The Zero Lower Bound, the Dual Mandate, and Unconventional Dynamics*

William T. Gavin Benjamin D. Keen
Alexander W. Richter Nathaniel A. Throckmorton

November 25, 2014

Abstract

This article examines monetary policy when it is constrained by the zero lower bound (ZLB) on the nominal interest rate. Our analysis uses a nonlinear New Keynesian model with technology and discount factor shocks that accounts for the expectational effects of falling to and remaining at the ZLB. Specifically, we investigate why technology shocks may have unconventional effects at the ZLB, what factors affect the likelihood of hitting the ZLB, and the tradeoffs a central bank faces under a dual mandate. We initially focus on a New Keynesian model without capital (Model 1) and then study that model with capital (Model 2). The advantage of including capital is that it introduces another mechanism for intertemporal substitution that strengthens the expectational effects of the ZLB. Three main findings emerge: (1) In Model 1, the choice of output target in the Taylor rule may reverse the effects of technology shocks when the ZLB binds; (2) When the central bank targets steady-state output in Model 2, a positive technology shock at the ZLB leads to more pronounced unconventional dynamics than in Model 1; and (3) In Model 1, the constrained linear solution is a decent approximation of the nonlinear solution, but meaningful differences exist between the solutions in Model 2.

Keywords: Monetary Policy; Zero Lower Bound; Nonlinear Solution Method; Capital
JEL Classifications: E31; E42; E58; E61

*Gavin, Research Division, Federal Reserve Bank of St. Louis, P.O. Box 442, St. Louis, MO (gavin@stls.frb.org); Keen, Department of Economics, University of Oklahoma, 308 Cate Center Drive, 437 Cate Center One, Norman, OK 73019 (ben.keen@ou.edu); Richter, Department of Economics, Auburn University, 0332 Haley Center, Auburn, AL 36849 (arichter@auburn.edu); Throckmorton, Department of Economics, College of William & Mary, Morton Hall 131, Williamsburg, VA (nathrockmorton@wm.edu). This paper previously circulated under the title "Global Dynamics at the Zero Lower Bound." We thank Klaus Adam, Javier Birchenall, Toni Braun, Brent Bundick, Ricardo Reis, and two anonymous referees for helpful comments on an earlier draft. We also thank seminar participants at the University of California, Santa Barbara, the Federal Reserve Bank of St. Louis, the Federal Reserve Bank of Dallas, and Renmin University in Beijing, China, and participants at the 2013 Midwest Economic Association meeting, the 2013 Midwest Macroeconomic meeting, and the 2013 Computational Economics and Finance meeting for helpful comments. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

1 INTRODUCTION

In the aftermath of the 2008 financial crisis, aggregate demand fell sharply. The Fed responded by lowering its policy rate to its zero lower bound (ZLB) by the end of the year. Six years after the crisis began, the Fed's target interest rate remains near zero and the economy is below potential.

Figure 1 shows the U.S. and Japanese interbank lending rates and employment-to-population percentages from 1990-2014. The U.S. policy rate (solid line) has varied between 8.3% and 0% since 1990 and has been held below 25 basis points since the end of 2008. During that time period, policymakers shifted their focus from inflation to the real economy, since the inflation rate has been at or below the Fed's inflation target. The Bank of Japan sharply lowered its policy rate in 1991 (dashed line), reaching 50 basis points in 1995. Since then it has remained between 0 and 50 basis points, while the employment-to-population percentage has fallen steadily from 62% to about 57.5%. The Japanese economy slightly rebounded in the mid-2000s, but after the financial crisis, the policy rate was cut and the employment-to-population percentage fell even further.

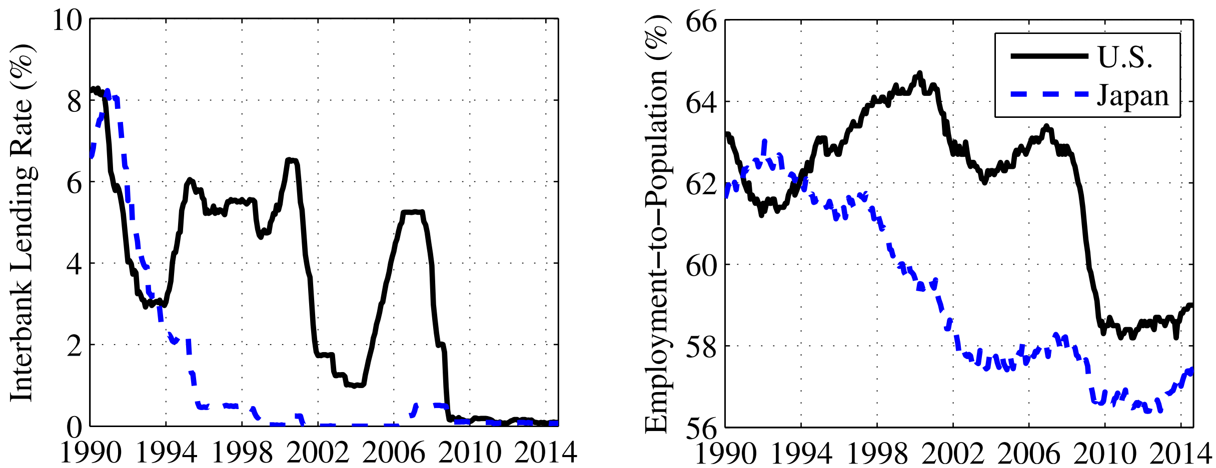

Figure 1: U.S. and Japanese interbank lending rates (left panel) and employment-to-population percentages (right panel). Sources: Federal Reserve, Bank of Japan, U.S. Bureau of Labor Statistics, and Statistics Bureau of Japan.

Over the last two decades, the Japanese economy has endured anemic growth in real GDP and slight deflation. Their experience has generated a significant amount of research on the effects of the Bank of Japan's zero interest rate policy [e.g., Braun and Waki (2006); Eggertsson and Woodford (2003); Hoshi and Kashyap (2000); Ito and Mishkin (2006); Krugman (1998); Posen (1998)]. Many arguments for avoiding the ZLB are motivated in part by the recent Japanese experience.

This article examines monetary policy when it is constrained by the ZLB on the nominal interest rate. Our analysis uses a nonlinear New Keynesian model with technology and discount factor shocks that captures the expectational effects of falling to and remaining at the ZLB. Discount factor shocks are a proxy for changes in demand that occurred during the Great Recession, while technology shocks account for changes in supply. Specifically, we investigate why technology shocks may have unconventional effects at the ZLB, what factors affect the likelihood of hitting the ZLB, and the tradeoffs a central bank faces under a dual mandate.

We initially focus on a New Keynesian model without capital and then study that model with capital to draw comparisons. In the model without capital, positive technology shocks may have unconventional effects at the ZLB, depending on which measure of output is targeted in the monetary policy rule. When the central bank targets steady-state output, positive technology shocks can cause output to decline when the ZLB binds. Those unconventional dynamics, however, nearly disappear when the central bank targets potential output, which is the level of output in our model with flexible prices. In that case, only large technology shocks reduce output when the ZLB binds.[1]

We primarily focus on the specification in which the central bank targets steady-state output because policymakers, in the short-to-medium term, assume potential output grows at a fairly constant rate [Basu and Fernald (2009)]. Various measures of potential output are often revised following incoming information about shocks, but the revisions occur well after the temporary economic effects from changes in monetary policy have dissipated. Furthermore, Orphanides (2003a,b) and Orphanides and van Norden (2002) document that historically neither the Fed nor statistical methods have been able to detect changes in potential output until long after they have occurred.

Most of the ZLB literature uses models without capital.[2] Capital, however, provides households with another margin to smooth consumption, which strengthens the expectational effects of the ZLB. Arbitrage implies the real interest rate equals the expected future real rental rate of capital. The decline in demand when the ZLB binds leads to a sharp reduction in the rental rate of capital. Therefore, households place increasing weight on the possibility of a lower future rental rate as the policy rate approaches zero, which causes sharper declines in the real interest rate before the ZLB binds. We also include capital adjustment costs to dampen investment volatility. That feature makes investment less attractive as a consumption smoothing mechanism, which causes a greater reduction in consumption and a larger increase in the real interest rate at the ZLB. When the central bank targets steady-state output, a positive technology shock at the ZLB produces more pronounced unconventional dynamics in our model with capital than in the model without capital.

We also evaluate how alternative monetary policy rules affect the likelihood of hitting the ZLB and the efficacy of stabilization policy. A policy rule based on a dual mandate is more likely to cause ZLB events when the central bank targets steady-state output in our model without capital. The opposite result occurs when the central bank targets potential output.[3] When technology is constant, an aggressive response by the central bank to steady-state output decreases the frequency of ZLB events in our model without capital but increases the frequency in our model with capital. Therefore, the frequency of ZLB events depends on (1) the measure of output the Fed targets; (2) the strength of the Fed's response to output; and (3) the sources of exogenous shocks in the model.

Any analysis of the ZLB is complicated by the kink that it imposes on the monetary policy rule. The literature has used a variety of techniques to address this problem. Many papers separate the problem into pre- and post-ZLB periods [e.g., Braun and Körber (2011); Braun and Waki (2006); Christiano et al. (2011); Eggertsson and Woodford (2003); Erceg and Lindé (2014); Gertler and Karadi (2011)]. With that approach, a specific sequence of shocks pushes the nominal interest rate

[1] Wieland (2014) uses structural VAR evidence to argue that these unconventional dynamics did not occur following the 2011 earthquake/tsunami in Japan or the recent oil supply shocks. Braun and Waki (2006) show that technology shocks generate unconventional dynamics at the ZLB in a log-linearized model with capital accumulation where monetary policy targets steady-state output. Using a nonlinear model with capital, Braun and Körber (2011) show that these unconventional dynamics may disappear if the expected duration at the ZLB is short enough.

[2] Braun and Körber (2011), Braun and Waki (2006), and Christiano (2004) are three notable exceptions.

[3] Several papers discuss optimal policy with a ZLB constraint and provide analysis of the welfare losses at the ZLB [Eggertsson and Woodford (2003); Günter et al. (2004); Jung et al. (2005); Nakov (2008); Werning (2011)]. For example, Adam and Billi (2006) find that it is optimal to reduce the nominal interest rate more aggressively in response to adverse shocks in models with a ZLB constraint, despite the welfare consequences that occur at the ZLB.

to zero. Each period, some positive probability exists that the nominal interest rate will exit the ZLB. Once that happens, the nominal interest rate can never fall back to zero. Those simplifying assumptions are made for computational tractability. The drawback is that if a shock causes the ZLB to bind in one period, the same shock will not cause the ZLB to bind in any future period.

Most studies of the ZLB linearize all of their equations with the exception of the monetary policy rule around their non-stochastic steady states. Such a procedure, however, can generate approximation errors. Braun et al. (2012) and Fernández-Villaverde et al. (2012) provide examples of the mistakes resulting from linearized models without capital evaluated at the ZLB. Braun et al. (2012) also argue that linearized models often lead to incorrect inferences about existence and uniqueness of the equilibrium and the local dynamics of the model. Our findings indicate the constrained linear model provides a good approximation of the nonlinear model without capital, but the errors are much larger in a model with capital.[4] In other words, the simulated moments and model predictions are different in the linearized model with capital than in the nonlinear model.

Our paper avoids the problems associated with linearization by obtaining the nonlinear solution to standard New Keynesian models that include an occasionally binding ZLB constraint on the nominal interest rate.[5] Rather than focus on specific sequences of shocks, we calculate the solution for all combinations of discount factor and technology shocks and then provide a thorough explanation of how dynamics change across the state space. Our nonlinear solution method emphasizes accuracy to capture important expectational effects of going to and returning from the ZLB.

The paper proceeds as follows. Section 2 outlines our models with and without capital. Section 3 describes the calibration and solution method, and sections 4 through 6 present the results. These sections report the model solutions across all technology and discount factor shocks, the dynamics at the ZLB, and the likelihood of hitting the ZLB. We also explain how the monetary policy rule impacts those results and provide a comparison between the New Keynesian models with and without capital. Lastly, we present new evidence that the solutions to the constrained linear and nonlinear models are significantly different in the model with capital. Section 7 concludes.

2 ECONOMIC MODELS

This section presents two New Keynesian models with Rotemberg (1982) price adjustment costs. Both models assume stochastic processes for the discount factor and technology, but they differ in their treatment of capital. Model 1 does not include capital while Model 2 does.

2.1 MODEL 1: BASELINE

A representative household chooses $\{c_t, n_t, b_t\}_{t=0}^{\infty}$ to maximize expected lifetime utility given by $E_0 \sum_{t=0}^{\infty} \tilde{\beta}_t[\log c_t - \chi n_t^{1+\eta}/(1+\eta)]$, where $1/\eta$ is the Frisch elasticity of labor supply, c_t is consumption, n_t is labor hours, b_t is a 1-period real bond, E_0 is an expectation operator conditional on information available in period 0, $\tilde{\beta}_0 \equiv 1$, and $\tilde{\beta}_t = \prod_{j=1}^{t} \beta_j$ for $t > 0$. β_j is a time-varying subjective discount factor that evolves according to

$$\beta_j = \bar{\beta}(\beta_{j-1}/\bar{\beta})^{\rho_\beta} \exp(\varepsilon_j), \tag{1}$$

[4] Braun and Waki (2010) show that the approximation error in a perfect-foresight version of a linear model with capital where monetary policy does not respond to output overstates the government spending multiplier.

[5] Several recent papers study the ZLB using nonlinear solution methods. Fernández-Villaverde et al. (2012) calculate the probabilities of ZLB events. Wolman (2005) shows the real effects of the ZLB depend on the policy rule and nominal rigidities. Gust et al. (2013) estimate the extent to which the ZLB constrained the central bank. Aruoba and Schorfheide (2013) and Mertens and Ravn (2014) show how the ZLB affects fiscal multipliers and Basu and Bundick (2012) and Nakata (2012) show the ZLB magnifies the effect of uncertainty on aggregate demand.

where $\bar{\beta}$ is the steady-state discount factor, $0 \leq \rho_\beta < 1$, and $\varepsilon_j \sim \mathbb{N}(0, \sigma_\varepsilon^2)$. Those choices are constrained by $c_t + b_t = w_t n_t + r_{t-1} b_{t-1}/\pi_t + d_t$, where $\pi_t = p_t/p_{t-1}$ is the gross inflation rate, w_t is the real wage rate, r_t is the gross nominal interest rate, and d_t are profits from intermediate firms. The optimality conditions to the household's problem imply

$$w_t = \chi n_t^\eta c_t, \tag{2}$$
$$1 = r_t E_t[\beta_{t+1}(c_t/c_{t+1})/\pi_{t+1}]. \tag{3}$$

The production sector consists of monopolistically competitive intermediate goods firms who produce a continuum of differentiated inputs and a representative final goods firm. Each firm $f \in [0,1]$ in the intermediate goods sector produces a differentiated good, $y_t(f)$, with identical technologies given by $y_t(f) = z_t n_t(f)$, where $n_t(f)$ is the level of employment used by firm f. z_t represents the level of technology, which is common across firms and follows

$$z_t = \bar{z}(z_{t-1}/\bar{z})^{\rho_z} \exp(\upsilon_t), \tag{4}$$

where $\bar{z}$ is steady-state technology, $0 \leq \rho_z < 1$, and $\upsilon_t \sim \mathbb{N}(0, \sigma_\upsilon^2)$. Each intermediate firm chooses its labor supply to minimize its operating costs, $w_t n_t(f)$, subject to its production function. The final goods firm purchases $y_t(f)$ units from each intermediate goods firm to produce the final good, $y_t \equiv [\int_0^1 y_t(f)^{(\theta-1)/\theta} di]^{\theta/(\theta-1)}$ according to a Dixit and Stiglitz (1977) aggregator, where $\theta > 1$ measures the elasticity of substitution between the intermediate goods. The optimality condition to the firm's profit maximization problem then yields the demand function for intermediate inputs given by $y_t(f) = (p_t(f)/p_t)^{-\theta} y_t$, where $p_t = [\int_0^1 p_t(f)^{1-\theta} di]^{1/(1-\theta)}$ is the price of the final good.

Following Rotemberg (1982), each firm faces a cost to adjusting its price, $adj_t(f)$, which emphasizes the negative effect that price changes can have on customer-firm relationships. Using the functional form in Ireland (1997), $adj_t(f) = \varphi[p_t(f)/(\bar{\pi} p_{t-1}(f)) - 1]^2 y_t/2$, the real profits of firm f are $d_t(f) = (p_t(f)/p_t) y_t(f) - w_t n_t(f) - adj_t(f)$, where $\varphi \geq 0$ scales the size of the adjustment costs and $\bar{\pi}$ is the steady-state gross inflation rate. Firm f chooses its price, $p_t(f)$, to maximize the expected discounted present value of real profits $E_t \sum_{k=t}^\infty \lambda_{t,k} d_k(f)$, where $\lambda_{t,t} \equiv 1$, $\lambda_{t,t+1} = \beta_{t+1}(c_t/c_{t+1})$ is the pricing kernel between periods t and $t+1$ and $\lambda_{t,k} \equiv \prod_{j=t+1}^k \lambda_{j-1,j}$. In a symmetric equilibrium, all firms make identical decisions and the optimality condition implies

$$\varphi\left(\frac{\pi_t}{\bar{\pi}} - 1\right)\frac{\pi_t}{\bar{\pi}} = (1-\theta) + \theta\Psi_t + \varphi E_t\left[\lambda_{t,t+1}\left(\frac{\pi_{t+1}}{\bar{\pi}} - 1\right)\frac{\pi_{t+1}}{\bar{\pi}}\frac{y_{t+1}}{y_t}\right], \tag{5}$$

where $\Psi_t = w_t/z_t$ is the real marginal cost. In the absence of price adjustment costs (i.e., $\varphi = 0$), $\Psi_t = (\theta-1)/\theta$, which is the inverse of a firm's markup of price over marginal cost.

Each period, the central bank sets the gross nominal interest rate according to

$$r_t = \max\{1, r^*(\pi_t/\pi^*)^{\phi_\pi}(y_t/y_t^*)^{\phi_y}\}, \tag{6}$$

where $\pi^* = \bar{\pi}$ is the inflation rate target and ϕ_π and ϕ_y are the policy responses to inflation and output.[6] The output target is either steady-state output, $y_t^* = \bar{y}$, or potential output, $y_t^* = y_t^n = (\chi\mu)^{-1/(1+\eta)} z_t$, which is the level of output when $\varphi = 0$. We also examine the case where $\phi_y = 0$.

[6] Although the policy rate cannot fall below zero, the same dynamics would occur if it is a small but positive value. The key point is that a lower bound exists, which prevents the Fed from adjusting the policy rate to adverse shocks.

