Online Appendix of

"Estimating a Nonlinear New Keynesian Model with the Zero Lower Bound for Japan"

Hirokuni Iiboshi Mototsugu Shintani Kozo Ueda

October 23, 2020

A. Methodology

A.1 Model Solution

We derive the rational expectations equilibrium of our model using the TL method. The model's equilibrium conditions are written as a vector-valued function, $f(\cdot)$, containing (ε_t, s_t):

$$\mathbb{E}[f(\varepsilon_{t+1}, s_{t+1}, \varepsilon_t, s_t)|\Omega_t] = 0,$$

where ε_t is a vector of exogenous variables, s_t is a vector of endogenous variables, and Ω_t is an information set of agents. In our study, we set $\varepsilon_t = \left(\epsilon_t^a,\, \epsilon_t^b,\, \epsilon_t^r\right)'$, $s_t = (y_t, y_{t-1}, c_t,$ $\pi_t, y_t^*, y_{t-1}^*,\ R_t, R_t^*, r_t^*, \mu_t^a, z_t^b\big)'$ where $y_t \equiv Y_t/A_t$, $c_t \equiv C_t/A_t$, and $y_t^* \equiv Y_t^*/A_t$.

Given function $f(\cdot)$, we can obtain a model's decision rules (or policy functions), $\Phi(\cdot)$, as a function of the state vector. The TL method locally approximates the time-invariant policy function at each node in the state space, $z_t \equiv \left(\mu_t^a,\, z_t^b,\, \epsilon_t^r, y_{t-1}, R_{t-1}^*\right)'$, that is,

$$\Phi(z_t) \simeq \hat{\Phi}(z_t).$$

We solve the rational expectations equilibrium by substituting $(y_t,\, y_t^*,\, \pi_t)' = \Phi(z_t)$ into the future variables of the function $f(\cdot)$. We discretize five grid points on each of the continuous state variables, which implies $3,125 (= 5^5)$ nodes in total.

The policy function iteration algorithm takes the following steps. Let $i \in \{0, \cdots, I\}$ denote the iterations of the algorithm and $n \in \{1, \cdots, N\}$ denote the nodes of the policy function, $\Phi(z_t)$.

1. For $i = 0$, we make an initial conjecture of the policy function, $\Phi^0(z_t)$, from the log-linearized model without the ZLB. To do so, we use Sims' (2002) `gensys` algorithm.

2. For iteration $i \in \{1, \cdots, I\}$ and node $n \in \{1, \cdots, N\}$, we execute the following procedures.

   (a) Solve for endogenous variables $\left(c_t, R_t, R_t^*, r^*, \mu_t^a, z_t^b\right)'$, given $\left(y_{t-1},\, y_{t-1}^*,\, \pi_{t-1}\right)' = \Phi(z_{t-1})$ under the ZLB.

   (b) Approximate the future variables $\{\, \mathbb{E}_t(y_{t+1}), \mathbb{E}_t(y_{t+1}^*), \mathbb{E}_t(\pi_{t+1})\}$ using a piecewise linear interpolation of the policy function $\Phi^{i-1}(z_t)$. Then, substitute the future variables into $\mathbb{E}_t[f(\cdot)|\Omega_t]$, where we employ Gauss-Hermite integration to approximate the conditional expectations with three nodes per shock based on Gust et al. (2017).

   (c) Use the nonlinear solver, Sims' `csolve`, to find the policy function, $\Phi^i(z_t)$, which minimizes the errors in intertemporal equations, that is, $E[f(\cdot)|\Omega_t] = 0$.

3. Define $\texttt{maxdist} = \texttt{max}(\, |y_n^i - y_n^{i-1}|,\, |y_n^{*,i} - y_n^{*,i-1}|,\, |\pi_n^i - \pi_n^{i-1}|)$. Repeat Step 2 until the policy function converges, say, to $\texttt{maxdist} < 10^{-4}$, for all nodes, n.

A.2 Estimation

To obtain draws from the posterior distribution of the parameters, θ, in a nonlinear DSGE model, we use the SMC^2 sampler combined with the *particle filter* instead of popular methods such as the PFMH algorithm. Because the PFMH algorithm cannot be parallelized when generating draws, they take a long time. By contrast, the SMC^2 method and particle filter can be used easily and may also more accurately approximate the posterior distribution. We employ a widely used particle filter, the so-called bootstrap particle filter, following Gust et al. (2017).

We explain the algorithms of the SMC^2 method and bootstrap particle filter following Herbst and Schorfheide (2015) and Fernández-Villaverde, Rubio-Ramirez, and Schorfheide (2016).

