Auburn University
Department of Economics
Working Paper Series

Global Dynamics at the Zero Lower Bound

William Gavin‡, Benjamin Keen*, Alexander Richter†, and Nathaniel Throckmorton**

‡FRB St. Louis, *University of Oklahoma, †*Auburn University*, ***Indiana University*

AUWP 2013-17

Global Dynamics at the Zero Lower Bound*

William T. Gavin   Benjamin D. Keen

Alexander W. Richter   Nathaniel A. Throckmorton

October 19, 2013

ABSTRACT

This article presents global solutions to standard New Keynesian models with a zero lower bound (ZLB) constraint on the nominal interest rate. Rather than focus on specific sequences of shocks, we provide the solution for all combinations of technology and discount factor shocks and a thorough explanation of how dynamics change across the state space. Our solution method emphasizes accuracy to capture important expectational effects of going to and returning from the ZLB, which commonly used solution methods based on specific sequences of shocks cannot capture. We focus on the New Keynesian model without capital, but we also study the model with capital, with and without capital adjustment costs. Capital adds another mechanism for intertemporal substitution, which strengthens the expectational effects of the ZLB and impacts dynamics even before the ZLB is hit. We also evaluate how monetary policy affects the likelihood of hitting the ZLB. A policy rule based on a dual mandate is more likely to cause ZLB events when the central bank places greater emphasis on output stabilization.

Keywords: Monetary Policy; Zero Lower Bound; Global Solution Method

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, Indiana University, 100 S. Woodlawn, Wylie Hall 105, Bloomington, IN 47405 (nathrock@indiana.edu). We thank Klaus Adam, Javier Birchenall, Toni Braun, Ricardo Reis, and an anonymous referee 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, 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.

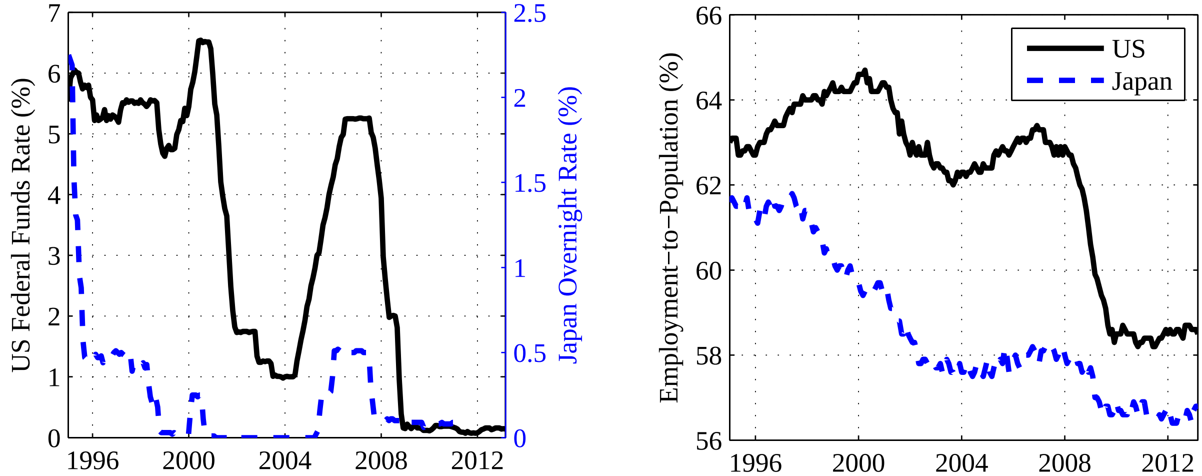

Figure 1: U.S. and Japanese interbank lending rates (left panel) and employment-to-population percentages (right panel). Sources: Federal Reserve Board of Governors, Bank of Japan, U.S. Bureau of Labor Statistics, and OECD.

1 INTRODUCTION

In the aftermath of the financial crisis, aggregate demand fell sharply. The Fed quickly responded by lowering its policy rate to its zero lower bound (ZLB) by the end of 2008. Five 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 rate and employment-to-population percentage from 1995-2012. The U.S. policy rate (solid line) has varied between 6.5 percent and 0 since 1995 and has been held below 25 basis points since the end of 2008. During this period, the inflation rate has been at or below the Fed's inflation target, which led policymakers to shift their focus from inflation to the real economy. The Bank of Japan sharply lowered its policy rate in 1990 (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 steadily fell from 62 percent to about 57.5 percent. The Japanese economy slightly rebounded in the mid-2000s, but following the financial crisis, the policy rate was cut and the employment-to-population percentage fell further.

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

A practical criticism is that a low nominal interest rate target may be misinterpreted by households. Bullard (2010) notes that attempting to stimulate the economy by promising to keep the interest rate at zero may backfire as inflation expectations may fall rather than rise. Del Negro et al. (2012) argue that recent promises to remain at the ZLB for an extended period have been interpreted as a signal that the central bank believes the economic outlook has worsened. These arguments suggest that people's expectations significantly affect the policy outcome. Schmitt-Grohé and Uribe (2012) show that when a central bank follows a Taylor rule, the consequences of hitting the ZLB may include moving to an undesirable low output/low inflation equilibrium.

Any ZLB analysis is complicated by the occasionally binding constraint on the monetary policy rule, which imposes a discontinuity in the policy functions. The literature has employed a variety of techniques to address this problem. Many studies log-linearize the equilibrium system, except the monetary policy rule, and solve either the deterministic model or the stochastic model based on specific sequences of shocks [Christiano et al. (2011); Eggertsson and Woodford (2003); Gertler and Karadi (2011)]. In these setups, the duration of the ZLB event is predetermined. Extensions of this work allow for stochastic ZLB events, but do not allow for recurring ZLB events [Braun and Waki (2006); Erceg and Linde (2010)]. Braun and Körber (2011) solve the nonlinear model, but use an extended shooting algorithm that still requires strong assumptions about future shocks.

There are three main drawbacks with these solution techniques. First, they violate the Lucas (1976) critique, which says that if policy changes, it is important to account for changes in expectations when studying the effects of the new policy. The sequences of shocks often used are very low probability events. Thus, when the ZLB is hit or continues to bind for several periods, the policy is virtually unaccounted for in the household's expectations. This has important implications for determinacy and dynamics [Richter and Throckmorton (2013)]. Second, using log-linearized models creates the potential for large approximation errors. Braun et al. (2012) and Fernández-Villaverde et al. (2012) provide explicit examples of the mistakes resulting from log-linearized models evaluated at the ZLB. Moreover, Braun et al. (2012) argue that log-linearized models often lead to incorrect inferences about existence of equilibrium, uniqueness, and local dynamics. Third, these methods prohibit Monte Carlo simulations of the model, which are necessary to study the conditional and unconditional probability distributions across alternative model specifications.

Our paper avoids these problems by solving for the global nonlinear solution to standard New Keynesian models that include an occasionally binding ZLB constraint on the nominal interest rate in the monetary policy rule.[1] Rather than focus on specific sequences of shocks, we provide the solution for all combinations of technology and discount factor shocks and a thorough explanation of how dynamics change across the entire state space. Our solution method emphasizes accuracy to capture important expectational effects of going to and returning from the ZLB, which commonly used solution methods based on specific sequences of shocks cannot capture.

In the variations of the New Keynesian model that we consider, we find that episodes at the ZLB are contractionary.[2] In the entire region of the state space where the ZLB binds, positive technology shocks, which would normally aid the recovery, have contractionary effects, which sharply contrasts with the findings of Braun and Körber (2011). At the ZLB, higher levels of technology increase the real interest rate, lower employment, and weaken aggregate demand, regardless of whether technology or discount factor shocks drive the nominal interest rate to zero. While no one believes interest rates fell to zero in December 2008 due to a series of positive technology shocks, our main interest is to learn how the economy reacts to technology shocks when the ZLB binds.

Much of the work on the ZLB uses models without capital.[3] We focus on the New Keynesian

[1] Recent papers that study the ZLB using global nonlinear solutions include Aruoba and Schorfheide (2013); Basu and Bundick (2012); Fernández-Villaverde et al. (2012); Gust et al. (2012); Mertens and Ravn (2013); Wolman (2005). The paper closest to ours is Fernández-Villaverde et al. (2012), which calculates the conditional and unconditional moments of ZLB events. Wolman (2005) shows that the real effects of the ZLB depend on the policy rule and nominal rigidities. Gust et al. (2012) estimates the extent to which the ZLB constrained the central bank's ability to stabilize the economy. Aruoba and Schorfheide (2013) and Mertens and Ravn (2013) show how the ZLB affects fiscal multipliers and Basu and Bundick (2012) show that the ZLB magnifies the effect of uncertainty on aggregate demand.

[2] There are some caveats to this conclusion. First, we have not explicitly modeled the Fed's unconventional policies, which seem to have kept deflation at bay. Second, Wieland (2012) 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.

[3] A notable exception is Christiano (2004), which generalizes Eggertsson and Woodford (2003) to include capital.

model without capital, but we also study the model with capital, with and without capital adjustment costs. Capital accumulation is a key feature because it gives households another margin to smooth consumption, which strengthens the expectational effects of the ZLB and impacts dynamics. Arbitrage implies that the real interest rate equals the expected future rental rate of capital. The sharp decline in demand when the ZLB binds leads to a sharp reduction in the rental rate of capital. Thus, the household places increasing weight on a lower rental rate as the ZLB nears, which leads to a sharp decline in the real interest rate even in states where the ZLB does not bind. Models that do not account for the expectational effects of the ZLB miss these dynamics. Capital adjustment costs make investment less attractive as a consumption smoothing mechanism, which causes a greater reduction in consumption and a larger increase in the real interest at the ZLB. Therefore, the presence of capital adjustment costs enhances the expectational effects of the ZLB, which alters dynamics in technology states that are even further from the ZLB.

We also evaluate how monetary policy affects the likelihood of encountering the ZLB. A policy rule based on a dual mandate is more likely to cause ZLB events when the central bank places greater emphasis on output stabilization. The policies that reduce the likelihood of hitting the ZLB also tend to deliver higher welfare.[4] The presence of capital increases the volatility of consumption and the nominal interest rate, decreasing the frequency of ZLB events for a given policy.

Section 2 briefly describes the alternative models. Section 3 describes the calibration and solution procedure, and sections 4 through 7 present the results. These sections report the complete model solutions across all technology and discount factor shocks, the dynamics at the ZLB, the likelihood of hitting the ZLB, and the welfare consequences of ZLB events. Section 8 concludes.

2 ECONOMIC MODELS

This section presents three alternative models. The baseline specification is a New Keynesian model with Rotemberg (1982) price adjustment costs. Model 1 assumes stochastic processes for the discount factor and technology but does not include capital. Models 2 and 3 incorporate capital accumulation into Model 1, and Model 3 also includes investment adjustment costs.

2.1 MODEL 1: BASELINE

A representative household chooses sequences $\{c_t, n_t, b_t\}_{t=0}^{\infty}$ to maximize expected lifetime utility, given by,

$$E_0 \sum_{t=0}^{\infty} \tilde{\beta}_t \left\{ \log c_t - \chi \frac{n_t^{1+\eta}}{1+\eta} \right\}, \tag{1}$$

where $1/\eta$ is the Frisch elasticity of labor supply, c_t is consumption of the final good, n_t is labor hours, $\tilde{\beta}_0 \equiv 1$, and $\tilde{\beta}_t = \prod_{i=1}^{t} \beta_i$ for $t > 0$. β_i is a time-varying subjective discount factor that evolves according to

$$\beta_i = \beta(\beta_{i-1}/\beta)^{\rho_\beta} \exp(\varepsilon_{\beta,i}),$$

where β is the stationary discount factor, $0 \leq \rho_\beta < 1$, and $\varepsilon_{\beta,i} \sim \mathbb{N}(0, \sigma_\beta^2)$. We normalize $\beta_{-1} = \beta$.

