Are Negative Supply Shocks Expansionary at the Zero Lower Bound? Inflation Expectations and Financial Frictions in Sticky-Price Models.

Johannes Wieland*

UC Berkeley

Job Market Paper
October 2012 – Version 1.0

For the latest version, **click here**

Abstract

Standard sticky-price models predict that temporary, negative supply shocks are expansionary at the zero lower bound (ZLB) because such shocks lower expected real interest rates and thus stimulate consumption. This paper tests that prediction with an earthquake and oil supply shocks, demonstrating that these shocks are contractionary at the ZLB despite also lowering expected real interest rates. Positive one-year inflation risk premia at the ZLB further indicate that investors want to insure against unanticipated inflation, which is inconsistent with the standard Euler equation framework and suggests that contractionary, negative supply shocks are quantitatively important over this horizon. These facts are rationalized in a model with financial frictions, where negative supply shocks reduce asset prices and net worth — this tightens balance sheet constraints at banks, so that borrowing spreads rise and consumption contracts. As such, ZLB episodes provide a unique opportunity to discriminate one class of models with financial frictions from standard sticky-price models. The model with financial frictions suggests that an intermediate range of policy multipliers is most plausible for forward guidance by the central bank and for fiscal stimulus.

*jwieland@econ.berkeley.edu.

I am grateful to Yuriy Gorodnichenko for invaluable support and guidance, as well as to Olivier Coibion, Martin Lettau, Maury Obstfeld, and David Romer for detailed comments that substantially improved this paper. I would also like to thank Dominick Bartelme, Gabriel Chodorow-Reich, Joshua Hausman, Hashmat Kahn, Matteo Maggiori, John Mondragon, Demian Pouzo, Andrés Rodríguez-Clare, Christina Romer, and Mu-Jeung Yang, as well as seminar participants at UC Berkeley and the NBER Summer Institute (EFCE) for very helpful advice and comments. All errors are my own.

"As some of us keep trying to point out, the United States is in a liquidity trap: [...] This puts us in a world of topsy-turvy, in which many of the usual rules of economics cease to hold. Thrift leads to lower investment; wage cuts reduce employment; even higher productivity can be a bad thing. And the broken windows fallacy ceases to be a fallacy: something that forces firms to replace capital, even if that something seemingly makes them poorer, can stimulate spending and raise employment."

Paul Krugman, 3^{rd} September 2011.[1]

1 Introduction

Does destroying productive capacity raise output when the zero lower bound (ZLB) binds? While this question may seem absurd, in fact it is a common prediction of many macroeconomic models: In these models, temporary negative supply shocks raise inflation expectations and lower expected real interest rates at the ZLB, which stimulates consumption and output. While some prominent economists have subscribed to this view and its policy implications (e.g., Eggertsson and Woodford [2003], Eggertsson and Krugman [2011], Eggertsson [2012]), there is wide disagreement over such a radical and unintuitive proposition.[2] In this paper I test that prediction with the Great East Japan Earthquake, oil supply shocks, and inflation risk premia, and find that negative supply shocks are contractionary at the ZLB. While this result poses a challenge for standard models, I also show that a simple extension of the standard model with credit frictions can accommodate these findings.

Although understanding the economic implications of the ZLB is critical to guiding policy decisions today, there is little empirical evidence regarding how an economy behaves when the central bank has exhausted its conventional ammunition. Policy makers therefore rely on dynamic stochastic general equilibrium (DSGE) models for advice.[3] However, these models were constructed to match the postwar data of major industrialized countries that, with the

---

1 http://krugman.blogs.nytimes.com/2011/09/03/broken-windows-ozone-and-jobs/

2 For example, John Cochrane has argued that negative supply shocks are a major contributor to the current crisis. See http://johnhcochrane.blogspot.com/2012/02/taylors-graphs.html.

3 While I focus on DSGE models in this paper, negative supply shocks are also expansionary in old Keynesian models at the ZLB (Romer [2011]).

exception of Japan, have not faced the ZLB constraint until very recently.[4] This model-based policy is controversial, because various DSGE vintages and calibrations can generate vastly-different policy conclusions at the ZLB that have not been verified in the data. This paper contributes to the debate by testing and rejecting a common prediction in DSGE models, narrowing the class of data-consistent models for policy analysis and the range of plausible policy multipliers.

Specifically, standard DSGE models — building on the work of Woodford [2003] — emphasize the importance of expectations in the propagation of shocks. With a standard Euler equation, consumption today is determined by expected future real interest rates and expected long-run consumption. Temporary, negative supply shocks will not affect long-run consumption, but will raise inflation expectations when prices are sticky. In normal times (i.e., outside the ZLB), the central bank will raise nominal interest rates enough that expected real interest rates rise (the Taylor principle) and consumption contracts. In contrast, nominal interest rates remain unchanged when the central bank is constrained by the ZLB, so that higher inflation expectations reduce real interest rates and consumption today expands. I call this mechanism the "expectations channel," since inflation expectations are the key mechanism by which supply shocks affect real output at the ZLB.

I test for the expectations channel by examining two pieces of evidence. First, I determine the macroeconomic impact of two negative supply shocks at the ZLB: the Japanese earthquake in 2011, and oil supply shocks. My results show that these negative supply shocks are contractionary, despite raising inflation expectations and lowering expected real interest rates. To the best of my knowledge, this is the first paper to test directly for the expansionary effects of negative supply shocks at the ZLB predicted by standard DSGE models. These results suggest that the expectations channel is trumped by other propagation mechanisms at the ZLB. I also provide evidence against a weaker interpretation of the

[4]For example, Christiano, Eichenbaum, and Evans [2005] show that their model can match economic behavior following monetary policy shocks and the Smets and Wouters [2007] model improves on the forecast performance of unrestricted VARs.

expectations channel. Because expected future nominal rates rise less at the ZLB than in normal times, supply shocks should be less contractionary at the ZLB; however, I document that oil shocks appear to be, if anything, *more* contractionary at the ZLB. This suggests that the expectations channel plays only a limited role in the propagation of such shocks.

Second, while it can be difficult to generalize from the economic impact of an earthquake and oil supply shocks to the whole universe of supply shocks, I argue that inflation insurance prices can signal whether negative supply shocks are generally contractionary at the ZLB. In the standard model, higher inflation increases consumption at the ZLB irrespective of the source of inflation. Thus, nominal assets become a hedge — they gain value in deflationary states when consumption is low. Conditional on the ZLB, the inflation risk premium should therefore be negative. In the data, however, the one-year inflation risk premium at the ZLB is typically positive, suggesting that investors want to insure against shocks that generate both high inflation and low consumption. This indicates that negative supply shocks are not only contractionary at the ZLB, but also constitute a significant fraction of aggregate inflation risk over the one-year horizon.

These findings constitute a puzzle for conventional theory, but I show that extending the standard model to allow for financial frictions reconciles the theory with the data.[5] Following Gertler and Kiyotaki [2010] and Cúrdia and Woodford [2009], my model features a balance sheet constraint on financial intermediaries that generates an *endogenous* spread between the borrowing and the deposit rate.[6] Because a negative supply shock reduces profits and share values, the net worth of financial intermediaries falls, tightening the balance sheet constraint. As a result the borrowing spread rises, which reduces consumption by borrowers such that negative supply shocks become contractionary at the ZLB. In addition, banks demand a positive inflation risk premium for holding nominal assets, because these lose value when net worth is low.

[5] In Section 5, I provide empirical support for this modeling choice. However, I do not mean to imply that other mechanisms are unimportant, and I discuss other (complementary) mechanisms in that section.

[6] This contrasts with the original financial accelerator models that focused on firm balance sheets (e.g., Bernanke and Gertler [1989] and Bernanke, Gertler, and Gilchrist [1999]).

While my model generates contractionary effects from negative supply shocks at the ZLB, in normal times it behaves similarly to a standard new Keynesian model, because the central bank's interest rate rule "smoothes out" the financial friction. Thus, a standard new Keynesian model can replicate broad macroeconomic data of the model with financial friction in normal times. However, at the ZLB this "estimated" new Keynesian model will fail to match the friction model qualitatively, because the new Keynesian model cannot generate the contractionary effects of negative supply shocks. Hence, matching key moments of the data in normal times does not ensure even a qualitative fit of moments at the ZLB. Caution should therefore be exercised in employing models for policy analysis that have been subjected to only this particular empirical verification. In contrast, the ZLB provides a unique testing ground, since with propagation mechanisms uninhibited by central bank actions, one can more easily discriminate among models.

Since the model with financial frictions is — unlike the standard new Keynesian model — consistent with the effects of supply shocks at the ZLB, I also investigate its implications for demand-side policies that have previously been derived in the new Keynesian framework. I find that the friction model narrows the range of policy multipliers at the ZLB compared to a standard model. Specifically, in the standard model with flexible real wages, both forward guidance by the central bank and fiscal stimulus significantly raise marginal costs and expected inflation, and are therefore very effective working through the expectations channel. However, in the friction model the rise in real wages lowers profits, asset prices, and net worth, such that higher borrowing spreads offset the strength of the expectations channel and dampen the policies' effectiveness. Conversely, when real wages are inflexible the expectations channel is weak because inflation expectations barely move, but the policies' effectiveness is amplified by the financial accelerator because profits and asset prices now rise. Thus, while demand-side policies remain effective, my model suggests an intermediate range of policy multipliers relative to standard new Keynesian models, so that policy makers should be cautious in expecting large positive outcomes through the expectations channel.

This paper is closely related to literature that has highlighted several "paradoxes" in standard macroeconomic models at the ZLB. First, according to the "Paradox of Thrift," a rise in the desire to save is self-defeating at the ZLB, because it reduces output so much that aggregate savings fall (Keynes [1936], Krugman [1998], Eggertsson and Woodford [2003], and Christiano [2004]). Second, according to the "Paradox of Flexibility," output volatility may rise at the ZLB when prices and wages are more flexible (e.g., Werning [2011], Eggertsson and Krugman [2011]).[7]

My empirical results concern primarily the "Paradox of Toil" (Eggertsson [2010]), whereby a temporary *increase* in desired labor supply at the ZLB *reduces* the equilibrium employment level in standard models. The ensuing decline in real wages lowers expected inflation and raises expected real interest rates at the ZLB, which reduces consumption, output, and employment through the expectations channel. Furthermore, with lower consumption today, wages fall even more — leading to a deflationary spiral (Christiano, Eichenbaum, and Rebelo [2011]). Following this logic, payroll tax cuts are contractionary at the ZLB because they lower expected inflation (Eggertsson [2011]), and allowing collusion among firms is expansionary because it raises expected inflation (Eggertsson [2012]).[8]

However, empirical validations of these paradoxes are scarce.[9] Closest in spirit to my paper is Bachmann, Berg, and Sims [2011]. They show that consumers with above-average inflation expectations have lower willingness to spend, which suggests that higher aggregate expected inflation may not be expansionary at the ZLB. This paper complements their

[7]While this paradox holds at the ZLB, it may also be a feature of normal times as emphasized by DeLong and Summers [1986] and Bhattarai, Eggertsson, and Schoenle [2012].

[8]The "Paradox of Toil" applies only to temporary shocks at the ZLB. Permanent positive supply shocks remain expansionary at the ZLB (Fernández-Villaverde, Guerrón-Quintana, and Rubio-Ramírez [2011b]).

[9]The exception is the literature on fiscal multipliers at the ZLB. Almunia, Bénétrix, Eichengreen, O'Rourke, and Rua [2010] and Gordon and Krenn [2010] estimate fiscal multipliers of around 2 for historical ZLB episodes, which is larger than typical estimates for normal times of around 1 (e.g., Hall [2009], Ramey [2011]). This is consistent with the theoretical ZLB literature (Christiano, Eichenbaum, and Rebelo [2011], Woodford [2011]), which argues that multipliers can significantly exceed 1. An emerging literature also estimates the cross-sectional effect of fiscal policy (e.g., Serrato and Wingender [2010], Nakamura and Steinsson [2011], Chodorow-Reich, Feiveson, Liscow, and Woolston [2012]). However, as emphasized in Nakamura and Steinsson [2011] and Farhi and Werning [2012], a theoretical model and additional assumptions are required in order to convert these local multipliers into aggregate multipliers.

reduced-form results by directly estimating the aggregate, general equilibrium effect of negative supply shocks and by constructing a model that can replicate the contractionary effects of negative supply shocks in the data.

The paper proceeds as follows. In Section 2, I show why standard sticky-price models predict that temporary, negative supply shocks are expansionary at the ZLB. In Section 3, I reject this prediction for the Great East Japan Earthquake and based on estimates from oil supply shocks. In Section 4, I calculate one-year inflation risk premia in the U.S., U.K., and Eurozone, and find that they are typically positive at the ZLB — which is inconsistent with the standard Euler equation framework, and suggests that contractionary, negative supply shocks are quantitatively important over this horizon. In Section 5 I build a model with credit frictions that can match these facts, and in Section 6 I investigate the model's implications for demand-side policies. I conclude in Section 7.

2 Implications from standard sticky-price models

The class of models in this section is a continuous-time adaptation of standard DSGE models that build on Woodford [2003]. Consider a representative agent with separable utility $U(t) = \frac{C(t)^{1-\gamma^{-1}}-1}{1-\gamma^{-1}} - \nu(L(t))$, where $C(t)$ is the consumption flow, γ is the intertemporal elasticity of substitution, and $\nu(L(t))$ is the disutility from work. This agent can purchase a risk-free nominal bond that pays an instantaneous nominal interest rate $i(t)$, as well as a complete set of Arrow-Debreu securities. Any optimal allocation for this agent must satisfy the standard Euler equation in continuous time,[10]

$$\mathbb{E}_t[d\ln C(t)] = \gamma[i(t) - \pi(t) - \varrho]dt, \tag{1}$$

[10]For ease of exposition I ignore the quadratic Ito-terms that arise in this set-up. For plausible calibrations these terms are negligible. In addition, this discussion can also be couched in terms of marginal utilities, where these terms are absent.

where $\pi(t)$ is the instantaneous inflation rate and ϱ the instantaneous discount rate. I ignore habits for ease of exposition, but all results carry over when there is a monotonic relationship between *current* consumption and *current* marginal utility. These results are also robust to including capital in the model, as shown in Appendix A.

Solving the Euler equation forward illustrates that today's consumption depends on the sum of expected real interest rates and expected consumption in the far future:[11]

$$\ln C(t) = -\gamma \mathbb{E}_t \int_0^\infty [\underbrace{i(t+s) - \pi(t+s)}_{\text{Expected Real Rates}} - \varrho] ds + \underbrace{\mathbb{E}_t \lim_{T\to\infty} \ln C(T)}_{\text{“Long-Run” Consumption}} . \tag{2}$$

Thus, consumption today is high relative to the long-run consumption level if the expected path of real interest rates is relatively low, and vice-versa.

I assume that the supply side can be represented by

$$\pi(t) = f(\ln C(t), u(t), \bullet), \quad \infty > f_1 > 0, \ \infty > f_2 > 0, \tag{3}$$

where $u(t)$ is a supply shifter (e.g., a negative technology shock) and f_x denotes the derivative of f with respect to the x^{th} argument.[12] Higher consumption today causes higher inflation, for example because higher consumption raises the real wage and therefore also the marginal cost of production. Similarly, an increase in $u(t)$ (e.g., a deterioration in productivity) also raises inflation. This supply side implicitly includes price or wage stickiness — when only a fraction of firms can adjust prices in response to a marginal cost shock, then these firms raise prices only a little to preserve their market share. The next group of firms will again raise their prices relative to the aggregate price level, so that the price level rises over time and inflation is positive. This contrasts with a flexible-price economy, in which higher marginal

[11] I use the "improper integral" notation to mean the continuous limit of the integral as time approaches infinity, $\int_z^\infty f(x)dx = \lim_{T\to\infty} \int_z^T f(x)dx$. This limit exists, for example, in the standard New Keynesian model. Even in cases where the limit does not exist, the following results hold for large enough T.

[12] For example, the standard new Keynesian Phillips Curve, $d\pi(t) = \varrho\pi(t) - \kappa[(\sigma+\nu)(\ln C(t) - \ln \bar{C}) + u(t)]$, can be rewritten as $\pi(t) = \frac{\kappa}{\rho_u+\varrho}[(\sigma+\nu)(\ln C(t) - \ln \bar{C}) + u(t)]$, assuming that $u(t)$ is the only state variable and follows an Ornstein-Uhlenbeck process with mean-reversion ρ_u.

costs generate expected deflation, because firms immediately raise prices to preserve the mark-up and then let them gradually fall as the marginal cost shock dissipates.

To analyze the effect of a negative supply shock on consumption — an increase in $u(t)$ — I first differentiate Equation (2):[13]

$$\frac{\mathrm{d}\ln C(t)}{\mathrm{d}u(t)} = -\gamma \frac{\mathrm{d}\mathbb{E}_t \int_0^\infty [i(t+s) - \pi(t+s)] ds}{\mathrm{d}u(t)} + \frac{\mathrm{d}\mathbb{E}_t \lim_{T\to\infty} C_T}{\mathrm{d}u(t)}.$$

This reveals that the supply shock can affect consumption by changing either real interest rates or long-run consumption. I restrict my attention to persistent, but temporary supply shocks, so that the second term on the right-hand side (RHS) is zero. Thus, a temporary, negative supply shock can be expansionary only if it lowers the expected sum of future real interest rates — the first term on the RHS.

Summary 1 *Given the Euler equation (1), a temporary, negative supply shock — an increase in $u(t)$ — is expansionary if and only if it lowers the sum of expected future real interest rates:*

$$\frac{d\mathbb{E}_t \int_0^\infty [i(t+s) - \pi(t+s)] ds}{du(t)} < 0.$$

A crucial determinant of the real interest rate response is the interest rate rule followed by the central bank. I assume that this takes a standard linear form subject to the zero bound constraint,

$$i(t) = \max\{\bar{i} + \phi_\pi(\pi(t) - \bar{\pi}) + \phi_y(\ln C(t) - \overline{\ln C}), 0\}, \tag{4}$$

where variables with a bar denote steady-state values. In standard new Keynesian models this rule encompasses as special cases both inflation targeting and the time-consistent optimal policy rule. One can also include interest-rate smoothing, but at the cost of notational complexity.

[13] I assume that all of these objects are differentiable with respect to $u(t)$.

Suppose first that the ZLB does not bind, and that the supply shifter $u(t)$ follows a first-order autoregressive process with mean-reversion ρ_u. This is equivalent to an autoregressive process with persistence $\hat{p}_u = \exp\{-\rho_u\}$ — i.e., the *larger* ρ_u is, the *less* persistent is the process. If $u(t)$ is the only state variable, then inflation and log consumption follow first-order autoregressive processes with mean-reversion $\rho_\pi = \rho_u$ and $\rho_c = \rho_u$ respectively. The change in consumption today from a supply shock is then given by

$$\frac{\mathrm{d}\ln C(t)}{\mathrm{d}u(t)} = -\gamma\frac{(\phi_\pi - 1)}{\rho_\pi}\frac{\mathrm{d}\pi(t)}{\mathrm{d}u(t)} - \gamma\frac{\phi_y}{\rho_c}\frac{\mathrm{d}\ln C(t)}{\mathrm{d}u(t)}, \tag{5}$$

where the terms on the RHS are respectively the change in real interest rates from higher inflation and from movements in consumption. Conditional on a supply shock, Equation (5) traces out a downward-sloping line in the $(\ln C(t), \pi(t))$-space with slope $\frac{\mathrm{d}\ln C(t)}{\mathrm{d}\pi(t)} = \frac{\phi_\pi - 1}{\rho_\pi[1+\gamma\phi_y/\rho_c]}$, which is plotted in Figure 1(a). Intuitively, under the policy rule (4) higher inflation from a supply shock translates into higher expected real interest rates, which reduces consumption today. Since this curve is derived from the consumer's FOC, I label it the "Aggregate Demand" (AD) curve as is common in the literature (Romer [2011], Eggertsson [2012]).

Using the supply side (3), we can substitute $\frac{\mathrm{d}\pi(t)}{\mathrm{d}u(t)} = \frac{\partial\pi(t)}{\partial\ln C(t)}\frac{\mathrm{d}\ln C(t)}{\mathrm{d}u(t)} + \frac{\partial\pi(t)}{\partial u(t)}$, and find that a negative supply shock (an increase in $u(t)$) lowers consumption when the ZLB does not bind:[14]

$$\frac{\mathrm{d}\ln C(t)}{\mathrm{d}u(t)} = \left[1 + \gamma\frac{(\phi_\pi - 1)}{\rho_\pi}\underbrace{\frac{\partial\pi(t)}{\partial\ln C(t)}}_{=f_1>0} + \gamma\frac{\phi_y}{\rho_c}\right]^{-1}\left[-\gamma\frac{(\phi_{pi} - 1)}{\rho_\pi}\underbrace{\frac{\partial\pi(t)}{\partial u(t)}}_{=f_2>0}\right] < 0.$$

This result is illustrated in Figure 1(a), which plots the supply side (3) as an upward-sloping AS curve because higher consumption generates higher inflation by raising marginal costs. A negative supply shock raises inflation *ceteris paribus*, so that the AS curve shifts left to

[14] I assume that (ϕ_π, ϕ_y) are such that the denominator in this expression is positive. This is a necessary and sufficient condition for determinacy in normal times.

AS_2. The central bank responds by raising real interest rates, after which consumption contracts along the AD curve. Thus, consistent with Summary 1, negative supply shocks are contractionary in normal times because they raise expected future real interest rates.

However, while in normal times the central bank raises nominal interest rates enough to also raise real rates, this may not occour when the central bank is constrained by the ZLB. To focus on this case, I first consider a supply shock that affects the economy only at the ZLB (Eggertsson and Woodford [2003]). At time t, the negative supply shock $u(t) = \Delta$ hits an economy at the ZLB. Both the exit from the ZLB and the negative supply shock are jointly determined by a Poisson process of intensity ρ_π^{ZLB}. With probability $\rho_\pi^{ZLB} dt$ (over an interval dt), the ZLB constraint no longer binds, the supply shock is set to zero, and the economy jumps back to the stochastic steady state. With probability $(1 - \rho_\pi^{ZLB} dt)$, the economy remains at the ZLB and the supply shock at Δ. This set-up ensures that the supply shock affects the economy only at the ZLB, so that there is no convolution of ZLB effects with normal times. In this case, the Euler equation yields the following:

$$\frac{\mathrm{d}\ln C(t)}{\mathrm{d}u(t)} = \gamma \frac{1}{\rho_\pi^{ZLB}} \frac{\mathrm{d}\pi(t)}{\mathrm{d}u(t)} \quad \Rightarrow \quad \frac{\mathrm{d}\ln C(t)}{\mathrm{d}\pi(t)} = \gamma \frac{1}{\rho_\pi^{ZLB}} > 0. \tag{6}$$

Compared to aggregate demand in normal times (Equation (5)) the persistence of inflation ρ_π^{ZLB} enters analogously. However, there is no change in nominal interest rates so that the terms involving the policy parameters ϕ_π and ϕ_y drop out. The resulting AD curve is now upward-sloping, as shown in Figure 1(b), because at the ZLB higher expected inflation lowers expected real interest rates so that consumption expands. The upward-shift in the AS curve from a negative supply shock therefore raises consumption,

$$\frac{\mathrm{d}\ln C(t)}{\mathrm{d}u(t)} = \left[1 - \gamma \frac{1}{\rho_\pi^{ZLB}} \underbrace{\frac{\partial \pi(t)}{\partial \ln C(t)}}_{=f_1>0} \right]^{-1} \left[\gamma \frac{1}{\rho_\pi^{ZLB}} \underbrace{\frac{\partial \pi(t)}{\partial u(t)}}_{=f_2>0} \right] > 0,$$

because — consistent with Summary 1 — expected real interest rates fall.[15]

[15]Determinacy at the ZLB requires that the denominator be positive. Results for the indeterminant case

Summary 2 *Consider an economy that satisfies the Euler equation* (1)*, the supply side* (3)*, and let monetary policy be given by* (4)*. Then, a temporary, negative supply shock — an increase in* $u(t)$ *— that does not persist after the ZLB binds is expansionary.*

The previous stochastic process implies that the supply shock does not affect the economy after the ZLB ceases to bind. However, if a negative supply shock persists beyond the duration of the ZLB, then the central bank may raise real interest rates at those dates, which will impact today's consumption through the expectations channel. I call these effects "spillovers."[16] For example, if exit from the ZLB is governed by an exogenous Poisson process with intensity λ, then the change in real interest rates is given by

$$\frac{\mathrm{d}\mathbb{E}_t \int_0^\infty [i(t+s) - \pi(t+s)]ds}{\mathrm{d}u(t)} = \underbrace{\frac{\mathrm{d}\mathbb{E}_t \int_0^\infty [\pi(t+s)]ds}{\mathrm{d}u(t)}}_{<0 \text{ (supply shock raises inflation)}} + \underbrace{\frac{\mathrm{d}\mathbb{E}_t \int_0^\infty (1-e^{-\lambda s})[i(t+s)]ds}{\mathrm{d}u(t)}}_{\text{Exit from ZLB: Spillovers}}.$$

If spillovers are zero (or negative), then negative supply shocks unambiguously lower real interest rates — a *sufficient* condition for negative supply shocks to be expansionary at the ZLB. However, if spillovers are large and positive, then a negative supply shock might raise expected real interest rates and be contractionary at the ZLB.[17] In fact, large negative supply shocks such as extreme natural disasters and wars will remain contractionary at the ZLB because of such spillover effects. These shocks significantly raise both marginal costs of production and inflation expectations, which in turn accelerates the exit from the ZLB and raises expected future real interest rates so that consumption contracts.

Long-term bond yields can reveal whether spillovers are important. According to the expectations hypothesis, long-term bond yields are an average of short-term rates, such that

are discussed in Benhabib, Schmitt-Grohé, and Uribe [2002].

[16] Woodford [2011] highlights the importance of spillovers for the effectiveness of fiscal policy. Swanson and Williams [2012] show that the long-end of the yield curve responds to economic news — that is, spillovers are a possibility at the ZLB.

[17] Spillovers also capture whether the central bank undoes the rise in prices following a negative supply shock. In this case, lower real interest rates at the ZLB must be offset by higher real interest rates in the future, so that a negative supply shock cannot be expansionary. These dynamics also occur in a monetary union, because the real exchange rate must be stationary (Nakamura and Steinsson [2011], Farhi and Werning [2012]). Thus, the ZLB paradoxes arise only in economies that issue their own currency.

higher expected future policy rates will be reflected in higher long-term bond yields today,

$$\frac{\mathrm{d}y(t,t+s)}{\mathrm{d}u(t)} = \frac{1}{s}\left[\underbrace{\int_0^s e^{-\lambda s}\frac{\mathrm{d}i(t+k)}{\mathrm{d}u(t)}dk}_{\text{ZLB} \Rightarrow =0} + \underbrace{\int_0^s (1-e^{-\lambda s})\frac{\mathrm{d}i(t+k)}{\mathrm{d}u(t)}dk}_{\text{Spillovers}}\right],$$

where $y(t, t+s)$ is the annualized yield of a bond with a maturity of s years. Conversely, if long-term bond yields do not rise, then there are no spillovers and negative supply shocks should be expansionary at the ZLB.

