Are Negative Supply Shocks Expansionary at the Zero Lower Bound?

Johannes F. Wieland*

University of California, San Diego

September 2014 – Version 2.6

For the latest version, **click here**

Abstract

The standard new Keynesian models predicts that temporary, negative supply shocks are expansionary at the zero lower bound (ZLB) because they raise inflation expectations and lower expected real interest rates, which stimulates consumption. This paper tests that prediction with oil supply shocks and the Great East Japan earthquake, demonstrating that negative supply shocks are contractionary at the ZLB despite also lowering expected real interest rates. A model with borrowing-constrained agents can match these findings for certain parameters, but simultaneously eliminates the channel by which the standard new Keynesian model generates fiscal multipliers above 1.

*University of California, San Diego, Economics Department, 9500 Gilman Dr. #0508 La Jolla, CA 92093-0508. Email: jfwieland@ucsd.edu. I am grateful to Yuriy Gorodnichenko for invaluable support and guidance, as well as to Olivier Coibion, Martin Lettau, Maurice Obstfeld, and David Romer for detailed comments that substantially improved this paper. I would also like to thank Dominick Bartelme, Gabriel Chodorow-Reich, James Hamilton, Joshua Hausman, Hashmat Kahn, Matteo Maggiori, John Mondragon, Demian Pouzo, Valerie Ramey, Andrés Rodríguez-Clare, Christina Romer, and Mu-Jeung Yang, as well as seminar participants at UC Berkeley, the NBER Summer Institute (EFCE), the Federal Reserve Bank of Boston, Chicago Booth, UC San Diego, HBS, UBC, LSE, LBS, CREI, the Federal Reserve Board, Michigan, and UC Santa Cruz 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, 3rd September 2011.

1 Introduction

How should policy makers combat recessions when monetary stabilization is constrained by the zero-lower bound (ZLB)? Some economists argue that demand-side policies, such as fiscal stimulus and forward guidance, are well-suited to address this challenge (e.g., Bernanke [2012], Eggertsson and Krugman [2011], Woodford [2012]), while others disagree that such policies are effective or desirable (e.g., Cochrane [2009], Taylor [2012]). Given a lack of clear empirical evidence around such rare events, proponents of demand-side policies have instead used new Keynesian models to support their policy recommendations. These models, originally designed and estimated to match normal times, typically predict that economies are governed by very different rules at the ZLB: wasteful government spending becomes very stimulative, and even negative supply shocks will be expansionary when the central bank is constrained. While opponents are skeptical whether these models describe the ZLB accurately enough to justify such bold policy choices, actual policy decisions in the U.S. and other countries have been based on these propositions (e.g, Temin and Wigmore [1990], Eggertsson [2012]).

This paper contributes to this debate in three ways. First, using the framework of Werning [2012], I establish conditions under which negative supply shocks are expansionary in the standard new Keynesian model when the ZLB constraint binds. Second, I show that oil supply shocks and the Great Japanese Earthquake are contractionary despite satisfying these conditions. Third, I incorporate agents with a nominal borrowing limit into an otherwise standard new Keynesian model, and show that this model can match my findings given certain parameters. In that case, the model also predicts fiscal multipliers below 1, unlike the standard new Keynesian model. These results suggest that the ZLB may not be such a radically different environment than asserted by Krugman and implied by the standard new Keynesian model.

I follow the set-up of Werning [2012] to analyze supply shocks in the standard new Keynesian model. An exogenous shock to the natural rate of interest causes the ZLB constraint binds up to period T. I show that any negative productivity shock that persists until $T_a < T$ raises output and consumption in the new Keynesian model. This occurs because sticky-price firms facing higher marginal cost gradually raise their prices, which raises expected inflation. At the ZLB, this lowers expected real interest rates which stimulates consumption and output. I call this mechanism the "inflation expectations channel," since inflation expectations are the key quantity through which temporary supply shocks affect real output at the ZLB. Next, I allow the productivity shock to be more persistent than the ZLB but finite in duration. In that case, the negative productivity shock is also expansionary if expected future nominal interest rates do not rise. I also derive additional restrictions under which these results apply for oil supply shocks in particular.

These predictions are not just a theoretical curiosity. A prominent result from the new Keynesian literature is that government spending multipliers are large at the ZLB (Christiano, Eichenbaum, and Rebelo [2011]; Woodford [2011]). But, as I show in Proposition 4, fiscal multipliers above 1 exist in the standard new Keynesian model if and only if there also exist negative supply shocks that are expansionary. Intuitively, in that model both fiscal shocks and supply shocks affect consumption at the ZLB only through the inflation expectation channel. Thus, by testing if negative supply shocks are expansionary at the ZLB, we test for the same mechanism whereby the standard new Keynesian model generates large fiscal multipliers.

Next, I determine the macroeconomic impact of two negative supply shocks at the ZLB: oil supply shocks and the Japanese earthquake in 2011. My results show that while inflation expectations rise and expected real interest rates fall, as predicted by the theory, these negative supply shocks are still contractionary overall. I also provide evidence against a weaker interpretation of the expectations channel. Because expected future nominal rates rise less at the ZLB, supply shocks should be less contractionary than in normal times. However, I estimate that oil supply shocks are *more* contractionary at the ZLB. This suggests that the expectations channel plays only a limited role in the propagation of such shocks.

These findings constitute a puzzle for the standard new Keynesian model. However, I show that adding a set of borrowing-constrained agents allows an otherwise standard new Keynesian model to match the data. The borrowing limit is nominal and binds until the ZLB periods ends. Further, these borrowers are detached from the labor market, preventing them from working their way out of the constraint. An increase in expected inflation reduces the borrowers' real purchasing power over the ZLB duration, and induces them to cut current consumption if the intertemporal elasticity of substitution (IES) is below 1. Further, for a sufficiently low IES, the cutback by borrowers exceeds the rise in consumption by the unconstrained agents, which allows the model to generate contractions to negative supply shocks at the ZLB. Simultaneously, fiscal multipliers fall below one, because the inflation expectations channel is now contractionary unlike in the standard new Keynesian model.

This paper is closely related to a recent literature that has explored how standard macroeconomic models make very different, even paradoxical, predictions at the ZLB. My empirical results concern primarily the "Paradox of Toil:" Eggertsson [2010b, 2011, 2012] documented cases where negative supply shocks are expansionary in standard new Keynesian models at the ZLB because they lower expected real interest rates.[1] However, subsequent work by Aruoba and Schorfheide [2013], Braun, Körber, and Waki [2012], Cochrane [2013], and Mertens and Ravn [2013] showed conditions under which negative supply shocks cause higher expected real interest rates in the new Keynesian model at the ZLB and are therefore contractionary. This paper contributes to this debate by establishing that negative supply shocks

[1]Other paradoxes include the "Paradox of Thrift" (Krugman [1998], Eggertsson and Woodford [2003], and Christiano [2004]) and the "Paradox of Flexibility" (Werning [2012], Eggertsson and Krugman [2011]).

are contractionary even in the "best-case" scenario where they reduce expected real interest rates.

This paper thus expands on earlier empirical tests of this paradoxical prediction. Mulligan [2010, 2012] argues that seasonal labor inflows do not appear to be contractionary and higher minimum wages do not appear to be expansionary in the data. However, Eggertsson [2010a] disputes that these are valid tests of the expectations channel, because the shocks are either forecastable or permanent and therefore cannot raise inflation expectations and lower real interest rates. Bachmann, Berg, and Sims [2011] show that consumers with above-average inflation expectations have lower willingness to spend even at the ZLB. However, in their cross-sectional analysis they cannot test whether raising aggregate inflation expectations is expansionary at the ZLB as predicted by the expectations channel. My analysis is robust to these critiques, as the supply shocks I use raise aggregate inflation expectations and lower expected real interest rates, yet still are contractionary.

This result also has important implications for the literature on fiscal multipliers. There is an ongoing debate whether fiscal multipliers at the ZLB are large or small (e.g., Cogan, Cwik, Taylor, and Wieland [2010]; Drautzburg and Uhlig [2011]; Dupor and Li [2014]; Christiano et al. [2011]; Woodford [2011]).[2] In both the standard new Keynesian model and the model with borrowing-constrained consumers, the inflation expectations channel generates fiscal multipliers are above 1 if and only if negative supply shocks are also expansionary. This suggests that when negative supply shocks are contractionary at the ZLB, then the inflation expectations channel is not a source of large fiscal multipliers.

This paper proceeds as follows. In Section 2, I show that temporary, negative supply shocks are expansionary at the ZLB in the standard new Keynesian model when they lower expected real interest rates. In Section 3, I reject this prediction for oil supply shocks and in Section 4 for the Great East Japan Earthquake. In Section 5, I construct a model that can match these findings. Section 6 concludes the paper.

[2]This literature largely focuses on positive aspects of fiscal policy. A notable exception is Werning [2012].

2 Predictions from the standard new Keynesian model

2.1 Model I adopt the framework of Werning [2012], a continuous-time adaption of the standard new Keynesian model of Woodford [2003]. The log-linearized equilibrium conditions conditions of the model are

$$dc(t) = \sigma^{-1}[i(t) - r(t) - \pi(t)]dt \tag{1}$$

$$d\pi(t) = \rho\pi(t)dt - \kappa\left\{\left[\sigma + \frac{\alpha(1-s_g)}{1-\alpha}\right]c(t) - \frac{1}{1-\alpha}a(t) + \frac{\alpha s_g}{1-\alpha}g(t)\right\}dt \tag{2}$$

$$i(t) = \max\{r(t) + \phi_\pi \pi(t), 0\}, \quad \phi_\pi > 1 \tag{3}$$

$$y(t) = (1-s_g)c(t) + s_g g(t) \tag{4}$$

where $c(t)$ is the log-linear deviation of consumption from steady-state, $i(t)$ is the nominal interest rate, $r(t)$ is the "natural rate of interest", $\pi(t)$ is the inflation rate, $a(t)$ the log-linear deviation of productivity from steady-state, and $g(t)$ the log-linear deviation of government spending from steady-state, and $y(t)$ is the log-linear deviation of output from steady-state. The paths for $\{r(t), a(t), g(t)\}$ are exogenously given and known. A bounded equilibrium of the model is a bounded sequence $\{c(t), \pi(t), i(t), y(t)\}$ such that equations (1)-(4) are satisfied given the sequence $\{r(t), a(t), g(t)\}$.

Equation (1) is the Euler equation, where consumption growth is proportional to the real interest rate $i(t) - \pi(t)$ net of the natural rate of interest $r(t)$. Equation (2) represents the new Keynesian Phillips curve, whereby current inflation is a function of expected future real marginal costs of production,

$$\pi(t) = \kappa\int_t^\infty e^{-\rho(s-t)}\left\{\left[\sigma + \frac{\alpha(1-s_g)}{1-\alpha}\right]c(s) - \frac{1}{1-\alpha}a(s) + \frac{\alpha s_g}{1-\alpha}g(s)\right\}ds.$$

Equation (4) is the national income accounting identity.

Equation (3) is the interest rate rule subject to the ZLB constraint. In contrast to

Werning [2012], this rule is not derived from an optimization problem, because I do not want to test empirical predictions conditional on optimal monetary policy. However, as $\phi_\pi \to \infty$ the policy rule approaches Werning's optimal policy without commitment, where inflation and the output gap are zero whenever the ZLB constraint does not bind.

That the central bank never deviate from the interest rate rule (3) implies the existence of a unique bounded equilibrium in the new Keynesian model.[3] I follow the existing literature by deriving and testing predictions based on this assumption. Cochrane [2013] studies the implication of adopting alternative equilibrium selection devices.

The positive parameters $\sigma^{-1}, \rho, \alpha, s_g$ are respectively the intertemporal elasticity of substitution, the discount rate, the decreasing returns in production and the steady-state share of government spending in output. The parameter $\kappa \geq 0$ is a non-linear function of deep model parameters and decreasing in the degree of price stickiness.

2.2 Supply shock I will compare two equilibria with different paths for productivity $a(t)$ but the same paths for the natural rate of interest $r(t)$ and government spending $g(t)$. In the first equilibrium productivity is always at steady-state, $a(t) = 0, \ \forall t$. In the second, productivity is below the steady-state for some time, $a'(t) < 0, \ \forall t < T_a$, and returns to steady state after a finite time $T_a < \infty$, $a'(t) = 0, \ \forall t \geq T_a$. As I show in the following sections, $T_a < \infty$ is the empirically relevant case. Note that T_a can be arbitrarily large. An example of this process is a step-function:

$$\begin{aligned} a'(t) &< 0 \qquad t < T_a, \\ a'(t) &= 0 \qquad t \geq T_a. \end{aligned}$$

In what follows I let the sequence $\{c(t), \pi(t), i(t), y(t)\}$ be the unique bounded equilibrium given the path $\{r(t), a(t), g(t)\}$ and the sequence $\{c'(t), \pi'(t), i'(t), y'(t)\}$ be the unique

[3] Cochrane [2011] criticizes the restriction to bounded equilibria that is universal in new Keynesian analysis.

bounded equilibrium given the path $\{r(t), a'(t), g(t)\}$. In both equilibria the Euler equation in integral form is satisfied,

$$c(t) = -\sigma^{-1} \int_t^{T_a} [i(s) - r(s) - \pi(s)] ds + c(T_a).$$

Because the model is forward looking and the paths for $\{a(t)\}$ and $\{a'(t)\}$ are identical after T_a, the endogenous variables are also identical after T_a, $\{c(t), \pi(t), i(t), y(t)\} = \{c'(t), \pi'(t), i'(t), y'(t)\}$ for all $t \geq T_a$. Differences in consumption before T_a then arise from differences in real interest rates,

$$c'(t) - c(t) = -\sigma^{-1} \int_t^{T_a} [i'(s) - \pi'(s) - (i(s) - \pi(s))] ds. \tag{5}$$

Thus, if expected future real interest rates are higher when productivity is lower, $\int_t^{T_a} [i'(s) - \pi'(s)] ds > \int_t^{T_a} (i(s) - \pi(s))] ds$, then consumption is lower, $c'(t) < c(t)$, and vice-versa.

2.3 Normal times I first consider the case where the ZLB constraint does not bind in either of the two equilibria. In that case, I obtain the conventional result that the equilibrium with lower productivity also features lower output and consumption.

Proposition 1 *If the ZLB never binds, $i(t) > 0,\ i'(t) > 0,\ \forall t$, then the negative productivity shock is contractionary*

$$\begin{aligned} y'(t) < y(t) \ \ \textit{and} \ \ c'(t) < c(t) \qquad & t < T_a, \\ y'(t) = y(t) \ \ \textit{and} \ \ c'(t) = c(t) \qquad & t \geq T_a. \end{aligned}$$

Proof See appendix A.1. ■

Higher marginal costs due to lower productivity causes sticky-price firms to gradually raising prices, which generates inflation. Because the central raises nominal interest rate more than one-for-one with inflation, real interest rates are higher when productivity is lower. Consumers respond to higher future real interest rates by reducing consumption today. In equilibrium the low-productivity economy therefore settles on a path of lower consumption and (typically) higher inflation.

2.4 Zero Lower Bound Following Werning [2012], I consider a scenario where the ZLB binds because the natural interest rate is negative up to a finite time T,

$$r(t) < 0 \qquad t < T,$$
$$r(t) \geq 0 \qquad t \geq T.$$

For simplicity and clarity I also assume that government spending is at steady-state,

$$g(t) = 0 \qquad \forall t.$$

Therefore, the path of $r(t)$ is the only disturbance in the first equilibrium (since $a(t) = 0$). For $t \geq T$, the interest rate rule then prescribes $i(t) = r(t) \geq 0$, which results in inflation and consumption at steady-state, $\pi(t) = c(t) = 0$ for all $t \geq T$. For $t < T$, this outcome is not feasible because the nominal interest rate would be negative, $i(t) = r(t) < 0$, which violates the ZLB constraint. Werning [2012] then shows that the unique equilibrium features deflation and low consumption while the ZLB binds, $c(t) < 0$ and $\pi(t) < 0$ for $t < T$. The depression arises because the ZLB constraint keeps the real interest rates above the natural rate of interest. This depresses consumption, which in turn leads to more deflation, higher real interest rates, and even lower consumption.

Next I compare this equilibrium outcome with the second equilibrium where productivity follows the path $a'(t)$. I denote the ZLB exit time for the second equilibrium by T'. For now I assume that the productivity shock vanishes before the economy exits the ZLB, $T_a < T' \leq T$. The low-productivity equilibrium then features higher consumption and output.

Proposition 2 *If $T_a < T' \leq T$, then the negative supply shock is expansionary*

$$y'(t) > y(t) \;\; \textit{and} \;\; c'(t) > c(t) \qquad t < T_a,$$
$$y'(t) = y(t) \;\; \textit{and} \;\; c'(t) = c(t) \qquad t \geq T_a.$$

Proof See appendix A.2. ■

As in normal times, the two equilibrium paths are identical for $t \geq T_a$ since the model is forward-looking and the exogenous variables follow the same path. For $t < T_a < T'$ the decline in productivity raises inflation just as in normal times. Higher inflation lowers real interest rates because nominal interest rates are fixed at zero for $t < T_a < T'$. Thus, the difference in consumption is now positively related to the difference in inflation expectations

$$c'(t) - c(t) = \sigma^{-1} \int_t^{T_a} [\pi'(s) - \pi(s)] ds.$$

Consumption and output rise because the decline in productivity raises expected inflation. I call this mechanism the "inflation expectations channel," since inflation expectations are the key quantity through which temporary supply shocks affect consumption and output in the standard new Keynesian model at the ZLB.

Appendix B shows that this proposition also holds in a non-linear model based on Rotemberg [1982]. This is an important check, because Braun et al. [2012] show that approximation errors induced by log-linearization can generate misleading comparative statics.[4]

The condition $T_a < T'$ rules out that supply shocks are more persistent than the ZLB constraint. If negative productivity shocks are more persistent than the negative natural interest rate shock then that assumption is violated. More subtly, large negative productivity shocks can endogenously shorten the ZLB duration T' such that it falls below T_a. In either

[4] Gavin, Keen, Richter, and Throckmorton [2013] use non-linear numerical methods to solve a medium-scale new Keynesian model with capital and find that negative productivity shocks are expansionary at the ZLB. Non-linearities are also the focus of Fernández-Villaverde, Gordon, Guerrón-Quintana, and Rubio-Ramírez [2012].

case the consumption response

$$c'(t) - c(t) = -\sigma^{-1} \int_{T'}^{T_a} [i'(s) - i(s)]ds + \sigma^{-1} \int_{t}^{T_a} [\pi'(s) - \pi(s)]ds,$$

becomes ambiguous. If expected nominal interest rates increase less than expected inflation then the decline in productivity raises consumption and vice-versa. It is difficult to directly establish $T' < T_a$ or $T' > T_a$, since this requires knowledge both of the supply shock $a(t)$ and the natural rate of interest $r(t)$. The following proposition therefore imposes restrictions on the nominal interest rate path, which may be easier to test.

Proposition 3 *If* $\int_{t_0}^{T_a} [i'(s) - i(s)]ds \leq 0$ *for* $t_0 < T'$ *then* $T_a < T' \leq T$ *and the negative supply shock is expansionary*

$$\begin{aligned} y'(t) > y(t) \ \ &\textit{and} \ \ c'(t) > c(t) \qquad t < T_a, \\ y'(t) = y(t) \ \ &\textit{and} \ \ c'(t) = c(t) \qquad t \geq T_a. \end{aligned}$$

Proof See appendix A.3. ■

Negative productivity shocks are expansionary in the standard new Keynesian model if they do not cause higher nominal interest rates. This was the case in Proposition 2 where the supply shock is less persistent than the ZLB. By contrast, if the productivity shock is more persistent than the ZLB, $T_a > T'$, the central bank raises nominal interest rates after T' in response to higher inflation. If the increase in nominal rates is sufficiently large, $\int_{T'}^{T_a} [i'(s) - i(s)]ds > \int_{t}^{T_a} [\pi'(s) - \pi(s)]ds$, then consumption and output will decline today. This channel should manifest itself in higher long-term bond yields today since they reflect expected future nominal rates. If we do not observe this increase, $\int_{t_0}^{T_a} [i'(s) - i(s)]ds \leq 0$, then we infer that $T_a < T'$ (based on the standard new Keynesian model) and the results from Proposition 2 apply.

2.5 Fiscal multipliers How serious should we take these predictions from the standard new Keynesian model? Exactly as serious as fiscal multipliers above 1 in the standard new Keynesian model. Both are caused by the same mechanism.

Take two exogenous sequences $\{r(t), a(t), g(t)\}$ and $\{r(t), a(t), g'(t)\}$, such that the second has strictly higher and finite cumulative government spending, $0 < \int_t^\infty [g'(s) - g(s)]ds < \infty$. Let $\{c(t), \pi(t), i(t), y(t)\}$ and $\{c'(t), \pi'(t), i'(t), y'(t)\}$ be associated unique bounded equilibria. Define the government spending multiplier as,

$$\mu_g(t) \equiv \frac{\int_t^\infty [y'(s) - y(s)]ds}{s_g \int_t^\infty [g'(s) - g(s)]ds}.$$

Consider next the exogenous sequence $\{r(t), \tilde{a}(t), g(t)\}$ and associated equilibrium $\{\tilde{c}(t), \tilde{\pi}(t), \tilde{i}(t), \tilde{y}(t)\}$. Let $\tilde{a}(t) = -\alpha s_g[g'(t) - g(t)] + a(t)$, which represents a negative productivity shock, $0 > \int_t^\infty [\tilde{a}(s) - a(s)]ds > -\infty$. Define the productivity multiplier as,

$$\mu_a(t) \equiv \frac{\int_t^\infty [\tilde{y}(s) - y(s)]ds}{\int_t^\infty [\tilde{a}(s) - a(s)]ds}.$$

Then fiscal multipliers above 1 exist if and only if there also exist negative productivity shocks that are expansionary.

Proposition 4 *Let $\alpha > 0$. Then $\mu_g(t) > 1$ if and only if there exists $\mu_a(t) < 0$.*

Proof See appendix A.4. ■

This is evident from equations (1)-(3). Both $a(t)$ and $g(t)$ enter only the Phillips curve (2) so that $g(t) > 0$ and $a(t) = -\alpha s_g g(t) < 0$ have the same effect on consumption and inflation. And since the fiscal multiplier is above 1 if and only if consumption increases, this can only occur if $a(t) < 0$ also raises consumption. Thus, by testing if negative supply shocks are expansionary at the ZLB, we test for the same mechanisms that generate large fiscal multipliers at the ZLB in the standard new Keynesian model.