The resource constraint is given by $c_t = y_t - adj_t \equiv y_t^{adj}$, where y_t^{adj} includes the value added by intermediate firms, which is their output minus quadratic price adjustment costs. A competitive equilibrium consists of sequences of quantities $\{c_t, n_t, b_t, y_t\}_{t=0}^{\infty}$, prices $\{w_t, r_t, \pi_t\}_{t=0}^{\infty}$, and exogenous variables $\{\beta_t, z_t\}_{t=0}^{\infty}$ that satisfy the household's and firm's optimality conditions (2), (3), and (5), the production function, $y_t = z_t n_t$, the monetary policy rule (6), the stochastic processes (1) and (4), the bond market clearing condition, $b_t = 0$, and the resource constraint.

2.2 MODEL 2: BASELINE WITH CAPITAL Model 2 adds capital accumulation to Model 1. The household chooses sequences $\{c_t, i_t, n_t, b_t\}_{t=0}^{\infty}$ to maximize the preferences in Model 1 subject to

$$c_t + i_t + \Phi(i_t/k_{t-1})k_{t-1} + b_t = w_t n_t + r_t^k k_{t-1} + r_{t-1} b_{t-1}/\pi_t + d_t, \tag{7}$$

$$k_t = (1-\delta)k_{t-1} + i_t, \tag{8}$$

where i_t is investment, k_t is capital, r_t^k is the real rental rate of capital, and $\Phi(\cdot)$ is a positive, increasing, and convex function that measures the cost of adjusting the capital stock. We assume $\Phi(x) = \phi(x-\delta)^2/2$, where ϕ controls the size of the adjustment cost. Although other papers utilize alternative specifications of capital/investment adjustment costs, we use this specification because it does not add another state variable to our model, which allows us to present the complete model solution. Optimality yields an equation for Tobin's q and a consumption Euler equation given by

$$q_t = 1 + \phi(i_t/k_{t-1} - \delta), \tag{9}$$

$$q_t = E_t\left[\beta_{t+1}\frac{c_t}{c_{t+1}}\left(r_{t+1}^k - \frac{\phi}{2}\left(\frac{i_{t+1}}{k_t} - \delta\right)^2 + \phi\left(\frac{i_{t+1}}{k_t} - \delta\right)\frac{i_{t+1}}{k_t} + (1-\delta)q_{t+1}\right)\right]. \tag{10}$$

Intermediate firm $f \in [0,1]$ produces a differentiated good, $y_t(f)$, according to $y_t(f) = z_t k_{t-1}(f)^\alpha n_t(f)^{1-\alpha}$, where $k_t(f)$ and $n_t(f)$ are the levels of capital and employment used by firm f. Each intermediate firm then chooses its inputs to minimize operating costs, $r_t^k k_{t-1}(f) + w_t n_t(f)$, subject to its production function, which yields a consolidated optimality condition given by

$$\alpha w_t n_t = (1-\alpha) r_t^k k_{t-1}. \tag{11}$$

The firm pricing equation (5) remains unchanged, except that $\Psi_t = w_t^{1-\alpha}(r_t^k)^\alpha/[z_t(1-\alpha)^{1-\alpha}\alpha^\alpha]$.

The resource constraint includes the output lost due to both price and capital adjustment costs and is given by $c_t + i_t + \Phi(i_t/k_{t-1})k_{t-1} = y_t^{adj}$. A competitive equilibrium consists of sequences of quantities $\{c_t, n_t, i_t, k_t, b_t, y_t\}_{t=0}^{\infty}$, prices $\{w_t, r_t^k, r_t, \pi_t, q_t\}_{t=0}^{\infty}$, and exogenous variables $\{\beta_t, z_t\}_{t=0}^{\infty}$ that satisfy the household's and firm's optimality conditions (2), (3), (5), (9), (10), and (11), the production function, $y_t = z_t k_{t-1}^\alpha n_t^{1-\alpha}$, the monetary policy rule (6), the stochastic processes (1) and (4), the capital law of motion (8), bond market clearing, $b_t = 0$, and the resource constraint.

3 CALIBRATION, SOLUTION METHOD, AND SIMULATION PROCEDURE

3.1 CALIBRATION We calibrate the models in section 2 at a quarterly frequency using common values in the monetary policy literature. The parameters are shown in table 1. The annual real interest rate is set to 2%, which implies a steady-state quarterly discount factor, $\bar{\beta}$, equal to 0.995. Those values correspond to the ratio of the federal funds rate to the percent change in the GDP deflator from 1983-2007. The Frisch elasticity of labor supply, $1/\eta$, is set to 3, which is consistent

Frisch Elasticity of Labor Supply	$1/\eta$	3	Inflation Coefficient: MP Rule	ϕ_π	1.5
Elasticity of Substitution between Goods	θ	6	Output Coefficient: MP Rule	ϕ_y	0.1
Rotemberg Adjustment Cost Coefficient	φ	59.11	Steady-State Technology	$\bar{z}$	1
Steady-State Labor	$\bar{n}$	0.33	Technology Persistence	ρ_z	0.9
Capital Depreciation Rate†	δ	0.025	Technology Shock Standard Deviation	σ_υ	0.0025
Cost Share of Capital†	α	0.33	Steady-State Discount Factor	$\bar{\beta}$	0.995
Capital Adjustment Cost†	ϕ	5.6	Discount Factor Persistence	ρ_β	0.8
Steady-State Inflation	$\bar{\pi}$	1.006	Discount Factor Standard Deviation	σ_ε	0.0025

Table 1: Baseline calibration. A † denotes a parameter that only applies to Model 2.

with Peterman (2012). The leisure preference parameter, χ, is calibrated so that steady-state labor equals $1/3$ of the available time. Capital's share of output, α, is set to 0.33 and the quarterly depreciation rate, δ, equals 2.5%. The capital adjustment cost parameter, ϕ, is set to 5.6, which follows Eberly (1997) and Erceg and Levin (2003). The elasticity of substitution between intermediate goods, θ, is set to 6, which corresponds to an average markup of price over marginal cost equal to 20%. The price adjustment cost parameter, φ, is set to 59.11, which is consistent with a Calvo (1983) price-setting specification where prices change on average once every four quarters.

The steady-state gross inflation rate, $\bar{\pi}$, is set to 1.006, which implies an annual inflation rate target of 2.4%. That value equals the average growth rate of the U.S. PCE chain-type price index from 1983-2007. In our baseline calibration, we set the coefficients on inflation and output in the monetary policy rule to 1.5 and 0.1, respectively, but we also consider other values.

The likelihood that the nominal interest rate falls to and remains at zero depends on both the parameters of the discount factor and technology processes. Richter and Throckmorton (2014) show a clear tradeoff exists between the persistence and the standard deviation of the stochastic shock processes. As the persistence of a process increases, the standard deviation of that shock must decline, otherwise our numerical algorithm will not converge to a minimum state variable (MSV) solution. The failure to converge occurs because the economy either remains at the ZLB too long when the shocks are very persistent or falls to the ZLB too frequently when the processes are highly volatile. We chose the discount factor and technology parameters so (1) They are constant across all models; (2) They generate ZLB events when simulating the model; and (3) They match the data as closely as possible. Specifically, we set the persistence of the discount factor, ρ_β, equal to 0.8 and the standard deviation of the shock, σ_ε, equal to 0.0025. Those values follow Fernández-Villaverde et al. (2012) who assume that a discount factor shock has a half life of about 3 quarters. Steady-state technology, $\bar{z}$, is normalized to 1, the persistence of the technology shock, ρ_z, is 0.9, and the standard deviation of the shock, σ_υ, equals 0.0025. In the data, deviations of log real GDP from trend are 1.85% per quarter and deviations of the log difference in the PCE price index are 0.29% from 1983-2007. The equivalent values in our models are smaller than is observed since additional real world shocks and sources of persistence are needed to match the data.

3.2 SOLUTION METHOD The model is solved using the policy function iteration algorithm described in Richter et al. (2014), which is based on the theoretical work on monotone operators in Coleman (1991). This solution method discretizes the state space and uses time iteration to solve for the updated decision rules until the tolerance criterion is met. We use piecewise linear interpolation to approximate future variables, since this approach more accurately captures the kink in the decision rules than continuous approximating functions, and then use Gauss-Hermite quadrature to numerically integrate. Those techniques capture the expectational effects of going to and returning to the ZLB. For a formal description of the numerical algorithm see appendix A.

Benhabib et al.'s (2001) finding that constrained New Keynesian models have two deterministic steady-state equilibria has generated considerable discussion in the literature about whether there are conditions in which a unique MSV solution exists in stochastic models with a ZLB constraint. Specifically, they find there are two nominal interest rate/inflation rate pairs consistent with the steady-state equilibrium system. In one steady state, the central bank meets its positive inflation target, whereas in the other steady state the economy experiences deflation. Richter and Throckmorton (2014) show that the numerical algorithm used in our paper converges to the inflationary equilibrium as long as there is a sufficient expectation of returning to a monetary policy rule that conforms to the Taylor principle.[7] The algorithm, however, never converges to the deflationary equilibrium because it is unstable. That is, the algorithm does not converge to the deflationary equilibrium after a shock pushes the economy away from that equilibrium, which is similar to findings in Christiano and Eichenbaum (2012) and Wolman (2005).

The intuition for how our algorithm behaves can be discerned from the simple three-equation linear New Keynesian model. We know determinacy in this model depends on whether the Taylor principle holds (i.e., the nominal interest rate moves more than one-for-one with inflation), assuming the fiscal authority ensures stable debt dynamics (i.e., passive fiscal policy). If the Taylor principle holds, our algorithm converges to the unique MSV solution that can be analytically derived. When the Taylor principle does not hold (i.e., passive monetary policy), our algorithm will not converge, even though the model has many solutions in this case. The only way our algorithm can locate these solutions is if a process for the sunspot shocks is explicitly written down.

The same rationale applies in our model with a ZLB constraint except that there are two types of sunspots. One type is analogous to the sunspots that occur when the Taylor principle does not hold. A pegged nominal interest rate is a special type of passive monetary policy, where the distribution of future shocks is truncated in a stochastic model. Thus, an occasionally binding ZLB constraint is similar to a Taylor rule that switches between an active and passive policy. As long as there is a sufficient expectation of returning to an active monetary policy, our algorithm will converge to the positive inflation equilibrium. If, on the other hand, the expectation of returning to the Taylor rule is not strong enough or the probability of returning to the ZLB is too high, then a stable inflationary equilibrium does not exist and our algorithm will diverge. That finding does not necessarily mean the model has no solutions. Instead, it could indicate many solutions exist, but that finding can not be observed without specifying a process for sunspot shocks.

The other type of sunspot shock is unique to a model with a ZLB constraint. The existence of both an inflationary and a deflationary steady state means the economy could fluctuate between the two steady states. Unless our model specifies how the economy switches between the infla-

[7] Davig and Leeper (2007) examine determinacy in a Fisherian economy that switches between active and passive policy. They prove that as long as one of the regimes satisfies the Taylor principle, the central bank can passively respond to inflation in the other regime and still have a determinate solution. Richter and Throckmorton (2014) show that the convergence region—the region of the parameter space where our algorithm converges to an MSV solution—is identical to the determinacy region Davig and Leeper (2007) derive. This exercise is informative because a model with an occasionally binding ZLB constraint is similar to a model with a monetary policy rule that switches between active and passive policy. Richter and Throckmorton (2014) also examine how the standard deviation of the stochastic processes affect whether the algorithm converges to the inflationary steady state in a model with a ZLB. They find that the boundary of the convergence region imposes a clear tradeoff between the expected frequency and average duration of ZLB events. Therefore, a model with a ZLB constraint produces the same intuition described in Davig and Leeper (2007). As long as the ZLB does not bind too frequently or for too long, our algorithm converges.

tionary and deflationary steady states, it does not converge to the deflationary steady state because that equilibrium is unstable. Since that mechanism is not present in our model, our algorithm is only able to converge to the inflationary steady state. Alternatively, our algorithm would fluctuate between those two steady states if we specify a Markov-switching process as in Aruoba and Schorfheide (2013). Appendix B shows the equilibria that our algorithm converges to in both a deterministic and stochastic version of Model 1. It also provides an example of how our algorithm behaves when the economy fluctuates between the deflationary and inflationary steady states.

Uncertainty exists about whether these sunspot shocks affect an economy with a ZLB constraint. Economists, for example, want to understand if the sunspot shocks are observed in the data or even reflect dynamics that are economically feasible. Although addressing those questions is important for future research, our analysis, like most macroeconomic research, is concentrated on examining solutions that converge to the inflationary steady state.[8]

3.3 SIMULATION PROCEDURE We simulate the model using draws from the distributions for the discount factor and technology shocks. The bounds of the state space are chosen to minimize extrapolation of the policy functions in the simulation. Figure 2 plots the distributions of the state variables and the nominal interest rate in a $500{,}000$ quarter simulation. The vertical axes show the frequency of each realization as a percent of the simulation length. Variables on the horizontal axes are shown as percent deviations from steady state, except the nominal interest rate which is a net percentage. The dashed lines represent the bounds of the state space. The solid lines denote the theoretical unconditional distributions scaled for comparison with the simulated distributions.

Figure 2a shows the unconditional distributions of technology, the discount factor, and the nominal interest rate.[9] The state space for technology lies within $\pm 2.5\%$ of its steady state, which is normalized to unity. The state space for the discount factor lies between $\pm 1.9\%$ of its steady state, which equals 0.995. Across these states, the quarterly net nominal interest rate is distributed over a range of 0% to 3.6%, with a large mass (5% of quarters) between 0 and 20 basis points. As we demonstrate below, model dynamics are very different when the policy rate lies in this interval.

Figure 2b shows the distribution of the discount factor and technology conditional on the ZLB binding. A high discount factor is the primary source of ZLB events, as indicated by the difference between its distribution conditional on the ZLB (bars) and its theoretical unconditional distribution (solid line). The conditional distribution for the discount factor is centered around 1% above steady state. A higher discount factor means households are more willing to postpone their consumption. Lower consumption pushes down inflation, which in turn causes the nominal interest rate to fall. If households are patient enough, then the nominal interest rate hits its ZLB. The nominal interest rate can also fall to zero when technology is sufficiently far above its steady state because higher supply leads to lower prices. Our result is consistent with the finding in Fernández-Villaverde et al. (2012) that high levels of technology are associated with low interest rates. Kiley (2003) uses U.S. data to show that periods of high labor productivity growth have been associated with low inflation and argues that result could be caused by the Fed's policy rule as our models suggest.

The Fed's policy rate has been at its effective ZLB since December 2008. Gust et al. (2013) show that most financial market participants expected the federal funds rate to remain below 25

---

[8] McCallum (2001) argues the deflationary equilibrium is not economically relevant since it is not E-stable. Christiano and Eichenbaum (2012) reiterate this point using a New Keynesian model with a ZLB constraint.

[9] In all of our results, a hat denotes percent deviation from the deterministic steady state (i.e., for some generic variable x in levels, $\hat{x}_t \equiv 100(x_t - \bar{x})/\bar{x}$) and a tilde denotes a net rate (i.e., for some gross rate x, $\tilde{x}_t = 100(x_t - 1)$).

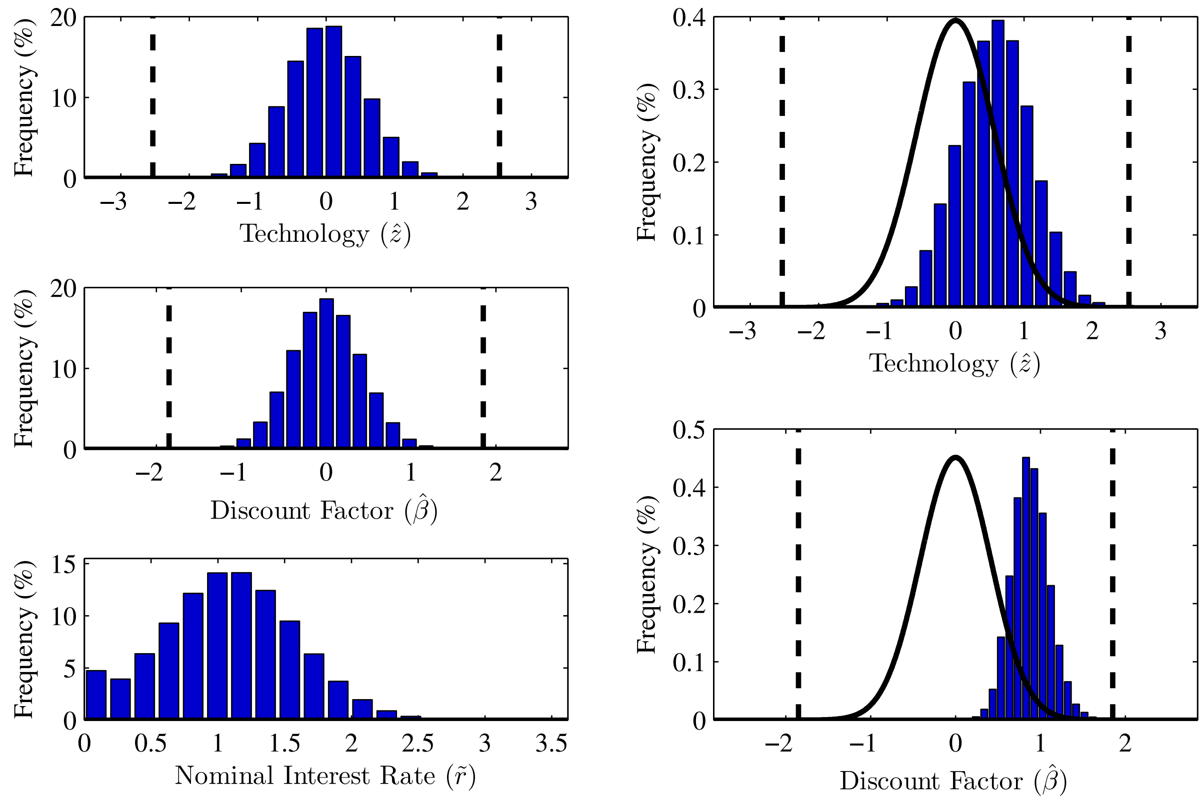

(a) Unconditional distribution

(b) Distribution conditional on the ZLB

Figure 2: Model 1 ($y_t^* = \bar{y}$) distributions as a percentage of a 500,000 quarter simulation. Each variable is in percent deviations from its steady-state value. The dashed lines are the bounds of the state space. The solid lines are the theoretical unconditional distributions of the state variables scaled for comparison with the conditional distributions.

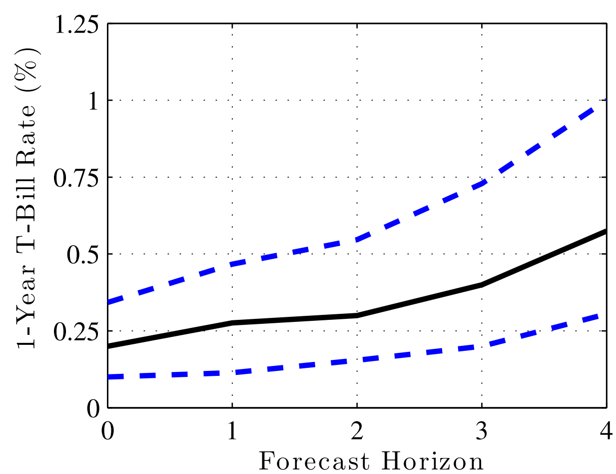

Figure 3: The median and $16/84$ percentiles of individual forecasts of the T-Bill rate in the first quarter of 2009 according to the SPF.

