbances in models with costs to adjusting investment as in Christiano, Eichenbaum, and Evans (2005). Thus, our results highlight the importance of estimating the nonlinear model in order to properly quantify the relative contributions of endogenous and exogenous sources of business cycle fluctuations.

We estimate the model using U.S. quarterly data on output, consumption, and investment as well as inflation and the nominal interest rate from 1983 through 2014. For this sample period, there is a single episode spanning from 2009 until the end of the sample in which the nominal interest rate is effectively at the lower bound. We examine the model's estimated shocks during this episode and characterize which shocks were important in causing the Great Recession and in driving the nominal rate to the lower bound.

For the episode in which the interest rate was at the lower bound, the two financial disturbances – the risk premium and MEI shocks – were key driving forces pushing the interest rate down and keeping it at the lower bound during and after the Great Recession. However, the risk premium shock is the relatively more important of the two financial shocks, especially in explaining the large contraction in output, consumption, and investment that occurred. The productivity shock played a role in contributing to the slow economic recovery and in explaining why inflation did not fall more sharply during the Great Recession. Monetary policy shocks contributed little to the dynamics of inflation, output, and the nominal interest rate over the course of this episode. However, this does not necessarily imply that monetary policy was unimportant, as our results highlight that the systematic

¹ See, for example, Aruoba, Cuba-Borda, and Schorfheide (2016) and Christiano, Eichenbaum, and Trabandt (2015).

portion of the monetary policy rule was critical to understanding the dynamics of the economy during and in the aftermath of the financial crisis.

Using the estimates of the shocks, the central question we address is how much the lower bound constrained the ability of monetary policy to stabilize the economy. We answer this question by comparing our estimated model in which the lower bound is imposed to the hypothetical case in which monetary policy can act in an unconstrained manner. Our mean estimate from this counterfactual experiment implies that about 30 percent of the sharp contraction in output during the Great Recession was due to the interest-rate lower bound. The zero lower bound played an even larger role in accounting for the slow recovery following this contraction, as U.S. GDP did not recover back to its pre-recessionary level until the end of 2012. In the absence of the ZLB constraint, our mean estimates imply that output would have recovered to its pre-recessionary level about a year earlier and that the ZLB explains a little more than half of the lower output over the 2008-2012 period. Moreover, the ZLB constraint was an important factor in contributing to inflation running persistently below the Federal Reserve's target over this period. An important caveat, associated with these estimates is that there is considerable uncertainty surrounding them.

Our paper is most closely related to Chung, Laforte, Reifschneider, and Williams (2012), who use counterfactual simulations in which policy is unconstrained to evaluate the cost of the ZLB in the context of a large-scale macroeconomic model. By comparing actual outcomes to the counterfactual outcomes of unconstrained simple policy rules, they argue that the ZLB was a significant constraint on monetary policy, contributing to lower inflation and a higher unemployment rate. We also perform a similar counterfactual exercise; however, our methodology differs significantly from theirs, including the fact that we estimate the policy rule used in our counterfactual analysis.

Our paper is also related to Christiano, Eichenbaum, and Trabandt (2015) and Campbell, Fisher, Justiniano, and Melosi (2016), who study the effectiveness of the Federal Reserve's forward guidance policies during the ZLB episode. Although this question is important in assessing the cost of the ZLB, these papers do not disentangle the effects of the ZLB from other frictions that contributed to the Great Recession as we do. Similarly, Wu and Xia (2016) also focus on evaluating the benefits of unconventional policy rather than the cost of the ZLB constraint. In doing so, they use a term structure model to construct a measure of the shadow or desired rate that policymakers would choose if they could ignore the ZLB constraint. We also construct such a measure, though our estimates are

derived from an estimated structural model.

The rest of the paper proceeds as follows. The following section presents the macroeconomic model that we estimate, while Section 3 discusses how the model is solved and estimated. The next three sections present the model’s results. Section 4 discusses the model’s parameter estimates and time series properties of the model. Section 5 examines the estimated shocks that took the economy to the lower bound and that account for the Great Recession, and Section 6 quantifies the effects of the zero lower bound constraint on the economy. Section 7 presents some conclusions.

2 The Model

The model is similar to Christiano, Eichenbaum, and Evans (2005) or Smets and Wouters (2007). The economy consists of a continuum of households, a continuum of firms producing differentiated intermediate goods, a perfectly competitive final goods firm, and a government in charge of fiscal and monetary policy. We now lay out the objectives and constraints of the different agents as well as the sources of uncertainty.

2.1 Firms

There is a continuum of monopolistically competitive firms producing differentiated intermediate goods. These goods are used as inputs by a perfectly competitive firm producing a single final good. The final good is produced by a representative, perfectly competitive firm with a constant-returns technology, $Y_t = \left(\int_0^1 Y_t(j)^{\frac{\varepsilon_p - 1}{\varepsilon_p}} dj \right)^{\frac{\varepsilon_p}{\varepsilon_p - 1}}$, where $Y_t(j)$ denotes intermediate good $j \in [0, 1]$, and $\varepsilon_p > 1$ is a constant elasticity of demand. Profit maximization implies that the final goods producer’s demand for good j is given by $Y_t(j) = \left(\frac{P_t(j)}{P_t} \right)^{-\varepsilon_p} Y_t$, where $P_t(j)$ denotes the price for intermediate good j and P_t denotes the aggregate price level.

The production function for good j is:

$$Y_t(j) = K_t(j)^\alpha (Z_t N_t(j))^{1-\alpha}, \tag{1}$$

where Z_t denotes the aggregate level of technology, and $K_t(j)$ and $N_t(j)$ denote the capital and labor

services hired by firm j , respectively. Technology is assumed to evolve according to:

$$Z_t = Z_{t-1} G_Z \exp(\epsilon_{Z,t}), \quad (2)$$

where $\epsilon_{Z,t} \sim iid N(0, \sigma_Z^2)$ denotes a random disturbance to technological growth that causes it to depart from a deterministic growth rate, G_Z .

Intermediate goods firm j sells its output in a monopolistically competitive market and sets its nominal price $P_t(j)$ subject to a quadratic cost to adjusting its price as in Rotemberg (1982). We depart from Christiano, Eichenbaum, and Evans (2005) and use Rotemberg contracts instead of Calvo contracts to model nominal price and wage stickiness, as this choice helps reduce the dimensionality of the model's nonlinear state space.² A firm's cost to adjusting its price in period t is given by:

$$\frac{\varphi_p}{2} \left(\frac{P_t(j)}{\tilde{\pi}_{t-1} P_{t-1}(j)} - 1 \right)^2 Y_t,$$

where the parameter $\varphi_p \geq 0$ governs the size of the adjustment cost and Y_t denotes aggregate output. A producer's price change is indexed to $\tilde{\pi}_{t-1} = \bar{\pi}^a \pi_{t-1}^{1-a}$ where $\bar{\pi}$ denotes the central bank's inflation target and $\pi_{t-1} = \frac{P_{t-1}}{P_{t-2}}$. The parameter a satisfies $1 \geq a \geq 0$ and determines the extent to which price indexation is tied to the central bank's inflation target or to the lagged inflation rate. The cost of price adjustment makes the problem of the intermediate goods producer dynamic; that is, it chooses $P_t(j)$ to maximize its expected present discounted value of profits:

$$E_0 \sum_{t=1}^{\infty} \beta^t \Lambda_t \left[\left(\frac{P_t(j)}{P_t} - mc_t \right) Y_t(j) - \frac{\varphi_p}{2} \left(\frac{P_t(j)}{\tilde{\pi}_{t-1} P_{t-1}(j)} - 1 \right)^2 Y_t \right],$$

taking mc_t , Y_t , $\tilde{\pi}_{t-1}$, and Λ_t as given. In equilibrium, the variable Λ_t is equal to the marginal value of an additional unit of real profits to a household in period t , reflecting the ownership of the firm by a household in the economy. The variable mc_t denotes real marginal costs, which are the same for all of the intermediate goods producers and are given by $mc_t = \frac{1}{\Phi} (r_t^k)^\alpha \left(\frac{W_t}{P_t} \right)^{1-\alpha}$, where $\frac{W_t}{P_t}$ and r_t^k denote the aggregate real wage and rental rate of capital services, respectively, and $\Phi \equiv \alpha^\alpha (1-\alpha)^{1-\alpha}$.

² With Rotemberg adjustment costs the model has twelve state variables. If, instead, we used Calvo contracts, there would be two additional state variables reflecting the dispersion of wages across households and prices across firms.

2.2 Households

We assume that there is a continuum of monopolistically competitive households indexed by $i \in [0, 1]$ supplying a differentiated labor service, $N_t(i)$. This allows us to model nominal wage rigidities in an analogous manner to the nominal price rigidities described earlier. Household i sells labor services to a representative employment agency producing a single labor input, N_t , which is in turn supplied to the intermediate goods firms in a perfectly competitive labor market. The Dixit-Stiglitz aggregator for labor services is given by:

$$N_t = \left(\int_0^1 N_t(i)^{\frac{\varepsilon_w - 1}{\varepsilon_w}} di \right)^{\frac{\varepsilon_w}{\varepsilon_w - 1}},$$

with $\varepsilon_w > 1$. Profit maximization by the employment agency taking as given the aggregate wage, W_t , and each household's wage, $W_t(i)$, yields the following set of demand schedules: $n_t(i) = \left(\frac{W_t(i)}{W_t} \right)^{-\varepsilon_w} N_t$, $\forall i$.

A household's preferences are given by:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left\{ \ln(C_t(i) - \gamma C_{t-1}(i)) - \psi_L \frac{N_t(i)^{1+\sigma_L}}{1+\sigma_L} - \varphi_t^w(i) \right\}, \quad (3)$$

where $C_t(i)$ denotes the consumption of household i , and $\varphi_t^w(i)$ is the loss in utility associated with adjusting the household's wage, $W_t(i)$. The parameter γ governs the importance of habits in consumption. Household i also experiences disutility from working, with the parameter ψ_L affecting the level of this disutility and σ_L affecting the Frisch elasticity of labor supply. Similar to the cost of adjusting prices, the wage adjustment cost is quadratic:

$$\varphi_t^w(i) = \frac{\varphi_w}{2} \left[\frac{W_t(i)}{\tilde{\pi}_t^w W_{t-1}(i)} - 1 \right]^2,$$

where $\varphi_w \geq 0$ governs the size of this cost. Wage contracts are indexed to productivity and inflation, since $\tilde{\pi}_t^w = G_Z \bar{\pi}^{a_w} (\exp(\epsilon_{Z,t}) \pi_{t-1})^{1-a_w}$ with $0 < a_w < 1$.

A household's budget constraint in period t is:

$$C_t(i) + I_t(i) + R_t^{-1} \frac{B_{t+1}(i)}{\eta_t P_t} + a(u_t(i)) \bar{K}_t(i) = \frac{W_t(i)}{P_t} n_t(i) + r_t^k K_t(i) + \frac{B_t(i)}{P_t} + \frac{D_t}{P_t} - T_t. \quad (4)$$

A household receives labor income from the employment agency and income from purchases of a risk-free

nominal bond, $B_t(i)$. Each household receives the same amount of dividends, D_t , from the economy's intermediate goods producers and pays the same amount of lump-sum taxes, T_t . Household i also owns capital, $\bar{K}_t(i)$, which it combines with its desired level of capital utilization, $u_t(i)$, to transform into capital services, $K_t(i) = u_t(i)\bar{K}_t(i)$. To utilize physical capital and transform it into capital services, a household incurs the cost $a(u_t(i))\bar{K}_t(i)$, where the function $a(u_t(i))$ is given by:

$$a(u_t(i)) = \frac{r^k}{\sigma_a} \{ \exp(\sigma_a(u_t(i) - 1)) - 1 \}. \quad (5)$$

A household purchases the final good and uses it for consumption, $C_t(i)$, or investment, $I_t(i)$. Nominal bonds are purchased at the price $\frac{1}{R_t\eta_tP_t}$, where R_t denotes the nominal interest rate. Following Smets and Wouters (2007) and Christiano, Eichenbaum, and Trabandt (2015), η_t is an exogenous disturbance to a household's return on the risk-free bond. Fisher (2015) discusses how an increase in η_t , to first order, can be interpreted as a rise in the demand for risk-free bonds. This risk premium shock affects the spread between the nominal rate and the return on risky assets, and is assumed to evolve according to:

$$\ln(\eta_t) = \rho_\eta \ln(\eta_{t-1}) + \epsilon_{\eta,t}, \quad (6)$$

where $\epsilon_{\eta,t} \sim iid N(0, \sigma_\eta^2)$.

A household's purchases of investment augment the physical stock of capital according to:

$$\bar{K}_{t+1}(i) = (1 - \delta)\bar{K}_t(i) + \mu_t(1 - S_t(i))I_t(i), \quad (7)$$

where δ is the depreciation rate of capital and μ_t is an exogenous disturbance to the marginal efficiency of transforming final goods into tomorrow's physical capital. This shock is assumed to evolve according to:

$$\ln(\mu_t) = \rho_\mu \ln(\mu_{t-1}) + \epsilon_{\mu,t}, \quad (8)$$

where $\epsilon_{\mu,t} \sim iid N(0, \sigma_\mu^2)$. As discussed in Justiniano, Primiceri, and Tambalotti (2011), this shock can be interpreted as a reduced-form way of capturing financial frictions that affect the transformation of aggregate savings into investment. In accumulating capital, households must also incur an adjustment cost, $S_t(i)$, to transform investment into capital available for production next period. Following

Christiano, Eichenbaum, and Evans (2005), these adjustment costs are given by:

$$S_t(i) = \frac{\varphi_I}{2} \left(\frac{I_t}{G_Z I_{t-1}} - 1 \right)^2. \quad (9)$$

2.3 Monetary and Fiscal Policies

The central bank must set the nominal rate R_t in accordance with its lower bound constraint:

$$R_t = \max [1, R_t^N], \quad (10)$$

where R_t^N denotes the notional or desired rate that the central bank would like to set in the absence of the constraint. The notional rate is set according to:

$$\ln\left(\frac{R_t^N}{R}\right) = \rho_R \ln\left(\frac{R_{t-1}^N}{R}\right) + (1 - \rho_R) \left[\gamma_\pi \ln\left(\frac{\pi_t}{\bar{\pi}}\right) + \gamma_x x_t^g + \gamma_g \ln\left(\frac{Y_t}{G_Z Y_{t-1}}\right) \right] + \epsilon_{R,t}, \quad (11)$$

with $0 < \rho_R < 1$, $\gamma_\pi \geq 0$, $\gamma_g \geq 0$, and $\gamma_x \geq 0$. In the above, $R = \beta^{-1} G_Z \bar{\pi}$ denotes the interest rate in the non-stochastic steady state, Y_t denotes aggregate output, and x_t^g denotes the output gap. The variable $\epsilon_{R,t}$ is an exogenous disturbance to the policy rule satisfying $\epsilon_{R,t} \sim iid N(0, \sigma_R^2)$.

The output gap is defined as a weighted average of the deviations of utilization and labor from their non-stochastic steady state values: $x_t^g = \alpha \log u_t + (1 - \alpha) \log(\frac{N_t}{N})$, where u_t is the aggregate level of capital utilization, whose non-stochastic steady state value is equal to one, and N denotes the non-stochastic steady state level of labor. This concept of the output gap departs from Smets and Wouters (2007), who use the natural level of output (i.e., the level of output in the absence of price and wage rigidities). Use of such a concept would considerably increase the size of the nonlinear state space needed to solve the model so our use of x_t^g in part reflects that it is computationally convenient. In addition, in practice a central bank may rely on estimates of an output gap concept more closely related to x_t^g than the natural level of output, and we find that when we use this concept of the output gap, it tracks reasonably well the measure constructed by the Congressional Budget Office (see Appendix E).

Government spending is determined exogenously as a time-varying fraction of output: $G_t = (1 - \frac{1}{g_t})Y_t$, where g_t evolves according to:

$$\ln(g_t) = (1 - \rho_g) \ln g + \rho_g \ln(g_{t-1}) + \epsilon_{g,t}, \quad (12)$$

where g denotes the steady state share of government spending and $\epsilon_{g,t} \sim iid N(0, \sigma_g^2)$. Finally, the government budget constraint is assumed to be satisfied on a period-by-period basis so that $G_t = T_t$.

2.4 Market Clearing

We focus on a symmetric equilibrium in which all intermediate goods producing firms and all households make the same decisions, and we denote aggregate variables by dropping a variable's dependence on i or j . In a symmetric equilibrium, the aggregate production function is given by:

$$Y_t = K_t^\alpha [Z_t N_t]^{1-\alpha}, \quad (13)$$

since $K_t = K_t(j)$, $N_t = N_t(j) \forall j$. The market clearing condition for the final good is:

$$C_t + I_t + G_t + \frac{\varphi_p}{2} \left[\frac{\pi_t}{\bar{\pi}_{t-1}} - 1 \right]^2 Y_t + a(u_t) \bar{K}_t = Y_t, \quad (14)$$

where we have used that $u_t = u_t(i)$, $C_t = C_t(i)$, and $I_t = I_t(i)$ for all i in a symmetric equilibrium.

3 Solution and Econometric Inference

3.1 Model Solution

The model is solved for a minimum state variable solution using a projection method similar to Christiano and Fisher (2000) and Judd, Maliar, Maliar, and Valero (2014), and the details of the method are provided in a Technical Appendix.³ Because we are estimating the model, it might be tempting to use a computationally-efficient solution algorithm that respects the nonlinearity in the Taylor rule but log-linearizes the remaining equilibrium conditions. This approach has been used by Bodenstein, Guerrieri, and Gust (2013) and Guerrieri and Iacoviello (2015) among others. Unfortunately, we find

³ The minimum state variable solution rules out equilibria with additional state variables such as a sunspot or those with nonstationary dynamics. See McCallum (1999) for a discussion. However, Benhabib, Schmitt-Grohe, and Uribe (2001) emphasize that there is a second steady state that is deflationary due to the imposition of the ZLB on the Taylor rule. Our solution algorithm in principle does not rule out dynamics that fluctuate around this deflationary steady state. But, in practice, we use the linearized solution as the initial guess for the solution algorithm and find that that the solution algorithm never converged to this unintended deflationary equilibrium: for every estimated parameter draw that we examined from the posterior distribution, the ergodic mean of inflation was greater than zero. Moreover, such deflationary outcomes do not characterize the U.S. data over our sample period. As discussed in Christiano and Eichenbaum (2012), this deflationary equilibrium is also not E-learnable.

that this approach performs poorly when applied to our model, as the linearized dynamics poorly approximate the nonlinear dynamics even for moderately-sized shocks. Instead, as discussed in the Technical Appendix, we develop a computationally-efficient algorithm that is easy to parallelize and ideal for solving a nonlinear, medium-scale DSGE model with an occasionally binding constraint. Using this solution algorithm, we show that there are differences between the nonlinear and linear versions of the model and that they have important implications for the estimates of the model’s parameters and the propagation of shocks. Another attractive feature of our solution method is that it takes into account the uncertainty about the likelihood that the economy will be at the lower bound, as it does not rely on perfect foresight or certainty equivalence.

3.2 Data and Estimation

The solution algorithm characterizes a transition equation for the model variables, s_t , as a function, Φ , of its past realization, s_{t-1} , and current innovations to the shocks, ϵ_t , which we write as

$$s_t = \Phi(s_{t-1}, \epsilon_t; \theta), \quad \epsilon_t \sim N(0, I), \quad (15)$$

where

$$\begin{aligned} s_t &= [\bar{k}_{t+1}, y_t, c_t, i_t, \pi_t, R_t^N, w_t, \eta_t, \mu_t, g_t, y_{t-1}, c_{t-1}, i_{t-1}], \\ \epsilon_t &= [\epsilon_{\eta,t}, \epsilon_{\mu,t}, \epsilon_{Z,t}, \epsilon_{g,t}, \epsilon_{R,t}], \text{ and} \\ \theta &= [\beta, \bar{\Pi}, g_z, \alpha, \rho_R, \gamma_{\Pi}, \gamma_g, \gamma_x, \gamma, \sigma_L, \sigma_a, \varphi_I, \varphi_p, \varphi_w, a, a_w, \rho_g, \rho_{\mu}, \sigma_{\eta}, \sigma_{\mu}, \sigma_Z, \sigma_g, \sigma_R]. \end{aligned}$$

Within the model variables s_t , \bar{k}_{t+1} , y_t , c_t , w_t , and i_t denote the level of the capital stock, output, consumption, real wage, and investment scaled by the level of technology, Z_t .

Our goal is to determine plausible values for the parameters, θ , for tracking and explaining the U.S. macroeconomic experience since the early 1980s. Specifically, we base the empirical analysis on quarterly U.S. data from 1983:Q1 to 2014:Q1, using five time series to estimate the model. As measures of real activity, we use per-capita GDP growth, per-capita consumption growth, and per-capita investment growth. Inflation is measured as the change in the log of the GDP deflator, while our measure of the short-term interest rate is given by the quarterly average of the daily three-month

U.S. Treasury bill (quarterly) rate, following Ireland (2011). Additional details about the data are provided in Appendix A.

We assume that the observables contain measurement error, so the relationship to our model variables is given by,

$$\begin{bmatrix} \text{Output Growth} \\ \text{Consumption Growth} \\ \text{Investment Growth} \\ \text{Inflation} \\ \text{Interest Rate} \end{bmatrix} = \begin{bmatrix} \ln(y_t) - \ln(y_{t-1}) + \ln(G_Z) + \epsilon_{Z,t} \\ \ln(c_t) - \ln(c_{t-1}) + \ln(G_Z) + \epsilon_{Z,t} \\ \ln(i_t) - \ln(i_{t-1}) + \ln(G_Z) + \epsilon_{Z,t} \\ \ln(\pi_t) \\ \ln(R_t) \end{bmatrix} + \text{measurement error}, \quad (16)$$

which we write compactly as,

$$y_t = \Psi(s_t, \epsilon_t; \theta) + u_t, \quad u_t \sim N(0, \Sigma_u). \quad (17)$$

We estimate θ using Bayesian methods. Herbst and Schorfheide (2015) and the references therein provide background on Bayesian estimation of (linear and nonlinear) DSGE models; here we give a brief overview and leave the details to the Technical Appendix. Stacking the five observed time series as $Y_{1:T}$, Bayesian inference is characterized by a posterior distribution over the parameters, $p(\theta|Y_{1:T})$, that is proportional to the product of the likelihood for $Y_{1:T}$ and prior distribution for θ ,⁴

$$p(\theta|Y_{1:T}) \propto \underbrace{p(Y_{1:T}|\theta)}_{\text{likelihood}} \times \underbrace{p(\theta)}_{\text{prior}}.$$

Constructing the posterior distribution is conceptually straightforward but computationally very challenging. The transition equation defined by equation (15) along with the observation equation (17) define a nonlinear state space model. Since the observables are only a subset of the model variables, we are required to integrate out the unobserved states to compute the likelihood. Unlike a linear (Gaussian) state space model, we cannot evaluate these integrals analytically via the Kalman filter. Instead, we use a (Sequential) Monte Carlo technique known as a particle filter. This simulation-based approach provides us with an estimate of the likelihood, $\hat{p}(Y_{1:T}|\theta)$, for a given parameter draw, θ , by

⁴ In a slight abuse of notation, Y_t denotes aggregate output at date t and $Y_{1:T}$ denotes the five observed time series from 1983:Q1 to 2014:Q1.

constructing discrete approximations (“particles”) to the distributions of the underlying states.⁵

We also resort to simulation-based methods to elicit draws from the posterior. Markov-chain Monte Carlo (MCMC) techniques construct a Markov chain $\{\theta^i\}_{i=1}^N$ whose invariant distribution coincides with the posterior distribution of interest, $p(\theta|Y_{1:T})$. We use the so-called random walk Metropolis-Hastings (RWMH) algorithm first used in DSGE model estimation by Schorfheide (2000) and Otrok (2001). A complication here is that, like Fernandez-Villaverde and Rubio-Ramirez (2007), we are using MCMC with a particle-filter-based estimate of the likelihood. We appeal to Andrieu, Doucet, and Holenstein (2010), who show formally that, despite the use of the particle filter, the Markov chain will still have the desired posterior distribution as its invariant distribution. For more details on the MCMC algorithm, including hyperparameters and run time, see the Technical Appendix.

