

1 Introduction

Are positive supply shocks contractionary in periods where monetary policy is constrained by the zero lower bound (ZLB)? The textbook New Keynesian (NK) model suggests that this is a possibility. The potential for this paradoxical result is driven by general equilibrium effects. Suppose that the supply shock is an increase in neutral productivity. Higher productivity lowers the natural rate of interest (i.e. the real interest rate consistent with a hypothetical equilibrium where prices are flexible). An inflation-targeting central bank can optimally respond by lowering nominal interest rates, which allows output to expand. If the central bank is constrained by the zero lower bound, however, nominal interest rates cannot fall. This results in a decrease in current and expected future inflation, which drives the equilibrium real interest rate up. The increase in the real interest rate chokes off demand, resulting in a smaller output increase than if policy were active. If the expected duration of the ZLB is long enough, the rise in the real interest rate can be sufficiently large that output declines.

In this paper we empirically test this prediction of the NK model using aggregate US data. Until very recently this has been a virtual impossibility, given the paucity of aggregate time series observations where the ZLB was binding. But there are now seven years of data (from the end of 2008 through the end of 2015) in which the effective Federal Funds Rate was at zero. We use [Fernald \(2014\)](#)'s utilization-adjusted total factor productivity (TFP) series, which we take to be a good measure of exogenous productivity. We then estimate impulse responses of output to changes in productivity, both inside and outside of the ZLB, using [Jordà \(2005\)](#)'s local projections method. The local projections method is a simple and robust way to estimate impulse responses, and can easily accommodate the kind of non-linearity induced by a binding ZLB.

In contrast to the predictions of the textbook NK model, we find that output responds more on impact to an increase in productivity at the ZLB than away from it. The difference in impact responses at and away from the ZLB is both economically and statistically significant. After a couple of quarters there is no difference at or away from the ZLB in the estimated response of output to a productivity shock. Our result is robust to a number of different variations on our baseline specification.

Since a decline in inflation is the mechanism by which the textbook NK model generates a smaller output response to a productivity shock at the ZLB, we also empirically examine the effects of productivity shocks on inflation, both inside and outside of the ZLB. Over the period 1984-2007, we find that there is no statistically significant effect of a productivity shock on inflation, which is consistent with an inflation targeting monetary policy regime. At the ZLB, in contrast, we find that productivity shocks have a large, negative effect on

inflation. This result is robust along a number of dimensions.

Taken together, our results for the effects of productivity shocks on output and inflation at the ZLB present something of a puzzle. A stronger deflationary effect of productivity shocks at the ZLB is consistent with the predictions of a textbook NK model. Since a deflationary response is the mechanism through which output responds less to a productivity shock at the ZLB, our finding that output responds more to a productivity shock at the ZLB is puzzling. Our results suggest a failing of the textbook NK model along some dimension.

While we focus on the effects of supply shocks at the ZLB, our work has implications for demand-side policies as well. [Christiano, Eichenbaum, and Rebelo \(2011\)](#) and others have argued that the government spending multiplier is significantly larger at the ZLB in comparison to normal times. A number of authors, most notably [Del-Negro and Giannoni \(2015\)](#), have noted that extended periods of anticipated monetary accommodation can be wildly expansionary. The mechanism by which demand shocks can have large effects at the ZLB is, in a sense, the mirror image of why supply shocks might have small effects. Demand shocks raise inflation, which pushes down real interest rates when nominal rates are constrained by the ZLB. Our empirical findings suggest that this “inflation channel” ([Dupor and Li \(2015\)](#)) does not seem to be operative at the ZLB in the way predicted by the theory, at least conditional on supply shocks. Our results therefore imply that caution ought to be in order when applying the basic intuition from the NK model to draw inferences about the likely effects of demand shocks.

Our work contributes to a burgeoning literature investigating the macroeconomic effects of supply shocks at the zero lower bound. [Eggertson \(2012\)](#) and [Eggertsson and Krugman \(2012\)](#) both argue that New Deal policies, which reduced the natural rate of output, were in fact expansionary due to the zero lower bound. On the other hand, [Wieland \(2015\)](#) uses the Great Japan Earthquake and global oil supply disruptions as exogenous supply shocks and finds that negative shocks are contractionary at the zero lower bound.¹ In a similar vein, [Cohen-Setton, Hausman, and Wieland \(2016\)](#) show that cartelization efforts exacerbated France’s Great Depression. These papers focus on shocks to aggregate supply which are different than neutral productivity shocks. We are unaware of any paper which studies the consequences of exogenous productivity shocks at the zero lower bound.

Our work also fits more broadly into a growing literature which empirically tests other predictions of the textbook NK model when the ZLB binds. [Bachmann, Berg, and Sims \(2015\)](#) find no evidence that consumer willingness to spend on durable goods is affected by

¹Also in the spirit of empirically testing predictions of New Keynesian models, [Mulligan \(2011\)](#) suggests that the behavior of some labor market variables over the 2008-2009 period were more consistent with models of flexible prices.

inflation expectations, either at or away from the ZLB. [Burke and Ozdagli \(2013\)](#) reach similar conclusions. In contrast, [D’Acunto, Hoang, and Weber \(2016\)](#) argue that a VAT increase in Germany which raised household inflation expectations was quite expansionary. Similar, [Ichiue and Nishiguchi \(2015\)](#) find that higher inflation expectations positively correlate with consumption spending for households in Japan. [Dupor and Li \(2015\)](#) find no evidence to support an important “expected inflation channel” for large fiscal multipliers at the ZLB. [Ramey and Zubairy \(2014\)](#) estimate state-dependent regression models similar to ours to study the magnitude of the fiscal multiplier, both across states of the business cycle as well as in periods where the ZLB binds. They find no evidence of a significantly larger multiplier during periods in which the ZLB binds.

2 Theory

Consider the textbook NK model. The two principal equations of the model are the linearized IS equation and a Phillips Curve:

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma} (i_t - \mathbb{E}_t \pi_{t+1} - r_t^f) \quad (1)$$

$$\pi_t = \gamma x_t + \beta \mathbb{E}_t \pi_{t+1}. \quad (2)$$

Here x_t is the output gap, defined as the log deviation of output, y_t , from its flexible price level, $x_t = y_t - y_t^f$. The nominal interest rate, expressed in absolute deviations from steady state, is i_t . The hypothetical real interest rate if prices were fully flexible is r_t^f . σ is the inverse elasticity of intertemporal substitution. The slope coefficient in the Phillips Curve is $\gamma = \frac{(1-\phi)(1-\phi\beta)}{\phi}(\sigma + \chi)$, where $\phi \in [0, 1)$ is the probability firms cannot adjust their price in a given period, $0 < \beta < 1$ is a subjective discount factor, and χ is the inverse Frisch elasticity of labor supply. The exogenous driving force in the model is log productivity, a_t , which obeys a stationary AR(1) process with $0 \leq \rho_a < 1$:

$$a_t = \rho_a a_{t-1} + \varepsilon_{a,t}, \quad \varepsilon_{a,t} \sim N(0, s^2). \quad (3)$$

In terms of exogenous productivity, the flexible price real interest rate and output can be solved for analytically as:

$$y_t^f = \frac{1 + \chi}{\sigma + \chi} a_t \quad (4)$$

$$r_t^f = \frac{\sigma(1 + \chi)(\rho_a - 1)}{\sigma + \chi} a_t. \quad (5)$$

To complete the model, it remains to specify a monetary policy rule. During normal times, we assume that the central bank follows a strict inflation target, adjusting the nominal interest rate so as to implement $\pi_t = 0$.² In order to approximate the effects of a binding zero lower bound on nominal interest rates in a tractable way, we consider a case in which the nominal interest rate is pegged at a fixed value for a deterministic number of periods, H . In particular, suppose that i_t is fixed at $1 - 1/\beta$ for H periods. Since i_t is the deviation of the nominal interest rate from steady state, and $1/\beta - 1$ is the steady state nominal interest rate, this means that the nominal interest rate is fixed at 0. After this period of time, agents in the economy expect the central bank to return to an inflation targeting regime, which requires that the nominal interest rate equal the natural rate of interest. Formally, monetary policy under this kind of peg is characterized by:³

$$\mathbb{E}_t i_{t+h} = \begin{cases} 1 - \frac{1}{\beta} & \text{if } h < H \\ \mathbb{E}_t r_{t+h}^f & \text{if } h \geq H. \end{cases} \quad (6)$$

We solve for the expected time path of output backwards. Starting in period $t + H$, agents will expect $\mathbb{E}_t i_{t+H} = \mathbb{E}_t r_{t+H}^f$, which implies that $\mathbb{E}_t \pi_{t+H} = \mathbb{E}_t x_{t+H} = 0$. This means that $\mathbb{E}_t x_{t+H-1} = -\frac{1}{\sigma} \left(1 - \frac{1}{\beta}\right) + \frac{1}{\sigma} \mathbb{E}_t r_{t+H-1}^f$ and $\mathbb{E}_t \pi_{t+H-1} = \gamma \mathbb{E}_t x_{t+H-1}$. This can then be iterated back to period t , yielding the expected time paths of the output gap and inflation. The expected path of output can then be recovered from the path of the output gap, given an exogenous path for productivity.

We parameterize the model as follows. We set $\phi = 0.75$, which implies a four quarter average duration between price changes. The elasticity of intertemporal substitution is $\sigma = 1$, and we assume preferences are linear over labor, so $\chi = 0$. The discount factor is $\beta = 0.99$. We assume that the persistence of the productivity shock is $\rho_a = 0.9$.

Consider a one unit positive shock to productivity. The impulse response of output under various different durations of the interest rate peg are shown in Figure 1. The solid blue line shows the response of output when $H = 0$, so that the central bank targets an inflation rate of zero in all periods. Given our parameterization of $\sigma = 1$, the impulse response of output is just equal to the impulse response of a_t . We consider three additional peg lengths of $H = 3$, $H = 6$, and $H = 10$. Given the absence of endogenous state variables in the model, after horizon H the response of output is identical to the inflation targeting case. One observes

²The results that follow are similar if the monetary authority instead follows a Taylor rule. See Appendix B for details.