3 Oil Supply Shocks

3.1 Are oil shocks like productivity shocks? The first empirical exercise uses oil supply shocks. This subsection establishes under what conditions oil supply shocks are also expansionary in the standard new Keynesian model. I first consider the case when oil is only used in production and then treat the case where oil is directly consumed. In appendix D, I also verify these predictions in a simple open economy model in which the home country imports oil. This is important because the countries in my sample are net importers of oil.

Consider an economy endowed with $O(t)$ units of oil each period and the production function of each firm is CES in oil and labor, $Y_i(t) = A(t)[(1-\xi)^{\frac{1}{\eta}}(N_i(t))^{\frac{\psi-1}{\psi}} + \xi^{\frac{1}{\psi}}(O_i(t))^{\frac{\psi-1}{\psi}}]^{\frac{\psi}{\psi-1}}$. Here ξ is the steady-state share of oil used in production, and ψ is the elasticity of substitution between oil and labor. The relative price of oil, $\frac{P^o_t}{P_t}$, adjusts instantaneously such that the oil market clears, $\int_0^1 O_{it}di = O_t$. In that model, the Phillips curve becomes

$$d\pi_t = \rho\pi_t - \kappa\left[\sigma + \frac{\psi^{-1}\xi}{1-\xi}\right]c_t + \kappa\left[1 + \frac{\psi^{-1}\xi}{1-\xi}\right]a(t) - \kappa\frac{\psi^{-1}\xi}{1-\xi}o(t)$$

with the other equations (1) and (3) unchanged. Since the Phillips curve is linear in oil supply and productivity, a negative oil supply shock has the same effect on consumption as a (scaled) negative productivity shock. Further, if a negative oil supply shock raises consumption, then it raises GDP even more. Thus, the results from Proposition 3 also apply to oil supply shocks in this model.

When oil is directly consumed we need to check an additional restriction. In that case expected inflation is

$$\int_t^T \pi(s)ds = \int_t^T \pi^y(s)ds + \gamma[p^o(T) - p^o(t)], \tag{6}$$

where π^y is inflation of the produced (sticky-price) good, p^o is the log-deviation of the real price of oil from steady-state, and γ is the share of oil in consumption. Since oil prices are flexible, a temporary decline in oil supply can introduce deflationary dynamics by raising current real oil prices above future real oil prices, $p^o(T) - p^o(t) < 0$. If the expected decline in real oil prices multiplied by its consumption share exceeds expected inflation in produced goods, $\int_t^T \pi^y(s)ds$, then the oil supply shock would reduce consumption at the ZLB. In appendix C, I show that we can check for this possibility by estimating the partial-equilibrium inflation response to an oil supply shock.

Proposition 5 *Consider a temporary oil supply shock,*

$$
\begin{aligned}
o'(t) &< 0 \qquad t < T_o, \\
o'(t) &\geq 0 \qquad t \geq T_o.
\end{aligned}
$$

Suppose the conditions in Proposition 3 are satisfied. If and only if the partial equilibrium response of inflation is positive, $[\pi'(t) - \pi(t)]|_{\{c(s)=0, s\geq t\}} > 0$, *then the oil supply shock is expansionary in the new Keynesian model*

$$
\begin{aligned}
y'(t) > y(t) \ \ \textit{and} \ \ c'(t) > c(t) \qquad t < T_o, \\
y'(t) = y(t) \ \ \textit{and} \ \ c'(t) = c(t) \qquad t \geq T_o.
\end{aligned}
$$

Proof See appendix C.2. ■

The partial-equilibrium inflation response $[\pi'(t) - \pi(t)]|_{\{c(s)=0, s\geq t\}}$ reveals whether oil shocks by themselves induce inflationary or deflationary dynamics. Note that this condition does not restrict the inflation response in general equilibrium. Rather, the restrictions imposed by the new Keynesian model imply that this condition also translates into a positive inflation response in general equilibrium, and thus an expansion of output and consumption at the ZLB. I will check this condition in the empirical analysis.

3.2 Are oil supply shocks "supply shocks" or "demand shocks"? In the new Keynesian literature (e.g., Woodford [2003]) shocks to the Phillips curve are typically labelled "supply shocks" and shocks to the Euler equation are labelled "demand shocks." However, this does not imply that oil supply shocks do not affect demand in the new Keynesian model. In any general equilibrium model, an oil supply shock will endogenously affect both the demand for goods and the supply of goods (since they have to be equal in equilibrium). Rather, the new Keynesian model restricts oil supply shocks to only affect demand through expectations of future real interest rates. When oil supply shocks raise expected real interest rates then demand (and output) declines in the model and vice-versa.

According to the model, the ZLB is then an ideal time to test this proposition. Subject to the restriction in Proposition 5, the model predicts an expansion to an oil supply shock, because real interest rates fall when the endogenous monetary policy response is constrained by the ZLB. Thus, this paper at least partially responds to Kilian's [2008] call for further empirical examination of the importance of endogenous policy responses to oil price shocks.[5]

3.3 Identification The challenge in uncovering oil supply shocks from production and price data is to separate demand from supply shocks. I follow Kilian's [2009] identification strategy. He estimates a VAR with three monthly variables: the growth in global oil production, $\Delta prod_t$, a measure of global economic activity, rea_t, and log real oil prices, rpo_t. I denote the data vector by $v(t) = (\Delta prod_t, rea_t, rpo_t)'$. Then the structural VAR representation is,

$$A_0 v_t = \alpha + \sum_{i=1}^{24} A_i v_t + \varepsilon_t.$$

Like Kilian [2009] I assume that A_0^{-1} is lower triangular. This assumption implies that oil supply does not respond to demand shocks within a month.[6] The structural shocks, ε_t, can then be recovered from the reduced-form errors, $e_t = A_0^{-1}\varepsilon_t$, using a Cholesky decomposition.

[5] [Kilian, 2008, p.893]: "[...] more recent studies have focused on the importance of endogenous policy responses to oil price shocks in the context of theoretical macroeconomic models [...]. The importance of this channel depends to a large extent on the definition of the counterfactual and on the modeling assumptions, making further empirical work on endogenous monetary policy responses all the more important."

[6] According to Kilian and Murphy [2011, p. 7], "it is widely accepted that this [short-run oil supply supply] elasticity is low on impact."

The sample is January 1973 through December 2012.

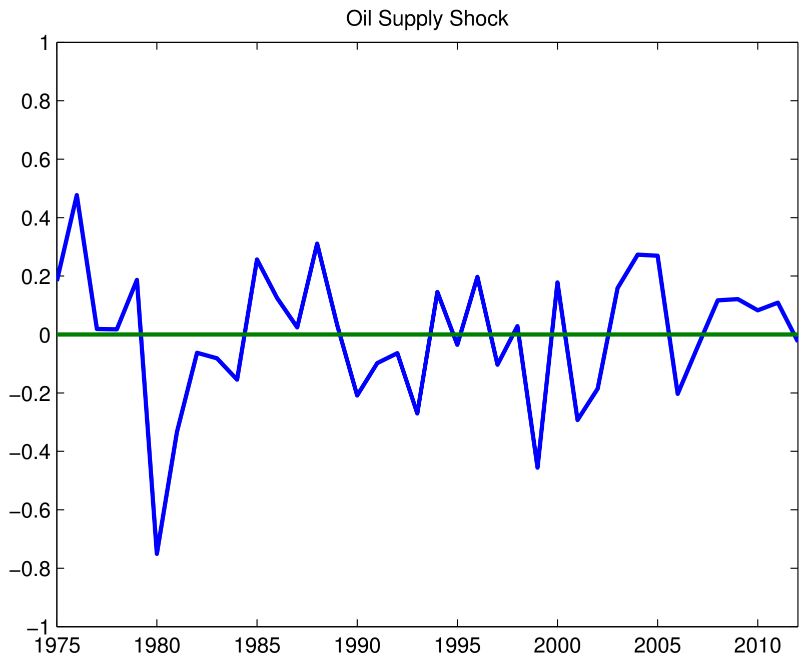

Figure 1 – Oil supply shocks derived using the methodology in Kilian [2009] aggregated to an annual frequency. Positive values show upward surprises, i.e., a positive oil supply shock.

Figure 1 plots the identified oil supply shocks aggregated to an annual frequency. As in Kilian [2009] a large negative oil supply shock in 1980 stands out. By contrast, oil supply shocks were on average positive during the financial crisis years. The VAR instead attributes the swing in oil prices around the financial crisis to demand shocks.

Similar to Kilian's 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 shocks capturing reductions in supply rather than current or expected negative demand shocks. Importantly, the estimated long-run response of real oil prices to an oil supply shock is essentially zero — the increase in real oil prices is less than 0.05% after five years and less than 0.01% after 20 years. This is consistent with the assumption that the oil supply shocks are temporary ($T_a < \infty$) as assumed in Propositions 2 and 3.

3.4 Analysis Since the oil supply shocks extend into the current ZLB episode, I include the U.S., the U.K., the Eurozone, Canada, Sweden, and Japan in my estimation. I exclude members of currency unions because standard new Keynesian theory predicts that they exhibit standard (non-paradoxical) dynamics at the ZLB (Nakamura and Steinsson [2011], Farhi and Werning [2012b]).

By pooling data across countries, I estimate the average treatment effect of oil supply shocks. While the magnitudes of these effects may vary across countries, there is no a priori reason why their *sign* should be different. Since I am interested in the sign of output and inflation responses, the cost of pooling data is likely small relative to the gain in power.

Table 1 – ZLB Dates

Country	Japan		U.S.	U.K.	Eurozone	Canada	Sweden
Start Date	10/1995	12/2008	12/2008	3/2009	1/2009	5/2009	7/2009
End Date	7/2006	-	-	-	-	5/2010	6/2010

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 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 G, I therefore show that my results are robust to excluding the Eurozone. In the same appendix, I also obtain similar results using projection methods that take differences in the tightness of the ZLB constraint across countries into account. However, note that the purpose of the cut-offs is to select states where central banks do not raise real interest rates in response to supply shocks. Conditional on the separation being accurate, the actual cut-off values are not important.

I first test if negative oil supply shocks have a positive on inflation expectations holding consumption and output fixed. Only then would the new Keynesian model predict an expansion at the ZLB as shown in Proposition 5. I construct a four-quarter-ahead inflation forecast from Consensus Economics forecasts. Consensus Economics publishes detailed quarter-by-quarter inflation forecasts every three months. To construct an inflation forecast at monthly frequency, I feed annual inflation forecasts, which are published by Consensus Economics every month, and the quarter-by-quarter forecasts into a Kalman filter, which "fills-in" the gaps when the quarterly forecasts are not published. This yields a monthly four-quarter-ahead inflation forecast. For instance, 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 — for an oil supply shock occurring in January, I only count price increases starting in April. The details of the Kalman filter estimation are relegated to Appendix E.

I regress changes in 4Q-ahead inflation expectations,[7] $\Delta\pi_t^{e,4}$, on lagged values, lagged oil shocks, lagged controls, and country-fixed-effects,

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

Since inflation forecasts are *published* at the beginning of the month,[8] 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]). To capture the partial equilibrium effect on inflation, I control for changes in industrial production, unemployment and long-term bond rates. Lag lengths are determined by the Akaike Information Criterion (AIC).[9]

Figure 2(a) plots the dynamic impulse response function (IRF) from Equation (7). The dashed lines are 95% confidence intervals based on Driscoll and Kraay [1998] standard errors, which are robust to heteroscedasticity, temporal dependence, and cross-country dependence. In Appendix F I prove that the standard asymptotic covariance matrix remains valid even

[7] I difference forecasts for same time interval that were made this period and last period.

[8] 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.

[9] Coibion [2012] shows that the AIC outperforms BIC in short samples.

though the oil shocks in Equation (7) are estimated. This result follows because the oil shocks are the OLS *residuals* of a first-stage. Unlike the first-stage predicted values, the residuals are asymptotically orthogonal to first-stage sampling uncertainty. In Figure 2(a) 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 (4 basis points) after twelve months, and tails off thereafter. In Appendix G, I obtain larger responses in inflation expectations when I use a less conservative estimate that includes changes in inflation expectations for the current quarter. Thus, for either of these measures, oil shocks raise inflation expectations for a given level of output, consistent with the requirement in Proposition 5.

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

$$\Delta y_{i,t} = \sum_{j=1}^{n} \beta_j \Delta y_{i,t-j} + \sum_{j=0}^{k} \gamma_j oil_{t-j} + \eta_i^y + \varepsilon_{i,t}^y. \tag{8}$$

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. I drop the lagged dependent variables in a robustness exercise.

In Figure 2(b) I plot the dynamic IRF of log IP 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 0.7%. In the following six months the IRF reverts back to zero, and becomes statistically insignificant. Thus, a negative oil supply shock causes a contraction in IP at the ZLB.

Estimating Equation (8) with the unemployment rate instead of IP provides corroborating evidence. The dynamic IRF in Figure 2(c) exhibits a statistically-significant increase in the unemployment rate by 0.1 percentage points after a one-standard-deviation negative

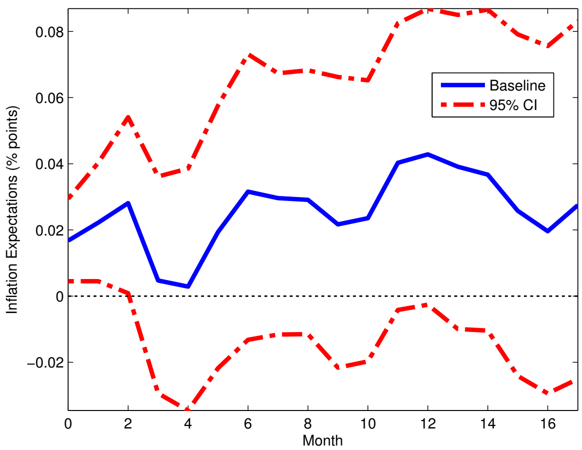

(a) 4Q Ahead Expected Inflation

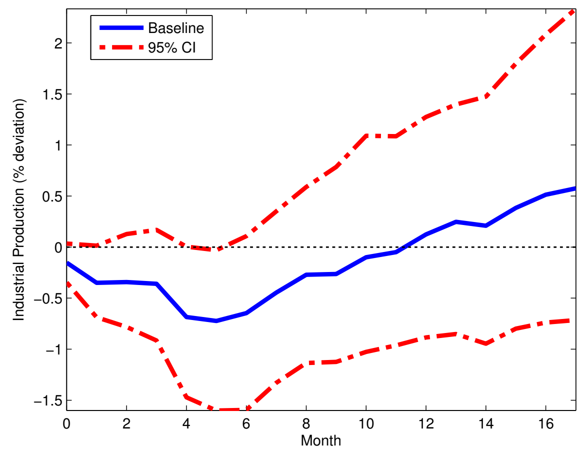

(b) Log Industrial Production

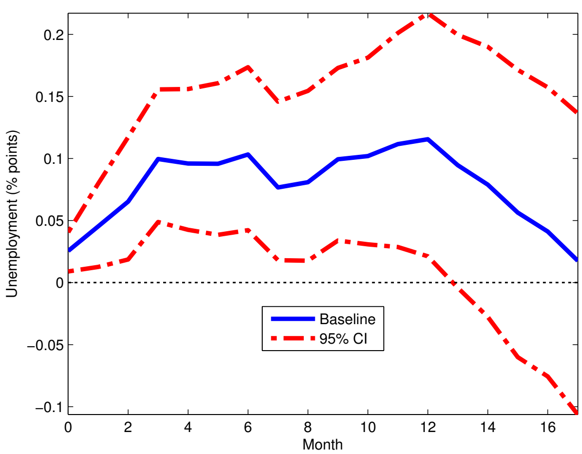

(c) Unemployment

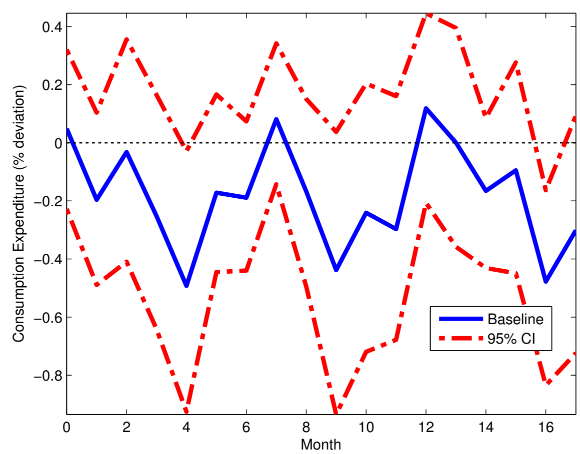

(d) Consumption Expenditures

Figure 2 – 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 Driscoll-Kraay covariance matrix. Lag lengths are set according to AIC: For the dependent variable they are set to 12 for expected inflation, 48 for industrial production, 36 for unemployment, and 36 for consumption. Lag lengths for oil shocks are set to 24 for each variable. See Section 3 for details.

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 simple new Keynesian model. To address this concern, I estimate Equation (8) using consumption expenditure data for Japan and the U.S., which is available at a monthly frequency. I use a 12-month difference for consumption, because the Japanese data is very volatile month-to-month. The IRF in Figure 2(d) is quite choppy because the Japanese consumption series is volatile, but nonetheless displays a statistically-significant decline in consumption expenditures four months after an oil shock.

These effects are inconsistent with Proposition 3, unless large increases in long-term nominal bond rates raise real rates. To assess whether long-term interest rates have risen, I estimate the impact of oil shocks on nominal bond yields at various maturities. As with consumption, I use 12-month differences to minimize the influence of noise in the data. Figure 3(a) shows that nominal bond yields at all maturities greater than or equal to five years exhibit a statistically significant *decline* after an oil supply shock at the ZLB.

Ten-year real bond yields are available only for a sub-sample. Bloomberg has such data for Japan from 4/2004-3/2009, for the U.K. from 7/1992-today, for the U.S. from 1/1997-today, for Sweden from 8/2008-today, and for Canada from 3/2000-today. For Japan I also construct a synthetic real bond from 10-year nominal bond yields and 10-year swap rates, which are available from 3/2007-today. For the overlap period, 3/2007-3/2009, I use synthetic real bond yields rather than actual real bond yields because the former behave less erratically. With these data, I estimate that real bond yields decline only slightly more than nominal bond yields. This suggests that long-run nominal rates largely capture the long-run real interest rate response to oil supply shocks (Figure 3(b)).[10] Thus, oil supply shocks have contractionary effects even though expected future real interest rates fall. This

[10]Expected real interest rates are a more accurate indicator of a central bank's stance, because a given nominal interest rate path is consistent with many different inflation paths (Werning [2012], Cochrane [2013]).

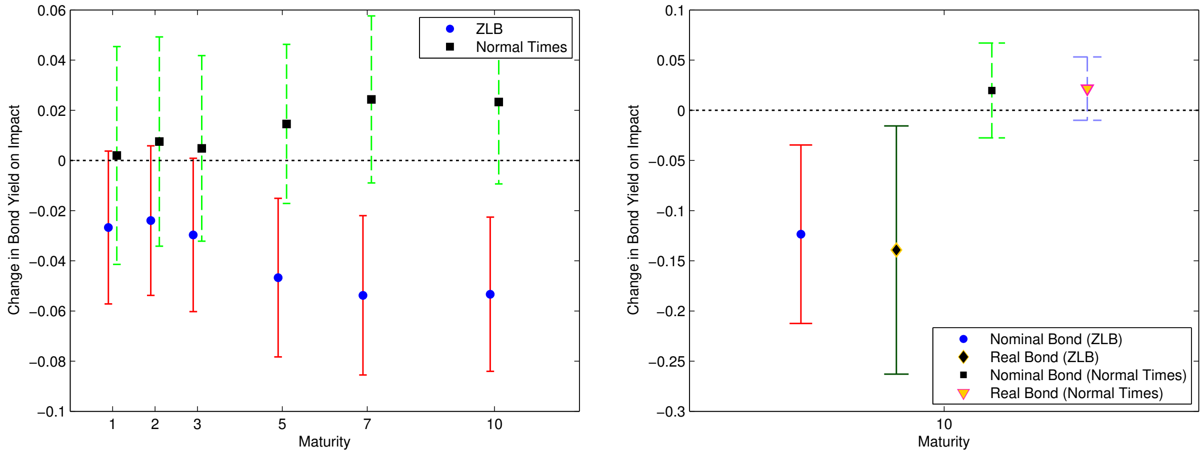

(a) Nominal Bond Yields

(b) 10-Year Nominal & Real Yields (sub-sample)

Figure 3 – *Left Panel:* Change in nominal bond yield on impact of an oil supply shock. Point estimates for the ZLB are solid circles, and point estimates for normal times are solid squares. Error bars show the 95% Driscoll-Kraay confidence interval around the point estimates. Estimated based on an autoregressive distributive lag equation, where the dependent variable enters with 48 lags and oil shocks with 24 lags based on the AIC criterion. *Right Panel:* Change in nominal and real ten-year bond yields on impact of an oil supply shock. Estimated on sub-samples where ten-year real bond yields are available. Error bars show the 95% Driscoll-Kraay confidence interval around the point estimates. Estimated based on an autoregressive distributive lag equation, where the dependent variable enters with 24 lags and oil shocks with 24 lags. See Section 3 for details.

constitutes a puzzle for the standard new Keynesian model, because temporary, negative supply shocks that lower expected real interest rates should be expansionary through the expectations channel (Proposition 3). For example, with log utility, the decline in ten-year nominal bond yields alone should generate a 0.5% rise in consumption following an oil supply shock. These results indicate that other propagation mechanisms trump the expectations channel for oil supply shocks, and that central banks respond by loosening future monetary policy.

3.5 Robustness If the oil supply shocks in Kilian [2009] are well-identified, then they would be orthogonal to the current state of the economy. Thus, excluding lagged dependent variables from Equations (7) and (8) should yield similar results. In Figure 4 I plot the IRFs from this estimation. The results for all variables 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.