Summary 3 *Spillovers from temporary, negative supply shocks at the ZLB into normal times will be reflected in higher long-term nominal interest rates. Without spillovers, negative supply shocks at the ZLB are expansionary in the baseline model.*

This summary motivates the empirical strategy in the following sections. I examine two temporary, negative supply shocks and show that they raise inflation expectations but not expected future nominal rates. I then test whether these shocks are indeed expansionary as predicted by the Euler equation (1).

3 Negative Supply Shocks at the ZLB

In this section, I consider two examples of negative supply shocks. The first is the 2011 earthquake in Japan, which serves as a clean natural experiment. The second example is oil supply shocks, the impact of which I estimate over a pooled sample including Japan, the U.S., the U.K., Canada, Sweden, and the Eurozone.

3.1 The Great East Japan Earthquake On March 11^{th} 2011, a magnitude-9.0 earthquake off the Japanese eastern coastline triggered a tsunami that caused extensive damage to Japanese residential and commercial structures, created an electricity shortage, and disrupted global supply chains. This was an exogenous negative supply shock — to produce the same quantity of output with less capital, producers had to incur higher costs. Consistent with this logic, the 2011 consensus inflation forecast for Japan rose by 0.3 percentage points, and the 2012 forecast by 0.2 percentage points.[18] This is plotted in Figure 2, which also shows that inflation expectations rose in the U.S., suggesting that the disruption of supply lines constituted a global negative supply shock.

While professional forecasters expected higher inflation following the earthquake, they revised output forecasts down, as shown in Figures 2(a) and 2(c). Relative to pre-earthquake predictions, Japanese output was forecasted to be about 1.2% lower during 2011, while the U.S. output forecast dropped by 0.3%. These annual forecasts also mask the severe output losses that occurred during the quarter of the earthquake, when Japanese real output declined at an annualized rate of 7.2%, and real consumption contracted by 4.4%. Japan recovered to its pre-earthquake peak only by the first quarter of 2012. Thus, this negative supply shock had contractionary effects despite raising inflation expectations at the ZLB.

To apply the results from the previous section, this supply shock must be temporary and must not significantly raise expected future nominal interest rates. A priori, the first condition seems reasonable — the capital stock destroyed by the earthquake will be rebuilt as the economy converges to its balanced growth path. In fact, the April survey's GDP growth forecast for 2012 was revised *upward*, making up half of the loss from the forecast revision for 2011, as shown in Figure 2(a). This suggests that the Japanese economy is catching up to its balanced growth path and that — at most — half of the decline in output could be due to reductions in permanent income. The second condition is also satisfied, as the yield on 10-year government bonds *fell* from 1.27% on March 10^{th} to 1.19% on March 14^{th}.[19]

As such, the earthquake was a temporary, negative supply shock that reduced expected future real interest rates. However, it was not expansionary, as predicted by the standard

[18] I use the February and April surveys as references, because the March 2011 Consensus Economics survey was published on March 15^{th}, only a few of days after the earthquake struck. Given the publishing lag, not all forecasts in the March edition reflect the impact of the earthquake. For example, the Morgan Stanley forecast in March is unchanged from the forecast in February.

[19] The earthquake occurred 15 minutes before the stock market closed. It is likely that the earthquake's devastating effects were not priced into the closing quote of 10-year government bonds (=1.25%) on that day. The 10-year bond yield remained below 1.27% until the end of March and by the end of May it had fallen to 1.15%.

consumption Euler equation, which emphasizes the expectations channel.[20] Thus, the actual evidence is inconsistent with a basic prediction of the framework in Section 2.

3.2 Oil Supply Shocks The economic impact of the earthquake in Japan illustrates that supply shocks can be contractionary at the ZLB, even when expected future real interest rates fall. In this section, I show that the effect of oil supply shocks is qualitatively similar.

The obvious challenge in uncovering oil supply shocks from production and price data is to separate demand from supply effects. I follow Kilian's [2009] identification assumption, that oil supply does not respond to demand shocks within a month; this allows oil supply shocks to be identified through a standard Cholesky decomposition.[21,22] Using this methodology, I extend the original oil supply shock series to December 2011. Similar to the original series, a one-standard-deviation negative oil supply shock raises real oil prices by 0.65% for about 12 months. This is consistent with the measure picking up reductions in supply rather than current or expected negative demand shocks. Importantly, oil supply shocks do *not* have a permanent effect on real oil prices in the VAR. Therefore, these constitute temporary shocks as required by the theory in Section 2.

Since this shock series extends into the current ZLB episode, I include the U.S., the U.K., the Eurozone, Canada, Sweden, and Japan in my estimation. I restrict the baseline estimation to dates when the ZLB binds, which are tabulated in Table 1. In practice the central bank interest rate floor is above zero, so dates are determined with the following

[20] The absence of precautionary savings motives in the baseline model is unlikely to be the source of this discrepancy. In the standard model the first order effect of a negative supply shock at the ZLB is to raise consumption today. Since variations around a higher consumption level are less costly, consumers are likely to reduce their precautionary savings, which further amplifies the first order effects. Thus, endogenous precautionary savings may in fact aggravate the empirical puzzle that the earthquake is not expansionary. Furthermore, uncertainty was only elevated for one month after the earthquake (as measured by the Nikkei VIX index), before returning to its historical average. So, precautionary motives from increased uncertainty cannot explain the persistence of the contraction.

[21] Kilian and Murphy [2011] find that the short-run supply elasticity is small, suggesting that the original Cholesky decomposition in Kilian [2009] will perform well in recovering the underlying supply shocks. However, I do not use their methododology, because they impose sign-restrictions on real activity.

[22] I have also used the narrative oil shock series in Cavallo and Wu [2006] and found that the responses of inflation expectations, unemployment, and industrial productions have the same sign as for the Kilian series. However, since the Cavallo and Wu [2006] shock series features only 12 shocks when the ZLB binds, the estimates are too imprecise to reject the null hypothesis, unlike for the Kilian series.

interest rate cut-offs: 0.5 for Japan before 1998, and 0.25 thereafter; 0.25 for the U.S., Canada, and Sweden; 0.5 for the U.K.; and 1.00 for the Eurozone. Unlike the other countries, Eurozone policy rates have fluctuated below this cut-off, so it is less clear that the ECB was constrained in responding to negative supply shocks. In Appendix B I therefore repeat my analysis excluding the Eurozone, and reach what essentially are unchanged results.

I first test whether negative oil supply shocks raise inflation expectations at the ZLB. These inflation expectations data are taken primarily from Consensus Economics, and supplemented by various national consensus forecasts. Consensus Economics provides annual inflation forecasts at a monthly frequency, and quarterly inflation forecasts at a quarterly frequency. I combine these two data sources in a Kalman filter to extract four-quarter-ahead inflation forecasts at a monthly frequency. For example, the March 4Q-ahead inflation forecast captures expected inflation from Q2 this year to Q2 next year. In April the window moves forward to Q3-Q3, in July to Q4-Q4, and in October to Q1-Q1. This is a conservative choice, since it excludes changes in expected inflation for the remainder of the quarter. The details of the data sources and the Kalman filter estimation are relegated to Appendix C.

In the baseline specification, I regress changes in 4Q-ahead inflation expectations[23] $\Delta\pi_t^{e,4}$ on lagged values, lagged oil shocks, and lagged controls:

$$\Delta\pi_t^{e,4} = \alpha + \sum_{j=1}^{12} \beta_j \Delta\pi_{t-j}^{e,4} + \sum_{j=1}^{24} \gamma_j oil_{t-1} + \sum_{j=1}^{24} \delta_j \Delta controls_{t-1} + \varepsilon_t^{\pi}. \tag{7}$$

Since inflation forecasts are *published* at the beginning of the month,[24] I exclude contemporaneous shocks from the regression because these are unlikely to be contained in the forecasters' information sets. I include lagged forecast revisions because they have been shown to predict current forecast revisions (Coibion and Gorodnichenko [2011]). I employ as additional controls changes in industrial production and long-term bond rates to proxy

[23] The differencing takes the moving window of inflation forecasts into account. For example, the change for April is calculated as the difference in the Q3-Q3 inflation forecast from April and the Q3-Q3 inflation forecast from March.

[24] In my sample, Consensus Economics forecasts are published between the 4^{th} and 15^{th}, with a median at the 11^{th} of the month. Forecasts contained in this survey are likely to be older.

for changes in economic activity. Thus, the regression estimates capture how an oil supply shock affects inflation expectations for a given level of output — that is, the partial derivative $f_2 = \partial\pi(t)/\partial u(t)$ in the Phillips curve (3). All regressions include country fixed effects, and lag lengths are determined by the Akaike Information Criterion (AIC).[25]

The baseline results are tabulated in Figure 3(a), along with 95% confidence intervals based on Driscoll and Kraay [1998] standard errors.[26] A one-standard-deviation negative oil supply shock has a positive and statistically-significant short-run effect on 4Q-ahead inflation expectations, which peaks at about $2/5^{th}$ of a standard deviation (5 basis points) after three months, and tails off thereafter. Again, this is a conservative estimate of changes in inflation expectations, because (by construction of the 4Q-ahead difference) it excludes effects of the oil shock in the current quarter. Including the current quarter generates larger responses in inflation expectations, as I show in Appendix B. Nonetheless, for either measure oil shocks raise inflation expectations for a given level of output, consistent with the discussion in Section 2.[27]

My first measure of monthly economic activity is industrial production (IP). I regress growth rates in industrial production on lagged values, and on both contemporaneous and lagged oil shocks:

$$\Delta y_t = \alpha + \sum_{j=1}^{48} \beta_j \Delta y_{t-j} + \sum_{j=0}^{24} \gamma_j oil_{t-j} + \varepsilon_t^y. \tag{8}$$

While lagged dependent variables are not necessary if the shocks are well-identified, including these terms sharpens the estimates and ensures that the oil shock coefficients do not capture dynamics of output induced by other shocks. In Appendix B, I show that very similar

[25] Coibion [2012] shows that the AIC performs better than BIC lag selection given the short samples employed here. Using BIC attenuates the results, but the estimates for inflation expectations and unemployment remain significantly different from zero. In Appendix B, I tabulate results for alternative lag lengths.

[26] Confidence intervals are constructed by Monte Carlo draws from the Driscoll-Kraay coefficient covariance matrix. These standard errors are robust to heteroscedasticity, temporal dependence, and general forms of cross-country dependence as the time dimension of the panel becomes large.

[27] Evidence from inflation swap rates also indicates that inflation expectations rise following an oil shock, and are not offset by lower future inflation expectations. For the sample when such data is available (2007-now for Japan, 2005-now for others), an oil shock raises the 10-year swap rate by 1.74 basis points, and when the ZLB binds by 1.1 basis points. However, neither estimate is significant at conventional levels.

impulse responses result when these terms are excluded, suggesting that the oil shocks are largely orthogonal to current economic conditions.

In Figure 3(b) I plot the impulse response function (IRF) of log industrial production from the estimated Equation (8). Following a one-standard-deviation oil supply shock, there is a marginally-significant decline in IP for about five months, with a peak decline of approximately 0.75%. In the following six months the IRF reverts back to zero, and is statistically insignificant. Thus, despite raising inflation expectations at the ZLB, a negative oil supply shock has contractionary effects on IP.

The same conclusion emerges when I estimate Equation (8) with the unemployment rate instead of IP (with the lag length for unemployment is set to 36, as recommended by the AIC). The IRF in Figure 3(c) exhibits a statistically-significant increase in the unemployment rate by 0.1 percentage points after a one-standard-deviation negative oil supply shock. This unemployment response begins to dissipate after eleven months, and becomes statistically insignificant twelve months after the shock.

Since the economic contraction from an oil shock may work through investment rather than consumption, it is not obvious that these results constitute a failure of the Euler equation. To address this, I use consumption expenditure data for Japan and the U.S., which is available at a monthly frequency. There is strong seasonality in the Japanese data, so I adjust it using the X-12 ARIMA filter from the U.S. Census Bureau website. I then estimate Equation (8) for consumption expenditure (with 36 lags of oil shocks, as recommended by AIC). The IRF in Figure 3(d) is quite choppy because this series is volatile, but nonetheless displays a statistically-significant decline in consumption expenditures four months after the oil shock hits. This evidence suggests that the contractionary effects of oil shocks at the ZLB also reflect cut-backs by consumers.

These effects are consistent with the standard Euler equation framework in Section 2 only if there is a large increase in long-term nominal bond rates (Summary 3) or if the oil shocks have permanent effects. The latter can be plausibly ruled out based on the temporary effect of supply shocks on real oil prices. In addition, the estimated impulse response functions are insignificant, and very close to zero after four years: -0.2% for IP, 0.04 percentage points for unemployment, and -0.15% for consumption.[28]

To assess whether long-term rates have risen, I estimate the impact of oil shocks on the 10-year bond rate. Column (1) of Table 2 shows that the bond rate significantly *falls* following an oil supply shock. The same result emerges when one uses fewer lags than implied by AIC (column 2), or a regression specification that down-weights outliers (column 3). Thus, spillovers of supply shocks into normal times cannot explain the contractionary effects of negative supply shocks. This constitutes a puzzle given the standard set-up in Section 2, because without spillovers a temporary, negative supply shock should be expansionary through the expectations channel (Summary 3). Since output contracts even though expected real rates decline, (at least) one assumption in this framework appears to be violated in the data.

These results indicate that other propagation mechanisms trump the expectations channel for negative supply shocks, and that monetary policy responds to the contractionary effects by loosening future monetary policy. However, it remains possible that the expectations channel mitigates the contraction at the ZLB, because a constrained central bank does not actively raise expected future nominal interest rates. I test this weaker hypothesis by estimating Equation (8) whenever the ZLB does not bind using data from 1985-2011, and I plot the resulting IRFs for these normal times in Figure 4.[29] Contrary to the previous reasoning, the average contractionary effects of an oil shock are much *weaker* in normal times than at the ZLB. However, the confidence intervals for the ZLB response are large and typically show substantial overlap with the IRF for normal times, so that one cannot reject the hypothesis that the responses are the same. Nonetheless, these outcomes do suggest that little support exists for this weaker prediction of the expectations channel.

[28] Long-run effects for IP and unemployment are even smaller when I employ HP-filtered data as in Appendix B.

[29] I apply 1985 as a cutoff for three reasons. First, this period broadly features a common monetary regime. Second, it excludes the major oil shocks around and before 1980, so that my estimates do not pick up non-linearities from these events. Third, since unemployment is very persistent in the late 1970s and early 1980s, this cut-off avoids possible non-stationary estimates. Estimates are similar for other cut-offs around 1985.

In Appendix B I show that these results are robust to controlling for inflation and spreads, dropping lagged dependent variables, excluding outliers, using HP-filtered IP and unemployment, using non-seasonally-adjusted consumption expenditures, and alternative lag lengths.

4 Inflation Risk Premia at the ZLB

The previous two sections have shown that two types of negative supply shocks — an earthquake, and oil supply shocks — have contractionary effects at the ZLB. In this section I present evidence that inflation risk premia are positive at the ZLB, which suggests that generic negative supply shocks are also contractionary during such episodes.

4.1 CCAPM Theory The inflation risk premium is determined by the conditional covariance of marginal utility and inflation,

$$rp_\pi(t) = \mathbb{C}\text{ov}_t\left(\frac{dV(t)}{V(t)}, d\pi(t)\right).$$

where $V(t)$ is the marginal utility of consumption. Thus, if marginal utility is high when inflation is high, then the inflation risk premium is positive. In this case, agents dislike nominal assets because they tend to be of low value just as the desire to consume is high. These assets must therefore pay the inflation risk premium to induce an agent to hold them. As in Section 2, I assume that the marginal utility of consumption is equal to $V(t) = \zeta C(t)^{-\gamma^{-1}}$, so that the inflation risk premium is proportional to the covariance of consumption and inflation:[30]

$$rp_\pi(t) = -\gamma^{-1}\mathbb{C}\text{ov}_t\left(\frac{dC(t)}{C(t)}, d\pi(t)\right).$$

[30] I can also allow for shocks to ζ. In an Euler equation framework, the risk premium for this shock will be positive in normal times and negative at the ZLB — with marginal costs given by $W(t)/P(t) = \chi L(t)^\eta/V(t)$, the increase in consumption from a positive ζ shock will raise real wages and trigger higher inflation if $\eta > 0$. In normal times real interest rates rise, so that higher inflation is associated with higher marginal utility. In contrast, at the ZLB higher inflation lowers real interest rates so that marginal utility is now lower.

Thus, the inflation risk premium is positive if inflation is associated with low consumption, because then the marginal utility of consumption is high. A positive inflation risk premium then signals that consumers perceive higher-than-expected inflation as bad. Conversely, if the inflation risk premium is negative, then higher inflation is considered good news.[31]

The determinants of the risk premium are the fundamental shocks in the economy. Consider two types of shock: demand shocks $dB_v(t)$ raise log consumption and inflation, $c_v(t) > 0$ and $\pi_v(t) > 0$, and supply shocks $dB_u(t)$ lower log consumption and raise inflation, $c_u(t) < 0$ and $\pi_u(t) > 0$ (subscripts denote partial derivatives). Assuming that these shocks are independent, the risk premium is given by,

$$rp_\pi(t) = -\gamma^{-1}\sigma_u^2 c_u(t)\pi_u(t) - \gamma^{-1}\sigma_v^2 c_v(t)\pi_v(t), \tag{9}$$

where σ_v^2 and σ_u^2 are the variances of the shocks. The first term is positive, since supply shocks imply that high inflation is associated with low consumption. The second term is negative, as demand shocks induce the opposite pattern. A priori, the sign of the inflation risk premium is ambiguous, but with data on the risk premium we can determine what kind of shocks are quantitatively more important. In particular, a positive risk premium suggests that investors primarily want to insure against contractionary, negative supply shocks.

However, if the theory in Section 2 is correct and a broad class of negative supply shocks are expansionary at the ZLB, then we would expect the inflation risk premium to be unambiguously negative in this state of the world — when $c_u^{ZLB}(t) > 0$, both supply shocks and demand shocks should generate a positive covariance of inflation and consumption. Intuitively, in a standard Euler equation framework inflation is *always* good at the ZLB, because it lowers real interest rates today so that consumption expands through the expectations channel. In contrast, if the standard theory does not hold and negative supply shocks are

[31]This intuition is similar to the Phillips Curve logic emphasized by Campbell, Sunderam, and Viceira [2012]. In their paper, they estimate a term structure model to capture time-varying risk premia in the U.S. In contrast, I show how the risk premium loads on different types of macroeconomic shocks in a CCAPM framework (with emphasis on the ZLB), and calculate short-term inflation risk premia for several countries.

predominantly contractionary at the ZLB, $c_u^{ZLB}(t) < 0$, then we may still observe a positive inflation risk premium. In fact, evidence of positive inflation risk premia at the ZLB suggests that contractionary, negative supply shocks (such as oil supply shocks) are quantitatively important at the ZLB.

Summary 4 *The standard theory predicts that inflation risk premia should be negative at the ZLB. If instead the inflation risk premium is positive, then this suggests that the standard model is incorrect and that negative supply shocks are typically contractionary at the ZLB.*

In Appendix D, I show that this intuition from the instantaneous inflation risk carries over to longer-horizon risk premia, and is robust to using Epstein-Zin preferences.[32]

4.2 Estimated Inflation Risk Premia I estimate inflation risk premia from one-year inflation swap contracts. These are financial instruments where one party pays the *ex ante* known swap rate and the other party pays the realized inflation rate over the contract duration. By comparing the inflation swap rate with *ex ante* inflation expectations, I derive the inflation risk premium over this horizon.

I focus on the one-year inflation risk premium in order to minimize the confounding of ZLB effects with normal times — that is, to minimize spillovers. In Appendix D, I show that the short end of the risk-premium curve loads more heavily on shocks that adjust quickly. For example, the one-year risk premium loading is maximized for shocks with a half-life of six months, and over this horizon there are unlikely to be significant spillovers from the ZLB to normal times. On the other hand, shocks with a half-life of more than two years receive less than half this weight, and are thus unlikely to dominate the inflation risk premium at this horizon. Evidence of positive one-year inflation risk premia at the ZLB is then suggestive of contractionary effects of negative supply shocks that are not due to spillovers.

[32] Piazzesi and Schneider [2006] show that Epstein-Zin utility is a key ingredient for rationalizing an upward sloping yield curve given the historical correlations of consumption and inflation. In addition, Rudebusch and Swanson [2012] show that a medium-scale New Keynesian model with Epstein-Zin preferences can broadly match macro facts as well as the size and variability of the term premium.

Due to the illiquidity of the Japanese one-year inflation swap market and the absence of matching inflation expectations data, I restrict my attention to the U.S., the U.K., and the Eurozone. For these countries, I match inflation swap rates with inflation expectations adjusted for the indexation lag. For example, in the U.K. the indexation lag is two months, so that a December inflation swap contract will pay the realized inflation rate from October in that year to October in the following year. This indexation lag allows an instantaneous payout at the time of settlement, given that official CPI data is published with delay. The exact specification of the indexation lag differs across countries, and I relegate the details of matching inflation swaps with inflation expectations to Appendix E.

With matched inflation swap rates and inflation expectations, I calculate the inflation risk premium as the simple difference of the two,[33]

$$rp_\pi(t, t+1) = s^\pi(t, t+1) - \pi^e(t, t+1),$$

where $s^\pi(t, t+1)$ is the one-year inflation swap rate and $\pi^e(t, t+1)$ is the appropriately-matched one-year expected inflation rate. I plot the resulting risk premia in Figure 5(a) for the U.K., in Figure 5(b) for the Eurozone, and in Figure 5(c) for the U.S. I do find positive risk premia for the U.K. and the Eurozone at the ZLB, whereas the inflation risk premia for the U.S. are consistently negative. The U.K. inflation risk premia at the ZLB is 25 basis points on average, and in fact is larger than during normal times — suggesting that investors are now more concerned about the contractionary effects of negative supply shocks (at least over the one-year horizon). A possible explanation is news in late April 2010 of a consumption tax hike to be implemented in January 2011, the uncertain macro-effects of which could have been a considerable risk to investors. In the Eurozone the estimated risk premia at the ZLB are consistently positive except for two quarters, and on average are equal to 18 basis points — again indicating the desire to insure against contractionary

[33]This calculation ignores the uncertainty over inflation, which overestimates the risk premium. However, this effect is likely to be small — a 0.5 percentage point standard deviation in the one-year inflation rate results in a correction of $1/4$ basis point assuming a normal distribution.

negative supply shocks.

One explanation for the persistent negative risk premia in the U.S. is a liquidity premium. It is well known that U.S. nominal treasuries are treated as "special" in the financial markets (Krishnamurthy and Vissing-Jorgensen [2012]). This makes TIPS relatively inexpensive (Fleckenstein, Longstaff, and Lustig [2010]), which reduces the Treasury-TIPS yield spread. Since this spread replicates the payoff of an inflation swap, by no-arbitrage the liquidity premium will reduce inflation swap rates and therefore also inflation risk premia. For example, Pflueger and Viceira [2011] estimate liquidity premia in the 10-year Treasury-TIPS spread on the order of 30 basis points in normal times and up to 150 basis points during the financial crisis. The post-Lehman decline in U.S. inflation risk premium is also much larger than in the U.K. and the Eurozone, which suggests that liquidity factors are more important in the U.S. Consistent with this logic, I apply a Pflueger and Viceira [2011] liquidity correction to the one-year risk premium, also displayed in Figure 5(c) (details of the liquidity correction are relegated to Appendix F). The liquidity-corrected inflation risk premium displays a much milder decline during the financial crisis, as well as more episodes with positive risk premia. This evidence suggests that factors other than the demand and supply shocks considered above drive the negative U.S. inflation risk premia at the ZLB. This warrants further investigation, but is beyond the scope of this paper.

In summary, evidence from the U.K. and the Eurozone indicates that investors want to insure against contractionary negative supply shocks even at the ZLB. This in turn suggests that negative supply shocks are not only contractionary at the ZLB, but are also quantitatively important to the marginal investor. Since the one-year risk premium loads primarily on transient shocks, these results are unlikely to be due to spillovers from the ZLB into normal times, and cannot be accommodated by the standard Euler equation framework.

5 A Model with Financial Frictions

The previous two sections examine two examples of contractionary negative supply shocks, and show that these shocks are a quantitatively-important part of inflation risk at the ZLB. Since these results are inconsistent with the standard Euler equation, I build a model with a more general Euler equation that can match the data.

Recall from Figure 1(b) that a negative, temporary supply shock raises consumption at the ZLB, because it lowers real interest rates. To overcome this prediction and match the data, I allow supply shocks to have *endogenous* demand effects through a channel other than expected inflation, as shown in Figure 6: if a negative supply shock directly affects the Euler equation such that the AD curve shifts to the left, then it can be contractionary even at the ZLB. To capture this possibility, I add a wedge $\vartheta(t)$ to the standard Euler equation. For example, this could represent a spread on interest rates faced by consumers over the safe rate set by the central bank. With this wedge, the consumption response to the supply shock $u(t)$ is the following:

$$\frac{\mathrm{d}\ln C(t)}{\mathrm{d}u(t)} = -\gamma \frac{\mathrm{d}\mathbb{E}_t \int_0^\infty [i(t+s) - \pi(t+s)]ds}{\mathrm{d}u(t)} - \gamma \frac{\mathrm{d}\mathbb{E}_t \int_0^\infty [\vartheta(t+s)]ds}{\mathrm{d}u(t)}.$$

To match the contraction of consumption in the data, this demand shifter $\vartheta(t)$ must be endogenous to the supply shock, and must offset the decline in real interest rates from higher expected inflation.

There are several plausible mechanisms that can generate this dependency. For example, if nominal wages are sticky and a large fraction of consumers are hand-to-mouth, then a negative supply shock can be contractionary because it lowers real income for these consumers. Another possibility is that negative supply shocks directly raise uncertainty in the economy, which will raise precautionary savings for a given real interest rate (i.e., raise $\vartheta(t)$). Fernández-Villaverde, Guerrón-Quintana, Rubio-Ramirez, and Uribe [2011a] and Basu and Bundick [2011] illustrate this mechanism, but treat the uncertainty shock as an exogenous impulse. In my model, a negative supply shock affects $\vartheta(t)$ by raising the borrowing costs of consumers. In particular, a negative supply shock lowers firm profits and asset values, which reduces net worth at banks. Banks will then reduce their lending, which is reflected in a higher spread on loans over the safe interest rate (i.e., a larger $\vartheta(t)$).