Measurement errors. As mentioned above, we include measurement error in the observation equation (16). One reason for its presence is feasibility: without measurement error—or with only very small measurement errors—our particle filter degenerates and the accuracy of the likelihood estimate becomes very poor.⁶ The measurement error is defined as

$$\Sigma_u = m_e \times \text{diag}(\mathbb{V}[Y_{1:T}]),$$

where we set $m_e = 0.25$ for the baseline parameter estimates of the model and $m_e = 0.1$ for a lower measurement error case. These values imply that, for each observable, the variance of its measurement error is 25 percent ($m_e = 0.25$) or 10 percent ($m_e = 0.1$) of its total variance over the estimation period, respectively. As discussed in the Technical Appendix, either of these values ensures that the particle filter estimates are stable, and we present results for two values of m_e to highlight how measurement error affects our estimates.

The inclusion of measurement error represents a departure from the majority—though not all—of the literature estimating linear DSGE models. In particular, it means that our estimates are more uncertain than they otherwise might be. The benefit of this approach, though, is that we are able to conduct inference using the full nonlinear structure of the model, including the zero lower bound. Given that the dynamics of the nonlinear model differ considerably from the linearized dynamics, we

⁵ There is a long literature on particle filtering—see Creal (2012) for a survey. We use a very basic version of the particle filter, the so-called bootstrap particle filter described in Gordon, Salmond, and Smith (1993) with some minor modifications. The Technical Appendix contains the details on the implementation of the particle filter.

⁶ Another important reason, as discussed later, is that measurement error can absorb some of the model misspecification.

think that it is a trade off worth making.

Parallelization. While most of the technical aspects are relegated to the Technical Appendix, it is worth mentioning some of the key innovations that we make in order to estimate a nonlinear, medium-scale DSGE model. The principal obstacle in our empirical exercise is time. It can take a long time to solve the model and evaluate the particle filter. Finally, the MCMC estimation strategy requires solving the model and evaluating the particle filter many times in an iterative fashion. For both the solution of the model and evaluation of the particle filter, we rely on extensive parallelization to reduce computational time. The parallelization of the projection method used to construct the model solution is straightforward. The parallelization of the particle filter, on the other hand, is more difficult, as the filter requires frequent communication amongst all particles. The Technical Appendix shows how to circumvent many of the problems associated with naive parallelization. With these modifications to the solution algorithm and particle filter, we leverage a large distributed computing environment to sample from our posterior quickly.⁷

4 Estimation Results

4.1 Parameter Estimates

Bayesian inference requires that we specify a prior distribution. Many of our marginal prior distributions are the same as in Justiniano, Primiceri, and Tambalotti (2011). The prior distributions for the Rotemberg adjustment cost parameters are specified to be consistent with prior beliefs that nominal wages are more rigid than prices. More details about the prior distributions are shown in Appendix B. We also fix several parameters. In particular, the elasticity of the demand for labor inputs, ϵ_w , and the elasticity of the demand for intermediate goods, ϵ_p , are not estimated and are fixed at 6. The depreciation rate, δ , is set to 0.025, and the non-stochastic steady state share of government spending, g , is set equal to 20 percent. The preference parameter ψ_L is normalized to one.

Estimating the nonlinear model is a challenging task, particularly for a model of this size, and this task is made more complicated because of the zero lower bound constraint, whose presence makes it difficult to solve the model for highly-persistent risk premium shocks. To simplify the estimation of the

⁷We could, in principal, enjoy further gains by modifying the MCMC algorithm along the lines of Smith (2011), who develops a Surrogate MCMC algorithm wherein the particle filter is not evaluated when the likelihood is predicted to be a low value. See Gust, López-Salido, and Smith (2012) for details.

model, we fix ρ_η at 0.85, a relatively high value that helps the model fit the data. In particular, this value is consistent with the micro-econometric evidence presented in Gilchrist and Zakrajšek (2012), whose measure of the excess bond premium has an autocorrelation of 0.84 over our sample period. Moreover, a value of $\rho_\eta = 0.85$ is not at odds with the aggregate data: as shown in Appendix B, the likelihood function is increasing in ρ_η ; however, as ρ_η gets close to one, it becomes more difficult to compute a solution to the model for a wide range of parameter configurations, considerably complicating the estimation of the model.⁸

Table 1 presents the means and the intervals bracketed by the 5th and 95th percentiles of the marginal posterior distributions of the parameters for the baseline value of the measurement error ($m_e = 0.25$). The mean estimates of the discount factor and the deterministic trend to the technology growth rate are just below one and 0.5 percent on a quarterly basis, respectively, implying a real rate of about 2.6 percent on an annualized basis in the non-stochastic steady state. The central bank’s inflation target is estimated to have a mean of about 2.5 percent, which is higher than the current 2 percent target of the Federal Reserve, reflecting that the sample mean of inflation from 1983-2014 is close to 2.5 percent.

For the policy rule, the estimates are largely in line with Del Negro, Eggertsson, Ferrero, and Kiyotaki (2011). The posterior mean for the interest-rate smoothing parameter, ρ_R , is near 0.7, while the posterior means for the policy-rate responses to inflation and output growth, γ_π and γ_g , are about 1.7 and 0.7, respectively. The policy rule also involves a smaller but statistically significant response of the short-term nominal rate to the output gap, as the posterior mean of γ_x equals 0.14.

The posterior mean of the adjustment cost parameter for investment, φ_I , is 3.7 with a 90 percent credible interval covering values from 2.2 to 5.2. This mean estimate is above the value of 3 estimated by Justiniano, Primiceri, and Tambalotti (2011), but below that of Smets and Wouters (2007). The posterior mean of the habit persistence parameter, γ , is equal to 0.7, and the 90 percent credible interval—covering values between 0.63 to 0.76—is relatively tight. The parameter governing the elasticity of labor supply, σ_L , has a mean of 2. The posterior mean of the elasticity of the rental cost of capital with respect to the utilization, σ_a , is 5.3 – similar to the estimates of Justiniano, Primiceri, and Tambalotti (2011) and Altig, Christiano, Eichenbaum, and Lindé (2011).

⁸ Whether it is possible to compute a solution for ρ_η close to one or not depends importantly on the other model parameters and in particular the adjustment costs for investment. Characterizing the parameter space for which we can find a solution to the model is an extremely difficult task given that the model is nonlinear and has many parameters.

The posterior mean of the price adjustment cost parameter, φ_p , is around 100 with a weight on lagged indexation, $1 - a$, equal to 0.56, implying a linearized slope coefficient for the New Keynesian Phillips curve of 0.07. As shown in Appendix F, such a slope coefficient corresponds to a frequency of price changes of slightly more than one year in a model with Calvo contracts. Nominal wages are estimated to be extremely rigid but, as shown in Appendix F, less so than in Justiniano, Primiceri, and Tambalotti (2011).

While the discussion so far has emphasized that most of the estimates are similar to those in the literature, this is not the case for the shock parameters. Most strikingly, the mean of the standard deviation of the innovation to this shock, σ_μ , is about half the value found by Justiniano, Primiceri, and Tambalotti (2011). This lower estimate primarily reflects our use of the nonlinear model rather than its linear approximation.⁹ To highlight this difference, the right panel of Figure 1 shows that the posterior distribution of σ_μ is considerably below the distribution derived from estimating the linearized version of the model. While the estimate of this parameter is the one that is most substantially affected by estimating the nonlinear model rather than the linear model, there are other notable differences in the estimates as well. The estimated volatilities of the risk premium shock, government spending shock, and technology shocks are all smaller using the nonlinear model, and the parameter estimates of the policy rule differ somewhat as well.¹⁰ Moreover, the left panel shows that the estimate of the adjustment cost on investment is higher in the nonlinear model, though the difference is not as dramatic as the estimate of the innovation variance of the MEI shock.

In Appendix D we show the parameter estimates imposing a lower measurement error ($m_e = 0.1$). Most of the parameter estimates are similar to the baseline case. However, there are a couple of important exceptions. The degree of wage and price rigidities is notably smaller than under the baseline estimates, and the estimated size of the technology shock is much larger.

4.2 The Nonlinear Propagation of Shocks

If the linear model approximated the nonlinear dynamics well, the two sets of parameter estimates would be the same. However, Figure 1 demonstrates that some of the parameters from these two

⁹ The differences in parameter estimates between the linear and nonlinear versions of the model would be larger if the measurement error were smaller. All other things equal, the estimated structural shocks would be larger in that case, causing the performance of the linear approximation to deteriorate further relative to the nonlinear solution.

¹⁰ See Appendix C for the estimates based on the linear version of our model.

versions of the model can differ substantially, and in this subsection we show that this result reflects that the linear approximation understates the endogenous propagation of shocks in the model. To do so, we use the mean estimates derived from the nonlinear model to simulate both the linear and nonlinear versions of the model and compare the impulse responses to the economy's two financial shocks.

In order to highlight the difference between the linear and nonlinear solutions, Figure 2 shows the effects of a one standard deviation increase in $\epsilon_{\eta,t}$, the innovation to the risk premium shock, using initial conditions in which the economy is moderately away from its non-stochastic steady state.¹¹ The solid dark lines display the responses using the nonlinear solution, and dashed red lines show the responses using the linear solution. By raising risky spreads, this shock reduces economic activity and inflation and also puts downward pressure on the nominal interest rate in both model versions. But, the fall in economic activity is notably larger in the nonlinear version of the model, especially for investment, where the difference is substantial. Because this larger fall in investment persists and because it also reduces the stock of capital, there is a larger and more persistent fall in output, consumption, and hours worked in the nonlinear version of the model. In contrast, the price of installed capital falls by less in the nonlinear version of the model than in the linear version. Overall, Figure 2 indicates that the linear solution poorly approximates the stronger propagation inherent in the nonlinear dynamics even for moderately-sized innovations in the shock.¹² Of course, these differences are magnified for larger shocks.

The stronger propagation inherent in the nonlinear dynamics is not driven by the ZLB but by the investment adjustment costs, which imply that optimal investment satisfies:

$$1 + q_t \mu_t \left[S_t + \varphi_I \left(\frac{I_t}{I_{t-1}} - 1 \right) \frac{I_t}{I_{t-1}} \right] = q_t \mu_t + \varphi_I E_t \left\{ q_{t+1} \mu_{t+1} m_{t+1} \left(\frac{I_{t+1}}{I_t} - 1 \right) \left(\frac{I_{t+1}}{I_t} \right)^2 \right\}, \quad (18)$$

where S_t is defined in equation (9), $m_{t+1} = \beta \frac{\Lambda_{t+1}}{\Lambda_t}$ denotes the stochastic discount factor, q_t denotes the price of newly installed capital, and for convenience we have abstracted from the trend in technological

¹¹ In particular, we choose initial conditions for the risk premium shock (η_0) so that it is one standard deviation above its unconditional mean while the lagged values of the other state variables are equal to their non-stochastic steady state values.

¹² The bottom right panel compares the approximation error associated with the Euler equation that determines the relationship between the price of investment and a household's optimal supply of investment, equation (18) along the shocked path. This error is about ten times larger using the linear solution than the nonlinear solution.

growth. The left-hand side of equation (18) reflects the cost of increasing investment by an extra unit and the right-hand side reflects the benefit. On the cost side, increasing investment in the current period means that the household foregoes consumption today and also incurs costs to transform the final good into new capital. On the benefit side, the extra investment boosts the amount of installed capital and the value to doing so reflects both the price of the installed capital, q_t , and the efficiency of the extra investment, μ_t . The benefits also reflect that raising investment today reduces the burden associated with incurring the adjustment costs in the future.

Although the quadratic nature of the adjustment costs make it difficult to see, equation (18) implies the supply of investment is an increasing function of q_t . All else equal, a reduction in the stochastic discount factor, either due to an increase in the real policy rate or the risky spread, tends to lower the supply of investment because the future benefits fall. This effect, however, is absent from the linearized counterpart of equation (18), because the linearized equation is approximated around the non-stochastic steady state where this effect is negligible. However, away from there, this effect will not necessarily be inconsequential, as Figure 2 demonstrates. In that case, the increase in the risky rate leads to a reduction in investment supply that contributes to a slump in investment. Furthermore, this shift in investment supply does not have to be large, because the elasticity of investment demand with respect to q_t is high meaning that relatively small changes in investment supply can induce relatively large movements in the level of investment.¹³ The linearized dynamics neglect the shift in investment supply, and hence tend to understate the drop in investment.

The linear dynamics also poorly capture the propagation of the MEI shock, whose estimated volatility is much smaller once we account for the stronger nonlinear propagation. Figure 3 shows the effects of a one standard deviation reduction in the innovation to the marginal efficiency of investment, $\epsilon_{\mu,t}$. This shock also produces a much larger fall in investment in the nonlinear version of the model than in the linear version. This lower investment in turn induces a larger fall in the stock of capital, which helps engender more persistent declines in output and hours worked in the nonlinear version of the model.

What are the implications of the stronger propagation of the nonlinear dynamics for the business cycle? As shown in Table 2, the nonlinear version of the model relies on its stronger internal propagation

¹³ The investment-demand relationship can be derived from the Euler equation for the capital stock which relates q_t to the rental rate of capital. The rental rate of capital, in turn, can be expressed to depend on investment using the capital accumulation equation. See the Technical Appendix for the details of this equation.

mechanism to generate output volatility in line with the observed data. To see this, the third column of the table, labeled ‘Nonlinear Constrained,’ shows the volatility of output, consumption, investment, and hours at business cycle frequencies simulating the nonlinear model using its parameter estimates. The standard deviations of output and consumption are near their empirical counterparts, while the volatilities of investment and hours are only somewhat lower than their empirical counterparts. If instead we used these same parameter estimates and simulated the linearized dynamics of the model, the linearized version of the model would substantially understate the volatility of investment, as shown in the final column. This reflects that, if the linearized model has to rely on the parameter estimates coming from the nonlinear model, the estimate of the volatility of the MEI shock is too small given its weaker endogenous propagation. Thus, it fails to generate enough volatility in output, hours, and especially in investment. For the linearized model to generate a realistic volatility of investment, it must rely on a much larger exogenous MEI shock, because it poorly approximates the nonlinear model’s dynamics. Hence, to reliably disentangle the endogenous and exogenous sources of business cycle fluctuations, one should estimate the nonlinear model rather than its linear approximation.

4.3 Time Series Properties of the Model

Figure 4 shows the implications of the estimated parameters by comparing the observed data on output growth, investment growth, consumption growth, inflation, and the nominal interest rate with the smoothed estimates produced by the model under the baseline value of the measurement error ($m_e = 0.25$). The figure also shows the 68 percent credible interval around these values.¹⁴

The model generally tracks the fluctuations in output, consumption, and investment growth, and generates contractions in these variables in all three of the recessions included in our sample period. For the Great Recession, the model accounts for the sharp falls in output and investment that occurred and modestly understates the fall in consumption. In addition, the model mimics well the slow growth in the variables that characterized the subsequent recovery. The model captures most of the low and medium frequency variation in quarterly inflation, displayed in the middle-right panel, though some of the high-frequency movements remain unexplained. For the Great Recession, the model accounts for a substantial part of the fall in inflation as well as its rebound.

¹⁴ To construct these series, we proceed as follows: For each draw $\theta \sim p(\theta|Y_{1:T})$, we sample from $p(S_{1:T}|Y_{1:T}, \theta)$ using the particle filter and report $h(S_{1:T})$ for the appropriate function $h(\cdot)$. The credible intervals include both uncertainty surrounding θ as well as uncertainty surrounding $S_{1:T}$.

The lower-left panel shows that the smoothed values for the nominal interest rate are close to the observed values and the lower-right panel displays the central bank's notional or desired path for the nominal interest rate, which provides a measure of the severity of the lower-bound constraint on actual monetary policy. From 2009 onwards, the notional interest rate is well below zero. It falls to about minus 5.5 percent in the first half of 2009, gradually moves up to about minus 1 percent at the end of 2012, and then hovers around that level through 2014:Q4.

Figure 5 displays the (smoothed) estimates of output growth and inflation for the model with baseline measurement error ($m_e = 0.25$) and compares them to the corresponding estimates from the model version with lower measurement error ($m_e = 0.1$). Although the two versions of the model produce similar estimates of output (as well as consumption and investment), the version with lower measurement error does a better job fitting the high frequency movements in inflation. This result is due to the more prominent role of technology shocks and the lower degree of price stickiness that produce larger and more transitory movements in inflation when $m_e = 0.1$.

In Figure 6, we examine the estimated model's implications for the probability of hitting the lower bound and the duration of a lower bound spell.¹⁵ The top panel displays the distribution of the probability of the nominal rate being at the zero lower bound by simulating data from the model using parameter draws from the posterior distribution. On average, the model implies that there is about a 4 percent probability of the nominal rate being at the ZLB. In comparison, in our sample from 1983:Q1-2014:Q1, the nominal rate was at the zero lower bound in the last 21 quarters of the 125 observations, yielding a probability near 17 percent. Although this is considerably larger than the mean estimate derived from the model, the top panel of Figure 6 also shows that estimates of this probability are disperse and that the distribution has a long right tail: about 15 percent of the draws have a probability of being at the ZLB more than 10 percent of the time and about 3 percent of the draws have a probability of being there more than 17 percent of the time.

The bottom panel of Figure 6 displays the distribution of the duration of lower bound episodes. The average duration for a lower bound spell is just over 3.5 quarters and the median is two. The distribution for the duration of a spell is skewed to the right and also has a long right tail. The ZLB spell at the end of our sample lasted twenty-one quarters, and spells with durations of twenty-one quarters or longer occur account for about $\frac{1}{2}$ percent of the total number of lower bounds spells. Accordingly, our

¹⁵ For each draw of parameters, we simulate 1,000 datasets each with 125 observations. We used 1,000 draws of parameters from the posterior distribution.

model implies that the long ZLB spell in the United States was very unlikely, and from the perspective of households and firms in the model, would have been difficult to predict ex ante. The result is in line with the view that the financial crisis and its effects on the macroeconomy were difficult to foresee.

5 What Accounts for the Great Recession?

Before quantifying how much the zero lower bound contributed to the Great Recession, we examine the estimated path of the shocks and their relative contributions to the Great Recession.

5.1 The Path of the Estimated Shocks

Figure 7 displays the smoothed estimates of the risk premium and the marginal efficiency of investment shocks.¹⁶ The vertical axes in these panels display the magnitudes of the risk premium and MEI shocks relative to their standard deviations, and thus, for example, a value of two denotes that the size of the shock is two standard deviations above its mean. The bottom left panel of Figure 7 shows the innovations ($\epsilon_{\eta,t}$) to the risk premium shock (η_t) and the bottom right panel displays the equity premium implied by the model.

At the beginning of 2008, the risk premium shock was already elevated and then there was a four standard deviation innovation in this shock when the financial crisis began. Although such an innovation is extremely unlikely, the financial crisis itself was an unprecedented event that was followed by the largest contraction in U.S. economic growth since the Great Depression.¹⁷ The model also requires a sequence of positive innovations in the risk premium shock from 2012:Q1 to 2013:Q2; however, these innovations are small, less than one standard deviation away from their zero mean.

Because the risk premium shock plays an important role in our analysis, it is interesting to compare it to movements in empirical measures of credit spreads during the financial crisis. The top left panel of Figure 7 also compares the estimated risk premium shock to the interest rate differential between BAA

¹⁶ The government spending and monetary policy shocks contributed less, according to the model. Smoothed estimates of the paths of these shocks are available in Appendix E. The fact that the government spending shock does not play a large role in the Great Recession does not necessarily mean that fiscal policy was unimportant during the episode. The government spending shock in our analysis is a reduced-form shock affecting the economy's resource constraint and reflects other factors besides changes in government spending such as movements in international trade that we do not have in the model.

¹⁷ To be consistent with the literature, we choose to estimate the model using normally-distributed shocks. However, with our methodology, it would be straightforward to allow for distributions with fatter tails or time-varying volatility and formally test different assumptions regarding the shock distributions. However, we viewed such an exercise as beyond the scope of this paper.

corporate bonds and 10-year Treasury bonds and a measure of credit spreads constructed by Gilchrist and Zakrajšek (2012), who abstract from the effect of the expected loss stemming from default on spreads and attempt to isolate the credit risk premium. Since the units of the estimated risk premium shock are not comparable to those of the other two measures of financial stress, it is more useful to focus on how movements in the risk premium shock correlate with these measures, both of which rose sharply in 2008 and declined thereafter. Although the measure constructed by Gilchrist and Zakrajšek (2012) falls relatively soon after the financial crisis, the BAA spread remains persistently elevated even after falling in 2009. Similar to the two empirical measures, the model’s estimate of the risk premium shock also rose sharply at the end of 2008 and fell in 2009. Even after falling in 2009, like the BAA spread, the risk premium shock remains persistently high. In sum, even though we do not use information on asset prices in our empirical approach, the model’s estimated risk premium displays broadly similar behavior as empirical measures of financial stress did at the time.

As can be seen by the upper right panel, low realizations of the MEI shock also contributed to weak activity at the time, as the MEI shock is more than one standard deviation below its mean in 2008 and in the first half of 2009. Thus, the model’s two financial shocks are consistent with the view that strains in financial markets triggered the economic slump. In subsequent periods, both shocks begin to move back toward their mean values, but the risk premium shock remains at a high level – 2 standard deviations above its mean – through the end of 2014. Moreover, the MEI shock’s recovery is followed by another fall in 2010 after which it recovers and falls again at the end of 2013.

The bottom right panel displays the model’s equity or finance premium, which we measure as the expected excess return on capital:

$$\frac{E_t R_{t+1}^k}{R_t} = \frac{\tilde{r}_{t+1}^k + q_{t+1}(1 - \delta)}{R_t q_t},$$

where $\tilde{r}_t^k = r_t^k u_t - a(u_t)$ is the user or rental cost of capital adjusted for utilization. The panel also shows an empirical measure of the equity premium from Adrian, Crump, and Moench (2012). The model’s estimate of the equity premium peaks at a level of about 11 percent at the end of 2008 and drops in 2009 but remains elevated, averaging about 6 percent through 2013. These movements are broadly consistent with the estimates of Adrian, Crump, and Moench (2012), though our model implies a larger jump in the equity premium at the end of 2008 and a less elevated equity premium overall.

The top panel of Figure 8 shows the difference in the level of technology beginning in 2008:Q1 from the level if it grew at its deterministic rate, G_Z , for the baseline estimates. Technology or total factor productivity was below trend during the Great Recession and fell further below trend through the end of the sample (2014:Q1). The dashed and dotted lines show the estimates in Christiano, Eichenbaum, and Trabandt (2015) and Fernald (2014), respectively. Despite our different econometric strategy from these authors, our estimates of total factor productivity are broadly similar to theirs, as total factor productivity, according to these measures, is also estimated to be below trend with the exception of a temporary increase in the measure constructed by Fernald (2014) in 2009.

The large risk premium shock as well as reductions in productivity and the marginal efficiency of investment are all important in accounting for the large drop in activity that occurred during the Great Recession with only a moderate disinflation. The latter two shocks put upward pressure on inflation, offsetting the significant downward pressure on inflation stemming from the risk premium shock. These results are broadly consistent with Christiano, Eichenbaum, and Trabandt (2015) and Del Negro, Giannoni, and Schorfheide (2015b) albeit with some differences. In particular, these authors introduce model features that allow financial frictions to break the link between prices and future marginal costs in the NK Philips curve and help mitigate the disinflationary pressure stemming from the risk premium shock.¹⁸

Figure 8 also compares the smoothed estimates of the technology shocks from 2008-2013 for the baseline value of the measurement error to the case with low measurement error (bottom panel). For the other shocks, the estimated paths are similar across model versions; however, the technology shock is notably different. For the model with lower measurement error, the level of technology is estimated to have fallen substantially more during the Great Recession and its aftermath: the level of technology at the end of 2013 is about 7 percent below trend, compared to only 3.5 percent below trend under the baseline estimates. Moreover, with lower measurement error, the estimated path of technology shocks lies well below the estimates in Fernald (2014).

Figure 8 highlights another reason for the inclusion of measurement error in the model. Specifically, the model is misspecified to the extent that it omits factors such as oil and commodity price movements that help explain high frequency fluctuations in inflation. Because of such misspecification, the model version with lower measurement error, as suggested by the estimates in Fernald (2014), overstates the

¹⁸ For an alternative explanation emphasizing the role of fiscal policy in mitigating the disinflationary pressure, see Bianchi and Melosi (2016).

role of technology shocks in accounting for the Great Recession; hence, we chose the model version with higher measurement error as the baseline with the implication that there is greater uncertainty surrounding our estimates of the cost of the ZLB.¹⁹

5.2 The Contribution of the Estimated Shocks to the Great Recession

Figure 9 shows how much of the model’s fit is attributable to individual shocks for output growth, investment growth, consumption growth, inflation, and both the observed and the notional nominal interest rates in the baseline model. More specifically, it displays the model’s dynamics if only one of the estimated shocks were present during the Great Recession and compares this path to the smoothed values that are generated using all five shocks (the solid line denoted as the baseline in the figure).²⁰ A clear picture emerges from the figure: the large contraction in output, consumption and investment growth is mostly explained by the risk premium shock. This shock also generates a somewhat larger fall in inflation than under the baseline path. However, as shown in the middle right panel, these disinflationary effects are offset by the upward pressure on inflation induced by the fall in the level of technology.