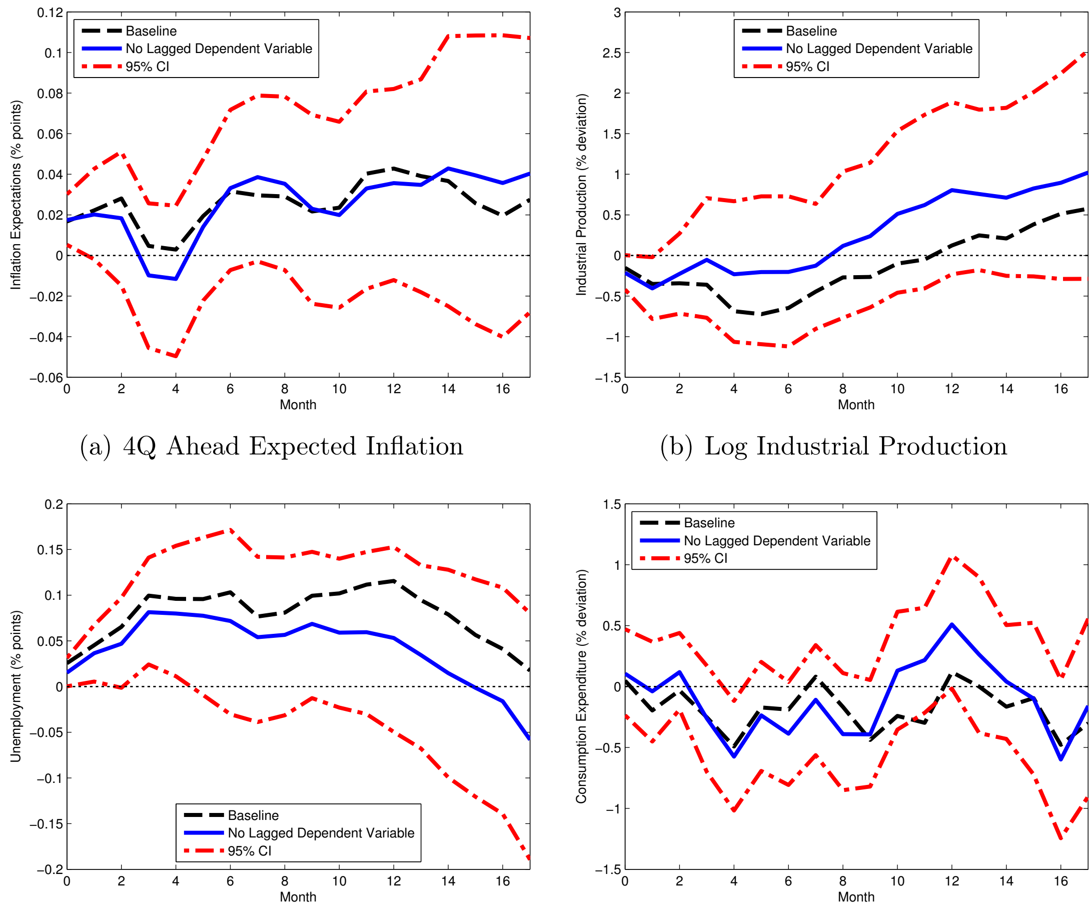

(a) 4Q Ahead Expected Inflation

(b) Log Industrial Production

(c) Unemployment

(d) Consumption Expenditures

Figure 4 – 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.

In Appendix G I establish further robustness of my findings. First, I restrict to estimation to Japanese data from before 2008, which demonstrates that the results are not driven by the recent recession or pooling data across countries. I also show that the results are robust to excluding the Eurozone, using projection methods, using HP-filtered IP and unemployment, and alternative lag lengths. In Appendix H, I corroborate my findings with an event study of the oil production disruptions from the Libyan civil war.[11]

[11] I am grateful to James Hamilton for this suggestion.

3.6 ZLB vs normal times Despite the negative results so far, it is possible that the expectations channel mitigates contractionary effects of oil supply shocks at the ZLB. If the central bank raises nominal interest rates in normal times but does not do so at the ZLB, then the standard new Keynesian model would predict a smaller contraction at the ZLB. Figure 3 shows that over the sample from 1986-2007, an oil supply shock raises nominal and real bond yields when the ZLB does not bind. I apply 1986 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 potential 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.[12] This evidence provides support for the notion that central banks are more aggressive in normal times.

I therefore estimate Equation (8) on normal-times data, to test if the more aggressive monetary stance causes larger contractionary effects in normal times. However, the IRFs Figure 5 show that, contrary to the previous reasoning, the average contractionary effects of an oil shock are much *stronger at the ZLB than in normal times.* The confidence intervals for the ZLB response typically overlap with the IRF for normal times, so one cannot reject that they are equal. But even an equal output response despite significantly different real interest behavior is evidence against this weaker prediction of the expectations channel.

3.7 Other interpretations Another explanation for my findings is that temporary oil supply shocks are contractionary, because they have permanent negative wealth effects. In appendix D I analyze wealth effects in a simple open-economy new Keynesian model with incomplete markets (see Farhi and Werning [2012b,a]). Output in the incomplete-market model is given by,

$$y^{IM}(t) = y^{CM}(t) + \delta_y^{IM}. \tag{9}$$

[12] For cut-offs in the 1970s and early 1980s nominal bond yields also fall in normal times, but by less than at the ZLB. The real effects of oil supply shocks are similar when these earlier data are included.

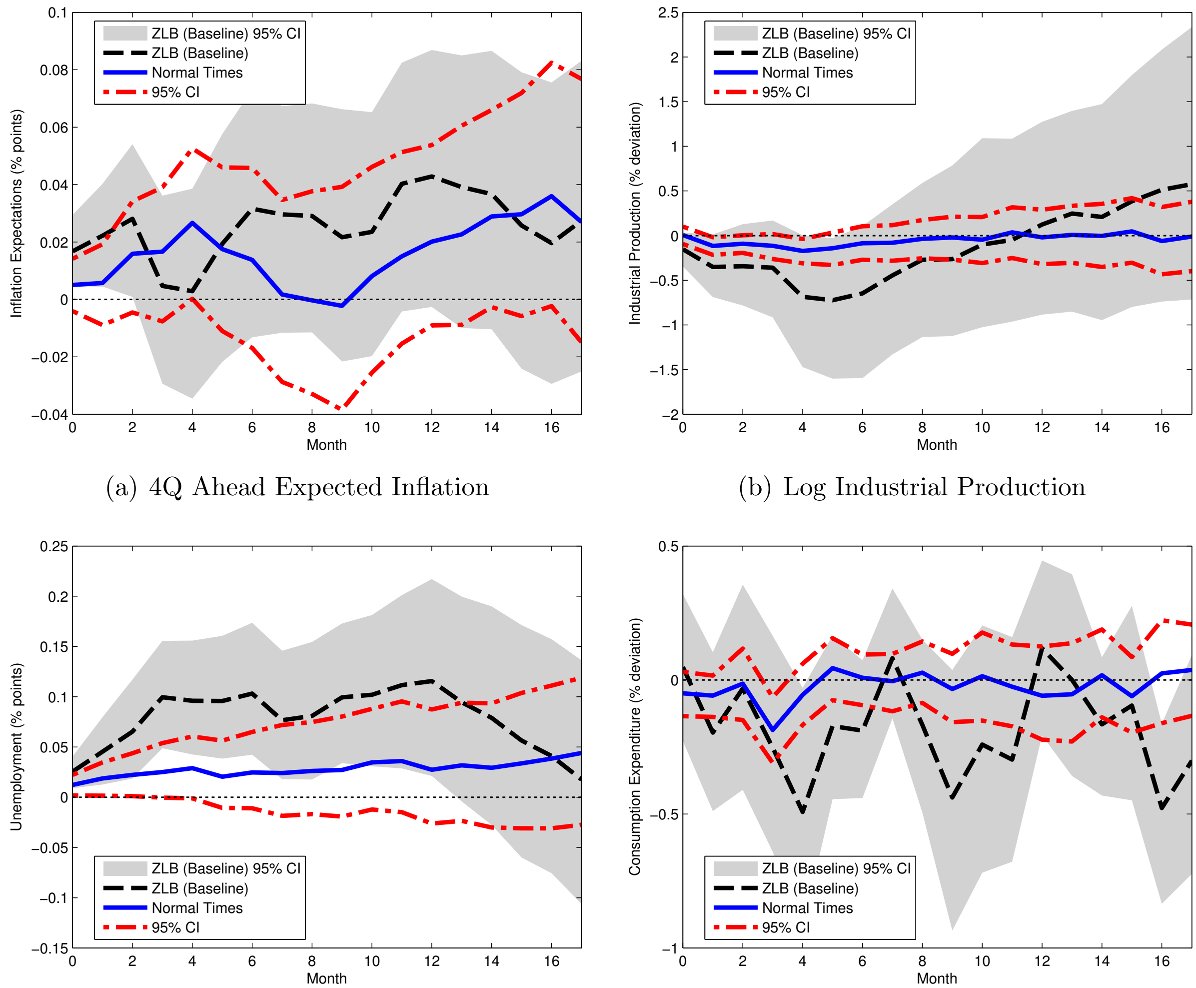

(a) 4Q Ahead Expected Inflation

(b) Log Industrial Production

(c) Unemployment

(d) Consumption Expenditures

Figure 5 – Impulse Response Functions to oil supply shocks 1985-2007: 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 Driscoll-Kraay covariance matrix. Inflation expectations data begin in 1995 because of data-constraints. See Section 3 for details.

Here $y^{CM}(t)$ is output in the analogous complete-market model, where foreign oil supply shocks have no wealth effects, and δ_y^{IM} captures the (time-invariant) wealth effect of the shock. Because negative supply shocks are expansionary in the complete market model, $y^{CM}(t) > 0$, a large negative value for δ_y^{IM} could produce my findings.

However, the model suggests at least three reasons for why wealth effects are an insufficient explanation for my results. First, appendix D shows that a decline in wealth reduces consumption but *raises* labor input and GDP for conventional parameters in the standard open-economy new Keynesian model. Thus, even theoretically it is not clear that $\delta_y^{IM} < 0$. Second, if $y^{CM}(t) > 0$ and $\delta_y^{IM} < 0$, then we should observe smaller contractionary effects in the short-run than in the long-run, $y(t) = y^{CM}(t) + \delta_y^{IM} > y(\infty) = \delta_y^{IM}$. By contrast, the oil supply shocks in Figure 2 have larger contractionary effects in the short-run. Thus, even if wealth effects are present, the data point towards a model where the complete-market effects are also negative, $y^{CM}(t) < 0$. Third, the contractionary effects in Figure 5 are larger at the ZLB than in normal times, whereas there is no a priori reason why the same shock should have larger negative wealth effects at the ZLB. In short, while I cannot rule out that some wealth effects are present, the evidence suggests that they are insufficient to account for the behavior of real variables following an oil supply shock.

4 The Great East Japan Earthquake

On Friday March 11^{th} 2011, a magnitude-9.0 earthquake off the Japanese eastern coastline triggered a tsunami that caused extensive damage to Japanese 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 Japanese consensus inflation forecast for 2011 and 2012 rose respectively by 0.3 percentage points and 0.2 percentage points, as shown in Figure 6.[13] In addition, the 10-year inflation swap rate rose from an average of -3 basis points from March 7^{th} through 10^{th} to an average of $+3$ from March 14^{th} through 18^{th}.

While professional forecasters and markets expected higher inflation following the earthquake, they revised output forecasts down, as shown in Figure 6(a). Japanese output for 2011 was forecasted to be about 1.2% below pre-earthquake predictions. 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, the earthquake had contractionary effects despite raising inflation expectations at the ZLB. I first determine whether this earthquake satisfies the conditions under which capital destruction is expansionary in the standard new Keynesian model. I then discuss to what extent demand-based explanations can account for the contraction.

Two conditions need to be checked: first, that the earthquake does not raise expected future nominal interest rates and, second, that this shock is of finite duration. The first condition is satisfied, as the yield on 10-year government bonds fell from an average of 1.30% from March 7^{th} through 10^{th} to 1.22% from March 14^{th} through 18^{th}. The 20- and 30-year bond yields also declined.[14] An important reason for the decline in nominal interest rates was a loosening of monetary policy in response to the earthquake and tsunami. I attribute this response to the earthquake in the sense that monetary easing would not have occurred otherwise. Of course, from the new Keynesian perspective it becomes even more puzzling that capital destruction was contractionary despite higher inflation expectations and looser monetary policy.

A priori, the second condition also seems reasonable — the capital stock destroyed by the earthquake will be rebuilt as the economy converges to its balanced growth path. Consistent with this prior, Brueckner [2014] estimates that it takes 20 years for the capital stock to

[13] I use February and April forecasts, because the March forecasts were released only shortly after the earthquake so that some of them were outdated (e.g., the Morgan Stanley forecast).

[14] This data is from Bloomberg.

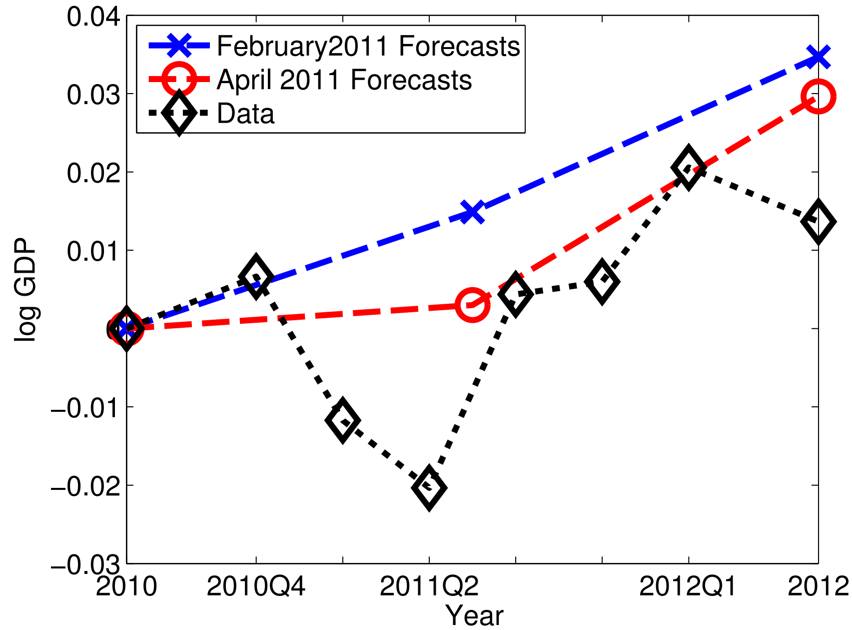

(a) Japan: GDP

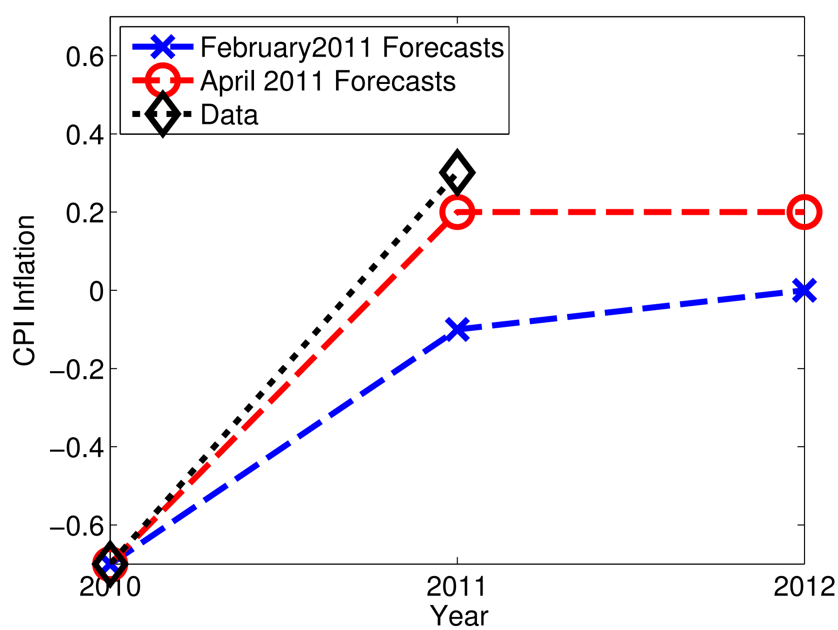

(b) Japan: Inflation

Figure 6 – 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 2012 and quarterly from 2010Q4 until 2012Q1. CPI data is annual year-on-year inflation. See Section 4 for details.

return to steady-state after an earthquake. Further, 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 6(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 change in Japan's nuclear power policy following the Fokushima failure disaster is perhaps the most salient sign of a persistent effect by the earthquake. This has raised the cost of energy in Japan and will likely continue to do so for some time. However, to the extent that these nuclear power plants would have been replaced by other forms of energy production at some finite time T_n, then the rise in energy prices is temporary, which constitutes exactly the kind of temporary supply shock that should be expansionary according to the standard new Keynesian model.

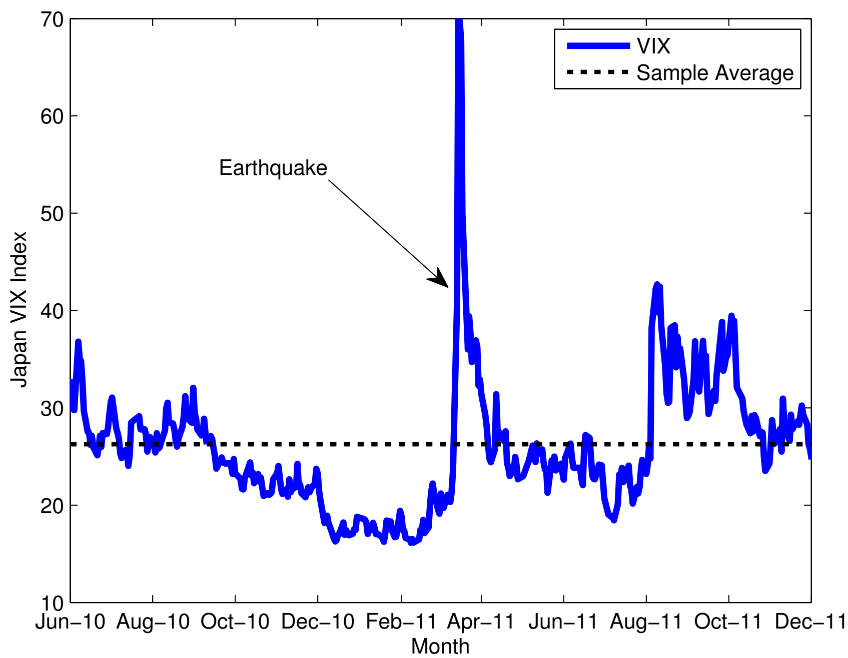

Figure 7 – Nikkei VIX index. Sample average is computed over January 2001 to July 2012. See Section 4 for details.

The standard new Keynesian model can then replicate contraction only if the earthquake simultaneously triggered a large and persistent negative demand shock. A prominent candidate is a rise in uncertainty associated with the potential consequences of the earthquake and tsunami. The earthquake was indeed associated with a rise in uncertainty as measured by the Nikkei VIX index shown in Figure 7. However, by April 2011 volatility returned to its historical average. This suggests that elevated uncertainty cannot explain why GDP in the second quarter of 2011 remained 3.5% below the pre-Earthquake trend (Figure 6(a)).

Thus, the large and persistent contraction in Japan after the earthquake appears to be inconsistent with the expansionary predictions of the standard new Keynesian model in Section 2.

5 A data-consistent model

The standard new Keynesian model imposes that supply shocks can only affect demand through real interest rates. I now loosen this restriction by introducing a fraction of borrowing-constraint agents with fixed nominal purchasing power. For certain parameters, this modification generates a contraction to negative supply shocks at the ZLB and, simultaneously, the fiscal multiplier will fall below 1.

As before, I assume that the natural rate of interest is strictly negative up to time T and positive thereafter. In equilibrium this causes the ZLB to bind up to time T. After time T the ZLB will cease to bind and the central bank will implement a zero-inflation target forever, $\pi(t) = 0$ for all $t \geq T$.

In contrast to the standard new Keynesian model, a fraction χ of all agents will be unemployed until time T and subject to a binding *nominal* borrowing limit B.[15] Given an initial date t_0, one can interpret B as the agents' remaining credit line plus unemployment benefits that accrue over $[t_0, T]$. I will call these agents "borrowers", denoted by b, but one can also interpret them as long-term unemployed. Limiting labor market access prevents the borrowers from working their way out of the borrowing limit. I can also allow for labor

[15] An earlier version of this paper introduced borrowing frictions in a more complex but complementary way to also match asset pricing moments.

market access so long as the agents nominal income is fixed, in which case B also includes labor income.

Over the interval $[t_0, T]$ the borrower optimally satisfies,

$$dc^b(t) = -\sigma^{-1}[r(t) + \pi(t)]dt \tag{10}$$

$$B = \int_{t_0}^{T} C^b(s)P(s)ds = \int_{t_0}^{T} C^b(s)P(t)e^{\int_{t_0}^{s}\pi(z)dz}ds. \tag{11}$$

Equation (10) is the borrowers' Euler equation, equation (11) is the borrowing limit, and $c^b(s) = \ln C^b(s) - \ln \bar{C}$ denotes the log-deviation from steady-state. Note that the approximations to this model are around a zero-inflation steady-state where all agents are identical.[16] The Euler equation now determines how a borrower allocates the available nominal funds B over $[t_0, T]$. After T these agents re-enter the labor market and are no longer subject to the borrowing constraint.

The solution to the borrowers problem at t_0 is,

$$C^b(t_0) = \left[\int_{t_0}^{T} e^{\int_{t_0}^{s}\pi(z)dz - \sigma^{-1}\int_{t_0}^{s}[\pi(z)+r(z)]dz}ds\right]^{-1}\frac{B}{P(t_0)}$$

The borrowers consumption at t_0 is increasing in the real borrowing limit $\frac{B}{P(t_0)}$ and decreasing in the natural rate of interest $r(t)$. The effect of expected inflation is ambiguous. Higher expected inflation induces the borrower to shift consumption forward per the Euler equation (10). However, it also reduces the borrower's expected real purchasing power which depresses consumption today. If the intertemporal elasticity of substitution σ^{-1} is less than one, then the second effect dominates and borrowers will reduce current consumption when they expect higher inflation. By contrast, in the standard new Keynesian model an increase in inflation expectations raised consumption at the ZLB.

[16] Farhi and Werning [2012b] use a similar trick to analyze open economies with incomplete market, where shocks can cause a permanent redistribution of wealth.

After normalizing the debt limit to $B \equiv \psi(T - t_0)P(t_0)\bar{C}^b$, the linearized solution is,

$$c^b(t_0) = \ln\psi - \int_{t_0}^{T} \frac{(T-s)}{(T-t_0)}\pi(s)ds + \sigma^{-1}\int_{t_0}^{T} \frac{(T-s)}{(T-t_0)}[\pi(s) + r(s)]ds. \tag{12}$$

The parameter $\psi < 1$ captures the tightness of the borrowing limit. I set it to a sufficiently low value such that the ZLB will always bind up to time T.

The borrowers' behavior is then completely described by the Euler equation (10) along with the boundary condition (12). These agents capture the notion that high prices depress consumption by reducing real income. This stands in contrast to the standard new Keynesian model, where supply shocks only affect demand through changes in real interest rates. In the context of oil supply shocks, the importance of income effects has been highlighted by Edelstein and Kilian [2009] and Kilian and Park [2009] among others.