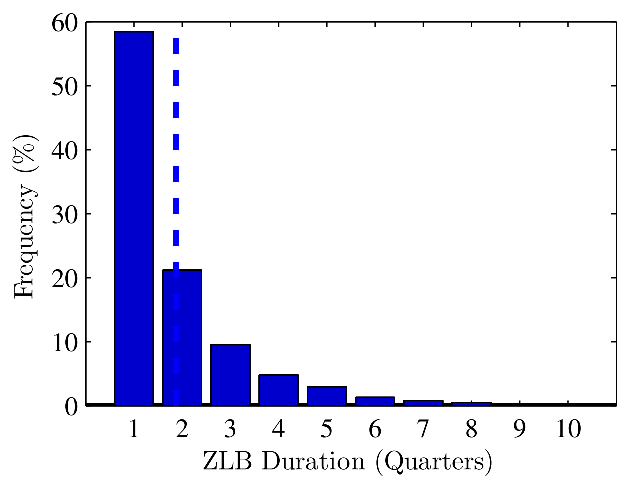

Figure 4: Model 1 ($y_t^* = \bar{y}$) ZLB event durations as a percentage of ZLB events in a 500,000 quarter simulation. The vertical dashed line is the average ZLB duration.

basis points for only a few quarters. For example, the median forecast in the first quarter of 2009 was below 25 basis points only until the third quarter of that year and gradually increased to 2.5% in 2012. The Survey of Professional Forecasters (SPF) conducted by the Philadelphia Fed asks its participants to forecast the 1-year T-Bill rate up to four quarters in the future. The median (solid line) and $16/84$ percentiles (dashed lines) of the individual forecasts in the first quarter of 2009 are shown in figure 3. The median forecast predicted the T-Bill rate would exceed 50 basis points while the 84th percentile predicted it would hit 1% within 1 year. Those forecasts indicate that people expected the ZLB to bind for just a few quarters even though the recession was quite severe.

Our model is calibrated to the average time the ZLB is expected to hold and not to the duration of the current ZLB episode in the U.S. With that being said, it is possible for longer ZLB events to occur in our framework. Figure 4, for example, shows the distribution of the length of each ZLB event as a percentage of the total number of those events in a $500{,}000$ quarter simulation of Model 1 ($y_t^* = \bar{y}$). The vertical dashed line indicates the average ZLB duration is 1.87 quarters. The longest ZLB event is 19 quarters, which is about the length of the current ZLB episode. ZLB events with a duration of 1, 2, and 3 quarters account for 58.4%, 21.2%, and 9.5%, respectively, of all ZLB events in the simulation. Therefore, our calibration of the stochastic processes produces a distribution of ZLB event durations that is similar to household expectations at the onset of the Great Recession. The calibration for Model 2 also yields a similar distribution of ZLB events.

4 MODEL 1: STATES OF THE ECONOMY, ECONOMIC DYNAMICS, AND THE ZLB

The New Keynesian model without capital, outlined in section 2.1, contains two state variables, the discount factor and technology. This section presents the complete solution to Model 1, key cross sections of that solution, impulse responses to technology shocks, and simulation statistics. We compare these results across alternative monetary policy rules. Each variable is shown in percent deviations from its steady state, except inflation and the interest rates, which are net percentages.

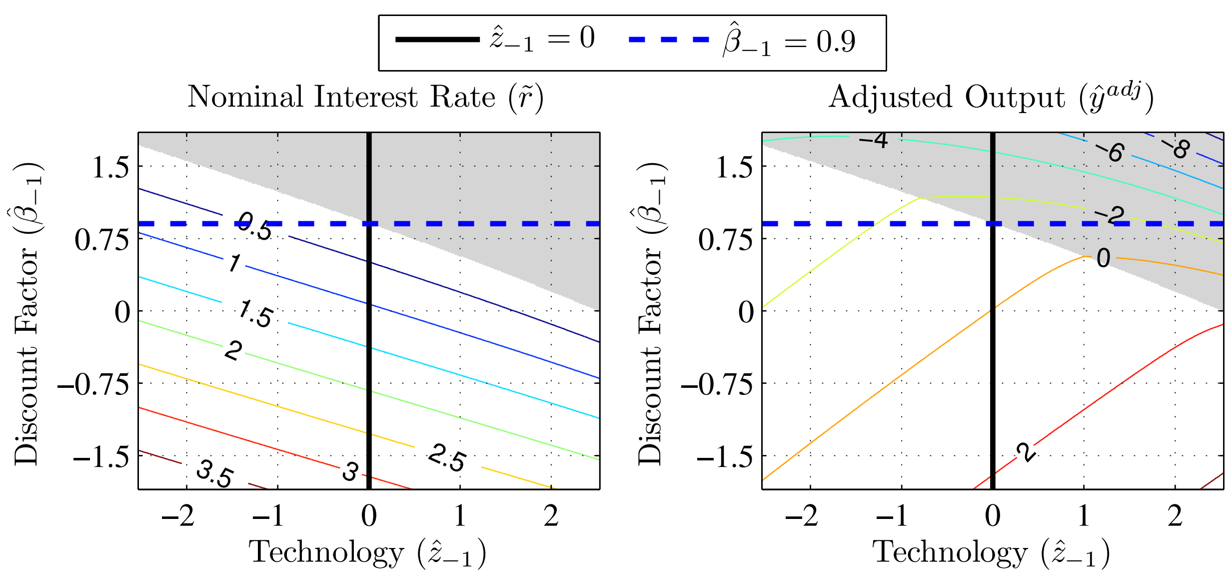

Figure 5: Model 1 ($y_t^* = \bar{y}$) decision rules as a function of the technology ($\hat{z}_{-1}$) and the discount factor ($\hat{\beta}_{-1}$) states. Each variable is in percent deviations from its deterministic steady state, except the nominal interest rate, which is a net percentage. The shaded region indicates where the ZLB binds.

Figure 5 shows three-dimensional contour plots of the net nominal interest rate and adjusted output over the entire state space. These plots provide a complete picture of the model solution for both variables when the central bank targets steady-state output ($y_t^* = \bar{y}$). The shaded areas represent the states of the economy where the net nominal interest rate, $\tilde{r}$, equals zero. Those areas reveal the nominal interest rate only hits the ZLB when either technology or the discount factor are unusually high. When the central bank targets steady-state output, a higher level of technology lowers inflation and the real interest rate when the ZLB does not bind. When the ZLB binds, higher technology continues to push down inflation, which forces up the real interest rate and causes demand to fall. Looking at the highest discount factor in figure 5, output exhibits the same unconventional response, even when technology is at or below its steady state. In fact, many studies assume an elevated discount factor is the cause of the current ZLB event in the U.S.

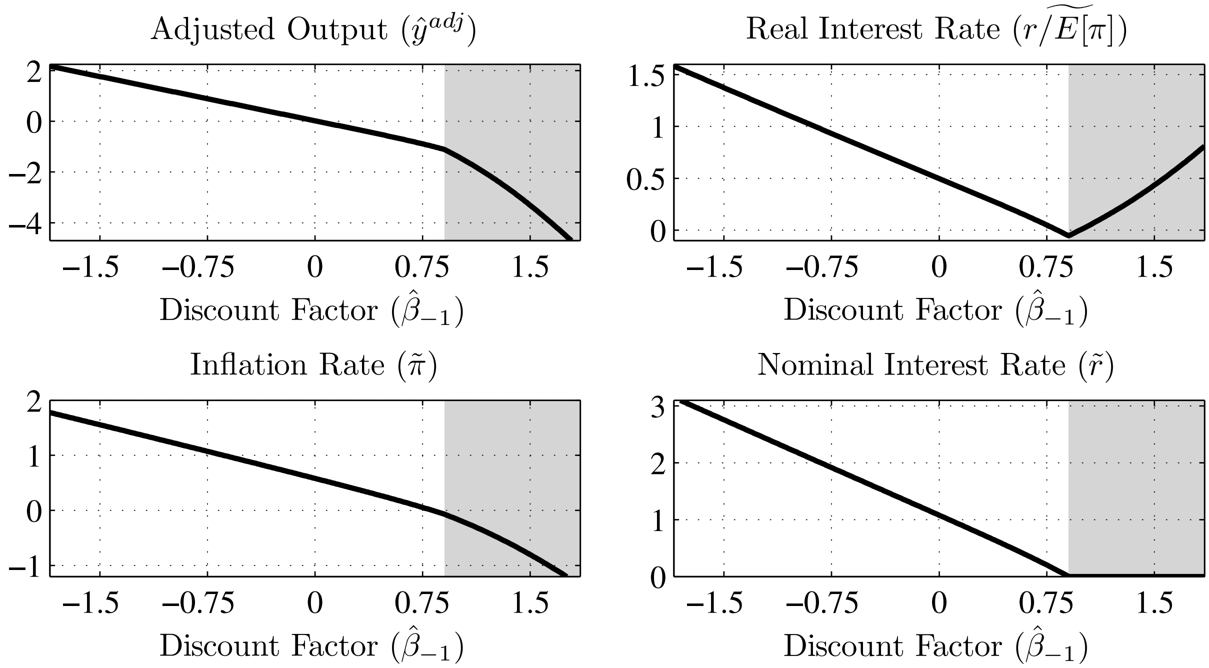

Figure 6: Model 1 ($y_t^* = \bar{y}$) decision rules as a function of the discount factor state ($\hat{\beta}_{-1}$). The technology state is fixed at its steady-state value ($\hat{z}_{-1} = 0$). Each variable is in percent deviations from its deterministic steady state, except inflation and the interest rates, which are net percentages. The shaded region indicates where the ZLB binds.

The contours in figure 5 are useful because they provide the solution for every combination of the two shocks, but they can be difficult to read. Therefore, we focus on specific cross sections of the state space. The solid line in figure 5 shows the cross section where the technology state is held constant at its steady state ($\hat{z}_{-1} = 0$). Two-dimensional representations of that cross section are shown in figure 6. The shaded region highlights where the ZLB binds, which begins when the discount factor is 0.9% above its steady state. A high discount factor indicates that households have a strong desire to save. Elevated savings depresses demand, which reduces output, inflation, and the nominal interest rate. At the ZLB, any further reduction in expected inflation is offset by an equal increase in the real interest rate. That higher real interest rate raises the cost of current consumption which further lowers demand in discount factor states where the ZLB binds.

The dashed line in figure 5 shows the cross section where the discount factor is held constant at 0.9% above its steady-state value ($\hat{\beta}_{-1} = 0.9$), which is the minimum value where the ZLB binds when technology is at its steady state. Figure 7 shows a two-dimensional representation of that cross section along with the same cross sections for different values of ϕ_y. The entire shaded region indicates where the ZLB binds when $\hat{z}_{-1} = 0$ and $\phi_y = 0$. Larger values of ϕ_y cause the ZLB to first bind in slightly higher technology states, as the darker shaded regions show. The unconventional response of the economy to a positive technology shock is smaller as the value of ϕ_y declines. With $\phi_y = 0.05$ ($\phi_y = 0$), the response of output is positive in technology states up to 0.75% (1.4%) above its steady state. Furthermore, in the high technology states where the economy does contract, output and inflation are more stable with a lower ϕ_y. For example, when $\phi_y = 0$, output never falls below its initial ZLB level ($\hat{y}^{adj} = -1.25\%$), even in the highest technology state. In contrast, when $\phi_y = 0.1$, output falls from -1% to -3% when technology increases from $\hat{z}_{-1} = 0$ to $\hat{z}_{-1} = 2.5$. It is clear from those results that a shorter expected duration at the ZLB can reverse the unconventional dynamics, since the expected duration of the ZLB increases in higher technology states. That finding is consistent with the conclusions of Braun and Körber (2011). It is also apparent that the monetary policy rule plays an important role in the dynamics at the ZLB since the slopes of the policy function differ greatly across the alternative values of ϕ_y.

To better understand those results, we begin by examining the region of the state space where the ZLB does not bind. In low technology states, workers are less productive and firms' per unit marginal cost of production is higher. Firms respond by raising prices and reducing their demand for labor. With less output available for consumption, the household wants to work more to moderate the decline in consumption. The higher labor supply dominates the drop in labor demand, so the equilibrium level of labor is higher and the real wage rate is lower. The household also believes technology will slowly return to its steady state and as a result, expects its future consumption to increase. Higher expected future consumption is reflected in an elevated real interest rate. A larger value of ϕ_y in technology states where the ZLB does not bind keeps output, labor, and the real wage rate closer to their steady states, but that additional stability comes at the expense of more inflation and a higher nominal interest rate. The real interest rate in that case is mostly unaffected.

The last area to consider are the technology states where the ZLB binds. In those states, higher technology continues to lower per unit production costs and firms react by lowering their prices. The additional decline in expected inflation when the nominal interest rate equals zero raises the real interest rate. The household reduces its consumption and increases it labor supply to capitalize on the higher returns which results in the paradox of thrift. Aggregate demand falls because everyone wants to save more at the higher real interest rate, but that is not possible in equilibrium. Thus, the lower demand reduces output until actual and desired savings are equal. Firms respond to the decrease in demand by further lowering prices and cutting labor demand. The drop in labor demand dominates the increase in labor supply, so that both total hours and the real wage decline. This is an example of the paradox of toil [Eggertsson (2010)]. At the ZLB, everyone wants to work more, but the higher real interest rate lowers demand, which causes firms to reduce employment.[10]

With a smaller response to the deviations from steady-state output, inflation is more stable in all technology states. Thus, the real interest rate rises less at the ZLB, which helps maintain household demand in high technology states. Higher labor demand raises equilibrium hours, which mitigates the decline in the real wage. In short, a tension exists at the ZLB between the supply-side effects

[10] The standard deviation of the stochastic processes affects the expected frequency and average duration of the ZLB. Appendix C shows that the qualitative effects of a larger ϕ_y in figure 7 are similar when $\rho_\beta = 0.75$.

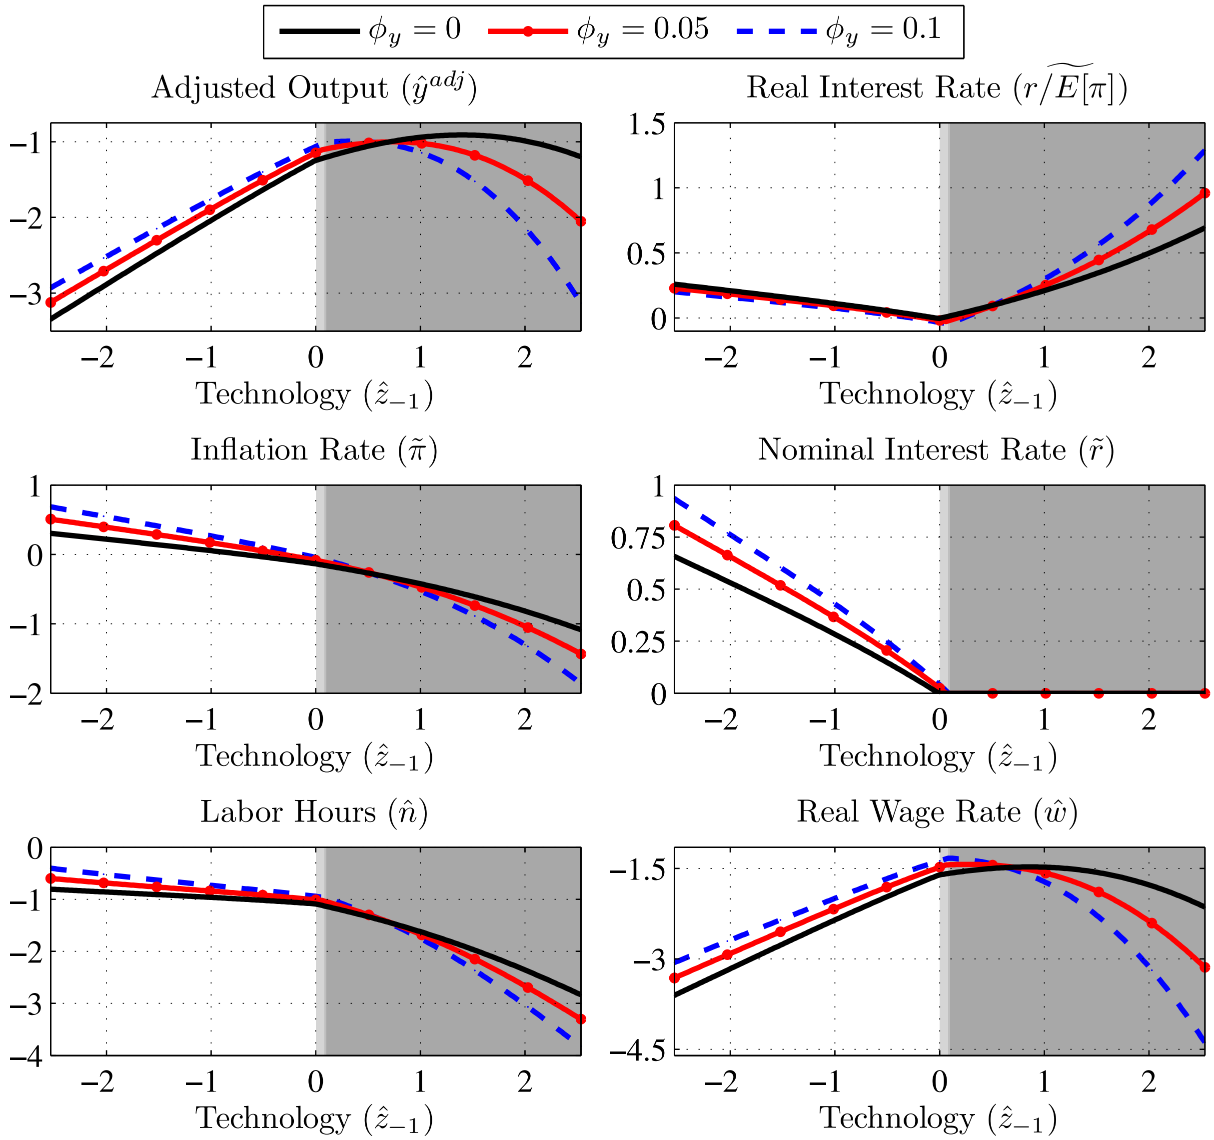

Figure 7: Model 1 ($y_t^* = \bar{y}$) decision rules as a function of the technology state ($\hat{z}_{-1}$). The discount factor state ($\hat{\beta}_{-1}$) is fixed at the minimum value that causes the ZLB to bind when $\hat{z}_{-1} = 0$ and $\phi_y = 0$. Each variable is in percent deviations from its deterministic steady state, except inflation and the interest rates, which are net percentages. The shaded region indicates where the ZLB binds for a given ϕ_y value.

of technology and the demand-side effects of the real interest rate. If the central bank responds less aggressively to the deviations from steady-state output when the ZLB does not bind, then the demand-side effects at the ZLB are weaker and both real and nominal variables are less volatile.

