Are Supply Shocks Contractionary at the ZLB? Evidence from Utilization-Adjusted TFP Data*

Julio Garín†
Claremont McKenna College

Robert Lester‡
Colby College

Eric Sims§
University of Notre Dame
& NBER

This Version: July 20, 2017

Abstract

The basic New Keynesian model predicts that positive supply shocks are less expansionary at the zero lower bound (ZLB) compared to periods of active monetary policy. We test this prediction empirically using Fernald (2014)'s utilization-adjusted total factor productivity series, which we take as a measure of exogenous productivity. In contrast to the predictions of the model, positive productivity shocks are estimated to be more expansionary at the ZLB compared to normal times. We find that there is no significant difference in the response of expected inflation to a productivity shock at the ZLB compared to normal times.

JEL Classification: E31; E43; E52.

Keywords: Zero Lower Bound; Interest Rates; Supply Shocks.

*We are grateful to Rüdiger Bachmann, Jesús Fernández-Villaverde, Yuriy Gorodnichenko, Bill Lastrapes, Ron Mau, seminar participants at the University of Maryland, Michigan State University, Claremont McKenna College, the 21st LACEA Meetings, and three anonymous referees for helpful comments and suggestions which have improved the paper. We also thank Régis Barnichon for kindly providing us with his code to estimate impulse responses via smooth local projections. The usual disclaimer applies.

†E-mail address: jgarin@cmc.edu.

‡E-mail address: rblester@colby.edu.

§E-mail address: esims1@nd.edu.

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 quarterly 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 a smoothed version of Jordà (2005)'s local projections method which has recently been proposed by Barnichon and Brownlees (2016). 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 and for several quarters thereafter to an increase in productivity when the ZLB binds in comparison to periods where it does not. The differences in output responses at and away from the ZLB are both economically and statistically significant. In particular, we find that output increases by more than three times as much to a positive productivity shock on impact when the ZLB binds than when it does not. We can easily reject the hypothesis of pointwise equality of impulse responses to the productivity shock across monetary regimes for several forecast horizons, and can reject pathwise equality of responses at all horizons. We also find similar results when focusing on labor market indicators such as total hours worked or the unemployment rate – a positive productivity shock is estimated to be significantly more expansionary when the ZLB binds than when it does not. This basic result is also robust to several different variations on our baseline empirical specification.

Since a decline in expected 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 expected inflation, both inside and outside of the ZLB. We consider several different measures of expected inflation, including the Michigan Survey of Consumers, the Survey of Professional Forecasters, and the Federal Reserve Bank of Cleveland's inflation expectations model. Outside of the ZLB we find that a positive productivity shock results in a mild decline in expected inflation. Over most forecast horizons, we find no significant difference in the estimated responses of expected inflation to a productivity shock when the ZLB binds versus when it does not. This is true regardless of which inflation expectations series we use.

Taken together, our empirical results are puzzling from the perspective of the basic NK model. Intertemporal substitution is the key economic mechanism in the model. When the nominal interest rate is fixed because of a binding ZLB, the real interest rate moves one-for-one in the opposite direction of expected inflation, and the behavior of current and expected real interest rates determines output. If positive productivity shocks were to raise expected inflation at the ZLB, it would be conceivable that output could rise more when the ZLB binds compared to when it does not. But this is not consistent with what we find – expected inflation falls after a positive productivity shock both when the ZLB binds and when it does not. Furthermore, while not typically statistically significant, in most specifications we find that expected inflation falls more when the ZLB binds compared to when it does not. Our empirical results therefore suggest some failing of the textbook NK model and its more complicated but closely related medium scale DSGE variant. Caution is of course in order when interpreting our results, as they are based on an admittedly small sample size (twenty-nine observations where the ZLB binds). But viewed in conjunction with other empirical evidence (some of which is discussed below), we feel that there is a compelling empirical case that some important ingredient is missing from the textbook NK model.

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 expected inflation, which pushes down real interest rates when nominal rates are constrained by the ZLB. Our empirical findings suggest that this "expected 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 productivity 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 (2017) 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 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 (2017) 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.

[1] Also 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 was more consistent with models of flexible prices.

2 Theory

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

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

$$\pi_t = \gamma x_t + \beta\, \mathbb{E}_t\, \pi_{t+1}. \tag{2}$$

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 $0 \leq \rho_a < 1$:

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

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

$$y_t^f = \frac{1+\chi}{\sigma+\chi} a_t \tag{4}$$

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

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

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

policy under this kind of peg is characterized by:[3]

$$\mathbb{E}_t\, i_{t+h} = \begin{cases} 1 - \frac{1}{\beta} & \text{if } h < H \\ \mathbb{E}_t\, r^f_{t+h} & \text{if } h \geq H. \end{cases} \tag{6}$$

We solve for the expected time path of output backwards. Starting in period $t + H$, agents will expect $\mathbb{E}_t\, i_{t+H} = \mathbb{E}_t\, r^f_{t+H}$, which implies that $\mathbb{E}_t\, \pi_{t+H} = \mathbb{E}_t\, x_{t+H} = 0$. This means that $\mathbb{E}_t\, x_{t+H-1} = -\frac{1}{\sigma}\left(1 - \frac{1}{\beta}\right) + \frac{1}{\sigma}\mathbb{E}_t\, r^f_{t+H-1}$ 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 $\rho_a = 0.9$.

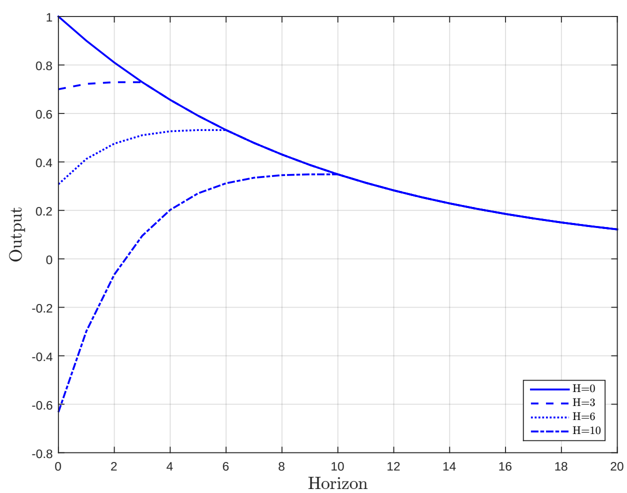

Figure 1: Response of Output to Productivity Shock as a Function of Duration of ZLB

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.

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

[3] It 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, $\mathbb{E}_t\, i_{t+h} = \mathbb{E}_t\, r^f_{t+h}$ in equilibrium, not as a policy rule that might hold out of equilibrium.

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 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 expected inflation. Dupor and Li (2015) have termed this the "expected inflation channel." In particular, a positive productivity shock lowers expected inflation when monetary policy is passive. The longer the nominal interest rate is pegged, the more expected 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 expected inflation to a productivity shock as a function of the duration of the interest rate peg, H.

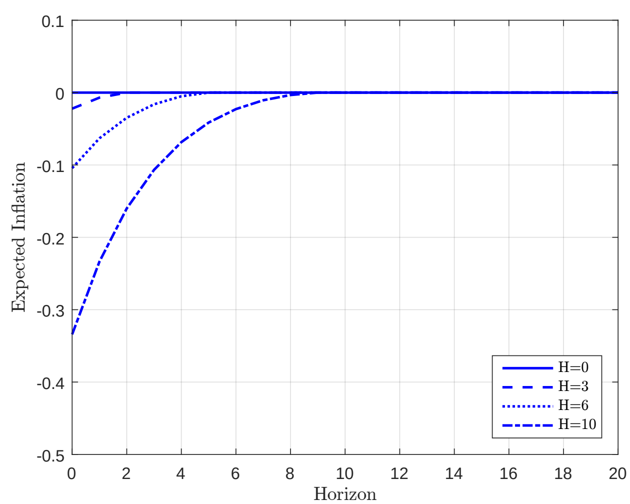

Figure 2: Response of Expected Inflation to Productivity Shock as a Function of Duration of ZLB

Notes: This figure plots the impulse responses of one period ahead expected 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.

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. Expected 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, expected inflation falls by less, not more, to the productivity shock.

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 expected inflation falls more. However, like Carlstrom, Fuerst, and Paustian (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 expected inflation response is positive, rather than negative.

[4] Carlstrom, Fuerst, and Paustian (2015) begin with a standard NK model and show that the responses of output and inflation to a deterministic period of forward guidance are exponentially increasing (decreasing for inflation) 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.

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. Given its central role in the transmission of productivity shocks at the ZLB in the textbook NK model, we also examine the effects of productivity shocks on expected inflation. 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 outlines our methodology for obtaining impulse responses at and away from the ZLB. Subsection 3.3 presents our main results concerning the effects of productivity shocks on output and expected inflation, both inside and outside the ZLB. Subsection 3.4 considers several robustness exercises. Subsection 3.5 studies whether a productivity shock has the same persistence during the ZLB period compared to the period prior and discusses whether and how the answer to this question affects the conclusions one can draw from our empirical 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:

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

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 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.

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. Output is measured as real GDP from the NIPA tables. We focus on the full available sample period, 1947Q2–2017Q1.

In terms of volatilities, the utilization-adjusted TFP is slightly less volatile than the conventional TFP series. Both the utilization-adjusted and conventional TFP series are less volatile than output growth. Utilization-adjusted TFP growth is not autocorrelated. Conventional TFP growth, in contrast, is positively autocorrelated, with an AR(1) correlation of 0.18. Output growth is also quite autocorrelated, with an AR(1) correlation of 0.37. The lower part of the table shows the correlation matrix for these variables. Utilization-adjusted TFP is positively correlated with output growth, though only mildly so with a correlation coefficient of 0.13. In contrast, the conventional TFP series is highly procyclical, with a correlation with output growth of 0.82.

Table 1: Empirical Moments

		$\Delta \ln A_t$	$\Delta \ln Y_t$	$\Delta \ln TFP_t$
Standard Deviation		0.0082	0.0095	0.0086
AR(1)		-0.029	0.373	0.177
	$\Delta \ln A_t$	1	0.125	0.441
Correlation	$\Delta \ln Y_t$		1	0.819
Matrix	$\Delta \ln TFP_t$			1

Notes: These moments are calculated for the period 1947Q2-2017Q1. 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.

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 shock variables. Under the null hypothesis that the utilization-adjusted TFP series is a measure of exogenous productivity, it should not be predictable from other exogenous shocks. We take four popular measures of macroeconomic shocks identified in the literature – Romer and Romer (2004) monetary policy shocks, Romer and Romer (2010) tax shocks, the defense news shock produced by Ramey (2011), and an exogenous oil price shock from Kilian (2008). These shocks are identified using either narrative methods or time series models. 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. These results are suggestive, but of course not dispositive, that Fernald's series can be treated as exogenous.

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

Measure/Variable	F statistic	p-value
Romer and Romer (2004) Monetary	0.6135	0.5436
Romer and Romer (2010) Tax	0.3112	0.7325
Ramey (2011) Gov. Spending	1.7614	0.1739
Kilian (2008) Oil	0.8699	0.4214

Notes: All the tests are performed using two lags using the maximum available sample period for each shock. The Romer and Romer (2004) monetary shock, the Romer and Romer (2010) tax shock, and the Ramey (2011) government spending shock were downloaded from Valerie Ramey's website in code to accompany Ramey (2016). The Kilian (2008) oil shock data were downloaded from Lutz Kilian's website. The sample periods for each shock are as follows: Romer and Romer (2004), 1983Q1–2007Q4; Romer and Romer (2010), 1947Q2–2007Q4; Ramey (2011), 1947Q2–2013Q4; Kilian (2008), 1971Q1–2004Q3.

3.2 Methodology

Our principal 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 (2017) are examples of two recent papers which have used the local projections method to estimate state-dependent fiscal multipliers.

The basic NK model outlined in Section 2 makes predictions about the relationship between productivity shocks and output, and this relationship relies upon the behavior of expected inflation. To that end, our local projection of interest to test the effect of a productivity shock on output both at and away from the ZLB is:

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

This regression would be estimated separately for different leads of log output, $\ln Y_{t+h}$, $h \geq 0$. Z_t is a dummy variable which takes on a value of $Z_t = 0$ when the economy is away from the ZLB and $Z_t = 1$ when the economy is at the ZLB. The period t value of log productivity, $\ln A_t$, is included on the right hand side, reflecting an assumption that utilization-adjusted TFP is exogenous with respect to output within a period. This variable is interacted with both $(1 - Z_t)$ and Z_t. The right hand side also includes a constant and p lags of log output, log productivity, and a measure of expected inflation, π_t^e, all of which are also interacted with both $(1 - Z_t)$ and Z_t. The estimated impulse response of output at horizon h to a shock to productivity in period t is given by $(1 - Z_t)\widehat{\beta}^h + Z_t\widehat{\beta}_z^h$. When the ZLB is not binding, the response is therefore given by $\widehat{\beta}^h$. When the ZLB binds, it is $\widehat{\beta}_z^h$. We are interested in whether $\widehat{\beta}^h = \widehat{\beta}_z^h$.

While the local projections methodology can more easily accommodate non-linearities than conventional vector autoregressions (VARs), in practice impulse responses estimated by local projections are often very choppy (Ramey 2016). For this reason, for our baseline analysis in the paper we adopt the smooth local projections approach proposed by Barnichon and Brownlees (2016). Intuitively, the smooth local projections estimator starts with a local projection like that given in (10), but imposes that impulse responses are smooth functions of the forecast horizon. We provide an outline of the smooth local projections estimator in Appendix D. There we also consider a simple Monte Carlo simulation exercise to demonstrate the suitability of the smooth local projections approach.

We estimate the smooth local projection over the sample period 1984Q1 through 2017Q1. The beginning date is chosen to coincide with conventional dating of the "Great Moderation." The end of the sample is mandated by current data availability. We consider the ZLB to be binding from 2008Q4 through 2015Q4. This leaves 29 observation where the ZLB binds and 104 observations where it does not (99 observations prior to the beginning of the ZLB period and 5 after it). We estimate the smooth local projections with $p = 2$ lags. We estimate the projections at horizons from $h = 0, \ldots, H = 8$. Following Jordà (2005) and Barnichon and Brownlees (2016), confidence intervals are constructed using Newey and West (1987) HAC standard errors.