³It is well-known that an exogenous interest rate peg results in equilibrium indeterminacy. We do not have such a problem in our setup, because policy after the ZLB is formulated in terms of an inflation target, $\mathbb{E}_t \pi_{t+h} = 0$ for $h \geq H$. With this inflation target, $\mathbb{E}_t i_{t+h} = \mathbb{E}_t r_{t+h}^f$ in equilibrium, not as a policy rule that might hold out of equilibrium.

that the output response on impact is smaller than the inflation targeting case for $H > 0$. Furthermore, the impact response of output is smaller the bigger is H . For H sufficiently large, output can actually decline on impact, as it does here in the case of $H = 10$. Given our parameterization of the model, the impact response of output is negative for $H \geq 8$.

The mechanism for the smaller output response for a longer duration of the interest rate peg lies in the response of inflation. Dupor and Li (2015) have termed this the “expected inflation channel.” In particular, a positive productivity shock lowers inflation when monetary policy is passive. The longer the nominal interest rate is pegged, the more inflation falls. A decline in expected inflation, coupled with a fixed nominal interest rate, results in an increase in the real interest rate. The higher real interest rate chokes off demand and results in a smaller increase in output. These effects can be seen in Figure 2, which is similar to Figure 1 but plots the response of inflation to a productivity shock as a function of the duration of the interest rate peg, H .




Figure 1: Response of Output to Productivity Shock as a Function of Duration of ZLB

Notes: This figure plots the impulse responses of output to a one percent increase in productivity for different durations of a pegged nominal interest rate at zero. $H = 0$ corresponds to the case where the central bank obeys a strict inflation target at all times.

In Appendix C, we show that these qualitative results also hold in a medium scale model with capital accumulation and several other frictions similar to Smets and Wouters (2007). A temporary rise in productivity leads to an increase in output away from the ZLB, but a decrease in output at the ZLB. The mechanism generating the decrease in output is more or less the same as in the basic NK model. Inflation decreases by more at the ZLB compared to

a Taylor rule, which puts upward pressure on the real interest rate and therefore works to limit demand. The dynamics of the natural rate of interest are somewhat different in the model with capital compared to the simple process shown in (5), as we discuss further in the Appendix. If the productivity shock is permanent, then output may respond more to it at the ZLB compared to normal times. But if that is the case, inflation falls by less, not more, to the productivity shock.




Figure 2: Response of Inflation to Productivity Shock as a Function of Duration of ZLB

Notes: This figure plots the impulse responses of inflation to a one percent increase in productivity for different durations of a pegged nominal interest rate at zero. $H = 0$ corresponds to the case where the central bank obeys a strict inflation target at all times.

An alternative approach to modeling the effects of the zero lower bound is to assume that the duration of a pegged nominal interest rate is stochastic, rather than deterministic as we have assumed. A stochastic duration of an interest rate peg is the approach taken, for example, in [Christiano, Eichenbaum, and Rebelo \(2011\)](#). In particular, one can assume that in each period there is a fixed probability, p , with $p \in [0, 1)$, that the nominal interest rate will remain at zero. The expected duration of the peg is then $1/(1 - p)$ periods. [Carlstrom, Fuerst, and Paustian \(2014\)](#) argue that a deterministic peg length provides much more reasonable results in a textbook NK model with government spending than does a stochastic peg. [Carlstrom, Fuerst, and Paustian \(2015\)](#) examine the effects of forward guidance in a textbook New Keynesian model. When the interest rate is pegged for a stochastic period of time, they find that there are sign reversals in the effects of forward guidance on current

inflation and output.⁴ In Appendix A, we show the impulse responses of output and inflation to a productivity shock for different expected durations of an interest rate peg. For moderate expected durations of the peg, our results are the same as in the main text – output responds less to a positive productivity shock the longer is the expected duration of the peg, and inflation falls more. However, like Carlstrom et al. (2015), if the expected duration of the peg is sufficiently long we find sign reversals, wherein output responds more to a productivity shock at the ZLB than under an inflation target and the inflation response is positive, rather than negative.

3 Empirical Analysis

In this section we empirically test the prediction that a positive productivity shock has a smaller, and potentially negative, effect on output when the zero lower bound binds in comparison to normal times. We measure productivity using Fernald (2014)’s quarterly series on utilization-adjusted total factor productivity. Subsection 3.1 describes this data series and makes the case that it can plausibly be considered an exogenous productivity series. Subsection 3.2 estimates the effects of productivity shocks on output, both inside and outside the ZLB. Subsection 3.3 does the same but for inflation. Subsection 3.4 considers several robustness exercises.

3.1 Data

For our empirical analysis it is critical that we observe a variable which accurately measures exogenous productivity. A traditional Solow residual is likely to be a poor measure of exogenous productivity because of factor hoarding. We therefore use the utilization-adjusted TFP series produced and provided by Fernald (2014). Formally, he assumes an aggregate production function of the form:

$$Y_t = A_t(z_t K_t)^{\alpha_t} (e_t L_t)^{1-\alpha_t} \quad (7)$$

where Y_t is output, K_t is physical capital, and L_t is aggregate labor hours. A_t is an exogenous productivity shifter. z_t denotes capital utilization and e_t labor effort. α_t is a potentially

⁴In Carlstrom et al. (2015), they begin with a standard NK model and show that the response of output and inflation to a deterministic period of forward guidance is exponentially increasing in the duration of low interest rates. However, when inflation indexation is introduced into the model, sign reversals begin to occur wherein output and inflation fall, rather than rise, with the anticipation of an extended period of low interest rates. When they consider a stochastic interest rate peg, they find sign reversals at modest durations of a peg even without backward indexation in the model.

time-varying capital’s share parameter. A traditional measure of TFP is log output less share-weighted capital and labor. In terms of (7) this can be written:

$$\begin{aligned}\ln TFP_t &= \ln Y_t - \alpha_t \ln K_t - (1 - \alpha_t) \ln L_t \\ &= \ln A_t + \ln u_t.\end{aligned}\tag{8}$$

Here $\ln u_t = \alpha_t \ln z_t + (1 - \alpha_t) \ln e_t$ is a composite utilization factor. Only if factor utilization is constant will a traditional TFP series correspond to the exogenous productivity concept in (7). Fernald (2014) uses the insights from Basu, Fernald, and Kimball (2006) and follow-up work from Basu, Fernald, Fisher, and Kimball (2013) to create an aggregate utilization series, which is used to “correct” a traditional TFP series. In other words:

$$\ln A_t = \ln TFP_t - \ln u_t.\tag{9}$$

We will denote the utilization-adjusted TFP series by A_t , the same symbol used to denote exogenous productivity in (7). The interested reader is referred to Fernald (2014) for more details on the construction of the utilization-adjusted TFP series.⁵

Table 1 presents some summary statistics on the log first difference of the utilization-adjusted TFP series. For point of comparison, we also show statistics on the log first difference of a traditional measure of TFP. In addition, we show moments for output growth and inflation, as these are the key aggregate variables used in our regression analysis below. Output is measured as real GDP from the NIPA tables, while inflation is the log first difference of the corresponding price deflator. We focus on the sample period 1984Q1–2015Q4, which is the sample period we use in our baseline regressions.

In terms of volatilities, the utilization-adjusted TFP series is actually more volatile than the conventional TFP series, both of which are slightly more volatile than output growth. Inflation is quite stable over the sample period in consideration. Output growth and inflation are both highly autocorrelated. The growth rate of utilization-adjusted TFP is not autocorrelated. The traditional TFP series is slightly positively autocorrelated. The lower part of the table shows the correlation matrix for these variables. Utilization-adjusted TFP is mildly positively correlated with output growth (0.2), though this correlation is small in absolute terms. In contrast, the conventional TFP series is quite strongly positively correlated

⁵Sims (2016) documents large differences in the time series properties of Fernald’s utilization-adjusted TFP series by vintage. In particular, there is a discrete change in the time series properties of the series which occurred beginning with vintages produced starting in the spring of 2014, based on an update to using the Basu et al. (2013) methodology for measuring factor utilization. Sims (2016) argues that the most recent vintages of Fernald’s data are better than earlier vintages. The vintage of Fernald’s data we use was downloaded in March of 2016, fully reflecting the updates to his series from the spring of 2014.

with output growth (0.7).⁶ The growth rate of the utilization-adjusted TFP series is mildly negatively correlated with inflation (-0.1), more strongly so than the correlation between the conventional TFP series and inflation (-0.05).

Table 1: Empirical Moments

		$\Delta \ln A_t$	$\Delta \ln Y_t$	π_t	$\Delta \ln TFP_t$
Standard Deviation		0.0069	0.0061	0.0006	0.0062
AR(1)		-0.008	0.441	0.615	0.111
Correlation Matrix	$\Delta \ln A_t$	1	0.228	-0.115	0.656
	$\Delta \ln Y_t$		1	0.157	0.673
	π_t			1	-0.050
	$\Delta \ln TFP_t$				1

Notes: These moments are calculated for the period 1984Q1–2015Q4. Δ denotes the first difference of the relevant variable. TFP_t corresponds to log output less share-weighted capital and labor as described in Equation (8). A_t refers to the measure provided by Fernald (2014) described in (9). Output, Y_t , is real GDP from the NIPA tables. Inflation, π_t is the log first difference of the corresponding price deflator.

The fact that the utilization-adjusted TFP series is more weakly correlated with output than a traditional TFP series suggests that the utilization-adjustment represents an improvement over a conventional growth accounting exercise. It does not, however, prove that Fernald’s series can be considered exogenous with respect to macroeconomic conditions. To go a step further, we conduct a sequence of pairwise Granger causality tests using the first log difference of Fernald’s utilization-adjusted TFP series and other macroeconomic shocks and variables. Under the null hypothesis that the utilization-adjusted TFP series is a measure of exogenous productivity, it should not be predictable from other variables. We take three popular measures of macroeconomic shocks identified in the literature – Romer and Romer (2004) monetary policy shocks, Romer and Romer (2010) tax shocks, and the defense news shock produced by Ramey (2011). We also consider the first difference of the log S&P 500 stock market index and the first log difference of a measure of oil prices. F statistics and p values from the pairwise Granger causality tests are presented in Table 2.

The results in Table 2 fail to reject the null hypothesis that any of the series in question do not Granger cause the log first difference of utilization-adjusted TFP. The p values for the exogenous policy shocks are particularly high. The p values for the S&P 500 index and the crude oil price are lower, but are still substantially higher than conventional significance levels. These results are suggestive, but of course not dispositive, that Fernald’s series can be treated as exogenous.

⁶Over a longer sample period, the utilization-adjusted TFP series is even more weakly correlated with output growth than presented here. In particular, over the period 1947Q2-2015Q4, the correlation between output growth and utilization-adjusted TFP is only 0.12.