We also examine the effects of technology shocks by computing generalized impulse response functions (GIRFs) of a policy shock. GIRFs provide a clear quantitative comparison between economic dynamics at and away from the ZLB. They are based on an average of model simulations where the realization of shocks is consistent with households' expectations over time. Figure 8 plots the generalized impulse responses to a 1% positive technology shock when the central bank targets steady-state output under two sets of initial conditions: (1) a non-ZLB case (solid line), where the discount factor remains at its steady state so that the nominal interest rate is above its ZLB; and (2) a ZLB case (dashed line), where the discount factor is set to its mean value over a $500{,}000$ quarter simulation under the condition that the ZLB binds and technology is at its steady state. To compute the GIRFs, we calculate the mean of $10{,}000$ simulations of the model conditional on a random shock in the first quarter. We then calculate a second mean from another set of $10{,}000$ simulations, but in this case the shock in the first quarter is replaced with a 1 standard deviation positive technology shock. We compute the percentage change (or difference for the interest rates and inflation) between the two means.[11] Those values are shown on the vertical axis.

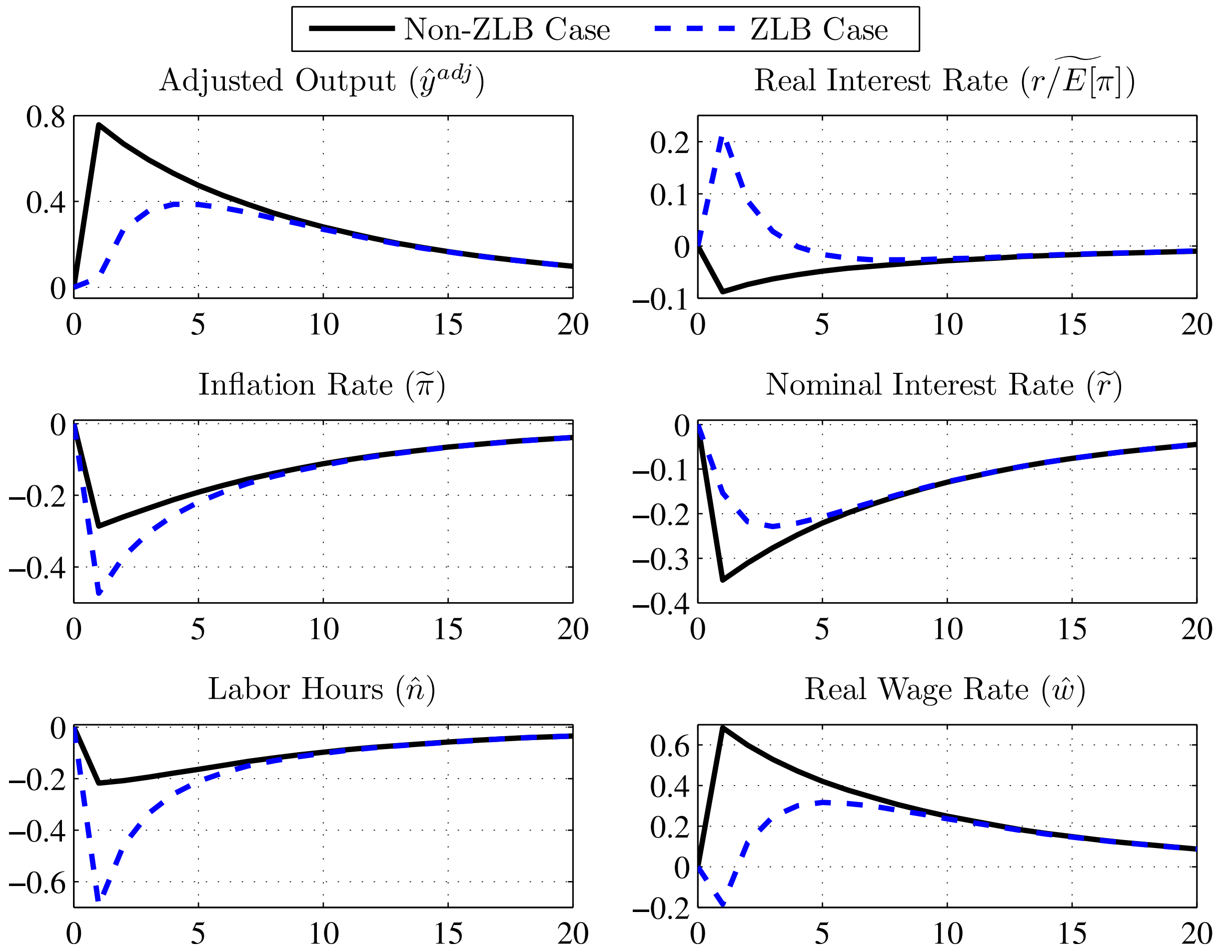

Figure 8: Model 1 ($y_t^* = \bar{y}$) GIRFs to a 1% positive technology shock. The steady-state case (solid line) is initialized at the model's steady state. The ZLB case (dashed line) is initialized at the average state vector conditional on the ZLB binding in a $500{,}000$ quarter simulation.

The impulse responses in the non-ZLB case are standard and follow the intuition from the decision rules. On average, a 1% positive technology shock increases adjusted output, lowers firms' per unit marginal cost of production, and causes inflation and the nominal interest rate to fall. According to the Taylor rule, the nominal interest rate falls more than the inflation rate, so the real interest rate declines, which increases consumption. The positive technology shock also raises productivity, which decreases the equilibrium level of labor and increases the real wage rate.

[11] The general procedure for computing GIRFs is outlined in Koop et al. (1996). See appendix D for details.

In the ZLB case, a 1% positive technology shock initially increases adjusted output by only 0.05% on average. That sluggish response occurs because the ZLB binds in 87% of the simulations after the positive technology shock, which means the nominal interest rate cannot fall by as much as it does in the non-ZLB case. The positive technology shock also lowers per unit production costs which helps to push down prices. With prices falling and the nominal interest rate stuck at zero, an increase in technology sharply raises the real interest rate. That spike then limits the increase in output and causes labor to fall further than in the non-ZLB case. Our results in figure 7, however, indicate that the responses of output and the real wage depend critically on the value of ϕ_y. If $\phi_y = 0$, a positive technology shock increases adjusted output more on impact than when $\phi_y > 0$ because the absence of a policy response to output limits the upward pressure on the real interest rate at the ZLB. From period 2 onward, adjusted output increases as the economy exits the ZLB due to the mean reversion in both technology and the discount factor. The nominal interest rate rises far enough above zero by period 8 that the ZLB case effectively mirrors the non-ZLB case. In both cases, technology returns to its steady state about 20 quarters after the initial shock.

Figure 9 plots the same cross section of the state space that is shown by the dashed line in figure 5 across three monetary policy rules: (1) The central bank does not respond to output ($\phi_y = 0$, solid line); (2) The central bank targets steady-state output ($y_t^* = \bar{y}$, $\phi_y = 0.1$, dashed line); and (3) The central bank targets potential output ($y_t^* = y_t^n$, $\phi_y = 0.1$, circle markers). The shaded region indicates where the ZLB binds, but the level of technology where that occurs depends on the policy rule. When $y_t^* = \bar{y}$ ($y_t^* = y_t^n$), the ZLB first binds when technology is 0.1% (0.25%) above its steady state. The most noteworthy difference among the policy rules is that higher technology states at the ZLB generate further increases in output and the real wage rate when the central bank targets potential output as opposed to a decline when it targets steady-state output. In addition, output falls in 49.4% of the simulations used to compute a GIRF initialized at the ZLB with a steady-state output target but only in 1.8% of the simulations with a potential output target.

Unlike steady-state output, which is constant, potential output positively co-moves with technology. When technology is below its steady state, potential output is lower so that the output gap is less negative than the gap with steady-state output. Since low technology drives up prices, inflation is positively related to the potential output gap but negatively related to the steady-state output gap. Therefore, inflation is more stable when the central bank targets potential output. When technology is above steady state, on the other hand, it lowers inflation and causes the real interest rate to rise at the ZLB. That higher real rate encourages the household to save more, which reduces demand. Lower demand dampens the upward pressure on output from the lower production costs. The effect that dominates depends on whether the real interest rate rises enough to offset the positive effects of higher technology. Given that the real interest rate is inversely related to the expected inflation rate at the ZLB, the real interest rate will rise less when the central bank targets potential output since inflation is more stable. Thus, output (the output gap) will be higher (lower) when the central bank targets potential output as opposed to steady-state output. Those results again demonstrate that both the expected duration of a ZLB event and the specification of the monetary policy rule can impact whether or not unconventional dynamics occur at the ZLB.

The Fed has recently preferred a policy rule that responds to inflation and output. Our findings above indicate that the output target qualitatively affects how the economy responds to high levels of technology. Next, we examine how the output target affects the likelihood of hitting the ZLB

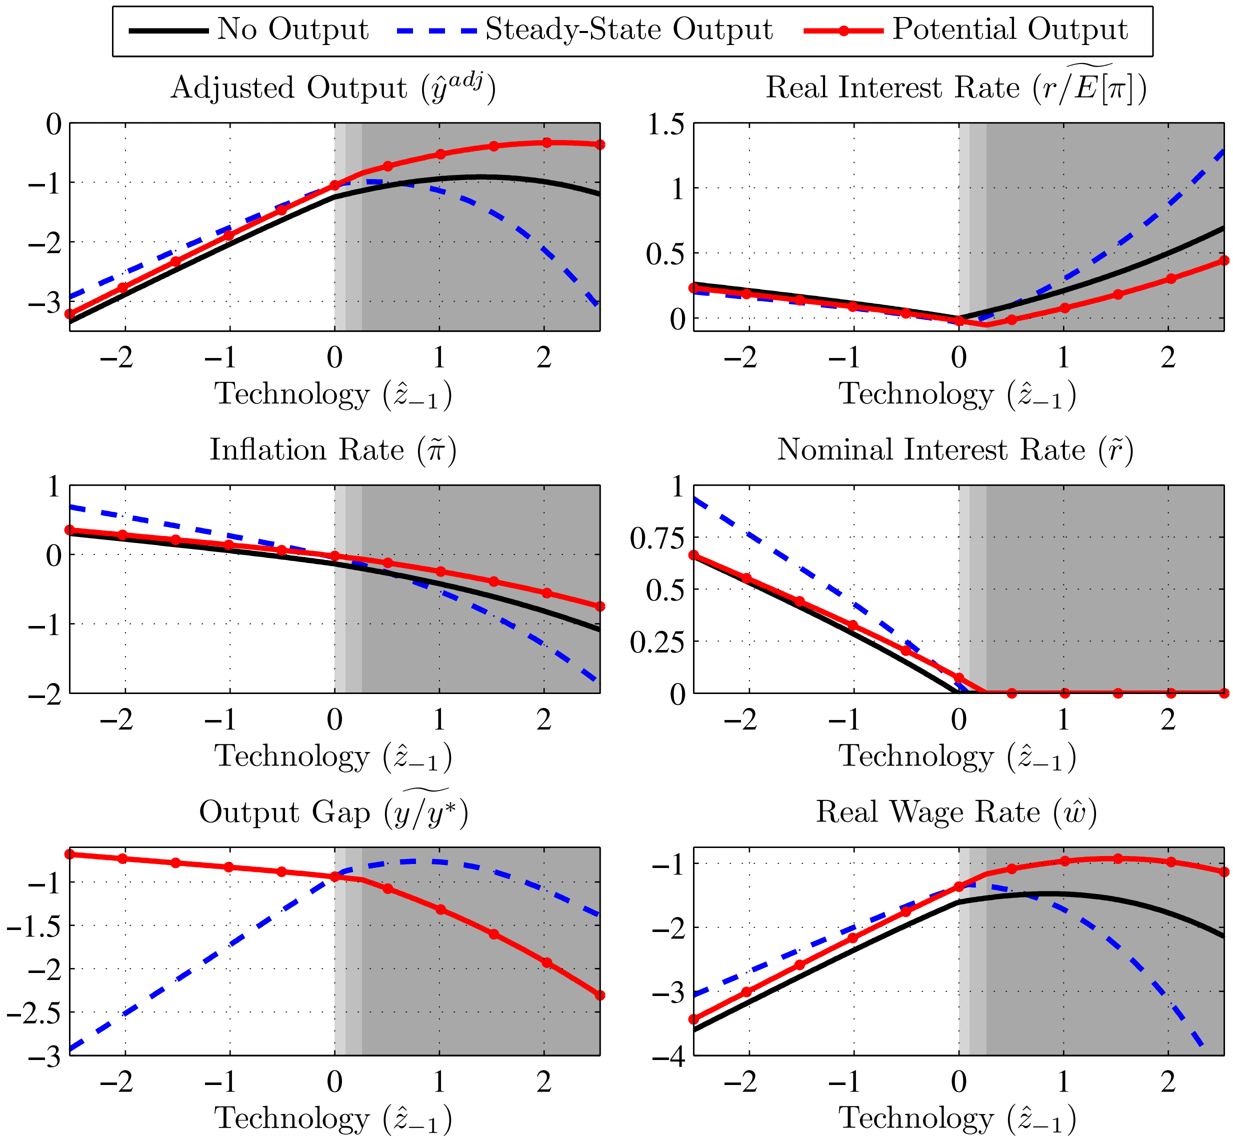

Figure 9: Model 1 decision rules as a function of the technology state ($\hat{z}_{-1}$). The discount factor is fixed at the minimum value that causes the ZLB to bind when $\hat{z}_{-1} = 0$ and $\phi_y = 0$. On the solid line, the central bank does not respond to output ($\phi_y = 0$). On the dashed line, it responds to deviations from steady-state output ($y_t^* = \bar{y}$) and on the line with circle markers it responds to deviations from potential output ($y_t^* = y_t^n$). Each variable is in percent deviations from its deterministic steady state, except inflation and the interest rates, which are net percentages. The shaded region indicates where the ZLB binds for a given specification of monetary policy.

using $500{,}000$ quarter simulations of the model.[12] Our main result is that a more aggressive steady-state (potential) output target will increase (decrease) the likelihood of hitting the ZLB.

Table 2a shows the effect of reducing the weight on output (ϕ_y) while holding the weight on inflation at $\phi_\pi = 1.5$. We begin with the original value in Taylor (1993), $\phi_y = 0.125$, and reduce it by increments of 0.025. With a steady-state output target ($y_t^* = \bar{y}$), the ZLB binds in 2.73% of the

[12]The underlying assumption on the behavior of the discount factor is different in table 2a and figure 7, so any variations in results can be traced to that difference. Specifically, the discount factor in figure 7 is set to the minimum value that causes the ZLB to bind when $z_{-1} = 0$, whereas table 2a is based on random discount factor shocks.

simulated quarters and has an average duration of 1.90 quarters when $\phi_y = 0.125$. These values monotonically decrease with ϕ_y and equal 2.33% and 1.81 quarters when $\phi_y = 0$. Decreasing the weight on the steady-state output target raises the volatility of output but has no meaningful effect on the volatility of inflation. Overall, there is not much of a tradeoff between the volatility of output and inflation. The results are reversed when the central bank targets potential output ($y_t^* = y_t^n$). Placing more weight on the deviations from potential output reduces the likelihood of hitting the ZLB and the standard deviations of output and inflation fall. Those results are consistent with Adam and Billi (2006). They show it is optimal for the central bank to aggressively reduce its policy rate after an adverse shock when it targets potential output. That policy reduces visits to the ZLB and welfare losses. Table 2b reports the results when we fix $\phi_y = 0.125$ and change the response to deviations of inflation from its target (ϕ_π). With $y_t^* = \bar{y}$, the probability of hitting the ZLB falls from 2.73% of the simulated quarters in the baseline case to 0.43% when $\phi_\pi = 3$. In addition, the standard deviations of output and inflation fall as ϕ_π increases. A higher ϕ_π reduces the volatility of output and inflation and decreases the likelihood of hitting the ZLB when $y_t^* = y_t^n$. In the baseline case, the ZLB binds in 1.56% of the quarters and 0.38% when $\phi_\pi = 3$.

	Steady-State Output ($y_t^* = \bar{y}$)				Potential Output ($y_t^* = y_t^n$)			
	ZLB Binds	Avg. ZLB	Std. Dev. (% of mean)		ZLB Binds	Avg. ZLB	Std. Dev. (% of mean)	
ϕ_y	% of quarters	Quarters	Output	Inflation	% of quarters	Quarters	Output	Inflation
0.125	2.73	1.90	0.6501	0.3326	1.56	1.72	0.6993	0.2800
0.100	2.56	1.87	0.6704	0.3308	1.67	1.73	0.7107	0.2908
0.075	2.45	1.86	0.6925	0.3311	1.80	1.75	0.7234	0.3025
0.050	2.38	1.84	0.7167	0.3335	1.95	1.77	0.7376	0.3152
0.025	2.33	1.82	0.7431	0.3379	2.13	1.79	0.7537	0.3293
0.000	2.33	1.81	0.7719	0.3447	2.33	1.81	0.7719	0.3447

(a) Volatility implications of alternative weights on the output gap (ϕ_y). The weight on inflation is $\phi_\pi = 1.5$.

	Steady-State Output ($y_t^* = \bar{y}$)				Potential Output ($y_t^* = y_t^n$)			
	ZLB Binds	Avg. ZLB	Std. Dev. (% of mean)		ZLB Binds	Avg. ZLB	Std. Dev. (% of mean)	
ϕ_π	% of quarters	Quarters	Output	Inflation	% of quarters	Quarters	Output	Inflation
1.500	2.73	1.90	0.6501	0.3326	1.56	1.72	0.6993	0.2800
1.750	1.40	1.72	0.6116	0.2582	0.99	1.62	0.6586	0.2291
2.000	0.95	1.62	0.5924	0.2160	0.74	1.56	0.6338	0.1950
2.250	0.72	1.59	0.5811	0.1866	0.59	1.53	0.6177	0.1700
2.500	0.58	1.55	0.5740	0.1645	0.50	1.51	0.6067	0.1508
3.000	0.43	1.51	0.5665	0.1333	0.38	1.48	0.5932	0.1231

(b) Volatility implications of alternative weights on the inflation gap (ϕ_π). The weight on the output gap is $\phi_y = 0.125$.

Table 2: Model 1: No capital, technology and discount factor shocks.