The representative household's choices are constrained by

$$c_t + b_t + \tau_t = w_t n_t + r_{t-1} b_{t-1}/\pi_t + d_t,$$

[4] Several papers discuss optimal policy with a ZLB constraint and provide analysis of the welfare losses at the ZLB [Adam and Billi (2006, 2007); Eggertsson and Woodford (2003); Jung et al. (2005); Nakov (2008); Werning (2012)].

where $\pi_t = p_t/p_{t-1}$ is the gross inflation rate, w_t is the real wage, τ_t is a lump-sum tax, b_t is a one-period real bond, r_t is the gross nominal interest rate, and d_t are profits from intermediate firms. Solving the household's utility maximization problem yields the following optimality conditions

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

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 $i \in [0,1]$ in the intermediate goods sector produces a differentiated good, $y_t(i)$, with identical technologies given by $y_t(i) = z_t n_t(i)$, where $n_t(i)$ is the level of employment used by firm i. z_t represents the level of technology, which is common across firms and follows

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

where $\bar{z}$ is steady-state technology, $0 \leq \rho_z < 1$, and $\varepsilon_{z,t} \sim \mathbb{N}(0, \sigma_z^2)$. Each intermediate firm chooses its labor supply to minimize its operating costs, $w_t n_t(i)$, subject to its production function.

Using a Dixit and Stiglitz (1977) aggregator, the representative final goods firm purchases $y_t(i)$ units from each intermediate goods firm to produce the final good, $y_t \equiv [\int_0^1 y_t(i)^{(\theta-1)/\theta} di]^{\theta/(\theta-1)}$, where $\theta > 1$ measures the elasticity of substitution between the intermediate goods. Maximizing profits for a given level of output yields the demand function for intermediate inputs given by $y_t(i) = (p_t(i)/p_t)^{-\theta} y_t$, where $p_t = [\int_0^1 p_t(i)^{1-\theta} di]^{1/(1-\theta)}$ is the price of the final good. Following Rotemberg (1982), each firm faces a cost to adjusting its price, which emphasizes the potentially negative effect that price changes can have on customer-firm relationships. Using the functional form in Ireland (1997), real profits of firm i are

$$d_t(i) = \left[\left(\frac{p_t(i)}{p_t}\right)^{1-\theta} - \Psi_t \left(\frac{p_t(i)}{p_t}\right)^{-\theta} - \frac{\varphi}{2}\left(\frac{p_t(i)}{\bar{\pi} p_{t-1}(i)} - 1\right)^2\right] y_t,$$

where $\varphi \geq 0$ determines the magnitude of the adjustment cost, Ψ_t is real marginal costs, and $\bar{\pi}$ is the steady-state gross inflation rate. Each intermediate goods firm chooses its price level, $p_t(i)$, to maximize the expected discounted present value of real profits $E_t \sum_{k=t}^{\infty} \lambda_{t,k} d_k(i)$, where $\lambda_{t,t} \equiv 1$, $\lambda_{t,t+1} = \beta_{t+1}(c_t/c_{t+1})^\sigma$ is the stochastic 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 intermediate goods firms make the same decisions and the optimality condition becomes

$$\varphi\left(\frac{\pi_t}{\bar{\pi}} - 1\right)\frac{\pi_t}{\bar{\pi}} = (1-\theta) + \theta\Psi_t + \varphi E_t\left[\lambda_{t,t+1}\left(\frac{\pi_{t+1}}{\bar{\pi}} - 1\right)\frac{\pi_{t+1}}{\bar{\pi}}\frac{y_{t+1}}{y_t}\right]. \tag{4}$$

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

Each period, the fiscal authority finances a constant level of discretionary spending, $\bar{g}$, by levying lump-sum taxes. The monetary authority sets policy according to

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

where π^* is the inflation rate target and ϕ_π and ϕ_y are the policy responses to inflation and output.[5]

In this paper, the output gap is defined as the deviation of output from its steady state. We use this measure because we believe policymakers, in the short-to-medium term, assume potential output grows at a relatively constant rate. Potential output measures are revised in the long run following incoming information about shocks, but the revisions occur well after the temporary economic effects from sticky prices have dissipated. In our model, a positive technology shock causes output to rise relative to its steady state and inflation to fall. For our baseline calibration, the lower inflation dominates the higher output leading to a lower nominal interest rate.

Alternatively, the output gap can be defined as the difference between actual output and the level of output in the absence of nominal frictions. Under this definition of the output gap, a positive technology shock would result in a negative output gap because price frictions would prevent actual output from rising as much as it would in the flexible price economy. Thus, the downward pressure on the nominal interest rate coming from low inflation would be reinforced by the additional downward pressure coming from a negative output gap.

The resource constraint is given by $c_t + \bar{g} = [1 - \varphi(\pi_t/\bar{\pi} - 1)^2/2]y_t = \tilde{y}_t$, where $\tilde{y}_t$ includes the value added by intermediate firms, which is their output minus quadratic price adjustment costs. Equilibrium is composed of the household's and firm's optimality conditions, the government's budget constraint, the bond market clearing condition ($b_t = 0$), and the resource constraint.

2.2 Model 2: Baseline with Capital

Models 2 adds capital accumulation to Model 1, but assumes a constant discount factor. Assuming, $\beta_t = \beta$ for all t, the household chooses sequences $\{c_t, k_t, i_t, n_t, b_t\}_{t=0}^{\infty}$ to maximize (1) subject to

$$c_t + i_t + b_t + \tau_t = w_t n_t + r_t^k k_{t-1} + r_{t-1} b_{t-1}/\pi_t + d_t, \tag{5}$$

$$k_t = (1-\delta)k_{t-1} + i_t, \tag{6}$$

where i_t is investment, k_t is the capital stock, and r_t^k is the real capital rental rate. The representative household's optimality conditions include (2), (3), and the consumption Euler equation, given by,

$$1 = \beta E_t\{(c_t/c_{t+1})(r_{t+1}^k + 1 - \delta)\}. \tag{7}$$

Each firm $i \in [0,1]$ in the intermediate goods sector produces a differentiated good, $y_t(i)$, with identical technologies given by $y_t(i) = z_t k_{t-1}(i)^\alpha n_t(i)^{1-\alpha}$, where $k_t(i)$ and $n_t(i)$ are the levels of capital and employment used by firm i. Every intermediate firm then chooses its capital and labor inputs to minimize its operating costs, $r_t^k k_{t-1}(i) + w_t n_t(i)$, subject to its production function. The firm pricing equation, (4), remains unchanged, except that the definition of the marginal cost changes. The aggregate resource constraint is now given by $c_t + i_t + \bar{g} = \tilde{y}_t$.

2.3 Model 3: Model 2 with Capital Adjustment Costs

Model 3 adds capital adjustment costs to Model 2. Following King and Wolman (1996), the budget constraint becomes

$$c_t + i_t + \Phi(i_i/k_{t-1})k_{t-1} + b_t + \tau_t = w_t n_t + r_t^k k_{t-1} + r_{t-1} b_{t-1}/\pi_t + d_t, \tag{8}$$

[5] Although we set the lower bound on the policy rate equal to zero, these same unconventional dynamics would occur if the bound was set to a small but positive value. The key is the existence of a lower bound, which prevents the Fed from responding to inflation. This is important because the Fed has not targeted a policy rate equal to zero.

Constant of Relative Risk Aversion	σ	1	Inflation Target	π^*	1.005
Frisch Elasticity of Labor Supply	$1/\eta$	3	Inflation Coefficient: MP Rule	ϕ_π	1.50
Elasticity of Substitution between Goods	θ	6	Output Coefficient: MP Rule	ϕ_y	0.125
Steady State Government Spending Share	$\bar{g}/\bar{y}$	0.2	Stationary Technology	$\bar{z}$	1
Rotemberg Adjustment Cost Coefficient	φ	58.25	Technology Persistence	ρ_z	0.80
Leisure Preference Parameter	χ	4.507	Technology Shock Standard Deviation	σ_z	0.012
Capital Depreciation Rate	δ	0.025	Stationary Discount Factor	β	0.99
Cost Share of Capital	α	0.33	Discount Factor Persistence	ρ_β	0.80
Capital Adjustment Cost	ϕ	5.6	Discount Factor Standard Deviation	σ_β	0.0025

Table 1: Baseline calibration

where $\Phi(\cdot)$ 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 measures the size of the adjustment cost. There are alternative specifications of adjustment costs used in the literature. We chose this specification because it does not add another state variable to our model, which allows us to present the entire model solution and easily compare the results from Model 3 to those from Model 2.

Once again, assuming $\beta_t = \beta$ for all t, the household chooses sequences $\{c_t, k_t, i_t, n_t, b_t\}_{t=0}^{\infty}$ to maximize (1) subject to (8) and (6). Optimality yields an equation for Tobin's q and a new consumption Euler equation, which replaces (7), given by,

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

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

The aggregate resource constraint is the same as in Model 2, except that both investment and output now include resources lost to capital adjustment costs.

3 Calibration and Solution Technique

The models in section 2 are calibrated at a quarterly frequency and the parameters are given in table 1. The risk-free real interest rate is set equal to 4 percent annually, which implies a stationary quarterly discount factor, β, equal to 0.99. We set the persistence of the discount factor, ρ_β, equal to 0.8 and the standard deviation of the shock, σ_β, equal to 0.0025. We follow Fernández-Villaverde et al. (2012) who chose these parameters so that a discount factor shock has a half life of about 3 quarters and an unconditional standard deviation of 0.42 percent. The Frisch elasticity of labor supply, $1/\eta$, is set to 3, which is consistent with estimates in 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 depreciation rate, δ, equals 2.5 percent per quarter.

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 percent. The costly price adjustment parameter, φ, is calibrated to 58.25, which is consistent with a Calvo (1983) price-setting specification where prices change on average once every four quarters.

In the policy sector, the steady-state gross inflation rate, $\bar{\pi}$, is set to 1.005, which implies an annual inflation rate target of 2 percent. The steady-state ratio of government spending to output is calibrated to 20 percent. In our baseline case, the coefficients on inflation and output in the policy rule are set to 1.5 and 0.125, which is consistent with Taylor (1993).

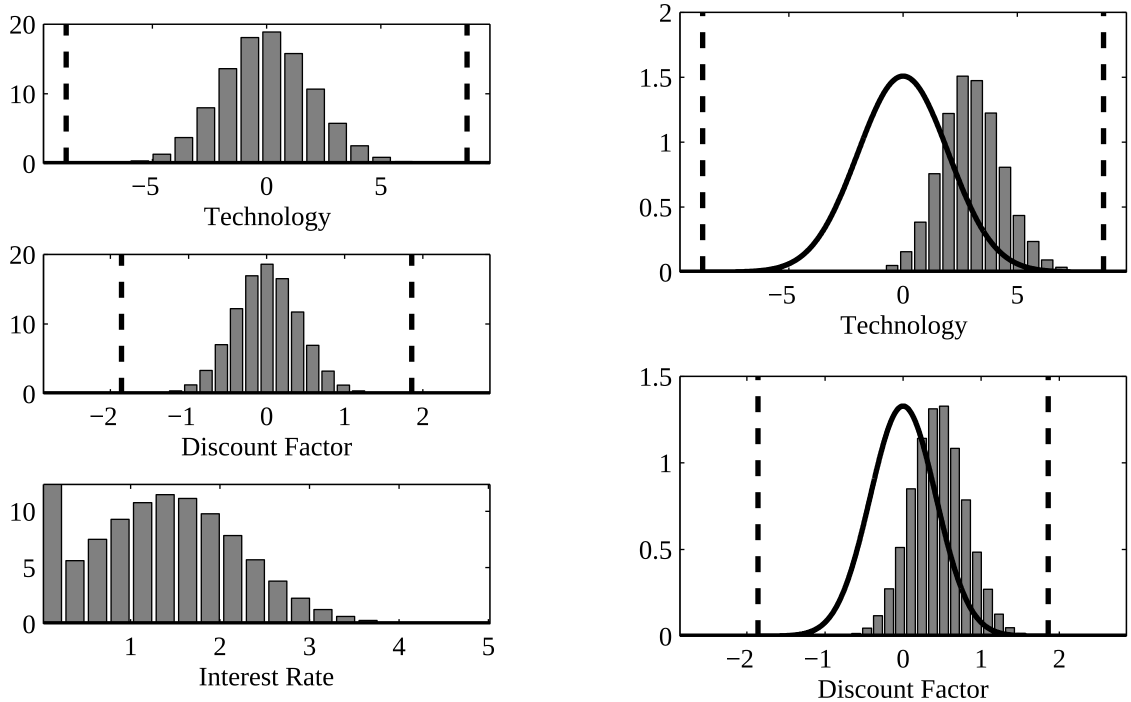

(a) Model 1: Unconditional

(b) Model 1: Conditional on ZLB

Figure 2: Model 1 distributions as a percentage of a 500,000 period simulation. The variables are in percent deviations from steady state. 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 distributions conditional on the ZLB.