The remainder of the model follows the standard new Keynesian model. The other $1-\chi$ agents, denoted by l for "lenders," supply labor to a competitive labor market and not subject to borrowing frictions. Thus, lenders satisfy the standard Euler equation,

$$dc^l(t) = \sigma^{-1}[i(t) - r(t) - \pi(t)]dt.$$

Firms produce output using decreasing-returns technology, $y(t) = a(t) + (1-\alpha)l(t)$, and are subject to Calvo pricing frictions. Letting the real wage be denoted by ω, then the new Keynesian Phillips curve is given by,

$$d\pi(t) = \rho\pi(t)dt - \kappa\left\{\omega(t) - \frac{1}{1-\alpha}a(t) - \frac{\alpha}{1-\alpha}y(t)\right\}dt.$$

I assume that lenders' labor supply is perfectly elastic, which implies that the real wage is $\omega(t) = \sigma c^l(t)$. For $t \geq T$, that equality also holds for borrowers as they re-enter the labor market. Finally, the national income accounting identity is given by,

$$y(t) = (1 - s_g)[\chi c^b(t) + (1 - \chi)c^l(t)] + s_g g(t).$$

This set-up is akin to Eggertsson and Krugman [2011], with the time before T corresponding to their short-run where borrowing limits are binding, and the time after T to their long-run when debt is paid off. However, the borrowing limit in Eggertsson and Krugman [2011] is real, whereas it is nominal in this model. At least in the short-run it is plausible that the borrowing limit is nominally sticky.

To solve the model we need to determine the boundary conditions at T. After T there are no more shocks, so $i(t) = r(t) = \rho$, inflation is zero, and output is constant. Further, the consumers' Euler equations imply that consumption is constant, $c^b(t) = c^b(t + s)$ and $c^c(t) = c^c(t + s)$ for $t \geq T,\ s \geq 0$. Because labor supply after T is perfectly elastic, borrwers will pay off debt by working harder rather than reducing their consumption,

$$\begin{aligned} c^b(T) &= 0 & \qquad c^l(T) &= 0 \\ L^b(T) &= \bar{L} + \rho\psi \frac{P(t_0)}{P(T)} & \qquad L^l(T) &= \bar{L} - \rho\psi \frac{\chi}{1-\chi} \frac{P(t_0)}{P(T)} \end{aligned}$$

Higher inflation from t_0 to T reduces the real value of debt, so that borrowers need to work less hard to repay their debt. However, unlike Eggertsson and Krugman [2011], borrowers cannot respond by raising their consumption at t_0 because they remain constrained by the nominal borrowing limit B.

As before we compare the equilibrium when productivity is low until T, $a'(t) < 0$ for $t < T$ and $a'(t) = 0$ for $t \geq T$, with the equilibrium when productivity is unchanged $a(t) = 0$. Then, the negative productivity shock is contractionary if the intertemporal elasticity of substitution σ^{-1} is sufficiently low.

Proposition 6 *Let $\chi > 0$. Then there exists $\tilde{\sigma}^{-1} \in (0, 1)$ such that for all $\sigma^{-1} < \tilde{\sigma}^{-1}$*

$$y'(t) > y(t) \quad \text{and} \quad c'(t) > c(t) \qquad t < T,$$
$$y'(t) = y(t) \quad \text{and} \quad c'(t) = c(t) \qquad t \geq T,$$

Proof See appendix A.5. ■

The productivity shock again raises expected inflation, which raises consumption by lenders and lowers consumption by borrowers if $\sigma^{-1} < 1$. By letting $\sigma^{-1} \to 0$ we reduce the response of lenders to lower real interest rates to zero, whereas the fall consumption by borrowers remains strictly negative. The proposition then follows by continuity. Thus, so long as the unconstrained agents do not respond much to real interest rates, the purchasing power effect will dominate and the negative supply shock will be contractionary.

The model also suggests possible explanations for the empirical observations that were puzzling from the standard new Keynesian perspective. First, oil supply shocks may be more contractionary at the ZLB because the share of unemployed borrowers is higher than in normal times. Second, if the ZLB duration is endogenous, then the contraction in spending by borrowers may induce the central bank to keep interest rates lower for longer. This channel could explain the decline in long-term bond yields.

In short, the model is (qualitatively) consistent with the empirical results unlike the standard new Keynesian model. But it also holds a cautionary message for fiscal activism at the ZLB.

Proposition 7 *Let $\alpha > 0$. Then for all $\sigma^{-1} < \tilde{\sigma}^{-1}$ the government spending multiplier μ_g for a sequence $g(t) > 0$, $t < T$ and $g(t) = 0$, $t \geq T$ is strictly below 1.*

Proof See appendix A.6. ■

Whereas high fiscal multipliers arose in the new Keynesian model because of higher expected inflation, in this model such an effect can be harmful because it reduces real purchasing power. When the purchasing power effect dominates ($\sigma^{-1} < \tilde{\sigma}^{-1}$) then negative supply shocks are contractionary and fiscal multipliers are below 1. Thus, getting the model to match the empirical evidence on supply shocks also eliminates the main mechanism by which the standard new Keynesian model generates large fiscal multipliers.

6 Conclusion

This paper shows that negative supply shocks are contractionary at the ZLB. This contradicts the prediction of the standard new Keynesian model, which has been the workhorse to study policy options and effects when the ZLB constraint is binding. In this model, negative supply shocks reduce expected real interest rates at the ZLB, which raises consumption today through standard intertemporal substitution. I examine two negative supply shocks — oil supply shocks and an earthquake — and show that these shocks are contractionary at the ZLB despite lowering expected future real interest rates. In short, contrary to Krugman's claim, the ZLB world may not be so "topsy-turvy" after all.

I then show that incorporating a set of borrowing-constrained consumers can resolve the conflict between data and the standard new Keynesian model. These agents' purchasing power is nominally fixed and, given a sufficiently low IES, higher expected inflation induces them to reduce consumption. When this cutback by borrowers dominates the increased consumption by unconstrained lenders, then the model can produce contractions following negative supply shocks at the ZLB. However, because the inflation expectations channel is now contractionary, fiscal multipliers also fall below one. Thus, by matching the data on supply shocks, the model also eliminates the main channel by which the standard new Keynesian model generates large fiscal multipliers. This suggests that policy makers should be cautious in expecting large positive outcomes from the inflation expectations channel.

References

Aruoba, S Boragan and Frank Schorfheide, "Macroeconomic Dynamics Near the ZLB: A Tale of Two Equilibria," 2013.

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

Bernanke, Ben, "Monetary Policy since the Onset of the Crisis," in "Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming," Vol. 31 2012.

Braun, R Anton, Lena Mareen Körber, and Yuichiro Waki, "Some unpleasant properties of log-linearized solutions when the nominal rate is zero," 2012.

Brueckner, Markus, "The Aggregate Capital-Output Elasticity: Evidence from an Instrumental Variables Approach," 2014.

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

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

Cochrane, John, "Fiscal stimulus, fiscal inflation, or fiscal fallacies?," *University of Chicago Booth School of Business. Manuscript, February*, 2009.

Cochrane, John H, "Determinacy and Identification with Taylor Rules," *Journal of Political Economy*, 2011, *119* (3), 565–615.

_ , "The New-Keynesian Liquidity Trap," 2013.

Cogan, John F, Tobias Cwik, John B Taylor, and Volker Wieland, "New Keynesian versus old Keynesian government spending multipliers," *Journal of Economic dynamics and control*, 2010, *34* (3), 281–295.

Coibion, O. and Y. Gorodnichenko, "Information rigidity and the expectations formation process: A simple framework and new facts," 2011.

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

Drautzburg, Thorsten and Harald Uhlig, "Fiscal stimulus and distortionary taxation," 2011.

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.

Dupor, Bill and Rong Li, "The Expected Inflation Channel of Government Spending in the Postwar US," 2014.

Edelstein, Paul and Lutz Kilian, "How sensitive are consumer expenditures to retail energy prices?," *Journal of Monetary Economics*, 2009, *56* (6), 766–779.

Eggertsson, Gauti, "A comment on Casey Mulligan's test of the Paradox of Toil," 2010.

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.

Farhi, Emmanuel and Ivan Werning, "Dealing with the Trilemma: Optimal Capital Controls with Fixed Exchange Rates," 2012.

— and Iván Werning, "Fiscal Multipliers: Liquidity Traps and Currency Unions," 2012.

Fernández-Villaverde, J., G. Gordon, P.A. Guerrón-Quintana, and J. Rubio-Ramírez, "Nonlinear Adventures at the Zero Lower Bound," 2012.

Gavin, William T, Benjamin D Keen, Alexander Richter, and Nathaniel Throckmorton, "Global dynamics at the zero lower bound," 2013.

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.

Kilian, Lutz, "The economic effects of energy price shocks," *Journal of Economic Literature*, 2008, pp. 871–909.

— and Cheolbeom Park, "The impact of oil price shocks on the us stock market," *International Economic Review*, 2009, *50* (4), 1267–1287.

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.

Mertens, Karel and Morten Ravn, "Fiscal Policy in an Expectations Driven Liquidity Trap," *mimeo*, 2013.

Mulligan, Casey B, "Simple Analytics and Empirics of the Government Spending Multiplier and Other "Keynesian" Paradoxes," 2010.

—, *The redistribution recession: How labor market distortions contracted the economy*, Oxford University Press, 2012.

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

Rotemberg, Julio J, “Monopolistic price adjustment and aggregate output,” *The Review of Economic Studies*, 1982, *49* (4), 517–531.

Taylor, John B, “Monetary Policy Rules Work and Discretion Doesn’t: A Tale of Two Eras,” *Journal of Money, Credit and Banking*, 2012, *44* (6), 1017–1032.

Temin, Peter and Barrie A Wigmore, “The end of one big deflation,” *Explorations in Economic history*, 1990, *27* (4), 483–502.

Werning, I., “Managing a Liquidity Trap: Monetary and Fiscal Policy,” 2012.

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

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

_ , “Methods of Policy Accommodation at the Interest-Rate Lower Bound,” 2012.

Not for Publication

A Proofs

A.1 Proof of Proposition 1 When the ZLB never binds, then the equilibrium conditions for $c(t), \pi(t)$ are

$$
\begin{aligned}
dc(t) &= \sigma^{-1}(\phi_\pi - 1)\pi(t)dt \\
d\pi(t) &= \rho\pi(t)dt - \kappa\left\{\left[\sigma + \frac{\alpha(1-s_g)}{1-\alpha}\right]c(t) - \frac{1}{1-\alpha}a(t) + \frac{\alpha s_g}{1-\alpha}g(t)\right\}dt
\end{aligned}
$$

This is a linear ODE that has a unique bounded solution. This system it is more conveniently expressed in matrix form,

$$
\begin{pmatrix} dc(t) \\ d\pi(t) \end{pmatrix} = M \begin{pmatrix} c(t) \\ \pi(t) \end{pmatrix} dt + Dv(t)dt
$$

where

$$
M \equiv \begin{pmatrix} 0 & \sigma^{-1}(\phi_\pi - 1) \\ -\kappa\left[\sigma + \frac{\alpha(1-s_g)}{1-\alpha}\right] & \rho \end{pmatrix}, \quad D \equiv \begin{pmatrix} 0 & 0 & 0 \\ 0 & \kappa\frac{1}{1-\alpha} & -\kappa\frac{\alpha s_g}{1-\alpha} \end{pmatrix}, \quad v(t) = \begin{pmatrix} r(t) \\ a(t) \\ g(t) \end{pmatrix}
$$

We solve the system using standard eigenvalue-eigenvector decomposition. The eigenvalues of M are,

$$
\lambda_{1,2} = \frac{\rho}{2} \pm \sqrt{\left(\frac{\rho}{2}\right)^2 - \kappa(\phi_\pi - 1)\left[1 + \frac{\alpha(1-s_g)}{\sigma(1-\alpha)}\right]},
$$

and satisfy $real(\lambda_1) > 0$ and $real(\lambda_2) > 0$ because $\phi_\pi > 1$. The associated matrix of eigenvectors is,

$$
Q \equiv \begin{pmatrix} \sigma^{-1}(\phi_\pi - 1) & \sigma^{-1}(\phi_\pi - 1) \\ \lambda_1 & \lambda_2 \end{pmatrix}.
$$

Define $z(t) \equiv Q^{-1}\begin{pmatrix} c(t) \\ \pi(t) \end{pmatrix}$, then

$$
dz(t) = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} z(t)dt + Q^{-1}Dv(t)dt.
$$

Solving forward and imposing boundedness on $z(t)$ yields,

$$
z(t) = -\begin{pmatrix} \int_t^\infty e^{-\lambda_1(s-t)}[Q^{-1}D]_{1\bullet}v(s)ds \\ \int_t^\infty e^{-\lambda_2(s-t)}[Q^{-1}D]_{2\bullet}v(s)ds \end{pmatrix},
$$

where the operator $[X]_{i\bullet}$ extracts the i^{th} row of the matrix X. The solution for consumption and inflation is then,

$$\begin{pmatrix} c(t) \\ \pi(t) \end{pmatrix} = -Q \begin{pmatrix} \int_t^\infty e^{-\lambda_1(s-t)}[Q^{-1}D]_{1\bullet}v(s)ds \\ \int_t^\infty e^{-\lambda_2(s-t)}[Q^{-1}D]_{2\bullet}v(s)ds \end{pmatrix}.$$

The differences in consumption and inflation due to the different paths $a'(t)$ and $a(t)$ are then given by,

$$\begin{pmatrix} c'(t) - c(t) \\ \pi'(t) - \pi(t) \end{pmatrix} = -Q \begin{pmatrix} \int_t^\infty e^{-\lambda_1(s-t)}[Q^{-1}D]_{12}(a'(s) - a(s))ds \\ \int_t^\infty e^{-\lambda_2(s-t)}[Q^{-1}D]_{22}(a'(s) - a(s))ds \end{pmatrix}.$$

Given the definitions of Q, D we find,

$$\begin{pmatrix} c'(t) - c(t) \\ \pi'(t) - \pi(t) \end{pmatrix} = \frac{\kappa}{(\lambda_2 - \lambda_1)(1-\alpha)} \begin{pmatrix} \sigma^{-1}(\phi_\pi - 1)\int_t^\infty [e^{-\lambda_1(s-t)} - e^{-\lambda_2(s-t)}][a'(s) - a(s)]ds \\ \int_t^\infty [\lambda_1 e^{-\lambda_1(s-t)} - \lambda_2 e^{-\lambda_2(s-t)}][a'(s) - a(s)]ds \end{pmatrix}.$$

Without loss of generality let $\lambda_2 - \lambda_1 > 0$. Then $e^{-\lambda_1(s-t)} - e^{-\lambda_2(s-t)} > 0$ for all $s > t$. Since $a'(s) - a(s) < 0$ for all $s < T_a$ and $a'(s) = a(s)$ for all $s \geq T_a$ it follows that,

$$\begin{aligned} c'(t) &< c(t), \quad t < T_a. \\ c'(t) &= c(t), \quad t \geq T_a. \end{aligned}$$

Finally, since $g(t)$ is the same in both equilibria, we have $y'(t) - y(t) = (1 - s_g)[c'(s) - c(s)]$ and

$$\begin{aligned} y'(t) &< y(t), \quad t < T_a. \\ y'(t) &= y(t), \quad t \geq T_a. \end{aligned}$$

as was to be shown.

A.2 Proof of Proposition 2 For $t < T_a$, the equilibrium conditions for $c(t), \pi(t)$ in integral form are

$$\begin{aligned} c(t) &= \sigma^{-1}\int_t^{T_a}[r(s) + \pi(s)]ds + c(T_a) \\ \pi(t) &= \kappa \int_t^{T_a} e^{-\rho(s-t)}\left\{\left[\sigma + \frac{\alpha(1-s_g)}{1-\alpha}\right]c(s) - \frac{1}{1-\alpha}a(s) + \frac{\alpha s_g}{1-\alpha}g(s)\right\}ds + e^{-\rho(T_a - t)}\pi(T_a), \end{aligned}$$

and for $c'(t), \pi'(t)$ they are,

$$\begin{aligned} c'(t) &= \sigma^{-1}\int_t^{T_a}[r(s) + \pi'(s)]ds + c'(T_a) \\ \pi'(t) &= \kappa \int_t^{T_a} e^{-\rho(s-t)}\left\{\left[\sigma + \frac{\alpha(1-s_g)}{1-\alpha}\right]c'(s) - \frac{1}{1-\alpha}a'(s) + \frac{\alpha s_g}{1-\alpha}g'(s)\right\}ds + e^{-\rho(T_a - t)}\pi'(T_a). \end{aligned}$$

For $t \geq T_a$ the paths $\{r(t), a(t), g(t)\}$ and $\{r(t), a'(t), g(t)\}$ are identical since $a(t) = a'(t) = 0$. Because the model is forward-looking it follows that the endogenous variables are equal $\{c(t), \pi(t), i(t), y(t)\} = \{c'(t), \pi'(t), i'(t), y'(t)\}, \quad \forall t \geq T_a$. Thus, both equilibria have the same boundary condition $c(T_a) = c'(T_a)$ and $\pi(T_a) = \pi'(T_a)$. Subtracting the conditions of the two equilibria then yields

$$c'(t) - c(t) = \sigma^{-1} \int_t^{T_a} [\pi'(s) - \pi(s)] ds$$
$$\pi'(t) - \pi(t) = \kappa \int_t^{T_a} e^{-\rho(s-t)} \left\{ \left[\sigma + \frac{\alpha(1-s_g)}{1-\alpha} \right] [c'(s) - c(s)] - \frac{1}{1-\alpha} [a'(s) - a(s)] \right\} ds.$$

By substituting for $\pi'(t) - \pi(t)$, we obtain a single condition on the difference in the consumption path across the two equilibria.

$$\begin{aligned} c'(t) - c(t) = & \frac{\kappa \left[\sigma + {\alpha(1-s_g)}/{1-\alpha}\right]}{\sigma\rho} \int_t^{T_a} (1 - e^{-\rho(s-t)}) [c'(s) - c(s)] \\ & - \frac{\kappa}{\sigma\rho(1-\alpha)} \int_t^{T_a} (1 - e^{-\rho(s-t)}) [a'(s) - a(s)] ds \end{aligned}$$

The proof now proceeds as in Werning [2012]. Let $z(t) \equiv c'(t) - c(t)$, $u(t) \equiv -\frac{\kappa}{\sigma\rho(1-\alpha)} \int_t^{T_a} (1 - e^{-\rho(s-t)}) [a'(s) - a(s)] ds > 0$, and $m(t) = \frac{\kappa[\sigma + \alpha(1-s_g)/1-\alpha]}{\sigma\rho} (1 - e^{-\rho(t)})$. Define the operator T as

$$T[z](t) = u(t) + \int_t^{T_a} m(s-t) z(s) ds.$$

The operator T maps the space of continuous functions on $(-\infty, T_a]$ onto itself. An equilibrium z^* is a fixed point to this problem. Werning [2012] shows that there exists such an equilibrium and it is unique. He further shows that starting with any bounded function z_0, the sequence defined by $z_n = T^n[z_0]$ converges to the equilibrium, $z_n \to z^*$.

Start with $z_0(t) = 0$ for all t and define $z_n = T^n[z_0]$. Since $u(t) > 0$ we have $z_1(t) = u(t) < 0$. Because the operator T is monotone, the sequence $\{z_n\}$ is increasing, $0 = z_0 < z_1 \leq ... \leq z_n \leq ....$ It follows that $z^* \geq z_1 > 0$. By the definition of $z(t)$ we have

$$\begin{aligned} c'(t) &> c(t), \quad t < T_a. \\ c'(t) &= c(t), \quad t \geq T_a. \end{aligned}$$

Finally, since $g(t)$ is the same in both equilibria, we have $y'(t) - y(t) = (1 - s_g)[c'(s) - c(s)]$ and

$$\begin{aligned} y'(t) &> y(t), \quad t < T_a. \\ y'(t) &= y(t), \quad t \geq T_a. \end{aligned}$$

as was to be shown.

A.3 Proof of Proposition 3 For $t < T_a$, the equilibrium conditions for $c(t), \pi(t)$ in integral form are

$$c(t) = \sigma^{-1} \int_t^{T_a} [r(s) + \pi(s)]ds - \sigma^{-1} \int_T^{T_a} i(s)ds + c(T_a)$$
$$\pi(t) = \kappa \int_t^{T_a} e^{-\rho(s-t)} \left\{ \left[\sigma + \frac{\alpha(1-s_g)}{1-\alpha}\right] c(s) - \frac{1}{1-\alpha} a(s) + \frac{\alpha s_g}{1-\alpha} g(s) \right\} ds + e^{-\rho(T_a - t)} \pi(T_a),$$

and for $c'(t), \pi'(t)$ they are,

$$c'(t) = \sigma^{-1} \int_t^{T_a} [r(s) + \pi'(s)]ds - \sigma^{-1} \int_{T'}^{T_a} i'(s)ds + c'(T_a)$$
$$\pi'(t) = \kappa \int_t^{T_a} e^{-\rho(s-t)} \left\{ \left[\sigma + \frac{\alpha(1-s_g)}{1-\alpha}\right] c'(s) - \frac{1}{1-\alpha} a'(s) + \frac{\alpha s_g}{1-\alpha} g(s) \right\} ds + e^{-\rho(T_a - t)} \pi'(T_a).$$

Unlike before, we now have to allow for positive nominal interest rates between the ZLB exit time (T or T') and the end of the productivity shock T_a.

For $t \geq T_a$ the paths $\{r(t), a(t), g(t)\}$ and $\{r(t), a'(t), g(t)\}$ are identical since $a(t) = a'(t) = 0$. Because the model is forward-looking it follows that the endogenous variables are equal $\{c(t), \pi(t), i(t), y(t)\} = \{c'(t), \pi'(t), i'(t), y'(t)\}, \ \ \forall t \geq T_a$. Thus, both equilibria have the same boundary condition $c(T_a) = c'(T_a)$ and $\pi(T_a) = \pi'(T_a)$. Subtracting the conditions of the two equilibria then yields

$$c'(t) - c(t) = \sigma^{-1} \int_t^{T_a} [\pi'(s) - \pi(s)]ds - \sigma^{-1} \int_{T'}^{T} i'(s)ds - \sigma^{-1} \int_T^{T_a} [i'(s) - i(s)]ds$$
$$\pi'(t) - \pi(t) = \kappa \int_t^{T_a} e^{-\rho(s-t)} \left\{ \left[\sigma + \frac{\alpha(1-s_g)}{1-\alpha}\right] [c'(s) - c(s)] - \frac{1}{1-\alpha} [a'(s) - a(s)] \right\} ds.$$

Relative to Proposition 2 there are two new terms. The first term, $\int_{T'}^{T} i'(s)ds$, captures the increase in nominal interest rates when the ZLB duration shortens from T to T'. The second term, $\int_T^{T_a} [i'(s) - i(s)]ds$, captures the increase in nominal interest rates that occurs when the decline in productivity persists beyond the ZLB, $T_a > T$.