5 MODEL 2: STATES OF THE ECONOMY AND THE ZLB

This section shows how our findings change when capital is incorporated into a New Keynesian model. In Model 1, the household can only smooth consumption by varying its labor supply. The presence of capital in Model 2 gives the household another margin to smooth consumption. The addition of another state variable, however, complicates the presentation of the complete solution to the model. We initially fix technology at its steady state, so the complete solution can be presented with contour plots. That allows us to focus on the dynamics created by the discount factor process, which is commonly used to generate ZLB events in the literature. Thus, this model initially contains two state variables—the discount factor and the endogenous capital stock. We then reintroduce the technology process to compare the dynamics between Models 1 and 2 in response to a technology shock. Also in this section, we focus on the dynamics when the central bank targets steady-state output since we believe it better reflects the behavior of actual monetary policy.[13]

Figure 10 shows the three-dimensional contour plots of the nominal interest rate, output, consumption, and investment over the entire state space. Presenting a complete picture of the solution is informative in models with an endogenous state variable like capital since it shows the interaction between the two states. The curvature of the ZLB (shaded) region is due to the quadratic capital adjustment costs. When capital is at its steady state ($\hat{k}_{-1} = 0$), the ZLB binds when $\hat{\beta}_{-1} = 1.22$. As capital rises, the nominal interest rate initially hits zero at lower values of the discount factor. In general, the qualitative behavior of consumption, inflation, and the nominal interest rate are similar to the model without capital. The household's ability to invest in capital, however, causes consumption to be less volatile and generates stronger expectational effects.

We focus our analysis on two cross sections of the contour map in figure 10. The endogeneity of capital makes selecting particular cross sections in Model 2 more difficult than in Model 1. In Model 1, the discount factor and technology states are independent; therefore, any one realization of the discount factor is just as likely regardless of the technology state. In Model 2, the capital and discount factor states are not independent, so the level of capital is likely below (above) its steady state when the discount factor is also below (above) its steady state.

Figure 11 shows two cross sections from the contour map in figure 10. The solid line is the cross section where the capital is fixed at its steady state ($\hat{k}_{-1} = 0$). The dashed line represents the cross section where capital increases with the discount factor along the diagonal of the state space ($\hat{k}_{-1} = \hat{k}_{diag}$). The darker (entire) shaded region indicates the area of the state space where the ZLB binds in the steady-state (diagonal) cross section. We begin by examining the behavior of the economy when the ZLB does not bind. Regardless of the capital state, a higher discount factor makes the household more patient, which increases their desire to invest in capital and to postpone consumption. The higher discount factor also encourages the household to supply more labor. As for the firms, the additional investment in capital pushes up the marginal product of labor, which encourages firms to raise their output and labor demand. The increase in output then puts downward pressure on firm prices so inflation falls. In equilibrium, labor increases and the real wage rate falls as the discount factor increases. Finally, a higher capital stock pushes down the marginal product of capital which causes its real rental rate to fall.

In the diagonal cross section where capital increases with the discount factor ($\hat{k}_{-1} = \hat{k}_{diag}$), the marginal product of capital falls as the discount factor rises which leads to a more rapid decline in the real rental rate. From the household's perspective, a lower rental rate makes investment less attractive as a consumption smoothing channel. The household responds by moderating both their decline in consumption and their increase in labor supply compared to the case in which capital is fixed at its steady state ($\hat{k}_{-1} = 0$). Those more modest responses in the real rental rate, investment,

[13]Potential output is defined as the level of output under flexible prices. In Model 1, potential output is exogenous and we can solve for it analytically, but in Model 2 it has no closed-form solution. It is impossible to numerically solve a flexible price model with a ZLB because sticky prices are necessary for an equilibrium to exist in our model when both monetary and fiscal policy are passive at the ZLB.

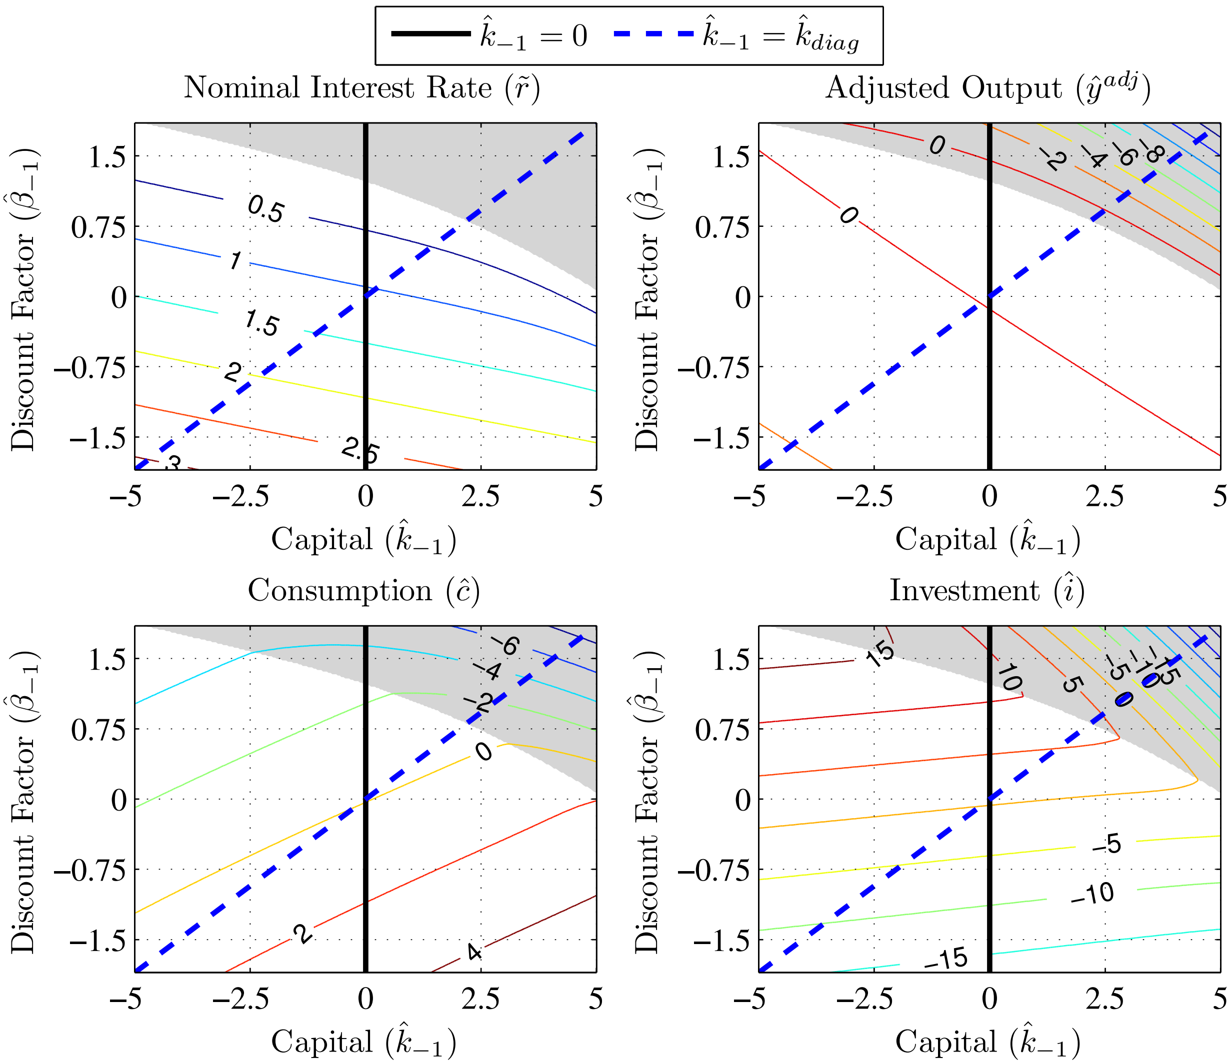

Figure 10: Model 2 ($y_t^* = \bar{y}$) decision rules as a function of the discount factor ($\hat{\beta}_{-1}$) and capital ($\hat{k}_{-1}$) states. Each variable is in percent deviations from its deterministic steady state, except the nominal interest rate, which is a net percentage. The shaded region indicates where the ZLB binds. The solid (black) and dashed (blue) lines correspond to cross sections of the state space, where $\hat{k}_{-1} = 0$ and $\hat{k}_{-1} = \hat{k}_{diag}$, respectively.

consumption, and labor are illustrated by their flatter decision rules when the ZLB does not bind.

In the diagonal (steady-state) cross section, the ZLB binds when the discount factor is more than 0.8% (1.2%) above its steady state. The qualitative properties of the decision rules when the ZLB binds are similar across both cross sections. The mechanism that distorts the economy in Model 2 is similar to Model 1. As the discount factor rises, both inflation and the real interest rate continue to fall. When the nominal rate hits zero, the real interest rate rises as inflation continues to fall.[14] That higher real rate further encourages households to postpone consumption and motivates them to supply more labor. Firms respond to the lower demand by further reducing their prices and sharply cutting their labor demand. That decline in labor demand dominates the increase in labor

[14]The rental rate of capital falls at the ZLB, but the household expects that the future rental rate will increase since they believe the discount factor will return to its steady state. That result is consistent with a rising real interest rate.

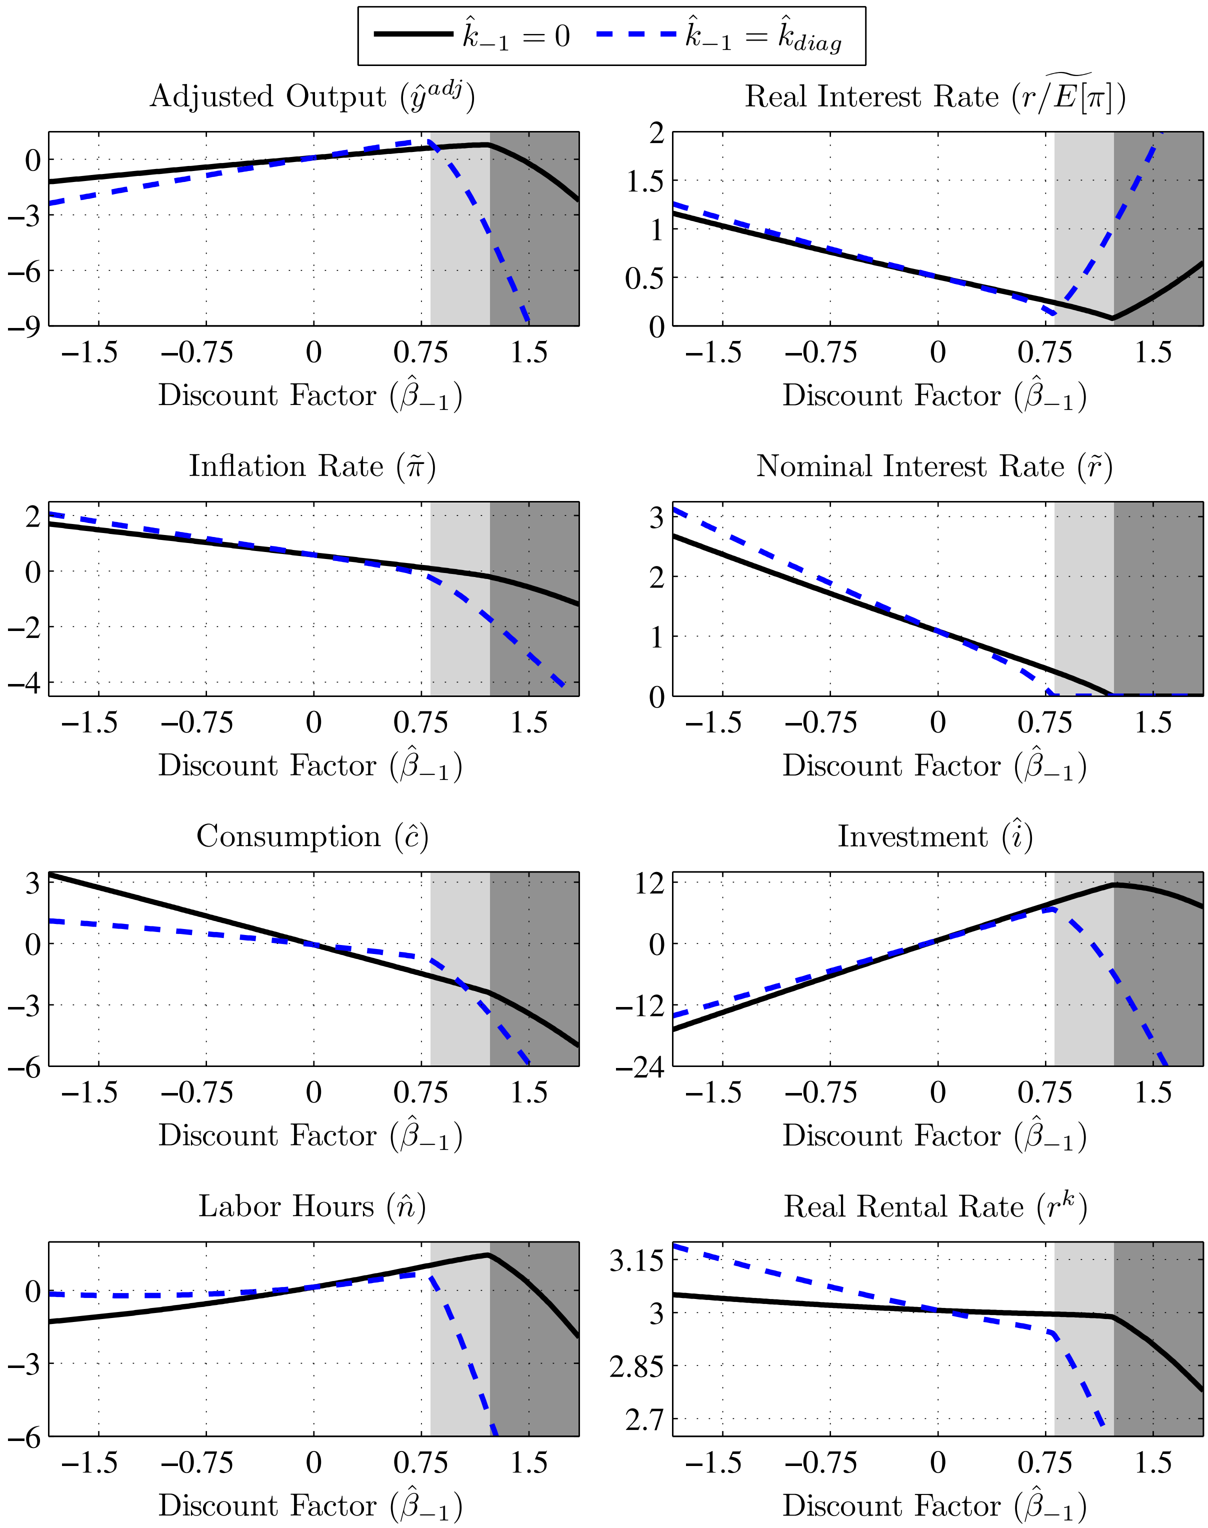

Figure 11: Model 2 ($y_t^* = \bar{y}$) decision rules as a function of the discount factor state ($\hat{\beta}_{-1}$). The solid line is the cross section of the state space where the capital state is fixed at its steady-state value ($\hat{k}_{-1} = 0$), and the dashed line is the diagonal cross section where the capital state changes with the discount factor state ($\hat{k}_{-1} = \hat{k}_{diag}$). Each variable is in percent deviations from its deterministic steady state, except inflation and the nominal interest rate, which are net percentages. The dark (entire) shaded region indicates where the ZLB binds when $\hat{k}_{-1} = 0$ ($\hat{k}_{-1} = \hat{k}_{diag}$).

supply so that both the equilibrium level of labor and the real wage rate fall. Lower consumption then pushes down output, which causes the household to reduce investment even more in order to further smooth its consumption. Thus, the paradoxes of toil and thrift both occur—despite the household wanting to work more to smooth consumption and save more to benefit from higher real interest rates, both hours and investment fall. Those findings demonstrate that our model with capital produces the same unconventional dynamics as the model without capital.

The Importance of Nonlinearities We apply the policy function iteration algorithm to log-linearized versions of Model 1 and Model 2, where the only nonlinearity is the ZLB constraint. This solution method is similar to the procedure employed in Nakov (2008), where linear splines are used to approximate the kink in the decision rules. We then compare the resulting linear decision rules to their nonlinear counterparts to demonstrate the importance of using the fully nonlinear model. The benefit of solving the nonlinear and linear models in this manner is that differences in the decision rules are entirely due to whether or not the model is linearized.

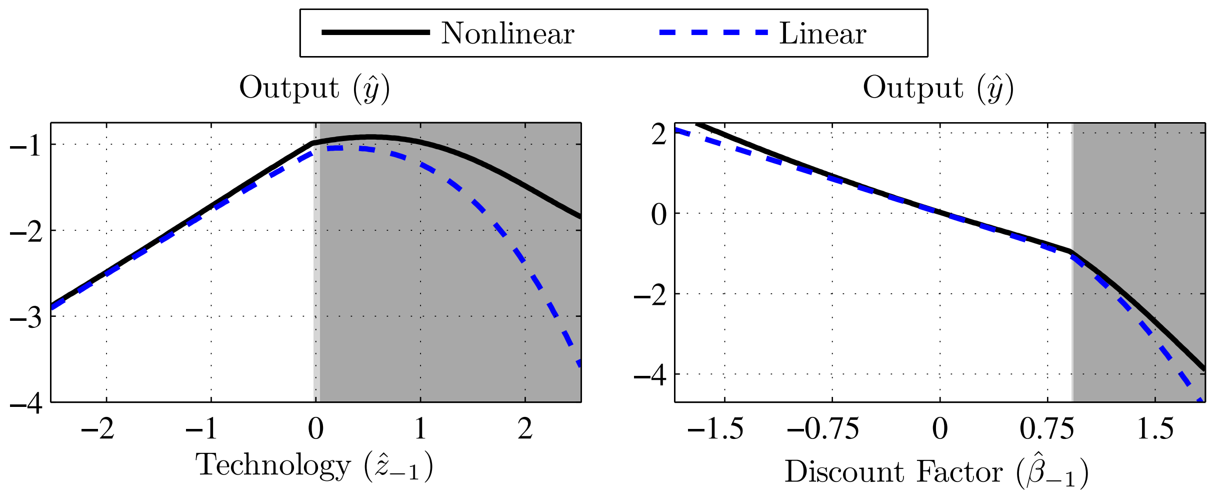

Figure 12: Model 1 ($y_t^* = \bar{y}$) decision rules as a function of the technology state (left panel) and the discount factor state (right panel). The solid line (dashed line) corresponds to the decision rules based on the constrained nonlinear (linear) model. Each variable is in percent deviations from its deterministic steady state. The dark (entire) shaded region indicates where the ZLB binds in the fully nonlinear (linear) model.

Figure 12 compares cross sections of the linear and nonlinear decision rules for output in Model 1 when the central bank targets steady-state output. The left (right) panel shows the decision rule as a function of the technology state (discount factor state). The linear decision rules are a fairly accurate approximation of the nonlinear decision rules for both the technology and the discount factor shocks so long as the ZLB does not bind. The linear and nonlinear decision rules for output then diverge as the economy moves deeper into the ZLB region for both shocks. That separation is initiated by the inability of monetary policy to compensate for growing price adjustment costs, which are different due to the linearization of the quadratic price adjustment cost function. Furthermore, the location of the ZLB kink is nearly identical for both the linear and nonlinear decision rules. Those results, in contrast to Fernández-Villaverde et al. (2012), indicate that the linear model provides a fairly good approximation of the nonlinear model without capital in most states. We will show, however, that decision rules from the linear model with capital are far less accurate because of the expectational effects caused by the ZLB constraint.