Steady-state technology, $\bar{z}$, is normalized to 1. The likelihood of hitting the ZLB depends critically on the parameters of the technology process (σ_z^2 and ρ_z). When we set these parameters to values typically used in quantitative New Keynesian and Neoclassical models, determinacy is not guaranteed on the entire state space of our models. A determinate solution requires that σ_z^2 and ρ_z are not too large for a given coefficient on the output gap in the policy rule.[6] Thus, we set $\rho_z = 0.8$ and σ_z between 1 and 1.3 percent per quarter, depending on the model, which pushes the standard deviation of z_t toward values that are common in the literature.

We solve the model using the policy function iteration algorithm described in Richter et al. (2013), 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 policy functions until the tolerance criterion is met. We use piecewise linear interpolation to approximate future variables that show up in expectations, since this approach more accurately captures the kink in the policy functions than continuous functions, and Gauss-Hermite quadrature to numerically integrate. These techniques capture the expectational effects of going to and returning to the ZLB.

The models are simulated using draws from the distributions for the discount factor and tech-

[6]The ZLB is equivalent to a fixed interest rate regime with a truncated distribution on the nominal interest rate. The solution is not determinate with an insufficient probability of returning to an interest rate rule that aggressively responds to inflation. Davig and Leeper (2007) discuss determinacy in linear models where the monetary coefficient follows a 2-state Markov process. Richter and Throckmorton (2013) discuss determinacy in a nonlinear model and show that determinacy imposes a clear tradeoff between the frequency and duration of ZLB events.

nology shocks. The state space is discretized to minimize extrapolation of the policy functions during the simulation. As an example, we plot the simulated distributions of the state variables for Model 1 in figure 2 and show that they are contained within the bounds of the state space. We simulate the model for 500,000 periods to obtain an accurate sample of ZLB events.

Panel (a) shows the unconditional distributions of technology, the discount factor, and the nominal interest rate. The state space for technology lies within ± 8.8 percent of the steady-state value, which is normalized to unity in our simulations. The state space of the discount factor lies between ± 1.9 percent of the steady state, which is equal to 0.99.[7] Over these states, the net nominal interest rate is distributed over a range of 0 to 5 percent, with a large mass (12.4 percent of the simulated quarters) between 0 and 25 basis points. The steady-state quarterly rate is 1.5 percent.

Panel (b) shows the distribution of the discount factor and technology conditional on the ZLB binding. When technology is high enough and the central bank follows a Taylor rule, the nominal interest rate hits its ZLB. Fernández-Villaverde et al. (2012) also find that high levels of technology are associated with low interest rates. This is because high levels of technology are associated with low inflation and low nominal interest rates. Kiley (2003) uses U.S. data to show that periods of high labor productivity growth have been associated with relatively low inflation and agues that this result could be caused by the Fed's policy rule.

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

This section shows the complete nonlinear solution to Model 1 as a function of the two state variables, the discount factor and technology. The monetary policy rule is based on Taylor's (1993) original specification with $\phi_\pi = 1.5$ and $\phi_y = 0.125$ when $r_t > 1$. All of the variables are given in percent deviations from their deterministic steady state, except inflation, expected consumption growth, and the (net) interest rates, which are presented in levels.

Figure 3 shows three-dimensional contour plots of the non-predetermined variables over the entire state space, which provides a complete picture of the model solution. The shaded areas represent the states of technology and the discount factor where the (net) nominal interest rate equals zero. This region illustrates that the nominal interest rate only hits the ZLB when either technology or the discount factor are unusually high. These maps are useful because they provide the solution for every possible combination of the shocks, but they also can be difficult to read. Thus, figure 4 plots two-dimensional representations of two alternative cross sections of the contour plots. In figures 3 and 4, the solid (black) line shows the cross section where the discount factor state is fixed at its stationary value ($\beta_{-1} = 0.99$) and the dashed (blue) line shows the cross section where the discount factor is held constant at 1, which is the minimum value where the ZLB binds when technology is at its steady state. In figure 4, the darker (entire) shaded region indicates where the ZLB binds when $\beta_{-1} = 0.99$ ($\beta_{-1} = 1$).

We begin by examining the cross section where the discount factor is fixed at 0.99. Let us initially consider the region of the state space where the ZLB does not bind. In this cross section, the ZLB does not bind in states where technology ranges between 8.8 percent below and 3.5 percent above its steady state. When technology is below its steady state, workers are less productive and

[7] These bounds are chosen so that they encompass 99.999 percent of the probability mass of the technology and discount factor distributions and to minimize extrapolation of the policy functions in simulations of the model. We also specify a very dense discretized state space, so the location of the kink in the policy function is accurate. This is particularly important since it affects the frequency and duration of ZLB events.

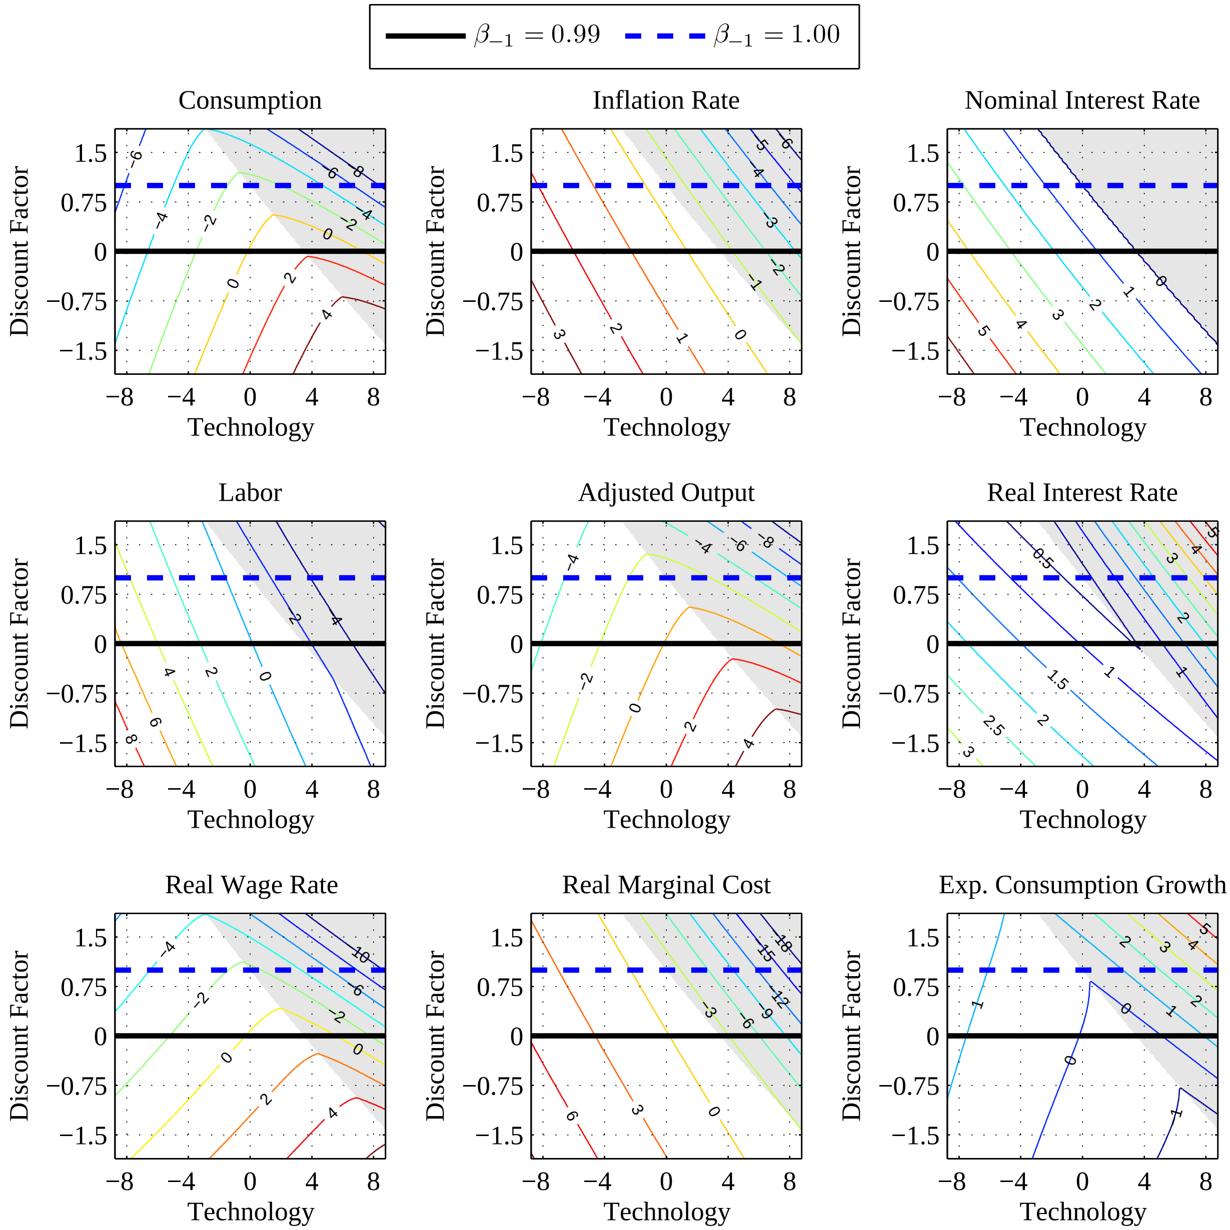

Figure 3: Model 1 non-predetermined variables as a function of technology and the discount factor states. All variables are in percent deviations from their deterministic steady state, except inflation, expected consumption growth, and the (net) interest rates, which are in levels. The shaded region indicates where the ZLB binds. The solid (black) and dashed (blue) horizontal lines correspond to the cross sections where $\beta_{-1} = 0.99$ and $\beta_{-1} = 1$.