The risk premium shock is also largely responsible for driving the nominal rate to the zero lower bound. By reducing both output growth and inflation, this shock pushes the notional interest rate well below zero (to about -5.5 percent), causing the interest rate to hit and then stay at the lower bound. A decline in the MEI shock in 2011 helped put downward pressure on the nominal interest rate in 2011 and 2012 and contributed significantly to the sharp decline in investment in late 2008.

While the technology shock plays an important role in moderating the decline in inflation during the ZLB episode, the model estimates imply that monetary policy shocks were not important during the ZLB episode.²¹ However, this does not necessarily imply that monetary policy is unimportant, because the systematic part of the rule and the lower bound constraint are crucial determinants of the

¹⁹ For a related approach in addressing model misspecification, see Canova (2014), who proposes a flexible, non-structural link between a DSGE model and the data. It would be preferable, of course, to incorporate features into the structural model that allow it to capture the high-frequency movements in inflation and we leave this task to future research.

²⁰ We compute the counterfactual with only one shock present as follows: Starting in 2007:Q4, we take the estimated state as an initial value, and simulate the economy forward feeding in the smoothed values of one of the shock processes, assuming that the variances of the other shocks are set to zero.

²¹ Monetary policy shocks are still identified at the lower bound, because these shocks can persistently lower the notional interest rate. As a result, agents anticipate that the actual nominal rate will be at the lower bound for longer, an expectation that affects current outcomes.

model's dynamics.

Appendix D presents the shock decomposition for the model version with lower measurement error. In that version, the risk premium shocks still play the predominant role in generating the large contraction in output but now technology shocks contribute noticeably to the declines in output and consumption growth at the end of 2008. Moreover, there is more upward pressure on inflation emanating from the technology shock in this version of the model.

6 How Costly Was the Zero Lower Bound?

6.1 The Contribution of the Lower Bound to the Great Recession

To determine the role of the zero lower bound constraint during the Great Recession, we compare the estimated model outcomes (in which the ZLB constraint is imposed) to the outcomes in a hypothetical scenario in which monetary policy is free to adjust the nominal interest rate in an unconstrained manner.²² Figure 10 compares the levels of output and the price level in these two scenarios under the posterior mean estimates of the baseline model. The left panel shows that in 2009:Q2 output was about 6 percent below its level in 2007:Q4. If monetary policy could have cut the nominal interest rate more aggressively, it would have helped to offset the contractionary effects of the shocks and output would have fallen by only 4 percent; thus, the zero lower bound accounted for about 30 percent of the sharp drop in output that occurred in 2009.²³

While the ZLB contributed significantly to the sharp fall in economic activity in 2009, it was an even more important factor in holding back activity during the subsequent recovery. Figure 10 shows that output did not recover back to its 2007:Q4, pre-recessionary level until 2012:Q4, five years later.²⁴ During those five years, the average level of output was 2.4 percent below its level in 2007:Q4. In comparison, average output over those years would have only been about 1.1 percent lower in the unconstrained scenario, as output would have been above its pre-recessionary level about a year earlier than in the estimated, constrained scenario. Thus, these estimates imply that the presence of the ZLB

²² The hypothetical scenario uses the same estimated initial conditions and time series for the shocks as the estimated model in which the constraint is imposed. These values are then used to simulate the model ignoring the ZLB constraint on the estimated rule.

²³ This is the mean estimate, and the 68 percent interval around it is wide, covering values between 6 to 46 percent.

²⁴ The baseline path does not exactly match the observed output series because of measurement error and recovers two quarters earlier than observed output.

accounted for a little over 50 percent of the lower output over the 2008-2012 period.

The right panel of Figure 10 shows that the ZLB constraint also contributed to a lower price level. Inflation averaged 1.45 percent per year from 2008:Q1 until the end of 2012, below the Federal Reserve's inflation target. Absent the zero lower bound constraint, the mean estimates indicate that inflation would have averaged around 1.75 percent per year, and hence was 0.3 percentage points lower, on average, over the 2008-2012 period. Overall, our results suggest that the interest-rate lower bound was a significant constraint on monetary policy that exacerbated the recession, inhibited the recovery, and contributed to inflation outcomes below the Federal Reserve's inflation target.

Figure 11 displays the difference in outcomes from the estimated, constrained scenario for output, investment, consumption, and the notional rate. For each variable, the solid line shows the point estimate of the effect of ZLB constraint and the shaded region displays the 68 percent credible region. As emphasized above, output is about 1.2 percent lower, on average, over the 2008-2012 period because of the constraint. Consumption would have been about 1 percent higher in absence of the constraint, and investment would have been 4 percent higher. Figure 11 also highlights an important caveat to interpreting these results: the estimates are subject to considerable uncertainty as the 68 percent credible region does not exclude the possibility that the estimated effects of the ZLB constraint were much smaller or much larger.

Figure 11 also shows the point estimates of the effect of the ZLB constraint in the model version with lower measurement error. The ZLB constraint accounts for about 15 percent of the sharp drop in output that occurred in 2009 and about 20 percent of the lower level of output over the 2009-2012 period. These smaller effects are driven by the more prominent role of technology shocks in accounting for the fall in output over this period. In particular, there is less of a role for monetary policy to stabilize the fall in output if more of its decline is driven by technology shocks. Hence, the zero lower bound is a less important constraint on monetary policy in this version of the model; however, as discussed earlier, there is evidence that this model version may overstate the role of technology shocks during this period and thus may understate the cost of the ZLB constraint.

6.2 The Contribution of the Lower Bound to the 2003-2004 Deflationary Scare

Our methodology can be used to examine episodes in which the realized nominal rate did not hit the lower bound but economic behavior was affected by the prospect that policy could become constrained.

One such episode in which uncertainty about the zero lower bound may have affected economic behavior was in 2003 when inflation was low and the Federal Open Market Committee (FOMC) was concerned about the possibility of further disinflation. This view was summarized by Alan Greenspan, the Chair of the FOMC at the time:

“.. [W]e face new challenges in maintaining price stability, specifically to prevent inflation from falling too low...[T]here is an especially pernicious, albeit remote, scenario in which inflation turns negative...engendering a corrosive deflationary spiral...” (Greenspan (2003))

Although this deflationary scenario never materialized and policy rates never reached the lower bound, this risk was an important consideration in 2003, and we use the estimated model to investigate the effect of the ZLB during this episode. In the model, this constraint has effects on inflation and economic activity, even if it never binds, because agents in the model take into account the entire probability distribution of future outcomes, including those in which the nominal interest rate obtains the zero lower bound, in making their current decisions. Accordingly, we follow the same approach as we used for the Great Recession and compare outcomes in the estimated, constrained model to the hypothetical scenario in which policy is unconstrained.

The left panel in Figure 12 shows the mean trajectories of output with and without the ZLB constraint, and the right panel shows the uncertainty faced by the agents about the future likelihood of being at the constraint. It shows that at the end of 2002, taking into the economy’s current state and the estimated policy rule, private sector agents believed there was a 12 percent chance that the nominal interest rate would fall to its lower bound during 2003. Because scenarios in which the ZLB binds also imply higher future real policy rates than if monetary policy could act in an unconstrained fashion, these scenarios are also characterized by downward shifts in the distributions of outcomes in aggregate spending and output; accordingly, the estimated mean outcome for output is lower in the constrained case than the unconstrained case. Figure 12 shows that the effects of this uncertainty on the estimated mean outcomes are relatively small: the estimated level of output is about 0.2 percent lower throughout 2003 and 2004 than in the unconstrained case. Still, the effect is a persistent one, and the results highlight that uncertainty about the course of monetary policy did have a tangible economic impact during this episode, an effect that is omitted *a priori* using the standard approach to estimate DSGE models.

6.3 Forward Guidance and the Estimated Policy Rule

In both the 2003 episode and during the Great Recession, forward guidance about the policy rate played a prominent role in FOMC communications. In late 2008, forward-guidance language was reintroduced into the FOMC statement, which indicated that “weak economic conditions are likely to warrant exceptionally low levels of the federal funds rate for some time.” This guidance grew more forceful and specific over time, and in late 2012, the FOMC made this forward-guidance more explicitly state contingent, linking the maintenance of the low funds rate to projected inflation and the level of the unemployment rate. We capture forward guidance through the state-contingent path of rates implied by our estimated interest-rate rule. The presence of the lagged notional rate, in particular, is an important element of a strategy that seeks to maintain future rates at a low level.²⁵ In particular, when the notional rate is negative, a higher coefficient on the lagged notional rate in the rule, all else equal, increases private sector beliefs that future policy rates will remain at the zero lower bound. Moreover, the further the lagged notional rate falls below zero, the more likely the future policy rate will remain at zero.

To investigate the role of the lagged notional rate in the policy rule, Figure 13 shows a counterfactual simulation of how the economy would have evolved during the Great Recession had policy be governed by an alternative rule that responds to the lagged actual policy rate rather than the lagged notional rate. Because the lagged notional rate fell sharply below zero during the recession and remained negative during the recovery, the estimated rule generates a policy rate path that is more accommodative than the alternative rule. As shown in the bottom left panel, the mean estimates for the path of the real interest rate are lower under the estimated rule than the rule that depends on the lagged actual rate. Accordingly, the fall in the mean estimates of output and prices during the Great Recession are smaller for the estimated rule than the alternative rule.

The estimated rule is more accommodative than the alternative rule, because the inclusion of the lagged notional rate generates the belief by private agents that the nominal rate will remain at the zero lower bound longer and under a wider set of circumstances. To highlight this, the bottom right panel shows the uncertainty faced by the agents in the model about whether the policy rate one-quarter

²⁵ In light of the evolving nature of the forward-guidance language in FOMC communications throughout the ZLB spell, it remains an open question as to how best to model it. An alternative way of modeling forward guidance is to allow for anticipated shocks in the monetary policy rule. See, for instance, Del Negro, Giannoni, and Patterson (2015a). Also, Keen, Richter, and Throckmorton (2015) study the effects of forward guidance modeled in this way in the context of a simple, nonlinear New Keynesian model.

ahead will be at the zero lower bound or not. In 2010:Q1, for example, under the estimated rule, private sector agents believed the probability that the nominal rate would ‘liftoff’ next quarter was very close to zero. In contrast, in the counterfactual using the alternative rule with lagged actual rate, agents in 2010:Q1 would have believed there was about at 15 percent chance that the policy rate would rise in 2010:Q2.

Under both of these rules the nominal interest rate is estimated to remain at the zero lower bound through the end of our sample, but the probability of a departure from the ZLB is notably higher under the alternative rule than the estimated rule. In 2014:Q1, for instance, the agents would have believed that there was about a 40 percent chance of liftoff next quarter under the alternative rule compared to only a 10 percent chance under the estimated rule. Thus, the simulation demonstrates that the inclusion of the lagged notional rate captures the forward-guidance language used by the FOMC at the time by validating private sector beliefs that policy is more likely to remain at the zero lower bound, thereby stimulating the economy.

7 Conclusions

In this paper, we estimated a nonlinear DSGE model in which the interest-rate lower bound is occasionally binding. This allowed us to quantify the size and nature of the disturbances that caused the Great Recession as well as pushed the nominal rate to the lower bound in late 2008 and kept it there throughout the economy’s long slump. Overall, our results suggest that the interest-rate lower bound was a significant constraint on monetary policy that exacerbated the recession, inhibited the recovery, and contributed to inflation outcomes below the Federal Reserve’s inflation target.

Our results also highlight the importance of estimating the nonlinear model instead of the linear approximation to it. Because the nonlinear model that we estimate has a stronger internal propagation mechanism than its linearized counterpart, we find that there are significant differences in the parameter estimates. In particular, estimates based on the linearized version overstate the importance of the exogenous shocks, and thus the estimation of the nonlinear model is essential in properly quantifying the relative contributions of endogenous and exogenous sources of business cycle fluctuations.

References

- ADRIAN, T., R. K. CRUMP, AND E. MOENCH (2012): “Regression-Based Estimation of Dynamic Asset Pricing Models,” Federal Reserve Bank of New York, Staff Report No. 493.
- ALTIG, D., L. CHRISTIANO, M. EICHENBAUM, AND J. LINDÉ (2011): “Firm-Specific Capital, Nominal Rigidities and the Business Cycle,” *Review of Economic Dynamics*, 14, 225–247.
- ANDRIEU, C., A. DOUCET, AND R. HOLENSTEIN (2010): “Particle Markov Chain Monte Carlo with Discussion,” *Journal of Royal Statistical Society Series B*, 72, 269–342.
- ARULAMPALAM, S., S. MASKELL, N. GORDON, AND T. CLAPP (2002): “A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking,” *IEEE Transactions on Signal Processing*, 50, 174–188.
- ARUOBA, B., P. CUBA-BORDA, AND F. SCHORFHEIDE (2016): “Macroeconomic Dynamics Near the ZLB: A Tale of Two Countries,” University of Maryland, manuscript.
- BENHABIB, J., S. SCHMITT-GROHE, AND M. URIBE (2001): “Monetary Policy and Multiple Equilibria,” *American Economic Review*, 91, 167–186.
- BIANCHI, F. AND L. MELOSI (2016): “Escaping the Great Recession,” Federal Reserve Bank of New York, Staff Report No. 493.
- BODENSTEIN, M., L. GUERRIERI, AND C. J. GUST (2013): “Oil Shocks and the Zero Bound on Nominal Interest Rates,” *Journal of International Money and Finance*, 32, 941–967.
- CAMPBELL, J. R., J. D. FISHER, A. JUSTINIANO, AND L. MELOSI (2016): “Forward Guidance and Macroeconomic Outcomes Since the Financial Crisis,” Federal Reserve Bank of Chicago Working Paper 2016-07.
- CANOVA, F. (2014): “Bridging DSGE Models and the Raw Data,” *Journal of Monetary Economics*, 67, 1–15.
- CAPPÉ, O., S. J. GODSILL, AND E. MOULINES (2007): “An Overview of Existing Methods and Recent Advances in Sequential Monte Carlo,” *Proceedings of the IEEE*, 95, 899–924.
- CHRISTIANO, L. AND M. EICHENBAUM (2012): “Notes on Linear Approximations, Equilibrium Multiplicity and E-learnability in the Analysis of the Zero Lower Bound,” Northwestern University, manuscript.
- CHRISTIANO, L., M. EICHENBAUM, AND M. TRABANDT (2015): “Understanding the Great Recession,” *American Economic Journal: Macroeconomics*, 7, 110–167.
- CHRISTIANO, L. AND J. FISHER (2000): “Algorithms for Solving Dynamic Models with Occasionally Binding Constraints,” *Journal of Economic Dynamics and Control*, 24, 1179–1232.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113, 1–45.
- CHUNG, H., J.-P. LAFORTE, D. REIFSCHNEIDER, AND J. WILLIAMS (2012): “Have We Underestimated the Likelihood and Severity of Zero Lower Bound Events?” *Journal of Money, Credit and Banking*, 44, 47–82.
- CREAL, D. (2012): “A Survey of Sequential Monte Carlo Methods for Economics and Finance,” *Econometric Reviews*, 31, 245–296.

- DEL NEGRO, M., G. EGGERTSSON, A. FERRERO, AND N. KIYOTAKI (2011): “The Great Escape? A Quantitative Evaluation of the Fed’s Liquidity Facilities,” Federal Reserve Bank of New York Staff Report 520.
- DEL NEGRO, M., M. P. GIANNONI, AND C. PATTERSON (2015a): “The Forward Guidance Puzzle,” FRB of New York Staff Report 574.
- DEL NEGRO, M., M. P. GIANNONI, AND F. SCHORFHEIDE (2015b): “Inflation in the Great Recession and New Keynesian Models,” *American Economic Journal: Macroeconomics*, 7, 186–196.
- DOUCET, A. AND A. M. JOHANSEN (2011): “A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later,” in *Handbook of Nonlinear Filtering*, ed. by D. Crisan and B. Rozovsky, Oxford University Press.
- FERNALD, J. (2014): “Productivity and Potential Output Before, During, and After the Great Recession,” NBER Macroeconomic Annual 2014.
- FERNANDEZ-VILLAYERDE, J. AND J. RUBIO-RAMIREZ (2007): “Estimating Macroeconomic Models: A Likelihood Approach,” *Review of Economic Studies*, 74, 1059–1087.
- FISHER, J. D. (2015): “On the Structural Interpretation of the Smets-Wouters “Risk Premium” Shock,” *Journal of Money, Credit and Banking*, 47, 511 – 517.
- GILCHRIST, S. AND E. ZAKRAJŠEK (2012): “Credit Spreads and Business Cycle Fluctuations,” *American Economic Review*, 102, 1692–1720.
- GORDON, N. J., D. J. SALMOND, AND A. F. SMITH (1993): “Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation,” in *IEE Proceedings F, Radar and Signal Processing*, IET, vol. 140, 107–113.
- GREENSPAN, A. (2003): “Testimony Before the Committee on Financial Services, U.S. House of Representatives,” Available at <https://www.federalreserve.gov/boarddocs/hh/2003/july/testimony.htm>.
- GUERRIERI, L. AND M. IACOVIELLO (2015): “OccBin: A Toolkit for Solving Dynamic Models with Occasionally Binding Constraints Easily,” *Journal of Monetary Economics*, 70, 22–38.
- GUST, C., D. LÓPEZ-SALIDO, AND M. E. SMITH (2012): “The Empirical Implications of the Interest-Rate Lower Bound,” Finance and Economics Discussion Papers, Federal Reserve Board, 2012 - 83.
- HERBST, E. AND F. SCHORFHEIDE (2015): *Bayesian Estimation of DSGE Models*, Princeton: Princeton University Press.
- IRELAND, P. N. (2011): “A New Keynesian Perspective on the Great Recession,” *Journal of Money, Credit, and Banking*, 43, 31–54.
- JUDD, K. L., L. MALIAR, AND S. MALIAR (2011): “Numerically Stable and Accurate Stochastic Simulation Approaches for Solving Dynamic Economic Models,” *Quantitative Economics*, 2, 173–210.
- JUDD, K. L., L. MALIAR, S. MALIAR, AND R. VALERO (2014): “Smolyak Method for Solving Dynamic Economic Models: Lagrange Interpolation, Anisotropic Grid and Adaptive Domain,” *Journal of Economic Dynamics and Control*, 44, 92–123.

- JUSTINIANO, A., G. PRIMICERI, AND A. TAMBALOTTI (2011): “Investment Shocks and the Relative Price of Investment,” *Review of Economic Dynamics*, 14, 101–121.
- KEEN, B. D., A. W. RICHTER, AND N. A. THROCKMORTON (2015): “Forward Guidance and the State of the Economy,” University of Oklahoma, manuscript.
- MCCALLUM, B. T. (1999): “Role of the Minimal State Variable Criterion in Rational Expectations Models,” *International Tax and Public Finance*, 6, 621–639.
- OTROK, C. (2001): “On Measuring the Welfare Costs of Business Cycles,” *Journal of Monetary Economics*, 47, 61–92.
- PITT, M. K., R. D. S. SILVA, P. GIORDANI, AND R. KOHN (2012): “On Some Properties of Markov Chain Monte Carlo Simulation Methods Based on the Particle Filter,” *Journal of Econometrics*, 171, 134–151.
- ROTEMBERG, J. (1982): “Sticky Prices in the United States,” *Journal of Political Economy*, 90, 1187–1211.
- SCHORFHEIDE, F. (2000): “Loss Function-Based Evaluation of DSGE Models,” *Journal of Applied Econometrics*, 15, 645–670.
- SMETS, F. AND R. WOUTERS (2007): “Shocks and Frictions in the U.S. Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 97, 586–606.
- SMITH, M. (2011): “Estimating Nonlinear Economic Models Using Surrogate Transitions,” Federal Reserve Board, manuscript.
- WU, J. C. AND F. D. XIA (2016): “Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound,” *Journal of Money, Credit and Banking*, 48, 253–291.

Table 1: Posterior Distribution of Parameters

Parameter	Mean	[05, 95]	Parameter	Mean	[05, 95]
Steady State					
$100(\beta^{-1} - 1)$	0.14	[0.06, 0.23]	$100(\bar{\Pi} - 1)$	0.61	[0.54, 0.68]
$100 \log(G_z)$	0.50	[0.46, 0.54]	α	0.19	[0.16, 0.22]
Policy Rule					
ρ_R	0.70	[0.59, 0.78]	γ_{Π}	1.67	[1.21, 2.14]
γ_g	0.73	[0.39, 1.07]	γ_x	0.14	[0.07, 0.24]
Endogenous Propagation					
γ	0.70	[0.63, 0.76]	σ_L	2.00	[1.01, 3.17]
σ_a	5.32	[3.78, 7.09]	φ_I	3.70	[2.24, 5.21]
φ_p	100.41	[65.10, 136.88]	$1 - a$	0.56	[0.36, 0.76]
φ_w	4420.49	[1693.15, 8356.34]	$1 - a_w$	0.51	[0.29, 0.72]
Exogenous Processes					
ρ_G	0.67	[0.29, 0.96]	$100\sigma_G$	0.15	[0.11, 0.20]
ρ_{μ_I}	0.80	[0.64, 0.92]	$100\sigma_{\mu_I}$	2.39	[1.43, 3.70]
$100\sigma_{\eta}$	0.44	[0.34, 0.54]	$100\sigma_Z$	0.56	[0.38, 0.80]
$100\sigma_R$	0.18	[0.14, 0.24]			

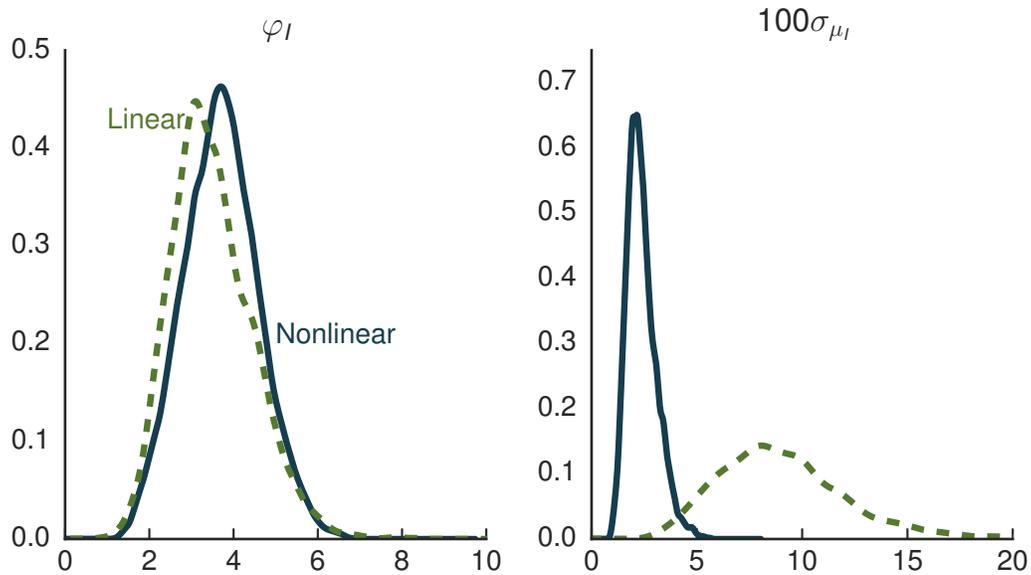
Notes. Table reports the mean, fifth, and ninety-fifth percentile of the posterior distribution estimated by pooling 4 MCMC chains with 50,000 draws each (including a 10,000 draw burn-in period.)

