Review of Economic Dynamics 41 (2021) 96–120

Contents lists available at ScienceDirect

Review of Economic Dynamics

www.elsevier.com/locate/red

Piecewise-linear approximations and filtering for DSGE models with occasionally-binding constraints ☆

S. Borağan Aruoba [a], Pablo Cuba-Borda [b], Kenji Higa-Flores [a], Frank Schorfheide [c,d,e,*], Sergio Villalvazo [c]

[a] Department of Economics, University of Maryland, College Park, MD 20742, United States of America
[b] Division of International Finance, Federal Reserve Board, 20th Street & Constitution Ave., NW, Washington, DC 20551, United States of America
[c] Department of Economics, University of Pennsylvania, 133 S. 36th Street, Philadelphia, PA 19104, United States of America
[d] CEPR, United Kingdom
[e] NBER, PIER, United States of America

ARTICLE INFO

Article history:
Received 15 January 2020
Received in revised form 9 October 2020
Available online 21 January 2021

JEL classification:
C5
E4
E5

Keywords:
Bayesian estimation
Effective lower bound on nominal interest rates
Nonlinear filtering
Nonlinear solution methods
Particle MCMC

ABSTRACT

We develop an algorithm to construct approximate decision rules that are piecewise-linear and continuous for DSGE models with an occasionally-binding constraint. The functional form of the decision rules allows us to derive a conditionally optimal particle filter (COPF) for the evaluation of the likelihood function that exploits the structure of the solution. We document the accuracy of the likelihood approximation and embed it into a particle Markov chain Monte Carlo algorithm to conduct Bayesian estimation. Compared with a standard bootstrap particle filter, the COPF significantly reduces the persistence of the Markov chain, improves the accuracy of Monte Carlo approximations of posterior moments, and drastically speeds up computations. We use the techniques to estimate a small-scale DSGE model to assess the effects of the government spending portion of the American Recovery and Reinvestment Act in 2009 when interest rates reached the zero lower bound.

1. Introduction

Dynamic stochastic general equilibrium (DSGE) models with financial frictions are widely used in central banks, by regulators, and in academia to study the effects of monetary and macroprudential policies and the propagation of shocks in the macro economy. The most recent vintage of these models involves occasionally-binding constraints arising from financial frictions and the effective lower bound (ELB) on nominal interest rates. In order for these models to be usable for

☆ We are thankful for helpful comments and suggestions from participants of the 2018 and 2019 MFM conferences, the 2019 conference of the Society for Nonlinear Dynamics, and the Alejandro Justiniano Memorial conference. Much of this paper was written while Aruoba and Schorfheide visited the Federal Reserve Bank of Philadelphia, whose hospitality they are thankful for. Higa-Flores and Villalvazo gratefully acknowledge financial support from the Becker Friedman Institute under the Macro Financial Modeling Project. Aruoba and Schorfheide gratefully acknowledge financial support from the National Science Foundation under Grant SES 1851634. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or any other person associated with the Federal Reserve System.
* Corresponding author.
E-mail addresses: aruoba@umd.edu (S.B. Aruoba), pablo.a.cubaborda@frb.gov (P. Cuba-Borda), kenjihf@umd.edu (K. Higa-Flores), schorf@ssc.upenn.edu (F. Schorfheide), vsergio@sas.upenn.edu (S. Villalvazo).

https://doi.org/10.1016/j.red.2020.12.003
1094-2025/

a quantitative analysis, they need to be solved numerically, and their parameters need to be estimated based on historical data.

Two types of solution approaches for models with occasionally-binding constraints have been used in the literature. The first group of solution algorithms can be broadly classified as global methods. Agents' decision rules (or value functions associated with optimization problems) are represented by a family of flexible functions—for example, Chebyshev polynomials—or by a discrete mapping on a finite state-space domain. The flexible functions are parameterized by coefficients that are chosen such that the resulting decision rules (approximately) satisfy the model's equilibrium conditions and solve the underlying intertemporal optimization problems. Examples of this approach include Christiano and Fisher (2000), Adam and Billi (2007), Fernández-Villaverde et al. (2015), Maliar and Maliar (2015), Nakata (2016), Gust et al. (2017), Aruoba et al. (2018), Mendoza and Villalvazo (2020), and Atkinson et al. (2020).

The second type of solution approaches are variants of the extended perfect-foresight path (EPFP) method that build on Fair and Taylor (1983). These algorithms rely on the assumption that, after H periods, the system reverts back to the steady state in which the constraint, say, is non-binding. With an initial guess about whether the constraint is binding in periods $t+h$, $h=1,\ldots,H$, it is possible to solve the dynamic system for the values of the endogenous variables. One can then compare the initial guess about the duration of the binding regime to the backward solution and iterate until consistency is achieved. Because the computations are based on the initial state, the previously described steps need to be repeated for every t in a multi-period simulation. Variants of this approach have been used in Eggertsson and Woodford (2003), Christiano et al. (2015), Guerrieri and Iacoviello (2015), Kulish et al. (2017), Holden (2019), and Boehl (2019). The Guerrieri and Iacoviello (2015) paper is accompanied by a popular model solution toolbox called OccBin that implements a variant of the EFPF approach. We will refer to OccBin in various instances throughout our paper.

Given the model solution, one then constructs a state-space representation for an estimable empirical model. The solution itself generates the state transition equations. A set of measurement equations can then be specified that links the state variables with the observables. Because the model solution is nonlinear, so is the state-space representation. Thus, a nonlinear filter is required to compute the likelihood function. For instance, in the context of DSGE models with an ELB constraint, Gust et al. (2017) and Aruoba et al. (2018) use a particle filter in combination with a global solution to construct likelihood functions. Guerrieri and Iacoviello (2017) use an EPFP solution for a model in which the number of observables equals the number of structural shocks and combine it with an inversion filter that essentially solves for the innovations as a function of the observables conditional on an initial state.

Against this backdrop, the contribution of our paper is to construct an alternative model solution that (i) is able to capture an important aspect of the nonlinear decision rule generated by an occasionally-binding constraint, (ii) can be solved quickly, and (iii) allows us to derive an accurate and fast filter for the evaluation of the likelihood function that exploits the structure of the solution. The procedure is efficient and can be run on a desktop computer in a reasonable amount of time. For instance, estimating the small-scale New Keynesian model in our empirical application using data from 1984 to 2018 takes about 18 hours on a single core.

The basic idea of the proposed solution method is to approximate agents' decision rules globally by piecewise-linear functions that are continuous but have a kink along the locus of the state space in which a constraint becomes binding. The coefficients of the decision rules are determined to ensure that the model's equilibrium conditions are (approximately) satisfied. The equilibrium conditions typically take the form of nonlinear expectational difference equations. We require that the (potentially transformed) state variables enter the occasionally-binding constraint linearly. For the remaining equilibrium conditions a (log)linearization is optional. In determining the decision rule coefficients, we take into account that the constraint could either be binding or non-binding in the next period. Thus, we are capturing precautionary behavior. Importantly, the decision rule coefficients only have to be computed once (as opposed to for each period t separately as in the EPFP approach). Compared with higher-order Chebyshev polynomials, the piecewise-linearity and continuity at the kink drastically reduce the number of coefficients that need to be determined and hence simplify computations.

Our solution method is motivated by the observation that more densely parameterized nonlinear decision rules look approximately piecewise-linear in a number of DSGE models. For instance, in Aruoba et al. (2018) we considered a New Keynesian DSGE model and stitched together higher-order Chebyshev polynomials along the locus in the state space where the ELB constraint becomes binding. We found that the decision rules on both sides of the kink are approximately linear. In the Online Appendix we solve a consumption-savings model with an occasionally-binding borrowing constraint and demonstrate that a global solution technique produces approximately piecewise-linear decision rules. Moreover, in Section 3 we solve a simplified version of the New Keynesian DSGE model with an ELB constraint and show that the piecewise-linear structure is exact.