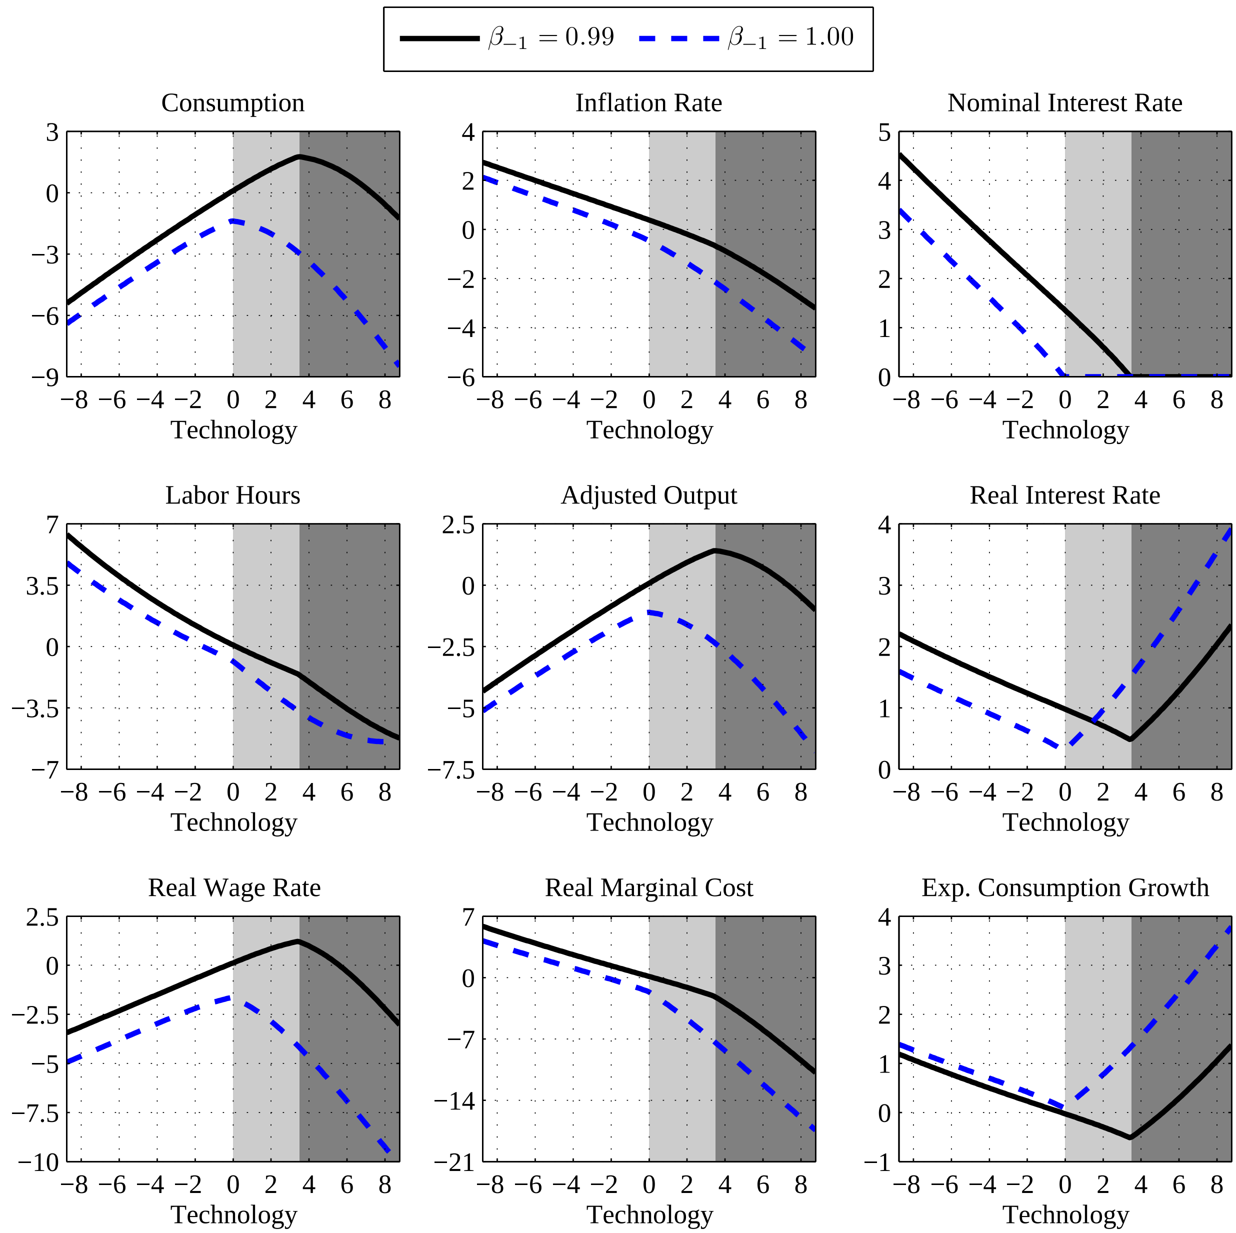

Figure 4: Model 1 non-predetermined variables as a function of the technology states. In the solid line the discount factor state (β_{-1}) is fixed at its deterministic steady state value and in the dashed line the discount factor state is fixed at 1. All variables are in percent deviations from their deterministic steady state, except inflation, expected consumption growth, and the (net) interest rates, which are in levels. The dark (entire) shaded region indicates where the ZLB binds when $\beta_{-1} = 0.99$ ($\beta_{-1} = 1$).

firms' per unit marginal cost of production is higher. At low levels of technology, firms have higher prices and lower output and labor demand. With less output available for consumption, the household works more to moderate the decline in consumption. The higher labor supply dominates the drop in labor demand so that the equilibrium level of labor is higher and the real wage is lower. The household also believes technology will gradually return to its steady state and, as a result, expects its future consumption to increase. Higher expected future consumption is reflected in an elevated *ex-ante* real interest rate. At higher levels of technology and before the ZLB binds, workers are more productive and firms choose lower prices and higher output. The household consumes more but also desires more leisure. In this part of the state space, the decline in the labor supply dominates the increase labor demand so that labor hours are lower and real wages are higher. The natural tendency for technology to return to its steady state means that households expect lower consumption growth in the future and observe a lower real interest rate.

Next consider the states where the ZLB binds (dark shaded region), which includes technology states that are more than 3.5 percent above the steady state. In this case, higher technology continues to push down per unit production costs and firms react by lowering their prices. The additional decline in expected inflation combined with a zero nominal interest rate forces the *ex-ante* real interest rate to rise. The household elects to sharply reduce its consumption and increase its labor supply to capitalize on those increased returns. Firms respond to the reduction in demand by further lowering their prices and decreasing their output and 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.

Now turn to the cross section where the discount factor is held constant at 1 (dashed line). A higher discount factor makes the household more patient, which reduces demand across the entire state space and causes the ZLB to bind at a lower technology state. The ZLB (entire shaded) region now includes all positive technology states. The main reason for showing this cross section is to highlight that the unconventional response of the economy to a positive technology shock at the ZLB does not depend on a high level of technology to drive the economy to its ZLB. The policy functions in this cross section display the same qualitative properties as the cross section where $\beta_{-1} = 0.99$. Looking at the highest discount factor shown in figure 3 ($\beta_{-1} = 1.008$), it is clear that the same dynamics continue to apply even when technology is below its steady state. Indeed, this is the area of the state space that is often considered in ZLB studies. If there was an even higher discount factor shock as modeled by Fernández-Villaverde et al. (2012), Christiano et al. (2011), and Schmitt-Grohé and Uribe (2012), then these same dynamics would appear.

Figure 5 compares the impulse responses to a one-time 1 percent positive technology shock under two cases—the baseline case (dashed line), which is initialized at the stochastic steady state with β at its deterministic steady state, and the ZLB case (solid line), where a sequence of discount factor shocks keep β constant and equal to 1 (i.e., the minimum value where the ZLB binds when technology is at its steady state). The horizontal dash-dotted lines are the stochastic steady-state values of inflation and the (net) interest rates, which differ from the deterministic steady state due to expectational effects of hitting the ZLB. In short, this exercise compares the conventional responses to a positive technology shock when the ZLB never binds to the responses based on a counterfactual where the ZLB always binds due to successive discount factor shocks. Intuitively, the series of discount factor shocks can be thought of as a persistent reduction in consumer confidence, an ongoing global savings glut, or a decision by the Fed to hold the policy rate at zero. The primary

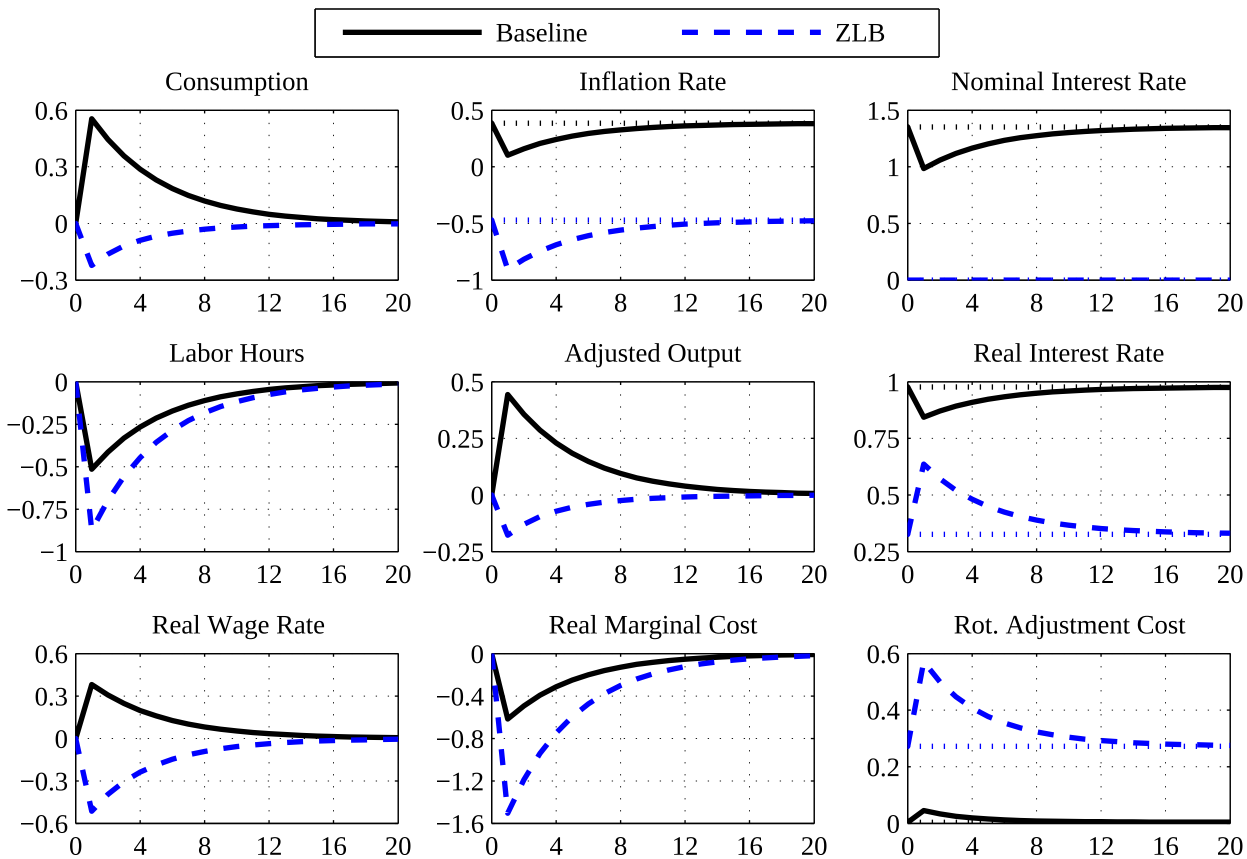

Figure 5: Model 1 impulse responses to a one-time 1 percent positive technology shock in period one. The baseline case (dashed line), where the ZLB does not bind, is initialized at the stochastic steady state with a discount factor equal to 0.99. In the ZLB case (solid line), the discount factor is constant and equal to one, so the ZLB always binds. The horizontal dash-dotted lines are the stochastic steady values of inflation and the (net) interest rates. The other variables are given in percent deviations from their respective stochastic steady state values.

advantage of looking at impulse responses over the policy functions is that they provide a clearer quantitative sense about how economic dynamics differ when the ZLB binds.

The results in the baseline case are standard and follow the intuition from the policy functions. A persistent technology shock lowers firms' per unit marginal cost of production, increases output, 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 there is also a decline in the *ex-ante* real rate, which increases consumption. A positive technology shock acts as a positive labor productivity shock, which decreases the equilibrium level of labor and raises real wages.

In the ZLB case, a positive technology shock has unconventional effects, as the policy functions predict. At the stochastic steady state with β fixed at 1, the higher discount factor imposes slight deflation. A positive technology shock leads to further deflation. With the nominal interest rate constrained at zero, the *ex-ante* real interest rate sharply rises. In response, consumption and labor both fall, but the effect on consumption and output is dampened relative to the baseline case. The greater deflation is associated with a rise in price adjustment costs. In both cases, the level of technology returns to its steady state about 20 quarters after the initial impact of the shock.

5 MODEL 2: STATES OF THE ECONOMY AND THE ZLB

This section shows how the model solution changes when capital accumulation is added to Model 1. In Model 1, the only way for the household to smooth consumption is by varying its labor supply. In Model 2, capital gives the household another margin to smooth consumption. This model contains two state variables—the lagged capital stock (k_{-1}), which is endogenous, and technology (z_0). Figure 6 shows the three-dimensional contour plots of the non-predetermined variables over the entire state space. Capturing a complete picture of the model solution is particularly important in models with an endogenous state variable. In Model 1, the discount factor and technology states are independent and, therefore, any one realization of the discount factor is just as likely at high and low technology states. In Model 2, the capital and technology states are not independent. At low (high) technology states, the capital state is most likely below (above) its steady-state value. The contour plots capture these endogenous dynamics and provide a complete picture of the model solution. In general, the patterns for consumption, inflation, and the nominal interest rate are qualitatively similar to Model 1. However, the household's ability to invest means consumption is less volatile and there are important expectational effects that are not present in Model 1.

