How Oil Shocks Propagate: Evidence on the Monetary Policy Channel*

Wataru Miyamoto† Thuy Lan Nguyen‡ Dmitriy Sergeyev§

March 11, 2026

Abstract

Using high-frequency responses of oil futures prices to prominent oil market news, we estimate the effects of oil supply news shocks when systematic monetary policy is switched off by the zero lower bound (ZLB) and when it is not (normal periods) in Japan, the United Kingdom, and the United States. We find that negative oil supply news shocks are less contractionary (and even expansionary) at the ZLB compared to normal periods. Inflation expectations increase during both periods, while the short nominal interest rates remain constant at the ZLB, pointing to the importance of monetary policy for oil shock propagation.

Keywords: oil price shocks, high-frequency identification, zero lower bound, systematic monetary policy

JEL Classification: E5, E7, G4

*The views expressed in this paper are those of the authors and do not represent the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System. Any errors are ours. We thank the editor and four anonymous referees, Regis Barnichon, Christiane Baumeister, Martin Bodenstein, James Cloyne, John Cochrane, Riccardo Degasperi, Laurent Ferrara, Yuriy Gorodnichenko, Oscar Jorda, Diego Kanzig, Neil Mehrotra, Dmitry Mukhin, Emi Nakamura, Maurice Obstfeld, Carolin Pflueger, Ricardo Reis, Jon Steinsson, Ludwig Straub, Ivan Werning, Dan Wilson, Christian Wolf, and seminar and conference participants at the Universitat Autonoma de Barcelona, Bocconi University, Texas A&M University, UC Berkeley, the 2023 Barcelona Summer Forum, Copenhagen Business School, 2023 CSEF-IGIER Symposium, Banque de France, 22nd Macro Finance Workshop at the University of Chicago, 2023 NBER Summer Institute, NYU Abu Dhabi, 2022 System Energy meeting, 2023 SED, Trinity College Dublin, 2023 T2M, Workshop on Emirical Macroeconomics for their helpful feedback and discussions. Chiara Bargellesi, Ethan Goode, Roberto Patino, and Aren Yalcin provided excellent research assistance.

†The University of Hong Kong, wataru.miyamoto1@gmail.com.

‡Federal Reserve Bank of San Francisco, thuylan.nguyen00@gmail.com.

§Bocconi University, Department of Economics, dmytro.sergeyev@unibocconi.it.

1 Introduction

Oil price changes are among the most prominent macroeconomic disturbances. They are large and sudden, especially when caused by wars involving oil producers. Oil price changes affect all countries at the same time because crude oil is an internationally traded commodity. Moreover, according to the conventional supply-side view, it is difficult for monetary policy to stabilize both output and inflation in the face of oil price changes because inflation and output move in opposite directions. Recent events, such as the oil price increase in 2021-22, the 2022 Russian invasion of Ukraine, and the burst of high inflation worldwide in 2021-22, reignited interest in exploring the impact and mechanisms of the effects of oil price changes.

The fact that most post-WWII recessions in the United States followed oil price spikes has inspired a vast empirical literature on oil shock effects (Hamilton, 1983; Kilian, 2009; Kilian and Murphy, 2012; Baumeister and Hamilton, 2019; Kanzig, 2021). This literature has documented substantial effects of oil shocks on macroeconomic variables, relying on a variety of econometric approaches to identify oil shocks. At the same time, how exactly oil shocks propagate is still unclear. Theoretical literature proposes aggregate supply channels, such as variations in the cost of production, which can be amplified by markup endogeneity (Rotemberg and Woodford, 1996) and production network nonlinearities (Baqaee and Farhi, 2019), as well as aggregate demand channels, such as the interaction between nominal price stickiness and real wage rigidity (Blanchard and Gali, 2009) and the presence of financially constrained households (Chan, Diz and Kanngiesser, 2022; Auclert, Monnery, Rognlie and Straub, 2023).

One aggregate demand channel stands out. It stems from the observation made in Bernanke, Gertler and Watson (1997) that most post-WWII recessions in the United States have followed monetary policy hikes. This raises the possibility that the estimated effects of oil shocks on real activity are also driven by systematic monetary policy reaction to these shocks. Such logic is central in standard New Keynesian models (Woodford, 2003). In particular, the central bank may tighten monetary policy following an oil price spike because the inflation rate increases, or inflation expectations go up, or because the central bank fears a future inflation rate increase. Such systematic monetary policy reactions can exacerbate the direct adverse effect of oil price spikes on output. Moreover, New Keynesian models can imply that oil price spikes are expansionary (Eggertsson, 2008; Bodenstein, Guerrieri and Gust, 2013), for example, if the central bank does not change the short-term nominal interest rate following an oil price spike. Intuitively, with higher inflation expectations and a constant nominal interest rate, the real interest rate drops, stimulating aggregate economic activity by increasing consumption and investment demand.

In this paper, we empirically evaluate the role of monetary policy in oil shock propagation. We leverage the fact that several countries have recently experienced zero lower bound (ZLB) episodes, when the short-term nominal interest rate—a standard conventional monetary policy tool—stayed close to zero.[1] This allows us to compare the effects of oil shocks during the ZLB periods, when monetary policy did not respond actively, and the normal periods, when the short-term nominal interest rate was not constrained by the ZLB. Specifically, we examine how Japan, the United Kingdom, and the United States responded to oil shocks between 1975:1 and 2019:12, differentiating between the ZLB and the normal periods.

Our primary focus is on Japan because it has the longest ZLB experience, which started at the end of 1995. Moreover, Japan experienced several business cycles during the ZLB periods, making it possible to average out the effects of potentially differential impacts of shocks during booms and busts. In addition, Japan was a net oil importer throughout our sample. Looking at the data in other countries, such as the United Kingdom and the United States, is helpful because the ZLB periods do not coincide, which alleviates the concern that some global factor could be responsible for results in all countries.

To extract exogenous and unexpected variation in oil prices, we build on the recent literature that uses high-frequency data to identify macroeconomic shocks (Kuttner, 2001; Gürkaynak, Sack and Swanson, 2004; Gertler and Karadi, 2015; Nakamura and Steinsson, 2018a). In particular, we follow Kanzig (2021), who uses the changes in oil futures prices in a tight window around the Organization of the Petroleum Exporting Countries (OPEC) production announcements. The series of oil futures price changes then becomes an external instrument in an oil market vector-autoregression (VAR), allowing us to estimate structural oil shocks, which we refer to as oil supply news shocks. These shocks change oil prices on impact and oil production gradually.

Using the state-dependent local projections method (Jorda, 2005), we estimate that oil supply news shocks are less contractionary, and at times even expansionary, in the ZLB compared to the normal periods. In particular, in our baseline specification using monthly data from Japan, industrial production increases by about 1 percent after one year during the ZLB periods following a shock that increases the oil price by 10 percent. Outside of the ZLB periods, industrial production falls by over 1 percent 12 months after the shock. The differences in the two responses are statistically significant at conventional levels. It is particularly remarkable that the increase in industrial production in response to the negative oil supply news shock during the ZLB periods is significantly different from

[1] In practice, the short-term nominal interest can go slightly below zero. This has led researchers to refer to the bound as the effective lower bound. In this paper, we will use the term ZLB to refer to the effective lower bound.

zero at some horizons. The unemployment rate in Japan exhibits similar patterns with the opposite sign: it decreases significantly in the ZLB periods, while it stays near zero or increases outside of the ZLB. Since Japan has long periods with a zero nominal interest rate, we also estimate the effects of oil supply shocks using quarterly macro variables. Consistent with the results from industrial production and the unemployment rate, oil supply news shocks cause real per capita gross domestic product (GDP), consumption, and investment to increase in the ZLB periods and decrease (or do not change) outside of the ZLB. However, with fewer observations in quarterly frequency, these estimates have larger standard errors.

To explore the relevance of the monetary policy channel, we estimate the responses of the Bank of Japan target interest rate (the Call rate) and 5-year nominal interest rates, realized inflation, and inflation expectations. In the ZLB periods, both the short-term interest rate and the 5-year rate do not react significantly to oil supply news shocks. In contrast, during the normal periods, both the short and 5-year rates increase considerably after an oil price increase, and these changes are statistically different from those during the ZLB periods. Moreover, inflation expectations react more during the ZLB. These results are consistent with the channel that works through ex ante real interest rate movements. At the same time, we do not observe a stronger reaction of realized inflation during the ZLB periods. In fact, realized inflation increases more in the normal periods.

Our estimated effects of oil supply news shocks in the United Kingdom and the United States are consistent with those in Japan. For example, in the United States, industrial production increases by nearly 0.8 percent 16 months after a shock that increases the oil price by 10 percent. Outside of the ZLB period, industrial production falls up to 0.9 percent. The difference in the two responses is statistically significant at conventional levels, and the ZLB response is significantly above zero at the 5 percent level for multiple horizons. The unemployment rate repeats this pattern with the opposite sign. The nominal interest rate responds more in the normal than in the ZLB periods, and both expected and realized inflation rates respond positively in the two periods. The results in the United Kingdom are qualitatively similar. These findings imply that the patterns obtained in Japan are not a particular feature of the Japanese economy.

We extend our main analysis in a number of ways. First, we consider alternative oil supply shocks obtained in two prominent recent studies: Kilian (2009) and Baumeister and Hamilton (2019). These papers rely on structural VAR identification of oil supply shocks. We find that the differences between the responses of the Japanese economy in the ZLB and the normal periods following oil supply shocks identified as in Kilian (2009) are not statistically significant. This is consistent with Wieland (2019), who also used these oil supply shocks to estimate the response of the Japanese economy during and outside of the ZLB. At the same time, the oil supply shocks identified as in Baumeister and Hamilton (2019) produce qualitatively similar results to our baseline findings. These differences can potentially be attributed to the fact that the oil supply shocks in Kilian (2009) do not move the oil price significantly, while they do in Baumeister and Hamilton (2019). Second, we estimate the responses of several other variables, such as output without oil production, exchange rates, and stock market prices, all of which exhibit differential responses across the ZLB and the normal periods. We further assess whether unconventional monetary policy responds to oil price shocks even during the ZLB period by examining shadow rate responses and by controlling for longer-term interest rates. We find that this channel is unlikely to explain the differential responses in Japan. We estimate the reaction of economic activity in countries that experienced near-constant nominal interest rate periods, such as Canada and the euro area, and countries that did not experience their own ZLB episodes, and we do not find evidence of differential responses. Third, we investigate subsample properties of our results and re-estimate results following a refined version of the oil supply news shock that corrects for a potential information revelation effect (Jarociński and Karadi, 2020; Degasperi, 2021). We find largely similar results. Fourth, we estimate alternative specifications, including a state-dependent proxy VAR and local projections in which the oil surprise series is used as the shock. We find evidence of state-dependent responses in both cases; however, the standard errors are substantially larger in the latter. Fifth, we measure the effects of the shocks in booms and recessions and find little difference. Finally, we discuss the sensitivity of our results to variations in the number of control variables used in our regressions.

The final part of the paper presents a stylized model based on Galí and Monacelli (2002), consisting of a small open economy and the rest of the world. In the model, oil supply shocks propagate through aggregate supply, aggregate demand (via monetary policy), and international spillovers that affect both supply and demand. We formally illustrate these transmission mechanisms and compare our empirical results with the model's numerical responses to oil supply shocks.

Related literature. Our paper contributes to the literature that studies the role of monetary policy in shock propagation and, in particular, a differential impact of shocks during the ZLB (or constant nominal interest rate periods) and the normal periods. On the aggregate demand side, Miyamoto, Nguyen and Sergeyev (2018) and Ramey and Zubairy (2018) estimated higher government spending multipliers in the ZLB periods than outside of it in Japan and the United States, respectively. Different from these papers, we focus on the shocks traditionally attributed to aggregate supply disturbances.

We contribute to the supply-side literature by analyzing state-dependent responses of several countries to oil shocks identified using high-frequency techniques. Bernanke, Gertler and Watson (1997) showed that most of the U.S. macroeconomic response to oil shocks is due to a systematic monetary policy.[2] Their analysis relies on the triangular VAR identification of oil shocks and a counterfactual simulation with a fixed policy interest rate in the spirit of Sims and Zha (2006). This method aims to remove systematic monetary policy effects but can be subject to the Lucas critique.[3] Unlike Bernanke, Gertler and Watson (1997), we directly estimate our results from the episodes with fixed interest rates. Garin, Lester and Sims (2019) find that *positive* total factor productivity shocks, measured by Fernald (2014), are more expansionary in the ZLB than in the normal periods in the United States. As mentioned above, Wieland (2019) finds no difference in Japanese output reactions to oil supply shocks identified as in Kilian (2009). Our results rely on the oil shocks identified using a high-frequency approach. Moreover, we use data from several countries, which is essential when studying relatively short ZLB periods. Unlike Garin, Lester and Sims (2019) and Wieland (2019), we find that shocks that depress aggregate supply increase output during the ZLB periods relative to the normal periods. Finally, our findings are consistent with the indirect evidence in Datta, Johannsen, Kwon and Vigfusson (2021). The authors estimate that oil prices and equity returns are more positively correlated in the recent past than before and that both variables become more responsive to macroeconomic news, consistent with the prediction of a model where the ZLB alters the economic environment.

Our paper also speaks to the literature on unconventional monetary policies. A constant short-term nominal interest rate, which we use to define ZLB episodes, does not logically imply that the medium- or long-term rates are constant. Swanson and Williams (2014b,a) estimate that, while the monetary policy rate is nearly constant during ZLB episodes, the medium- and long-term rates continue to respond to macroeconomic news. Such behavior can be caused by an active use of unconventional monetary policies, such as quantitative easing and forward guidance. We do not find economically and statisti-

[2] Hamilton and Herrera (2004) show that the quantitative importance of monetary policy in propagating oil shocks in Bernanke, Gertler and Watson (1997) depends on the lag order of the estimated VAR. Kilian and Lewis (2011) argue that Volcker's 1979 tightening is crucial for estimating a strong response by the Fed to oil price shock in Bernanke, Gertler and Watson (1997).

[3] Wolf and McKay (2022) extend the Sims and Zha (2006) methodology by incorporating expectation effects using information on how economic variables respond to news shocks. Using this approach, Bjørnland, Cross and Hölz (2025) use several measures of oil shocks (Kanzig, 2021; Degasperi, 2021; Baumeister and Hamilton, 2019), together with the instruments of Bauer and Swanson (2023) to identify monetary policy shocks, and construct counterfactual responses under a policy scenario in which the policy rate does not react endogenously. They find that, absent a systematic monetary policy response, oil shocks have a weaker contractionary effect on U.S. industrial production when identified as in Kanzig (2021), and are even expansionary when identified as in Degasperi (2021) or Baumeister and Hamilton (2019). This approach, however, relies on additional structural assumptions, such as linearity and full-information rational expectations, which we avoid.

cally significant responses of longer-term nominal interest rates to oil supply news shocks in Japan during the ZLB. Debortoli, Galí and Gambetti (2020) estimate a time-varying-coefficient structural VAR driven by shocks identified through the long-run and sign restrictions. They find no evidence for differential responses of macro variables to identified macroeconomic shocks. In contrast, we use a high-frequency identification of oil supply news shocks and find significantly different responses of macro variables during the ZLB and normal periods.

The rest of the paper is organized as follows. Section 2 describes our empirical methodology and data. Sections 3 and 4 present the main results for Japan, the United Kingdom, and the United States. Section 5 examines robustness. Section 6 introduces a stylized model and compares its predictions with the empirical results. Section 7 concludes. Online Appendix Sections A-F provide data sources, additional figures, and model details and extensions.

2 Measurement, Data, and ZLB Definition

This section describes our empirical specification, details the data used for estimation, and concludes by defining ZLB periods.

2.1 Measurement of Oil Shock Effects

We measure the effects of oil supply news shocks in two steps. First, we define a daily oil supply surprise series and use it as an external instrument in an oil market VAR to identify the shocks. Second, we use these identified shocks to estimate their effects in ZLB and normal periods using state-dependent local projections.

Oil supply news shocks identification. To isolate exogenous and unanticipated changes in oil prices, we follow Kanzig (2021) and use a high-frequency approach, which has been used widely in monetary economics (Kuttner, 2001; Gürkaynak, Sack and Swanson, 2004; Gertler and Karadi, 2015; Nakamura and Steinsson, 2018a). The main idea relies on the observation that the oil market is dominated by a few large players, one of which—the Organization of Petroleum Exporting Countries (OPEC)—is responsible for about half of the world's oil production. OPEC is an intergovernmental organization that coordinates its oil production during conference meetings that end with public announcements about oil supply changes. Asset prices, particularly oil futures prices, react immediately to these announcements. Hence, looking at oil futures price changes in a narrow window of one day around OPEC announcements reveals unanticipated and arguably exogenous news about the future oil supply.[4]

The daily oil supply surprises are constructed by the first principle component of the daily price changes of the West Texas Intermediate (WTI) crude oil futures with maturities from one to twelve months in a narrow window around the OPEC announcements. These futures contracts are the most liquid contracts with the longest sample size, starting in 1983. The daily first principle component is summed up to a monthly frequency if there were more than one OPEC meeting during a particular month. As Kanzig (2021), we will refer to this monthly series as *the oil supply surprise series*.

The information contained in the OPEC announcements summarized by the oil supply surprise series is only a subset of all news about oil supply changes affecting the oil market. In addition, this series can be subject to a measurement error. As a result, we do not use the series as a shock.[5] Instead, like Kanzig (2021), following Stock and Watson (2012) and Mertens and Ravn (2013), we employ the oil supply surprise series as an external instrument—a variable that is correlated with the shock of interest but not the other shocks—in a world oil market VAR.

Formally, consider a VAR(p) process of the form

$$Y_t = C + \sum_{l=1}^{p} B_l Y_{t-l} + \mathcal{U}_t,$$

where p is the number of lags, Y_t is an $n \times 1$ vector of endogenous variables, $\mathcal{U}_t$ is an $n \times 1$ vector of reduced-form errors with the variance-covariance matrix Σ, C is an $n \times 1$ vector of constants, $\{B_l\}$ are $n \times n$ matrices of coefficients. The reduced-form errors are related to an $n \times 1$ vector of uncorrelated structural shocks $\mathcal{E}_t$ linearly as $\mathcal{U}_t = S\mathcal{E}_t$, where S is an unknown $n \times n$ matrix. If we order the oil supply news shocks to be the first element of $\mathcal{E}_t$, the goal is to identify the first column of matrix S, which we denote as $S_{1:n,1}$.

Let z_t be an external instrument, that is, the oil supply surprise series in our application. If it is a valid instrument, it must be correlated with the shock of interest, i.e., $\mathbb{E}\left[z_t \mathcal{E}_{1,t}\right] \neq 0$, where $\mathbb{E}\left[\cdot\right]$ is the expectations operator, and uncorrelated with the other shocks, i.e., $\mathbb{E}\left[z_t \mathcal{E}_{k,t}\right] = 0$ for $k = 2, 3, \ldots, n$. Under these assumptions, $S_{1:n,1}$ can be identified up to a scale. We can write

$$\widetilde{S}_{2:n,1} \equiv \mathbb{E}\left[z_t \mathcal{U}_{2:n,t}\right] / \mathbb{E}\left[z_t \mathcal{U}_{1,t}\right],$$

[4]The size of the window is motivated by the fact that the exact time of the announcement is unavailable in the initial part of our sample.

[5]A high-frequency identification of monetary policy shocks faces a similar challenge (Nakamura and Steinsson, 2018a). In Section 5.4, we present the results based on the specification where the oil supply surprise series is used directly as a shock.

and the whole vector $S_{1:n,1}$ equals $\left(x, x\widetilde{S}'_{2:n,1}\right)'$, where x is a normalization constant which can take any value. We normalize x so that the shock $\mathcal{E}_{1,t}$ increases oil price on impact by 10 percent. Having computed $S_{1:n,1}$, we obtain the oil supply news shock $\widehat{\mathcal{E}}_{1,t}$ from the estimated reduced-form error terms $\widehat{\mathcal{U}}_t$. We will refer to these shocks as *oil supply news shocks* or *oil shocks*, for brevity. Importantly, to be able to estimate $\widehat{\mathcal{E}}_{1,t}$, the structural shock $\mathcal{E}_{1,t}$ has to be invertible, which would allow us to recover these shocks from current and past values of the variables entering the oil market VAR (Nakamura and Steinsson, 2018b; Stock and Watson, 2018).[6]

The state-dependent effects of oil supply news shocks. Finally, to evaluate the impact of oil supply news shocks on a variable of interest y_t during and outside of the ZLB periods, we estimate a series of regressions at each horizon h from $h = 0$ to H of the form

$$\begin{aligned} y_{t+h} - y_{t-1} =& \mathbb{I}_{t-1} \cdot \left[\alpha^y_{A,h} + \beta^y_{A,h}\widehat{\mathcal{E}}_{1,t} + \psi^y_{A,h}(L)x_{t-1}\right] \\ &+ (1 - \mathbb{I}_{t-1}) \cdot \left[\alpha^y_{B,h} + \beta^y_{B,h}\widehat{\mathcal{E}}_{1,t} + \psi^y_{B,h}(L)x_{t-1}\right] + \epsilon^y_{t,h}, \end{aligned} \tag{1}$$

where $\mathbb{I}_{t-1}$ is the indicator variable that takes the value of one if the economy is in the ZLB in period $t-1$, and zero otherwise, and the subscripts A and B indicate the ZLB and the normal periods, respectively, and the controls $\psi^y_{A,h}(L)x_{t-1}$ are *lags* of the left-hand side variable of interest, the unemployment rate, and the measure of oil supply news shock $\widehat{\mathcal{E}}_{1,t}$.[7] Then, $\beta^y_{A,h}$ and $\beta^y_{B,h}$ are the responses of y_t in horizon h to an oil supply news shocks in the ZLB and normal periods, respectively.

We use 12 lags with monthly data and 4 lags with quarterly data.[8] In some cases, there are additional controls, which we specify separately. In principle, the estimation does not require the addition of these controls when the data are infinite. In practice, however, our sample is far from infinite. In this case, it is helpful to increase the signal-to-noise ratio of the oil supply news shocks by removing the variation predicted by past values of the variable of interest and the variation attributed to past oil supply news shocks. Additionally, to purge any potential remaining endogeneity in the oil shocks, as the oil market VAR does not include country-specific variables, we include controls for the country-specific unemployment rate. Although a larger set of controls could be

[6]In Section 3.1, we present evidence supporting the assumption that the shock $\mathcal{E}_{1,t}$ is invertible.

[7]The variables of interest, such as industrial production, real GDP, real aggregate consumption, and real aggregate investment are in logs, while nominal interest rates, the unemployment rate, the inflation rate, inflation expectations, and trade balance to GDP ratio are in levels, that is, we do not apply logs.

[8]The information criteria suggest a smaller number of lags, but the literature, such as Montiel-Olea and Plagborg-Moller (2020), recommends including a larger number of lags in the controls.

included, we use this more limited set in the baseline given the sample size.[9]

Because we include several lags of the dependent variables, our specification corresponds to what Montiel-Olea and Plagborg-Moller (2020) refer to as a lag-augmented local projection. Accordingly, we report heteroskedasticity-robust standard errors (henceforth, "robust" standard errors). We also compute Newey–West standard errors to account for potential serial correlation. The results are similar to the robust standard errors and do not alter our conclusions.

The standard errors on the coefficients $\beta^{y}_{A,h}$ and $\beta^{y}_{B,h}$ may require adjustment because the oil supply news shock $\widehat{\mathcal{E}}_{1,t}$ is a generated regressor. However, Pagan (1984) shows that such an adjustment is unnecessary when testing whether the coefficients on generated regressors are zero, since in that case the additional uncertainty from generating them does not affect the asymptotic distribution of the test statistic.[10] Consistent with this result, we report heteroskedasticity-robust standard errors. As an additional robustness check, we also compute bootstrap standard errors for the local projections that account for the uncertainty in the recovered oil supply news shocks.[11]

The state-dependent specification allows us to test the null hypothesis that the responses are identical in the two periods, that is, $\beta^{y}_{A,h} = \beta^{y}_{B,h}$. More specifically, we rewrite specification (1) in the form:

$$y_{t+h} - y_{t-1} = \alpha^{y}_{C,h} + \beta^{y}_{C,h}\widehat{\mathcal{E}}_{1,t} + \psi^{y}_{C,h}(L)x_{t-1} + \mathbb{I}_{t-1} \cdot [\alpha^{y}_{D,h} + \beta^{y}_{D,h}\widehat{\mathcal{E}}_{1,t} + \psi^{y}_{D,h}(L)x_{t-1}] + \epsilon^{y}_{t+h},$$

and test the null hypothesis that $\beta^{y}_{D,h} = 0$. We report the estimates of $\beta^{y}_{D,h}$ and the corresponding robust standard errors.

In our analysis, we assume that the regimes (ZLB vs. normal) are independent of the shocks. In principle, however, the occurrence of the zero lower bound can depend on shocks. In the Japanese context, a single large shock (e.g., the collapse of the Japanese housing and stock market bubble in the early 1990s) could have brought the natural real interest rate down substantially for a prolonged period forcing the Bank of Japan to keep its policy rate at zero. The fact that Japan's normal and ZLB periods are highly persistent, with only two switches over the sample period, suggests that the ZLB and normal regimes are unlikely to be driven by the frequent, small shocks we analyze.

[9] In Section 5.4, we show that alternative set of controls (e.g., adding more country-specific variables, such as lags of the interest rate and inflation rate, adding all four variables used in the oil VAR, or removing all controls except for the oil supply news shocks) and lag length structure have some quantitative impact on our results but do not change our qualitative conclusions.

[10] Coibion and Gorodnichenko (2012) argue that the generated regressor problem is likely to be small when the generated regressor is a residual rather than a fitted value from a regression.

[11] Another potential concern for our approach is that in finite samples, the extracted shock from the proxy VAR is potentially correlated with the other structural shocks, and thus, can bias the results. We argue that the bias is small through a series of Monte Carlo simulation exercises in Appendix E.

Finally, we note that our two-step approach assumes that ZLB episodes in Japan and other countries do not coincide with a structural break in the oil supply news shocks identified from the oil VAR. We adopt this approach for three reasons. First, changes in Japanese monetary policy are unlikely to have a significant effect on the global oil market, as Japanese oil consumption represents only a small share of world oil production.[12] Second, the two-step approach allows us to use a common oil supply shock across specifications and countries. Finally, robustness exercises reported in Appendix Section D.2, where we estimate a state-dependent VAR to identify oil supply news shocks, yield broadly similar results.

2.2 Data

We estimate the four-variable oil market VAR using the log levels of the following monthly variables. The real oil price is the WTI spot oil price deflated by the U.S. consumer price index. For world industrial production, we use the Baumeister and Hamilton (2019) index, which covers the OECD countries plus six major economies (Brazil, China, India, Indonesia, the Russian Federation, and South Africa). Together, these countries jointly represent 75 percent of world's GDP. World oil inventories are OECD petroleum stocks from Kilian and Murphy (2014). World oil production comes from the U.S. Energy Information Administration.[13] The detailed description of the data sources is in Appendix A. Because of the availability of oil futures data, the oil supply surprise series covers a shorter period between 1983:4 and 2019:12, but the external instrument VAR uses data between 1974:1 and 2019:12 to estimate oil supply news shocks. The VAR includes a constant term and has 12 lags.

We use monthly and quarterly macro data for three advanced countries—Japan, the United Kingdom, and the United States—that have experienced a sufficiently long period where the nominal interest rate is at the ZLB. The sample period is between 1975:1 and 2019:12, which corresponds to the sample of the extracted oil supply news shock $\widehat{\mathcal{E}}_{1,t}$. We exclude the COVID pandemic by dropping data after 2019:12.

[12] Japanese oil consumption, in million barrels per day, rose from 4.9 in 1974 to a peak of 5.6 in 1996, before falling to 3.8 in 2019, according to the International Energy Statistics database of the U.S. Energy Information Administration (EIA). Over the same period, world oil production increased from 55.7 in 1974 to 71.4 in 1996 and to 100.6 in 2019.

[13] Unlike Kanzig (2021), we do not directly use the U.S. industrial production and inflation rate in the oil market VAR. This is because the United States can exhibit a state-dependent response of these two variables.

2.3 ZLB Definition

We define the ZLB periods as those when the short-term nominal interest rate—the standard instrument of conventional monetary policy—was close enough to zero that it did not respond substantially to macroeconomic developments. In particular, the ZLB periods for Japan are October 1995 to June 2006 and January 2009 to December 2019. This definition is similar to Wieland (2019), who follows the previous literature on the timing of the ZLB spell in Japan. This period of the ZLB in Japan coincides with a 0.5 percent cutoff before 1998 and 0.25 percent after that.[14] We define the ZLB periods in the United Kingdom and the United States as the months during which each country's short-term nominal interest rate is at or below 0.5 percent. The ZLB period in the United States was between 2008:11 and 2016:11.[15] According to our definition, the United Kingdom experienced two ZLB episodes: 2009:4–2010:9 and 2012:2–2018:7.[16] Finally, as mentioned in the introduction, we rely only on the stance of the conventional monetary policy to define ZLB episodes. We estimate the responses of longer-term and shadow rates to analyze if these alternative policies managed to overcome the ZLB constraint.

3 Oil Supply News Shocks and their Effects in Japan

This section presents the estimates of the oil supply news shocks and their effects on the Japanese economy when the nominal interest rate is at the zero lower bound and outside of it. Specifically, we present the results for real monthly variables, including industrial production and the unemployment rate, then quarterly variables, such as real GDP, real private consumption, and real investment per capita. Finally, we examine the New Keynesian mechanism by estimating the responses of the nominal interest rate and inflation expectations.

3.1 Oil Supply News Shocks

The oil supply surprise series, plotted on the left panel of Figure 1, illustrates that OPEC announcements had a sizable impact on the oil futures prices in both directions throughout our sample. This series is also a relevant instrument for the external instrument VAR:

[14] We also estimate the effects of oil supply news shocks on Japanese economy in the ZLB period when the Japanese nominal interest rate is below the 0.5 percent cutoff. The results are similar.

[15] This 0.5 percent cutoff results in a ZLB period similar to Ramey and Zubairy (2018). Lowering the cutoff to 0.25 percent shrinks the ZLB period to 2008:12-2015:12, however, the estimation results remain qualitatively similar.

[16] The U.K. interest rate rose just above 0.5 percent during 2010:10-2012:1. Our results do not materially change if we define this period as ZLB as well.

Figure 1: The oil supply surprise series and oil supply news shock

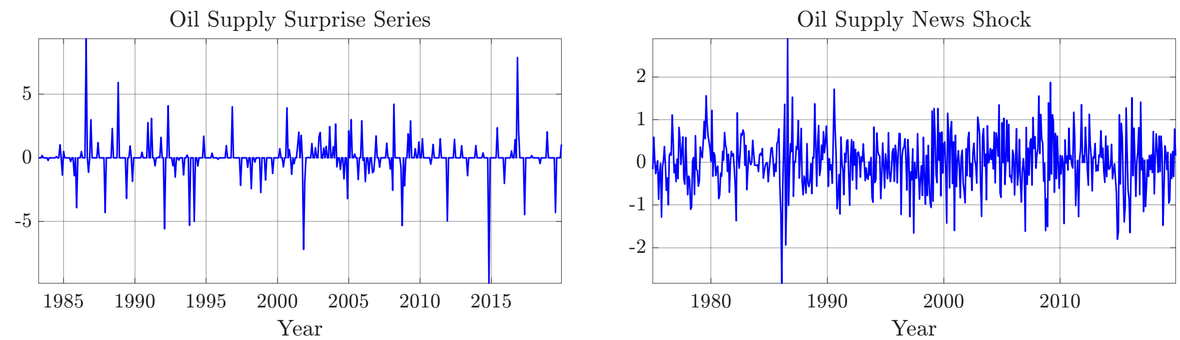

Notes: The left panel plots the oil supply surprise series from 1983:4 to 2019:12. The right panel presents the oil supply news shock $\widehat{\mathcal{E}}_{1,t}$ from 1975:1 to 2019:12, estimated using a four-variable oil VAR instrumented by the oil supply surprise series. The details are in Section 2.1.

The first-stage regression of the oil price residuals from the VAR on the oil news surprise series, yields an F-statistic of 19.9, well over a recommended threshold of 10 (Montiel Olea, Stock and Watson, 2021). The instrument explains 3.6 percent of the variation in the oil price residual. Furthermore, we find evidence supporting the assumption of the invertibility of the oil supply news shocks. Specifically, we test whether the oil supply surprise series jointly Granger causes the four variables in the oil VAR (Giannone and Reichlin, 2006; Forni and Gambetti, 2014; Plagborg-Møller and Wolf, 2022; Forni, Gambetti and Ricco, 2024). The *p*-value for the Wald test of no Granger causality null hypothesis is 0.90.[17]

The right panel of Figure 1 plots the estimate of the oil supply news shocks $\widehat{\mathcal{E}}_{1,t}$ from 1975:1 to 2019:12. There is no noticeable difference between the different parts of the sample in terms of the size and frequency of the shocks. There was a relatively large negative shock in February 1986, followed by a positive shock of similar size in August 1986.

Before turning to the analysis of the Japanese economy, in Figure 2, we plot the responses of the real oil prices to the oil supply news shock. We normalize the responses such that the oil supply news shocks increase the real oil price on impact by 10 percent in both the normal and the ZLB periods. We use the ZLB definition for Japan and generate these responses by estimating equation (1) with the real oil price as the variable of interest. There is some evidence that the shock in the ZLB periods leads to a more persistent increase in the real oil price than that in the normal periods. The differences between the real oil price responses in the ZLB periods and those in the normal periods are statistically

[17]This result also holds when we test for Granger causality using first-differences, a shorter sample from 1983:4 to 2017:12, and a 6-variable VAR as in Kanzig (2021) instead of the 4-variable VAR as in our baseline specification.

Figure 2: Impulse responses of real oil prices

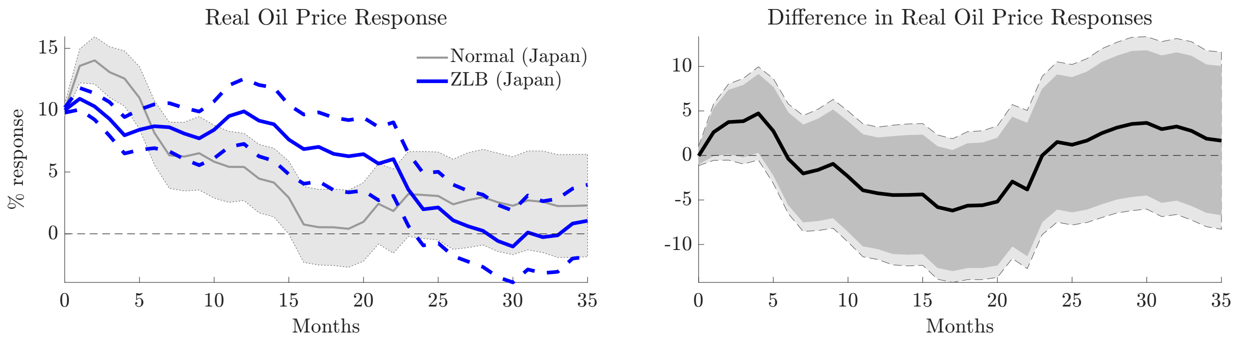

Notes: The left panel plots the impulse response functions and one-standard-deviation robust confidence bands for the real oil price in the ZLB (thick blue line) and in the normal (thin gray line) periods. The Japanese ZLB definition is used to break the full sample into the ZLB and normal periods. The responses are normalized to increase the oil price by 10 percent in both periods on impact. The right panel plots the differences in the responses in the ZLB and normal periods along with the 90 percent (dark gray area) and 95 percent (light gray area) confidence bands.

different at the 10 percent level in the first few months after the shock, as seen from the right panel of Figure 2. However, the responses in the two subperiods are qualitatively similar.[18]

3.2 Real Economy

The top left panel of Figure 3 shows the responses of industrial production (a volume measure) at horizons up to 36 months after an oil supply news shock that increases real oil prices, along with the one-standard-deviation confidence bands in both the normal and the ZLB periods. To compare the effects of oil supply news shocks across normal and ZLB periods, we normalize the estimated responses of industrial production (and subsequently other variables) so that the shocks raise the real oil price by 10 percent on impact in both periods. In normal times, industrial production slowly falls, declining over 1 percent one year after the shock. This decline reverts around two years after the shock. In the ZLB periods, the response of industrial production increases to nearly 1 percent one year after the shock, and it is statistically different from zero at the 5 percent level at horizons between 10 and 13 months. The difference between the two responses of industrial production in the ZLB and the normal periods are in the bottom left panel of Figure 3. This difference is statistically significant at the 10 percent level at horizons between 10 and 17 months and at the 5 percent level at horizons between 12 and 13 months after the shock.[19]

[18]A higher persistence of the oil supply news shocks in the ZLB period implies a larger disruption in aggregate economic activity under the conventional supply-side view. Our estimation below presents evidence to the contrary.

[19]To compare our findings with the existing literature, which typically does not distinguish between ZLB and normal periods, we estimate linear local projections that abstract from regime differences. Appendix

Figure 3: Japanese industrial production and unemployment rate impulse responses

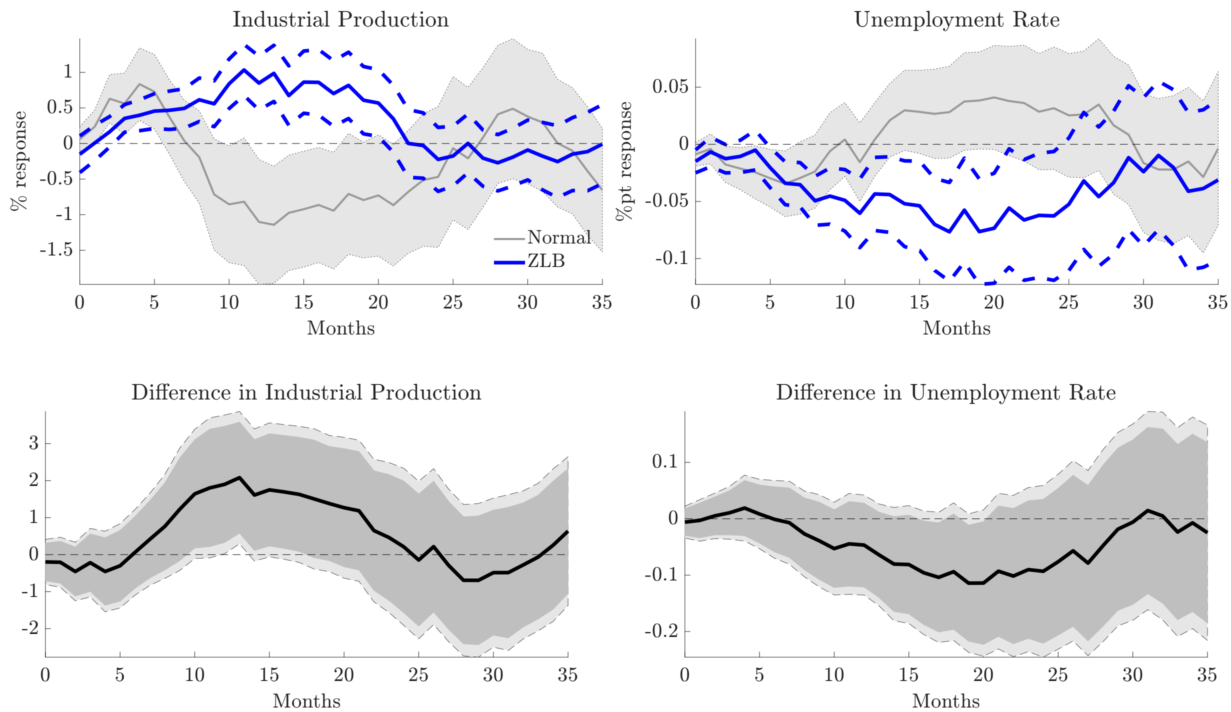

Notes: Each figure in the top panel plots the impulse response functions and one-standard-deviation robust confidence bands of industrial production and the unemployment rate in Japan in the ZLB (thick blue line) and in the normal (thin gray line) periods. The lower panels present the differences in the responses of industrial production and the unemployment rate between the normal and ZLB periods, with the 95 percent (light gray area) and 90 percent (dark gray area) confidence bands.

Oil supply news shocks also lead to different labor market dynamics in the ZLB and the normal periods. The top right panel of Figure 3 plots the responses of the unemployment rate to an oil supply news shock in both the normal and the ZLB periods. Initially, the responses of the unemployment rate to an oil supply news shock in both periods are negative and not significantly different from zero. Then, the unemployment rate increases up to 0.03 percentage points during normal times at the 14-month horizon, while it decreases 0.06 percentage points in the ZLB periods a year after the shock. This decrease is statistically significant at the 10 percent significance level at horizons between 6 and 11 months, as well as horizons 16, 17, 19 and 20 months after the shock. We test the differences in the responses in the normal and ZLB periods, and the differences are statistically significant at the 10 percent significance level for some horizons, such as 16, 17, and 19 months after the shock, as depicted in the bottom-right panel of Figure 3. As

Figure B.1 reports the results. For the United States, the responses are consistent with Kanzig (2021). For Japan, industrial production increases at short horizons and the unemployment rate declines, though neither effect is statistically significant at longer horizons. These results differ from Peersman and Van Robays (2012), who find a contraction in Japanese GDP following oil supply shocks. The discrepancy may reflect our longer ZLB sample for Japan.

shown in Appendix Figure B.2, the results are similar if we use Newey-West standard errors. Also, the results are similar if we use bootstrapped standard errors to account for the uncertainty in the oil supply news shocks as in Appendix Figure B.3.