Table 2: Standard Deviations of Aggregate Variables at Business Cycle Frequencies

Variable	Data	Model Versions		
		Nonlinear Constrained	Nonlinear Unconstrained	Linear
Output	1.13	1.15	1.10	0.88
Consumption	0.78	0.77	0.75	0.81
Investment	4.96	4.01	3.92	1.59
Hours	1.80	1.34	1.29	1.04

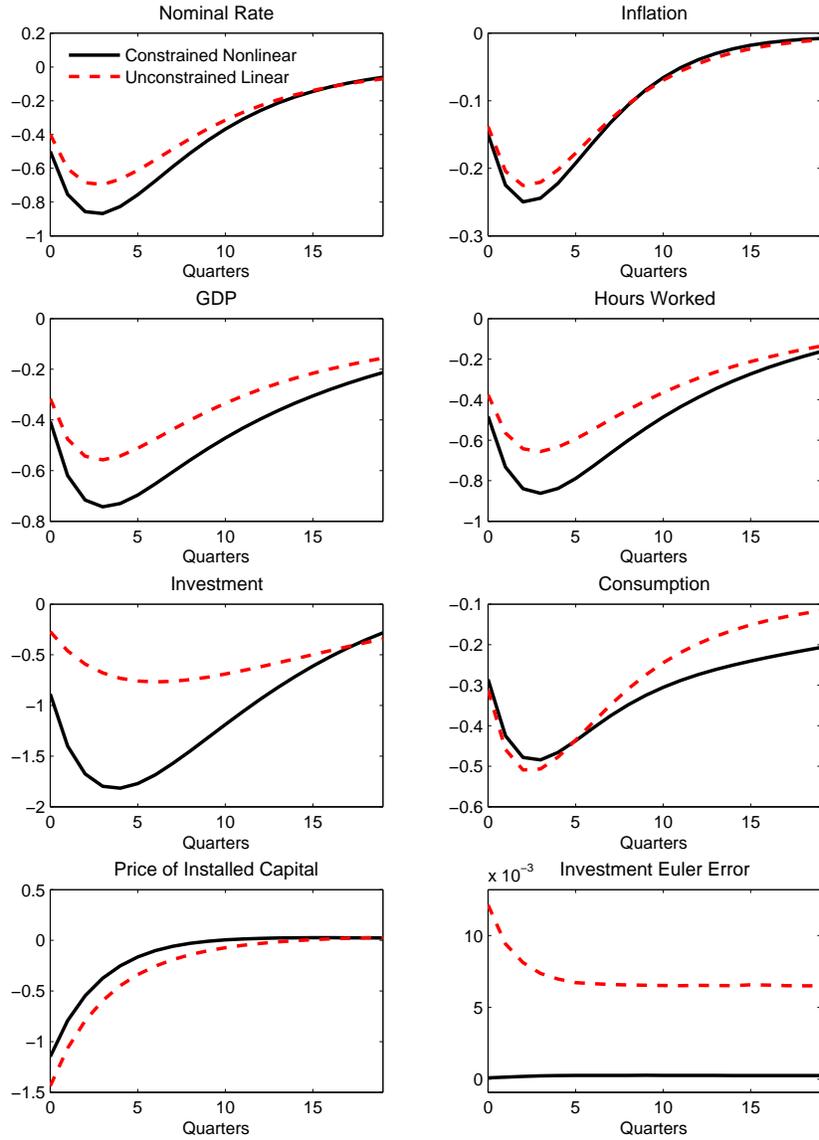
Notes. The model series are constructed from 1000 draws from the posterior distribution. For each parameter draw, a data series with 125 observations is simulated, and the HP-filter is applied to the series with a smoothing parameter of 1600. Simulations for all three model versions use the parameters estimated using the nonlinear, constrained version of the model.

Figure 1: A Comparison of the Posterior Estimates: Nonlinear versus Linear



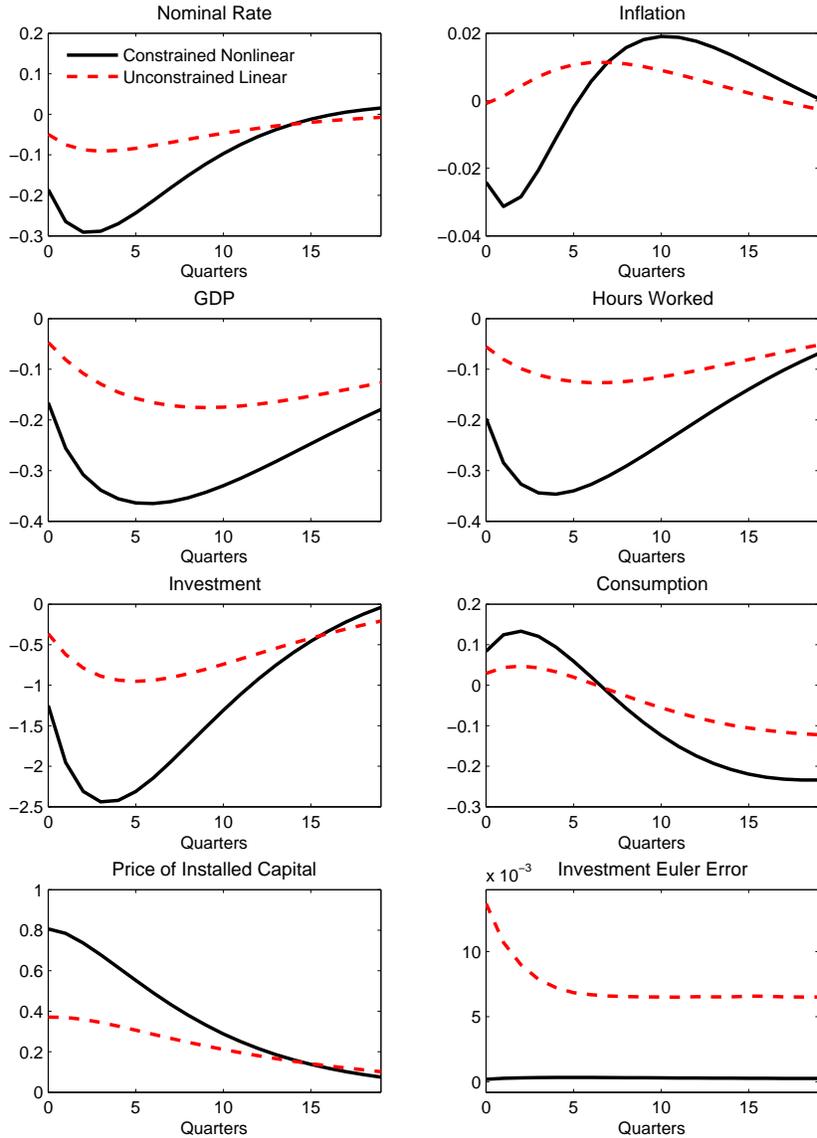
Notes. Figure shows the posterior density functions (kernel density estimate) for φ_I (left panel) and $100\sigma_{\mu}$ (right panel) under the nonlinear (solid lines) and linearized (dashed lines) versions of the model.

Figure 2: Response to an Exogenous Increase in the Risk Premium



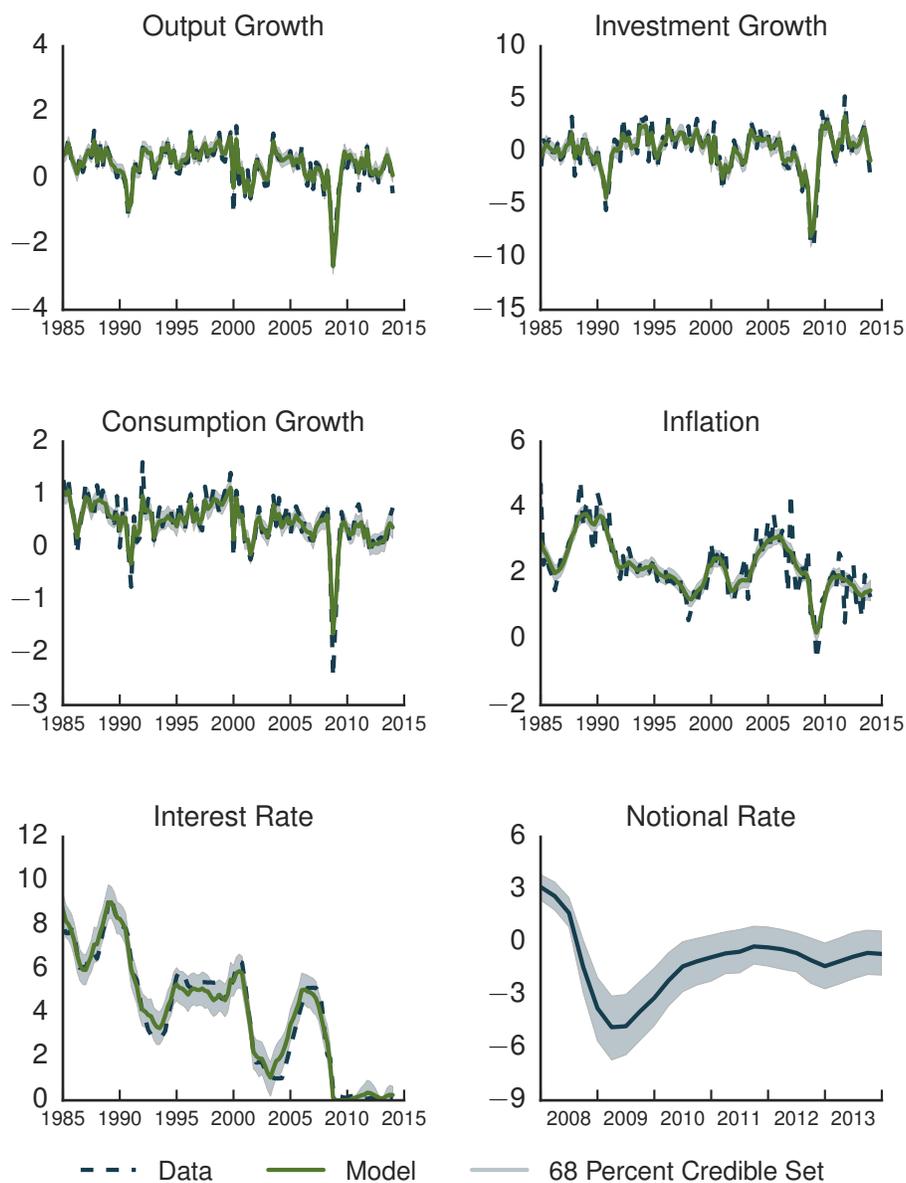
Notes. Figure reports the difference between the mean values of a shocked path, which introduces a one-standard deviation innovation in the risk premium shock at date 1, $\epsilon_{\eta,1}$, and the mean values of a path in which this innovation does not occur. Both paths are constructed using the posterior mean parameter values and initial conditions such that $\ln(\eta_0)$ is one standard deviation above its non-stochastic steady-state, and all other variables are at their non-stochastic steady states. The dashed line shows the responses for the unconstrained linear model, while the solid line shows the responses for the constrained nonlinear model. The lower right panel with the investment Euler error shows the mean absolute values of the errors along the shocked path at each date.

Figure 3: Response to a Fall in Investment Efficiency



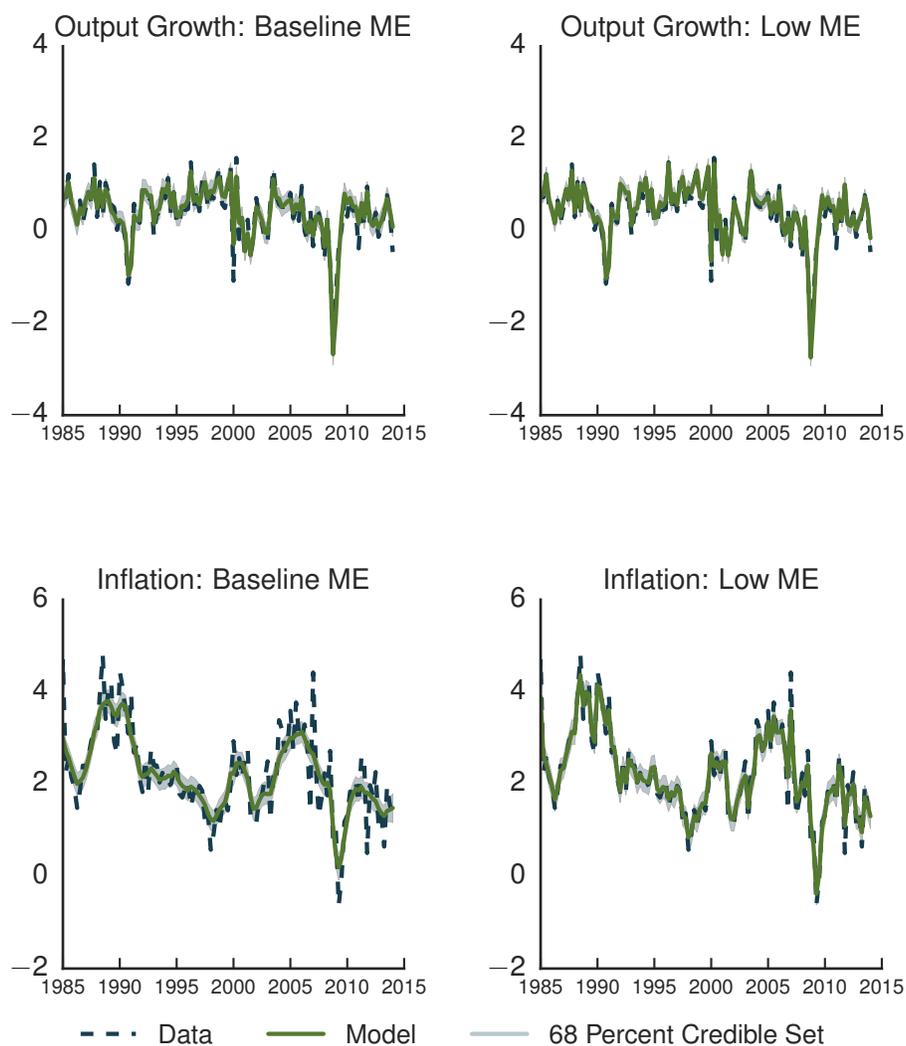
Notes. Figure reports the difference between the mean values of a shocked path, which introduces a one-standard deviation innovation in the MEI shock at date 1, $\epsilon_{\mu,1}$, and the mean values of a path in which this innovation does not occur. Both paths are constructed using the posterior mean parameter values and initial conditions such that $\ln(\mu_0)$ is one standard deviation above its nonstochastic steady-state, and all other variables are at their nonstochastic steady states. The dashed line shows the responses for the unconstrained linear model, while the solid line shows the responses for the constrained nonlinear model. The lower right panel with the investment Euler error shows the mean absolute values of the errors along the shocked path at each date.

Figure 4: Smoothed Estimates of Model Objects



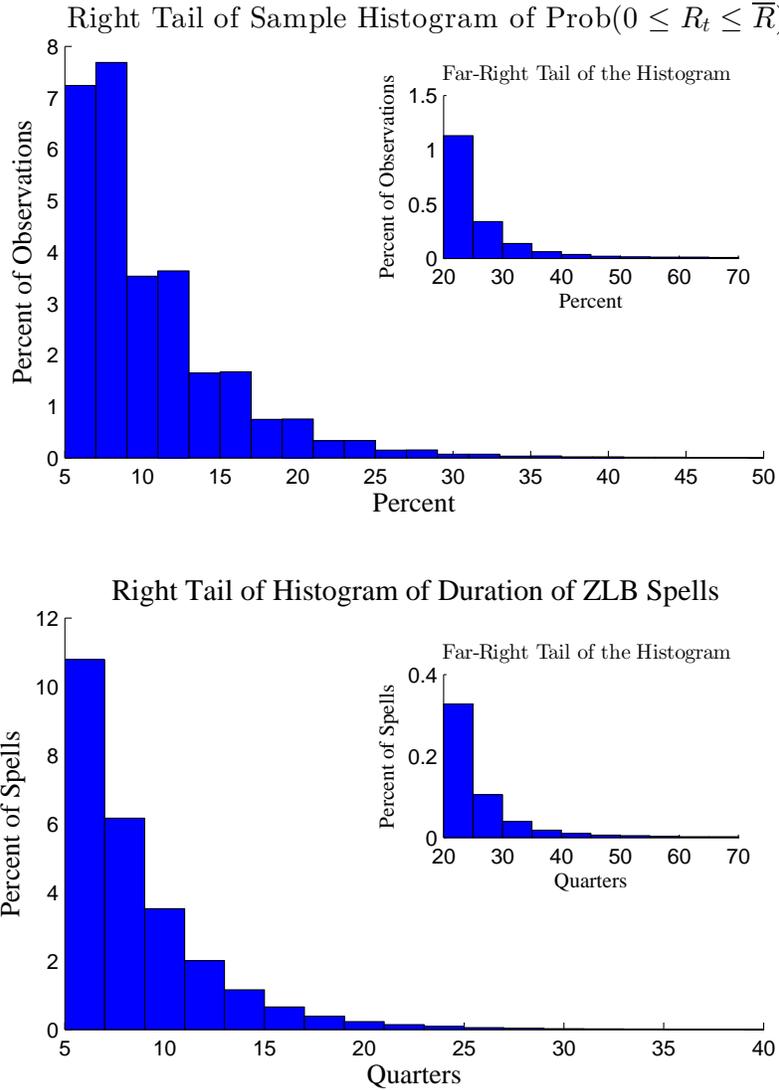
Notes. Figure shows the time series of the means (solid lines) and 68% bands (shaded regions) of the smoothed distributions of model variables, as well as their data counterparts over the estimation period. The bottom right panel shows estimates of the shadow (notional) rate from 2008:Q1 onwards.

Figure 5: Model Objects For Different Values of Measurement Error



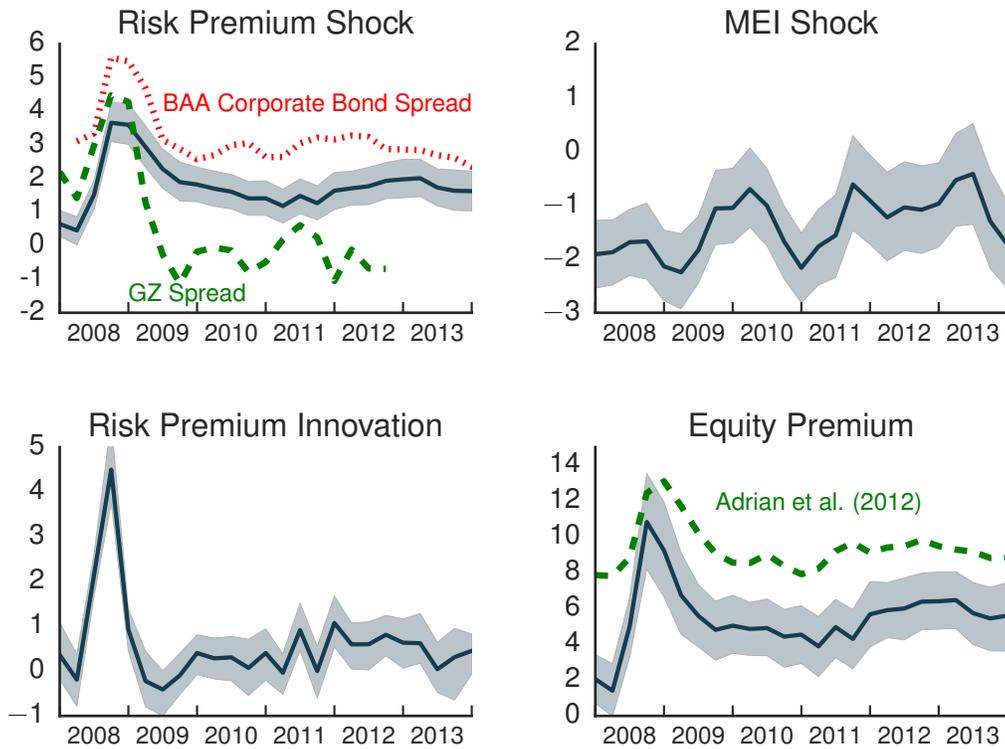
Notes. The four panels show the time series of the means (solid lines) and 68% bands (shaded regions) of the smoothed distributions of model variables, as well as their data counterparts over the estimation period. Baseline ME corresponds to the estimated model in which $m_e = 0.25$, and Low ME corresponds $m_e = 0.1$.

Figure 6: Distribution of the Probability and Duration of Being at the Zero Lower Bound



Notes. We draw 1000 times from our posterior distribution. For each draw, we simulate the economy for 1,000,000 time periods. For the top panel, we break the 1,000,000 time periods into bins of 125 time periods, compute the fraction of time the economy is at the ZLB, and report that distribution. For these simulations, the bottom panel computes the duration of ZLB spells using each of the 1,000,000 time periods.

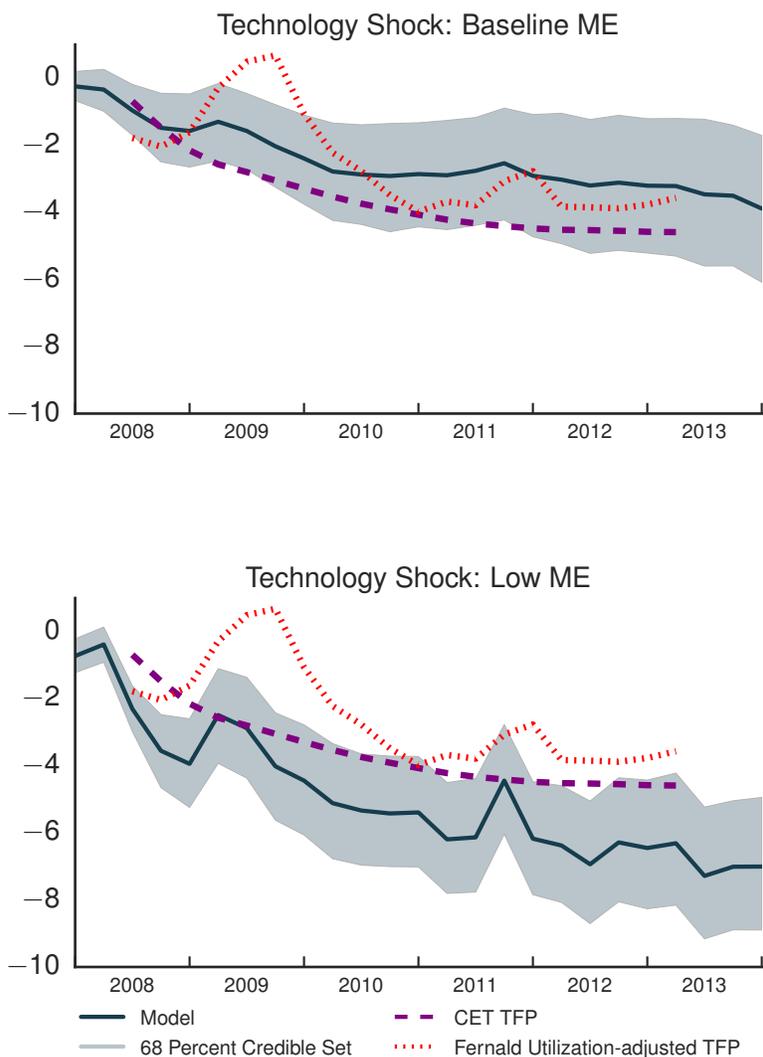
Figure 7: The Path of the Estimated Shocks and Equity Premium During the Great Recession



Notes. Figure shows the time series of the mean (solid line) and 68% bands (shaded region) of the smoothed distributions of $\ln(\eta_t)$ (top left panel), $\ln(E_t R_{t+1}^k / R_t)$, the model-implied equity premium (bottom right panel), $\epsilon_{\eta,t}$ (bottom left panel), and $\ln(\mu_t)$ (top right panel). Both $\ln(\eta_t)$ and $\ln(\mu_t)$ are normalized by their unconditional standard deviations.

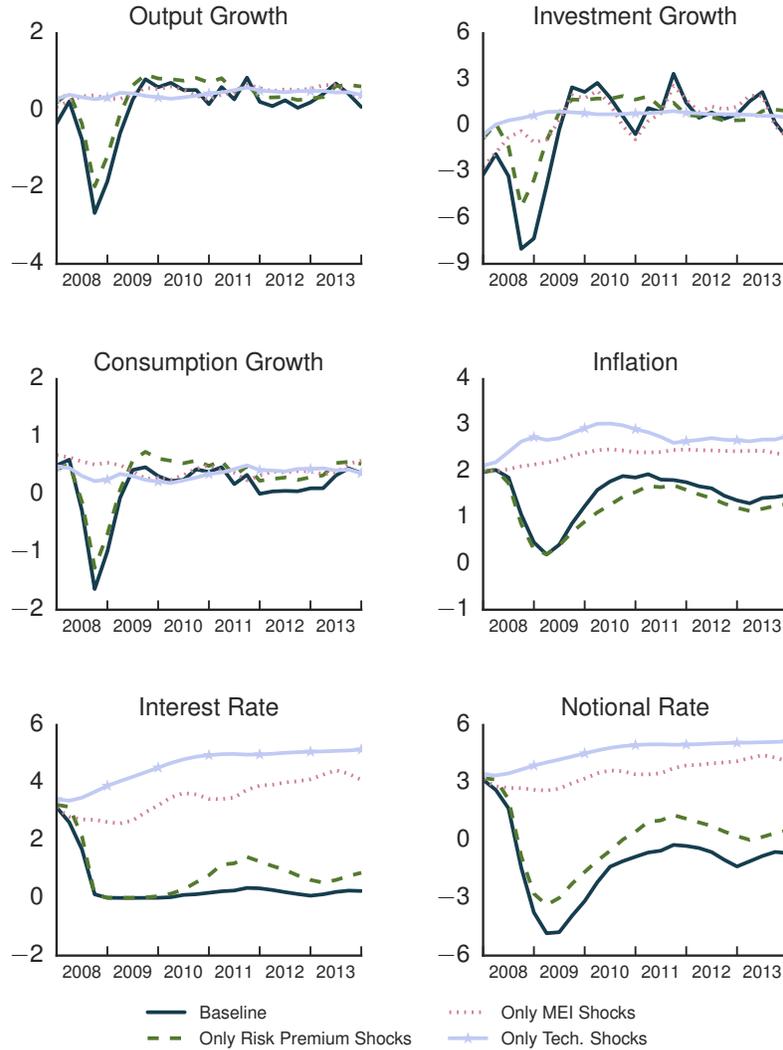
The top left panel also includes the standardized estimates of the credit spread from Gilchrist and Zakrajsek (2012) and BAA spread, while the bottom right panel shows an estimate of the equity premium, computed as in Adrian, Crump, and Moench (2012).

