WORKING PAPERS
RESEARCH DEPARTMENT

WORKING PAPER NO. 12-10
NONLINEAR ADVENTURES AT THE ZERO LOWER BOUND

Jesús Fernández-Villaverde
University of Pennsylvania

Grey Gordon
University of Pennsylvania

Pablo Guerrón-Quintana
Federal Reserve Bank of Philadelphia

Juan F. Rubio-Ramírez
Duke University

May 4, 2012

RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA
Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/research-and-data/

Nonlinear Adventures at the Zero Lower Bound*

Jesús Fernández-Villaverde† Grey Gordon‡ Pablo Guerrón-Quintana§

Juan F. Rubio-Ramírez¶

May 4, 2012

Abstract

Motivated by the recent experience of the U.S. and the Eurozone, we describe the quantitative properties of a New Keynesian model with a zero lower bound (ZLB) on nominal interest rates, explicitly accounting for the nonlinearities that the bound brings. Besides showing how such a model can be efficiently computed, we find that the behavior of the economy is substantially affected by the presence of the ZLB. In particular, we document 1) the unconditional and conditional probabilities of hitting the ZLB; 2) the unconditional and conditional probabilty distributions of the duration of a spell at the ZLB; 3) the responses of output to government expenditure shocks at the ZLB, 4) the distribution of shocks that send the economy to the ZLB; and 5) the distribution of shocks that keep the economy at the ZLB.

Keywords: Zero lower bound, New Keynesian models, Nonlinear solution methods.

JEL classification numbers: E30, E50, E60.

[*]We thank Michael Woodford for useful comments. We also thank Larry Christiano for explaining to us several aspects about his work on the topic. Keith Kuester provided valuable comments on earlier drafts. Beyond the usual disclaimer, we must note that any views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Atlanta, the Federal Reserve Bank of Philadelphia, or the Federal Reserve System. Finally, we also thank the NSF for financial support. This paper is available free of charge at http://www.philadelphiafed.org/researchand-data/publications/working-papers/.

[†]University of Pennsylvania, NBER, CEPR, and FEDEA, <jesusfv@econ.upenn.edu>.

[‡]University of Pennsylvania, <greygordon@gmail.com>.

[§]Federal Reserve Bank of Philadelphia, <pablo.guerron@phil.frb.org>.

[¶]Duke University, Federal Reserve Bank of Atlanta, and FEDEA, <juan.rubio-ramirez@duke.edu>.

1. Introduction

How does the economy behave when it faces the zero lower bound (ZLB) on the nominal interest rate? To answer this basic question in macroeconomics, we build a New Keynesian model with the ZLB and solve it nonlinearly using a projection method and rational expectations. Besides showing how the model can be efficiently computed, this paper documents 1) the unconditional and conditional probabilities of hitting the ZLB; 2) the unconditional and conditional probability distributions of the duration of a spell at the ZLB; 3) the distributions of shocks that send the economy to the ZLB and the ones that keep it there; 4) the responses of output to government expenditure shocks at the ZLB; and 5) the impulse response functions of the economy at and away from the ZLB.

Our investigation is motivated by the recent U.S. and Eurozone experience of (nearly) zero short-term nominal interest rates after 2008. This experience has rekindled the interest in understanding the theoretical and quantitative effects of the ZLB in dynamic stochastic general equilibrium (DSGE) models. These effects are consequential for the conduct of policy as illustrated, for instance, by Eggertsson and Woodford (2003)—for monetary policy—and Christiano, Eichenbaum, and Rebelo (2011)—for fiscal policy.

However, the quantitative analysis of the ZLB is complicated by the essential nonlinearity that it generates. Along the equilibrium path of the model, the nominal interest rate is a function of the states of the economy and government policy. Since, in the typical model, the nominal interest rate will be zero for some states of the economy and government policies but not for others, the function that determines interest rates must be nonlinear. This nonlinearity is an obstacle both for traditional linearizations techniques and for the increasingly popular higher-order perturbations. These methods, by construction, cannot handle the ZLB in a model with rational expectations.

The literature has tried to get around this problem using different approaches. We highlight three of the most influential. Eggertsson and Woodford (2003) log-linearize the equilibrium relations except for the ZLB constraint, which is retained in its exact form. They then solve the corresponding system of linear expectational difference equations with the linear inequality constraint. To obtain numerical results, Eggertsson and Woodford assume a 2-state Markov chain with an absorbing state for an exogenous disturbance that drives the dynamics of the model.[1] Christiano, Eichenbaum, and Rebelo (2011) linearize the two models they use: a simple model without capital and a model with capital. In the first model, they assume that there is a 2-state Markov chain with an absorbing state for the discount factor.

[1] Eggertsson and Woodford also derive analytical results, such an optimal monetary policy, for more general disturbance processes.

The economy starts with a high discount factor that sends the nominal interest rate to zero and, in every period, there is a probability that the discount factor will move to its long run value and never change again. In the model with capital, the path for the discount factor (and hence, for the economy, conditional on some fiscal policy) is deterministic. Braun and Körber (2010) employ a variant of extended shooting, which solves the model forward for an exogenously fixed number of periods (after which the economy reverts to the steady state) given some initial conditions and assuming that in these periods, shocks are set to zero.

Although a big step toward understanding the ZLB, the existing solutions have made simplifying assumptions that may have unexpected implications. First, linearizing equilibrium conditions such as the Euler equations may hide some nonlinear interactions between the ZLB and the policy functions of the agents. Second, linear approximations provide a poor description of the economy during deep recessions such as the 2007-2009 one. Hence, our approach is a step forward to analyze the role of policy during contractions afflicted by the ZLB. Third, dynamics driven by exogenous variables that follow 2-state Markov chains with absorbing states or have deterministic paths imply the durations of spells at the ZLB have very simple expectations and variances. When we solve the model nonlinearly, however, these expectations and variances change substantially over time. This is key because the shape of the distribution of future events and the uncertainty it implies have material consequences for quantities important to policy analysis such as the size of the fiscal multiplier.

We address these nonlinear interactions and the time-varying distributions of future events by computing a fully nonlinear New Keynesian model with a ZLB and rational expectations. We incorporate monetary and fiscal policy and a set of shocks that push the model into the ZLB. In that way, the agents in the model have the right expectations about the probability of hitting the ZLB. To solve the model, we approximate the equilibrium functions of consumption, inflation, and one auxiliary variable that controls the pricing decision of firms as a nonlinear function of the states. More concretely, we employ Chebyshev polynomials as a basis for the projection of these endogenous variables on the states of the model. Since we are working in a high-dimensional space, the coefficients are solved for with the Smolyak collocation method laid out in Krueger, Kubler, and Malin (2011). This method significantly reduces the burden of the curse of dimensionality. However, to maintain tractability and provide easier insight, we will specify a relatively simple model with only five state variables. Once we have the nonlinear functions for consumption, inflation, and the one auxiliary variable, we exploit the rest of the equilibrium conditions of the model to solve for all other variables of interest. In particular, we find the nominal interest rate from the exact Taylor rule with a ZLB with its kink at zero holding exactly. This feature of our solution method is crucial because it allows us to circumvent the difficulty in directly approximating a non-differentiable function. Next, we simulate the model to calculate its moments, asking questions such as what is the effect of fiscal policy when the economy is at the ZLB. Our research finds inspiration in Wolman (1998), who solved a simpler sticky prices model using finite elements. The main difference is that our model considers a larger set of shocks and it is closer to the ones that have been taken successfully to the data and employed in applied policy analysis. Thus, we can offer quantitative assessments of issues such as the size of fiscal multipliers, which are at the heart of the current policy discussion.[2]

Our main findings are:

1. The dynamics of the economy near or at the ZLB are poorly approximated by a linear solution.

2. The economy spends 5.53 percent of the time at the ZLB (one of each 18 quarters), and the average duration of a spell at the ZLB is 2.06 quarters with a variance of 3.33. In our simulation, the duration of spells at the ZLB may last up to one decade.

3. The expected duration and the variance of the number of additional periods at the ZLB is time-varying. Furthermore, the number of expected periods and its variance grows with the number of periods already spent at the ZLB. This has been a point overlooked in the past, but it is important for the assessment of fiscal policy.

4. Positive shocks to the discount factor or to productivity raise the probability of the economy hitting the ZLB. Shocks on monetary or fiscal policy have little effect on this probability. While most previous literature has focused on shocks to the discount factor, we highlight the importance of *positive* shocks to productivity as a force pushing the economy to the ZLB.

5. The ZLB is associated with output, consumption, and inflation below their steady-state values; although, given the rich stochastic structure of the model, it is not always the case. In our simulation, we observe quarters with positive inflation, while the ZLB binds.

6. When the economy is hit by a discount factor shock that sends it to the ZLB for an average of 4 quarters, the multiplier of government expenditure at impact is around 3 times larger than when the economy is outside the bound.

[2] After the first draft of this paper was completed, we became aware of two recent attempts at solving a New Keynesian model with ZLB nonlinearly. Judd, Maliar, and Maliar (2011) use a cluster-grid algorithm that shares some similarities with our solution method. However, the main thrust of the paper is computational, and they solve only the New Keynesian model as a brief illustration of their method, without a complete analysis of its numerical properties. Tanaka (2012) uses a variation of the time iteration algorithm proposed by Coleman (1991), but he has only one state variable and focuses only on some aspects of the model.

7. Since the model is nonlinear, we need to distinguish between the marginal multiplier (the response of the economy to an infinitesimal innovation in government expenditure) and the average multiplier (the response of the economy to an innovation in government expenditure of size x; see, for the same distinction, Erceg and Lindé (2010)). The marginal multiplier is substantially larger than the average one.

8. When the economy is hit by a shock that sends it to the ZLB for an average of 2 quarters, the multiplier is only around twice as large as when the economy is outside the bound. Further, the multiplier is close to 1.

9. Increases on the variance of the additional number of periods at the ZLB lower the fiscal multiplier.

The rest of this paper is organized as follows. In section 2, we present a baseline New Keynesian model and we calibrate it to the U.S. economy. In section 3, we show how we compute the model nonlinearly. In section 4, we compare the solution that we obtain with the resulting one from a linearization. Section 5 presents our quantitative findings and section 6 concludes. A technical appendix offers some further details on the solution of the model.

2. The Model

Our investigation is built around a baseline New Keynesian model, which has often been used to frame the discussion of the effects of the ZLB. The structure of the economy is as follows. A representative household consumes, saves, and supplies labor. The final output is assembled by a final good producer who uses as inputs a continuum of intermediate goods manufactured by monopolistic competitors. The intermediate good producers rent labor from the household. Also, these intermediate good producers can only change prices following Calvo's rule. Finally, there is a government that fixes the one-period nominal interest rate, sets taxes, and consumes. We have four shocks: one to the time discount factor, one to technology, one to monetary policy, and one to fiscal policy.

2.1. Household

There is a representative household that maximizes a lifetime utility function separable in consumption c_t and hours worked l_t,

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \left(\prod_{i=0}^{t} \beta_i \right) \left\{ \log c_t - \psi \frac{l_t^{1+\vartheta}}{1+\vartheta} \right\},$$

where ϑ is the inverse of the Frisch labor supply elasticity and β_t is the discount factor. The discount factor follows a law of motion

$$\beta_{t+1} = \beta^{1-\rho_b} \beta_t^{\rho_b} \exp\left(\sigma_b \varepsilon_{b,t+1}\right) \text{ where } \varepsilon_{b,t+1} \sim \mathcal{N}(0,1).$$

Thus, the discount factor fluctuates around its mean β with a persistence ρ_b and innovations $\varepsilon_{b,t+1}$. We normalize $\beta_0 = \beta$.

The household trades Arrow securities (which we do not discuss further to save on notation) and a nominal government bond b_t that pays a nominal gross interest rate of R_t. Then, given a price level p_t, the household's budget constraint is

$$c_t + \frac{b_{t+1}}{p_t} = w_t l_t + R_{t-1}\frac{b_t}{p_t} + T_t + \digamma_t$$

where w_t is the real wage, T_t is a lump-sum transfer, and $\digamma_t$ are the profits of the firms in the economy. Finally, the household, as all agents in the economy, has rational expectations about the evolution of the variables in the model.

The first-order conditions of this problem are

$$\frac{1}{c_t} = \mathbb{E}_t \left\{ \beta_{t+1} \frac{1}{c_{t+1}} \frac{R_t}{\Pi_{t+1}} \right\}$$

and

$$\psi l_t^{\vartheta} c_t = w_t.$$

2.2. The final good producer

There is one final good y_t produced using intermediate goods y_{it} and the technology

$$y_t = \left(\int_0^1 y_{it}^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}} \tag{1}$$

where ε is the elasticity of substitution. Final good producers are perfectly competitive and maximize profits subject to the production function (1), taking as given all intermediate goods prices p_{ti} and the final good price p_t. The input demand functions associated with this problem are

$$y_{it} = \left(\frac{p_{it}}{p_t}\right)^{-\varepsilon} y_t \qquad \forall i,$$

and the price level is

$$p_t = \left(\int_0^1 p_{it}^{1-\varepsilon} di \right)^{\frac{1}{1-\varepsilon}}.$$

2.3. Intermediate good producers

Each intermediate firm produces differentiated goods using $y_{it} = A_t l_{it}$, where l_{it} is the amount of the labor input rented by the firm. Productivity A_t follows

$$A_t = A^{1-\rho_a} A_{t-1}^{\rho_a} \exp(\sigma_a \varepsilon_{a,t}) \text{ where } \varepsilon_{a,t} \sim \mathcal{N}(0,1)$$

where A is the average productivity level. Therefore, the real marginal cost of the intermediate good producer, common across all firms, is $mc_t = \frac{w_t}{A_t}$.

The monopolistic firms engage in infrequent price setting á la Calvo. In each period, a fraction $1-\theta$ of intermediate good producers reoptimize their prices p_{it}. All other firms keep their old prices. Thus, prices are set to solve

$$\max_{p_{it}} \mathbb{E}_t \sum_{\tau=0}^{\infty} \theta^{\tau} \left(\prod_{i=0}^{\tau} \beta_{t+i} \right) \frac{\lambda_{t+\tau}}{\lambda_t} \left(\frac{p_{it}}{p_{t+\tau}} - mc_{t+\tau} \right) y_{it+\tau}$$
$$\text{s.t. } y_{it} = \left(\frac{p_{it}}{p_t} \right)^{-\varepsilon} y_t$$

where λ_{t+s} is the Lagrangian multiplier for the household in period $t+s$.