Figure 4: Japanese quarterly variable impulse responses

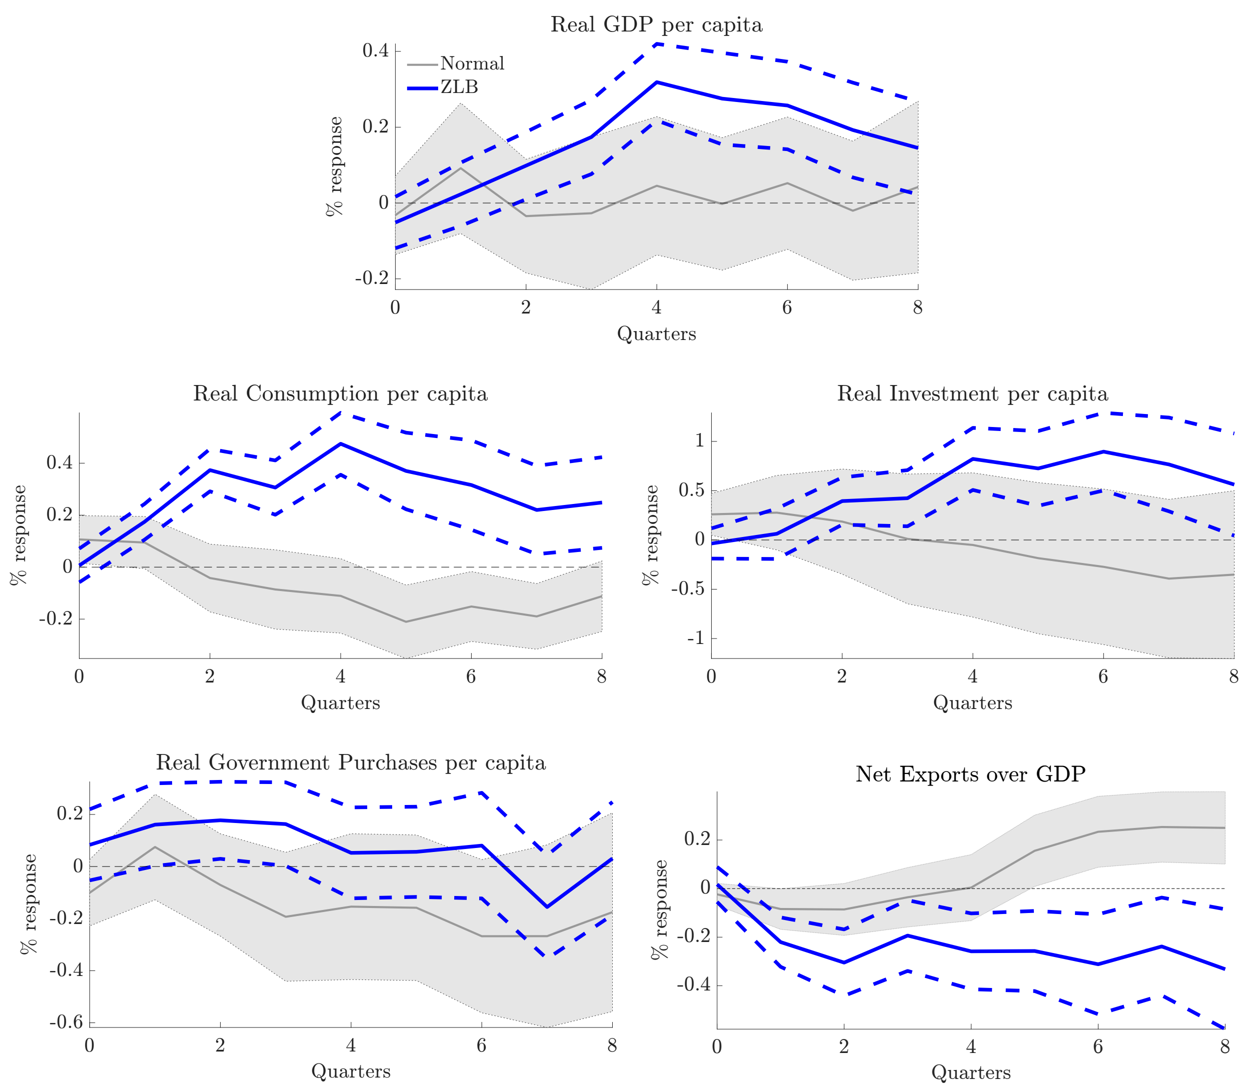

Notes: Each figure plots the impulse response functions and one-standard-deviation robust confidence bands of real GDP, real consumption, real investment, real government purchases (all per capita), and nominal trade balance over nominal GDP in Japan to an oil supply news shock in the ZLB (thick blue line) and the normal (thin gray line) periods. The estimated specification is in equation (1). The sample is from 1975:Q1 to 2019:Q4. The differences in the impulse responses across the ZLB and the normal periods are plotted in Figure C.1.

Long ZLB spells in Japan allow us to perform estimation using quarterly data. Figure 4 plots the impulse responses of real per capita GDP, private consumption, and private investment, government purchases, and net exports in Japan to an oil supply news shock. We sum all monthly shocks within a quarter to obtain quarterly oil supply news shocks that we use in equation (1). The controls are identical to those described in Section 2.1, and the number of lags at quarterly frequency is four. Consistent with the monthly real variables, we observe that real GDP, consumption, and investment per capita responses increase in the ZLB period, and they are larger than in the normal periods. In particular, GDP virtually does not change on impact but increases above 0.3 percent after one year at the ZLB. It remains roughly constant in the normal periods. The differences are less precisely estimated, likely due to fewer observations at the quarterly frequency. Appendix Figure C.1 plots the difference in responses in the ZLB and normal periods.[20]

Interestingly, the point estimate for government purchase responses are also higher in the ZLB periods than outside of it. However, the standard errors are substantial. Finally, the trade balance exhibits a differential response across the two states as well.[21] The difference between the responses in the ZLB and normal periods is negative a year after the shock. This is consistent with the fact that aggregate consumption and investment increase at the ZLB relative to the normal period, which can drive an increase in demand for imports.[22]

3.3 Interest Rates, Inflation, and Inflation Expectations

One possible explanation for the observed differences in the responses of real macro variables in Japan to oil supply news shocks is that the monetary policy stance during these periods differs. To inspect this mechanism, we estimate the responses of nominal interest rates, inflation, and inflation expectations to an oil supply news shock in the same specification described in Section 2.1. When estimating the responses of nominal interest rates, we include 12 lags of CPI inflation in the control set, as inflation is typically an important factor that central banks consider when setting monetary policy.[23]

The top left panel of Figure 5 plots the impulse response of the short-term nominal interest rate to an oil supply news shock. In the ZLB periods, the response of the nominal interest rate stays near zero and is precisely estimated, as one would expect from the ZLB periods. In normal times, the short-term nominal interest rate does not change on impact, but then it gradually increases by 0.5 percentage points one year after the shock. Appendix Figure C.2 plots the difference in the responses of the short-term nominal interest rate, as well as of the other variables discussed in this subsection. The short-term nominal interest rate difference is significant at conventional levels 2 months and up to

[20] From now on, we present the difference in responses in the ZLB and normal periods in Appendix C.

[21] To estimate the response of the real trade balance, we normalize the change in the nominal trade balance by lagged nominal GDP. That is, we use $(NX_{t+h} - NX_{t-1})/Y_{t-1}$ as the left-hand variable.

[22] Section 5.2 presents the responses of the real exchange rate.

[23] In a robustness check, we remove 12 lags of the CPI inflation rate from controls. The results are similar.

14 months after the shock. In addition, the response in the normal periods is significantly different from zero at the 5 percent significance level from horizons of 2 to 14 months.

As mentioned earlier, the longer-term rates can vary even when the short-term nominal rate is constant at zero. The top right panel of Figure 5 presents the response of the 5-year nominal government bonds yield. This longer rate behaves qualitatively similar to the short rate. In the normal periods, this longer rate rises less than the short rate. As shown in Appendix Figure C.2, the differences are significant up to 4 months after the shock.

Figure 5: Interest rates, inflation, and inflation expectations responses in Japan

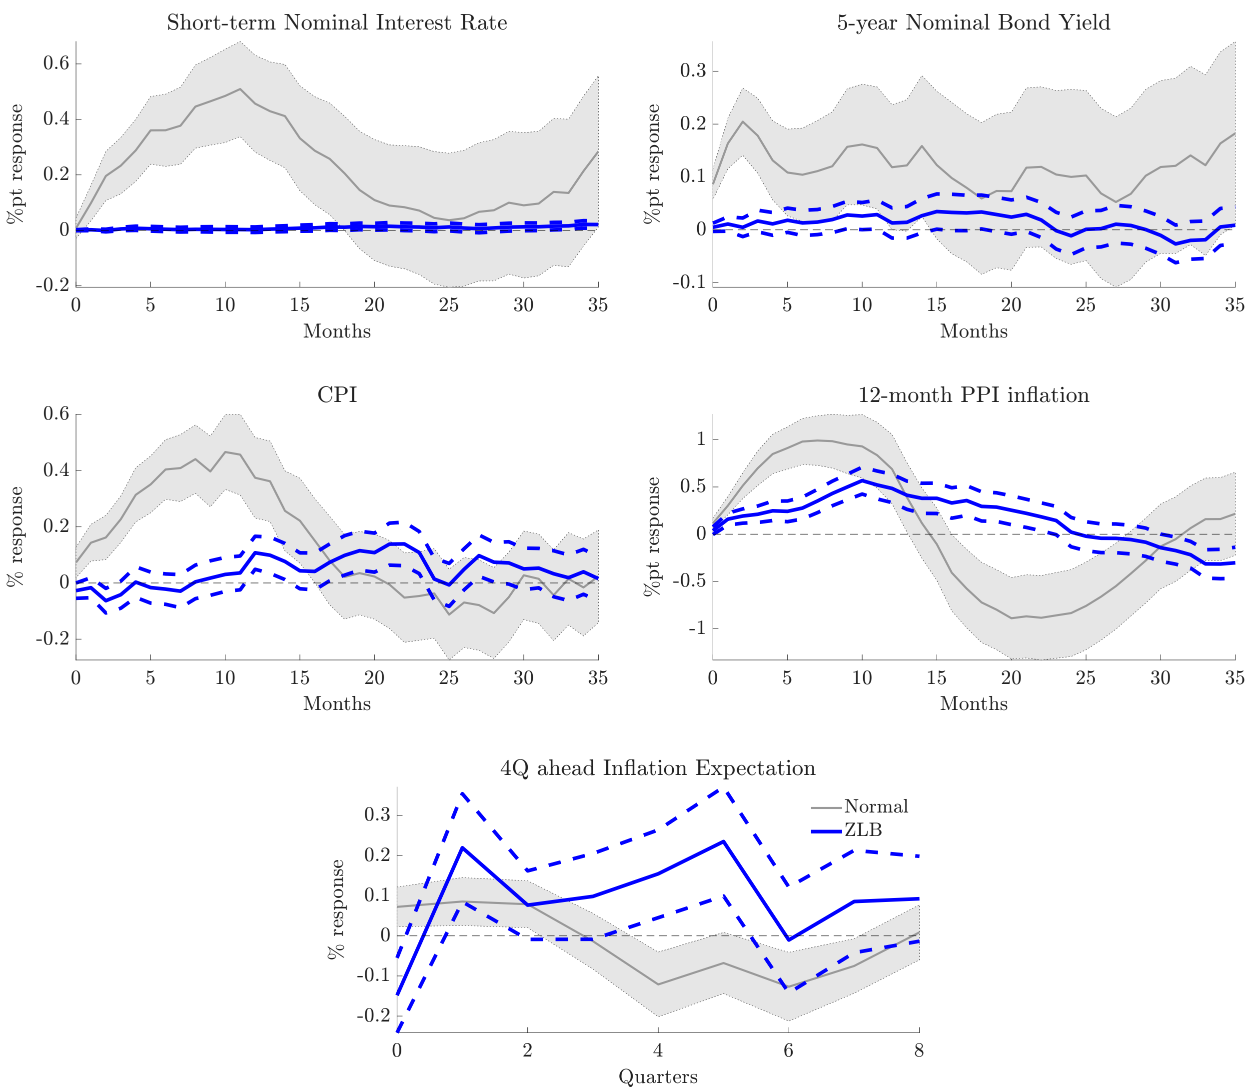

Notes: The figures plot the impulse response functions and one-standard-deviation robust confidence bands of the short-term nominal interest rate, 5-year nominal interest rate, CPI inflation rate, PPI inflation rate, and CPI inflation expectations in Japan in the ZLB (thick blue line) and in the normal (thin gray line) periods. The responses are to an oil supply news shock. The differences in responses across ZLB and the normal periods are in Figure C.1.

The reaction of the CPI inflation rate to the oil shock is plotted in the middle left panel of Figure 5. The inflation rate, measured as the 12-month percentage change in the consumer price index, increases in the normal periods about 0.4 percentage points after 8 months from the shock and reverts after 17 months. In the ZLB periods, the inflation rate does not react at the beginning and shows a moderate increase of 0.14 percentage points 22 months after the shock. The middle right panel of the same figure presents the reaction of the producer price index (PPI) inflation rate. The behavior is qualitatively similar to the behavior of the CPI inflation rate. However, the response of the PPI inflation rate during the ZLB periods is more pronounced than the response of the CPI inflation rate.[24]

The bottom panel of Figure 5 demonstrates the impulse responses of the one-year-ahead inflation expectations at quarterly frequency. We use the Japanese Center for Economic Research (JCER) CPI inflation forecast to proxy for inflation expectations.[25] Inflation expectations increase to 0.24 percentage points in the ZLB one and five quarters after the shock, with the increase at the five-quarter horizon statistically significant at the 10 percent level. The inflation expectations response in the normal periods is above zero on impact and becomes negative after three quarters. The differences of the responses between the ZLB and normal periods are significant at the 4- and 5-quarter horizons. The responses of the short-term rates and expected inflation suggest that the ex ante real ante interest rate falls during the ZLB periods but increases in the normal periods.

4 Oil Supply News Shocks Effects in the U.K. and U.S.

Since Japan is not the only country that has experienced ZLB episodes, we now investigate the effects of oil supply news shocks in the United Kingdom and the United States, both of which faced a near-zero interest rate in the recent past. The results provide additional evidence that oil supply news shocks have differential impacts during and outside of the ZLB periods, suggesting that the results using Japanese data are not specific only to Japan or to the period when Japan was in a liquidity trap.

[24]The diminished reaction of price indices to oil shocks during the ZLB, relative to the normal period, aligns with Watanabe and Watanabe (2017) finding that goods level price inflation in Japan has shown reduced variability since the mid-1990s.

[25]Alternative inflation expectations datasets are too short to estimate the effects of oil shocks on inflation expectations in the normal periods in Japan. For example, the Consensus Economics Forecast data for Japan started only in 1989.

4.1 Real Economy

Panel (a) of Figure 6 plots the impulse responses of industrial production to an oil supply news shock, and the difference is in Appendix Figure C.3. Both in the United States and United Kingdom, the difference between the industrial production responses in the ZLB and normal periods becomes positive and statistically significant a few months after the shock. Moreover, in these countries, the industrial production responses become positive and statistically different from zero in the ZLB period several months after the shock. For example, in the United States, industrial production increases by nearly 0.8 percent 16 months after the shock. In contrast, industrial production falls in the normal period: the responses are negative and statistically different from zero at several horizons. These results are qualitatively consistent but quantitatively and statistically more pronounced than those for Japan.[26]

Figure 6: Industrial production and unemployment rate responses in the U.S. and U.K.

(a) Industrial production

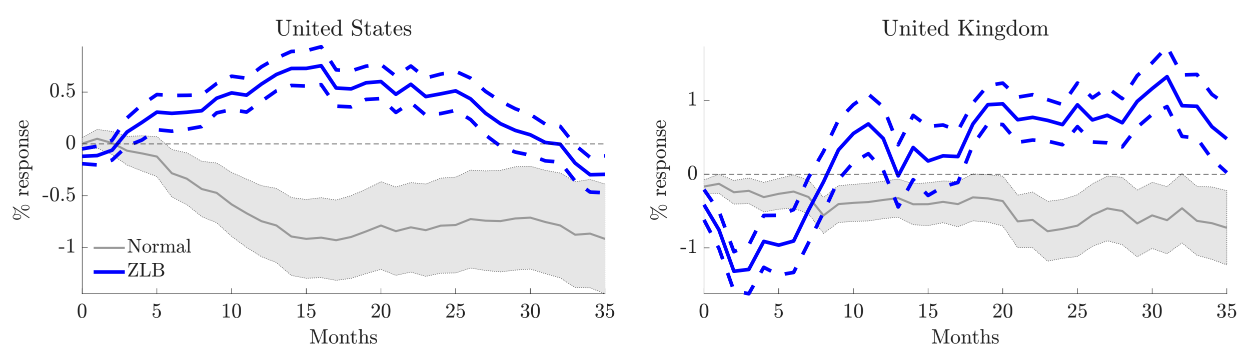

(b) Unemployment rate

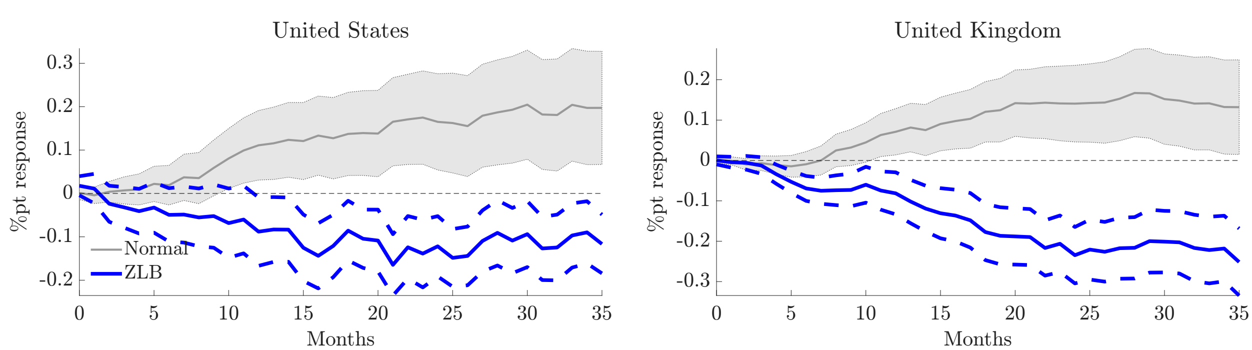

Notes: Each figure plots the impulse response functions and one-standard-deviation robust confidence bands. Panel (a) shows industrial production responses in the United Kingdom and the United States in the ZLB (thick blue line) and in the normal (thin gray line) periods. Panel (b) plots unemployment rate responses in the United Kingdom and the Unites States in the ZLB (thick blue line) and in the normal (thin gray line) periods. Figure C.3 plots the differences in the responses of these variables in the ZLB and outside of it.

[26] Focusing on manufacturing industrial production, which is discussed in detail in Section 5.2 and presented in Appendix Figure B.8, the U.K.'s ZLB response does not show an initial decline.

The labor market responses, which we plot in Panel (b) of Figure 6, mirror the responses of industrial production. Specifically, the unemployment rate in the United States and the United Kingdom falls after a negative oil supply news shock that increases the oil price in the ZLB period a few months after the shock. These ZLB responses are statistically different from zero and from those in the normal periods at several horizons longer than 12 months.

4.2 Interest Rates, Inflation, and Inflation Expectations

To examine the monetary policy channel in the United States and the United Kingdom, we estimate the effects of oil supply news shocks on nominal variables. Panel (a) of Figure 7 shows that short-term rates remain at zero during the ZLB but rise in normal times. The estimates are noisier than for Japan: the differences, plotted in Figure C.4, are not statistically significant in the United States and are significant only at 3–5 month horizons in the United Kingdom. Five-year yields respond in both regimes, with no significant differences across periods. Inflation responses are similar across regimes in the United States. In the United Kingdom, they are initially alike, but inflation falls markedly during the ZLB after one year, with differences significant at longer horizons.

We close this section by describing the behavior of inflation expectations. We focus on households' inflation expectations because this group of economic agents is responsible for consumption choices. In addition, recent literature, such as Weber *et al.* (2022) and Candia *et al.* (2023), documents that inflation expectations of firm managers are closer to those of households than of professional forecasters. In the bottom panel of Figure 7, we plot the responses of the median one-year-ahead inflation expectations in the United States using the Michigan Survey of Consumers data and in the United Kingdom using the Bank of England/Ipsos Inflation Attitudes Survey data. The former survey collects data monthly, while the latter is a quarterly survey. As in the case of regressions with quarterly Japanese data, we construct quarterly oil supply news shocks by adding up the shocks within the same quarter. In both countries, we observe an increase in inflation expectations on impact in both the ZLB and normal periods. In the United States, the response in the normal periods is larger than that in the ZLB periods: the reaction in the normal periods reaches 0.23 percentage points compared to 0.1 percentage point in the ZLB periods.[27] In contrast, the response in the ZLB periods in the United Kingdom is larger than the response in the normal periods: the ZLB largest responses is 0.3 percentage points, while it is 0.12 percentage point in the normal period. The difference between the

[27] Figure B.4 shows the response of one-year ahead inflation expectations by professional forecasters using the data from Consensus Economics. There is also evidence of a stronger reaction of the expectations in the normal period, however, the difference is smaller than in the case of the U.S. consumers.

Figure 7: Interest rates and inflation in the United States and the United Kingdom

(a) Short-term nominal interest rate

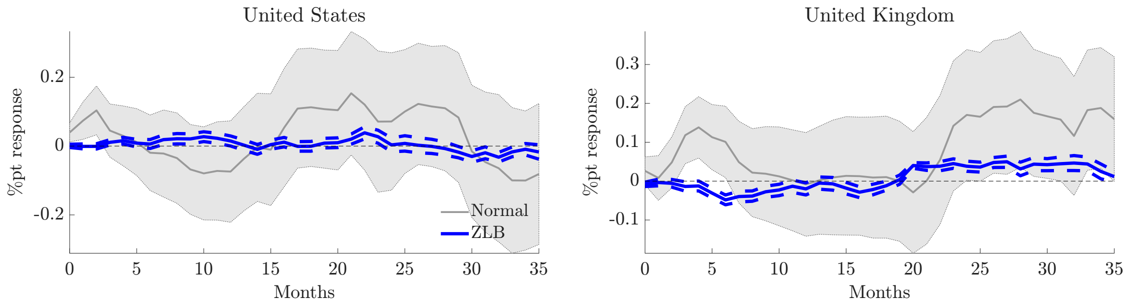

(b) 5-year nominal government bond yield

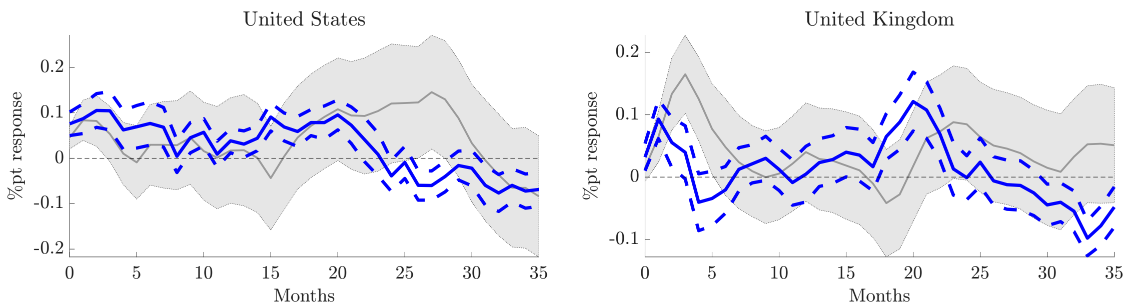

(c) CPI inflation rate

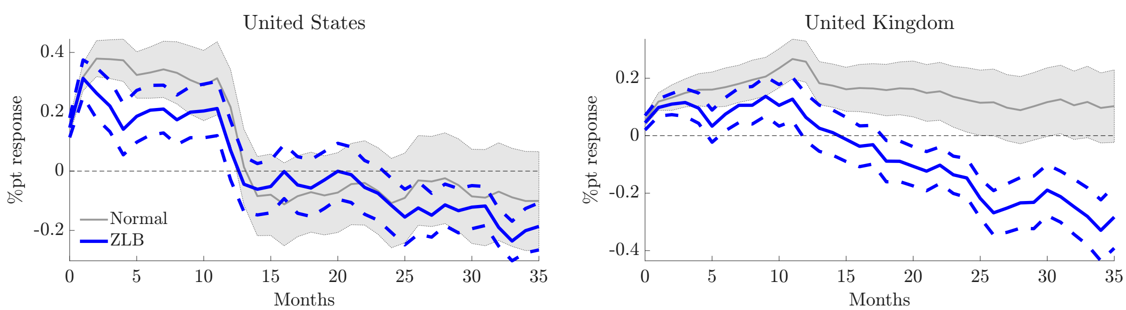

(d) Median household inflation expectations

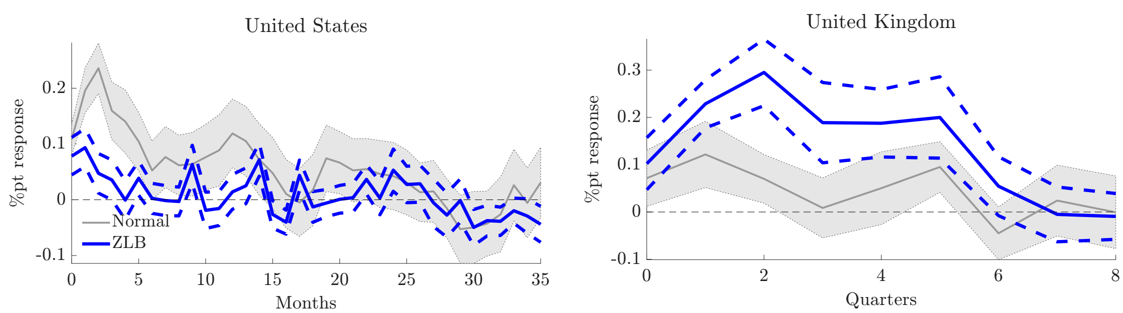

Notes: Each figure plots the impluse response functions and one-standard-deviation robust confidence bands the central bank interest rate, the 5-year nominal yield on government bonds, the CPI inflation rate, and household inflation expectations responses to the oil supply shock in the United States and United Kingdom in the ZLB (thick blue line) and in the normal (thin gray line) periods. The U.S. inflation expectations are 12-month-ahead inflation expectations from the Michigan Survey of Consumers measured at monthly frequency. The U.K. inflation expectations are 4-quarter-ahead inflation expectations from the Bank of England/Ipsos Inflation Attitudes Survey measured at quarterly frequency. Figure C.4 plots the differences in the responses of these variables in the ZLB and outside of it.

two responses are statistically significant at the 2-quarter horizon.[28]

Overall these results parallel those in Japan in Section 3.3. Specifically, the reaction of the nominal interest rates and expected inflation indicate a differential reaction of the ex ante real interest rate in the ZLB and normal period.

5 Other Shocks, Variables, Countries, and Specifications

Our results for Japan align with those for the United States and the United Kingdom, suggesting that the less contractionary, or even expansionary, effects of oil supply news shocks at the ZLB are not unique to Japan. In this section, we extend the analysis by considering alternative oil shock measures, additional outcome variables, other economies (near and far from the ZLB), and several robustness checks, including collapsing the two-step procedure into a single state-dependent proxy VAR, varying the sample, excluding key episodes, and modifying controls.

5.1 Other Oil Shocks

We begin by analyzing the effects of oil supply shocks identified using alternative methods to demonstrate the importance of shock identification for our findings. This analysis clarifies why our baseline results differ from those of Wieland (2019), who finds no differences in the responses of Japanese industrial production and the unemployment rate to oil supply shocks.

We focus on two prominent identification strategies for oil supply shocks. The first one is the oil supply shocks identified in Kilian (2009), also used in Wieland (2019). We will refer to these oil supply shocks as *the Kilian shocks*. The second set of oil supply shocks is from Baumeister and Hamilton (2019), which we will call *the BH shocks*. To identify oil supply shocks in a structural VAR, Kilian (2009) assumes that there is no short-run reaction of the oil supply to changes in the real oil price.[29] In contrast, Baumeister and Hamilton (2019) do not restrict the oil supply elasticity to be small and estimate it to be a much larger number of 0.15. Baumeister and Hamilton (2019)'s identification strategy of oil supply shocks recovers a series of oil supply shocks that is substantially more important in accounting for the historical oil price movements than Kilian (2009)'s

[28] We also estimate responses of U.S. inflation expectations using the Survey of Professional Forecasters and find the patterns similar to the Consensus Forecast. The Consensus data for Japan and the United Kingdom begin in 1989 and 2004, respectively. For Japan, we find a positive response of CPI inflation expectations during the ZLB, consistent with our baseline results. For the United Kingdom, using data from 2004–2019, we also estimate a positive response over the full sample.

[29] Kilian and Murphy (2012) later modified this assumption with an alternative that the short-run price elasticity of the oil supply is small or, more specifically, less than 0.0258.

identified oil supply shocks. We use these two oil supply shocks in the same way as in our baseline specification described in Section 2.1.

We find that the responses of Japanese industrial production and unemployment rate to the Kilian shocks in the normal and ZLB periods are not different, as plotted in the top row of Appendix Figure B.5. This result is consistent with Wieland (2019).[30] In contrast, the responses of industrial production and the unemployment rate to the BH shocks, shown in the bottom row of Appendix Figure B.5, are qualitatively similar to our baseline estimation using oil supply news shocks. In particular, industrial production increases in the ZLB periods but decreases in the normal periods, and the unemployment rate decreases in the ZLB periods but increases in the normal periods. Differences in unemployment responses to BH shocks between the ZLB and normal periods are significant at the 10 percent level at horizons of 1 and 2 months. Industrial production responses differ significantly at the 10 percent level at horizons of 20 and 21 months after the shock (see Figure C.5).

What may explain the different effects of these shocks? As shown in the top-right panel of Figure B.6, the Kilian shocks generate mostly statistically insignificant responses of the real oil price. Moreover, Kilian (2009) finds that oil supply shocks are not an important driver of historical oil price fluctuations in his shock decomposition exercise (see the top panel of Figure 4). In addition, the top-middle panel of Figure 3 in Kilian (2009) shows that oil supply shocks do not affect U.S. industrial production in an economically or statistically significant way. This could account for the absence of differential real responses between the ZLB and non-ZLB periods in Japan, as shown in Figure B.5. At the same time, the response of the real oil price to the BH shock (the bottom right panel of Figure B.6) is more significant than that to the Kilian shock, but still not as precise as the response of the real oil price to the oil supply news shocks identified in Kanzig (2021) (the top left panel of Figure B.6). Baumeister and Hamilton (2019) show that their oil supply shocks contribute meaningfully to historical oil price dynamics (their Figure 9) and affect world industrial production (their Figure 8).

We close this section by presenting the responses of the Japanese economy to another oil shock estimated in Kilian (2009). This shock is ordered last in a triangular identification of his three-variable VAR, making it effectively a residual to the oil supply and global aggregate demand shocks, which come first and second. Kilian (2009) refers to this shock as the oil-market specific demand shock. When normalized to increase the oil price, this

[30]Note that our responses are not identical to the responses in Figure 6 (panels B and C) in Wieland (2019). There are three important differences between our and his estimation. First, we compute the Kilian shocks using the 1973:2-2019:12 sample instead of the 1973:2-2015:9 sample in Wieland (2019). Second, following our analysis in Section 3, we start our sample in 1975:1 instead of starting in 1986:1 as in Wieland (2019). Finally, we also normalize the responses so that they correspond to a 10% increase in the oil price in both the normal and the ZLB periods.

shock reduces U.S. GDP and, at the same time, increases the U.S. CPI, unlike his oil supply shock that does not change U.S. GDP and CPI. As shown in Figure B.7, where we use this shock to estimate the responses in Japan, the results are qualitatively similar to our baseline.

5.2 Other Variables

Production without oil. Our analysis used a measure of industrial production that includes mining, such as oil extraction. This raises a concern that oil price variations can mechanically affect the mining sector's output. This problem can be potentially more pronounced in the United States, where the importance of oil production has grown in the last two decades following the shale oil extraction boom. To address this concern, we present the responses of industrial production in the United States, the United Kingdom, and Japan only for the manufacturing sector in Panel (a) of Figure B.8 and for industrial production when we remove oil production for the United States in Panel (b). Our main findings remain qualitatively unchanged.

Stock market response. If oil supply news shocks lead to different responses in aggregate economic activity during and outside of the ZLB, it is logical to expect this pattern to be reflected in financial markets. To investigate this possibility, we estimate the response of the U.S. stock market represented by the S&P 500 index.[31] The left panel of Figure B.9 plots the reaction of the S&P 500 index estimated using specification (1) on monthly data. The results confirm our previous finding that the U.S. economy responds positively to the shocks at the ZLB and negatively in the normal periods. The positive response is significant at almost all horizons at the 5 percent significance level, while the negative response is only significant at this level at the 13-month horizon.

We complement the monthly analysis with daily stock market data, which allow us to use the oil supply surprise series directly to assess the role of monetary policy in oil shock propagation. We estimate a state-dependent specification that projects the log change in the S&P 500 index onto the oil supply surprise series, controlling for ten lags of the index. The sample includes all OPEC press conference announcements between January 1, 1983, and December 31, 2019. As shown in the right panel of Figure B.9 the daily results align with the monthly evidence and our earlier findings on aggregate activity. The differential stock market responses based on the oil supply surprise series reinforce

[31] We primarily look at the US stock market for two reasons. First, it is the biggest stock market in the world. Second, the recent literature, documenting the presence of the global financial cycle (Rey, 2015; Miranda-Agrippino and Rey, 2022), notes that financial market developments are correlated across countries.

our baseline results using the oil supply news shock.

Exchange rates. Another financial variable that can exhibit different responses during the ZLB and the normal periods is exchange rates. This can be due to the differential response of trade balance and monetary policy we estimated earlier.[32] Figure B.10 shows that the real exchange rate depreciates in the ZLB period and appreciates in the normal periods during the first two years after the shock, although the responses in the normal periods are not statistically significant at conventional levels. This reaction is in line with a relative tightening of monetary policy during normal times.

Unconventional monetary policy. During the ZLB periods, central banks often use unconventional monetary policies, such as forward guidance (announcements about future conventional monetary policy) or quantitative easing (purchases of long-term private and public financial assets). As a result, central banks can potentially respond to oil price shocks even during the ZLB periods. One way to assess this possibility is to look at longer-term nominal interest rates, which we did in Sections 3.3 and 4.2. We did not find a significant reaction of long-term rates to oil supply news shocks in Japan in the ZLB periods. An alternative assessment tool is a measure of monetary policy stance summarized by the shadow rate (Wu and Xia, 2016; Krippner, 2020). The shadow rate is a synthetic short-term nominal interest rate that is estimated to be consistent with the part of the yield curve that is above zero. By construction, the shadow rate closely follows the actual short-term nominal interest rate when the ZLB constraint does not bind, but it may go negative when the ZLB constraint is active. Panel (a) of Figure B.11 presents the impulse responses of the shadow rate to an oil supply news shock in the ZLB periods in the United States using the Wu and Xia (2016) updated shadow rate. To be consistent with the estimation of the interest rate responses in Sections 3.3 and 4.2, we add 12 lags of the CPI inflation rate as controls. We estimate projections on the sample from 1990:1 to 2019:12 and report the ZLB shadow rate responses only.[33] The shadow rate falls by 0.1 percentage point, and the fall is significant at horizons 2-3 and 21-22 months after an oil price hike. This pattern is consistent with the expansion of real economic activity during the ZLB we obtained earlier. One reason why the shadow rate might fall after the oil price spike is that the inflation rate is typically below target during ZLB episodes, potentially mak-

[32] Most macro models emphasize the role of trade balance and monetary policy in determining exchange rates. However, the exchange rate disconnect puzzle (Obstfeld and Rogoff, 2000)—the absence of correlation between exchange rate and macro fundamentals—is still standing. See Engel and Wu (2023)for recent progress in identifying the effects of macro variables on the exchange rate.

[33] The normal period response of the shadow rate is qualitatively similar to the response of the short-term nominal rate in Figure 7 but nosier at longer horizons.

ing central bankers less concerned about inflation and more concerned about a possible output decline following oil price spikes.

We also explore the response of an alternative measure of the shadow rates computed by Krippner (2020).[34] As plotted in Panel (b) of Figure B.11, the responses for the United States are not statistically different from zero. Similarly, the responses of the shadow rates in Japan and the United Kingdom are not statistically different from zero. These results suggest the inherent uncertainty across alternative shadow rate measures.

We perform two additional robustness checks using the shadow rate as a control. First, we include 12 lags of the shadow rate at the ZLB and of the short-term rate in normal times; Appendix Figure B.12 shows similar results. Second, following Holm, Paul and Tischbirek (2021), we add 12 leads of the shadow rate, which mechanically shuts down unconventional policy effects; Appendix Figure B.13 yields comparable responses.

5.3 Other Countries

Canada and the euro area. Japan, the United States, and the United Kingdom experienced clear episodes when the interest rate was almost constant due to the ZLB. Two more economies experienced approximately constant interest rate episodes: Canada and the euro area. Figure B.14 plots their short-term nominal interest rates. We define the "ZLB period" (or, more precisely, an approximately constant-rate period) to be between 2009:4 and 2017:4 in Canada and between 2009:07 and 2010:09, 2010:11 (just one month), and between 2012:01 and 2019:12 in the euro area.[35] It is important to note that the interest rates continued changing even when they were approximately constant in both economies, which is why it is reasonable to expect noisier results.

Figure B.15 shows no clear difference in the responses of industrial production and unemployment between the "ZLB" (i.e., approximately constant rates period) and normal periods in Canada and the euro area. In Canada, although the estimates are not statistically significant, the point estimates are consistent with our findings for Japan, the United States, and the United Kingdom. In the euro area, industrial production rises more and unemployment falls more in normal times than during the ZLB in the first months after

[34] Wu and Xia (2016) estimate a Gaussian affine shadow-rate term structure model and use the implied shadow short rate as a measure of monetary policy stance. Krippner (2020), while based on the same Black (1995) framework, adopts a different term-premium decomposition and pricing approach, typically yielding less negative shadow rates.

[35] In Canada, the short-term nominal interest rate never approached zero but remained at 1.25 percent from September 2010 to December 2014. As this window is too short for our baseline specification with 12 lags, we extend it to April 2009–August 2010 and January 2015–April 2017, both periods featuring only two policy rate changes. For the euro area, we define the constant-rate period as months when the EONIA rate was below 0.5 percent. During this time, EONIA fluctuated within a narrow range (0.25 to −0.45 percent) and displayed prolonged flat spells.

the shock. As shown in Figure B.11, the shadow policy rate in the euro area increases during the ZLB period, which may help explain why the euro area differs from Japan, the United States, and the United Kingdom.

Placebo countries. We next present the responses of four other countries—Mexico, Sweden, the Republic of Korea, and Norway—that did not experience ZLB during the Global Financial Crisis. We chose them among the OECD countries and ensured they are located in different parts of the world. Figure B.16 shows that neither industrial production nor unemployment rates respond differently in the normal and the ZLB periods based on the U.S. definition of the ZLB period.[36]

5.4 Other Specifications

State-dependent oil proxy VAR. Our baseline uses a two-step approach: we first extract oil supply news shocks from an oil proxy VAR without state dependence and then use these shocks in state-dependent local projections. To address the concern that a structural break in the oil VAR may coincide with Japan's monetary policy regimes, we estimate a state-dependent world oil VAR, where the state is defined by Japan's ZLB, and re-extract the shocks (see Appendix D.1). Using these shocks in the local projections yields results qualitatively similar to our baseline (Appendix Figure D.1).

We also extend the state-dependent oil proxy VAR by adding each variable of interest one at a time, effectively collapsing the two-step procedure into a single estimation. Appendix D.2 shows that, even in this specification, the differences in industrial production and the unemployment rate in Japan remain statistically significant at the 10 percent level.

Subperiod analysis. The advantage of showing results from several countries that experienced ZLB episodes is that these episodes may not coincide, providing evidence from different subperiods. In practice, however, the ZLB episodes in the United States and the United Kingdom occur almost simultaneously and partly overlap with the ZLB periods in Japan. To avoid the potential spillover of the United States during its ZLB period to other countries, we estimate the reaction of the Japanese industrial production and the unemployment rate in the normal and the ZLB periods when we shorten the sample to before the start of the ZLB period in the United States. Specifically, we choose the ending

[36] According to Norges Bank, there was a break in the Norwegian Labour Force Survey (LFS) in 2006, meaning that unemployment figures before and after 2006 are not directly comparable. Moreover, the 2006 data fluctuate substantially from month to month and are therefore rarely used; instead, staff primarily focus on developments in the trend. For this reason, we do not present the unemployment rate for Norway.

date for the sample to be June 2006, corresponding to the end of the first ZLB sub-period in Japan. The results, plotted in Figure B.17, qualitatively repeat the patterns for the full sample estimation in Figure 3.

Additionally, there may be structural breaks unrelated to the ZLB. For example, the shale oil revolution may have substantially altered the oil industry. In 2011, the U.S. became a net exporter of refined petroleum products, a development that could affect the dynamics of the oil market.[37] We present impulse responses to the oil supply news shock recovered from the linear oil proxy VAR estimated over the 1974:1–2006:6 sample.[38] Figure B.18 shows that the results are qualitatively similar to our baseline, although standard errors are somewhat larger due to the shorter ZLB subsample.

The first half of 2008 was marked by the world plunging into the Global Financial Crisis, with recession in Japan starting in March 2008. At the same time, the oil price continued rising until July 2008, partly due to rising oil demand in China and a stable oil supply in oil-producing countries (Hamilton, 2009). Such a coincidence may have an important effect on the estimation of the oil supply news shocks on economic activity. To investigate this possibility, we re-estimate the responses of Japanese industrial production and the unemployment rate using specification (1), where we not only interact the ZLB and normal-period dummies with the control variables but also introduce another dummy that equals one during 2008:3–2008:7 and interact it with all control variables. Figure B.19 presents the responses of Japanese variables during the ZLB and normal periods, and we do not add the responses during 2008:3-2008:7 to the figure. The responses of industrial production and the unemployment rate in the ZLB period remain virtually unchanged compared with our baseline results. However, the point estimate responses in the normal periods become less contractionary than before, consistent with the hypothesis that the developments in 2008 may be important for generating a negative relationship between economic activity and the oil supply news shocks. This observation can also be seen in Figure B.17 where the point estimates in the normal periods are less contractionary than in our baseline results. Finally, if we were to assign the 2008:3-2008:7 period to the ZLB episode, we would estimate a weaker expansionary effect of the shocks in Japan during the ZLB episode.[39]

The United States responded to the large oil price swings of the 1970s by introducing rationing of gasoline. To check if these episodes affect our U.S. results, we estimate

[37] At the same time, Baumeister and Hamilton (2023) suggest a possible shift in the global oil industry around 2005, coinciding with Saudi Arabia reaching peak oil production. None of these dates coincides with the Japanese ZLB period.

[38] The choice of the end date is a compromise between having a break not too much later than 2005 and not having ZLB period too short.

[39] We thank Christiane Baumeister for pointing out the importance of the developments in the first half of 2008.

the responses of U.S. industrial production and the unemployment rate with the sample starting in 1986:1, following Ramey and Vine (2011). Figure B.20 shows that our results are robust to shortening of the sample. The only notable difference is that the normal period responses become more persistent than our baseline results in Figure 6.

Controls. In Section 2.1, we argued in favor of using controls, such as a lagged dependent variable, lagged oil shocks, and the lagged unemployment rate, to enhance the signal-to-noise ratio of the oil supply news shocks. To verify this claim, we present results obtained without these controls. Appendix Figure B.21 displays the impulse response functions of Japanese industrial production and the unemployment rate, where we have omitted controls $\psi^y_{A,h}(L)x_{t-1}$ and $\psi^y_{B,h}(L)x_{t-1}$, except for the lags of the left-hand side variable of interest from equation (1). While the point estimates are similar to those with controls, the standard errors increase slightly. Furthermore, Appendix Figure B.22 demonstrates that our findings hold up when we add additional country-specific controls, such as the lagged inflation rate and the lagged 5-year nominal government bond yield.[40]

Finally, adding 12 lags of each of the four variables in the oil market VAR, or including fossil fuel imports as controls in the baseline specification, does not alter our findings (see Appendix Figures B.23 and B.24).

Information effect. In monetary economics, a high-frequency approach to identifying monetary policy shocks is subject to a criticism that it cannot distinguish between a true shock and the information revelation about the state of the economy by a central bank (Nakamura and Steinsson, 2018a). The same criticism may apply to the announcements by OPEC. For example, OPEC may have better information about the state of the global oil market than market participants themselves. To check whether our results are driven by the information revelation about the state of the economy, we use recent insights from the monetary shocks identification literature (Cieslak and Schrimpf, 2019; Cieslak and Pang, 2021; Jarociński and Karadi, 2020) that were applied to OPEC announcements in Degasperi (2021). The idea is to separate the movements in the oil supply surprise series on the days of the OPEC announcements into those that positively co-move with the stock market, and are likely to represent the information revelation effect, and those that co-move negatively, and are likely to proxy true oil supply surprises. We use the S&P 500 stock price index to measure the stock market reaction. The advantage of using the

[40] The results are not sensitive to using one or two-year nominal interest rates, or controlling additionally for the country-specific shadow rate in the ZLB period and short-term nominal interest rate in the normal period.