To solve the nonlinear filtering problem, we develop a conditionally optimal particle filter (COPF). A particle filter is a stochastic algorithm that approximates the distribution of a vector of hidden states s_t conditional on the sequence of observations $Y_{1:t}$ available in time t by a swarm of M particle values and weights $\{s_t^j, W_t^j\}_{j=1}^M$. Because of its stochastic structure, repeated runs of the filter generate a distribution of likelihood values. An important property of the particle filter is that the average likelihood across repeated runs is equal to the exact likelihood (unbiasedness). The tuning of the particle filter determines the precision of the approximation. A key step in the specification of the algorithm is the mutation of time $t-1$ particle values into time t particle values. We show how, in the case of a piecewise-linear DSGE model solution, the mutation step can be executed optimally, conditional on the stage $t-1$ particle values.

In a sequence of numerical illustrations based on a small-scale New Keynesian DSGE model with an ELB constraint, we document important properties of our solution algorithm and the COPF likelihood approximation. We show that, compared to a naive bootstrap particle filter (BSPF), which mutates particle values by simulating the model solution forward, our COPF drastically reduces the variance of the likelihood approximation holding the runtime fixed. In practice, this allows users to run the COPF with far fewer particles than the BSPF, which in turn speeds up the computations. When we embed the more accurate COPF into a random walk Metropolis–Hastings (RWMH) algorithm, we are able to significantly reduce the persistence of the resulting Markov chain and therefore improve the accuracy of Monte Carlo approximations of moments of the posterior distribution.

A key feature of our paper is that it integrates model solution, likelihood approximation, and Bayesian estimation. There are a few papers that assess the interplay of existing model solution and likelihood evaluation techniques in Monte Carlo experiments. The ones most closely related to our work are Cuba-Borda et al. (2019) and Atkinson et al. (2020).[1] Cuba-Borda et al. (2019) take a simple consumption-savings model subject to a borrowing constraint and illustrate that less accurate solution methods affect inference even when the inversion filter is available. They also show that, as one increases the measurement error variance in the BSPF, the likelihood misspecification becomes more problematic, making it harder to retrieve the parameter values that govern the data generating process (DGP). In their setting, measurement error and solution approximation error make it difficult for the econometrician to identify the model regime that generates the data, and this incorrect classification of regimes leads to a bias in parameter estimates. In our empirical application, one of the observed time series allows us to exactly identify the regime, and we modify the COPF to capture this feature.

Atkinson et al. (2020) compare the performance of a fully nonlinear solution and a variant of the BSPF for estimation, with the approximated solution using OccBin and the inversion filter. They simulate data from a DSGE model that includes more frictions and shocks than the model used for estimation. As such, their estimated model is misspecified with respect to the DGP. Their results show that the nonlinear approach performs slightly better than the OccBin approach, but the differences are small. Moreover, relative to the pseudo-true parameters, the estimates from both approaches show biases in some key parameters, such as the degree of price rigidities. Since the OccBin-inversion filter approach can be scaled up easily and is faster, they argue that building a bigger and less misspecified model using this approach may be preferable. Similarly, our method offers scalability, even without multicore processing or distributed computing, and allows for more general model structures and state-space representations than the inversion filter.

Based on U.S. data from 1984 to 2018 on output growth, inflation, interest rates, and the government-spending to GDP ratio, we estimate a small-scale DSGE model using our proposed piecewise-linear and continuous (PLC) solution in combination with the COPF. From the estimated model, we compute dollar-for-dollar government spending multipliers associated with the increase in government spending that was part of the 2009 American Recovery and Reinvestment Act (ARRA). The counterfactual output levels are computed by lowering the exogenous government spending process in the model by an amount that is commensurable to the ARRA intervention, setting expansionary policy shocks during that period to zero, and keeping all other exogenous processes at their historical levels. We find that the ex-post multiplier associated with the combined fiscal and monetary policy intervention in 2009 and 2010, when the United States was at the ELB, was around one over a one- to two-year horizon.

The remainder of the paper is organized as follows. Section 2 describes the small-scale New Keynesian DSGE model with ELB constraint used in the subsequent analysis. In Section 3 we solve a simplified version of the New Keynesian model and show that the resulting decision rules are piecewise-linear and continuous. We also provide a comparison to the OccBin solution. In Section 4, we describe how to impose continuity on piecewise-linear decision rules and derive a canonical form for the DSGE model solution. Section 5 discusses how the decision rule coefficients are determined to approximately satisfy the model's equilibrium conditions. The COPF is derived in Section 6. Section 7 presents some numerical experiments to document the accuracy of the likelihood approximation through the COPF, and Section 8 contains the empirical analysis. Finally, Section 9 concludes. Derivations and further implementation details are provided in the Online Appendix. The Appendix also contains a section that shows how to solve a consumption-savings model with an occasionally-binding borrowing constraint using the techniques proposed in this paper and compares our PLC solution to an "exact" solution and a solution constructed with OccBin.

2. A prototypical New Keynesian DSGE model

We will illustrate our solution and filtering methods based on a prototypical New Keynesian DSGE model. The model is identical to the one used in Aruoba et al. (2018). Variants of this model have been widely studied in the literature, and its properties are discussed in detail in Woodford (2003). To make this paper self-contained and introduce some important notation, we briefly describe the preferences and technologies of the agents in Section 2.1 and summarize the equilibrium conditions in Section 2.2.

[1] Boehl (2019) combines his model solution, which is a variant of the EPFP, with a variant of an ensemble Kalman filter. His paper presents an application but does not focus on accuracy comparisons of solution and estimation methods.

2.1. Preferences and technologies

Households. Households derive utility from consumption C_t relative to an exogenous habit stock and disutility from hours worked H_t.[2] The households maximize

$$\mathbb{E}_t\left[\sum_{s=0}^{\infty}\beta^s d_{t+s}\left(\frac{(C_{t+s}/A_{t+s})^{1-\tau}-1}{1-\tau}-\frac{H_{t+s}^{1+1/\eta}}{1+1/\eta}\right)\right], \tag{1}$$

subject to the budget constraint

$$P_tC_t+T_t+B_t=P_tW_tH_t+R_{t-1}B_{t-1}+P_tD_t+P_tSC_t.$$

Here β is the discount factor, d_t is an exogenous shock to the discount factor, $1/\tau$ is the intertemporal elasticity of substitution, and η is the Frisch labor supply elasticity. P_t is the price of the final good. The households receive the real wage W_t in exchange for labor services. B_t is the quantity of nominal bonds, which pay gross interest R_t. Furthermore, the households receive profits D_t from the firms and pay lump-sum taxes T_t. SC_t is the net cash inflow from trading a full set of state-contingent securities.

Firms. The final-goods producers generate aggregate output Y_t by aggregating intermediate goods $Y_t(j)$, $j\in[0,1]$. Under the assumption of perfect competition and free entry, the demand for the intermediate inputs and the price of the aggregate final good are given by

$$Y_t(j)=\left(\frac{P_t(j)}{P_t}\right)^{-1/\nu}Y_t \quad\text{and}\quad P_t=\left(\int_0^1 P_t(j)^{\frac{\nu-1}{\nu}}dj\right)^{\frac{\nu}{\nu-1}}, \tag{2}$$

respectively. We define inflation as $\pi_t=P_t/P_{t-1}$.