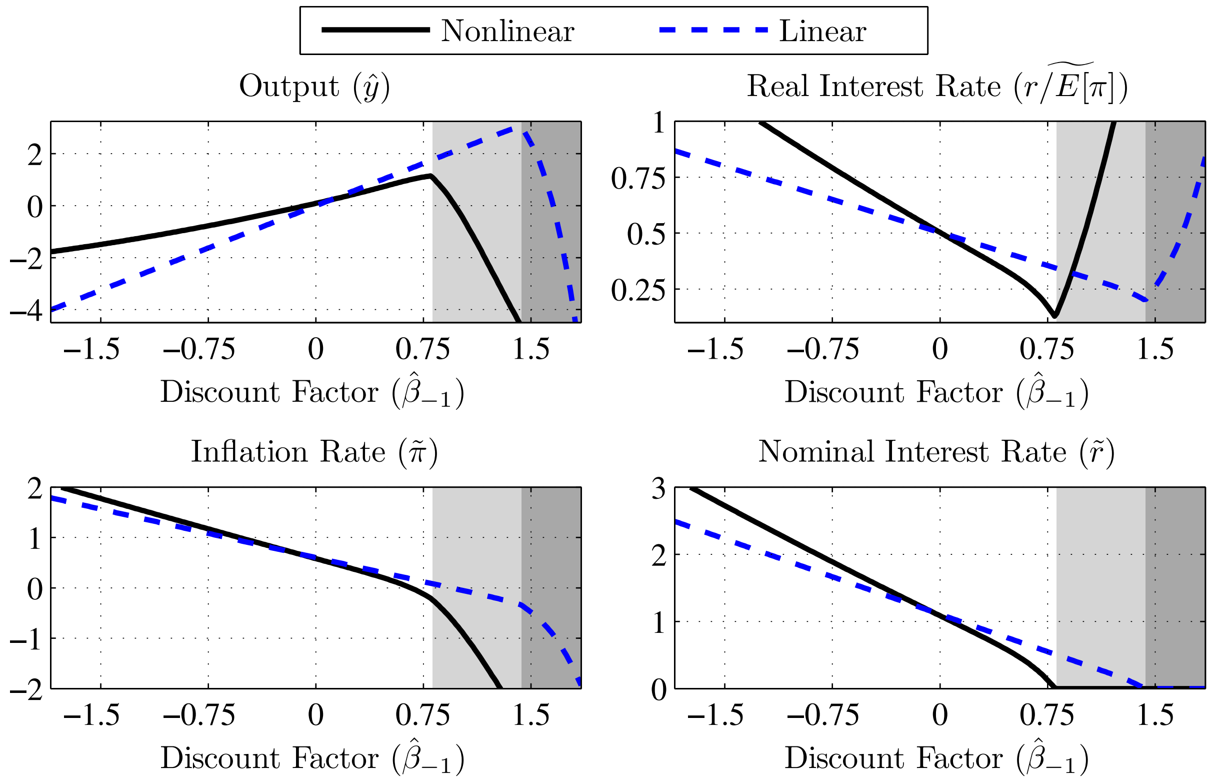

Figure 13: Model 2 ($y_t^* = \bar{y}$) decision rules as a function of the discount factor state ($\hat{\beta}_{-1}$). The capital state changes with the discount factor state ($\hat{k}_{-1} = \hat{k}_{diag}$). The solid line (dashed line) corresponds to the decision rules based on the constrained nonlinear (linear) model. Each variable is in percent deviations from its deterministic steady state, except inflation and the interest rates, which are net percentage. The dark (entire) shaded region indicates where the ZLB binds in the fully nonlinear (linear) model.

Figure 13 compares the linear and nonlinear decision rules in Model 2 as a function of the discount factor state. The ZLB first binds in the linear (nonlinear) model when the discount factor is 1.4% (0.8%) percent above its steady state. That difference has two important implications. One, simulations of the linear model indicate that the economy hits the ZLB far less frequently than in the nonlinear model. In our baseline calibration, the ZLB never binds in the linear model, while it binds in 1.15% of the time in the nonlinear model. That result is one reason why ZLB studies that linearize the model must specify much larger shocks to generate ZLB events. Two, the decision rules differ between the linear and nonlinear model when the ZLB does not bind, because the expectational effects of visiting the ZLB are weaker in the linear model. Thus, the linear model cannot accurately quantify the effects of discount factor shocks even when the ZLB does not bind. The reason the linear and nonlinear models generate such different results in the model with capital is because Model 2 has two assets—capital and bonds, while Model 1 has only bonds. Arbitrage implies that the expected rates of return on investment and bonds are identical, which means the expected future rental rate of capital equals the current real interest rate.

6 MODEL 1 AND MODEL 2 COMPARISONS

This section shows that capital qualitatively and quantitatively affects dynamics at the ZLB. We compare GIRFs in models with and without capital. We evaluate GIRFs because a cross section requires assumptions about how the capital state in Model 2 co-moves with the exogenous state variables. To conduct such an experiment, technology is stochastic with the same parameter values in both models. We also assume that the central bank targets steady-state output ($y_t^* = \bar{y}$) and set $\phi_y = 0.025$ in both models. A small weight on ϕ_y is necessary to make a direct comparison because our numerical algorithm does not converge for higher values of ϕ_y in Model 2.

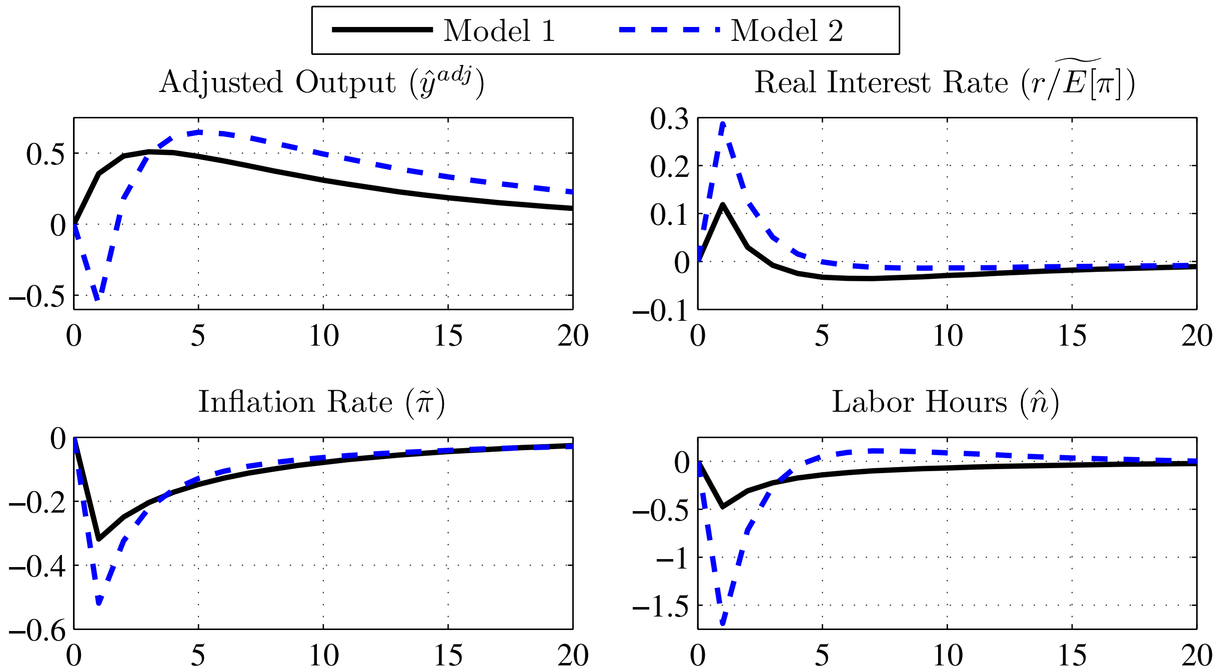

Figure 14: Comparison of the GIRFs to a 1% positive technology shock in Model 1 (solid line) and in Model 2 (dashed line). The initial state vector is equal to the average state vector in a 500,000 quarter simulation conditional on the ZLB binding. The central bank targets steady-state output ($y_t^* = \bar{y}$).

Figure 14 plots GIRFs to a 1% positive technology shock in Model 1 (solid line) and in Model 2 (dashed line). In this experiment, the discount factor is initially set to its mean value that causes the ZLB to bind in a 500,000 quarter simulation of the model where technology shocks are set to zero (Model 1: $\hat{\beta}_{-1} = 1$ and Model 2: $\hat{\beta}_{-1} = 1.4$). The unconventional dynamics are not present in Model 1 when ϕ_y is small because the positive supply-side effects of higher technology dominate the negative demand-side effects of a higher real interest rate. As ϕ_y increases, the adverse demand-side effects overcome the beneficial supply-side effects so that output and labor hours both decline in response to a positive technology shock at the ZLB. Output, in contrast, declines on impact in Model 2 even when ϕ_y is small. The responses of the real interest rate, inflation, and labor are qualitatively the same in both models but quantitatively larger in Model 2. From period 2 onward, technology and the discount factor mean revert and the GIRFs from both models converge.

Next, we show the impact of capital on the volatility of output and inflation for a range of values for ϕ_y. Model 1 and Model 2 are simulated for 500,000 quarters under the assumption that the central bank targets steady-state output. We fix technology at its steady state to examine a broader range of values for ϕ_y. Table 3 shows the effect of reducing the weight on the output gap (ϕ_y) while holding the weight on inflation (ϕ_π) at 1.5. The value of ϕ_y is initially set slightly below the original Taylor (1993) specification, $\phi_y = 0.1$, and is reduced by increments of 0.025.

	Model 1				Model 2			
	ZLB Binds	Avg. ZLB	Std. Dev. (% of mean)		ZLB Binds	Avg. ZLB	Std. Dev. (% of mean)	
ϕ_y	% of quarters	Quarters	Output	Inflation	% of quarters	Quarters	Output	Inflation
0.100	1.20	1.63	0.4972	0.2769	1.15	1.87	0.4005	0.2979
0.075	1.29	1.64	0.5168	0.2878	0.35	1.66	0.4127	0.2654
0.050	1.39	1.65	0.5382	0.2997	0.16	1.50	0.4271	0.2473
0.025	1.51	1.66	0.5615	0.3126	0.07	1.41	0.4421	0.2304
0.000	1.64	1.68	0.5870	0.3268	0.03	1.46	0.4581	0.2133

Table 3: Volatility implications of alternative weights on the output gap. Comparison between Model 1 and Model 2. The only stochastic component in both models is discount factor shocks. $\phi_\pi = 1.50$, $\rho_\beta = 0.80$, and $\sigma_\beta = 0.0025$.

Our results show that capital decreases the frequency of ZLB events for every value of ϕ_y, which is consistent with Christiano (2004). A larger value of ϕ_y in Model 1 decreases the likelihood and shortens the average duration of ZLB events. At first glance, that result appears to contradict our Model 1 findings when the central bank targets steady-state output (see table 2a, columns 2 and 3). Since technology is fixed at its steady state, the steady-state output and the potential output are equal. Thus, the qualitative result is consistent with the findings from Model 1 when the central bank targets potential output (see table 2a, columns 6 and 7). Model 2, however, does generate two opposite results from Model 1. One, a higher ϕ_y increases the likelihood and average duration of ZLB events in Model 2. Two, both output and inflation volatility decline as ϕ_y increases in Model 1, whereas a tradeoff exists between lower output volatility and higher inflation volatility in Model 2. Those differences are important since many central banks have emphasized output stabilization since the end of the Great Recession.

7 Conclusion

This paper examines monetary policy when it is constrained by the ZLB using models with and without capital. Our analysis focuses on discount factor shocks since they are viewed as the likely cause of many ZLB events. Nevertheless, we also examine technology shocks because they are an important source of aggregate fluctuations in many dynamic models. We use these models to analyze why technology shocks at the ZLB may have unconventional effects, what factors influence the likelihood of hitting the ZLB, and the tradeoffs faced by the central bank under a dual mandate.

Three main findings emerge: (1) A positive technology shock can generate lower consumption, labor, and output when the ZLB binds and the central bank targets steady-state output. Those unconventional dynamics usually disappear when the central bank targets potential output; (2) When the central bank targets steady-state output in the model with capital, a positive technology shock at the ZLB generates more contractionary dynamics than in the model without capital; and (3) The constrained linear model provides a good approximation of the nonlinear model without capital, but differences exist between the solutions in the model with capital.

In spite of the large amount of research on the ZLB, many important questions remain. For example, do the medium- to long-run benefits of returning to normal policy outweigh the short-run costs of a higher policy rate? What are the benefits of forward guidance and quantitative easing in a dynamic model that accounts for expectational effects? Lastly, how does the impact of discount factor and technology shocks change as a result of those polices? Answering those questions and others requires careful treatment of expectations when examining dynamics near or at the ZLB.

REFERENCES

ADAM, K. AND R. M. BILLI (2006): "Optimal Monetary Policy under Commitment with a Zero Bound on Nominal Interest Rates," *Journal of Money, Credit and Banking*, 38, 1877–1905.

ARUOBA, S. AND F. SCHORFHEIDE (2013): "Macroeconomic Dynamics Near the ZLB: A Tale of Two Equilibria," NBER Working Paper 19248.

BASU, S. AND B. BUNDICK (2012): "Uncertainty Shocks in a Model of Effective Demand," NBER Working Paper 18420.

BASU, S. AND J. G. FERNALD (2009): "What Do We Know and Not Know About Potential Output?" *Federal Reserve Bank of St. Louis Review*, 91, 187–214.

BENHABIB, J., S. SCHMITT-GROHE, AND M. URIBE (2001): "The Perils of Taylor Rules," *Journal of Economic Theory*, 96, 40–69.

BRAUN, R. A. AND L. M. KÖRBER (2011): "New Keynesian Dynamics in a Low Interest Rate Environment," *Journal of Economic Dynamics and Control*, 35, 2213–2227.

BRAUN, R. A., L. M. KÖRBER, AND Y. WAKI (2012): "Some Unpleasant Properties of Log-Linearized Solutions When the Nominal Rate is Zero," FRB Atlanta Working Paper 2012-5a.

——— (2013): "Small and Orthodox Fiscal Multipliers at the Zero Lower Bound," FRB Atlanta Working Paper 2013-13.

BRAUN, R. A. AND Y. WAKI (2006): "Monetary Policy During Japan's Lost Decade," *The Japanese Economic Review*, 57, 324–344.

——— (2010): "On the Size of the Fiscal Multiplier When the Nominal Interest Rate is Zero," Mimeo, University of Tokyo.

CALVO, G. A. (1983): "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12, 383–398.

CHRISTIANO, L. J. (2004): "The Zero-Bound, Low Inflation, and Output Collapse," Mimeo, Northwestern University.

CHRISTIANO, L. J. AND M. EICHENBAUM (2012): "Notes on Linear Approximations, Equilibrium Multiplicity and E-learnability in the Analysis of the Zero Lower Bound," Manuscript, Northwestern University.

CHRISTIANO, L. J., M. EICHENBAUM, AND S. REBELO (2011): "When Is the Government Spending Multiplier Large?" *Journal of Political Economy*, 119, 78–121.

COLEMAN, II, W. J. (1991): "Equilibrium in a Production Economy with an Income Tax," *Econometrica*, 59, 1091–1104.

DAVIG, T. AND E. M. LEEPER (2007): "Generalizing the Taylor Principle," *American Economic Review*, 97, 607–635.

DIXIT, A. K. AND J. E. STIGLITZ (1977): "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 67, 297–308.

EBERLY, J. C. (1997): "International Evidence on Investment and Fundamentals," *European Economic Review*, 41, 1055–1078.

EGGERTSSON, G. AND M. WOODFORD (2003): "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 34(1), 139–235.

EGGERTSSON, G. B. (2010): "The Paradox of Toil," FRB New York Staff Report 433.

ERCEG, C. AND J. LINDÉ (2014): "Is There A Fiscal Free Lunch In A Liquidity Trap?" *Journal of the European Economic Association*, 12, 73–107.

ERCEG, C. J. AND A. T. LEVIN (2003): "Imperfect Credibility and Inflation Persistence," *Journal of Monetary Economics*, 50, 915–944.

FERNÁNDEZ-VILLAVERDE, J., G. GORDON, P. GUERRÓN-QUINTANA, AND J. F. RUBIO-RAMÍREZ (2012): "Nonlinear Adventures at the Zero Lower Bound," NBER Working Paper 18058.

GERTLER, M. AND P. KARADI (2011): "A Model of Unconventional Monetary Policy," *Journal of Monetary Economics*, 58, 17–34.

GÜNTER, C., O. ATHANASIOS, AND W. VOLKER (2004): "Price Stability and Monetary Policy Effectiveness when Nominal Interest Rates are Bounded at Zero," *The B.E. Journal of Macroeconomics: Advances*, 4, 1–25.

GUST, C., D. LÓPEZ-SALIDO, AND M. E. SMITH (2013): "The Empirical Implications of the Interest-Rate Lower Bound," CEPR Discussion Paper 9214.

HOSHI, T. AND A. KASHYAP (2000): "The Japanese Banking Crisis: Where Did It Come From and How Will It End?" in *NBER Macroeconomics Annual 1999, Volume 14*, MIT Press, Cambridge, 129–212.

IRELAND, P. N. (1997): "A Small, Structural, Quarterly Model for Monetary Policy Evaluation," *Carnegie-Rochester Conference Series on Public Policy*, 47, 83–108.

ITO, T. AND F. S. MISHKIN (2006): "Two Decades of Japanese Monetary Policy and the Deflation Problem," in *Monetary Policy under Very Low Inflation in the Pacific Rim*, University of Chicago Press, vol. 15, 131–202.

JUNG, T., Y. TERANISHI, AND T. WATANABE (2005): "Optimal Monetary Policy at the Zero-Interest-Rate Bound," *Journal of Money, Credit and Banking*, 37, 813–35.

KILEY, M. T. (2003): "Why Is Inflation Low When Productivity Growth Is High?" *Economic Inquiry*, 41, 392–406.

KOOP, G., M. H. PESARAN, AND S. M. POTTER (1996): "Impulse Response Analysis in Nonlinear Multivariate Models," *Journal of Econometrics*, 74, 119–147.

KRUGMAN, P. R. (1998): "It's Baaack: Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity*, 29, 137–206.

MCCALLUM, B. T. (2001): "Inflation Targeting and the Liquidity Trap," NBER Working Paper 8225.

MERTENS, K. AND M. O. RAVN (2014): "Fiscal Policy in an Expectations Driven Liquidity Trap," *Review of Economic Studies*, forthcoming.

NAKATA, T. (2012): "Uncertainty at the Zero Lower Bound," Finance and Economics Discussion Series, 2013-09.

NAKOV, A. (2008): "Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate," *International Journal of Central Banking*, 4, 73–127.

ORPHANIDES, A. (2003a): "Monetary Policy Evaluation with Noisy Information," *Journal of Monetary Economics*, 50, 605–631.

——— (2003b): "The Quest for Prosperity Without Inflation," *Journal of Monetary Economics*, 50, 633–663.