The solution for the firm's pricing problem has a recursive structure in two auxiliary variables $x_{1,t}$ and $x_{2,t}$ that satisfy $\varepsilon x_{1,t} = (\varepsilon - 1) x_{2,t}$ and have laws of motion

$$x_{1,t} = \frac{1}{c_t} mc_t y_t + \theta \mathbb{E}_t \beta_{t+1} \Pi_{t+1}^{\varepsilon} x_{1,t+1}$$

and

$$x_{2,t} = \frac{1}{c_t} \Pi_t^* y_t + \theta \mathbb{E}_t \beta_{t+1} \Pi_{t+1}^{\varepsilon-1} \frac{\Pi_t^*}{\Pi_{t+1}^*} x_{2,t+1} = \Pi_t^* \left(\frac{1}{c_t} y_t + \theta \mathbb{E}_t \beta_{t+1} \frac{\Pi_{t+1}^{\varepsilon-1}}{\Pi_{t+1}^*} x_{2,t+1} \right)$$

in which $\Pi_t^* = \frac{p_t^*}{p_t}$ is the ratio between the optimal new price (common across all firms that reset their prices) and the price of the final good. Also, by the properties of Calvo pricing, inflation dispersion is given by $1 = \theta \Pi_t^{\varepsilon-1} + (1-\theta)(\Pi_t^*)^{1-\varepsilon}$.

2.4. The government

The government sets the nominal interest rates according to $R_t = \max[Z_t, 1]$ where

$$Z_t = R^{1-\rho_r} R_{t-1}^{\rho_r} \left[\left(\frac{\Pi_t}{\Pi} \right)^{\phi_\pi} \left(\frac{y_t}{y} \right)^{\phi_y} \right]^{1-\rho_r} m_t.$$

The variable Π represents the steady-state target level of inflation and R is the steady-state nominal gross return of bonds (equal to Π divided by β). The term m_t is a random shock to monetary policy that follows $m_t = \exp(\sigma_m \varepsilon_{m,t})$ with $\varepsilon_{m,t}$ distributed according to $\mathcal{N}(0,1)$.

This policy rule is the maximum of two terms. The first term, Z_t, follows a conventional Taylor rule that depends on the lagged interest rate, the deviation of inflation with respect to its target, and the output gap ($R^{1-\rho_r}$ appears on the right-hand side because we move R from the left-hand side of the equation). The second term is the ZLB: the gross nominal interest rate cannot be lower than 1.

Beyond the open market operations, the lump-sum transfers also finance a stream of government expenditures g_t with

$$g_t = s_{g,t} y_t \text{ and}$$
$$s_{g,t} = s_g^{1-\rho_g} s_{g,t-1}^{\rho_g} \exp(\sigma_g \varepsilon_{g,t}) \text{ where } \varepsilon_{g,t} \sim \mathcal{N}(0,1). \tag{2}$$

Because of Ricardian irrelevance, the timing of these transfers is irrelevant, so we set $b_t = 0$ period by period. The evolution of g_t is expressed in terms of its share of output, $s_{g,t}$.[3]

2.5. Aggregation

Aggregate demand is given by $y_t = c_t + g_t$. By well-known arguments, aggregate supply is

$$y_t = \frac{A_t}{v_t} l_t$$

where:

$$v_t = \int_0^1 \left(\frac{p_{it}}{p_t} \right)^{-\varepsilon} di$$

is the aggregate loss of efficiency induced by price dispersion of the intermediate goods. By the properties of Calvo pricing,

$$v_t = \theta \Pi_t^{\varepsilon} v_{t-1} + (1-\theta) (\Pi_t^*)^{-\varepsilon}.$$

2.6. Equilibrium

The definition of equilibrium in this model is standard. This equilibrium is given by the sequence $\{y_t, c_t, l_t, mc_t, x_{1,t}, x_{2,t}, w_t, \Pi_t, \Pi_t^*, v_t, R_t, Z_t, \beta_t, A_t, m_t, g_t, b_t, s_{g,t}\}_{t=0}^{\infty}$ determined by:

[3] While it could happen that $s_{g,t} > 1$, our calibration of s_g and σ_g is such that this will happen with a negligible probability, so we can ignore this possibility (which never occurs in our simulations).

- the first-order conditions of the household

$$\frac{1}{c_t} = \mathbb{E}_t \left\{ \frac{\beta_{t+1}}{c_{t+1}} \frac{R_t}{\Pi_{t+1}} \right\}$$
$$\psi l_t^{\vartheta} c_t = w_t$$

- profit maximization

$$mc_t = \frac{w_t}{A_t}$$
$$\varepsilon x_{1,t} = (\varepsilon - 1)\, x_{2,t}$$
$$x_{1,t} = \frac{1}{c_t} mc_t y_t + \theta \mathbb{E}_t \beta_{t+1} \Pi_{t+1}^{\varepsilon} x_{1,t+1}$$
$$x_{2,t} = \Pi_t^* \left(\frac{1}{c_t} y_t + \theta \mathbb{E}_t \beta_{t+1} \frac{\Pi_{t+1}^{\varepsilon-1}}{\Pi_{t+1}^*} x_{2,t+1} \right)$$

- government policy

$$R_t = \max\left[Z_t, 1\right]$$
$$Z_t = R^{1-\rho_r} R_{t-1}^{\rho_r} \left[\left(\frac{\Pi_t}{\Pi} \right)^{\phi_\pi} \left(\frac{y_t}{y} \right)^{\phi_y} \right]^{1-\rho_r} m_t$$
$$g_t = s_{g,t} y_t$$
$$b_t = 0$$

- inflation evolution and price dispersion

$$1 = \theta \Pi_t^{\varepsilon-1} + (1-\theta) \left(\Pi_t^*\right)^{1-\varepsilon}$$
$$v_t = \theta \Pi_t^{\varepsilon} v_{t-1} + (1-\theta) \left(\Pi_t^*\right)^{-\varepsilon}$$

- market clearing

$$y_t = c_t + g_t$$
$$y_t = \frac{A_t}{v_t} l_t$$

- and the stochastic processes

$$\beta_{t+1} = \beta^{1-\rho_b} \beta_t^{\rho_b} \exp\left(\sigma_b \varepsilon_{b,t+1}\right)$$
$$A_t = A^{1-\rho_a} A_{t-1}^{\rho_a} \exp\left(\sigma_a \varepsilon_{a,t}\right)$$
$$m_t = \exp(\sigma_m \varepsilon_{m,t})$$
$$s_{g,t} = s_g^{1-\rho_g} s_{g,t-1}^{\rho_g} \exp\left(\sigma_g \varepsilon_{g,t}\right).$$

The steady state of the model is given by $\{y, c, l, mc, x_1, x_2, w, \Pi, \Pi^*, v, R, Z, \beta, A, m, g, b, s_g\}$, where we have eliminated subindexes to denote a steady-state value of the variable. See the appendix for details.

2.7. Calibration

We calibrate the model to standard choices in the literature. We start with the preferences parameters. We set $\beta = 0.994$ to match a real interest rate at the steady state of roughly 2.5 percent on an annual basis, $\vartheta = 1$ to deliver a Frisch elasticity of 1 (in the range of the numbers reported when we consider both the intensive and the extensive margin of labor supply), and $\psi = 1$, a normalization of the hours worked that is nearly irrelevant for our results. As it is common in the New Keynesian literature, we set the Calvo parameter $\theta = 0.75$, which implies an average duration of the price of 4 quarters, and the elasticity of substitution among good varieties to $\varepsilon = 6$, which implies a steady-state markup of 20 percent (Christiano, Eichenbaum, and Evans, 2005, and Eichenbaum and Fisher, 2007).

The parameters in the Taylor rule are conventional: $\phi_\pi = 1.5$, $\phi_y = 0.25$, $\Pi = 1.005$, and, to save on the dimensionality of the problem, $\rho_r = 0$ (this is also the case in Christiano, Eichenbaum, and Rebelo, 2011). Then:

$$R_t = \max\left[\frac{R}{\Pi^{\phi_\pi} y^{\phi_y}} \Pi_t^{\phi_\pi} y_t^{\phi_y} m_t, 1\right].$$

For fiscal policy, $s_g = 0.2$ so that government expenditures in steady state account for 20 percent of output, which is close to the average of government consumption in the U.S.

We set, for the discount factor, $\rho_b = 0.8$ and $\sigma_b = 0.0025$. Thus, the preference shock has a half-life of roughly 3 quarters and an unconditional standard deviation of 0.42 percent. With these values, our economy hits the ZLB with a frequency consistent with values previously reported (see section 5 below for details). For the technology process, we set $A = 1$, $\rho_a = 0.9$, and $\sigma_a = 0.0025$. These numbers are lower than those used in Cooley and Prescott (1995) to reflect the lower volatility of productivity in the last two decades. The volatility of the monetary shock is set to $\sigma_m = 0.0025$, which is in the range of values reported in Guerrón-Quintana (2010). Finally, for the government expenditure shock, we have somewhat smaller values than Christiano, Eichenbaum, and Rebelo (2011) by setting $\rho_g = 0.8$ and $\sigma_g = 0.0025$. This last value is half of that estimated in Justiniano and Primiceri (2008). We pick those numbers to avoid numerical problems associated with very persistent fiscal shocks when we hit the economy with a large fiscal expansion. Since we will document that fiscal shocks are less important than preference shocks in sending the economy to the ZLB, this lower persistence is not terribly important.

3. Solution of the Model

Given the previous calibration, our model has five state variables: price dispersion, v_{t-1}; the time-varying discount factor, β_t; productivity, A_t; the monetary shock, m_t; and the government expenditure share, $s_{g,t}$. Then, we define the vector of state variables:

$$\mathbb{S}_t = (\mathbb{S}_{1,t}, \mathbb{S}_{2,t}, \mathbb{S}_{3,t}, \mathbb{S}_{4,t}, \mathbb{S}_{5,t}) = (v_{t-1}, \beta_t, A_t, m_t, s_{g,t}).$$

Note that we have one endogenous state variable and four exogenous ones. For convenience, we also define the vector evaluated at the steady state as $\mathbb{S}_{ss} = (v, \beta, 1, 1, s_g)$.

Solving nonlinear models with five state variables that are not Pareto efficient is challenging. For instance, it is computationally expensive to apply value function iteration techniques or any projection method using the tensor product of one-dimensional polynomials. Instead, we use a projection method—Smolyak's algorithm—in which the number of grid points and the associated number of terms of the approximating polynomial do not grow exponentially in the dimensionality of the problem. This projection method judiciously chooses grid points and polynomials so that their number coincide and so that the function to be approximated is interpolated at the grid points. Although the algorithm is described in detail in the appendix, its basic structure is as follows.

Let the equilibrium functions for c_t, Π_t, and $x_{1,t}$ be written as functions of the states by

$$\begin{aligned} c_t &= f^1(\mathbb{S}_t) \\ \Pi_t &= f^2(\mathbb{S}_t) \\ x_{1,t} &= f^3(\hat{\mathbb{S}}_t) \end{aligned}$$

where $f^i : \mathbb{R}^5 \to \mathbb{R}$ for $i \in \{1, 2, 3\}$ and let $f = (f^1, f^2, f^3)$. If we had access to f, we could find the values of the remaining endogenous variables using the equilibrium conditions. Since

f is unknown, we approximate c_t, Π_t, and $x_{1,t}$ as

$$\begin{aligned} c_t &= \widehat{f}^1(\mathbb{S}_t) \\ \Pi_t &= \widehat{f}^2(\mathbb{S}_t) \\ x_{1,t} &= \widehat{f}^3(\mathbb{S}_t) \end{aligned}$$

where $\widehat{f} = \left(\widehat{f}^1, \widehat{f}^2, \widehat{f}^3\right)$ are polynomials.

To compute $\widehat{f}$, we first define a hypercube for the state variables (we extrapolate when we need to move outside the hypercube in the simulations) and rely on Smolyak's algorithm to obtain collocation points (grid points) within the hypercube. We guess the values of $\widehat{f}$ at those collocation points, which implicitly define $\widehat{f}$ over the entire state space. Then, treating $\widehat{f}$ as the true time $t+1$ functions, we exploit the equilibrium conditions to back out the current c_t, Π_t, and $x_{1,t}$ values at the collocation points and check that those values coincide with the ones implied by $\widehat{f}$. If not, we update our guess until convergence. In our application of Smolyak's algorithm, we use polynomials of up to degree 4.

Once we have found $\widehat{f}$, from the equilibrium conditions we get an expression for the second auxiliary variable,

$$x_{2,t} = \frac{\varepsilon}{\varepsilon - 1} x_{1,t},$$

output,

$$y_t = c_t + g_t = c_t + s_g s_{g,t} y_t = \frac{1}{1 - s_g s_{g,t}} c_t,$$

and for reset prices,

$$\Pi_t^* = \left(\frac{1 - \theta \Pi_t^{\varepsilon - 1}}{1 - \theta}\right)^{\frac{1}{1-\varepsilon}}.$$

Then, we get the evolution of price dispersion,

$$v_t = \theta \Pi_t^{\varepsilon} v_{t-1} + (1 - \theta) \left(\Pi_t^*\right)^{-\varepsilon},$$

labor,

$$l_t = \frac{v_t}{A_t} y_t,$$

wages,

$$w_t = \psi l_t^{\vartheta} c_t,$$

and marginal costs,

$$mc_t = \frac{w_t}{A_t}.$$

Now we are ready to determine the interest rate:

$$R_t = \max\left[\frac{R}{\Pi^{\phi_\pi} y^{\phi_y}} \Pi_t^{\phi_\pi} y_t^{\phi_y} m_t, 1\right].$$

Thus, with this procedure, we get a set of functions $h_t = h(\mathbb{S}_t)$ for the additional variables in our model $h_t \in \{y_t, l_t, mc_t, x_{2,t}, w_t, \Pi_t^*, v_t, R_t\}$.

An important advantage of our procedure is that we solve nonlinearly for c_t, Π_t, and $x_{1,t}$ and only later for the other variables exploiting the remaining equilibrium conditions. Hence, conditional on c_t, Π_t, and $x_{1,t}$, our method deals with the kink of R_t at 1 in the Taylor rule without any approximation: R_t comes from a direct application of the Taylor rule. Also, the functions for c_t, Π_t, and $x_{1,t}$ are derived from equations that involve expectations, which smooth out any possible kinks created by the ZLB.

4. Comparison between Log-linear and the Nonlinear Solution

Equipped with our nonlinear solution, we can address many important questions. In this section, we start by comparing the decision rules that we just obtained with the ones resulting from a log-linear approximation when we ignore the ZLB, only imposing it ex-post in the simulation (that is, when the Taylor rule would like to set $R_t < 1$, we instead set $R_t = 1$ and continue). We illustrate, in this way, the important differences that the presence of the ZLB generates in the dynamics of the model and motivate the need for a full nonlinear approach.