U.S. stock market index is that it may not react strongly to monetary policy in Japan. Appendix Figure B.25 presents the results for Japan based on the oil supply news shocks that were estimated using a part of the oil supply surprise series that generates a negative comovement between the U.S. stock market and the real oil price. The results are qualitatively similar to the baseline results, although the differences in the responses of the unemployment rate are not as significant as our baseline result.[41]

Expansions and recessions. Macroeconomic variables may respond differently to shocks depending on other states of the economy. Since ZLB episodes often coincide with recessions, our results could reflect differences between recessions and expansions rather than the ZLB itself (Auerbach and Gorodnichenko, 2013). To examine this, we re-estimate equation (1) using the same controls but replace the ZLB dummy with a recession indicator based on the Cabinet Office of Japan's business cycle dating. If our earlier findings were driven solely by cyclical conditions, we should observe systematic differences between recessions and expansions. Appendix Figure B.26 shows the responses, and Figure C.8 plots their difference. The unemployment responses do not differ significantly at any horizon. Industrial production differs at a few horizons (1–2 and 9–13 months), but the sign of the difference changes over time. Overall, these results do not seem to explain the persistent differences between the ZLB and normal periods that we find in our baseline estimation.[42]

Doing without external instruments. In our baseline analysis, we assumed that the four-variable oil market VAR spans all the relevant information to estimate the series of the oil supply news shocks $\widehat{\mathcal{E}}_{1,t}$, an assumption for which we could find supporting evidence. Nevertheless, for completeness, we now show the responses using an alternative strategy, where we directly use the oil supply surprise series z_t as a shock. More precisely, we first regress the series on twelve lags of four variables from the oil VAR and extract the residuals.[43] We then estimate specification (1), with the shocks being the residuals

[41] One can object to even using the U.S. stock market by arguing that the U.S. economy entered the ZLB in 2008 and the S&P 500 changed the sign of its co-movement with the oil supply series after that, as we previously illustrated. To address this concern, we re-estimate the results in Figure B.25 by removing a period when the U.S. short-term nominal interest rate was constant. Specifically, we choose a sample between 1975:1 and 2006:6. The main finding that the shock that increases the real oil price is less contractionary remains true even in this case.

[42] It is sometimes argued that the entire ZLB experience in Japan coincides with a long recession. We cannot separate the effects of this long recession from a constant interest rate period. However, following Section 6.2 in Miyamoto, Nguyen and Sergeyev (2018), we note that in modern macro models that generate different responses in booms and busts, a crucial indicator is labor market tightness, which did not exhibit drastically different behavior before and during the ZLB.

[43] This step is not necessary and does not change the results. However, by doing so, we make sure that the oil supply surprise series is not predicted by other oil variables. This step also relates the exercise to

just obtained. The advantage of this approach is that it does not require the invertibility assumption. However, we can only use the sample where the oil supply surprise series is available, namely 1983:4-2019:12. The left panel of Appendix Figure B.27 shows that industrial production in Japan responds differentially in the ZLB and normal periods, consistent with our baseline results. The estimates are, however, much noisier, which is expected because the sample is shorter and shocks identified in high frequency tend to be small. The magnitudes of the point estimates are larger than those in our baseline estimation, which can be related to the fact that, unlike in our baseline estimation, the response of the real oil price peaks at higher than 10 percent level after several months following the shock. The right panel of Appendix Figure B.27 presents the responses of the unemployment rate. The standard errors are large, preventing us from making any statistical statements. However, the point estimates paint a similar picture to our baseline results.[44]

6 Oil Shocks in an Open Economy New Keynesian Model

This section examines a standard Open Economy New Keynesian model with oil shocks with the following two goals. First, we formalize the intuition that oil market shocks that increase oil price can sometimes be expansionary. Second, we present a stylized quantitative exercise aimed at quantifying the model-implied effects of oil price spikes on macro variables under constant nominal interest rate monetary policy.

6.1 The Model

Our model is a version of Galí and Monacelli (2002) extended with oil intermediate inputs used by firms in the spirit of Bodenstein, Guerrieri and Gust (2013).[45] It consists of a small open home country and a foreign country representing the world economy. The small open economy is intended to approximate Japan. Each country has a representative household that consumes home and foreign goods with home-biased preferences, saves via complete domestic and international asset markets, and supplies labor to domestic

the internal instrument approach (Ramey, 2011; Plagborg-Møller and Wolf, 2021), where the instrument is added to the VAR and ordered first.

[44]We also estimate a version of our main specification in which we replace the oil supply news shock with the oil price instrumented by the oil supply surprise series. The industrial production responses in Japan are qualitatively similar to our baseline results, although the standard errors are substantially larger. We do not find any difference in the unemployment rate responses. This result reflects a power problem: identifying the effects of oil supply shocks, proxied by oil surprise series, at all horizons is difficult without imposing additional assumptions (Nakamura and Steinsson, 2018b).

[45]Galí and Monacelli (2002) is a working paper of the classical Gali and Monacelli (2005) paper. The key difference between the two is that the working paper features two countries, a small open economy and the world economy, while the published version has a continuum of small open-economies.

firms. Monopolistically competitive firms produce differentiated goods from domestic labor and oil using a constant-returns-to-scale production function. These firms set their nominal prices in domestic currencies and change them infrequently a la Calvo. Only foreign country receives a stochastic endowment of oil in each period. All goods, including oil, cannot be stored. In addition, all goods are traded globally, implying the law of one price.[46]

The foreign country—the world economy—can be described by the following log-linearized equilibrium conditions, which we derive in Appendix F. Aggregate consumption follows the Euler equation:

$$\mathbb{E}_t\left[\widehat{c}^*_{t+1}\right] - \widehat{c}^*_t = \frac{1}{\sigma}\left(i^*_t - \mathbb{E}_t\left[\pi^*_{F,t+1}\right] - \iota\right), \tag{2}$$

where asterisks denote foreign country variables, $\widehat{c}^*_t$ is the log deviation of world consumption from its steady state, i^*_t is the nominal interest rate, $\pi^*_{F,t+1}$ is both a consumer price index (CPI) and producer price index (PPI) inflation rate because the world economy is closed, ι is the steady-state nominal interest rate, and σ is the inverse of the intertemporal elasticity of substitution. The world variables satisfy the Phillips curve

$$\pi^*_{F,t} = \kappa\zeta^*\widehat{c}^*_t + \beta\mathbb{E}_t\left[\pi^*_{F,t+1}\right] + \kappa\psi^*_o\widehat{r}_t, \tag{3}$$

where ζ^* and ψ^*_o are composite parameters that represent the sensitivities of firm's marginal costs to consumption $\widehat{c}^*_t$ and the real oil price $\widehat{r}_t$, which is expressed in units of foreign goods, κ is the sensitivity of inflation to changes in the marginal cost.[47] The main difference between a standard closed economy model and this world economy is the last term in the Phillips curve equation, $\kappa\psi^*_o\widehat{r}_t$, which relates changes in the real oil price and inflation. The foreign central bank sets the interest rate according to the Taylor rule

$$i^*_t = \iota + \varphi^*_\pi\pi^*_{F,t}, \tag{4}$$

where the inflation target is zero, the effective lower bound is not explicitly introduced, and φ^*_π is non-negative. Finally, the equilibrium real oil price is determined in the world's oil market:

$$\widehat{r}_t = \phi_c\widehat{c}^*_t - \phi_o\widehat{o}^*_t, \tag{5}$$

[46] This stylized model can be extended in many realistic dimensions that we do not pursue here. It can feature non-rational expectations (Milani, 2009), incomplete markets within (Chan *et al.*, 2022; Auclert *et al.*, 2023) and across (Itskhoki and Mukhin, 2021) countries, local or dollar currency pricing (Gopinath and Itskhoki, 2022), production and storage of oil (Pindyck, 2001), and many more.

[47] The composite coefficients $\kappa, \zeta^*, \psi^*_o$, are expressed in equations (F.24), (F.27), and (F.28) in Appendix F through preferences and production function parameters.

where $\widehat{o}_t^*$ is the oil endowment and ϕ_c, ϕ_o are positive.[48] This equation implies that an increase in the world economic activity or a decrease in the oil supply lead to an increase in the real oil price.

Home country consumption dynamics follows a similar Euler equation

$$\mathbb{E}_t\left[\widehat{c}_{t+1}\right] - \widehat{c}_t = \frac{1}{\sigma}\left(i_t - \mathbb{E}_t\left[\pi_{t+1}\right] - \iota\right), \tag{6}$$

while the home-produced goods inflation rate obeys the following Phillips curve

$$\pi_{H,t} = \kappa\zeta_H\widehat{c}_t + \beta\mathbb{E}_t\left[\pi_{H,t+1}\right] + \kappa\psi_o\widehat{r}_t + \kappa\zeta_F\widehat{c}_t^* + \kappa\psi_q\widehat{q}_t, \tag{7}$$

where, in addition, to the channels present in the Phillips curve of the foreign country in equation (3), the term $\kappa\zeta_F\widehat{c}_t^*$ incorporates the effect of global aggregate demand on domestic inflation, and the last term shows the influence of the real exchange rate $\widehat{q}_t$ on the marginal costs of home firms and, hence, the inflation rate. The real exchange rate is defined as the price of the foreign consumption basket in units of the home consumption basket, so an increase in $\widehat{q}_t$ is a real depreciation. When the global aggregate demand drops, foreign residents purchase fewer home goods, reducing marginal costs of home firms and, hence, the inflation rate. In addition, when the home currency depreciates in real terms ($\widehat{q}_t$ increases), oil becomes more expensive to buy at home for any level of the real oil price $\widehat{r}_t$ denominated in foreign goods. As a result, home firms' marginal costs and the inflation rate go up.

The central bank at home follows the Taylor rule

$$i_t = \iota + \varphi_\pi\pi_{H,t}, \tag{8}$$

where $\varphi_\pi \geq 0$. The link between the CPI and PPI inflation rates is

$$\pi_t = \pi_{H,t} + \frac{1-\Omega}{\Omega}\Delta\widehat{q}_t, \tag{9}$$

where $\Omega \in [0,1]$ is the share of home goods in home household spending, which represents the degree of home bias. The second term expresses the effect of real depreciation of domestic currency on the consumption basket inflation.

With complete financial markets, there is perfect risk sharing across countries, so the real exchange rate is proportional to the difference in consumption in the two countries:

$$\widehat{q}_t = \sigma\left(\widehat{c}_t - \widehat{c}_t^*\right). \tag{10}$$

[48] The composite coefficients ϕ_c and ϕ_o are in equations (F.25) and (F.26).

The production of home goods is a sum of home and foreign demand for home goods

$$\widehat{y}_{H,t} = \frac{\Omega}{\Omega + 1 - \Omega^*}\widehat{c}_t + \frac{1-\Omega^*}{\Omega + 1 - \Omega^*}\widehat{c}_t^* + \frac{\Omega\,(1-\Omega) + (1-\Omega^*)}{\Omega + 1 - \Omega^*} \cdot \frac{\gamma_n}{\Omega}\widehat{q}_t. \tag{11}$$

where $1 - \Omega^*$ is the proportion of consumption that foreign households allocate to home goods relative to the size of home country.[49] The last two terms express two forces that take into account that oil shocks affect home country through the world economy. The presence of $\widehat{c}_t^*$ captures the spillover of world consumption demand to domestic goods. The term featuring the real exchange rate $\widehat{q}_t$ measures the strength of the expenditure switching effect between foreign to home goods, which crucially depends on the elasticity of substitution between home and foreign goods γ_n.

We look into the consequences of a one-time unanticipated negative global oil supply shock $\widehat{o}_0^* = \widehat{o}^* < 0$ that reverts back to its steady-state value of zero with probability $1 - \alpha \in [0, 1]$. There are no other sources of uncertainty.

Instead of explicitly representing the zero lower bound on the nominal interest rate, we assume that the nominal interest rate stays constant following an oil price shock. This assumption retains the key property of the ZLB that the interest rate does not respond to developments in the economy. At the same time, it avoids the need to introduce additional notation. See Nakamura and Steinsson (2014) for a similar treatment in the context of fiscal multipliers.

Oil supply news shocks. We note that in our empirical results, oil supply news shocks affect the oil price immediately but the actual oil supply gradually change with a lag of several months. In this stylized model, where oil cannot be stored and the price of oil is determined in the spot market, oil supply news can only affect current oil price through other forward-looking macro variables, such as consumption. Because of this, we choose to focus on the realized oil supply shocks in this model. From the perspective of a small open economy, the distinction between oil supply news and oil supply shocks may not be crucial when the key manifestation of the shock is an increase in the oil price. Moreover, in our numerical exercise, we employ quarterly calibration, which brings the supply news shocks closer to the realized supply shocks. Finally, in Appendix Section F.2, we demonstrate that our simple model extended with oil inventories have the same implications about consumption and inflation responses to oil supply news shocks as those in the baseline model with respect to unanticipated oil supply shocks.

[49] As Appendix F formally shows, we assume that foreign demand for home goods is proportional to the size of the home country. Because we consider a limit when the size of the home country tends to zero (the small open economy assumption), both home production and foreign demand for home goods shrink to zero. However, their ratio stays positive in the limit.

Figure 8: The World Economy

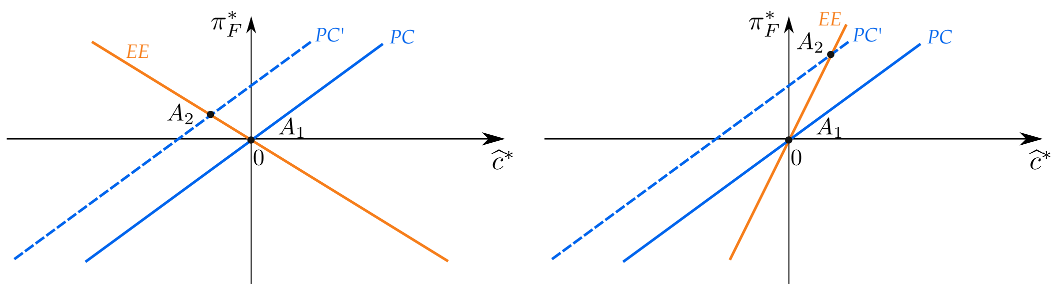

Notes: Both panels show the response of the world economy to a decline in oil supply when monetary policy is active (the left panel) and when it does not respond following an oil shock (the right panel).

6.2 The Effects of Oil Shocks

We start by presenting the responses of foreign goods production, which equals consumption, and the inflation rate. We look for a unique bounded Markov equilibrium where endogenous variables $\{\widehat{c}_t^*, \pi_{F,t+1}^*, i_t^*, \widehat{r}_t\}$ are linear functions of oil supply $\widehat{o}_t^*$. The Phillips curve and the Euler equation (together with the Taylor rule) are

$$\pi_F^* = \frac{\kappa\left(\zeta^* + \psi_o^* \phi_c\right)}{1-\beta\alpha}\widehat{c}^* + \frac{\kappa\psi_o^*\phi_o}{1-\beta\alpha}\left(-\widehat{o}^*\right), \tag{12}$$

$$\widehat{c}^* = -\frac{\varphi_\pi^* - \alpha}{\left(1-\alpha\right)\sigma}\pi_F^*. \tag{13}$$

Figure 8 plots the two equations under the assumption that monetary policy is active, that is, $\varphi_\pi^* > 1$ (the left panel), and when monetary policy does not change following an oil shock, that is, $\varphi_\pi^* = 0$ (the right panel). The upward-sloping Euler equation on the right panel is due to the fact that higher inflation in the absence of the response of the nominal interest rate reduces the real interest rate and increases aggregate consumption. For the unique bounded solution to exist, the Euler equation should be steeper than the Phillips curve when they have slopes with the same sign.[50]

A decline in the oil supply results in an upward shift in the Phillips curve in both panels of Figure 8. Intuitively, a decline in the oil supply increases the real oil price, which, in turn, pushes up the inflation rate.

Under active monetary policy, the central bank increases the nominal interest rate such that the real interest rate rises, reducing aggregate output. At the same time, under the constant nominal interest rate policy, the real rate falls, stimulating the world economy.

Similarly to the world economy, we can represent the Phillips curve and Euler equa-

[50] Formally, we focus on the parameters that satisfy $(1-\beta\alpha)(1-\alpha)\sigma + \kappa\left(\varphi_\pi^* - \alpha\right)\left(\zeta^* + \psi_o^*\phi_c\right) > 0$. This inequality always holds when the monetary policy is active. When $\varphi_\pi^* = 0$, however, the inequality restricts the parameter space.

Figure 9: The Small-Open Economy

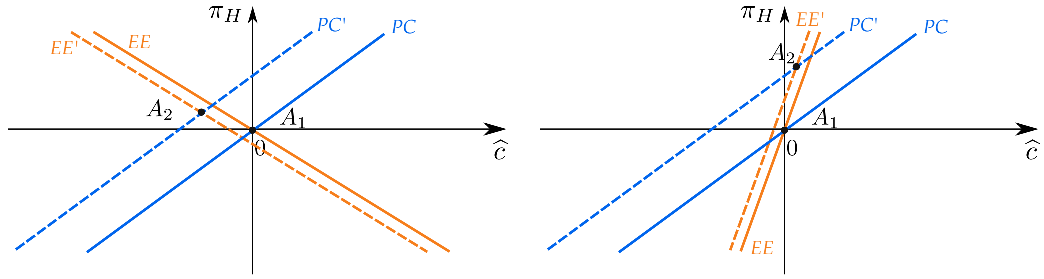

Notes: Both panels show the response of the small-open economy to a decline in oil supply when the monetary policy is active (the left panel) and when it does not respond following the oil shock (the right panel). Both panels assume that the world economy's output declines following an increase in the oil price.

tion of the home economy as

$$\pi_H = \frac{\kappa\left(\zeta_H + \psi_q \sigma\right)}{1-\alpha\beta}\widehat{c} + \frac{\kappa}{1-\alpha\beta}\psi_o \widehat{r} + \frac{\kappa\left(\zeta_F - \psi_q \sigma\right)}{1-\alpha\beta}\widehat{c}^*, \tag{14}$$

$$\widehat{c} = -\Omega \frac{\varphi_\pi - \alpha}{\sigma\left(1-\alpha\right)}\pi_H + \left(1-\Omega\right)\widehat{c}^*. \tag{15}$$

The two equations closely resemble the world equations (12) and (13), with the exception of additional terms. When the oil supply decreases, real oil prices increase, pushing down global demand $\widehat{c}^*$ when the world monetary policy tightens, the case we will focus on. The last term in the Phillips curve encapsulates two global forces. First, a slowdown in global demand leads domestic marginal costs and inflation to fall. Second, to keep the marginal utility of consumption equalized in both countries as required by equation (10), the home currency depreciates in real terms. This makes the purchase of oil from the global market costlier and consequently increases the marginal costs of production at home. The net effect of these two forces is ambiguous, and it depends on the degree of openness of the home country. In Figure 9, we plot equation (14) when monetary policy in the rest of the world actively responds to the oil shock. In the graph, we assume that the effect of the oil supply decline on the real oil price and the real exchange rate is stronger than its effect on foreign demand for home goods. Therefore, a decrease in oil supply shifts the PC to the left.

In addition, when there is a decline in global economic activities, the last term in the Euler equation implies that the EE line shifts to the left, as plotted in Figure 9. The reason is that a decline in global economic activity, which is expected to revert in the future, appreciates home currency in real terms over time. This, in turn, reduces home inflation and, hence, increases the real interest rate for any level of the nominal interest rate. As a result, domestic aggregate demand falls.

In equilibrium, as demonstrated in Figure 9, a decline in oil supply that increases the oil price can also be expansionary when the leftward shift in the Euler equation is small enough, which is the case with a sufficiently high degree of home bias, i.e., when the coefficient Ω is close to one.[51]

Table 1: Calibration Parameters

$\beta = 0.99$	Discount factor
$\sigma^{-1} = 1$	Intertemporal elasticity of substitution
$\varphi^{-1} = 0.5$	Frisch elasticity of labor supply
$\gamma_n = 0.5$	Elasticity of substitution between home and foreign goods
$\Omega = 0.8$	Share of home goods in home consumption basket
$\alpha = 0.58$	Persistence of oil shocks
$\varphi_\pi = \varphi_\pi^* = 1.5$	Taylor rule coefficient on inflation
$\gamma_y = 0.3$	Elasticity between labor and oil in production
$\omega_{oy} = 0.057$	share of oil in production
$1 - \theta = 0.25$	Probability of price adjustment
Composite parameters	
$\kappa = (1-\theta)(1-\beta\theta)/\theta = 0.086$	$\zeta^* = (1-\omega_{oy})(\sigma+\varphi)/(1+\omega_{oy}\gamma_y\varphi) = 2.735$
$\psi_o^* = \psi_o = \frac{(1+\gamma_y\varphi)\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}} = 0.088$	$\phi_c = \{1+\gamma_y[\varphi+(1-\omega_{oy})\sigma]\}/[(1-\omega_{oy})\gamma_y] = 6.656$
$\phi_o = \frac{1+\omega_{oy}\gamma_y\varphi}{(1-\omega_{oy})\gamma_y} = 3.656$	$\psi_q = \omega_{oy} + \frac{1-\Omega}{\Omega} + \frac{(1-\omega_{oy})\varphi\gamma_y}{(1+\gamma_y\varphi\omega_{oy})}\left(\omega_{oy} + \frac{\gamma_n}{\gamma_y}\frac{1-\Omega^2(1-\omega_{oy})}{\Omega}\right) = 0.790$
$\zeta_H = \frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}\left(\sigma + \varphi\frac{\Omega}{\Omega+1-\Omega^*}\right) = 2.288$	$\zeta_F = \frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}\varphi\frac{1-\Omega^*}{\Omega+1-\Omega^*} = 0.448$
$\Omega^* = \Omega - \omega_{oy}/(1-\omega_{oy}) = 0.740$	$\iota = -\log\beta = 0.01$

Numerical exercise. To illustrate the quantitative importance of the highlighted channels, we calibrate the model by choosing the values used in the recent macro literature. The quarterly calibration parameters are listed in Table 1.[52] Most of the parameters are taken from Gali and Monacelli (2005). The exceptions are the share of oil in production, which come from Bodenstein *et al.* (2013). The persistence of oil shocks is set to 0.58 (half life of 1.3 quarters or 3.8 months) to match the maximum point estimate decline in Japanese aggregate consumption by 0.2 percent in the normal period following a shock that increases oil prices by 10 percent. The elasticity of substitution between home and foreign goods consumption is set to 0.5 and between labor and oil in production to 0.3. We chose these somewhat lower than usual values to highlight the short-run effects of oil supply shocks. The weight of home goods in the home consumption basket is $\Omega = 0.8$, which is set to approximate this value in Japan.[53]

[51] Note that for the bounded solution to be unique, the following condition must hold $\sigma(1-\alpha\beta)(1-\alpha) + \kappa(\varphi_\pi - \alpha)(\zeta_H + \psi_q\sigma)\Omega > 0$. This condition always holds when the central bank actively responds to the oil shock, i.e., $\varphi_\pi > 1$. However, when $\varphi_\pi = 0$, the above condition restricts the set of model parameters.

[52] We use quarterly calibration because we estimate the response of consumption and GDP in Japan using quarterly variables.

[53] This value is lower than the one typically used for the U.S., but higher that that for euro area countries.

Under these parameters and a shock to oil supply that increases oil prices by 10 percent, world consumption falls by 0.18 percent under active monetary policy and increases by 1 percent when the nominal interest rate is constant. The home economy's consumption falls 0.18 percent when monetary policy is active both at home and abroad, while home consumption increases 0.55 percent when monetary policy does not respond at home but responds in the rest of the world. This is compared to the maximum estimated response of Japanese consumption at the ZLB of 0.47 percent, presented in Figure 4. This stylized numerical exercise captures our empirical result that the ZLB response is larger than the absolute value of the response in the normal period.

The model allows us to look at other variables. For example, home production increases by 0.56 percent when home monetary policy does not respond (but foreign monetary policy does), and it falls by 0.18 percent when home monetary policy is active. Furthermore, using equation (11), we can decompose the 0.56 percent change in output into 0.42 percent due to a higher aggregate demand at home, -0.04 percent due to a lower foreign demand, and 0.18 percent due to expenditure switching, suggesting the importance of the international spillovers of the oil price variations.

Optimal monetary policy. We conclude by discussing the monetary policy implications of the model. In our baseline model, with complete markets, producer-currency pricing, and subsidies offsetting monopolistic distortions, the only friction is price stickiness. In a closed-economy setting such as our model's world economy, optimal monetary policy stabilizes output at its natural level, implying inflation targeting (Gali, 2008).

In a multi-country setting, optimal policy depends on the degree of monetary cooperation across countries. Under cooperation, central banks close domestic output gaps, consistent with the international "divine coincidence" (Corsetti, Dedola and Leduc, 2010). Without coordination, policy choices may exhibit a terms-of-trade externality, as central banks have incentives to manipulate relative prices. Faia and Monacelli (2008) numerically show that, for productivity shocks, non-cooperative policy deviates only slightly from the cooperative benchmark. Although oil shocks differ from productivity shocks because they enter countries' budget constraints, if oil's share in a country's spending is small, a continuity argument suggests similar conclusions in our setting.

Our empirical results show that inflation responds to oil supply news shocks in normal periods. In the context of our model, this may indicate suboptimal policy. However, in richer frameworks with additional frictions, such inflation responses could be consistent with optimal policy.

7 Conclusion

This paper presents new evidence on the effects of oil supply news shocks when the nominal interest rate is not responsive due to the ZLB constraint. We focus on Japan, which has the longest spell at the ZLB, and supplement with evidence from the United Kingdom and the United States, which have also experienced considerable periods with fixed interest rates. We find that oil supply shocks are less contractionary in the ZLB period than in the periods outside of the ZLB in Japan, the United Kingdom, and the United States. Notably, we document that these economies expand following an oil price increase during the ZLB periods. In addition, we find that inflation expectations rise after an oil supply news shock, while the nominal interest rate remains at zero during the ZLB and increases in normal periods. This suggests that the monetary policy channel can play a significant role in the transmission of oil shocks.

References

ALQUIST, R. and KILIAN, L. (2010). What do we learn from the price of crude oil futures? *Journal of Applied econometrics*, **25** (4), 539–573.

AUCLERT, A., MONNERY, H., ROGNLIE, M. and STRAUB, L. (2023). *Managing an Energy Shock: Fiscal and Monetary Policy*. Tech. rep., Mimeo, Stanford.

AUERBACH, A. J. and GORODNICHENKO, Y. (2013). Fiscal multipliers in recession and expansion. In A. Alesina and F. Giavazzi (eds.), *Fiscal Policy after the Financial crisis, 2*, University of Chicago Press, pp. 63–98.

BAQAEE, D. R. and FARHI, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, **87** (4), 1155–1203.

BAUER, M. D. and SWANSON, E. T. (2023). A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, **37** (1), 87–155.

BAUMEISTER, C. and HAMILTON, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, **109** (5), 1873–1910.

— and — (2023). *Uncovering disaggregated oil market dynamics: A full-information approach to granular instrumental variables*. Tech. rep.

BERNANKE, B. S., GERTLER, M. and WATSON, M. (1997). Systematic monetary policy and the effects of oil price shocks. *Brookings papers on economic activity*, **1997** (1), 91–157.

BJØRNLAND, H. C., CROSS, J. L. and HÖLZ, J. (2025). Re-visiting the relationship between oil prices and monetary policy cama working paper 19/2025 april 2025.

BLACK, F. (1995). Interest rates as options. *the Journal of Finance*, **50** (5), 1371–1376.

BLANCHARD, O. and GALI, J. (2009). The macroeconomic effects of oil shocks: Why are the 2000s so different from the 1970s? In J. Gali and M. Gertler (eds.), *International Dimensions of Monetary Policy.*, University of Chicago Press.

BODENSTEIN, M., GUERRIERI, L. and GUST, C. J. (2013). Oil shocks and the zero bound on nominal interest rates. *Journal of International Money and Finance*, **12**, 941–967.

CANDIA, B., COIBION, O. and GORODNICHENKO, Y. (2023). The macroeconomic expectations of firms. In *Handbook of Economic Expectations*, Elsevier, pp. 321–353.

Chan, J., Diz, S. and Kanngiesser, D. (2022). Energy prices and household heterogeneity: Monetary policy in a gas-tank. *Available at SSRN 4255158*.

Cieslak, A. and Pang, H. (2021). Common shocks in stocks and bonds. *Journal of Financial Economics*, **142** (2), 880–904.

— and Schrimpf, A. (2019). Non-monetary news in central bank communication. *Journal of International Economics*, **118**, 293–315.

Coibion, O. and Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, **120** (1), 116–159.

Corsetti, G., Dedola, L. and Leduc, S. (2010). Optimal monetary policy in open economies. In *Handbook of monetary economics*, vol. 3, Elsevier, pp. 861–933.

Datta, D. D., Johannsen, B. K., Kwon, H. and Vigfusson, R. J. (2021). Oil, equities, and the zero lower bound. *American Economic Journal: Macroeconomics*, **13** (2), 214–253.

Debortoli, D., Galí, J. and Gambetti, L. (2020). On the empirical (ir) relevance of the zero lower bound constraint. *NBER Macroeconomics Annual*, **34** (1), 141–170.

Degasperi, R. (2021). Identification of expectational shocks in the oil market using opec announcements. *manuscript, University of Warwick*.

Eggertsson, G. B. (2008). Can tax cuts deepen recessions? *Federal Reserve Bank of New York, December*.

Engel, C. and Wu, S. P. Y. (2023). Liquidity and exchange rates: An empirical investigation. *The Review of Economic Studies*, **90** (5), 2395–2438.

Faia, E. and Monacelli, T. (2008). Optimal monetary policy in a small open economy with home bias. *Journal of Money, credit and Banking*, **40** (4), 721–750.

Fernald, J. (2014). *A Quarterly, Utilization-Adjusted Series on Total Factor Productivity*. Tech. rep., Federal Reserve Bank of San Francisco.

Forni, M. and Gambetti, L. (2014). Sufficient information in structural vars. *Journal of Monetary Economics*, **66**, 124–136.

—, — and Ricco, G. (2024). Generalised external-instrument svar analysis, working paper, dated 15 September 2024.

Gali, J. (2008). Monetary policy, inflation, and the business cycle: An introduction to the new keynesian framework.

GALÍ, J. and MONACELLI, T. (2002). Monetary policy and exchange rate volatility in a small open economy.

GALI, J. and MONACELLI, T. (2005). Monetary policy and exchange rate volatility in a small open economy. *The Review of Economic Studies*, **72** (3), 707–734.

GARIN, J., LESTER, R. and SIMS, E. (2019). Are supply shocks contractionary at the zlb? evidence from utilization-adjusted tfp data. *Review of Economics and Statistics*, **101** (1), 160–175.

GERTLER, M. and KARADI, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, **7** (1), 44–76.

GIANNONE, D. and REICHLIN, L. (2006). Does information help recovering structural shocks from past observations? *Journal of the European Economic Association*, **4** (2-3), 455–465.

GOPINATH, G. and ITSKHOKI, O. (2022). Dominant currency paradigm: A review. *Handbook of International Economics*, **6**, 45–90.

GÜRKAYNAK, R. S., SACK, B. P. and SWANSON, E. T. (2004). Do actions speak louder than words? the response of asset prices to monetary policy actions and statements. *The Response of Asset Prices to Monetary Policy Actions and Statements (November 2004)*.

HAMILTON, J. D. (1983). Oil and the macroeconomy since world war ii. *Journal of political economy*, **91** (2), 228–248.

— (2009). Causes and consequences of the oil shock of 2007-08. *Paper presented at Brookings Institution*.

— and HERRERA, A. M. (2004). Comment: oil shocks and aggregate macroeconomic behavior: the role of monetary policy. *Journal of Money, credit and Banking*, pp. 265–286.

HOLM, M. B., PAUL, P. and TISCHBIREK, A. (2021). The transmission of monetary policy under the microscope. *Journal of Political Economy*, **129** (10), 2861–2904.

HURWICZ, L. (1950). Statistical inference in dynamic economic models. *Chap. Generalization of the Concept of Identification, edited by TC Koopmans*, **10**, 245–257.

ITSKHOKI, O. and MUKHIN, D. (2021). Exchange rate disconnect in general equilibrium. *Journal of Political Economy*, **129** (8), 2183–2232.

JAROCIŃSKI, M. and KARADI, P. (2020). Deconstructing monetary policy surprises-the role of information shocks. *American Economic Journal: Macroeconomics*, **12** (2), 1–43.

JENTSCH, C. and LUNSFORD, K. G. (2022). Asymptotically valid bootstrap inference for proxy svars. *Journal of Business & Economic Statistics*, **40** (4), 1876–1891.

JORDA, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, **95** (1), 161–182.

KANZIG, D. (2021). The macroeconomic effects of oil supply news: Evidence from opec announcements. *American Economic Review*, **111** (4), 1092–1125.

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

— and LEWIS, L. T. (2011). Does the fed respond to oil price shocks? *The Economic Journal*, **121** (555), 1047–1072.

— and MURPHY, D. (2012). Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market var models. *Journal of the European Economic Association*, **10** (5), 1166–1188.

— and MURPHY, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied econometrics*, **29** (3), 454–478.

KRIPPNER, L. (2020). A note of caution on shadow rate estimates. *Journal of Money, Credit and Banking*, **54** (4), 951–962.

KUTTNER, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the fed funds futures market. *Journal of monetary economics*, **47** (3), 523–544.

MERTENS, K. and RAVN, M. O. (2013). The dynamic effects of personal and corporate income tax changes in the united states. *American economic review*, **103** (4), 1212–47.

MICHAILLAT, P. and SAEZ, E. (2021). Resolving new keynesian anomalies with wealth in the utility function. *Review of Economics and Statistics*, **103** (2), 197–215.

MILANI, F. (2009). Expectations, learning, and the changing relationship between oil prices and the macroeconomy. *Energy Economics*, **31** (6), 827–837.

MIRANDA-AGRIPPINO, S. and REY, H. (2022). The global financial cycle. In *Handbook of international economics*, vol. 6, Elsevier, pp. 1–43.

MIYAMOTO, W., NGUYEN, T. L. and SERGEYEV, D. (2018). Government spending multipliers under the zero lower bound: Evidence from japan. *American Economic Journal: Macroeconomics*, **10** (3), 247–277.

MONTIEL-OLEA, J. L. and PLAGBORG-MOLLER, M. (2020). Local projection inference is simpler and more robust than you think. *Econometrica*, **89** (4), 1789–1823.

MONTIEL OLEA, J. L., STOCK, J. H. and WATSON, M. W. (2021). Inference in structural vector autoregressions identified with an external instrument. *Journal of Econometrics*, **225** (1), 74–87, themed Issue: Vector Autoregressions.

NAKAMURA, E. and STEINSSON, J. (2014). Fiscal Stimulus in a Monetary Union: Evidence from US Regions. *American Economic Review*, **104** (3), 753–92.

— and — (2018a). High-frequency identification of monetary non-neutrality: the information effect. *The Quarterly Journal of Economics*, **133** (3), 1283–1330.

— and — (2018b). Identification in macroeconomics. *Journal of Economic Perspectives*, **32** (3), 59–86.

OBSTFELD, M. and ROGOFF, K. (2000). The six major puzzles in international macroeconomics: is there a common cause? *NBER macroeconomics annual*, **15**, 339–390.

PAGAN, A. (1984). Econometric issues in the analysis of regressions with generated regressors. *International economic review*, pp. 221–247.

PEERSMAN, G. and VAN ROBAYS, I. (2012). Cross-country differences in the effects of oil shocks. *Energy Economics*, **34** (5), 1532–1547.

PINDYCK, R. S. (2001). The dynamics of commodity spot and futures markets: a primer. *The energy journal*, **22** (3).

PLAGBORG-MØLLER, M. and WOLF, C. K. (2021). Local projections and vars estimate the same impulse responses. *Econometrica*, **89** (2), 955–980.

— and WOLF, C. K. (2022). Instrumental variable identification of dynamic variance decompositions. *Journal of Political Economy*, **130** (8), 2164–2202.

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

— and VINE, D. J. (2011). Oil, automobiles, and the us economy: How much have things really changed? *NBER Macroeconomics Annual*, **25** (1), 333–368.

— and ZUBAIRY, S. (2018). Government spending multipliers in good times and in bad: Evidence from us historical data. *Journal of Political Economy*, **126** (2), 850–901.

REY, H. (2015). *Dilemma not trilemma: the global financial cycle and monetary policy independence*. Tech. rep., National Bureau of Economic Research.

ROTEMBERG, J. and WOODFORD, M. (1996). Imperfect competition and the effects of energy price increases on economic activity. *Journal of Money, Credit and Banking*, **28** (4), 550–77.

SIMS, C. A. and ZHA, T. (2006). Does monetary policy generate recessions? *Macroeconomic Dynamics*, **10** (2), 231–272.

STOCK, J. H. and WATSON, M. W. (2012). *Disentangling the Channels of the 2007-2009 Recession*. Tech. rep., National Bureau of Economic Research.

— and — (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal*, **128** (610), 917–948.

SWANSON, E. T. and WILLIAMS, J. C. (2014a). Measuring the effect of the zero lower bound on medium-and longer-term interest rates. *American economic review*, **104** (10), 3154–3185.

— and — (2014b). Measuring the effect of the zero lower bound on yields and exchange rates in the uk and germany. *Journal of International Economics*, **92**, S2–S21.

UNALMIS, D., UNALMIS, I. and UNSAL, D. F. (2012). *On the Sources and Consequences of Oil Price Shocks: the Role of Storage*. Tech. rep., IMF.

WATANABE, K. and WATANABE, T. (2017). Price rigidity at near-zero inflation rates: Evidence from japan. *CARF Working Papers*.

WEBER, M., D'ACUNTO, F., GORODNICHENKO, Y. and COIBION, O. (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives*, **36** (3), 157–184.

WIELAND, J. F. (2019). Are negative supply shocks expansionary at the zero lower bound? *Journal of Political Economy*, **127** (3), 973–1007.

WOLF, C. K. and MCKAY, A. (2022). *What Can Time-Series Regressions Tell Us About Policy Counterfactuals?* Tech. rep., National Bureau of Economic Research.

WOODFORD, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press Princeton, NJ.

WU, J. C. and XIA, F. D. (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking*, **48** (2-3), 253–291.

Online Appendix

A Data Details

1. **Industrial production**. Monthly industrial production data for all countries (i.e., Japan, UK, US, as well as Canada, the euro area, Mexico, the Republic of Korea, and Sweden) are taken from Haver Analytics. The series name is "Industrial production excluding construction, seasonally adjusted 2015=100." For all countries, the data are from 1974:1 to 2019:12. Non-oil and manufacturing industrial production for the US are downloaded from FRED from 1974:1 to 2019:12. The UK's manufacturing industrial production is downloaded from the Office for National Statistics website in December 2023 and is available from 1970:01 to 2019:12. Japan manufacturing industrial production is downloaded from Ministry of Economy, Trade and Industry (METI) website and is available from 1978:01 to 2019:12.

2. **Unemployment rate**. The unemployment rate data for all countries are taken from the OECD database. The series name is "Unemployment rate (monthly), Total, All persons."

3. **Inflation rate**. CPI data for all countries are taken from the OECD database. The subject of the series is "CPI: 01-12 All items, not seasonally adjusted, 2015=100." Core CPI data (2015=100, not seasonally adjusted) for all countries are downloaded from Haver Analytics, which populates data from the OECD Major Economic Indicator (MEI) database. The codename for the series is PCXG_N023. We seasonally adjusted core CPI using X-12-ARIMA in Stata.

4. **Inflation expectations.** The inflation expectations for Japan come from the Japan Center for Economic Research (JCER). The US inflation expectations are 12-month ahead median inflation expectations from the Michigan Survey of Consumers measured at monthly frequency. The UK inflation expectations are 4-quarter ahead median inflation expectations from the Bank of England/Ipsos Inflation Attitudes Survey measured at quarterly frequency.

5. **Short-term nominal interest rate.**

    (a) For Japan, the short-term nominal interest rate is the uncollateralized overnight call rate from July 1985. Prior to that date, we use the collateralized overnight call rate. Both data series are taken from the Bank of Japan website.

    (b) For the United States and the United Kingdom, the short-term interest rate is the overnight interbank rate downloaded from Haver Analytics, with a ticker FRUO, which populates data from OECD MEI.

    (c) For the Euro Area, the short-term nominal interest rate is the EONIA (Euro Overnight Index Average) from the ECB Data Portal.

6. **Long-term nominal interest rate.** For all countries, the long-term nominal interest rate is the 5-year government bond yield. Data are taken from the GFD.

7. **Shadow rate.** The U.S. shadow rate series is from Wu and Xia (2016). The updated version of the series was downloaded from Jing Cynthia Wu personal website: `https://sites.google.com/view/jingcynthiawu/`. We also use the shadow rate developed by Krippner (2020) for Japan, the United State, United Kingdom, and the euro area.

8. **Nominal and real exchange rates.** The nominal and real exchange rates for all countries are taken from the BIS, end of period.
9. **Oil VAR data.** All four variables for the VAR are taken from Kanzig (2021).

B Additional Figures and Tables

Figure B.1: Responses in linear specification

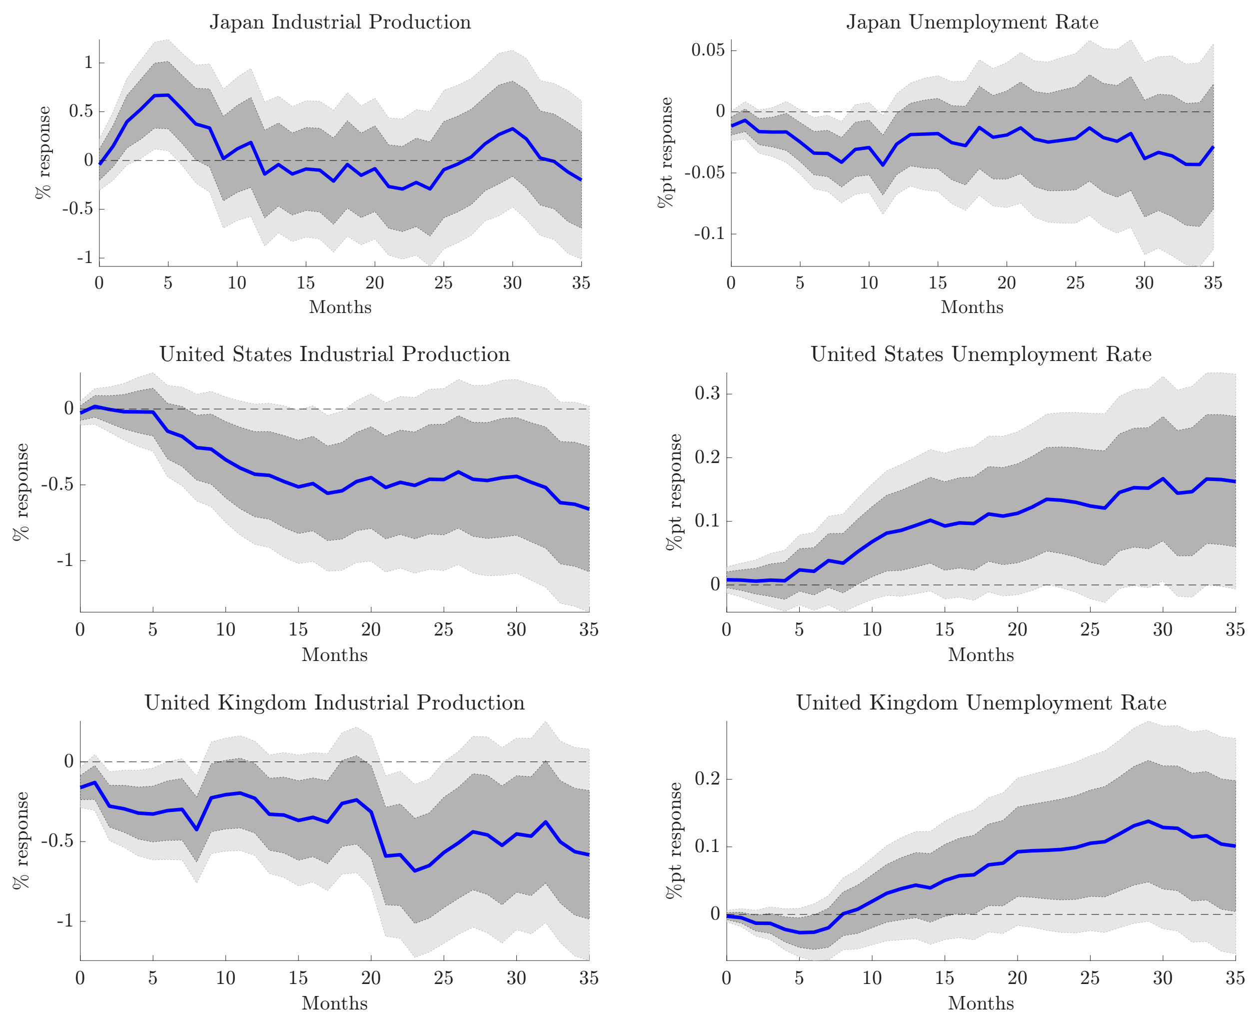

Notes: Each figure plots the impulse response functions, along with one-standard-deviation and 90% robust confidence bands, for the responses of industrial production and the unemployment rate in Japan, the United States, and the United Kingdom in a specification similar to our baseline but where we do not differentiate between ZLB and normal periods. The estimation is over the 1974:1-2019:12 sample.