While other mechanisms may be important in explaining my empirical findings, I focus on credit frictions as determinants of the friction $\vartheta(t)$, for four reasons. First, credit frictions have arguably played a prominent role in ZLB episodes such as the Japanese lost decade, in the current crisis, and the Great Depression. For example, Friedman [2000] argues that "*the enormous loan losses and balance sheet erosion that nearly all Japanese banks have sustained during this period, and the resulting impairment of their ability to carry out ordinary credit creation activities, have been both a consequence and a cause of Japan's prolonged economic stagnation.*" Regarding the recent crisis, both Hall [2010] and Woodford [2010] argue that frictions to financial intermediation are crucial for understanding the output collapse following the bankruptcy of Lehman Brothers, while Bernanke [1983] is a classic reference on credit frictions in the Great Depression.

Second, a negative oil supply shock tightens credit standards for firms and households in Japan (Figure 7). I regress credit standards (reported in the Japanese loan officer survey) on four of its own lags, on oil shocks, and on recession/recovery dummies.[34] The tightening of credit availability persists for about two years for both firms and households, and (consistent with this interpretation) credit tightens more for small firms than for medium firms, which in turn are more affected than large firms. These results indicate that a decline in credit supply is contributing to the contraction following negative supply shocks.[35]

Third, empirically-reasonable credit frictions are sufficient to rationalize my empirical findings. In Section 5.4, I estimate the impact of utilization-adjusted technology shocks

[34] I conduct the regression in levels because the data is already reported as changes in credit conditions.

[35] I also find evidence that long-term loan and discount rates in Japan rise following a negative oil supply shock, whereas short-term loan and discount rates are approximately unchanged. Thus, an increase in loan spreads appears to offset the decline in policy rates, so that consumers face higher borrowing costs following a negative oil supply shock.

from Fernald [2009] on asset spreads, and show that the calibrated model is quantitatively consistent with these estimates. The calibrated credit frictions in turn suffice to generate contractionary effects from negative supply shocks at the ZLB.

Fourth, the model with financial frictions can generate positive inflation risk premia at the ZLB. This is a non-trivial result. For example, assets in the hand-to-mouth and sticky-wage model are priced by the stochastic discount factor of unconstrained agents. Since these unconstrained agents follow a standard Euler equation, their inflation risk premium is negative at the ZLB, because higher inflation always raises their consumption in this state. An advantage of my set-up is that a single microfounded mechanism generates both contractionary effects from negative supply shocks and positive inflation risk premia at the ZLB.

5.1 Households The friction model features a continuum of households of two types, b and l, which split into borrowers and lenders in equilibrium. I follow Cúrdia and Woodford [2009] to deal with a number of technical issues that arise in this set-up. In particular, I assume that a household type is stochastically re-selected according to a Poisson process with intensity δ. A household hit by a re-selection shock draws type l with probability p and type b with probability $1-p$, so that these probabilities are also the frequency of each type. In addition, whenever a new (but not necessarily different) type is drawn, an insurance contract resets household wealth to the average wealth level.[36] This intermittent consumption insurance across households guarantees that all households of the same type behave identically.

Each household can deposit at a financial intermediary at a nominal interest rate $i(t)$, or borrow at a nominal interest rate $i(t)+\omega(t)$, where ω is a spread between the borrowing and the deposit rate. I assume that households cannot engage in any other financial contract, so that financial intermediaries price all risky assets. Intertemporal optimization by lenders

[36] The optimal insurance contract is contingent on an agent's behavior before the contract is invoked, but the reset of a household's wealth to the average level is the equilibrium outcome. See Cúrdia and Woodford [2009] for details.

and borrowers generates an Euler equation for each type,

$$\mathbb{E}_t dV^l(t) = [\varrho - i(t) + \pi(t)]V^l(t)dt - \delta(1-p)[V^b(t) - V^l(t)]dt,$$
$$\mathbb{E}_t dV^b(t) = [\varrho - i(t) + \pi(t) - \omega(t)]V^b(t)dt - \delta p[V^l(t) - V^l_b(t)]dt,$$

where $V^x = \zeta^x/C^x$ is the marginal utility of type $x \in \{l, b\}$ assuming log-utility, and ζ^x is a constant marginal utility shifter. The marginal utility of borrowers and lenders will differ, because they face different interest rates due to the spread $\omega(t)$, which enters only a borrower's Euler equation. In addition, agents take into account that they change their type with intensity δ, which generates the additional capital gain terms in both Euler equations. Setting $\zeta^b > \zeta^l$ will make consumption by borrowers more valuable, which induces these households to borrow, provided the borrowing spread is not too large.

Households supply perfectly-substitutable labor in a competitive market. The common real wage, $W(t)/P(t)$, is then equal to the marginal rate of substitution between consumption and leisure for both households,

$$\frac{\chi^x L^x(t)^\eta}{V^x(t)} = \frac{W(t)}{P(t)}, \tag{10}$$

where η is the inverse Frisch elasticity, and χ^x is a scalar such that steady-state labor supply is equal to 1 for both household types — $\bar{L}^x = 1$.

5.2 Firms The supply side of the model features standard new Keynesian price setting: a continuum of monopolistically-competitive firms produces output with a linear production function, $y(i,t) = \bar{A}e^{a(t)}l(i,t)$, where $a(t)$ is an exogenous technology process and $l(i,t)$ is labor input. The technology shock follows an Ornstein-Uhlenbeck process with mean-reversion ρ_a and volatility σ_a,

$$da(t) = -\rho_a a(t)dt + \sigma_a dB_a(t).$$

Given constant demand elasticity of ε across all varieties and a Calvo intensity α, the producers will choose current prices to maximize their share price. Rather than characterize this non-linear problem, I simplify the analysis by restricting the supply side to first-order effects, which eliminates price dispersion as a state variable. Thus, aggregate output is linear in labor, $Y(t) = e^{a(t)}L(t)$, and the continuous-time New Keynesian Phillips Curve is given by

$$\mathbb{E}_t d\pi(t) = \tilde{\varrho}\pi(t)dt - \kappa(\ln W(t)/P(t) - \ln \overline{W/P})dt + \kappa a(t)dt,$$

where $\tilde{\varrho}$ is the discount rate of bankers (who own firm shares), and $\ln \overline{W/P}$ is the real wage in the stochastic steady state. Intuitively, when marginal costs are above steady state, inflation is high and falling — that is, $\mathbb{E}_t d\pi(t) < 0$. The real profits of firms are the output net of labor costs, $F(t) = Y(t) - W(t)L(t)/P(t)$, which are distributed to financial intermediaries according to share ownership.

Finally, I assume that the central bank follows a simple interest rate rule subject to the ZLB constraint,

$$i(t) = \max\{i_0 + \phi_\pi \pi(t) + \phi_y(\ln Y(t) - a(t)), 0\},$$

where i_0 is set such that the trend inflation rate in the stochastic steady state is zero, and $\ln Y(t) - a(t)$ is the output gap.

5.3 Banking Sector The banking sector will drive the key results in this section. I assume that each *lending* household consists of a continuum of agents of two types. As in Gertler and Kiyotaki [2010], Gertler and Karadi [2011], and Maggiori [2011], I let a fraction f be workers and a fraction $1 - f$ be bankers. While workers earn labor income, each banker manages a financial intermediary. There is complete consumption risk sharing among these household members, and an exogenous Poisson process switches the roles of $\lambda(1-f)dt$ workers and $\lambda f dt$ bankers in each instant. When bankers switch to workers, they pay back their accumulated net worth to the household, while workers who become bankers are endowed with start-up funds. This ensures that bankers will not accumulate enough net worth to void the financial frictions in this section. To the same effect, a household may not deposit funds at its own intermediaries, and must instead spread its assets evenly over banks managed by other households. Finally, I assume that when a household becomes a borrower, its bank is taken over by another household that has become a lender. This implies that net worth is evaluated using the discount factors of lenders, which simplifies the optimization problem.

Only bankers are able to hold risky assets and to intermediate between households. Each banker's balance sheet is given by

$$Q(t)s(t) + b(t) = n(t) + d(t),$$

where the lower-case variables are bank-specific deposits $d(t)$, net worth $n(t)$, real loans to borrowers $b(t)$, and shares $s(t)$. $Q(t)$ is the real price of shares, which pay firm profits $F(t)$ as dividends. While bankers have specific financial skills, they are subject to an incentive compatibility constraint — they can walk away and divert a fraction θ of all assets, at which point the intermediary is wound down and depositors recover the remaining assets. To ensure that bankers do not walk away, the value of the bank must exceed the value of divertible assets. As in Gertler and Kiyotaki [2010], I assume that this incentive constraint is always binding in equilibrium.[37] I can therefore write the optimization problem of the lender directly as

$$\begin{aligned} \mathcal{V}^B(t) &= \max_{\{s(t+s),b(t+s),d(t+s)\}_{s=0}^{\infty}} \int_{s=0}^{\infty} \lambda e^{-(\lambda+\varrho)t} V^l(t+s)n(t+s)ds \\ \text{s.t. } dn(t) &= s(t)[dQ(t) + F(t)] + [i(t) - \pi(t) + \omega(t)]b(t) - [i(t) - \pi(t)]d(t) \\ \mathcal{V}^B(t) &= \theta[Q(t)s(t) + b(t)] \;\forall t, \end{aligned}$$

where $\mathcal{V}^B(t)$ is the value of the financial intermediary.

[37]There are subsets of the state space where this assumption is violated, but in my calibrations these occur with negligible probability in the stationary equilibrium. Thus, explicitly incorporating these states is unlikely to affect the qualitative and quantitative features of the model. This contrasts with Maggiori [2011], who explicitly incorporates when the incentive constraint binds, because in his set-up there is a significant probability of reaching the boundaries of the state space.

Following Gertler and Kiyotaki [2010] and Maggiori [2011], I guess and verify that the value of the intermediary is a linear function of bank-specific net worth, $\mathcal{V}^B(t) = \Omega(t)n(t)$, where $\Omega(t)$ is the marginal value of net worth. Intuitively, the value function is linear because an atomistic intermediary with twice the net worth can double its value by scaling up the starting portfolio while also satisfying the incentive constraint. Given this solution, the leverage of a financial intermediary is equal to $\phi(t) = \Omega(t)/\theta$. Thus, banks are allowed to lever up more when their marginal value of net worth is high, because then the bank's value is high and the incentives to steal from other households are mitigated.

The optimality condition for stock holdings of banks determines the excess return of stocks over bonds,

$$\mathbb{E}_t\left(\frac{dQ(t)+F(t)}{Q(t)} - r(t)\right) = -\mathbb{E}_t\frac{dQ(t)}{Q(t)}\frac{dV^l(t)}{V^l(t)} - \mathbb{E}_t\frac{dQ(t)}{Q(t)}\frac{d\Omega(t)}{\Omega(t)} + \theta\frac{\mu(t)}{\Omega(t)}, \tag{11}$$

where $\mu(t)$ is the shadow value of the incentive constraints for banks. The RHS shows that stocks can earn a positive excess return for three reasons. First, if stock prices and marginal utility have negative covariance, this generates a risk premium because stocks are worth less when agents want to consume more (the first term in the equation). Second, if stock prices and the marginal value of net worth have negative covariance, then this also generates a positive risk premium because banks would rather not hold assets that lose value when the marginal value of net worth is high (the second term). Third, stocks may earn an additional excess return because the incentive compatibility constraint prevents banks from purchasing more stock even when the latter are relatively cheap (the third term). A standard CCAPM model would incorporate only the first term, while Maggiori [2011] features both the first and second terms. The third term is specific to this model because the incentive constraint binds in equilibrium.

The FOC for lending to borrowers in turn pins down the spread of the borrowing rate over the deposit rate:

$$\omega(t) = \theta \frac{\mu(t)}{\Omega(t)}. \tag{12}$$

Because lending to borrowers is risk-free, the spread depends only on the tightness of the incentive constraint, and not on any covariances. If the incentive compatibility constraint is very tight due to low net worth, then $\mu(t)$ is high and the spread is large, as depositors do not permit increased lending by banks.[38] The endogenous demand effects of supply shocks will unfold through this channel. As I discuss in more detail below, a negative supply shock lowers profits for reasonable calibrations, because the increase in marginal costs swamps the impact on profits from changes in aggregate sales. With lower profits, asset prices decline — which reduces net worth, tightens the incentive constraint, and raises the spread on borrowers.

This set-up can also match positive inflation risk premia at the ZLB because assets are priced by banks, which face the additional balance sheet risk. In particular, the instantaneous inflation risk premium is given by

$$rp(t) = \mathbb{C}\mathrm{ov}_t \left(\frac{dV^l(t)}{V^l(t)}, d\pi(t) \right) + \mathbb{C}\mathrm{ov}_t \left(\frac{d\Omega(t)}{\Omega(t)}, d\pi(t) \right),$$

where the first term is the familiar CCAPM risk from Section 4, and the new second term captures the risk of inflation on the balance sheets of banks. Since negative supply shocks lower asset prices and net worth — and thus raise the marginal value of net worth together with inflation — this second term will be positive even at the ZLB. This contrasts with the first term, which is negative at the ZLB because negative supply shocks raise both inflation and (lender) consumption in the standard model. If the balance-sheet risk is sufficiently strong, then it can offset the negative CCAPM risk premium and generate positive inflation risk premia as in the data.

Plugging the FOCs back into the Bellman equation shows that the guess is verified if the

[38] This variation dominates the variation in $\Omega(t)$ because the latter is close to 1, whereas $\mu(t)$ is close to zero.

marginal value of net worth satisfies the following equation:

$$\lambda\frac{1-\Omega(t)}{\Omega(t)}dt+\mathbb{E}_t\frac{d\Omega(t)}{\Omega(t)}+\mathbb{E}_t\frac{d\Omega(t)}{\Omega(t)}\frac{d\Lambda^l(t)}{\Lambda^l(t)}+\delta(1-p)\left(\frac{V^b(t)}{V^l(t)}-1\right)dt+\mu(t)dt=0.$$

For example, $\Omega(t)=1$ is the unique bounded solution when there are no frictions, $\mu(t)=0$ and $V^b(t)/V^l(t)=1$. Intuitively, in a frictionless world the marginal value of inside funds should equal the marginal value of outside funds (=1). However, if the banks' incentive constraint binds, then $\mu(t)>0$ and the marginal value of net worth will be higher, $\Omega(t)>1$.

In addition, the banking sector adds two state variables to the standard model. First, aggregate net worth is the sum of net worth from surviving banks and newly-founded banks,

$$\begin{aligned}dN(t)=&[i(t)-\pi(t)-\lambda]N(t)dt+\left(\frac{dQ(t+s)+F(t+s)}{Q_t}-i(t)-\pi(t)\right)S(t)Q(t)dt\\&+\omega(t)B(t)dt+dN^{new}(t),\end{aligned}$$

where uppercase letters denote aggregate variables. New banks are endowed with net worth equal to a fraction of existing shares $dN^{new}(t)=\nu Q(t)$, while existing banks accumulate net worth through deposit interest on net worth and through excess returns on stocks and lending. The second new state variable is aggregate real borrowing,

$$dB(t)=[-\delta+i(t)-\pi(t)+\omega(t)]B(t)dt+(1-p)[C^b(t)-W^b(t)]dt-\delta(1-p)Q(t)dt,$$

which rises when the consumption of borrowers $C^b(t)$ exceeds their wage rate $W^b(t)$, and declines when inflation reduces the real value of nominal debt. In addition, the insurance agency nets out debt and assets for a fraction δ of borrowers and lenders, and transfers real resources equal to $\delta(1-p)Q(t)$ over to new borrowing households.

5.4 Calibration The calibrated parameters are listed in Table 3. The household parameters are taken from Cúrdia and Woodford [2009]. The share of lenders is $p=0.5$, the type-switching intensity is $\delta=0.1$, the discount rate is $\varrho=0.0125$, and the steady-state borrowing spread is $\bar{\omega} = 0.02$. The discount rate is lower than in typical calibrations to make the ZLB bind for small shocks. The marginal utility shifters are $\zeta^l = 0.72$ and $\zeta^b = 1.34$, which implies that steady-state debt equals 80% of GDP as in their calibration.

I calibrate the elasticity of substitution across goods at $\varepsilon = 40$, to generate a 2.5% steady-state profit share consistent with the evidence in Basu and Fernald [1997]. This value is relatively high, but its only function is to match the profit share. The slope of the Phillips Curve is separately parameterized at $\kappa = 0.15$, based on estimates from Altig, Christiano, Eichenbaum, and Lindé [2011].[39] Erceg and Lindé [2010] show that DSGE models require very rigid prices and wages at the ZLB to match the small decline in inflation in the recent crisis, which suggests that the estimate from Altig et al. [2011] — which was estimated on data for normal times — is likely to be an upper bound. Since lower values for κ weaken the inflation expectations channel, a calibration along the lines of Erceg and Lindé [2010] will strengthen my results.

The average banking horizon λ^{-1} is set to 10 years, and the steady-state leverage ratio to $\bar{\phi} = 2$, similar to Gertler and Kiyotaki [2010]. As in their calibration I choose a relatively high Frisch elasticity of $\eta = 2$ to compensate for the lack of labor-market frictions. The fraction of divertible assets, $\theta = 0.7$, is determined by steady-state borrowing spread, and the infusion of net worth to new banks, $\nu = 0.05$, is determined by the steady-state leverage ratio. The interest rate rule parameters for inflation and the output gap are set to $\phi_\pi = 1.5$ and $\phi_y = 0.5$, as is common in the literature.

To validate the strength of the financial accelerator, I determine whether the model is quantitatively consistent with the behavior of asset market spreads in the data. Consistent with the interpretation of $a(t)$ in the model, I use the utilization-adjusted TFP series for the U.S. from Fernald [2009].[40] These shocks are approximately uncorrelated, so I set $\rho_a = 0$

[39] Altig, Christiano, Eichenbaum, and Lindé [2011] estimate the parameter $\tilde{\kappa} = 0.014$ in a discrete time model at quarterly frequency. Using $\tilde{\kappa} = (1 - \tilde{\alpha})(1 - \beta\tilde{\alpha})/\tilde{\alpha}$ I back out the implied price stickiness $\tilde{\alpha}$ and calculate its continuous-time counterpart $\alpha = (1 - \tilde{\alpha}^4)$. The continuous-time parameter κ is then set to $\kappa = \alpha(\alpha + i_0 + \bar{\phi}\bar{\omega}) \approx 0.15$.

[40] Verifying with technology shocks is a closer match with $a(t)$ in the model than using the oil supply shock series. However, since the Fernald-series is only available for the U.S., I was unable to use it in the earlier

for this exercise. I use the TED spread as benchmark, because in the model interbank loans are subject to the same friction as consumer loans (see Gertler and Kiyotaki [2010]). In Figure 8(a) I report the effect of a 1% negative TFP shock on the TED spread, estimated from an ADL equation with 12 lags. The TFP shock raises the TED spread by up to 24 basis points after eight quarters. I compute the corresponding model statistic by linearizing the model around the non-stochastic steady state. The IRF is plotted in Figure 8(b), and shows that a 1% permanent TFP shock raises the borrowing-lending spread by 25 basis points on impact, which subsequently declines to zero. This suggests that the credit frictions in this calibration are empirically reasonable.

5.5 Computational Strategy To evaluate the model's fit at the ZLB, I must determine its non-linear solution. To make the solution and mechanics of the model as intuitive as possible, I use a short-cut to reduce the state space from the three state variables $a(t)$, $N(t)$, and $B(t)$ to only the variable $a(t)$. Specifically, I impose that the marginal value of net worth is inversely related to asset prices,

$$\frac{\Omega(t)}{\bar{\Omega}} = \left(\frac{Q(t)}{\bar{Q}}\right)^{-\xi}. \tag{13}$$

where $\xi > 0$. For example, with this short-cut I can plot consumption and other variables as global functions of the supply shock, which is more transparent than illustrating local results through partial derivatives or impulse-response functions. In addition, this assumption greatly simplifies the computation of the problem.[41]

Importantly, the economics of the problem makes Equation (13) a very good approximation for a judiciously-chosen ξ, which I discuss below. In particular, Equation (13) captures the following logic in the three-state-variable model. A decline in asset prices reduces net

analysis.

[41] Fernández-Villaverde, Gordon, Guerrón-Quintana, and Rubio-Ramírez [2012] solve a five-state-variable model at the ZLB in discrete time using the Smoliak algorithm. However, in my model two state variables follow endogenous processes (as opposed to only one in their model), which significantly complicates the iterative solution procedure.

worth more than proportionally because banks are levered up. Without any change in asset-holdings, leverage at banks $\phi(t)$ will be higher. But, in equilibrium, banks are only permitted a higher leverage if the marginal value of net worth is high, as $\phi(t) = \Omega(t)/\theta$. Thus, banks have to cut back on their asset holdings, which will raise expected excess returns on assets and the marginal value of net worth. Note that the change in asset prices is still endogenous — the marginal value of net worth will decline only if a negative supply shock lowers asset prices. The parameter ξ controls the strength of this financial accelerator. Higher values of ξ accentuate the financial friction, because a given change in asset prices results in a larger change in the marginal value of net worth. Thus, the simplification in Equation (13) can capture the logic of a more complicated model without sacrificing the economics of the problem.

I calibrate ξ at a value such that the one-state-variable model closely mimics the three-state-variable model. First, I determine the linearized solution for the three-state-variable model, and calculate the IRF for output following a technology shock. Second, I set ξ such that the one-state-variable model best matches this IRF. In my baseline calibration, I set $\rho_a = 0.69$ and $\sigma_a = 0.005$, which results in $\xi = 0.6$. This persistence implies a half-life of one year, which is similar to the transient nature of the oil shocks in Section 3.2.

For this calibration I also calculate the non-linear solution to the three-state-variable model using a projection grid, and find that the approximation (13) performs very well in normal times and at the ZLB (Appendix G.1). Since this simplification performs well, I compute the one-state-variable model and compare the model output with a standard two-equation new Keynesian model.[42] The details of the non-linear projection methods are relegated to Appendix G.2.

5.6 Results Figure 9(a) plots output as a function of the technology shifter $a(t)$ for four cases: the standard new Keynesian model without ZLB, the standard model with ZLB,

[42]The discount rate for the standard new Keynesian economy is adjusted to make the ZLB bind for the same shock size. Otherwise, the calibration is identical.

the one-state-variable model with credit frictions but without ZLB (the "friction model"), and the one-state-variable model with both credit frictions and ZLB. Without ZLB, the standard model has an (almost) linear solution, in which consumption rises in response to technology shocks. Intuitively, the central bank lowers real interest rates in response to the deflation generated by a positive supply shock (Figure 9(b)), which raises consumption today. However, when the ZLB binds, the paradox of toil applies in the standard new Keynesian model. In this case, the decline in inflation from improved technology raises real interest rates, so that consumption today contracts. Conversely, if we start in ZLB territory, then reductions in technology are expansionary because they raise expected inflation and lower real interest rates.[43]

Aggregate consumption in the friction model is more responsive than in the standard model because of the financial accelerator. To see this, note that profits are primarily determined by variation in marginal costs. Linearized profits are given by

$$\hat{F}(t) = \hat{Y}(t)(1 - \overline{MC}) - \bar{Y}\widehat{MC}(t), \tag{14}$$

where the first term reflects variation in real output (sales), $\hat{Y}(t)$, and the second term reflects variations in marginal costs, $\widehat{MC}(t)$. Empirically, profit shares are small, so that real marginal costs $\overline{MC}$ are close to 1. Thus, the first term is close to zero, and variations in profits are dominated by variations in marginal cost rather than aggregate output.[44] Since the technology shock directly lowers marginal costs, higher profits ensue, asset prices rise (Figure 10(a)), and the marginal value of net worth falls (Figure 10(b)) — which is reflected in a lower borrowing spread (Figure 10(c)). Hence, outside the ZLB region both real interest rates and spreads fall in response to a positive technology shock, so that consumption rises more than in the baseline model where only the former effect is present. Note that

[43]In this setting, negative supply shocks are expansionary at the ZLB even though there are spillovers. A negative supply shock at the ZLB reduces the time to reach the stochastic steady state, and thus raises expected future nominal interest rates.

[44]For example, in models of rational inattention this fact makes firms pay much more attention to technology shock rather than to monetary policy shocks (Mackowiak and Wiederholt [2009]).

this financial accelerator is moderate in size: it increases the responsiveness of aggregate consumption by approximately 31% outside the ZLB region.

When the ZLB binds in the friction model, a positive supply shock is expansionary — unlike in the baseline model. In this case the decline in real borrowing rates from a lower borrowing spread raises consumption by borrowers, which dominates the fall in consumption by lenders from higher deposit rates. The financial accelerator now has a much larger effect on the equilibrium outcome, as is evident in Figure 9(a): output in the friction model is almost 3% higher than in the baseline model when the ZLB constraint binds and TFP is 5% above steady state. Thus, financial accelerator effects that are moderately sized in normal times can have a significant impact at the ZLB.

The intuition for this difference is similar to that in Christiano, Eichenbaum, and Rebelo [2011]: in normal times, a positive supply shock lowers borrowing spreads, so that the central bank lowers real interest rates *less* in the friction model than in the standard new Keynesian model. Thus, the central bank smooths out the financial accelerator through its interest rate policy. At the ZLB, however, a positive technology shock generates a deflationary spiral in the standard model: lower expected inflation raises real interest rates and lowers consumption, which in turn lowers expected inflation, and so forth. In the friction model, the decline in borrowing spreads arrests this downward spiral and thus has a large impact on equilibrium outcomes. The absence of active stabilization policy therefore has dramatically different effects in these two models.

In turn, these results suggest that matching moments of the data in normal times does not guarantee even a qualitative fit at the ZLB. For example, estimating a standard New Keynesian model on simulated data of the friction model in normal times will closely replicate the latter's output, inflation, and interest-rate behavior. However, this estimated new Keynesian model generates expansionary effects to negative supply shocks at the ZLB, and therefore cannot match the friction model even qualitatively during these times. On the other hand, while it is difficult to distinguish between models during normal times, the reverse is true at the ZLB because propagation mechanisms are not attenuated by the central bank's interest rate rule. Thus, the ZLB provides a good testing ground for models.

This calibration of the friction model can also rationalize positive inflation risk premia at the ZLB. In Figure 10(d), I plot the instantaneous inflation risk premium for the four models. In the standard model the inflation risk premium becomes negative when the ZLB binds, because higher inflation is associated with higher consumption. However, in the friction model the instantaneous inflation risk premium is always positive, even at the ZLB. Intuitively, higher inflation coincides with lower net worth and a higher marginal value of net worth in banks; banks are thus unwilling to hold assets that load on inflation unless they pay a positive risk premium. This balance-sheet risk premium dominates the CCAPM inflation risk premium in this calibration, which allows the model to match the data in the previous section.

While the model can generate contractionary effects from negative supply shocks at the ZLB, these are attenuated at the ZLB compared to normal times. In contrast, my estimates for oil shocks indicate that negative supply shocks are more contractionary at the ZLB (Figure 4). However, there are a number of plausible extensions that are likely to correct this shortcoming. For example, ZLB episodes tend to be severe recessions, where a larger-than-usual proportion of agents is likely to be credit-constrained. Thus, the financial friction is likely to be more important in such episodes, which will amplify the contractionary effects of negative supply shocks. Similarly, if negative supply shocks raise uncertainty then this will be more costly at the ZLB where consumption is low, and will lead to larger cutbacks by consumers. Thus, a more detailed model of the ZLB will likely be able to match this fact in the data, but at the cost of additional complexity.