Table 2: H_0 : Alternative Measure does not Granger Cause Utilization-Adjusted TFP

Measure/Variable	F statistic	p value
Romer and Romer (2004)	0.5579	0.5761
Romer and Romer (2010)	0.4075	0.6665
Ramey (2011)	0.9763	0.3798
Δ S&P 500 Index	1.3549	0.2618
Δ OIL	1.8525	0.1612

Notes: All the tests are performed using two lags. Romer and Romer (2004) the sock is the quarterly average of their monthly shock measure (mnemonic RESID) for the period 1984Q1–1996Q4. For Romer and Romer (2010)’s measure we use “Exogenous Tax Changes” based on the present value of tax changes relative to nominal GDP (mnemonic EXOGE PDVRATIO), for the period 1984Q1–2007Q4. Ramey (2011)’s series is “Defense news as a % of lagged nominal GDP” for the period 1984Q1–2013Q4. Δ S&P 500 Index is the quarterly average of monthly growth rate of close price adjusted for dividends and splits for the period 1984Q1–2015Q4 divided by the consumer price index. OIL is the log first difference of the BLS’s producer price index for domestic crude oil (series ID WPU0561) for the period 1984Q1–2015Q4.

3.2 Productivity Shocks and Output: At and Away from the ZLB

Our objective is to estimate the impulse response of output to a productivity shock, both inside and outside of periods where the ZLB binds. To that end, we estimate a state-dependent regression model using Jordà (2005)’s local projection method. This is more robust to misspecification than a traditional VAR and it is straightforward to adapt to a non-linear setting. Auerbach and Gorodnichenko (2013) and Ramey and Zubairy (2014) are examples of two recent papers which have used the local projections method to estimate state-dependent fiscal multipliers.

We estimate a sequence of regressions of the form:

$$\ln Y_{t+h} = \alpha^h + \beta^h \ln A_t + \sum_{s=1}^p \gamma_s^h \ln Y_{t-s} + \sum_{s=1}^p \phi_s^h \ln A_{t-s} + Z_t \left[\alpha_z^h + \beta_z^h \ln A_t + \sum_{s=1}^p \gamma_{z,s}^h \ln Y_{t-s} + \sum_{s=1}^p \phi_{z,s}^h \ln A_{t-s} \right] + u_{t+h}. \quad (10)$$

This regression is estimated separately for different forecast horizons, $h \geq 0$. In words, the log level of output at a lead of h horizons is regressed on a constant, the period t level of productivity, and p lags of the log levels of output and productivity. That the period t value of productivity appears on the right hand side reflects an assumption that the utilization-adjusted TFP series is exogenous with respect to output within a period. Lags of output and utilization-adjusted TFP are included in the regression to partial out any

predictable movements in adjusted TFP. Z_t is a dummy variable taking on a value of 1 in periods where the zero lower bound binds. To allow for different effects at the zero lower bound, we include an interaction term which allows all coefficients to differ during periods in which the ZLB binds. Outside of the zero lower bound, the estimate of the impulse response of output at horizon h to a change in productivity is given by β^h . At the zero lower bound the response is given by $\beta^h + \beta_z^h$.

As a baseline, we estimate these regressions over the sample period 1984Q1 through 2015Q4. The beginning date is chosen to coincide with conventional dating of the “Great Moderation.” The end of the sample leaves 29 observations where the zero lower bound binds (2008Q4 through 2015Q4), or about 23 percent of the sample period. In accordance with the Akaike Information Criterion, we estimate the regression with $p = 3$ lags. Inference is conducted using [Newey and West \(1987\)](#) HAC standard errors.




Figure 3: Results from Baseline Regression

Notes: This figure shows the effects of A_t on output at various horizons. The dotted line is when the ZLB binds or $Z_t = 1$. The blue line with ‘x’ markers is when $Z_t = 0$. The shaded area bands represent the 95 percent confidence interval about the no ZLB case using [Newey and West \(1987\)](#) HAC standard errors. The red line with ‘*’ markers shows the response when the ZLB binds.

We graphically display the results from the baseline model in Figure 3. The blue line with ‘x’ markers shows the estimated impulse response of output to a one unit productivity shock outside of the ZLB; the shaded gray region corresponds to the 95 percent confidence interval associated with this response. Output increases on impact and continues to rise for several quarters thereafter. The red line with ‘*’ markers plots the estimated response when the zero lower bound binds. Contrary to the theory outlined in Section 2, output is

estimated to respond by nearly four times as much to a productivity shock on impact at the ZLB than outside of it. Hence, not only does the economy expand after a productivity increase at the ZLB, the response is actually amplified in periods where nominal interest rates are zero. After a couple of periods the response at the ZLB converges to the response outside of the ZLB.

Table 3: Standard Errors and p -values on β_z^h

Horizon	β_z^h	S.E.	p -value
$h = 0$	0.540	0.160	0.001
$h = 1$	0.445	0.294	0.067
$h = 2$	0.047	0.379	0.451
$h = 3$	-0.009	0.363	0.509
$h = 4$	0.147	0.404	0.359
$h = 5$	0.039	0.389	0.461
$h = 6$	0.069	0.399	0.431
$h = 7$	-0.226	0.471	0.684
$h = 8$	-0.184	0.494	0.645

Notes: This table shows the estimates of β_z^h at different forecast leads, as well as the corresponding standard errors and p values. Standard errors are computed using [Newey and West \(1987\)](#).

In [Table 3](#) we show the estimates, standard errors, and p values for the coefficient β_z^h by forecast horizon. The p values can be interpreted as a test of the hypothesis that $\beta_z^h = 0$ (equivalently, that there is no statistically significant difference between the response of output in normal times compared to the ZLB). On impact, the p value is close to zero, so the hypothesis can be rejected. At a lead of one period, the hypothesis that the responses are the same can be rejected at the 10 percent level, but not at a 5 percent significance level. The estimates of β_z^h are statistically insignificant from zero at leads of $h = 2$ and higher.

3.3 Inspecting the Mechanism

As discussed in [Section 2](#), in the NK model output responds less to a productivity shock at the ZLB because of an effect of inflation on the real interest rate. In particular, inflation falls after a productivity shock when the central bank is unable to adjust nominal rates, which drives the real rate up and crowds out demand, resulting in a smaller increase in output.

In this subsection, we empirically investigate the response of inflation to a productivity shock, both inside and outside of the ZLB. We estimate a sequence of regressions similar to

(10), but with inflation in place of output:

$$\begin{aligned} \pi_{t+h} = & \alpha^h + \beta^h \ln A_t + \sum_{s=1}^p \gamma_s^h \pi_{t-s} + \sum_{s=1}^p \phi_s^h \ln A_{t-s} \\ & + Z_t \left[\alpha_z^h + \beta_z^h \ln A_t + \sum_{s=1}^p \gamma_{z,s}^h \pi_{t-s} + \sum_{s=1}^p \phi_{z,s}^h \ln A_{t-s} \right] + u_{t+h}. \end{aligned} \quad (11)$$

We measure inflation as the annualized percentage change in the GDP price deflator. We again estimate these regressions with $p = 3$ lags. The estimated impulse responses of inflation both inside and outside of the ZLB are shown in Figure 4. The blue line with ‘x’ markers denotes the response outside of the ZLB, while the red line with the ‘*’ markers is the response at the ZLB. The shaded gray region is the 95 percent confidence interval for the response estimated outside of the ZLB.

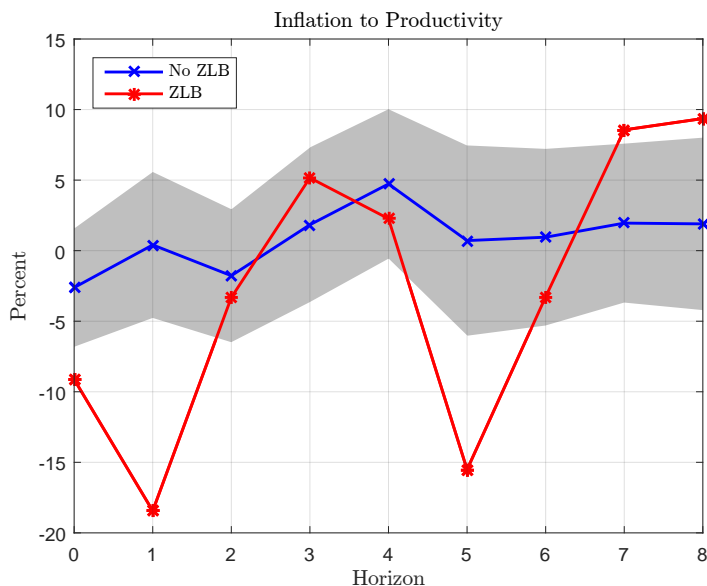


Figure 4: Results from Baseline Regression, Inflation

Notes: This figure shows the effects of A_t on inflation at various horizons. The dotted line is when the ZLB binds or $Z_t = 1$. The blue line with ‘x’ markers is when $Z_t = 0$. The shaded area bands represent the 95 percent confidence interval about the no ZLB case using [Newey and West \(1987\)](#) HAC standard errors. The red line with ‘*’ markers shows the response when the ZLB binds.

During normal times, the estimated response of inflation to a productivity shock is negative on impact but statistically indistinguishable from zero. The response is also insignificant at all subsequent forecast horizons. This response is consistent with a strict inflation target, as assumed in the theoretical model in Section 2. At the ZLB, in contrast, the inflation response is negative on impact, much more so than during periods where the ZLB does not bind. While the estimated response is quite choppy, at most forecast horizons the inflation

response at the ZLB lies below the estimated response in normal times.

Table 4 shows the coefficient estimates of β_z^h , along with standard errors and p values. A test of the hypothesis that $\beta_z^h = 0$ amounts to a test of the hypothesis that inflation response at horizon h differs at the ZLB in comparison to normal times. The response of inflation is significantly different at horizons $h = 1$ and $h = 5$.

In summary, our results present something of a puzzle in light of the predictions of the textbook NK model. The mechanism through which output ought to respond less to a productivity shock at the ZLB – namely, a decrease in current and expected inflation – is present in the data. Yet, in contrast to the theory, the output response to a productivity shock at the ZLB is amplified rather than dampened. Put together, these findings suggest some basic flaw in the textbook NK model.