Figure 8: The Path of the Estimated Technology Shock



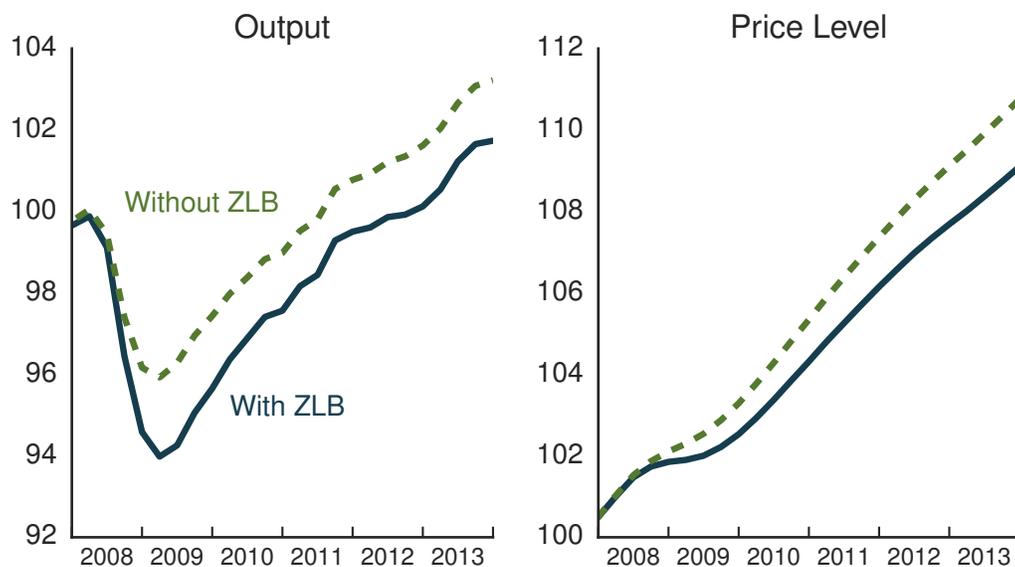
Notes. The two panels show the time series of the mean (solid line) and 68% bands (shaded regions) of the smoothed distribution of the difference between the smoothed estimate of $\ln(Z_t)$ and its forecasted trajectory in 2007:Q4 under alternative assumptions for the measurement error. The panels also show the utilization-adjusted estimate of TFP (red dotted line) of Fernald (2014) and the calibrated TFP measure (purple dashed line) used in Christiano, Eichenbaum, and Trabandt (2015), both taken from Christiano, Eichenbaum, and Trabandt (2015). Baseline ME corresponds to the estimated model in which $m_e = 0.25$, and Low ME corresponds $m_e = 0.1$.

Figure 9: The Contribution of the Estimated Shocks to the Great Recession



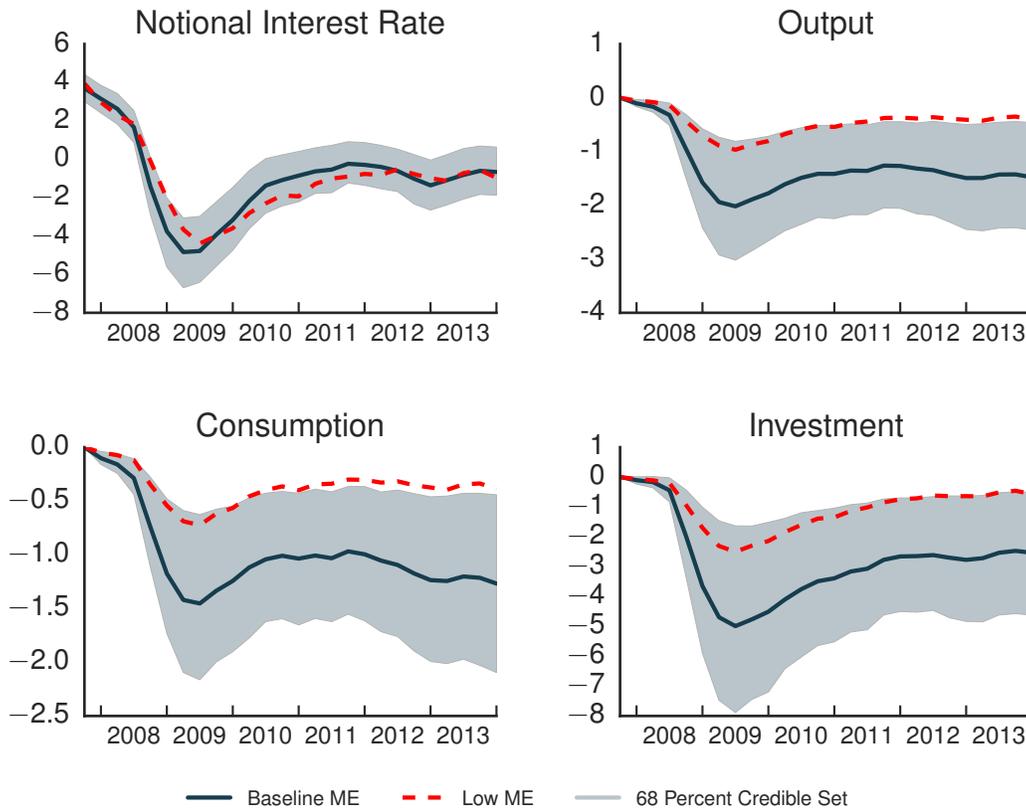
Notes. Figure shows counterfactual trajectories of output growth (upper left panel), investment growth (upper right panel), consumption growth (middle left panel), inflation (middle right panel), the interest rate (bottom left panel), and the notional rate (bottom right panel). The trajectories are computed using smoothed estimates of the states in 2007:Q4 as initial conditions and the smoothed shock estimates from 2008:Q1 to 2014:Q1 for only liquidity shocks (dashed lines), only MEI shocks (dotted line), only technology shocks (line with triangles), and all of the structural shocks (solid line).

Figure 10: The Contribution of the Zero Lower Bound to the Great Recession



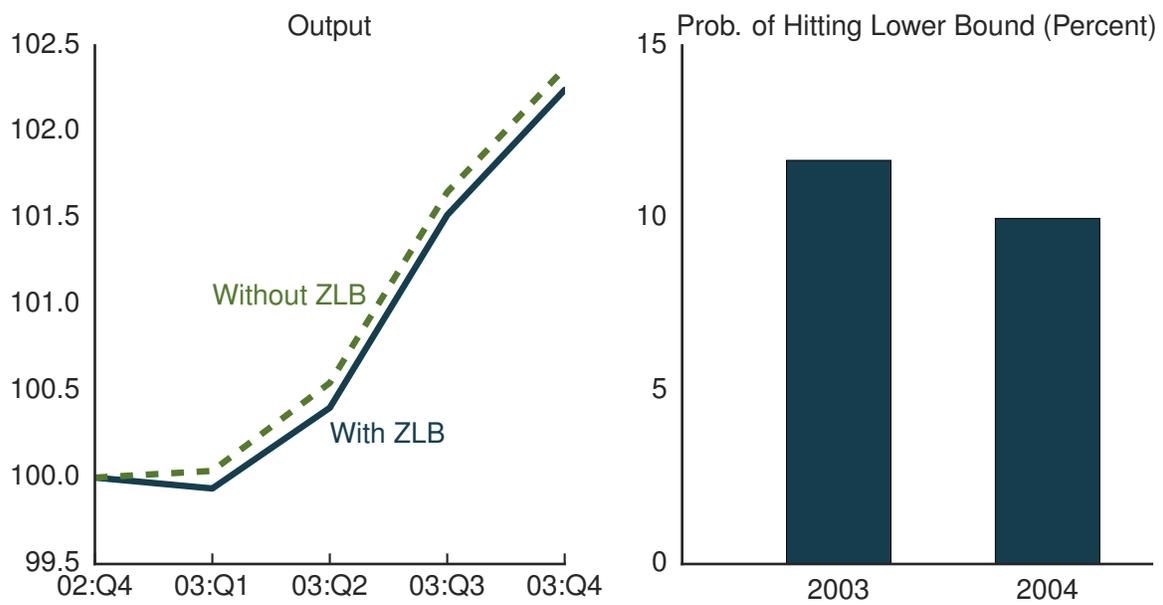
Notes. Figure shows the posterior mean trajectories of the levels of output (left panel) and prices (right panel). The solid lines show the effects imposing the ZLB constraint and the dashed lines show the effects without this constraint. The simulations use the smoothed estimates of the states in 2007:Q4 as initial conditions and the smoothed shock estimates from 2008:Q1 to 2014:Q1 to construct the paths. Each series is normalized to 100 in 2007:Q4.

Figure 11: The Effect of the Lower Bound on Aggregate Spending



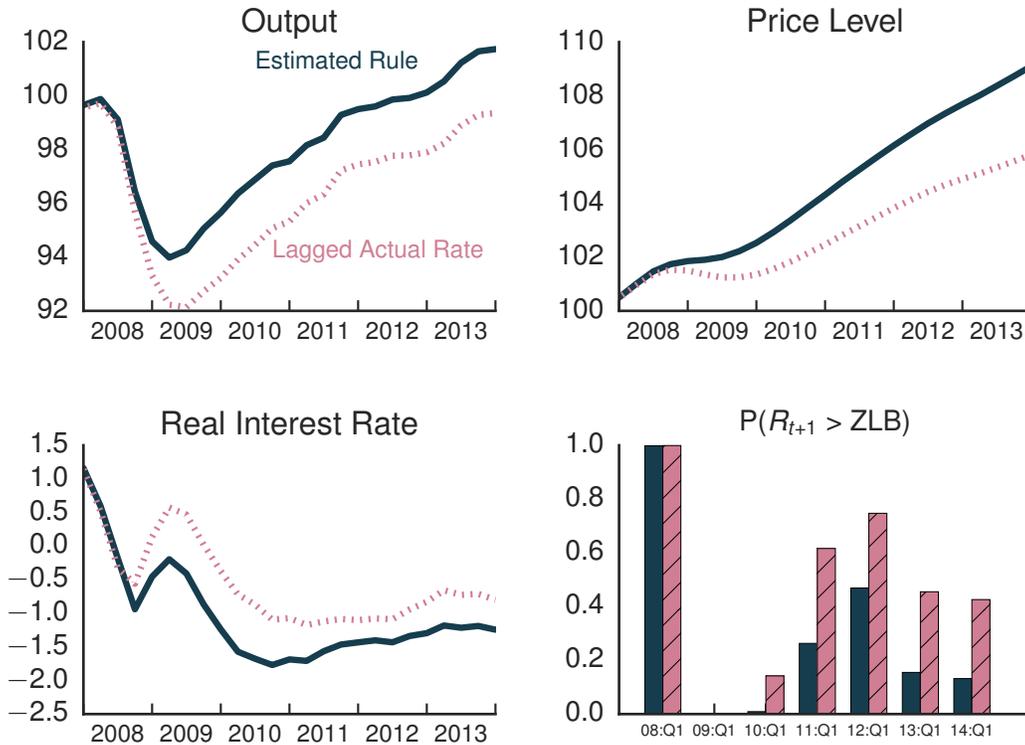
Notes. Figure shows the time series of the mean (solid line) and 68% bands (shaded region) of the distribution of differences for of the level of output (top right panel), consumption (bottom left panel), and investment (bottom right panel) simulated with and without the zero lower bound respectively. The top left panel shows the mean estimated path of the notional rate along with 68% bands. The dashed red lines correspond to the mean estimates of the effect of the ZLB constraint in the model with lower measurement error ($m_e = 0.1$).

Figure 12: The 2003–2004 Episode



Notes. Figure shows the posterior mean trajectories of the levels of output (left panel). The solid line shows the effects imposing the ZLB constraint and the dashed line shows the effects without this constraint. The simulations use the smoothed estimates of the states in 2002:Q4 as initial conditions and the smoothed shock estimates from 2003:Q1 to 2004:Q4 to construct the paths. The level of output is normalized to 100 in 2002:Q4. The right panel shows the probabilities, conditional on the posterior mean parameter estimates, of hitting the ZLB in 2003 and 2004, given 2002:Q4 initial condition (posterior mean of smoothed estimates).

Figure 13: The Effect of Inertia in the Policy Rule



Notes. Figure shows the posterior mean trajectory of the level of output (top left panel), prices (top right panel), and the real interest rate (bottom left panel) under the estimated monetary policy rule (solid lines) and the rule with the lagged actual rate (dotted lines). The simulation uses the smoothed estimates of the states in 2007:Q4 as initial conditions and the smoothed shock estimates from 2008:Q1 to 2014:Q1 to construct the paths. The level of output and prices are normalized to 100 in 2007:Q4.

The bottom right panel shows the probability at each point in time—according to the agents in the model—that the interest rate next quarter will be away from the ZLB. In computing the probabilities, we use the mean parameter estimates and the mean estimates of the smoothed states.

Appendix

A Data

Output growth is measured by quarter-to-quarter changes in the log of real GDP (chained 2005 dollars, seasonally adjusted, converted to per capita terms using the civilian non-institutional population ages 16 and over). Non-durable consumption is measured by personal consumption expenditures, and investment corresponds to fixed private investment in the National Income and Product Accounts. The inflation rate is measured as the quarter-to-quarter change in the log of the GDP deflator, seasonally adjusted. The short-term nominal interest rate is measured by quarterly averages of daily readings on the three-month U.S. Treasury bill rate, converted from an annualized yield on a discount basis to a quarterly yield to maturity. The three-month T-bill rate tracks the federal funds rate closely over our sample period, and at the end of the sample, after the FOMC established a target range from 0 to 25 basis points for the federal funds rate, the quarterly average federal funds rate and three-month T-bill rate were within a few basis point of each other.

B Prior Distribution of the Parameters

Table B.1 shows the values for the fixed parameters in the estimation. Table B.2 displays the prior distribution for the estimated parameters.

Table B.1: Fixed Parameters

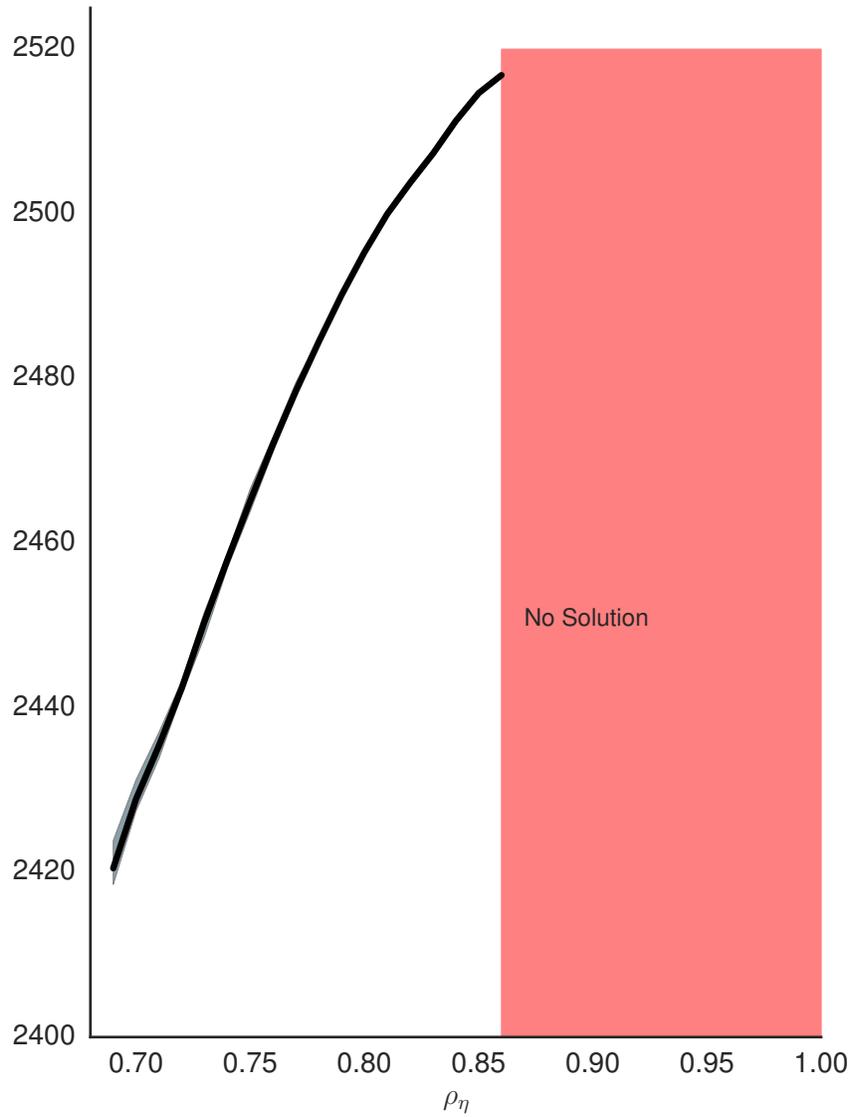
Parameter	Value	Description
δ	0.025	Depreciation of capital stock.
g	1.25	Steady state government spending ($G/Y = 0.2$)
$1/(\varepsilon_p - 1)$	0.20	Steady state net price markup.
$1/(\varepsilon_w - 1)$	0.20	Steady state net wage markup.
ψ_L	1	Disutility of Labor.
ρ_η	0.85	Persistence of the Liquidity Shock.

Table B.2: Prior Distribution

Parameter	Dist.	Para(1)	Para(2)	Parameter	Dist.	Para(1)	Para(2)
Steady State							
$100(\beta^{-1} - 1)$	Gamma	0.25	0.10	$100(\bar{\Pi} - 1)$	Normal	0.62	0.10
$100 \ln(G_z)$	Normal	0.50	0.03	α	Normal	0.30	0.05
Monetary Policy Rule							
ρ_R	Beta	0.60	0.20	γ_{Π}	Normal	1.70	0.30
γ_g	Normal	0.40	0.30	γ_x	Normal	0.40	0.30
Endogenous Propagation							
γ	Beta	0.60	0.10	σ_L	Gamma	2.00	0.75
σ_a	Gamma	5.00	1.00	φ_I	Gamma	4.00	1.00
φ_p	Normal	100.00	25.00	$1 - a$	Beta	0.50	0.15
φ_w	Normal	3000.00	5000.00	$1 - a_w$	Beta	0.50	0.15
Exogenous Processes							
ρ_g	Beta	0.60	0.20	ρ_{μ}	Beta	0.60	0.20
$100\sigma_g$	Inv. Gamma	0.33	2.00	$100\sigma_{\mu}$	Inv. Gamma	0.33	2.00
$100\sigma_Z$	Inv. Gamma	0.33	2.00	$100\sigma_R$	Inv. Gamma	0.33	2.00
$100\sigma_{\eta}$	Inv. Gamma	0.33	2.00				

Notes: Para (1) and Para (2) correspond to the mean and standard deviation of the Beta, Gamma, and Normal distributions and to the upper and lower bounds of the support for the Uniform distribution. For the Inverse (Inv.) Gamma distribution, Para (1) and Para (2) refer to s and ν , where $p(\sigma|\nu, s) \propto \sigma^{-\nu-1} e^{-\nu s^2/2\sigma^2}$.

Figure B.1: Likelihood Function of ρ_η



Notes. Figure shows the mean estimates (black line) of the likelihood function from the particle filter, as well as 90% bands (grey region), as a function of ρ_η , with all other parameters held fixed at their posterior mean level. The red shaded region shows the area for which the model does not solve.

C Parameter Estimates from the Linearized Version of the Model

Table C.1: Posterior Distribution of Parameter Estimates from the Linearized Version of the Model

Parameter	Mean	[05, 95]	Parameter	Mean	[05, 95]
Steady State					
$100(\beta^{-1} - 1)$	0.16	[0.07, 0.26]	$100(\bar{\Pi} - 1)$	0.64	[0.56, 0.72]
$100 \ln(G_z)$	0.50	[0.46, 0.55]	α	0.20	[0.17, 0.23]
Policy Rule					
ρ_R	0.77	[0.69, 0.83]	γ_π	1.78	[1.34, 2.21]
γ_g	0.57	[0.27, 0.87]	γ_x	0.07	[0.01, 0.13]
Endogenous Propagation					
γ	0.76	[0.68, 0.83]	σ_L	2.05	[1.03, 3.36]
σ_a	5.27	[3.79, 7.04]	φ_I	3.51	[2.08, 5.22]
φ_p	102.25	[64.01, 142.54]	$1 - a$	0.64	[0.44, 0.82]
φ_w	5102.38	[2025.78, 9502.62]	$1 - a_w$	0.55	[0.33, 0.76]
Exogenous Processes					
ρ_g	0.70	[0.37, 0.94]	$100\sigma_g$	0.15	[0.12, 0.20]
ρ_μ	0.73	[0.61, 0.84]	$100\sigma_\mu$	9.12	[4.84, 14.80]
$100\sigma_\eta$	0.50	[0.31, 0.73]	$100\sigma_Z$	0.50	[0.31, 0.76]
$100\sigma_R$	0.17	[0.13, 0.21]			

Notes. The table reports the mean, fifth, and ninety-fifth percentiles of the posterior distribution under the linearized model. The model was estimated using a Sequential Monte Carlo algorithm tailored towards linearized DSGE models; see Herbst and Schorfheide (2015).

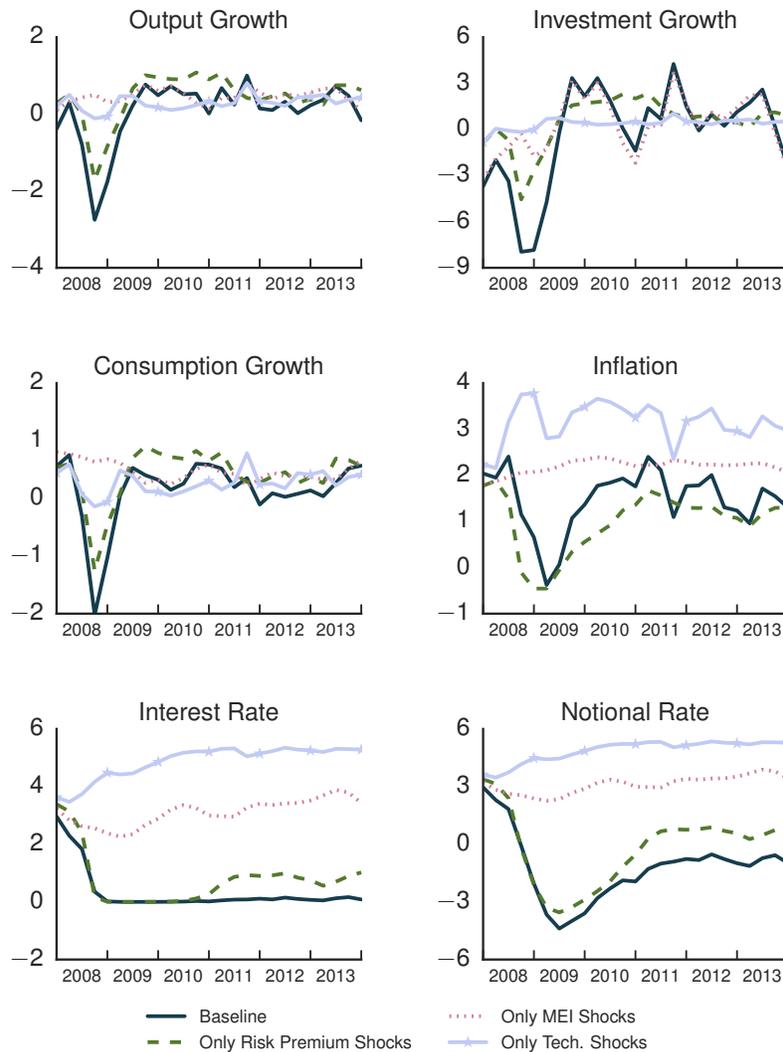
D Estimates Using Low Measurement Error

Table D.1: Posterior Distribution of Parameter Estimates Using Low Measurement Error ($m_e = 0.1$)

Parameter	Mean	[05, 95]	Parameter	Mean	[05, 95]
Steady State					
$100(\beta^{-1} - 1)$	0.14	[0.07, 0.23]	$100(\bar{\Pi} - 1)$	0.62	[0.55, 0.70]
$100\log(G_z)$	0.50	[0.44, 0.54]	α	0.18	[0.16, 0.21]
Policy Rule					
ρ_R	0.78	[0.71, 0.82]	γ_Π	1.60	[1.21, 2.01]
γ_g	0.67	[0.40, 0.95]	γ_x	0.24	[0.11, 0.46]
Endogenous Propagation					
γ	0.67	[0.61, 0.72]	σ_L	2.00	[0.95, 3.30]
σ_a	5.64	[4.04, 7.55]	φ_I	3.95	[2.74, 5.42]
φ_p	77.41	[38.31, 116.64]	$1 - a$	0.15	[0.07, 0.26]
φ_w	1287.68	[275.41, 3100.21]	$1 - a_w$	0.34	[0.16, 0.53]
Exogenous Processes					
ρ_G	0.71	[0.39, 0.97]	$100\sigma_G$	0.17	[0.13, 0.23]
ρ_{μ_I}	0.72	[0.54, 0.86]	$100\sigma_{\mu_I}$	3.81	[2.68, 5.43]
$100\sigma_\eta$	0.42	[0.33, 0.52]	$100\sigma_Z$	0.94	[0.73, 1.20]
$100\sigma_R$	0.16	[0.13, 0.19]			

Notes. The table reports the mean, fifth, and ninety-fifth percentiles of the posterior distribution for the model with low measurement error using an MCMC chain of length 75,000 after a burn in period of 5,000 draws.