Output is measured as real GDP from the NIPA accounts and productivity is the aforementioned utilization-adjusted TFP series provided by Fernald (2014). A number of different measures of aggregate expected inflation exist, all with different pros and cons. A substantial empirical literature studies the properties of survey-based measures of expected inflation, for instance those collected by the Michigan Survey of Consumers and the Survey of Professional Forecasters. Still other work infers inflation expectations from differences in yields between real and nominal bonds. The Federal Reserve Bank of Cleveland employs a statistical model to measure expected inflation that relies both on survey-based measures and Treasury bond yields. As a benchmark measure, we use the median of the expected inflation series from the Michigan Survey of Consumers. The wording of the question used in the Michigan Survey is as follows. Survey respondents are asked "By about what percent do you expect prices to go (up/down) on the average, during the next 12 months?" Respondents then give a quantitative point forecast. Desirable features of the Michigan Survey are its relatively long time series, its high frequency, and that it probes actual economic decision-makers about their inflation expectations. Undesirable features are that it is ambiguous as to what measure of prices agents are forecasting and that the time frame over which the expectation is given is fixed at a year (whereas in the baseline NK model one quarter ahead expected inflation is what is relevant). These data are available at a monthly frequency going back to 1978. To match the quarterly frequency of the utilization-adjusted TFP series, we aggregate the expected inflation series to a quarterly frequency by averaging across months within a quarter.

3.3 Results

We graphically display the results from the baseline smooth local projection 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 90 percent confidence interval associated with this response. Output increases mildly on impact and continues to rise for several quarters thereafter. The units of the response can be interpreted as an elasticity. Hence, outside of the ZLB on impact we estimate that a one percent increase in productivity leads to a 0.2 percent increase in output. The solid red line with '*' markers plots the estimated response when the zero lower bound binds; the dotted red lines with the '*' markers correspond to the lower and upper bounds of the 90 percent confidence interval. Contrary to the theory outlined in Section 2, output is estimated to respond by roughly four times as much to a productivity shock on impact at the ZLB compared to normal times.

The point estimate for the output response at the ZLB is larger than away from the ZLB for about six periods, and the confidence intervals for the two responses do not overlap for the first four forecast horizons.

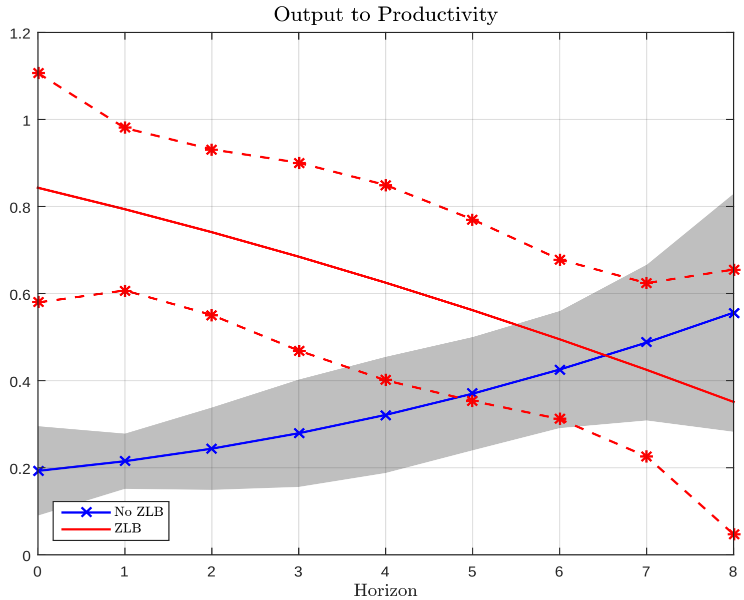

Figure 3: Results from Baseline SLP, Output

Notes: This figure shows the estimated impulse response of output to a one unit productivity shock at various horizons. The blue line with 'x' markers covers the case when the ZLB does not bind (i.e. when $Z_t = 0$). The shaded area bands represent the 90 percent confidence interval about the no ZLB case. The solid red line is when the ZLB binds or $Z_t = 1$. The dotted red lines with the '*' markers represent the 90 percent confidence interval about the ZLB case.

Table 3 shows point estimates and standard errors for the estimated responses away from the ZLB (β^h) and at the ZLB (β_z^h) for several different horizons. Standard errors are below the point estimates in parentheses. The third column presents test statistics of the null hypothesis of equality of the IRFs at each forecast horizon. p-values are presented below the test statistics in brackets. p-values are computed by comparing the test statistic to an F distribution with 1 numerator degree of freedom and $T - p - h - K$ denominator degrees of freedom, where T is the sample size and K is the number of regressors.[5] The null hypothesis of pointwise equality can be rejected at better than the 1 percent percent significance level or better for horizons $h = 0, 1, 2$ and at the 5 percent significance level for horizon $h = 3$. Beyond this horizon, one cannot reject that the point estimates are the same.

[5]For the standard local projection methodology, the maximum available sample size is used in each individual regression. This means that $T - p - h$ observations are available after accounting for lags (p) and leads of the left hand side variable (h).

Table 3: Tests of Equality of Output IRFs at and Away from ZLB

Horizon	β^h	β_z^h	H_0: $\beta^h = \beta_z^h$	H_0: Pathwise Equality
$h = 0$	0.193	0.843	8.738***	10.355***
	(0.08)	(0.20)	[0.00]	[0.00]
$h = 1$	0.215	0.794	14.248***	22.535***
	(0.05)	(0.15)	[0.00]	[0.00]
$h = 2$	0.244	0.741	9.025***	23.545***
	(0.07)	(0.15)	[0.00]	[0.00]
$h = 3$	0.279	0.685	4.389**	23.557***
	(0.10)	(0.17)	[0.04]	[0.00]
$h = 4$	0.321	0.625	2.231	23.721***
	(0.10)	(0.17)	[0.14]	[0.00]
$h = 5$	0.370	0.562	1.002	24.056***
	(0.10)	(0.16)	[0.32]	[0.00]
$h = 6$	0.426	0.495	0.154	24.084***
	(0.11)	(0.14)	[0.70]	[0.00]
$h = 7$	0.488	0.425	0.090	24.101***
	(0.14)	(0.15)	[0.76]	[0.00]
$h = 8$	0.556	0.351	0.415	24.091***
	(0.21)	(0.24)	[0.52]	[0.00]

Notes: This table shows the estimates of β^h and β_z^h at different forecast leads, as well as the corresponding standard errors in parentheses. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels for the pointwise and pathwise hypotheses tests in the right two columns, and p-values are presented in brackets.

The final column in Table 3 considers the null hypothesis of *pathwise* equality of the estimated impulse response functions. Rather than testing the equality of IRFs across regimes *at* a forecast horizon, here we test equality of IRFs across regimes *up to* a particular forecast horizon. For example, in the $h = 1$ row, the pointwise null hypothesis is $H_0: \ \beta^1 = \beta_z^1$, whereas the pathwise test involves a joint null hypothesis of $H_0: \ \beta^0 = \beta_z^0 \ \& \ \beta^1 = \beta_z^1$. The pathwise test requires computing covariances of estimators across different forecasting regressions. For this reason, in computing these tests we restrict the sample size to be constant across forecasting regressions instead of varying with the forecasting horizon (see the discussion in Footnote 5). In this system of seemingly unrelated regressions, the right hand side variables are identical. The pathwise test statistic would ordinarily be identical to the pointwise statistic at $h = 0$, but it differs slightly here because the sample sizes are different. Nevertheless, we can still reject pathwise equality on impact at better than the 1 percent significance level. Furthermore, we can easily reject pathwise equality at all subsequent horizons at greater than the 1 percent significance level.

Our baseline responses and statistical tests are based on the smooth local projections technique of Barnichon and Brownlees (2016). To be sure that the smoothing of the responses is not driving any of our results, in Figure 4 we show estimated responses of output to a productivity shock using the standard local projections technique. The figure is otherwise constructed and labeled similarly to Figure 3. The same basic picture emerges regardless of whether the responses are smoothed or not – output responds significantly more to a productivity shock on impact and at short forecast horizons when the ZLB binds than when it does not. Aside from the estimated responses being somewhat smoother, the main difference between Figures 3 and 4 is that the confidence intervals are tighter when using the smooth local projection as opposed to the standard local projection. This arises naturally because of efficiency gains associated with imposing more structure in the estimation (in particular, imposing the smoothness of responses across different horizons).

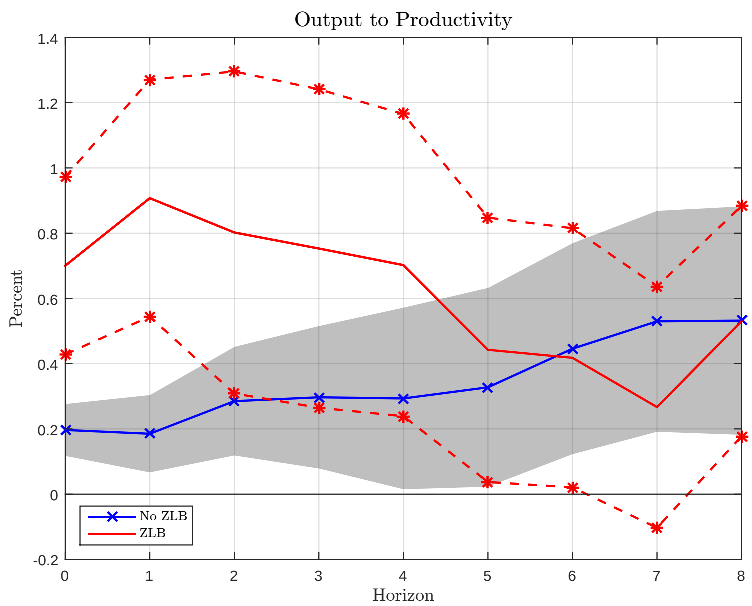

Figure 4: Results from Baseline LP, Output

Notes: This figure is similar to the one plotted in Figure 3, but estimates the responses using the conventional local projection as opposed to a smooth local projection.

In the basic NK model, the behavior of expected inflation is the key mechanism by which the output responses to a productivity shock differ at and away from the ZLB. To that end, we estimate a version of (10), but place expected inflation, rather than output, on the left hand side. This regression allows us to estimate the impulse respone of expected inflation to a productivity shock across monetary policy regimes. Formally:

$$\pi^e_{t+h} = (1 - Z_t)\left[\alpha^h_\pi + \beta^h_\pi \ln A_t + \sum_{s=1}^{p} \gamma^h_{s,\pi} \ln Y_{t-s} + \sum_{s=1}^{p} \phi^h_{s,\pi} \ln A_{t-s} + \sum_{s=1}^{p} \theta^h_{s,\pi} \pi^e_{t-s}\right]$$
$$+Z_t\left[\alpha^h_{z,\pi} + \beta^h_{z,\pi} \ln A_t + \sum_{s=1}^{p} \gamma^h_{z,s,\pi} \ln Y_{t-s} + \sum_{s=1}^{p} \phi^h_{z,s,\pi} \ln A_{t-s} + \sum_{s=1}^{p} \theta^h_{z,s,\pi} \pi^e_{t-s}\right] + u_{t+h}. \quad (11)$$

Similarly to the output projection, β^h_π measures the impulse response of expected inflation at horizon h to a productivity shock when the ZLB does not bind, while $\beta^h_{z,\pi}$ measures the response when the ZLB does bind. The specifics of this local projection are the same as when using output on the left hand side, and we focus on results using smooth local projections. Figure 5 plots the estimated response.

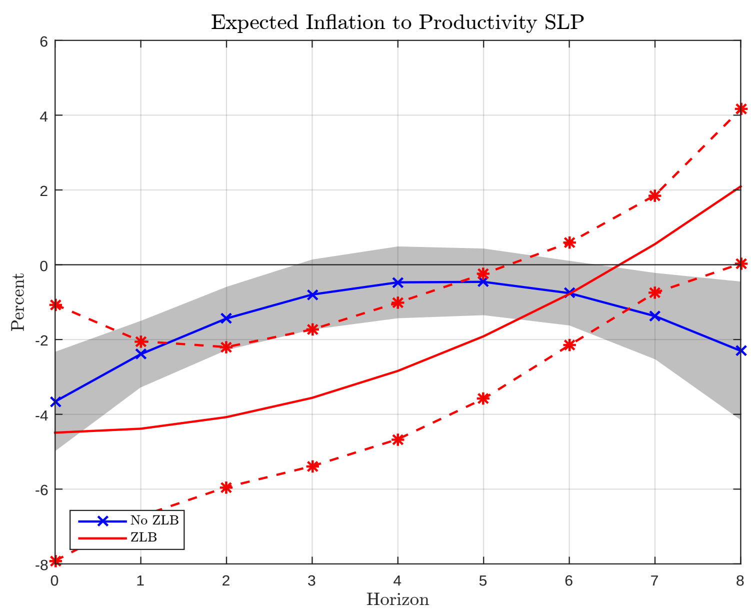

Figure 5: Results from Baseline SLP, Expected Inflation

Notes: This figure plots the estimated impulse response of expected inflation to a one unit productivity shock, both at and away from the ZLB. See the notes to Figure 3.

Outside of the ZLB period, expected inflation falls in response to a productivity shock, as one might expect. The expected inflation series is expressed at a quarerly percentage rate. Hence, one can interpret the units of the impulse response as suggesting that a one percent increase in productivity results in expected inflation falling by about 0.04 percentage points at a quarterly frequency, or roughly 0.16 percent annualized.[6] The response ceases to be

[6]The coefficients in the local projection, β^h_π and $\beta^h_{z,\pi}$, measure the response to a change in log productivity

statistically different from zero after about one year. At the ZLB (red line), we also estimate that expected inflation falls in response to a positive productivity shock. The point estimate on impact is slightly more negative than away from the ZLB, but this difference is neither economically nor statistically significant. For most subsequent horizons, the point estimate of the response at the ZLB lies below the response away from the ZLB, although these differences are not economically large and for the most part are statistically insignificant.