We can solve for the second term using the solution to Proposition 1 since the ZLB does not bind after T,

$$\int_T^{T_a} [i'(s) - i(s)]ds = -\frac{\kappa \phi_\pi}{(\lambda_2 - \lambda_1)(1-\alpha)} \int_T^{T_a} [e^{-\lambda_1 (s-t)} - e^{-\lambda_2 (s-t)}][a'(s) - a(s)]ds \geq 0.$$

Note that the inequality is strict if $T < T_a$. Thus, if productivity is lower for $t > T$, then inflation will be higher and the central bank sets higher nominal interest rates.

The first term must also be weakly positive due to the ZLB constraint,

$$\int_{T'}^{T} i'(s)ds \geq 0.$$

Intuitively, if the economy exits earlier then nominal interest rates must be higher. The inequality is also strict when $T' < T$.

Imposing the restriction on interest rates, we obtain

$$0 \geq \int_{t_0}^{T_a} [i'(s) - i(s)]ds = \int_{T'}^{T} i'(s)ds + \int_{T}^{T_a} [i'(s) - i(s)]ds \geq 0,$$

since $i(t) = 0$ for $t < T$ and $i'(t) < 0$ for $t < T'$. This restriction implies $T' = T$ and $T < T_a$ for otherwise expected future nominal interest rates are strictly positive. Thus the assumption for Proposition 2 are satisfied and the results follow.

A.4 Proof of Proposition 4 **Only if:** By definition of output, $y(t) = (1-s_g)c(t) + s_g g(t)$

$$\mu_g(t) = 1 + \frac{(1-s_g)\int_t^\infty [c'(s) - c(s)]ds}{s_g \int_t^\infty [g'(s) - g(s)]ds}.$$

Therefore $\mu_g(t) > 1$ implies $\frac{\int_t^\infty [c'(s)-c(s)]ds}{\int_t^\infty [g'(s)-g(s)]ds} > 0$.

Rewrite (2) as,

$$d\pi(t) = \rho\pi(t)dt - \kappa \left\{ \left[\sigma + \frac{\alpha(1-s_g)}{1-\alpha} \right] c(t) + v(t) \right\} dt$$

where $v(t) = -\frac{1}{1-\alpha}a(t) + \frac{\alpha s_g}{1-\alpha}g(t)$. Given the sequences $\{r(t), g'(t), a(t)\}$ and $\{r(t), g(t), \tilde{a}(t)\}$, we note that $\{v'(t)\} = \{\tilde{v}'(t)\}$ given the definition of $\tilde{a}(t)$. Since $a(t)$ and $g(t)$ do not enter elsewhere in equations (1)-(3), the solutions for consumption, inflation and nominal interest rates are identical, $\{c'(t), \pi'(t), i'(t)\} = \{\tilde{c}(t), \tilde{\pi}(t), \tilde{i}(t)\}$. It then follows that

$$\begin{aligned}
& \frac{\int_t^\infty [c'(s) - c(s)]ds}{\int_t^\infty [g'(s) - g(s)]ds} > 0 \\
\iff & \frac{\int_t^\infty [\tilde{c}(s) - c(s)]ds}{\int_t^\infty [g'(s) - g(s)]ds} > 0 \\
\iff & \frac{\int_t^\infty [\tilde{c}(s) - c(s)]ds}{-\alpha s_g \int_t^\infty [g'(s) - g(s)]ds} < 0 \\
\iff & \frac{\int_t^\infty [\tilde{c}(s) - c(s)]ds}{\int_t^\infty [\tilde{a}(s) - a(s)]ds} < 0 \\
\iff & \frac{(1-s_g)\int_t^\infty [\tilde{c}(s) - c(s)]ds}{\int_t^\infty [\tilde{a}(s) - a(s)]ds} < 0 \\
\iff & \frac{\int_t^\infty [\tilde{y}(s) - y(s)]ds}{\int_t^\infty [\tilde{a}(s) - a(s)]ds} < 0
\end{aligned}$$

If: Straightforward reversal of previous steps.

A.5 Proof of Proposition 6 Over $[t_0, T]$ the economy is described by

$$dc^b(t) == \sigma^{-1}[\pi(t) + r(t)]dt$$
$$dc^l(t) == \sigma^{-1}[\pi(t) + r(t)]dt$$
$$d\pi(t) = \rho\pi(t)dt - \kappa\left\{\left[\sigma + \frac{\alpha(1-s_g)(1-\chi)}{1-\alpha}\right]c^l(s) + \frac{\alpha(1-s_g)\chi}{1-\alpha}c^b(s) - \frac{1}{1-\alpha}a(t) + \frac{\alpha s_g}{1-\alpha}g(t)\right\}dt$$

with boundary conditions at T given by

$$c^b(T) = \ln\psi - \sigma^{-1}\int_{t_0}^{T}[\pi(s) + r(s)]ds - (1-\sigma^{-1})\int_{t_0}^{T}\frac{(T-s)}{(T-t_0)}\pi(s)ds + \sigma^{-1}\int_{t_0}^{T}\frac{(T-s)}{(T-t_0)}r(s)ds$$
$$c^l(T) = 0$$
$$\pi(T) = 0$$

This is a linear ODE with three boundary conditions. Because the first boundary condition is endogenous a solution is not necessarily guaranteed. I characterize conditions under which an equilibrium exists below.

This system it is more conveniently expressed in matrix form,

$$\begin{pmatrix} dc^b(t) \\ dc^l(t) \\ d\pi(t) \end{pmatrix} = M \begin{pmatrix} c^b(t) \\ c^l(t) \\ \pi(t) \end{pmatrix} dt + Dv(t)dt$$

where

$$M \equiv \begin{pmatrix} 0 & 0 & -\sigma^{-1} \\ 0 & 0 & -\sigma^{-1} \\ -\kappa^b & -\kappa^l & \rho \end{pmatrix}, \qquad D \equiv \begin{pmatrix} -\sigma^{-1} & 0 & 0 \\ -\sigma^{-1} & 0 & 0 \\ 0 & \kappa\frac{1}{1-\alpha} & -\kappa\frac{\alpha s_g}{1-\alpha} \end{pmatrix}, \qquad v(t) = \begin{pmatrix} r(t) \\ a(t) \\ g(t) \end{pmatrix}$$
$$\kappa^b = \kappa\frac{\alpha(1-s_g)\chi}{1-\alpha}, \qquad \kappa^l = \kappa\left[\sigma + \frac{\alpha(1-s_g)(1-\chi)}{1-\alpha}\right]$$

We solve the system using standard eigenvalue-eigenvector decomposition. The eigenvalues of M are,

$$\lambda_1 = 0 \qquad \lambda_{2,3} = \frac{\rho}{2} \pm \sqrt{\left(\frac{\rho}{2}\right)^2 + \sigma^{-1}(k^l + b^b)},$$

and satisfy $\lambda_2 < 0$ and $\lambda_3 > 0$. The associated matrix of eigenvectors is,

$$Q \equiv \begin{pmatrix} k^l & -\sigma^{-1} & -\sigma^{-1} \\ -k^b & -\sigma^{-1} & -\sigma^{-1} \\ 0 & \lambda_2 & \lambda_3 \end{pmatrix}.$$

Its inverse is given by,

$$Q^{-1} = \frac{1}{-(\lambda_3 - \lambda_2)\sigma^{-1}(k^l + b^b)} \begin{pmatrix} -\sigma^{-1}(\lambda_3 - \lambda_2) & \sigma^{-1}(\lambda_3 - \lambda_2) & 0 \\ k^b \lambda_3 & k^l \lambda_3 & \sigma^{-1}(k^l + b^b) \\ -k^b \lambda_2 & -k^l \lambda_2 & -\sigma^{-1}(k^l + b^b) \end{pmatrix}.$$

Define $z(t) \equiv Q^{-1} \begin{pmatrix} c^b(t) \\ c^l(t) \\ \pi(t) \end{pmatrix}$, then

$$dz(t) = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} z(t)dt + Q^{-1}Dv(t)dt.$$

Solving forward up to time T yields,

$$z(t) = \begin{pmatrix} e^{-\lambda_1(T-t)} z_1(T) \\ e^{-\lambda_2(T-t)} z_2(T) \\ e^{-\lambda_3(T-t)} z_3(T) \end{pmatrix} - \begin{pmatrix} \int_t^T e^{-\lambda_1(s-t)} [Q^{-1}D]_{1\bullet} v(s)ds \\ \int_t^T e^{-\lambda_2(s-t)} [Q^{-1}D]_{2\bullet} v(s)ds \\ \int_t^T e^{-\lambda_3(s-t)} [Q^{-1}D]_{3\bullet} v(s)ds \end{pmatrix},$$

where the operator $[X]_{i\bullet}$ extracts the i^{th} row of the matrix X. The boundary conditions in the z-plane are

$$z(T) = \frac{1}{-(\lambda_3 - \lambda_2)\sigma^{-1}(k^l + b^b)} \begin{pmatrix} -\sigma^{-1}(\lambda_3 - \lambda_2) \\ k^b \lambda_3 \\ -k^b \lambda_2 \end{pmatrix} c^b(T).$$

Rotating back to the consumption-inflation plane we obtain,

$$\begin{aligned} \begin{pmatrix} c^b(t) \\ c^l(t) \\ \pi(t) \end{pmatrix} &= \frac{1}{-(\lambda_3 - \lambda_2)\sigma^{-1}(k^l + b^b)} \begin{pmatrix} -k^l \sigma^{-1}(\lambda_3 - \lambda_2) - \sigma^{-1} k^b [\lambda_3 e^{-\lambda_2(T-t)} - \lambda_2 e^{-\lambda_3(T-t)}] \\ k^b \sigma^{-1}(\lambda_3 - \lambda_2) - \sigma^{-1} k^b [\lambda_3 e^{-\lambda_2(T-t)} - \lambda_2 e^{-\lambda_3(T-t)}] \\ k^b \lambda_2 \lambda_3 [e^{-\lambda_2(T-t)} - e^{-\lambda_3(T-t)}] \end{pmatrix} c^b(T) \\ &+ \frac{1}{\lambda_3 - \lambda_2} \int_t^T \begin{pmatrix} \sigma^{-1}[\lambda_3 e^{-\lambda_3(s-t)} - \lambda_2 e^{-\lambda_2(s-t)}] \\ \sigma^{-1}[\lambda_3 e^{-\lambda_3(s-t)} - \lambda_2 e^{-\lambda_2(s-t)}] \\ \lambda_2 \lambda_3 [e^{-\lambda_3(s-t)} - e^{-\lambda_2(s-t)}] \end{pmatrix} r(s)ds \\ &+ \frac{1}{\lambda_3 - \lambda_2} \int_t^T \begin{pmatrix} \sigma^{-1}[e^{-\lambda_3(s-t)} - e^{-\lambda_2(s-t)}] \\ \sigma^{-1}[e^{-\lambda_3(s-t)} - e^{-\lambda_2(s-t)}] \\ \lambda_2 e^{-\lambda_2(s-t)} - \lambda_3 e^{-\lambda_3(s-t)} \end{pmatrix} \left[\frac{1}{1-\alpha} a(s) - \frac{\alpha s_g}{1-\alpha} g(s) \right] ds, \end{aligned}$$

It remains to determine the endogenous boundary condition. This a fixed point of the solution to inflation and borrowers' terminal consumption,

$$\begin{aligned}
c^b(T) &= \ln\psi - \int_{t_0}^{T} \theta(s)\pi(s)ds + \sigma^{-1}\int_{t_0}^{T} \frac{(T-s)}{(T-t_0)}r(s)ds \\
\pi(t) &= \varphi\mu(t)c^b(T) + m(t)
\end{aligned}$$

where

$$\begin{aligned}
\theta(s) &= \sigma^{-1} + (1-\sigma^{-1})\frac{(T-s)}{(T-t_0)} \geq 0 \\
\varphi &= -\frac{k^b\lambda_2\lambda_3}{(\lambda_3-\lambda_2)\sigma^{-1}(k^l+b^b)} \geq 0 \\
\mu(t) &= e^{-\lambda_2(T-t)} - e^{-\lambda_3(T-t)} \geq 0 \\
m(t) &= \frac{1}{\lambda_3-\lambda_2}\int_t^T \lambda_2\lambda_3[e^{-\lambda_3(s-t)} - e^{-\lambda_2(s-t)}]r(s)ds \\
&\quad + \frac{1}{\lambda_3-\lambda_2}\int_t^T [\lambda_2 e^{-\lambda_2(s-t)} - \lambda_3 e^{-\lambda_3(s-t)}]\left[\frac{1}{1-\alpha}a(s) - \frac{\alpha s_g}{1-\alpha}g(s)\right] ds < 0.
\end{aligned}$$

Note that $m(s)$ is increasing in $r(s)$, decreasing in $a(s)$, and increasing in $g(s)$.

The solution to the fixed point problem is,

$$\begin{aligned}
\int_{t_0}^{T} \theta(s)\pi(s)ds &= \frac{\varphi\int_{t_0}^{T}\mu(s)\theta(s)ds}{1+\varphi\int_{t_0}^{T}\mu(s)\theta(s)ds}\left[\ln\psi + \sigma^{-1}\int_{t_0}^{T}\frac{(T-s)}{(T-t_0)}r(s)ds\right] \\
&\quad + \frac{1}{1+\varphi\int_{t_0}^{T}\mu(s)\theta(s)ds}\int_{t_0}^{T}\theta(s)m(s)ds
\end{aligned}$$

which implies,

$$\begin{aligned}
\pi(t) &= \frac{\varphi\mu(t)}{1+\varphi\int_{t_0}^{T}\mu(s)\theta(s)ds}\left[\ln\psi + \sigma^{-1}\int_{t_0}^{T}\frac{(T-s)}{(T-t_0)}r(s)ds - \int_{t_0}^{T}\theta(s)m(s)ds\right] + m(t) \\
c^b(T) &= \frac{1}{1+\varphi\int_{t_0}^{T}\mu(s)\theta(s)ds}\left[\ln\psi + \sigma^{-1}\int_{t_0}^{T}\frac{(T-s)}{(T-t_0)}r(s)ds - \int_{t_0}^{T}\theta(s)m(s)ds\right]
\end{aligned}$$

For the economy to remain at the ZLB we require that $\pi(t) < 0$ for $t \in [t_0, T]$. A sufficiently low value for $\psi < 1$ guarantees this outcome and therefore the existence of an equilibrium.

Next we consider the case where $\sigma^{-1} \to 0$. It is straightforward to show that,

$$
\begin{aligned}
c^l(t) &= 0 \\
c^b(t) = c^b(T) &= \frac{1}{1+\varphi \int_{t_0}^T \mu(s)\theta(s)ds}\left[\ln\psi - \int_{t_0}^T \theta(s)m(s)ds\right] \\
y(t) &= (1-s_g)g(s) + \chi c^b(t)
\end{aligned}
$$

where $\varphi = -\frac{k^b\lambda_2\lambda_3}{(\lambda_3-\lambda_2)\kappa} \geq 0$.

Given a negative productivity shock, the difference in consumption is given by,

$$
c^{b'}(t) - c^b(t) = c^b(T) = -\frac{1}{1+\varphi \int_{t_0}^T \mu(s)\theta(s)ds}\left[\int_{t_0}^T \theta(s)[m'(s)-m(s)]ds\right]
$$

Since $\theta(s) > 0$ and $m(s)$ is decreasing in $a(s)$, it follows that a decline in $a(s)$ causes a decline in the borrower's consumption and thus output. Because the decline in output is strict, by continuity, we know there exists a $\tilde{\sigma}^{-1} > 0$ such that for all $\sigma^{-1} < \tilde{\sigma}^{-1}$ the negative productivity shock is contractionary. This proves Proposition 6.

A.6 Proof of Proposition 7 Follow the same steps as in the proof for Proposition 6. We then note that government spending enters analogously to productivity but with an opposite sign.

B Proposition 2 in a non-linear model

B.1 Non-linear Phillips curve I consider a new Keynesian model with Rotemberg [1982] price-stickiness. Firms produce with linear technology, $Y_i(t) = A(t)L_i(t)$, where $A(t)$ is productivity and $L_i(t)$ is labor input. Their demand function is standard CES $Y_i(t) = \left(\frac{P_i(t)}{P_t}\right)^{-\theta} Y(t)$, where where $P_i(t)$ is firm i's price, P_t is the aggregate price, and $Y(t)$ is aggregate output. The profit maximization condition is then given by,

$$\max_{\{d(t)\}} \int_t^{\infty} e^{-\rho(s-t)} C(s)^{-\sigma} \left[(1+\tau)\frac{P_i(s)}{P(s)}Y_i(s) - \frac{W(s)}{A(s)}Y_i(s) - \frac{\gamma}{2}\left(d_i(t)\right)^2 Y(s)\right]$$
$$s.t. \qquad d\ln P_i(s) = d_i(s)dt \;\; \forall s \geq t,$$

where $W(s)$ is the real wage and $(1+\tau)(\theta-1) = \theta$ undoes the steady-state monopoly distortion. The last term incorporates costs of price adjustments. As in Rotemberg [1982], these are quadratic in the change of the log price. I multiply the costs of price adjustment by $Y(t)$ to ensure that they do not become negligible as the economy grows. The price adjustment costs are rebated lump-sum to households. Thus, while they affect the incentives of firms, they do not constitute lost resources.

Denote the co-state variable by $q(t)$. The FOC for the present-value Hamiltonian are:

$$0 = -\gamma d_i(t) C(s)^{-\sigma} Y(t) + q_i(t),$$
$$\frac{dq(t)}{dt} - \rho q(t) = \left[-(1+\tau)(1-\theta)\left(\frac{P_i(t)}{P(t)}\right)^{1-\theta} - \theta\frac{W(t)}{A(t)}\left(\frac{P_i(t)}{P(t)}\right)^{-\theta}\right] Y(t)C(t)^{-\sigma}.$$

To simplify the analysis I specialize to the case where there is no government so $C(t) = Y(t)$ and I set $\sigma = 1$.[17] In the symmetric equilibrium it follows that $\pi(t) = d(t) = \gamma^{-1}q(t)$ and,

$$d\pi(t) = \rho\pi(t)dt - \frac{\theta}{\gamma}\left[\frac{W(t)}{A(t)} - 1\right]dt$$

The household labor supply decision pins down the real wage,

$$W(t) = \chi C(t).$$

This implies the Phillips curve,

$$d\pi(t) = \rho\pi(t)dt - \frac{\theta}{\gamma}\left[\chi\frac{C(t)}{A(t)} - 1\right]dt \tag{13}$$

This Phillips curve is exact. A linear approximation to this function yields equation (2) with $\kappa = \frac{\theta}{\gamma}$, $\sigma = 1$, $\alpha = 0$, and $s_g = 0$.

B.2 Results The Euler equation (1) is also exact in continuous time. We therefore solve for a bounded equilibrium that satisfies equations (13), (1) and (3) given an exogenous path

[17] Because price adjustment costs are an intermediate input they do not count towards GDP $Y(t)$.

$\{r(t), A(t)\}$.

We consider the same path for the natural rate of interest $r(t)$ as before,

$$\begin{aligned} r(t) &< 0 \qquad t < T, \\ r(t) &\geq 0 \qquad t \geq T. \end{aligned}$$

We then compare the equilibrium associated with $A(t) = \bar{A}$ for all t with the equilibrium associated with $A'(t) < \bar{A}$ for $t < T_a$ and $A'(t) = \bar{A}$ for $t \geq T_a$.

Proposition 8 *If $T_a < T' \leq T$, then the negative supply shock is expansionary in the non-linear model*

$$\begin{aligned} Y'(t) > Y(t) \;\; &\textit{and} \;\; C'(t) > C(t) \qquad t < T_a, \\ Y'(t) = Y(t) \;\; &\textit{and} \;\; C'(t) = C(t) \qquad t \geq T_a. \end{aligned}$$

B.3 Proof The equilibrium conditions for $C(t)$ and $\pi(t)$ in integral form are,

$$\begin{aligned} C(t) &= \left[e^{\int_t^T [-i(s)+r(s)+\pi(s)]ds}\right] C(T) \\ \pi(t) &= \frac{\theta}{\gamma} \int_t^T e^{-\rho(s-t)} \left[\frac{C(s)}{\bar{C}} \frac{\bar{A}}{A(s)} - 1\right] ds + e^{-\rho(T-t)} \pi(T), \end{aligned}$$

After T the central bank insulates the economy, so $i(t) = i'(t) = r(t) \geq 0$, $\pi(T) = \pi'(T) = 0$ and $C(T) = C'(T) = \bar{C} = \frac{\bar{A}}{\chi}$. Before T the ZLB binds and $i(t) = 0$. Substituting for $\pi(t)$ then yields,

$$\begin{aligned} \frac{C(t)}{\bar{C}} &= \left[e^{\int_t^T \left\{r(s) + \frac{\theta}{\gamma} \int_s^T e^{-\rho(z-s)} \left[\frac{C(z)}{\bar{C}} \frac{\bar{A}}{A(z)} - 1\right] dz\right\} ds}\right] \\ &= \left[e^{\int_t^T r(s)ds + \frac{\theta}{\gamma\rho} \int_t^T (1-e^{-\rho(s-t)}) \left[\frac{C(s)}{\bar{C}} \frac{\bar{A}}{A(s)} - 1\right] ds}\right]. \end{aligned}$$

Define $Z(t) = \frac{C(z)}{\bar{C}}$ and the operator T,

$$T[Z](t) = \left[e^{\int_t^T r(s)ds + \frac{\theta}{\gamma\rho} \int_t^T (1-e^{-\rho(s-t)}) \left[Z(s) \frac{\bar{A}}{A(s)} - 1\right] ds}\right]$$

Because the operator is non-linear in Z, the results from Werning [2012] do not immediately follow.

Note that $\pi(t) \leq -r(t)$ for $t < T$ is a necessary condition for an equilibrium. Otherwise $\phi_\pi \pi(t) > \pi(t) > -r(t) > 0$ implies that the ZLB would not bind, $i(t) = r(t) + \phi_\pi \pi(t) > 0$. It follows that $e^{\int_t^T [r(s)+\pi(s)]ds} \leq 1$ for all t. This also restricts the process $\{A(t)\}$,

$$\frac{\theta}{\gamma\rho} \int_t^T (1 - e^{-\rho(s-t)}) \left[\frac{\bar{A}}{A(s)} - 1\right] ds < -\int_t^T r(s)ds, \qquad \forall t < T.$$

If this condition was violated for some t, then the economy would exit the ZLB at $t < T_a$. Thus, this condition must be satisfied under the assumption $T_a < T' < T$ in the proposition.

In any equilibrium we must therefore have $0 \leq Z(t) \leq 1$. Fix an interval $[\hat{t}, T]$. Let $F \subseteq C([\hat{t}, T])$ be the space of continuous functions that map the $[\hat{t}, T]$ interval onto the unit interval, $f : [\hat{t}, T] \to [0, 1]$ for all $f \in F$. Then the operator T maps F onto itself, $T : F \to F$.