In summary, calibrating the friction model with standard parameters can qualitatively rationalize contractionary effects of negative supply shocks and positive inflation risk premia at the ZLB. In this model, negative supply shocks lower net worth of banks, which translates into higher borrowing spreads — allowing negative supply shocks to have endogenous (negative) demand effects. In addition, banks demand a positive inflation risk premium at the ZLB, because high inflation coincides with a high marginal value of internal funds.

6 Policy Implications

In the previous section, I document that the friction model is more successful at replicating the contractionary effects of negative supply shocks at the ZLB than the baseline model. Hence, this class of models may be more appropriate for policy analysis at the ZLB than standard models, which omit endogenous credit spreads. In this section, I therefore explore how the friction model propagates demand-side policies. I consider two popular examples: forward guidance, and fiscal stimulus through government spending. I model the former through a disturbance ε_i in the Taylor rule,

$$i(t) = \max\{i_0 + \phi_\pi \pi(t) + \phi_y(\ln Y(t) - a(t)) + \varepsilon_i(t), 0\},$$

which follows an autoregressive process, $d\varepsilon_i(t) = -\rho_i \varepsilon_i(t)dt + \sigma_i dB_i(t)$. This correlated shock is a simple way to capture persistent deviations from the normal policy rule. Thus, a negative value for ε_i at the ZLB implies that interest rates are kept low for an extended period, which has a stimulative effect by reducing long-term bond yields. I set the mean-reversion of the shock as $\rho_i = 0.9$, which is equivalent to a quarterly persistence of 0.8 as is typically assumed in the literature (although results are similar for other plausible values of ρ_i).

Similarly, I model government spending as $G(t) = \bar{G}e^{g(t)}$, where $\bar{G}$ is the steady-state value of government spending and $g(t)$ also follows an autoregressive process $dg(t) = -\rho_g g(t)dt + \sigma_g dB_g(t)$. I set the mean-reversion for these processes to $\rho_g = 0.69$, and the share of government spending in output to $\bar{G}/\bar{Y} = 0.2$.[45] To focus on the first-order effects of these policies, I calibrate the volatilities of $dB_i(t)$ and $dB_g(t)$ as negligible.

[45]With $\rho_g = 0.69$, government spending process has the same half-life as an AR(1) process with quarterly persistence of 0.85 — midway between Christiano, Eichenbaum, and Rebelo [2011] and Woodford [2011]. I adjust the marginal utility shifters ζ such that the debt/GDP ratio equals 80%, as in the baseline model.

Figure 11(a) illustrates the effectiveness of forward guidance by plotting the percentage change in output on impact for each basis point reduction in the 10-year bond yield at various durations of the ZLB. In both the standard model and the friction model, forward guidance raises output and becomes increasingly more effective the longer the ZLB binds. This occurs because at longer durations a one-basis-point reduction in long-term bond rates raises expected inflation more and thus lowers expected real interest rates more. To see this, consider the first-round effects of forward guidance on marginal cost — that is, holding inflation fixed for now. I show that these first-round effects of forward guidance are larger at longer durations of the ZLB, which generates larger second-round increases in inflation expectations, larger second-round reductions in real interest rates, and so larger increases in output.

I normalize today's time to $T_0 = 0$, so that the time of exit, T_1, is the expected duration of the ZLB in Figure 11(a). Note that from T_1 to $T_2 = 10$, nominal interest rates have to fall a total of 10 basis points to reduce the 10-year bond yield by one point, $\int_{T_1}^{T_2} \Delta i(t)dt = 0.1$. This implies that, conditional on exit, the increase in consumption $\Delta \ln C(T_1) = -\gamma \int_{T_1}^{T_2} \Delta i(t)dt \equiv \bar{\Delta} > 0$ is independent of T_1 according to the Euler equation (1). This derivation uses the fact that the shock ε_i has largely dissipated when $T_2 - T_1 \geq 5$, so that $\int_{T_2}^{\infty} \Delta i(t)dt \approx 0$.[46] This is illustrated in Figure 12, where the increase in consumption due to forward guidance is the same at two different exit dates, T_1 and T_1'.

It remains to establish what happens before and after the exit date. Before the exit date nominal interest rates are unchanged at zero, and inflation is held fixed for now. Since there is no change in real interest rates, consumption is constant given the standard Euler equation, $\Delta \ln C(T < T_1) = \bar{\Delta}$. After the exit date T_1, the monetary policy shock dissipates at rate ρ_i. Thus, nominal interest rates increase and consumption declines at rate ρ_i, as shown in Figure 12. In summary, consumption rises by $\bar{\Delta}$ before T_1 and by less than $\bar{\Delta}$ after T_1, so that for higher T_1 — a longer duration at the ZLB — there is a larger average

[46] Accordingly, the results are similar for using longer maturities, such as $T_2 = 20$ or $T_2 = 30$.

increase in consumption. Since real wages are proportional to consumption, this implies a larger increase in expected marginal costs for higher T_1, which generates a larger second-round increase in expected inflation and a greater second-round decline in real interest rates. Thus, forward guidance is more effective when the ZLB binds for a long time.

In the friction model, this policy is less effective than in the standard model. The rise in marginal cost from higher wages reduces profit margins at firms, which depresses asset prices, lowers net worth, and raises the borrowing spread. Thus, in the friction model any basis point decline in the 10-year yield has only about 45-90% of its potency in the standard model (Figure 11(a)).

In practice one may be skeptical of large changes in marginal cost (in particular real wages) from forward guidance at the ZLB. For example, Erceg and Lindé [2010] argue that prices and wages in new Keynesian models need to be very sticky to match the small decline in inflation during the 2007-2009 recession. Thus, in Figure 11(a) I also plot the impact of forward guidance when real wages are fixed, rather than as given by Equation (10). In this scenario, forward guidance in the friction model is *more* effective than in the standard model — with fixed marginal cost, higher output raises profits and asset prices, resulting in lower borrowing spreads. With fixed real wages, however, forward guidance loses much of its potency because it no longer triggers increases in inflation from higher marginal cost. Thus, expected future real interest rates decline by less and output expands by less. In summary, the financial friction narrows the range of forward guidance "multipliers," and significantly dampens the effectiveness of forward guidance for calibrations where standard models indicate large gains.

The results for fiscal policy are similar. Fiscal stimulus is more effective for longer durations of the ZLB, because the spillovers to normal times become smaller. In the friction model the impact fiscal multiplier is approximately 10% smaller than in the standard new Keynesian model when wages are flexible (Figure 11(b)). As with forward guidance, fiscal stimulus at the ZLB raises real wages, which reduces profits and raises the borrowing spread through the financial accelerator. The financial friction thus dampens the effectiveness of fiscal policy. Conversely, the impact fiscal multipliers in the friction model with fixed real wages are slightly larger than in the standard model when the ZLB lasts beyond four years. In this case there are few spillovers to normal times, so the increase in profits raises asset prices and lowers borrowing spreads. For shorter ZLB durations, an increase in government spending also raises expected future nominal rates, which reduces asset prices today and raises borrowing spreads. Thus, when the standard model suggests multipliers in excess of 1, the friction model generates lower multipliers. However, if baseline model multipliers are close to 1, then the friction model raises the multiplier.

Note that my fiscal multipliers are small compared to most of the ZLB literature (Christiano, Eichenbaum, and Rebelo [2011], Woodford [2011]). The main reason for this is that prices and wages in my model are relatively rigid, which weakens the effect of government spending on real interest rates. Erceg and Lindé [2010] show that calibrations with very sticky prices and wages generate a fiscal multiplier close to 1 even when the ZLB lasts four years (see their Figure 6), whereas calibrating the standard new Keynesian model with more responsive prices and wages (e.g., $\kappa = 0.25$ and $\eta = 0.5$) can generate multipliers in excess of 3. Second, there is no investment in the model, and the investment response raises the fiscal multiplier as in Christiano, Eichenbaum, and Rebelo [2011]. Third, the preferences I use do not have labor-consumption complementarities, which can also boost the multiplier. However, while these aspects will affect the absolute magnitudes of the fiscal multiplier, they are unlikely to change the finding that the friction model generates an intermediate range of multipliers at the ZLB relative to the baseline model.

To summarize, the friction model typically generates a narrower range of policy multipliers at the ZLB than a standard model. When the baseline multipliers are high due to the expectations channel, they are dampened in the friction model because the large increase in marginal cost depresses firm profits and asset prices. On the other hand, when the expectations channel is weak because marginal costs are rigid, then the financial friction boosts the effectiveness of demand-side policies. Thus, in the context of this model, an intermediate range of policy multipliers appears most plausible when the central bank is constrained by the ZLB.

7 Conclusion

This paper shows that negative supply shocks are contractionary at the ZLB. What may sound tautological in fact contradicts the prediction of many macroeconomic models with a standard Euler equation and sticky prices. These models emphasize how negative supply shocks at the ZLB reduce expected real interest rates, which raises consumption today through standard intertemporal substitution. In this paper, I examine two examples of negative supply shocks — an earthquake and oil supply shocks — and show that these shocks are contractionary at the ZLB despite lowering expected future real interest rates. In addition, evidence of positive one-year inflation risk premia at the ZLB is inconsistent with the standard Euler equation framework with sticky prices, and suggests that contractionary, negative supply shocks are quantitatively important over this horizon. I rationalize these empirical observations by allowing supply shocks to have endogenous, negative demand effects. In my model, the supply shock reduces the net worth of banks, which raises borrowing costs, and thus reduces consumption by borrowers. In addition, the desire of banks to insure against balance-sheet contractions from negative supply shocks generates a positive inflation risk premium. Since the model provides a more data-consistent description of the ZLB, I examine its implication for demand-side policies and find that the model favors an intermediate range of policy multipliers.

The theoretical literature on the ZLB stresses how the standard models behave differently from how they do in normal times, primarily because of the inflation expectations channel. My result — that supply shocks remain contractionary at the ZLB despite lowering expected future real interest rates — suggests that more research is needed to empirically validate and quantify this mechanism. This is particularly important because policy discussions on forward guidance and fiscal stimulus appear to be based on the premise that large benefits can accrue through this channel, whereas my model — which matches the data — attenuates large multipliers. Pooling exogenous aggregate shocks across countries and using financial data appear to be promising avenues to verify such general-equilibrium predictions. Furthermore, it is important to determine why the standard models may give misleading predictions, as in the case of negative supply shocks. While this paper has presented evidence that credit constraints appear to be important, other mechanisms such as precautionary savings motives or hand-to-mouth consumers may also play a role. Understanding the relative quantitative importance of these channels at the ZLB is likely to be a promising area for future research.

References

Almunia, M., A. Bénétrix, B. Eichengreen, K. O'Rourke, and G. Rua, "From Great Depression to great credit crisis: similarities, differences and lessons," *Economic Policy*, 2010, *25* (62), 219–265.

Altig, D., L.J. Christiano, M. Eichenbaum, and J. Lindé, "Firm-specific capital, nominal rigidities and the business cycle," *Review of Economic Dynamics*, 2011, *14* (2), 225–247.

Bachmann, R., T.O. Berg, and E.R. Sims, "Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence," 2011.

Basu, S. and B. Bundick, "Uncertainty Shocks in a Model of Effective Demand," 2011.

— and J.G. Fernald, "Returns to scale in US production: Estimates and implications," *Journal of political economy*, 1997, *105* (2), 249–283.

Benhabib, J., S. Schmitt-Grohé, and M. Uribe, "Avoiding liquidity traps," *Journal of Political Economy*, 2002, *110* (3), 535–563.

Bernanke, B. and M. Gertler, "Agency Costs, Net Worth, and Business Fluctuations," *The American Economic Review*, 1989, *79* (1), 14–31.

Bernanke, B.S., "Nonmonetary Effects of the Financial Crisis in the Propagation of the Great Depression," *The American Economic Review*, 1983, *73* (3), 257–276.

—, M. Gertler, and S. Gilchrist, "The financial accelerator in a quantitative business cycle framework," *Handbook of macroeconomics*, 1999, *1*, 1341–1393.

Bhattarai, S., G. Eggertsson, and R. Schoenle, "Is Increased Price Flexibility Stabilizing? Redux," 2012.

Campbell, J.Y., A. Sunderam, and L.M. Viceira, “Inflation bets or deflation hedges? the changing risks of nominal bonds,” Technical Report, National Bureau of Economic Research 2012.

Cavallo, M. and T. Wu, “Measuring oil-price shocks using market-based information,” 2006.

Chodorow-Reich, Gabriel, Laura Feiveson, Zachary Liscow, and William Gui Woolston, “Does State Fiscal Relief during Recessions Increase Employment? Evidence from the American Recovery and Reinvestment Act,” *American Economic Journal: Economic Policy*, April 2012, *4* (3), 118–45.

Christiano, L., M. Eichenbaum, and S. Rebelo, “When Is the Government Spending Multiplier Large?,” *Journal of Political Economy*, 2011, *119* (1), 78–121.

Christiano, L.J., “The zero bound, low inflation, and output collapse,” *Northwestern University*, 2004.

_ , M. Eichenbaum, and C.L. Evans, “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 2005, *113* (1).

Coibion, O. and Y. Gorodnichenko, “Information rigidity and the expectations formation process: A simple framework and new facts,” Technical Report, National Bureau of Economic Research 2011.

Coibion, Olivier, “Are the Effects of Monetary Policy Shocks Big or Small?,” *American Economic Journal: Macroeconomics*, April 2012, *4* (2), 1–32.

Cúrdia, V. and M. Woodford, *Credit frictions and optimal monetary policy* number 278 2009.

DeLong, J.B. and L. Summers, “Is increased price flexibility stabilizing?,” *American Economic Review*, 1986, *76* (5), 1031–1044.

Driscoll, J.C. and A.C. Kraay, “Consistent covariance matrix estimation with spatially dependent panel data,” *Review of Economics and Statistics*, 1998, *80* (4), 549–560.

Eggertsson, G.B., “The paradox of toil,” *Staff Reports*, 2010.

_ , “What Fiscal Policy Is Effective at Zero Interest Rates?,” *NBER Macroeconomics Annual*, 2011, *25* (1), 59–112.

_ , “Was the New Deal Contractionary?,” *The American Economic Review*, 2012, *102* (1), 524–555.

_ and M. Woodford, “The Zero Bound on Interest Rates and Optimal Monetary Policy.,” *Brookings Papers on Economic Activity*, 2003, (1), 139–234.

_ and P. Krugman, “Debt, deleveraging, and the liquidity trap: a Fisher-Minsky-Koo approach,” *Federal Reserve Bank of New York, unpublished manuscript, February*, 2011.

Erceg, C.J. and J. Lindé, *Is there a fiscal free lunch in a liquidity trap?*, Centre for Economic Policy Research, 2010.

Farhi, E. and I. Werning, “Fiscal Multipliers: Liquidity Traps and Currency Unions,” 2012.

Fernald, J., “A quarterly, utilization-adjusted series on total factor productivity,” *Manuscript, Federal Reserve Bank of San Francisco*, 2009.

Fernández-Villaverde, J., G. Gordon, P.A. Guerrón-Quintana, and J. Rubio-Ramírez, “Non-linear Adventures at the Zero Lower Bound,” Technical Report, National Bureau of Economic Research 2012.

_ , P. Guerrón-Quintana, J.F. Rubio-Ramirez, and M. Uribe, “Risk matters: The real effects of volatility shocks,” *The American Economic Review*, 2011, *101* (6), 2530–2561.

_ , P.A. Guerrón-Quintana, and J. Rubio-Ramírez, “Supply-side policies and the zero lower bound,” Technical Report, National Bureau of Economic Research 2011.

Fleckenstein, M., F.A. Longstaff, and H. Lustig, “Why does the Treasury issue TIPS? The TIPS–Treasury bond puzzle,” Technical Report, National Bureau of Economic Research 2010.

Friedman, B., “Japan now and the United States then: Lessons from the parallels,” *Japan’s Financial Crisis and Its Parallels to US Experience*, 2000, pp. 37–56.

Gertler, M. and N. Kiyotaki, “Financial intermediation and credit policy in business cycle analysis,” *Handbook of Monetary Economics*, 2010, *3*, 547.

_ and P. Karadi, “A model of unconventional monetary policy,” *Journal of Monetary Economics*, 2011, *58* (1), 17–34.

Gordon, R.J. and R. Krenn, “The end of the Great Depression 1939-41: policy contributions and fiscal multipliers,” Technical Report, National Bureau of Economic Research 2010.

Gürkaynak, R.S., B. Sack, and J.H. Wright, “The US Treasury yield curve: 1961 to the present,” *Journal of Monetary Economics*, 2007, *54* (8), 2291–2304.

Gurkaynak, R.S., B. Sack, and J.H. Wright, “The tips yield curve and inflation compensation,” *American Economic Journal: Macroeconomics*, 2010, *2* (1), 70–92.

Hall, R.E., “By How Much Does GDP Rise If the Government Buys More Output?,” *Brookings Papers on Economic Activity*, 2009, *2*, 183–250.

_ , “Why does the economy fall to pieces after a financial crisis?,” *The Journal of Economic Perspectives*, 2010, *24* (4), 3–20.

Hamilton, J.D., *Time series analysis*, Vol. 2, Cambridge Univ Press, 1994.

Keynes, J.M., *The general theory of interest, employment and money*, London: Macmillan, 1936.

Kilian, L., “Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market,” *The American Economic Review*, 2009, *99* (3), 1053–1069.

_ and D.P. Murphy, "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," 2011.

Krishnamurthy, A. and A. Vissing-Jorgensen, "The aggregate demand for treasury debt," *Journal of Political Economy*, 2012, *120* (2), 233–267.

Krugman, P.R., "It's baaack: Japan's slump and the return of the liquidity trap," *Brookings Papers on Economic Activity*, 1998, *1998* (2), 137–205.

Lucas, R.E. and E.C. Prescott, "Investment under uncertainty," *Econometrica: Journal of the Econometric Society*, 1971, pp. 659–681.

Mackowiak, B. and M. Wiederholt, "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, 2009, *99* (3), 769–803.

Maggiori, M., "Financial Intermediation, International Risk Sharing, and Reserve Currencies," 2011.

Nakamura, E. and J. Steinsson, "Fiscal stimulus in a monetary union: Evidence from US regions," Technical Report, National Bureau of Economic Research 2011.

Pflueger, C.E. and L.M. Viceira, "An Empirical Decomposition of Risk and Liquidity in Nominal and Inflation-Indexed Government Bonds," Technical Report, National Bureau of Economic Research 2011.

Piazzesi, M. and M. Schneider, "Equilibrium Yield Curves," *NBER Macroeconomics Annual*, 2006, pp. 389–472.

Ramey, V.A., "Identifying Government Spending Shocks: It's all in the Timing," *The Quarterly Journal of Economics*, 2011, *126* (1), 1–50.

Romer, D., "Short-run fluctuations," 2011.

Rudebusch, G.D. and E.T. Swanson, "The Bond Premium in a DSGE Model with Long-Run Real and Nominal," *American Economic Journal: Macroeconomics*, 2012, *4* (1), 105–143.

Serrato, J.C.S. and P. Wingender, "Estimating Local Fiscal Multipliers," *University of California at Berkeley, mimeo*, 2010.

Sims, C.A., "Solving linear rational expectations models," *Computational Economics*, 2002, *20* (1), 1–20.

Smets, F. and R. Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *The American Economic Review*, 2007, *97* (3), 586–606.

Swanson, M.T. and J.C. Williams, "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates," 2012.

Werning, I., "Managing a Liquidity Trap: Monetary and Fiscal Policy," Technical Report, National Bureau of Economic Research 2011.

Woodford, M., *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, 2003.

_ , "Financial intermediation and macroeconomic analysis," *The Journal of Economic Perspectives*, 2010, *24* (4), 21–44.

_ , "Simple Analytics of the Government Expenditure Multiplier," *American Economic Journal: Macroeconomics*, 2011, *3* (1), 1–35.

8 Figures & Tables

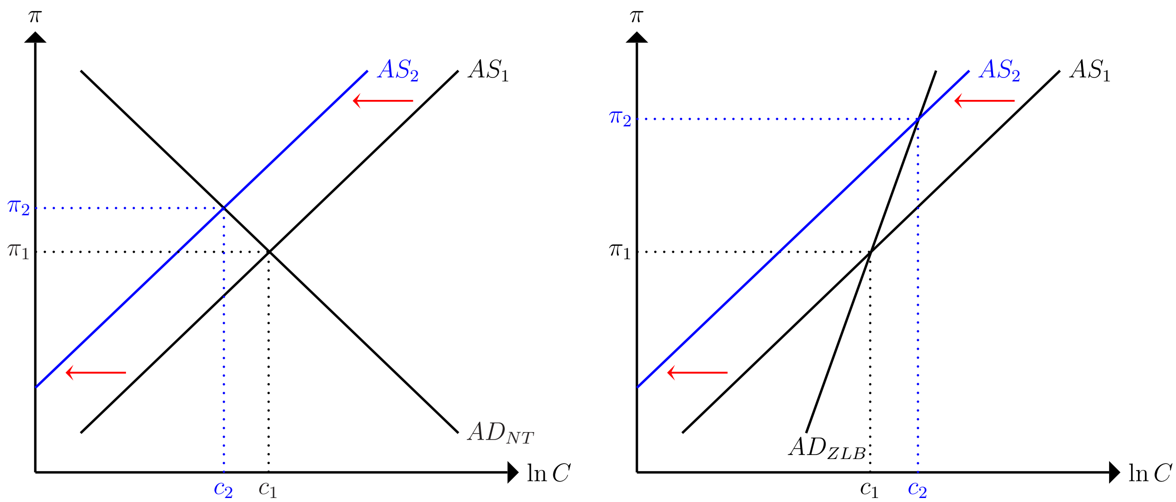

(a) Normal Times (b) ZLB

Figure 1: Impact of negative supply shock in Euler equation framework of Section 2. Left panel shows impact for normal times when the central bank is unconstrained. Right panel shows impact when the zero lower bound binds.

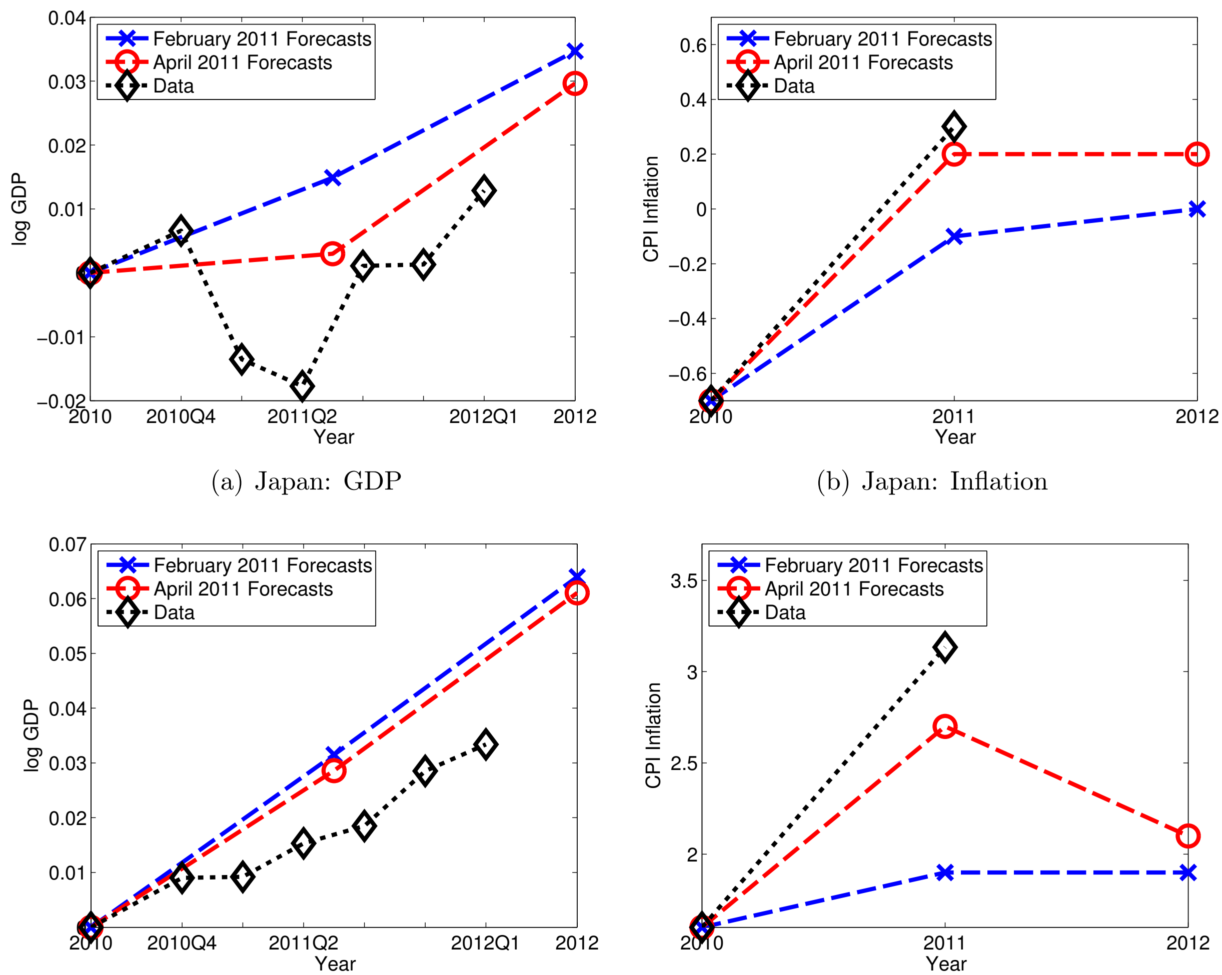

(a) Japan: GDP

(b) Japan: Inflation

(c) U.S.: GDP

(d) U.S.: Inflation

Figure 2: Consensus Economics forecasts from before Japanese Great Earthquake (February 2011) and after (April 2011). Forecasts are for annual GDP and year-on-year inflation. GDP data is annual for 2010 and quarterly from 2010Q4 until 2012Q1. CPI data is annual year-on-year inflation. See Section 3.1 for details.