Figure B.2: Japan industrial production and unemployment rate with Newey-West HAC standard errors

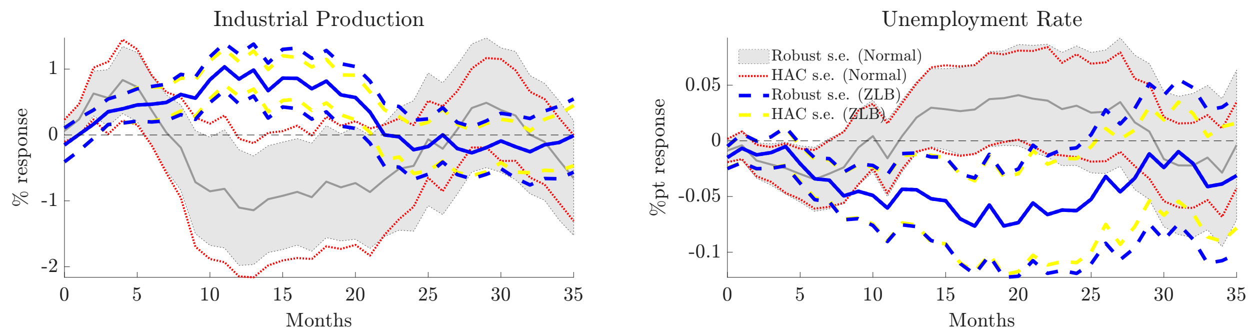

Notes: Each figure plots the impulse response functions and one-standard-deviation confidence bands of the responses of Japanese industrial production and the unemployment rate to the oil supply news shock. The red dotted lines and the yellow dashed lines are the confidence bands constructed from the Newey-West HAC standard errors, while the blue dashed and grey dotted lines show the robust confidence bands, which we report throughout the paper.

Figure B.3: Japan industrial production and unemployment rate with bootstrapped standard errors

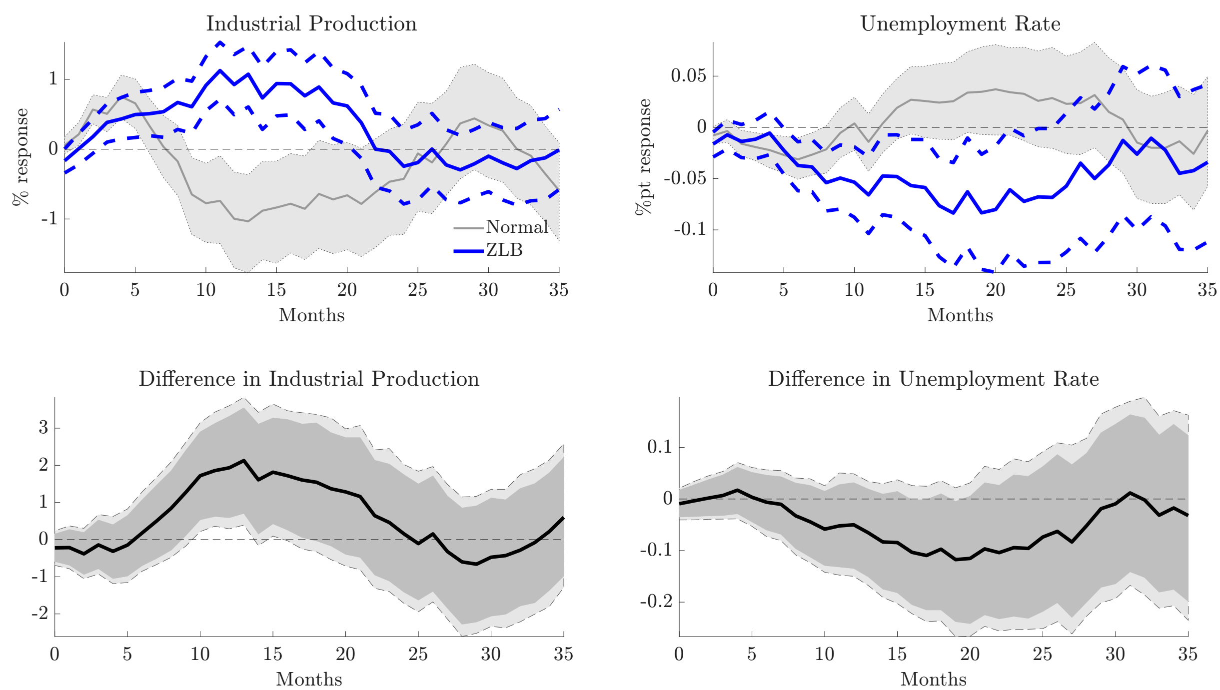

Notes: The top panel plots the impulse response functions and one-standard-deviation confidence bands of the responses of Japanese industrial production and the unemployment rate to the oil supply news shock. The lower panel plots the differences with 90 and 95% confidence bands. We implement a moving block bootstrap to account for the uncertainty of the extracted shock from the world oil VAR. We first estimate our two-step approach and obtain parameter estimates for the proxy VAR and the local projections (LP). To compute the standard errors for the local projections estimates, we first generate a series of oil supply news shocks $\widehat{\mathcal{E}}^{*}_{1,t}$ using the moving-block bootstrap (MBB) from the proxy VAR, then for each generated shocks, we generate the left-hand side variables from the LP using residual-based MBB given the estimated parameters of the LP and the generated shocks $\widehat{\mathcal{E}}^{*}_{1,t}$. We obtain similar results with the residual-based wild bootstrap and the iid bootstrap for the LP.

Figure B.4: Responses of inflation expectations in the United States in the ZLB and normal periods

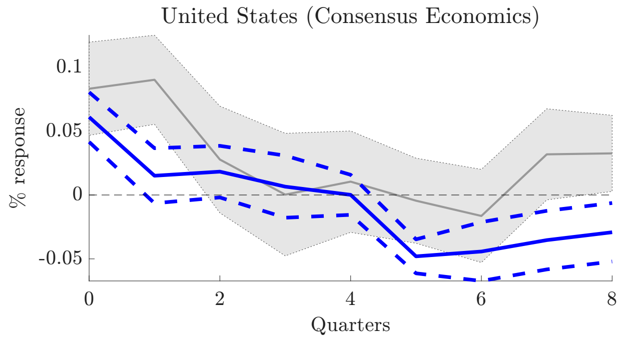

Notes: The figure plots the impulse response function and one-standard-deviation confidence bands for the professional forecasters' (Consensus Economics) inflation expectations in the Unites States in the ZLB (thick blue line) and normal (thin gray line) periods.

Figure B.5: Japan industrial production and unemployment rate: alternative oil supply shocks

(a) Kilian (2009) oil supply shocks

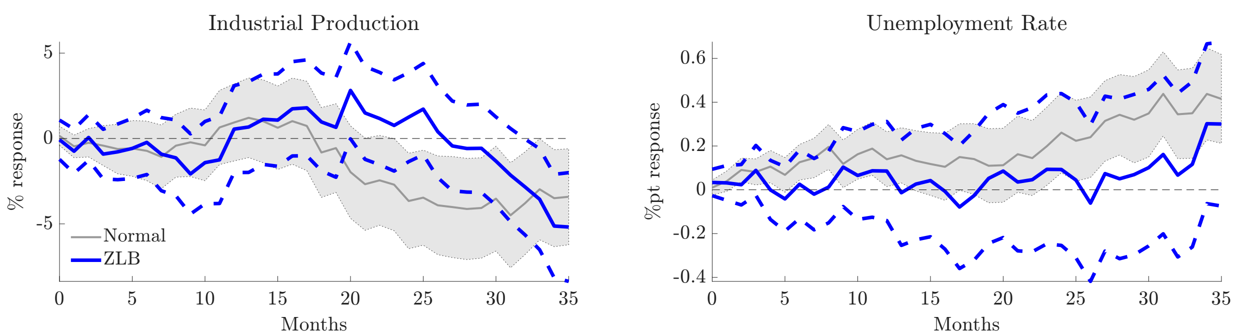

(b) Baumeister and Hamilton (2019) oil supply shocks

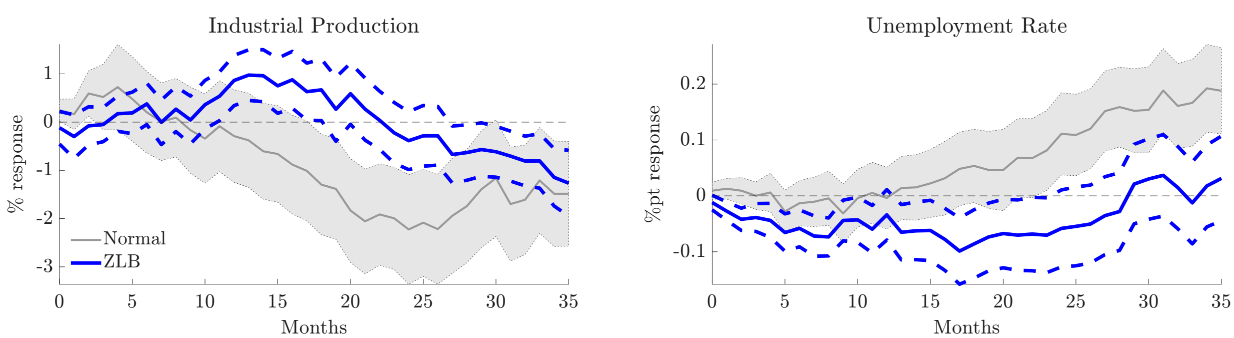

Notes: Each figure plots the impulse response functions and 68 percent confidence bands of industrial production and the unemployment rate in Japan to oil shocks identified in Kilian (2009) and Baumeister and Hamilton (2019). Figure C.5 plots the differences in the responses of these variables in the ZLB and outside of it.

Figure B.6: Real oil price response to different identified oil shocks

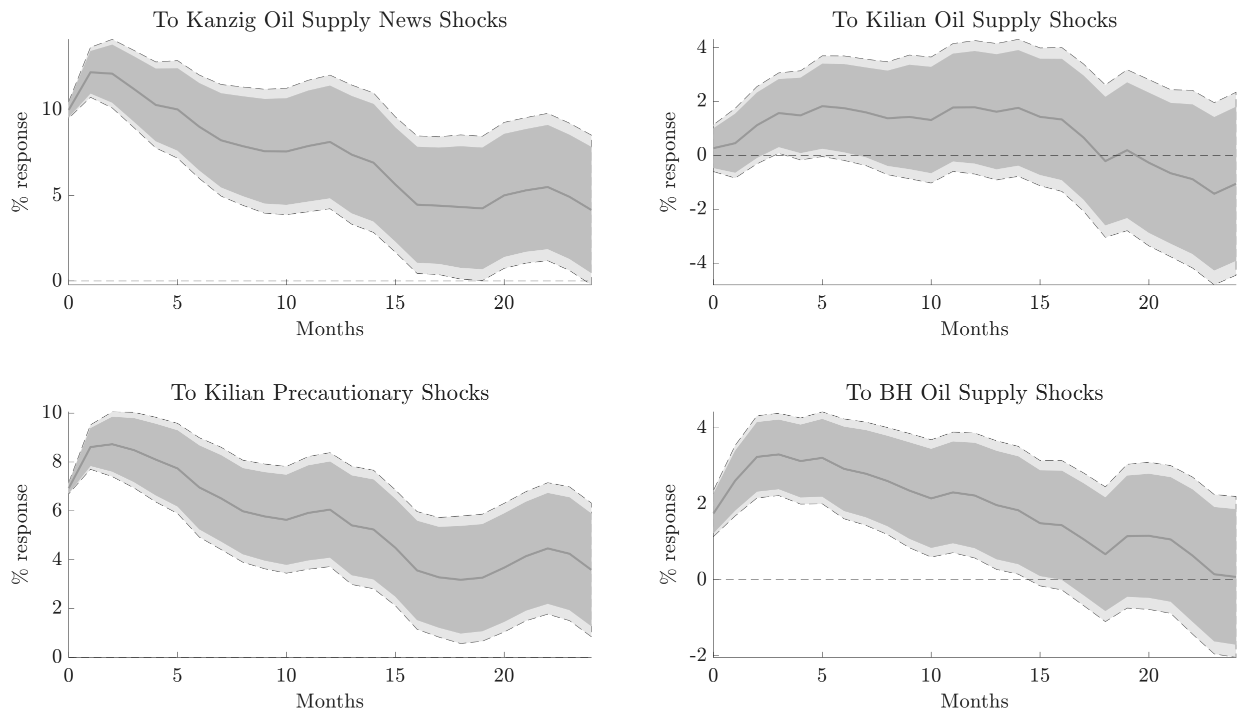

Notes: Each figure plots the impulse response functions and 95 and 90 percent robust confidence bands for the responses of oil price to oil shocks identified in Kanzig (2021) (oil supply news shock), Kilian (2009) (oil supply and oil-market specific demand shocks), and Baumeister and Hamilton (2019) (oil supply shocks). In the case of the Kanzig (2021) shocks, we use the whole sample of 1974:1 to 2019:12 to estimate the response of oil prices.

Figure B.7: Japan impulse responses to Kilian (2009) oil-market specific demand shocks

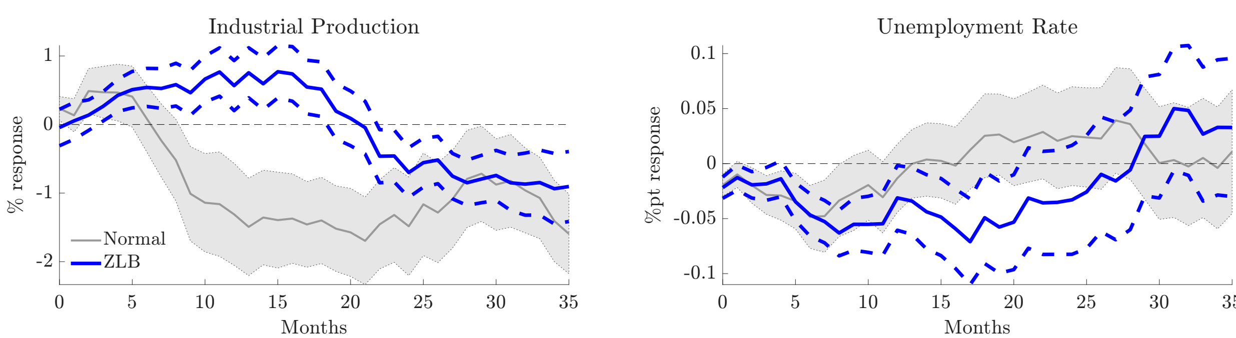

Notes: Each figure plots the impulse response functions and 68 percent robust confidence bands for the responses of Japanese industrial proudction and unemployment rate to oil price shocks identified in Kilian (2009) oil-market specific demand shocks. Figure C.6 plots the differences in the responses of these variables in the ZLB and outside of it.

Figure B.8: Non-oil output responses in Japan, the United Kingdom, and the United States

(a) Manufacturing Industrial Production

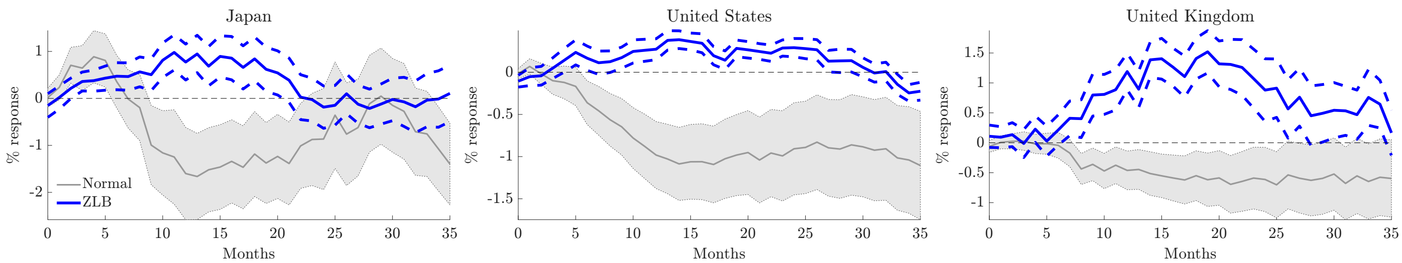

(b) Non-Oil Industrial Production

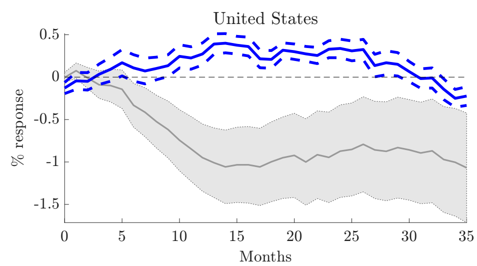

Notes: This figure plots the impulse response functions and one-standard-deviation robust confidence bands of the responses of manufacturing industrial production in Japan, the Unites States, and the United Kingdom, and non-oil industrial production in the United States (lower panel) in the ZLB (thick blue line) and outside of it (think gray line) to an oil supply news shock. Formally, we estimate equation (**??**) using data from 1975:1 to 2019:12.

Figure B.9: The S&P 500 responses to oil supply news shocks and surprises

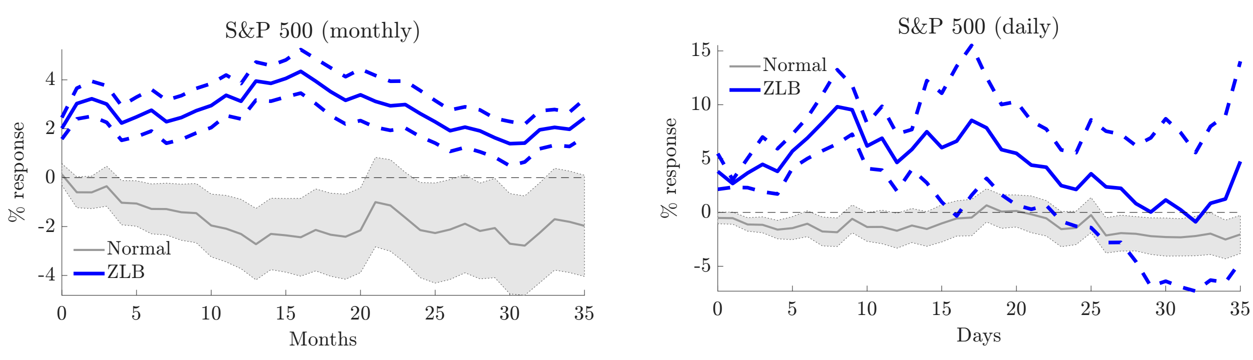

Notes: This figure plots the impulse response functions and one-standard-deviation robust confidence bands of the responses of the S&P 500 index to the oil supply news shock (left) and to the oil supply surprise series (right). The left panel presents the results based on the specification (**??**) using monthly data, while the right panel uses daily data. The daily estimation uses 10 lags of the daily stock price as the only controls.

Figure B.10: Real exchange rate responses

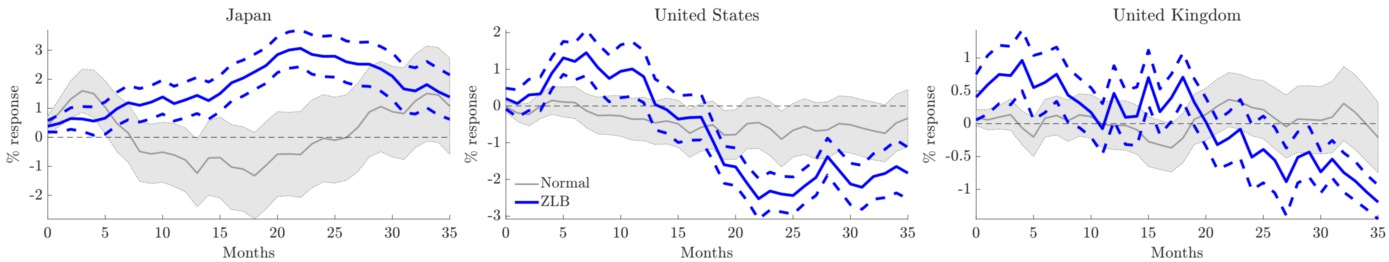

Notes: Each figure plots the impulse response functions and one-standard-deviation robust confidence bands for the real effective exchange rates of Japan, the United States, and the United Kingdom. An increase in the real exchange rate corresponds to a depreciation. The corresponding differences are shown in Figure C.7.

Figure B.11: Shadow rate impulse responses in the ZLB period

(a) US shadow rate constructed by Wu and Xia (2016)

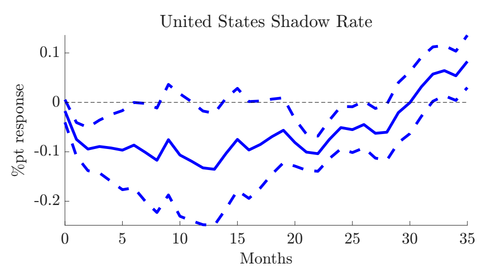

(b) Shadow rates constructed by Krippner (2020)

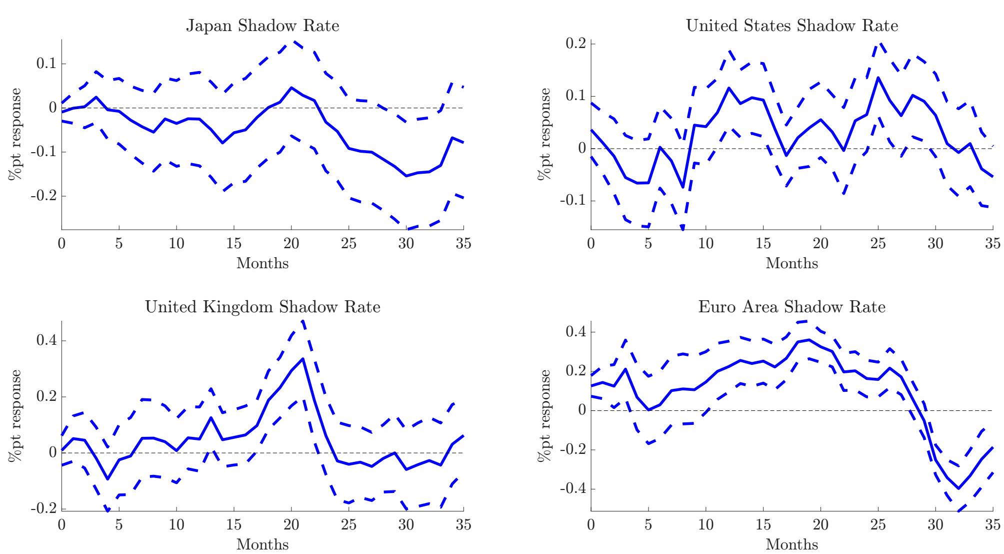

Notes: This figure plots the impulse response function and one-standard-deviation robust confidence bands of the responses of the shadow interest rates in the Unites States, Japan, the United Kingdom and the euro area in the ZLB.

Figure B.12: Impulse responses with Krippner (2020) shadow rates as controls

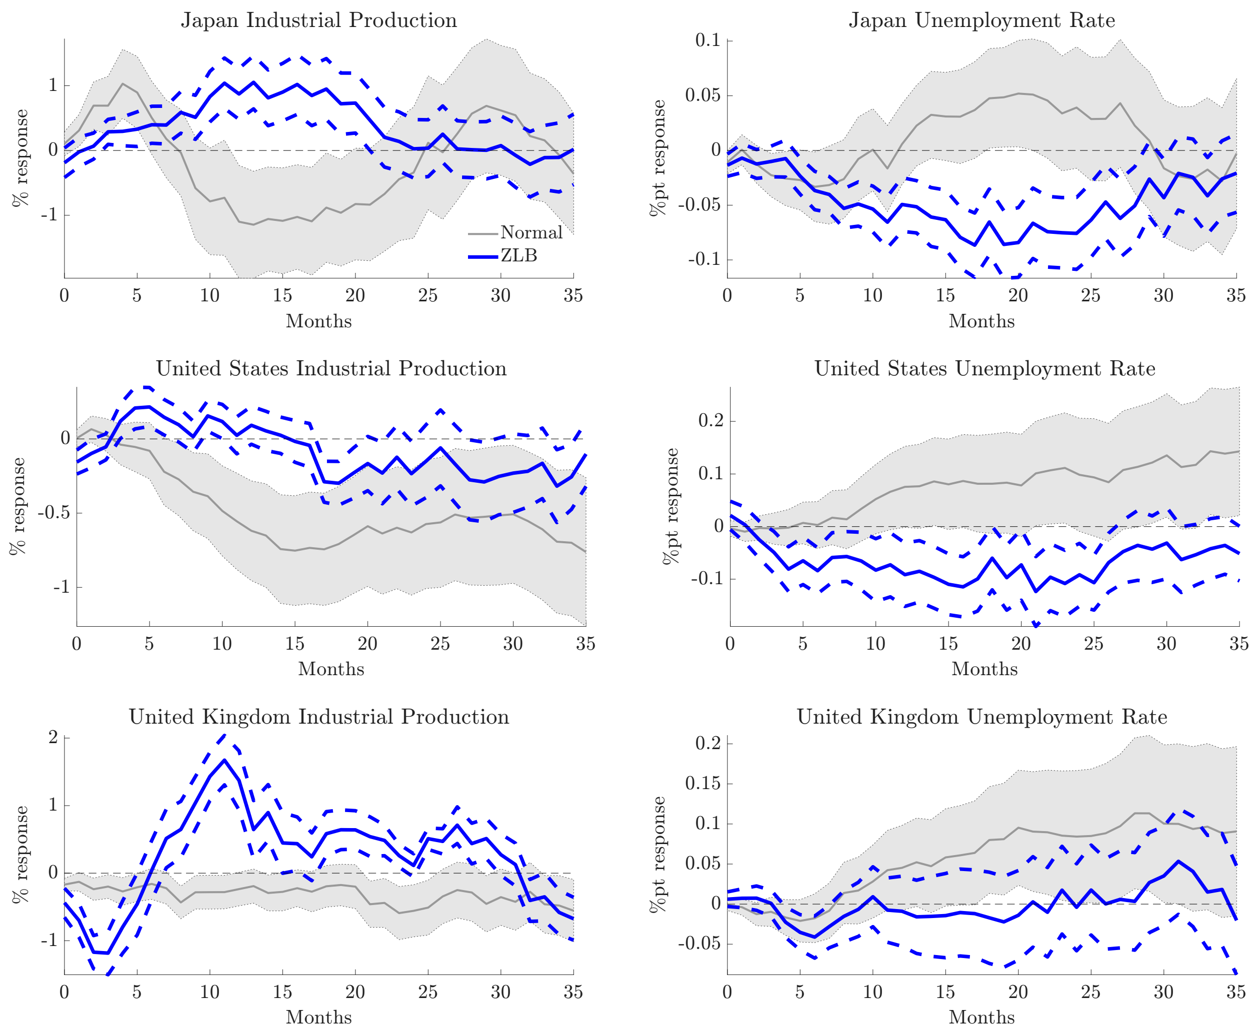

Notes: Each figure plots the impulse response functions and one-standard-deviation HAC confidence bands for the responses of industrial production and the unemployment rate in Japan, the U.S., and U.K., adding controls for the shadow rate at the ZLB and for short-term interest rates in normal times to the baseline specification.

Figure B.13: Impulse responses with leads of the differences between shadow rate and nominal rate

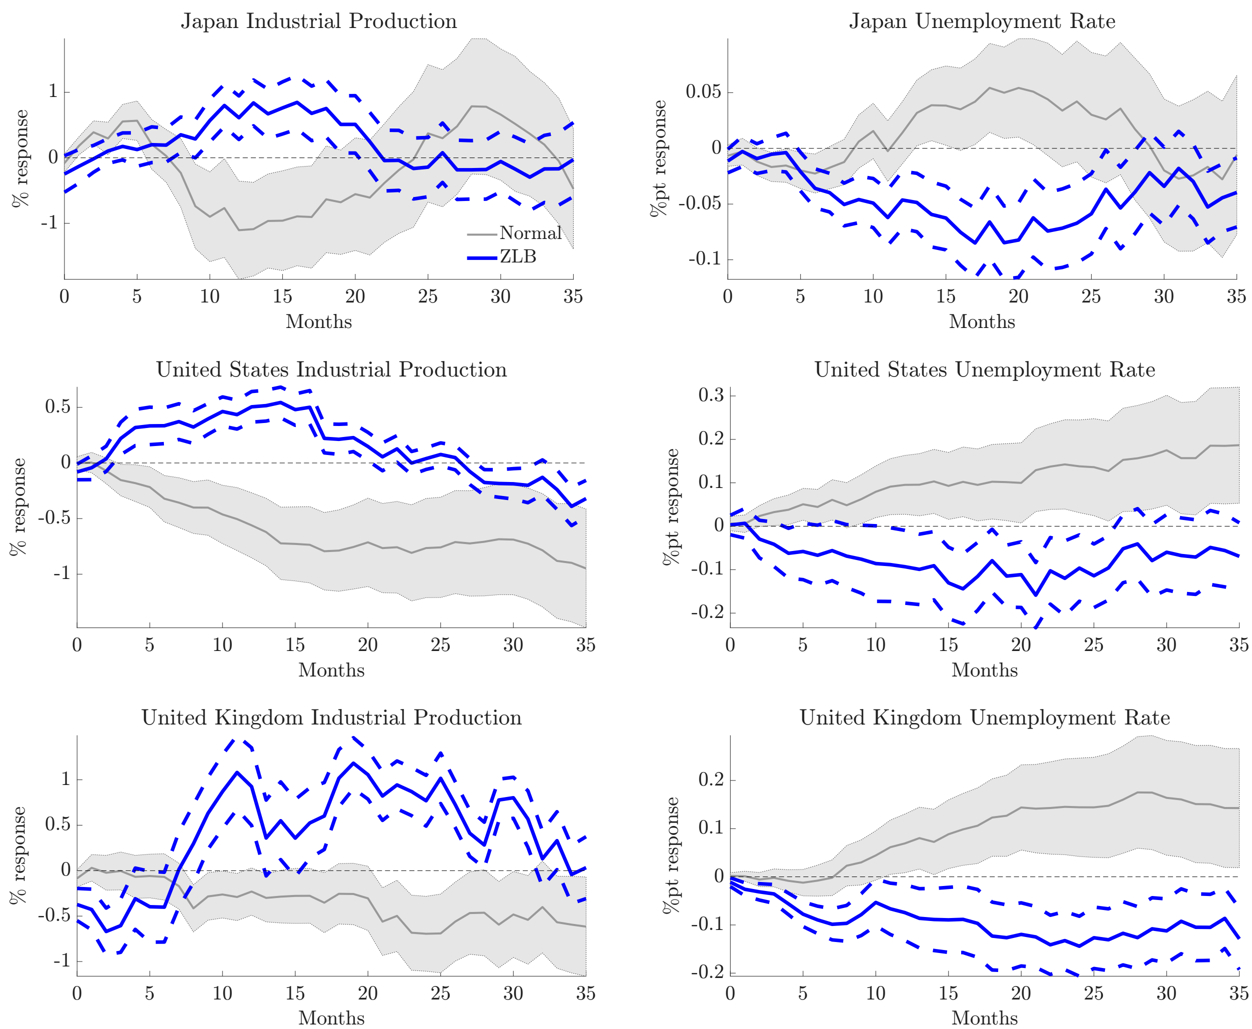

Notes: Each figure plots the impulse response functions and one-standard-deviation HAC confidence bands of the responses of industrial production and the unemployment rate adding controls for 12 leads of the differenece between the Krippner (2020) shadow rate and nominal rate in the ZLB period in Japan, the U.S., and U.K..

Figure B.14: The short-term nominal interest rates in Canada and the euro area

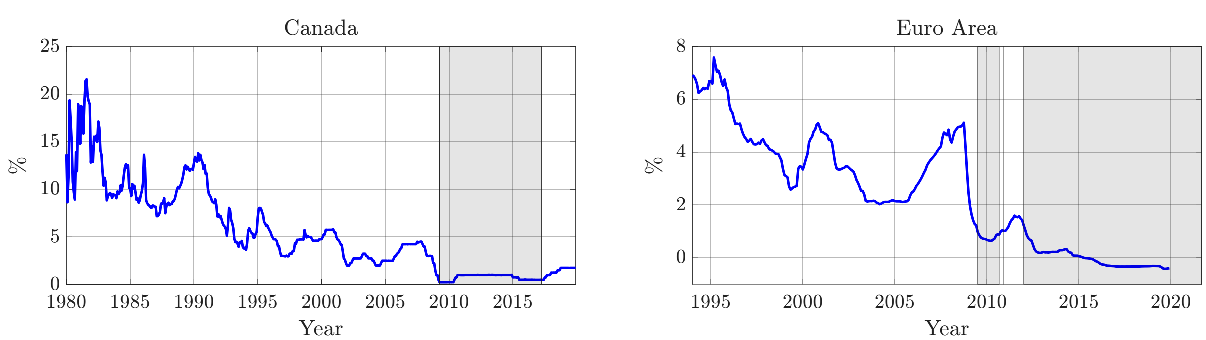

Notes: Each panel presents central bank's policy rate for each country. The shaded areas are the zero lower bound periods defined in Section 5.3.

Figure B.15: Industrial production and unemployment rate responses in Canada and the euro area

(a) Industrial production

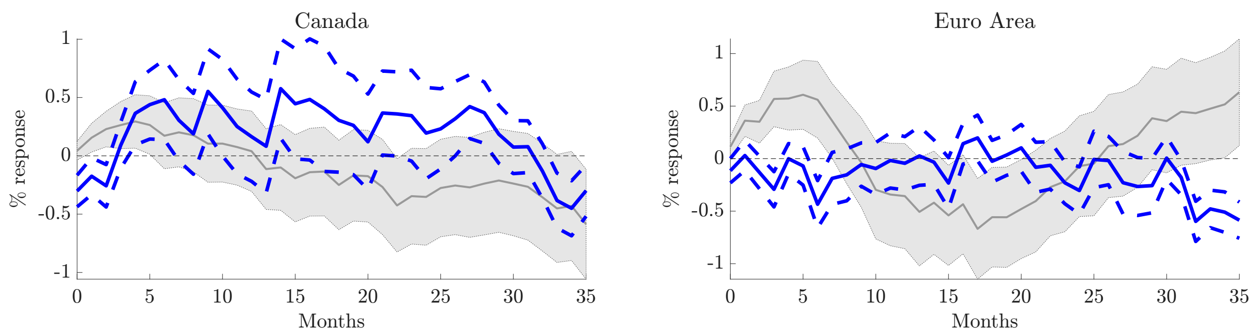

(b) Unemployment rate

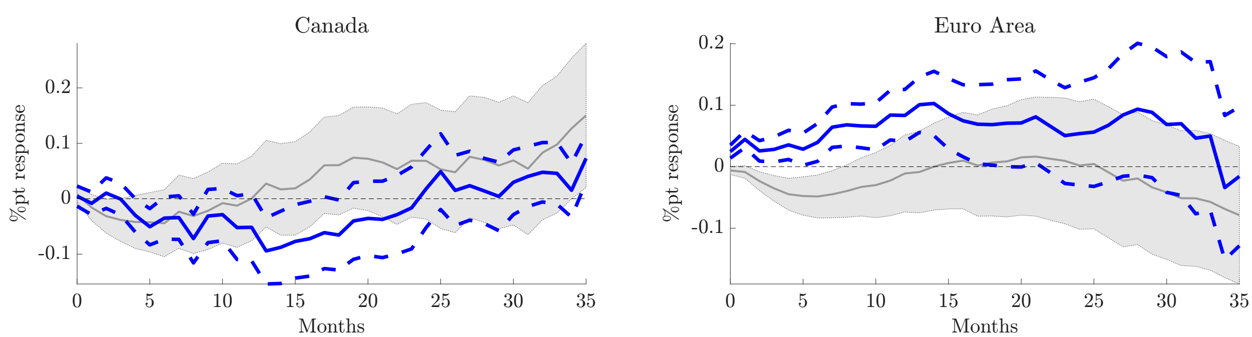

Notes: Each figure plots the impulse response functions and one-standard-deviation confidence bands of the responses. Panel (a) shows industrial production responses in Canada and the euro area in the ZLB (thick blue line) and in the normal (thin gray line) periods. Panel (b) plots unemployment rate responses in Canada and the Euro Area in the ZLB (thick blue line) and in the normal (thin gray line) periods.

Figure B.16: Responses in other countries: Korea, Mexico, Sweden, and Norway

(a) Republic of Korea

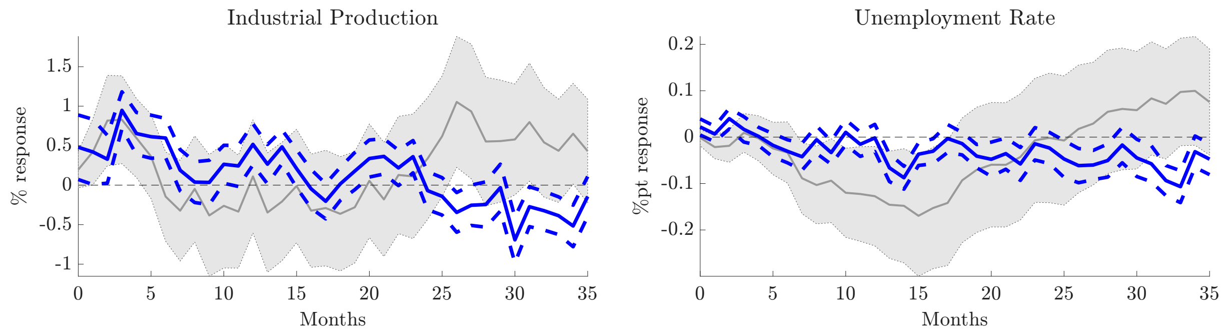

(b) Mexico

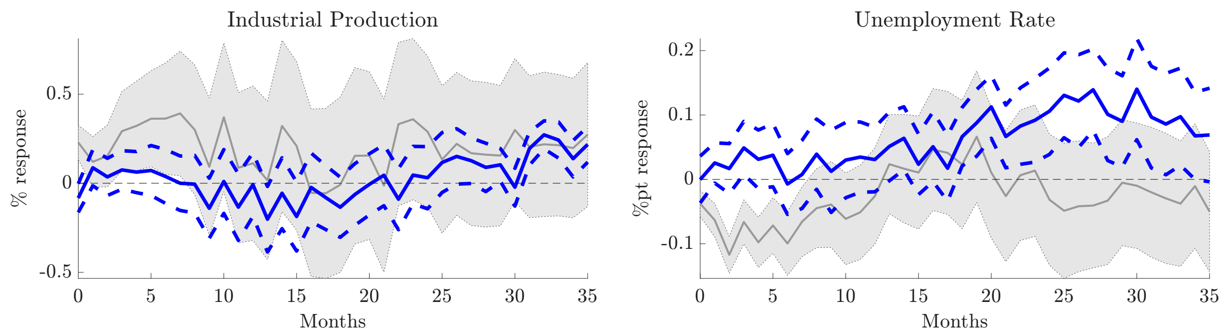

(c) Sweden

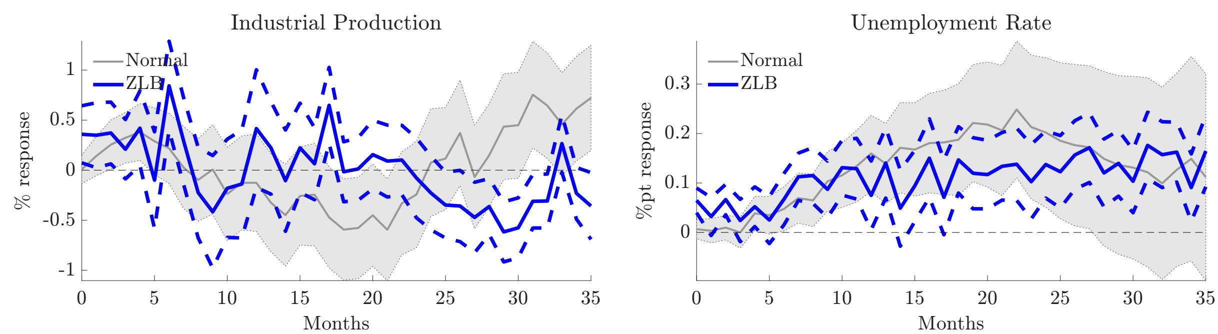

(d) Norway

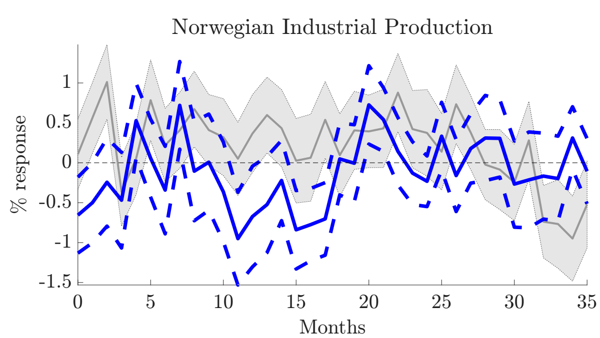

Notes: Each figure plots the impulse response function and one-standard-deviation robust confidence bands of the responses of industrial production (left panels) and the unemployment rate (right panels) in Republic of Korea, Mexico, Sweden, and Norway in the ZLB (thick blue line) and in the normal (thin gray line) periods, where the ZLB period is defined based on the behavior of the short term nominal interest rate in the US.