Our first step is to compute the log-linearized decision rules for our model given our calibration. For example, the log-linear decision rule for consumption is given by:

$$\log c_t - \log c = \alpha_1 (\log v_{t-1} - \log v) + \alpha_2 (\log \beta_t - \log \beta) + \alpha_3 \log A_t + \alpha_4 \log m_t + \alpha_5 (\log s_{g,t} - \log s_g).$$

To compare this decision rule with our nonlinear solution, we plot in figure 4.1 both decision rules along the $\log \beta_t$-axis, while we keep the other state variables at their steady-state values. The decision rule has an overall negative slope because a value of the discount factor above its steady state induces the household to save more, which lowers the real interest rate. Hence, if sufficiently large and positive innovations buffet the discount factor, the economy may be pushed to the ZLB. We observe this urgency to save more and consume less in the nonlinear policy. Consumption plunges for shocks that push the discount factor 0.7 percent or higher than its steady-state value. This is not a surprise. We calibrated β to be 0.994. Then, when the discount factor increases 0.6 percent (0.006 in our graph), β_t becomes bigger than 1 and future consumption yields more utility than current expenditure. Our projection method is perfectly capable to respond to that feature of the model and capture the change of optimal consumption.[4] In contrast, the log-linear approximation, since it is built around the value of β_t in the steady state, predicts a stable decline in consumption. This difference becomes stronger as β_t moves into higher values. If we expanded the horizontal scale, we would observe a bigger difference between the two decision rules. At the same time, the two policies almost coincide for values of the discount factor close to or below its steady state value, where the log-linear approximation is valid.[5]

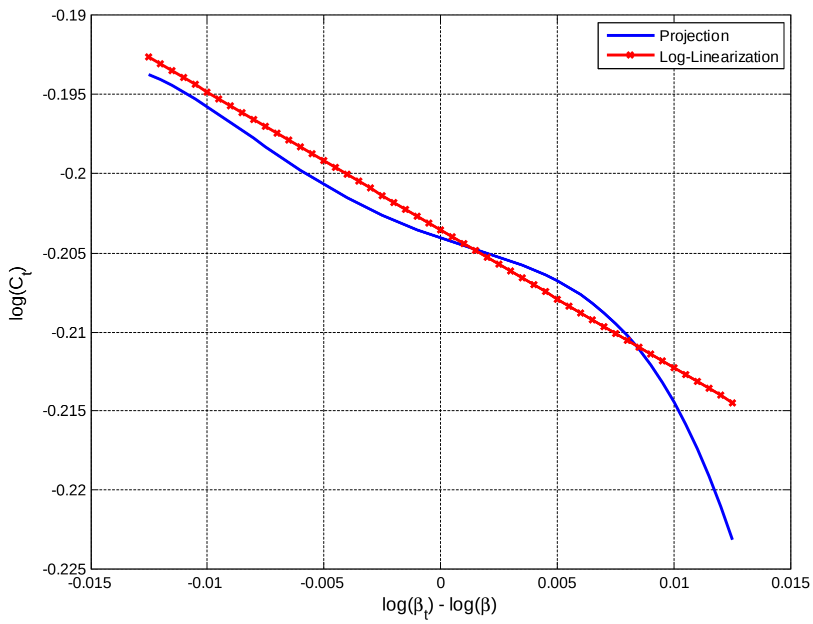

Figure 4.1: Log-linear versus nonlinear decision rules for consumption

The policy function for inflation (figure 4.2) confirms the same pattern as in the previous plot. As the discount factor becomes bigger, there is less aggregate demand today, and this pushes inflation down as those firms that reset their prices do not raise them as much as they would otherwise do. In the nonlinear solution, this phenomenon of contraction-deflation becomes acuter as β increases and the economy moves toward the ZLB.

In conclusion, nonlinearities matter when we talk about the ZLB. This point will be even clearer with the findings of the next section.

[4] A technical point: the projection method does not deliver a decision rule that is globally concave. Since we have price rigidities, there is no theoretical result that implies that this should be the case.

[5] There is also a small difference in levels dues to precautionary behavior, which is not captured by the log-linear approximation.

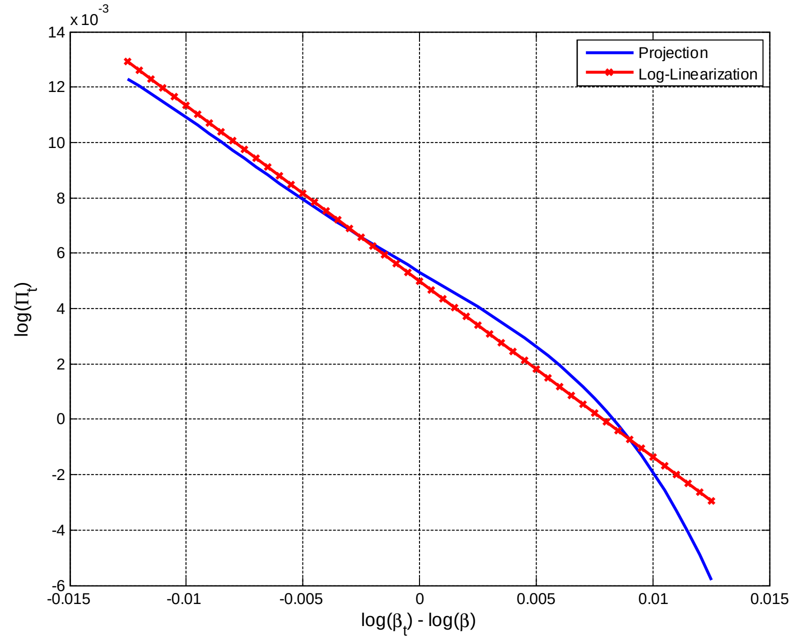

Figure 4.2: Log-linear versus nonlinear equilibrium functions for inflation

5. Results

We now present the quantitative results from the model. First, we report the probabilities of being at the ZLB. Second, we analyze the distribution of the duration of a spell at the ZLB. Third, we show the distribution of shocks that send and keep the economy at the ZLB. Fourth, we analyze the behavior of the endogenous variables at the ZLB. Fifth, we analyze in detail one spell at the ZLB. Sixth, we analyze the size of the fiscal multipliers. We close by talking about the impulse-response functions (IRFs) of the model.

5.1. Probability of being at the ZLB

Our first exercise is to derive the probability that the economy is at the ZLB in any particular period. This step is important because, in our model, the entry and exit from the ZLB is endogenous. The spells are caused by the interaction of the exogenous state variables (which have a continuous distribution over many values) and the endogenous state. Thus, we compute the integral

$$\Pr\left(\{R_t = 1\}\right) = \int \mathcal{I}_{\{1\}}\left[R\left(\mathbb{S}_t\right)\right]\mu\left(\mathbb{S}_t\right)d\mathbb{S}_t \tag{3}$$

for all $t \geq 1$, where $\mathcal{I}_B [f(x)] : \mathbb{R}^n \rightarrow \{0,1\}$, is the indicator function of a set $B \subset \mathbb{R}^k$ and $f : \mathbb{R}^n \rightarrow \mathbb{R}^k$, defined as

$$\mathcal{I}_B [f(x)] = \begin{cases} 1 \text{ if } f(x) \in B \\ 0 \text{ if } f(x) \notin B \end{cases}$$

for any n, and where $\mu(\mathbb{S}_t)$ is the unconditional distribution of the states of the economy. We further define $\mathcal{I}_B [x]$ as $\mathcal{I}_B [\iota(x)]$, where ι is the identity map.

Given our solution from section 3, we simulate the model for $T = 300,000$ periods starting at $\mathbb{S}_{ss}$. We define the sequence of simulated states as $\{\mathbb{S}_{i,ss}\}_{i=1}^{T}$, where $\mathbb{S}_{i,ss}$ is the value of the state vector in period i of the simulation when the initial condition is $\mathbb{S}_{ss}$.[6] Using this sequence, we follow Santos and Peralta-Alva (2005) to approximate equation (3) by

$$\Pr(\{R_t = 1\}) \simeq \widehat{\Pr}(\{R_t = 1\}) = \frac{\sum_{i=1}^{T} \mathcal{I}_{\{1\}} [R(\mathbb{S}_{i,ss})]}{T}. \tag{4}$$

From equation (4) we learn that our economy is at the ZLB during 5.53 percent of quarters, or approximately 1 out of each 18 quarters, a noticeable share. However, since this percentage was a criterion for the calibration, we do not emphasize it.[7]

The probability of being at the ZLB is not independent of the history of states. To show this, we calculate the probability of being at least one extra period at the ZLB conditional on having been at the ZLB for exactly s periods as

$$\Pr(\{R_t = 1\} \mid \{R_{t-1} = 1, \ldots, R_{t-s} = 1, R_{t-s-1} > 1\}) =$$
$$\int \mathcal{I}_{\{1\}} [R(\mathbb{S}_t)] \, \mu(\mathbb{S}_t | \{R_{t-1} = 1, \ldots, R_{t-s} = 1, R_{t-s-1} > 1\}) \, d\mathbb{S}_t$$

for all $t \geq 3$ and $t - 2 \geq s \geq 1$.

If we let $\mathbb{R}_{t-1,s} = \{R_{t-1} = 1, \ldots, R_{t-s} = 1, R_{t-s-1} > 1\}$ for all $t \geq 3$ and $t - 2 \geq s \geq 1$, then the above expression can be rewritten as

$$\Pr(\{R_t = 1\} | \mathbb{R}_{t-1,s}) = \int \mathcal{I}_{\{1\}} [R(\mathbb{S}_t)] \, \mu(\mathbb{S}_t | \mathbb{R}_{t-1,s}) \, d\mathbb{S}_t \tag{5}$$

[6] We have only one endogenous state variable: price dispersion, v_{t-1}. The mean of its ergodic distribution and its steady-state value are nearly identical. The means of the ergodic distribution and the steady-state values of the four exogenous state variables are, by construction, the same. Hence, we do not need a burn-in period at the start of our simulation.

[7] This percentage of quarters is close to the findings of Reifschneider and Williams (1999), who use the FRB/US model to argue that the U.S. economy will spend 5 percent of the quarters at the ZLB (defined somewhat more loosely by them as a FFR below 25 basis points at an annual basis) when the monetary authority has the same inflation target as in our calibration. In more recent work, Chung *et al.* (2010) have recomputed this number with the latest observations to be 6 percent.

for all $t \geq 3$ and $t-2 \geq s \geq 1$ and where $\mu(\mathbb{S}_t|\mathbb{R}_{t-1,s})$ is the distribution of the states conditional on $\mathbb{R}_{t-1,s}$ implied by our model.

For $s \geq 1$, let $\{\mathbb{S}_{i_{k_s},ss}\}_{k_s=1}^{M_s}$ be the M_s elements of the sequence $\{\mathbb{S}_{i,ss}\}_{i=1}^{T}$ such that $R(\mathbb{S}_{i_{k_s}-1,ss}) = 1, \ldots, R(\mathbb{S}_{i_{k_s}-s,ss}) = 1$, and $R(\mathbb{S}_{i_{k_s}-s-1,ss}) > 1$ for all $k_s \in \{1, \ldots, M_s\}$. We then approximate (5) by

$$\Pr(\{R_t = 1\}|\mathbb{R}_{t-1,s}) \simeq \widehat{\Pr}(\{R_t = 1\}|\mathbb{R}_{t-1,s}) = \sum_{k_s=1}^{M_s} \frac{\mathcal{I}_{\{1\}}[R(\mathbb{S}_{i_{k_s},ss})]}{M_s}. \tag{6}$$

for $M_s > 0$. When M_s is 0, we say $\widehat{\Pr}(\{R_t = 1\}|\mathbb{R}_{t-1,s}) = 0$.

Table: 5.1 Probability of being at least one extra period at the ZLB

	1	2	3	4	5	6	7	8	9	10
$\widehat{\Pr}(\{R_t = 1\}\|\mathbb{R}_{t-1,s})$	44%	54%	60%	62%	65%	60%	65%	67%	65%	60%

Table 5.1 reports the values of equation (6) for s from 1 to 10. The probability of being at least one extra period at the ZLB increases from 44 percent conditional on having been at the ZLB for 1 period to 67 percent for 8 periods and decreases thereafter.[8] That is, during the first two years of a spell at the ZLB, the longer we have been at the ZLB, the more likely it is we stay there next period. The intuition is simple: if we have been at the ZLB for, say, 4 quarters, it is because the economy has suffered a string of bad shocks. Thus, the economy will linger longer on average around the state values that led to the ZLB in the first place. The fact that the probabilities in table 5.1 are not constant over time already points out a feature of our model that will play a key role in the next subsections: the expectation and variance of the time remaining at the ZLB are time-varying.

5.2. How long are the spells at the ZLB?

Another important question is to determine the length of the spells of the economy at the ZLB. To answer that question, first, we compute the expected number of consecutive periods that the economy will be at the ZLB right after entering it. In mathematical terms:

$$\begin{aligned}
\mathbb{E}(\mathbb{T}_t|\mathbb{R}_{t-1,1}) &= \sum_{j=1}^{\infty} (j \Pr(\{\mathbb{T}_t = j\}|\mathbb{R}_{t-1,1})) \\
&= \sum_{j=1}^{\infty} \left(j \left(\int \mathcal{I}_{\{j\}}[\mathbb{T}_t]\, \mu(\mathbb{S}_t, \mathbb{S}_{t+1}, \ldots |\mathbb{R}_{t-1,1})\, d\mathbb{S}_t d\mathbb{S}_{t+1} \ldots \right)\right)
\end{aligned} \tag{7}$$

[8] We have few simulations with spells at the ZLB of durations 9 or longer, so the numbers at the end of the table are subject to more numerical noise than those at the start.

for all $t \geq 3$, where $\mathbb{R}_{t-1,1}$ describes the event "just entering the ZLB", $\mu(\mathbb{S}_t, \mathbb{S}_{t+1}, \ldots | \mathbb{R}_{t-1,1})$ is the distribution of the sequence of future states conditional on $\mathbb{R}_{t-1,s}$ implied by our model, and $\mathbb{T}_t = \mathbb{T}(\mathbb{S}_t, \mathbb{S}_{t+1}, \ldots) = \min\{\omega : \omega \in \{1, 2, \ldots\} \text{ and } R_{t+\omega-1} = R(\mathbb{S}_{t+\omega-1}) > 1\}$ for all $t \geq 3$ is the first period outside the ZLB as a function of the future path of state variables. If $\mathbb{T}_t = 1$, then $R_t > 1$ for that path. In the same way, if $\mathbb{T}_t = 2$, then $R_t = 1$ and $R_{t+1} > 1$ for that path of state variables.