Table 4: Tests of Equality of Expected Inflation IRFs at and Away from ZLB

Horizon	β_{π}^{h}	$\beta_{z,\pi}^{h}$	H_0: $\beta_{\pi}^{h} = \beta_{z,\pi}^{h}$	H_0: Pathwise Equality
$h = 0$	-3.658	-4.385	0.047	0.049
	(1.05)	(2.66)	[0.77]	[0.83]
$h = 1$	-2.390	-4.385	1.047	2.850**
	(0.71)	(1.82)	[0.31]	[0.06]
$h = 2$	-1.4366	-4.074	2.708	4.852***
	(0.67)	(1.45)	[0.10]	[0.00]
$h = 3$	-0.797	-3.559	2.974*	4.863***
	(0.74)	(1.42)	[0.09]	[0.00]
$h = 4$	-0.470	-2.838	2.164	5.024***
	(0.76)	(1.42)	[0.14]	[0.00]
$h = 5$	-0.458	-1.913	0.971	5.944***
	(0.71)	(1.30)	[0.33]	[0.00]
$h = 6$	-0.759	-0.782	0.003	5.953***
	(0.67)	(1.06)	[0.98]	[0.00]
$h = 7$	-1.374	0.553	1.996	5.961***
	(0.91)	(1.01)	[0.16]	[0.00]
$h = 8$	-2.302	2.094	4.128**	7.016***
	(1.45)	(1.60)	[0.04]	[0.00]

Notes: This table shows estimates, standard errors, and test-statistics when expected inflation is the outcome variable. See also the notes to Table 3.

Table 4 is structured similarly to Table 3 but presents coefficients, standard errors, and test-statistics when expected inflation is on the left hand side. Although $\beta_{z,\pi}^{H}$ is smaller than

of 1 (or 100 percent). Hence, the effect of a 1 percent increase in productivity equals (1/100) times the units of the response on the vertical axis in Figure 5.

β^h at most horizons, pointwise equality can only be rejected at $h = 3$ and $h = 8$, at the 10 and 5 percent significance levels, respectively. Pathwise equality of the responses can be rejected over most horizons at conventional significance levels. This should not necessarily be surprising – the pathwise hypothesis test is a much more demanding test than the pointwise tests. For example, if there were no covariance between coefficient estimates at different leads, h, then the pathwise test statistic would be the sum of the pointwise test statistics. Our interpretation of the evidence presented in Figure 5 and Table 4 is that while there is some evidence to suggest that expected inflation falls more to a productivity shock at the ZLB in comparison to normal times, the evidence is not nearly as compelling as for the differences in the estimated output responses. Quite differently than for the case of output, the responses of expected inflation at and away from the ZLB appear qualitatively rather similar.

For the sake of completeness, we also show the expected inflation responses at and away from the ZLB in the case of a standard local projection. These are shown in Figure 6. Qualitatively these responses are very similar to Figure 5 (although the estimated responses are much choppier in the standard local projection with expected inflation than with output). Because the confidence bands are significantly wider with the local projection instead of the smooth local projection, it would not be possible to reject the pointwise equality of the estimated responses at and away from the ZLB at any horizon.

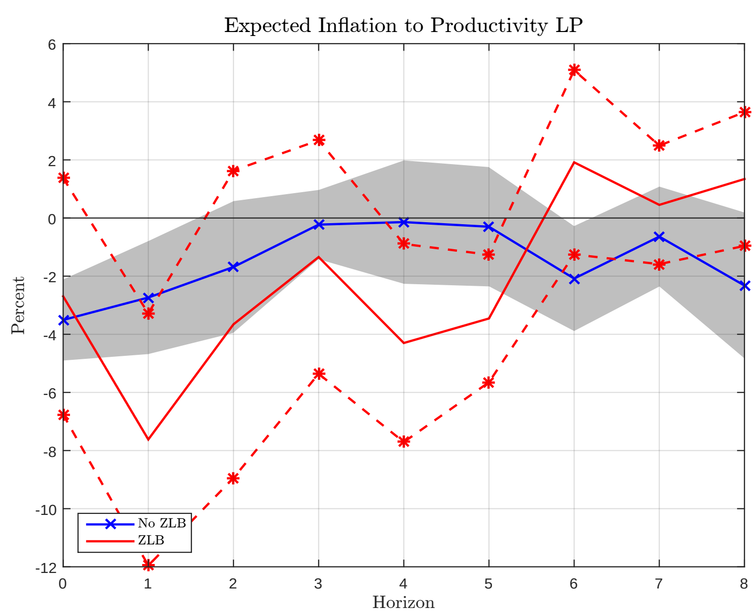

Figure 6: Results from Baseline LP, Expected Inflation

Notes: This figure is similar to the one plotted in Figure 5, but estimates the responses using the conventional local projections as opposed to smooth local projections.

3.4 Robustness

We find that output responds significantly more to a productivity shock when monetary policy is constrained by the ZLB than otherwise. Further, the response of expected inflation to a productivity shock is qualitatively similar across monetary regimes (although point estimates are generally more negative at the ZLB). These results stands in contrast to the predictions of a textbook New Keynesian model. In this subsection we explore the robustness of these results along several different dimensions.

We begin by considering alternative measures of output or economic activity more generally. As alternative measures to real GDP, we use the Industrial Production (IP) index, total hours per capita in the non-farm business sector, and the civilian unemployment rate. For the IP and unemployment series, the underlying monthly data are converted to quarterly by averaging observations within the quarter. The sample period for the estimation is the same as our baseline case. Results are displayed in Figure 7.

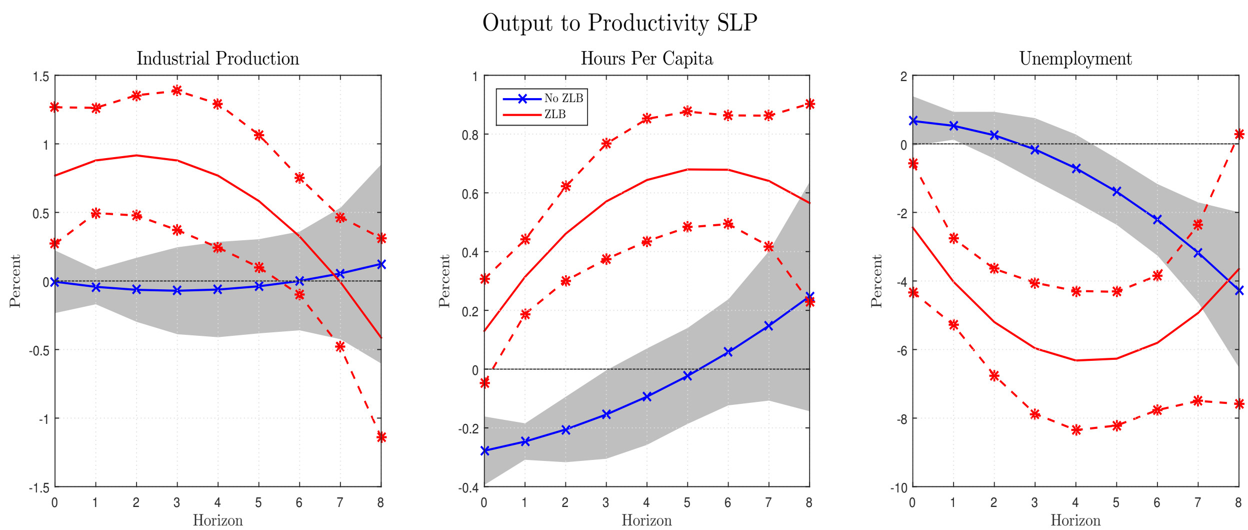

Figure 7: Alternative Measures of Economic Activity

Notes: This figure plots the estimated impulse responses of economic activity to a productivity shock for three alternative measures of economic activity – log industrial production (left panel), log hours per capita (middle panel), and the log unemployment rate (right panel). See the notes to Figure 3.

The responses depicted in Figure 7 convey a similar message to our baseline results using real GDP. In particular, a positive productivity shock is estimated to be more expansionary on impact and for several quarters thereafter when the ZLB binds compared to when it does not. Somewhat surprisingly, we estimate no economically or statistically significant response of IP to a productivity shock away from the ZLB. But when the ZLB binds, we find that IP increases significantly on impact and for several quarters thereafter before reverting back towards zero.

Because a large literature has investigated the effects of technology shocks on labor market variables (e.g. Shapiro and Watson 1988; Blanchard and Quah 1989; Gali 1999; Christiano, Eichenbaum, and Vigfusson 2004; Basu, Fernald, and Kimball 2006), it is of particular interest to examine how such variables react to a productivity shock both inside and away from the ZLB. Consistent with much of this literature, we find that hours worked initially declines after a productivity improvement outside of the ZLB before turning positive after a little more than a year. The estimated impact decline in hours is statistically significant. Similarly, outside of the ZLB we find that the unemployment rate increases on impact and for a few quarters thereafter when productivity increases, although this response is not statistically significant. In other words, away from the ZLB a productivity shock seems to be contractionary for labor market variables. When the ZLB binds, in contrast, we find that hours worked increases on impact, albeit statistically insignificantly so. The response continues to grow and quickly turns significant. Likewise, we estimate that the unemployment rate declines (significantly) on impact when productivity increases.

We next consider alternatives measures of expected inflation in our local projections. Our baseline measure of expected inflation is the median point estimate from the Michigan Survey of Consumers. We also consider the mean (across households within a period) from the Michigan Survey, the one year ahead inflation expectation from the Federal Reserve Bank of Cleveland, the median SPF forecast for one period ahead CPI inflation, and the median of one period ahead GDP deflator inflation from the SPF.[7] Impulse responses obtained from smooth local projections using these alternative series are shown in Figure 8.

Estimated impulse responses of all four alternative measures of expected inflation are qualitatively similar to our baseline analysis making use of the median expectation from the Michigan Survey of Consumers. In normal times (blue lines with 'x' markers) expected inflation falls on impact and reverts back to zero after a few horizons. Estimated responses at the ZLB are insignificantly different from the responses during normal times for all four series over most horizons. Point estimates suggest that expected inflation falls more at the ZLB for the first few quarters after a productivity shock when using the mean from the Michigan Survey and the Cleveland Fed survey. When using either SPF forecast, in contrast, we find that expected inflation falls less on impact at the ZLB in comparison to normal times, though we must stress that these differences are statistically insignificant.

[7]Details on the Cleveland Fed Survey can be found here. The survey provides forecasts of CPI inflation at forecast horizons of one through thirty years and is available monthly beginning in 1982. We focus on the one year inflation expectation. For the SPF data, we can access one quarter and one year ahead forecasts of both CPI and GDP deflator inflation. We focus on one quarter ahead, as this frequency most closely matches the theoretical model of Section 2. Our results are nearly identical when using one year ahead forecasts.

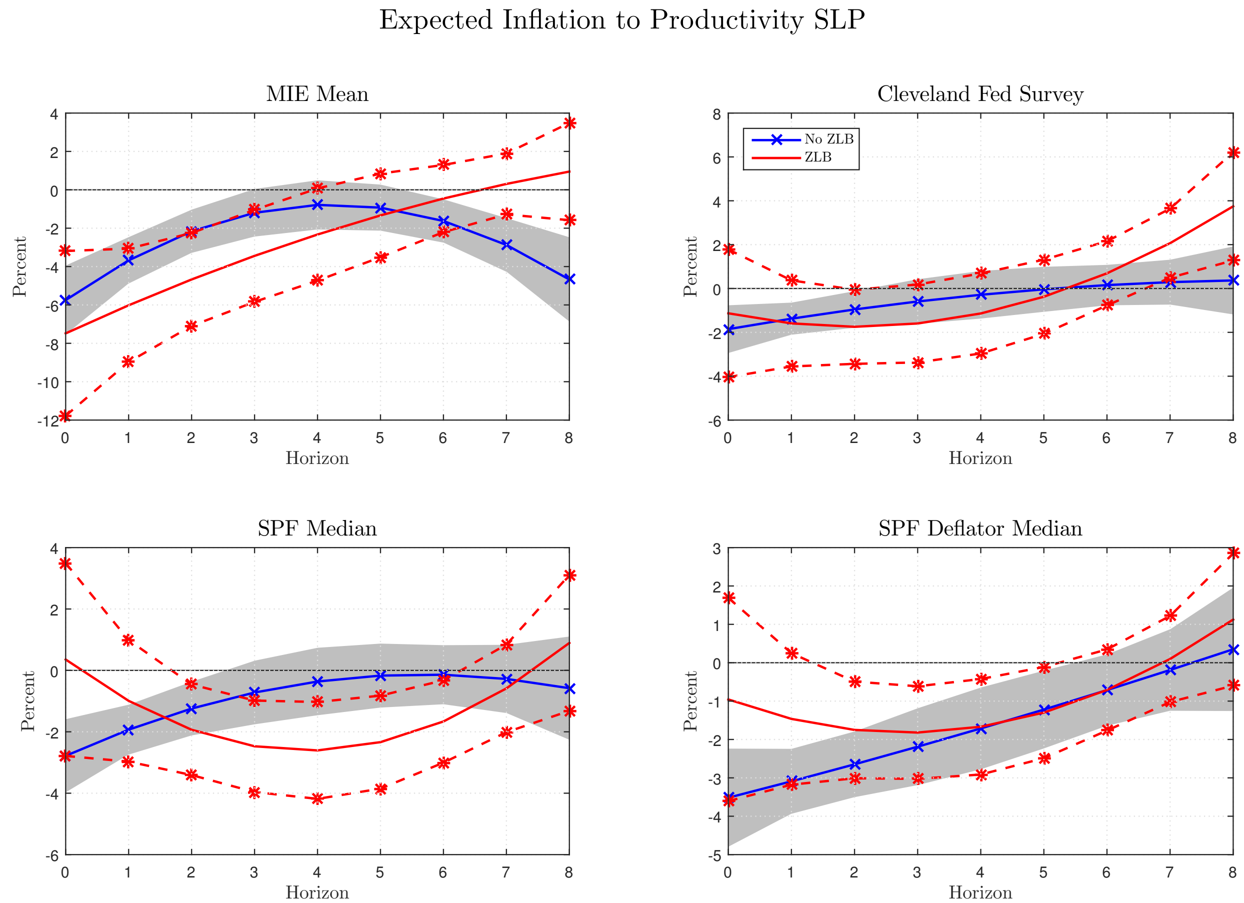

Figure 8: Alternative Measures of Expected Inflation

Notes: This figure plots the estimated impulse responses of expected inflation to a productivity shock using alternative measures of expected inflation – the mean expectation from the Michigan Survey of Consumers (upper left plot), the Cleveland's Fed expected inflation series (upper right), the median one quarter ahead SPF forecast of CPI inflation (lower left), and the median SPF forecast of GDP deflator inflation (lower right). See also the notes to Figure 5.