Table 4: Standard Errors and p -values on β_z^h

Horizon	β_z^h	S.E.	p -value
$h = 0$	-6.507	8.408	0.220
$h = 1$	-18.848	7.761	0.008
$h = 2$	-1.543	11.572	0.447
$h = 3$	3.325	4.789	0.244
$h = 4$	-2.460	4.987	0.311
$h = 5$	-16.229	8.940	0.036
$h = 6$	-4.226	5.049	0.202
$h = 7$	6.606	5.673	0.123
$h = 8$	7.466	6.082	0.111

Notes: This table shows the estimates of β_z^h at different forecast leads, as well as the corresponding standard errors and p values. These are for regressions with inflation in place of output. Standard errors are computed using [Newey and West \(1987\)](#).

In Section 2 we discussed the possibility of a sign reversal under a stochastic peg whereby output would increase more at the ZLB than away from it. One rationalization for our results for the output effects of a productivity shock at the ZLB is that the duration of the peg is stochastic and individuals expect that it will persist for a sufficiently long duration. Sign reversals in the theoretical model apply both to the responses of output and inflation – if output responds more to a productivity shock because of a binding ZLB, then the inflation response ought to be positive. We find that the inflation response to a productivity shock is negative, both away from the ZLB and more strongly so at the ZLB. Hence, our empirical results for output and inflation jointly taken together cannot be rationalized by the type of sign reversals in the standard NK model emphasized by [Carlstrom et al. \(2015\)](#).

Other possible reconciliations of our results with the theory include adding capital to the model or considering permanent, rather than transitory, productivity shocks. In the textbook model, a non-stationary productivity shock would result in the natural rate of interest rising, rather than declining, after a positive shock to productivity. This can easily be seen in (5) with $\rho_A > 1$. If this were the case, the model would predict that both output and inflation would increase more after a productivity shock at the ZLB than under an inflation target. While this would be consistent with our empirical estimate of the output response to a productivity shock at the ZLB, it is inconsistent with our empirical results for inflation, where we find that inflation falls significantly after a productivity shock at the ZLB. The inclusion of capital in the model alters the dynamics of the natural rate of interest, which we discuss further in Appendix C. There, we show that a medium scale model with capital accumulation and several additional frictions produces responses to a stationary productivity shock at and away from the ZLB that are qualitatively similar to the textbook model without capital. If the productivity shock is permanent, then output may respond more at the ZLB than away from it, but this coincides with inflation falling less at the ZLB, rather than more as we find empirically. Hence, it is the joint dynamics of output and inflation to a productivity shock at the ZLB that are difficult to square with the basic NK theory, not the responses of output or inflation in isolation.

3.4 Robustness

We consider a number of robustness checks on our baseline result. Robustness checks for the output regressions are shown in different panels of Figure 5. First, we consider robustness to sample period. We consider the following two alternative sample periods – 1947Q2 - 2015Q4 and 1960Q1 - 2015Q4. The responses are shown in the upper left panel of the figure. Because β_z^h is identified only off of the most recent data, the estimate of the impulse response of output at the zero lower bound is not affected by longer sample periods. The estimated response outside of the ZLB is fairly similar across samples, though it is a bit smaller at all horizons when using the longest sample period. This means that the difference between the impact response at the ZLB compared to away from the ZLB is even larger with the longer sample compared to our baseline results.

One might worry that our result that output responds by more to a productivity shock at the ZLB compared to normal times is in actuality driven by the fact that the most recent ZLB period coincides with the height of the Great Recession. Put differently, it may be the case that output responds by more to a productivity shock in a time of recession, and that our ZLB dummy variable is simply proxying for periods of recession. We address this concern

by augmenting our baseline regression to include a recession dummy variable, interacted in an analogous way to the ZLB dummy interaction. Our regression specification is shown below:

$$\begin{aligned} \ln Y_{t+h} = & \alpha^h + \beta^h \ln A_t + \sum_{s=1}^p \gamma_s^h \ln Y_{t-s} + \sum_{s=1}^p \phi_s^h \ln A_{t-s} \\ & + Z_t \left[\alpha_z^h + \beta_z^h \ln A_t + \sum_{s=1}^p \gamma_{z,s}^h \ln Y_{t-s} + \sum_{s=1}^p \phi_{z,s}^h \ln A_{t-s} \right] \\ & + R_t \left[\alpha_r^h + \beta_r^h \ln A_t + \sum_{s=1}^p \gamma_{r,s}^h \ln Y_{t-s} + \sum_{s=1}^p \phi_{r,s}^h \ln A_{t-s} \right] + u_{t+h}. \end{aligned} \quad (12)$$

In this specification R_t is a dummy variable taking on a value of one in periods identified by the NBER as recessions. We do find that output responds more to a productivity shock in a time of recession compared to normal times. Nevertheless, we still find that output responds by more on impact to a productivity shock at the ZLB than away from it. This response is shown in the middle panel of Figure 5. Relative to our baseline, though larger on impact the response of output at the ZLB is estimated to be less persistent compared to normal times.

Our baseline regression specification is in the levels of the variables. This specification is robust to cointegration between output and productivity. We consider an alternative specification in which variables appear in growth rates. Let Δ denote the first difference operator. We estimate:

$$\begin{aligned} \ln Y_{t+h} - \ln Y_{t-1} = & \alpha^h + \beta^h \Delta \ln A_t + \sum_{s=1}^p \gamma_s^h \Delta \ln Y_{t-s} + \sum_{s=1}^p \phi_s^h \Delta \ln A_{t-s} \\ & + Z_t \left[\alpha_z^h + \beta_z^h \Delta \ln A_t + \sum_{s=1}^p \gamma_{z,s}^h \Delta \ln Y_{t-s} + \sum_{s=1}^p \phi_{z,s}^h \Delta \ln A_{t-s} \right] + u_{t+h}. \end{aligned} \quad (13)$$

In this specification the cumulative growth rate of output over h horizons is regressed on the current growth rate of the adjusted TFP series and lags of the growth rates of the adjusted TFP series and output. The estimated responses inside and outside of the ZLB are shown in the upper right panel of the first row of Figure 5. Qualitatively our results are the same with this alternative specification. On impact, output responds by more to the productivity shock at the ZLB than outside of it. The main difference relative to our baseline result is that the output response at the ZLB is less persistent.

We consider two additional robustness checks. The lower panel of the bottom row of the figure uses different numbers of lags. The estimated responses are qualitatively insensitive to the number of lags in the regression. The right panel in the lower row includes a deterministic time trend in our baseline regression. We consider both a linear time trend and a quadratic

trend. As one can see in the figure, the inclusion of a time trend has little noticeable effect on the results.

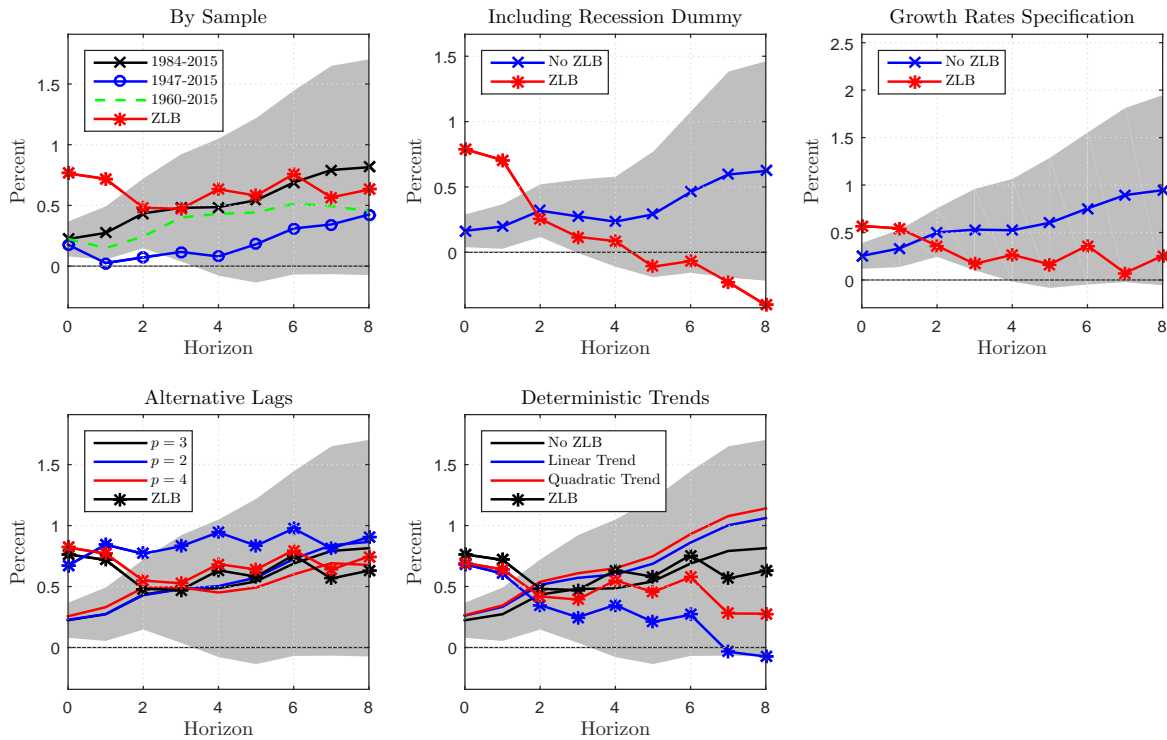


Figure 5: IRFs of Output to Productivity

Notes: This figure shows estimated impulse responses of output to a productivity shock, both inside and outside of the ZLB, for various different variations on our baseline regression model.

Next, we consider the same robustness checks for the response of inflation to a productivity shock. The estimated impulse responses for these robustness checks are depicted graphically in Figure 6. In the growth rates specification inflation enters in levels, but utilization-adjusted TFP enters in log first differences. The upper left panel shows responses estimated over different time horizons. The response of inflation at the ZLB is the same across samples, since the ZLB sample is fixed across specifications. Interestingly, one observes that inflation responds more negatively to a productivity shock outside of the ZLB regime in samples that extend farther back in time. In particular, the impact decline in inflation is largest over the sample that extends back to 1947. These findings are consistent with the analysis in [Gali, Lopez-Salido, and Valles \(2003\)](#), who estimate the effects of technology shocks identified from a VAR on inflation in the pre- and post-Volcker periods. In their analysis, there is no significant response of inflation to a positive productivity shock in the post-Volcker period, whereas in the earlier sample the decline in inflation after a productivity improvement is quantitatively large and statistically significant. The other panels of the figure confirm the

robustness of our main result for inflation. In all of these different specification, inflation falls more on impact at the ZLB than outside of it.