To approximate the expectation in (7), take $\left\{\mathbb{S}_{i_{k_1},ss}\right\}_{k_1=1}^{M_1}$ and calculate

$$\mathbb{E}(\mathbb{T}_t|\mathbb{R}_{t-1,1}) \simeq \widehat{\mathbb{E}}(\mathbb{T}_t|\mathbb{R}_{t-1,1}) = \sum_{j=1}^{T-2}\left(j\left(\sum_{k_1=1}^{M_1}\frac{\mathcal{I}_{\{j\}}\left[\widehat{\mathbb{T}}\left(\mathbb{S}_{i_{k_1}+1,ss}, \ldots, \mathbb{S}_{T,ss}\right)\right]}{M_1}\right)\right) \tag{8}$$

where we approximate the probability

$$\Pr(\{\mathbb{T}_t = j\}|\mathbb{R}_{t-1,1}) = \int \mathcal{I}_{\{j\}}[\mathbb{T}_t]\,\mu(\mathbb{S}_t, \mathbb{S}_{t+1}, \ldots|\mathbb{R}_{t-1,1})\, d\mathbb{S}_t d\mathbb{S}_{t+1}\ldots$$

for all $t \geq 3$ and $T - 2 \geq j \geq 1$, by

$$\Pr(\{\mathbb{T}_t = j\}|\mathbb{R}_{t-1,1}) \simeq \widehat{\Pr}\left(\left\{\widehat{\mathbb{T}}_t = j\right\}|\mathbb{R}_{t-1,1}\right) = \sum_{k_1=1}^{M_1}\frac{\mathcal{I}_{\{j\}}\left[\widehat{\mathbb{T}}\left(\mathbb{S}_{i_{k_1},ss}, \ldots, \mathbb{S}_{T,ss}\right)\right]}{M_1}$$

where $\widehat{\mathbb{T}}_t = \widehat{\mathbb{T}}(\mathbb{S}_t, \ldots, \mathbb{S}_T) = \min\{\omega : \omega \in \{1, 2, \ldots, T - t + 1\} \text{ and } R_{t+\omega-1} = R(\mathbb{S}_{t+\omega-1}) > 1\}$ for all $t \geq 3$. The expression (8) equals 2.06 quarters in our simulation. That is, on average, we are at the ZLB for slightly more than half a year after entering it.

We highlight two points. First, given our simulation of T periods, we can be at the ZLB at most $T - 2$ (we need a period before entry and one period of exit of the ZLB). Second, we approximate $\mathbb{T}\left(\mathbb{S}_{i_{k_1},ss}, \mathbb{S}_{i_{k_1}+1,ss}, \ldots\right)$ by $\widehat{\mathbb{T}}\left(\mathbb{S}_{i_{k_1},ss}, \ldots, \mathbb{S}_{T,ss}\right)$. But this approximation has no effect as long as $R(\mathbb{S}_{T,ss}) > 1$, similar to what happens in our simulation.

We can also compute the variance of the number of consecutive periods at the ZLB right after entering it. Again, in mathematical terms:

$$\mathrm{Var}(\mathbb{T}_t|\mathbb{R}_{t-1,1}) = \sum_{j=1}^{\infty}\left((j - \mathbb{E}(\mathbb{T}_t|\mathbb{R}_{t-1,1})^2 \Pr(\{\mathbb{T}_t = j\}|\mathbb{R}_{t-1,1})\right)$$

$$= \sum_{j=1}^{\infty}\left((j - \mathbb{E}(\mathbb{T}_t|\mathbb{R}_{t-1,1})^2\left(\int \mathcal{I}_{\{j\}}[\mathbb{T}_t]\,\mu(\mathbb{S}_t, \mathbb{S}_{t+1}, \ldots|\mathbb{R}_{t-1,1})\, d\mathbb{S}_t d\mathbb{S}_{t+1}\ldots\right)\right)$$

for all $t \geq 3$, that we approximate by:

$$\text{Var}\left(\mathbb{T}_t|\mathbb{R}_{t-1,1}\right) \simeq \widehat{\text{Var}}\left(\mathbb{T}_t|\mathbb{R}_{t-1,1}\right) = \sum_{j=1}^{T-2}\left(\left(j - \widehat{\mathbb{E}}\left(\mathbb{T}_t|\mathbb{R}_{t-1,1}\right)\right)^2 \sum_{k_1=1}^{M_1} \frac{\mathcal{I}_{\{j\}}\left[\widehat{\mathbb{T}}\left(\mathbb{S}_{i_{k_1},ss},\ldots,\mathbb{S}_{T,ss}\right)\right]}{M_1}\right).$$

By applying the previous steps, we find that the variance is 3.33 quarters. Since the number of consecutive periods at the ZLB is bounded below by 1 quarter (one period in our model), a variance of 3.33 (standard deviation of 1.8 quarters) denotes a quite asymmetric distribution of consecutive periods at the ZLB just after entering it, with a long right tail. We revisit this issue below.

Now, we would like to see how the conditional expectations and variance depend on the number of periods already at the ZLB. First, we calculate the expected number of additional periods at the ZLB conditional on having already been at the ZLB for s periods:

$$\mathbb{E}\left(\mathbb{T}_t - 1|\mathbb{R}_{t-1,s}\right) = \sum_{j=1}^{\infty}\left((j-1)\Pr\left(\{\mathbb{T}_t = j\}|\mathbb{R}_{t-1,s}\right)\right) =$$
$$\sum_{j=1}^{\infty}\left((j-1)\left(\int \mathcal{I}_{\{j\}}\left[\mathbb{T}_t\right]\mu\left(\mathbb{S}_t,\mathbb{S}_{t+1},\ldots|\mathbb{R}_{t-1,s}\right)d\mathbb{S}_t d\mathbb{S}_{t+1}\ldots\right)\right)$$

for all $t \geq 3$ and $t-2 \geq s \geq 1$, where we are subtracting 1 from $\mathbb{T}_t$ because we are calculating the number of additional periods at the ZLB. This expression can approximated by:

$$\widehat{\mathbb{E}}(\mathbb{T}_t - 1|\mathbb{R}_{t-1,s}) = \sum_{j=1}^{T-s-1}\left((j-1)\sum_{k_s=1}^{M_s}\frac{\mathcal{I}_{\{j\}}\left[\widehat{\mathbb{T}}\left(\mathbb{S}_{i_{k_s},ss},\ldots,\mathbb{S}_{T,ss}\right)\right]}{M_s}\right).$$

We build the analogous object for the conditional variance:

$$\text{Var}\left(\mathbb{T}_t - 1|\mathbb{R}_{t-1,s}\right) = \sum_{j=1}^{\infty}\left(j - \mathbb{E}\left(\mathbb{T}_t - s|\mathbb{R}_{t-1,s}\right)\right)^2 \Pr\left(\{\mathbb{T}_t = j\}|\mathbb{R}_{t-1,s}\right) =$$
$$\sum_{j=1}^{\infty}\left(\left(j - \mathbb{E}\left(\mathbb{T}_t|\mathbb{R}_{t-1,s}\right)\right)^2\left(\int \mathcal{I}_{\{j\}}\left[\mathbb{T}_t\right]\mu\left(\mathbb{S}_t,\mathbb{S}_{t+1},\ldots|\mathbb{R}_{t-1,s}\right)d\mathbb{S}_t d\mathbb{S}_{t+1}\ldots\right)\right).$$

approximated by

$$\widehat{\text{Var}}\left(\mathbb{T}_t|\mathbb{R}_{t-1,s}\right) = \sum_{j=1}^{T-s-1}\left(\left(j - \widehat{\mathbb{E}}\left(\mathbb{T}_t|\mathbb{R}_{t-1,s}\right)\right)^2 \sum_{k_s=1}^{M_s}\frac{\mathcal{I}_{\{j\}}\left[\widehat{\mathbb{T}}\left(\mathbb{S}_{i_{k_s},ss},\ldots,\mathbb{S}_{T,ss}\right)\right]}{M_s}\right).$$

Recall that, given our simulation of T periods, we can be at the ZLB at most $T - s - 1$ additional periods (we need s periods before entry and one period of exit of the ZLB).

These two statistics are reported in table 5.2. The expected number of additional periods increases with the number of periods already at the ZLB (at least until 8 quarters). This is not surprising after having seen table 5.1. The variance grows monotonically, from 2.70 to 5.57. Hence, the agents have more uncertainty forecasting the periods remaining at the ZLB as the time at the ZLB accumulates (at least for the first 10 quarters).

Table 5.2: Number of additional periods at the ZLB

s	1	2	3	4	5	6	7	8	9	10
$\widehat{\mathbb{E}}(\mathbb{T}_t - 1 \vert \mathbb{R}_{t-1,s})$	1.06	1.41	1.62	1.72	1.79	1.75	1.91	1.92	1.86	1.86
$\widehat{\text{Var}}\,(\mathbb{T}_t - 1 \vert \mathbb{R}_{t-1,s})$	2.70	3.39	3.78	4.00	4.22	4.35	4.64	4.87	5.16	5.57

The results in table 5.2 reflect how the distribution of additional periods at the ZLB is time-varying. This is an important difference of our paper with respect to the literature. Previous papers have considered two cases. In the first one, there is an exogenous variable that switches in every period with some constant probability; the typical change is the discount factor going from low to high. Once the variable has switched, there is a perfect foresight path that exists from the ZLB (see, for instance, Eggertsson and Woodford, 2003). Although this path may depend on how long the economy has been at the ZLB, there is no uncertainty left once the variable has switched. In the second case, researchers have looked at models with perfect foresight (see, for example, sections IV and V of Christiano, Eichenbaum, and Rebelo, 2011). Even if fiscal policy influences how long we are at the ZLB, there is no uncertainty with respect to events, and the variance is degenerate.

The number of additional periods at the ZLB and the beliefs that agents have about them will play a key role later when we talk about the fiscal multiplier. The multiplier will depend on the uncertainty regarding how many more periods the economy will be at the ZLB. Having more uncertainty will, in general, raise the multiplier. However, the result cannot be ascertained in general because of the skewness of the distribution. By construction, we cannot have less than 0 additional periods at the ZLB. Thus, if the expectation is, for instance, 1.79 quarters and the variance 4.22 (column $s = 5$ in table 5.2), we must have a large right tail and a concentrated mass on the (truncated) left tail. A higher variance can be caused by movements in different parts of this asymmetric distribution and, thus, has complex effects on the fiscal multiplier (see section 5.7 for details).

5.3. What shocks do take us to the ZLB?

We would like to know whether a shock of a given size increases or decreases the probability of hitting the ZLB in the next I periods. More concretely, we would like to compute:

$$\Pr\left(\cup_{i=1}^{I}\left\{R_{t+i}=1\right\} \mid \left\{\mathbb{S}_{j,t}=A_j\right\}\right) =$$
$$\int \max\left\{\mathcal{I}_{\{1\}}\left[R\left(\mathbb{S}_{t+1}\right)\right],\ldots,\mathcal{I}_{\{1\}}\left[R\left(\mathbb{S}_{t+I}\right)\right]\right\}\mu\left(\mathbb{S}_{t+1},\ldots,\mathbb{S}_{t+I}\right|\left\{\mathbb{S}_{j,t}=A_j\right\}\right)d\mathbb{S}_{t+1}\ldots d\mathbb{S}_{t+I} \quad (9)$$

for $A_j \in \mathbb{R}$ and where $\mathbb{S}_{j,t}$ is the j^{th} element of $\mathbb{S}_t$. We perform this analysis only for the exogenous states variables, that is $j \in \{2,\ldots,5\}$.

To approximate (9) given our solution, we simulate the model $N = 10,000$ times for I periods starting at $\mathbb{S}_{ss}$ for all states, but for the j^{th} element of $\mathbb{S}_t$, which we start at A_j. For each simulation $n \in \{1,\ldots,N\}$, we call the I periods long sequence of states $\left\{\mathbb{S}_{n,i,ss}^{A_j}\right\}_{i=1}^{I}$. Then, we set $I = 10$ and approximate the probability by:

$$\widehat{\Pr}\left(\cup_{i=1}^{I}\left\{R_{t+i}=1\right\} \mid \left\{\mathbb{S}_{j,t}=A_j\right\}\right) = \sum_{n=1}^{N}\frac{\max\left(\mathcal{I}_{\{1\}}\left[R\left(\mathbb{S}_{n,1,ss}^{A_j}\right)\right],\ldots,\mathcal{I}_{\{1\}}\left[R\left(\mathbb{S}_{n,I,ss}^{A_j}\right)\right]\right)}{N}.$$

Table 5.3 reports these probabilities when the initial exogenous states are set, one at a time (corresponding to each column), to 0, ± 1, and ± 2 standard deviation innovations away from the steady state. In other words, $A_j = \mathbb{S}_{j,ss} \pm \{0,1,2\}\,\sigma_j$ for the discount factor, technology, monetary, and government expenditure. The third row (labeled 0 std) shows a baseline scenario in which all the exogenous states are set at their steady-state values. In this baseline, the probability of getting to the ZLB in the next two and a half years is 15 percent (all the entries in the row are equal, since we are describing the same event). When we have a one-standard-deviation positive innovation to the discount factor, that probability goes up to 20 percent, and when we have a two-standard-deviation innovation to 29 percent. When β_t is higher than in the steady state, the household is more patient and the interest rate that clears the good market is low. Hence, it is easy for the economy to be pushed into the ZLB. The reverse results occur when the innovation is negative: the probability of entering into a ZLB falls to 12 percent (one-standard-deviation) or 9 percent (two-standard deviation).