ORPHANIDES, A. AND S. VAN NORDEN (2002): "The Unreliability of Output-Gap Estimates in Real Time," *The Review of Economics and Statistics*, 84, 569–583.

PETERMAN, W. B. (2012): "Reconciling Micro and Macro Estimates of the Frisch Labor Supply Elasticity," Finance and Economics Discussion Series, 2012-75.

POSEN, A. (1998): *Restoring Japan's Economic Growth*, Peterson Institute.

RICHTER, A. W. AND N. A. THROCKMORTON (2014): "The Zero Lower Bound: Frequency, Duration, and Numerical Convergence," *B.E. Journal of Macroeconomics,* forthcoming.

RICHTER, A. W., N. A. THROCKMORTON, AND T. B. WALKER (2014): "Accuracy, Speed and Robustness of Policy Function Iteration," *Computational Economics*, 44, 445–476.

ROTEMBERG, J. J. (1982): "Sticky Prices in the United States," *Journal of Political Economy*, 90, 1187–1211.

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

TAYLOR, J. B. (1993): "Discretion Versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

WERNING, I. (2011): "Managing a Liquidity Trap: Monetary and Fiscal Policy," NBER Working Paper 17344.

WIELAND, J. (2014): "Are Negative Supply Shocks Expansionary at the Zero Lower Bound?" Mimeo, University of California-Berkeley.

WOLMAN, A. L. (2005): "Real Implications of the Zero Bound on Nominal Interest Rates," *Journal of Money, Credit and Banking*, 37, 273–96.

A NUMERICAL ALGORITHM

A formal description of the numerical algorithm begins by writing the model compactly as

$$E[f(\mathbf{v}_{t+1}, \mathbf{w}_{t+1}, \mathbf{v}_t, \mathbf{w}_t)|\Omega_t] = 0,$$

where f is vector-valued function that contains the equilibrium system, $\mathbf{v}$ is a vector of exogenous variables, $\mathbf{w}$ is a vector of endogenous variables, and $\Omega_t = \{M, P, \mathbf{z}_t\}$ is the household's information set in period t, which contains the structural model, M, its parameters, P, and the state vector, $\mathbf{z}$. In Model 1, $\mathbf{v} = \mathbf{z} = (\beta, z)$ and $\mathbf{w} = (c, \pi, y, n, w, r)$. In Model 2 with both stochastic processes, $\mathbf{v} = (\beta, z)$, $\mathbf{z} = (k, \beta, z)$, and $\mathbf{w} = (c, \pi, y, n, w, r, k, i, r^k, q)$.

Policy function iteration approximates the vector of decision rules, Φ, as a function of the state vector, $\mathbf{z}$. The time-invariant decision rules for the exogenous model are

$$\underbrace{\Phi(\mathbf{z}_t)}_{\text{True RE Solution}} \approx \underbrace{\hat{\Phi}(\mathbf{z}_t)}_{\substack{\text{Approximating}\\ \text{Function}}}.$$

We iterate on $\Phi = (c, \pi)$ for Model 1 and $\Phi = (n, \pi, i)$ for Model 2 so that we can easily solve for future variables that enter the household's expectations using f. Each continuous state variable in $\mathbf{z}$ is discretized into N^d points, where $d \in \{1, \ldots, D\}$ and D is the dimension of the state space. The discretized state space is represented by a set of unique D-dimensional coordinates (nodes). In general, we set the bounds of continuous stochastic state variables to encompass 99.999% of the probability mass of the distribution. We specify 251 grid points for each continuous state variable and use 31 Gauss-Hermite weights for each continuous shock. Those techniques minimize extrapolation and ensure that the location of the kink in the decision rules is accurate.

The following outline summarizes the policy function algorithm we employ for our models. Let $i \in \{0, \ldots, I\}$ index the iterations of the algorithm and $n \in \{1, \ldots, \Pi_{d=1}^{D} N^d\}$ index the nodes.

1. Obtain initial conjectures for the policy functions on each node from the log-linear model without a ZLB constraint. The initial conjectures are $\hat{c}_0$ and $\hat{\pi}_0$ for Model 1 and $\hat{n}_0$, $\hat{\pi}_0$, and $\hat{i}_0$ for Model 2. We use Sims (2002) `gensys.m` program to obtain these conjectures.

2. For $i \in \{1, \ldots, I\}$, implement the following steps:

   (a) On each node, solve for $\{y_t, r_t\}$ given $\hat{c}_{i-1}(\mathbf{z}_t^n)$ and $\hat{\pi}_{i-1}(\mathbf{z}_t^n)$ in Model 1 and given $\hat{n}_{i-1}(\mathbf{z}_t^n)$, $\hat{\pi}_{i-1}(\mathbf{z}_t^n)$, and $\hat{i}_{i-1}(\mathbf{z}_t^n)$ in Model 2 with the ZLB imposed.

   (b) Linearly interpolate (c_{t+1}, π_{t+1}) in Model 1 and $(n_{t+1}, \pi_{t+1}, i_{t+1})$ in Model 2 given $\{\varepsilon_{t+1}^m\}_{m=1}^M$. Each of the M values ε_{t+1}^m are Gauss-Hermite quadrature nodes. We use Gauss-Hermite quadrature to numerically integrate, since it is accurate for normally distributed shocks. We use piecewise linear interpolation to approximate future variables, since this approach more accurately captures the kink in the decision rules than continuous approximating functions such as cubic splines or Chebyshev polynomials.[15]

[15] Aruoba and Schorfheide (2013) use a linear combination of two Chebyshev polynominals—one that captures the dynamics when the ZLB binds and one that captures the dynamics when the Taylor principle holds. This approach is more accurate than using one Chebyshev polynomial, but there is no guarantee that it will accurately locate the kink. Moreover, Chebyshev polynomials can lead to large approximation errors due to extrapolation. With linear interpolation, a dense state space will lead to more predictable extrapolation and more accurately locate the kink. See Richter et al. (2014) for a comparison of these two solution methods in a New Keynesian model with a ZLB constraint.

(c) We use the nonlinear solver, `csolve.m`, to minimize the Euler equation errors. On each node, numerically integrate to approximate the expectation operators,

$$\mathbb{E}\left[f(\mathbf{x}_{t+1}^m, \mathbf{x}_t^n)|\Omega_t\right] \approx \frac{1}{\pi}\sum_{m=1}^{M} f(\hat{\mathbf{x}}_{t+1}^m, \hat{\mathbf{x}}_t^n)\phi(\varepsilon_{t+1}^m),$$

where $\mathbf{x} \equiv (\mathbf{v}, \mathbf{w})$ and ϕ are the respective Gauss-Hermite weights. The superscripts on $\mathbf{x}$ indicate which realizations of the state variables are used to compute expectations. The nonlinear solver searches for $\hat{c}_i(\mathbf{z}_t^n)$ and $\hat{\pi}_i(\mathbf{z}_t^n)$ in Model 1 and $\hat{n}_i(\mathbf{z}_t^n)$, $\hat{\pi}_i(\mathbf{z}_t^n)$, and $\hat{i}_i(\mathbf{z}_t^n)$ in Model 2 so that the Euler equation errors are less than 10^{-4} on each node.

3. Define $\text{maxdist}_i \equiv \max\{|\hat{c}_i - \hat{c}_{i-1}|, |\hat{\pi}_i - \hat{\pi}_{i-1}|\}$ in Model 1 and $\text{maxdist}_i \equiv \max\{|\hat{n}_i - \hat{n}_{i-1}|, |\hat{\pi}_i - \hat{\pi}_{i-1}|, |\hat{i}_i - \hat{i}_{i-1}|\}$ in Model 2. Repeat the steps in item 2 until one of the following conditions is satisfied:

   - If for all n, $\text{maxdist}_i < 10^{-13}$ for 10 consecutive iterations, then the algorithm converged to a MSV solution. In Model 1, since the state is composed of only exogenous variables, the solution is bounded so long as the decisions rules are positive and finite. In Model 2, simulations of the model must not be explosive.
   - Otherwise, we say the algorithm is non-convergent for one of the following reasons:
     - $i = I = 500{,}000$ (Algorithm times out)
     - For all n and any i, $\hat{\pi}_i < 0.5$, or for any n, $\hat{c}_i < 0$ in Model 1 or $\hat{n}_i < 0$ in Model 2 (Approximating functions drift)
     - Define $\text{dir}_i = \text{maxdist}_i - \text{maxdist}_{i-1}$. For all n, $\text{dir}_i \geq 0$ and $\text{dir}_i - \text{dir}_{i-1} \geq 0$ for 50 consecutive iterations (Algorithm diverges).

The same criteria is used to generate the results in Richter and Throckmorton (2014).

B CONVERGENCE PATHS AND STEADY-STATE EQUILIBRIA

Figure 15 illustrates that the model has two steady states, which is consistent with Benhabib et al. (2001). The left panel highlights the intersections of the consumption Euler equation and the monetary policy rule in steady state (circles). Those two steady states are

$$\begin{aligned} \bar{r} &= \bar{\pi}/\bar{\beta}, && \text{(Consumption Euler Equation)} \\ \bar{r} &= \max\{1, r^*(\bar{\pi}/\pi^*)^{\phi_\pi}\}, && \text{(Interest Rate Rule)} \end{aligned}$$

which results in two steady-state inflation rates:

$$\bar{\pi} = \begin{cases} \pi^* & \text{when } \bar{r} = r^* \\ \bar{\beta} & \text{when } \bar{r} = 1 \end{cases}.$$

When combined with the household's first-order condition for labor and the steady-state resource constraint, the firm pricing equation yields the steady-state value of consumption as a function of the steady-state inflation rate:

$$\bar{c} = \left(\frac{1}{\theta\chi}\left((1-\beta)\varphi\left(\frac{\bar{\pi}}{\pi^*} - 1\right)\frac{\bar{\pi}}{\pi^*} - (1-\theta)\right)\left(1 - \frac{\varphi}{2}\left(\frac{\bar{\pi}}{\pi^*} - 1\right)^2\right)^\eta\right)^{1/(1+\eta)}. \tag{12}$$

Since the model contains two steady-state inflation rates, consumption also has two steady-state values, which are shown in the right panel of figure 15. In every figure in this section, the inflation rate is shown as a net percentage and consumption is in percent deviations from its steady-state value when the inflation rate is positive.

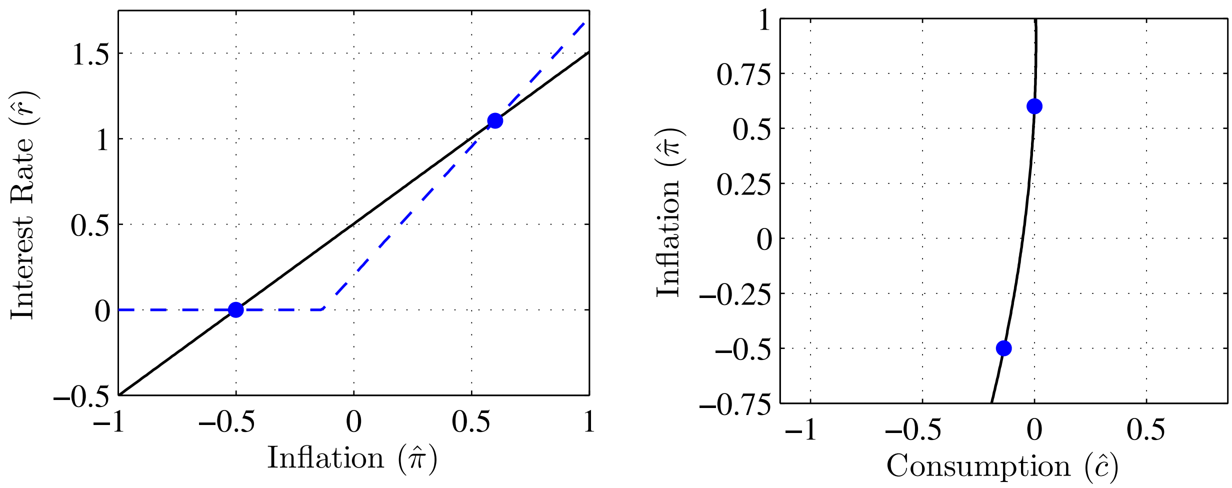

Figure 15: Model 1 steady states (circles). The left panel shows the consumption Euler equation (black line) and the interest rate rule (dashed line) in steady state. The right panel shows the firm pricing equation in steady state.

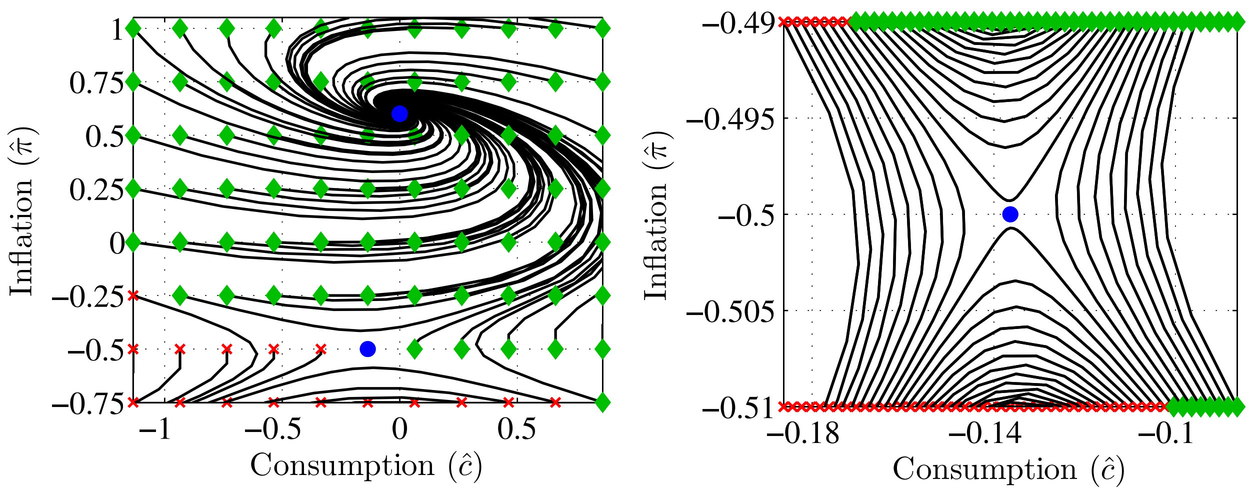

Figure 16: Convergence paths to the steady-state equilibria (circles) in the deterministic version of Model 1. A diamond denotes an initial conjecture that converges to the positive inflation steady state, and a cross denotes an initial conjecture that asymptotically converges to a corner solution where there is no consumption.

Figure 16 shows the convergence paths of consumption and inflation when our algorithm is initialized at different points (diamonds and crosses). The convergence paths correspond to the values of consumption and inflation in each iteration of our numerical algorithm. The left panel shows the model's two steady states (circles). Inflation is positive in one steady state and negative in the other steady state. The right panel shows that our model may converge to the deflationary steady state via a saddle path that runs from the northwest and southeast as shown in Benhabib et al. (2001).

That saddle path is not explicitly shown in the right panel because the algorithm only converges to the deflationary steady state if the initial conjecture is exactly equal to one of the values on either side of the stable manifold. Obtaining that precise convergence path is not possible without the analytical equation for the saddle path because any numerical algorithm is based on an approximation of the true rational expectations solution. Thus, our numerical algorithm only converges to the deflationary steady state if the distance between the initial conjecture and the deflationary steady state is less than the tolerance criterion, 10^{-10}.[16] The unstable manifolds point away from the deflationary steady state toward the southwest. The policy function iteration algorithm converges to the inflationary steady state so long as the initial conjectures for inflation and consumption are to the northeast (diamonds) of the stable manifolds of the deflationary steady state. Initial conjectures in the southwest (crosses) yield paths that are unstable because they asymptotically approach a corner solution where consumption is equal to 0.

Figure 17 shows the convergence paths for consumption and inflation when the discount factor state, β_{-1}, does not equal the deterministic steady-state value, $\bar{\beta}$. The right (left) panel displays the convergence paths for $\beta_{-1} = 0.9975$ ($\beta_{-1} = 0.9925$), which is above (below) $\bar{\beta} = 0.995$. In both cases, the inflationary steady state is stable, but there is no evidence that the algorithm converges to the deflationary steady state as in figure 16. Furthermore, our findings provide no indication that there is a saddle path to the deflationary steady state.

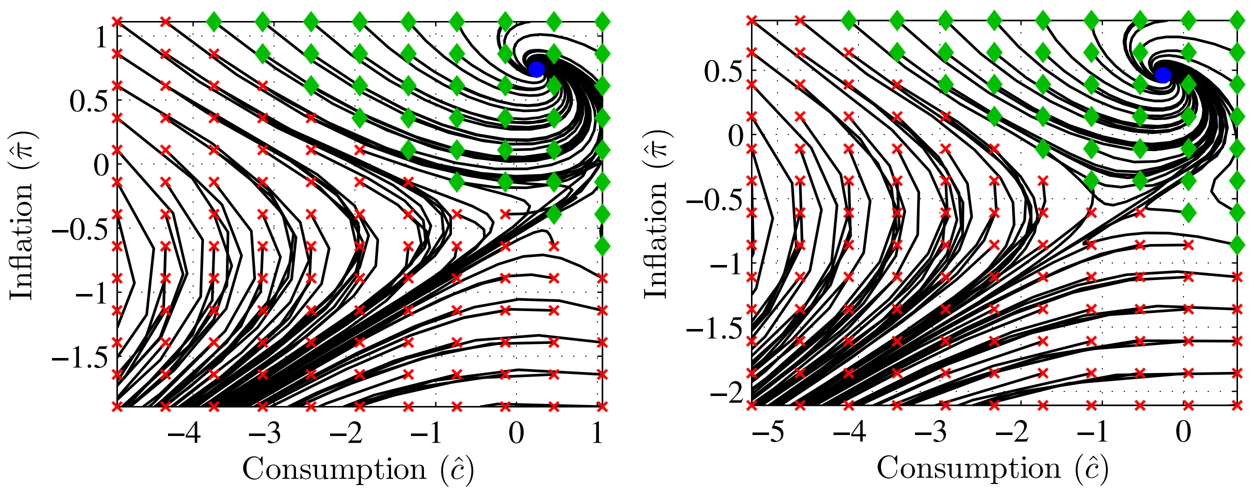

(a) $\beta_{-1} = 0.9925 < \bar{\beta}$

(b) $\beta_{-1} = 0.9975 > \bar{\beta}$

Figure 17: Convergence paths to the steady-state equilibria (circles) in the deterministic version of Model 1. A diamond denotes an initial conjecture that converges to the positive inflation steady state, and a cross denotes an initial conjecture that asymptotically converges to a corner solution where there is no consumption.