A.2.1 Algorithm of the Sequential Monte Carlo Squared

Suppose that ϕ_n for $n = 0, \cdots, N_\phi$ is a sequence that slowly increases from zero to one. We define a sequence of bridge distributions, $\{\pi_n(\theta)\}_{n=0}^{N_\phi}$, that converge to the target posterior distribution for $n = N_\phi$ and $\phi_n = 1$ as

$$\pi_n(\theta) = \frac{[p(\mathbf{Y}|\theta)]^{\phi_n} p(\theta)}{\int [p(\mathbf{Y}|\theta)]^{\phi_n} p(\theta) d\theta}, \quad \text{for } n = 0, \cdots, N_\phi,\ \phi_n \uparrow 1,$$

where $p(\theta)$ and $p(\mathbf{Y}|\theta)$ are the prior density and likelihood function, respectively. We adopt the likelihood tempering approach that generates the bridge distributions, $\{\pi_n(\theta)\}_{n=0}^{N_\phi}$, by taking the power transformation of $p(\mathbf{Y}|\theta)$ with parameter ϕ_n (i.e., $[p(\mathbf{Y}|\theta)]^{\phi_n}$).

The SMC^2 method with likelihood tempering has the following steps. Let $i \in \{1, \cdots, N_\theta\}$ denote the particles of the parameter sets, θ^i, and $n \in \{0, \cdots, N_\phi\}$ denote the stage of the algorithm. Herbst and Schorfheide (2015) recommend a convex tempering schedule in the form of $\phi_n = (n/N_\phi)^\lambda$ with $\lambda = 2$ for a small-scale DSGE model.

1. Initialization

   (a) Set the initial stage as $n = 0$ and draw the initial particles of parameters θ_0^i from the prior distribution $p(\theta)$.

   (b) Set the weight of each particle in the initial stage as $W_0^i = 1$ for $i = 1, \cdots, N_\theta$.

   Then, for stage $n \in \{1, \cdots, N_\phi\}$ and particle $i \in \{1, \cdots, N_\theta\}$, we repeat Steps 2 to 4.

2. Correction. Calculate the normalized weight, $\tilde{W}_n^i$, for each particle as

$$\tilde{W}_n^i = \frac{\tilde{w}_n^i W_{n-1}^i}{\frac{1}{N}\sum_{i=1}^N \tilde{w}_n^i W_{n-1}^i} \text{ for } i = 1, \cdots, N_\theta,$$

   where $\tilde{w}_t^i$ is an incremental weight derived from

$$\tilde{w}_n^i = [p(\mathbf{Y}|\theta_{n-1}^i)]^{\phi_n - \phi_{n-1}},$$

   and the likelihood, $\hat{p}(\mathbf{Y}|\theta)$, is approximated from the particle filter, as explained in the next subsection.

   The correction step is a classic importance sampling step in which particle weights are updated to reflect the stage n distribution, $\pi_n(\theta)$. Because this step does not change the particle value, we can skip this step only by calculating the power transformation of $p(\mathbf{Y}|\theta)$ with parameter ϕ_n.

3. Selection (Resampling).

   (a) Calculate an effective particle sample size, $\widehat{ESS_n}$, which is defined as

$$\widehat{ESS_n} = N_\theta / \left(\frac{1}{N_\theta} \sum_{i=1}^{N_\theta} (\tilde{W}_n^i)^2 \right).$$

   (b) If $\widehat{ESS_n} < N_\theta/2$, then resample particles $\{\hat{\theta}_n\}_{i=1}^{N_\theta}$ by multinomial resampling and set $W_n^i = 1$.

   (c) Otherwise, let $\hat{\theta}_n^i = \theta_{n-1}^i$ and $W_n^i = \tilde{W}_n^i$.

4. Mutation. Propagate the particles $\{\hat{\theta_n^i}, W_n^i\}$ via the random walk MH algorithm with the proposal density,

$$\vartheta | \hat{\theta_n}^i \sim N\left(\hat{\theta_n}^i, c_n^2 \Sigma\left(\hat{\theta_n}\right)\right),$$

where $N(\cdot)$ is the nominal distribution and $\Sigma\left(\hat{\theta_n}\right)$ denotes the covariance matrix of parameter $\hat{\theta_n}$ for all the particles $i \in \{1, \cdots, N_\theta\}$ in the n-th stage. To keep the acceptance rate around 25%, we set a scaling factor c_n for $n > 2$ as

$$c_n = c_{n-1} f(A_{n-1}),$$

where A_n represents the acceptance rate in the mutation step in the n-th stage and function $f(x)$ is given by

$$f(x) = 0.95 + 0.10 \frac{e^{16(x-0.25)}}{1 + e^{16(x-0.25)}}.$$

5. For the final stage of $n = N_\phi$, calculate the final importance sampling approximation of the posterior estimator, $E_\pi[h(\theta)]$, as

$$h_{N_\phi, N_\theta} = \sum_{i=1}^{N_\theta} h(\theta^i_{N_\phi}) W^i_{N_\phi}.$$

In the final stage, the approximated marginal likelihood of the model is also obtained as a by-product. It can be shown that

$$P_{SMC}(Y) = \prod_{n=1}^{N_\phi} \left(\frac{1}{N_\theta} \sum_{i=1}^{N_\theta} \tilde{w}^i_n W^i_{n-1} \right)$$

converges almost surely to $p(Y)$ as the number of particles $N_\theta \to \infty$.