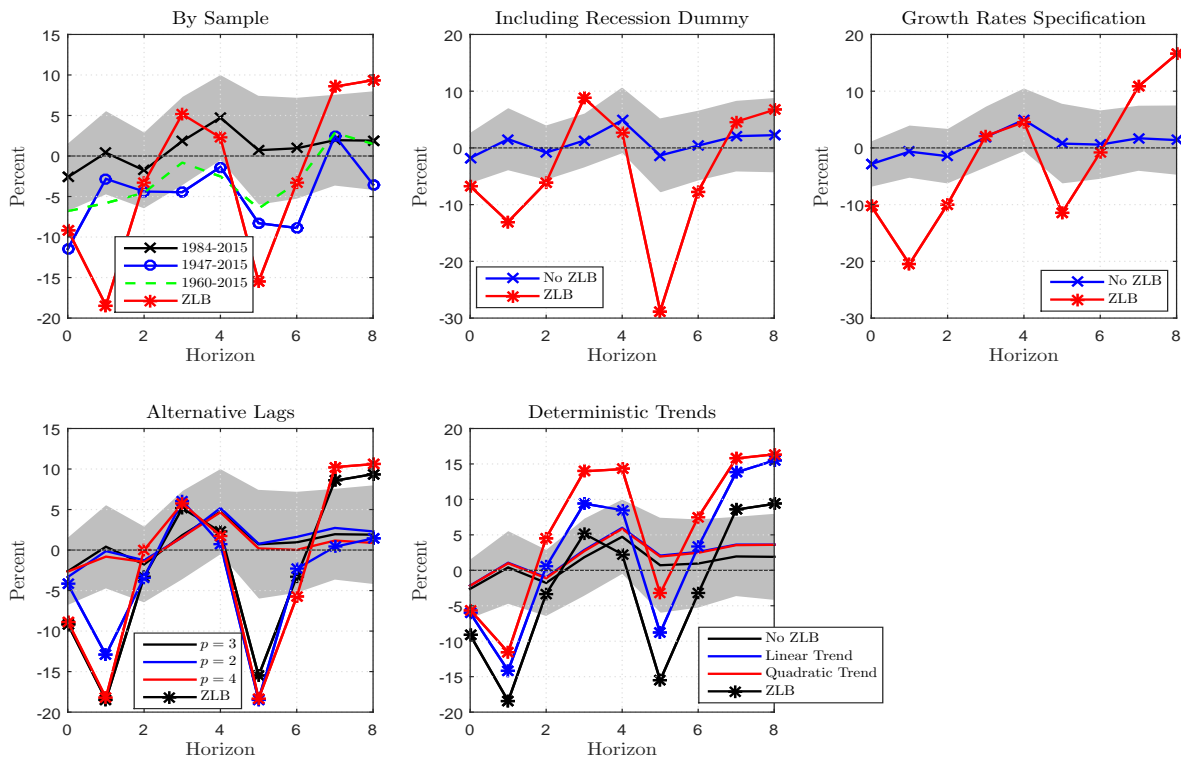


Figure 6: IRFs of Inflation to Productivity

Notes: This figure shows estimated impulse responses of inflation to a productivity shock, both inside and outside of the ZLB, for various different variations on our baseline regression model.

4 Conclusion

The textbook New Keynesian model predicts that the output response to a productivity shock is smaller when the nominal interest rate is constrained by the zero lower bound compared to periods where monetary policy is active. The mechanism by which this happens is that a positive productivity shock results in a decrease in inflation, which drives up the real interest rate when the nominal rate is constrained by the ZLB.

Contrary to the theory, we show that positive productivity shocks increase output more in periods where the zero lower bound binds in comparison to normal times. For inflation, we find results that are broadly consistent with the textbook NK model – inflation does not react significantly to a productivity shock in normal times, but falls significantly at the ZLB. Taken together, these results present something of a puzzle, and suggest that some important ingredient is missing from the textbook model.

One potential solution to reconcile theory and evidence is to discard the representative

agent assumption. [Wieland \(2015\)](#) shows that when a fraction of households are borrowing constrained, a negative supply shock at the ZLB is contractionary provided the elasticity of substitution is sufficiently small. In this regard, heterogeneous agent New Keynesian models as discussed in [Kaplan, Moll, and Violante \(2016\)](#) are promising. Another alternative is to build on [Bundick \(2015\)](#), whose results suggest that allowing history dependence in the central bank's policy rule could reconcile the basic model with our empirical results. Similarly, [Fernández-Villaverde \(2014\)](#) points out that the introduction of capital into the basic framework generates a positive output response to a productivity shock at the ZLB provided the shock is sufficiently transitory. This result depends on the magnitude of the investment adjustment cost in the model, which governs the extent to which the marginal product of capital is related to the real interest rate. Finally, [Fernández-Villaverde, Guerrón-Quintana, and Rubio-Ramírez \(2014\)](#) show that increases in future productivity at the ZLB act like a demand shock and therefore raise current output more than in an active monetary regime. They argue that productivity enhancing policies (e.g. deregulation) are likely to change productivity more in the future compared to the present. In other words, such shocks are more like demand shocks (due to the forward-looking nature of consumption and investment) than supply shocks.

In the meantime, since our empirical results are difficult to square with the textbook theory, caution seems to be in order when advocating for policies such as forward guidance and fiscal stimulus, both of which are predicted to be highly expansionary when policy is constrained by the ZLB.

References

- AUERBACH, A. AND Y. GORODNICHENKO (2013): “Fiscal Multipliers in Recession and Expansion,” in *Fiscal Policy After the Financial Crisis*, ed. by A. Alesina and F. Giavazzi, University of Chicago Press, 63–98.
- BACHMANN, R., T. BERG, AND E. SIMS (2015): “Inflation Expectations and Readiness to Spend at the Zero Lower Bound: Cross-Sectional Evidence,” *American Economic Journal: Economic Policy*, 7, 1–35.
- BASU, S., J. FERNALD, J. FISHER, AND M. KIMBALL (2013): “Sector-Specific Technical Change,” Tech. rep.
- BASU, S., J. G. FERNALD, AND M. S. KIMBALL (2006): “Are Technology Improvements Contractionary?” *American Economic Review*, 96, 1418–1448.
- BUNDICK, B. (2015): “Real Fluctuations at the Zero Lower Bound,” Tech. rep.
- BURKE, M. AND A. OZDAGLI (2013): “Household Inflation Expectations and Consumer Spending: Evidence from Panel Data,” Working Paper Series 13-25, Federal Reserve Bank of Boston.
- CARLSTROM, C., T. FUERST, AND M. PAUSTIAN (2014): “Fiscal Multipliers under an Interest Rate Peg of Deterministic versus Stochastic Duration,” *Journal of Money, Credit, and Banking*, 76, 1293–1312.
- (2015): “Inflation and Output in New Keynesian Models with a Transient Interest Rate Peg,” *Journal of Monetary Economics*, 76, 230–243.
- CHRISTIANO, L., M. EICHENBAUM, AND S. REBELO (2011): “When Is the Government Spending Multiplier Large?” *Journal of Political Economy*, 119, 78 – 121.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113, 1–45.
- COHEN-SETTON, J., J. K. HAUSMAN, AND J. F. WIELAND (2016): “Supply-Side Policies in the Depression: Evidence from France,” Working Paper 22140, National Bureau of Economic Research.
- D’ACUNTO, F., D. HOANG, AND M. WEBER (2016): “Unconventional Fiscal Policy, Inflation Expectations, and Consumption Expenditure,” Tech. rep.

- DEL-NEGRO, M. AND M. GIANNONI (2015): “The Forward Guidance Puzzle,” Staff Report 574, Federal Reserve Bank of New York.
- DUPOR, B. AND R. LI (2015): “The Expected Inflation Channel of Government Spending in the Postwar US,” *European Economic Review*, 74, 35–56.
- EGGERTSON, G. B. (2012): “Was the New Deal Contractionary?” *American Economic Review*, 102, 524–555.
- EGGERTSSON, G. B. AND P. KRUGMAN (2012): “Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach,” *The Quarterly Journal of Economics*, 127, 1469–1513.
- FERNALD, J. (2014): “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity,” Tech. Rep. 2012-19, Federal Reserve Bank of San Francisco.
- FERNÁNDEZ-VILLAYERDE, J. (2014): “Discussion of “Can Structural Reforms Help Europe?” by Gauti Eggertsson, Andrea Ferrero, and Andrea Raffo,” *Journal of Monetary Economics*, 61, 23 – 31, carnegie-Rochester-NYU Conference Series on “Fiscal Policy in the Presence of Debt Crises” held at the Stern School of Business, New York University on April 19-20, 2013.
- FERNÁNDEZ-VILLAYERDE, J., P. GUERRÓN-QUINTANA, AND F. J. RUBIO-RAMÍREZ (2014): “Supply-Side Policies and the Zero Lower Bound,” *IMF Economic Review*, 62, 248–260.
- GALI, J., J. D. LOPEZ-SALIDO, AND J. VALLES (2003): “Technology shocks and monetary policy: assessing the Fed’s performance,” *Journal of Monetary Economics*, 50, 723–743.
- ICHIUE, H. AND S. NISHIGUCHI (2015): “Inflation Expectations and Consumer Spending at the Zero Lower Bound: Micro Evidence,” *Economic Inquiry*, 53, 1086–1107.
- JORDÀ, O. (2005): “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 95, 161–182.
- KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2016): “Monetary Policy According to HANK,” Working Paper 21897, National Bureau of Economic Research.
- LASEEN, S. AND L. E. SVENSSON (2011): “Anticipated Alternative Instrument-Rate Paths in Policy Simulations,” *International Journal of Central Banking*, 7, 1–35.