Figure B.17: Industrial production and unemployment rate responses in Japan: 1995:10-2006:6

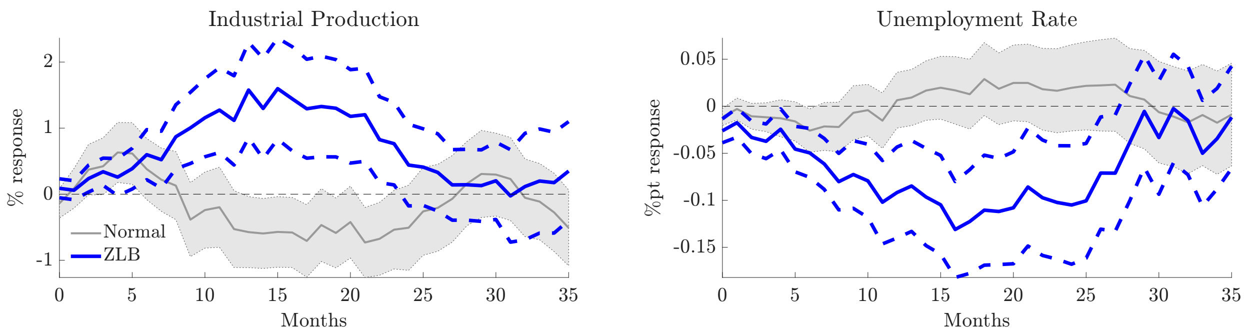

Notes: The impulse responses and one-standard-deviation robust confidence bands of industrial production and unployment rate to an oil supply news shock in the ZLB (think blue line) and normal (thin gray line) periods. The local projections estimation period is from 1975:1 to 2006:6.

Figure B.18: Impulse responses to VAR shocks estimated up to 2006:6

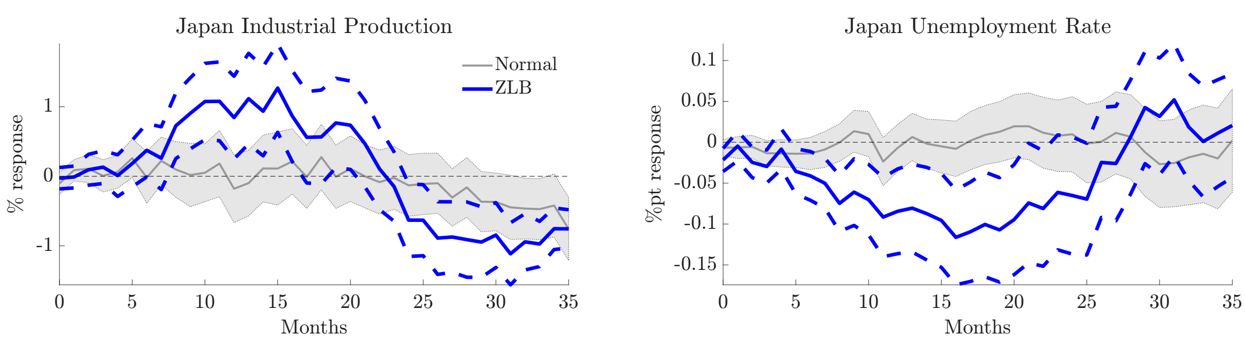

Notes: Each figure plots the impulse response functions and one-standard-deviation robust confidence bands of the responses of Japanese industrial production and the unemployment rate when the shocks are extracted from an oil proxy VAR estimated over the 1974:1-2006:6 sample, to avoid possible shifts in global oil market.

Figure B.19: Industrial production and unemployment rate responses in Japan when 2008:3-2008:7 is dummied out

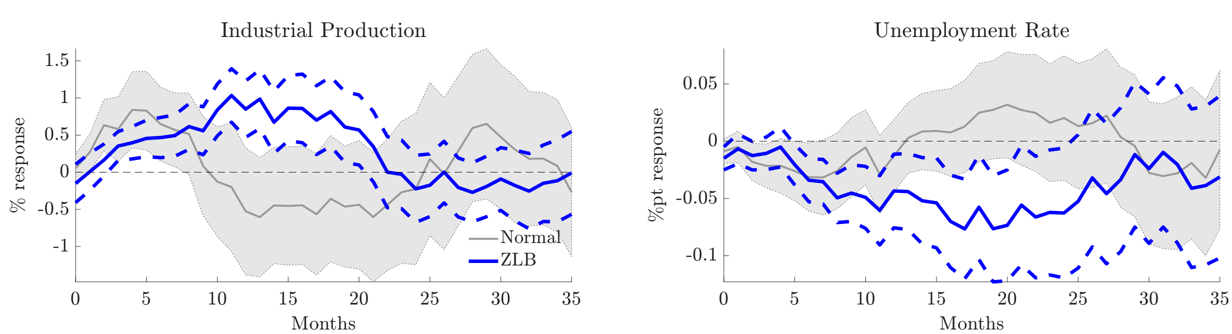

Notes: The figures plot impulse response functions and one-standard-deviation robust confidence bands of industrial production (left) and unployment rate (right) to an oil supply news shock in the ZLB (thick blue line) and normal (thin gray line) periods. The results based on specification when the 2008:3-2008:7 period is dummied out.

Figure B.20: Industrial production and unemployment rate responses in the US: the 1986:1-2019:12 sample

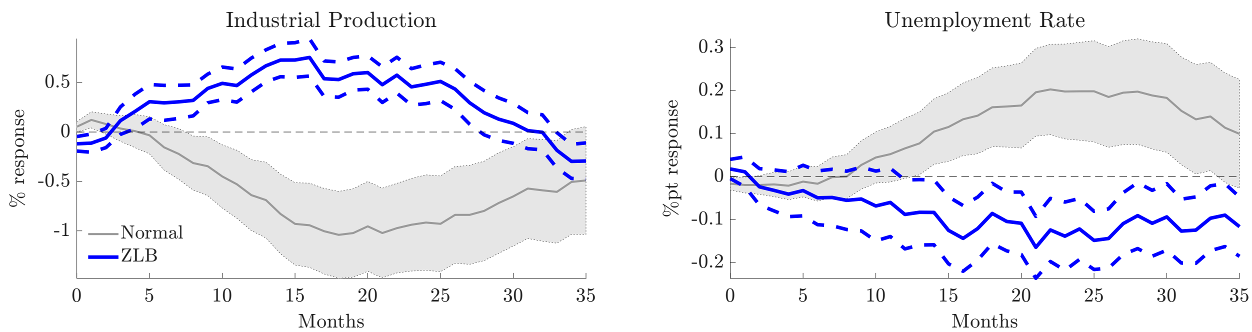

Notes: The figures plot the impulse response functions and one-standard-deviation robust confidence bands of industrial production (left) and unployment rate (right) to an oil supply news shock in the ZLB (thick blue line) and normal (thin gray line) periods in the United States. The sample is from 1986:1 to 2019:12.

Figure B.21: Japanese industrial production and unemployment rate responses: no controls

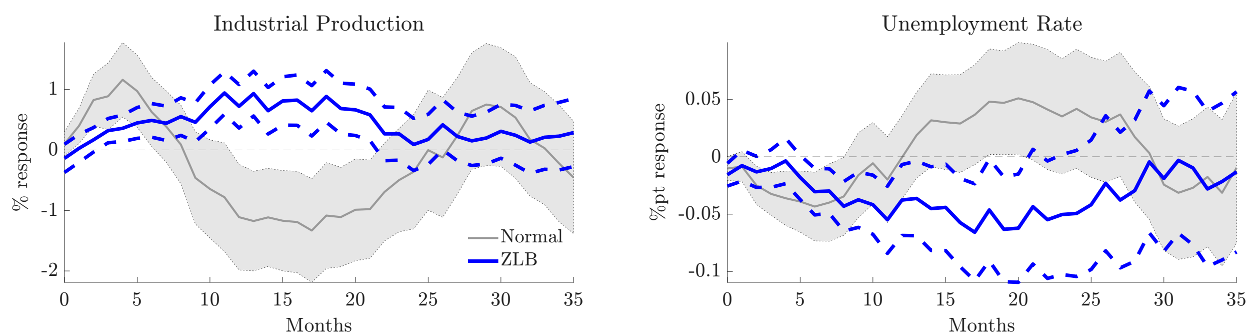

Notes: Each figure plots the impulse response functions and one-standard-deviation robust confidence bands of the responses of Japanese industrial proudction and unemployment rate to an oil supply news shock in the ZLB (thick blue line) and the normal (thin gray line) periods. The controls only include a constant and lags of dependent variable, unlike the baseline specification that features the unemployment rate and the lags of the oil supply news shocks.

Figure B.22: Japanese industrial production and unemployment rate responses: additional controls

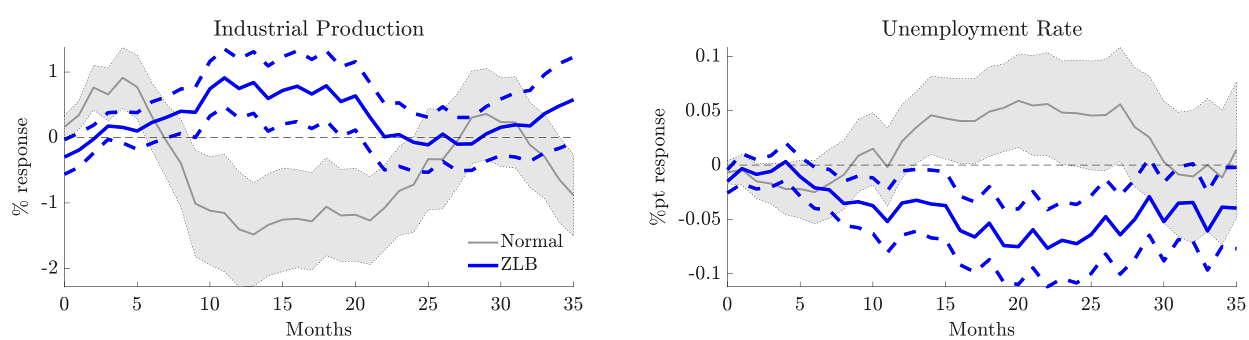

Notes: Each figure plots the impulse response functions and one-standard-deviation robust confidence bands of the responses of Japanese industrial proudction and unemployment rate to an oil supply news shock in the ZLB (thick blue line) and the normal (thin gray line) periods. In addition to the baseline specification control variables (the lagged left-hand variable, unemployment rate, and the oil supply news shocks), we add two nominal variables to the list of controls: the lags of inflation rate and lags of 5-year nominal yield on government bonds.

Figure B.23: Japanese industrial production and unemployment rate responses: additional oil VAR controls

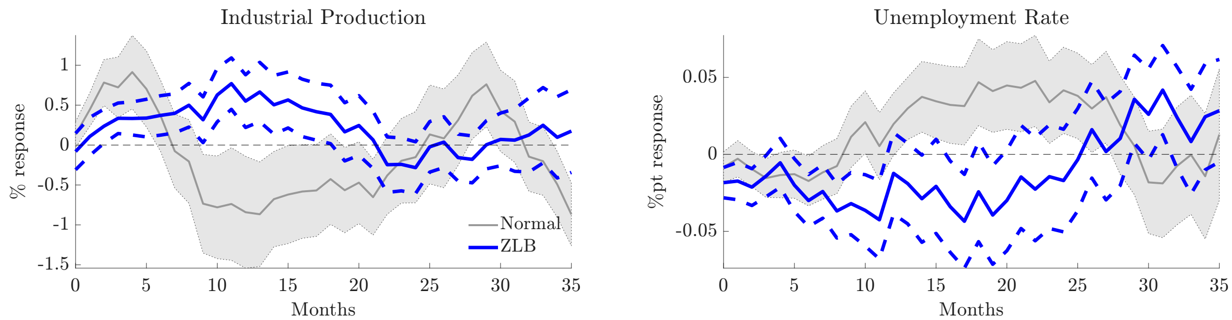

Notes: Each figure plots the impulse response functions and one-standard-deviation robust confidence bands of the responses of Japanese industrial proudction and unemployment rate to an oil supply news shock in the ZLB (thick blue line) and the normal (thin gray line) periods. In addition to the baseline specification control variables (the lagged left-hand variable, unemployment rate, and the oil supply news shocks), we add 12 lags for each of the variables in the oil market VAR interacted with the ZLB indicator.

Figure B.24: Impulse responses with fossil fuel imports as controls

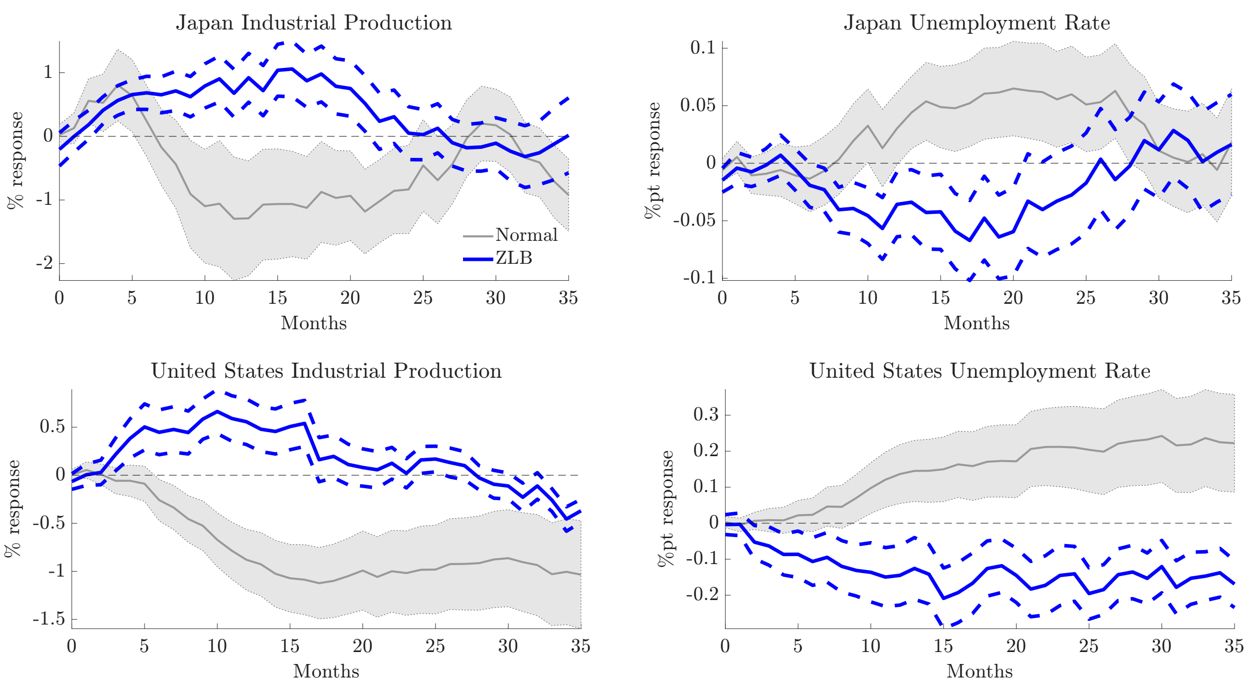

Notes: Each figure plots the impulse response functions and one-standard-deviation HAC confidence bands for the responses of industrial production and the unemployment rate in Japan, with lags of fossil fuel imports added to the baseline specification.

Figure B.25: Industrial production and unemployment rate responses in Japan (without info revelation)

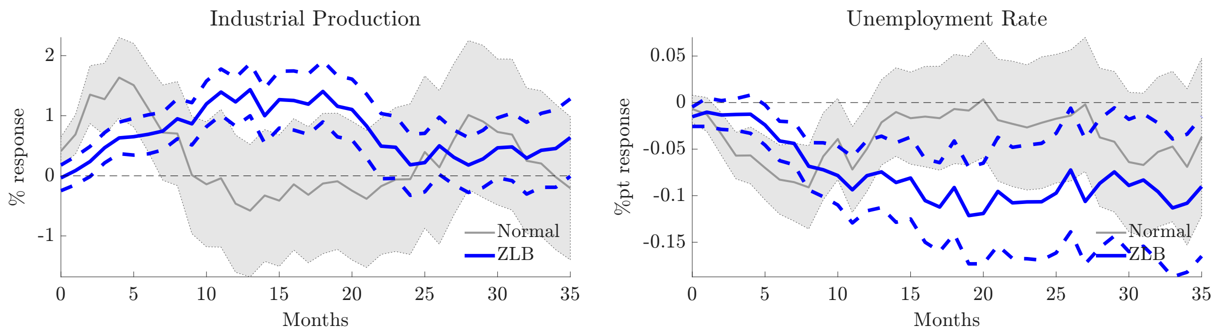

Notes: Each figure plots the impulse response functions and one-standard-deviation confidence bands of industrial production (left) and unployment rate (right) to a modified oil supply news shock in the ZLB (thick blue line) and the normal (thin gray line) periods. A modified oil supply news shock was obtained by using only a part of the oil supply surprise series that generates a negative comovement between oil prices and the U.S. stock market index (Degasperi, 2021).

Figure B.26: Japanese industrial production and unemployment rate responses in booms and busts

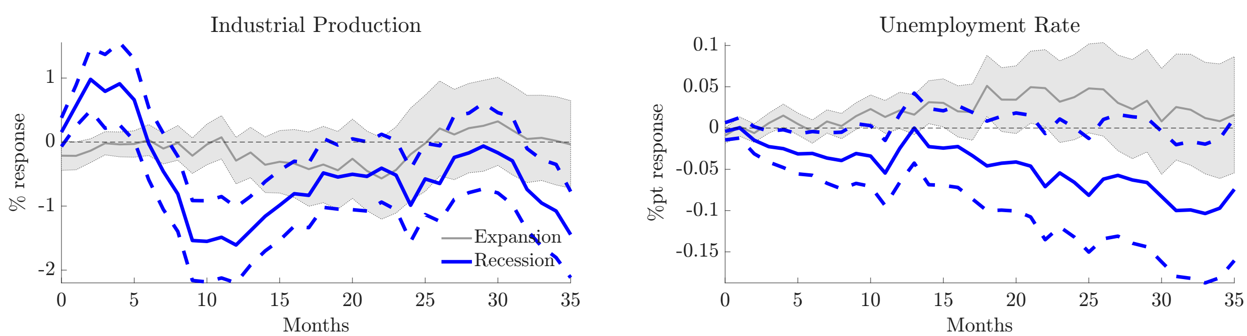

Notes: Each figure plots the impulse response functions and one-standard-deviation confidence bands of the responses of Japanese industrial production and the unemployment rate in recessions and booms following an oil supply news shock. The difference is plotted in Figure C.8.

Figure B.27: Japanese Industrial Production and Unemployment Rate Responses to purified oil surprise series

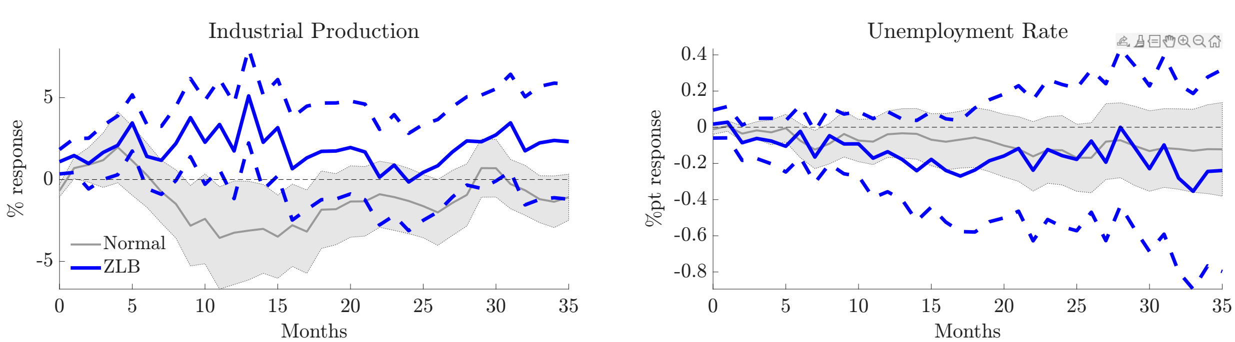

Notes: Each figure plots the IRFs and 68% robust confidence bands of the responses of Japanese industrial proudction and unemployment rate to an oil supply news shock. Oil supply shocks are residuals from regressing the oil supply surpise series on twelve lags of four variables from the world oil VAR.

Figure B.28: Japan responses with real oil prices deflated by Japanese CPI

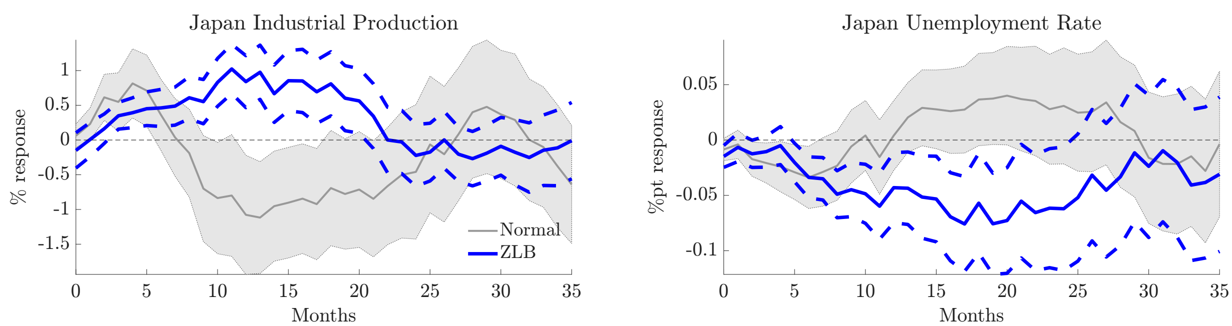

Notes: Each figure plots the impulse response functions and one-standard-deviation HAC confidence bands of the responses of Japanese industrial production and the unemployment rate when real oil prices are nominal oil prices deflated by the Japanese CPI.

C Additional Figures: Differences in Estimates

Figure C.1: Difference in quarterly impulse responses in Japan

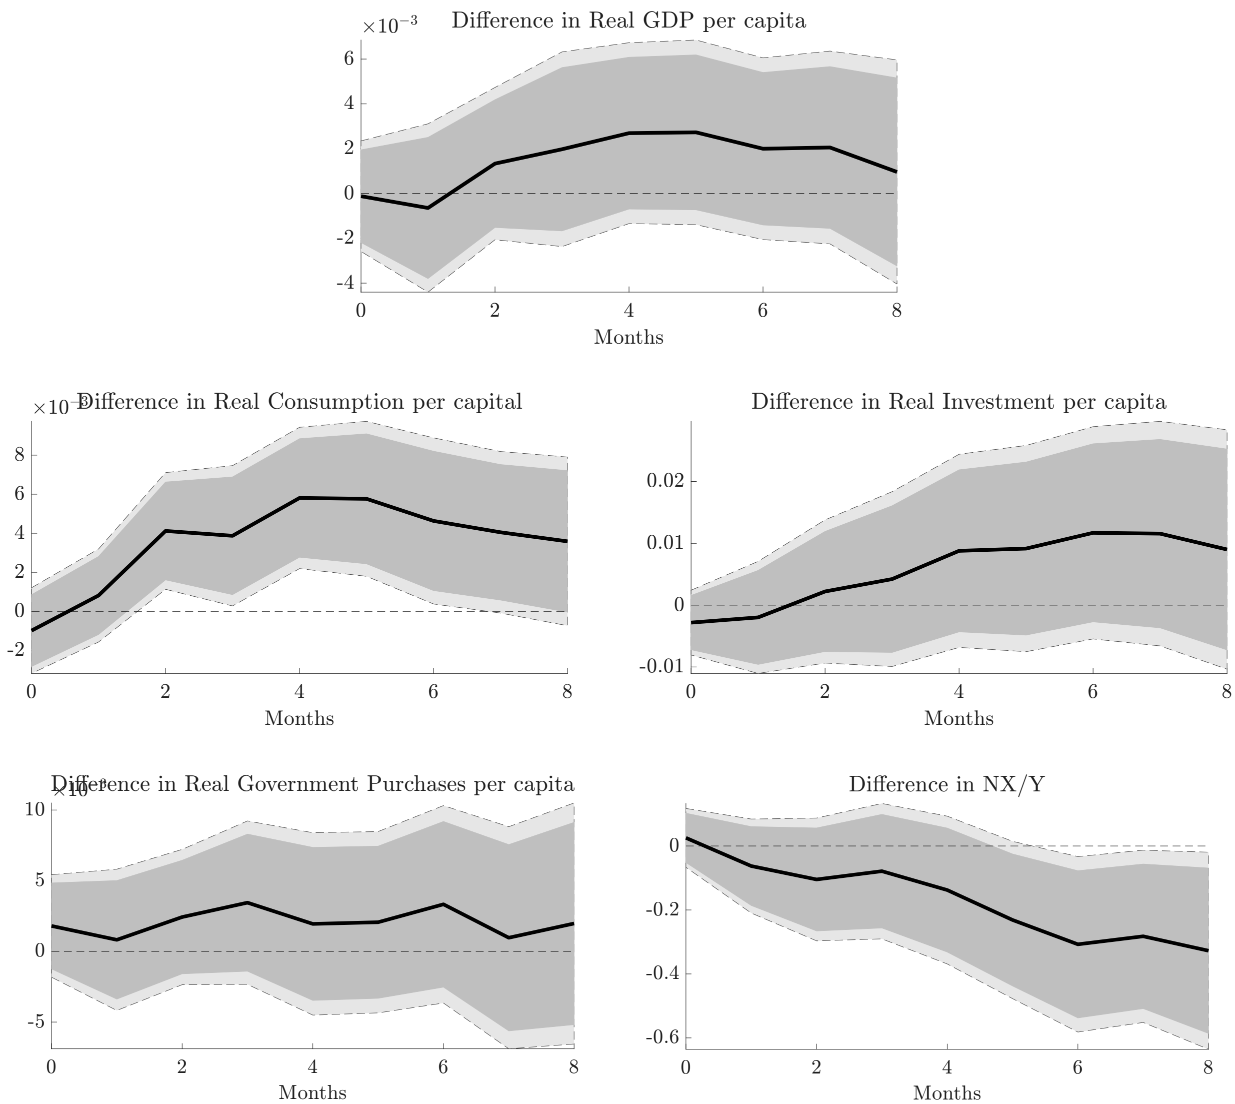

Notes: Each figure plots the difference between the responses of real GDP, real consumption, real investment, real government purchases (all per capita), and nominal trade balance over nominal GDP in Japan to the oil supply news shock across the ZLB and in the normal periods. The estimated specification is in equation presented in Section 2.1. The sample is from 1975Q1 to 2019Q4. Each plot also presents the 90-percent (dark grey) and 95-percent (light grey) confidence bands. The levels of the impulse responses are plotted in Figure 4.

Figure C.2: Difference in interest rates, inflation, and inflation expectations responses in Japan

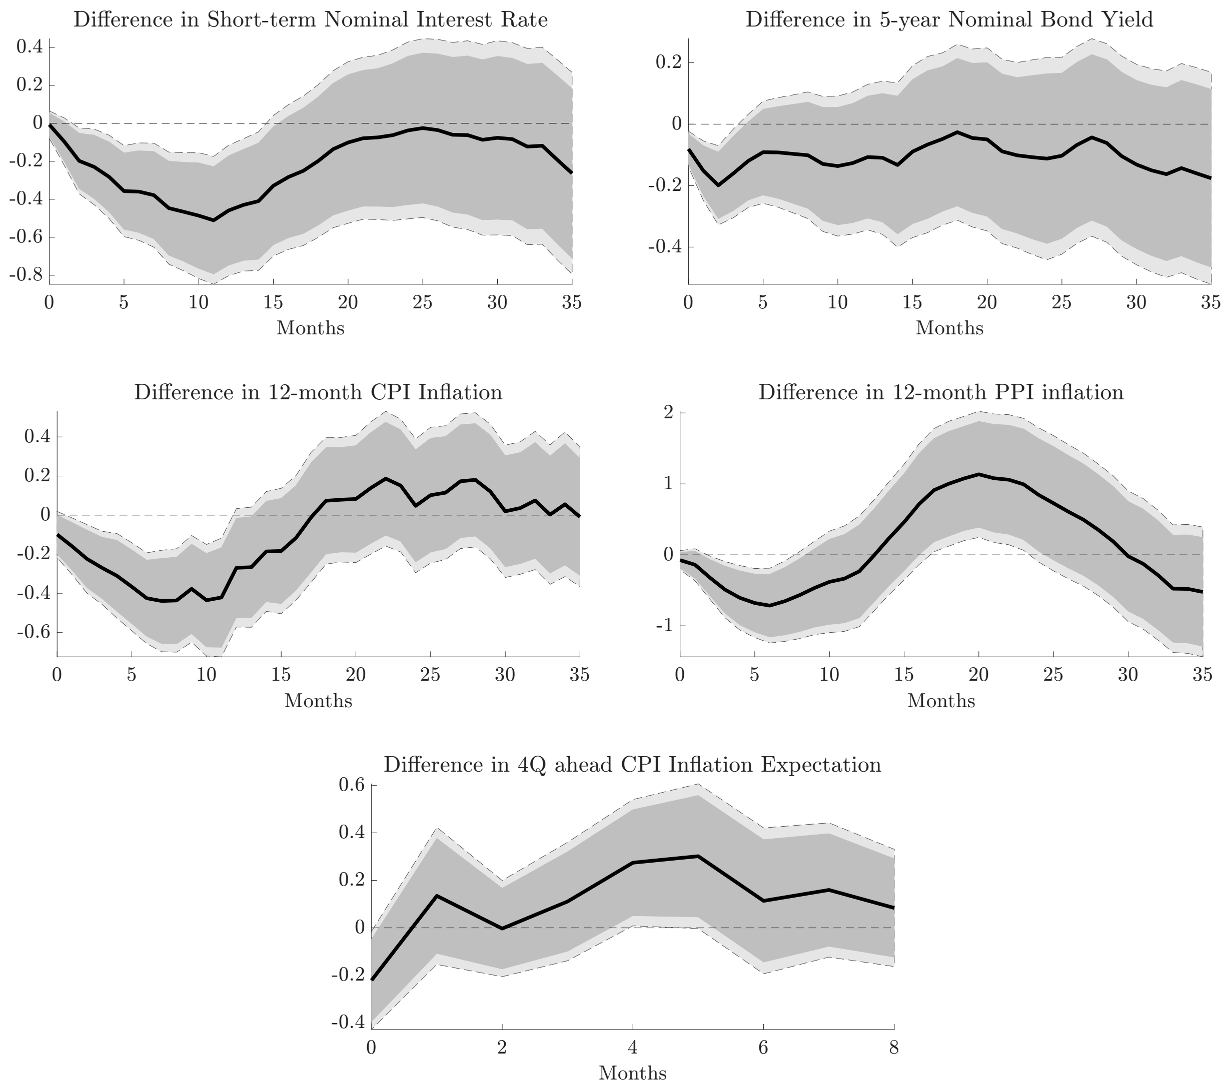

Notes: The figures plot the difference in impulse responses of the short-term nominal interest rate, 5-year nominal interest rate, CPI inflation rate, PPI inflation rate, and CPI inflation expectations between normal and ZLB periods in Japan. The errors bands are 90- and 95-percent confidence intervals. The corresponding impulse responses in the ZLB and normal periods are in Figure 5.

Figure C.3: Difference in industrial production and unemployment rate responses in the United States and the United Kingdom

(a) Difference in the industrial production responses

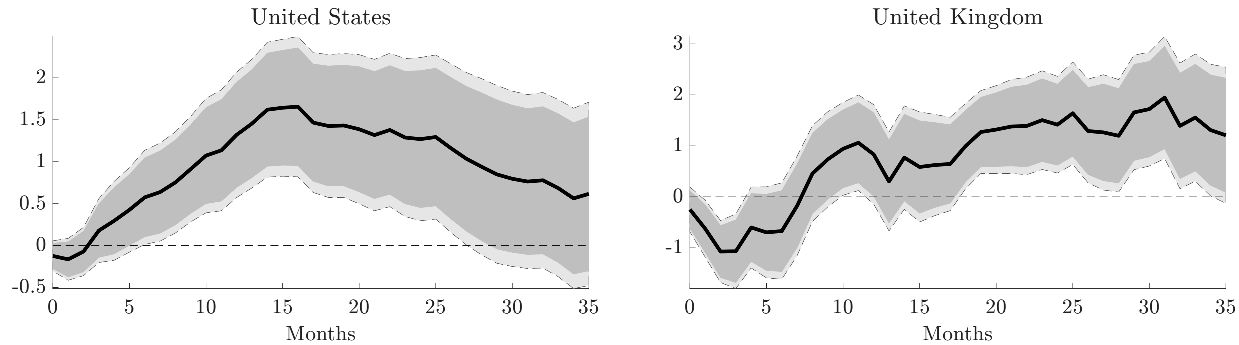

(b) Difference in the unemployment rate responses

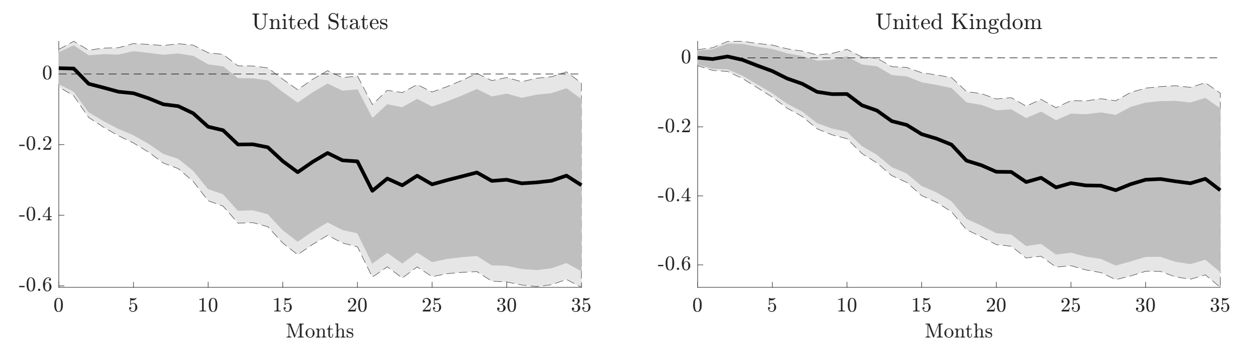

Notes: The figures plot the differences in the impulse responses of industrial production (top panel) and the unemployment rate (bottom panel) in the United States and the United Kingdom. The error bands represent 90 and 95 percent confidence intervals. The corresponding impulse responses in the ZLB and normal periods are shown in Figure 6.

Figure C.4: Difference in interest rates and inflation in the United States and the United Kingdom

(a) Short-term nominal interest rate difference

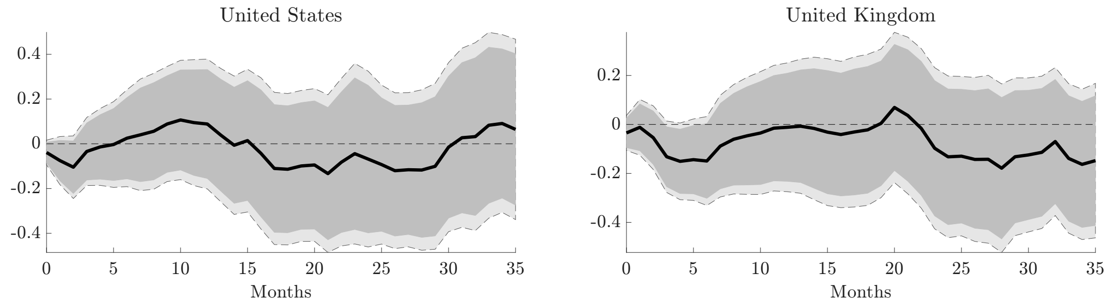

(b) 5-year nominal government bond yield difference

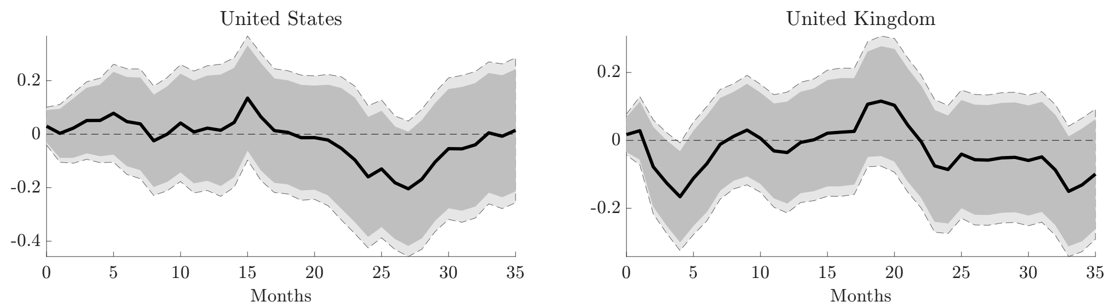

(c) CPI inflation rate difference

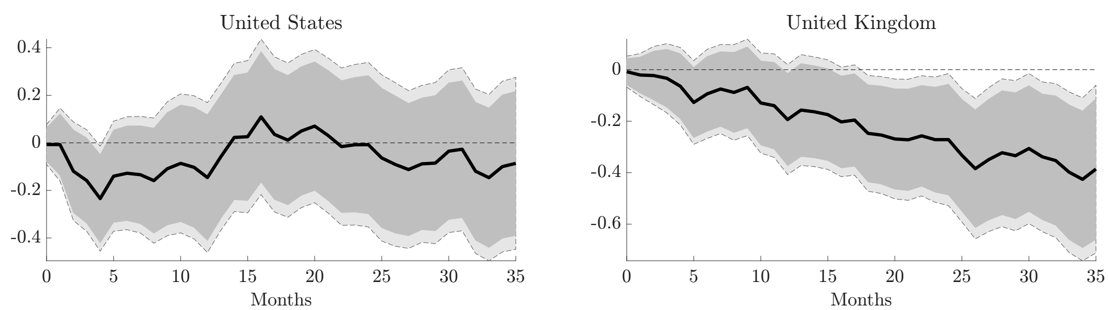

(d) Household inflation expectations difference

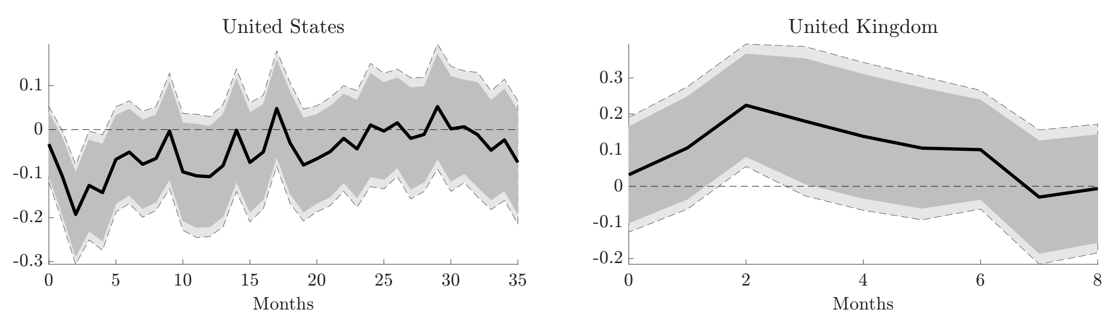

Notes: Each figure plots the mean, 90- and 95-percent confidence intervals of the difference in the responses of the central bank nominal interest rate, the 5-year nominal yield on government bonds, and the CPI inflation, and the household inflation expectations in the Unites States and the United Kingdom in the ZLB (thick blue line) and in the normal (thin grey line) periods. All the units on thy Y axes are percentage points. The levels of the impulse responses are plotted in Figure 7.

Figure C.5: Difference in responses of Japan industrial production and unemployment rate: different oil supply shocks

(a) Kilian (2009) oil supply shocks

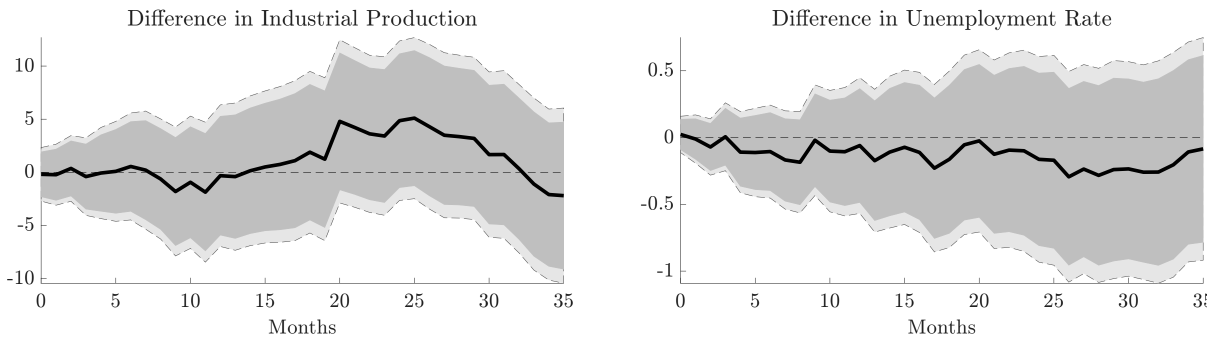

(b) Baumeister and Hamilton (2019) oil supply shocks

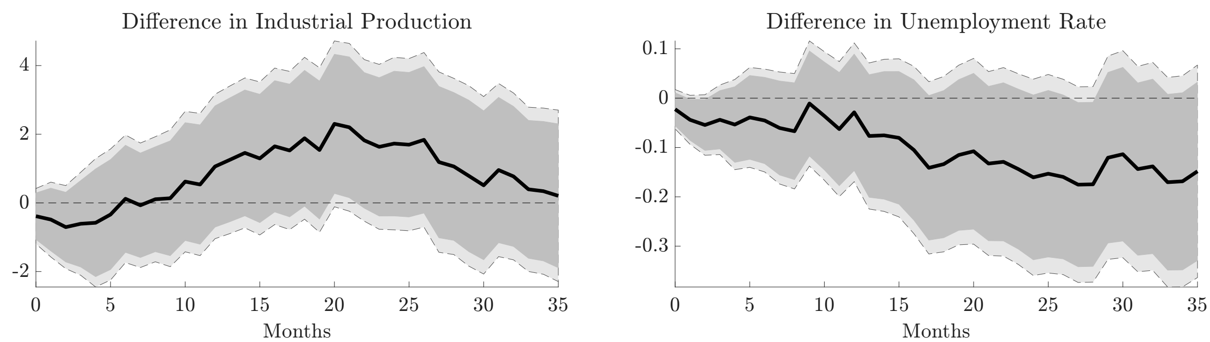

Notes: Each figure plots the mean, 90- and 95-percent confidence intervals of the difference in the responses of industrial production and the unemployment rate in Japan to oil shocks identified as in Kilian (2009) and Baumeister and Hamilton (2019). The levels of the impulse responses are plotted in Figure B.5.

Figure C.6: Difference in responses of Japan industrial production and unemployment rate: oil-market specific demand shock

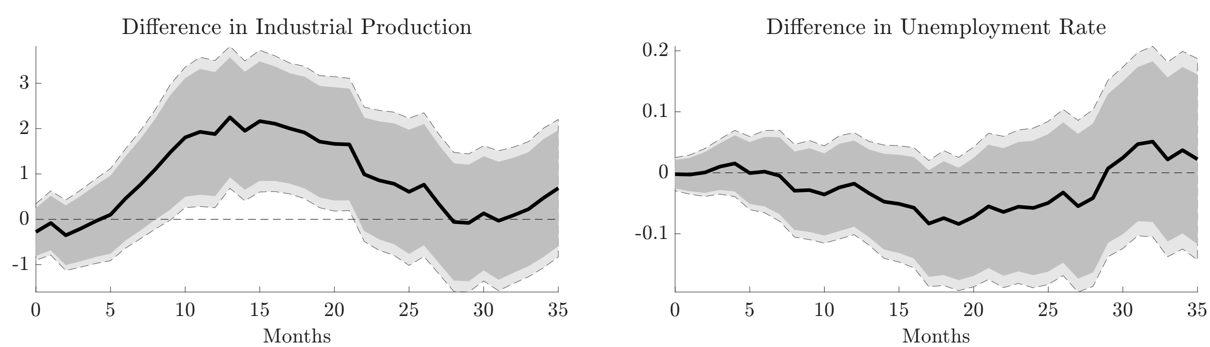

Notes: Each figure plots the mean, 90 and 95 percent confidence intervals of the difference in the responses of industrial production and the unemployment rate in Japan to oil-market specific demand shocks identified as in Kilian (2009). The levels of the impulse responses are plotted in Figure B.7.