We can show that T is a contraction for some $\hat{t}$.

$$
\begin{aligned}
|T[V](t) - T[W](t)| &= \left| e^{\int_t^T r(s)ds + \frac{\theta}{\gamma\rho}\int_t^T (1-e^{-\rho(s-t)})\left[V(s)\frac{\bar{A}}{A(s)}-1\right]ds} - e^{\int_t^T r(s)ds + \frac{\theta}{\gamma\rho}\int_t^T (1-e^{-\rho(s-t)})\left[W(s)\frac{\bar{A}}{A(s)}-1\right]ds} \right| \\
&\leq \left| \int_t^T r(s)ds + \frac{\theta}{\gamma\rho}\int_t^T (1 - e^{-\rho(s-t)}) \left[V(s)\frac{\bar{A}}{A(s)} - 1\right] ds \right. \\
&\quad \left. - \int_t^T r(s)ds - \frac{\theta}{\gamma\rho}\int_t^T (1 - e^{-\rho(s-t)}) \left[W(s)\frac{\bar{A}}{A(s)} - 1\right] ds \right| \\
&\leq \frac{\theta}{\gamma\rho}\int_t^T (1 - e^{-\rho(s-t)}) \left| [V(s) - W(s)]\frac{\bar{A}}{A(s)} \right| ds \\
&\leq \frac{\theta}{\gamma\rho} \left| \frac{\bar{A}}{A(s)} \right| |T - t||V(s) - W(s)|ds
\end{aligned}
$$

Pick $\hat{t}$ such that $\frac{\theta}{\gamma\rho}\left|\frac{\bar{A}}{A(s)}\right| |T - \hat{t}| < 1$. Then the operator T is a contraction under the sup-norm. It follows that a unique equilibrium Z^* exists over $[\hat{t}, T]$ and it is the limit of the sequence, $Z_n = T^n Z_0$, where $Z_0 \in F$. Using $Z(\hat{t})$ as a unique boundary condition, we can repeat the above steps for a $\tilde{t}$, and find a contraction over $[\tilde{t}, \hat{t}]$. It follows that a unique equilibrium exists over $[t, T]$ for arbitrary t.

The analogous results to Proposition 2 then follow from the fact that the operator is monotonic. Let C^* be the equilibrium given the sequence $\{r(t), A(t)\}$.

C A model with oil consumption and production

C.1 Model Agents maximize the stream of utility,

$$\int_0^\infty e^{-\rho t}\left[\frac{C(t)^{1-\sigma}}{1-\sigma} - \chi N(t)\right],$$

where $C(t)$ is domestic consumption and $N(t)$ is labor. The inverse of the intertemporal elasticity of substitution is σ, the inverse of the Frisch elasticity is ν, and χ is a parameter that determines steady-state labor supply.

Domestic consumption $C(t)$ is an aggregate of a produced good and consumed oil,

$$C(t) = \left[(1-\gamma)^{\frac{1}{\zeta}} N_i(t)^{\frac{\zeta-1}{\zeta}} + \gamma^{\frac{1}{\zeta}} O^c(t)^{\frac{\zeta-1}{\zeta}}\right]^{\frac{\zeta}{\zeta-1}}.$$

The produced good is a combination of individual varieties,

$$C^y(t) = \left(\int_0^1 C_i(t)^{\frac{\varepsilon-1}{\varepsilon}} di\right)^{\frac{\varepsilon}{\varepsilon-1}}.$$

Asset holdings of the risk-free bond $D(t)$ evolve according to

$$D(t) = i(t)D(t) - P(t)C(t) + W(t)N(t) + \Pi(t),$$

where $i(t)$ is the nominal interest rate, $W(t)$ the wage rate, and $\Pi(t)$ are profits from firms.

Firms produce output $Y_i(t)$ of variety i according to a CES technology,

$$Y_i(t) = A(t)\left[(1-\xi)^{\frac{1}{\psi}} N_i(t)^{\frac{\psi-1}{\psi}} + \xi^{\frac{1}{\psi}} O_i^y(t)^{\frac{\psi-1}{\psi}}\right]^{\frac{\psi}{\psi-1}}$$

where $O_i^y(t)$ is oil input, ξ the share of oil in production, and ψ its elasticity of substitution with labor input. Firms face standard Calvo pricing frictions.

The real price of oil $\frac{P^o(t)}{P(t)}$ adjusts such that the oil market clears,

$$\int_0^1 O_i^y(t)di + O^c(t) = O(t)$$

The log-linearized equilibrium conditions for this model are,

$$\begin{aligned}
dc(t) &= \sigma^{-1}[i(t) - r(t) - \pi(t)]dt \\
\pi(t) &= \pi^y(t) + \gamma dp^o(t) \\
d\pi^y(t) &= \rho\pi^y(t)dt - \kappa\, mc(t)dt \\
i(t) &= \max\{r(t) + \phi_\pi \pi(t), 0\}, \quad \phi_\pi > 1 \\
mc(t) &= M_c c(t) - M_o o(t) - M_a a(t) \\
p^o(t) &= Q_c c(t) - Q_o o(t) - Q_a a(t) \\
y(t) &= c(t) + \gamma\zeta p^o(t)
\end{aligned}$$

where p^o is the real price of oil in terms of the produced good and the constants M_x, Q_x are positive,

$$M_c = \frac{\sigma(1-\xi)(1-\gamma)[\psi\xi(1-\gamma)+\zeta\gamma(1-\xi)] + [\gamma+(1-\gamma)\xi](1-\gamma)\xi}{(1-\gamma)(1-\xi)(\psi\xi(1-\gamma)+\zeta\gamma)}$$
$$Q_c = \frac{\gamma+(1-\gamma)\xi+\sigma\psi\xi(1-\gamma)(1-\xi)}{(1-\gamma)(1-\xi)(\psi\xi(1-\gamma)+\zeta\gamma)}$$
$$M_o = \frac{[\gamma+(1-\gamma)\xi]\xi(1-\gamma)}{(1-\gamma)(1-\xi)(\psi\xi(1-\gamma)+\zeta\gamma)}$$
$$Q_o = \frac{\gamma+(1-\gamma)\xi}{(1-\gamma)(1-\xi)(\psi\xi(1-\gamma)+\zeta\gamma)}$$
$$M_a = \frac{(1-\gamma)\gamma\zeta(1-\xi)+(1-\gamma)^2\xi[\psi+\xi(1-\psi)]}{(1-\gamma)(1-\xi)(\psi\xi(1-\gamma)+\zeta\gamma)}$$
$$Q_a = \frac{(1-\gamma)\xi}{(1-\gamma)(1-\xi)(\psi\xi(1-\gamma)+\zeta\gamma)}$$

C.2 Proof of Proposition 5 From the conditions in Proposition 3 it follows that $T_o < T$. For $t < T_o$, the equilibrium conditions for $c(t), \pi(t)$ in integral form are

$$c(t) = \sigma^{-1}\int_t^{T_o}[r(s)+\pi^y(s)]ds + \sigma^{-1}\gamma[p^o(T)-p^o(t)] + c(T_o)$$
$$\pi^y(t) = \kappa\int_t^{T_o} e^{-\rho(s-t)}[Q_c c(s) - Q_o o(s) - Q_a a(s)]ds + e^{-\rho(T_o-t)}\pi(T_o),$$

and for $c'(t), \pi'(t)$ they are,

$$c'(t) = \sigma^{-1}\int_t^{T_o}[r(s)+\pi^{y'}(s)]ds + \sigma^{-1}\gamma[p^{o'}(T)-p^{o'}(t)] + c'(T_o)$$
$$\pi^{y'}(t) = \kappa\int_t^{T_o} e^{-\rho(s-t)}[Q_c c'(s) - Q_o o'(s) - Q_{a'} a(s)]ds + e^{-\rho(T_o-t)}\pi'(T_o),$$

For $t \geq T_o$ the paths $\{r(t), a(t), o(t)\}$ and $\{r(t), a(t), o'(t)\}$ are identical since $o(t) = o'(t) = 0$. Because the model is forward-looking it follows that the endogenous variables are equal $\{c(t), \pi(t), i(t), y(t), p^o(t)\} = \{c'(t), \pi'(t), i'(t), y'(t), p^o(t)\}, \quad \forall t \geq T_o$. Thus, both equilibria have the same boundary condition $c(T_o) = c'(T_o)$ and $\pi(T_o) = \pi'(T_o)$. Subtracting the conditions of the two equilibria then yields

$$c'(t) - c(t) = \sigma^{-1}\int_t^{T_a}[\pi^{y'}(s)-\pi^y(s)]ds - \sigma^{-1}\gamma[p^{o'}(t)-p^o(t)]$$
$$\pi^{y'}(t) - \pi^y(t) = \kappa\int_t^{T_a} e^{-\rho(s-t)}\{Q_c[c'(s)-c(s)] - Q_o[o'(s)-o(s)]\}ds.$$

Substituting for real oil prices yields,

$$c'(t) - c(t) = \sigma^{-1} \int_t^{T_a} [\pi^{y'}(s) - \pi^y(s)]ds - \sigma^{-1}\gamma Q_c[c'(t) - c(t)] + \sigma^{-1}\gamma Q_o[o'(s) - o(s)]$$
$$\pi^{y'}(t) - \pi^y(t) = \kappa \int_t^{T_a} e^{-\rho(s-t)}\{Q_c[c'(s) - c(s)] - Q_o[o'(s) - o(s)]\}ds.$$

By substituting for $\pi'(t) - \pi(t)$, we obtain a single condition on the difference in the consumption path across the two equilibria.

$$\begin{aligned} c'(t) - c(t) =& \frac{\kappa M_c}{\rho[\sigma + \gamma Q_c]} \int_t^{T_a} (1 - e^{-\rho(s-t)})[c'(s) - c(s)] - \frac{\kappa M_o}{\rho[\sigma + \gamma Q_c]} \int_t^{T_a} (1 - e^{-\rho(s-t)})[o'(s) - o(s)] \\ &+ \frac{\gamma Q_o}{\rho[\sigma + \gamma Q_c]}[o'(s) - o(s)] \end{aligned}$$

The proof now proceeds as in Werning [2012] or, identically, section A.2. Define

$$v(t) = -\frac{\kappa M_o}{\rho[\sigma + \gamma Q_c]} \int_t^{T_a} (1 - e^{-\rho(s-t)})[o'(s) - o(s)] + \frac{\gamma Q_o}{\rho[\sigma + \gamma Q_c]}[o'(s) - o(s)],$$

then the oil supply shock is expansionary if and only if $v(t) > 0$. The expected inflation response of the shock is,

$$\begin{aligned} \pi'(t) - \pi(t) &= \int_t^{T_a} [\pi'(s) - \pi(s)]ds - \gamma[p^{o'}(t) - p^o(t)] \\ &= v(t) + \frac{\kappa M_c}{\rho[\sigma + \gamma Q_c]} \int_t^{T_a} (1 - e^{-\rho(s-t)})[c'(s) - c(s)]. \end{aligned}$$

Thus, holding consumption fixed,

$$[\pi'(t) - \pi(t)]|_{\{c(s)=0, s\geq t\}} = v(t).$$

Hence, the oil supply shock is expansionary if and only if

$$[\pi'(t) - \pi(t)]|_{\{c(s)=0, s\geq t\}} > 0.$$

D Correspondence of productivity shocks and oil supply shocks in open economy models

I consider the case of a small open economy that imports oil for the purpose of production. It pays for these imports by exporting the produced output. Home agents maximize the stream of utility,

$$\int_0^\infty e^{-\rho t}\left[\frac{C(t)^{1-\sigma}}{1-\sigma}-\chi N(t)\right],$$

where $C(t)$ is domestic consumption and $N(t)$ is labor. The inverse of the intertemporal elasticity of substitution is σ and χ is a parameter that determines steady-state labor supply.

Domestic consumption $C(t)$ is an aggregate of individual varieties,

$$C(t)=\left(\int_0^1 C_i(t)^{\frac{\varepsilon-1}{\varepsilon}}di\right)^{\frac{\varepsilon}{\varepsilon-1}}.$$

Home asset holdings of the risk-free bond $D(t)$ evolve according to

$$D(t)=i(t)D(t)-P(t)C(t)+W(t)N(t)+\Pi(t),$$

where $i(t)$ is the nominal interest rate, $W(t)$ the wage rate, and $\Pi(t)$ are profits from firms.

Firms produce output $Y_i(t)$ of variety i according to a CES technology,

$$Y_i(t)=A(t)\left[(1-\xi)^{\frac{1}{\psi}}N_i(t)^{\frac{\psi-1}{\psi}}+\xi^{\frac{1}{\psi}}O_i(t)^{\frac{\psi-1}{\psi}}\right]^{\frac{\psi}{\psi-1}}$$

where $O_i(t)$ is oil input, ξ the share of oil in production, and ψ its elasticity of substitution with labor input. Firms face standard Calvo pricing frictions.

The foreign economy is large relative to the domestic economy and described by the equations in the previous section. It exports oil to the foreign economy in exchange for oil. Its consumption bundle is given by

$$C^*(t)=[(1-\alpha)^{\frac{1}{\eta}}(C^{F*}(t))^{\frac{\eta-1}{\eta}}+\alpha^{\frac{1}{\eta}}(C^{H*}(t))^{\frac{\eta-1}{\eta}}]^{\frac{\eta}{\eta-1}}$$

where α is the share of foreign goods in consumption and η the elasticity of substitution among home and foreign goods. Because the foreign economy is large, the share of home goods in foreign consumption is close to zero $\alpha^*\to 0$. I denote foreign quantities with a $*$, i.e. $C^*(t)$ is foreign consumption. The foreign economy is endowed with a stochastic supply of oil $O^*(t)$ and the real price of oil in foreign goods $\frac{P^{o*}(t)}{P^*(t)}$ will adjust such that the oil market clears,

$$\int_0^1 O_i^*(t)di=O^*(t)$$

Because the home economy is small, it has no influence on the oil market and thus no influence on the real price of oil denoted in foreign currency.

I assume that the law of one price holds. I then define the following relative prices. The real exchange rate $Q(t)$ is equal to

$$Q(t) = \frac{\mathcal{E}(t)P^*(t)}{P(t)},$$

where $\mathcal{E}(t)$ is the nominal exchange rate. The terms of trade,

$$S(t) = \frac{P^O(t)}{P(t)},$$

are equal to the relative price of oil.

The central bank follows a standard interest rate rule subject to the zero bound constraint.

$$i(t) = \max\{\rho + \phi\pi(t), 0\}.$$

The net foreign asset position of the home economy (denominated in home currency) evolves according to

$$\mathrm{d}NFA(t) = (P(t)C^{H*}(t) - P^O(t)O(t)) + i(t)NFA(t),$$

The workings of international financial markets play an important role in open-economy models. I therefore consider two polar cases. First, international financial markets are complete and provide full insurance. Second, there are no international financial markets and thus no borrowing or lending across countries.

D.1 Case 1: complete international financial markets When financial markets are complete, domestic and foreign consumption are related by the Backus-Smith condition

$$C(t) = \Theta C^*(t)Q(t)^{\frac{1}{\sigma}},$$

where Θ is the relative Pareto weight.

The log-linearized equations of the model are (domestic economy only)

$$
\begin{aligned}
dc(t) &= \sigma^{-1}[i(t) - r(t) - \pi(t)]dt \\
d\pi(t) &= \rho\pi(t)dt - \kappa\; mc(t)dt \\
i(t) &= \max\{r(t) + \phi_\pi \pi(t), 0\}, \quad \phi_\pi > 1 \\
c(t) &= c^*(t) + \sigma^{-1} q(t) \\
\omega(t) &= \sigma c(t) \\
y(t) &= a(t) + (1-\xi)n(t) + \xi o(t) \\
y(t) &= (1-\xi)c(t) + \xi c^{h*}(t) \\
mc(t) &= \omega(t) - \psi^{-1}(y(t) - n(t)) + (\psi^{-1} - 1)a(t) \\
o(t) &= n(t) - \psi(p^o(t) - \omega(t)) \\
p^o(t) &= p^{o*}(t) + q(t) \\
c^{h*}(t) &= c^*(t) + \eta q(t)
\end{aligned}
$$

The log-linearized equations of the model can be reduced to

$$
\begin{aligned}
dc(t) &= \sigma^{-1}[i(t) - r(t) - \pi(t)]dt \\
d\pi(t) &= \rho\pi(t)dt - \kappa\{(1-\xi)\sigma c(t) - a(t) + \xi(p^{o*}(t) + q(t))\}dt \\
i(t) &= \max\{r(t) + \phi_\pi \pi(t), 0\}, \quad \phi_\pi > 1 \\
c(t) &= c^*(t) + \sigma^{-1} q(t)
\end{aligned}
$$

The foreign economy can be written as:

$$
\begin{aligned}
dc^*(t) &= \sigma^{-1}[i^*(t) - r^*(t) - \pi^*(t)]dt \\
d\pi^*(t) &= \rho\pi^*(t)dt - \kappa\{(1-\xi)\sigma c^*(t) - a(t) + \xi p^{o*}(t)\}dt \\
i^*(t) &= \max\{r^*(t) + \phi_\pi \pi^*(t), 0\}, \quad \phi_\pi > 1
\end{aligned}
$$

Assuming that both countries face the same path $\{r(s), a(s), o^*(s)\}$, then the unique solution is,

$$
\begin{aligned}
c^*(t) &= c(t) \\
\pi^*(t) &= \pi(t) \\
q(t) &= 0
\end{aligned}
$$

Thus, a negative oil supply shock is expansionary at home if and only if it is expansionary in the foreign economy. Since the foreign economy behaves just like the closed economy in the text, the same results apply.

D.2 Case 2: incomplete financial markets While the equilibrium under complete markets featured $\Theta = 1$, in the equilibrium under incomplete markets $NFA(0)$ is given. As in Farhi and Werning [2012b] the incomplete market allocation is the sum of two components

– the complete market allocation, denoted CM, and an additional term, denoted δ_x^{IM} for variable x,

$$c^{IM}(t) = c^{CM}(t) + \delta_c^{IM}, \quad \pi^{IM}(t) = \pi^{CM}(t) + \delta_\pi^{IM}, \quad o^{IM}(t) = o^{CM}(t) + \delta_o^{IM}, \quad n^{IM}(t) = n^{CM}(t) + \delta_n^{IM}.$$

The δ^{IM}-terms are constant because the home economy is forward looking and thus instantaneously adjusts to the new wealth level.

With zero initial net financial assets, $NFA(0) = 0$, I use the evolution of the net financial asset position to solve for the new Pareto weight θ. Define $N\tilde{F}A(t) = \frac{C(s)^{-\sigma}}{P(s)} NFA(t)$ as real financial assets in utility units. Given the no-Ponzi scheme condition, $N\tilde{F}A(t)$ must satisfy,

$$N\tilde{F}A(t) = \int_s^\infty \frac{C(s)^{-\sigma}}{P(s)} \left(C^{h*}(s) - \frac{P^O(t)}{P(t)} O(t) \right) ds.$$

Log-linearizing this condition and using $N\tilde{F}A(0) = 0$ we obtain,

$$\int_0^\infty e^{-\rho s} c^{h*}(s) ds = \int_0^\infty e^{-\rho s} [p^o(t) + o(t)](s) ds,$$

which states that trade must be balanced in the long-run. We use this condition to solve for $\theta = \ln \Theta$ that satisfies,

$$c(t) = \theta + c^*(t) + \frac{1}{\sigma} q(t).$$

Substituting the linearized equation into the international budget constraint yields a solution for θ,

$$\theta = \frac{1}{\xi(1-\xi)(1-\sigma\eta) + \sigma(\eta-1)(1-\xi) + \sigma\psi(1-\xi)} \int_0^T [c^{*,CM}(s) - p^{o*,CM}(t) - o^*(t)] ds$$

Given θ, we can calculate the incomplete market component of the home allocation as follows

$$\begin{aligned} \delta_c^{IM} &= \xi\theta, & \delta_y^{IM} &= -\xi(1-\xi)(\sigma\eta - 1)\theta, & \delta_\pi^{IM} &= 0, \\ \delta_q^{IM} &= -(1-\xi)\sigma\theta, & \delta_n^{IM} &= -\xi[(1-\xi)(\sigma\eta-1) + \sigma\psi]\theta. \end{aligned}$$

If $\sigma\eta \geq 1$ then a wealth transfer has opposite effects on consumption than on GDP ($= n$). Take a negative wealth transfer, $\theta < 0$, for example. This causes the real exchange rate to depreciate, $\delta_q^{IM} > 0$. If foreign import demand is sufficiently elastic (high η) and/or the depreciation is large enough (high σ), then the export channel dominates the decline in domestic consumption from lower wealth and gross output rises. Further the depreciation raises the real import price of oil which causes a substitution towards labor. Thus, GDP also rises when gross output rises.

The condition $\sigma\eta \geq 1$ is satisfied in standard parameterizations of these models. For example, Ferrero, Gertler, and Svensson [2008] set $\sigma = 1$ and $\eta = 2$, while Bodenstein, Erceg, and Guerrieri [2011] set $\sigma = 1$ and $\eta = 1.5$. Obstfeld and Rogoff [2005] argue that $\eta = 2$ is a reasonable calibration balancing micro and macro estimates. However, they suggest that micro estimates (which imply higher elasticities) are likely less biased, and thus also experiment with higher values of η. If we set $\eta \geq 2$, then any intertemporal elasticity of substitution $\sigma^{-1} \leq 2$ will satisfy this condition. Bayesian estimation of macro models, typically produce estimates in that range (e.g., Smets and Wouters [2007]). Thus, at least in these standard parameterizations of the model, wealth effects do not produce the correct co-movement to explain my findings in section 3.

E 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.

E.1 Data sources The data sources and their frequency are tabulated in Appendix Table 2.

Table 2 – Inflation Expectations Data Sources

Country	Source	Reference Price	Forecast Frequency	
			Q4-Q4 Forecasts	Annual Forecast
Japan	Consensus Economics	CPI	Quarterly	Monthly
U.S.	Consensus Economics	CPI	Quarterly	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.

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

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

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 $\mathbb{E}_t\pi_t^{-1,0} = \pi_{t-1}^{0,1}$. 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, $\mathbb{E}_t\pi_t^{7,8} = \zeta\pi_{t-1}^{7,8}$. 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 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 8 I plot the baseline impulse response function from Section 3 using inflation expectations both filters. Inflation expectations rise less than in the baseline, but are still positive consistent with the oil shocks capturing a contraction in supply.