We begin by examining the behavior of the economy when the ZLB does not bind. Regardless of the capital state, higher technology states are associated with a lower marginal cost and lower inflation. Firms increase their production and labor demand. With more output available to divide between consumption and investment, both variables increase. To smooth its consumption across time, the household reduces its labor supply and increases its investment in capital. Whether higher technology states increase or decrease the equilibrium level of labor when the ZLB does not bind depends on how the capital state co-moves with the technology state.

When capital is held fixed at its steady-state value, the increase in labor demand dominates the decrease in labor supply, causing the equilibrium level of labor and the real wage rate to rise. Alternatively, if the capital state rises with technology, the decrease in labor supply dominates the increase in labor demand, causing the real wage rate to rise and the equilibrium level of labor to fall. These two alternative cross sections of the contour plots are shown in figure 6 and in figure 7. The solid (black) line shows the cross section where capital is held fixed at its steady-state value and the dashed (blue) line shows the cross section where capital increases along the diagonal of the state space. In figure 7, the darker (entire) shaded region indicates the area of the state space where the ZLB binds in the steady-state (diagonal) cross section.

The differences between the steady-state and diagonal cross sections are shown in figure 7. In the diagonal cross section where the capital state increases with the technology state, the marginal product of capital is lower in higher technology states. From the household's perspective, this makes investment less attractive as a consumption smoothing channel. The household responds by increasing consumption and decreasing labor supply more than in the cross section where the capital state is held fixed at its steady-state value. This different behavior is apparent from the slopes of the investment, consumption, and labor policy functions. The policy function for the rental rate of capital is also qualitatively different between these cross sections. In the steady-state cross section, higher levels of technology and labor raise the marginal product of capital and the rental rate of capital due to complementarity. In the diagonal cross section, higher technology states are associated with more rapid increases in the capital stock and declining labor. The negative effects of capital and labor dominate the positive effects from technology so that the marginal product of capital and the capital rental rate decline when the ZLB does not bind.

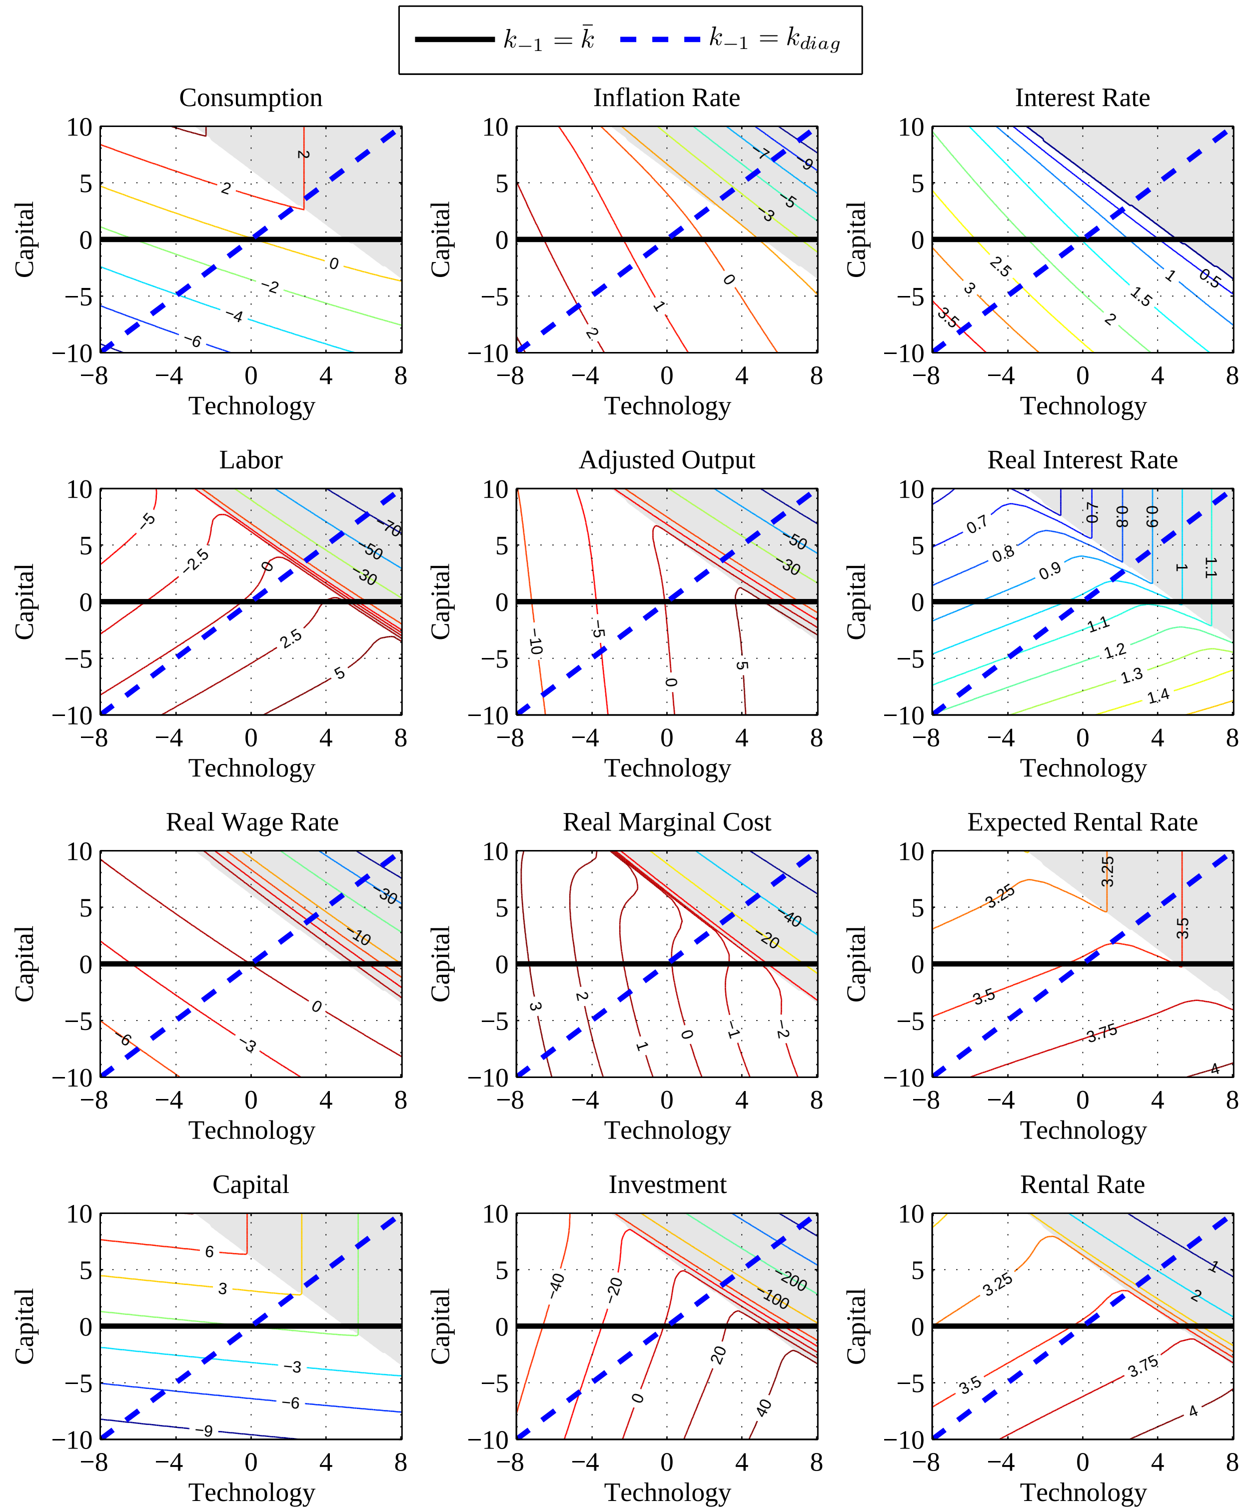

Figure 6: Model 2 non-predetermined variables as a function of capital and the technology states. The solid line indicates the cross section of the state space with capital in steady state, and the dashed line indicates the diagonal cross section where capital positively co-moves with technology in the state space. All variables are in percent deviations from their deterministic steady state, except inflation, the expected rental rate, and the (net) interest rates, which are in levels. The shaded region indicates where the ZLB binds.