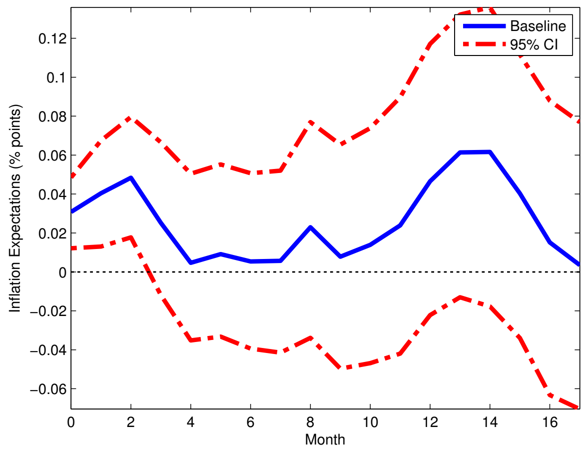

(a) 4Q Ahead Expected Inflation

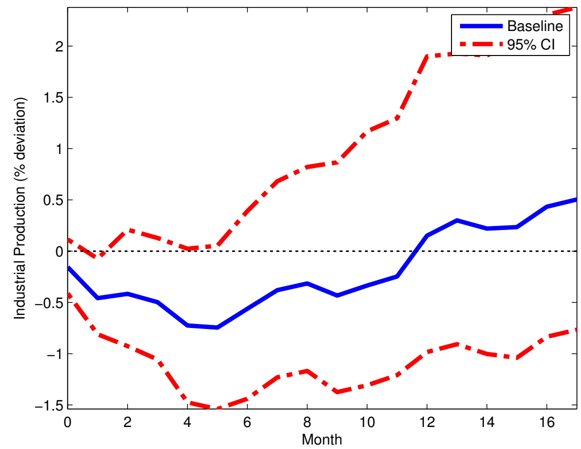

(b) Log Industrial Production

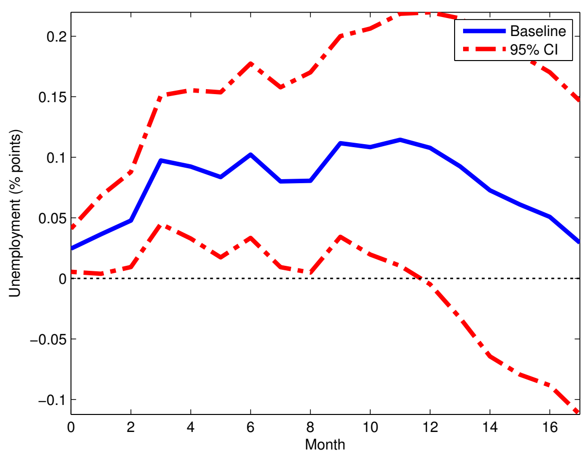

(c) Unemployment

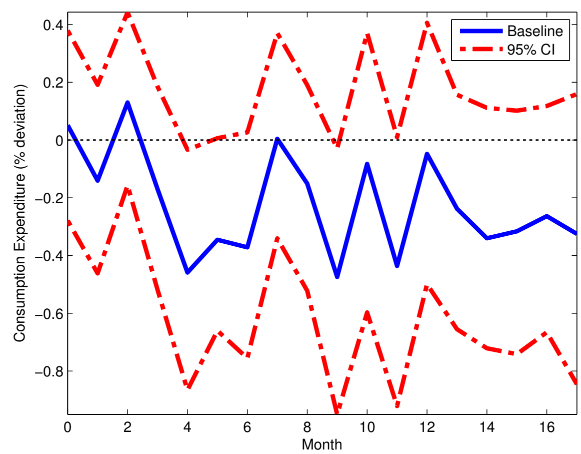

(d) Consumption Expenditures

Figure 3: Impulse Response Functions to negative oil supply shocks. IRFs are constructed from autoregressive distributive lag estimates in changes or growth rates and aggregated to levels. 95% confidence intervals are derived by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix. Lag lengths are according to AIC: For the dependent variable they are set to 12 for expected inflation, 48 for industrial production, and 36 for unemployment and consumption. Lag lengths for oil shocks are set to 24 for expected inflation, industrial production, and unemployment, and to 36 for consumption. See Section 3.2 for details.

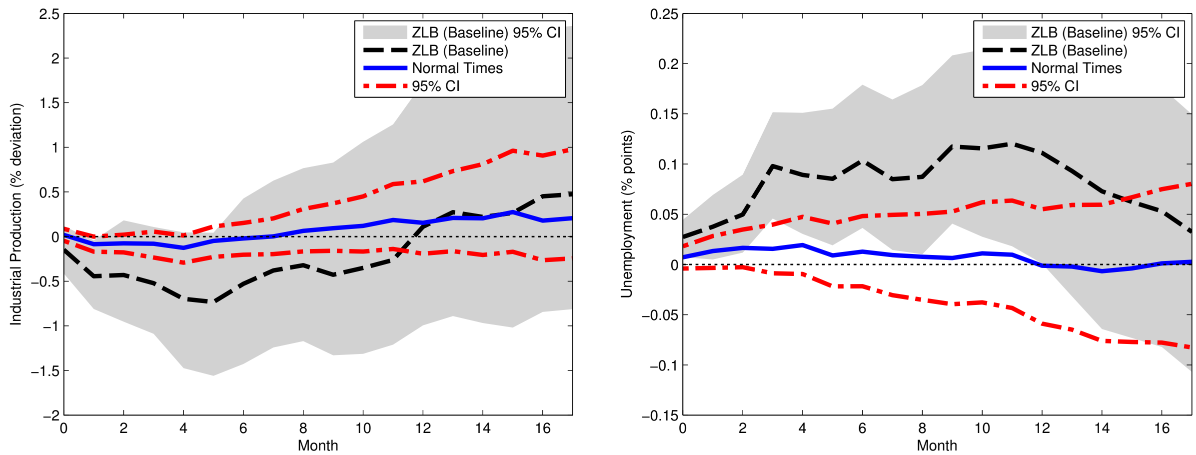

(a) Log Industrial Production

(b) Unemployment

Figure 4: Impulse Response Functions to oil supply shocks 1985-present: When ZLB binds ("Baseline") and when policy rates are unconstrained ("Normal Times"). IRFs are constructed from autoregressive distributive lag estimates in changes or growth rates and aggregated to levels. 95% confidence intervals are derived by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix. See Section 3.2 for details.

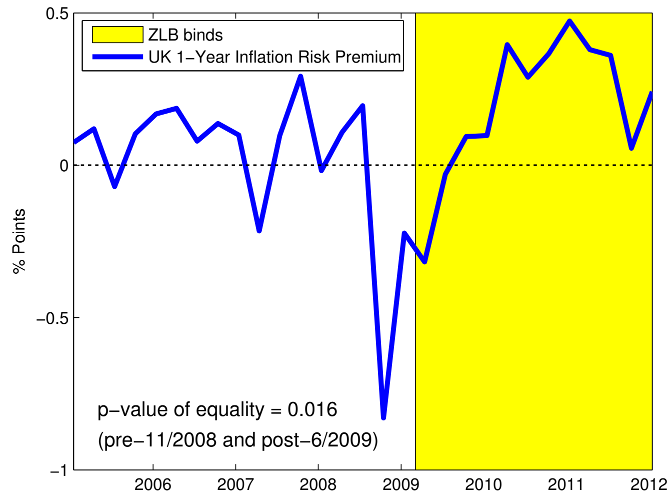

(a) Excess Return on U.K. Inflation Swaps.

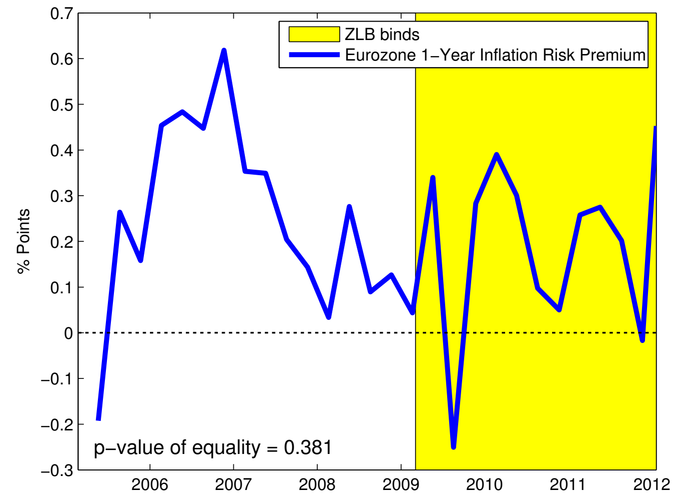

(b) Excess Return on Eurozone Inflation Swaps.

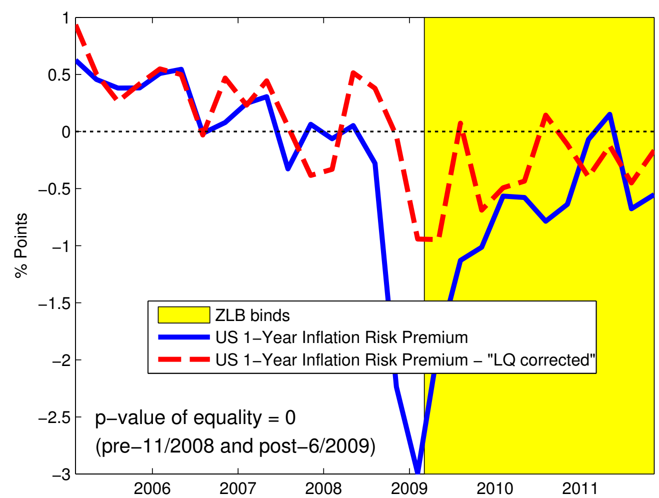

(c) Excess Return on U.S. Inflation Swaps.

Figure 5: Inflation Risk Premia for the U.K., Eurozone, and the U.S. Calculated as the difference in inflation swap rates and expected inflation rates. For the U.K. and the U.S. monthly risk premia are averaged over a quarter to account for seasonalities. Liquidity-corrected ("LQ-corrected") excess returns for the U.S. are calculated as described in Appendix F. Shaded areas indicate time periods were the ZLB binds as defined in Table 1. P-values are calculated over the sub-samples in brackets (U.K. and U.S.) or over pre- and post-ZLB (Eurozone) based on Newey-West standard errors. For the U.S. p-values are calculated only for non-LQ-corrected data. See Section 4 for details.

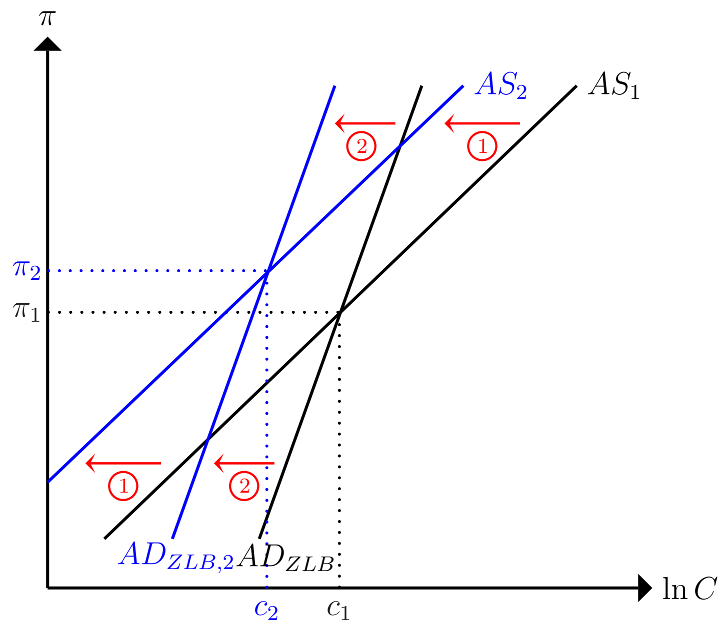

Figure 6: Impact of negative supply shock at ZLB when these shocks have endogenous negative demand effects, as in the friction model of Section 5. The standard shift in the AS is marked by "1" and the endogenous demand effect - the induced shift in the AD curve - is marked by "2."

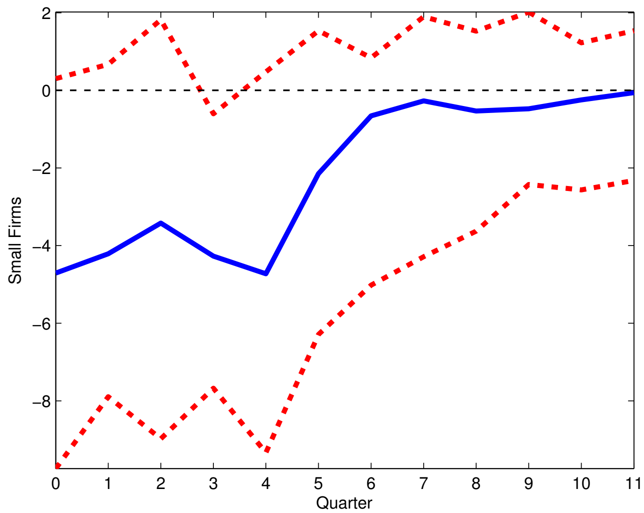

(a) Credit Standards: Small Firms.

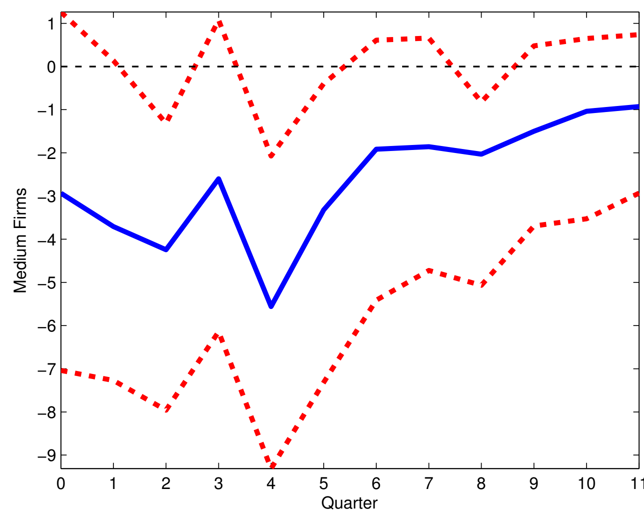

(b) Credit Standards: Medium Firms.

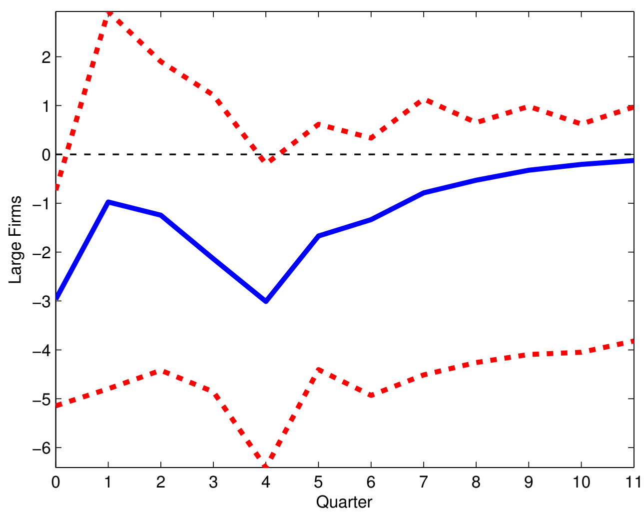

(c) Credit Standards: Large Firms.

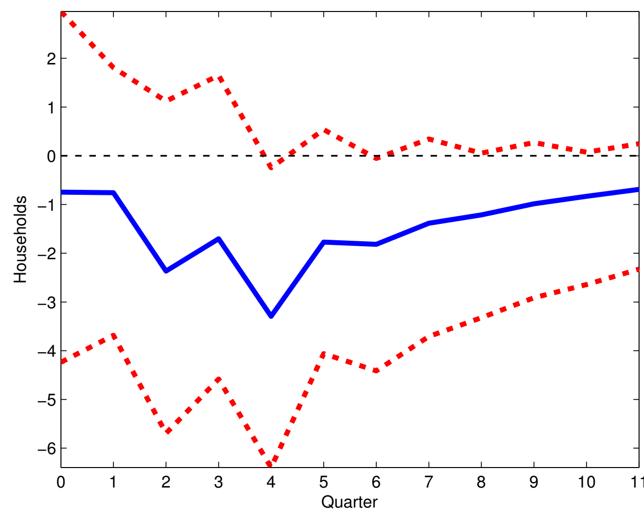

(d) Credit Standards: Households.

Figure 7: Impact of oil shocks on credit standards in Japan. Negative IRFs indicate a tightening of credit standards. 95% confidence intervals are derived by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix. Units are calculated as (percentage of respondents selecting “eased considerably” + percentage of respondents selecting “eased somewhat” * 0.5) -(percentage of respondents selecting “tightened considerably” + percentage of respondents selecting “tightened somewhat” * 0.5). See Section 5 for details.

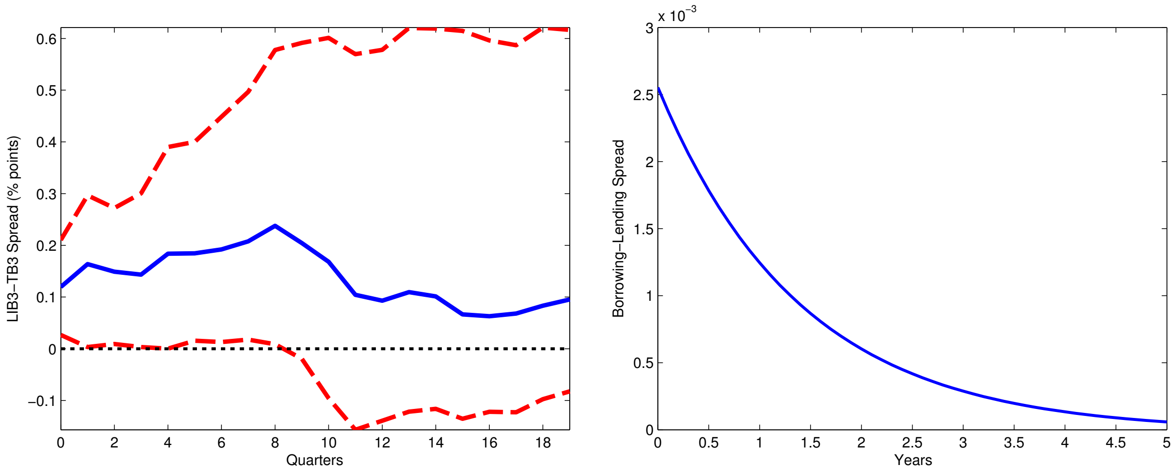

(a) TED Spread. (b) Model Borrowing Spread.

Figure 8: Left Panel: Impact of 1% permanent, negative TFP Shock on TED Spread (3 month LIBOR - 3 month Treasury Bill). Estimated using an ADL equation with 12 lags of the TED spread and 12 lags of TFP shock. Newey-West standard error bands are drawn from the Variance-Covariance matrix using Monte-Carlo methods. Right Panel: Model impact of a 1% permanent, negative TFP shock on the borrowing-lending spread given the baseline calibration. See Section 5.4 for details.

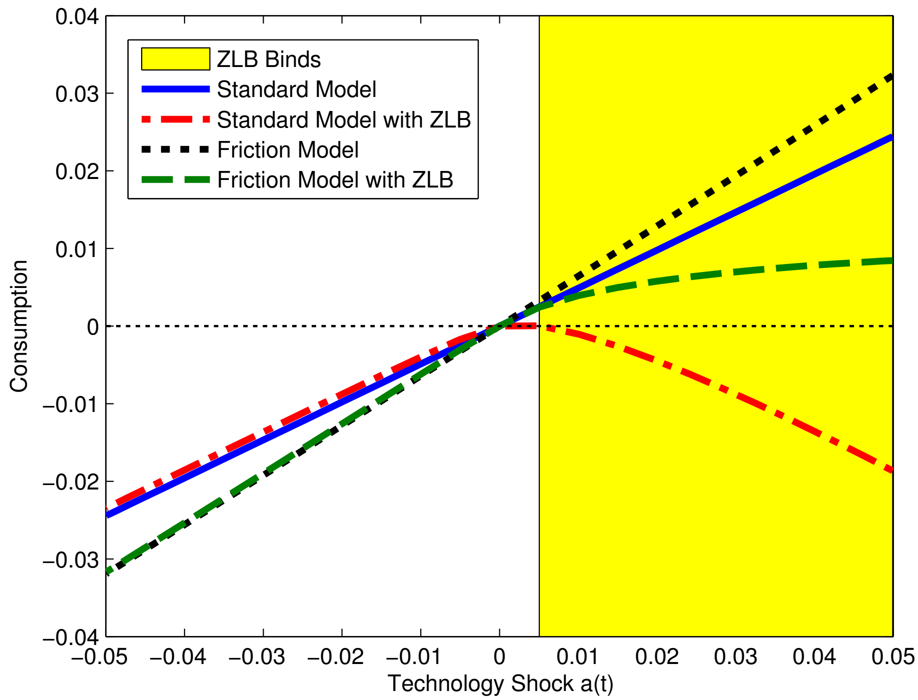

(a) Log Aggregate Consumption (Output).

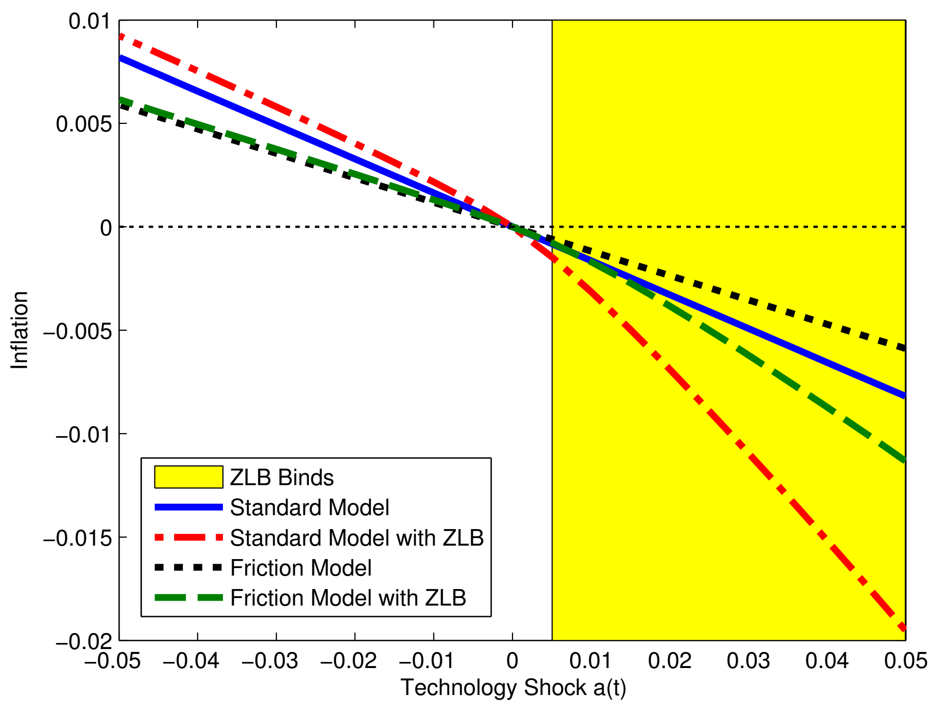

(b) Inflation.

Figure 9: Model solutions for log aggregate consumption $\ln C(a(t))$ and inflation $\pi(a(t))$ as a function of the technology shifter $a(t)$. Shaded areas are intervals of the state space where the ZLB binds. "Standard Model" denotes the standard New Keynesian model, and "Standard Model with ZLB" the same model explicitly incorporating the ZLB constraint. The "Friction Model" denotes the model outlined in Section 5, where borrowing agents face a time-varying credit spread. See Section 5 for details.

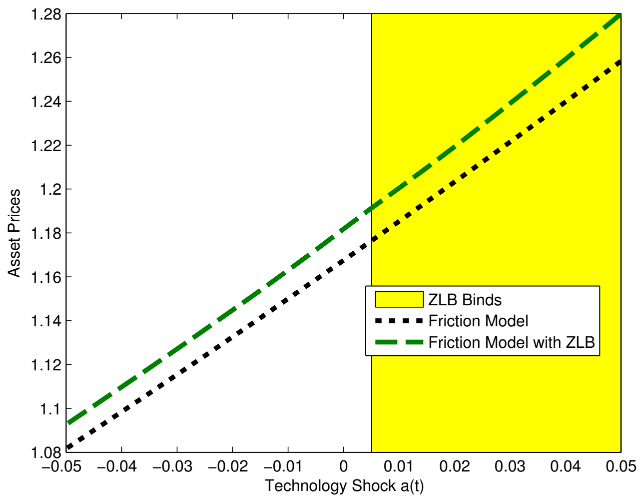

(a) Asset Prices.

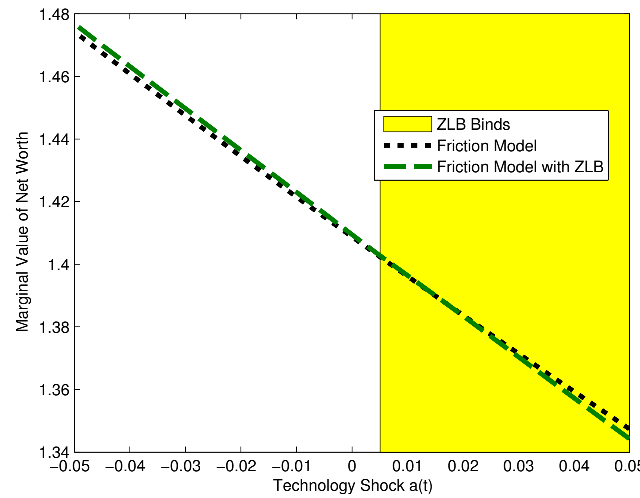

(b) Marginal Value of Net Worth.

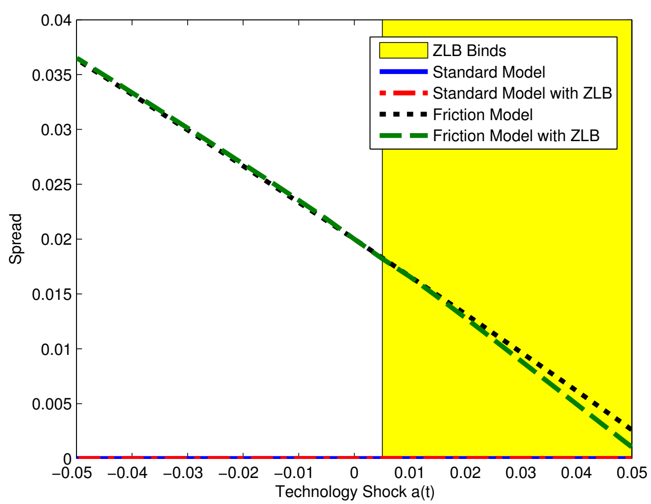

(c) Borrowing Spread.

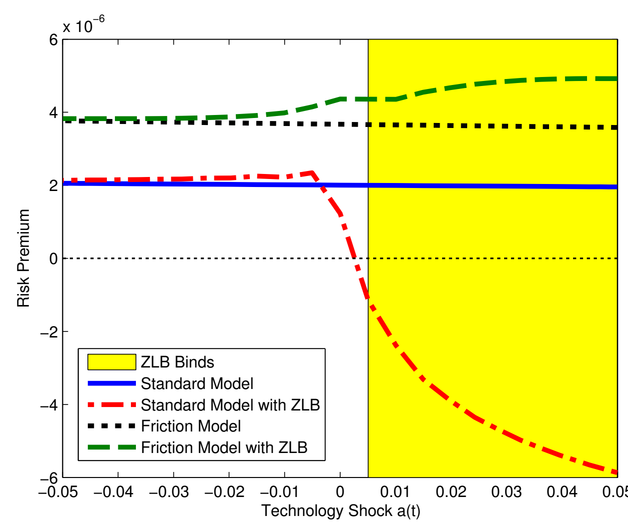

(d) Instantaneous Inflation Risk Premium.