Interestingly, we also find that positive productivity shocks raise the probability of being at the ZLB. There are two mechanisms at work here. First, higher productivity means lower inflation: firms charge a mark-up over their marginal cost and, with higher productivity, costs are going down. Since the monetary authority responds to lower inflation by lowering the nominal interest rate even more (the coefficient in the Taylor rule is $\phi_\pi = 1.5$), this puts the economy closer to the ZLB. Second, when productivity is high, the real interest rate is low. In the absence of capital, the extra output must be consumed for markets to clear. The real interest rate goes down to overcome the desire of the representative household to smooth consumption. This lower real interest rate translates, through the working of the Taylor rule, into a lower nominal interest rate. Often, the ZLB prevents the necessary reduction in the nominal interest rate caused by these two mechanisms.[9]

Table 5.3: $\widehat{\Pr}\left(\cup_{i=1}^{I}\{R_{t+i}=1\}\,|\,\{\mathbb{S}_{j,t}=A_j\}\right)$

	β_t	A_t	m_t	$s_{g,t}$
+2 std	29%	27%	15%	15%
+1 std	20%	21%	14%	15%
0 std	15%	15%	15%	15%
-1 std	12%	11%	14%	15%
-2 std	9%	7%	16%	16%

Understanding the economics behind this point is important. During the 2007-2009 recession in the U.S., many observers argued that the good behavior of productivity was compelling proof that the problems of the economy came from an insufficient aggregate demand. Our previous finding reinterprets the situation. Far from being just a problem of demand, high productivity may have aggravated the problem of the ZLB (see, for a related point, Eggertsson, 2010). Finally, monetary and fiscal policy shocks have a volatility that is too small to make much of a difference with respect to the probability of hitting the ZLB (and the Taylor rule undoes much of their impact).[10]

Understanding which shocks lead us to the ZLB is particularly relevant for policymakers who need to assess the conditions of the economy in order to respond to them. Consequently, we can also learn about the distribution of states conditional on being at the ZLB, that is:

$$\mu\left(\mathbb{S}_t|\,\{R_t=1\}\right). \tag{10}$$

[9] The first mechanism still exists with endogenous capital. The second mechanism will switch signs because, with capital, a higher productivity increases the desire to invest. With endogenous capital, the strength of each mechanism depends on parameter values. In Smets and Wouters (2007), to use a well-known example, the nominal interest rate goes down after a positive technological shock (see their figure 7, p. 602). Therefore, it is still the case that in their model good productivity shocks increase the risk of hitting the ZLB.

[10] Fernández-Villaverde, Guerrón-Quintana, and Rubio-Ramírez (2011) argue that higher future productivity raises current output when the economy is at the ZLB. The mechanism, far from contradicting the one in the main text, is actually quite close. A higher future productivity generates a wealth effect that increases the desire to consume today. In normal times, such desire is undone by a rise in the real interest rate to ensure market clearing. When the economy is at the ZLB, instead, it translates into higher demand and higher output. The key here is to note that current or future productivity increases have different effects on the nominal interest rate and on inflation.

Let $\{\mathbb{S}_{i_k,ss}\}_{k=1}^{M}$ be the M elements subsequence of the sequence $\{\mathbb{S}_{i,ss}\}_{i=1}^{T}$ such that $R(\mathbb{S}_{i_k,ss}) = 1$ for all $k \in \{1,\ldots,M\}$. Then, we approximate (10) by:

$$\mu(\{\mathbb{S}_t \in A\} \,|\, \{R_t = 1\}) \simeq \frac{\sum_{k=1}^{M} \mathcal{I}_A(\mathbb{S}_{i_k,ss})}{M} \tag{11}$$

for any set $A \subset \mathbb{R}^5$. Since $\mathbb{S}_t \in \mathbb{R}^5$, it is hard to represent equation (11) graphically. Instead, in figure 5.3 we represent the four individual marginal distributions

$$\mu(\{\mathbb{S}_{j,t} \in A_j\} \,|\, \{R_t = 1\}) \simeq \frac{\sum_{k=1}^{M} \mathcal{I}_{A_i}(\mathbb{S}_{j,i_k,ss})}{M}$$

for any set $A_j \subset \mathbb{R}$, where $\mathbb{S}_{j,t}$ is the j^{th} element of $\mathbb{S}_t$ and $\mathbb{S}_{j,i_k,ss}$ is the j^{th} element of $\mathbb{S}_{i_k,ss}$ (again, we drop the distribution for v_{t-1}, since it has a less straightforward interpretation). To facilitate comparison, the states are expressed as deviations from their steady state values and we plot, with the red continuous line, the unconditional distribution of the state.

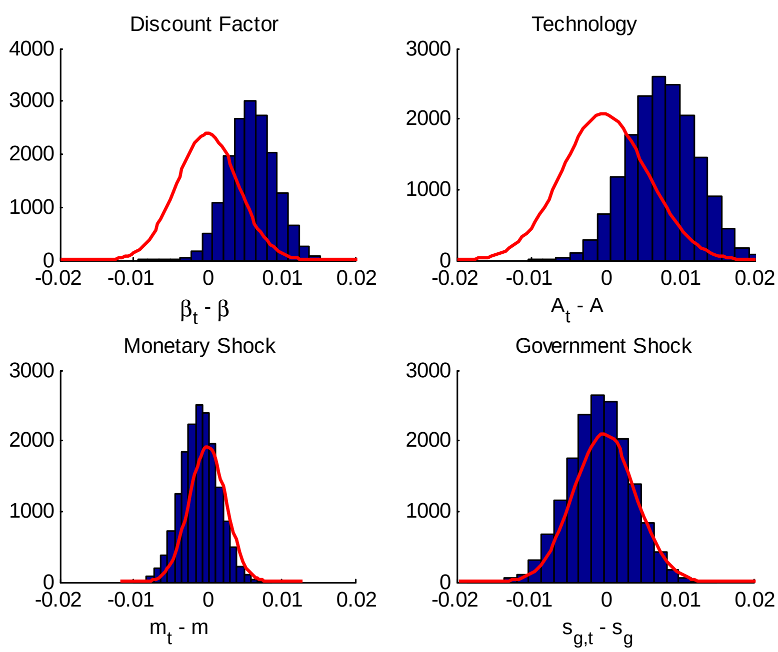

Figure 5.3: Distribution of exogenous states, unconditional and conditional on being at the ZLB

In figure 5.3 we see a pattern similar to the one in table 5.3: high discount factors and high productivity are associated with the ZLB (the mass of the distributions conditional on being at the ZLB of these two shocks is to the right of the unconditional distribution) while the fiscal and monetary shocks are nearly uncorrelated. The same information appears in table 5.4, where we report the mean and the standard deviation of the four exogenous variables, unconditionally (this is just our calibration of these processes) and conditional at being at the ZLB. The ZLB is associated with high discount factors and high productivities, but it is not correlated with either monetary or fiscal policy.

Table 5.4: Unconditional and conditional moments of (log of) exogenous states

	Mean (%)					Std (%)			
log of	β	A	m	s_g		β	A	m	s_g
Unconditional	0.00	0.00	0.00	0.00		0.41	0.58	0.25	0.42
At ZLB	0.60	0.78	-0.10	-0.07		0.32	0.47	0.25	0.42

We just analyzed the conditional marginal distribution of every individual exogenous state. We can also study the bivariate conditional distributions of the exogenous states:

$$\mu\left((\mathbb{S}_{i,t}, \mathbb{S}_{j,t}) \,|\, \{R_t = 1\}\right) \tag{12}$$

where $\mathbb{S}_{j,t}$ is the j^{th} element of $\mathbb{S}_t$ and $i, j \in \{2, \ldots, 5\}$. We approximate (12) by:

$$\mu\left(\{(\mathbb{S}_{i,t}, \mathbb{S}_{j,t}) \in A_i \times A_j\} \,|\, \{R_t = 1\}\right) \simeq \frac{\sum_{k=1}^{M} \mathcal{I}_{A_i \times A_j}\left((\mathbb{S}_{i,i_k,ss}, \mathbb{S}_{j,i_k,ss})\right)}{M}. \tag{13}$$

Figure 5.4 plots the four most interesting of these bivariate distributions. In the top left panel we see how the high productivity shock and the high preference shock comove at the ZLB. We are at the ZLB precisely because these two shocks are simultaneously high. In the top right panel we see a smaller comovement between high discount factors and the monetary policy shock. Not surprisingly, in the bottom left panel we observe again the lack of comovement between the monetary policy shock and productivity. Finally, the bottom right panel shows there is little comovement between government expenditure and productivity.

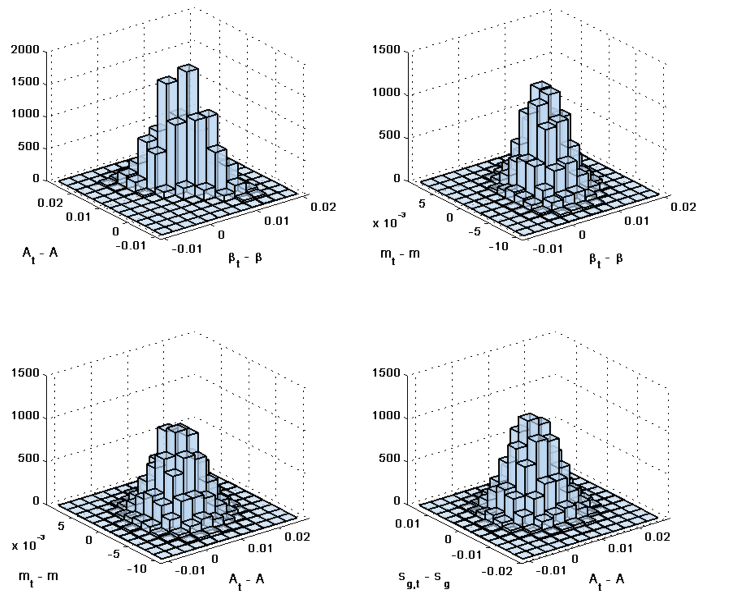

Figure 5.4: Conditional bivariate distribution of exogenous states conditional on being at the ZLB

5.4. Endogenous variables at the ZLB

We now document how the endogenous variables behave while the economy is at the ZLB. In order to analyze this issue, figure 5.5 compares the distributions of consumption, output, and inflation. The first row represents the unconditional distribution of the three variables. In particular, $\mu\left(variable_t\right)$ for $variable_t \in \{c, y, \Pi\}$ to be approximated by:

$$\mu\left(\{variable_t \in A\}\right) \simeq \frac{\sum_{i=1}^{T} \mathcal{I}_A\left(variable\left(\mathbb{S}_{i,ss}\right)\right)}{T}$$

for any set $A \subset \mathbb{R}$. Here we do not condition on being at the ZLB because we use the unconditional distributions as a reference. The second row shows the same distribution conditional of being at the ZLB, that is $\mu\left(variable_t \middle| \{R_t = 1\}\right)$, approximating it by:

$$\mu\left(\{variable_t \in A\} \,|\, \{R_t = 1\}\right) \simeq \frac{\sum_{k=1}^{M} \mathcal{I}_A\left(variable\left(\mathbb{S}_{i_k,ss}\right)\right)}{M}.$$

The third row conditions on being at the ZLB for four periods $\mu\left(variable_t|\mathbb{R}_{t-1,4}\right)$, approximating it by:

$$\mu\left(\{variable_t \in A\}\left|\mathbb{R}_{t-1,4}\right.\right) \simeq \frac{\sum_{k_4=1}^{M_4} \mathcal{I}_A\left(variable\left(\mathbb{S}_{i_{k_4},ss}\right)\right)}{M_4}.$$

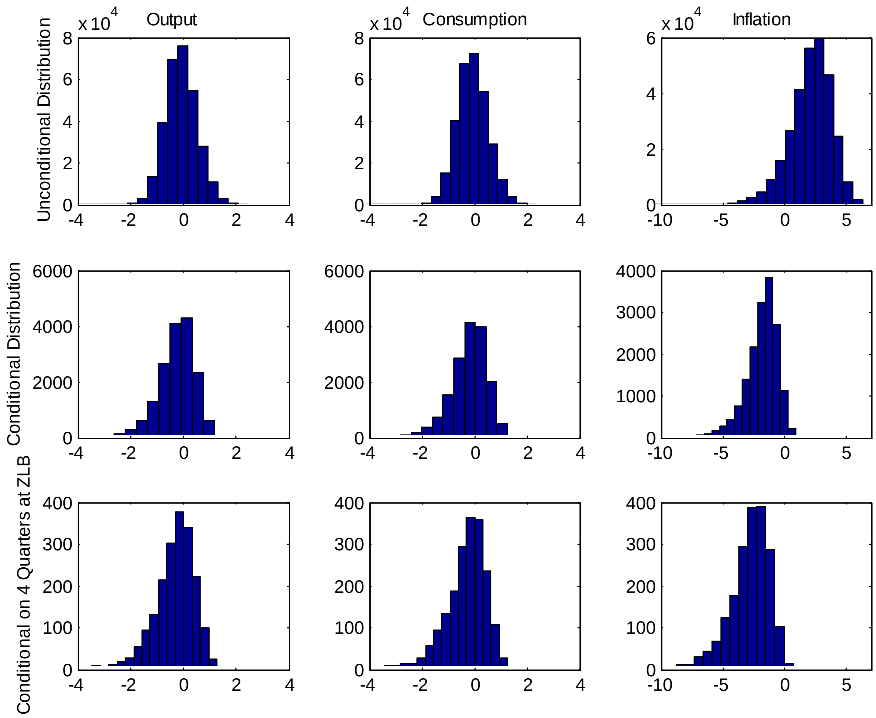

Figure 5.5: Distributions of endogenous variables

Figure 5.5 tells us that the three distributions are negatively skewed: the ZLB is associated with low consumption, output, and inflation. When at the ZLB, consumption is on average 0.23 percent below its steady-state value, output 0.25 percent below its steady-state value, and inflation is -1.89 percent (in annualized terms), 4 percentage points less than its average at the unconditional distribution (2.1 percent). As shown in the third row, the distributions are even more skewed if we condition on being at the ZLB for four periods. The average values also get more negative. Consumption is on average 0.34 percent below its steady-state value, output is 0.32 percent below its steady state, and inflation -2.83 percent (in annualized terms). We should not read the previous numbers as suggesting that the ZLB is a mild illness: in some bad events, the ZLB makes output drop 8.6 percent, consumption falls 8.2 percent, and inflation is -13.5 percent. Given our rich stochastic structure, a spell at the ZLB is not always associated with deflation. Out of the 300,000 simulations, our economy is at the ZLB in 16,588 periods, and in 562 of those, inflation is positive (around 3 percent of the time[11]), with an annualized average value of 0.29 percent.

5.5. Autopsy of a spell at the ZLB

The longest spell in our simulation lasts 25 quarters. This spell is triggered by a sharp spike in the discount rate combined with a positive productivity shock and a negative monetary policy shock. Since we can think about our discount factor shock as a proxy for demand shocks (as those caused by debt overhangs), this situation resembles the U.S. data from 2008-2009 when poor behavior of output was combined with good performance of productivity. Therefore, it is worthwhile to analyze this spell at the ZLB in great detail.