Intermediate good j is produced by a monopolist who has access to the production technology

$$Y_t(j)=A_tH_t(j), \tag{3}$$

where A_t is an exogenous productivity process that is common to all firms and $H_t(j)$ is the firm-specific labor input. Intermediate-goods-producing firms face quadratic price adjustment costs of the form

$$AC_t(j)=\frac{\phi}{2}\left(\frac{P_t(j)}{P_{t-1}(j)}-\bar{\pi}\right)^2Y_t(j),$$

where ϕ governs the price stickiness in the economy and $\bar{\pi}$ is a baseline rate of price change that does not require the payment of any adjustment costs. In our quantitative analysis, we set $\bar{\pi}=\pi_*$, where π_* is the target inflation rate of the central bank. Firm j chooses its labor input $H_t(j)$ and the price $P_t(j)$ to maximize the present value of future profits

$$\mathbb{E}_t\left[\sum_{s=0}^{\infty}\beta^s Q_{t+s|t}\left(\frac{P_{t+s}(j)}{P_{t+s}}Y_{t+s}(j)-W_{t+s}H_{t+s}(j)-AC_{t+s}(j)\right)\right]. \tag{4}$$

Here, $Q_{t+s|t}$ is the time t value to the household of a unit of the consumption good in period $t+s$, and is treated as exogenous by the firm.

Government policies. Monetary policy is described by an interest rate feedback rule. Because the ELB constraint is an important part of our analysis we introduce it explicitly as follows:

$$R_t=\max\left\{1,\ R_t^*e^{\sigma_R\epsilon_{R,t}}\right\},\quad R_t^*=\left[r\pi_*\left(\frac{\pi_t}{\pi_*}\right)^{\psi_1}\left(\frac{Y_t}{\gamma Y_{t-1}}\right)^{\psi_2}\right]^{1-\rho_R}R_{t-1}^{\rho_R}. \tag{5}$$

Here R_t^* is the systematic part of monetary policy which reacts to an inflation gap and an output growth gap, r is the steady-state real interest rate, π_* is the target-inflation rate, γ is the long-run growth rate of the economy, and $\epsilon_{R,t}$ is a monetary policy shock.

The government consumes a stochastic fraction of aggregate output. We assume that government spending evolves according to

[2] The habit stock is proxied by the level of technology A_t, which ensures that the economy evolves along a balanced growth path. Since we will not focus on it in the subsequent analysis, we do not make a money-holding motive, such as valuing transaction services from real money balances, explicit in the description of the environment. Such a motive is necessary to make the ELB a relevant constraint in a model like this.

$$G_t = \left(1 - \frac{1}{g_t}\right) Y_t \tag{6}$$

where g_t is an exogenous process. The government levies a lump-sum tax T_t (or provides a subsidy if T_t is negative) to finance any shortfalls in government revenues (or to rebate surpluses). Its budget constraint is given by

$$P_t G_t + M_{t-1} + R_{t-1} B_{t-1} = T_t + M_t + B_t. \tag{7}$$

Exogenous shocks. The model economy is perturbed by four exogenous processes. Aggregate productivity evolves according to

$$\ln A_t = \ln \gamma + \ln A_{t-1} + \ln z_t, \text{ where } \ln z_t = \rho_z \ln z_{t-1} + \sigma_z \epsilon_{z,t}. \tag{8}$$

Thus, on average, the economy grows at the rate γ, and z_t generates exogenous stationary fluctuations of the technology growth rate around this long-run trend. We assume that the government spending shock follows the AR(1) law of motion

$$\ln g_t = (1 - \rho_g) \ln g_* + \rho_g \ln g_{t-1} + \sigma_g \epsilon_{g,t}. \tag{9}$$

The shock to the discount factor evolves according to

$$\ln d_t = \rho_d \ln d_{t-1} + \sigma_d \epsilon_{d,t}. \tag{10}$$

The monetary policy shock $\epsilon_{R,t}$ is assumed to be serially uncorrelated. We stack the four innovations into the vector $\epsilon_t = [\epsilon_{z,t}, \epsilon_{g,t}, \epsilon_{d,t}, \epsilon_{R,t}]'$ and assume that $\epsilon_t \sim iidN(0, I)$.

2.2. Equilibrium conditions

Because the exogenous productivity process has a stochastic trend, it is convenient to characterize the equilibrium conditions of the model economy in terms of detrended consumption $c_t \equiv C_t / A_t$ and detrended output $y_t \equiv Y_t / A_t$.

It is well known that the New Keynesian model features multiple equilibria. In one of the equilibria, the so-called targeted-inflation equilibrium, the endogenous variables fluctuate around the steady state in which inflation equals the value targeted by the central bank. Another important equilibrium is the so-called deflation equilibrium where the economy fluctuates around the so-called deflation steady state in which nominal interest rates are zero.[3] In the remainder of the paper we mostly focus on the targeted-inflation equilibrium, though we also discuss a deflation equilibrium in Section 3. The former is essentially the equilibrium that arises in linearized New Keynesian DSGE models, adjusted for the presence of the ELB constraint. The corresponding steady state is given by

$$\pi_*, \quad r_* = \frac{\gamma}{\beta}, \quad R_* = r_* \pi_*, \quad y_* = \left[(1-\nu) g_*^{\tau}\right]^{\frac{1}{\tau + 1/\eta}}, \quad c_* = \frac{y_*}{g_*}. \tag{11}$$

Without loss of generality, for any variable x_t we can define the percentage deviations from this steady state as $\hat{x}_t = \ln x_t - \ln x_*$. Using this notation we can substitute x_t by $x_* e^{\hat{x}_t}$ and obtain[4]:

$$\begin{aligned}
1 &= \mathbb{E}_t \left[e^{\hat{R}_t - \hat{\pi}_{t+1} - \hat{z}_{t+1} + \hat{d}_{t+1} - \hat{d}_t - \tau(\hat{c}_{t+1} - \hat{c}_t)} \right] \\
&\left(1 - \nu - \chi_h c_*^{\tau} y_*^{\frac{1}{\eta}} e^{\tau \hat{c}_t + \frac{1}{\eta}\hat{y}_t}\right) + \nu\phi\pi_*(\pi_* e^{\hat{\pi}_t} - \bar{\pi}) e^{\hat{\pi}_t} - \frac{\phi}{2}\left(\pi_* e^{\hat{\pi}_t} - \bar{\pi}\right)^2 \\
&= \nu\phi\beta\pi_* \mathbb{E}_t \left[\left(e^{-\tau(\hat{c}_{t+1} - \hat{c}_t) + \hat{d}_{t+1} - \hat{d}_t + \hat{y}_{t+1} - \hat{y}_t + \hat{\pi}_{t+1}} \right) \left(\pi_* e^{\hat{\pi}_{t+1}} - \bar{\pi} \right) \right] \\
1 &= \left[\frac{1}{e^{\hat{g}_t}} - g_* \pi_*^2 \frac{\phi}{2} \left(e^{\hat{\pi}_t} - 1 \right)^2 \right] e^{\hat{y}_t - \hat{c}_t}.
\end{aligned} \tag{12}$$

The first equation is the consumption Euler equation, the second equation is the optimality condition for the firms, and the third equation is the aggregate resource constraint. The monetary policy rule can be expressed as

$$\hat{R}_t = \max\left\{ (1 - \rho_R)\left[\psi_1 \hat{\pi}_t + \psi_2 (\hat{y}_t - \hat{y}_{t-1} + \hat{z}_t)\right] + \rho_R \hat{R}_{t-1} + \sigma_R \epsilon_{R,t}, -\ln(r\pi_*) \right\}. \tag{13}$$

More compactly, the equilibrium conditions of the DSGE models can be expressed as a system of expectational difference equations of the form

$$\mathbb{E}_t\left[\mathcal{R}(\hat{y}_t, \hat{c}_t, \hat{\pi}_t, \hat{R}_t, \hat{y}_{t+1}, \hat{c}_{t+1}, \hat{\pi}_{t+1}, \hat{R}_{t+1}, \ldots)\right] = 0, \tag{14}$$

where $\mathcal{R}(\cdot)$ captures residuals in the equilibrium conditions.