- MULLIGAN, C. B. (2011): “Simple Analytics and Empirics of the Government Spending Multiplier and Other Keynesian Paradoxes,” *The B.E. Journal of Macroeconomics*, 11, 1–47.
- NEWBY, W. AND K. WEST (1987): “A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703–708.
- RAMEY, V. A. (2011): “Identifying Government Spending Shocks: It’s all in the Timing,” *The Quarterly Journal of Economics*, 126, 1–50.
- RAMEY, V. A. AND S. ZUBAIRY (2014): “Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data,” Working Paper 20719, National Bureau of Economic Research.
- ROMER, C. D. AND D. H. ROMER (2004): “A New Measure of Monetary Shocks: Derivation and Implications,” *American Economic Review*, 94, 1055–1084.
- (2010): “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks,” *American Economic Review*, 100, 763–801.
- SIMS, E. (2016): “Differences in Quarterly Utilization-Adjusted TFP by Vintage, with an Application to News Shocks,” Working Paper 22154, National Bureau of Economic Research.
- SMETS, F. AND R. WOUTERS (2007): “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *The American Economic Review*, 97, 586–606.
- WIELAND, J. F. (2015): “Are Negative Supply Shocks Expansionary at the Zero Lower Bound?” Tech. rep.

A Stochastic Peg

As an alternative to what we do in the main text, which is to approximate the effects of a binding zero lower bound with an interest rate peg of deterministic duration, in this Appendix we consider the case in which the duration of the interest rate peg is stochastic. This is the approach taken in the small-scale NK model in [Christiano et al. \(2011\)](#), for example. A stochastic peg length has the advantage that it permits clean closed form solutions, which is not the case for the deterministic peg case. A downside of the stochastic peg case is that it can result in counterintuitive “sign flips” in which the effect of a natural rate shock on output flips sign for a sufficiently long expected duration of the peg.

As in the main text, suppose that the current nominal interest rate equals zero, so in deviation terms we have $i_t = 1 - 1/\beta$. With probability $1 - p$, in period $t + 1$ the central bank returns to an inflation target, which implies that $\mathbb{E}_t \pi_{t+1} = \mathbb{E}_t x_{t+1} = 0$ and $\mathbb{E}_t i_{t+1} = \mathbb{E}_t r_{t+1}^f$. With probability p , the nominal interest rate in period $t + 1$ remains at zero. The probability of returning to the strict inflation target in any subsequent period, conditional on arriving in that period with the nominal rate still at zero, is fixed at p . We solve for analytic solutions for π_t and x_t using the method of undetermined coefficients. Using the expression mapping a_t into r_t^f from the text, (4), we can write these solutions as:

$$\pi_t = \frac{\gamma}{-\sigma(1 - \beta p)(1 - p) + p\gamma} \left(1 - \frac{1}{\beta}\right) + \frac{\gamma}{\sigma(1 - \beta p \rho_a)(1 - p \rho_a) - p \rho_a \gamma} \frac{\sigma(1 + \chi)(\rho_a - 1)}{\sigma + \chi} a_t \quad (\text{A.1})$$

$$x_t = \frac{1 - \beta p}{-\sigma(1 - \beta p)(1 - p) + p\gamma} \left(1 - \frac{1}{\beta}\right) + \frac{1 - \beta p}{\sigma(1 - \beta p \rho_a)(1 - p \rho_a) - p \rho_a \gamma} \frac{\sigma(1 + \chi)(\rho_a - 1)}{\sigma + \chi} a_t. \quad (\text{A.2})$$

Since $x_t = y_t - y_t^f$, and $y_t^f = \frac{1 + \chi}{\sigma + \chi} a_t$, this implies that output can be written:

$$y_t = \frac{1 - \beta p}{-\sigma(1 - \beta p)(1 - p) + p\gamma} \left(1 - \frac{1}{\beta}\right) + \left[1 + \frac{\sigma(\rho_a - 1)(1 - \beta p)}{\sigma(1 - \beta p \rho_a)(1 - p \rho_a) - p \rho_a \gamma}\right] \frac{1 + \chi}{\sigma + \chi} a_t. \quad (\text{A.3})$$

In [Figure A1a](#) we plot impulse responses of output to a productivity shock for two different levels of p : $p \in \{2/3, 4/5\}$. This corresponds to expected durations of three and five quarters, respectively. For point of comparison, we also show the case in which the central bank targets a zero inflation rate in every period, in which case $y_t = y_t^f$. We assume that $\rho_a = 0.90$, $\sigma = 1$, $\chi = 0$, $\beta = 0.99$, and $\phi = 0.75$. Clearly, as the expected duration of the ZLB increases, the less output expands on impact in response to a productivity shock. For sufficiently long forecast horizons the responses are not affected much by the value of p , because in expectation the economy will have likely exited the ZLB. Differently than the deterministic peg case, the response under a stochastic peg only asymptotically approaches the flexible price responses,

whereas in the deterministic peg the responses lie on top of one another after the peg period.




Figure A1: Responses to a Productivity Shock as a Function of Duration of ZLB From A Stochastic Peg

Notes: These figures plots the impulse responses of output (right) inflation (left) to a one percent increase in productivity for different values of α . The time period is a quarter.

Figure A1b is similar to Figure A1a, except we plot the inflation response for different expected durations of the peg. As in the main text, the longer is the expected duration of the peg, the more inflation falls on impact.

As documented in Carlstrom et al. (2014), for sufficiently high values of p , the sign of the effect of a productivity shock on output and inflation can flip. For the values of the other parameters we have chosen, this sign flip occurs at about $p = 0.83$, or an expected duration of the peg of about six quarters. The sign flips do not occur in the deterministic duration case considered in the text. It is important to reiterate that the sign flip applies to both the output and inflation response – if output responds more to the productivity shock at the ZLB, then the inflation response is positive at the ZLB, rather than negative.

B Taylor Rule

The IS equation and Phillips Curve (PC) are the same as in the main text:

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma} (i_t - \mathbb{E}_t \pi_{t+1} - r_t^f) \quad (\text{B.1})$$

$$\pi_t = \gamma x_t + \beta \mathbb{E}_t \pi_{t+1}. \quad (\text{B.2})$$

Outside of the ZLB, the interest rate rule is given by

$$i_t = \phi_\pi \pi_t.$$

The process for the natural rate of interest is the same. We first solve for the policy functions of the output gap and inflation outside the ZLB and then consider the ZLB. Substitute the interest rate rule into Equation B.1:

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma} (\phi_\pi \pi_t - \mathbb{E}_t \pi_{t+1} - r_t^f).$$

Guess $x_t = \theta_1 a_t$ and $\pi_t = \theta_2 a_t$. After some algebraic manipulations:

$$\begin{aligned} x_t &= \frac{(1 - \beta\rho)\sigma^{-1}}{(1 - \beta\rho)(1 - \rho) + \frac{\gamma}{\rho}(\phi_\pi - \rho)} \frac{1 + \chi}{\sigma + \chi} (\rho - 1) a_t \\ \pi_t &= \frac{\gamma\sigma^{-1}}{(1 - \beta\rho)(1 - \rho) + \frac{\gamma}{\rho}(\phi_\pi - \rho)} \frac{1 + \chi}{\sigma + \chi} (\rho - 1) a_t. \end{aligned}$$

The peg runs for H periods and lifts in period H . This implies that $\mathbb{E}_t x_{t+H} = \theta_1 \rho^H a_t$ and $\mathbb{E}_t \pi_{t+H} = \theta_2 \rho^H a_t$. Since the interest rate is now constrained at 0, this means $\mathbb{E}_t i_{t+h} = 1 - 1/\beta$ for all $h < H$. In period $H - 1$ we have:

$$\begin{aligned} \mathbb{E}_t x_{t+H-1} &= \mathbb{E}_t x_{t+H} + \frac{1}{\sigma} \left(1 - \frac{1}{\beta} + \mathbb{E}_t \pi_{t+H} + \mathbb{E}_t r_{t+H-1}^f \right) \\ &= \theta_1 \rho^H a_t + \frac{1}{\sigma} \left(1 - \frac{1}{\beta} + \theta_2 \rho^H a_t + \delta \rho^{H-1} a_t \right) \end{aligned}$$

where $\delta = \frac{1+\chi}{\chi+\sigma}(\rho - 1)$. Substitute the last expression into Equation B.2:

$$\begin{aligned} \mathbb{E}_t \pi_{t+H-1} &= \gamma \mathbb{E}_t x_{t+H-1} + \beta \theta_2 \rho^H a_t \\ &= \gamma \left[\theta_1 \rho^H a_t + \frac{1}{\sigma} \left(1 - \frac{1}{\beta} + \theta_2 \rho^H a_t + \delta \rho^{H-1} a_t \right) \right] + \beta \theta_2 \rho^H a_t. \end{aligned}$$

We continue to iterate back to period t .

With the exception of ϕ_π , which is set to 1.5, the rest of the parameterization is identical to the one in the main text. We consider peg lengths of $H \in \{0, 3, 6, 10\}$. Figure B1 presents the results. Note that when $H = 0$ output does not increase by as much as productivity. This is the consequence of price stickiness and it is exactly what inflation targeting avoids, namely non-zero values of the output gap and inflation. Also note that the fall in output for longer peg lengths is significantly bigger than when the unconstrained rule is inflation targeting. Relatedly, output can decline on impact for shorter durations of the peg than in the strict

inflation targeting case.

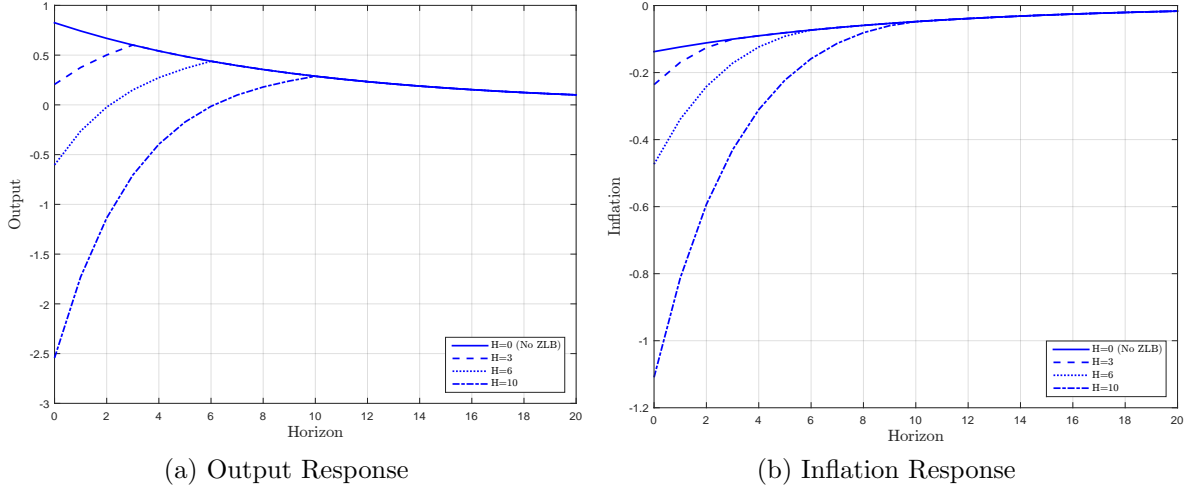


Figure B1: Responses to a Productivity Shock as a Function of Duration of ZLB From a Taylor Rule

Notes: These figures plots the impulse responses of output (right) inflation (left) to a one percent increase in productivity for different durations of a pegged nominal interest rate at zero. $H = 0$ corresponds to the case where the central bank obeys a Taylor Rule.