Figure 10: Model solutions for asset prices, marginal value of net worth, the borrowing spread and the inflation risk premium as a function of the technology shifter $a(t)$. Shaded areas are intervals of the state space where the ZLB binds. “Standard Model” denotes the standard New Keynesian model, and “Standard Model with ZLB” the same model explicitly incorporating the ZLB constraint. The “Friction Model” denotes the model outlined in Section 5, where borrowing agents face a time-varying credit spread. See Section 5 for details.

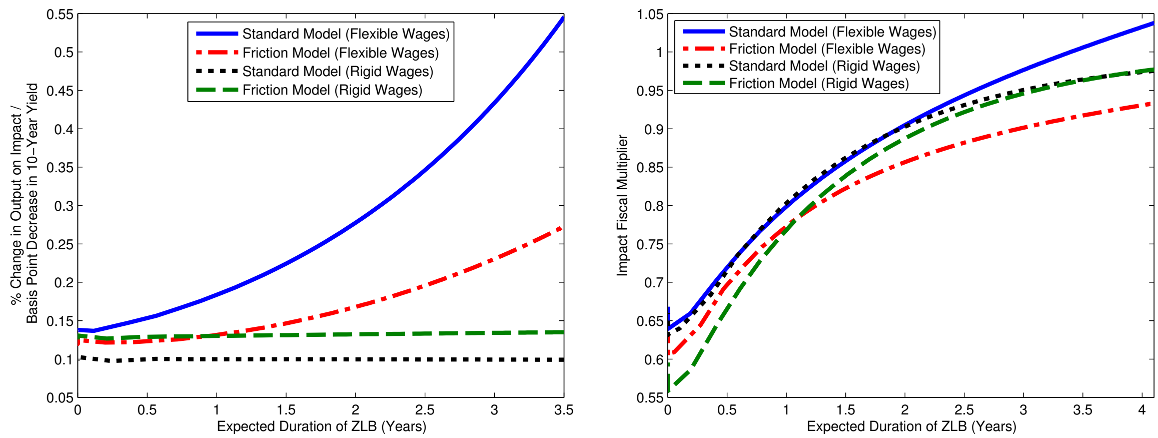

(a) Forward Guidance.

(b) Fiscal Multiplier.

Figure 11: Percentage change in output on impact for each basis point change in the 10-year bond rate (left panel) and impact fiscal multiplier (right panel). "Standard Model (Flexible Wages)" denotes the standard New Keynesian model with competitive labor market, "Friction Model (Flexible Wages)" denotes the model outlined in Section 5 with a competitive labor market, "Standard Model (Rigid Wages)" denotes the standard New Keynesian model with real wages fixed at their steady-state value, and "Friction Model (Rigid Wages)" denotes the model outlined in Section 5 with real wages fixed at their steady-state value. See Section 6 for details.

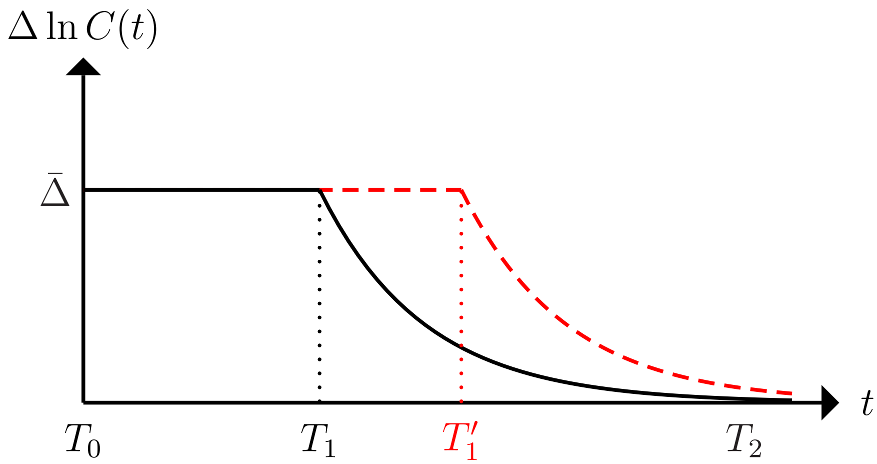

Figure 12: Percentage change in consumption from forward guidance at various exit dates when inflation is held fixed. T_1 and T_1' denote two alternative exit dates. T_0 is the date at which forward guidance is "announced." T_2 is the maturity date of the long-term bond whose yields are reduced by one basis point as a result of forward guidance.

Table 1: ZLB Dates

Country	Start Date	End Date
Japan	October 1995	July 2006
	December 2008	-
U.S.	December 2008	-
U.K.	March 2009	-
Eurozone	January 2009	-
Canada	May 2009	May 2010
Sweden	July 2009	June 2010

Table 2: 10-Year Bond Yields

Dependent variable	10 YR Bond Rate		
Specification	Many Lags	One Lag	Robust
Oil Shock	-0.0465^{**}	-0.0188	-0.0369^{**}
	(0.0182)	(0.0187)	(0.0164)
Lagged Bond Rate	-0.0545	0.0385	0.0047
	(0.0962)	(0.0806)	(0.0496)
Dep. Var. Lags	36	1	1
Country FE	Yes	Yes	Yes
Observations	294	294	294

(1) & (2): Driscoll-Kraay standard errors in parentheses

(3): Stata rreg implementation

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Parameterisation of the Model

Parameter	Definition	Value	Source
Households			
ϱ	Discount Rate	0.012	Own calculations
η	Labor Supply Elasticity	2	Gertler and Kiyotaki [2010]
p	Share of Lenders	0.5	Cúrdia and Woodford [2009]
δ	Type-switching Intensity	0.1	Cúrdia and Woodford [2009]
ζ^l	Marginal Utility Shifter (Lender)	0.7	Cúrdia and Woodford [2009][1]
ζ^b	Marginal Utility Shifter (Borrower)	1.34	Cúrdia and Woodford [2009][1]
Firms			
ε	Elasticity of Substitution (Goods)	40	Basu and Fernald [1997][2]
κ	Slope of Phillips Curve	0.15	Altig et al. [2011]
Banking Sector			
λ	Payout Intensity	0.1	Gertler and Kiyotaki [2010]
ξ	Elasticity of Marginal Value of Net Worth w.r.t. Asset Prices	0.6	Own calculations
$\bar{\phi}$	Steady-State Leverage	2	Gertler and Kiyotaki [2010][3]
θ	Fraction of Divertable Assets	0.70	Gertler and Kiyotaki [2010][4]
ν	Infusion of Net Worth	0.05	Gertler and Kiyotaki [2010]
Government			
ϕ_π	Interest Rate Rule Coeff. (Inflation)	1.5	Standard Parameters
ϕ_y	Interest Rate Rule Coeff. (Output Gap)	0.5	Standard Parameters
Shocks			
ρ_a	Persistence of Technology Shock	0.5	Own calculations
σ_a	SD of Technology Shock	0.005	Own calculations

[1] Set to match Debt/GDP ratio of 80% as in their calibration.

[2] Set to match steady-state economic profit share of 2.5% as in their estimates.

[3] Their calibration has $\bar{\phi} = 4$, but $\bar{\phi} = 2$ better matches their steady-state marginal value of net worth.

[4] Set to match Steady-State Leverage.

A Supply Shocks in a Model with Capital

In standard models without capital, consumption equals output, so if negative supply shock raises consumption at the ZLB, then this shock must be expansionary. Adding investment and capital to the model complicates the algebra, although typically negative supply shocks remain expansionary at the ZLB in this setting. The key intuition is that higher consumption raises the return on capital, while higher expected inflation lowers the return on nominal bonds; both effects make investment more lucrative.

I first consider a model with adjustment costs on capital. Let $q(t)$ be the price of capital, δ its rate of depreciation, and $r^k(t) = \alpha Y(t)/K(t)$ its instantaneous rate of return, assuming a Cobb-Douglas production function. Then the arbitrage condition between nominal government bonds and capital is given by,

$$i(t) - \pi(t) = \frac{r^k(t) - \delta + \mathbb{E}_t dq(t)}{q(t)} \tag{15}$$

where $q(t)$ is the price (incremental value) of capital. With adjustment costs, the price of capital today will jump to satisfy the arbitrage condition. For example, if the RHS exceeds the LHS, then $q(t)$ jumps up and is expected to decline over time, $dq(t) < 0$. Both effects will reduce the RHS until no arbitrage opportunity remains. Another (equivalent) way to view $q(t)$ is as the expected discounted return to capital,

$$q(t) = \int_0^\infty e^{\int_0^s [i(t+j) - \pi(t+j)]dj} [r^k(t+s) - \delta] ds. \tag{16}$$

Higher returns to capital imply a higher value for $q(t)$ as do lower real returns on nominal bonds. I will assume that the price of capital, $q(t)$, is increasing in investment, which is satisfied by both the Lucas and Prescott [1971] and the Christiano, Eichenbaum, and Evans [2005] adjustment costs. Thus, if the current price of capital rises/falls, then investment increases/decreases.

Consider now a negative supply shock that raises expected inflation for a given level of output as in Section 2. In normal times consumption falls, because the central bank satisfies the Taylor principle, which reduces the return on capital, $r^k(t)$. In addition, the return on nominal bonds $i(t) - \pi(t)$ rises due to the Taylor principle, making investment even less attractive. Both effects will reduce the price of capital, $q(t)$, and thus investment today. Thus, a negative supply shock remains contractionary in normal times in models with investment.

At the ZLB both effects are reversed, as the supply shock raises consumption and the return to capital, $r^k(t)$, while lowering the return on nominal bonds, $-\pi(t)$.[47] Both forces raise the price of capital $q(t)$ as the representative agent increases her investments. Since both investment and consumption rise, the negative supply shock is expansionary at the ZLB. I summarize these results in the following proposition.

Summary 5 *Consider an economy that satisfies the Euler equation* (1)*, the supply side* (3)*, the arbitrage equation* (15)*, let monetary policy be given by* (4)*, and let the price of capital* $q(t)$

[47] I assume that the supply shock only affects the economy while the ZLB binds, as in Section 2.

be increasing in investment. Then a temporary, negative supply shock - an increase in $u(t)$ - that does not persist after the ZLB binds is expansionary, and increases both consumption and investment.

This shows that the results for the consumption-only case are robust to adding a standard investment sector. However, this category excludes the popular accelerator models, both of old and new vintage. In these models, the arbitrage condition (15) no longer holds. Instead, investment is constrained by the net worth $N(t)$ of entrepreneurs (or banks). Denote leverage by $\phi(t)$, then by definition

$$N(t)\phi(t) = q(t)K(t). \tag{17}$$

These models then provide a theory of how $\phi(t)$ is determined. The typical theory of leverage assumes that it is increasing in the expected excess return on capital, $\frac{\partial\phi(t)}{\partial[r^k(t)-\delta-i(t)+\pi(t)]} \geq 0$ (e.g., Bernanke and Gertler [1989] and Gertler and Kiyotaki [2010]). Intuitively, higher expected excess returns on capital reduce the probability of default among entrepreneurs (or banks), which mitigates the financial friction in the economy and permits higher leverage. Since excess returns fall in normal times following a temporary negative supply shock, leverage will fall and so will the price of investment $q(t)$. This again implies that investment falls in normal times. At the ZLB we found that excess returns rise, which by similar logic, will boost investment in an accelerator models. Note that the change in capital prices will further affect net worth, which boosts the quantitative response both at the ZLB and in normal times. Thus, the classic financial frictions on the investment side do not overturn the qualitative results derived above, but magnify the quantitative response.

The single exception is when the supply shock *directly* wipes out some net worth.[48] In this case a negative supply shock is more contractionary in normal times as both leverage and net worth fall. At the ZLB, the effect is now ambiguous - if the rise in leverage exceeds the fall in net worth then negative supply shocks remain expansionary in this model. However, if the opposite is true, then the LHS of Equation (17) falls, $q(t)$ must fall and investment contracts. The supply shock may nevertheless still be expansionary because consumption unambiguously rises. However, if a supply shock has large, direct negative effects on net worth, then the aggregate output effects may obscure differential consumption and investment responses. In this case, testing for expansionary negative supply shocks at the ZLB and testing for the strength of the expectations channel in the Euler equation are not identical. However testing for the consumption response alone remains a valid test for the strength of expectations channel even in this class of models.

B Robustness: Oil Shocks

This appendix investigates the robustness in economic responses to oil shocks. In the following sections I conduct the following checks: letting 4-quarter ahead inflation expectations be measured from the current quarter, excluding the Eurozone, adding controls for inflation and spreads, dropping lagged dependent variables, excluding outliers, using hp-filtered data, using non-seasonally-adjusted consumption, and various lag lengths for oil shocks.

[48]Remember that the standard effect of supply shocks works through $q(t)$, which affects net worth. The condition here is that the shock must destroy some fraction of net worth directly.

B.1 4Q-Expectations Measured from Current Quarter In Section 3.2 the change in inflation expectations was calculated ignoring the current quarter. This, if an oil shock occurs in April of a given year, I only measured the change in inflation expectations from July onwards. In this subsection, I calculate the 4-quarter ahead inflation expectation starting in the current quarter. For example, for April this will be from Q2 this year until Q2 next year, rather than from Q3 this year to Q3 next year as in the baseline. Thus, this measure will also capture if the oil shock raises inflation expectations from April until June. On the other hand, at the end of a quarter, say in June, it will also capture higher inflation expectations for past two months, thus potentially overstating the total increase in inflation expectations from today. I plot the resulting IRF from Equation (7) in Appendix Figure 13. As expected, it displays a larger increase in inflation expectations than the baseline since the latter is a conservative estimate.

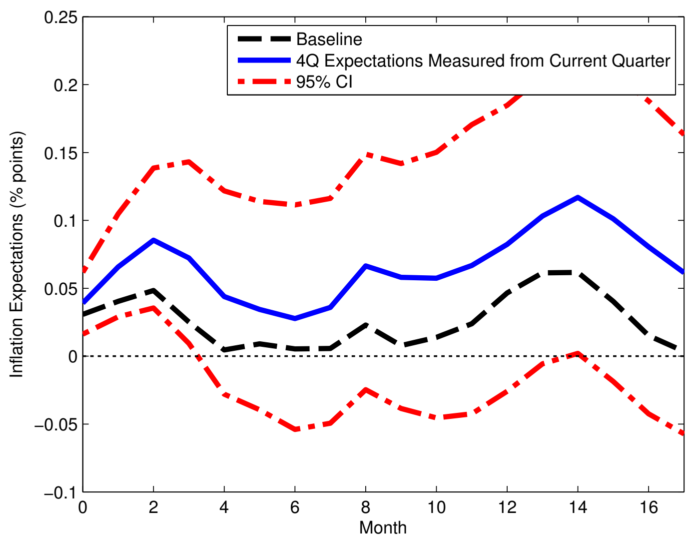

Figure 13: Impulse Response Functions to Oil Shocks for 4-quarter ahead inflation expectations measured from the *current* quarter. "Baseline" refers to 4-quarter ahead inflation expectations measured from the *next* quarter. Two-standard-error confidence intervals are constructed by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix.

B.2 Eurozone I exclude the Eurozone from the sample and repeat the earlier analysis. Since I did not have data on consumption expenditures for the Eurozone, this plot is unchanged and not reported. For the remaining variables the IRFs are essentially identical.

B.3 Additional Controls I include contemporaneous values and twelve lags of changes in inflation and changes in corporate bond spreads in Equation (8). The resulting IRFs for oil shocks are plotted in (Appendix Figure 15(c)) and are very similar to the baseline results in Section 3.2. This suggests that the results are not driven by small-sample covariance of oil shocks with financial shocks or inflation dynamics.

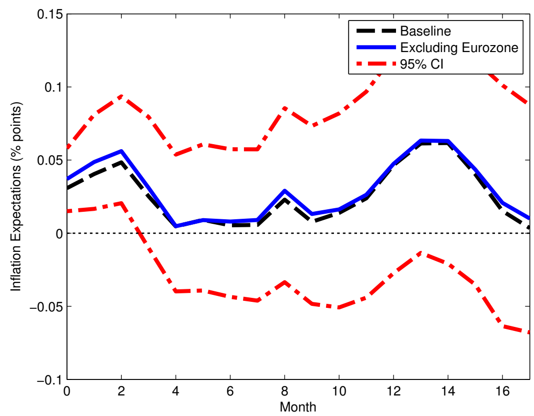

(a) 4Q Ahead Expected Inflation

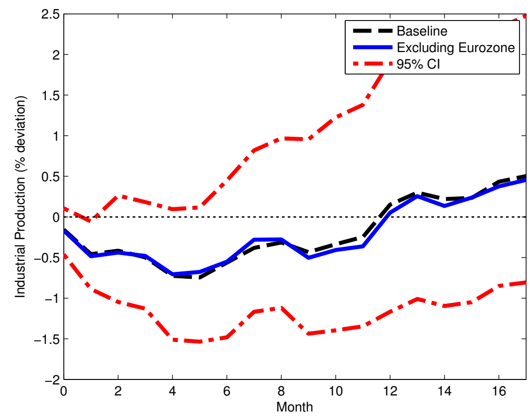

(b) Log Industrial Production

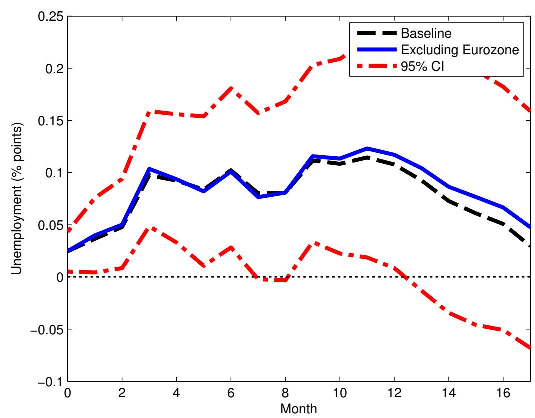

(c) Unemployment

Figure 14: Impulse Response Functions to Oil Shocks *excluding the Eurozone from the sample*. Two-standard-error confidence intervals are constructed by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix.

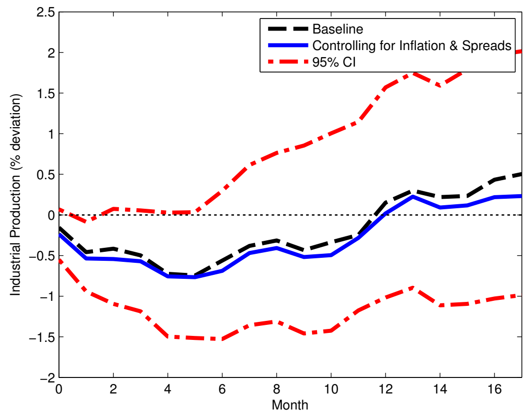

(a) Log Industrial Production

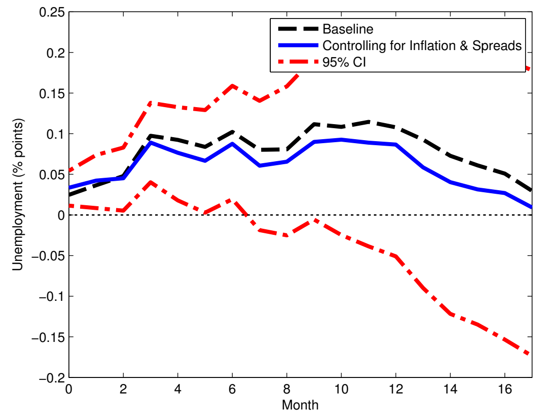

(b) Unemployment

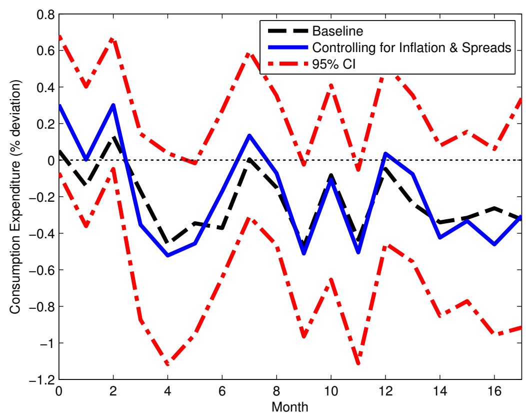

(c) Consumption Expenditures

Figure 15: Impulse Response Functions to Oil Shocks *controlling for inflation and corporate bond spreads*. Two-standard-error confidence intervals are constructed by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix.

B.4 Lagged Dependent Variables If the oil supply shocks in Kilian [2009] are well-identified, then they would be orthogonal to the current state of the economy, so that excluding lagged dependent variables from Equations (7) and (8) should yield similar results. In (Appendix Figure 16) I plot the IRFs from this estimation. The results for inflation expectations, unemployment, and industrial production are very similar to the baseline results, which suggests that the oil shock is not picking up dynamics induced by other shocks. If anything, the weakened responses for IP and unemployment indicate that negative oil supply shocks are correlated with improving economic dynamics.

More troubling is that consumption expenditure expand after a negative oil supply shock when I exclude the lagged dependent variables, whereas in the baseline model they contracted (Appendix Figure 16(d)). However, when I use non-seasonally-adjusted consumption, then the estimated response without lagged dependent variables become more contractionary relative to the baseline (Appendix Figure 16(e)). With lagged dependent variables included both series agree that negative oil supply shocks reduce consumption expenditures (Figure 3(d) and Appendix Figure 20(a)), which suggests that the noise in consumption expenditure series makes it necessary to control for the lagged state of the economy. Thus, excluding lagged dependent variables does not appear to be informative in this instance.

B.5 Outliers Outliers are determined by jackknifed residuals in the estimated ADLs. Residuals that exceed the 1% critical value are then removed from the sample and the ADLs are re-estimated. The resulting IRFs are plotted in Appendix Figure 17, together with the baseline results. Typically, the baseline IRF lies within the 95% confidence interval of the outlier-corrected IRF. Quantitatively, the only significant difference is the IRF for industrial production, which now displays less mean-reversion.

B.6 HP-filtered data For industrial production and unemployment I also estimate equation (8) using HP-filtered data. The resulting IRFs and 95% confidence intervals are plotted in Appendix Figure 18. As in the case for outliers the baseline IRF lies within two standard deviations of the HP-filtered IRF. The magnitudes are somewhat smaller, but the contractionary effects of oil shocks remain significant.

I also check if oil shocks affect the economy differently at the ZLB compared to normal times using HP-filtered data. The IRFs in Appendix Figure 19 display similar results as in the main text, namely that average contractionary effects are stronger at the ZLB, but that uncertainty over these effects is large.

B.7 NSA consumption In Appendix Figure 20 I plot the IRF for non-seasonally adjusted consumption expenditures in Japan. As with the X-12 adjusted series, the IRF displays a statistically significant contraction in consumption expenditures following an oil shock. This result is also robust to excluding outliers. Thus, the measured contraction in consumption is not due to the particular seasonal adjustment mechanism.

B.8 Lag Lengths I investigate the sensitivity of my results to changes in lag lengths on the oil shocks. (I already checked for sensitivity to lagged dependent variable lag length above.) In Appendix Table 4 I tabulate the values of the IRF one month and four months

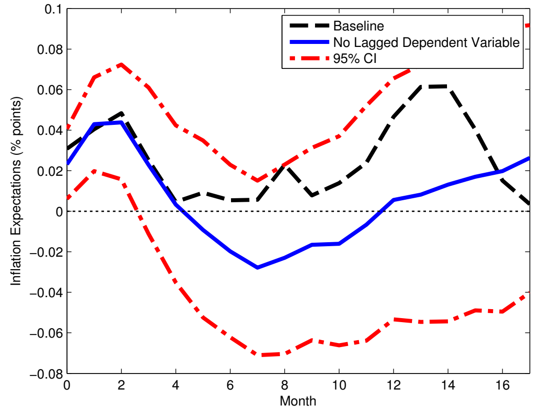

(a) 4Q Ahead Expected Inflation

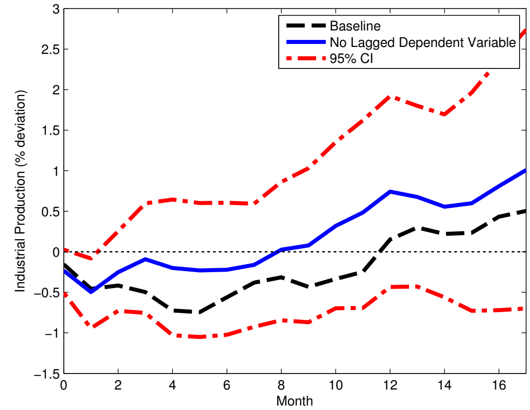

(b) Log Industrial Production

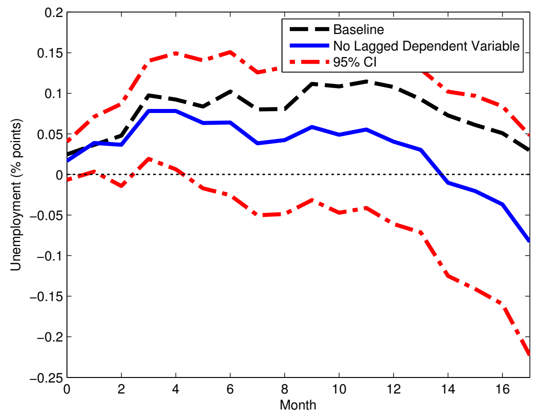

(c) Unemployment

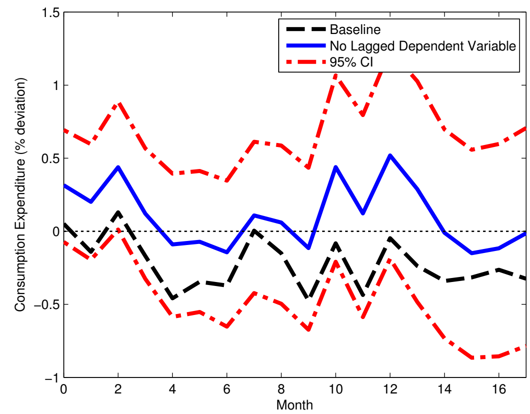

(d) Consumption Expenditures (SA)

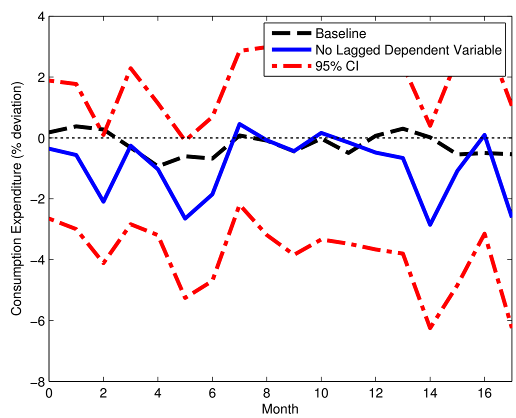

(e) Consumption Expenditures (Non-SA)

Figure 16: Impulse Response Functions to Oil Shocks *excluding lagged dependent variables* in Equations (7) and (8). Two-standard-error confidence intervals are constructed by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix.