A.2.2 Algorithm of the Particle Filter

Suppose that a state-space representation for the nonlinear DSGE model consists of

$$\mathbf{Y}_t = \Psi(s_t,\ \theta) + u_t, \quad u_t \sim N(0,\ \Sigma_u),$$

$$s_t = \Phi(s_{t-1},\ \varepsilon_t,\ \theta), \quad \varepsilon_t \sim N(0,\ \Sigma_\varepsilon),$$

where $\mathbf{Y}_t$ and s_t denote the observable and state variables, respectively. In our study, we set $\mathbf{Y}_t = (\log(y_t/y_{t-1}),\ \pi_t,\ R_t)'$ and $s_t = \left(y_t,\ y_{t-1},\ c_t,\ \pi_t,\ y^*_t,\ y^*_{-1},\ R_t,\ R^*_t,\ r^*_t,\ \mu^a_t,\ z^b_t\right)'$. A measurement error vector, u_t, and an exogenous shock vector, $\varepsilon_t = \left(\epsilon^a_t,\ \epsilon^b_t,\ \epsilon^r_t\right)'$, follow a normal distribution with covariance matrixes Σ_u and Σ_ε, respectively. The nonlinear policy function $\Phi(s_{t-1},\ \varepsilon_t,\ \theta)$ is derived in Appendix A.1, while the function $\Psi(s_t, \theta)$ represents the linkage between $\mathbf{Y}_t$ and s_t.

The particle filter algorithm is as follows. Let $j \in \{0, \cdots, N_S\}$ denote the index for the particles of the state variables and exogenous shocks.

1. For period $t = 0$, draw the N_S initial particles of the state variables in period 0, say $s^j_{0|0}$, from $s^j_{0|0} = \Phi(s,\ \varepsilon_0,\ \theta)$ with $\varepsilon_0 \sim N(0,\ \Sigma_\varepsilon)$, where $s = \Phi(s,\ 0,\ \theta)$.

2. For period $t \in \{1, \cdots, T\}$ and particle $j \in \{1, \cdots, N_S\}$, take the following three steps.

   (a) Forecasting the state variables: $s^j_{t|t-1}$. Generate N_S particles of the shock vector ε^j_t from $N(0,\ \Sigma_\varepsilon)$. Using the nonlinear policy function, we obtain N_S particles of the forecasts of the state variables corresponding to the shocks generated above:

   $$s^j_{t|t-1} = \Phi(s^j_{t-1|t-1}, \varepsilon^j_t, \theta).$$

   (b) Forecasting the observable variables. Calculate the approximated predictive density of y^{obs}_t given by

   $$p(\mathbf{Y}_t|\mathbf{Y}_{1:t-1},\ \theta) \simeq \frac{1}{N_S}\sum_{j=1}^{N_S} w^j_t,$$

   where w^j_t is the normal predictive density of particle j measured from $\Psi(s^j_{t|t-1}, \theta)$ and the covariance matrix of the measurement error Σ_u in period t, say,

   $$w^j_t = (2\pi)^{-N_y/2}|\Sigma_u|^{-1/2}\exp\left\{-\frac{1}{2}\left(\mathbf{Y}_t - \Psi(s^j_{t|t-1}, \theta)\right)'\Sigma_u^{-1}\left(\mathbf{Y_t} - \Psi(s^j_{t|t-1}, \theta)\right)\right\},$$

   where N_y is the dimension of y_t.

   (c) Updating the state variables: $s^j_{t|t}$. Resample N_S particles of the state variables from a multinomial distribution. That is,

   $$s^j_{t|t} = \text{resample out of } (s^1_{t|t-1}, \cdots s^j_{t|t-1}, \cdots s^{N_S}_{t|t-1}) \text{ with probabilty } (w^j_t/\Sigma w^j_t).$$

3. For the final period $t = T$, collect all the predictive densities of y_t from periods 1 to T calculated above. Using these densities, the log likelihood of the model is approximated as

$$\log p(\mathbf{Y}_{1:t}|\theta) \simeq \sum_{t=1}^{T}\log\left(\frac{1}{N_S}\sum_{j=1}^{N_S} w^j_t\right).$$

B. Comparisons of Notional Interest Rates

Figure 1 shows developments in the notional interest rate $R^*_t - 1$ in the Exit Condition Model, the Notional Rate Model, and the Nominal Rate Model. Of these models, only the Notional Rate Model gives the notional interest rate a role in lowering expectations for future interest rates, because it influences the interest rate in the next period through inertia. The figure shows that the notional interest rate in the Notional Rate Model is the lowest.

C. Duration of the Zero-Rate Policy Based on the Shadow-Rate Model

Using the estimation results of Ueno (2017), we calculate the expected duration being at the ZLB. Specifically, we use the fixed extended and variable extended models in Ueno (2017). Considering that these models assume a nonstationary I(1) process for interest rates, we generate a relatively longer sequence of interest rates, as long as 10^7, to calculate the expected duration. Figure 2 shows the expected duration of being at the ZLB.