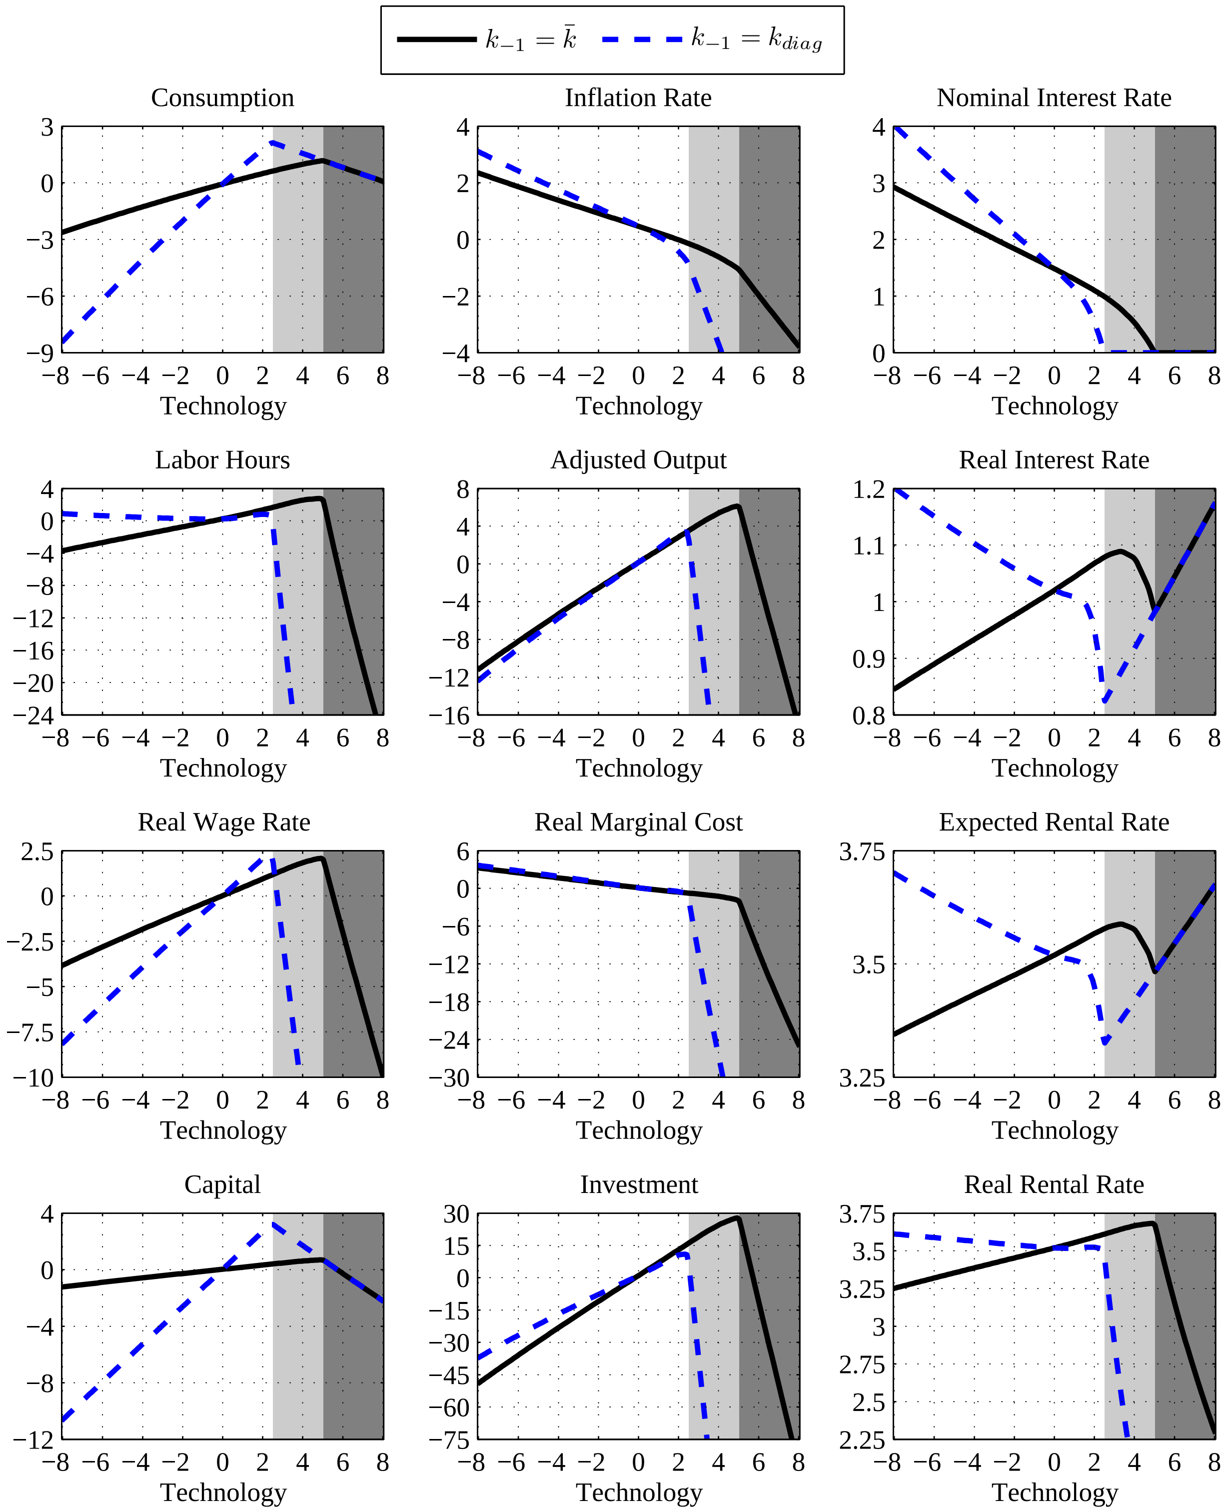

Figure 7: Model 2 non-predetermined variables as a function of technology. The solid line indicates the cross section of the state space with capital in steady state, and the dashed line indicates the diagonal cross section where capital positively co-moves with technology in the state space. All variables are in percent deviations from their deterministic steady state, except inflation, the expected rental rate, and the (net) interest rates, which are in levels. The dark (entire) shaded region indicates where the ZLB binds when $k_{-1} = \bar{k}$ ($k_{-1} = k_{diag}$).

Another difference between these two cross sections is the behavior of the *ex-ante* real interest rate. Unlike Model 1, which only has one asset, Model 2 has two assets—capital and bonds. Arbitrage implies that the expected rates of return on investment and bonds are equal. Thus, the expected future rental rate of capital positively co-moves with the current *ex-ante* real interest rate. In the steady-state (diagonal) cross section, the real interest rate rises (falls) in higher technology states, because the household expects technology to return to its steady-state level, as they did in Model 1. However, it is interesting that in both cross sections, the *ex-ante* real interest rate falls in technology states that are high, but not high enough for the ZLB to bind. In these states, the household places substantially more weight on the shocks that cause the ZLB to bind. Thus, the unconventional dynamics that occur at the ZLB cause the household to expect the rental rate of capital to fall and consumption growth to slow. Both of these effects cause the real interest rate to fall before the ZLB is hit. In short, the household and firms anticipate the economic contraction at the ZLB by reducing investment and employment. These declines lead to sharper reductions in inflation, the nominal interest rate, and the real interest rate well before the ZLB it hit.

In the steady-state (diagonal) cross section, the ZLB binds when technology is more than 2.5 (5.0) percent above steady state. The qualitative properties of the policy functions when the ZLB binds are nearly identical across all possible cross sections. The mechanism that distorts the economy is essentially the same as Model 1. As the real marginal cost continues to decline in higher technology states, inflation falls. With the nominal interest rate pegged at zero, the *ex-ante* real interest rate rises.[8] When the household's demand falls, both consumption and investment decrease. Firms respond to the lower demand by further reducing their prices and sharply cutting their labor demand, which causes the equilibrium level of labor and the real wage to fall. Lower consumption and investment pushes down output, despite the high technology state. As output falls, the household further reduces its investment to smooth consumption. Thus, the paradoxes of toil and saving both occur—despite the household wanting to work more to smooth consumption and save more to benefit from higher real interest rates, output contracts and both employment and investment fall. These results demonstrate that Model 2 faces the same unconventional dynamics as Model 1.

6 MODEL 3: STATES OF THE ECONOMY AND THE ZLB

The rapid adjustment in capital and investment shown in figures 6 and 7 is at odds with the data and motivates us to add capital adjustment costs to Model 2. Given our adjustment costs specification, Model 3 contains the same state variables as Model 2—the lagged capital stock and technology. The complete solution to the model is shown in figure 8. The curvature of the policy functions in states where the nominal interest rate is near the ZLB illustrates strong expectational effects.

A comparison of figures 6 and 8 reveals that the behavior of consumption and the real interest rate are noticeably different in Models 2 and 3. First, consumption decreases with technology, but is independent of the capital stock when the ZLB binds in Model 2. In Model 3, consumption declines more if either technology or the capital stock increases. Second, the real interest rate is (mostly) an increasing function of technology and a decreasing function of capital in states when the ZLB does not bind in Model 2. In Model 3, it is a decreasing function of technology and the capital stock. These differences imply that the dynamics in Model 3 are closer to Model 1.

[8] Despite that fact that the rental rate of capital falls sharply at the ZLB, the household expects the rental rate to increase since they expect technology to return to its steady state. This is consistent with a rising real interest rate.

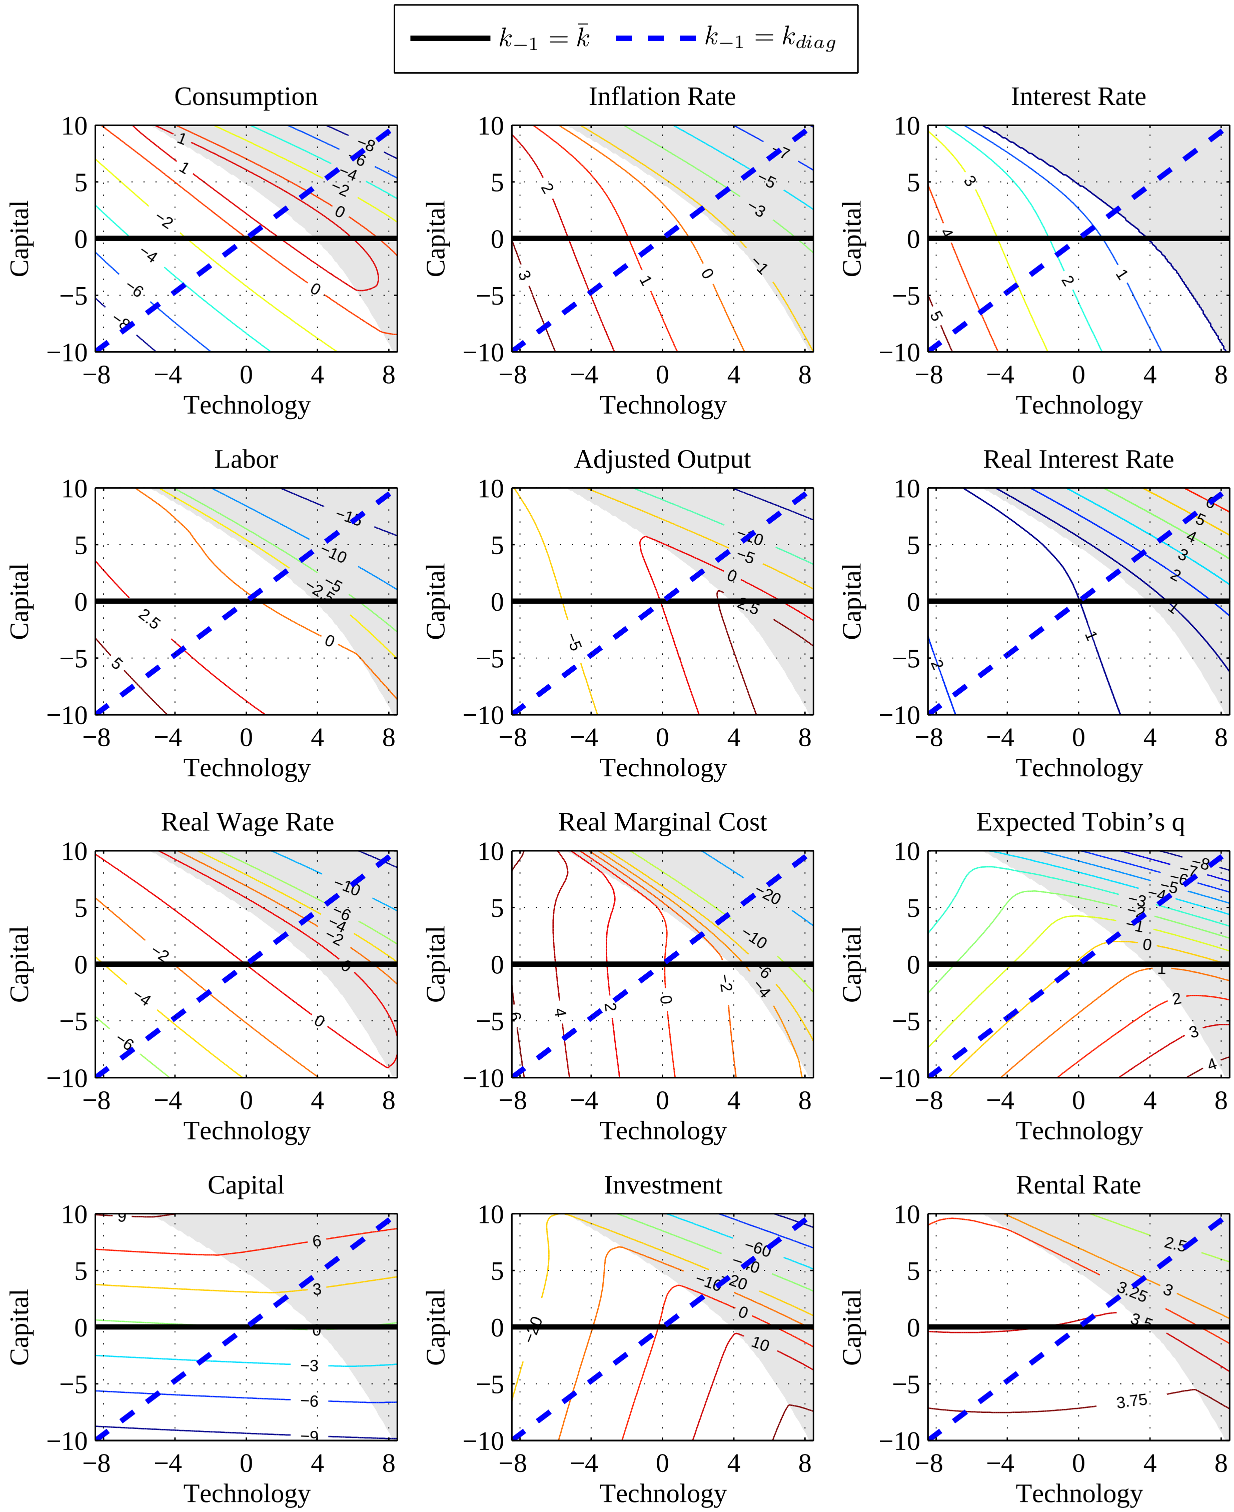

Figure 8: Model 3 non-predetermined variables as a function of capital and the technology states. The solid line indicates the cross section of the state space with capital in steady state, and the dashed line indicates the diagonal cross section where capital positively co-moves with technology in the state space. All variables are in percent deviations from their deterministic steady state, except inflation, the expected rental rate, and the (net) interest rates, which are in levels. The shaded region indicates where the ZLB binds.