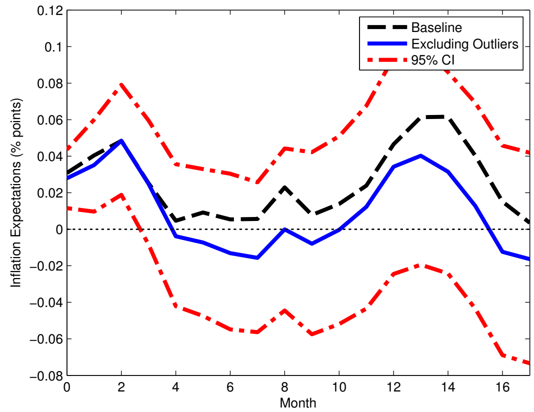

(a) 4Q Ahead Expected Inflation

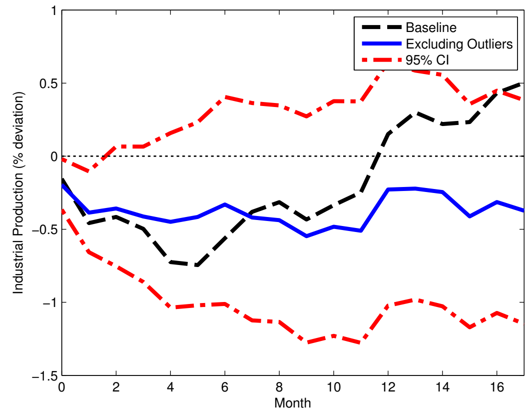

(b) Log Industrial Production

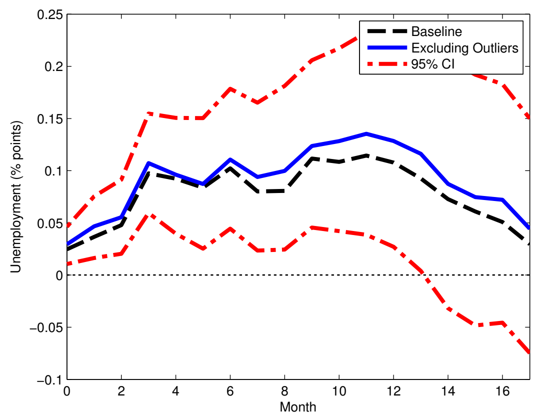

(c) Unemployment

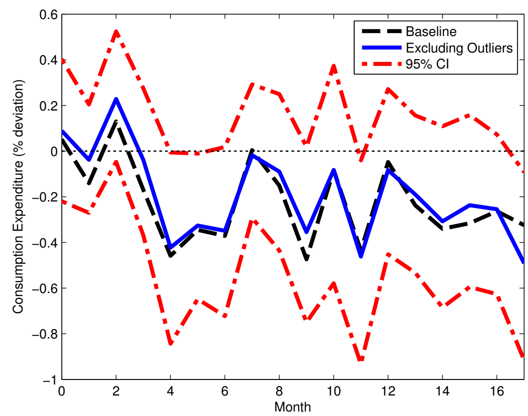

(d) Consumption Expenditures

Figure 17: Impulse Response Functions to Oil Shocks *excluding outliers* based on jacknifed residual with 1% critical value. Two-standard-error confidence intervals are constructed by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix.

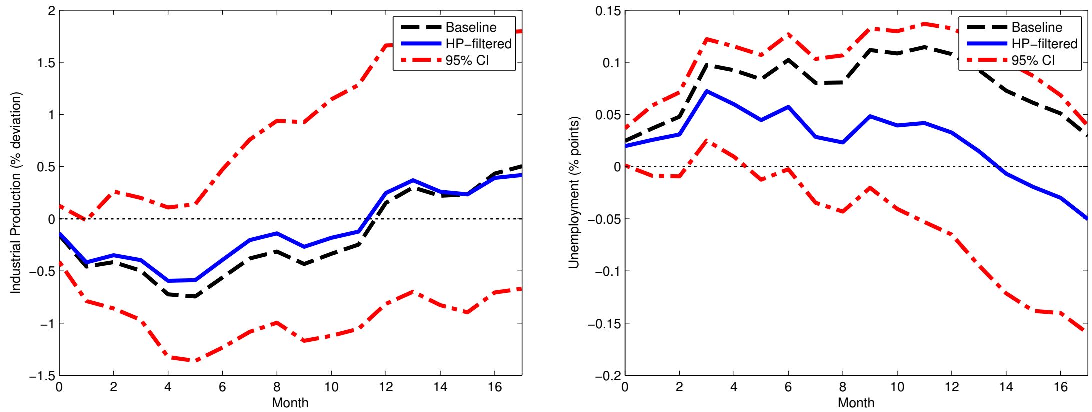

(a) Log Industrial Production

(b) Unemployment

Figure 18: Impulse Response Functions to Oil Shocks for *HP-filtered* variables. Two-standard-error confidence intervals are constructed by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix.

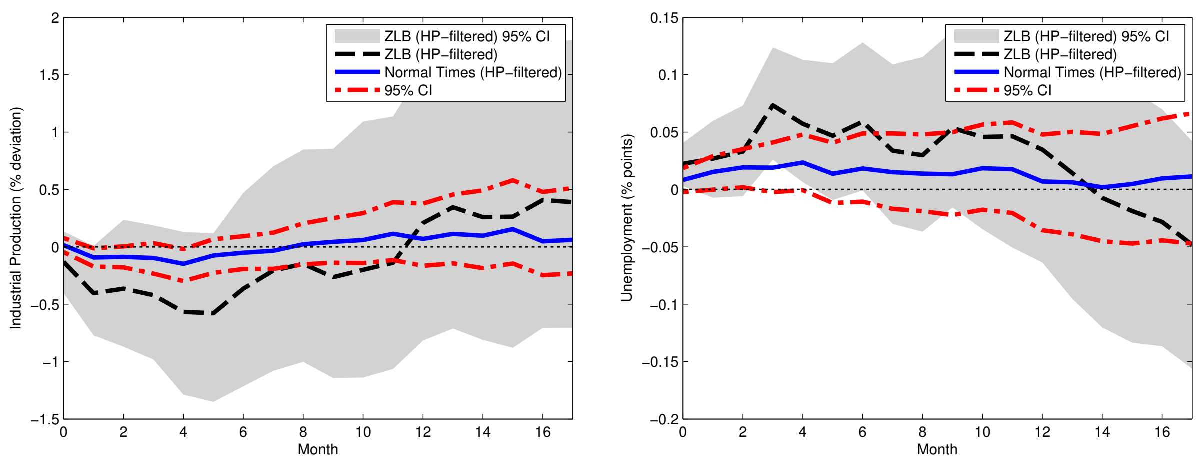

(a) Log Industrial Production

(b) Unemployment

Figure 19: Impulse Response Functions to oil supply shocks over 1984-now when ZLB binds ("Baseline") and when policy rates are unconstrained ("Normal Times") for *HP-filtered* variables. IRFs are constructed from autoregressive distributive lag estimates in changes or growth rates and aggregated to levels. 95% confidence intervals are derived by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix. See Section 3.2 for details.

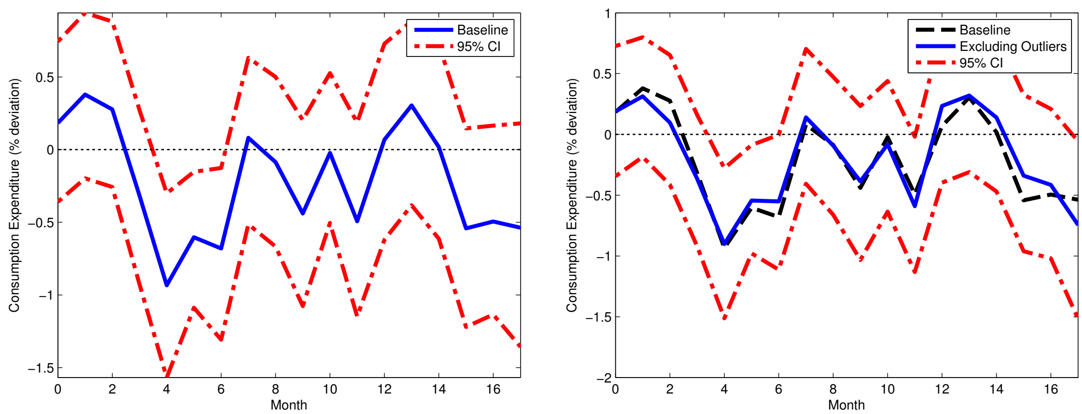

(a) Baseline (b) Excluding Outliers

Figure 20: Impulse Response Functions of NSA consumption for Japan. Two-standard-error confidence intervals are constructed by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix.

after the shock for 12, 24, and 36 lags of oil shocks. The results are very similar across all these lag lengths and similar to the baseline. In particular, inflation, IP, and unemployment show significant responses one month out, and unemployment and consumption expenditures display significant movements after four months as in the baseline.

Table 4: Impact of Oil Shocks on Macroeconomic Variables for Various Lags of Oil Shocks in Equations (7) and (8).

Lags	Dependent Variable							
	Inflation Expectations		Industrial Production		Unemployment		Consumption Expenditure	
	1 Month	4 Months	1 Month	4 Months	1 Month	4 Months	1 Month	4 Months
12	0.027 [0.009,0.044]	0.006 [-0.029,0.043]	-0.462 [-0.838,-0.023]	-0.717 [-1.459,-0.002]	0.032 [-0.000,0.062]	0.096 [0.032,0.155]	-0.31 [-0.618,0.012]	-0.411 [-0.808,-0.027]
24	**0.031** **[0.012,0.048]**	**0.025** **[-0.012,0.067]**	**-0.444** **[-0.811,-0.042]**	**-0.698** **[-1.474,0.047]**	**0.038** **[0.005,0.069]**	**0.089** **[0.030,0.151]**	-0.239 [-0.593,0.115]	-0.483 [-0.898,-0.046]
36	0.032 [0.010,0.056]	0.039 [0.001,0.082]	-0.498 [-0.894,-0.078]	-0.766 [-1.502,-0.064]	0.038 [0.002,0.074]	0.086 [0.030,0.145]	**-0.141** **[-0.462,0.191]**	**-0.459** **[-0.866,-0.033]**
N	299	299	299	299	299	299	163	163

95% Confidence intervals in brackets based on Driscoll-Kraay s.e. "1 Month" and "4 Months" are the impact and associated 95% CI at 1 month and 4 months after the shock. Baseline specifications are marked in bold font. Dependent variable lags are as in the baseline regressions (12 for inflation, 48 for IP, 36 for unemployment and consumption).

C Inflation Expectations & Kalman Filter

This appendix details the inflation expectations data sources and the Kalman filter used to infer four-quarter ahead inflation expectations from available data.

C.1 Data sources The data sources and their frequency are tabulated in Appendix Table 5.

Table 5: Inflation Expectations Data Sources

Country	Source	Reference Price	Forecast Frequency	
			Q4-Q4 Forecasts	Annual Forecast
Japan	Consensus Economics	CPI	Quarterly	Monthly
U.S.	Blue-Chip Economic Indicators	CPI	Monthly	Monthly
U.K.[1]	Consensus Economics	RPIX	Quarterly	Monthly
	HM Treasury	RPIX	-	Monthly
	HM Treasury	RPI	-	Monthly
Eurozone	Consensus Economics	CPI	Quarterly	Monthly
Sweden	Consensus Economics	CPI	Quarterly	Monthly
Canada	Consensus Economics	CPI	Quarterly	Monthly

[1] RPIX forecasts are used in oil shock regressions, and RPI forecast to derive inflation risk premia.

C.2 Kalman Filter As in Hamilton [1994], the state space representation of the system is given by,

$$\xi_t = \mathbf{F}\xi_{t-1} + \mathbf{M} + \mathbf{v}_t$$
$$\mathbf{y}_t = \mathbf{H}_t'\xi_t + \mathbf{w}_t$$

where the first equation is the evolution of the state and the second equation is the observation equation. Time units are monthly. The state vector ξ_t consists of lagged and future one-quarter-ahead inflation forecasts. Let $\pi_t^{s,s+x}$ denote the time t inflation forecast from quarter s to quarter $s+x$. For example $\pi_t^{0,4}$ denotes the 4Q-ahead inflation forecast from the current quarter $s = 0$. I measure inflation forecast in log points, so that 4Q-ahead inflation forecasts can be calculated as,

$$\pi_t^{0,4} = \pi_t^{0,1} + \pi_t^{1,2} + \pi_t^{2,3} + \pi_t^{3,4}.$$

The state vector ξ_t is then equal to,

$$\xi_t = \begin{pmatrix} \pi_t^{-1,0} & \pi_t^{0,1} & \pi_t^{1,2} & \dots & \pi_t^{7,8} \end{pmatrix}'.$$

The transition matrix $\mathbf{F}$ is defined as follows:

$$\mathbf{F} = \begin{cases} I_9 & \text{if quarter unchanged} \\ \begin{pmatrix} \mathbf{0} & I_8 \\ \mathbf{0} & \zeta \end{pmatrix} & \text{if new quarter begins} \end{cases}$$

Thus, if there is no change in the quarter, the state space evolves (in expectations) as $\pi_t^{s,s+x} = \pi_{t-1}^{s,s+x}$, which is consistent with rational expectations. Whenever a new quarter begins, the inflation forecast from two quarters ago $\pi_{t-1}^{-1,0}$ drops out of the state space and gets replaced with the inflation forecast from last quarter $\pi_t^{-1,0} = \pi_{t-1}^{0,1}$. The same procedure is repeated for all inflation forecasts in the state space. The 1Q-ahead forecast seven quarters from today, $\pi_t^{7,8}$ was not in the lagged state space. It is therefore set to be proportional to its counterpart from last quarter, $\pi_t^{7,8} = \zeta\pi_{t-1}^{7,8}$. I capture the mean of these new expectations in the vector $\mathbf{M}$, which equals

$$\mathbf{M} = \begin{cases} \mathbf{0} & \text{if quarter unchanged} \\ \begin{pmatrix} \mathbf{0} \\ (1-\zeta)\mu \end{pmatrix} & \text{if new quarter begins} \end{cases}$$

I allow errors in the state equation to be correlated to be capture the notion that shocks affect inflation expectations for several quarters. Specifically, I let the (i, j) element of $\mathbf{v}_t$ be equal to,

$$v_{ij} = \rho^{|i-j|}\sigma_\xi^2.$$

In the observation equation, I let the observation vector be

$$\mathbf{y}_t = \begin{pmatrix} \tilde{\pi}_t^{y1} & \tilde{\pi}_t^{y0} & \tilde{\pi}_t^{-1,0} & \tilde{\pi}_t^{0,1} & \tilde{\pi}_t^{1,2} & \dots & \tilde{\pi}_t^{7,8} \end{pmatrix}'.$$

where $\tilde{\pi}_t^{s,s+x}$ denotes the observed inflation forecasts from quarter s to quarter $s + x$. The first two elements of $\mathbf{y}_t$ are annual forecasts for next year (first element) and the current year (second element), which are a tent-shaped function of the 1Q-ahead inflation forecasts in the state space. As inflation becomes published, I remove it from the current year forecast to ensure that it can be expressed in terms of the state vector.

The observation matrix $\mathbf{H}_t'$ will be time varying depending on what forecasts are available. Annual forecasts are available at monthly frequency, but quarterly forecasts are typically only available every 3 months. Nevertheless, this can be easily handled by the Kalman filter. I also allow for white noise errors in the observation equation (correlations are already built into the state space equation). In particular, I let the observation error for annual and quarterly forecasts be σ_a^2 and σ_q^2 respectively, so that

$$w_{ii} = \begin{cases} \sigma_a^2 & \text{if } i \leq 2 \\ \sigma_q^2 & \text{if } i > 2 \end{cases}$$

The parameters estimated in the Kalman filter are $\mathbf{\Omega} = (\rho, \zeta, \mu, \sigma_\xi^2, \sigma_a^2, \sigma_q^2)$. I use the standard estimation procedures outlined in Hamilton [1994]. With the estimated parameters, I calculate smoothed estimates for the state vector $\hat{\xi}_t$ and derive the (conservative) 4Q-ahead inflation forecast

$$\hat{\pi}_t^{1,5} = \hat{\pi}_t^{1,2} + \hat{\pi}_t^{2,3} + \hat{\pi}_t^{3,4} + \hat{\pi}_t^{4,5}.$$

The results are robust to using the unsmoothed (one-sided) inflation expectations. In Appendix Figure 21 I plot the baseline impulse response function from Section 3.2 using inflation expectations both filters. The results are nearly identical.

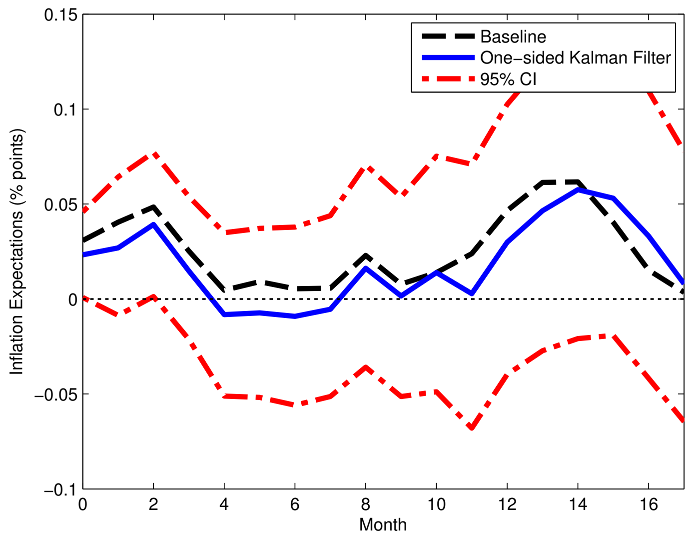

Figure 21: Impulse Response Functions to Oil Shocks for inflation expectations derived from smoothed estimates ("Baseline") and for inflation expectations derived from the one-sided filter. Two-standard-error confidence intervals are constructed by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix.

The basic set-up is slightly modified for the U.S. and the U.K. In the U.S. I use both Consensus Economics and Blue Chip forecasts, which have correlated errors in the observation equation. In the U.K. I also use HM Treasury forecasts to determine RPI inflation. While I use RPIX forecasts in the oil shock estimates, I need the RPI estimates to calculate inflation risk premia. In this case the state space includes two additional variable that capture the difference between the current 4Q-ahead RPI and RPIX forecasts, $s_t^{z,z+4}$ and $s_t^{z+4,z+8}$, where z is the first quarter of this year. I then calculate RPI forecasts by $rpi_t^{s,s+x} = rpix_t^{s,s+x} + w_1 s_t^{z,z+4} + (1 - w_1) s_t^{z+4,z+8}$ where w_1 is the fraction of quarters in $(s, s+x)$ that lie in the current year.

D Inflation Risk Premia

This appendix illustrates that the intuition from instantaneous inflation risk premia carries over to longer-horizon risk premia, and is robust to using Epstein-Zin preferences.

D.1 s-year Inflation Risk Premia The 1 year risk premium in the model is harder to sign than the instantaneous risk premium. The problems in calculating this quantity are twofold: First, when the economy randomly switches between the ZLB and normal times, then inflation and marginal utility will be non-linear functions of the state variables, which makes it impossible to characterize their distributions analytically. This was not a problem before, because the derivatives $c_u(t)$ and $c_v(t)$, which multiply normal shocks, are fixed over a small time interval. Second, these derivatives must now be solved for within the context of the model, whereas before I didn't have to specify why $c_u(t)$ is positive or negative at the ZLB.

To make analytical progress on the first issue, I will make the following assumptions: First, inflation is a linear function of the supply and demand shifters, $\pi(t) = \pi_u u(t) + \pi_v v(t)$. Second, the interest rate rule is given by, $i(t) = r + \phi\pi(t)$. Third, I assume that the current state of the world - ZLB or normal times - persists forever.[49] Fourth, the demand and supply shocks follow mean-reverting Ornstein-Uhlenbeck processes, $dx(t) = -\rho_x x(t) + \sigma_x dB_x(t)$, for $x \in \{u, v\}$. Under these assumptions marginal utility and consumption are linear functions of the state variables.

Given these assumptions, I can calculate an explicit solution for the s-period inflation risk premium,[50]

$$rp_\pi(t, t+s) = \frac{(1-e^{-\rho_u s})^2}{2\rho_u^2 s}\left[-\frac{\sigma_u^2 \pi_u c_u}{\gamma}\right] - \frac{(1-e^{-\rho_v s})^2}{2\rho_v^2 s}\left[\frac{\sigma_v^2 \pi_v c_v}{\gamma}\right] \tag{18}$$

where $c_u = C_u(t)/C(t)$ and $c_v = C_v(t)/C(t)$ are constant parameters. The solution is similar to the instantaneous risk premium in Equation (9).[51] Intuitively, when the solution to the model is linear, the s-year risk premium strings together a set of constant instantaneous risk premia, while accounting for mean-reversion in the state variables.

The effect of changing persistence on the risk premium is non-monotonic because the terms π_x and c_x are themselves functions of ρ_x for $x \in u, v$. With a vertical long-run Phillips curve, the instantaneous effect of a permanent supply shock $u(t)$ on inflation is zero, because consumption rises instantaneously such that marginal costs are unchanged. Thus, $\pi_u = \tilde{\pi}_u \rho_u$ where $\tilde{\pi}_u$ is non-zero when $\rho_u = 0$. Similarly, because demand shocks do not affect long-run consumption, $c_v = \tilde{c}_v \rho_v$, where $\tilde{c}_v$ is non-zero when $\rho_v = 0$. Substituting these expressions back into (18) yields,

$$rp_\pi(t, t+s) = \frac{(1-e^{-\rho_u s})^2}{2\rho_u s}\left[-\frac{\sigma_u^2 \tilde{\pi}_u c_u}{\gamma}\right] - \frac{(1-e^{-\rho_v s})^2}{2\rho_v s}\left[\frac{\sigma_v^2 \pi_v \tilde{c}_v}{\gamma}\right]. \tag{19}$$

The term multiplying the squared brackets, $(1-e^{-\rho_x s})^2/2\rho_x s$, is equal to zero when there is instant mean-reversion $\rho_x \to \infty$ or when there is complete persistence $\rho_x = 0$. Intuitively, a shock with instant mean-reversion affects neither consumption nor inflation because these are forward-looking objects. Similarly, permanent shocks do not affect the risk premium

[49] In each case I assume that the equilibrium is determinate.

[50] Technically, this calculates a *price level* risk premium rather than an *inflation* risk premium. However, I follow the literature that has settled on the latter term.

[51] It can be shown that the partial derivative of $rp_\pi(t, t+s)$ with respect to s evaluated at $s = 0$ equals $rp_\pi(t)$.

because their inflation (supply shock) or consumption (demand shock) response is zero. The maximum loading occurs when $1 + \rho_x^* s = \exp(\rho_x^* s)$, which implies that the short-end loads more heavily on shocks that mean-revert quicker, $\partial \rho_x^* / \partial s < 0$. For example, this term is maximized at $s = 1$ when $\rho_x \approx 1.25$, which implies a half-life of 6 months. Shocks with a half-life of 2 years only receive half the weight.

When the ZLB persists forever, there are no spillovers into normal times and supply shocks are contractionary in the baseline model, $c_u(t) > 0$ (Proposition 3). This is problematic, because it will make inflation risk premia unambiguously negative at the ZLB, so that empirical evidence of positive risk premia becomes hard to interpret. To address this issue, I allow supply shocks $dB_u(t)$ to be correlated with demand shocks $dB_v(t)$. If this correlation is negative, then negative supply shocks may be (on average) contractionary at the ZLB because they coincide with negative demand shocks, and this can rationalize a positive risk premia at the ZLB. To simplify the algebra, I assume that the demand and supply shifters are identically distributed, $\sigma_u = \sigma_v = \sigma_x$ and $\rho_u = \rho_v = \rho_x$, and that the correlation between the Brownian Motions is given by φ.

Given these assumptions, the s-period inflation risk premium is given by,

$$rp_\pi(t, t+s) = \frac{(1 - e^{-\rho_x s})^2}{2\rho_x^2 s} \left[-\frac{\sigma_x^2 \pi_u (c_u + \varphi c_v)}{\gamma} - \frac{\sigma_x^2 \pi_v (c_v + \varphi c_u)}{\gamma} \right] \tag{20}$$

In normal times, the sign of risk premium in Equation (20) is ambiguous when $\varphi = 0$, because negative supply shocks are contractionary and will generate positive risk premia, whereas demand shocks will generate negative risk premia. Thus, if contractionary supply shocks are the most important source of inflation volatility, then the inflation risk premium will be positive.

At the ZLB, negative supply shocks are expansionary in standard models when there are no spillovers, so that the risk premium is unambiguously negative if $\varphi = 0$. However, if the correlation φ is negative, then negative supply shocks will be (on average) paired with negative demand shocks, which allows negative co-movement between inflation and consumption, effectively mimicking the behavior of contractionary negative supply shocks. This will be the case when $c_u + \varphi c_v < 0$, making the first term in brackets positive. If, in addition, supply shocks dominate inflation volatility, $\pi_u \gg \pi_v$, then the overall inflation risk premium is positive. Thus, the conclusions are qualitatively similar to the instantaneous case, which is summarized as follows.[52]

Summary 6 *In the standard CCAPM, s-year inflation risk premia are positive, in normal times or at the ZLB, if and only if negative supply shocks have contractionary effects on average, and this constitutes a significant fraction of aggregate inflation risk.*

D.2 Epstein-Zin Inflation Risk Premia I keep the previous simplifying assumption and consider a consumer with Epstein-Zin utility with IES equal to 1 and risk aversion γ^{-1}. Then

[52] Another possibility to rationalize positive inflation risk premia at the ZLB is to make positive demand shocks deflationary so that $c_v(\pi_v + \varphi \pi_u) < 0$. In fact, if inflation risk premia are positive and $c_u + \varphi c_v < 0$, then it must also be the case that $c_v(\pi_v + \varphi \pi_u) < 0$, because $\varphi < 0$ and $\pi_u \gg \pi_v$. This is the consequence of assuming a statistical correlation between demand and supply shocks: if negative supply shocks are paired with negative demand shocks, then positive demand shocks must be paired with positive supply shocks.

the continuous time limit of the recursive definition of the value function implies,

$$\mathbb{E}_t d\ln V(t) = -\varrho \ln C(t) + \varrho \ln V(t) + \frac{(1-\gamma^{-1})}{2}\sigma_v^2, \tag{21}$$

where σ_v^2 is the variance of $\ln V(t)$, which is constant in this linear set-up. Solving forward and substituting $\ln C(t) = c_u u(t) + c_v v(t)$ yields,

$$d\ln V(t) - \mathbb{E}_t d\ln V(t) = \frac{\varrho}{\varrho+\rho_u} c_u \sigma_u dB_u(t) + \frac{\varrho}{\varrho+\rho_v} c_v \sigma_v dB_v(t). \tag{22}$$

The continuous time limit of the SDF is in turn given by,

$$d\ln\Lambda(t) - \mathbb{E}_t d\ln\Lambda(t) = -[d\ln C(t) - \mathbb{E}_t d\ln C(t)] + (1-\gamma^{-1})[d\ln V(t) - \mathbb{E}_t d\ln V(t)], \tag{23}$$

where the second term captures the utility recursion, i.e. deviations from CRRA. Substituting the solution from the previous equation and for consumption yields,

$$d\ln\Lambda(t) - \mathbb{E}_t d\ln\Lambda(t) = -c_u\left[1 + \frac{(\gamma^{-1}-1)\varrho}{\varrho+\rho_u}\right]\sigma_u dB_u(t) - c_v\left[1 + \frac{(\gamma^{-1}-1)\varrho}{\varrho+\rho_v}\right]\sigma_v dB_v(t), \tag{24}$$

so that the price of inflation risk is given by,

$$rp_t = -c_u\pi_u\left[1 + \frac{(\gamma^{-1}-1)\varrho}{\varrho+\rho_u}\right]\sigma_u^2 - c_v\pi_v\left[1 + \frac{(\gamma^{-1}-1)\varrho}{\varrho+\rho_v}\right]\sigma_v^2. \tag{25}$$

This is similar to Equation (9) except for the additional term in square brackets involving γ^{-1}. Standard calibration set $\gamma^{-1} \geq 6$, so that EZ utility amplify risk premia, but do not change their sign. Intuitively, EZ utility also punishes assets that lose value when there is bad news about the future. Due to persistence of shocks in my setting ($\rho_x < \infty$), bad news today is also bad news for the future, so risk premia get amplified (relative to the log utility case) when risk aversion exceeds one.