[3] See, for instance, Benhabib et al. (2001), Aruoba and Schorfheide (2016) and Aruoba et al. (2018) for a discussion of multiplicity of equilibria in this model.

[4] Introducing $\hat{x}_t$ does not imply that we are log-linearizing all of the equilibrium conditions. It is foremost a reparameterization. However, in our model it happens to be the case that the consumption Euler equation and (abstracting from the max operator) the monetary policy rule are log-linear.

3. Solving a simplified version of the DSGE model

In order to highlight some important features of the proposed solution method, we first consider a highly simplified version of the DSGE model introduced in Section 2 that can be (almost) solved analytically.[5] In particular, we will show that the PLC form emerges as the exact solution to the simplified DSGE model. We also discuss the multiplicity of solutions and provide a comparison to OccBin.

On the model described in the previous section we impose the parameter restrictions $\tau = 1$, $\gamma = 1$, $\eta = \infty$, $g_* = 1$, $\bar{\pi} = \pi_*$, $\psi_1 = \psi$, $\psi_2 = 0$, $\rho_R = 0$, $\sigma_z = 0$, $\sigma_g = 0$, and $\rho_d = 0$. We log-linearize the equilibrium conditions (except for the ELB constraint) around the targeted inflation steady state in (11) and regard the resulting equations as *the model* to be solved. Some details of the calculations are relegated to the Online Appendix.

Equilibrium conditions. The log-linearized equilibrium conditions for the simplified model are of the form

$$\begin{aligned} -\hat{c}_t &= \mathbb{E}[\hat{d}_{t+1}] - \hat{d}_t - \mathbb{E}[\hat{c}_{t+1}] + \hat{R}_t - \mathbb{E}[\hat{\pi}_{t+1}] \\ \hat{\pi}_t &= \beta \mathbb{E}[\hat{\pi}_{t+1}] + \kappa \hat{c}_t \\ \hat{R}_t &= \max\left\{\psi \hat{\pi}_t + \sigma_R \epsilon_{R,t},\ -\ln(r\pi_*)\right\}. \end{aligned} \tag{15}$$

Under parameterizations in which monetary policy is active, i.e., $\psi > 1$, the model has typically two stationary solutions in which $(\hat{R}_t, \hat{c}_t, \hat{\pi}_t)$ are independently and identically distributed (iid) over time. These correspond to the targeted-inflation and the deflation equilibria defined in the previous section.

In both equilibria expected consumption and inflation are time-invariant and can be replaced by $\mu_c = \mathbb{E}_t[\hat{c}_{t+1}]$ and $\mu_\pi = \mathbb{E}_t[\hat{\pi}_{t+1}]$. Thus, setting the expected value of the residual function to zero, and conducting a few basic algebraic manipulations, we obtain:

$$\begin{aligned} \hat{R}(\epsilon_{d,t}, \epsilon_{R,t}) &= \max\left\{\frac{1}{1+\psi\kappa}\left[\psi\kappa\mu_c + \psi(\kappa+\beta)\mu_\pi + \psi\kappa\sigma_d\epsilon_{d,t} + \sigma_R\epsilon_{R,t}\right],\ -\ln(r\pi_*)\right\} \\ \hat{c}(\epsilon_{d,t}, \epsilon_{R,t}) &= -\hat{R}(\epsilon_{d,t}, \epsilon_{R,t}) + \mu_c + \mu_\pi + \sigma_d\epsilon_{d,t} \\ \hat{\pi}(\epsilon_{d,t}, \epsilon_{R,t}) &= -\kappa\hat{R}(\epsilon_{d,t}, \epsilon_{R,t}) + \kappa\mu_c + (\kappa+\beta)\mu_\pi + \kappa\sigma_d\epsilon_{d,t}. \end{aligned} \tag{16}$$

Here we replaced $\hat{R}_t$, $\hat{c}_t$, and $\hat{\pi}_t$ by decision rules that are time-invariant functions of the state variables $(\epsilon_{d,t}, \epsilon_{R,t})$.

Constructing a solution. In order to solve the system (16) we need to find constants μ_c and μ_π such that

$$\mu_c = \mathbb{E}[c(\epsilon_{d,t}, \epsilon_{R,t})] \quad \text{and} \quad \mu_\pi = \mathbb{E}[\pi(\epsilon_{d,t}, \epsilon_{R,t})].$$

Because of the max operator in the monetary policy rule, this requires the computation of the mean of a truncated Normal random variable. We rotate the vector of innovations $[\epsilon_{d,t}, \epsilon_{R,t}]'$ to separate the component that enters the monetary policy rule from a second component that is orthogonal:

$$\eta_{1,t} = \frac{1}{\sigma_\eta}(\psi\kappa\sigma_d\epsilon_{d,t} + \sigma_R\epsilon_{R,t}), \quad \eta_{2,t} = \frac{1}{\sigma_\eta}(\sigma_R\epsilon_{d,t} - \psi\kappa\sigma_d\epsilon_{R,t}), \quad \sigma_\eta = \sqrt{(\psi\kappa\sigma_d)^2 + \sigma_R^2}. \tag{17}$$

By construction, the innovations $\eta_{1,t}$ and $\eta_{2,t}$ are also $N(0,1)$.

Using the expression for $\eta_{1,t}$ in (17) we can rewrite the interest rate rule as

$$R(\eta_{1,t}) = \max\left\{\frac{1}{1+\psi\kappa}\left[\psi\kappa\mu_c + \psi(\kappa+\beta)\mu_\pi + \sigma_\eta\eta_{1,t}\right],\ -\ln(r\pi_*)\right\}. \tag{18}$$

Define the cutoff value

$$\bar{\eta}_1 = -\frac{1}{\sigma_\eta}\left[(1+\psi\kappa)\ln(r\pi_*) + \psi\kappa\mu_c + \psi(\kappa+\beta)\mu_\pi\right] \tag{19}$$

such that $R(\eta_{1,t}) = -\ln(r\pi_*)$ whenever $\eta_{1,t} \le \bar{\eta}_1$. Using the formula for the mean of a truncated standard normal random variable, we obtain

$$\begin{aligned} \mathbb{E}[R(\eta_{1,t})] = &-\Phi_N(\bar{\eta}_1)\ln(r\pi_*) \\ &+\frac{1}{1+\psi\kappa}\left[\big(1-\Phi_N(\bar{\eta}_1)\big)\big(\psi\kappa\mu_c + \psi(\kappa+\beta)\mu_\pi\big) + \sigma_\eta\phi_N(\bar{\eta}_1)\right]. \end{aligned} \tag{20}$$

[5] A similar model was solved in Mendes (2011).