D. Impulse Response Functions

Figures 3, 4, and 5 show the IRFs given the monetary policy shock, the discount factor (preference) shock, and the technology shock, respectively. The sizes of the shocks are 0.025%, 0.15%, and 0.05%, respectively. For each type of shock, we calculate the IRFs to both positive and negative shocks. Furthermore, we show the IRFs in different models: the Exit Condition Model, the Notional Rate Model, the Nominal Rate Model, and the Model without the ZLB. The IRFs are conditional on states in two historical periods, 1985:1Q and 2013:2Q. The latter is a period in which the ZLB constrains the economy and π_t is slightly below the estimated exit condition $\overline{\pi}$.

Monetary policy shock

Figure 3 shows that, in 1985:1Q, the IRFs are symmetric to the sign of the monetary policy shock irrespective of the models. In 2013:2Q, the IRFs are asymmetric except for the Model without the ZLB. The positive monetary shock yields a bigger decrease in π_t than the negative monetary policy shock increases π_t. Moreover, the positive monetary policy shock has a larger negative effect on π_t in the Nominal Rate Model than in the Notional Rate Model. In the former model, the experience of prolonged recessions does not tie the hand of the central bank and, hence, the positive monetary policy shock leads to an immediate increase in R_t, which decreases π_t. In response to the negative monetary policy shock, π_t increases in the Notional Rate Model, because it incorporates a promise to continue the zero-rate policy in the future. The monetary policy in period t depends on R^*_{t-1}, which can take negative values. The negative monetary policy shock in period t decreases R^*_t, serving to lower future nominal interest rates. This promise increases π_t in the current period.

Discount factor (preference) shock

Figure 4 shows that, in 1985:1Q, the IRFs are symmetric to the sign of the shock, and their patterns are almost the same in the four models. A positive (negative) shock to Z^b_t increases

(decreases) the weight on the current period's utility, thereby increasing (decreasing) y_t, π_t, R_t, and R_t^* in the current period.

It is important to emphasize that the IRFs in 2013:2Q are not much different from those in 1985:1Q. The ZLB constrains the economy and R_t hardly decreases in response to the negative shock in the three models that take account of the ZLB. However, this inaction of R_t hardly influences π_t. The only difference is that π_t converges to zero a little more stagnantly in response to the negative shock than it does in response to the positive shock.

This result is in a sharp contrast to those of existing studies that point out significant changes in policy implications due to the ZLB, for example, an increase in government purchase multiplier or a fall in employment by the cut in the labor tax (see Eggertsson 2011). However, our result is in line with that of Boneva, Braun, and Waki (2016), who find that implications for fiscal policy at the ZLB are not very different when they employ a nonlinear solution method, instead of a loglinearized solution method with the ZLB.

Technology shock

Figure 5 shows the IRFs to the technology shock. In 1985:1Q, the IRFs are symmetric to the sign of the shock. A positive shock to μ_t^a increases π_t in period 1 slightly, but then decreases π_t below zero from period 2, mainly because of the consumption habit. To understand this change, let us look at the log-linearized equations described by equations (11) to (14) in the main text. On the one hand, equation (13) shows that the positive shock increases r_t^*, which contributes to increasing $y_t - y_t^*$ and π_t through equations (11) and (12). On the other hand, equation (14) suggests that the positive shock decreases the natural level of output defined by $y_t^* = Y_t^*/A_t$ owing to the consumption habit ($h > 0$), which functions to decrease y_t with a lag. Thus, $y_t - y_t^*$ is expected to fall below zero, which decreases π_t. In period 1, the former effect dominates, yielding a slight increase in $y_t - y_t^*$ and π_t. From period 2, the latter effect dominates and both $y_t - y_t^*$ and π_t decrease below zero. Accordingly, R_t and R_t^* decrease. We confirm that the positive shock increases $y_t - y_t^*$ and π_t rather than decreases them, when there is no habit formation ($h = 0$).

In 2013:2Q, when the ZLB constrains the economy, the IRFs of R_t are asymmetric, but those of π_t are almost symmetric and hardly different from those in 1985:1Q. Although the ZLB prevents R_t from decreasing in response to the positive technology shock, this does not decrease π_t much, except for a slightly slower convergence. This is consistent with the result of the IRFs to the discount factor shock. The ZLB itself also does not change the implications for the supply-side shock.

E. Comparisons of Moments

In this section, we examine the validity of the model by comparing the moments of key economic variables with data. In particular, we look at hours worked (l_t), real wage (W_t), and the real labor cost ($W_t l_t$), which are not used for the estimation. The data for l_t and W_t are extended from those used in Sugo and Ueda (2008): they are from the Ministry of Health, Labour and Welfare "Monthly Labor Survey." The effect of *jitan*, a decrease in the number of statutory workdays per week, is adjusted. Considering the non-stationarity, we detrend Y_t, W_t, and $W_t l_t$ using the HP filter with $\lambda = 1,600$. As for the moments, we calculate standard deviations and correlation coefficients.

Table 1 shows the results. First, regarding standard deviations, we express them by a ratio to those of Y_t. The estimated standard deviations of l_t based on the Exit Condition Model are higher than those of data, but of the same order. However, the Exit Condition Model yields standard deviations of W_t and $W_t l_t$ larger by one order of magnitude than the actual data. Because our model does not embed frictions in the labor market (e.g., wage stickiness), the real wage is adjusted in an excessively volatile manner.