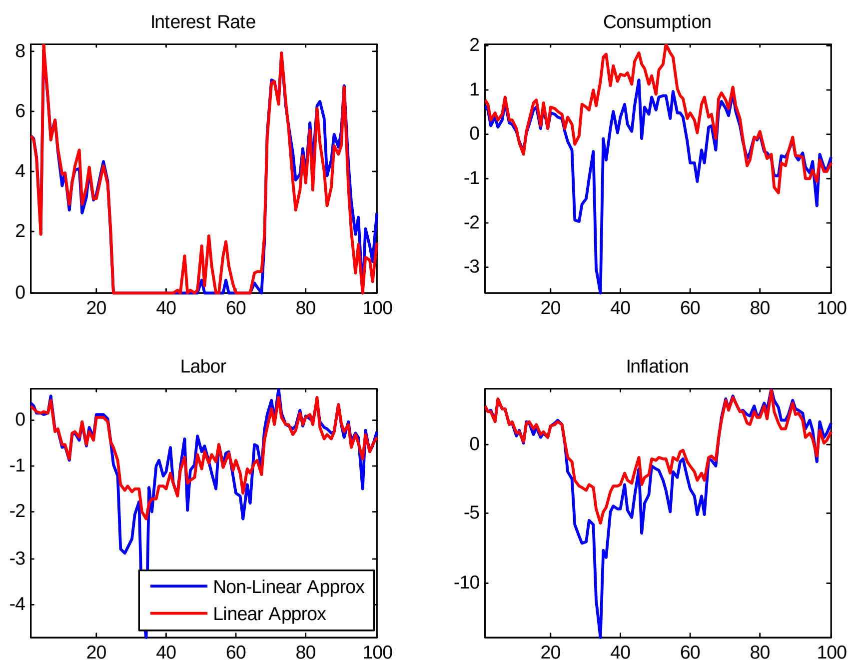

Figure 5.6: Spell at the ZLB

Figure 5.6 plots the evolution of key endogenous variables during the spell. The top left panel is the interest rate; the top right panel consumption; the bottom left panel hours worked;

[11]Since the recent spell of the U.S. economy at the ZLB has been associated with mild inflation, one can use this 3 percent as an account of that observation. In addition, we could include further mechanisms (such as an exogenous or endogenous mark-up shock) to increase this 3 percent. All that we need for the ZLB to induce negative outcomes is that inflation does not go up enough to bring the real interest rate down as much as the economy would do with completely flexible prices.

and the bottom right panel inflation (the last three variables in deviations with respect to the steady state). In each panel there are two lines, one for the nonlinear approximation and one (to appreciate the differences implied by solution methods) for the log-linear approximation.

The spell starts at period 25 and lasts until period 50. We plot some periods before and after to provide a frame for the reader for comparison. Several points are worth highlighting. First, the ZLB is associated with low consumption, less hours, and with deflation. Second, even if the economy is out of the ZLB by period 50, it is still quite close to it for many quarters (up to period 68), which shows that a model such as ours can generate (although admittedly with low probability) a "lost decade" of recession and deflation.[12] Third, the linear dynamics departs from the nonlinear one in a significant way when we are at (or close to) the ZLB. In particular, the recession is deeper and the deflation acuter. Also, consumption goes down, while in the linearized world it goes up. Thus, a policy maker looking at the linearized world would misread the situation. This divergence in the paths for consumption is not a surprise because the nonlinear policy functions that we computed in the previous section curved down with respect to the linear approximations when at (or close to) the ZLB.

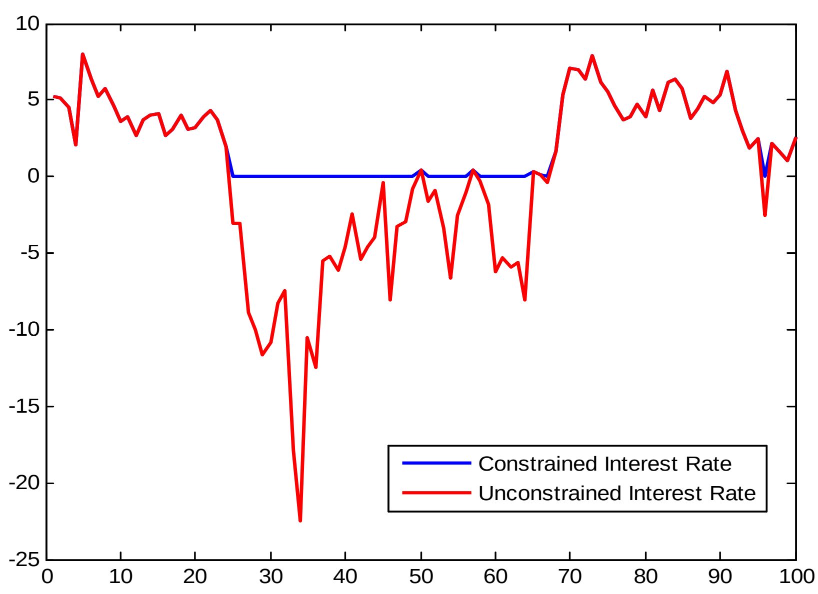

Figure 5.7: Constrained versus unconstrained interest rate

[12]We could modify the model to incorporate mechanisms such as financial frictions, correlation among shocks (such as the shocks to productivity and to preferences), or fat tails in the innovations to the exogenous variables that increase the persistence of the spells at the ZLB.

Figure 5.7 helps to understand the role of the ZLB. We plot the same path of the interest rate (from the nonlinear approximation) as in figure 5.6, but we add $Z_t - 1$, the unconstrained net interest rate that the model would have computed in the absence of the bound.[13] During the spell, the economy would require a negative interest rate (as large as -22.5 percent) to induce the households to consume enough in response to the (temporarily) high discount factor and productivity shock. Since a negative interest rate is precluded, the economy must contract to reduce the desired level of savings.

5.6. The size of the fiscal multiplier

Our next set of results corresponds to the size of the fiscal multiplier. The ZLB is a situation in which conventional monetary policy has lost its power as further reductions of the nominal interest rate are not possible. Even if other types of monetary policy (such as announcing a price level target) are still feasible, the difficulties in applying them without full commitment has led many macroeconomists to refocus their attention on fiscal policy, particularly because, at the ZLB, the fiscal multiplier may be large; as forcefully argued by Woodford (2011) and Christiano, Eichenbaum, and Rebelo (2011).

We analyze fiscal policy with the help of our model. In figure 5.8, we plot (in blue lines) the size of the government multiplier when the economy is outside the ZLB. Since the model is nonlinear, the multiplier depends on the point where we make our computation. A natural candidate is the unconditional mean of the states. Thus, the multiplier is computed as follows:

1. We find the unconditional means of output, y_1, and government spending, G_1.

2. Starting at those, we increase government consumption at impact by 1 percent, 10 percent, 20 percent, and 30 percent (that is, if government consumption is 1, we raise it to either 1.01, 1.1, 1.2, or 1.3). As time goes by, government consumption follows its law of motion (2) back to its average level.

3. For each increase in government spending, we simulate the economy. Let $y_{g,t}$ and $G_{g,t}$ denote the new simulated paths where the subindex g indicates the economy has been buffeted by a government spending shock.

4. We compute the multiplier:
$$\frac{y_{g,t} - y_1}{G_{g,1} - G_1}.$$

[13] To facilitate the interpretation, in figure 5.2 we compute $Z_t - 1$ period by period; that is, without considering that, in the absence of the ZLB, the economy would have entered in the period with a different set of state variable values (caused by the different past equilibrium dynamics triggered precisely by the absence of the bound).

We calculate the response of the economy to different increases of government consumption because, when we solve the model nonlinearly, the *marginal multiplier* (the multiplier when government consumption goes up by an infinitesimal amount) is different from the *average multiplier* (the multiplier when government consumption goes up by a discrete number). We approximate the marginal multiplier with the multiplier that increases government consumption by 1 percent. The other three higher increases give us an idea of how the average multiplier changes with the size of the increment in government consumption. When the nominal interest rate is positive, though, figure 5.8 tells us that the difference between the marginal and the average multiplier is small. This is just another manifestation of the near linearity of the model in that situation.

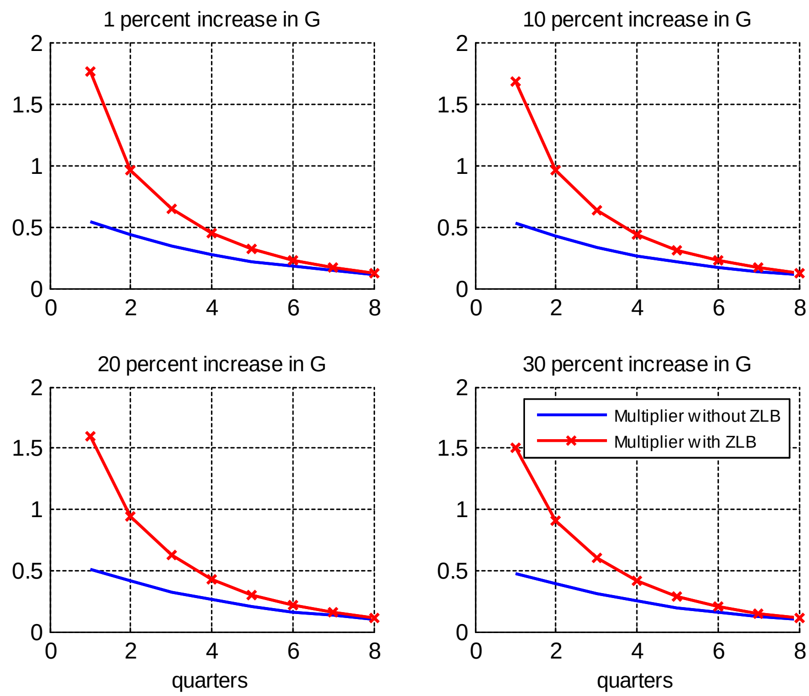

Figure 5.8: Government spending multiplier when the ZLB lasts 4 quarters on average

Outside the ZLB, the value of the multiplier at impact is around 0.5 (going from 0.54 when the increment is 1 percent to 0.47 when the increment is 30 percent), a relatively small number, but not unusual in New Keynesian models with a low labor elasticity, without liquidity-constrained households, and with a Taylor rule that responds to the output gap.[14] This value is also consistent with the analytic insights in Woodford (2011), who derives multipliers less than 1 when the economy is outside the ZLB and has parameter values commonly found

[14]In the log-linearized version of the model, the multiplier at impact is 0.53. Just to illustrate our argument in the main text, if we eliminate the response of the Taylor rule to the output gap, the multiplier from the log-linear approximation would be 0.68.

in empirically estimated medium-sized DSGE models. The observation that the multiplier decreases with the size of the increase in government consumption suggests that we must be careful when we read the empirical evidence of fiscal multipliers, as their estimated value may depend on the size of the observed changes in government consumption.

Figure 5.8 also displays the multiplier when the economy is at the ZLB (red crossed lines). The multiplier is computed as follows:

1. We start at the unconditional mean of states except that, to force the economy into the ZLB, we set the discount factor 2 percent above its unconditional mean (at 1.014 instead of the calibrated mean 0.994). This shock sends the economy to the ZLB, on average, for 4 consecutive quarters in the absence of any additional shock during those periods. We call these unconditional means, y_1^{zlb}, and government spending, G_1^{zlb}.

2. With these states, we simulate the economy and store the time paths for output, y_t^{zlb}, and government spending, G_t^{zlb}.

3. With the same states, except that we consider an increase in government consumption at impact of 1 percent, 10 percent, 20 percent, and 30 percent relative to its unconditional mean, we simulate the economy. As time goes by, government consumption follows its law of motion. Let $y_{g,t}^{zlb}$ and $G_{g,t}^{zlb}$ denote the new simulated paths.

4. We compute the multiplier:
$$\frac{y_{g,t}^{zlb} - y_t^{zlb}}{G_{g,1}^{zlb} - G_1^{zlb}}.$$

Now the multiplier of government consumption at impact is much larger: 1.76 for increases of government consumption of 1 percent, 1.68 for increases of 10 percent, 1.59 for increases of 20 percent, and 1.50 for increases of 30 percent. That is, the new multiplier is around 3 times larger than when outside the ZLB.

We now discuss some details of our exercise. We raise β_t to 1.014. Given our calibration, this raise corresponds to a roughly 8 standard deviations shock to the discount factor. This is a rare event that pushes us deeply into the ZLB territory. We picked it to generate a spell at the ZLB of average duration of 4 quarters (although, remember that this is just an average duration and that the economy actually faces a whole distribution of durations with shorter and longer durations possible). We select an average duration of 4 quarters to be close to other papers in the literature. This choice is consequential because in our economy the characteristics of the shock that bring you to the ZLB matter for the size of the multiplier (this is not the case in Christiano, Eichenbaum, and Rebelo, 2011).

Several patterns are apparent from the red lines in figure 5.8. First, the multiplier at the ZLB is significantly above one, albeit smaller than the values reported in the literature (Christiano, Eichenbaum, and Rebelo, 2011). Second, as in the previous case, the multiplier declines as the size of the government shock increases. This is because, as the government shock increases, the expected number of periods at the ZLB decreases for a given discount factor shock. In our experiment, a 30 percent government shock pushes the economy, on average, out of the ZLB upon impact.

The natural question is to ask why the multiplier is smaller than the one reported in the literature, for instance by Christiano, Eichenbaum, and Rebelo (2011) in a related environment. Our exercise is somewhat different from previous ones. First, in our model there is a rich stochastic structure: in any period the economy is hit by several shocks. Agents do not have perfect foresight of future events and the distribution of the number of periods at the ZLB is not degenerate. Second, the government expenditure follows its law of motion (2) regardless of whether or not the economy is at the ZLB. So, for instance, if a large government consumption shock pushes the economy out of the ZLB right away, government consumption would still be high for many quarters, lowering the multiplier.[15] Our experiment is a nice complement to other exercises in the literature and provides answers that policy-makers are interested in: neither governments nor agents know in real life for how many periods the economy will be at the ZLB, and government expenditure is difficult to change once stimulus plans have been approved.

Table 5.5: Unconditional and conditional moments of exogenous states

σ_b	Expected periods at ZLB	Standard dev. of periods at ZLB	Multiplier
0.0027	4	2.04	1.97
0.0025	4	2.08	1.76
0.00175	4	2.17	1.57

Let us now analyze in detail how the nondegenerate distribution of the number of periods at the ZLB that the agents face affects our results. To do this, we solve three versions of the model with three levels of volatility for the discount factor shock (column 1 of table 5.5). In each of the three cases, we increase—at time zero—the discount factor so that, on average, the economy stays at the ZLB for 4 periods.[16] We report the standard deviation of the periods at the ZLB (third column) and the impact fiscal multiplier (fourth column). To compute these numbers (and for each version of our model), we simulate the economy $5,000$ times. As we

[15] Christiano, Eichenbaum, and Rebelo (2011) also show a similar result.

[16] The size of the innovation is different in each case because of σ_b, which affects both the impact of the current and future innovation and the behavior of agents.

reduce σ_b, the multiplier falls, while the standard deviation of the number of periods at the ZLB grows. Because of the left censuring of the spell duration, even if the expected durations of the ZLB spell are constant across the three experiments, a higher standard deviation of the number of periods at the ZLB means a higher probability of ZLB spell durations shorter than four periods. For instance, this probability increases four percentage points from $\sigma_b = 0.0027$ to $\sigma_b = 0.00175$. Since, in these events, government expenditure pushes interest rates up, the total effect of additional government expenditure on aggregate demand today is much lower.