Figure 18 compares the convergence paths of consumption and inflation when $\beta_{-1} = \bar{\beta}$ for the deterministic (left panel) and stochastic (right panel) versions of Model 1. In the stochastic model, the household forms expectations over a range of future realizations of β. The paths in the stochastic model differ from the deterministic model in two important ways. One, fewer initial

[16] Christiano and Eichenbaum (2012) describe the deflationary steady state as a pencil standing on its tip. If the agent's belief is incorrect by 10^{-9}, then the equilibrium falls apart (i.e., the pencil falls over). Initial conjectures for consumption and inflation that slightly deviate from the deflationary steady state mean the algorithm either converges to the inflationary steady state or diverges away from the deflationary steady state.

conjectures converge to the inflationary steady state in the stochastic model. That decline occurs because expectations are formed over future values of β that are in the region where the paths of consumption and inflation diverge from the inflationary steady state. In other words, additional instability is created by forming expectations over values of β that put the algorithm on an unstable path. Two, the unstable deflationary steady state is no longer present in the stochastic model. If the stochastic model is initialized at values of consumption and inflation on either stable manifold in the deterministic model, then many realizations of future β result in paths that move away from the inflationary steady state.

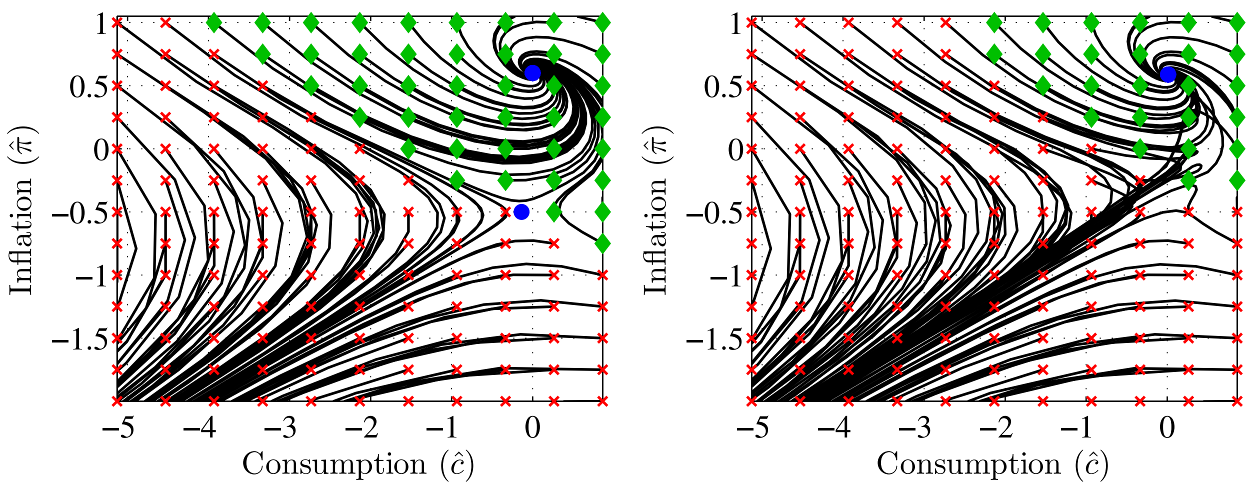

(a) Perfect Foresight Model 1

(b) Stochastic Model 1

Figure 18: Convergence paths to steady states (circles) for the perfect foresight and stochastic models when $\beta_{-1} = \bar{\beta}$. A diamond denotes an initial conjecture that converges to the positive inflation steady state, and a cross denotes an initial conjecture that asymptotically converges to a corner solution where there is no consumption.

To further understand these convergence paths, we also examine a Markov-switching specification of the stochastic model [Eggertsson and Woodford (2003)]. This exercise demonstrates that expectational effects in the stochastic model destabilize the deflationary steady state that is present in the deterministic model. In this model, the equilibrium at time t is determined by a 2-state Markov chain with transition matrix $\Pr\{s_t = j | s_{t-1} = i\} = p_{ij}$, $i, j \in \{1, 2\}$. The two equilibria are the inflationary and deflationary steady states of the deterministic model shown in figure 15:

$$(\bar{\pi}, \bar{r}) = \begin{cases} (\pi^*, r^*) & \text{for } s_t = 1 \\ (\bar{\beta}, 1) & \text{for } s_t = 2 \end{cases}.$$

Figure 19 shows the paths of consumption and inflation beginning from their initial conjectures, $(\hat{c}_0, \hat{\pi}_0)$, for each state, s. In both panels, the inflationary steady state is perfectly absorbing, $p_{11} = 1$, which is often assumed in the ZLB literature [e.g., Braun et al. (2013); Christiano et al. (2011); Eggertsson and Woodford (2003)]. The left panel shows that when the initial conjecture for $s = 2$ is equal to the deflationary steady state and its state is perfectly absorbing (i.e., $p_{22} = 1$) the algorithm converges to the deflationary steady state. The right panel, however, shows that if $p_{22} = 0.99$ (i.e., there is a small probability of leaving the deflationary steady state), then only the inflationary steady state remains. Conditional on starting in state 2, even the smallest possibility of returning to the inflationary steady state makes the deflationary steady state unstable. The deflationary equilibrium satisfies the steady-state system of equations of the deterministic model, but it no longer satisfies the stochastic and dynamic system of equations due to expectational effects.

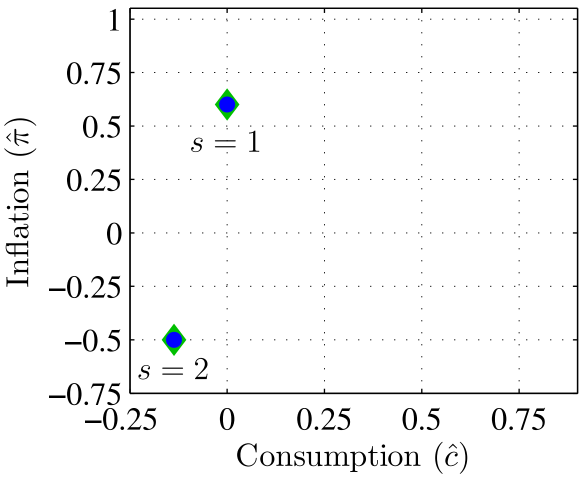

(a) $(p_{11}, p_{22}) = (1, 1)$

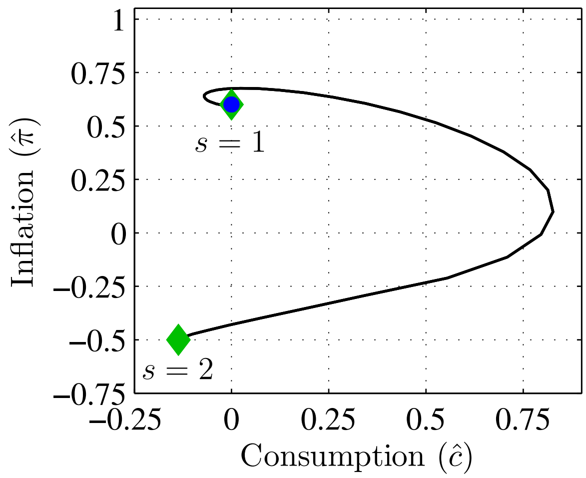

(b) $(p_{11}, p_{22}) = (1, 0.99)$

Figure 19: Convergence paths in a version of Model 1 that switches between the steady states. A diamond denotes an initial conjecture that converges to a steady state. $s = 1$ ($s = 2$) is the inflationary (deflationary) steady state.

C Sensitivity Analysis

The parameters of the stochastic processes impact where the ZLB binds in the state space, the slope of the policy functions, and model dynamics. As an example, figure 20 compares the Model 1 policy functions with $\rho_\beta = 0.8$ and $\rho_\beta = 0.75$ when technology is held at its steady-state value, $z_{-1} = 0$. A more persistent discount factor process makes the ZLB bind in lower discount factor states. When $\rho_\beta = 0.8$, the nominal interest rate is stuck at its ZLB whenever $\beta_{-1} > 0.9$, whereas when $\rho_\beta = 0.75$, the ZLB does not bind unless $\beta_{-1} > 1.3$. This result means that households will expect that the nominal interest rate will hit its ZLB less frequently and in situations when it does hit its ZLB, they will expect it to bind for fewer quarters. When $\rho_\beta = 0.75$, the average ZLB event is only 1.3 quarters, compared to 1.9 quarters when $\rho_\beta = 0.8$. In states of the economy where the ZLB binds, the dynamics are virtually identical to the dynamics that occur when $\rho_\beta = 0.8$. At the ZLB, the economy is far more sensitive to shocks that affect demand because the central bank can no longer blunt the effects of those shocks by lowering the nominal interest rate. Thus, the slopes of the policy functions for inflation and output are steeper and the real interest rate rises sharply at the ZLB, regardless of when it first binds.

Figure 21 reproduces figure 7 when $\rho_\beta = 0.75$. These results are informative because they show that increasing the weight on the steady-state output target does not affect our qualitative results. The decline in output at the ZLB for a given value of ϕ_π is less severe because the household expects that the nominal interest rate will rise in the near future and that the central bank will be able to stabilize output and inflation. A higher value of ϕ_y, however, reduces the positive effect of technology shocks on output and in unusually high technology states, these shocks can reduce output. Those findings are consistent with our results in section 4. If we increase the value of ρ_β, the unconventional effects of positive technology shocks at the ZLB are even more pronounced.

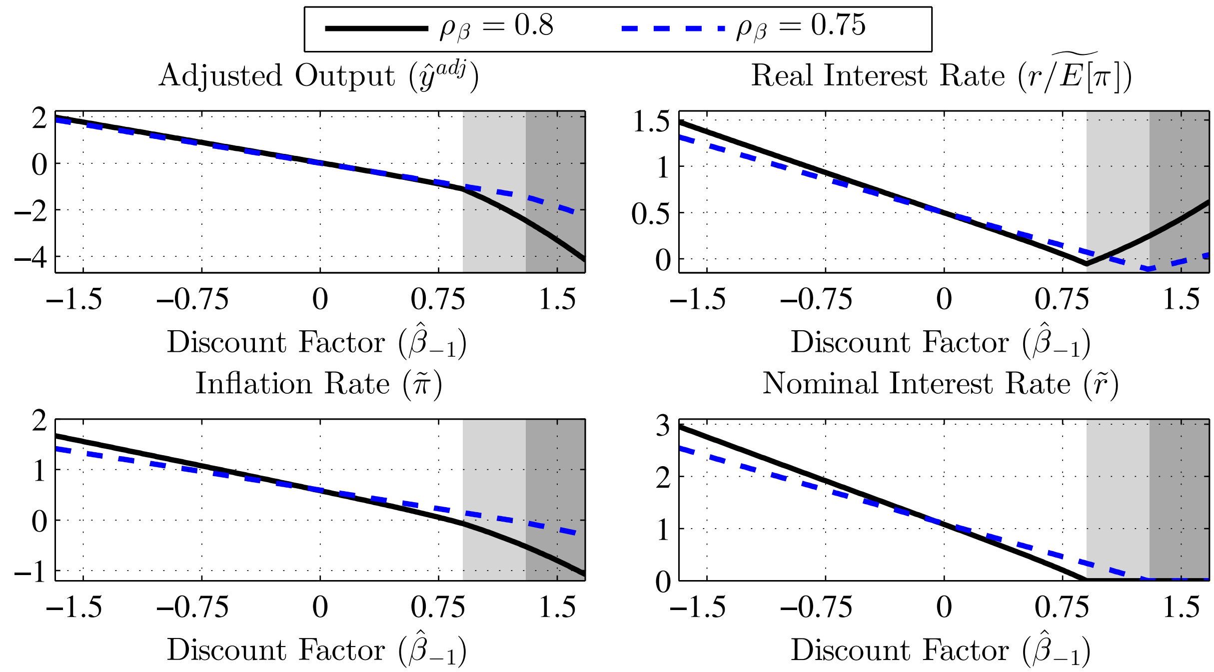

Figure 20: Model 1 ($y_t^* = \bar{y}$) decision rules as a function of the discount factor state ($\hat{\beta}_{-1}$) when the persistence is high (solid line) and low (dashed line). The technology state is fixed at its steady-state value ($\hat{z}_{-1} = 0$). Each variable is in percent deviations from its deterministic steady state, except inflation and the interest rates, which are net percentages. The light (dark) shaded region indicates where the ZLB binds when $\rho_\beta = 0.8$ ($\rho_\beta = 0.75$).

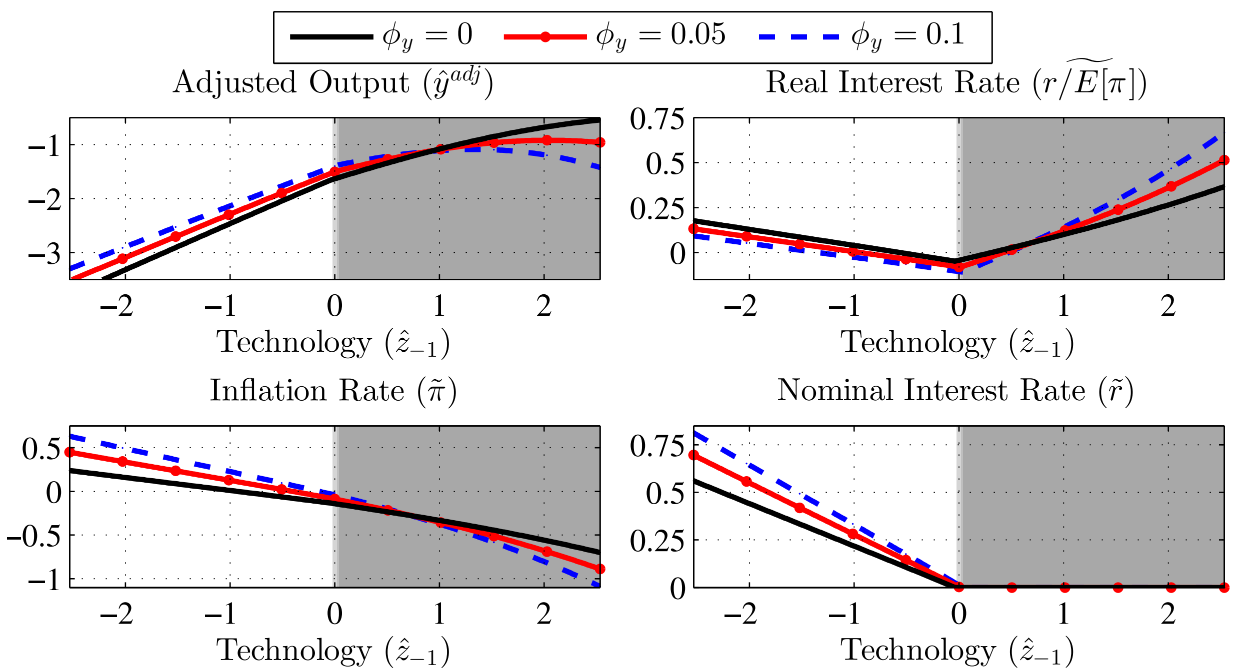

Figure 21: Model 1 ($y_t^* = \bar{y}$) decision rules as a function of the technology state ($\hat{z}_{-1}$). The discount factor persistence is $\rho_\beta = 0.75$. The discount factor state ($\hat{\beta}_{-1}$) is fixed at the minimum value that causes the ZLB to bind when $\hat{z}_{-1} = 0$ and $\phi_y = 0$. Each variable is in percent deviations from its deterministic steady state, except inflation and the interest rates, which are net percentages. The shaded region indicates where the ZLB binds for a given ϕ_y value.

D GENERALIZED IMPULSE RESPONSE FUNCTIONS

The GIRFs are based on the average of $10{,}000$ Monte Carlo simulations of the model. The advantage of this approach is that the realization of shocks are consistent with the household's expectation that the stochastic processes will mean revert when the GIRF is initialized away from the model's stochastic steady state. The general procedure for calculating GIRFs is laid out in Koop et al. (1996). We apply the following steps to our models:

1. Find the state vector at which to initialize each case:

   (a) *Non-ZLB Case*: Simulate the model without shocks until it converges to its stochastic steady state, $\mathbf{z}_0^{ss}$.

   (b) *ZLB Case*: Simulate the model for $500{,}000$ quarters using random draws of discount factor shocks. The initial state vector is the average state vector conditional on the ZLB binding, $\mathbf{z}_0^{zlb}$. The average discount factor when the ZLB binds is 1% above steady state.

2. Draw random shocks to technology and the discount factor, $\{\varepsilon_{z,t}, \varepsilon_{\beta,t}\}_{t=0}^{N}$, from their independent normal distributions. Simulate each case for R different draws of the sequence of shocks beginning at the alternative initial state vectors, $\mathbf{z}_0^{ss}$ and $\mathbf{z}_0^{zlb}$. This yields R equilibrium paths for each case, $\{\mathbf{x}_t^j(\mathbf{z}_0^{ss})\}_{t=0}^{N}$ and $\{\mathbf{x}_t^j(\mathbf{z}_0^{zlb})\}_{t=0}^{N}$, where $j \in \{1, 2, \dots, R\}$. We set $N = 20$ and $R = 10{,}000$.

3. Using the same R draws of shocks from step 2, replace the technology shock in period one with a 1% shock (i.e., set $\varepsilon_{z,1} = 0.01$ for all $j \in \{1, 2, \dots, R\}$). Simulate each case with these alternate sequences of shocks. This yields R equilibrium paths for each case, $\{\mathbf{x}_t^j(\mathbf{z}_0^{ss}, \varepsilon_{z,1})\}_{t=0}^{N}$ and $\{\mathbf{x}_t^j(\mathbf{z}_0^{zlb}, \varepsilon_{z,1})\}_{t=0}^{N}$.

4. Average across the R simulations from step 2 and step 3 to obtain average paths given by

$$\bar{\mathbf{x}}_t(\mathbf{z}_0^{ss}) = \frac{1}{R}\sum_{j=1}^{R}\mathbf{x}_t^j(\mathbf{z}_0^{ss}), \qquad \bar{\mathbf{x}}_t(\mathbf{z}_0^{ss}, \varepsilon_{z,1}) = \frac{1}{R}\sum_{j=1}^{R}\mathbf{x}_t^j(\mathbf{z}_0^{ss}, \varepsilon_{z,1}),$$
$$\bar{\mathbf{x}}_t(\mathbf{z}_0^{zlb}) = \frac{1}{R}\sum_{j=1}^{R}\mathbf{x}_t^j(\mathbf{z}_0^{zlb}), \qquad \bar{\mathbf{x}}_t(\mathbf{z}_0^{zlb}, \varepsilon_{z,1}) = \frac{1}{R}\sum_{j=1}^{R}\mathbf{x}_t^j(\mathbf{z}_0^{zlb}, \varepsilon_{z,1}).$$

5. The difference between the two average paths for each case is a GIRF. In our figures, a variable in either case with a hat is calculated as $100(\bar{\mathbf{x}}_t(\mathbf{z}_0^{s}, \varepsilon_{z,1})/\bar{\mathbf{x}}_t(\mathbf{z}_0^{s}) - 1)$ and with a tilde is calculated as $100(\bar{\mathbf{x}}_t(\mathbf{z}_0^{s}, \varepsilon_{z,1}) - \bar{\mathbf{x}}_t(\mathbf{z}_0^{s}))$, where $s \in \{ss, zlb\}$.