C Medium Scale Model

In this Appendix we show that the theoretical results derived in Section 2 continue to hold in a medium scale model similar to [Smets and Wouters \(2007\)](#) and [Christiano, Eichenbaum, and Evans \(2005\)](#). Specifically, we include: capital accumulation, investment adjustment costs, variable capital accumulation, nominal price and wage rigidities, partial wage and price indexation, and habit formation.

As the model is fairly standard, we only present the first order conditions characterizing the equilibrium of the model:

$$\lambda_t = (C_t - bC_{t-1})^{-1} - \beta b \mathbb{E}_t (C_{t+1} - bC_t)^{-1} \quad (\text{C.1})$$

$$\lambda_t = \beta(1 + i_t) \mathbb{E}_t \lambda_{t+1} (1 + \pi_{t+1})^{-1} \quad (\text{C.2})$$

$$\lambda_t R_t = \mu_t [\delta_1 + \delta_2 (u_t - 1)] \quad (\text{C.3})$$

$$\mu_t = \beta \mathbb{E}_t [\lambda_{t+1} R_{t+1} u_{t+1} + (1 - \delta(u_{t+1})) \mu_{t+1}] \quad (\text{C.4})$$

$$\lambda_t = \mu_t \left[1 - \frac{\kappa}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right) - \kappa \left(\frac{I_t}{I_{t-1}} - 1 \right) \frac{I_t}{I_{t-1}} \right] + \beta \mathbb{E}_t \mu_{t+1} \kappa \left(\frac{I_{t+1}}{I_t} - 1 \right) \left(\frac{I_{t+1}}{I_t} \right)^2 \quad (\text{C.5})$$

$$w_t^\# = \frac{\epsilon_w}{\epsilon_w - 1} \frac{f_{1,t}}{f_{2,t}} \quad (\text{C.6})$$

$$f_{1,t} = \psi \left(\frac{w_t}{w_t^\#} \right)^{\epsilon_w(1+\chi)} N_t^{1+\chi} + \beta \theta_w \mathbb{E}_t \left(\frac{w_{t+1}^\#}{w_t^\#} \right)^{\epsilon_w(1+\chi)} \left(\frac{(1+\pi_t)^{\zeta_w}}{1+\pi_{t+1}} \right)^{-\epsilon_w(1+\chi)} f_{1,t+1} \quad (\text{C.7})$$

$$f_{2,t} = \lambda_t \left(\frac{w_t}{w_t^\#} \right)^{\epsilon_w} N_t + \beta \theta_w \mathbb{E}_t \left(\frac{w_{t+1}^\#}{w_t^\#} \right)^{\epsilon_w} \left(\frac{(1+\pi_t)^{\zeta_w}}{1+\pi_{t+1}} \right)^{1-\epsilon_w} f_{2,t+1} \quad (\text{C.8})$$

$$R_t = \alpha m c_t A_t \tilde{K}_t^{\alpha-1} N_t^{1-\alpha} \quad (\text{C.9})$$

$$w_t = (1-\alpha) m c_t A_t \tilde{K}_t^\alpha N_t^{-\alpha} \quad (\text{C.10})$$

$$\frac{1+\pi_t^\#}{1+\pi_t} = \frac{\epsilon_p}{\epsilon_p - 1} \frac{x_{1,t}}{x_{2,t}} \quad (\text{C.11})$$

$$x_{1,t} = \lambda_t m c_t Y_t + \beta \theta_p \mathbb{E}_t (1+\pi_t)^{-\zeta_p \epsilon_p} (1+\pi_{t+1})^{\epsilon_p} x_{1,t+1} \quad (\text{C.12})$$

$$x_{2,t} = \lambda_t Y_t + \beta \theta_p (1+\pi_t)^{\zeta_p(1-\epsilon_p)} E_t (1+\pi_{t+1})^{\epsilon_p-1} x_{2,t+1} \quad (\text{C.13})$$

$$K_{t+1} = \left[1 - \frac{\kappa}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2 \right] I_t + [1 - \delta(u_t)] K_t \quad (\text{C.14})$$

$$\delta(u_t) = \delta_0 + \delta_1(u_t - 1) + \frac{\delta_2}{2}(u_t - 1)^2 \quad (\text{C.15})$$

$$Y_t = C_t + I_t \quad (\text{C.16})$$

$$\tilde{K}_t = u_t K_t \quad (\text{C.17})$$

$$Y_t v_t^p = A_t \tilde{K}_t^\alpha N_t^{1-\alpha} - F \quad (\text{C.18})$$

$$v_t^p = (1+\pi_t)^{\epsilon_p} \left[(1-\theta_p)(1+\pi_t^\#)^{-\epsilon_p} + \theta_p(1+\pi_{t-1})^{-\epsilon_p \zeta_p} v_{t-1}^p \right] \quad (\text{C.19})$$

$$(1+\pi_t)^{1-\epsilon_p} = (1-\theta_p)(1+\pi_t^\#)^{1-\epsilon_p} + \theta_p(1+\pi_{t-1})^{\zeta_p(1-\epsilon_p)} \quad (\text{C.20})$$

$$w_t^{1-\epsilon_w} = (1-\theta_w) w_t^{\#, 1-\epsilon_w} + \theta_w \left[\frac{(1+\pi_{t-1})^{\zeta_w}}{1+\pi_t} w_{t-1} \right]^{1-\epsilon_w} \quad (\text{C.21})$$

$$i_t = (1-\rho_i)i + \rho_i i_{t-1} + (1-\rho_i) [\phi_\pi(\pi_t - \pi_t^*) + \phi_y(\ln Y_t - \ln Y_{t-1})] \quad (\text{C.22})$$

$$1+r_t = (1+i_t) \mathbb{E}_t (1+\pi_{t+1})^{-1} \quad (\text{C.23})$$

$$\ln A_t = \rho_A \ln A_{t-1} + \varepsilon_{A,t}. \quad (\text{C.24})$$

In these equations λ_t is the Lagrange multiplier on the flow budget constraint and μ_t is the Lagrange multiplier on the accumulation equation. (C.1) defines λ_t in terms of the marginal utility of consumption. (C.2) is the Euler equation for bonds, which prices the nominal interest rate, i_t . (C.3) is the first order condition for capital utilization. The optimality

condition for the choice of K_{t+1} is (C.4), and the FOC for investment is given by (C.5). Optimal wage-setting for updating households is characterized by (C.6)–(C.8), where $w_t^\#$ is the reset real wage, which is common across updating households. Cost-minimization by firms defines the capital-labor ratio and real marginal cost in (C.9)–(C.10). Optimal price-setting for updating firms is characterized by (C.11)–(C.13), where $\pi_t^\# = P_t^\# / P_{t-1} - 1$ is the reset inflation rate, which is common across updating firms. The aggregate production function is given by (C.18). v_t^p is a measure of price dispersion which can be written recursively as in (C.19). The evolution of aggregate inflation is given by (C.20) and the aggregate real wage evolves according to (C.21). The real interest rate is defined in the Fisher relationship, (C.23). Monetary policy during normal times is characterized by a Taylor rule, given in (C.22). The exogenous process for productivity is given by (C.24), where the non-stochastic level of productivity is normalized to unity. The model is solved by linearization about a zero inflation non-stochastic steady state.

To approximate the effects of a binding zero lower bound, we augment the Taylor rule with monetary policy news shocks as in [Laseen and Svensson \(2011\)](#). Then, conditional on a productivity shock, we solve for the values of these news shocks so as to keep the nominal interest rate fixed (in expectation) for a specified period of time. The model is parameterized as follows: $\beta = 0.995$, $\epsilon_p = \epsilon_w = 11$, $\psi = 6$, $\delta_2 = 0.05$, $\alpha = 1/3$, $\chi = 1$, $\theta_w = \theta_p = 0.66$, $\zeta_p = \zeta_w = 0$, $\delta_0 = 0.025$, $b = 0.7$, $\phi_\pi = 1.5$, $\phi_y = 0.2$, $\rho_i = 0.8$, $\kappa = 4$, and $\rho_A = 0.95$. δ_1 is chosen to be consistent with steady state utilization of 1, and F is chosen so that profits are zero in the steady state. We consider a shock to productivity of one percent.

Figure C1 plots impulse responses of output and inflation to a positive productivity shock. The solid lines are the responses outside the ZLB, the dashed lines are responses when the nominal interest rate is pegged for a duration of eight quarters.

Here, we see exactly the same behavior (qualitatively) as in the model without capital. At the ZLB, output declines, inflation falls by more, and the real interest rate rises. Note that the real interest rate rises on impact in both normal times as well as at the ZLB. This is different than the model without capital, where the real interest rate declines after a temporary productivity improvement. Outside of the ZLB, the impact increase in the real interest rate is small and quickly turns negative. What accounts for the difference relative to the textbook model is the presence of capital. A positive productivity shock raises the marginal product of capital, which works to put upward pressure on the real interest rate. In the model with investment adjustment costs, this effect is small and only temporary. Without investment adjustments costs, the impact rise in the real interest rate outside of the ZLB is much stronger and more persistent.




Figure C1: Response to Productivity Shock, with and without ZLB

C.1 Permanent Productivity Shocks

Now suppose that exogenous productivity evolves according to a non-stationary stochastic process instead of a mean-reverting process. In particular:

$$\ln g_t = (1 - \rho_A) \ln g_A + \rho_A \ln g_{t-1} + \varepsilon_{g,t} \quad (\text{C.25})$$

where $\ln g_t = \ln A_t - \ln A_{t-1}$ and g_A denotes the steady state gross growth rate of productivity. If $\rho_A = 1$, productivity obeys a random walk with drift. If $\rho_A > 0$, then the growth rate of productivity follows a stationary AR(1) process. This specification introduces stochastic trends into the model. Most variables need to be detrended to be rendered stationary. Define $X_t = A_t^{\frac{1}{1-\alpha}}$ as the trend factor. Define $\widehat{H}_t = H_t/X_t$ for generic variable H_t . Exceptions will be λ_t and μ_t , for which the stationary transformations will be $\widehat{\lambda}_t = \lambda_t X_t$. The real and nominal interest rates, the inflation rate, capital utilization, and labor hours will all be stationary without need for transformation.