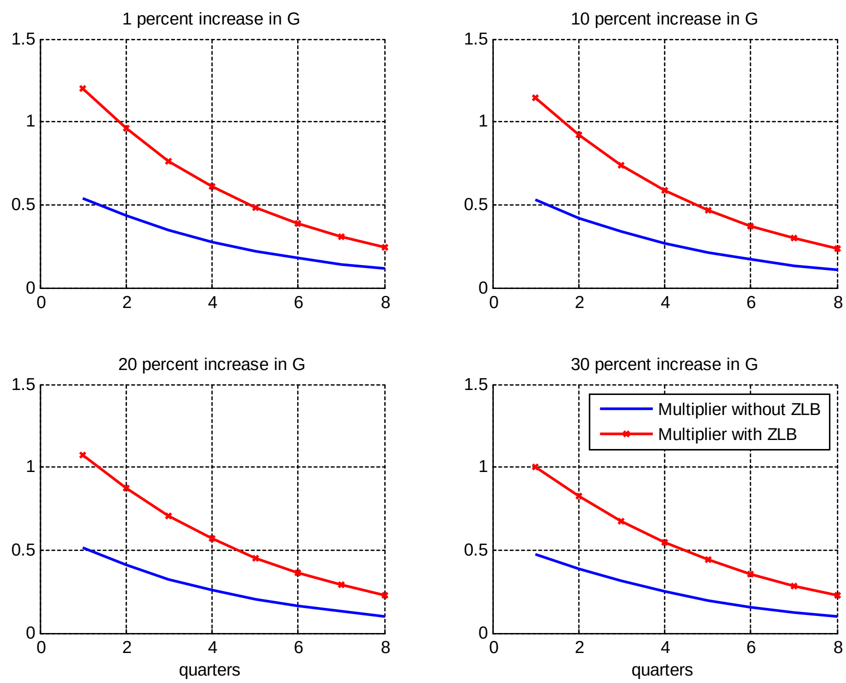

Figure 5.9: Government spending multiplier when the ZLB lasts 2 quarters on average

To evaluate how the expected number of periods at the ZLB affects the fiscal multiplier, figure 5.9 shows the results when the discount factor is initially set to be 1.2 percent above its unconditional mean. This sends the economy to the ZLB, in the absence of additional shocks, on average for 2 quarters. The multiplier outside the ZLB is the same as before. But the multiplier at the ZLB is now smaller at only around 2 times larger than in normal times. For an increment of 1 percent of government expenditure, the impact multiplier is 1.20, and it goes down to 1.00 when the increment is 30 percent. Figure 5.9 illustrates, thus, how the size of the fiscal multiplier depends on the expected duration of the spell at the ZLB without fiscal shocks (Christiano, Eichenbaum, and Rebelo, 2011, document this point as well). The intuition comes directly from the Euler equation forwarded several periods. The longer we are expected to stay at the ZLB, the lower is the current demand for consumption (a lower consumption in the future imposes, by optimality, a lower consumption today).

5.7. Endogenous durations of spells at the ZLB

We complete our analysis by reporting the IRFs of the model (computed as we did for the fiscal multipliers). The IRFs highlight how the duration of the spell at the ZLB is endogenous. Figure 5.10 shows the IRFs for an increase of government expenditure of 1 percent. The IRFs are computed as percentage deviations from the unconditional mean when we are outside the ZLB (blue line) and as percentage deviations from the trajectories without government spending when the ZLB binds (red crossed line). At the ZLB, Figure 5.10 is again computed under a demand shock that sends the economy to the ZLB, on average, for 4 consecutive quarters.

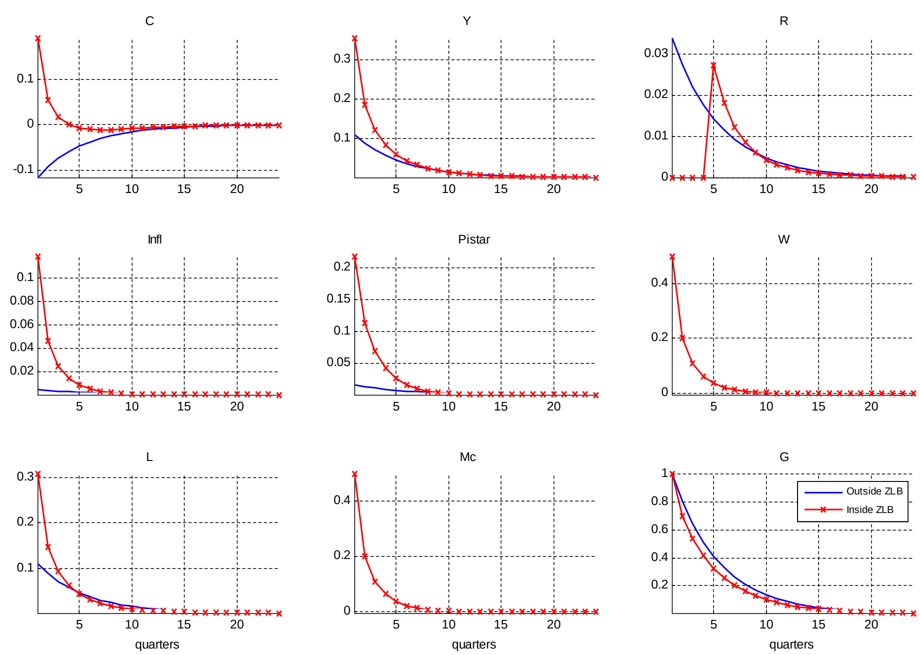

Figure 5.10: IRFs to a 1% government spending increase when the ZLB lasts 4 quarters in expectation

Under normal circumstances (blue line), government spending crowds out consumption. This is a direct consequence of a multiplier less than 1. Additional government spending puts upside pressure on prices forcing the nominal interest rate to go up. Real interest rates follow suit, since the central bank follows the Taylor principle. Households take advantage of high interest rates and defer consumption for the future. Ultimately, output (and labor) rise above its unconditional mean, but by less than the increase in government spending. In contrast, when the economy is at the ZLB (red crossed line), consumption and output simultaneously go up in response to higher government spending. With the nominal interest rate locked at zero and rising inflation, the real rate is below the value prevailing in the absence of the government stimulus. Hence, households prefer to reduce savings and consume more. Note that, when the government spending shock is small it does not affect the expected duration of the ZLB spell (4 periods).

Figure 5.11 shows that a government shock of 10 percent pushes the economy out of the ZLB, on average, after three quarters. This is one period less than the expected duration without the government shocks. This shorter expected duration implies a smaller response of consumption today and, therefore, a smaller effect of government spending at impact. Yet the multiplier is still above one because the shock does not preclude the economy from immediately escaping the ZLB.

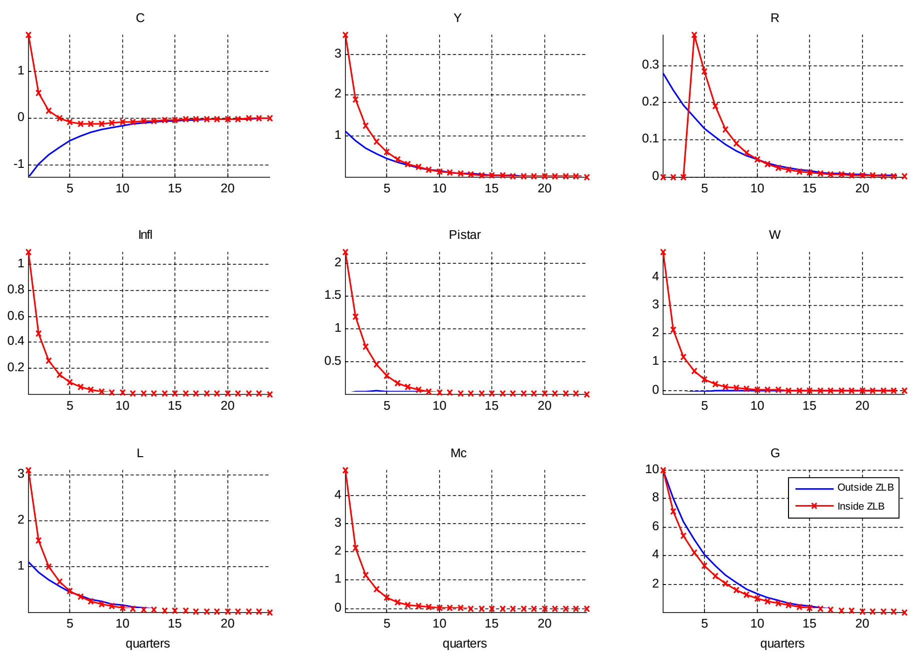

Figure 5.11: IRFs to a 10% government spending increase when the ZLB lasts 4 quarters in expectation

Finally, for a sufficiently large government shock, the economy is expected to escape the ZLB (figure 5.12) at impact. The rise in interest rates compresses consumption although not to the levels when the economy is unconstrained (blue line). The multiplier for this case is close to one but larger than when computed without the ZLB.

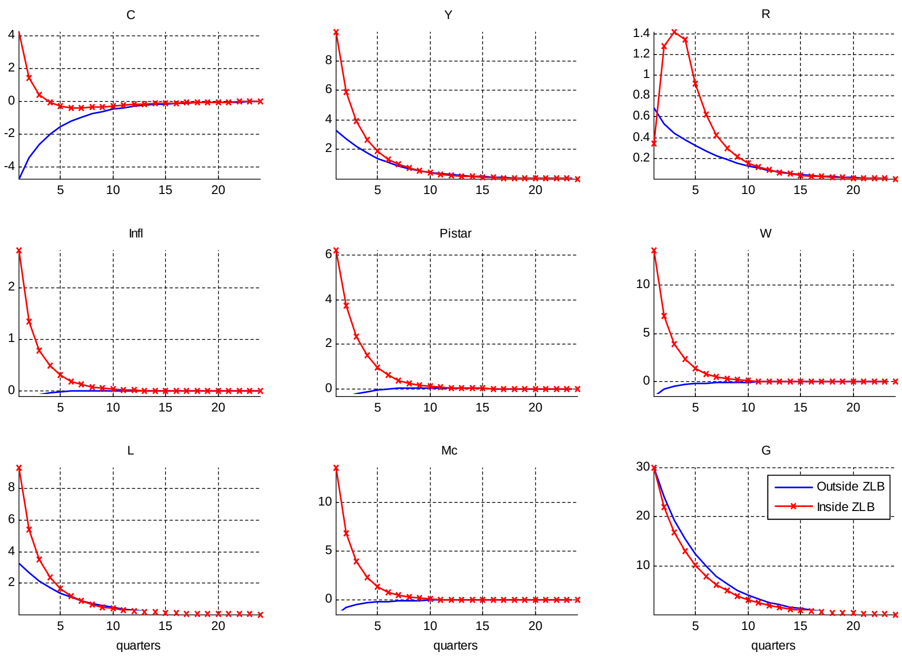

Figure 5.12: IRFs to a 30% government spending increase when the ZLB lasts 4 quarters in expectation

6. Conclusion

Our paper is a step toward a more complete understanding of economies in which the ZLB is a relevant constraint in the implementation of policy. We have demonstrated that a full nonlinear solution delivers important insights, for instance with respect to the role of fiscal policy at the ZLB and the conditional expectations of agents.

Several lines of future research lay ahead of us. First, we can have richer models that are closer to the ones that are used by central banks for practical policy advising (for instance, Christiano, Eichenbaum, and Evans, 2005). These models will have both forces that make the ZLB less important (endogenous capital) and forces that make it stronger (habit persistence, financial frictions, etc.). It is an open question to assess which of these will dominate. Second, we could extend our model to include a richer set of fiscal policy instruments and ask, for instance, about the differences in responses of reductions in distortionary taxes when we are at the ZLB and when we are not. Since the ZLB is a situation in which households are already saving too much, reductions in taxes financed through debt (with its associated quasi-Ricardian effect of higher savings) may have little bite. Third, we can mix our solution with a particle filter to build the likelihood function for estimation as in Fernández-Villaverde and Rubio-Ramírez (2007). Estimating this class of models with a full-likelihood approach is promising. We can use the information in the data about parameters describing preferences and technology to evaluate the behavior of the model at the ZLB, something that more reduced-form models, such as Structural Vector Autoregressions, would have a harder time doing because the U.S. has only been at the ZLB once since World War II. If the U.S. and other European economies stay at the ZLB in the close future or the prospects of coming back to it are sufficiently high, these next steps seem to be a high priority for applied macroeconomists.

References

[1] Braun, R.A. and L.M. Körber (2010). "New Keynesian Dynamics in a Low Interest Rate Environment." *Mimeo*, Federal Reserve Bank of Atlanta.

[2] Christiano, L., 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.

[3] Christiano, L., M. Eichenbaum, and S. Rebelo (2011). "When is the Government Multiplier Large?" *Journal of Political Economy* 119, 78-121.

[4] Chung, H., J-P. Laforte, D. Reifschneider, and J. C. Williams (2010). "Have We Underestimated the Probability of Hitting the Zero Lower Bound?" *Mimeo*, Board of Governors.

[5] Coleman, W. J. (1991) "Equilibrium in a Production Economy with an Income Tax." *Econometrica* 59, 1091-1104.

[6] Cooley, T. and E.C. Prescott (1995). Economic Growth and Business Cycles. *Frontiers of Business Cycle Research*, edited by Thomas Cooley. Princeton University Press.

[7] Eggertsson, G.B. (2010). What Fiscal Policy is Effective at Zero Interest Rates? *NBER Macroeconomic Annual*.

[8] Eggertsson, G.B. and M. Woodford (2003). "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity* 34, 139-235.

[9] Eichenbaum, M. and J. Fisher (2007). "Estimating the Frequency of Price Reoptimization in Calvo-style Models." *Journal of Monetary Economics* 54, 2032-2047.

[10] Erceg, C. and J. Lindé (2010). "Is There a Fiscal Free Lunch in a Liquidity Trap?" *CEPR Discussion Paper* 7624.

[11] Fernández-Villaverde, J. and J. Rubio-Ramírez (2007). "Estimating Macroeconomic Models: A Likelihood Approach." *Review of Economic Studies* 74, 1059-1087.

[12] Fernández-Villaverde, J., P. Guerrón-Quintana, and J. Rubio-Ramírez (2011). "Supply-Side Policies and the Zero Lower Bound." *Working Papers* 11-47, Federal Reserve Bank of Philadelphia.