Here $\Phi_N(\cdot)$ and $\phi_N(\cdot)$ are the cumulative density function (cdf) and the probability density function (pdf) of a standard Normal random variable. Taking expectations of the second and third equation in (16) and substituting out $\mathbb{E}[\hat{R}(\epsilon_{d,t}, \epsilon_{R,t})] = \mathbb{E}[R(\eta_{1,t})]$ using (20) leads to the following nonlinear system of equations:

$$\mu_c = \frac{1-\Phi_N(\bar{\eta}_1)}{1+\psi\kappa}\Big[\mu_c + (1-\psi\beta)\mu_\pi\Big] - \frac{\sigma_\eta \phi_N(\bar{\eta}_1)}{1+\psi\kappa} + \Phi_N(\bar{\eta}_1)\Big[\ln(r\pi_*) + \mu_c + \mu_\pi\Big]. \tag{21}$$

$$\mu_\pi = \frac{1-\Phi_N(\bar{\eta}_1)}{1+\psi\kappa}\Big[\kappa\mu_c + (\kappa+\beta)\mu_\pi\Big] - \frac{\kappa\sigma_\eta \phi_N(\bar{\eta}_1)}{1+\psi\kappa} + \Phi_N(\bar{\eta}_1)\Big[\kappa\ln(r\pi_*) + \kappa\mu_c + (\kappa+\beta)\mu_\pi\Big],$$

where $\bar{\eta}_1$ is given in (19). Conditional on $\bar{\eta}_1$ the system is linear in (μ_c, μ_π) which means that it can be reduced to a single nonlinear equation in terms of $\bar{\eta}_1$ that needs to be solved numerically. This equation typically has two solutions, which generate the targeted-inflation and deflation equilibrium, respectively.

Properties of the solution. *(i) The decision rules are piecewise-linear and continuous for each* (μ_c, μ_π). Plugging the expression for $\hat{R}(\epsilon_{d,t}, \epsilon_{R,t}) = R(\eta_{1,t})$ in (20) into the consumption and inflation decision rules in (16) and using the relationship between the ϵ_t's and η_t's in (17), we can rewrite the decision rules for consumption and inflation as follows:

$$c(\eta_{1,t}, \eta_{2,t}) = \begin{cases} \frac{1}{1+\psi\kappa}\Big[\mu_c + (1-\psi\beta)\mu_\pi\Big] + \big(\gamma_{d,1} - \frac{1}{1+\psi\kappa}\big)\sigma_\eta\eta_{1,t} + \gamma_{d,2}\sigma_\eta\eta_{2,t} & \text{if } \eta_{1,t} > \bar{\eta}_1 \\ \ln(r\pi_*) + \mu_c + \mu_\pi + \gamma_{d,1}\sigma_\eta\eta_{1,t} + \gamma_{d,2}\sigma_\eta\eta_{2,t} & \text{otherwise} \end{cases} \tag{22}$$

$$\pi(\eta_{1,t}, \eta_{2,t}) = \begin{cases} \frac{1}{1+\psi\kappa}\Big[\kappa\mu_c + (\kappa+\beta)\mu_\pi\Big] + \kappa\big(\gamma_{d,1} - \frac{1}{1+\psi\kappa}\big)\sigma_\eta\eta_{1,t} + \kappa\gamma_{d,2}\sigma_\eta\eta_{2,t} & \text{if } \eta_{1,t} > \bar{\eta}_1 \\ \kappa\ln(r\pi_*) + \kappa\mu_c + (\kappa+\beta)\mu_\pi + \kappa\gamma_{d,1}\sigma_\eta\eta_{1,t} + \kappa\gamma_{d,2}\sigma_\eta\eta_{2,t} & \text{otherwise} \end{cases}$$

Here, the γ constants are functions of the structural parameters obtained from (17) by re-arranging the equations to express the ϵ_t's as a function of the η_t's. The ELB becomes binding at the locus in the state space defined by

$$\bar{\eta}_1 = \eta_{1,t} = \frac{1}{\sigma_\eta}(\psi\kappa\sigma_d\epsilon_{d,t} + \sigma_R\epsilon_{R,t}). \tag{23}$$

The second innovation, $\eta_{2,t}$, does not enter the policy rule and therefore cannot push the economy toward the ELB. The decision rules are piecewise-linear functions of the innovations $\eta_{1,t}$ and $\eta_{2,t}$. The slope coefficients associated with $\eta_{1,t}$ change when the economy hits the ELB at $\eta_{1,t} = \bar{\eta}_1$, whereas the slope coefficients for $\eta_{2,t}$ do not change. We verify in the Online Appendix that the decision rules are also continuous at $\eta_{1,t} = \bar{\eta}_1$, meaning that for each $\eta_{2,t}$ and for each sequence $\eta^n_{1,t} \longrightarrow \bar{\eta}_1$ as $n \longrightarrow \infty$:

$$\big|c(\eta^n_{1,t}, \eta_{2,t}) - c(\bar{\eta}_1, \eta_{2,t})\big| \longrightarrow 0 \quad \text{and} \quad \big|\pi(\eta^n_{1,t}, \eta_{2,t}) - \pi(\bar{\eta}_1, \eta_{2,t})\big| \longrightarrow 0. \tag{24}$$

Because the transformation between $(\epsilon_{d,t}, \epsilon_{R,t})$ and $(\eta_{1,t}, \eta_{2,t})$ is continuous, the continuity result also holds for the decision rules $\hat{c}(\epsilon_{d,t}, \epsilon_{R,t})$ and $\hat{\pi}(\epsilon_{d,t}, \epsilon_{R,t})$, expressed in terms of the original state variable, as $(\psi\kappa\sigma_d\epsilon^n_{d,t} + \sigma_R\epsilon^n_{R,t})/\sigma_\eta \longrightarrow \bar{\eta}_1$.

(ii) The law of motion given by the interest rate rule (18) and the consumption and inflation decision rules (22) is coherent and complete for each (μ_c, μ_π) *and each realization of the innovations* $(\eta_{1,t}, \eta_{2,t})$. The system takes the form of a linear simultaneous equations model with regime switches. The concepts of coherency and completeness were introduced by Gourieroux et al. (1980) and more recently studied in the context of ELB applications by Mavroeidis (2020) and Ascari and Mavroeidis (2020). Coherency requires that given any innovation $(\eta_{1,t}, \eta_{2,t})$ there exists a solution to the system of equations. Completeness refers to the uniqueness of that solution. Define

$$R^{(1)}_t = \frac{1}{1+\psi\kappa}\Big[\psi\kappa\mu_c + \psi(\kappa+\beta)\mu_\pi + \sigma_\eta\eta_{1,t}\Big] \quad \text{and} \quad R^{(2)}_t = -\ln(r\pi_*).$$

Coherency and completeness require that: (a) $\eta_{1,t} > \bar{\eta}_1$ implies $R_t = R^{(1)}_t$ is a solution to the maximization on the right-hand side of (20), whereas $R_t = R^{(2)}_t$ is not. Likewise, it is required that (b) $\eta_{1,t} < \bar{\eta}_1$ implies $R_t = R^{(2)}_t$ is a solution to the maximization on the right-hand side of (20), whereas $R_t = R^{(1)}_t$ is not. Coherency and completeness follows directly from the linearity of $R^{(1)}_t$ with respect to $\eta_{1,t}$ and the definition of $\bar{\eta}_1$ in (19) and require no further restrictions on the domain of the innovations.

(iii) Solutions for μ_c *and* μ_π. The means μ_c and μ_π together with the cutoff value $\bar{\eta}_1$ are determined by the nonlinear system of equations (19) and (21). To understand the properties of the nonlinear system, assume that $\psi > 1$, i.e., monetary policy is active, and $\sigma_\eta = 0$, i.e., there is no uncertainty. First, suppose we start with the conjecture that $\bar{\eta}_1 = -\infty$. Then (21) simplifies to

$$\mu_c = \frac{1}{1+\psi\kappa}\Big[\mu_c + (1-\psi\beta)\mu_\pi\Big], \quad \mu_\pi = \frac{1}{1+\psi\kappa}\Big[\kappa\mu_c + (\kappa+\beta)\mu_\pi\Big]$$

which is solved by

$$\mu_c = 0, \quad \mu_\pi = 0. \tag{25}$$

Recall that the system was expressed in deviations from the targeted-inflation steady state. Thus if $\mu_c = \mu_\pi = 0$, then the means of consumption and inflation are equal to the steady state. The analysis is completed by noting that (19) implies that indeed $\bar{\eta}_1 = -\infty$ as initially assumed.