We next consider a number of additional robustness checks, focusing first on output responses. We revert to assuming that the outcome variable of interest is real GDP, but consider several different specifications of the local projection. Results are shown graphically in Figure 9.

One might worry that our result that output responds 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 is conceivable that output responds more to a productivity shock during a recession, and that our ZLB dummy variable is simply proxying for periods of recession. We address this concern by augmenting our baseline regression specification, (10), with a dummy variable equaling one in periods defined by the NBER to be recessions (and zero otherwise) as well as an interaction between this dummy variable and the period t value of productivity. Responses under this specification are shown in the left panel of the upper row of Figure 9. This figure is qualitatively similar to our baseline specification, and if anything the differences between the estimated output responses at and away from the ZLB are even larger and more statistically significant than in our baseline specification.

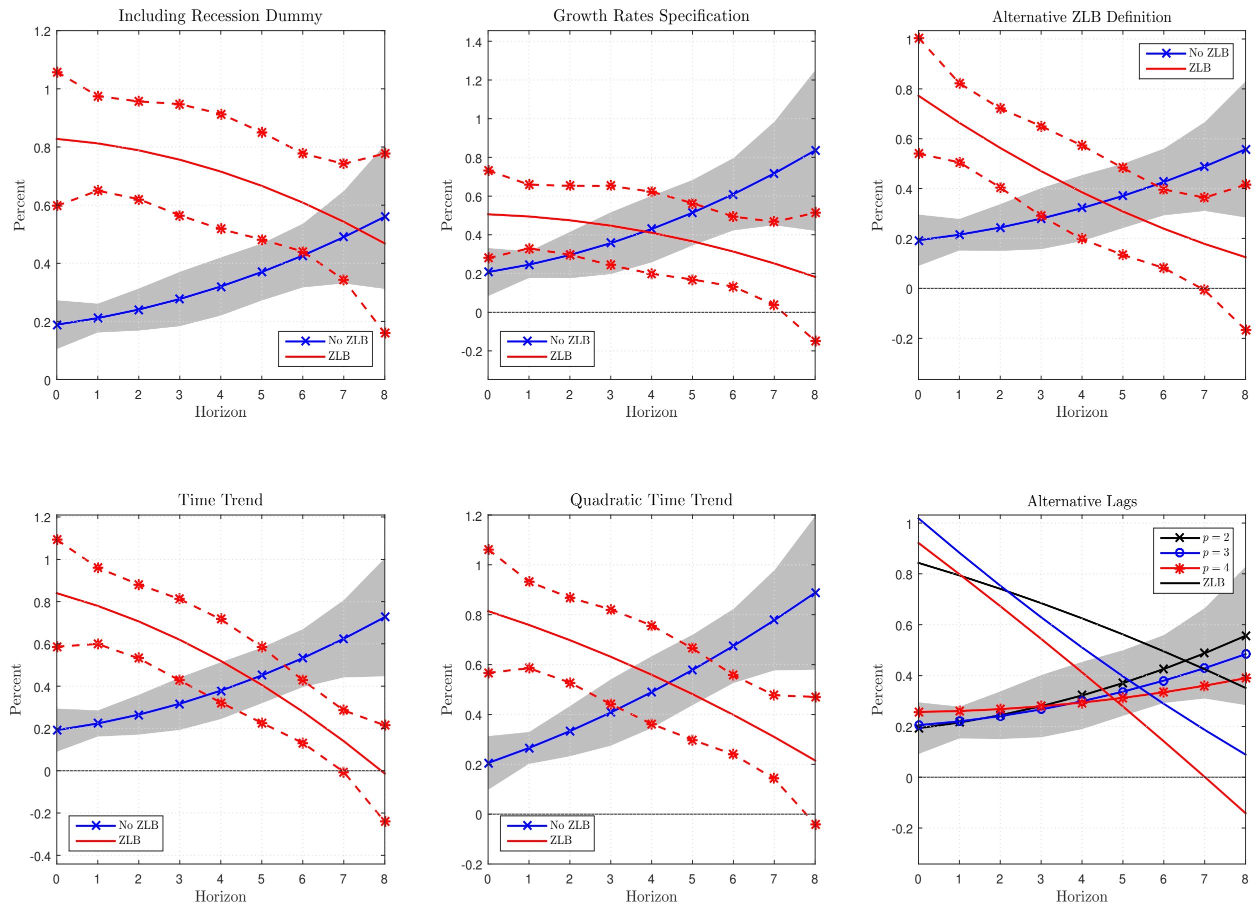

Figure 9: Alternative Specifications, Output

Notes: This figure plots estimated impulse responses of output to a productivity shock in various different specifications of our smooth local projection. See the notes to Figure 3. To improve readability, when considering alternative lag lengths (lower right plot) we only show the confidence interval around the estimated response away from the ZLB in the baseline specification with $p = 2$ lags.

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 output and productivity appear in log first differences rather than log levels (the expected inflation series continues to enter in levels). For example, the regressor of interest on the right hand side becomes $\Delta \ln A_t$ instead of $\ln A_t$, and lags of both productivity and output appear in first differences as well. The left hand side variable is $\ln Y_{t+h} - \ln Y_{t-1}$, which is equivalent to the cumulative sum of growth rates of output from horizons 0 through h.[8] The estimated responses are shown in the upper middle panel of the first row of Figure

[8] An alternative would be to use the period over period growth rate of output on the left hand side, i.e.

9. The estimated response outside of the ZLB is virtually identical to our baseline result from estimation in levels. The estimated impact effect and the response for several quarters thereafter is again larger at the ZLB than in normal times, albeit the differences between responses at and away from the ZLB are not as statistically significant as in the levels specification.

In the left and middle panels of the bottom row we revert to including the variables in levels, but include a deterministic linear time trend (left panel) and a quadratic time trend (middle panel) in the regressions. There is virtually no effect on the shape of the estimated response of output in normal times. Point estimates on impact and for a few quarters thereafter of the output response at the ZLB are similar to our baseline specification and the differences between the responses at and away from the ZLB are highly statistically significant. One noticeable difference, apparent in most of the figures where the left hand side variable is output but especially so in these two, is that the response at the ZLB, while larger on impact, seems to be less persistent in comparison to normal times. We will return to a discussion of this feature in Section 3.5 below.

In our estimation we suppose that the ZLB binds from the final quarter of 2008 to the final quarter of 2015. While the effective Federal Funds Rate was below one hundred basis points for virtually all of the fourth quarter of 2008, the target range for the Funds Rate did not formally hit zero until December. Conversely, while the target Funds Rate included zero for most of the fourth quarter of 2015, a movement to a target range of twenty-five to fifty basis points took place near the end of the quarter. The right panel of the upper row examines the sensitivity of our analysis to a slight modification of our ZLB definition. In particular, we suppose that $Z_t = 1/3$ for the final quarter of 2008 (equal to the number of months as a fraction of the total months in the quarter were the ZLB was binding) and $Z_t = 2/3$ for the final quarter of 2015. Output still responds substantially more on impact and for a few quarters at the ZLB compared to normal times. The different definition of the ZLB does result in the output response being somewhat less persistent than in our baseline analysis.

The final panel of the bottom row of Figure 9 considers robustness to the assumed number of lags of the control variables in the local projections. The number of lags has a limited effect on the estimated response of output to a productivity shock during normal times. The estimated response in the ZLB regime on impact and for a few quarters thereafter is similar across different lag lengths. The estimated responses in the ZLB regime do seem to be less persistent the more lags are included in the projection.

$\Delta \ln Y_{t+h} = \ln Y_{t+h} - \ln Y_{t+h-1}$, and take the cumulative sum of estimated growth rate responses up to horizon h to recover the level response. Our specification instead directly estimates the level response.

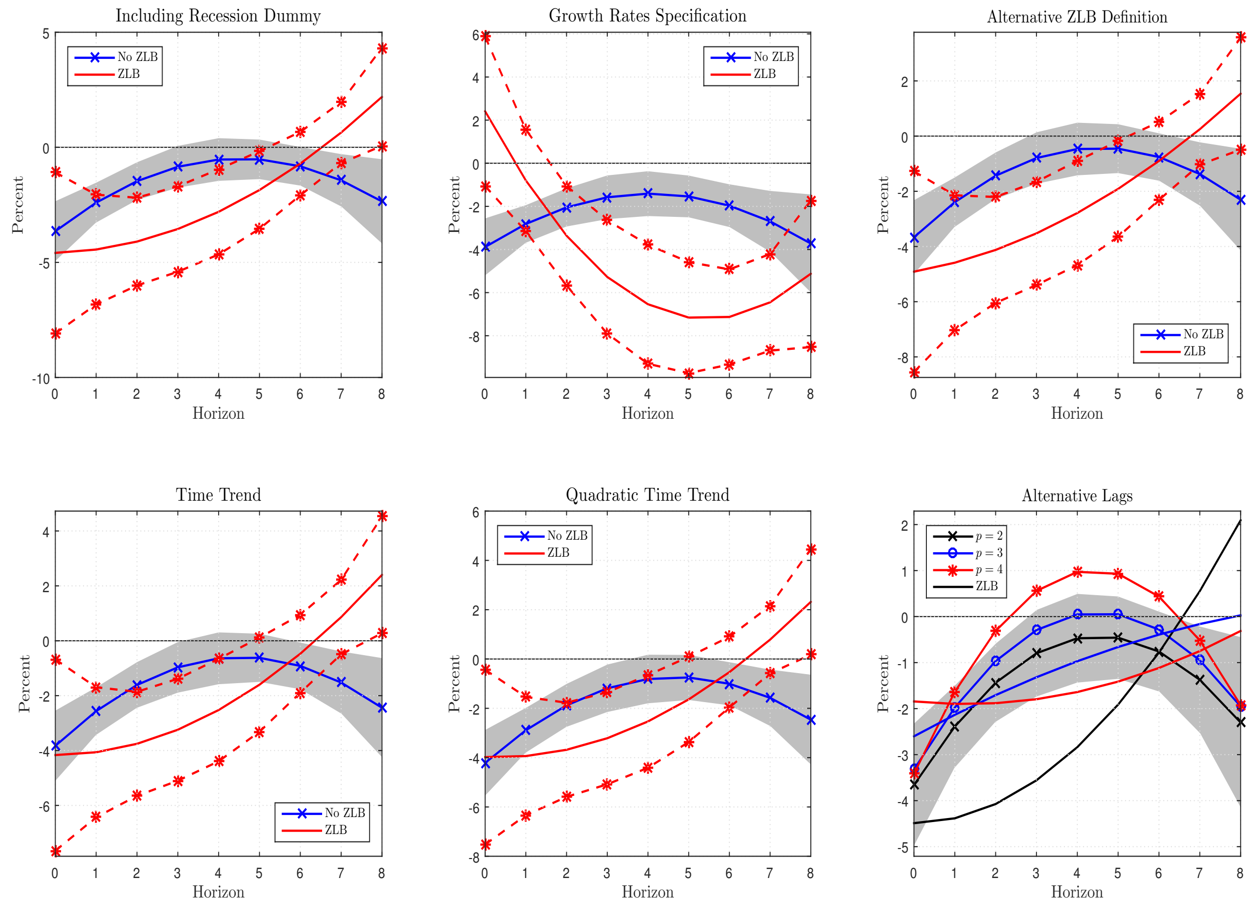

Figure 10: Alternative Specifications, Expected Inflation

Notes: This figure plots estimated impulse responses of expected inflation to a productivity shock in various different specifications of our smooth local projection. See the notes to Figure 3. To improve readability, when considering alternative lag lengths (lower right plot) we only show the confidence interval around the estimated response away from the ZLB in the baseline specification with $p = 2$ lags. For the responses generated in this figure, we use the median expectation from the Michigan Survey of Consumers as our inflation expectations series.

We also consider the same set of robustness checks when expected inflation is the left hand side variable of interest. For these robustness checks, we use the median of the Michigan Survey of Consumers as our inflation expectations series. Impulse responses are shown below in Figure 10. The figure is structured similarly to Figure 9. In most specifications our results are very similar to the baseline results depicted in Figure 5. In all specifications, expected inflation falls when productivity increases outside of the ZLB. For most specifications the estimated response when the ZLB binds is qualitatively similar, with expected inflation falling by slightly more on impact and the first several horizons thereafter but being somewhat less persistent. An exception is the growth rates specification, where output and utilization-adjusted TFP appear in first differences on the right hand side of our local projection. In this specification the point estimate suggests that expected inflation increases slightly on impact when productivity increases, albeit this is not statistically significant. Furthermore, after a couple of horizons the estimated expected inflation response at the ZLB lies significantly below the estimated response when the ZLB does not bind.

3.5 Is the Productivity Response to a Productivity Shock the Same at the ZLB?

We have documented that output, and a variety of similar measures, seems to react significantly more positively to a positive productivity shock when the ZLB binds than when it does not. In contrast, we find little qualitative difference in the response of expected inflation to a productivity shock at and away from the ZLB. These results are inconsistent with the predictions of a textbook NK model, the simplest version of which predicts that output ought to increase less (or decrease) and expected inflation ought to fall more when productivity increases at the ZLB in comparison to a more active monetary policy regime.

In standard forward-looking macroeconomic models, how output and expected inflation react to a productivity (or any other) exogenous shock depends not only on the stance of monetary policy, but also on the persistence of the shock itself. In our state-dependent local projection model, one would hope that the persistence of the productivity shock is similar across monetary regimes. Nothing in our estimation, however, imposes this. Indeed, a close inspection of the estimated responses of output reveals that the persistence of the productivity shock may in fact not be the same across regimes. Whereas we find that the estimated output response grows with the forecast horizon in normal times, at the ZLB the response of output seems to be declining with the forecast horizon. This feature is somewhat evident in our baseline specification (see Figure 3), but is perhaps even more so in several of the robustness checks considered in Subsection 3.4 (see, e.g., Figure 9).

We therefore examine whether the dynamic paths of productivity in response to a productivity shock are estimated to be similar across monetary regimes. We estimate a smooth local projections version of (10), but with $\ln A_{t+h}$ on the left hand side in place of leads of output. Results are shown in Figure 11 below:

It is visually apparent that the persistence of the productivity impulse response is quite different across monetary regimes. Whereas the estimated response away from the ZLB is more or less consistent with a random walk process, at the ZLB the productivity response is quickly mean-reverting. This finding does not emerge as a consequence of smoothing and seems to be robust to several different versions of our local projection.