Figure D.1: The Contribution of the Estimated Shocks to the Great Recession
 Low Measurement Error ($m_e = 0.1$)



Notes. Figure shows counterfactual trajectories of output growth (upper left panel), investment growth (upper right panel), consumption growth (middle left panel), inflation (middle right panel), the interest rate (bottom left panel), and the notional rate (bottom right panel). The trajectories are computed using smoothed estimates of the states in 2007:Q4 as initial conditions and the smoothed shock estimates from 2008:Q1 to 2014:Q1 for only liquidity shocks (dashed lines), only MEI shocks (dotted line), only technology shocks (line with triangles), and all of the structural shocks (solid line).

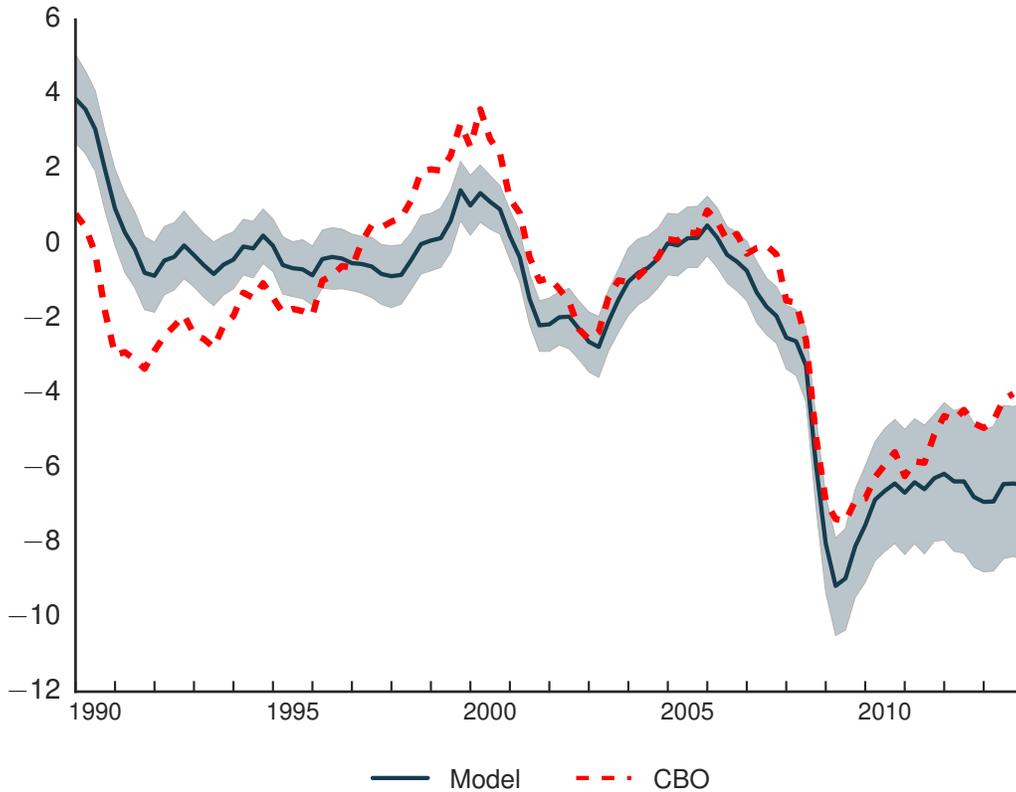
E The Paths of the Output Gap and Additional Shocks

Figure E.1 plots the posterior median and 68 percent pointwise credible sets for model-implied output gap,

$$x_t^g = \alpha \log u_t + (1 - \alpha) \log \left(\frac{N_t}{N} \right),$$

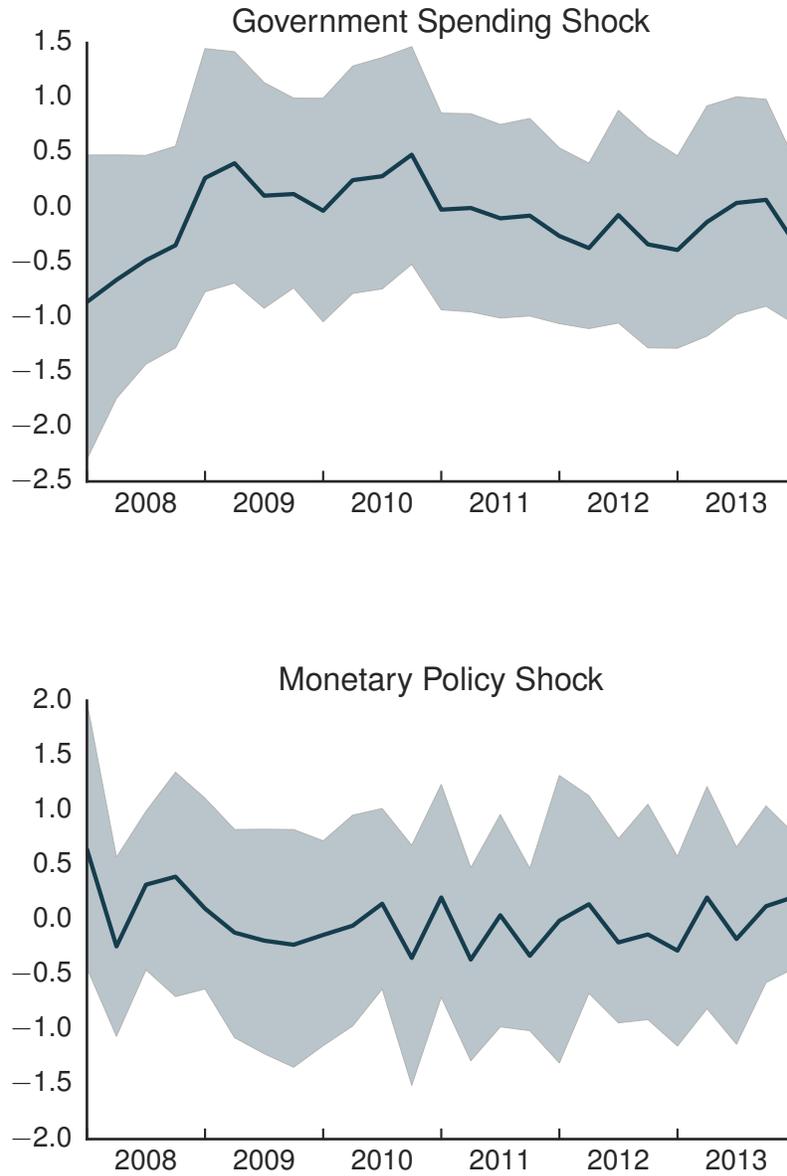
along with the Congressional Budget Office’s (CBO) estimate of the output gap from February 2014. This vintage of the CBO output gap is chosen to be consistent with the end of our sample. The model’s estimated output gap shares the same general features as the one estimated by the CBO, though the CBO output gap often lies outside of the 68 percent bands associated with the model estimate. In particular, the CBO’s output gap closes more quickly in the aftermath of the Great Recession.

Figure E.1: The Path of the Estimated Output Gap



Notes. Figure shows the time series of the mean (solid line) and 68% bands (shaded region) of the smoothed distributions of the model’s output gap. The red dashed line shows the February 2014 vintage of the CBO output gap, constructed using real potential GDP from the CBOs February 2014 report, The Budget and Economic Outlook: 2014 to 2024, and real GDP from Table 1.1.6 of the BEAs February 2014 release of National Accounts (NIPA) data.

Figure E.2: The Path of Two Estimated Shocks During the Great Recession



Notes. Figure shows the time series of the mean (solid line) and 68% bands (shaded region) of the smoothed distributions of $\ln(g_t)$ (top panel) and $\epsilon_{R,t}$ (bottom panel). Both are normalized by their unconditional standard deviations.

F Rotemberg and Calvo Estimates of Nominal Rigidities

Under Calvo, the slope of the price inflation equation is given by:

$$\kappa_p^{Calvo} = \frac{(1 - \xi_p \beta)(1 - \xi_p)}{\xi_p(1 + \iota_p \beta)},$$

where ξ_p corresponds to the probability of changing prices and ι_p represents the degree of price indexation. Evaluating this expression at $\xi_p = 0.787$, $\iota_p = 0.131$, and $\beta = 0.9986$, the posterior medians estimated by Justiniano, Primiceri, and Tambalotti (2011) yields $\kappa_p^{Calvo} = 0.0512$. Under Rotemberg contracts, the slope coefficient is:

$$\frac{\varepsilon_p - 1}{(1 + \beta(1 - a))\varphi_p}$$

Setting $\beta = 0.9986$, $(1 - a) = \iota_p = 0.131$, and $\varepsilon_p - 1 = \frac{1}{0.180}$ as in Justiniano, Primiceri, and Tambalotti (2011) implies a parameter of price adjustment cost of:

$$\varphi_p = \frac{\varepsilon_p - 1}{(1 + \beta(1 - a))\kappa_p^{Calvo}} \approx 93.5.$$

In comparison, our posterior mean estimate of this parameter is 100, and 93.5 is well within the 90 percent credible band.

Regarding the slope of the wage inflation equation, under Calvo, the expression is given by:

$$\kappa_w^{Calvo} = \frac{(1 - \xi_w \beta)(1 - \xi_w)}{\xi_w(1 + \beta)(1 + \nu(1 + \frac{1}{\lambda_w}))}.$$

If this expression is evaluated at the posterior medians estimated by Justiniano, Primiceri, and Tambalotti (2011) in which $\xi_w = 0.777$, $\lambda_w = 0.144$ (which corresponds to $\varepsilon_w - 1 = \frac{1}{\lambda_w}$, $\varepsilon_w \simeq 7.9$), and $\nu = 4.492$, then $\kappa_w^{Calvo} \approx 0.0009$, which implies wage inflation responds very little to labor market slack.

The equivalent slope coefficient using Rotemberg contracts is given by:

$$\frac{\varepsilon_w - 1}{\varphi_w} mc \frac{(1 - \alpha)}{\frac{c}{y}},$$

where the steady state marginal costs is given by: $mc = \frac{\varepsilon_p - 1}{\varepsilon_p}$. To relate this to Calvo contracts, we use:

$$\varphi_w = \frac{\varepsilon_w - 1}{\kappa_w^{Calvo}} \frac{(1 - \alpha)mc}{\frac{c}{y}}.$$

The estimates of Justiniano, Primiceri, and Tambalotti (2011) imply that $\frac{(1 - \alpha)}{\frac{c}{y}} \simeq 1.5$ and $mc \simeq 0.84$ implying extremely high nominal wage adjustment costs as $\frac{\varphi_w}{1000} \approx 9936$. This estimate of wage adjustment costs is outside the upper end of the 90 percent credible set and is substantially higher than our posterior mean estimate.

Technical Appendix

This appendix describes the equilibrium conditions of the model. It uses those equilibrium conditions to characterize the solution as time-invariant functions. We approximate the nonlinear solution using a computationally efficient algorithm that is easily parallelizable and well-suited for handling an occasionally binding constraint. With this algorithm, we then describe the particle filtering algorithm for estimating the likelihood of a nonlinear state space system, highlighting the modifications made to facilitate estimation, including its parallelization. Finally, we lay out the particle filter Metropolis-Hastings algorithm used in conjunction with the particle filter to elicit draws from the posterior distribution.

1 Equilibrium Conditions

In a symmetric equilibrium, optimization by the firms in the economy implies:

$$\left[\frac{\pi_t}{\tilde{\pi}_{t-1}} - 1 \right] \frac{\pi_t}{\tilde{\pi}_{t-1}} = \beta E_t \left\{ \frac{\Lambda_{t+1}}{\Lambda_t} \left[\frac{\pi_{t+1}}{\tilde{\pi}_t} - 1 \right] \frac{\pi_{t+1}}{\tilde{\pi}_t} \frac{Y_{t+1}}{Y_t} \right\} + \frac{\varepsilon_p}{\varphi_p} \left\{ mc_t - \frac{\varepsilon_p - 1}{\varepsilon_p} \right\}, \quad (1.1)$$

$$(1 - \alpha) mc_t = \frac{W_t N_t}{P_t Y_t}, \quad (1.2)$$

$$P_t r_t^k = \frac{\alpha}{1 - \alpha} \frac{W_t N_t}{u_t \bar{K}_t}, \quad (1.3)$$

where the indexation term for price changes is given by:

$$\tilde{\pi}_{t-1} = \bar{\pi}^a \pi_{t-1}^{1-a}. \quad (1.4)$$

The aggregate production function is given by:

$$Y_t = (u_t \bar{K}_t)^\alpha (Z_t N_t)^{1-\alpha}. \quad (1.5)$$

Household optimization for consumption, bond holdings, and wages implies:

$$\Lambda_t = [C_t - \gamma C_{t-1}]^{-1} - \gamma \beta E_t [C_{t+1} - \gamma C_t]^{-1}, \quad (1.6)$$

$$\Lambda_t = \beta R_t \eta_t E_t \{ \Lambda_{t+1} \pi_{t+1}^{-1} \}, \quad (1.7)$$

$$\left[\frac{\pi_{w,t}}{\tilde{\pi}_{w,t}} - 1 \right] \frac{\pi_{w,t}}{\tilde{\pi}_{w,t}} = \beta E_t \left\{ \left[\frac{\pi_{w,t+1}}{\tilde{\pi}_{w,t+1}} - 1 \right] \frac{\pi_{w,t+1}}{\tilde{\pi}_{w,t+1}} \right\} + N_t \Lambda_t \varepsilon_w \varphi_w^{-1} \left\{ \psi_L \frac{N_t^{\sigma_L}}{\Lambda_t} - \frac{\varepsilon_w - 1}{\varepsilon_w} \frac{W_t}{P_t} \right\}. \quad (1.8)$$

In the above, wage inflation is defined by:

$$\pi_{w,t} = \frac{W_t}{W_{t-1}} \quad (1.9)$$

and the indexation term for wage changes is given by:

$$\tilde{\pi}_{w,t} = G_Z \bar{\pi}^{a_w} (\exp(\varepsilon_{Z,t}) \pi_{t-1})^{1-a_w}. \quad (1.10)$$

A household's optimal choice of physical capital, investment, and utilization imply:

$$q_t = \beta E_t \left\{ \frac{\Lambda_{t+1}}{\Lambda_t} \left[r_{t+1}^k u_{t+1} - a(u_{t+1}) + (1 - \delta)q_{t+1} \right] \right\}, \quad (1.11)$$

$$1 = q_t \mu_t \left(1 - \frac{\varphi_I}{2} \left(\frac{I_t}{G_Z I_{t-1}} - 1 \right)^2 - \varphi_I \left(\frac{I_t}{G_Z I_{t-1}} - 1 \right) \frac{I_t}{G_Z I_{t-1}} \right) + \quad (1.12)$$

$$\beta \varphi_I E_t \left\{ q_{t+1} \frac{\Lambda_{t+1}}{\Lambda_t} \mu_{t+1} \left(\frac{I_{t+1}}{G_Z I_t} - 1 \right) \frac{I_{t+1}^2}{G_Z I_t^2} \right\},$$

$$r_t^k = r^k \exp(\sigma_a(u_t - 1)), \quad (1.13)$$

where r^k denotes the rental cost of capital in the non-stochastic steady state and the utilization cost is given by:

$$a(u_t) = \frac{r^k}{\sigma_a} \{ \exp(\sigma_a(u_t - 1)) - 1 \}. \quad (1.14)$$

The capital stock evolves according to:

$$\bar{K}_{t+1} = (1 - \delta)\bar{K}_t + \mu_t \left\{ 1 - \frac{\varphi_I}{2} \left(\frac{I_t}{G_Z I_{t-1}} - 1 \right)^2 \right\} I_t. \quad (1.15)$$

The central bank's desired or notional interest rate is given by:

$$\ln \left(\frac{R_t^N}{R} \right) = \rho_R \ln \left(\frac{R_{t-1}^N}{R} \right) + (1 - \rho_R) \left[\gamma_\pi \ln \left(\frac{\pi_t}{\bar{\pi}} \right) + \gamma_x x_t^g + \gamma_g \ln \left(\frac{Y_t}{G_Z Y_{t-1}} \right) \right] + \epsilon_{R,t}, \quad (1.16)$$

where $R = \beta^{-1} G_Z \bar{\pi}$ denotes the steady state nominal rate. The output gap is given by:

$$x_t^g = \alpha \ln(u_t) + (1 - \alpha) (\ln(N_t) - \ln(N)), \quad (1.17)$$

where N denotes the non-stochastic steady state value of hours worked. The nominal interest rate satisfies the zero lower bound (ZLB) constraint so that:

$$R_t = \max(1, R_t^N). \quad (1.18)$$

The economy's resource constraint is:

$$C_t + I_t + G_t + \frac{\varphi_p}{2} \left[\frac{\pi_t}{\bar{\pi}_{t-1}} - 1 \right]^2 Y_t + a(u_t) \bar{K}_t = Y_t \quad (1.19)$$

The economy's disturbances evolve according to:

$$Z_t = Z_{t-1} G_Z \exp(\epsilon_{Z,t}), \quad \epsilon_{Z,t} \sim iid N(0, \sigma_Z^2), \quad (1.20)$$

$$\ln(\eta_t) = \rho_\eta \ln(\eta_{t-1}) + \epsilon_{\eta,t}, \quad \epsilon_{\eta,t} \sim iid N(0, \sigma_\eta^2), \quad (1.21)$$

$$\ln(\mu_t) = \rho_\mu \ln(\mu_{t-1}) + \epsilon_{\mu,t}, \quad \epsilon_{\mu,t} \sim iid N(0, \sigma_\mu^2), \quad (1.22)$$

$$\ln(g_t) = (1 - \rho_g) \ln g + \rho_g \ln(g_{t-1}) + \epsilon_{g,t}, \quad \epsilon_{g,t} \sim iid N(0, \sigma_g^2), \quad (1.23)$$

where $g_t = \frac{1}{1 - \frac{G_t}{Y_t}}$ and g denotes its value in the non-stochastic steady state. The monetary policy disturbance satisfies $\epsilon_{R,t} \sim iid N(0, \sigma_R^2)$.

1.1 Stationary Representation of Equilibrium Conditions

The random walk in the technology shock, expression (1.20), implies that some of the real variables will be non-stationary. The system of equations can be transformed so that it becomes stationary. The transformed variables are denoted as the lower case of a variable. So, $y_t = \frac{Y_t}{Z_t}$, $c_t = \frac{C_t}{Z_t}$, $i_t = \frac{I_t}{Z_t}$, $w_t = \frac{W_t}{P_t Z_t}$, $\bar{k}_{t+1} = \frac{\bar{K}_{t+1}}{Z_t}$, and $\lambda_t = Z_t \Lambda_t$. It is also convenient to define $G_{Z,t} = G_Z \exp(\epsilon_{Z,t})$.

The stationary versions of the price and wage inflation equations are:

$$\left[\frac{\pi_t}{\tilde{\pi}_{t-1}} - 1 \right] \frac{\pi_t}{\tilde{\pi}_{t-1}} = \beta E_t \left\{ \frac{\lambda_{t+1}}{\lambda_t} \left[\frac{\pi_{t+1}}{\tilde{\pi}_t} - 1 \right] \frac{\pi_{t+1} y_{t+1}}{\tilde{\pi}_t y_t} \right\} + \frac{\varepsilon_p}{\varphi_p} \left\{ mc_t - \frac{\varepsilon_p - 1}{\varepsilon_p} \right\}, \quad (1.24)$$

$$\left[\frac{\pi_{w,t}}{\tilde{\pi}_{w,t}} - 1 \right] \frac{\pi_{w,t}}{\tilde{\pi}_{w,t}} = \beta E_t \left\{ \left[\frac{\pi_{w,t+1}}{\tilde{\pi}_{w,t+1}} - 1 \right] \frac{\pi_{w,t+1}}{\tilde{\pi}_{w,t+1}} \right\} + N_t \lambda_t \varepsilon_w \varphi_w^{-1} \left\{ \psi_L \frac{N_t^{\sigma_L}}{\lambda_t} - \frac{\varepsilon_w - 1}{\varepsilon_w} w_t \right\}, \quad (1.25)$$

where the price and wage indexation terms are given by equations (1.4) and (1.10).

The stationary representations for the consumption Euler equation and its associated Lagrange multiplier can be written as:

$$\lambda_t = V_{\lambda,t} \equiv \beta \eta_t R_t E_t \left\{ \frac{\lambda_{t+1}}{G_{Z,t+1}} \pi_{t+1}^{-1} \right\}, \quad (1.26)$$

$$\lambda_t = \left[c_t - \gamma \frac{c_{t-1}}{G_{Z,t}} \right]^{-1} - \gamma \beta V_{c,t}, \quad (1.27)$$

where $V_{c,t}$, is defined as:

$$V_{c,t} \equiv E_t \left\{ \frac{1}{G_{Z,t+1}} \left[c_{t+1} - \gamma \frac{c_t}{G_{Z,t+1}} \right]^{-1} \right\}. \quad (1.28)$$

The stationary representations of the optimal conditions for capital and investment supply are given by:

$$q_t = V_{q,t} \equiv \beta E_t \left\{ \frac{\lambda_{t+1}}{\lambda_t G_{Z,t+1}} \left[r_{t+1}^k u_{t+1} - a(u_{t+1}) + (1 - \delta) q_{t+1} \right] \right\}, \quad (1.29)$$

$$1 = q_t \mu_t \left(1 - \varphi_I \left(\frac{i_t}{i_{t-1}} \exp(\epsilon_{Z,t}) - 1 \right) \frac{i_t}{i_{t-1}} \exp(\epsilon_{Z,t}) \right) + V_{i,t}, \quad (1.30)$$

where $V_{i,t}$ is defined as:

$$V_{i,t} \equiv \beta \varphi_I E_t \left\{ q_{t+1} \mu_{t+1} \frac{\lambda_{t+1}}{\lambda_t} \left(\frac{i_{t+1}}{i_t} \exp(\epsilon_{Z,t+1}) - 1 \right) \left(\frac{i_{t+1}}{i_t} \right)^2 \exp(\epsilon_{Z,t+1}) \right\} - q_t \mu_t \frac{\varphi_I}{2} \left(\frac{i_t}{i_{t-1}} \exp(\epsilon_{Z,t}) - 1 \right)^2. \quad (1.31)$$

The stationary representation for the optimal utilization of capital is unchanged and given by equation (1.13). The expression for the utilization cost is also unchanged and given by equation (1.14).