Second, suppose we start from the conjecture that $\bar{\eta}_1 = +\infty$. Then (21) simplifies to

$$\mu_c = \ln(r\pi_*) + \mu_c + \mu_\pi, \quad \mu_\pi = \kappa \ln(r\pi_*) + \kappa\mu_c + (\kappa + \beta)\mu_\pi,$$

which implies

$$\mu_c = -\frac{1}{\kappa}(1-\beta)\ln(r\pi_*), \quad \mu_\pi = -\ln(r\pi_*). \tag{26}$$

Substituting the means into (19), we obtain $\bar{\eta}_1 = (\psi - 1)\ln(r\pi_*)/\sigma_\eta > 0$. Thus, as $\sigma_\eta \longrightarrow 0$, $\bar{\eta}_1 \longrightarrow +\infty$ as required. In this case, the system is in the so-called deflation steady state and the ELB constraint is always binding.

Once we allow for uncertainty, $\sigma_\eta > 0$, then the means in (25) and (26) no longer solve the system of equations (19) and (21). However, for values of σ_η that are not "too large," one can obtain solutions that are "close" to the ones derived above. Due to the nonlinearity of the $\Phi_N(\cdot)$ and $\phi_N(\cdot)$ functions, these solutions can only be computed numerically.

Comparison to OccBin solution. It is instructive to compare the above solution to the one generated by OccBin. OccBin requires the choice of a reference regime. Because subsequently we focus on the targeted inflation equilibrium in which the ELB is non-binding with high probability, we impose that the ELB is non-binding in the reference regime. Consider a generic period t. The OccBin solution is based on the assumption that for $\tau \geq T$ the economy will be in the reference regime. Given the lack of dynamics in the simple model, we can choose $T = t + 1$ and assume that the regime will remain in the targeted-inflation steady state: $\hat{R}_{t+1} = 0$, $\hat{c}_\tau = 0$, and $\hat{\pi}_\tau = 0$. For period t, the algorithm draws the shocks $(\epsilon_{d,t}, \epsilon_{R,t})$ and solves the system (15) under the restrictions that $\hat{d}_{t+1} = 0$, $\mathbb{E}_t[\hat{c}_{t+1}] = 0$ and $\mathbb{E}_t[\hat{\pi}_{t+1}] = 0$. The solution is identical to (18) and (22) with $\mu_c = \mu_\pi = 0$ imposed. Thus, in the context of our stylized model the difference between our solution, which happens to be exact, and the OccBin solution is that the latter does not take into account the uncertainty about the regime in period $t + 1$.

Summary. We draw the following conclusions from the analysis of the simplified DSGE model. First, the log-linearized model has static solutions that are exact and in which the decision rules for consumption and inflation are piecewise-linear and continuous. We use this as a motivation for subsequently considering a class of approximate solutions with PLC decision rules for richer nonlinear DSGE models with occasionally-binding constraints. Second, although we approximated the equilibrium conditions around the targeted-inflation steady state, because of the nonlinearity generated by the occasionally-binding constraint, for $\sigma_\eta > 0$ the decision rules do not pass through the steady state around which the model was initially approximated. In fact, because the decision rules are flexible enough to have unrestricted intercepts, we can also generate the deflation equilibrium which is far away from the targeted-inflation steady state. Third, an important difference between our proposed PLC solution and the OccBin solution is that our decision rule coefficients capture uncertainty about the future.

4. PLC decision rules and the canonical form

In the analysis of the simplified model in the previous section, the PLC decision rules emerged from the analytical solution of the model. For more elaborate DSGE models, we will parameterize a family of piecewise-linear decision rules and then impose coefficient restrictions that guarantee that the decision rules are continuous at the kink, where the constraint changes from being slack to being binding. The remaining free coefficients of the decision rules can then be used to (approximately) satisfy the equilibrium conditions of the model by setting the residual functions (close) to zero. A discussion of how to do this numerically is deferred to Section 5.

4.1. PLC decision rules

Let $\mathbb{X} = [x_1, X_2'] \in \mathcal{X}$ be an $n \times 1$ vector of non-redundant state variables. We assume that $\mathbb{X}$ also contains a constant. Here x_1 is one particular element of $\mathbb{X}$ that enters the characterization of the locus of points in the state space at which the constraint becomes binding and the decision rules have their kink. The reason for separating out one of the $\mathbb{X}$ elements will become clear below. Let Y denote a $k \times 1$ vector of control variables. As we make explicit below, we assume Y depends on x_1 and X_2 linearly where the coefficients may depend on whether the constraint is binding or not.

We assume that there is a linear(ized) scalar-valued function $h(x_1, X_2, Y)$ that determines whether the constraint is binding:

$$h(x_1, X_2, Y) = \begin{cases} > 0 & \text{if constraint is non-binding (n)} \\ \leq 0 & \text{if constraint is binding (b)} \end{cases}. \tag{27}$$

The $h(\cdot)$ function may depend on the state variables (x_1, X_2) and some of the elements in Y. Because the function is assumed to be linear, we write it as

$$h(x_1, X_2, y) = \gamma_1 x_1 + \gamma_2' X_2 + \gamma_Y' Y. \tag{28}$$

The γ's are not free coefficients. They are obtained from the equilibrium conditions of the DSGE model. In the simplified model of Section 3 we set $x_1 = 1$, $X_2 = [\epsilon_d, \epsilon_R]'$, and $y = [\hat{c}, \hat{\pi}]'$ such that $h(x_1, X_2, y) = \ln(r\pi_*) + \epsilon_R + \psi\hat{c}$ with $\gamma_1 = \ln(r\pi_*)$, $\gamma_2' = [0,\ 1]'$, and $\gamma_Y' = [0,\ \psi]$.

We define the kink function $x_1 = \ell(X_2)$ such that $[\ell(X_2), X_2']' \in \mathcal{X}$ characterizes the locus of points in the state space for which $h\big(\ell(X_2), X_2, Y(\ell(X_2), X_2)\big) = 0$, that is, the constraint is just binding. Here $Y(\cdot)$ denotes the assumed piecewise-linear decision rules for the control variables. The linearity of $h(\cdot)$ in (28) and the assumed piecewise-linearity of the decision rules for y imply that $\ell(X_2)$ is a linear function and we parameterize it as

$$\ell(X_2) = \delta' X_2, \tag{29}$$

where δ is a $(n-1)\times 1$. The δ coefficients will be determined as functions of the decision rule coefficients and the coefficients of the constraint function $h(\cdot)$. So far we have not yet made a determination whether the constraint is slack if $x_1 < \delta' X_2$. In the simplified model the equilibrium kink function is given by (23).

Returning to the control variables, we assume the decision rules for each y^i are of the piecewise-linear form

$$Y^i(x_1, X_2) = \begin{cases} \alpha^i_{1,1} x_1 + \alpha^i_{1,2}{}' X_2 & \text{if } x_1 \geq \ell(X_2) \\ \alpha^i_{2,1} x_1 + \alpha^i_{2,2}{}' X_2 & \text{if } x_1 < \ell(X_2) \end{cases} \quad i = 1, \ldots, k\,, \tag{30}$$

where each decision rule has $2n$ unknown coefficients. The decision rules are exactly linear if $\alpha^i_{1,1} = \alpha^i_{2,1}$ and $\alpha^i_{1,2} = \alpha^i_{2,2}$. The specification in (30) makes the benefit of using the kink function $\ell(\cdot)$ clear: given the state variables x_1 and X_2, we can easily determine on which side of the constraint we need to be, even when the constraint contains some control variables. In the simplified model the equilibrium decision rules are given by (22). Once we replace the η_t's by the ϵ_t's we obtain the same form as (30), where the α coefficients were determined such that the decision rules satisfy the equilibrium conditions of the model.