[13] Guerrón-Quintana, P. (2010). "The Implications of Inflation in an Estimated New-Keynesian Model." *Working Papers* 10-2, Federal Reserve Bank of Philadelphia.

[14] Judd, K.L., L. Maliar, and S. Maliar (2011). "A Cluster-Grid Algorithm: Solving Problems with High Dimensionality." *Mimeo*, Hoover Institution.

[15] Justiniano A. and G.E. Primiceri (2008). "The Time-Varying Volatility of Macroeconomic Fluctuations." *American Economic Review* 98, 604-641.

[16] Krueger, D., F. Kubler, and B. Malin (2011). "Solving the Multi-Country Real Business Cycle Model Using a Smolyak-Collocation Method." *Journal of Economic Dynamics and Control* 35, 229-239.

[17] Reifschneider, D. and J. C. Williams (1999). "Three Lessons for Monetary Policy in a Low Inflation Era." *Mimeo*, Board of Governors.

[18] Santos, M.S. and A. Peralta-Alva (2005). "Accuracy of Simulations for Stochastic Dynamic Models." *Econometrica* 73, 1939-1976.

[19] Smets, F. and R. Wouters (2007). "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97, 586-606.

[20] Tanaka, T. (2012). "Optimal Fiscal and Monetary Policy with Occasionally Binding Zero Bound Constraints." *Mimeo*, New York University.

[21] Wolman, A.L. (1998). "Staggered Price Setting and the Zero Bound on Nominal Interest Rates." *Economic Quarterly, Federal Reserve Bank of Richmond* Fall, 1-24.

[22] Woodford, M. (2011). "Simple Analytics of the Government Expenditure Multiplier." *American Economic Journal: Macroeconomics* 3, 1–35.

7. Appendix: Solving the Model

We explain now in further detail how we compute the model using a projection method. First, however, we need to find the steady state as this centers the projection and serves as the point around which we log-linearize the model in the comparison exercise of section 4.

7.1. Steady state

The steady state is defined by the following equations (where now Π is not a variable but a parameter, the inflation target):

$$R = \frac{\Pi}{\beta}$$

$$\psi l^{\vartheta} c = w$$

$$mc = w$$

$$\varepsilon x_1 = (\varepsilon - 1) x_2$$

$$x_1 = \frac{1}{c} mcy + \beta\theta\Pi^{\varepsilon} x_1$$

$$x_2 = \left(\frac{\Pi^*}{c} y + \beta\theta\Pi^{\varepsilon-1} x_2 \right)$$

$$g = s_g y$$

$$b = 0$$

$$1 = \theta\Pi^{\varepsilon-1} + (1-\theta) (\Pi^*)^{1-\varepsilon}$$

$$v = \theta\Pi^{\varepsilon} v + (1-\theta) (\Pi^*)^{-\varepsilon}$$

$$y = c + g$$

$$y = \frac{1}{v} l.$$

To solve these equations, we first get the variables:

$$\Pi^* = \left(\frac{1-\theta\Pi^{\varepsilon-1}}{1-\theta}\right)^{\frac{1}{1-\varepsilon}}$$
$$v = \frac{1-\theta}{1-\theta\Pi^{\varepsilon}}\left(\Pi^*\right)^{-\varepsilon}$$
$$\frac{g}{y} = s_g$$
$$\frac{c}{y} = (1-s_g)$$
$$x_2 = \frac{1}{1-s_g}\frac{\Pi^*}{1-\beta\theta\Pi^{\varepsilon-1}}$$

$$x_1 = \frac{\varepsilon-1}{\varepsilon}x_2$$
$$mc = (1-s_g)\left(1-\beta\theta\Pi^{\varepsilon}\right)x_1$$

and

$$w = mc.$$

Setting $\psi = 1$, which just selects a normalization for l, one has

$$\frac{l}{y} = v$$
$$\frac{l}{c} = \frac{l}{y}\frac{y}{c}$$
$$l = \left(\frac{w}{\psi}\frac{l}{c}\right)^{\frac{1}{1+\vartheta}}.$$

Lastly, c, g, and y are determined using $l/y, c/y, g/y$, and l.

7.2. The projection method

Recall that the state variables of the model are price dispersion, v; the discount factor, β_t; productivity, A_t; the monetary shock, m_t; and the government expenditure share, $s_{g,t}$. Recall that we have defined a vector,

$$\mathbb{S}_t = \left(v_{t-1}, \beta_t, A_t, m_t, s_{g,t}\right),$$

where the first state, price dispersion, is endogenous, while the other four are exogenous. Define a new vector,

$$\hat{\mathbb{S}}_t = (v_{t-1}, \log \beta_t, \log A_t, \log m_t, \log s_{g,t}),$$

so that the exogenous states are in logs. Then, we can write the equilibrium functions that characterize the dynamics of the model as:

$$\begin{aligned} \log c_t &= f^1(\hat{\mathbb{S}}_t) \\ \log \Pi_t &= f^2(\hat{\mathbb{S}}_t) \\ \log x_{1,t} &= f^3(\hat{\mathbb{S}}_t) \end{aligned}$$

where $f = (f^1, f^2, f^3)$. Since f is unknown, to solve the model, we approximate $\log c_t$, $\log \Pi_t$, and $\log x_{1,t}$ as

$$\begin{aligned} \log c_t &= \widehat{f}^1(\hat{\mathbb{S}}_t) \\ \log \Pi_t &= \widehat{f}^2(\hat{\mathbb{S}}_t) \\ \log x_{1,t} &= \widehat{f}^3(\hat{\mathbb{S}}_t) \end{aligned}$$

using the Smolyak collocation method laid out in Krueger, Kubler, and Malin (2011). Define $\widehat{f} = \left(\widehat{f}^1, \widehat{f}^2, \widehat{f}^3\right)$.

To construct $\widehat{f}$, three steps are required.[17] First, one must specify both a level of approximation μ and bounds on a hypercube. Given these two, the method specifies a finite number n of collocation points $\{\hat{\mathbb{S}}_i\}_{i=1}^n$ within the hypercube. Second, one must provide the values of f at each of these n points. Given these values, the method constructs polynomial coefficients that implicitly define $\hat{f}$. Third, given the coefficients, the method provides a way to evaluate $\hat{f}$ at any point (inside or outside of the hypercube).

The approximation $\hat{f}$ satisfies several important properties:

1. $\hat{f}$ agrees with f at each $\hat{\mathbb{S}}_i$.

2. If f is composed of polynomials that have degree μ or belong to a certain subset of the complete set of polynomials of degree 2^μ, then $\hat{f}$ agrees with f everywhere. For our level of approximation, $\mu = 2$, when f is composed of linear combinations of the polynomials x_j^4,x_j^3, x_j^2, x_j, 1, $x_j^2 x_k^2$, $x_j^2 x_k$, $x_j x_k$ for $j, k = 1, \ldots, d$, where d is the dimension of $\mathbb{S}$, it will be reproduced exactly. See Barthelmann, Novak, and Ritter (2000) for details.

3. If f is continuous, the approximation $\hat{f}$ is nearly optimal in a certain sense.

[17] The exposition of the method here closely follows Gordon (2011).

To define the hypercube, we must choose bounds for the 5 state variables. For the endogenous state variable v_{t-1}, we choose [1.000,1.005]. For each exogenous one, we choose bounds equal to the steady-state value plus or minus three unconditional standard deviations. In this way, we cover 99.7 percent of the mass of each process. For domain values outside of the hypercube (which sometimes occur in the simulation and when computing expectations), we extrapolate.

To compute expectations, we use a degree 11 monomial formula from Stroud (1971, p. 323).[18] The formula provides k innovations $\{(\varepsilon_{j,b}, \varepsilon_{j,a}, \varepsilon_{j,m}, \varepsilon_{j,g})\}_{j=1}^{k}$ and weights $\{w_j\}_{j=1}^{k}$ such that, for any function of interest h,

$$\mathbb{E}_t h(\hat{\mathbb{S}}_{t+1}) \approx \sum_{j=1}^{k} w_j h(\hat{\mathbb{S}}_{j,t+1})$$

where $\hat{\mathbb{S}}_{j,t+1}$ is found using the $j-$th innovations and the time t information $\hat{\mathbb{S}}_t$ and v_t.[19] It is called a degree 11 rule because if h is a polynomial with degree less than or equal to 11, then the approximation is exact. We also check the solution using a Monte Carlo with 100,000 draws, which generates a slightly different solution. Specifically, the values at the collocation points differ by at most 0.0006 for $\log c$, 0.0004 for $\log \Pi$, and 0.0018 for $\log x_1$.

To solve for $\hat{f}$, we use a time-iteration procedure as follows:

1. Guess on the values of $\log c$, $\log \Pi$, and $\log x_1$ at each collocation point $\mathbb{S}_i$. This implicitly gives an approximating function $\widehat{f}$ that is defined over the entire state space.

2. Treating the approximations as the true time $t+1$ functions, compute the value of the optimal time t functions at each collocation point $\hat{\mathbb{S}}_i$, $i = 1, \ldots, n$, in the following way:

   a. Fix an i. This gives the state today, $(v_{t-1}, \beta_t, A_t, m_t, s_{g,t})$.

   b. Guess on Π_t.

   c. Calculate Π_t^* from Π_t.

   d. Calculate v_t from the law of motion for price dispersion using Π_t and Π_t^*.

[18] See Judd (1998) for a discussion of monomial rules and formulas for degree 3 and degree 5 rules. Krueger, Kubler, and Malin (2011) use a degree 5 rule. We use a degree 11 formula because it has been recently documented that accurately evaluating expectations in dynamic equilibrium models is very important for overall accuracy (see, in a slightly different context, Judd, Maliar, and Maliar, 2011). To apply the formula, some slight adjustments must be made, as the weighting function is not exactly the Gaussian density.

[19] Specifically, $\hat{\mathbb{S}}_{j,t+1} = (v_t, (1-\rho_b)\log\tilde{\beta} + \rho_b \log \beta_t + \sigma_b \varepsilon_{j,b}, \rho_a \log A_t + \sigma_a \varepsilon_{j,a}, \sigma_m \varepsilon_{j,m}, \rho_g \log s_{g,t} + \sigma_g \varepsilon_{j,g})$. Note that while v_t is part of $\hat{\mathbb{S}}_{t+1}$, it is a function of only v_{t-1}, Π_t, Π_t^*, and parameters.

e. Using v_t, the state today, the innovations and weights from the monomial rule, and the time $t+1$ functions, compute the expectations

$$\mathbb{E}_t\left\{\frac{\beta_{t+1}}{c_{t+1}}\frac{1}{\Pi_{t+1}}\right\}, \mathbb{E}_t\left\{\beta_{t+1}\Pi_{t+1}^{\varepsilon}x_{1,t+1}\right\}, \text{ and } \mathbb{E}_t\left\{\beta_{t+1}\frac{\Pi_{t+1}^{\varepsilon-1}}{\Pi_{t+1}^{*}}x_{2,t+1}\right\}.$$

Here, Π_{t+1}^* and $x_{2,t+1}$ are calculated from Π_{t+1} and $x_{1,t+1}$.

f. Using these expectations and the guess on Π_t, calculate all remaining time t values:

1. c_t/y_t from the government expenditure shock.
2. $x_{2,t}$ from its intertemporal equation.
3. $x_{1,t}$ from $x_{2,t}$.
4. mc_t from the intertemporal equation for x_1.
5. The unique (c_t, R_t) pair using a guess-and-verify approach.[20] This is done in the following way:
    1. First, assume $R_t > 1$ and compute c_t using the Euler equation, the Taylor rule, and the consumption-output ratio.
    2. Second, check whether Z_t as a function of c_t is greater than 1. If it is, c_t is optimal, and R_t is equal to Z_t. If it is not, then R_t equals 1, and c_t is given by the Euler equation.
6. y_t from c_t and $s_{g,t}$.
7. l_t from aggregate supply as $l_t = y_t v_t / A_t$.
8. w_t from the labor-leisure choice condition.

g. To check whether the initial guess on Π_t was correct, compute $\hat{mc}_t = w_t/A_t$. If $|\hat{mc}_t - mc_t|$ is close, then stop. Otherwise, make a new guess on Π_t and go to step c.[21]

[20] To see that this is unique, consider the following. First, defining the value of the expectation

$$\mathbb{E}_t\left\{\frac{\beta_{t+1}}{c_{t+1}}\frac{1}{\Pi_{t+1}}\right\}$$

as κ, from the Euler equation one has $R_t = 1/(c_t\kappa)$. Second, the Taylor rule with the zero lower bound can be manipulated to say $R_t = \max\{\lambda c_t^{\phi_y}, 1\}$ where

$$\lambda = R(\Pi_t/\Pi)^{\phi_\pi} y^{-\phi_y}(1 - s_{g,t}s_g)^{-\phi_y} m_t.$$

From these two equations, it is easy to see there exists a unique (c_t, R_t) pair.

[21] Actually, we use a slightly different stopping criteria. In particular, we use the Matlab routine `fzero` to find a zero of the residual function $r(\Pi_t) = \hat{mc}_t(\Pi_t) - mc_t(\Pi_t)$. The routine brackets the equilibrium value of Π_t in an interval $[a, b]$ and progressively shrinks the bracket until $b - a < 2 \times 10^{-16}$ apart.

3. Having found time t equilibrium values for $\log c, \log \Pi$, and $\log x_1$ at each of the collocation points, check how different they are from the $t+1$ values. If the new values differ by less than 10^{-6}, stop. If they differ by more, use the time t values to update the guesses and go to step 2.

We tried several alternative procedures but found this one to be the most reliable. We encountered convergence problems for highly persistent shocks or large shocks. Making appropriate initial guesses for the functions is important when the ZLB binds often. To construct good guesses, we first solved the model for i.i.d. discount factor shocks—in which case the lower bound does not bind—and then progressively increased the persistence to its benchmark value.

References

[1] Barthelmann, V., E. Novak, and K. Ritter (2000). "High Dimensional Polynomial Interpolation on Sparse Grids," *Advances in Computational Mathematics* 12, 273-288.

[2] Gordon, G. (2011). "Computing Dynamic Heterogeneous-Agent Economies: Tracking the Distribution." *PIER Working Paper* 11-018, University of Pennsylvania.

[3] Judd, K.L. (1998). *Numerical Methods in Economics.* MIT Press.

[4] 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.

[5] Krueger, D., F. Kubler, and B. Malin (2011). "Solving the Multi-country Real Business Cycle Model Using a Smolyak-Collocation Method." *Journal of Economic Dynamics and Control* 35, 229-239.

[6] Stroud, A.H. (1971). *Approximate Calculation of Multiple Integrals.* Englewood Cliffs, NJ: Prentice-Hall.