After solving for output, the stationary representation for the resource constraint is:

$$y_t = A_{y,t}^{-1} \left(c_t + i_t + a(u_t) \frac{\bar{k}_t}{G_{Z,t}} \right), \quad (1.32)$$

where $A_{y,t}$ is given by:

$$A_{y,t} = \frac{1}{g_t} - \frac{\varphi_p}{2} \left[\frac{\pi_t}{\tilde{\pi}_{t-1}} - 1 \right]^2. \quad (1.33)$$

The stationary representation for the real wage is given by:

$$w_t = \frac{\pi_{w,t} w_{t-1}}{G_{Z,t} \pi_t}. \quad (1.34)$$

The stationary representation of the production function can be expressed to determine hours worked:

$$N_t = \left(\frac{u_t \bar{k}_t}{G_{Z,t}} \right)^{\frac{\alpha}{\alpha-1}} y_t^{\frac{1}{1-\alpha}}. \quad (1.35)$$

Real marginal cost and the rental rate of capital are given by:

$$mc_t = \frac{w_t N_t}{(1-\alpha) y_t}, \quad (1.36)$$

$$r_t^k = \frac{\alpha}{1-\alpha} \frac{G_{Z,t} w_t N_t}{u_t \bar{k}_t}. \quad (1.37)$$

The stationary representation of the capital accumulation equation is:

$$\bar{k}_{t+1} = (1-\delta) \frac{\bar{k}_t}{G_{Z,t}} + \mu_t \left[1 - \frac{\varphi_I}{2} \left(\frac{i_t}{i_{t-1}} \exp(\epsilon_{Z,t}) - 1 \right)^2 \right] i_t. \quad (1.38)$$

The notional rate is given by:

$$\ln \left(\frac{R_t^N}{R} \right) = \rho_R \ln \left(\frac{R_{t-1}^N}{R} \right) + (1-\rho_R) \left[\gamma_\pi \ln \left(\frac{\pi_t}{\bar{\pi}} \right) + \gamma_x x_t^g + \gamma_g \ln \left(\frac{y_t \exp(\epsilon_{Z,t})}{y_{t-1}} \right) \right] + \epsilon_{R,t} \quad (1.39)$$

where the output gap, x_t^g , is defined in equation (1.17) and the nominal rate is defined in equation (1.18).

1.2 Steady State

In the non-stochastic steady state, equation (1.31) implies $V_i = 0$ and equation (1.30) implies $q = 1$. Similarly, in the non-stochastic steady state, $\pi = \tilde{\pi} = \bar{\pi}$, $\pi_w = G_Z \bar{\pi}$, and $u = 1$. The other steady state

relationships include:

$$mc = \frac{\varepsilon_p - 1}{\varepsilon_p}, \quad (1.40)$$

$$mc = \frac{wN}{(1 - \alpha)y}, \quad (1.41)$$

$$\lambda c = (1 - G_Z^{-1}\gamma)^{-1} (1 - \beta G_Z^{-1}\gamma), \quad (1.42)$$

$$R = \beta^{-1} G_Z \bar{\pi}, \quad (1.43)$$

$$\frac{i}{\bar{k}} = 1 - G_Z^{-1}(1 - \delta), \quad (1.44)$$

$$y = (G_Z^{-1}\bar{k})^\alpha N^{1-\alpha}, \quad (1.45)$$

$$r^k = \beta^{-1} G_Z - 1 + \delta, \quad (1.46)$$

$$\varepsilon_w \psi_L \frac{N^{\sigma_L}}{\lambda c} = (\varepsilon_w - 1) \frac{w}{c}, \quad (1.47)$$

$$r^k = \frac{\alpha}{1 - \alpha} \frac{\omega N}{\bar{k}/G_Z}, \quad (1.48)$$

$$y = g(c + i). \quad (1.49)$$

From equations (1.46) and (1.48) it follows that:

$$\bar{k} = \frac{\alpha}{1 - \alpha} \frac{G_Z \omega N}{(\beta^{-1} G_Z - 1 + \delta)}.$$

Using equation (1.41) to replace wN yields the following expression for the capital-output ratio:

$$\frac{\bar{k}}{y} = \frac{(\varepsilon_p - 1)}{\varepsilon_p} \frac{\alpha G_Z}{(\beta^{-1} G_Z - 1 + \delta)}.$$

The steady state capital to output ratio and equation (1.44) can be combined to determine the investment-output ratio:

$$\frac{i}{y} = (1 - G_Z^{-1}(1 - \delta)) \frac{\bar{k}}{y} = \frac{(\varepsilon_p - 1)}{\varepsilon_p} \frac{(1 - G_Z^{-1}(1 - \delta)) \alpha G_Z}{(\beta^{-1} G_Z - 1 + \delta)}.$$

From equation (1.49) it follows:

$$\frac{c}{y} = \frac{1}{g} - \frac{i}{y},$$

where g is a parameter fixed to match the sample average of the ratio of government spending to output, i.e., $g = \frac{1}{1 - \bar{g}}$.

Combining equations (1.47) and (1.42) yields:

$$\psi_L \left(1 - \frac{\gamma}{G_Z}\right) N^{1+\sigma_L} \frac{c}{y} = \frac{\varepsilon_w - 1}{\varepsilon_w} \frac{\omega N}{y} \left(1 - \frac{\gamma}{G_Z} \beta\right).$$

Using expression (1.41) to rewrite the above expression gives:

$$\psi_L \left(1 - \frac{\gamma}{G_Z}\right) N^{1+\sigma_L} \frac{c}{y} = \frac{\varepsilon_w - 1}{\varepsilon_w} mc(1 - \alpha) \left(1 - \frac{\gamma}{G_Z} \beta\right).$$

This expression can be used to determine steady state hours:

$$N = \left[\frac{\frac{\varepsilon_w - 1}{\varepsilon_w} mc(1 - \alpha) \left(1 - \frac{\gamma}{G_Z} \beta\right)}{\psi_L \left(1 - \frac{\gamma}{G_Z}\right) \frac{c}{y}} \right]^{\frac{1}{1+\sigma_L}}. \quad (1.50)$$

2 Solution Algorithm

This section discusses the characterization of the model's solution and then how to approximate it.

2.1 Characterizing the Solution

The model's equilibrium conditions are written as time-invariant functions that depend on the minimum state vector, $(\mathbb{X}_{t-1}, \tau_t)$, where

$$\mathbb{X}_{t-1} = (\bar{k}_t, c_{t-1}, i_{t-1}, w_{t-1}, R_{t-1}^N, \pi_{t-1}, y_{t-1}), \quad (2.1)$$

$$\tau_t = (\eta_t, \mu_t, \epsilon_{Z,t}, \epsilon_{R,t}, g_t). \quad (2.2)$$

We denote the endogenous state vector, \mathbb{X}_{t-1} , as depending on date $t - 1$ variables, because \bar{k}_t is determined at date $t - 1$.

It is convenient to define the time invariant functions, $V_\pi(\mathbb{X}_{t-1}, \tau_t) \equiv V_{\pi,t}$ and $V_w(\mathbb{X}_{t-1}, \tau_t) \equiv V_{w,t}$ using equations (1.24) and (1.25):

$$V_{\pi,t} = \beta E_t \left\{ \frac{\lambda_{t+1}}{\lambda_t} \left[\frac{\pi_{t+1}}{\tilde{\pi}_t} - 1 \right] \frac{\pi_{t+1} y_{t+1}}{\tilde{\pi}_t y_t} \right\} + \frac{\varepsilon_p}{\varphi_p} \left\{ mc_t - \frac{\varepsilon_p - 1}{\varepsilon_p} \right\}, \quad (2.3)$$

$$V_{w,t} = \beta E_t \left\{ \left[\frac{\pi_{w,t+1}}{\tilde{\pi}_{w,t+1}} - 1 \right] \frac{\pi_{w,t+1}}{\tilde{\pi}_{w,t+1}} \right\} + N_t \lambda_t \varepsilon_w \varphi_w^{-1} \left\{ \psi_L \frac{N_t^{\sigma_L}}{\lambda_t} - \frac{\varepsilon_w - 1}{\varepsilon_w} w_t \right\}. \quad (2.4)$$

Given these two functions, equations (1.24) and (1.25) can be rewritten as:

$$\pi(\mathbb{X}_{t-1}, \tau_t) = \frac{\tilde{\pi}_{t-1}}{2} \left(1 + \sqrt{1 + 4V_{\pi,t}} \right) \quad (2.5)$$

$$\pi_w(\mathbb{X}_{t-1}, \tau_t) = \frac{\tilde{\pi}_{w,t}}{2} \left(1 + \sqrt{1 + 4V_{w,t}} \right) \quad (2.6)$$

where $\pi(\mathbb{X}_{t-1}, \tau_t) \equiv \pi_t$ and $\pi_w(\mathbb{X}_{t-1}, \tau_t) \equiv \pi_{w,t}$.

More broadly, the decision rules for the remainder of the endogenous state variables in \mathbb{X}_{t-1} can be determined given the vector of functions:

$$V(\mathbb{X}_{t-1}, \tau_t) = (V_\pi(\mathbb{X}_{t-1}, \tau_t), V_w(\mathbb{X}_{t-1}, \tau_t), V_c(\mathbb{X}_{t-1}, \tau_t), V_i(\mathbb{X}_{t-1}, \tau_t), \\ V_\lambda(\mathbb{X}_{t-1}, \tau_t), V_q(\mathbb{X}_{t-1}, \tau_t), V_u(\mathbb{X}_{t-1}, \tau_t)), \quad (2.7)$$

where $V_u(\mathbb{X}_{t-1}, \tau_t) \equiv u_t$ is a function determining the equilibrium value for the utilization of capital.

Using these functions, we determine the remaining endogenous variables as follows. The Lagrange multiplier on the household's budget constraint, $\lambda(\mathbb{X}_{t-1}, \tau_t)$, can be determined using equation (1.26) and $V_\lambda(\mathbb{X}_{t-1}, \tau_t)$. Equation (1.27) then determines the decision rule for consumption, $c(\mathbb{X}_{t-1}, \tau_t)$ as:

$$c(\mathbb{X}_{t-1}, \tau_t) = \gamma \frac{c_{t-1}}{G_{Z,t}} + (\lambda_t + \gamma\beta V_{c,t})^{-1}. \quad (2.8)$$

Taking $V_{i,t}$ as given and applying the quadratic equation to expression (1.30), equilibrium investment is given by:

$$i(\mathbb{X}_{t-1}, \tau_t) = \frac{i_{t-1}}{2 \exp(\epsilon_{Z,t})} \left[1 + \sqrt{1 + \frac{4}{\varphi_I} \left(1 - \frac{1 - V_{i,t}}{q_t \mu_t} \right)} \right]. \quad (2.9)$$

Given the functions for c_t , i_t , π_t and the vector of functions, $V(\mathbb{X}_{t-1}, \tau_t)$, equilibrium output, $y(\mathbb{X}_{t-1}, \tau_t)$, is determined using expression (1.32). Equation (1.34) determines the real wage, $w(\mathbb{X}_{t-1}, \tau_t)$, using the functions for π_t and $\pi_{w,t}$, while equation (1.35) determines hours worked, $N(\mathbb{X}_{t-1}, \tau_t)$ using the functions for y_t and u_t . Equations (1.36)-(1.39) can be used to determine the functions, $mc(\mathbb{X}_{t-1}, \tau_t)$, $r^k(\mathbb{X}_{t-1}, \tau_t)$, $k(\mathbb{X}_{t-1}, \tau_t) \equiv \bar{k}_{t+1}$, and $R^N(\mathbb{X}_{t-1}, \tau_t)$. Accordingly, given $V(\mathbb{X}_{t-1}, \tau_t)$, we can define the vector of functions that define the decision rules, $\mathbb{X}_t = g_X(\mathbb{X}_{t-1}, \tau_t)$, where

$$g_X(\mathbb{X}_{t-1}, \tau_t) = (k(\mathbb{X}_{t-1}, \tau_t), c(\mathbb{X}_{t-1}, \tau_t), i(\mathbb{X}_{t-1}, \tau_t), w(\mathbb{X}_{t-1}, \tau_t), R^N(\mathbb{X}_{t-1}, \tau_t), \pi(\mathbb{X}_{t-1}, \tau_t), y(\mathbb{X}_{t-1}, \tau_t)). \quad (2.10)$$

We use the functions, $V(\mathbb{X}_{t-1}, \tau_t)$, to determine the decisions rule though we do not approximate $V(\mathbb{X}_{t-1}, \tau_t)$, directly, because they inherit a kink associated with the zero lower bound constraint on the nominal rate. Instead, we follow the methodology described in Gust, López-Salido, and Smith (2012), which builds on Christiano and Fisher (2000). Specifically, we approximate functions, $V_{l,i}(\mathbb{X}_{t-1}, \tau_t)$, that are smoother and easier to approximate by specifying:

$$V_l(\mathbb{X}_{t-1}, \tau_t) = V_{l,1}(\mathbb{X}_{t-1}, \tau_t) \mathbb{I}(\mathbb{X}_{t-1}, \tau_t) + V_{l,2}(\mathbb{X}_{t-1}, \tau_t) (1 - \mathbb{I}(\mathbb{X}_{t-1}, \tau_t)). \quad (2.11)$$

for $l \in \{\pi, w, c, i, \lambda, q, u\}$ and $j = 1, 2$ and where $\mathbb{I}(\mathbb{X}_{t-1}, \tau_t)$ is defined by:

$$\mathbb{I}(\mathbb{X}_{t-1}, \tau_t) = 1 \text{ if } R(\mathbb{X}_{t-1}, \tau_t) > 0 \quad (2.12)$$

$$= 0 \text{ otherwise.} \quad (2.13)$$

In the above, $R(\mathbb{X}_{t-1}, \tau_t) = \max(1, R_1^N(\mathbb{X}_{t-1}, \tau_t))$ where $R_1^N(\mathbb{X}_{t-1}, \tau_t)$ denotes the value of the notional rate derived from evaluating the functions $V_{l,1}(\mathbb{X}_{t-1}, \tau_t)$ and using expression (1.39). (For each variable, we use $j = 1$ to denote the function associated with the regime with a positive nominal rate and $j = 2$ to denote that function associated with the ZLB regime; similarly, $g_{X,j}(\mathbb{X}_{t-1}, \tau_t)$ denotes the vector of regime-specific decision rules.)

The functions, $V_{l,j}(\mathbb{X}_{t-1}, \tau_t)$, satisfy the residual functions, $\nu_{l,j}(\mathbb{X}_{t-1}, \tau_t)$, for $l \in \{\pi, w, c, i, \lambda, q, u\}$

and $j = 1, 2$:

$$\nu_{\lambda,1}(\mathbb{X}_{t-1}, \tau_t) = V_{\lambda,1}(\mathbb{X}_{t-1}, \tau_t) - \beta \eta_t R_t E_t \left\{ \frac{\lambda(\mathbb{X}_t, \tau_{t+1})}{G_{Z,t+1}} \pi^{-1}(\mathbb{X}_t, \tau_{t+1}) \right\} = 0, \quad (2.14)$$

$$\nu_{\lambda,2}(\mathbb{X}_{t-1}, \tau_t) = V_{\lambda,2}(\mathbb{X}_{t-1}, \tau_t) - \beta \eta_t E_t \left\{ \frac{\lambda(\mathbb{X}_t, \tau_{t+1})}{G_{Z,t+1}} \pi^{-1}(\mathbb{X}_t, \tau_{t+1}) \right\} = 0,$$

$$\nu_{c,j}(\mathbb{X}_{t-1}, \tau_t) = V_{c,j}(\mathbb{X}_{t-1}, \tau_t) - E_t \left\{ \frac{1}{G_{Z,t+1}} \left[c(\mathbb{X}_t, \tau_{t+1}) - \gamma \frac{c_j(\mathbb{X}_{t-1}, \tau_t)}{G_{Z,t+1}} \right]^{-1} \right\} = 0, \quad (2.15)$$

$$\nu_{q,j}(\mathbb{X}_{t-1}, \tau_t) = V_{q,j}(\mathbb{X}_{t-1}, \tau_t) - \beta E_t \left\{ \frac{\tilde{\lambda}_j}{G_{Z,t+1}} \left[\tilde{r}^k(\mathbb{X}_t, \tau_{t+1}) + (1 - \delta)q(\mathbb{X}_t, \tau_{t+1}) \right] \right\} = 0, \quad (2.16)$$

$$\begin{aligned} \nu_{i,j}(\mathbb{X}_{t-1}, \tau_t) &= V_{i,j}(\mathbb{X}_{t-1}, \tau_t) - \beta \varphi_I E_t \left\{ q(\mathbb{X}_t, \tau_{t+1}) \mu_{t+1} \tilde{\lambda}_j (\tilde{i}_j \exp(\epsilon_{Z,t+1}) - 1) \tilde{i}_j^2 \exp(\epsilon_{Z,t+1}) \right\} \\ &\quad - q_j(\mathbb{X}_{t-1}, \tau_t) \mu_t \frac{\varphi_I}{2} \left(\frac{i_j(\mathbb{X}_{t-1}, \tau_t)}{i_{t-1}} \exp(\epsilon_{Z,t}) - 1 \right)^2 = 0, \end{aligned} \quad (2.17)$$

$$\begin{aligned} \nu_{\pi,j}(\mathbb{X}_{t-1}, \tau_t) &= V_{\pi,j}(\mathbb{X}_{t-1}, \tau_t) - \beta E_t \left\{ \tilde{\lambda}_j \left[\frac{\pi(\mathbb{X}_t, \tau_{t+1})}{\tilde{\pi}_t} - 1 \right] \frac{\pi(\mathbb{X}_t, \tau_{t+1})}{\tilde{\pi}_t} \frac{y(\mathbb{X}_t, \tau_{t+1})}{y_j(\mathbb{X}_{t-1}, \tau_t)} \right\} + \\ &\quad \frac{\varepsilon_p}{\varphi_p} \left\{ m c_j(\mathbb{X}_{t-1}, \tau_t) - \frac{\varepsilon_p - 1}{\varepsilon_p} \right\} = 0, \end{aligned} \quad (2.18)$$

$$\nu_{w,j}(\mathbb{X}_{t-1}, \tau_t) = V_{w,j}(\mathbb{X}_{t-1}, \tau_t) - \beta E_t \left\{ \left[\frac{\pi_w(\mathbb{X}_t, \tau_{t+1})}{\tilde{\pi}_{w,t+1}} - 1 \right] \frac{\pi_w(\mathbb{X}_t, \tau_{t+1})}{\tilde{\pi}_{w,t+1}} \right\} + \quad (2.19)$$

$$N_j(\mathbb{X}_{t-1}, \tau_t) \lambda_j(\mathbb{X}_{t-1}, \tau_t) \varepsilon_w \varphi_w^{-1} \left\{ \psi_L \frac{N_j(\mathbb{X}_{t-1}, \tau_t)^{\sigma_L}}{\lambda_j(\mathbb{X}_{t-1}, \tau_t)} - \frac{\varepsilon_w - 1}{\varepsilon_w} w_j(\mathbb{X}_{t-1}, \tau_t) \right\} = 0,$$

$$\nu_{u,j}(\mathbb{X}_{t-1}, \tau_t) = V_{u,j}(\mathbb{X}_{t-1}, \tau_t) - 1 - \frac{1}{\sigma_a} \log \left(\frac{r_j^k(\mathbb{X}_{t-1}, \tau_t)}{r^k} \right) = 0, \quad (2.20)$$

where

$$\begin{aligned} \tilde{r}^k(\mathbb{X}_t, \tau_{t+1}) &= r^k(\mathbb{X}_t, \tau_{t+1}) u(\mathbb{X}_t, \tau_{t+1}) - a(u(\mathbb{X}_t, \tau_{t+1})), \\ \tilde{\lambda}_j &\equiv \tilde{\lambda}_j(\mathbb{X}_t, \mathbb{X}_{t-1}, \tau_t, \tau_{t+1}) = \frac{\lambda(\mathbb{X}_t, \tau_{t+1})}{\lambda_j(\mathbb{X}_{t-1}, \tau_t)}, \\ \tilde{i}_j &\equiv \tilde{i}_j(\mathbb{X}_t, \mathbb{X}_{t-1}, \tau_t, \tau_{t+1}) = \frac{i(\mathbb{X}_t, \tau_{t+1})}{i_j(\mathbb{X}_{t-1}, \tau_t)}, \end{aligned}$$

and $\mathbb{X}_t = g_{X,j}(\mathbb{X}_{t-1}, \tau_t)$.

Because the functions, $V_l(\mathbb{X}_{t-1}, \tau_t)$ depend directly on the nominal interest rate, we expect them to have a kink or non-differentiability. By contrast, the counterpart functions, $V_{l,j}(\mathbb{X}_{t-1}, \tau_t)$, that are

indexed by the interest-rate regime do not depend on the current indicator function and thus are more likely to be smooth. The regime-specific functions still depend on a secondary effect that the kink in the nominal rate next period has on the expectations of future variables. This secondary effect enters through evaluating the decision rule in the next period (i.e., $X_{t+1} = g_X(\mathbb{X}_t, \tau_{t+1})$) but it does not affect \mathbb{X}_t which depends on $g_{X,j}(\mathbb{X}_t, \tau_t)$. Following the arguments in Christiano and Fisher (2000), the secondary effects of the kink on the regime-specific functions should be small because of the presence of the expectations operator, which involves summing over the future states of τ and acts to smooth out the regime-specific functions. While our approach of using relatively smooth functions is similar to Christiano and Fisher (2000), our approach is more general as it does not require that we parameterize functions that depend only on future variables.

2.2 Approximating the Solution

We approximate the functions, $V_{l,j}(\mathbb{X}_{t-1}, \tau_t)$, as follows:

$$V_{l,j}(\mathbb{X}_{t-1}, \tau_t) \approx \sum_{k=1}^{N^\tau} T(\varphi(\mathbb{X}_{t-1})) a_{l,j,k} \Gamma_k(\tau_t), \quad (2.21)$$

where $T(\varphi(\mathbb{X}_{t-1}))$ is a 1×17 vector constructed from an anisotropic Smolyak method. Specifically, $T(\varphi(\mathbb{X}_{t-1}))$ includes a constant and the first and second degree Chebyshev polynomials for each variable in \mathbb{X}_{t-1} . It also includes the third and fourth degree Chebyshev polynomials for investment. We augment investment with higher-order polynomials, because we found that in practice this helped reduce the size of the residual error, $\nu_{q,j}(\mathbb{X}_{t-1}, \tau_t)$. (See Judd, Maliar, Maliar, and Valero (2014) for a discussion of anisotropic Smolyak methods.)

For each state variable in \mathbb{X}_{t-1} , we use $\varphi_f : [\underline{\mathbb{X}}_f, \overline{\mathbb{X}}_f] \rightarrow [-1, 1]$ for $f = 1, 2, \dots, 7$, where $\varphi_f(\mathbb{X}_{t-1,f}) = \frac{2(\mathbb{X}_{t-1,f} - \underline{\mathbb{X}}_f)}{\overline{\mathbb{X}}_f - \underline{\mathbb{X}}_f}$ and f indexes one of the state variable in \mathbb{X}_{t-1} . So $\varphi(\mathbb{X}_{t-1})$ is given by:

$$\varphi(\mathbb{X}_{t-1}) = [\varphi_1(\mathbb{X}_{t-1,1}), \dots, \varphi_7(\mathbb{X}_{t-1,7})],$$

and $\overline{\mathbb{X}}_f$ and $\underline{\mathbb{X}}_f$ are maximum and minimum values of each state variable chosen to encompass a wide interval.

In equation (2.21), $a_{l,j,k}$ is a 17×1 vector of parameters for $l \in \{\lambda, c, q, i, \pi, w, u\}$, $j = 1, 2$, and $k = 1, 2, \dots, N^\tau$. The function $\Gamma_k(\tau_t)$ is the product of univariate piecewise linear basis functions for each variable in τ_t . These basis functions use evenly-spaced breakpoints which are also the interpolation nodes. We use 3 breakpoints for each shock except η_t for which we use 7 so that $N^\tau = 567$. We could in principle use a sparse grid to construct $\Gamma_k(\tau_t)$ as we did for $T(\varphi(\mathbb{X}_{t-1}))$.