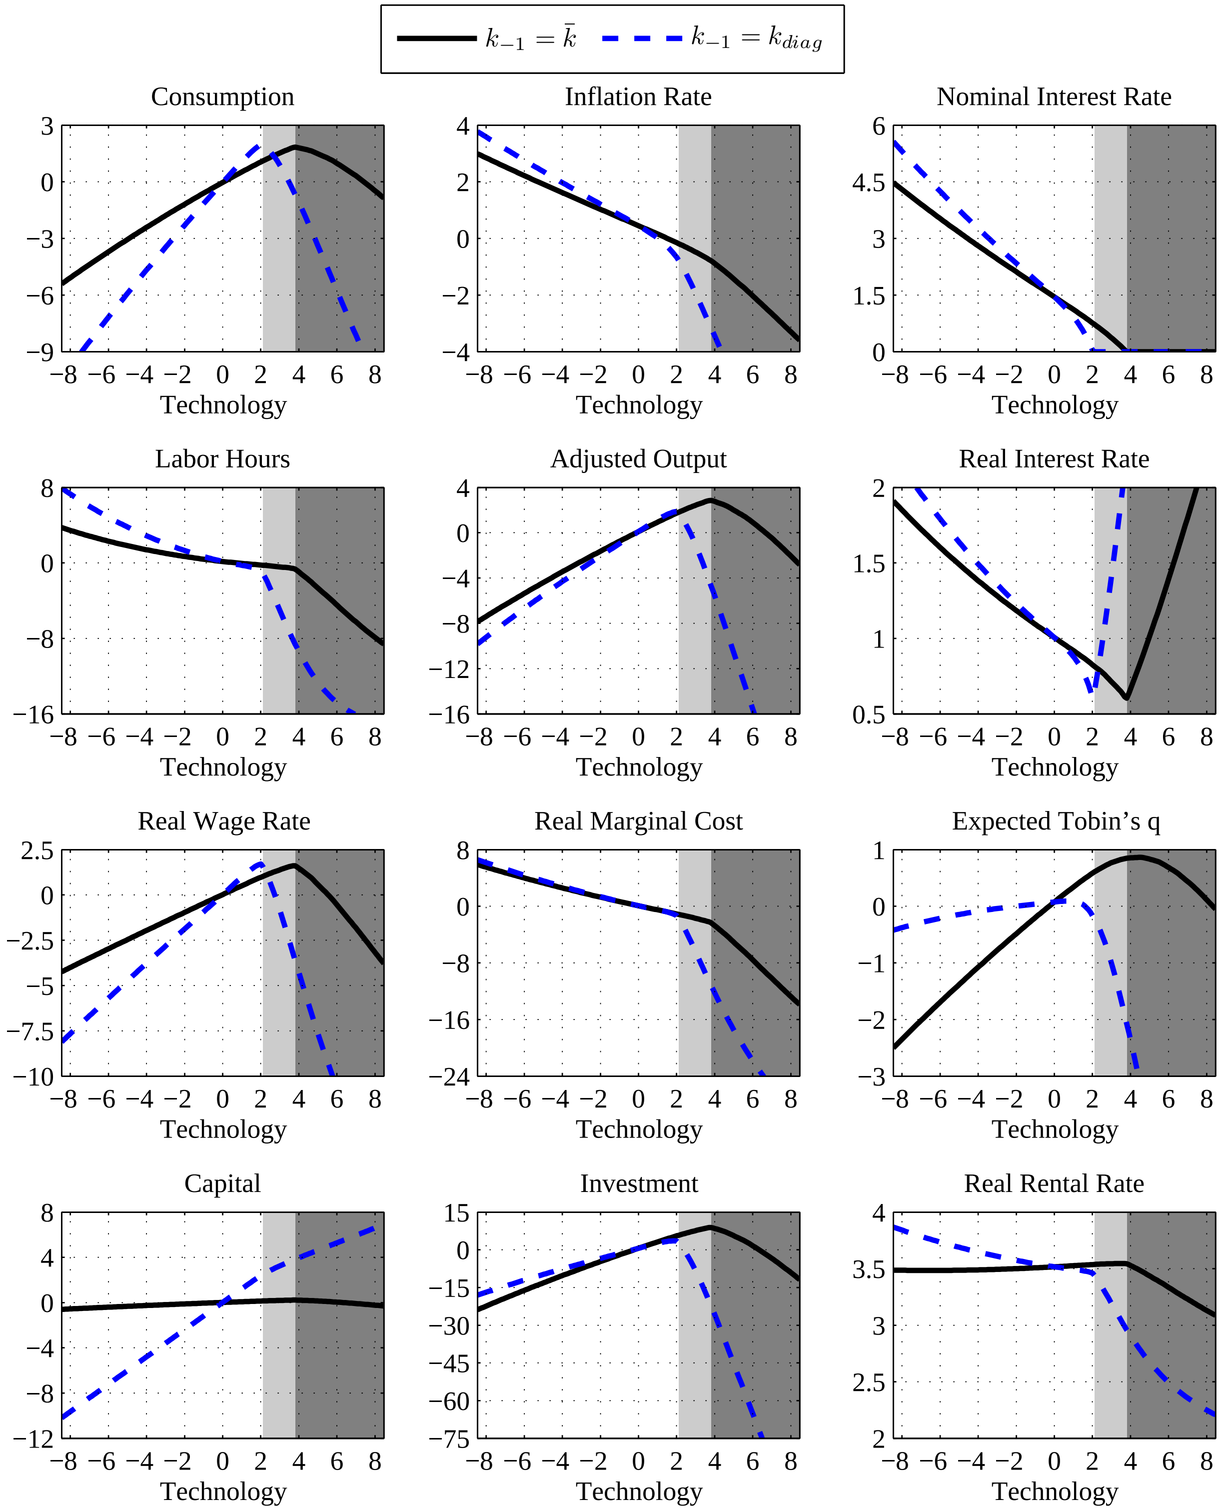

Figure 9: Model 3 non-predetermined variables as a function of technology. The solid line indicates the cross section of the state space with capital in steady state, and the dashed line indicates the diagonal cross section where capital positively co-moves with technology in the state space. All variables are in percent deviations from their deterministic steady state, except inflation, the expected rental rate, and the (net) interest rates, which are in levels. The dark (entire) shaded region indicates where the ZLB binds when $k_{-1} = \bar{k}$ ($k_{-1} = k_{diag}$).

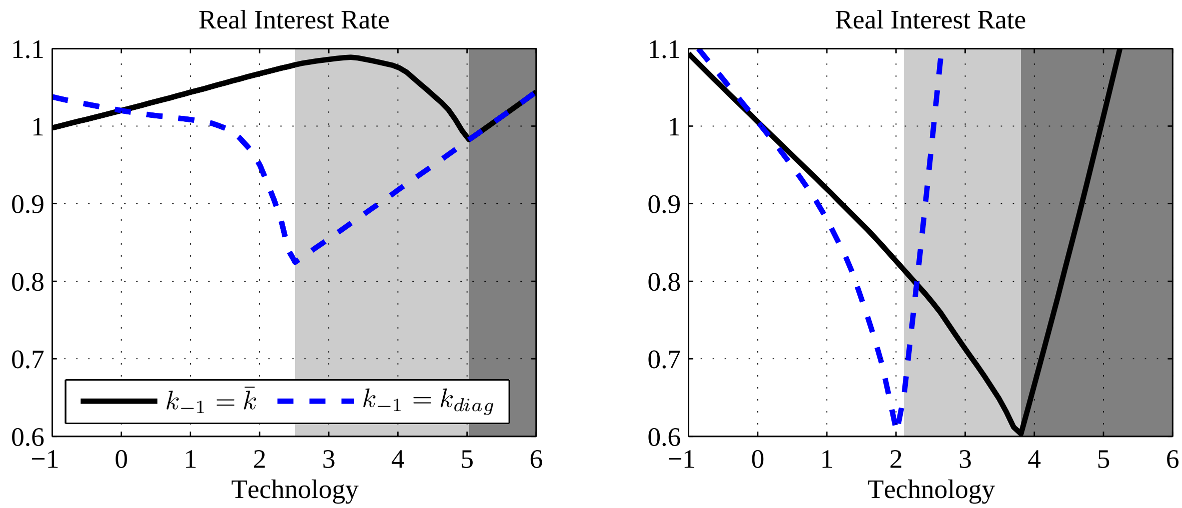

(a) Model 2: No capital adjustment costs

(b) Model 3: With capital adjustment costs

Figure 10: Comparison of the real interest rate near and at the ZLB between Models 2 and 3

The presence of capital adjustment costs in Model 3 reduces the volatility of capital and investment across the technology state, which means the policy functions are less variable in the alternative cross sections. Figure 9 plots the same cross sections of the Model 3 solution that are shown for Model 2 in figure 7. In states when the nominal interest rate is far from the ZLB, the dynamics are similar to Model 1. As technology increases, firms' per unit marginal cost declines, reducing inflation and increasing labor demand. Consumption and investment both increase, the labor supply decline, equilibrium hours fall, and the real wage rises.

In technology states where the nominal interest rate is near the ZLB, the dynamics of Model 3 are closer to those in Model 2 than Model 1, because of strong expectational effects. To understand why, we need to know how dynamics change when the ZLB binds. When capital is fixed at its steady-state value (solid line), the ZLB binds (dark shaded region) in technology states that are more than 3.8 percent above its steady state. At the ZLB, agents prefer to save more as the real interest rate rises with the technology state, but capital adjustment costs make investment less attractive as a consumption smoothing mechanism. This means consumption falls further and the real interest rate increases more at the ZLB, relative to Model 2. In the alternative cross section (dashed line), where the capital state increases with the technology state, the ZLB binds (entire shaded region) in technology states that are more than 2.1 percent above steady state. In this cross section, the unconventional dynamics at the ZLB are even more stark.[9]

Figure 10 plots the real interest rate in Models 2 and 3 and zooms in on the technology states just before and after the ZLB binds. The presence of capital adjustment costs in Model 3 makes investment less attractive as a consumption smoothing mechanism. That causes the real interest rate to rise faster in Model 3 when the technology state is high enough that the ZLB binds. The household and firms expect these unconventional dynamics to occur at the ZLB. Since the dynamics are more dramatic in a model with capital adjustment costs, the expectational effects are stronger in Model 3 than in Model 2. This result has two important implications for our analysis. First, the magnitude of the the real interest rate decline is greater in Model 3 when ZLB does not

[9]This cross section is less likely than in Model 2, since capital adjustment costs decrease the volatility of capital.

	ZLB Binds	ZLB Spells		Std. Dev. (% of mean)		
ϕ_y	% of quarters	Average	Longest	Output	Inflation	Nom. Int. Rate
0.125	8.417	2.21	21	0.97	0.68	0.82
0.100	6.667	2.11	20	1.01	0.64	0.80
0.075	5.941	2.06	20	1.04	0.61	0.79
0.050	5.416	2.03	20	1.08	0.59	0.78
0.025	4.993	2.01	20	1.12	0.57	0.78
0.000	4.637	1.98	20	1.16	0.55	0.77

(a) Model 1: No capital, technology and discount factor shocks. $\phi_\pi = 1.50$, $\rho_z = 0.80$, $\sigma_z = 0.01191$, $\rho_\beta = 0.80$, and $\sigma_\beta = 0.0025$.

0.125	2.902	1.59	10	2.45	0.61	0.53
0.100	0.113	1.42	5	2.60	0.42	0.38
0.075	0.005	1.44	3	2.81	0.34	0.33
0.050	0.000	0.00	0	3.05	0.26	0.28
0.025	0.000	0.00	0	3.34	0.19	0.25
0.000	0.000	0.00	0	3.69	0.17	0.26

(b) Model 2: Capital, only technology shocks. $\phi_\pi = 1.50$, $\rho_z = 0.80$, and $\sigma_z = 0.01094$.

Table 2: Volatility implications of a dual mandate.

bind. Second, the real interest rate in Model 3 falls at a faster rate as the nominal rate approaches the ZLB than it does in Model 2. These results demonstrate that the expectational effects of the ZLB are more pronounced and the model dynamics are more significant when real frictions, such as capital adjustment costs, are added to the model. The existence of these expectational effects is important because most medium-scale dynamic models that are used for policy analysis include several types of real frictions, in addition to the one we study in this section.

7 THE LIKELIHOOD OF HITTING THE ZLB AND WELFARE

Some economists argue the ZLB constrains the central bank's ability to achieve its employment and output targets [Blanchard et al. (2010); Coibion et al. (2012); Reifschneider and Williams (2000); Summers (1991); Williams (2009)]. They recommend raising the inflation target, but doing so is harmful in a New Keynesian economy. Others, using linear approximations, have shown that adopting a price level target can reduce the likelihood of ZLB events, even at a zero inflation target.