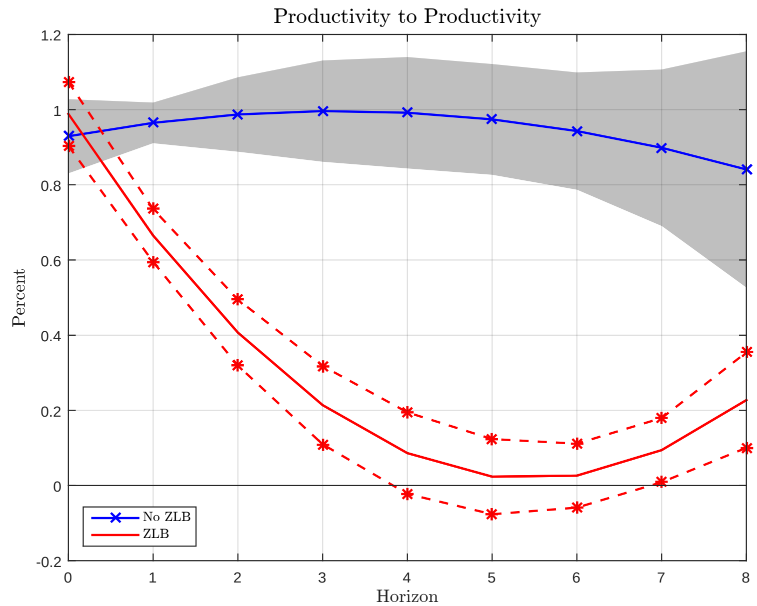

Figure 11: Baseline SLP, Productivity

Notes: This figure plots the estimated impulse response of productivity to a one unit productivity shock at various forecast horizons. See also the notes to Figure 3.

At first pass the results visually conveyed in Figure 11 would seem to call into question the validity of our empirical analysis. The simple theoretical model giving rise to testable empirical predictions holds everything in the model fixed but the stance of monetary policy. If, in addition to monetary policy being constrained by the ZLB, the persistence of productivity shocks has changed, is there any merit to our analysis? We believe that there is, and argue that in fact the dampened persistence of the productivity shock during the ZLB regime provides an even more compelling test than would be the case had there been no change in the persistence of the shock.

In a textbook NK model, a binding ZLB has much stronger impacts on the responses of output and expected inflation to a productivity shock the less persistent that shock is. We touch on this point in Appendix C when discussing the medium scale NK model, and a similar point is also emphasized in Wieland (2015). Here in the body of the paper we make this point referencing the textbook NK model developed in Section 2. We compute impulse responses of output and expected inflation to a productivity shock with different peg lengths ($H = 0$, $H = 3$, $H = 6$, and $H = 10$). We do so for two different levels of persistence of the productivity shock, $\rho_a = 0.97$ and $\rho_a = 0.75$. Responses are plotted in Figure 12. The left column plots the output (upper row) and expected inflation (bottom row) responses when $\rho_a = 0.97$, and the right column does the same when $\rho_a = 0.75$.

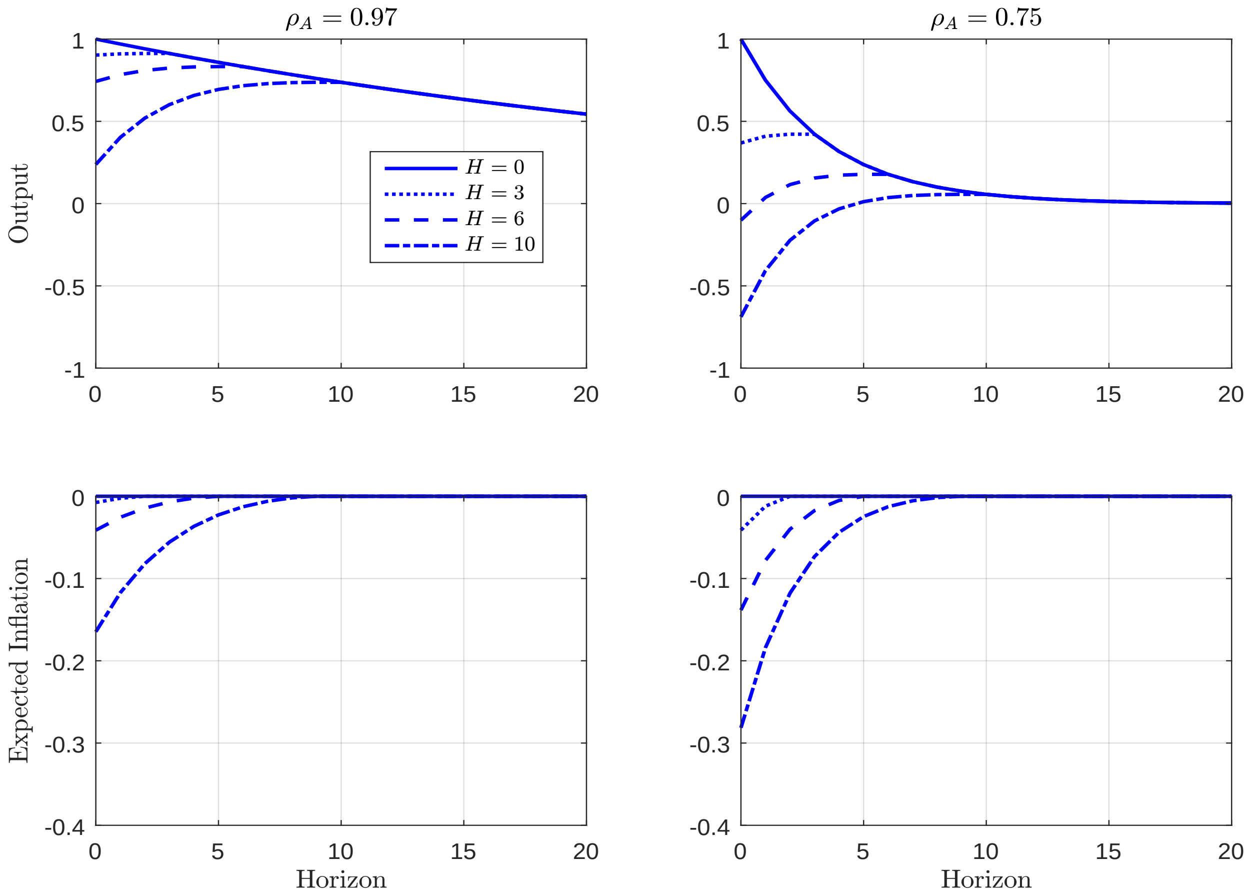

Figure 12: Output and Expected Inflation Responses to a Productivity Shock as a Function of Duration of ZLB and Persistence of Shock

Notes: This figure plots the impulse responses of output and one period ahead expected inflation to a positive one unit productivity shock in the textbook NK model of Section 2. Different marked lines correspond to different periods of an interest rate peg as indicated in the legend. The left column considers an autoregressive parameter in the productivity process of $\rho_a = 0.97$, while the right column shows responses to a less persistent productivity shock with $\rho_a = 0.75$.

As documented in Section 2, regardless of the value of ρ_a, output reacts less to the productivity shock the longer is the duration of the interest rate peg. What is noteworthy in Figure 12 is that the effect of the peg on the output response is substantially larger when the productivity shock is less persistent. For example, when the interest rate is pegged for six periods, output jumps up by about 0.75 percent on impact when $\rho_a = 0.97$, but actually falls slightly on impact when $\rho_a = 0.75$. Similarly, when $H = 10$, output increases by 0.25 percent when $\rho_a = 0.97$, but declines by 0.6 percent on impact when $\rho_a = 0.75$. We observe the same pattern in the expected inflation responses plotted in the lower panel of Figure 12. In particular, for a given peg length, expected inflation falls more the less persistent is the productivity shock. The intuition for these findings is fairly straightforward. With forward-looking agents, when the productivity shock is very persistent there is a fairly large increase in demand. With a given increase in supply, this results in a reasonably-sized increase in output and only a small decline in current and expected inflation, which means that there is little resulting change in the real interest rate even though the nominal rate is fixed. When the shock is not very persistent, in contrast, aggregate demand increases very little in response to a productivity shock. With a given an increase in supply, current and expected inflation must fall more substantially, which drives up the real interest rate and depresses demand, potentially resulting in a decline in output.

Viewed through the lens of the NK model, then, a less persistent productivity shock in conjunction with a binding ZLB ought to result in an even smaller increase (or bigger decline) in output and a larger decline in expected inflation than if the persistence of the shock were held fixed across monetary regimes. Put another way, our findings that output responds more to a productivity shock at the ZLB, and that the response of expected inflation is roughly the same as in normal times, are even *more* inconsistent with the predictions of the textbook NK model when the persistence of the productivity shock is lower. Rather than invalidating our empirical results, the fact that the persistence of productivity shocks has evidently been lower in the ZLB period compared to earlier times actually makes our empirical results even more puzzling from the perspective of the textbook NK model.

4 Discussion

The empirical evidence presented and discussed in Section 3 is inconsistent with the predictions of a textbook NK model, as well as with its more complicated cousin, the medium scale DSGE model á la Christiano, Eichenbaum, and Evans (2005) or Smets and Wouters (2007). We view our paper as contributing to a growing literature which points out predictions of these models which do not seem to hold up to empirical scrutiny. While our objective is not to propose a new theoretical framework, we use this section to discuss relevant papers which, in our view, offer promising alterations to the basic NK framework which might help make the model more consistent with our empirical findings.

One avenue which we find particularly promising is to replace the assumption of sticky prices and/or wages with sticky or imperfect information, for example as proposed in Mankiw and Reis (2003). In their model, firms are free to set prices each period and do so optimally given available information, but can only update their information sets sporadically. This setup implies that the current inflation rate depends on past expectations of current economic conditions, which contrasts with the textbook NK model wherein inflation is purely forward-looking.[9] The extreme forward-looking nature of inflation lies at the heart of the basic

[9] Sticky price models can, of course, be tweaked to make inflation less forward-looking. But these tweaks, a very common one of which is the assumption of indexation to lagged inflation, are often both theoretically

model's prediction that positive supply shocks might be contractionary at the ZLB. Indeed, Kiley (2016) shows that positive technology shocks are expansionary at the ZLB in a sticky information economy.[10] Empirical evidence in support of models of informational rigidities more generally is provided in, for example, Coibion and Gorodnichenko (2012), Coibion and Gorodnichenko (2015), and Coibion, Gorodnichenko, and Kamdar (2017).

A related avenue that may prove helpful in reconciling the basic model with the evidence is to discard the assumption of rational expectations altogether. Gabaix (2016) replaces rational expectations with bounded rationality and shows that this can resolve a number of anomalies in the NK model. Angeletos and Lian (2016) show that relaxing the assumption of common knowledge makes expectations of economic outcomes (such as inflation) stickier relative to the baseline NK model. Along similar lines, García-Schmidt and Woodford (2015) use a model in which agents iteratively update their expectations to study the robustness of the so-called "Neo-Fisherian" predictions of the baseline NK model. In the context of our empirical results, these and related departures from the rational expectations benchmark will serve to make inflation expectations stickier (in a way consistent with much of our empirical analysis making use of inflation expectations data), thereby dampening the "expected inflation channel" which lies at the heart of many of the NK model's paradoxical predictions at the ZLB.

Another promising twist of the model is to assume some type of financial friction so as to make some fraction of agents borrowing constrained or to incorporate ex-post heterogeneity via market incompleteness. The introduction of borrowing constraints is discussed for example in Wieland (2015). For financially constrained agents, total spending cannot exceed some exogenous nominal debt limit. When productivity increases there are offsetting effects at the ZLB. On the one hand expected inflation decreases, which pushes current output down given a fixed nominal interest rate. On the other hand, lower inflation raises the debt limit of the constrained household, thereby allowing them to increase spending which pushes current output up. Wieland (2015) shows that for a sufficiently low rate of intertemporal substitution, the second channel dominates the first and output can rise when productivity increases. Recent examples of incorporating non-trivial heterogeneity and market incompleteness into the NK model include Kaplan, Moll, and Violante (2016) and McKay, Nakamura, and Steinsson (2016). Whereas the textbook NK model operates almost entirely through intertemporal substitution, versions of the model with heterogeneous agents and liquidity constraints work much more in-line with "old Keynesian" intuition. In particular, a higher than average fraction of constrained households during the ZLB period might help rationalize why we

and empirically unattractive.

[10] An additional desirable feature of sticky information is that the "sign reversals" associated with interest rate pegs, which we discuss in Appendix A and which Carlstrom, Fuerst, and Paustian (2015) call "pathological," are not present when monetary non-neutrality arises from sticky information instead of sticky prices.

empirically find that output reacts more strongly to a productivity shock at the ZLB.

Finally, one could argue that the zero lower bound has not in actuality imposed much of a constraint on monetary policy. While central banks around the world have been constrained in setting short term policy rates, they have aggressively moved along other margins in an to attempt to influence economic outcomes. One such margin involves promises about future policy. Bundick (2015) shows that if monetary policy follows a "history dependent" rule the perverse effects of productivity shocks at the ZLB can be mitigated. Even though a central bank cannot adjust short term nominal rates in response to a productivity improvement, it can promise to lower future short term rates in such a way as to stimulated expected inflation instead of allowing it decline. Another possibility is to allow for a more intricate central bank balance sheet so as to analyze the term and risk structure of interest rates. Along these lines, Chen, Cúrdia, and Ferrero (2012) build a model with debt instruments of different maturities and segmented markets to study the implications of large scale asset purchases. While in their theoretical model quantitative easing has limited effects, in a more empirically-oriented approach Wu and Xia (2016) argue that the Federal Reserve's unconventional policies during the ZLB period were even more aggressive than historical policies prior to the Great Recession. Relatedly, Wu and Zhang (2016) formalize the notion of a "shadow rate" for the Federal Funds Rate as a summary statistic for unconventional monetary policies. They argue that the ZLB is not a constraint on monetary policy, and in their model productivity shocks are expansionary even when the notional ZLB binds. The potential for effective (or, in the case of Wu and Xia 2016, hyper-effective) unconventional monetary policies at the ZLB could go a long way in helping to square our results with the predictions of the textbook NK model.