Our solution strategy involves finding the matrix of coefficients, a^* such that:

$$\nu_{l,j}(\mathbb{X}_m, \tau_k; a^*) = 0, \quad (2.22)$$

for $j = 1, 2$, and $l \in \{\pi, w, c, i, \lambda, q, u\}$. In equation (2.22), \mathbb{X}_m denotes that the vector of state variables is evaluated at each point, $m = 1, 2, \dots, 17$, on an appropriately constructed Smolyak grid using Chebyshev extrema for the unidimensional grid points. Also, each exogenous state variable is

evaluated at one of its equally-spaced breakpoints, $k = 1, 2, \dots, N^\tau$. The matrix a^* consists of:

$$a^* = [a_{\lambda,1,1}^*, \dots, a_{\lambda,1,N^\tau}^*, a_{\lambda,2,1}^*, \dots, a_{\lambda,2,N^\tau}^*, a_{c,1,1}^*, \dots, a_{u,2,N^\tau}^*],$$

and is $17 \times 14N^\tau$. With $N^\tau = 567$, the matrix a^* has 134,946 elements. We find even though this a very large number of coefficients, we are able to solve and estimate the model because, as discussed below, there are substantial gains to parallelizing the solution algorithm.

Before discussing how we determine a^* , it is important to note that we use Gauss-Hermite integration to approximate the conditional expectations operator in $\nu_{i,j}(\mathbb{X}_{t-1}, \tau_t)$. We use 3 nodes per shock and construct the multidimensional integral as a tensor product of the one-dimensional nodes so that there are 243 nodes in total. It may be possible to speed up the solution algorithm further using the monomial rule discussed in Judd, Maliar, and Maliar (2011) to approximate the multidimensional integrals.

2.3 Parallelization of Solution Algorithm

For an initial guess of a^* , we can evaluate the decision rule, $g_X(\mathbb{X}_t, \tau_t)$, and conditional expectations operators in equations (2.14)-(2.20). With these expressions in hand, we use the fixed point algorithm described in Judd, Maliar, Maliar, and Valero (2014) to update our guess for a^* . Given that our polynomial approximation is linear in these coefficients, updating these coefficients involves only trivial calculations and avoids using a numerical routine for solving nonlinear equations. Moreover, updating these coefficients is easily parallelizable, as updating each of the 17×1 vectors, $a_{i,j,k}$, involves a relatively small set of calculations that is independent from the calculations necessary to update the other coefficients. Using a Message Passing Interface, we can distribute this updating step for $a_{i,j,k}$ across processors and make the necessary calculation independently. Figure 2.1 shows the runtime for solving the model 100 times (in log seconds) against the number of processors used in the procedure. With 1 processor, the model takes about three and half hours to solve 100 times, or a little over 2 minutes on average. With 300 processors, the model takes about 3 minutes to solve 100 times, or about 1.8 seconds on average.

3 Particle Filter

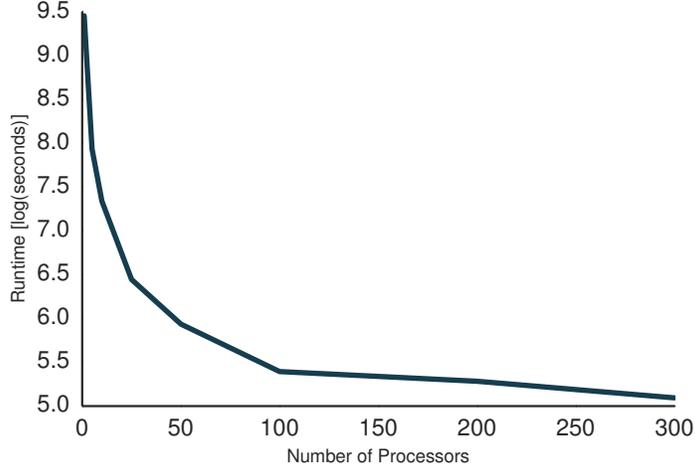
This section describes the particle filter used to estimate the likelihood given a set of parameters. The literature on particles filters is vast: surveys and tutorials can be found, for instance, in Arulampalam, Maskell, Gordon, and Clapp (2002), Cappé, Godsill, and Moulines (2007), Doucet and Johansen (2011), and Creal (2012). The presentation of the algorithm is adapted from Herbst and Schorfheide (2015).

The starting point is the nonlinear state space model:

$$s_t = \Phi(s_{t-1}, \epsilon_t; \theta), \quad \epsilon_t \sim N(0, I), \tag{3.1}$$

$$y_t = \Psi(s_t, \epsilon_t; \theta) + u_t, \quad u_t \sim N(0, \Sigma_u), \tag{3.2}$$

Figure 2.1: Effects of Parallelization: Model Solution



where

$$\begin{aligned}
 s_t &= [\bar{k}_{t+1}, y_t, c_t, i_t, \pi_t, R_t^N, w_t, \eta_t, \mu_t, g_t, y_{t-1}, c_{t-1}, i_{t-1}], \\
 \epsilon_t &= [\epsilon_{\eta,t}, \epsilon_{\mu,t}, \epsilon_{Z,t}, \epsilon_{g,t}, \epsilon_{R,t}], \text{ and} \\
 \theta &= [\beta, \bar{\Pi}, g_z, \alpha, \rho_R, \gamma_{\Pi}, \gamma_g, \gamma_x, \gamma, \sigma_L, \sigma_a, \varphi_I, \varphi_p, \varphi_w, a, a_w, \rho_g, \rho_{\mu}, \sigma_{\eta}, \sigma_{\mu_I}, \sigma_Z, \sigma_g, \sigma_R].
 \end{aligned}$$

The nonlinear transition equations, $\Phi(s_{t-1}, \epsilon_t; \theta)$, are formed using the nonlinear approximation to the decision rules, $g_X(\mathbb{X}_t, \tau_t)$, described in Section 2 and the transition equations for the shocks:

$$\tau_t = \Omega_1 \tau_{t-1} + \epsilon_t, \quad (3.3)$$

where Ω_1 is a 5×5 diagonal matrices whose diagonal elements are the AR(1) coefficients of the shocks.

The particle filter recursively produces discrete approximations to the distribution of the states, s_t , condition on time $t - 1$ information (*forecasting distribution*) and t information (*updated distribution*). We generically refer to the set of the tuples $\{s_t^j, W_t^j\}_{j=1}^M$, which approximates $s_t|Y_{1:t}$, as particles, where M denotes the number of particles in the approximation. The object s_t^j references a point in the state space and W_t^j denotes the weight associated with that point. Note that $\sum_j^M W_t^j = M$.

The recursive formulation is useful for helping to understand the particle filter. Given a time $t - 1$ particle approximation $\{s_{t-1}^j, W_{t-1}^j\}_{j=1}^M$, obtain a t time particle approximation, roughly speaking, as follows. First, simulate forward the proposed particles s_{t-1}^j to obtain particles s_t^j , and second, re-weight these particles using the new data, y_t . The first step is known as forecasting and the second updating. We use the simplest particle filter—with some modifications described in Section 3.1—known as the Bootstrap particle filter (BSPF), which was first used in Gordon, Salmond, and Smith (1993). The key idea is that forecasted states come by drawing ϵ_t^j and iterating the state equation (3.1) forward to obtain: $s_t^j = \Phi(s_{t-1}^j, \epsilon_t^j; \theta)$. A nice by-product of the BSPF is that the likelihood function can be computed directly from the weights. The details are produced in Algorithm 1.

Algorithm 1 (Bootstrap Particle Filter)

1. **Initialization.** Simulate the initial particles from the distribution $s_0^j \stackrel{iid}{\sim} p(s_0)$ and set $W_0^j = 1$, $j = 1, \dots, M$.

2. **Recursion.** For $t = 1, \dots, T$:

(a) **Forecasting s_t .** Propagate the period $t - 1$ particles $\{s_{t-1}^j, W_{t-1}^j\}$ by iterating the state-transition equation forward:

$$\tilde{s}_t^j = \Phi(s_{t-1}^j, \epsilon_t^j; \theta), \quad \epsilon_t^j \sim F_\epsilon(\cdot; \theta). \quad (3.4)$$

An approximation of $\mathbb{E}[h(s_t)|Y_{1:t-1}, \theta]$ is given by

$$\hat{h}_{t,M} = \frac{1}{M} \sum_{j=1}^M h(\tilde{s}_t^j) W_{t-1}^j. \quad (3.5)$$

(b) **Forecasting y_t .** Define the incremental weights

$$\tilde{w}_t^j = p(y_t | \tilde{s}_t^j, \theta). \quad (3.6)$$

The predictive density $p(y_t | Y_{1:t-1}, \theta)$ can be approximated by

$$\hat{p}(y_t | Y_{1:t-1}, \theta) = \frac{1}{M} \sum_{j=1}^M \tilde{w}_t^j W_{t-1}^j. \quad (3.7)$$

The incremental weights take the form

$$\tilde{w}_t^j = (2\pi)^{-n/2} |\Sigma_u|^{-1/2} \exp \left\{ -\frac{1}{2} (y_t - \Psi(\tilde{s}_t^j, t; \theta))' \Sigma_u^{-1} (y_t - \Psi(\tilde{s}_t^j, t; \theta)) \right\}, \quad (3.8)$$

where n here denotes the dimension of y_t .

(c) **Updating.** Define the normalized weights

$$\tilde{W}_t^j = \frac{\tilde{w}_t^j W_{t-1}^j}{\frac{1}{M} \sum_{j=1}^M \tilde{w}_t^j W_{t-1}^j}. \quad (3.9)$$

An approximation of $\mathbb{E}[h(s_t)|Y_{1:t}, \theta]$ is given by

$$\tilde{h}_{t,M} = \frac{1}{M} \sum_{j=1}^M h(\tilde{s}_t^j) \tilde{W}_t^j. \quad (3.10)$$

(d) **Selection.** Define the Effective Sample Size as:

$$\widehat{ESS}_t = M / \left(\frac{1}{M} \sum_{j=1}^M (\tilde{W}_t^j)^2 \right).$$

Let $\rho_t = \mathbf{1}_{\{ESS_t < M/2\}}$. Case (i): If $\rho_t = 1$ resample the particles via multinomial resampling. Let $\{s_t^j\}_{j=1}^M$ denote M iid draws from a multinomial distribution characterized by support points and weights $\{\tilde{s}_t^j, \tilde{W}_t^j\}$ and set $W_t^j = 1$ for $j = 1, \dots, M$. Case (ii): If $\rho_t = 0$, let $s_t^j = \tilde{s}_t^j$ and $W_t^j = \tilde{W}_t^j$ for $j = 1, \dots, M$. An approximation of $\mathbb{E}[h(s_t)|Y_{1:t}, \theta]$ is given by

$$\bar{h}_{t,M} = \frac{1}{M} \sum_{j=1}^M h(s_t^j) W_t^j. \quad (3.11)$$

3. Likelihood Approximation. The approximation of the log-likelihood function is given by

$$\ln \hat{p}(Y_{1:T}|\theta) = \sum_{t=1}^T \ln \left(\frac{1}{M} \sum_{j=1}^M \tilde{w}_t^j W_{t-1}^j \right). \quad (3.12)$$

3.1 Adaption

While the BSPF is extremely simple to implement, it can perform poorly in practice, particularly in the presence of outliers. The forecast distribution can be highly mismatched with the updated distribution, which manifests itself as an extremely uneven distribution of the weights, and thus imprecise estimates of the likelihood function. To avoid degeneration of the particle filter in $t=2008:Q4$, we adapt the innovations of the proposal distribution, similar to Aruoba, Cuba-Borda, and Schorfheide (2016).

Details. For the BSPF, \tilde{s}_t is generated by drawing $\epsilon_t^j \sim N(0, I)$ and using equation (3.1), and has density $p(\tilde{s}_t^j | s_{t-1}^j; \theta)$. Under a generic PF, we construct \tilde{s}_t by sampling from an arbitrary distribution with density $g_t(\tilde{s}_t^j | s_{t-1}^j; \theta)$. Specifically, we simulate ϵ_t^j instead from $N(\mu, \Sigma)$, to elicit a proposal from $g_t(\tilde{s}_t^j | s_{t-1}^j; \theta)$.

When using this proposal distribution, the weights in the particle filter must be adjusted by a factor:

$$\kappa = \frac{p(\tilde{s}_t^j | s_{t-1}^j; \theta)}{g_t(\tilde{s}_t^j | s_{t-1}^j; \theta)}. \quad (3.13)$$

When applying a change of variable formula to represent $p(\tilde{s}_t^j | s_{t-1}^j; \theta)$ and $g_t(\tilde{s}_t^j | s_{t-1}^j; \theta)$, both densities contain the same Jacobian. This term drops out from the multiplicative ratio in (3.13), and it is easy to deduce that

$$\kappa = \frac{\exp \left\{ -\frac{1}{2} \epsilon_t^{j'} \epsilon_t^j \right\}}{|\Sigma|^{-1/2} \exp \left\{ (\epsilon_t^j - \mu)' \Sigma^{-1} (\epsilon_t^j - \mu) \right\}}, \quad (3.14)$$

with n the dimensionality of y_t . Unlike Aruoba, Cuba-Borda, and Schorfheide (2016), we do not use a gridsearch algorithm to generate μ and Σ . Instead, after extensive experimentation, for $t=2008:Q4$, we set

$$\mu = [3, 0, 0, 0, 0]' \text{ and } \Sigma = 1.2 \times I.$$

For every other time period we set $\mu = 0$ and $\Sigma = I$; i.e., we use the standard BSPF.

4 Parallelization of the Particle Filter

Parallelization of the particle filter is difficult because of the required communication during the resampling phase. We construct a particle filter adapted to a distributed computing environment. In this section, we sketch out some key aspects of the filter.

Suppose that we have K processors²⁶ and we are using M total particles in the particle filter. Let $M_{local} = M/K$, assume that M_{local} is an integer, and let $s_t^{i,k}$ denote the i th particle on the k th processor and similarly for $W_t^{i,k}$. It is obvious that the forecasting step can be done in parallel across the K processors. To update, we use a Message Passing Interface to aggregate the weights. The key part of the parallelized particle filter is selection and rebalancing. First, we resample every period ($\rho_t = 1$, for all t), but we only resample (using systematic resampling) among the M_{local} particles on each processor. Instead of a sample with uniform weights, we are left with a sample which is evenly weighted *on a given processor*, with each particle having weight:

$$W_t^{i,k} = M_{local}^{-1} \sum_{j=1}^{M_{local}} \tilde{W}_t^{j,k}.$$

We account for the fact that distribution of total weight across processors can become uneven using the following procedure. Let:

$$\alpha_k = \frac{\sum_{i=1}^{M_{local}} W_t^{i,k}}{\sum_{j=1}^K \sum_{i=1}^{M_{local}} W_t^{i,j}}, \quad (4.1)$$

where α_k is the mass of the particle distribution located on processor k . Define the effective number of processors as

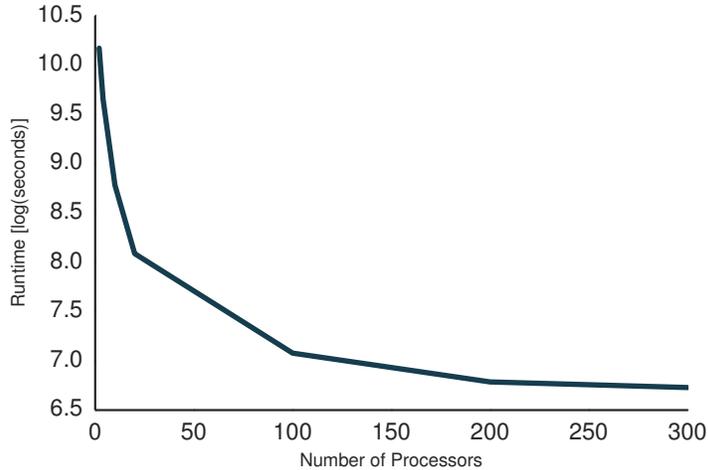
$$EP_t = \frac{1}{\sum_{k=1}^K \alpha_k^2}. \quad (4.2)$$

If $EP_t < K$, then reshuffle the particles among the processors by first ranking the processors according to α_k . Then assign each processor a partner in reverse order: the k with the largest α_k with the k with the smallest α_k , and so on. To rebalance the weights across particles, have each partner exchange $M_{exchange}$ ($< M_{local}$) particles with one another deterministically (i.e., use the first $M_{exchange}$ particles).

This procedure helps ensure that particles will not degenerate on any one processor, while still reducing the size of the resampling problem. Moreover, parallelization allows one to use extremely large number of particles without performance degradation because of memory constraints. Figure 4.1 demonstrates the effectiveness of parallelization. by plotting the run time for 100 evaluations of the particle filter on the vertical axis and for different values of K in Figure 4.1. As seen in the Figure, using two processors, it takes about 7 hours for 100 evaluations of the particle filter, or 4 minutes and 12 seconds per evaluation. When using 300 processors, it takes only about 13 minutes. While the returns to parallelization are not as large using as for the solution algorithm, they are still substantial.

²⁶This is shorthand for whatever the lowest unit of instruction is.

Figure 4.1: Effects of Parallelization: Particle Filter



5 Stability Analysis

It is crucial for the particle filter to produce *stable*—i.e., with low variance—estimates of the likelihood. When used within an MCMC algorithm, particle-filter-based estimators of the likelihood with high variance will tend to get stuck at a single value after an improbably high likelihood estimate. This means that the chain generated by the MCMC algorithm will converge very slowly (or not at all). In this section, we repeatedly apply the particle filter on a single representative parameter draw to show that our filter produces stable estimates.

In this section, we repeatedly apply the particle filter to a single, representative parameter value to highlight the stability of our particle filter. We also examine the contribution of the adaption discussed in Section 3.1. The parameter values use are given in Table 5.1.

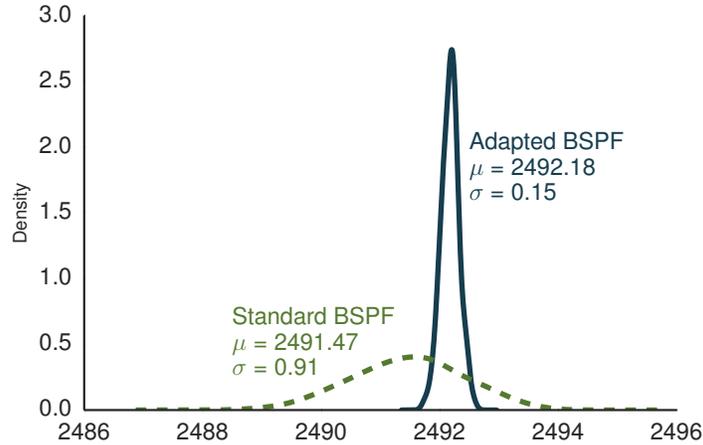
Table 5.1: Parameter Values for Stability Analysis

Parameter	Value	Parameter	Value	Parameter	Value
$100(\beta^{-1} - 1)$	0.11	γ	0.69	ρ_g	0.53
$100(\bar{\Pi} - 1)$	0.60	σ_L	1.27	$100\sigma_g$	0.14
$100\ln(G_z)$	0.54	σ_a	4.60	ρ_{μ_I}	0.72
α	0.21	φ_I	2.79	$100\sigma_{\mu_I}$	2.33
ρ_R	0.77	φ_p	61.32	$100\sigma_\eta$	0.45
γ_Π	1.18	$1 - a$	0.59	$100\sigma_Z$	0.56
γ_g	0.91	φ_w	4615.94	$100\sigma_R$	0.15
γ_x	0.20	$1 - a_w$	0.69		

We apply the particle filter $N_{sim} = 100$ using $M = 500,000$ particles. (In the actual estimation, we use $M = 1,500,000$.) The sampling distribution of standard BSPF and the adapted BSPF are shown

in Figure 5.1. Both estimators are fairly stable, with standard deviations less than one, the rule-of-thumb given in Pitt, Silva, Giordani, and Kohn (2012). However, the adapted BSPF has a standard deviation of 0.15 about six times less than the standard BSPF, indicating that there is substantial gain to adapting the particle filter during the Great Recession period.

Figure 5.1: Sampling Distribution of Log Likelihood Function Plus Log Prior



6 Particle Filter Metropolis-Hastings

In this section we describe the particle filter Metropolis-Hastings (PFMH) algorithm for generating a Markov Chain that converges to the posterior distribution of interest. The general algorithm to construct $\{\theta^i\}_{i=0}^N$ is given in Algorithm 2.

Algorithm 2 (PFMH Algorithm) For $i = 1$ to N :

1. Draw ϑ from a density $q(\vartheta|\theta^{i-1})$.
2. Set $\theta^i = \vartheta$ with probability

$$\alpha(\vartheta|\theta^{i-1}) = \min \left\{ 1, \frac{\hat{p}(Y|\vartheta)p(\vartheta)/q(\vartheta|\theta^{i-1})}{\hat{p}(Y|\theta^{i-1})p(\theta^{i-1})/q(\theta^{i-1}|\vartheta)} \right\}$$

and $\theta^i = \theta^{i-1}$ otherwise. The likelihood approximation $\hat{p}(Y|\vartheta)$ is computed using Algorithm 1.

As in Fernandez-Villaverde and Rubio-Ramirez (2007), we use the random walk variant of the PFMH algorithm, which amounts to:

$$q(\cdot|\theta^{i-1}) = N(\theta^{i-1}, c\hat{\Sigma}). \tag{6.1}$$

The key choices in this algorithm are the matrix $\hat{\Sigma}$ and scaling factor c . We set

$$\hat{\Sigma} = \text{diag}(\mathbb{V}[\{\theta^j\}_{j=1}^{N_{\text{tuning}}}]),$$

the estimated variance from a tuning run with $N_{tuning} = 5000$ ²⁷ and we set c to ensure a reasonable acceptance rate. Table 6.1 gives an overview of the hyperparameter choices we made for the PFMH algorithm. Setting $c = 0.20$ yields us an average acceptance rate of about 27 percent. To obtain 50,000 draws, the algorithm took about 10 days.

Table 6.1: PFMH details

Object	Description	Value
N	Length of Chain	50000
K	Number of Processors	336
M	Number of Particles in Adapted BSPF	1500000
$\hat{\Sigma}$	Proposal Variance	Tuning Run
c	Scaling Factor	0.2
Acceptance Rate		0.27
Run Time		10 days

Notes: We run 4 chains each of length 50000.

Table 6.2 reproduces Table 1 in the main text, while also providing the standard deviation of the posterior means across the 4 runs.

7 Computational Environment

We performed all computations at the High Performance Computing (HPC) Cluster maintained at the Federal Reserve Board. The project is coded in **Fortran**, and compiling using the **Intel Fortran Compiler** (version: 13.1.0 20130121), including the Math Kernel Library. The distributed aspects of the computation (i.e., parallelization) use **MPICH**, an implementation of the Message Passing Interfact standard—the servers are connected using **Infiniband**, a high-speed, low-latency connection.

²⁷The proposal variance for the tuning run came from an estimation the linearized version of the model.

Table 6.2: Posterior Distribution – Nonlinear Model

Parameter	Mean	[05, 95]	Parameter	Mean	[05, 95]
Steady State					
$100(\beta * *-1 - 1)$	0.14 (0.01)	[0.06, 0.23]	$100(\bar{\Pi} - 1)$	0.61 (0.01)	[0.54, 0.68]
$100 \log(G_z)$	0.50 (0.00)	[0.46, 0.54]	α	0.19 (0.00)	[0.16, 0.22]
Policy Rule					
ρ_R	0.70 (0.01)	[0.59, 0.78]	γ_{Π}	1.67 (0.04)	[1.21, 2.14]
γ_g	0.73 (0.02)	[0.39, 1.07]	γ_x	0.14 (0.01)	[0.07, 0.24]
Endogenous Propagation					
γ	0.70 (0.01)	[0.63, 0.76]	σ_L	2.00 (0.05)	[1.01, 3.17]
σ_a	5.32 (0.08)	[3.78, 7.09]	φ_I	3.70 (0.15)	[2.24, 5.21]
φ_p	100.41 (1.09)	[65.10, 136.88]	$1 - a$	0.56 (0.00)	[0.36, 0.76]
φ_w	4420.49 (170.02)	[1693.15, 8356.34]	$1 - a_w$	0.51 (0.02)	[0.29, 0.72]
Exogenous Processes					
ρ_G	0.67 (0.04)	[0.29, 0.96]	$100\sigma_G$	0.15 (0.00)	[0.11, 0.20]
ρ_{μ_I}	0.80 (0.02)	[0.64, 0.92]	$100\sigma_{\mu_I}$	2.39 (0.16)	[1.43, 3.70]
$100\sigma_{\eta}$	0.44 (0.02)	[0.34, 0.54]	$100\sigma_Z$	0.56 (0.02)	[0.38, 0.80]
$100\sigma_R$	0.18 (0.01)	[0.14, 0.24]			

Notes. Table reports the mean, fifth, and ninety-fifth percentile of the posterior distribution estimated by pooling 4 MCMC chains with 50,000 draws each (including 10,000 draw burn-in period.) The number in parentheses is the standard deviation of the mean across the 4 runs.