This section examines the likelihood of hitting the ZLB in Models 1 and 2 using 500,000 quarter simulations of the models. The results are not strictly comparable across models because they are based on different assumptions about the shocks, but they provide a qualitative indication for how the frequency and duration of ZLB events differ. Our main result is that policymakers can reduce the likelihood of hitting the ZLB by de-emphasizing the dual mandate. This is accomplished by either lowering the weight on the output gap (ϕ_y) or raising the weight on inflation (ϕ_π).

Table 2 shows the effect of reducing the weight on output while holding the weight on inflation constant at $\phi_\pi = 1.5$. We begin with the original Taylor (1993) specification, $\phi_y = 0.125$, and reduce this coefficient by increments of 0.025. In Model 1, the ZLB binds in 8.4 percent of the quarters in our simulation when $\phi_y = 0.125$. This value monotonically falls with ϕ_y and equals 4.6 percent when $\phi_y = 0$. The longest ZLB event is fairly stable; when $\phi_y = 0.125$, is it 21 quarters

ϕ_π	ZLB Binds % of quarters	ZLB Spells Average	Longest	Std. Dev. (% of mean) Output	Inflation	Nom. Int. Rate
1.500	8.417	2.21	21	0.97	0.68	0.82
1.750	3.763	1.91	20	1.08	0.50	0.74
2.000	2.811	1.82	16	1.16	0.42	0.71
2.250	2.287	1.78	16	1.22	0.37	0.69
2.500	1.941	1.74	16	1.26	0.32	0.67
3.000	1.533	1.69	15	1.34	0.26	0.64

(a) Model 1: No capital, technology and discount factor shocks. $\phi_y = 0.125$, $\rho_z = 0.80$, $\sigma_z = 0.01191$, $\rho_\beta = 0.80$, and $\sigma_\beta = 0.0025$.

1.500	2.902	1.59	10	2.45	0.61	0.53
1.750	0.006	1.65	3	2.62	0.36	0.33
2.000	0.000	2.00	2	2.73	0.29	0.27
2.250	0.000	2.00	2	2.81	0.25	0.22
2.500	0.000	1.00	1	2.87	0.21	0.20
3.000	0.001	1.67	2	2.95	0.17	0.16

(b) Model 2: Capital, only technology shocks. $\phi_y = 0.125$, $\rho_z = 0.80$, and $\sigma_z = 0.01094$.

Table 3: Volatility implications of alternative weights on the inflation gap.

and when $\phi_y \leq 0.1$ it is 20 quarters. Technology shocks create a tradeoff between the volatility of output and inflation. We find that reducing the weight on the output gap raises output volatility by about 20 percent, but reduces inflation volatility by 19 percent.

In Model 2, capital accumulation makes the nominal interest rate less volatile and reduces the likelihood of hitting the ZLB, which is only about 3 percent in the baseline calibration even though the standard deviation of the technology shock is about 1.1 percent per quarter. Once again, reducing the weight on the output gap reduces the likelihood of ZLB events. If the weight is at or below 0.05, the ZLB never binds in our simulation. The longest episode at the ZLB is 10 quarters, half the length in Model 1. Reducing the weight from 0.125 to 0 raises the standard deviation of output by 51 percent and reduces the standard deviation of inflation by about 72 percent.

Table 3 reports the results when we fix $\phi_y = 0.125$ and change the weight on the inflation gap. The longest and average ZLB events are more sensitive to rasing the weight on inflation than to lowering the weight on the output gap. In Model 1, the probability of hitting the ZLB is 8.4 percent when $\phi_\pi = 1.5$ but only 1.5 percent when $\phi_\pi = 3$. The longest event falls from 21 to 15 quarters. Also, raising the weight on the inflation gap raises the volatility of output by 38 percent and reduces inflation variability by 62 percent. In Model 2, increasing the weight on inflation from 1.5 to 3.0 raises output volatility by about 20 percent and reduces the standard deviation of inflation by 72 percent. Overall, the average ZLB event is longer in Model 1 than in Model 2.

Finally, we show how welfare changes as the dual mandate is de-emphasized. Following Schmitt-Grohé and Uribe (2007), the welfare cost associated with any policy is the fraction of consumption goods a household must give up under policy 1 to be indifferent between policies 1 and 2. Specifically, we solve for a value of λ that satisfies

$$W(c_t^2, n_t^2) = W((1-\lambda)c_t^1, n_t^1),$$

ϕ_y	$\bar{\lambda}_1(\%)$	$\bar{\lambda}_2(\%)$	ϕ_π	$\bar{\lambda}_1(\%)$	$\bar{\lambda}_2(\%)$
0.125	0.0000	0.0000	1.50	0.0000	0.0000
0.100	−0.0116	−0.0537	1.75	−0.0400	−0.0668
0.075	−0.0183	−0.0727	2.00	−0.0545	−0.0809
0.050	−0.0242	−0.0867	2.25	−0.0622	−0.0885
0.025	−0.0295	−0.0942	2.50	−0.0664	−0.0931
0.000	−0.0344	−0.0895	3.00	−0.0701	−0.0979

Table 4: Welfare implications of a dual mandate. The subscript represents the model number. The welfare measure represents the percentage of consumption goods that must be forgone in the baseline parameterization (given in the first row) to equate utility with the alternative parameterization. Thus a negative (positive) number is a welfare gain (loss). Calculations are based on an average of 1,000 simulations, each 2,500 quarters long.

where $W(c_t^\ell, n_t^\ell) \equiv E_0 \sum_{t=0}^{T-1} \beta^t u(c_t^\ell, n_t^\ell)$ is lifetime utility under policy ℓ. That value of λ is

$$\lambda = 1 - \exp\left\{ \frac{1-\beta}{1-\beta^T} \left(W(c_t^2, n_t^2) - W(c_t^1, n_t^1) \right) \right\}.$$

In table 4, we show the welfare gains from de-emphasizing the dual mandate. Columns 1-3 display the results when we hold $\phi_\pi = 1.5$ and decrease the weight on ϕ_y to 0 by 0.025 increments. The measure, $\bar{\lambda}_j(\%)$, represents the percentage of consumption goods that must be forgone in the baseline parameterization (first row) to equate utility to the alternative parameterization. Thus, a negative (positive) number is a welfare gain (loss) under parameterization j. Calculations are based on an average of $1{,}000$ Monte Carlo simulations, each $2{,}500$ quarters long.

The first column lists the value of ϕ_y and the next two columns show the results for Model 1 and Model 2. For Model 1, the highest welfare is achieved with no weight on the output gap. For Model 2, with the inflation parameter fixed at 1.5, the highest utility is achieved when we reduced the weight by a factor of 5, from 0.125 to 0.025. Columns 4-6 also de-emphasize the dual mandate by increasing the weight on inflation by increments of 0.25 while fixing the weight on output at 0.125. In every case, a larger weight on inflation leads to higher utility. These results are not surprising because the key distortion in this model is the nominal price rigidity. Any policy that reduces this distortion raises welfare.

8 CONCLUSION

This paper calculates global nonlinear solutions to standard New Keynesian models with and without capital to study how the ZLB affects economic dynamics. It has become well-known in the literature that the various short cuts for dealing with the discontinuity in the monetary policy rule do not accurately capture expectational effects and can lead to inaccurate conclusions about determinacy and dynamics. More specifically, we provide the solution for all combinations of technology and discount factor shocks and a thorough explanation of how dynamics change across the state space. This work makes a compelling argument that it is important to accurately capture expectations and provides additional evidence about how these shocks affect dynamics.

Our analysis focuses on technology shocks since they are an important source of aggregate fluctuations in many dynamic models. The key result in our models is that a positive technology shock generates lower inflation, consumption, and labor when the ZLB binds. While expectational effects are present in the model without capital, capital adds another consumption smoothing mechanism, which strengthens the expectational effects of the ZLB and impacts dynamics even before the ZLB binds. Capital adjustment costs reduce the degree of consumption smoothing, which increases the contraction at the ZLB and further strengthens expectational effects.

Given these unconventional dynamics, we also revisit policy issues. We show that policymakers can reduce the frequency and duration of ZLB events by de-emphasizing the dual mandate and confirm that doing so increases welfare. Despite the large volume of work 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 nominal interest rate? What are the benefits of forward guidance and quantitative easing in a dynamic model that accounts for expectational effects, and how do these policies change the effects of technology shocks? Answering these questions and others will require careful treatment of expectations and is the subject of ongoing research.

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.

——— (2007): "Discretionary Monetary Policy and the Zero Lower Bound on Nominal Interest Rates," *Journal of Monetary Economics*, 54, 728–752.

ARUOBA, S. AND F. SCHORFHEIDE (2013): "Macroeconomic Dynamics Near the ZLB: A Tale of Two Equilibria," Manuscript, University of Pennsylvania.

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

BLANCHARD, O., G. DELL'ARICCIA, AND P. MAURO (2010): "Rethinking Macroeconomic Policy," *Journal of Money, Credit and Banking*, 42, 199–215.

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.

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

BULLARD, J. (2010): "Seven Faces of 'The Peril'," *Federal Reserve Bank of St. Louis Review*, 92(5), 339–352.

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

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

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

COIBION, O., Y. GORODNICHENKO, AND J. WIELAND (2012): "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?" *Review of Economic Studies*, 79, 1371–1406.

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.

DEL NEGRO, M., M. GIANNONI, AND C. PATTERSON (2012): "The Forward Guidance Puzzle," FRB New York Staff Report 574.

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. (2010): "The Paradox of Toil," Federal Reserve Bank of New York Staff Report 433.

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.

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

ERCEG, C. J. AND J. LINDE (2010): "Is There a Fiscal Free Lunch in a Liquidity Trap?" International Finance Discussion Paper 1003.

FERNÁNDEZ-VILLAVERDE, J., G. GORDON, P. A. GUERRÓN-QUINTANA, AND J. 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.

GUST, C., D. LÓPEZ-SALIDO, AND M. E. SMITH (2012): "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.

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.

KING, R. G. AND A. L. WOLMAN (1996): "Inflation Targeting in a St. Louis Model of the 21st Century," *Federal Reserve Bank of St. Louis Review*, 78(6), 83–107.

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.

LUCAS, R. J. (1976): "Econometric Policy Evaluation: A Critique," *Carnegie-Rochester Conference Series on Public Policy*, 1, 19–46.

MERTENS, K. AND M. O. RAVN (2013): "Fiscal Policy in an Expectations Driven Liquidity Trap," Manuscript, Cornell university.

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.

PETERMAN, W. B. (2012): "Reconciling Micro and Macro Estimates of the Frisch Labor Supply Elasticity," Manuscript, Federal Reserve Board of Governors.

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

REIFSCHNEIDER, D. AND J. C. WILLIAMS (2000): "Three Lessons for Monetary Policy in a Low-Inflation Era," *Journal of Money, Credit, and Banking*, 32, 936–966.

RICHTER, A. W. AND N. A. THROCKMORTON (2013): "The Zero Lower Bound: Frequency, Duration, and Determinacy," Manuscript, Auburn University.

RICHTER, A. W., N. A. THROCKMORTON, AND T. B. WALKER (2013): "Accuracy, Speed and Robustness of Policy Function Iteration," *Forthcoming, Computational Economics*.

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

SCHMITT-GROHÉ, S. AND M. URIBE (2007): "Optimal Simple and Implementable Monetary and Fiscal Rules," *Journal of Monetary Economics*, 54, 1702–1725.

——— (2012): "The Making Of A Great Contraction With A Liquidity Trap and A Jobless Recovery," NBER Working Paper 18544.

SUMMERS, L. (1991): "How Should Long-Term Monetary Policy Be Determined? Panel Discussion," *Journal of Money, Credit and Banking*, 23, 625–31.

TAYLOR, J. B. (1993): "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39(1), 195–214.

WERNING, I. (2012): "Managing a Liquidity Trap: Monetary and Fiscal Policy," NBER Working Paper 17344.

WIELAND, J. (2012): "Are Negative Supply Shocks Expansionary at the Zero Lower Bound? Inflation Expectations and Financial Frictions in Sticky-Price Models," Manuscript, UC Berkeley.

WILLIAMS, J. C. (2009): "Heeding Daedalus: Optimal Inflation and the Zero Lower Bound," *Brookings Papers on Economic Activity*, 40, 1–49.

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