Figure C.7: Difference in the real effective exchange rate responses

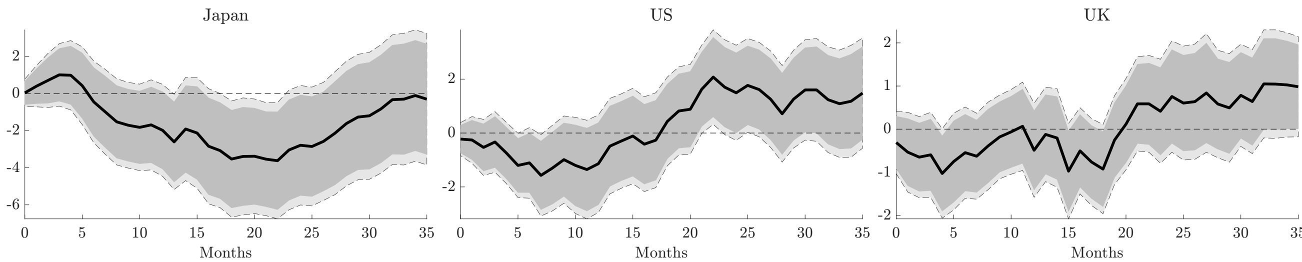

Notes: Notes: Each figure plots the mean, 90 and 95 percent confidence intervals of the difference in the responses of the real effective exchange rate in Japan, the United States, and the United Kingdom in the ZLB and normal periods. The levels of the impulse responses are plotted in Figure B.10.

Figure C.8: Difference in Industrial Production and Unemployment Rate: booms and busts

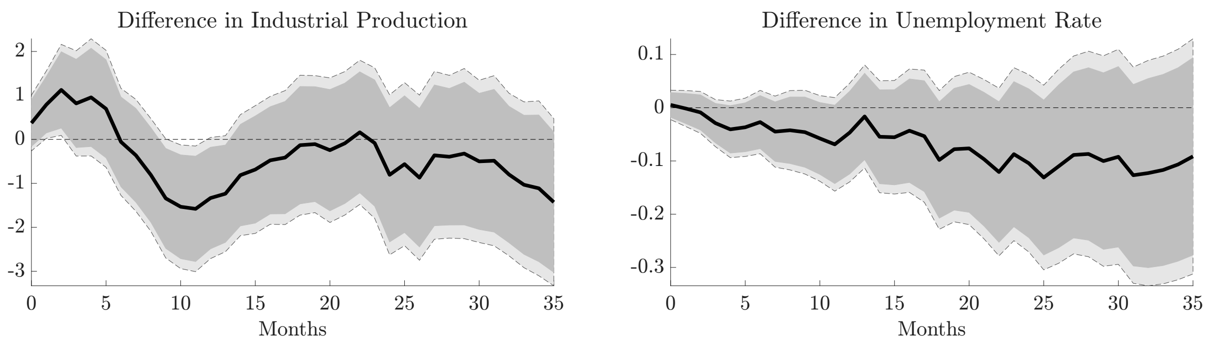

Notes: Notes: Each figure plots the mean, 90- and 95-percent confidence intervals of the difference in the responses of the Japanese industrial production and the unemployment rate in recessions and booms following an oil supply news shock. The levels of the impulse responses are plotted in Figure B.26.

D State-Dependent Proxy VAR

D.1 Oil Market Variables

We estimated a state-dependent world oil VAR where the state is based on the Japanese ZLB definition, our primary focus in this paper. Based on the results of the estimation, we extract the oil supply news shocks. To estimate a state-dependent structural world oil VAR, we first estimated a reduced form VAR of the form

$$Y_t = \mathbb{I}_{t-1} \cdot \left(C_A + \sum_{l=1}^{p} B_{A,l} Y_{t-l} \right) + (1 - \mathbb{I}_{t-1}) \cdot \left(C_B + \sum_{l=1}^{p} B_{B,l} Y_{t-l} \right) + \mathcal{U}_t, \tag{D.1}$$

where the subscripts A and B denote the ZLB and normal periods (i.e., states), and $\mathbb{I}_{t-1}$ is the ZLB dummy in period $t-1$, Y_t is a 4×1 vector of oil variables, and $\mathcal{U}_t$ is a 4×1 vector of the reduced-form residuals. Then, we estimated the state-dependent impact vector, which, as before, can only be identified up to scale:

$$\begin{aligned} \widetilde{S}^{A}_{2:n,1} &= \mathbb{E}\left[z_t \mathcal{U}_{2:n,t} \mathbb{I}_{t-1}\right] / \mathbb{E}\left[z_t \mathcal{U}_{1,t} \mathbb{I}_{t-1}\right], \\ \widetilde{S}^{B}_{2:n,1} &= \mathbb{E}\left[z_t \mathcal{U}_{2:n,t} \left(1 - \mathbb{I}_{t-1}\right)\right] / \mathbb{E}\left[z_t \mathcal{U}_{1,t} \left(1 - \mathbb{I}_{t-1}\right)\right], \end{aligned}$$

where z_t is the oil surprise news series constructed in Kanzig (2021), the subscript notation "*2:n,1*" means that we extract all but the first element of the first column of a matrix, and the subscript notation "*2:n,t*" indicates that we extract all but the first element of a vector in period t. Finally, we extract the oil supply news shock $\widehat{\mathcal{E}}_{1,t}$ in each period. We reduce the number of lags from 12 to 6 to mitigate overfitting arising from the larger number of parameters implied by the two-regime specification. Figure D.1 plots the results where the shocks used in the local projection is extracted from the state-dependent VAR. The ZLB response is somewhat larger in magnitude. The fact that these responses are quite similar to our baseline is consistent with our previous observation that the extracted oil supply shocks from linear and state-dependent proxy VARs are strongly correlated, the Pearson correlation coefficient is 0.81.

Figure D.1: Impulse responses to state-dependent proxy VAR shocks

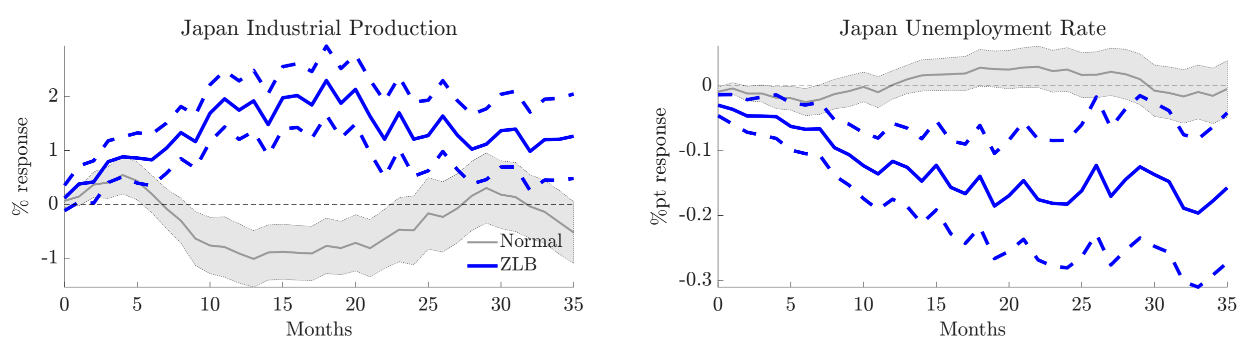

Notes: Each figure plots the impulse response functions and one-standard-deviation robust confidence bands of the responses of Japanese industrial production and the unemployment rate when the shocks are extracted from a state-dependent oil proxy VAR. The responses are obtained using our baseline local projection specification.

D.2 Oil Market and Japanese Variables

This section extends the four-variable state-dependent oil proxy VAR by adding Japanese industrial production and the Japanese unemployment rate, one at a time.

Before we present our results, we note that with only oil variables in the state-dependent proxy VAR, the standard first-stage F-statistic for the proxy in the ZLB period is 3.32 (the heteroskedasticity-robust first-stage F-statistic is 2.11), while it is 26.37 (11.62 when heteroskedasticity robust) in the normal period. The ZLB value is below the recommended threshold of 10 (Montiel Olea, Stock and Watson, 2021), indicating that the instrument is weak in the ZLB period. To assess how this affects inference, we compute Anderson–Rubin (AR) confidence sets using the moving-block bootstrap (MBB) in addition to the standard moving-block bootstrap confidence bands (Jentsch and Lunsford, 2022).

Figure D.2 plots the impulse responses of these variables in both the normal and ZLB periods, together with 68 percent confidence sets computed using the MBB and MBB AR procedures. In the normal period, the MBB and MBB–AR confidence sets are virtually identical, as expected when the instrument is not weak. In the ZLB period, when the instrument is weak, Figure D.2 shows that the Anderson–Rubin confidence sets differ somewhat from those obtained using the moving-block bootstrap alone.

Figure D.2: 2-state Proxy VAR IRFs for Macro Variables

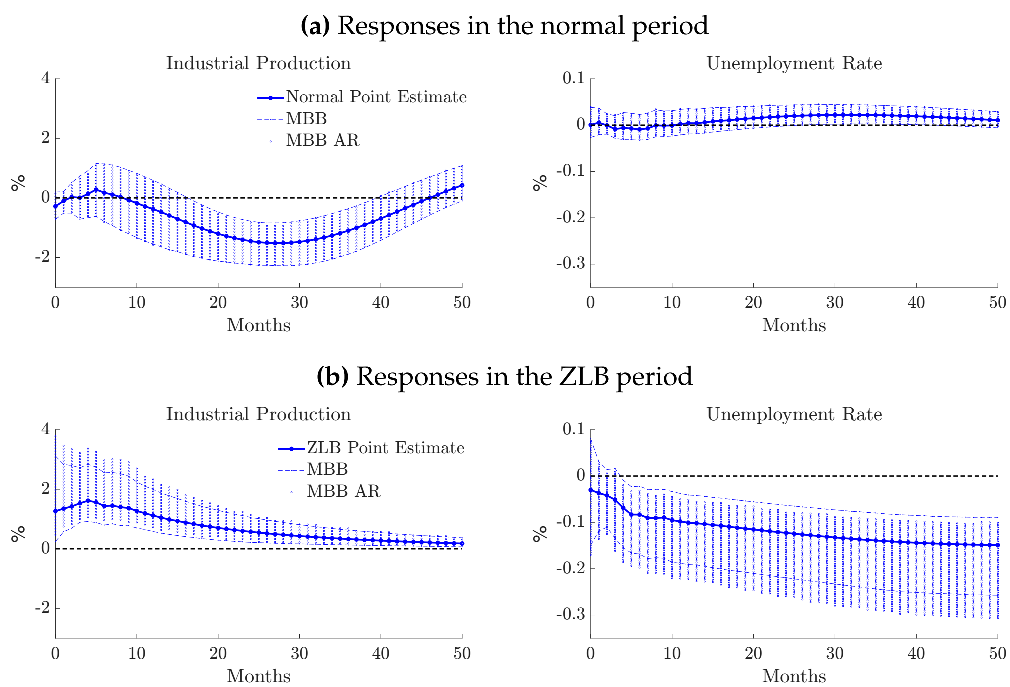

Notes: The figures plot the impulse response functions and confidence sets for industrial production and the unemployment rate in normal period (top panel) and the the ZLB period (bottom panel) in Japan. The responses are to an oil supply news shock. The error bands consist of (i) 68 percent confidence sets computed using the moving-block bootstrap with a block length of 12 (dashed contour lines) and (ii) 68 percent Anderson-Rubin confidence sets computed using the moving-block bootstrap.

Figure D.2 qualitatively reproduces the findings obtained using our benchmark specification. For example, in the normal period, the point estimate of the response of industrial production declines, while the unemployment rate increases. In contrast, during the ZLB period, industrial production increases and the unemployment rate declines.

Figure D.3 presents the differences in the responses of industrial production and the unemployment rate between the ZLB and normal periods, together with 68- and 90-percent confidence intervals computed using the moving-block bootstrap. The difference in impulse responses across the two states becomes statistically different from zero approximately two to three years after the shock at the 90-percent confidence level. Overall, we conclude that the state-dependent proxy VAR analysis yields a qualitatively similar picture to our baseline results.

Figure D.3: 2-state Proxy VAR IRFs difference across ZLB and normal periods

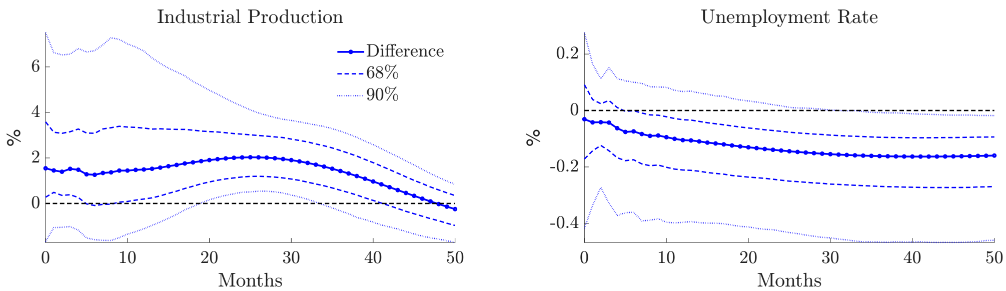

Notes: The figures plot the difference between impulse response functions in the ZLB and normal periods. The responses are to an oil supply news shock. The error bands are 68 and 90 percent confidence sets computed using moving block boostrap with the block length of 12.

E Local Projection and Proxy VAR: a Small Sample Bias

In finite samples, the extracted shock from the proxy VAR is potentially correlated with the other structural shocks, and thus, can bias the results. We next present a series of Monte Carlo simulation exercises where we assess the extent to which this may bias our results in small samples. We start with a generic 2-variable VAR with two lags example to fix ideas, and then we present a 5-variable oil VAR example based on our estimates in the paper. We argue that the bias is small.

2-variable VAR. Consider the following 2-variable proxy VAR with two lags based on Jentsch and Lunsford (2022):

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix} = A_1 \begin{bmatrix} y_{1,t-1} \\ y_{2,t-1} \end{bmatrix} + A_2 \begin{bmatrix} y_{1,t-2} \\ y_{2,t-2} \end{bmatrix} + \begin{bmatrix} u_{1,t} \\ u_{2,t} \end{bmatrix},$$

where the reduced-form innovation and the structural shocks are related as follows

$$\begin{bmatrix} u_{1,t} \\ u_{2,t} \end{bmatrix} = H \begin{bmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \end{bmatrix},$$

and also assume that there is a variable m_t that will be used as an instrument for the shock $\epsilon_{1,t}$

$$m_t = \psi \epsilon_{1,t} + \nu_t,$$

where the parameter ψ indexes the strength of the instrument. All the innovations are standard normal, that is, $\epsilon_{1,t}, \epsilon_{2,t}$, and ν_t are independent and follow $\mathcal{N}(0,1)$. We assume the following parameter values for A_1, A_2, H

$$A_1 = \begin{pmatrix} 0.44 & 0.66 \\ -0.11 & 1.32 \end{pmatrix}, \; A_2 = \begin{pmatrix} -0.18 & 0 \\ -0.18 & -0.09 \end{pmatrix}, \; H = \begin{pmatrix} 0.707 & 0.707 \\ -0.259 & 0.966 \end{pmatrix}.$$

This parameters generate a hump-shaped responses of endogenous variables y_{1t} and y_{2t} in response to shock $\epsilon_{1,t}$. We will assume three levels of the instrument strength: $\psi = \{1, 0.32, 0.17\}$, where $\psi = 1$ corresponds to the case of a strong instrument, and $\psi = 0.17$ is the case of a weak instrument.

In what follows, we will generate 450 periods of data and drop the initial 200 periods. The remaining 250 periods approximately correspond to the combined length of two ZLB periods (260 months in Japan) or the combined lengths of two normal periods (266 months in Japan). We will generate these data 2,000 times. We will compare three sets of results. First, we will estimate the results based on our updated benchmark procedure by estimating a proxy VAR, extracting the shock $\widehat{\epsilon}_{1,t}$, and by running an OLS regressions of $y_{i,t+h}$ on $\widehat{\epsilon}_{1,t}$ and controls of the following form

$$y_{i,t+h} = c + \gamma \widehat{\epsilon}_{1,t} + \beta_{11} y_{1,t-1} + \beta_{12} y_{2,t-1} + \beta_{21} y_{1,t-2} + \beta_{22} y_{2,t-2} + \xi^i_{t,h},$$

for every horizon $h = 0, 1, \ldots, H$ and variable $i = 1, 2$, and where $\xi^i_{t,h}$ represents the residual. Second, we estimate a standard 2-variable proxy VAR with 2 lags and where m_t is used as an external instrument. This approach avoids the need to generate a shock $\widehat{\epsilon}_{1,t}$ and, hence, avoids a bias potentially present in the first OLS procedure. Third, we will generate the true impulse response to the shock $\epsilon_{1,t}$ using our knowledge of the data generating process and its parameter values. Figure E.1 presents the results. Specifically, it shows the true impulse response in solid yellow, the proxy VAR impulse response in red solid-dotted line together with 16% and 84% quantiles of the responses across our Monte Carlo simulations in red dashed-dotted line, and the OLS local projections in blue solid-crossed line together with 16% and 84% quantiles of the responses across our Monte Carlo simulations.

Figure E.1: Monte Carlo Impulse Responses: LP OLS vs. Proxy VAR vs. Truth (2-variable VAR)

(a) Strong instrument

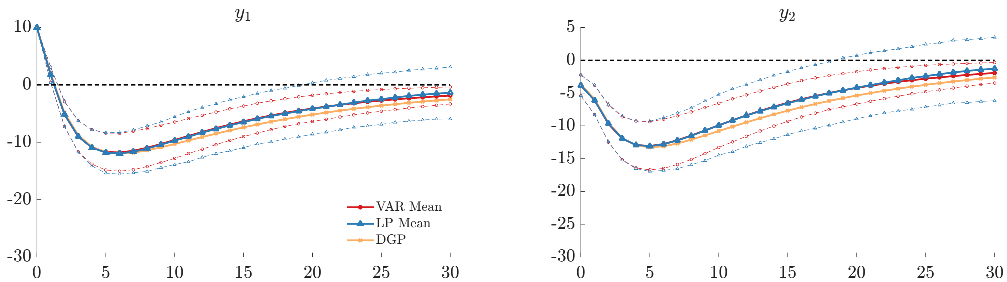

(b) Intermediate strength of the instrument

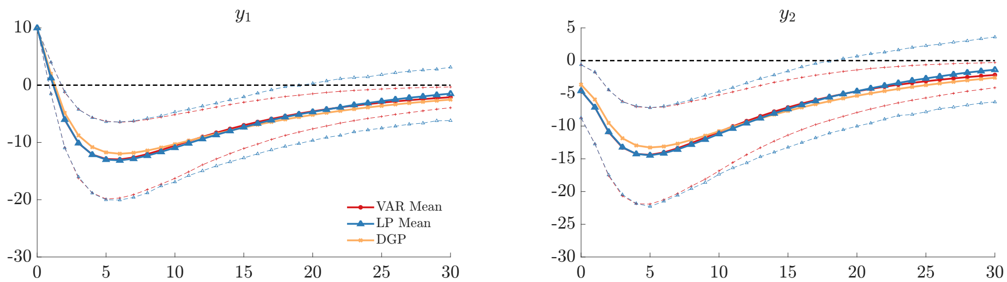

(c) Weak instrument

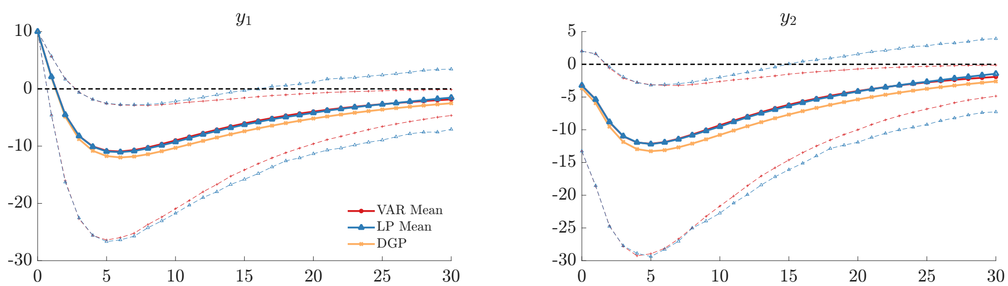

We note several features in Figure E.1. First, the difference between the cross-simulation mean of the proxy VAR (red line), where we do not generate shocks $\widehat{\epsilon}_{1,t}$, and OLS impulse responses (blue line), where we do generate shocks $\widehat{\epsilon}_{1,t}$, is very small under all levels of the instrument strength and both variables. This suggests that the bias due to correlation between $\widehat{\epsilon}_{1,t}$ and $\epsilon_{2,t}$ is small. Second, while both the proxy VAR and OLS impulse responses are very close to each other, they are not identical to the true impulse response (although the difference is not large either). This bias stems from a classical finite-sample least squares bias in time series (Hurwicz, 1950). These results show that the bias associated with generated shock $\widehat{\epsilon}_{1,t}$ is considerably smaller than the bias associated with estimating autoregressive coefficients in small samples.

5-variable VAR. The above results are specific to the two-variable VAR example and the particular parameter values used. To verify whether the same conclusion holds in an empirical model that captures out data, we repeat the exercise by first estimating a five-variable proxy VAR with 12 lags (and without structural breaks) using four oil variables (real oil price, oil production, oil inventories, world production) and Japanese industrial data. We then treat the estimated coefficients as the true parameters in the data-generating process and perform the same Monte Carlo experiment again.

Figure E.2 presents the results. As in the previous example, the cross-simulation mean of impulse responses from the proxy VAR and OLS LP are very close to each other, but both are biased towards zero compared to the true impulse response. This again shows that the bias due to our use of the shock $\widehat{\epsilon}_{1,t}$ is quite small.

Figure E.2: Monte Carlo Impulse Responses: LP OLS vs. Proxy VAR vs. Truth (5-variable VAR)

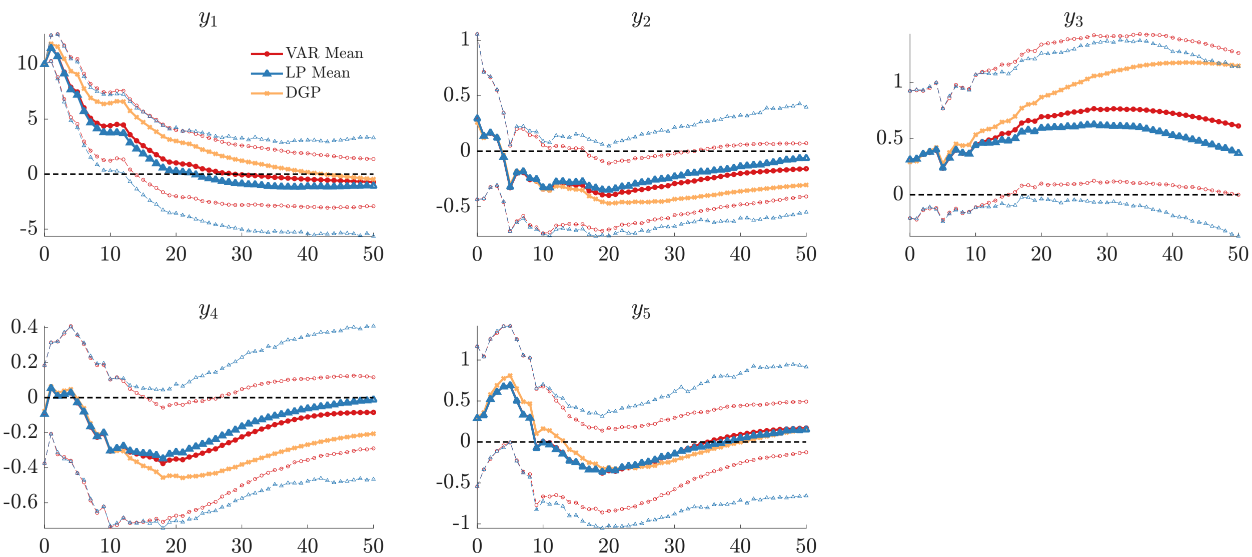

F Model Details

We consider several models with oil. In the first section, we describe a standard open-economy New Keynesian model extended with demand for oil by firms. We then subsequently add anticipated shocks into the model with oil inventories.

F.1 Open-economy New Keynesian Model with Oil

This section presents a standard open-economy New Keynesian model extended with demand for oil by firms. The goal of this section is to illustrate the direct effects of oil shocks on aggregate supply, systematic monetary policy, and the world economy.

There are two countries home (H) and foreign (F). We denote variables of foreign country with an asterisk. Each country is populated by measure n and $n^* \equiv 1-n$ residents. All quantity variables are expressed in per capita terms. Monopolistically competitive producers in home country produce varieties with indexes i on the interval $[0, n)$, while foreign firms on the interval $[n, 1]$.

We describe the agent in home country. Foreign country agents description is symmetric.

F.1.1 Households

At home, there is a representative household with preferences represented by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t, N_t), \tag{F.1}$$

where

$$U(C, N) \equiv \frac{C^{1-\sigma}}{1-\sigma} - \Theta \frac{N^{1+\varphi}}{1+\varphi},$$

and N_t is labor supply, and C_t is consumption index[54]

$$C_t \equiv \left[\omega^{\frac{1}{\gamma_n}} C_{H,t}^{\frac{\gamma_n - 1}{\gamma_n}} + (1-\omega)^{\frac{1}{\gamma_n}} C_{F,t}^{\frac{\gamma_n - 1}{\gamma_n}} \right]^{\frac{\gamma_n}{\gamma_n - 1}}.$$

The indices of home and foreign goods consumption $C_{H,t}$ and $C_{F,t}$ by home residents are

$$C_{H,t} \equiv \left[\left(\frac{1}{n}\right)^{\frac{1}{\epsilon}} \int_0^n C_{H,t}(i)^{\frac{\epsilon-1}{\epsilon}} di \right]^{\frac{\epsilon}{\epsilon-1}},$$

and

$$C_{F,t} \equiv \left[\left(\frac{1}{1-n}\right)^{\frac{1}{\epsilon}} \int_n^1 C_{F,t}(i)^{\frac{\epsilon-1}{\epsilon}} di \right]^{\frac{\epsilon}{\epsilon-1}},$$

where $i \in [0, 1]$ is the index of individual good variety, ϵ is the elasticity of substitution across varieties, γ_n is the elasticity of substitution between home and foreign goods, $1-\omega \equiv (1-n)(1-\Omega)$ is a share

[54] The consumption index in foreign country is symmetric to home consumption index and equals: $C_t^* \equiv [(\omega^*)^{1/\gamma_n} (C_{F,t}^*)^{(\gamma_n-1)/\gamma_n} + (1-\omega^*)^{1/\gamma_n} (C_{H,t}^*)^{(\gamma_n-1)/\gamma_n}]^{\gamma_n/(\gamma_n-1)}$, where ω^* is a share of foreign goods in foreign consumption basket.

of foreign-good expenditures in total expenditures, and $\Omega \in [0,1)$ is an index of home bias. The value of $\Omega = 0$ corresponds to the case of no home bias, while $\Omega > 0$ is required for non-zero degree of home bias. For the foreign economy, $1-\omega^* \equiv n(1-\Omega^*)$ is the share of home-good expenditures in total expenditures.

The household maximizes its preferences by choosing a plan $\{\{C_{F,t}(i)\}, \{C_{H,t}(i)\}, C_{F,t}, C_{H,t}, C_t, N_t, B_{H,t+1}, B_{F,t+1}\}$ subject to the flow budget constraint

$$\int_0^n P_{H,t}(i)C_{H,t}(i)di + \mathcal{E}_t \int_n^1 P^*_{F,t}(j)C_{F,t}(j)dj + \mathbb{E}_t\left[M_{t,t+1}B_{H,t+1}\right] \\ +\mathbb{E}_t\left[\mathcal{E}_t M^*_{t,t+1}B_{F,t+1}\right] \leq B_{H,t} + \mathcal{E}_t B_{F,t} + W_t N_t + P_{O,t}O_t + \Pi_t - T_t, \tag{F.2}$$

where $P_{H,t}(i)$ and $P^*_{F,t}(j)$ are the prices of home and foreign varieties denoted in home and foreign currencies, respectively, $\mathcal{E}_t$ is the nominal exchange rate in units of domestic currency per foreign currency (an increase in $\mathcal{E}_t$ is a depreciation of home currency), $P_{O,t}$ is the oil price in units of home currency, $M_{t,t+1}$ and $M^*_{t,t+1}$ are prices of domestic and foreign state-contingent securities in home and foreign currencies, W_t is nominal wage at home, Π_t represent nominal profits, and T_t is nominal lump-sum taxes. The holdings of home and foreign state-contingent securities are $B_{H,t+1}$ and $B_{F,t+1}$. O_t represents the oil endowment at home.

Price indexes, terms of trade, and real exchange rate. We define the price index of home and foreign produced goods as $P_{H,t} \equiv [n^{-1}\int_0^n P_{H,t}(i)^{1-\epsilon}di]^{1/(1-\epsilon)}$ and $P^*_{F,t} \equiv [(1-n)^{-1}\int_n^1 P^*_{F,t}(i)^{1-\epsilon}di]^{1/(1-\epsilon)}$, and the home CPI is

$$P_t \equiv \left[\omega P_{H,t}^{1-\gamma_n} + (1-\omega)\left(\mathcal{E}_t P^*_{F,t}\right)^{1-\gamma_n}\right]^{1/(1-\gamma_n)}, \tag{F.3}$$

In addition, we assume that the law of one price holds for all goods. We define the terms of trade as the relative price of imported goods $S_t \equiv P_{F,t}/P_{H,t}$ and the real exchange rate as

$$Q_t \equiv \mathcal{E}_t P^*_t / P_t \tag{F.4}$$

The real oil price as the nominal oil price expressed in home currency divided by the foreign producers price index $P_{F,t}$ expressed in home currency:

$$R_t \equiv \frac{P_{O,t}}{P_{F,t}}.$$

Optimality conditions. Consumption

$$C_{H,t}(i) = \frac{1}{n}\left(\frac{P_{H,t}(i)}{P_{H,t}}\right)^{-\epsilon} C_{H,t}, \tag{F.5}$$

$$C_{F,t}(i) = \frac{1}{1-n}\left(\frac{P^*_{F,t}(i)}{P^*_{F,t}}\right)^{-\epsilon} C_{F,t}, \tag{F.6}$$

where

$$C_{H,t} = \omega\left(\frac{P_{H,t}}{P_t}\right)^{-\gamma_n} C_t, \tag{F.7}$$

$$C_{F,t} = (1-\omega)\left(\frac{\mathcal{E}_t P^*_{F,t}}{P_t}\right)^{-\gamma_n} C_t, \tag{F.8}$$

Labor supply

$$-\frac{U_2(C_t, N_t)}{U_1(C_t, N_t)} = \frac{W_t}{P_t}. \tag{F.9}$$

where $U_1(C_t, N_t)$ and $U_2(C_t, N_t)$ are partial derivatives with respect to the first and the second arguments.
Assets

$$M_{t,t+1} = \beta \frac{U_1(C_{t+1}, N_{t+1})}{U_1(C_t, N_t)} \cdot \frac{P_t}{P_{t+1}}, \tag{F.10}$$

$$M^*_{t,t+1} = \beta \frac{U_1(C_{t+1}, N_{t+1})}{U_1(C_t, N_t)} \cdot \frac{P_t}{P_{t+1}} \cdot \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t}. \tag{F.11}$$

Foreign households. Conditions symmetric to (F.5)-(F.11) hold for foreign country. In particular, the first-order condition with respect to $B^*_{H,t}$ imply

$$M_{t,t+1} = \beta \frac{U_1(C^*_{t+1}, N^*_{t+1})}{U_1(C^*_t, N^*_t)} \cdot \frac{P^*_t}{P^*_{t+1}} \cdot \frac{1/\mathcal{E}_{t+1}}{1/\mathcal{E}_t}. \tag{F.12}$$

Combining equations (F.10) and (F.12), we get

$$\frac{U_1(C_{t+1}, N_{t+1})}{U_1(C^*_{t+1}, N^*_{t+1})} Q_{t+1} = \frac{U_1(C_t, N_t)}{U_1(C^*_t, N^*_t)} Q_t \equiv \phi_q, \tag{F.13}$$

which implies that the ratio of marginal utility of consumption in the two countries multiplied by the real exchange rate is constant over time, which we denote as ϕ_q.

Log-linearization. We next log-linearize around the steady state where all relative goods prices are 1 and all prices and quantities are constant. Section F.1.8 will provide more details about steady state. For now, we obtain

$$\widehat{c}_{H,t} = -\gamma_n \left(\widehat{p}_{H,t} - \widehat{p}_t\right) + \widehat{c}_t,$$
$$\widehat{c}_{F,t} = -\gamma_n \left(\widehat{e}_t + \widehat{p}^*_{F,t} - \widehat{p}_t\right) + \widehat{c}_t,$$

where hats denote log-deviations from steady state values.

The labor supply

$$\sigma \widehat{c}_t + \varphi \widehat{n}_t = \widehat{w}_t - \widehat{p}_t.$$

The Euler equation

$$\mathbb{E}_t \left[\widehat{c}_{t+1}\right] - \widehat{c}_t = \frac{1}{\sigma} \left(i_t - \mathbb{E}_t \left[\pi_{t+1}\right] - \iota\right),$$

where $i_t \equiv -\log \mathbb{E}_t[M_{t,t+1}]$ is the safe short-term nominal interest rate and $\iota \equiv -\log \beta$.

The international risk sharing condition—a log-linear version of the second equality in equation (F.13)—is

$$\widehat{c}^*_t + \frac{1}{\sigma} q_t = \widehat{c}_t.$$

Total consumption is

$$\widehat{c}_t = \omega \widehat{c}_{H,t} + (1-\omega) \widehat{c}_{F,t}.$$

F.1.2 Price indices

The log-linearization of price indexes in home and foreign countries are

$$\widehat{p}_t = \omega \widehat{p}_{H,t} + (1-\omega)\widehat{p}_{F,t} = \widehat{p}_{H,t} + (1-\omega)\left(\widehat{p}_{F,t} - \widehat{p}_{H,t}\right),$$
$$\widehat{p}^*_t = \omega^* \widehat{p}^*_{F,t} + (1-\omega^*)\widehat{p}^*_{H,t} = \widehat{p}^*_{F,t} + (1-\omega^*)\left(\widehat{p}^*_{H,t} - \widehat{p}^*_{F,t}\right).$$

The real exchange rate is

$$\widehat{q}_t = \widehat{e}_t + \widehat{p}^*_t - \widehat{p}_t = (\omega + \omega^* - 1)\left(\widehat{p}_{F,t} - \widehat{p}_{H,t}\right).$$

Home relative prices are

$$\begin{aligned}
\widehat{p}_{H,t} - \widehat{p}_t &= -\frac{1-\omega}{\omega+\omega^*-1}\widehat{q}_t, \\
\widehat{p}_{F,t} - \widehat{p}_t &= \frac{\omega}{\omega+\omega^*-1}\widehat{q}_t, \\
\widehat{p}_{O,t} - \widehat{p}_t &= \widehat{p}_{O,t} - \widehat{p}_{F,t} + \frac{\omega}{\omega+\omega^*-1}\widehat{q}_t.
\end{aligned}$$

Foreign relative prices are

$$\begin{aligned}
\widehat{p}^*_{F,t} - \widehat{p}^*_t &= \frac{1-\omega^*}{\omega+\omega^*-1}\widehat{q}_t, \\
\widehat{p}^*_{H,t} - \widehat{p}^*_t &= -\frac{\omega^*}{\omega+\omega^*-1}\widehat{q}_t, \\
\widehat{p}^*_{O,t} - \widehat{p}^*_t &= \widehat{p}^*_{O,t} - \widehat{p}^*_{F,t} + \frac{1-\omega^*}{\omega+\omega^*-1}\widehat{q}_t.
\end{aligned}$$

F.1.3 Firms

A home producer of variety i combines labor $N_t(i)$ and oil inputs $O_{Y,t}(i)$ to produce good i according to the CES production function of the form

$$Y_t(i) = \left[(1-\omega_{oy})^{\frac{1}{\gamma_y}}\left(AN_t(i)\right)^{\frac{\gamma_y-1}{\gamma_y}} + \omega_{oy}^{\frac{1}{\gamma_y}} O_{Y,t}(i)^{\frac{\gamma_y-1}{\gamma_y}}\right]^{\frac{\gamma_y}{\gamma_y-1}},$$

where A_t is labor productivity, the intra-temporal elasticity of substitution γ_y is positive, and $\omega_{oy} \in [0,1]$ is the share of oil in production. The firm is free to optimize its inputs every period. Because of the constant-elasticity-of-substitution assumption about the form of the production function, the *nominal* marginal cost of production does not depend on output and equals

$$MC_t = \left[(1-\omega_{oy})\left(\frac{W_t}{A}\right)^{1-\gamma_y} + \omega_{oy} P_{O,t}^{1-\gamma_y}\right]^{\frac{1}{1-\gamma_y}}, \tag{F.14}$$

and the optimal labor and oil choices are

$$O_{Y,t}(i) = \omega_{oy} \left(\frac{P_{O,t}}{MC_t} \right)^{-\gamma_y} Y_t(i), \tag{F.15}$$

$$N_t(i) = (1 - \omega_{oy}) \left(\frac{W_t / A}{MC_t} \right)^{-\gamma_y} \frac{Y_t(i)}{A}. \tag{F.16}$$

The firm resets its price infrequently in the spirit of Calvo. When allowed, a firm chooses the reset price $P^r_{H,t}(i)$ to maximizes an expected discounted sum of future profits

$$\mathbb{E}_t \left[\sum_{k=0}^{\infty} \theta^k M_{t,t+k} \left(P^r_{H,t}(i) - (1+\tau) MC_{t+k} \right) Y_{H,t+k|t}(i) \right]$$

where $1 - \theta \in [0, 1]$ is the probability of price reset, $Y_{H,t+k|t}(i)$ is the output in period $t + k$ conditional on the price set in period t, and τ is the government's proportional tax. The price is set in producer's currency and the firm does not differentiate between domestic and foreign consumers. In other words, the law of one price holds for individual varieties.

The demand for variety i is

$$Y_{H,t+k|t}(i) = nC_{H,t+k|t}(i) + (1-n) C^*_{H,t+k|t}(i) = \left(\frac{P^r_{H,t}(i)}{P_{H,t+k}} \right)^{-\epsilon} \left(C_{H,t+k} + \frac{1-n}{n} C^*_{H,t+k} \right),$$

where $C_{H,t+k}$ and $C^*_{H,t+k}$ are the home and foreign demand for home goods.

Firm's optimal choice of its reset price requires

$$\mathbb{E}_t \left[\sum_{k=0}^{\infty} \theta^k M_{t,t+k} Y_{H,t+k|t}(i) \left(\frac{P^r_{H,t}(i)}{P_{H,t-1}} - \frac{\epsilon}{\epsilon - 1} (1+\tau) \frac{MC_{t+k}}{P_{H,t+k}} \cdot \frac{P_{H,t+k}}{P_{H,t-1}} \right) \right] = 0, \tag{F.17}$$

where we deflated the expression by $P_{H,t-1}$.

Log-linearization. We log-linearize around a steady state with constant prices and quantities and where all relative prices equal one. The marginal costs are

$$\widehat{mc}_t = (1 - \omega_{oy}) \widehat{w}_t + \omega_{oy} \widehat{p}_{O,t}.$$

The demand for labor and oil are

$$\widehat{n}_t = -\gamma_y \left(\widehat{w}_t - \widehat{mc}_t \right) + \frac{nC_H}{nC_H + (1-n) C^*_H} \widehat{c}_{H,t} + \frac{(1-n) C^*_H}{nC_H + (1-n) C^*_H} \widehat{c}^*_{H,t}$$
$$\widehat{o}_{Y,t} = -\gamma_y \left(\widehat{p}_{O,t} - \widehat{mc}_t \right) + \frac{nC_H}{nC_H + (1-n) C^*_H} \widehat{c}_{H,t} + \frac{(1-n) C^*_H}{nC_H + (1-n) C^*_H} \widehat{c}^*_{H,t},$$

The optimal reset price satisfies

$$p^r_{H,t} - p_{H,t-1} = (1 - \beta\theta) \left(\widehat{mc}_t - \widehat{p}_{H,t} \right) + \pi_{H,t} + \beta\theta \mathbb{E}_t \left[p^r_{H,t+1} - p_{H,t} \right]. \tag{F.18}$$

F.1.4 Government

The home government consists of fiscal and monetary authorities. The fiscal authority undoes monopolistic competition distortion by subsidizing production. Specifically, it sets the production tax $\tau = -1/\epsilon$ and runs a balanced budget each period.

The monetary authority sets the short-term safe nominal interest rate $i_t \equiv -\log \mathbb{E}_t[M_{t,t+1}]$ according to the Taylor rule that reacts to the producers price inflation rate $\pi_{H,t} \equiv \log(P_{H,t}/P_{H,t-1})$. Formally,

$$i_t = \iota + \varphi_\pi \pi_{H,t}, \tag{F.19}$$

where $\iota \equiv -\log \beta$ and $\varphi_\pi \geq 0$. We do not explicitly introduce the zero lower bound on the interest rate here.

The foreign government acts analogous to home government. In particular, the central bank set the interest rate according to

$$i_t^* = \iota + \varphi_\pi^* \pi_{H,t}^*, \tag{F.20}$$

where $\varphi_\pi^* \geq 0$.

F.1.5 Oil

Oil is supplied as endowment that equals O_t and O_t^* in per capita terms, and these endowments change randomly every period. Households own oil and it cannot be stored. The oil prices $P_{O,t}$ and $P_{O,t}^*$ are perfectly flexible and the law of one price applies to oil.

F.1.6 Market Clearing

Local markets. The home labor supply equals home labor demand

$$N_t = \int_0^n \left(1 - \omega_{oy}\right) \left(\frac{W_t/A}{MC_t}\right)^{-\gamma_y} \frac{Y_t(i)}{A} di.$$

The market for every home variety $i \in [0, n)$ clears

$$Y_{H,t}(i) = nC_{H,t}(i) + (1-n)C_{H,t}^*(i). \tag{F.21}$$

The analogoes conditions hold for foreign country.