4.2. Imposing continuity on piecewise-linear decision rules

We now turn to imposing continuity on the decision rules at the kink, which means we impose the restriction that the two parts of each decision rule are equal to each other along the kink. Doing so will restrict a subset of the unknown α and δ coefficients. Continuity at $x_1 = \delta' X_2$ requires that for each $i = 1, ..., k$

$$\alpha^i_{1,1}\delta' X_2 + \alpha^i_{1,2}{}' X_2 = \alpha^i_{2,1}\delta' X_2 + \alpha^i_{2,2}{}' X_2 \quad \forall X_2,$$

which generates $(n-1)$ restrictions for each i:

$$\alpha^i_{1,1}\delta' + \alpha^i_{1,2}{}' = \alpha^i_{2,1}\delta' + \alpha^i_{2,2}{}'. \tag{31}$$

Next, we impose restrictions that make the $\ell(\cdot)$ and $Y(\cdot)$ functions consistent with the constraint in (27). The condition $h[g(X_2), X_2, Y\big(\ell(X_2), X_2\big)] = 0$, which represents the kink in terms of the $h(\cdot)$ function, can be written as

$$\gamma_1\delta' X_2 + \gamma_2' X_2 + \sum_{i=1}^{k} \gamma^i_Y (\alpha^i_{1,1}\delta' X_2 + \alpha^i_{1,2}{}' X_2) = 0 \quad \forall X_2,$$

which leads to another set of $(n-1)$ restrictions:

$$\gamma_1\delta' + \gamma_2' + \sum_{i=1}^{k} \gamma^i_Y (\alpha^i_{1,1}\delta' + \alpha^i_{1,2}{}') = 0. \tag{32}$$

Counting all unknowns and restrictions, we have $k(n+1)$ degrees of freedom.[6] Let us assume the coefficients $\alpha^i_{1,1}$, $\alpha^i_{1,2}$, and $\alpha^i_{2,1}$ for each decision rule are free and collect them in the vector ϑ of size $k(n+1)$

$$\vartheta = [\alpha^1_{1,1}, \ldots, \alpha^k_{1,1}, \alpha^1_{1,2}{}', \ldots, \alpha^k_{1,2}{}', \alpha^1_{2,1}, \ldots, \alpha^k_{2,1}]'. \tag{33}$$

[6] There are $2n$ α coefficients for each decision rule and $(n-1)$ δ coefficients, which yield $2nk + n - 1$ unknowns. With $(n-1)$ restrictions for each decision rule as derived in (31) and the $(n-1)$ restrictions in (32), we get $(k+1)(n-1)$ restrictions. Subtracting the number of restrictions from the number of unknowns, we get $k(n+1)$.

In other words, we treat all the decision rule coefficients for the "1" regime and the coefficient in front of x_1 in the "2" regime as free. The remaining decision rule coefficients in the "2" regime, $\alpha^i_{2,2}$, $i = 1, \ldots, k$, as well as all of the δ coefficients are determined as functions of these free coefficients, which we now turn to. In our application the choice of which coefficients go in to ϑ is driven by numerical considerations. The "1" regime corresponds to the ELB not being binding. In the targeted-inflation equilibrium, a good starting value for the numerical procedure that is used to determine ϑ is given by the decision-rule coefficients of a log-linear approximation that ignores the ELB constraint.

Conditional on ϑ, we can rewrite (32) as

$$\underbrace{\left(\gamma_1 + \sum_{i=1}^{k} \gamma_Y^i \alpha_{1,1}^i\right)}_{a(\vartheta)} \delta'(\vartheta) + \underbrace{\left(\gamma_2 + \sum_{i=1}^{k} \gamma_Y^i \alpha_{1,2}^i\right)}_{-B'(\vartheta)} = 0$$

and solve for δ as

$$\delta'(\vartheta) = \frac{1}{a(\vartheta)} B'(\vartheta), \tag{34}$$

where $a(\vartheta)$ is a scalar and $B(\vartheta)$ is a $(n-1)$-dimensional vector. By combining (34) with (31) we obtain an expression for the constrained decision rule coefficients $\alpha^i_{2,2}$:

$$\alpha^i_{2,2}{}'(\vartheta) = (\alpha^i_{1,1} - \alpha^i_{2,1}) \left(\frac{1}{a(\vartheta)} B(\vartheta)\right) + \alpha^i_{1,2}{}'. \tag{35}$$

The last step is to determine which part of the decision rule in (30) corresponds to the part of the state space where the constraint is slack and which part where the constraint is binding. Take $h\big(x_1, X_2, Y(x_1, X_2)\big)$ for some x_1 and X_2. Let us derive how its sign depends on the sign of $(x_1 - \delta(\vartheta)' X_2)$. First, assume $x_1 > \delta'(\vartheta) X_2$, then

$$\begin{aligned} h\big[x_1, X_2, Y(x_1, X_2)\big] &= \gamma_1 x_1 + \gamma_2' X_2 + \sum_{i=1}^{k} \gamma_Y^i (\alpha^i_{1,1} x_1 + \alpha^i_{1,2}{}' X_2) \\ &= \underbrace{\left(\gamma_1 + \sum_{i=1}^{k} \gamma_Y^i \alpha^i_{1,1}\right)}_{c(\vartheta)} x_1 + \underbrace{\left(\gamma_2' + \sum_{i=1}^{k} \gamma_Y^i \alpha^i_{1,2}{}'\right)}_{D'(\vartheta)} X_2, \end{aligned} \tag{36}$$

where $c(\vartheta)$ is a scalar and $D(\vartheta)'$ is a $(n-1)$-dimensional vector, which can be evaluated given model parameters and free decision-rule coefficients. Collecting the δ terms, we notice that the restriction in (32) implies $c(\vartheta)\delta'(\vartheta) + D'(\vartheta) = 0$, or $D'(\vartheta) = -c(\vartheta)\delta'(\vartheta)$. Using this on (36), we get

$$h[x_1, X_2, Y(x_1, X_2)] = c(\vartheta)[x_1 - \delta'(\vartheta) X_2]. \tag{37}$$

Since we assumed $x_1 > \delta'(\vartheta) X_2$ above in the derivations, we conclude that $h(\cdot) > 0$ if and only if $c(\vartheta) > 0$. In other words, if $c(\vartheta) > 0$, then the "1" regime in (30) corresponds to the constraint being slack.

4.3. Example: the full New Keynesian model

We now adapt the generic notation so far to the New Keynesian model with ELB constraint described in Section 2. We partition the state space as $x_{1,t} = \hat{R}_{t-1}$ and $X_{2,t} = [1, \hat{y}_{t-1}, \hat{z}_t, \hat{d}_t, \hat{g}_t, e_{R,t}]'$, and we have $n = 7$.[7] As for the choice of control variables to approximate, we have a few options. We approximate the decision rules $\hat{\pi}(x_1, X_2)$ and $\hat{y}(x_1, X_2)$ directly and let the remaining control variables $\hat{c}$ and $\hat{R}$ follow exactly from the equilibrium conditions. Specifically, given x_1, X_2, $\hat{y}_t = \hat{y}(x_1, X_2)$ and $\hat{\pi}_t = \hat{\pi}(x_1, X_2)$, $\hat{c}(x_1, X_2)$ follows from solving the aggregate resource constraint for $\hat{c}_t$. $\hat{R}(x_1, X_2)$ is obtained by plugging $\hat{y}(x_1, X_2)$ and $\hat{\pi}(x_1, X_2)$ into the monetary policy rule in (13). Thus, we set $Y(\cdot) = \big[\hat{y}(\cdot), \hat{\pi}(\cdot)\big]'$ and $k = 2$.