We think all of the different lines of research discussed in this section are promising. The key thread running through our discussion in this section is that it is possible to perturb the assumptions of a standard NK model in such a way as to deliver quite different results when the ZLB binds. However, the model we discuss in Section 2 and its medium-scale cousin are quite prevalent in policy institutions. Our discussion in this section suggests that considering some of these extensions/modifications would be worthwhile in light of our empirical results.

5 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 expected inflation, which drives up the real interest rate when the nominal rate is constrained by the ZLB. The higher real interest rate chokes off demand and limits the output response to a positive supply shock.

We test these predictions of the textbook NK model in the data. We estimate a state-dependent empirical model according to the Barnichon and Brownlees (2016) smoothed version of Jordà (2005)'s local projection method. We measure productivity shocks using Fernald (2014)'s utilization-adjusted total factor productivity series. Contrary to the predictions of the model, we find that output and related measures respond significantly more to positive productivity shocks in periods where the ZLB binds compared to periods where it does not. Further, we estimate that there is little significant difference in the response of a variety of measures of expected inflation to a productivity shock across monetary regimes.

As our empirical results are inconsistent with the predictions of the textbook model, caution seems to be in order when using the model to make predictions about the economic consequences of alternative policies (such as forward guidance or fiscal stimulus) when the ZLB binds. In contrast, more research into alternative model specifications such as those discussed in Section 4 seems desirable and likely to be fruitful.

References

ANGELETOS, G.-M. AND C. LIAN (2016): "Forward Guidance without Common Knowledge," Working Paper 22785, National Bureau of Economic Research.

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 Econonomic Journal: Economic Policy*, 7, 1–35.

BARNICHON, R. AND C. BROWNLEES (2016): "Impulse Response Estimation by Smooth Local Projections," Discussion Papers DP11726, CEPR Discussion Paper Series.

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.

BLANCHARD, O. J. AND D. QUAH (1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review*, 79, 655–673.

BUNDICK, B. (2015): "Real Fluctuations at the Zero Lower Bound," Tech. rep.

BURKE, M. AND A. OZDAGLI (2013): "Household Inflation Expectatons 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.

CHEN, H., V. CÚRDIA, AND A. FERRERO (2012): "The Macroeconomi Effects of Large-Scale Asset Purchase Programmes," *Economic Journal*, 122, F289–F315.

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.

Christiano, L. J., M. Eichenbaum, and R. Vigfusson (2004): "The Response of Hours to a Technology Shock: Evidence Based on Direct Measures of Technology," *Journal of the European Economic Association*, 2, 381–395.

Cohen-Setton, J., J. K. Hausman, and J. F. Wieland (2017): "Supply-Side Policies in the Depression: Evidence from France," *Journal of Money, Credit and Banking*, 49, 273–317.

Coibion, O. and Y. Gorodnichenko (2012): "What Can Survey Forecasts Tell Us About Information Rigidities?" *Journal of Political Economy*, 120, 116–159.

——— (2015): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 105, 2644–2678.

Coibion, O., Y. Gorodnichenko, and R. Kamdar (2017): "The Formation of Expectations, Inflation and the Phillips Curve," *Journal of Economic Literature*, forthcoming.

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," *Euoprean 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.

Eilers, P. H. and B. D. Marx (1996): "Flexible Smoothing with B-Splines and Penalties," *Statistical Science*, 11, 89–121.

Fernald, J. (2014): "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," Tech. Rep. 2012-19, Federal Reserve Bank of San Francisco.

Gabaix, X. (2016): "A Behavioral New Keynesian Model," Working Paper 22954, National Bureau of Economic Research.

Gali, J. (1999): "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, 89, 249–271.

García-Schmidt, M. and M. Woodford (2015): "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis," Working Paper 21614, National Bureau of Economic Research.

Ichiue, H. and S. Nishiguchi (2015): "Inflation Expectations and 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.

Kiley, M. T. (2016): "Policy paradoxes in the New Keynesian model," *Review of Economic Dynamics*, 21, 1 – 15.

Kilian, L. (2008): "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the US Economy?" *Review of Economics and Statistics*, 90, 216–240.

Laseen, S. and L. E. Svensson (2011): "Anticipated Alternative Instrument-Rate Paths in Policy Simulations," *International Journal of Central Banking*, 7, 1–35.

Mankiw, N. G. and R. Reis (2003): "Sticky Information: A Model of Monetary Nonneutrality and Structural Slumps," in *Knowledge Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, ed. by J. S. Phillipe Aghion, Romain Frydman and M. Woodford, Princeton, NJ: Princeton University Press, 64–86.

McKay, A., E. Nakamura, and J. Steinsson (2016): "The Power of Forward Guidance Revisited," *American Economic Review*, 106, 3133–3158.

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.

Newey, 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.

——— (2016): "Macroeconomic Shocks and Their Propagation," in *Handbook of Macroeconomics, Volume 2*, ed. by H. Uhlig and J. B. Taylor, Radarweg 29, PO Box 211, 1000 AE Amsterdam, The Netherlands: North-Holland, chap. 2, 71–162.

Ramey, V. A. and S. Zubairy (2017): "Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data," *Journal of Political Economy*, Forthcoming.

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.

Shapiro, M. and M. Watson (1988): "Sources of Business Cycles Fluctuations," in *NBER Macroeconomics Annual 1988, Volume 3*, National Bureau of Economic Research, Inc, 111–156.

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.

Wu, J. C. and F. D. Xia (2016): "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit and Banking*, 48, 253–291.

Wu, J. C. and J. Zhang (2016): "A Shadow Rate New Keynesian Model," Working Paper 22856, National Bureau of Economic Research.

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, Eichenbaum, and Rebelo (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 $\mathbb{E}_t i_{t+1} = \mathbb{E}_t r^f_{t+1}$. 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^f_t 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 $x_t = y_t - y^f_t$, and $y^f_t = \frac{1+\chi}{\sigma+\chi}a_t$, 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 $y_t = y^f_t$. 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, output expands less 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.

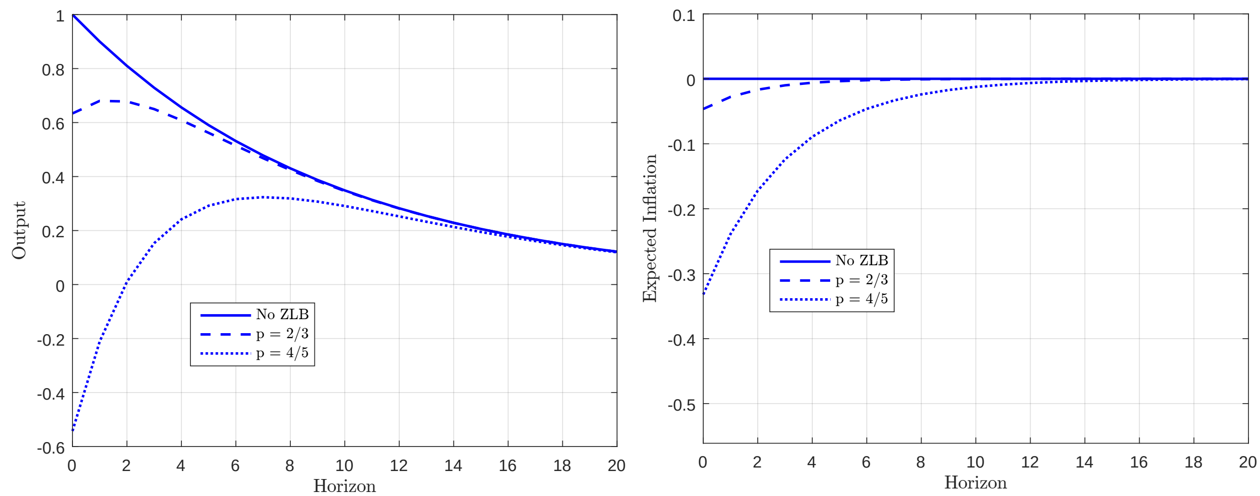

(a) Output Response

(b) Expected Inflation Response

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 (left) expected inflation (right) to a one percent increase in productivity for different values of p. The time period is a quarter.

Figure A1b is similar to Figure A1a, except we plot the expected 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, Fuerst, and Paustian (2014), for sufficiently high values of p, the sign of the effect of a productivity shock on output and expected 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 expected inflation response – if output responds more to the productivity shock at the ZLB, then the expected 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:

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

$$\pi_t = \gamma x_t + \beta\, \mathbb{E}_t\, \pi_{t+1}. \tag{B.2}$$

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:

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

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.

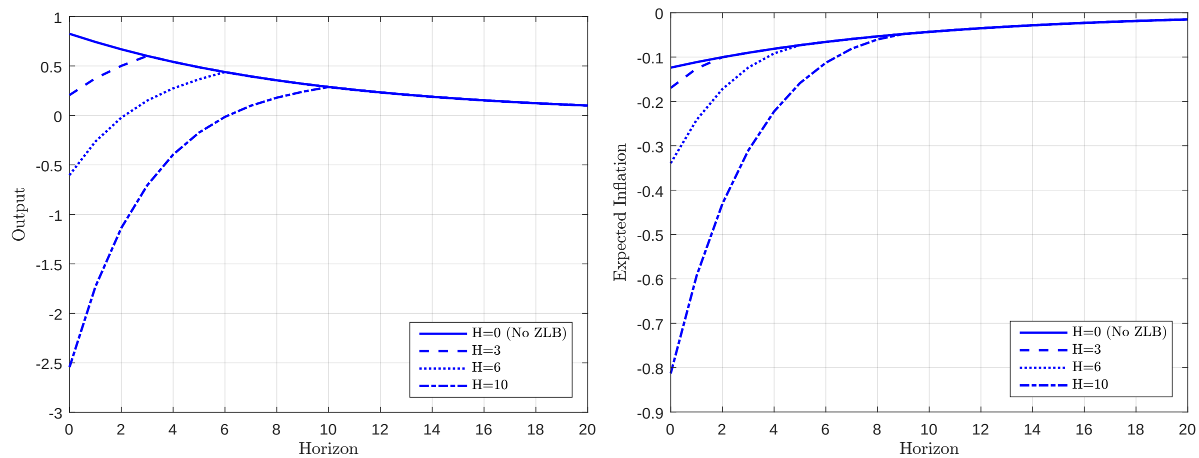

(a) Output Response (b) Expected Inflation Response

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 (left) and expected inflation (right) 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:

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

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

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

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

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

$$w_t^{\#} = \frac{\epsilon_w}{\epsilon_w - 1}\frac{f_{1,t}}{f_{2,t}} \tag{C.6}$$

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

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

$$R_t = \alpha mc_t A_t \widetilde{K}_t^{\alpha-1} N_t^{1-\alpha} \tag{C.9}$$

$$w_t = (1-\alpha) mc_t A_t \widetilde{K}_t^{\alpha} N_t^{-\alpha} \tag{C.10}$$

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

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

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

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

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

$$Y_t = C_t + I_t \tag{C.16}$$

$$\widetilde{K}_t = u_t K_t \tag{C.17}$$

$$Y_t v_t^p = A_t \widetilde{K}_t^{\alpha} N_t^{1-\alpha} - F \tag{C.18}$$

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

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

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

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

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

$$\ln A_t = \rho_A \ln A_{t-1} + \varepsilon_{A,t}. \tag{C.24}$$

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 expected inflation to a positive productivity shock. For completeness we also plot the impulse responses of the nominal rate and the real interest rate. 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, expected 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.

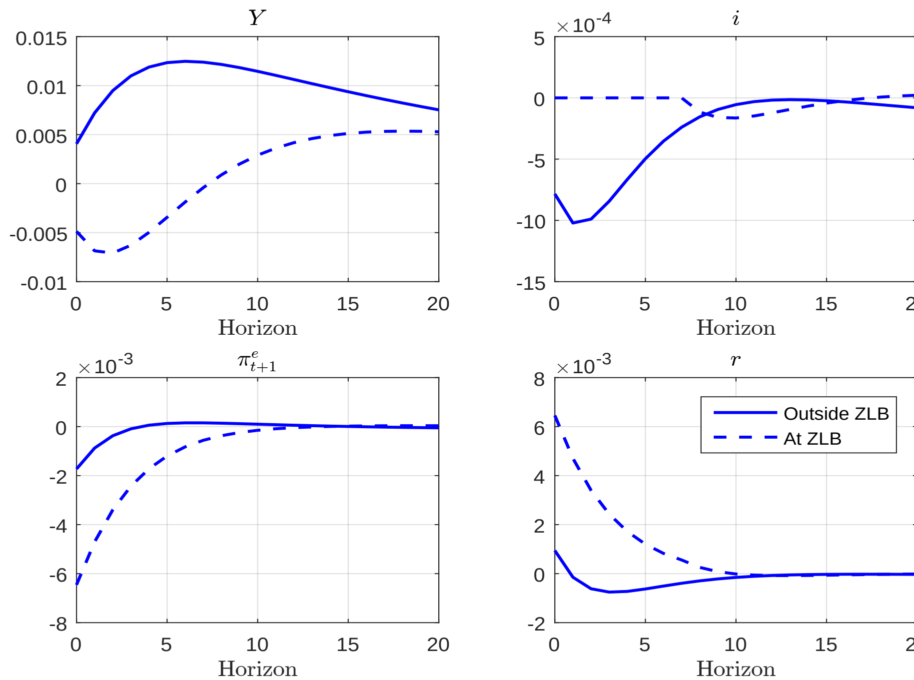

Figure C1: Response to Productivity Shock, with and without ZLB

Notes: This figure plots impulse responses of output, the nominal interest rate, one period ahead expected inflation, and the real interest rate to a productivity shock in the medium scale DSGE model described in this appendix. The solid blue lines are responses when policy is governed by a Taylor rule, while the dashed blue lines are responses when the nominal interest rate is pegged for eight periods.