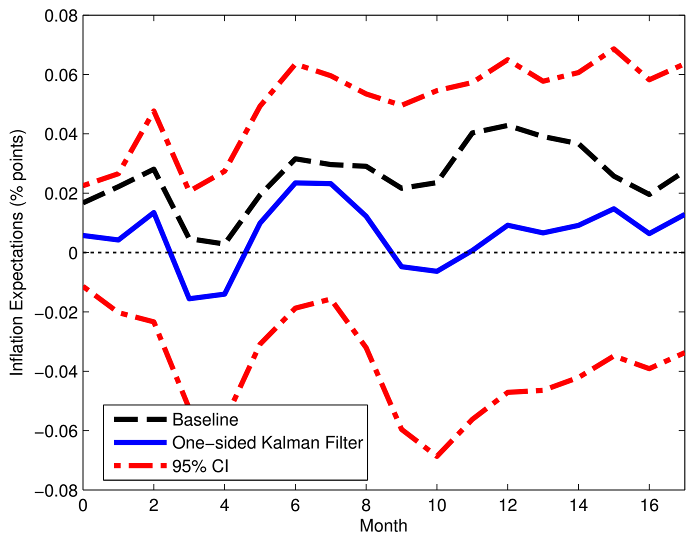

Figure 8 – Impulse Response Functions to Oil Shocks for inflation expectations derived from smoothed estimates ("Baseline") and for inflation expectations derived from the one-sided filter. Two-standard-error confidence intervals are constructed by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix.

The basic set-up is slightly modified for the U.K., where I also use HM Treasury forecasts to determine RPI inflation expectations. 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.

F Oil Shocks: Standard Error Correction

This section derives the standard error correction when the generated regressor is the residual of a first stage. I consider a univariate setting, but with more cumbersome notation, this derivation can also be extended to a multivariable setting where the same results obtain (see Wooldridge [2001] p.139-142). The true model is given by

$$y = x\beta + u,$$

where y is the outcome variable (e.g., unemployment), x is the true oil supply shock, and u is a residual. The oil supply shock x is unobserved but can be obtained in a first stage

$$z = w\delta + x.$$

With z (global oil supply) and w (lagged real economic activity) known, δ is estimated by OLS,

$$\hat{\delta} = (w'w)^{-1}w'z$$

Thus, the estimated residuals are

$$\hat{x} = (I - P_w)z.$$

where $P_w = w(w'w)^{-1}w'$. Note that these residuals are orthogonal to w by definition of OLS.

We can rewrite the second stage as

$$y = \hat{x}\beta + (x - \hat{x})\beta + u,$$

and we are interested in the distribution of $\hat{\beta} = (\hat{x}'\hat{x})^{-1}\hat{x}'y$ – the estimated impact of an oil supply shock on unemployment. Substituting this into the equation, rearranging and multiplying by $\sqrt{N}$ yields,

$$\begin{aligned}
\sqrt{N}(\hat{\beta} - \beta) &= (\hat{x}'\hat{x})^{-1}\{\frac{1}{\sqrt{N}}\hat{x}'[(x - \hat{x})\beta + u]\} \\
&= (\hat{x}'\hat{x})^{-1}\{[\underbrace{\frac{1}{N}\hat{x}'w\beta}_{\equiv G} \underbrace{\sqrt{N}(\hat{\delta} - \delta)}_{\text{sampling uncertainty}} + \frac{1}{\sqrt{N}}\hat{x}'u]\}
\end{aligned}$$

Focussing on the first term we obtain

$$G = \beta\frac{1}{N}\hat{x}'w = 0,$$

because $\hat{x}$ (the estimated oil shock) is orthogonal to w (lagged real activity) by construction. Thus, by Slutsky Theorem the sampling uncertainty in the first stage does not affect the asymptotic variance in the second stage. Hence no correction is necessary.

More generally, so long as one is interested in the hypothesis $\beta = 0$ sampling uncertainty also does not affect the asymptotic variance under the null.

G Oil Shocks: Robustness

This appendix investigates the robustness in economic responses to oil shocks. In the following sections I conduct the following checks: excluding post-2006 data, letting 4-quarter ahead inflation expectations be measured from the current quarter, excluding the Eurozone, dropping lagged dependent variables, excluding outliers, using HP-filtered data, and using various lag lengths for oil shocks.

G.1 Excluding post-2006 data In Appendix Figure 9 I plot the IRFs when the financial crisis is dropped from the sample. Thus, these IRFs are constructed based only on Japanese data. The results are very similar to the baseline, so the contractionary effects of oil supply shocks are not drive by particularities of recent events, or by pooling of the data. Figure Appendix Figure 10 then compares normal times in Japan with the pre-2008 ZLB IRFs. As in Figure 5, oil supply shocks are more contractionary at the ZLB than in normal times. Thus, these results are not driven by a country-composition bias.

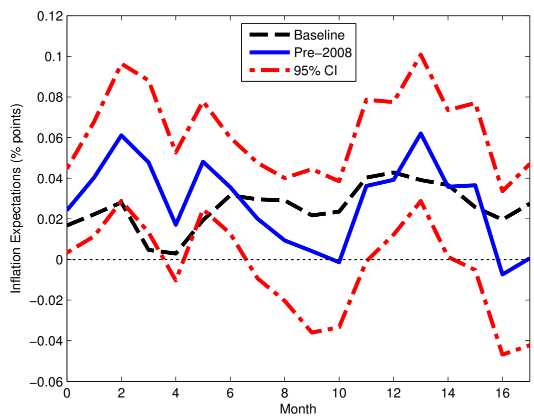

(a) 4Q Ahead Expected Inflation

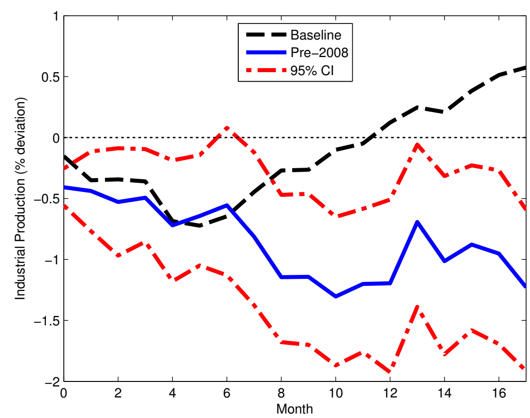

(b) Log Industrial Production

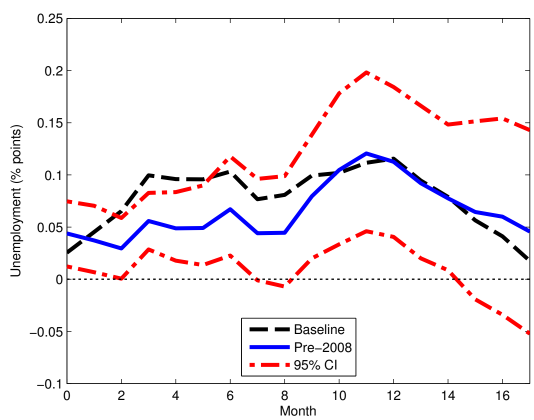

(c) Unemployment

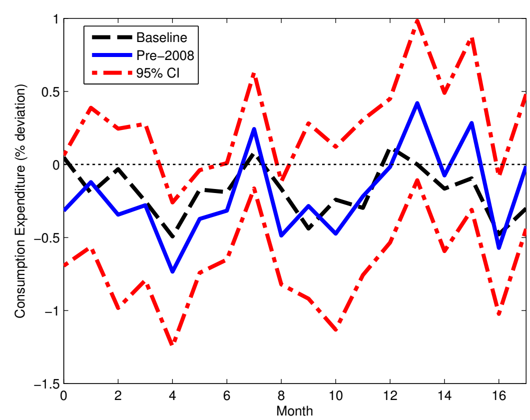

(d) Consumption Expenditures

Figure 9 – Impulse response functions to oil supply shocks *excluding post-2006 data.* Two-standard-error confidence intervals are constructed by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix.

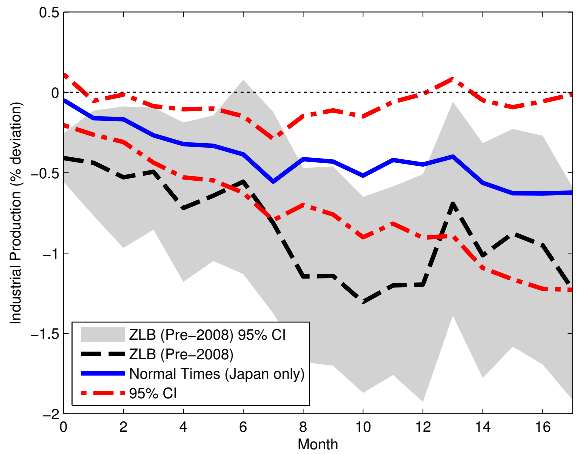

(a) Log Industrial Production

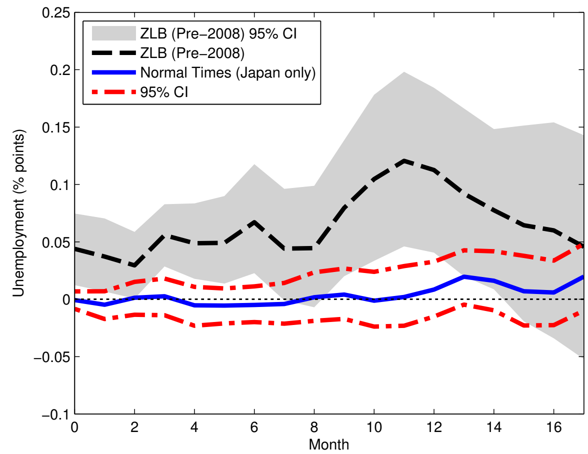

(b) Unemployment

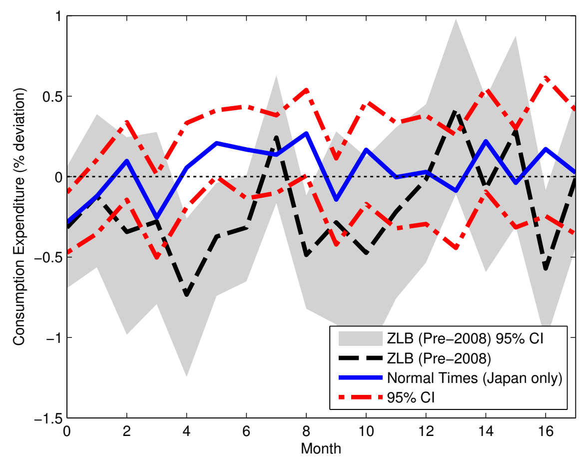

(c) Consumption Expenditures

Figure 10 – Impulse response functions to oil supply shocks *1981-2007 in Japan*: 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 Driscoll-Kraay covariance matrix. IRFs for inflation expectations are not reported because these data begin in 1995. See Section 3 for details.

G.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 in Appendix Figure 11 are essentially identical.

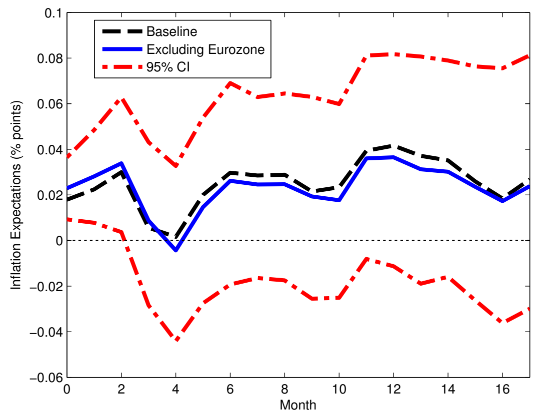

(a) 4Q Ahead Expected Inflation

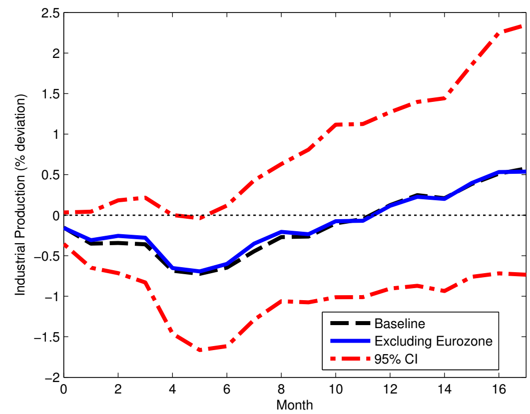

(b) Log Industrial Production

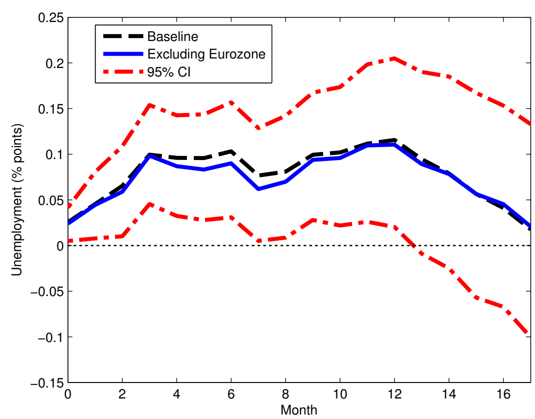

(c) Unemployment

Figure 11 – Impulse response functions to oil supply 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.

G.3 Non-linear Projection As an alternative to separating time-periods into discrete ZLB and non-ZLB regimes, I let the state vary continuously, $s_t = F(z_t)$, where $0 \leq F(z) \leq 1$ captures the degree to which the ZLB binds. As indicator of the severity of the ZLB constraint, I let z_t be the centered 3-month moving average of the two-year bond yield. In the event that this value is negative, I set it to 1bp. Relative to shorter-maturity yields, focussing on the two-year bond yield has the advantage that it takes into account differences in the expected duration of the ZLB.

Following Auerbach and Gorodnichenko [2012], I then define F as

$$F(z_t) = \frac{\exp\left\{-\gamma \frac{\ln(z_t)-c}{\sigma_{\ln z}}\right\}}{1+\exp\left\{-\gamma \frac{\ln(z_t)-c}{\sigma_{\ln z}}\right\}}. \tag{14}$$

The parameter γ captures the smoothness between regimes. When $\gamma = 0$ the regime is constant, whereas if $\gamma \to \infty$ the regime switches discretely between 0 (if $\ln(z_t) < c$) and 1 (if $\ln(z_t) > c$). The parameter c determines the switching point – when $\ln(z_t) = c$ then $F = 0.5$ for all $\gamma < \infty$. I set $c = 0$, which implies that regimes switch when the two-year bond yield is around 1%. This value captures the notion that Japan entered the ZLB in late 1995 and that the ZLB started binding in the U.S. in December 2012. In contrast, setting the cut-off at 0.5% reduces the ZLB sample essentially to Japan. I let $\gamma = 10$, which generates a sharp contrast between the regime above and below the threshold, similar to the discrete regimes in Table 1. However, other values in the range from 1.5 to 10 yield similar results.

The path of the regime by country is plotted in Appendix Figure 12 from 1993 onwards (for all previous dates $F(z(t)) < 0.00004$). While the categorization is similar to the dates in Table 1, there are some notable differences: First, up until the very end of the sample, the ZLB constraint binds more strongly in Japan than elsewhere. Second, for all countries except Japan, there is significant variation in the tightness of the ZLB constraint since the end 2008. Third, the ZLB constraint is less tight in Sweden and Canada throughout. Fourth, episodes such as the early 2000s in the U.S. have some semblance to the ZLB regime.

Appendix Figure 13 displays the IRFs from the projection method for the ZLB and normal times. For comparison, the baseline IRFs for the ZLB and normal times are also plotted. The main results in the text also obtain using the projection method: First, negative supply shocks are contractionary at the ZLB. Second, they are more contractionary than in normal times. If anything, the projection method further accentuates these differences.

One difference between the baseline estimates and the projection method is the behavior of inflation expectations. Inflation expectations implied by the projection method are lower in the first nine months and then rise above the baseline estimates. These differences arise due to a few observations in early 2008, when a series of positive supply shocks first coincides with a rise in inflation expectations (in early 2008) and then a decline in inflation expectations (in late 2008). Excluding these observations yields very similar results to the baseline inflation expectations response, while barely affecting the behavior of real variables.

G.4 4Q-Expectations Measured from Current Quarter In Section 3 the change in inflation expectations was calculated ignoring the current quarter. Thus, 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 14. As expected,

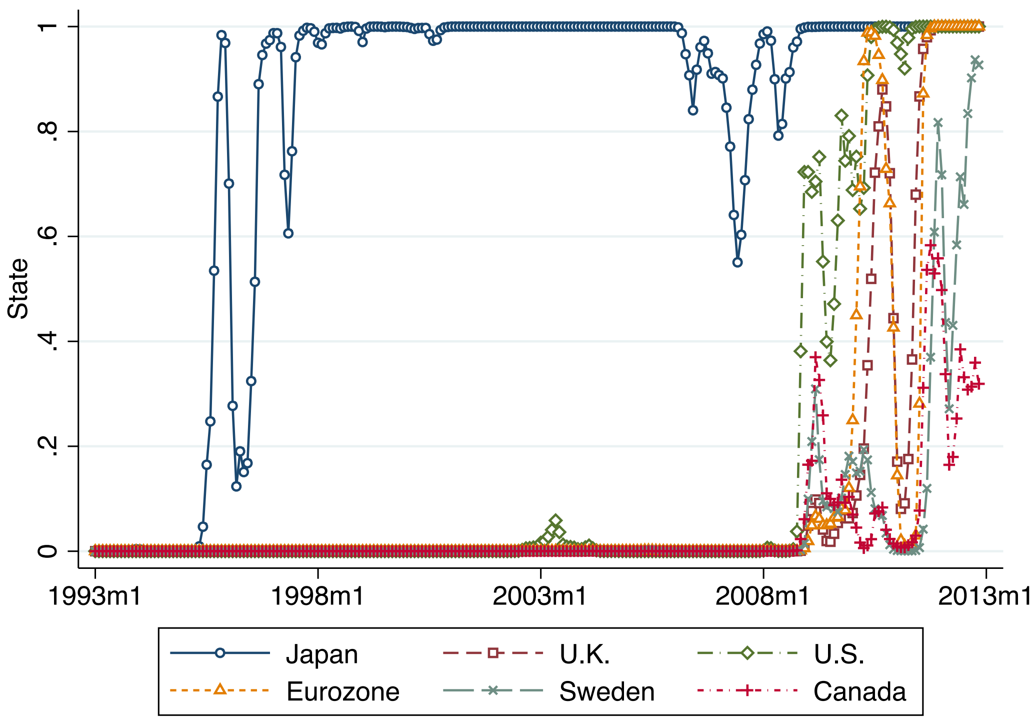

Figure 12 – Regime based on Equation (14) where z_t is the centered 13-month moving average of the two-year bond yield, with parameters $\gamma = 10$ and $c = 0$. A value of 1 denotes the ZLB regime whereas a value of 0 denotes normal times.

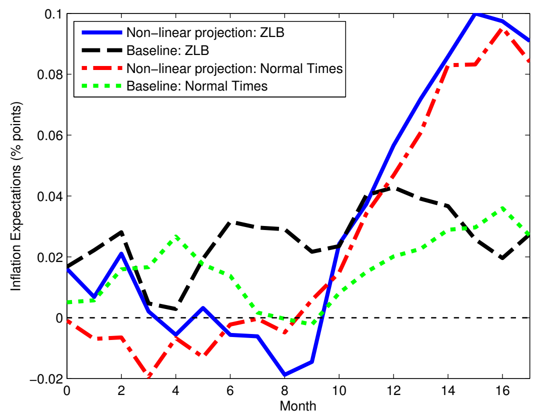

(a) 4Q Ahead Expected Inflation

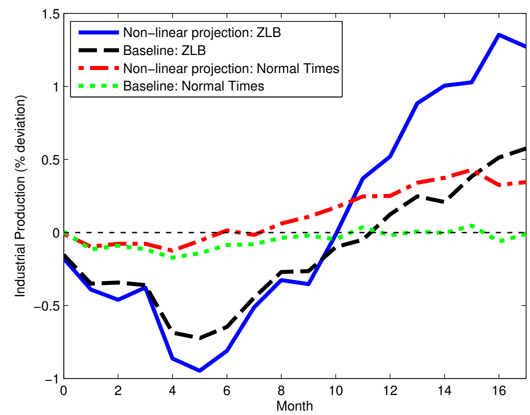

(b) Log Industrial Production

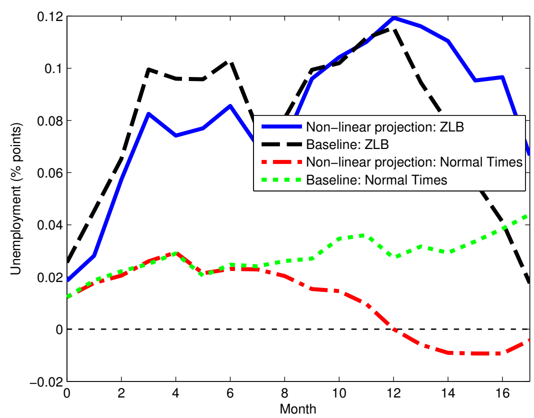

(c) Unemployment

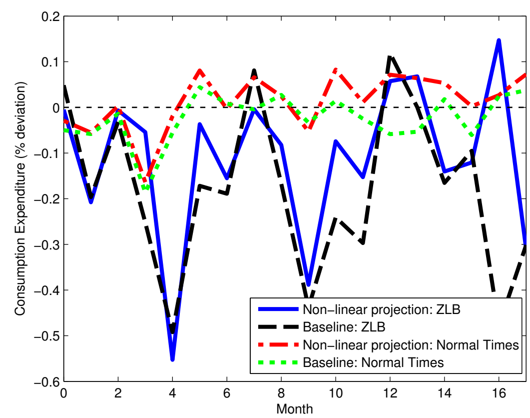

(d) Consumption Expenditures

Figure 13 – Impulse response functions to oil supply shocks based on non-linear projection with regime given by equation (14). Baseline estimates are based on dates from Table 1.

it displays a larger increase in inflation expectations than the baseline since the latter is a conservative estimate.

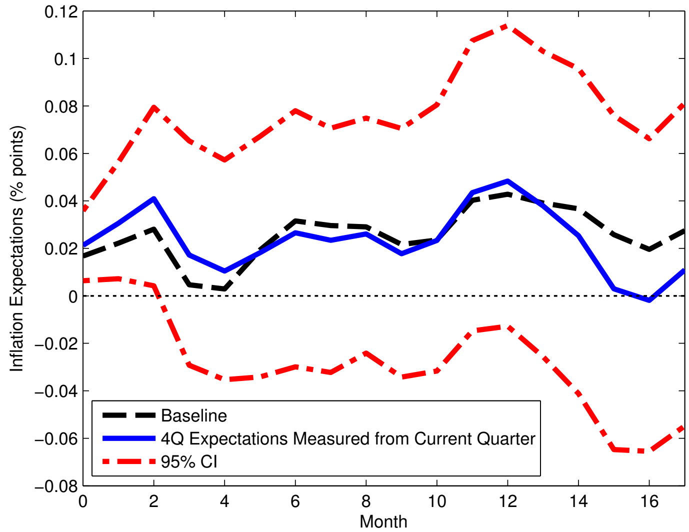

Figure 14 – Impulse Response Functions to Oil Shocks for 4-quarter ahead inflation expectations measured from the *current* quarter. "Baseline" refers to 4-quarter ahead inflation expectations measured from the *next* quarter. Two-standard-error confidence intervals are constructed by Monte-Carlo draws from a normal distribution with variance equal to the estimated covariance matrix.