The full set of detrended equilibrium conditions are presented below:

$$\widehat{\lambda}_t = (\widehat{C}_t - b g_{X,t}^{-1} \widehat{C}_{t-1})^{-1} - \beta \mathbb{E}_t b (\widehat{C}_{t+1} g_{X,t+1} - b \widehat{C}_t)^{-1} \quad (\text{C.26})$$

$$\widehat{\lambda}_t = \beta (1 + i_t) \mathbb{E}_t \widehat{\lambda}_{t+1} g_{X,t+1}^{-1} (1 + \pi_{t+1})^{-1} \quad (\text{C.27})$$

$$\widehat{\lambda}_t = \widehat{\mu}_t [\delta_1 + \delta_2(u_t - 1)] \quad (\text{C.28})$$

$$\widehat{\mu}_t = \beta \mathbb{E}_t g_{X,t+1}^{-1} [\widehat{\lambda}_{t+1} R_{t+1} u_{t+1} + (1 - \delta(u_{t+1})) \widehat{\mu}_{t+1}] \quad (\text{C.29})$$

$$\begin{aligned} \widehat{\lambda}_t = \widehat{\mu}_t & \left[1 - \frac{\kappa}{2} \left(\frac{\widehat{I}_t}{\widehat{I}_{t-1}} g_{X,t} - g_X \right) - \kappa \left(\frac{\widehat{I}_t}{\widehat{I}_{t-1}} g_{X,t} - g_X \right) \frac{\widehat{I}_t}{\widehat{I}_{t-1}} g_{X,t} \right] \\ & + \beta \mathbb{E}_t g_{X,t+1}^{-1} \widehat{\mu}_{t+1} \kappa \left(\frac{\widehat{I}_{t+1}}{\widehat{I}_t} g_{X,t+1} - g_X \right) \left(\frac{\widehat{I}_{t+1}}{\widehat{I}_t} g_{X,t+1} \right)^2 \end{aligned} \quad (\text{C.30})$$

$$f_{1,t} = \psi \left(\frac{\widehat{w}_t}{\widehat{w}_t^\#} \right)^{\epsilon_w(1+\chi)} N_t^{1+\chi} + \beta \theta_w \mathbb{E}_t \left(\frac{\widehat{w}_{t+1}^\#}{\widehat{w}_t^\#} g_{X,t+1} \right)^{\epsilon_w(1+\chi)} \left(\frac{(1 + \pi_t)^{\zeta_w}}{1 + \pi_{t+1}} \right)^{-\epsilon_w(1+\chi)} f_{1,t+1} \quad (\text{C.31})$$

$$\widehat{f}_{2,t} = \widehat{\lambda}_t \left(\frac{\widehat{w}_t}{\widehat{w}_t^\#} \right)^{\epsilon_w} N_t + \beta \theta_w \mathbb{E}_t g_{X,t+1}^{-1} \left(\frac{\widehat{w}_{t+1}^\#}{\widehat{w}_t^\#} g_{X,t+1} \right)^{\epsilon_w} \left(\frac{(1 + \pi_t)^{\zeta_w}}{1 + \pi_{t+1}} \right)^{1-\epsilon_w} \widehat{f}_{2,t+1} \quad (\text{C.32})$$

$$\widehat{w}_t^\# = \frac{\epsilon_w}{\epsilon_w - 1} \frac{f_{1,t}}{\widehat{f}_{2,t}} \quad (\text{C.33})$$

$$R_t = \alpha g_{X,t}^{1-\alpha} m c_t \widehat{K}_t^\alpha N_t^{1-\alpha} \quad (\text{C.34})$$

$$\widehat{w}_t = (1 - \alpha) m c_t g_{X,t}^{-\alpha} \widehat{K}_t^\alpha N_t^{-\alpha} \quad (\text{C.35})$$

$$\frac{1 + \pi_t^\#}{1 + \pi_t} = \frac{\epsilon_p}{\epsilon_p - 1} \frac{x_{1,t}}{x_{2,t}} \quad (\text{C.36})$$

$$x_{1,t} = \widehat{\lambda}_t m c_t \widehat{Y}_t + \beta \theta_p \mathbb{E}_t (1 + \pi_t)^{-\zeta_p \epsilon_p} (1 + \pi_{t+1})^{\epsilon_p} x_{1,t+1} \quad (\text{C.37})$$

$$x_{2,t} = \widehat{\lambda}_t \widehat{Y}_t + \beta \theta_p (1 + \pi_t)^{\zeta_p(1-\epsilon_p)} \mathbb{E}_t (1 + \pi_{t+1})^{\epsilon_p - 1} x_{2,t+1} \quad (\text{C.38})$$

$$\widehat{K}_{t+1} = \left[1 - \frac{\kappa}{2} \left(\frac{\widehat{I}_t}{\widehat{I}_{t-1}} g_{X,t} - g_X \right)^2 \right] \widehat{I}_t + [1 - \delta(u_t)] g_{X,t}^{-1} \widehat{K}_t \quad (\text{C.39})$$

$$\widehat{Y}_t = \widehat{C}_t + \widehat{I}_t \quad (\text{C.40})$$

$$\widehat{K}_t = u_t \widehat{K}_t \quad (\text{C.41})$$

$$\widehat{Y}_t v_t^p = g_{X,t}^{-\alpha} \widehat{K}_t^\alpha N_t^{1-\alpha} - F \quad (\text{C.42})$$

$$v_t^p = (1 + \pi_t)^{\epsilon_p} \left[(1 - \theta_p) (1 + \pi_t^\#)^{-\epsilon_p} + \theta_p (1 + \pi_{t-1})^{-\epsilon_p \zeta_p} v_{t-1}^p \right] \quad (\text{C.43})$$

$$(1 + \pi_t)^{1-\epsilon_p} = (1 - \theta_p) (1 + \pi_t^\#)^{1-\epsilon_p} + \theta_p (1 + \pi_{t-1})^{\zeta_p(1-\epsilon_p)} \quad (\text{C.44})$$

$$\widehat{w}_t^{1-\epsilon_w} = (1 - \theta_w) \widehat{w}_t^{\#, 1-\epsilon_w} + \theta_w \left[\frac{(1 + \pi_{t-1})^{\zeta_w}}{1 + \pi_t} \widehat{w}_{t-1} g_{X,t}^{-1} \right]^{1-\epsilon_w} \quad (\text{C.45})$$

$$i_t = (1 - \rho_i)i + \rho_i i_{t-1} + (1 - \rho_i)[\phi_\pi(\pi_t - \pi_t^*) + \phi_y(\ln \widehat{Y}_t - \ln \widehat{Y}_{t-1})] \quad (\text{C.46})$$

$$1 + r_t = (1 + i_t) \mathbb{E}_t(1 + \pi_{t+1})^{-1} \quad (\text{C.47})$$

$$\ln g_t = (1 - \rho_A) \ln g_A + \rho_A \ln g_{t-1} + \varepsilon_{g,t} \quad (\text{C.48})$$

$$g_{X,t} = g_t^{\frac{1}{1-\alpha}}. \quad (\text{C.49})$$

The equilibrium conditions are the same as in the case with a mean-reverting productivity process, only re-written according to stationary transformations. $g_{X,t}$ is the gross growth rate of the trend factor, and is given by (C.49). With the exception of g_A and ρ_A , the model is parameterized as before.

C.1.1 IRFs in the Permanent Shock Case

We produce IRFs in the permanent shock case for different values of ρ_A . We assume that $g_A = 1$. Our results are not affected by this assumption. First, consider $\rho_A = 0$, so that productivity is an exact random walk. The IRFs of output, the nominal interest rate, the inflation rate, and the real interest rate are shown below in Figure C2. The solid lines depict the responses when policy is governed by the Taylor rule, while the dashed lines represent responses when the nominal interest rate is pegged for eight quarters.




Figure C2: Response to Productivity Shock, with and without ZLB, $\rho_A = 0$

These responses are qualitatively different than the textbook model presented in the text, as well as compared to the responses from the medium scale model when productivity obeys a stationary stochastic process. In particular, output increases by more at the ZLB in comparison to when monetary policy is governed by a Taylor rule. This pattern of response would seem to be consistent with the empirical results presented in Section 3. However, different than our empirical results, in the model inflation falls by less at the ZLB compared to normal times. As we discuss in Section 3.3, empirically it is the joint behavior of output and inflation in response to a productivity shock that is difficult to square with the theory, not the response of output in isolation.

Next, consider higher levels of persistence (so that the shock is correlated in growth rates). First, we set $\rho_A = 0.1$. Qualitatively, the responses are similar to the random walk case shown in Figure C2, though the differences between the interest rate peg and Taylor rule case are quantitatively smaller.




Figure C3: Response to Productivity Shock, with and without ZLB,
 $\rho_A = 0.1$

Lastly, we increase the persistence of the productivity process further, setting $\rho_A = 0.4$. The responses are shown in Figure C4. we are back in the case where output responds less (and inflation falls more). The results are presented in Figure C4. In terms of the impact responses of output and inflation, these responses differ from the random walk case and are more similar to the textbook model and the medium scale model with a stationary

productivity shock. Output rises by less, and inflation falls by more, on impact after a positive shock to the growth rate of productivity.




Figure C4: Response to Productivity Shock, with and without ZLB, $\rho_A = 0.4$

In summary, the basic logic of the model in Section 2 is preserved when capital and additional real and nominal frictions are added to the model. Output rises by less and inflation falls by more at the ZLB than outside of it when the productivity shock is persistent but stationary. When the productivity process is an exact random walk, on the other hand, output and inflation increase by more under an interest rate peg than under a Taylor rule. If the productivity process is sufficiently persistent in growth rates, output again rises by less on impact, and inflation falls by more, to a positive productivity shock. Though there is not a robust prediction from the medium scale model on the sign of the effect of a binding ZLB on the response of output, what is robust is the joint behavior of output and inflation. If inflation falls by more at the ZLB, then output increases by less, and vice-versa. This pattern is not consistent with the empirical responses of output and inflation to a productivity shock which we identify in the data.