Next, we investigate correlation coefficients between the variables associated with output (Y_t or $\Delta \log Y_t$) and the variables associated with labor. The correlation coefficient concerning l_t is almost the same between the data and the Exit Condition Model. However, again, the Exit Condition Model performs very poorly in explaining the correlation coefficient concerning W_t. For example, the model suggests $+0.7$ for the correlation coefficient between Y_t and W_t, while it is almost zero for the data. Nevertheless, the correlation coefficient concerning their product $W_t l_t$ is almost the same between the data and the Exit Condition Model; the correlation coefficient between Y_t and $W_t l_t$ is around $+0.6$.

The other models, including that without the ZLB, perform no better than the Exit Condition Model does. In particular, the Linear Model, which neglects all nonlinearity, yields even larger standard deviations for the variables associated with labor. However, there is no significant difference among the models that incorporate the ZLB.

F. Further Estimation Results

Table 2 shows further estimation results. We estimate the model without the ZLB nonlinearly but ignoring the ZLB. This is estimated using the same approach as in our benchmark model, and thus the computational costs are equally high.

Next, we check the robustness of our estimation using an alternative measure of output, that is, the output gap. We estimate the same model using either the output gap instead of growth of real GDP or using both the output gap and growth of real GDP. Table 2 shows that the parameter estimates are almost the same.

Table 1: Comparisons of Moments

(1) Standard deviations relative to $\sigma(Y)$

	$\sigma(l)$	$\sigma(W)$	$\sigma(Wl)$
Data	0.45	0.36	0.48
Exit condition model	0.76	5.27	6.02
Notional rate model	0.78	5.18	5.94
Nominal rate model	0.89	5.35	6.22
Model w/o ZLB	0.95	4.60	5.53
Linear model	1.30	6.37	7.65

(2) Correlation coefficients

	Y, l	Y, W	Y, Wl	dY, l	dY, W	dY, Wl
Data	0.69	-0.05	0.60	0.20	0.07	0.24
Exit condition model	0.51	0.69	0.67	0.33	0.47	0.45
Notional rate model	0.43	0.64	0.62	0.23	0.41	0.39
Nominal rate model	0.49	0.66	0.64	0.34	0.46	0.45
Model w/o ZLB	0.56	0.72	0.69	0.43	0.50	0.49
Linear model	0.45	0.52	0.51	-0.15	0.02	-0.01

Note: Output (Y), real wage (W), and real labor cost (Wl) are detrended by the HP filter with $\lambda = 1,600$, and dY represents the growth rate of output ($\Delta \log Y$).

Table 2: Further Estimation Results

	Benchmark		Model w/o ZLB		Output GAP		Output Growth and Gap	
	Mean	(95% low, high)	Mean	(95% low, high)	Mean	(95% low, high)	Mean	(95% low, high)
Parameter								
σ	1.548	(1.507, 1.621)	1.108	(1.069, 1.138)	1.709	(1.148, 1.997)	1.616	(1.301, 1.765)
h	0.641	(0.602, 0.665)	0.317	(0.311, 0.352)	0.507	(0.445, 0.629)	0.597	(0.575, 0.655)
γ_a	-0.046	(-0.055, -0.03)	0.020	(0.019, 0.024)	0.044	(-0.053, 0.076)	-0.012	(-0.042, 0.009)
ω	3.922	(3.833, 4.086)	3.748	(3.539, 3.768)	4.051	(3.569, 4.293)	3.962	(3.697, 4.091)
κ	0.051	(0.05, 0.053)	0.049	(0.046, 0.049)	0.044	(0.04, 0.053)	0.048	(0.046, 0.051)
r^*	0.054	(0.039, 0.08)	0.147	(0.146, 0.151)	0.205	(0.061, 0.269)	0.105	(0.053, 0.14)
π^*	0.325	(0.292, 0.348)	0.251	(0.236, 0.253)	0.283	(0.216, 0.346)	0.186	(0.161, 0.198)
$\bar{\pi}$	0.339	(0.3, 0.366)	–	–	0.297	(0.22, 0.376)	0.216	(0.188, 0.23)
ρ_r	0.394	(0.268, 0.483)	0.629	(0.614, 0.732)	0.477	(0.219, 0.592)	0.540	(0.391, 0.605)
ψ_π	2.070	(1.872, 2.251)	1.880	(1.866, 1.984)	2.131	(1.665, 2.396)	1.907	(1.748, 2.168)
ψ_y	0.137	(0.132, 0.142)	0.078	(0.076, 0.087)	0.127	(0.111, 0.172)	0.098	(0.087, 0.13)
ρ_a	0.437	(0.403, 0.474)	0.204	(0.164, 0.212)	0.737	(0.516, 0.839)	0.334	(0.277, 0.443)
ρ_b	0.163	(0.094, 0.278)	0.355	(0.209, 0.377)	0.484	(0.377, 0.65)	0.722	(0.673, 0.827)
σ_a	1.567	(1.49, 1.674)	1.484	(1.343, 1.584)	1.664	(0.995, 3.159)	1.833	(1.592, 2.131)
σ_b	4.407	(4.047, 4.685)	2.781	(2.745, 2.968)	2.781	(2.234, 3.176)	2.781	(3.026, 4.627)
σ_r	1.849	(1.773, 1.966)	1.547	(1.248, 1.583)	1.438	(0.891, 2.162)	1.151	(0.993, 1.436)
Likelihood								
	-204.722		-467.963		-232.938		-412.533	