G.5 Inflation Expectations Measured from Inflation Swap Rates This subsection demonstrates that market-based inflation expectations also rise following an oil supply shock by using inflation swap rates. A downside of using inflation swap data is that it is not available before 2004 in the U.S., the Eurozone and the U.K. and not before 2007 in Japan and Sweden. Since this significantly constrains the sample, I estimate Equation (7) without controlling for the bond rate, industrial production, or unemployment. The resulting estimates for γ_0 at various maturities are plotted in Appendix Figure 15. I find that an oil supply shock raises short-term inflation expectations, while long-term inflation expectations are close to zero. Thus, these results show that oil supply shocks raise inflation expectations for about a year.

G.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 16. 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 17 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.

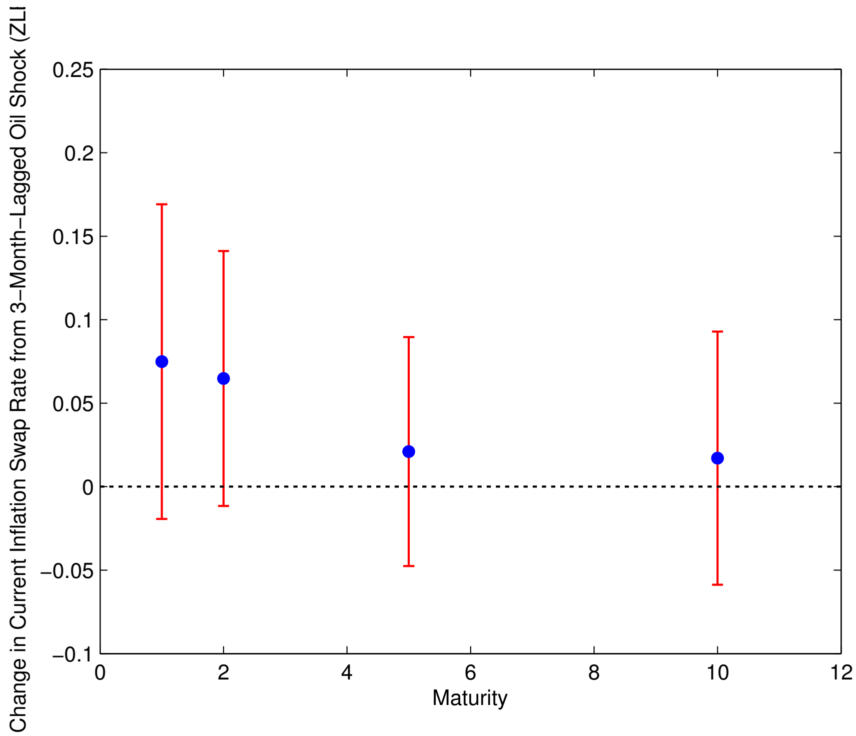

Figure 15 – Contemporaneous change in the inflation swap rate after an oil supply shock has hit the economy (γ_0 in Equation (7)). Point estimates are solid circles. Error bars show the 95% Driscoll-Kraay confidence interval around the point estimates. Estimated based on an autoregressive distributive lag equation, where the dependent variable enters with 12 lags and oil shocks enter with 12 lags. See Section 3 for details.

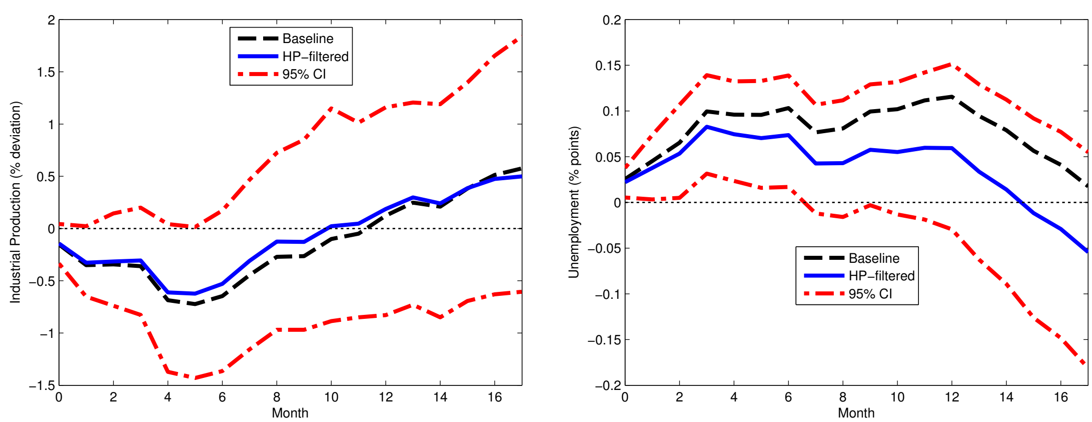

(a) Log Industrial Production

(b) Unemployment

Figure 16 – 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.

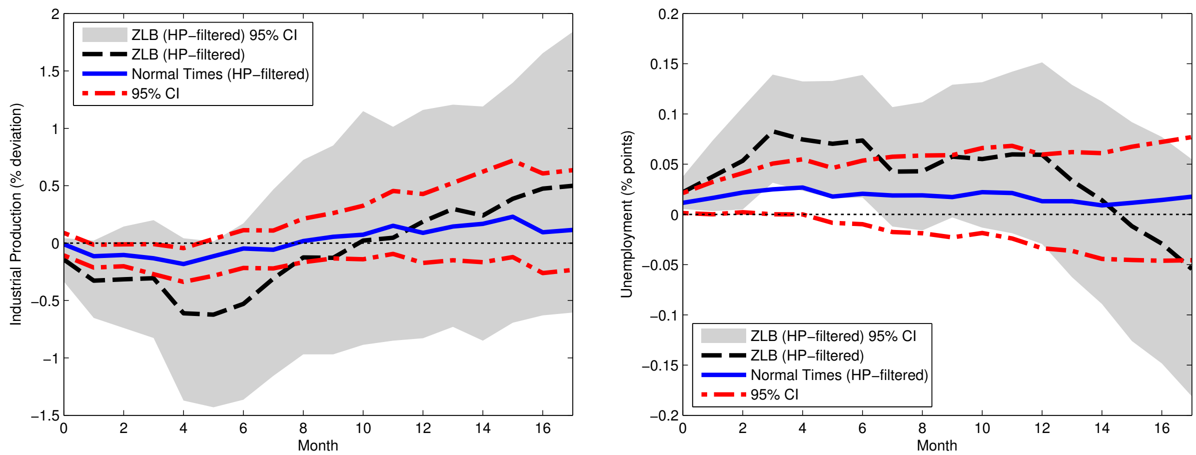

(a) Log Industrial Production (b) Unemployment

Figure 17 – 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 for details.

G.7 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 3 I tabulate the values of the IRF one month and four months 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 3 – 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	5 Months	1 Month	5 Months	1 Month	5 Months	1 Month	5 Months
12	0.022	0.010	-0.315	-0.753	0.040	0.096	-0.266	-0.280
	[0.002,0.043]	[-0.030,0.048]	[-0.635,0.070]	[-1.623,-0.014]	[0.014,0.065]	[0.036,0.163]	[-0.586,0.031]	[-0.571,0.043]
24	**0.022**	**0.019**	**-0.350**	**-0.723**	**0.045**	**0.096**	**-0.196**	**-0.172**
	[0.004,0.040]	**[-0.022,0.058]**	**[-0.688,0.014]**	**[-1.600,-0.030]**	**[0.013,0.080]**	**[0.038,0.161]**	**[-0.489,0.104]**	**[-0.445,0.166]**
36	0.025	0.027	-0.353	-0.698	0.045	0.088	-0.173	-0.235
	[-0.001,0.051]	[-0.013,0.065]	[-0.671,-0.010]	[-1.534,-0.075]	[0.012,0.080]	[0.034,0.149]	[-0.439,0.103]	[-0.498,0.096]

95% Confidence intervals in brackets based on Driscoll-Kraay s.e. "1 Month" and "5 Months" are the impact and associated 95% CI at 1 month and 5 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).

H Oil Shocks: Event Study

This appendix provides corroborating evidence for Section 3. Rather than constructing oil shocks based on a statistical model as in Kilian [2009], in this section I use an event study methodology around the Libyan civil war. Before the civil war, Libya produced approximately 2% of global oil supply. The beginning of the civil war is typically dated on February 15^{th} 2011. Foreign intervention officially commenced on March 19^{th}. This conflict caused a significant contraction in Libyan oil production, which by April had declined by almost 90% relative to pre-war levels. Thus, the Libyan civil war constituted a relatively large and exogenous shock to global oil production. Consistent with this interpretation the Kilian [2009] series also displays a 1.4 standard deviation oil shocks over February 2011 and March 2011. This suggests that the Kilian [2009] series does indeed successfully identify exogenous supply shocks. However, while this oil shock is plausibly exogenous, to the extend that foreign governments increase military spending it may be correlated with a positive demand shock. Thus, this event study will be *biased against* finding contractionary effects from negative supply shocks.

To determine the effects of this oil shock on expected inflation and real economic activity I proceed as in Section 4. The timing of the oil supply disruptions associated with Libyan civil war are less precise than the timing of the Japanese earthquake – it could plausibly be dated on February 15^{th} 2011 or the beginning of foreign intervention on March 19^{th}. Thus, I use both these dates in the event study that comprises the U.S., the U.K., and the Eurozone. In Appendix Figure 18 I compare pre-February 15^{th} forecasts with post-February 15^{th} forecasts, and in Appendix Figure 19 I compare pre-March 19^{th} forecasts with post-March 19^{th} forecasts. Both figures also display ex-post data.

In all six cases there are significant increases in expected inflation, and even higher ex-post inflation outcomes, consistent with there being a negative supply shock. In addition, in all countries ex-post real output was below the February/March 2011 forecast, and in four out of six cases output forecasts were revised downwards. These negative comovements between expected inflation and output, as well as actual inflation and output, suggest that the Libyan civil war and associated oil supply disruption did indeed constitute have contractionary effects at the ZLB. This is remarkable since this war likely triggered greater military spending, i.e., a positive demand shock. Thus, this event study corroborates the findings from Section 3 – that oil supply shocks are contractionary at the ZLB despite lowering real interest rates.

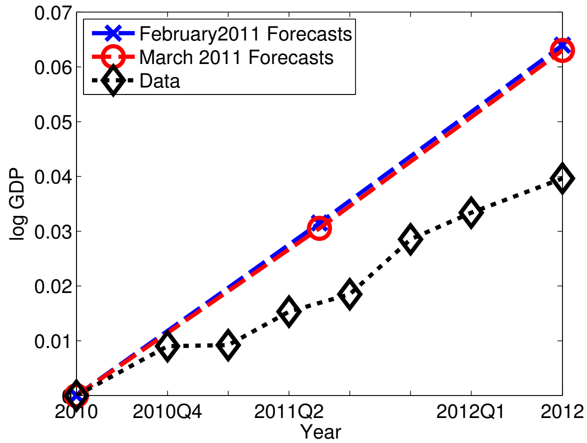

(a) U.S.: GDP

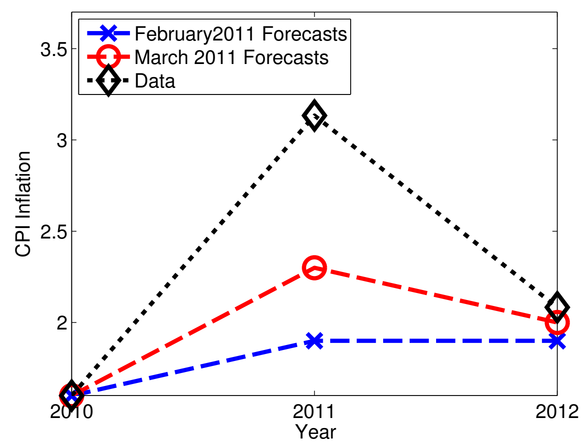

(b) U.S.: Inflation

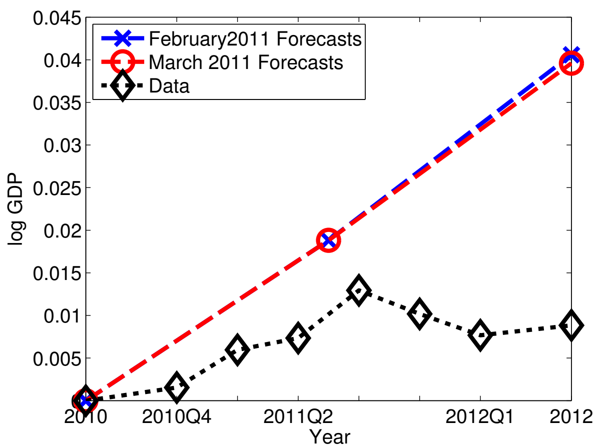

(c) U.K.: GDP

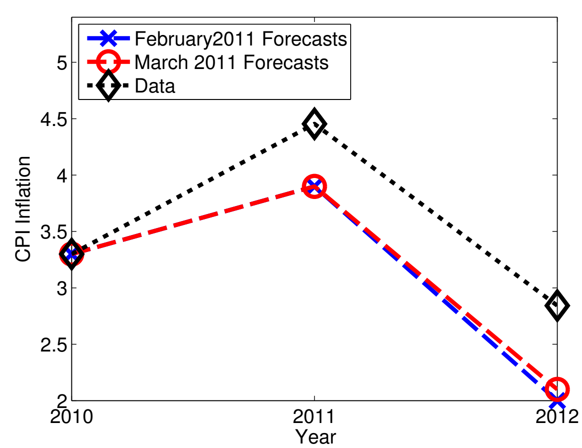

(d) U.K.: Inflation

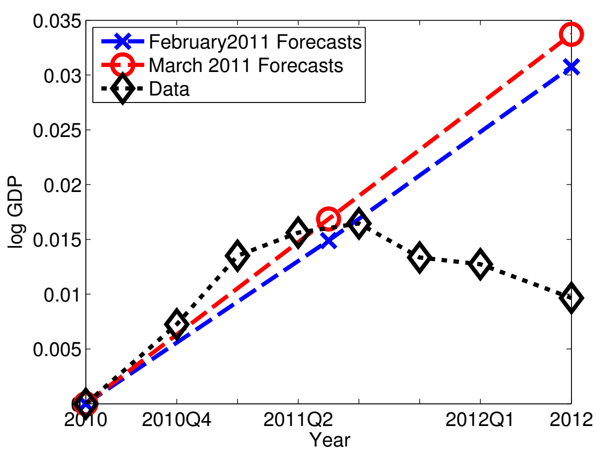

(e) Eurozone: GDP

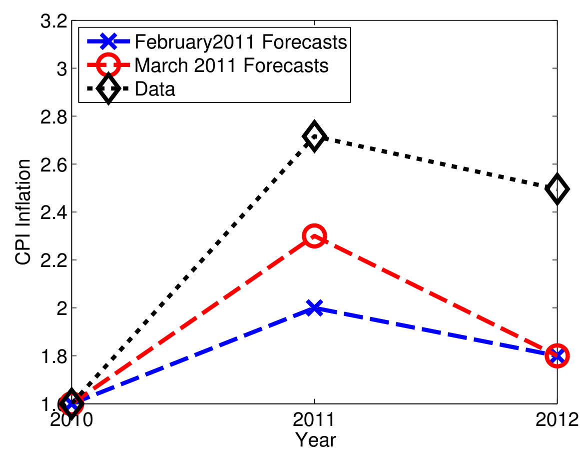

(f) Eurozone: Inflation

Figure 18 – Consensus Economics forecasts from before the Libyan uprising (February 2011) and after (March 2011). Forecasts are for annual GDP and year-on-year inflation. GDP data is annual for 2010 and 2012 and quarterly from 2010Q4 until 2012Q1.

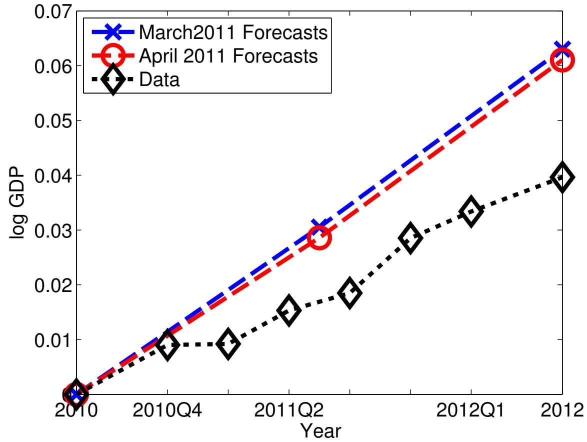

(a) U.S.: GDP

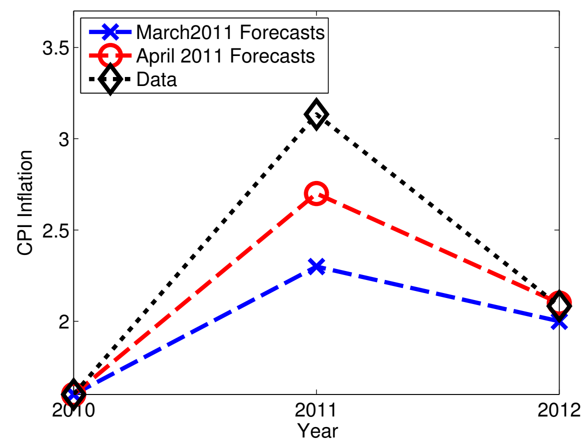

(b) U.S.: Inflation

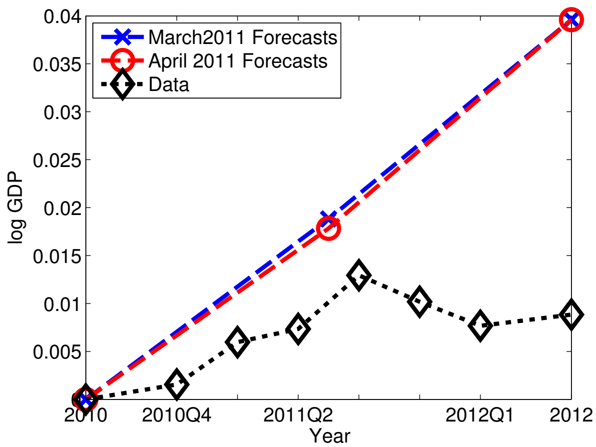

(c) U.K.: GDP

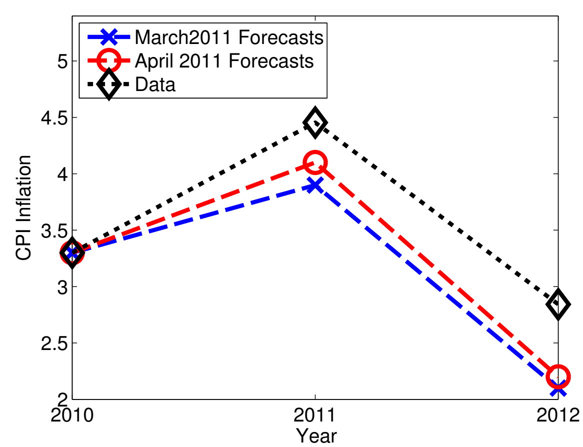

(d) U.K.: Inflation

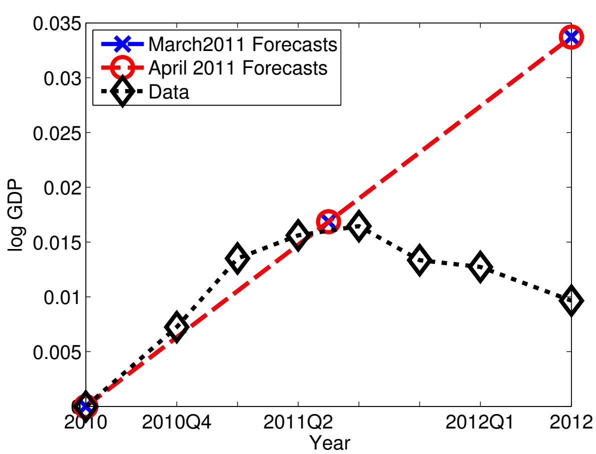

(e) Eurozone: GDP

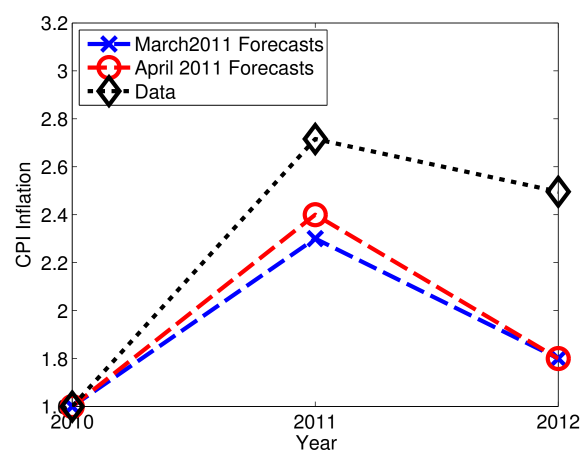

(f) Eurozone: Inflation

Figure 19 – Consensus Economics forecasts from before foreign intervention in the Libyan civil war (March 2011) and after (April 2011). Forecasts are for annual GDP and year-on-year inflation. GDP data is annual for 2010 and 2012 and quarterly from 2010Q4 until 2012Q1.

Appendix References

Auerbach, Alan and Yuriy Gorodnichenko, "Fiscal multipliers in recession and expansion," in "Fiscal Policy after the Financial Crisis," University of Chicago Press, 2012.

Bodenstein, Martin, Christopher J Erceg, and Luca Guerrieri, "Oil shocks and external adjustment," *Journal of International Economics*, 2011, *83* (2), 168–184.

Farhi, Emmanuel and Iván Werning, "Fiscal Multipliers: Liquidity Traps and Currency Unions," 2012.

Ferrero, Andrea, Mark Gertler, and Lars EO Svensson, "Current account dynamics and monetary policy," Technical Report, National Bureau of Economic Research 2008.

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

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.

Obstfeld, Maurice and Kenneth S Rogoff, "Global current account imbalances and exchange rate adjustments," *Brookings papers on economic activity*, 2005, *2005* (1), 67–146.

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.

Wooldridge, Jeffrey M, *Econometric analysis of cross section and panel data*, MIT press, 2001.