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:

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

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 $\widehat{\lambda}_t = \lambda_t X_t$ and $\widehat{\mu}_t = \mu_t X_t$. 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:

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

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

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

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

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

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

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

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

$$R_t = \alpha g_{X,t}^{1-\alpha} mc_t \widehat{\widehat{K}}_t^{\alpha} N_t^{1-\alpha} \tag{C.34}$$

$$\widehat{w}_t = (1-\alpha) mc_t g_{X,t}^{-\alpha} \widehat{\widehat{K}}_t^{\alpha} N_t^{-\alpha} \tag{C.35}$$

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

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

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

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

$$\widehat{Y}_t = \widehat{C}_t + \widehat{I}_t \tag{C.40}$$

$$\widehat{\widehat{K}}_t = u_t \widehat{K}_t \tag{C.41}$$

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

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

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

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

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

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

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

$$g_{X,t} = g_t^{\frac{1}{1-\alpha}}. \tag{C.49}$$

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.

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 this version of the model expected inflation falls by less at the ZLB compared to normal times. As we discuss in Section 3, empirically it is the joint behavior of output and expected inflation in response to a productivity shock that is difficult to square with the theory, not the response of output in isolation.

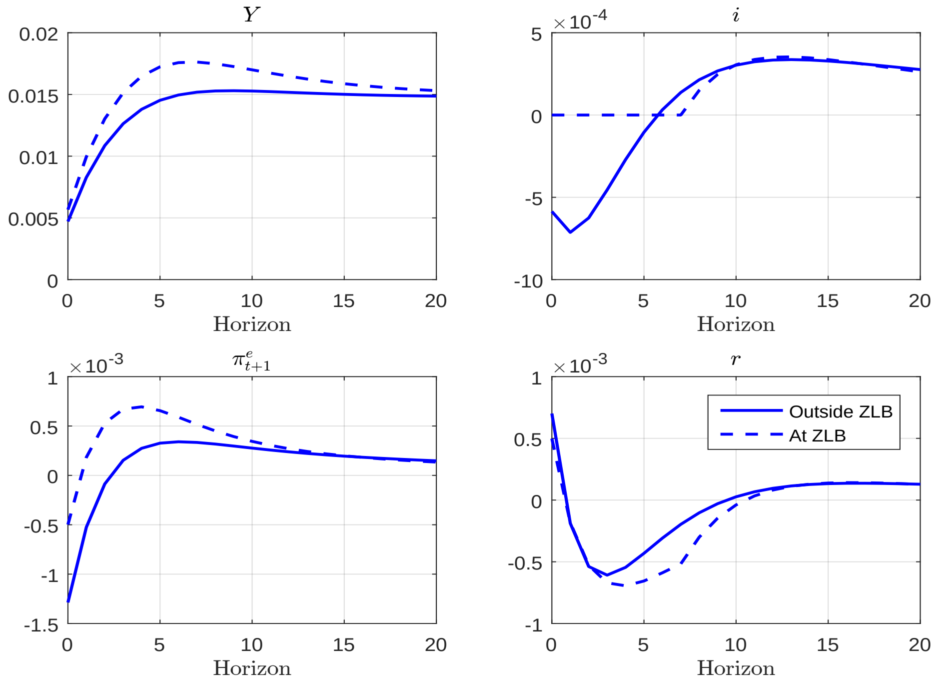

Figure C2: Response to Productivity Shock, with and without ZLB, Permanent Shock with $\rho_A = 0$

Notes: This figure is similar to Figure C1, but considers the case when productivity follows an exact random walk.

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.

Lastly, we increase the persistence of the productivity process further, setting $\rho_A = 0.4$. The responses are shown in Figure C4. Here we are back in the case where output responds less (and expected 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.

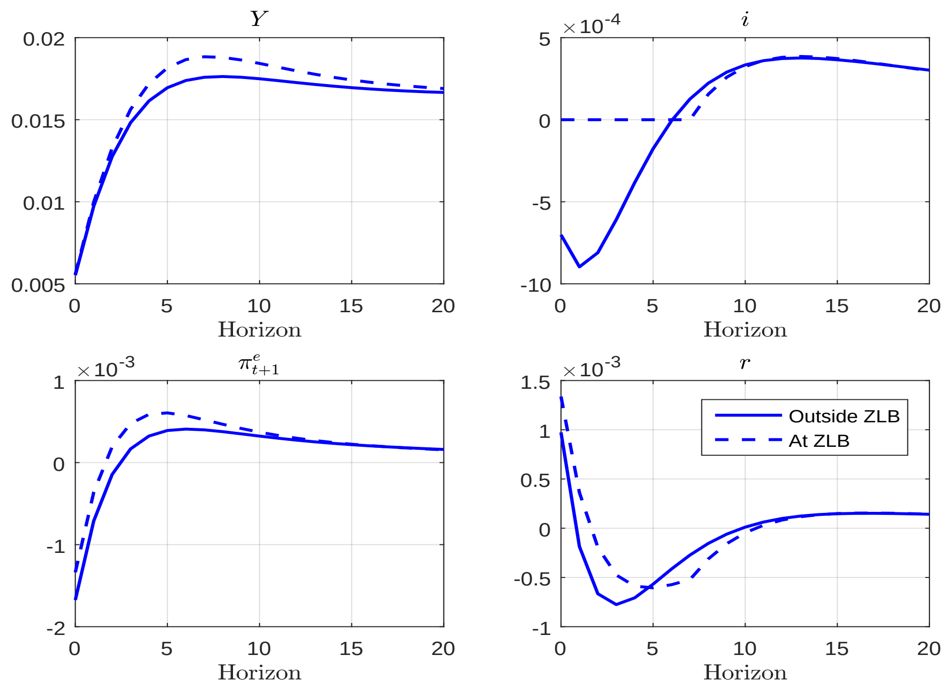

Figure C3: Response to Productivity Shock, with and without ZLB, Permanent Shock with $\rho_A = 0.1$

Notes: This figure is similar to Figure C2, but considers the case when productivity follows a persistent AR(1) process in the growth rate with $\rho_A = 0.1$.

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 expected 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 expected 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 expected 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 expected inflation. If expected 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 expected inflation to a productivity shock which we identify in the data.

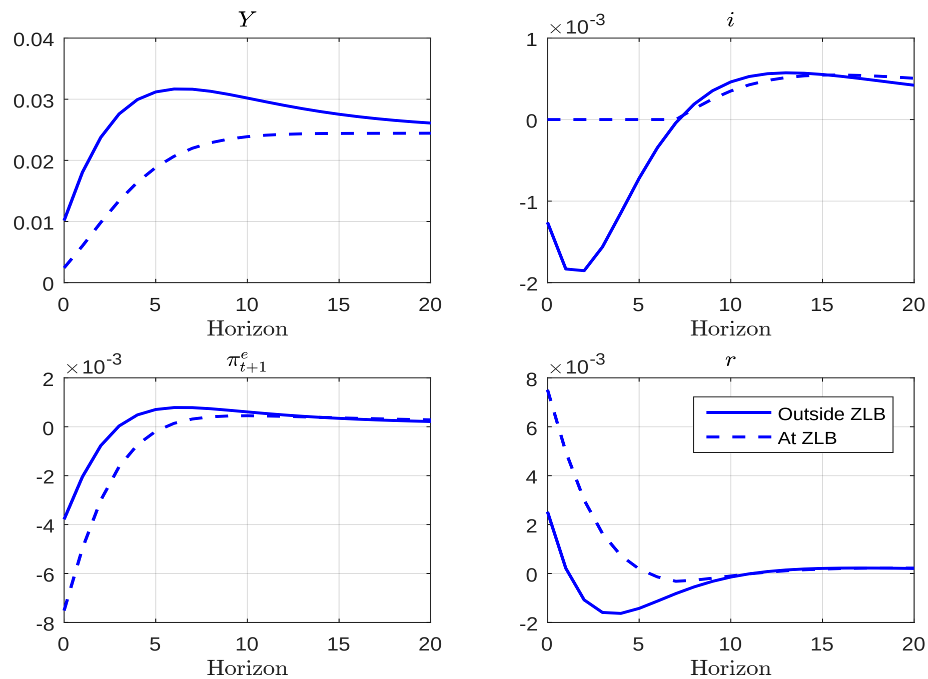

Figure C4: Response to Productivity Shock, with and without ZLB, Permanent Shock with $\rho_A = 0.4$

Notes: This figure is similar to Figure C2, but considers the case when productivity follows a persistent AR(1) process in the growth rate with $\rho_A = 0.4$.

D Smooth Local Projections

Here we briefly outline the smooth local projections (SLP) methodology of Barnichon and Brownlees (2016). In describing the SLP methodology, we abstract from state-dependence. It is straightforward to modify the model to include state-dependence. Suppose that one is interested in the following local projection:

$$Y_{t+h} = \alpha^h + \beta^h X_t + \sum_{i=1}^{N} \gamma_i^h W_{i,t} + u_{t+h}. \tag{D.1}$$

In (D.1), X_t is the regressor of interest, and β^h measures the estimated impulse response at forecast horizon $h \geq 0$. We sometimes also refer to β^h as the "dynamic multiplier." $W_{i,t}$, for $i = 1, \ldots, N$, are control variables (for example lags of Y_t and X_t). The sample size is T periods, so t runs from $t = 1, \ldots, T$. We consider responses out to some horizon $H \geq 0$.

SLP makes use of penalized B-spline smoothing following Eilers and Marx (1996). Let $\mathbf{B}$ be a $(H+1) \times K$ matrix. We follow Barnichon and Brownlees (2016) in using a cubic B-spline, so $K = 3 + (H+1)$. Let $B_k(h)$, for $k = 1, \ldots, K$ and $h = 0, \ldots, H$, denote the k^{th}

column of the h^{th} row of $\mathbf{B}$. As in Barnichon and Brownlees (2016), we only smooth the dynamic multiplier of interest, although it is straightforward to smooth all coefficients. The local projection (D.1) can be approximated as:

$$Y_{t+h} \approx \alpha^h + \sum_{k=1}^{K} b_k B_k(h) X_t + \sum_{i=1}^{N} \gamma_i^h W_{i,t} + u_{t+h}. \tag{D.2}$$

Here, b_k, $k = 1, \ldots, K$, are a set of scalar parameters to be estimated. These parameters are the coefficients on $B_k(h)X_t$ for $k = 1, \ldots, K$. Eilers and Marx (1996) provide Matlab code for computing $\mathbf{B}$, the B-spline base matrix. We follow the details provided in Barnichon and Brownlees (2016).[11] The impulse response coefficient of interest is $\beta^h \approx \sum_{k=1}^{K} b_k B_k(h)$. Once estimates of $\widehat{b_k}$ are obtained, given the B-spline base matrix, it is straightforward to recover the impulse response coefficient.

The B-spline approximation of the local projection, (D.2), can be written in matrix form. We employ the following notation. θ is a vector of parameters:

$$\theta = \begin{pmatrix} b_1 & \ldots & b_k & \alpha^h & \gamma_1^h & \ldots & \gamma_N^h \end{pmatrix}'. \tag{D.4}$$

Let $\widetilde{\mathbf{Y}}$ be a $T(H+1) \times 1$ vector, where the $1 + (t-1)(H+1) + h$ entry equals Y_{t+h} for $t = 1, \ldots, T$ and $h = 0, \ldots, H$. Entries which would require observations of Y_j where $j > T$ are left blank. After removing blank entries, define $\mathbf{Y}$ as the $\left(T(H+1) - \frac{H(H+1)}{2}\right) \times 1$ vector of non-blank elements of $\widetilde{\mathbf{Y}}$. If $H = 0$, this is just a vector of Y_t observations running from $t = 1, \ldots, T$. If $H = 1$, then $\mathbf{Y} = \begin{bmatrix} Y_t & Y_{t+1} & Y_{t+1} & Y_{t+2} & \ldots & Y_{t+T-1} & Y_{t+T} & Y_{t+T} \end{bmatrix}'$ and so on for larger values of H.

Similarly, let $\widetilde{\mathbf{X}}$ be a $T(H+1) \times K$ matrix, where the $(1 + (t-1)(H+1) + h, k)$ element is $B_k(h)X_t$ for $t = 1, \ldots, T$, $h = 0, \ldots, H$, and $k = 1, \ldots, K$. Entries which would correspond to observations of Y_j with $j > T$ are dropped. The remaining $\left(T(H+1) - \frac{H(H+1)}{2}\right) \times K$ entries are collected into the matrix $\mathbf{X}_\beta$. Let $\widetilde{\mathbf{X}}_\alpha$ be a $T(H+1) \times (H+1)$ matrix. The $(1 + (t-1)(H+1) + h, j+1)$ element is 1 for all $t = 1, \ldots, T$, $h = 0, \ldots, H$, and $j = 0, \ldots, H$ if $j = H$ and zero otherwise. Similarly, elements of this matrix which would correspond to values of the left hand side variable outside of the sample are dropped, leaving $\mathbf{X}_\alpha$ as a $\left(T(H+1) - \frac{H(H+1)}{2}\right) \times (H+1)$ matrix corresponding to the constant in the projection. Finally, let $\mathbf{X}_{\gamma_i}$ be similarly defined as a $\left(T(H+1) - \frac{H(H+1)}{2}\right) \times (H+1)$ matrix corresponding

[11] For example, with $H = 2$, our $\mathbf{B}$ matrix is:

$$\mathbf{B} = \begin{bmatrix} 0.1667 & 0.6667 & 0.1667 & 0 & 0 & 0 \\ 0 & 0.1667 & 0.6667 & 0.1667 & 0 & 0 \\ 0 & 0 & 0.1667 & 0.6667 & 0.1667 & 0 \end{bmatrix}. \tag{D.3}$$

to control variables $i = 1, \ldots N$. Horizontally stacking these matrixes together yields a $\left(T(H+1) - \frac{H(H+1)}{2}\right) \times (K + (H+1)(N+1))$ matrix:

$$\mathbf{X} = \left[\begin{array}{cccc} \mathbf{X}_\beta & \mathbf{X}_\alpha & \mathbf{X}_{\gamma_1} & \ldots \mathbf{X}_{\gamma_N} \end{array} \right]. \tag{D.5}$$

In matrix notation, (D.2) can therefore be written:

$$\mathbf{Y} = \mathbf{X}\theta + \mathbf{u}. \tag{D.6}$$

The $\left(T(H+1) - \frac{H(H+1)}{2}\right) \times 1$ vector $\mathbf{u}$ is a vector of prediction errors.