Figure 1: Notional Nominal Interest Rate R_t^*

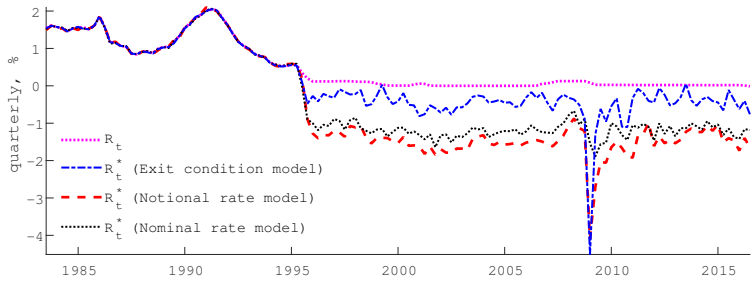

Figure 2: Duration of Being at the ZLB

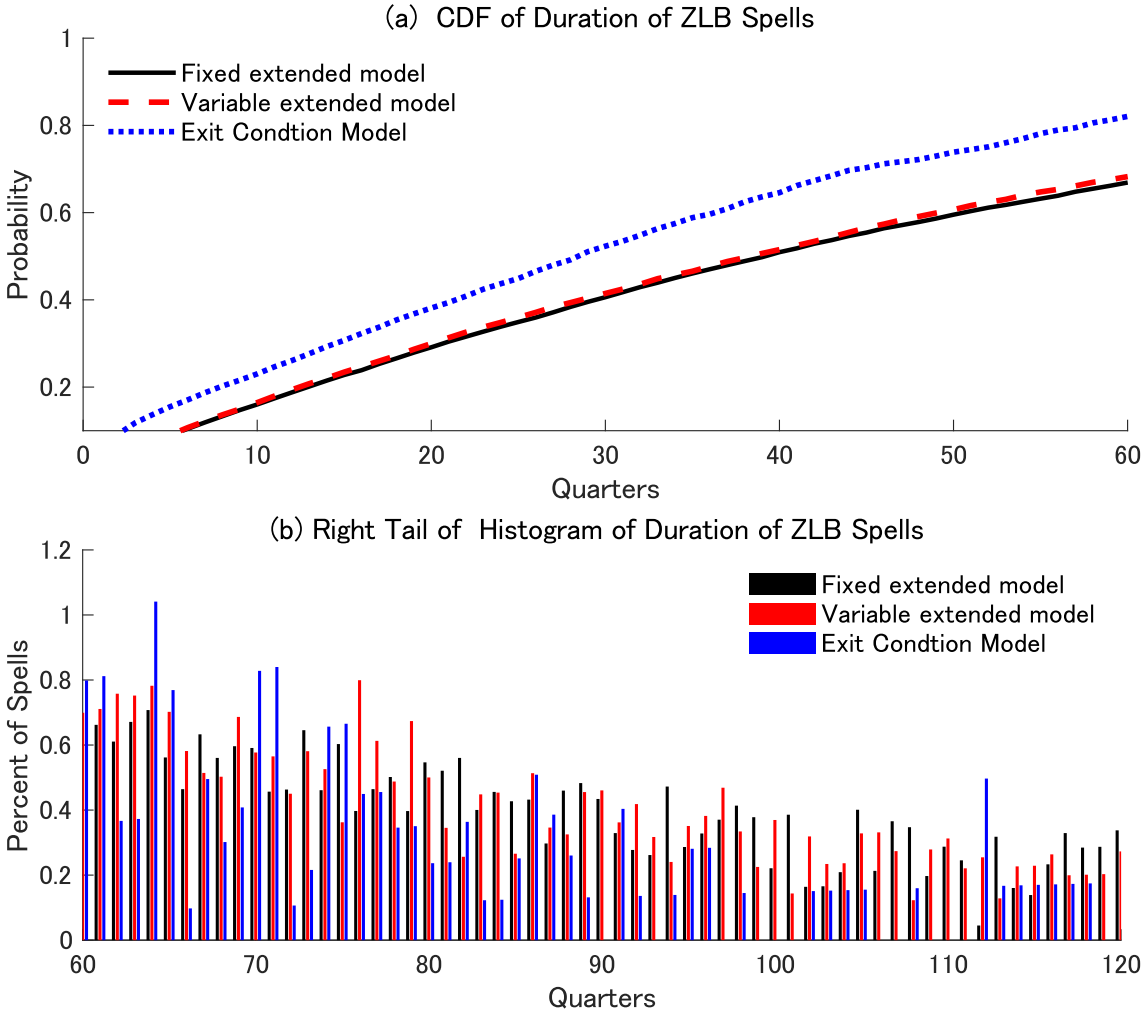