E Swap Rate - Inflation Expectation Matching

Appendix Table 6 illustrates the reference price levels, inflation lags, and inflation expectations sources for inflation swap rates. I match inflation expectations to risk premia as follows: For the Eurozone the ECB releases quarterly 12 month ahead consensus inflation expectations exactly when the previous month's inflation data is released. For example, in January the forecast horizon is from December to December. The appropriate match for the inflation swap is the October to October expected inflation rate. I therefore add (known) CPI inflation of the previous two months to the forecast, while proportionally downweighting the original 12 month forecast. This procedure is appropriate, if the last two months of the 12 month inflation forecast have a similar inflation rate than the first 10 months. For the U.K. and the U.S. I use the estimates obtained from the Kalman filter procedure in the previous section. For the U.K., I match the December, January, and February swap rates to the Q4-Q4 forecast, which is appropriate given the two month lag, and average over these months to account for seasonal effects. For the U.S. the matched forecast is calculated as the weighted average of the 4-quarter forecasts originating 2 and 3 months back.

Table 6: Swap Rate - Inflation Expectations Matching

Country	Reference Price	Inflation Lag	Inflation Expectations Source
U.S.	CPI	2-3 Months	Kalman filter (Blue-Chip)
U.K.	RPI	2 Months	Kalman filter (HM Treasury)
Eurozone	CPI	3 Months	ECB

F Liquidity Correction

The liquidity correction follows Pflueger and Viceira [2011]. I calculate risk premia $rp_\pi(t, t+1)$ for the United States according to equation (4.2). Then I regress risk premia on a set of liquidity proxies $\mathbf{X}_t$,

$$rp_\pi(t, t+1) = \alpha + \beta \mathbf{X}_t + \varepsilon_t.$$

I use the following variables as proxies: The two-year swap spread, a five year agency spread, and the spread between 1-year swap rates and 1-year inflation compensation from the TIPS-Treasury spread. The first two are used in Pflueger and Viceira [2011] and also enter the regression quadratically. The third variable is calculated using swap data and yield curve estimates from Gürkaynak, Sack, and Wright [2007] and Gurkaynak, Sack, and Wright [2010]. With an estimate $\hat{\beta}$ I calculate "liquidity-adjusted" risk premia as

$$\tilde{rp}_\pi(t, t+1) = rp_\pi(t, t+1) - \hat{\beta} \mathbf{X}_t.$$

G Model Solutions

This appendix examines the computational properties of the model. In the first subsection, I examine how closely equation (13) captures the solution in the linearized model with three state variables. In the second subsection, I show that the results are robust to more mean-reversion in the shock process. In the third subsection, I lay out the computational strategy for the non-linear model with one state variable and compute Euler equation residuals.

G.1 Accuracy of State Space Reduction This section examines how closely equation (13) captures the solution to the three state variable model in Section 5. First, the model equation of the three-state-variable model and the one-state-variable model are linearized and solved using the Matlab code described in Sims [2002]. This code is available for download on

Chris Sims' website. Appendix Figure 22 displays the resulting impulse response functions of the standard New Keynesian model, the three state variable model, and the one state variable friction model to a 1% TFP shocks. Overall, the models display very similar behavior in the responses of output, consumption, inflation, and the banking variables. While the dynamics of the interest rate spread are somewhat different in the two models, what matters for a borrower's consumption choice is the integral over the spread IRF. The integrals are very similar across the two models so there is little difference in consumption behavior of borrowers. Thus the one-state-variable model delivers essentially the same results, but greatly simplifies the non-linear computation.

Second, I solve the three-state-variable model using projection methods as outlined in the following section. Since the state space is significantly larger in this model I only use polynomials up to degree 5 for each state variable, for a total of $5^3 = 125$ polynomials. While this grid is coarse, it can provide insight if the one-state-variable model also provides a good approximation to the three-state-variable model at the ZLB. In addition, in order to solve this model I have to raise the discount rate slightly to $\varrho = 0.01275$ (compared to $\varrho = 0.0125$ in the baseline) and lower the volatility to $\sigma_a = 0.0001$ (compared to $\sigma_a = 0.005$ in the baseline). In Appendix Figure 23 I plot IRFs from a 1% positive technology shock without the ZLB. As in the linearized case, the IRFs of the one-state-variable model and the three-state-variable model are very close. Appendix Figure 24 shows that equation (13) also provides a very good approximation at the ZLB. In particular, output expands by similar magnitudes in the one-state-variable model and the three-state-variable model.

G.2 Accuracy of Non-linear Solution The model is computed by iterating over the FOCs, while using Chebychev polynomials to approximate equilibrium objects. I use polynomials up to degree 33 in the computation. These polynomials automatically implement the boundary condition that the solution must be bounded. The initial guess is the (linear) solution to the standard New Keynesian model. Convergence is fast, taking about 20 seconds using the baseline calibration.

The residuals from the two Euler equations (lender and borrower) as well as from the Phillips Curve are plotted in Appendix Figure 25. The errors in the Phillips Curve are of trivial magnitude (10^{-14}) and can be ignored. Similarly, the Euler equation errors ignoring the ZLB are very small (10^{-13}). The ZLB introduce larger errors because the approximation will not perfectly capture when the ZLB starts to bind. Thus, the errors are largest at the border of the ZLB and non-ZLB regions. However, the maximum errors are still small: at most 3.5 basis points in the standard model and 1.2 basis points in the friction model.

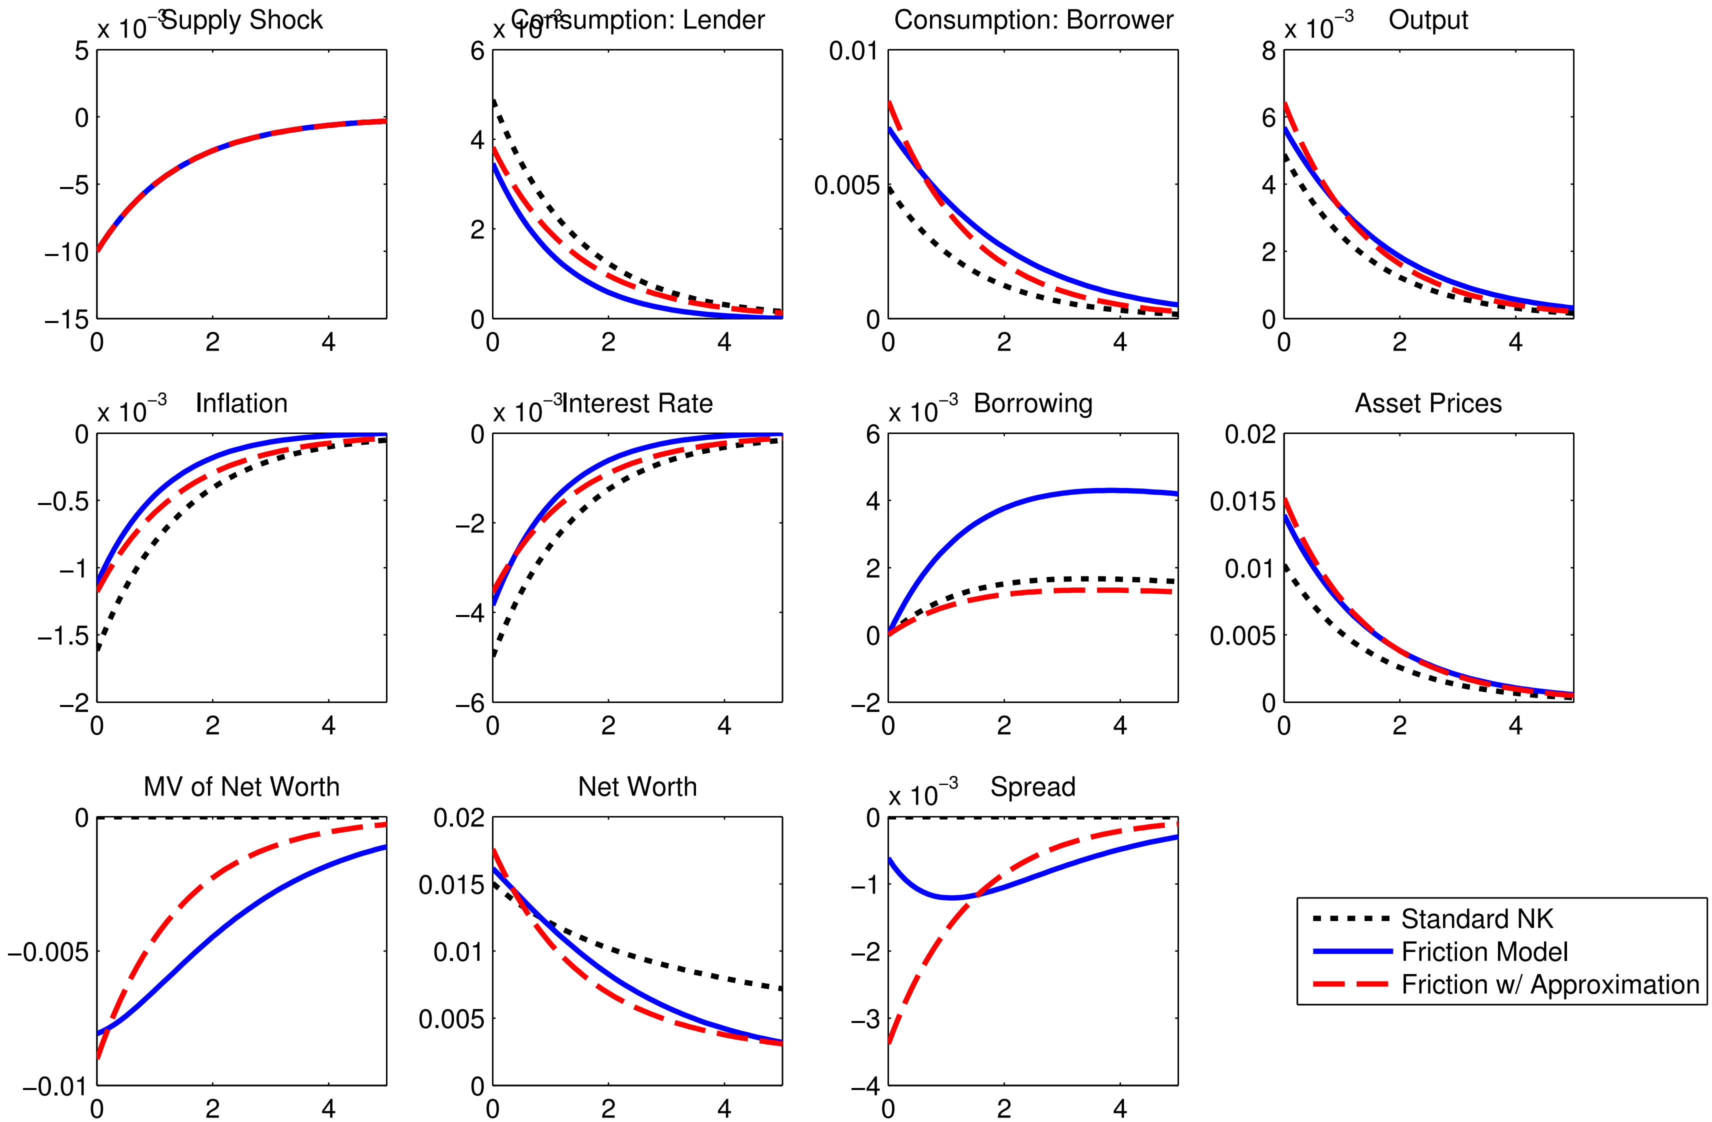

Figure 22: IRFs of Standard NK model, Friction Model with Three State Variables, and Friction Model with One State Variable to a 1% TFP shock. IRFs are calculated from linearized solution to each model.

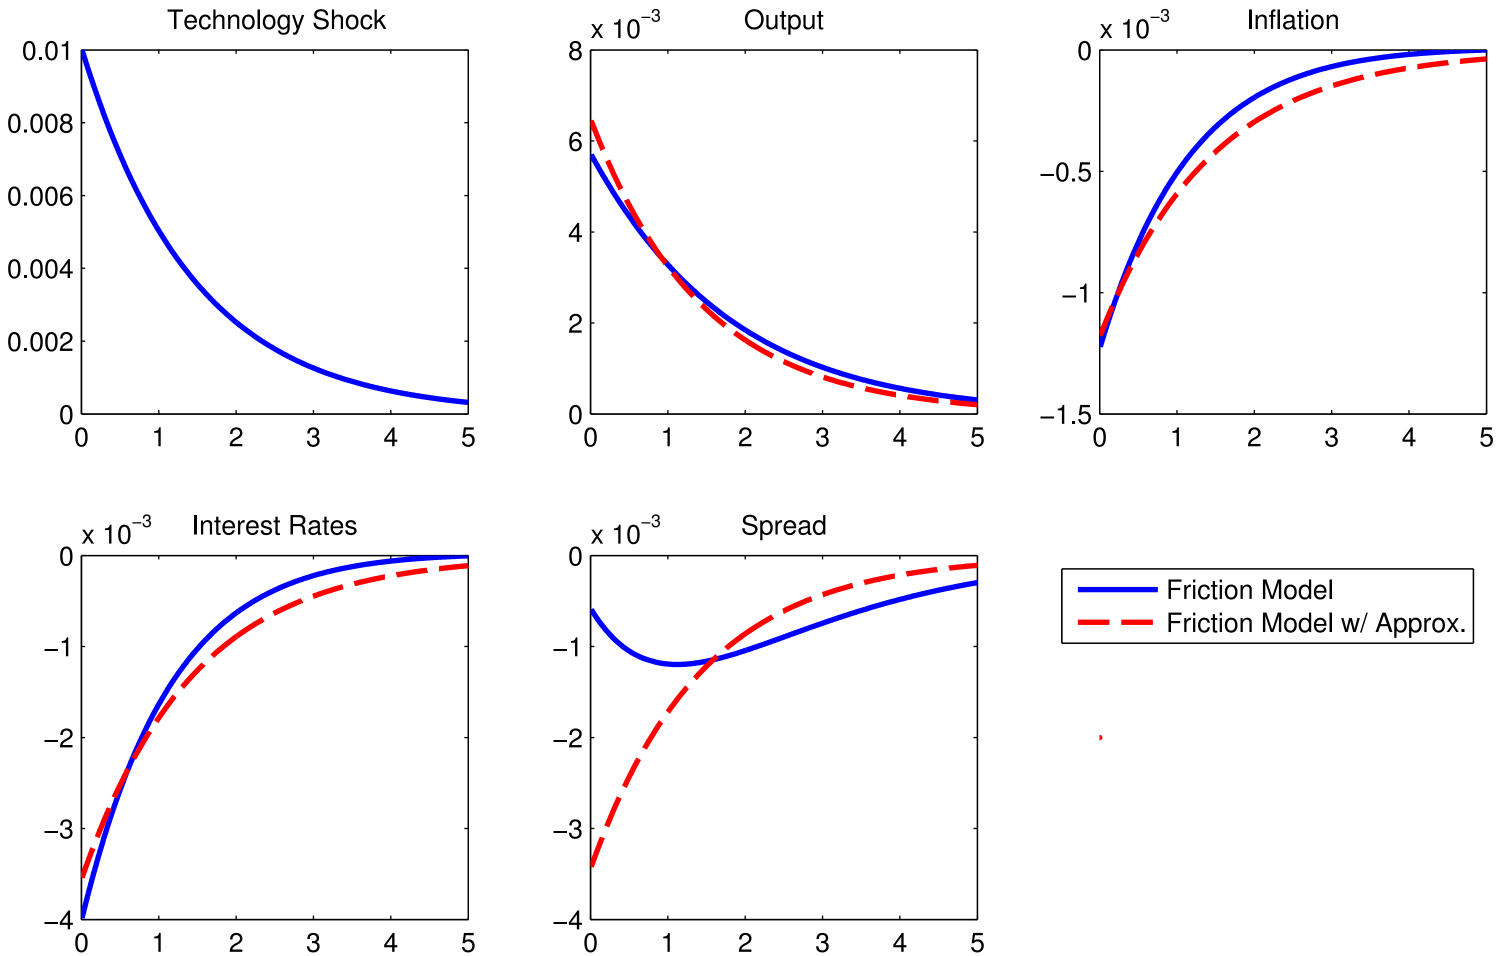

Figure 23: IRFs of Friction Model with Three State Variables and Friction Model with One State Variable to a 1% TFP shock assuming that ZLB does not bind. IRFs are calculated from non-linear model using projection methods.

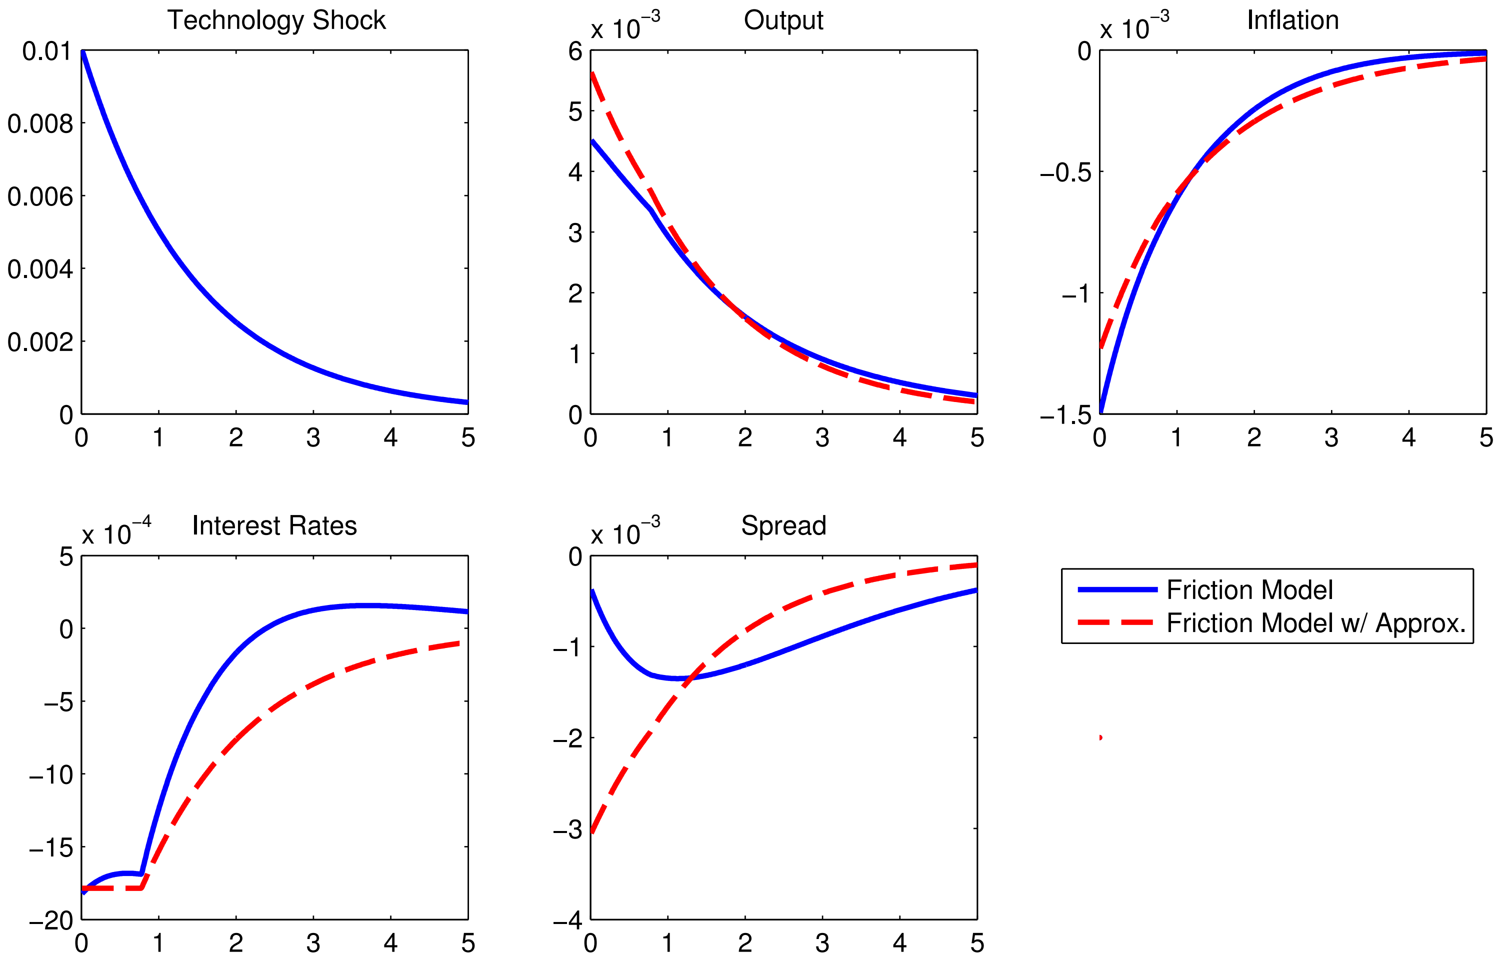

Figure 24: IRFs of Friction Model with Three State Variables and Friction Model with One State Variable to a 1% TFP shock when the ZLB binds. IRFs are calculated from non-linear model using projection methods.

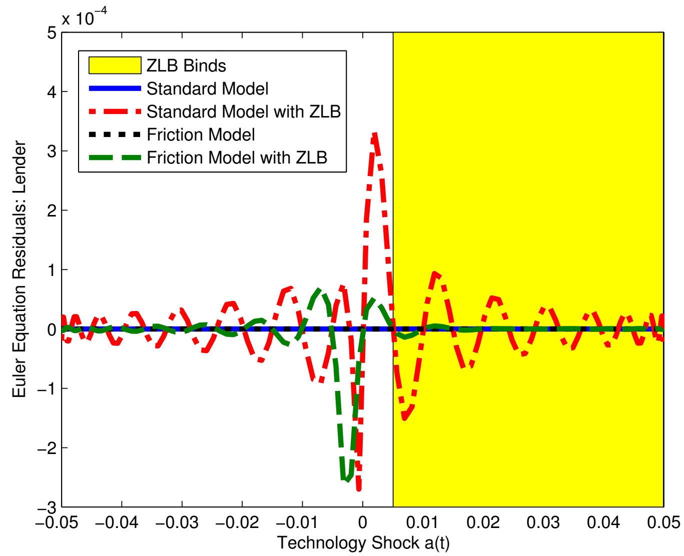

(a) Euler Equation: Lender

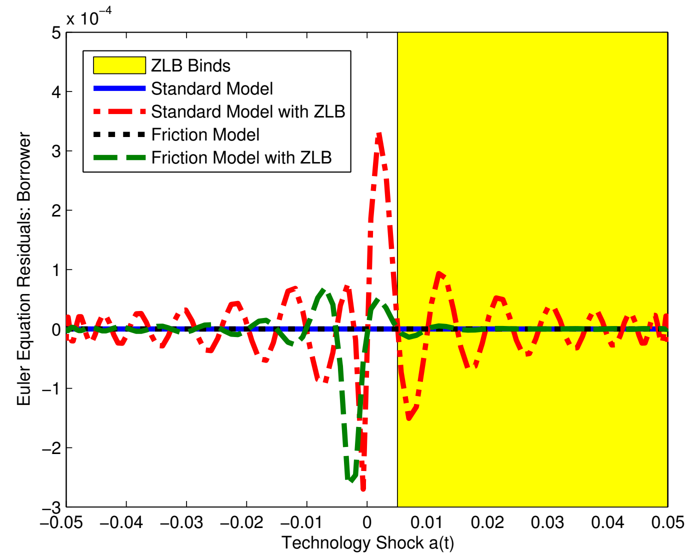

(b) Euler Equation: Borrower

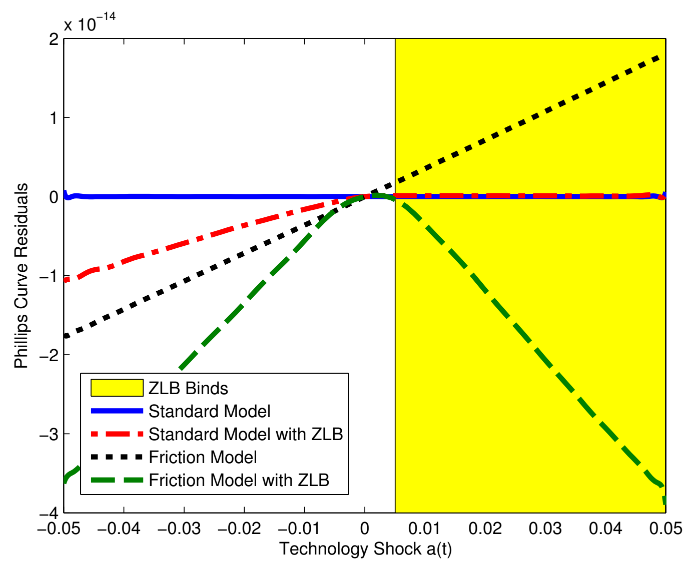

(c) Phillips Curve

Figure 25: Residuals of dynamic equations in non-linear solution.