Following Barnichon and Brownlees (2016) and the smoothing literature, we estimate the parameter vector θ using generalized ridge estimation. This estimator minimizes the penalized residual sum of squares and is given by:

$$\begin{aligned}\widehat{\theta} &= \arg\min_{\theta} \quad (\mathbf{Y} - \mathbf{X}\theta)' (\mathbf{Y} - \mathbf{X}\theta) + \lambda\theta'\mathbf{P}\theta \\ &= (\mathbf{X}'\mathbf{X} + \lambda\mathbf{P})^{-1} \mathbf{X}'\mathbf{Y}. \end{aligned} \tag{D.7}$$

In (D.7) λ is a scalar penalty parameter and $\mathbf{P}$ is a $(K + (H+1)(N+1)) \times (K + (H+1)(N+1))$ penalty matrix. We set $\lambda = 100$, though our results are robust to different values. The penalty matrix $\mathbf{P}$ is chosen in such a way as to allow one to shrink the dynamic multiplier, $\beta^h \approx \sum_{k=1}^{K} b_k B_k(h)$, to a polynomial of given order, r. Let $\mathbf{D}$ be an identity matrix of dimension $K \times K$. Then let $\mathbf{D}_r$ denote the rth difference of $\mathbf{D}$, which is a matrix of dimension $(K - r) \times K$.[12] We then let the first $K \times K$ elements of $\mathbf{P}$ be $\mathbf{D}_r'\mathbf{D}_r$, while all remaining elements are zero. For our work in the paper, we follow Barnichon and Brownlees (2016) in choosing $r = 3$. This has the effect of shrinking the estimated impulse response to a polynomial of order $r - 1 = 2$ to the extent to which λ is large.

Given an estimate of $\widehat{\theta}$ from (D.7), it is straightforward to recover the estimates of the dynamic multipliers, $\widehat{\beta}^h \approx \sum_{k=1}^{K} \widehat{b}_k B_k(h)$. Given a vector of prediction errors, $\widehat{\mathbf{u}} = \mathbf{Y} - \mathbf{X}\widehat{\theta}$, we construct a Newey-West estimator for the variance-covariance matrix of $\widehat{\theta}$. Since the impulse response functions are just linear combinations of the elements of $\widehat{\theta}$, it is straightforward to recover standard errors for these response coefficients.

[12] For example, if $r = 3$ and $H = 2$, so that $K = 6$, then $\mathbf{D}_r$ is:

$$\mathbf{D}_r = \left[\begin{array}{cccccc} -1 & 3 & -3 & 1 & 0 & 0 \\ 0 & -1 & 3 & -3 & 1 & 0 \\ 0 & 0 & -1 & 3 & -3 & 1 \end{array} \right]$$

D.1 Monte Carlo Illustration

To study the suitability of the SLP methodology, we consider a simple Monte Carlo exercise. Suppose that we have a data generating process which follows an exact VAR(1):

$$\begin{bmatrix} X_t \\ Y_t \end{bmatrix} = \begin{pmatrix} 0.7 & 0.1 \\ 0.8 & 0.2 \end{pmatrix} \begin{bmatrix} X_{t-1} \\ Y_{t-1} \end{bmatrix} + \begin{pmatrix} 1 & 0 \\ 0.2 & 1 \end{pmatrix} \begin{bmatrix} \varepsilon_{X,t} \\ \varepsilon_{Y,t} \end{bmatrix}, \quad \varepsilon_{X,t} \sim N(0,1),\ \varepsilon_{Y,t} \sim N(0,1) \tag{D.8}$$

In the data generating process we have imposed a recursive structure, wherein the shock to X_t can affect Y_t on impact but not vice-versa.

We generate $N = 1,000$ different data samples with $T = 300$ observations each. The simulations start at initial values of $X_{t-1} = Y_{t-1} = 0$, so we drop the first 100 simulated periods in each simulation to eliminate the influence of initial conditions. This leaves $1,000$ data sets with 200 observations each. On each simulated data set, we seek to estimate the impulse response of Y_t to a shock to $\varepsilon_{X,t}$. We do so by three different methods. In the first, we estimate a VAR(1) on simulated data and impose a Choleski ordering with Y_t ordered second to recover the impulse response. In the second, we estimate a local projection of the form:

$$Y_{t+h} = \alpha^h + \beta^h X_t + \gamma_Y^h Y_{t-1} + \gamma_X^h X_{t-1} + u_{t+h} \tag{D.9}$$

While the local projection in (D.9) looks similar to the second equation of a VAR system, note that we estimate $H + 1$ separate forecasting regressions of this sort. For the third estimation approach, we estimate a smooth local projection of (D.9), using penalty value $\lambda = 100$ and $r = 3$, as we do in the text.

Figure D1 summarizes the results. Solid lines correspond to the true impulse response of Y_t to $\varepsilon_{X,t}$ in the data generating process. Dashed black lines are the responses obtained from estimating a VAR(1) on simulated data. Dotted black lines are the responses estimated from a conventional local projection. Solid blue lines are the responses obtained from estimating a smooth local projection on simulated data.

The upper left plot shows the impulse responses averaged across the $N = 1,000$ different simulations. There is little noticeable difference among the different methodologies for obtaining the impulse response – all do qualitatively well. If anything, the SLP and LP methodologies perform better than the VAR (as evidenced by being less downward-biased at longer forecast horizons).

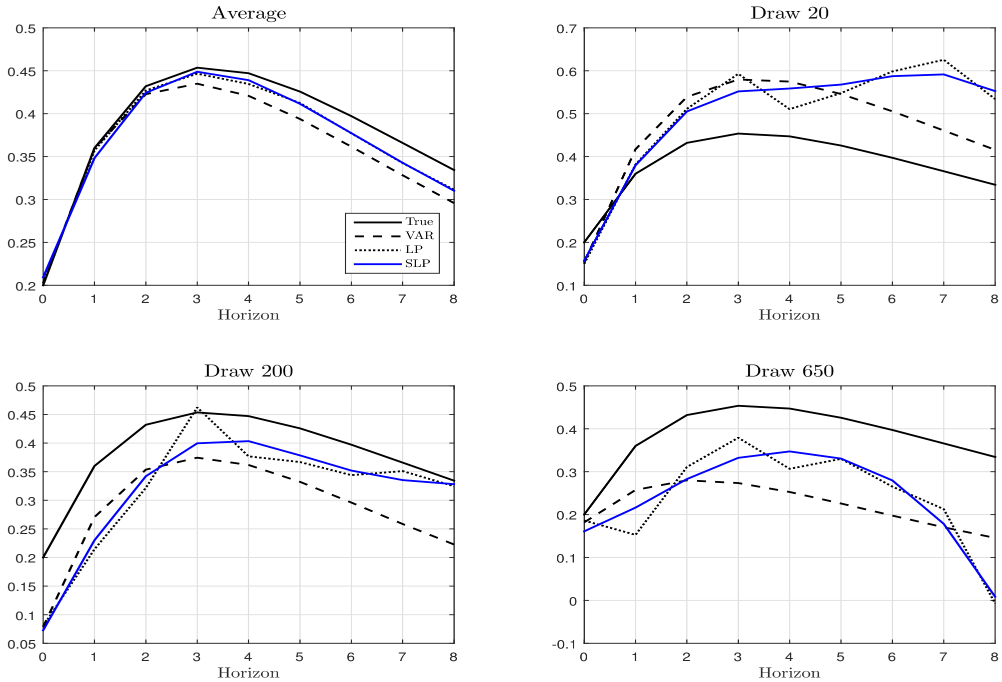

Figure D1: Estimated Impulse Responses from Monte Carlo Exercise

Notes: This figure plots the true response of Y_t to $\varepsilon_{X,t}$ in the data generating process, (D.8). Dashed lines correspond to responses obtained from estimating a VAR(1), dotted lines the responses from a local projection (D.9), and solid blue lines the responses obtained from estimating a smooth local projection version of (D.9). The upper left panel shows the average of estimated responses across 1000 different simulations, while the remaining three plots in the figure correspond to the denoted draw of simulated data.

The remaining three plots in the figure (upper right and the lower row) show estimated responses obtained from particular simulated data sets. In particular, we show responses on the 20th, 200th, and 650th draws of the $N = 1,000$ different simulations. These draws are arbitrarily chosen but serve to illustrate the potential benefits of the SLP methodology. Qualitatively, all the different methodologies capture the true model impulse response fairly well. Nevertheless, on particular draws of data it is evident that the impulse responses obtained via a conventional local projection are quite "choppy." The responses from the SLP are essentially just smoothed versions of those responses.

We next examine the roles played by r and λ. Recall from above that the penalized estimation is designed to push the estimated impulse response function to a polynomial of order r in the limit as λ gets sufficiently large. Figure D2 plots the true impulse response of Y_t to $\varepsilon_{X,t}$ in the model as the solid line. Similarly to Figure D1, there are four panels, with the upper left plotting average responses, the upper right plotting estimated responses from the 20th draw of data, and the bottom row plotting estimated responses from the 200th and 650th draws of data in the simulation. The dashed black, dotted black, and solid blue lines refer, respectively, to the SLP estimate of the impulse response for values of $r \in \{2, 3, 4\}$. These responses are obtained holding λ fixed at 100. Visually, there is not much difference between the SLP responses with different values of r, though one can observe that the estimated responses are somewhat less smooth with larger values of r. This is intuitive – as r is bigger, one is pushing the estimated impulse response to a polynomial of higher order, which allows for more non-linearity.

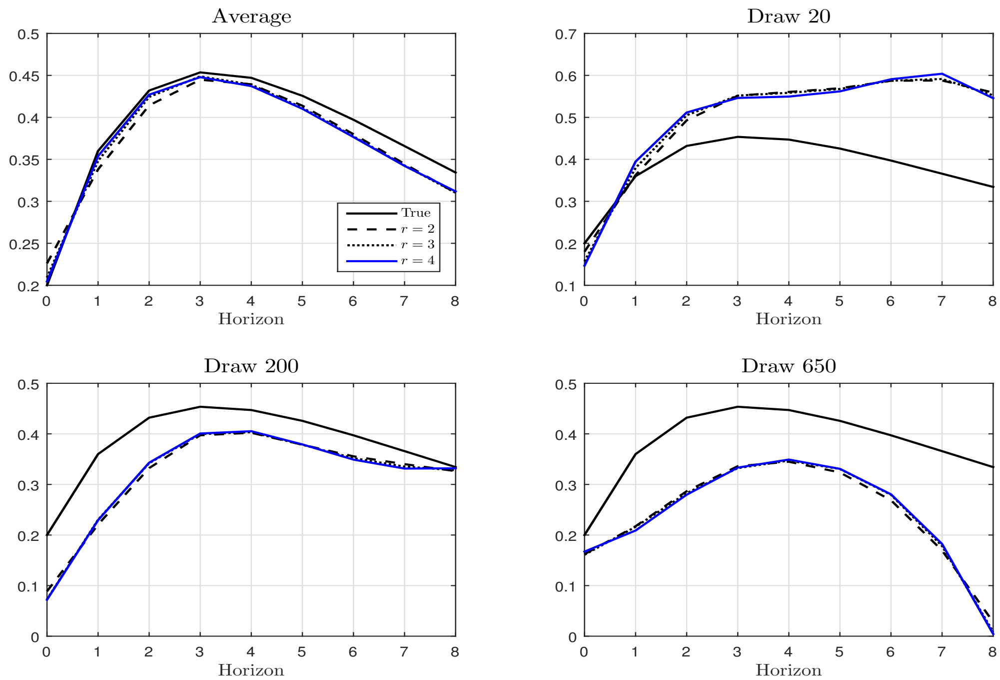

Figure D2: Effect of Varying r

Notes: This figure is constructed similarly to Figure D1, but only plots responses obtained from smooth local projections for different values of r. Solid lines show the true response from the data generating process.

Figure D3 is similar to Figure D2, but instead considers the effects of varying λ holding r fixed at its baseline value of 3. We consider three different values of λ: 1, 100, and 10,000. Here there are more noticeable differences than is evident when varying r. In particular, when λ is lower, the estimated SLP responses are relatively more choppy, whereas when λ is larger the responses are smoother. This again makes sense, as λ measures the penalty for the estimated responses differing from a polynomial of order r.

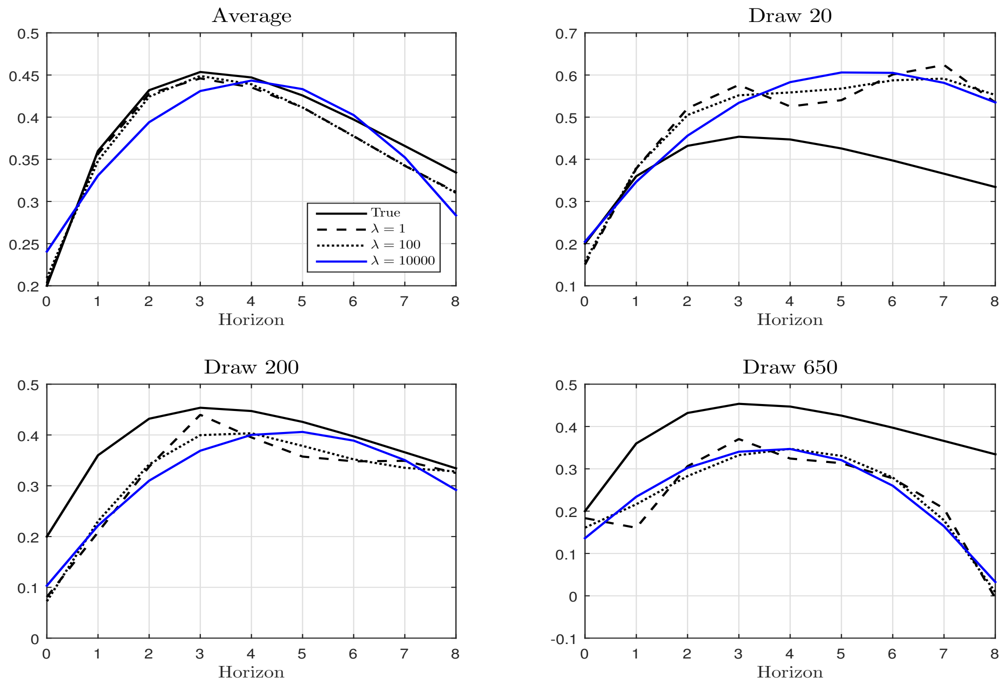

Figure D3: Effect of Varying λ

Notes: This figure is constructed similarly to Figure D2, but only plots responses obtained from smooth local projections for different values of λ. Solid lines show the true response from the data generating process.