Global markets. In equilibrium, all markets clear. Specifically, all state-contingent asset markets clear

$$n\left(B_{H,t} + \mathcal{E}_t B_{F,t}\right) + (1-n)\left(B_{H,t}^* + \mathcal{E}_t B_{F,t}^*\right) = 0.$$

The oil market clears

$$nO_{Y,t} + (1-n)\, O_{Y,t}^* = nO_t + (1-n)O_t^*.$$

and its log-linear approximation is

$$nO_Y \widehat{o}_{Y,t} + (1-n)O_Y^* \widehat{o}_{Y,t}^* = nO\widehat{o}_t + (1-n)O^* \widehat{o}_t^*.$$

F.1.7 Aggregation

The law of motion of the price index

$$\Pi_{H,t}^{1-\epsilon} = \theta + (1-\theta)\left(\frac{P_{H,t}^{r}}{P_{H,t-1}}\right)^{1-\epsilon},$$
$$\left(\Pi_{F,t}^{*}\right)^{1-\epsilon} = \theta + (1-\theta)\left(\frac{P_{F,t}^{*,r}}{P_{F,t-1}^{*}}\right)^{1-\epsilon}.$$

The price dispersion in home country is

$$\Delta_{H,t} \equiv \frac{1}{n}\int_0^n \left(\frac{P_{H,t}(i)}{P_{H,t}}\right)^{-\epsilon} di = \theta\left(\frac{P_{H,t-1}}{P_{H,t}}\right)^{-\epsilon}\Delta_{H,t-1} + (1-\theta)\left(\frac{P_{H,t}^{r}}{P_{H,t}}\right)^{-\epsilon}.$$

Aggregate goods market clearing in home and foreign countries

$$\left[(1-\omega_{oy})^{\frac{1}{\gamma_y}}(AN_t)^{\frac{\gamma_y-1}{\gamma_y}} + \omega_{oy}^{\frac{1}{\gamma_y}} O_{Y,t}^{\frac{\gamma_y-1}{\gamma_y}}\right]^{\frac{\gamma_y}{\gamma_y-1}} = \left(C_{H,t+k} + \frac{1-n}{n}C_{H,t+k}^{*}\right)\Delta_{H,t},$$
$$\left[(1-\omega_{oy})^{\frac{1}{\gamma_y}}(A^{*}N_t^{*})^{\frac{\gamma_y-1}{\gamma_y}} + \omega_{oy}^{\frac{1}{\gamma_y}} \left(O_{Y,t}^{*}\right)^{\frac{\gamma_y-1}{\gamma_y}}\right]^{\frac{\gamma_y}{\gamma_y-1}} = \left(C_{F,t+k}^{*} + \frac{n}{1-n}C_{F,t+k}\right)\Delta_{F,t}^{*}.$$

Value added in terms of produced goods

$$Z_{H,t} = Y_{H,t} - \frac{P_{O,t}}{P_{H,t}}O_{Y,t} = Y_{H,t} - R_t\frac{P_{F,t}}{P_{H,t}}O_{Y,t},$$
$$Z_{F,t}^{*} = Y_{F,t}^{*} - \frac{P_{O,t}^{*}}{P_{F,t}^{*}}O_{Y,t}^{*} = Y_{F,t}^{*} - R_tO_{Y,t}^{*}.$$

where $Y_{H,t}$ and $Y_{F,t}^{*}$ is the demand for home and foreign-produced goods.

Trade balance in units of domestically produced goods over steady state output

$$\frac{NX_t}{Y_H} = \frac{1}{Y_H}\left(Y_{H,t} + \frac{P_{O,t}}{P_{H,t}}O_t - \frac{P_t}{P_{H,t}}C_t - \frac{P_{O,t}}{P_{H,t}}O_{Y,t}\right),$$
$$\frac{NX_t^{*}}{Y_F^{*}} = \frac{1}{Y_F^{*}}\left(Y_{F,t}^{*} + \frac{P_{O,t}^{*}}{P_{F,t}^{*}}O_t^{*} - \frac{P_t^{*}}{P_{F,t}^{*}}C_t^{*} - \frac{P_{O,t}^{*}}{P_{F,t}^{*}}O_{Y,t}^{*}\right).$$

This formulas takes into account that the countries trade in goods and oil.

Log-linearization. Home and foreign inflation rates are

$$\pi_{H,t} = (1-\theta)\left(\widehat{p}_{H,t}^{r} - \widehat{p}_{H,t-1}\right),$$
$$\pi_{F,t}^{*} = (1-\theta)\left(\widehat{p}_{F,t}^{*,r} - \widehat{p}_{F,t-1}^{*}\right).$$

Combining the last two equations with the optimal reset price equation at home F.18 and a similar equation abroad, we get two standard Phillips curves

$$\pi_{H,t} = \frac{(1-\theta)(1-\beta\theta)}{\theta}\left(\widehat{mc}_t - \widehat{p}_{H,t}\right) + \beta\mathbb{E}_t\left[\pi_{H,t+1}\right],$$
$$\pi^*_{F,t} = \frac{(1-\theta)(1-\beta\theta)}{\theta}\left(\widehat{mc}^*_t - \widehat{p}^*_{F,t}\right) + \beta\mathbb{E}_t\left[\pi^*_{F,t+1}\right].$$

Goods market clearing

$$(1-\omega_{oy})^{\frac{1}{\gamma_y}}\left(\frac{AN}{C_H + \frac{1-n}{n}C^*_H}\right)^{\frac{\gamma_y-1}{\gamma_y}}\widehat{n}_t + \omega_{oy}^{\frac{1}{\gamma_y}}\left(\frac{O_Y}{C_H + \frac{1-n}{n}C^*_H}\right)^{\frac{\gamma_y-1}{\gamma_y}}\widehat{o}_{Y,t}$$
$$= \frac{nC_H}{nC_H + (1-n)\,C^*_H}\widehat{c}_{H,t} + \frac{(1-n)\,C^*_H}{nC_H + (1-n)\,C^*_H}\widehat{c}^*_{H,t},$$
$$(1-\omega_{oy})^{\frac{1}{\gamma_y}}\left(\frac{A^*N^*}{C^*_F + \frac{n}{1-n}C_F}\right)^{\frac{\gamma_y-1}{\gamma_y}}\widehat{n}^*_t + \omega_{oy}^{\frac{1}{\gamma_y}}\left(\frac{O^*_Y}{C^*_F + \frac{n}{1-n}C_F}\right)^{\frac{\gamma_y-1}{\gamma_y}}\widehat{o}^*_{Y,t}$$
$$= \frac{(1-n)\,C^*_F}{(1-n)\,C^*_F + nC_F}\widehat{c}^*_{F,t} + \frac{nC_F}{(1-n)\,C^*_F + nC_F}\widehat{c}_{F,t},$$

Value added

$$\widehat{z}_{H,t} = \frac{1}{1-\omega_{oy}}\widehat{y}_{H,t} - \frac{\omega_{oy}}{1-\omega_{oy}}\left(\widehat{r}_t + \widehat{p}_{F,t} - \widehat{p}_{H,t} + \widehat{o}_{y,t}\right),$$
$$\widehat{z}^*_{F,t} = \frac{1}{1-\omega_{oy}}\widehat{y}^*_{F,t} - \frac{\omega_{oy}}{1-\omega_{oy}}\left(\widehat{r}_t + \widehat{o}^*_{y,t}\right).$$

The trade balance is

$$\widehat{nx}_t = \widehat{y}_{H,t} + \frac{O-O_Y}{Y_H}\left(\widehat{r}_t + \frac{1}{\omega+\omega^*-1}\widehat{q}_t\right) + \frac{O\widehat{o}_t - O_Y\widehat{o}_{Y,t}}{Y_H} - \frac{C}{Y_H}\left(\frac{1-\omega}{\omega+\omega^*-1}\widehat{q}_t + \widehat{c}_t\right), \tag{F.22}$$
$$\widehat{nx}^*_t = \widehat{y}^*_{F,t} + \frac{O^*-O^*_Y}{Y^*_F}\widehat{r}_t + \frac{O^*\widehat{o}^*_t - O^*_Y\widehat{o}^*_{Y,t}}{Y^*_F} - \frac{C^*}{Y^*_F}\left(-\frac{1-\omega^*}{\omega+\omega^*-1}\widehat{q}_t + \widehat{c}^*_t\right). \tag{F.23}$$

F.1.8 Steady State

We assume that in steady state all prices are constant and the relative prices equal one:

$$\frac{P_O}{P} = \frac{P_F}{P} = \frac{P_H}{P} = 1.$$

This implies that the real exchange rate is

$$Q = \frac{\mathcal{E}P^*}{P} = \mathcal{E}\frac{\left[(1-\omega^*)\left(P_H/\mathcal{E}\right)^{1-\gamma_n} + \omega^*\left(P^*_F\right)^{1-\gamma_n}\right]^{1/(1-\gamma_n)}}{\left[\omega P_H^{1-\gamma_n} + (1-\omega)\left(\mathcal{E}P^*_F\right)^{1-\gamma_n}\right]^{1/(1-\gamma_n)}} = 1$$

We get

$$-\frac{U_2\left(C,N\right)}{U_1\left(C,N\right)} = A$$

$$
\begin{aligned}
e^{-\iota} &= \beta, \\
\frac{W}{P} &= A, \\
\frac{MC}{P} &= \frac{W}{AP} = 1, \\
C_H &= \omega C, \\
C_F &= (1-\omega) C, \\
Y_H &= C_H + \frac{1-n}{n} C_H^*, \\
Y_H &= \left[(1-\omega_{oy})^{\frac{1}{\gamma_y}} (AN)^{\frac{\gamma_y - 1}{\gamma_y}} + \omega_{oy}^{\frac{1}{\gamma_y}} O_Y^{\frac{\gamma_y - 1}{\gamma_y}} \right]^{\frac{\gamma_y}{\gamma_y - 1}}, \\
AN &= \left(1-\omega_{oy}\right) Y_H, \\
O_Y &= \omega_{oy} Y_H, \\
C &= Y_H + O - O_Y.
\end{aligned}
$$

Observe that that the last equation is the flow budget constraint of the country. Also note that if there were no oil, then we would have $C = Y_H$ and $C^* = C_F^* + \frac{n}{1-n} C_F$ and $C_H = \omega C$ and $C_F = (1-\omega) C$ and $C_F^* = \omega^* C^*$ and $C_H^* = (1-\omega^*) C^*$. Combining these together, we get $(1-\omega) nC = (1-n)(1-\omega^*) C^*$ and $(1-n)(1-\omega^*) C^* = n(1-\omega) C$. The last two equations are identical. They determine the ratio: $C/C^* = (1-n)/n \cdot (1-\omega^*)/(1-\omega) = (1-\Omega^*)/(1-\Omega)$. How can we determine steady state consumption level? The Backus-Smith condition $C = \phi_q^{-\frac{1}{\sigma}} C^* Q^{\frac{1}{\sigma}}$ implies $C/C^* = \phi_q^{-\frac{1}{\sigma}} = (1-\Omega^*)/(1-\Omega)$.

Simplifying we get 3 equations and 2 unknowns C, N as functions of parameters and foreign consumption C^*. The first two equations below unambiguously determine C and AN, while the third equation is a constraint on the parameter Θ.

$$
\begin{aligned}
nC &= \frac{(1-n)(1-\omega^*)\left(1-\omega_{oy}\right) C^* + nO}{1-\left(1-\omega_{oy}\right)\omega}, \\
AN &= \left(1-\omega_{oy}\right)\left[\omega C + \frac{1-n}{n}(1-\omega^*) C^*\right], \\
-\frac{U_2(C,N)}{U_1(C,N)} &= \frac{\Theta N^{\varphi}}{C^{-\sigma}} = A,
\end{aligned}
$$

We can write symmetric equations for the foreign economy. Specifically,

$$
(1-n) C^* = \frac{n(1-\omega)\left(1-\omega_{oy}\right) C + (1-n) O^*}{1-\left(1-\omega_{oy}\right)\omega^*},
$$

$$
N^* = \frac{\left(1-\omega_{oy}\right)\left(\omega^* C^* + \frac{n}{1-n}(1-\omega) C\right)}{A}.
$$

For the Backus-Smith condition to be satisfied, the constant ϕ_q has to take a certain value.

$$
\frac{C}{C^*} = \phi_q^{-\frac{1}{\sigma}}.
$$

Solving jointly for C and C^*, we get:

$$nC=\frac{\left[1-\left(1-\omega_{oy}\right)\omega^*\right]nO+\left(1-\omega_{oy}\right)\left(1-\omega^*\right)\left(1-n\right)O^*}{\omega_{oy}\left\{1+\left(1-\omega_{oy}\right)\left[1-\left(\omega^*+\omega\right)\right]\right\}}$$

and

$$\left(1-n\right)C^*=\frac{\left(1-\omega\right)\left(1-\omega_{oy}\right)nO+\left[1-\omega\left(1-\omega_{oy}\right)\right]\left(1-n\right)O^*}{\omega_{oy}\left\{1+\left(1-\omega_{oy}\right)\left[1-\left(\omega^*+\omega\right)\right]\right\}}.$$

The steady state share of home goods domestic demand in home good demand

$$\frac{nC_H}{nC_H+\left(1-n\right)C_H^*}=\frac{n\omega C}{n\omega C+\left(1-n\right)\left(1-\omega^*\right)C^*}$$

And

$$\frac{nC_F}{nC_F+\left(1-n\right)C_F^*}=\frac{n\left(1-\omega\right)C}{n\left(1-\omega\right)C+\left(1-n\right)\omega^*C^*}$$

Consumption normalization. We further normalize steady state to have unit consumption in home and foreign countries: $C=C^*=1$.This normalization requires the following restriction on the parameters

$$1=\left(1-\omega_{oy}\right)\left[1+\left(1-n\right)\left(\Omega-\Omega^*\right)\right]+O.$$
$$1=\left(1-\omega_{oy}\right)\left[1+n\left(\Omega^*-\Omega\right)\right]+O^*,$$

Small open economy limit. In the case when n goes to zero, we get

$$\left(1-\Omega^*\right)=\left(1-\Omega\right)+\frac{\omega_{oy}-O}{1-\omega_{oy}}.$$
$$O^*=\omega_{oy}.$$

In this case, we have

$$\frac{nC_H}{nC_H+\left(1-n\right)C_H^*}=\frac{\Omega}{1-O}\left(1-\omega_{oy}\right)=\Omega\frac{1-O^*}{1-O}=\frac{\Omega}{1+\Omega-\Omega^*}.$$

As a benchmark case, we assume that home country does not produce oil, that is, $O=0$**.** In this case,

$$\left(1-\Omega^*\right)=\left(1-\Omega\right)+\frac{\omega_{oy}}{1-\omega_{oy}}.$$
$$O^*=\omega_{oy},$$

and the share of domestic and foreign consumption are

$$\frac{nC_H}{nC_H+\left(1-n\right)C_H^*}=\Omega\left(1-\omega_{oy}\right),$$
$$\frac{nC_F}{nC_F+\left(1-n\right)C_F^*}=0.$$

F.1.9 Equilibrium

Unknowns (27).

$$\begin{aligned}\textbf{Home quantities (6): } & \widehat{c}_t, \widehat{c}_{H,t}, \widehat{c}_{F,t}, \widehat{o}_{C,t}, \widehat{o}_{Y,t}, \widehat{n}_t, \\ \textbf{Home prices (6):} & \widehat{p}_t, \widehat{p}_{H,t}, \widehat{w}_t, i_t, \widehat{mc}_t, \pi_{H,t} \\ \textbf{Foreign quantities (6): } & \widehat{c}^*_t, \widehat{c}^*_{H,t}, \widehat{c}^*_{F,t}, \widehat{o}^*_{C,t}, \widehat{o}^*_{Y,t}, \widehat{n}^*_t, \\ \textbf{Foreign prices (6):} & \widehat{p}^*_t, \widehat{p}^*_{F,t}, \widehat{w}^*_t, i^*_t, \widehat{mc}^*_t, \pi^*_{F,t} \\ \textbf{International prices (3): } & \widehat{e}_t, \widehat{q}_t, \widehat{p}_{O,t}\end{aligned}$$

Home conditions (11 equations). Households

$$\begin{aligned}\widehat{c}_{H,t} &= -\gamma_n \left(\widehat{p}_{H,t} - \widehat{p}_t\right) + \widehat{c}_t, \\ \widehat{c}_{F,t} &= -\gamma_n \left(\widehat{e}_t + \widehat{p}^*_{F,t} - \widehat{p}_t\right) + \widehat{c}_t, \\ \sigma\widehat{c}_t + \varphi\widehat{n}_t &= \widehat{w}_t - \widehat{p}_t, \\ \mathbb{E}_t\left[\widehat{c}_{t+1}\right] - \widehat{c}_t &= \frac{1}{\sigma}\left(i_t - \mathbb{E}_t\left[\pi_{t+1}\right] - \iota\right).\end{aligned}$$

The Price index

$$\widehat{p}_t = \omega\widehat{p}_{H,t} + (1-\omega)\left(\widehat{e}_t + \widehat{p}^*_{F,t}\right).$$

Firms

$$\begin{aligned}\widehat{mc}_t &= (1-\omega_{oy})\widehat{w}_t + \omega_o\widehat{p}_{O,t}, \\ \widehat{n}_t &= -\gamma_y\left(\widehat{w}_t - \widehat{mc}_t\right) + \frac{nC_H}{nC_H + (1-n)\,C^*_H}\widehat{c}_{H,t} + \frac{(1-n)\,C^*_H}{nC_H + (1-n)\,C^*_H}\widehat{c}^*_{H,t}, \\ \widehat{o}_{Y,t} &= -\gamma_y\left(\widehat{p}_{O,t} - \widehat{mc}_t\right) + \frac{nC_H}{nC_H + (1-n)\,C^*_H}\widehat{c}_{H,t} + \frac{(1-n)\,C^*_H}{nC_H + (1-n)\,C^*_H}\widehat{c}^*_{H,t}, \\ \pi_{H,t} &= \frac{(1-\theta)(1-\beta\theta)}{\theta}\left(\widehat{mc}_t - \widehat{p}_{H,t}\right) + \beta\mathbb{E}_t\left[\pi_{H,t+1}\right].\end{aligned}$$

Goods market clearing

$$\left(1-\omega_{oy}\right)\widehat{n}_t + \omega_{oy}\widehat{o}_{Y,t} = \frac{nC_H}{nC_H + (1-n)\,C^*_H}\widehat{c}_{H,t} + \frac{(1-n)\,C^*_H}{nC_H + (1-n)\,C^*_H}\widehat{c}^*_{H,t}.$$

Government

$$i_t = \iota + \varphi_\pi\pi_{H,t}.$$

Foreign condition (11 equations). Households

$$\begin{aligned}
\widehat{c}_{H,t}^{*} &= -\gamma_n \left(\widehat{p}_{H,t} - \widehat{e}_t - \widehat{p}_t^{*}\right) + \widehat{c}_t^{*}, \\
\widehat{c}_{F,t}^{*} &= -\gamma_n \left(\widehat{p}_{F,t}^{*} - \widehat{p}_t^{*}\right) + \widehat{c}_t^{*}, \\
\sigma \widehat{c}_t^{*} + \varphi \widehat{n}_t^{*} &= \widehat{w}_t^{*} - \widehat{p}_t^{*}, \\
\mathbb{E}_t \left[\widehat{c}_{t+1}^{*}\right] - \widehat{c}_t^{*} &= \frac{1}{\sigma}\left(i_t^{*} - \mathbb{E}_t\left[\pi_{t+1}^{*}\right] - \iota\right).
\end{aligned}$$

The price index

$$\widehat{p}_t^{*} = \omega^{*} \widehat{p}_{F,t}^{*} + (1-\omega^{*}) \widehat{p}_{H,t}^{*}.$$

Firms

$$\begin{aligned}
\widehat{mc}_t^{*} &= (1-\omega_{oy}) \widehat{w}_t^{*} + \omega_{oy} \left(\widehat{p}_{O,t} - \widehat{e}_t\right), \\
\widehat{n}_t^{*} &= -\gamma_y \left(\widehat{w}_t^{*} - \widehat{mc}_t^{*}\right) + \frac{nC_F}{nC_F + (1-n)\,C_F^{*}} \widehat{c}_{F,t} + \frac{(1-n)\,C_F^{*}}{nC_F + (1-n)\,C_F^{*}} \widehat{c}_{F,t}^{*} \\
\widehat{o}_{Y,t}^{*} &= -\gamma_y \left(\widehat{p}_{O,t}^{*} - \widehat{mc}_t^{*}\right) + \frac{nC_F}{nC_F + (1-n)\,C_F^{*}} \widehat{c}_{F,t} + \frac{(1-n)\,C_F^{*}}{nC_F + (1-n)\,C_F^{*}} \widehat{c}_{F,t}^{*} \\
\pi_{F,t}^{*} &= \frac{(1-\theta)(1-\beta\theta)}{\theta} \left(\widehat{mc}_t^{*} - \widehat{p}_{F,t}^{*}\right) + \beta \mathbb{E}_t\left[\pi_{F,t+1}^{*}\right]
\end{aligned}$$

Goods market clearing

$$(1-\omega_{oy})\widehat{n}_t^{*} + \omega_{oy} \widehat{o}_{Y,t}^{*} = \frac{(1-n)\,C_F^{*}}{(1-n)\,C_F^{*} + nC_F} \widehat{c}_{F,t}^{*} + \frac{nC_F}{(1-n)\,C_F^{*} + nC_F} \widehat{c}_{F,t}.$$

Government

$$i_t^{*} = \iota + \varphi_\pi^{*} \pi_{F,t}^{*}.$$

International conditions (3 equations).

$$\begin{aligned}
\widehat{c}_t - \widehat{c}_t^{*} &= \frac{\widehat{q}_t}{\sigma}, \\
\widehat{q}_t &= \widehat{e}_t + \widehat{p}_t^{*} - \widehat{p}_t, \\
nO_Y \widehat{o}_{Y,t} + (1-n)\,O_Y^{*} \widehat{o}_{Y,t}^{*} &= nO\widehat{o}_t + (1-n)O^{*}\widehat{o}_t^{*}.
\end{aligned}$$

F.1.10 Euler Equations, Phillips Curves, Oil Price, and Risk Sharing

This section reduces all the equilibrium equations to only six: two Euler equations, two Phillips curves, the Backus-Smith condition and the equilibrium on a global oil market. The unknowns are consumption $(\widehat{c}_t, \widehat{c}_t^{*})$, inflation $(\pi_{H,t}, \pi_{F,t}^{*})$, the real exchange rate q_t and the real price of oil expressed in units of *foreign* goods $\widehat{r} \equiv \widehat{p}_{O,t} - \widehat{p}_{F,t} = \widehat{p}_{O,t}^{*} - \widehat{p}_{F,t}^{*}$.

Euler equations. Rewrite the Euler equation as

$$\mathbb{E}_t\left[\widehat{c}_{t+1}\right] - \widehat{c}_t = \frac{1}{\sigma}\left(i_t - \mathbb{E}_t\left[\pi_{H,t+1}\right] - \iota\right) - \frac{1}{\sigma} \cdot \frac{1-\omega}{\omega + \omega^{*} - 1} \mathbb{E}_t\left[\Delta \widehat{q}_{t+1}\right].$$

The foreign Euler equation is

$$\mathbb{E}_t\left[\widehat{c}_{t+1}^*\right]-\widehat{c}_t^*=\frac{1}{\sigma}\left(i_t^*-\mathbb{E}_t\left[\pi_{F,t+1}^*\right]-\iota\right)+\frac{1}{\sigma}\cdot\frac{1-\omega^*}{\omega+\omega^*-1}\mathbb{E}_t\left[\Delta\widehat{q}_{t+1}\right].$$

Consumption. Rewrite consumption of domestic goods through overall consumption and prices at home

$$\widehat{c}_{H,t}=\frac{1-\omega}{\omega+\omega^*-1}\gamma_n\widehat{q}_t+\widehat{c}_t,$$

and abroad

$$\widehat{c}_{H,t}^*=\frac{\omega^*}{\omega+\omega^*-1}\gamma_n\widehat{q}_t+\widehat{c}_t^*.$$

And consumption of foreign goods

$$\widehat{c}_{F,t}=-\frac{\omega}{\omega+\omega^*-1}\gamma_n\widehat{q}_t+\widehat{c}_t,$$

and

$$\widehat{c}_{F,t}^*=-\frac{1-\omega^*}{\omega+\omega^*-1}\gamma_n\widehat{q}_t+\widehat{c}_t^*.$$

Production. Domestic production of home goods

$$\begin{aligned}\widehat{y}_{H,t}=&\frac{n\omega C}{n\omega C+(1-n)\,(1-\omega^*)C^*}\widehat{c}_{H,t}+\frac{(1-n)\,(1-\omega^*)\,C^*}{n\omega C+(1-n)\,(1-\omega^*)C^*}\widehat{c}_{H,t}^*\\=&\widehat{c}_t^M+\frac{n\omega C\,(1-\omega)+(1-n)\,(1-\omega^*)C^*\omega^*}{n\omega C+(1-n)\,(1-\omega^*)C^*}\cdot\frac{\gamma_n}{\omega+\omega^*-1}\widehat{q}_t.\end{aligned}$$

where

$$\widehat{c}_t^M\equiv\frac{n\omega C}{n\omega C+(1-n)\,(1-\omega^*)\,C^*}\widehat{c}_t+\frac{(1-n)\,(1-\omega^*)\,C^*}{n\omega C+(1-n)\,(1-\omega^*)\,C^*}\widehat{c}_t^*.$$

Foreign production of foreign goods

$$\begin{aligned}\widehat{y}_{F,t}^*=&\frac{nC_F}{nC_F+(1-n)\,C_F^*}\widehat{c}_{F,t}+\frac{(1-n)\,C_F^*}{nC_F+(1-n)\,C_F^*}\widehat{c}_{F,t}^*\\=&\widehat{c}_t^{M*}-\frac{n\omega\,(1-\omega)\,C+(1-n)\,(1-\omega^*)\,\omega^*C^*}{n\,(1-\omega)\,C+(1-n)\,\omega^*C^*}\cdot\frac{\gamma_n}{\omega+\omega^*-1}\widehat{q}_t\end{aligned}$$

where

$$\widehat{c}_t^{M,*}\equiv\frac{n\,(1-\omega)\,C}{n\,(1-\omega)\,C+(1-n)\,\omega^*C^*}\widehat{c}_t+\frac{(1-n)\,\omega^*C^*}{n\,(1-\omega)\,C+(1-n)\,\omega^*C^*}\widehat{c}_t^*.$$

Home marginal costs. Replace real wage (using the labor supply equation) from the labor demand and the equation for the marginal costs

$$\widehat{mc}_t-\widehat{p}_t=(1-\omega_{oy})\,(\sigma\widehat{c}_t+\varphi\widehat{n}_t)+\omega_{oy}\,(\widehat{p}_{O,t}-\widehat{p}_t)\,,$$

$$
\begin{aligned}
\widehat{n}_t =& -\frac{\gamma_y}{1+\gamma_y\varphi}\sigma\widehat{c}_t + \frac{\gamma_y}{1+\gamma_y\varphi}\left(\widehat{mc}_t - \widehat{p}_t\right) \\
&+ \frac{1}{1+\gamma_y\varphi}\left(\frac{nC_H}{nC_H+(1-n)\,C_H^*}\widehat{c}_{H,t} + \frac{(1-n)\,C_H^*}{nC_H+(1-n)\,C_H^*}\widehat{c}^*_{H,t}\right).
\end{aligned}
$$

Solve for equilibrium labor and the marginal costs

$$
\begin{aligned}
\widehat{mc}_t - \widehat{p}_t =& \frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}\left[\sigma\widehat{c}_t + \varphi\left(\frac{nC_H}{nC_H+(1-n)\,C_H^*}\widehat{c}_{H,t} + \frac{(1-n)\,C_H^*}{nC_H+(1-n)\,C_H^*}\widehat{c}^*_{H,t}\right)\right] \\
&+ \frac{1+\gamma_y\varphi}{1+\gamma_y\varphi\omega_{oy}}\omega_{oy}\left(\widehat{p}_{O,t} - \widehat{p}_t\right).
\end{aligned}
$$

and

$$
\begin{aligned}
\widehat{n}_t =& -\frac{\sigma\gamma_y\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}\widehat{c}_t + \frac{\gamma_y\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}\left(\widehat{p}_{O,t} - \widehat{p}_t\right) \\
&+ \frac{1}{1+\gamma_y\varphi\omega_{oy}}\left(\frac{nC_H}{nC_H+(1-n)\,C_H^*}\widehat{c}_{H,t} + \frac{(1-n)\,C_H^*}{nC_H+(1-n)\,C_H^*}\widehat{c}^*_{H,t}\right).
\end{aligned}
$$

Replace the CPI $\widehat{p}_t$ with $\widehat{p}_{H,t}$ in the marginal cost equation above at home

$$
\begin{aligned}
\widehat{mc}_t - \widehat{p}_{H,t} =& \frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}\left[\sigma\widehat{c}_t + \varphi\left(\frac{nC_H}{nC_H+(1-n)\,C_H^*}\widehat{c}_{H,t} + \frac{(1-n)\,C_H^*}{nC_H+(1-n)\,C_H^*}\widehat{c}^*_{H,t}\right)\right] \\
&+ \frac{1+\gamma_y\varphi}{1+\gamma_y\varphi\omega_{oy}}\omega_{oy}\left(\widehat{p}_{O,t} - \widehat{p}_{F,t}\right) + \frac{\left(1+\gamma_y\varphi\right)\omega_{oy} + (1-\omega)\left(1-\omega_{oy}\right)}{\left(1+\gamma_y\varphi\omega_{oy}\right)\left(\omega+\omega^*-1\right)}\widehat{q}_t.
\end{aligned}
$$

Rewrite the marginal costs

$$
\widehat{mc}_t - \widehat{p}_{H,t} = \frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}\left(\sigma\widehat{c}_t + \varphi\widehat{c}_t^M\right) + \psi_o\left(\widehat{p}_{O,t} - \widehat{p}_{F,t}\right) + \psi_q\widehat{q}_t,
$$

where

$$
\begin{aligned}
\psi_o \equiv& \frac{\left(1+\gamma_y\varphi\right)\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}, \\
\psi_q \equiv& \frac{\left(1+\gamma_y\varphi\right)\omega_{oy} + (1-\omega)\left(1-\omega_{oy}\right)}{\left(1+\gamma_y\varphi\omega_{oy}\right)\left(\omega+\omega^*-1\right)} \\
&+ \frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}} \cdot \frac{1}{\omega+\omega^*-1} \cdot \frac{n\omega\,(1-\omega)\,C + (1-n)\,(1-\omega^*)\omega^*C^*}{n\omega C + (1-n)\,(1-\omega^*)C^*}\varphi\gamma_n
\end{aligned}
$$

$$
\psi_q = \frac{\left(1+\gamma_y\varphi\right)\omega_{oy} + (1-\omega)\left(1-\omega_{oy}\right)}{\left(1+\gamma_y\varphi\omega_{oy}\right)\left(\omega+\omega^*-1\right)} + \frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}} \cdot \frac{1}{\omega+\omega^*-1} \cdot \frac{n\omega\,(1-\omega)\,C + (1-n)\,(1-\omega^*)\,\omega^*C^*}{nC_H + (1-n)\,C_H^*}\varphi\gamma_n.
$$

Foreign marginal costs. A similar expression for foreign country marginal costs is

$$\begin{aligned}\widehat{mc}_t^* - \widehat{p}_t^* =& \frac{1-\omega_{oy}}{1+\omega_{oy}\gamma_y\varphi}\left[\sigma\widehat{c}_t^* + \varphi\left(\frac{nC_F}{nC_F+(1-n)\,C_F^*}\widehat{c}_{F,t} + \frac{(1-n)\,C_F^*}{nC_F+(1-n)\,C_F^*}\widehat{c}_{F,t}^*\right)\right] \\ &+ \omega_{oy}\frac{1+\varphi\gamma_y}{1+\omega_{oy}\gamma_y\varphi}\left(\widehat{p}_{O,t}^* - \widehat{p}_t^*\right)\end{aligned}$$

and replace $\widehat{p}_t^*$

$$\begin{aligned}\widehat{mc}_t^* - \widehat{p}_{F,t}^* =& \frac{1-\omega_{oy}}{1+\omega_{oy}\gamma_y\varphi}\left[\sigma\widehat{c}_t^* + \varphi\left(\frac{nC_F}{nC_F+(1-n)\,C_F^*}\widehat{c}_{F,t} + \frac{(1-n)\,C_F^*}{nC_F+(1-n)\,C_F^*}\widehat{c}_{F,t}^*\right)\right] \\ &+ \omega_{oy}\frac{1+\varphi\gamma_y}{1+\omega_{oy}\gamma_y\varphi}\left(\widehat{p}_{O,t}^* - \widehat{p}_{F,t}^*\right) - \frac{1-\omega_{oy}}{1+\omega_{oy}\gamma_y\varphi}\cdot\frac{1-\omega^*}{(\omega+\omega^*-1)}\widehat{q}_t\end{aligned}$$

Foreign country marginal costs are

$$\widehat{mc}_t^* - \widehat{p}_{F,t}^* = \frac{1-\omega_{oy}}{1+\omega_{oy}\gamma_y\varphi}\left(\sigma\widehat{c}_t^* + \varphi\widehat{c}_t^{M,*}\right) + \psi_o^*\left(\widehat{p}_{O,t}^* - \widehat{p}_{F,t}^*\right) + \psi_q^*\widehat{q}_t$$

where

$$\begin{aligned}\psi_o^* \equiv& \frac{\omega_{oy}\left(1+\varphi\gamma_y\right)}{1+\omega_{oy}\gamma_y\varphi} = \psi_o, \\ \psi_q^* \equiv& -\frac{1-\omega_{oy}}{1+\omega_{oy}\gamma_y\varphi}\cdot\frac{1-\omega^*}{\omega+\omega^*-1} \\ &-\frac{1-\omega_{oy}}{1+\omega_{oy}\gamma_y\varphi}\cdot\frac{1}{\omega+\omega^*-1}\cdot\frac{n\omega C_F+(1-n)\,(1-\omega^*)\,C_F^*}{nC_F+(1-n)\,C_F^*}\varphi\gamma_n.\end{aligned}$$

Note that ψ_q^* and ψ_q are not symmetric. This is because in both equations for the marginal costs the real oil price is expressed in terms of foreign goods.

Phillips curves. The home Phillips curve is

$$\pi_{H,t} = \kappa\left\{\frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}\left[\sigma\widehat{c}_t + \varphi\widehat{c}_t^M\right] + \psi_o\widehat{r}_t + \psi_q\widehat{q}_t\right\} + \beta\mathbb{E}_t\left[\pi_{H,t+1}\right],$$

and the foreign Phillips curve is

$$\pi_{F,t}^* = \kappa\left\{\frac{1-\omega_{oy}}{1+\omega_{oy}\gamma_y\varphi}\left[\sigma\widehat{c}_t^* + \varphi\widehat{c}_t^{M,*}\right] + \psi_o^*\widehat{r}_t + \psi_q^*\widehat{q}_t\right\} + \beta\mathbb{E}_t\left[\pi_{F,t+1}^*\right],$$

where

$$\kappa \equiv \frac{(1-\theta)(1-\beta\theta)}{\theta}. \tag{F.24}$$

Oil demand. The home firms demand for oil

$$\begin{aligned}\widehat{o}_{Y,t} =&\gamma_y \frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}\sigma\widehat{c}_t - \frac{1-\omega_{oy}}{1+\omega_{oy}\gamma_y\varphi}\gamma_y\widehat{r}_t + \frac{1+\varphi\gamma_y}{1+\gamma_y\varphi\omega_{oy}}\widehat{c}_t^M \\ &+\left\{\left[(\omega+\omega^*-1)\,\psi_q-1\right]\gamma_y + \frac{n\omega C\,(1-\omega)+(1-n)\,(1-\omega^*)C^*\omega^*}{n\omega C+(1-n)\,(1-\omega^*)C^*}\gamma_n\right\}\frac{\widehat{q}_t}{\omega+\omega^*-1}.\end{aligned}$$

The foreign firm demand for oil

$$\begin{aligned}\widehat{o}^*_{Y,t} =&\gamma_y \frac{1-\omega_{oy}}{1+\omega_{oy}\gamma_y\varphi}\sigma\widehat{c}^*_t - \frac{1-\omega_{oy}}{1+\omega_{oy}\gamma_y\varphi}\gamma_y\widehat{r}_t + \frac{1+\gamma_y\varphi}{1+\omega_{oy}\gamma_y\varphi}\widehat{c}_t^{M*} \\ &-\frac{1}{1+\omega_{oy}\gamma_y\varphi}\left(\gamma_n\frac{n\omega C_F+(1-n)\,(1-\omega^*)\,C^*_F}{nC_F+(1-n)\,C^*_F}+\gamma_y\left(1-\omega_{oy}\right)(1-\omega^*)\right)\frac{\widehat{q}_t}{\omega+\omega^*-1}\end{aligned}$$

Oil market. Oil market equilibrium condition

$$nO_Y\widehat{o}_{Y,t}+(1-n)O^*_Y\widehat{o}^*_{Y,t} = nO\widehat{o}_t+(1-n)O^*\widehat{o}^*_t,$$

which determines $\widehat{r}_t = \widehat{p}^*_{O,t}-\widehat{p}^*_{F,t}$.

F.1.11 A Small-Open Economy Limit

We now take the limit as the size of home economy approaches zero and the size of the foreign economy approaches one, that is, $n \to 0$ and $n^* = 1-n \to 1$. Taking into account the following definitions

$$\begin{aligned}1-\omega &= (1-n)\,(1-\Omega)\,,\\ 1-\omega^* &= n\,(1-\Omega^*)\,,\end{aligned}$$

we have that in the limit the fraction of domestic goods expenditure

$$\begin{aligned}\omega &= \Omega,\\ \omega^* &= 1.\end{aligned}$$

This implies that

$$\begin{aligned}\frac{1-\omega^*}{\omega+\omega^*-1} &= \frac{n\,(1-\Omega^*)}{1-(1-n)\,(1-\Omega)-n\,(1-\Omega^*)} = 0,\\ \frac{nC_H}{nC_H+(1-n)\,C^*_H} &= \frac{\Omega C}{\Omega C+(1-\Omega^*)\,C^*},\\ \frac{nC_F}{nC_F+(1-n)\,C^*_F} &= 0,\\ \frac{n\omega C_F+(1-n)\,(1-\omega^*)\,C^*_F}{nC_F+(1-n)\,C^*_F} &= 0.\end{aligned}$$

The relative prices of home goods in units of domestic and foreign consumption baskets are

$$\begin{aligned}
\widehat{p}_{H,t} - \widehat{p}_t &= -\frac{1-\omega}{\omega+\omega^*-1}\widehat{q}_t = -\frac{1-\Omega}{\Omega}\widehat{q}_t, \\
\widehat{p}^*_{H,t} - \widehat{p}^*_t &= -\frac{1}{\Omega}\widehat{q}_t.
\end{aligned}$$

The world economy. The world equilibrium consists of six unknowns $\left(\widehat{y}^*_F, \widehat{c}^*_t, \pi^*_t, \pi^*_{F,t}, \widehat{r}_t, i^*_t\right)$ and six equations are

$$\begin{aligned}
\widehat{r}_t &= \phi_c \widehat{c}^*_t - \phi_o \widehat{o}^*_t, \\
\pi^*_{F,t} &= \kappa\zeta^*\widehat{c}^*_t + \beta\mathbb{E}_t\left[\pi^*_{F,t+1}\right] + \kappa\psi^*_o\widehat{r}_t, \\
\mathbb{E}_t\left[\widehat{c}^*_{t+1}\right] - \widehat{c}^*_t &= \frac{1}{\sigma}\left(i^*_t - \mathbb{E}_t\left[\pi^*_{t+1}\right] - \iota\right), \\
i^*_t &= \iota + \varphi^*_\pi \pi^*_{F,t}, \\
\pi^*_t &= \pi^*_{F,t}, \\
\widehat{y}^*_{F,t} &= \widehat{c}^*_t.
\end{aligned}$$

where $\widehat{r}_t \equiv \widehat{p}^*_{O,t} - \widehat{p}^*_{F,t}$ and

$$\phi_o = \frac{1+\omega_{oy}\gamma_y\varphi}{\left(1-\omega_{oy}\right)\gamma_y}, \tag{F.25}$$

$$\phi_c = \frac{1+\gamma_y\left[\varphi + \left(1-\omega_{oy}\right)\sigma\right]}{\left(1-\omega_{oy}\right)\gamma_y}, \tag{F.26}$$

$$\kappa \equiv \frac{(1-\theta)(1-\beta\theta)}{\theta},$$

$$\zeta^* \equiv \frac{1-\omega_{oy}}{1+\omega_{oy}\gamma_y\varphi}(\sigma+\varphi), \tag{F.27}$$

$$\psi^*_o = \frac{\left(1+\gamma_y\varphi\right)\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}. \tag{F.28}$$

We took into account that $\psi^*_q = 0, \widehat{c}^{M*}_t = \widehat{c}^*_t$.

Substituting oil price into the Phillips curve, we get

$$\pi^*_{F,t} = \kappa\left(\zeta^* + \psi^*_o\phi_c\right)\widehat{c}^*_t + \beta\mathbb{E}_t\left[\pi^*_{F,t+1}\right] - \kappa\psi^*_o\phi_o\widehat{o}^*_t,$$

The value added is

$$\widehat{z}^*_{F,t} = \frac{\left(1-\omega_{oy}\phi_c\right)\widehat{c}^*_t - \omega_{oy}\left(1-\phi_o\right)\widehat{o}^*_t}{1-\omega_{oy}}.$$

Trade balance is

$$\widehat{nx}^*_t = 0.$$

Small-open economy. The SOE block consists of six unknowns $(\widehat{c}_t, \widehat{y}_{H,t}, \pi_t, \pi_{H,t}, i_t, q_t)$ and six equations

$$
\begin{aligned}
\mathbb{E}_t\left[\widehat{c}_{t+1}\right]-\widehat{c}_t &= \frac{1}{\sigma}\left(i_t-\mathbb{E}_t\left[\pi_{t+1}\right]-\iota\right), \\
\pi_{H,t} &= \kappa\zeta_H\widehat{c}_t+\beta\mathbb{E}_t\left[\pi_{H,t+1}\right]+\kappa\psi_o\widehat{r}_t+\kappa\zeta_F\widehat{c}_t^*+\kappa\psi_q\widehat{q}_t, \\
\pi_t &= \pi_{H,t}+\frac{1-\Omega}{\Omega}\Delta\widehat{q}_t, \\
i_t &= \iota+\phi_\pi\pi_{H,t}, \\
\widehat{q}_t &= \sigma\left(\widehat{c}_t-\widehat{c}_t^*\right), \\
\widehat{y}_{H,t} &= \frac{\Omega}{\Omega+1-\Omega^*}\widehat{c}_t+\frac{1-\Omega^*}{\Omega+1-\Omega^*}\widehat{c}_t^*+\frac{\Omega\left(1-\Omega\right)+\left(1-\Omega^*\right)}{\Omega+1-\Omega^*}\cdot\frac{\gamma_n}{\Omega}\widehat{q}_t.
\end{aligned}
$$

where

$$
\begin{aligned}
\zeta_H &= \frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}\left(\sigma+\varphi\frac{\Omega}{\Omega+1-\Omega^*}\right), \\
\zeta_F &= \frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}\varphi\frac{1-\Omega^*}{\Omega+1-\Omega^*}, \\
\psi_q &= \frac{\left(1+\gamma_y\varphi\right)\omega_{oy}+\left(1-\Omega\right)\left(1-\omega_{oy}\right)}{\left(1+\gamma_y\varphi\omega_{oy}\right)\Omega}+\frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}\cdot\frac{1}{\Omega}\cdot\frac{\Omega\left(1-\Omega\right)+\left(1-\Omega^*\right)}{\Omega+1-\Omega^*}\varphi\gamma_n, \\
\psi_o &= \frac{\left(1+\gamma_y\varphi\right)\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}},
\end{aligned}
$$

and we took into account the fact that $\widehat{c}_t^M=\widehat{c}_t\Omega/\left(\Omega+1-\Omega^*\right)+\widehat{c}_t^*\left(1-\Omega^*\right)/\left(\Omega+1-\Omega^*\right)$.[55] Note that the Phillips curve has three additional terms compared to the standard closed-economy formulation. The term $\kappa\psi_o\widehat{r}_t$ shows that higher real oil price increase the marginal cost and, hence, inflation. The term $\kappa\zeta_F\widehat{c}_t^*$ is due to the fact that the world aggregate demand affects the demand for home products and increases the cost of production. The term $\kappa\psi_q\widehat{q}_t$ reflects the fact that oil prices are quoted in units of foreign goods. This implies that absence any chance in the real oil price (in units of foreign goods), a real depreciation of domestic currency (an increase in $\widehat{q}_t$) acts to increase the oil price in units of home goods, which, in turn, increases the marginal cost and inflation. Finally, note that $\zeta_H+\zeta_F=\zeta^*$.