Figure 3: Impulse Responses to a Monetary Policy Shock

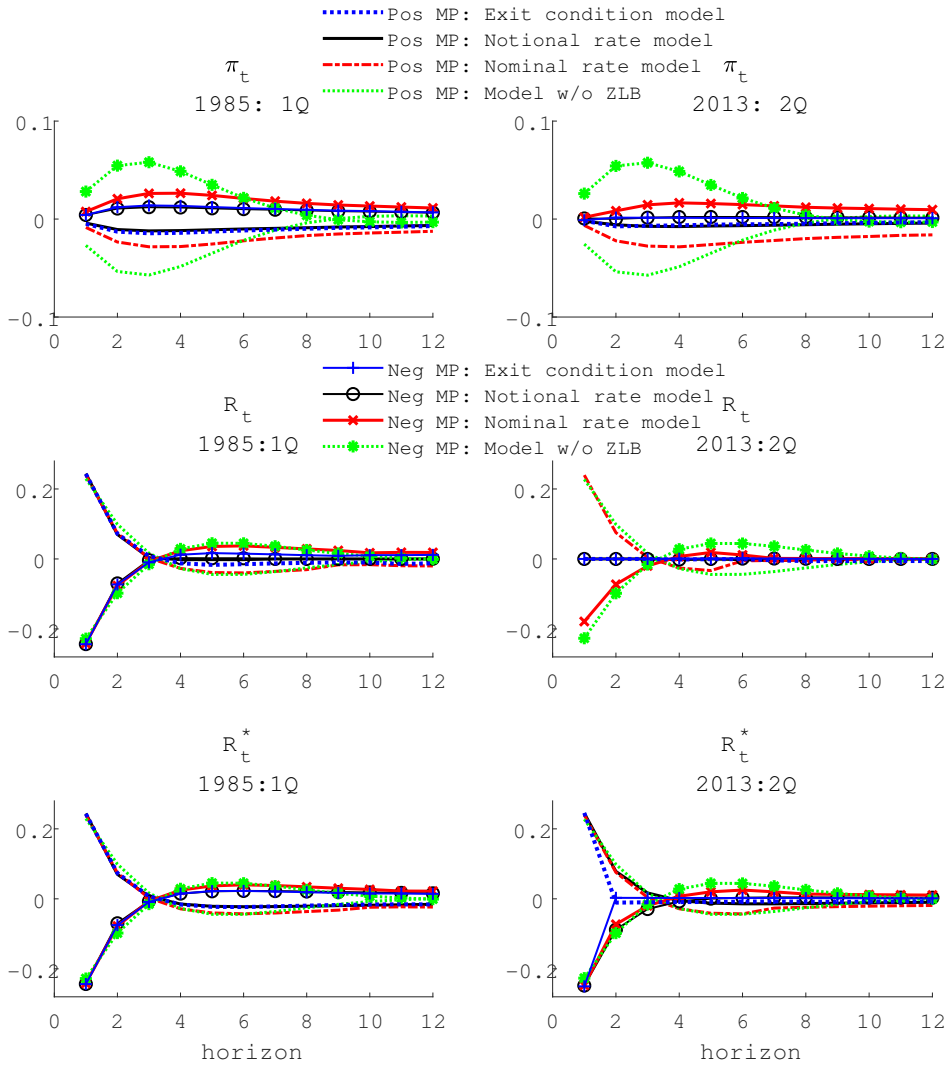

Note: "Pos" and "Neg" represent positive (tightening) and negative (easing) monetary policy shocks, respectively.