Substituting away the real exchange rate from the Euler equation and the Phillips curve, we obtain

$$
\begin{aligned}
\mathbb{E}_t\left[\widehat{c}_{t+1}\right]-\widehat{c}_t &= \frac{\Omega}{\sigma}\left(i_t-\mathbb{E}_t\left[\pi_{H,t+1}\right]-\iota\right)+\left(1-\Omega\right)\left\{\mathbb{E}_t\left[\widehat{c}_{t+1}^*\right]-\widehat{c}_t^*\right\}, \\
\pi_{H,t} &= \kappa\left(\zeta_H+\psi_q\sigma\right)\widehat{c}_t+\beta\mathbb{E}_t\left[\pi_{H,t+1}\right]+\kappa\psi_o\widehat{r}_t+\kappa\left(\zeta_F-\psi_q\sigma\right)\widehat{c}_t^*,
\end{aligned}
$$

This form of the Euler equation implies that the home aggregate demand is affected by the rest of the world because it changes the relative price of imported goods. A booming world economy has two opposing effects on the Phillips curve. On the one hand, this boom increases demand for home goods and pushes up inflation. On the other hand, a booming world economy appreciates the home currency in real terms and makes oil cheaper at home. The net effect of these two forces is captured by the coefficient

$$
\zeta_F-\psi_q\sigma=\frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}\left[\frac{1-\Omega^*}{\Omega+1-\Omega^*}\varphi-\frac{1-\Omega}{\Omega}\left(1+\varphi\gamma_n\right)\sigma\right].
$$

[55]The formula for ψ_q can be further expressed as $\psi_q=\omega_{oy}+(1-\Omega)/\Omega+[\omega_{oy}+[1/\Omega-(1-\omega_{oy})\Omega]\gamma_n/\gamma_y]\gamma_y\varphi(1-\omega_{oy})/(1+\gamma_y\varphi\omega_{oy})$ by taking account of the fact that $\Omega^*=\Omega-\omega_{oy}/(1-\omega_{oy})$.

In general, the sign of this coefficient is ambiguous. However, under the empirical relevant parameters $\Omega = \Omega^* = 2/3, \sigma = 1, \varphi = 2$, this coefficient is negative.

Oil demand is

$$\begin{aligned}\widehat{o}_{Y,t} =& \gamma_y \frac{1-\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}} \left(\sigma\widehat{c}_t - \widehat{r}_t\right) + \frac{1+\varphi\gamma_y}{1+\gamma_y\varphi\omega_{oy}} \left(\frac{\Omega}{\Omega+1-\Omega^*}\widehat{c}_t + \frac{1-\Omega^*}{\Omega+1-\Omega^*}\widehat{c}_t^*\right) \\ &+ \left\{\left(\Omega\psi_q - 1\right)\gamma_y + \frac{\Omega\left(1-\Omega\right)+\left(1-\Omega^*\right)}{\Omega+1-\Omega^*}\gamma_n\right\}\frac{\widehat{q}_t}{\Omega}.\end{aligned}$$

Trade balance is

$$\widehat{nx}_t = \widehat{y}_{H,t} - \omega_{oy}\widehat{r}_t - \left(\omega_{oy} + \frac{1-\Omega}{\Omega+1-\Omega^*}\right)\frac{\widehat{q}_t}{\Omega} - \omega_{oy}\widehat{o}_{Y,t} - \frac{1}{\Omega+1-\Omega^*}\widehat{c}_t.$$

F.1.12 Neo-classical Effects

When prices are completely flexible, the Phillips curve in foreign country implies

$$\widehat{c}_t^* = \Gamma_O^{c,*}\phi_o\widehat{o}_t^*.$$

where

$$\Gamma_O^{c,*} \equiv \frac{\psi_o^*}{\zeta^* + \psi_o^*\phi_c} = \frac{\frac{\omega_{oy}}{1-\omega_{oy}}\left(1+\varphi\gamma_y\right)}{\sigma+\varphi+\frac{\omega_{oy}}{1-\omega_{oy}}\left(1+\varphi\gamma_y\right)\frac{1-\gamma_y\omega_{oy}\sigma+\gamma_y(\varphi+\sigma)}{\gamma_y\left(1-\omega_{oy}\right)}}.$$

The consumption effect of oil supply change

$$\begin{aligned}\Gamma_O^{c,*}\phi_o =& \frac{\frac{\omega_{oy}}{1-\omega_{oy}}\left(1+\varphi\gamma_y\right)}{\frac{\left(1-\omega_{oy}\right)\gamma_y}{\left(1+\omega_{oy}\gamma_y\varphi\right)}\left[\left(\sigma+\varphi\right)+\frac{\omega_{oy}}{1-\omega_{oy}}\left(1+\varphi\gamma_y\right)\frac{\left(1-\gamma_y\omega_{oy}\sigma\right)+\gamma_y(\varphi+\sigma)}{\left(1-\omega_{oy}\right)\gamma_y}\right]} \\ =& \frac{\omega_{oy}\left(1+\varphi\gamma_y\right)}{\omega_{oy}+\gamma_y\left(\sigma+\varphi-\omega_{oy}\sigma\right)}.\end{aligned}$$

The oil price

$$\begin{aligned}\widehat{r}_t =& -\frac{(\sigma+\varphi)}{\sigma+\varphi+\frac{\omega_{oy}}{1-\omega_{oy}}\left(1+\varphi\gamma_y\right)\frac{\left(1-\gamma_y\omega_{oy}\sigma\right)+\gamma_y(\varphi+\sigma)}{\left(1-\omega_{oy}\right)\gamma_y}}\phi_o\widehat{o}_t^* \\ =& -\frac{(\sigma+\varphi)\left(1-\omega_{oy}\right)}{\omega_{oy}+\gamma_y\left(\sigma+\varphi-\omega_{oy}\sigma\right)}\widehat{o}_t^*.\end{aligned}$$

A change in output normalized by a change in oil price

$$\frac{\widehat{c}_t^*}{\widehat{r}_t} = \frac{\frac{\omega_{oy}\left(1+\varphi\gamma_y\right)}{\omega_{oy}+\gamma_y\left(\sigma+\varphi-\omega_{oy}\sigma\right)}}{-\frac{(\sigma+\varphi)\left(1-\omega_{oy}\right)}{\omega_{oy}+\gamma_y\left(\sigma+\varphi-\omega_{oy}\sigma\right)}} = -\frac{\omega_{oy}}{1-\omega_{oy}}\cdot\frac{1+\varphi\gamma_y}{\sigma+\varphi}.$$

F.1.13 A Two-state Markov Shock

With probability α, the shock persists, with the remaining probability it goes away. Oil shock $\widehat{o}_t^* \in \{\widehat{o}^*, 0\}$,where $\widehat{o}^* < 0$.

The World Economy We start from the world economy. Using the equation for the oil price in equilibrium, we obtain the following Phillips curve

$$\pi_F^* = \frac{\kappa\left(\zeta^* + \psi_o^* \phi_c\right)}{1-\beta\alpha}\widehat{c}^* + \frac{\kappa\psi_o^*\phi_o}{1-\beta\alpha}\left(-\widehat{o}^*\right), \tag{F.29}$$

The Euler equation is

$$\widehat{c}^* = -\frac{\left(\varphi_\pi^* - \alpha\right)\pi_F^*}{\left(1-\alpha\right)\sigma}. \tag{F.30}$$

The response of the inflation rate that solves the above two equations is

$$\frac{d\pi_F^*}{d\left(-\widehat{o}^*\right)} = \phi_o \frac{\frac{\kappa\psi_o^*}{(1-\beta\alpha)}}{1 + \frac{\kappa(\varphi_\pi^*-\alpha)}{(1-\beta\alpha)(1-\alpha)\sigma}\left(\zeta^* + \psi_o^*\phi_c\right)} > 0. \tag{F.31}$$

There are two notable features here. First, the effect on inflation is non-negative (when the denominator is positive). Second, when the shock is permanent, inflation response is zero. The response of inflation rate when oil shock is such that the oil price increases by one percent is

$$\pi_F^* = \frac{\frac{\kappa\psi_o^*}{1-\beta\alpha}}{1 + \frac{\kappa(\varphi_\pi^*-\alpha)}{(1-\beta\alpha)(1-\alpha)\sigma}\zeta^*}.$$

The response of consumption (and output) that solve the Euler equation (combined with the monetary policy rule) and the Phillips curve is

$$\frac{d\widehat{c}^*}{d\left(-\widehat{o}^*\right)} = -\phi_o \frac{\frac{\kappa(\varphi_\pi^*-\alpha)}{(1-\beta\alpha)(1-\alpha)}\psi_o^*}{\sigma + \frac{\kappa(\varphi_\pi^*-\alpha)}{(1-\beta\alpha)(1-\alpha)}\left(\zeta^* + \psi_o^*\phi_c\right)} = \frac{\left(\zeta^* + \psi_o^*\phi_c\right)\frac{\kappa(\varphi_\pi^*-\alpha)}{(1-\beta\alpha)(1-\alpha)}}{\sigma + \left(\zeta^* + \psi_o^*\phi_c\right)\frac{\kappa(\varphi_\pi^*-\alpha)}{(1-\beta\alpha)(1-\alpha)}}\Gamma_y^{nc,*} \lesseqgtr 0. \tag{F.32}$$

where $\Gamma_y^{nc,*} = -\phi_o\psi_o^* / \left(\zeta^* + \psi_o^*\phi_c\right)$ is the neo-classical response. There are several notable observations here. When the Taylor rule response to the inflation rate is strong enough, that is, $\varphi_\pi^* - \alpha > 0$, the aggregate consumption unambiguously falls following a hike in oil prices. This is because of the increase in the real interest rate and the fall in demand for oil consumption. Second, when the shock is permanent: $d\widehat{c}^*/d(-\widehat{o}^*) = -\phi_o\psi_o^*/(\zeta^* + \psi_o^*\phi_c) < 0$. Third, when goods prices are completely sticky, i.e., $\kappa = 0$, or the central bank targets a fixed real rate, i.e., $\varphi_\pi^* = \alpha$, the response of consumption is zero: $d\widehat{c}^*/d(-\widehat{o}^*) = 0$. Fourth, the response is smaller compared to the case of completely flexible prices when $\varphi_\pi^* - \alpha > 0$.

The reaction of the oil price is

$$\widehat{r} = \frac{1 + \frac{\kappa(\varphi_\pi^*-\alpha)}{(1-\beta\alpha)(1-\alpha)\sigma}\zeta^*}{1 + \frac{\kappa(\varphi_\pi^*-\alpha)}{(1-\beta\alpha)(1-\alpha)\sigma}\left(\zeta^* + \psi_o^*\phi_c\right)}\phi_o\left(-\widehat{o}^*\right). \tag{F.33}$$

Note that an oil supply decline unambiguously raises oil price when $\varphi_\pi^* - \alpha > 0$. The size of the oil shock that increases the oil price by one percent equals

$$\phi_o\left(-\widehat{o}^*\right)=\frac{1+\frac{\kappa(\varphi_\pi^*-\alpha)}{(1-\beta\alpha)(1-\alpha)\sigma}\left(\zeta^*+\psi_o^*\phi_c\right)}{1+\frac{\kappa(\varphi_\pi^*-\alpha)}{(1-\beta\alpha)(1-\alpha)\sigma}\zeta^*}.$$

The log-deviation of the consumption level from its steady state c^* following the oil supply shock that increases oil price by one percent, $\widehat{r}=1$, is

$$\widehat{c}^*=-\frac{\frac{\kappa(\varphi_\pi^*-\alpha)}{\sigma(1-\beta\alpha)(1-\alpha)}\zeta^*}{1+\frac{\kappa(\varphi_\pi^*-\alpha)}{(1-\beta\alpha)(1-\alpha)\sigma}\zeta^*}\cdot\frac{\psi_o^*}{\zeta^*}. \tag{F.34}$$

When prices are flexible, i.e., $\kappa\to\infty$,we have $\widehat{c}^*=-\psi_o^*/\zeta^*$, while when they are completely rigid, consumption does not respond, i.e., $\widehat{c}^*=0$.

ZLB. All the above formulas can be applied to the case of the liquidity trap or inelastic interest rates by setting $\varphi_\pi^*=0$. We have

$$\begin{aligned}
\frac{d\pi_F^*}{d\left(-\widehat{o}^*\right)}&=\phi_o\frac{\frac{\kappa\psi_o^*}{1-\beta\alpha}}{1-\frac{\kappa\alpha}{(1-\beta\alpha)(1-\alpha)\sigma}\left(\zeta^*+\psi_o^*\phi_c\right)}>0,\\
\frac{d\widehat{c}^*}{d\left(-\widehat{o}^*\right)}&=-\phi_o\frac{-\frac{\alpha\kappa\psi_o^*}{(1-\beta\alpha)(1-\alpha)\sigma}}{1-\frac{\alpha\kappa}{(1-\beta\alpha)(1-\alpha)\sigma}\left(\zeta^*+\psi_o^*\phi_c\right)}>0,\\
\frac{d\widehat{r}}{d\left(-\widehat{o}^*\right)}&=\phi_o\frac{1-\frac{\alpha\kappa\zeta^*}{(1-\beta\alpha)(1-\alpha)\sigma}}{1-\frac{\alpha\kappa}{(1-\beta\alpha)(1-\alpha)\sigma}\left(\zeta^*+\psi_o^*\phi_c\right)}>0.
\end{aligned}$$

For the equilibrium to be unique, we assume that

$$1-\frac{\alpha\kappa}{\sigma\left(1-\beta\alpha\right)\left(1-\alpha\right)}\left(\zeta^*+\psi_o^*\phi_c\right)>0.$$

The inflation rate response is

$$\begin{aligned}
&\left.\frac{d\pi_F^*}{d\left(-\widehat{o}^*\right)}\right|_{\text{ZLB}}-\left.\frac{d\pi_F^*}{d\left(-\widehat{o}^*\right)}\right|_{\text{normal}}\\
=&\frac{\kappa\left(\zeta^*+\psi_o^*\phi_c\right)\varphi_\pi^*\phi_o\left(1-\alpha\right)\kappa\psi_o^*\sigma}{\left[\left(1-\beta\alpha\right)\left(1-\alpha\right)\sigma-\alpha\kappa\left(\zeta^*+\psi_o^*\phi_c\right)\right]\left[\left(1-\beta\alpha\right)\left(1-\alpha\right)\sigma+\kappa\left(\zeta^*+\psi_o^*\phi_c\right)\left(\varphi_\pi^*-\alpha\right)\right]}>0.
\end{aligned}$$

The consumption response is

$$\left.\frac{d\widehat{c}^*}{d\left(-\widehat{o}^*\right)}\right|_{\text{ZLB}}-\left.\frac{d\widehat{c}^*}{d\left(-\widehat{o}^*\right)}\right|_{\text{normal}}=\phi_o\frac{\frac{\kappa\varphi_\pi^*}{(1-\beta\alpha)(1-\alpha)}\psi_o^*\sigma}{\left[\sigma+\frac{-\alpha\kappa}{(1-\beta\alpha)(1-\alpha)}\left(\zeta^*+\psi_o^*\phi_c\right)\right]\left[\sigma+\frac{\kappa(\varphi_\pi^*-\alpha)}{(1-\beta\alpha)(1-\alpha)}\left(\zeta^*+\psi_o^*\phi_c\right)\right]}.$$

The absolute response

$$\left.\frac{d\widehat{c}^*}{d\left(-\widehat{o}^*\right)}\right|_{\text{ZLB}} = \phi_o \frac{\frac{\alpha\kappa}{(1-\beta\alpha)(1-\alpha)} \cdot \frac{(1+\gamma_y\varphi)\omega_{oy}}{1+\gamma_y\varphi\omega_{oy}}}{\sigma + \frac{\kappa(0-\alpha)}{(1-\beta\alpha)(1-\alpha)}\left(\zeta^* + \psi_o^*\phi_c\right)}$$

The SOE Economy The Euler equation

$$\widehat{c} = -\Omega\frac{(\varphi_\pi - \alpha)}{\sigma(1-\alpha)}\pi_H + (1-\Omega)\,\widehat{c}^*.$$

The Phillips curve

$$\pi_H = \frac{\kappa\left(\zeta_H + \psi_q\sigma\right)}{1-\alpha\beta}\widehat{c} + \frac{\kappa}{1-\alpha\beta}\psi_o\widehat{r} + \frac{\kappa}{1-\alpha\beta}\left(\zeta_F - \psi_q\sigma\right)\widehat{c}^*.$$

The inflation rate response is

$$\pi_H = \frac{\psi_o\frac{\kappa}{1-\alpha\beta}}{1+\frac{\kappa(\varphi_\pi-\alpha)}{(1-\alpha)(1-\alpha\beta)\sigma}\left(\zeta_H + \sigma\psi_q\right)\Omega}\widehat{r} + \frac{\frac{\kappa}{1-\alpha\beta}\left[\zeta_H(1-\Omega) + \zeta_F - \psi_q\sigma\Omega\right]}{1+\frac{\kappa(\varphi_\pi-\alpha)}{\sigma(1-\alpha\beta)(1-\alpha)}\left(\zeta_H + \psi_q\sigma\right)\Omega}\widehat{c}^*,$$

The first term represents the reaction of the economy to the world oil price, while the second term represents the reaction to a change in the global economic activity. Specifically, a higher demand in the foreign economy increases the demand for home products and, at the same time, appreciates domestic real exchange rate making it less costly to produce conditional on an unchanged real oil price in the units of foreign goods.

The response of aggregate consumption is

$$\widehat{c} = \frac{1-\Omega+\frac{\kappa(\varphi_\pi-\alpha)}{\sigma(1-\alpha\beta)(1-\alpha)}\left(\zeta_H+\psi_q\sigma\right)\Omega\frac{\psi_q\sigma-\zeta_F}{\zeta_H+\psi_q\sigma}}{1+\frac{\kappa(\varphi_\pi-\alpha)}{\sigma(1-\alpha\beta)(1-\alpha)}\left(\zeta_H+\psi_q\sigma\right)\Omega}\widehat{c}^* - \frac{\frac{(\varphi_\pi-\alpha)\kappa}{\sigma(1-\alpha)(1-\alpha\beta)}\left(\zeta_H+\sigma\psi_q\right)\Omega}{1+\frac{\kappa(\varphi_\pi-\alpha)}{(1-\alpha)(1-\alpha\beta)\sigma}\left(\zeta_H+\sigma\psi_q\right)\Omega}\cdot\frac{\psi_o}{\zeta_H+\sigma\psi_q}\widehat{r}.$$

The response of aggregate consumption conditional on the oil supply shock that increases oil price by one percent

$$\widehat{c} = \frac{1-\Omega+\Omega\frac{\kappa(\varphi_\pi-\alpha)}{\sigma(1-\alpha\beta)(1-\alpha)}\left(\psi_q\sigma-\zeta_F\right)}{1+\frac{\kappa(\varphi_\pi-\alpha)}{\sigma(1-\alpha\beta)(1-\alpha)}\left(\zeta_H+\psi_q\sigma\right)\Omega}\widehat{c}^* - \frac{\frac{(\varphi_\pi-\alpha)\kappa}{\sigma(1-\alpha)(1-\alpha\beta)}\left(\zeta_H+\sigma\psi_q\right)\Omega}{1+\frac{\kappa(\varphi_\pi-\alpha)}{(1-\alpha)(1-\alpha\beta)\sigma}\left(\zeta_H+\sigma\psi_q\right)\Omega}\cdot\frac{\psi_o}{\zeta_H+\sigma\psi_q},$$

where $\widehat{c}^*$ is from equation (F.34).

We note that when $\varphi_\pi = \varphi_\pi^*$, we get $\widehat{c} = \widehat{c}^*$.

Home production is

$$\widehat{y}_H = \frac{\Omega}{\Omega+1-\Omega^*}\widehat{c} + \frac{1-\Omega^*}{\Omega+1-\Omega^*}\widehat{c}^* + \frac{\Omega(1-\Omega)+(1-\Omega^*)}{\Omega+1-\Omega^*}\cdot\frac{\gamma_n}{\Omega}\widehat{q}.$$

This formula clearly illustrates that the oil shock affects the non-oil production through three distinct channels corresponding to three terms in the formula: (i) a change in domestic aggregate demand; (ii) a change in foreign aggregate demand; (iii) an expenditure switching effect from foreign to domestic goods.

We replace the real exchange rate and domestic consumption

$$\begin{aligned}\widehat{y}_H &= \frac{\Omega}{\Omega+1-\Omega^*}\widehat{c}_t + \frac{1-\Omega^*}{\Omega+1-\Omega^*}\widehat{c}_t^* + \frac{\Omega\left(1-\Omega\right)+\left(1-\Omega^*\right)}{\Omega+1-\Omega^*}\cdot\frac{\gamma_n}{\Omega}\sigma\left(\widehat{c}-\widehat{c}^*\right)\\ &= \frac{\Omega}{\Omega+1-\Omega^*}\left(1+\frac{\Omega\left(1-\Omega\right)+\left(1-\Omega^*\right)}{\Omega}\cdot\frac{\gamma_n\sigma}{\Omega}\right)\widehat{c} + \frac{1-\Omega^*}{\Omega+1-\Omega^*}\left(1-\frac{\Omega\left(1-\Omega\right)+\left(1-\Omega^*\right)}{1-\Omega^*}\cdot\frac{\gamma_n\sigma}{\Omega}\right)\widehat{c}^*.\end{aligned}$$

The response of the real exchange rate conditional on the oil supply shock that increases oil price by one percent is

$$\widehat{q} = \sigma\left(\widehat{c}-\widehat{c}^*\right) = -\sigma\frac{\frac{\kappa(\varphi_\pi-\varphi_\pi^*)}{\sigma(1-\beta\alpha)(1-\alpha)}\psi_o\Omega}{1+\frac{\kappa(\varphi_\pi-\alpha)}{\sigma(1-\alpha\beta)(1-\alpha)}\left(\zeta_H+\psi_q\sigma\right)\Omega}\cdot\frac{1}{1+\frac{\kappa(\varphi_\pi^*-\alpha)}{(1-\beta\alpha)(1-\alpha)\sigma}\zeta^*}.$$

The last expression implies that the real exchange rate does not respond when the home and foreign central banks respond to domestic inflation in the same way, that is, $\varphi_\pi = \varphi_\pi^*$. When, however, home country is at the ZLB while foreign country actively responds to oil shock, i.e., $\varphi_\pi - \varphi_\pi^* < 0$, the real exchange rate depreciates $\widehat{q}_t > 0$.

For uniqueness of bounded ZLB solution, the following condition must hold

$$1-\frac{\alpha\kappa}{\sigma\left(1-\alpha\beta\right)\left(1-\alpha\right)}\left(\zeta_H+\psi_q\sigma\right)\Omega > 0.$$

F.2 Open-economy New Keynesian Model with Oil Inventories

In this section, we present an extension of our baseline model that incorporates oil inventories. We maintain all the assumptions of the baseline model, except that we introduce an additional type of agent: oil inventory firms.

F.2.1 Inventories

Oil inventory firms are modeled following the framework of Unalmis, Unalmis and Unsal (2012). There is a continuum of oil storage firms indexed by $j \in [0,1]$, operating in the foreign (world) economy. These firms purchase oil on the spot market, store it, and sell it in the subsequent period. Each storage firm maximizes its expected profits from holding oil, including any additional (psychological) benefits associated with inventory holdings. Specifically, an oil inventory firm j chooses the level of oil inventories $I_{t+1}^*\left(j\right)$ to solve

$$\max \mathbb{E}_t\left[M_{t,t+1}^*P_{O,t+1}^*I_{t+1}^*\left(j\right)\left(1-\delta\right)\right] - P_{O,t}^*I_{t+1}^*\left(j\right) + P_t^*S\left(I_{t+1}^*\left(j\right)\right),$$

where $M_{t,t+1}^*$ is the price of state-contingent securities in the foreign country, $P_{O,t}^*$ denotes the nominal price of oil in the foreign country, P_t^* is the price level in the foreign country, I_{t+1}^* is the inventories at the beginning of period $t+1$ that a firm chooses at time t, and $S\left(I_{t+1}^*\left(j\right)\right)$ represents the convenience of holding inventories—a common assumption in the commodity literature (Alquist and Kilian, 2010)—net of the costs. Since the net benefit of holding oil $P_t^*S\left(I_{t+1}^*\left(j\right)\right)$ is psychological, it will not enter any of the clearing conditions.[56] We assume that all the ex-post profits that the oil inventory firm make are rebated to the foreign households.

[56]The assumption that storing oil entails psychological benefits and costs echoes a modeling device commonly used in the sticky-price literature, where the cost of adjusting prices is modeled as entering the utility of firm managers (Michaillat and Saez, 2021).

The functional form of $S\left(I_{t+1}^*(j)\right)$ is given by:

$$S\left(I_{t+1}^*\right) \equiv \left[\phi_1^{O^*}\left(\frac{I_{t+1}^*}{I^*}\right) - \frac{\phi_2^{O^*}}{2}\left(\frac{I_{t+1}^*}{I^*}\right)^2\right] O_S^*,$$

where O_S^* is the steady state sales of oil (use of oil) in the foreign country, $\phi_1^{O^*}$ and $\phi_2^{O^*}$ are parameters related to the benefit and the cost of holding inventories, respectively, and δ is the depreciation of oil inventories.

Given the symmetric structure, we have $I_{t+1}^*(j) = I_{t+1}^*$. Aggregate inventory evolves according to

$$I_{t+1}^* = (1-\delta) I_t^* + O_t^* - O_{S,t}^*,$$

where O_t^* is the oil supply in the foreign country. First order condition with respect to I_{t+1}^* leads to

$$\frac{O_S^*}{(I^*)^2} I_{t+1}^* = \frac{1}{\phi_2^{O^*}}\left(\phi_1^{O^*}\frac{O_S^*}{I^*} + \mathbb{E}_t\left[M_{t,t+1}^*\frac{P_{O,t+1}^*}{P_t^*}(1-\delta)\right] - \frac{P_{O,t}^*}{P_t^*}\right).$$

In steady state, we have

$$I^* = \frac{I^{*2}}{O_S^*}\cdot\frac{\phi_1^{O^*}\frac{O_S^*}{I^*} + \beta(1-\delta) - 1}{\phi_2^{O^*}} \quad \Leftrightarrow \quad \phi_1^{O^*} = \phi_2^{O^*} - \frac{I^*}{O_S^*}\left[\beta(1-\delta) - 1\right].$$

F.2.2 Anticipated Shocks

We introduce anticipated (news) shocks in addition to unanticipated shocks. Let $e_t^{O^*}$ denote an unanticipated shock, and let $v_{1,t-1}^{O^*}$ denote a shock anticipated at time t−1 that materializes at time t (i.e., a one-period-ahead anticipated shock). We assume that both shocks are i.i.d. The stochastic process for oil production is given by

$$\widehat{o}_t^* = \rho_{O^*}\widehat{o}_{t-1}^* + v_{1,t-1}^{O^*} + e_t^{O^*}.$$

F.2.3 Equilibrium Conditions

Denote, as before, $\widehat{r}_t$ to be the deviation from steady state of the real oil price in units of foreign goods. The log-linearized oil storage firms optimality condition and the market clearing for oil are given by:

$$\phi_2^{O^*}\frac{O_S^*}{I^*}\widehat{I}_{t+1}^* - \beta(1-\delta)\left(\mathbb{E}_t\left[\widehat{r}_{t+1}\right] + \mathbb{E}_t\left[\pi_{F,t+1}^*\right]\right) = -\beta(1-\delta)\phi_\pi^*\pi_{F,t}^* - \widehat{r}_t$$

$$\frac{I^*}{O_S^*}\widehat{I}_{t+1}^* = (1-\delta)\frac{I^*}{O_S^*}\widehat{I}_t^* + \frac{O^*}{O_S^*}\widehat{o}_t^* + \left(\frac{1}{\phi_o}\widehat{r}_t - \frac{\phi_c}{\phi_o}\widehat{c}_t^*\right),$$

where we replaced the oil demand $\widehat{o}_{S,t}^*$ in the second equation using

$$\widehat{r}_t = \phi_c\widehat{c}_t^* - \phi_o\widehat{o}_{S,t}^*.$$

The remaining equilibrium conditions are the same as before. Specifically, the full set of equilibrium conditions is

- World economy

$$
\begin{aligned}
\widehat{o}^*_{t+1} &= \rho_{O^*}\widehat{o}^*_t + v^{O^*}_{1,t} + e^{O^*}_{t+1}, \\
\mathbb{E}_t\left[\widehat{c}^*_{t+1}\right] - \widehat{c}^*_t &= \frac{1}{\sigma}\left(\phi^*_\pi \pi^*_{F,t} - \mathbb{E}_t\left[\pi^*_{F,t+1}\right]\right), \\
\pi^*_{F,t} &= \beta \mathbb{E}_t\left[\pi^*_{F,t+1}\right] + \kappa\left(\zeta^*\widehat{c}^*_t + \psi^*_o \widehat{r}_t\right), \\
\beta(1-\delta)\phi^*_\pi \pi^*_{F,t} + \widehat{r}_t &= \beta(1-\delta)\left(\mathbb{E}_t\left[\widehat{r}_{t+1}\right] + \mathbb{E}_t\left[\pi^*_{F,t+1}\right]\right) - \phi^{O^*}_2 \frac{O^*_S}{I^*}\widehat{I}^*_{t+1}, \\
\frac{I^*}{O^*_S}\widehat{I}^*_{t+1} &= (1-\delta)\frac{I^*}{O^*_S}\widehat{I}^*_t + \frac{O^*}{O^*_S}\widehat{o}^*_t + \left(\frac{1}{\phi_o}\widehat{r}_t - \frac{\phi_c}{\phi_o}\widehat{c}^*_t\right), \\
v^{O^*}_{1,t+1} &= e^{O^*}_{1,t+1}
\end{aligned}
$$

- Home country

$$
\begin{aligned}
\mathbb{E}_t\left[\widehat{c}_{t+1}\right] - \widehat{c}_t &= \frac{1}{\sigma}\left(\phi_\pi \pi_{H,t} - \mathbb{E}_t\left[\pi_{H,t+1}\right]\right) + \frac{1-\Omega}{\Omega}\left(\mathbb{E}_t\left[\widehat{c}^*_{t+1}\right] - \mathbb{E}_t\left[\widehat{c}_{t+1}\right]\right) - \frac{1-\Omega}{\Omega}\left(\widehat{c}^*_t - \widehat{c}_t\right), \\
\pi_{H,t} &= \beta \mathbb{E}_t\left[\pi_{H,t+1}\right] + \kappa\left[\zeta_H \widehat{c}_t + \psi_o \widehat{r}_t + \zeta_F \widehat{c}^*_t - \psi_q \sigma\left(\widehat{c}^*_t - \widehat{c}_t\right)\right].
\end{aligned}
$$

F.2.4 Model Calibration

We set the depreciation of oil $\delta = 0.01$, following Unalmis *et al.* (2012). To target inventory-supply ratio, we set $I^*/O^*_S = 0.61$. Then, $O^* = O^*_S + \delta I^*$ pins down O^*. Since $\phi^{O^*}_1 = \phi^{O^*}_2 - (\beta(1-\delta) - 1)I^*/O^*_S$, we need to calibrate $\phi^{O^*}_1$ or $\phi^{O^*}_2$. We experiment with three values of $\phi^{O^*}_2 \in \{2, 1, 0.5\}$. A high value of $\phi^{O^*}_2$ means it is costly to adjust inventories, so the model would be closer to our model without inventories. We set $\phi^{O^*}_2 = 2$ as a baseline below, which leads to a "moderate" deviation from our baseline. Similar to the baseline, we set $\rho_{O^*} = 0.58$, so that the price of oil does not become too persistent.

F.2.5 Oil Shock Effects

Panel (a) of Figure F.3 plots the impulse responses to an unanticipated oil supply shock normalized to increase the oil price by 10 percent on impact in the model without inventories (baseline model). In both the model with and without inventories, an unanticipated disruption to oil supply leads to an immediate increase in oil prices. The responses of consumption and inflation, both in normal times and during the ZLB, are similar across the two models.

Panel (b) of Figure F.3 presents the responses to an anticipated oil supply shock normalized to increase the oil price by 10 percent in its peak in the model without inventories. In the model without oil inventories, an anticipated disruption to oil supply does not raise the oil price; instead, it lowers it on impact. The intuition is that, absent storage, the expectation of lower future oil supply reduces expected future consumption, which—through consumption smoothing—reduces current consumption. This lowers current oil demand and, consequently, the current oil price. This prediction is opposite to what we observe in the data.

We next examine the model with inventories and different values of $\phi^{O^*}_2 \in \{2, 1, 0.5\}$. Compared to the baseline responses, the smaller adjustment cost for inventories lead to larger on impact response of oil prices but smaller responses of home consumption and inflation. The results are shown in Figure F.4.

Next, Figure F.5 compares the responses of the two models (baseline and extended) when they are hit by two different sequences of shocks. In the model with inventories, the economy is hit by a news shock that

Figure F.3: Responses in the model with and without inventories

(a) Current oil supply shock

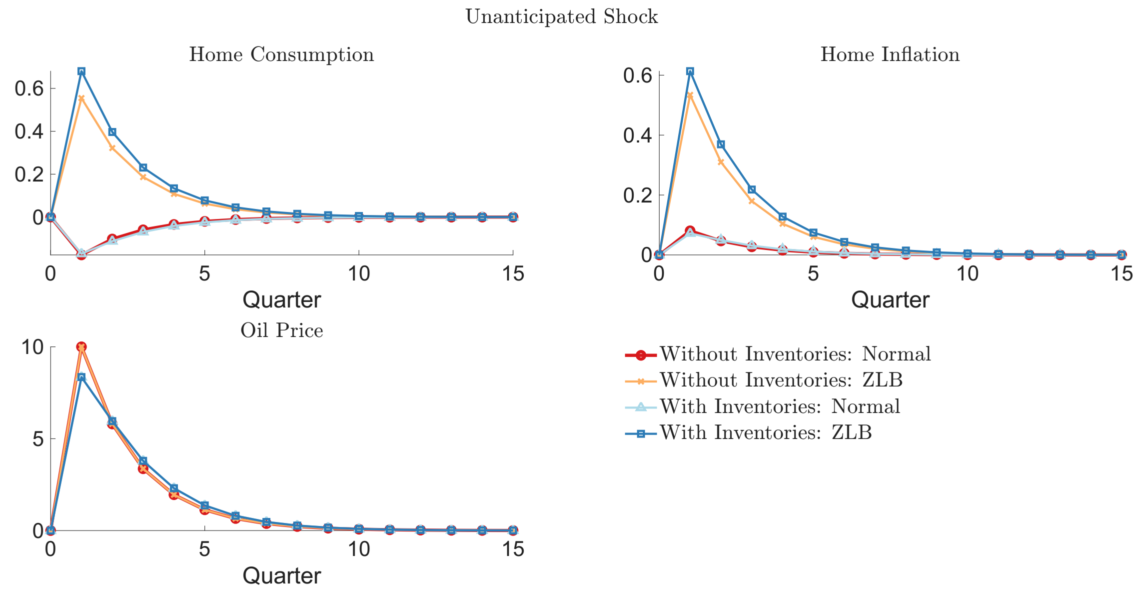

(b) Oil supply news shock

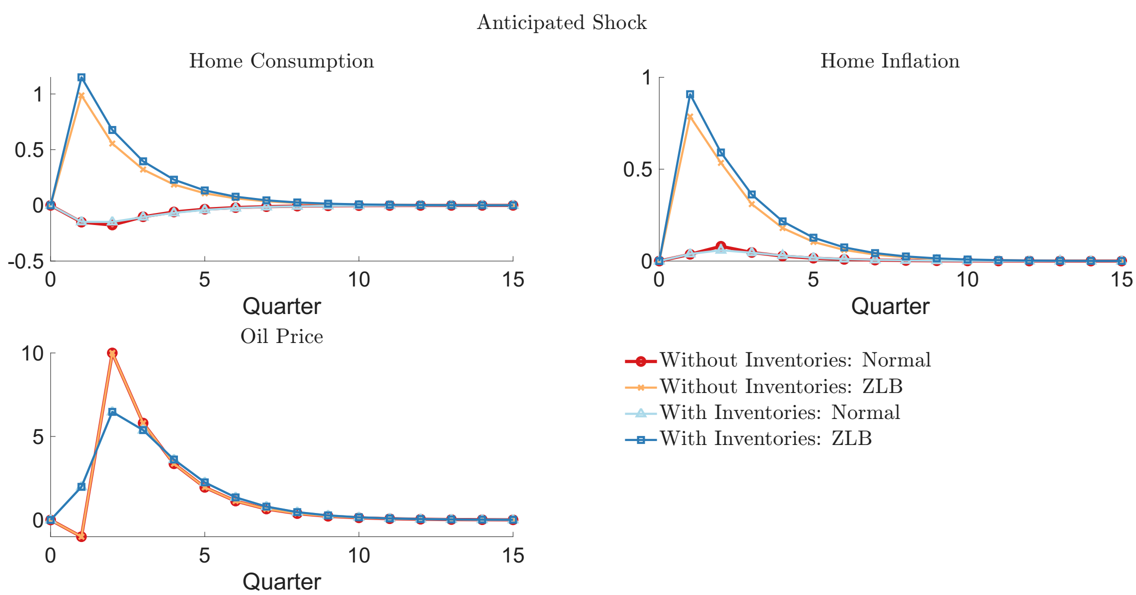

Notes: Each figure plots the impulse response functions to oil shock. The top panel plots the responses to an unanticipated oil supply shock that decays exponentially. The bottom panel plots the resposnes to an anticipated oil supply shock that is learned in period 1, materializes in period 2, and decays exponentially thereafter. The label ”Without Inventories” refers to the responses in our baseline model, while ”With Inventories” refers to the responses in the extended model with oil storage, described in the Appendix to the main text. The foreign (world) economy is assumed to conduct active monetary policy.

Figure F.4: Responses in the model with inventories and different adjustment costs

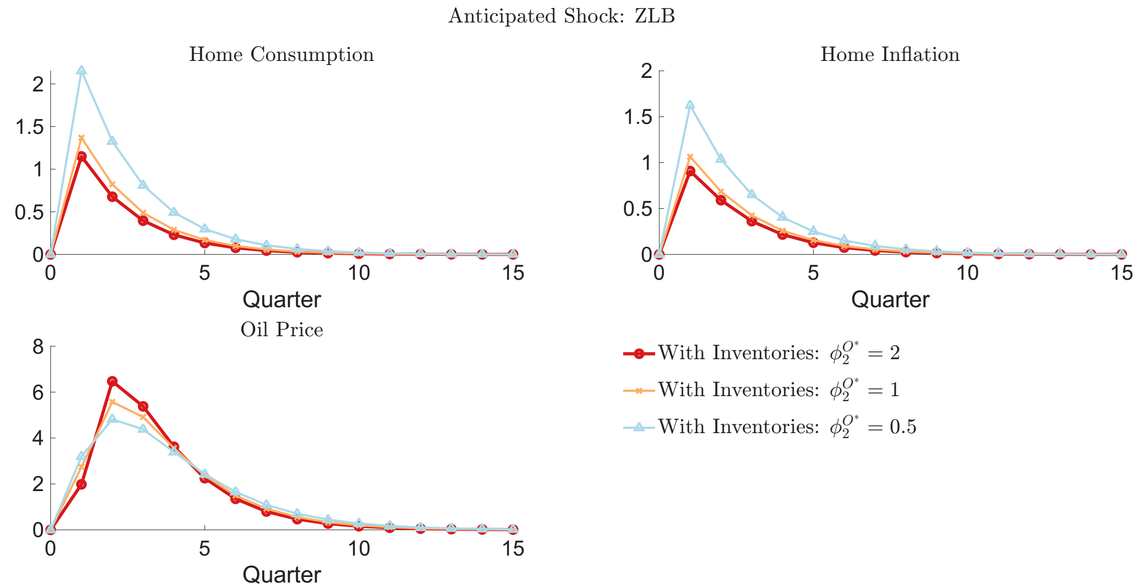

Notes: Each figure plots the impulse response functions in the model with inventories for different values of $\phi_2^{O^*}$.

is learned in period 1, materializes in period 2, and then decays exponentially. We calibrate this shock so that the oil price increases by 10 percent in its peak. In the model without inventories, we instead construct a sequence of oil supply shocks, which are revealed and start in period 1, that replicates exactly the oil price path generated in the model with inventories. The bottom right panel of Figure F.5 presents the evolution of two oil supply sequences.

Figure F.5 shows that the responses of home-country variables are identical across the two models. In particular, home consumption and inflation coincide. The reason is that, for the small open economy, what primarily matters is the path of the oil price, not the underlying source of the shock.

Interestingly, the responses of the world economy are also identical across the two models (see the bottom left panel of Figure F.5). This implies that the spillovers from the world economy to the small open economy are the same in both cases. The reason is that introducing oil storage firms in our environment adds only one additional equilibrium condition linking oil prices and inventories; all other equilibrium equations remain unchanged. Consequently, different oil supply shocks (bottom middle panel of Figure F.5) that generate the same oil price path lead to identical responses of oil consumption and other macroeconomic variables in the world economy.

These results highlight that oil supply news shocks in a model with inventories can generate responses that closely resemble those produced by contemporaneous oil supply shocks in the model without inventories. This observation motivates our use of the baseline model with contemporaneous oil supply shocks in the main text.

Figure F.5: Responses of models with and without oil inventories with matched oil price dynamics

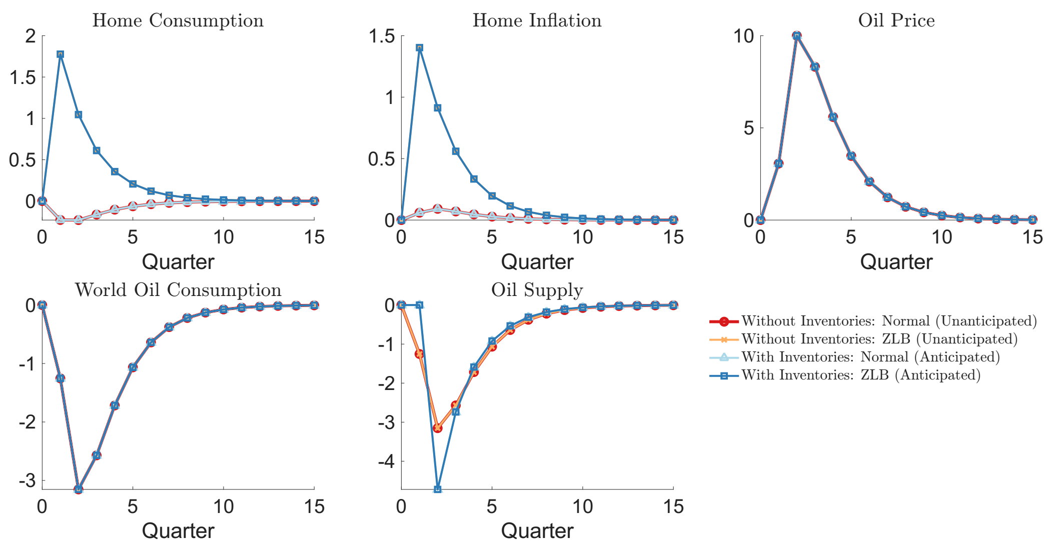

Notes: Each figure plots the impulse response functions. The label "Without Inventories" refers to the responses in our baseline model, while "With Inventories" refers to the responses in the extended model with oil storage, described in the Appendix to the main text. In the model with inventories, the shock consists of a sequence of oil supply changes that begins in period 2 and decays exponentially thereafter; the entire sequence is learned in period 1. In the model without inventories, the sequence of shocks is announced and begins in period 1 and is chosen to replicate the same oil price dynamics. The foreign (world) economy is assumed to conduct active monetary policy.