Figure 4: Impulse Responses to a Discount Factor Shock

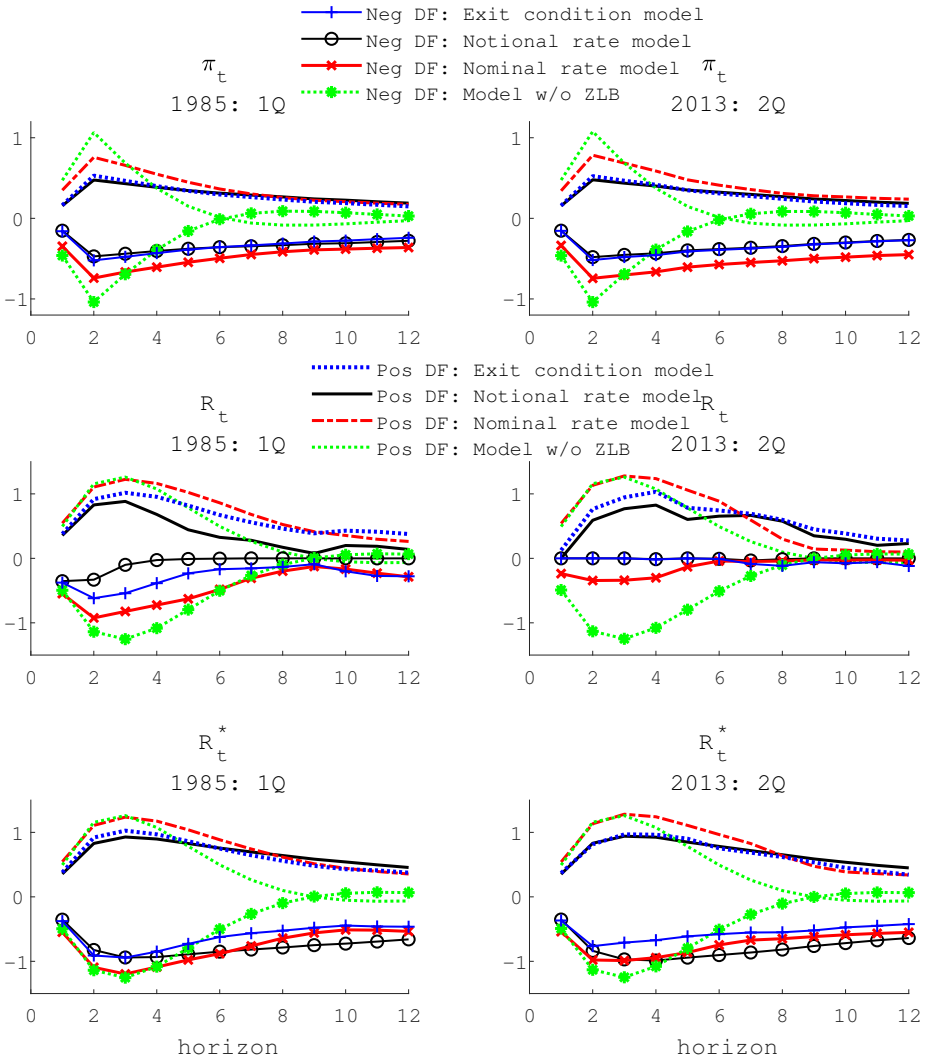

Note: "Pos" and "Neg" represent positive and negative discount factor shocks, respectively.

Figure 5: Impulse Responses to a Technology Shock

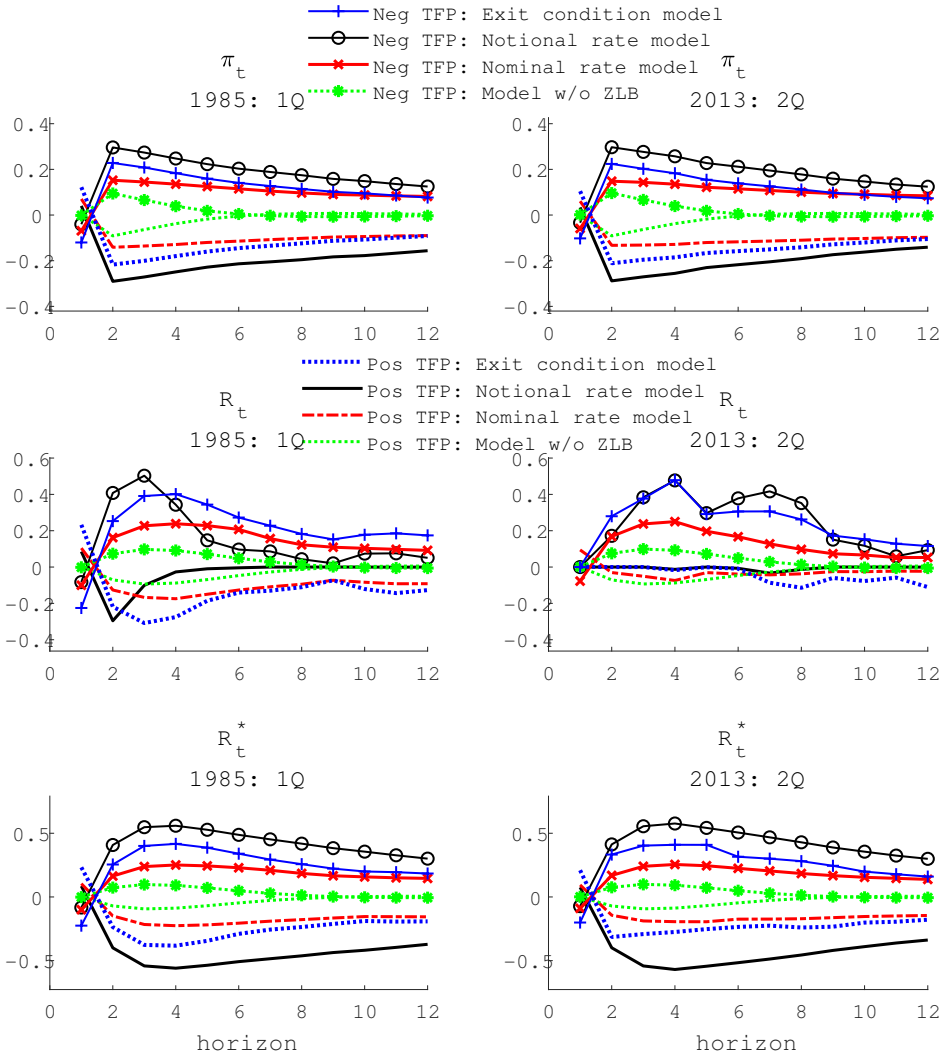

Note: "Pos" and "Neg" represent positive and negative technology shocks, respectively.