The ELB constraint can be written in terms of the variables defined so far as $\hat{R}_t + \ln(r_* \pi_*) \geq 0$, which leads to the $h(\cdot)$ function

$$h\big(x_{1,t}, X_{2,t}, Y_t(\cdot)\big) = (1-\rho)\big[\psi_1 \hat{\pi}(\cdot) + \psi_2(\hat{y}(\cdot) - \hat{y}_{t-1} + \hat{z}_t)\big] + \rho_R \hat{R}_{t-1} + e_{R,t} + \ln(r_* \pi_*). \tag{38}$$

Thus, the γ coefficients in (28) are

[7] In principle, one of the other state variables could have been chosen as x_1. However, we found it natural to use the lagged interest rates, because all else being equal, higher lagged interest rates move the economy away from the ELB constraint.

$$\gamma_1 = \rho_R,\ \gamma_2' = [\ln(r_*\pi_*),\ -(1-\rho_R)\psi_2,\ (1-\rho_R)\psi_2,\ 0,\ 0],\ \gamma_Y' = [(1-\rho_R)\psi_1,\ (1-\rho_R)\psi_2],$$

and we can write $c(\vartheta)$ in (36) as

$$c(\vartheta) = \rho_R + (1-\rho_R)\psi_1\alpha^{\pi}_{1,1} + (1-\rho_R)\psi_2\alpha^{y}_{1,1}.$$

If $\alpha^{\pi}_{1,1}$ and $\alpha^{y}_{1,1}$ are both positive, then $c(\vartheta)$ is also positive because the remaining structural parameters are positive under standard parameterizations. Thus, we label the "1" regime as the regime where the ELB is slack, "n" (non-binding), and the "2" regime as the "b" (binding) regime. We check that $c(\vartheta)$ is indeed positive every time we solve the model.

4.4. Canonical form

The final step in preparing the model solution for filtering is to cast the solution in the following canonical form:

$$s_t = \begin{cases} \Phi_0(n) + \Phi_1(n)s_{t-1} + \Phi_\eta(n)\eta_t & \text{if } \eta_{1,t} < \zeta(s_{t-1}) \\ \Phi_0(b) + \Phi_1(b)s_{t-1} + \Phi_\eta(b)\eta_t & \text{otherwise,} \end{cases} \tag{39}$$

which is a VAR for s_t with endogenous regime switching. The innovations η_t are a function of the structural innovations ϵ_t. This transformation is done so that the first element of η_t, $\eta_{1,t}$, is a linear combination of structural shocks that determines whether the constraint is binding in period t. Precise definitions of η_t, the impact matrices $\Phi_\eta(\cdot)$ and the threshold function $\zeta(\cdot)$ will be provided below. Recall that we used the rotated η_t shocks also in the construction of the solution to the simplified DSGE model in Section 3 and its decision rules (22) were written in the canonical form.

Equation (39) will serve as a transition equation in a state-space model. Thus, the vector s_t needs to include all variables (whether or not they are directly approximated) that are necessary for the construction of the measurement equations and all variables necessary to determine the transition of such variables, which are all the state variables. The canonical form resembles a regime-switching VAR with a "binding" (b) and "non-binding" (n) regime. However, the regime transition is not determined by an exogenous Markov process. Instead, it is determined by the realization of the shock innovations.[8] Whether the coherency and completeness conditions are satisfied – recall that we showed in Section 3 that they are satisfied for the simplified New Keynesian model – depends on the structure of the $\Phi(\cdot)$ matrices and is model specific. The construction of the canonical form for the New Keynesian DSGE model of Section 2 is outlined in the Online Appendix.

4.5. Measurement equations

The key requirement for the conditionally optimal particle filter that is developed in Section 6 is that the conditional mean function (given s_{t-1}) of the observables is piecewise-linear. This is guaranteed if the state-transition equation has the canonical form (39) and the measurement equation is linear in s_t as in

$$y^o_t = A_0 + A_s s_t + u_t, \quad u_t \sim N(0, \varsigma\Sigma_u), \tag{40}$$

where y^o_t is the vector of observable variables, u_t is a vector of measurement errors, and the constant ς allows us to scale the measurement error covariance.

The small-scale New Keynesian DSGE model is typically estimated using output growth $y^o_{gr,t}$, inflation π^o_t, and interest rates R^o_t. Starting from $s_t = \left[\hat{y}_t, \hat{\pi}_t, \hat{R}_t, \hat{z}_t, \hat{d}_t, \hat{g}_t, e_{R,t}\right]'$, we define the augmented vector $\tilde{s}_t = [s_t', \hat{y}_{t-1}]'$ and add the trivial equation $\hat{y}_{t-1} = \hat{y}_{t-1}$ to the canonical form in (39). Because the $\hat{y}_{t-1}$ identity is linear, the structure of the canonical form is preserved. Assuming that output growth is measured in quarter-on-quarter percentages, and inflation and interest rates are measured in annualized percentages, we obtain

$$\begin{aligned} y^o_{gr,t} &= 100\ln(\gamma) + 100(\hat{y}_t - \hat{y}_{t-1} + \hat{z}_t) + \sigma_{u,y}u_{y,t} \\ \pi^o_t &= 400\ln(\pi_*) + 400\hat{\pi}_t + \sigma_{u,\pi}u_{\pi,t} \\ R^o_t &= 400\ln(R_*) + 400\hat{R}_t + \sigma_{u,R}u_{R,t}. \end{aligned} \tag{41}$$

In addition, we will include a measure of the consumption-output ratio, constructed from data on government spending G_t. The consumption output ratio is $C_t/Y_t = 1 - G_t/Y_t$. We define cy^o_t as linearly detrended $100 \cdot \ln(1 - G_t/Y_t)$. Because in our model $1 - G_t/Y_t = 1/g_t$, we obtain the additional measurement equation

$$cy^o_t = -100\ln g_* - 100\hat{g}_t + \sigma_{u,c}u_{c,t}. \tag{42}$$

[8] Aruoba et al. (2020) estimate a structural VAR that takes the form of (39). Chen (2017) and Bianchi and Melosi (2017) use an exogenous regime-switching process to characterize the ELB dynamics. Such models can be solved using the tools proposed by Farmer et al. (2011). Benigno et al. (2016) endogenize the regime-switching probability in a model of financial crisis, but, unlike in our paper, the transition from one to the other regime remains partly decoupled from the realization of the fundamental shocks.

Thus, we are treating the exogenous process $\hat{g}_t$ as observed in the estimation. Because the law of motion of $\hat{g}_t$ is linear, the PLC structure of the empirical model is maintained.

5. PLC model solution

In this section, we describe how the free coefficients ϑ in the PLC decision rules, as defined in (33), are determined. We first discuss the equilibrium conditions that define the objective function that we will minimize to find the optimal ϑ coefficients. We then describe the choice of the solution grid, the integration method and the optimization.

Equilibrium conditions. More formally, let us rewrite the equilibrium conditions in (14) generically as

$$\mathcal{H}[f_0(\cdot), \mathbb{X}] = 0, \quad \forall\, \mathbb{X} \in \mathcal{X}, \tag{43}$$

where $f_0(\mathbb{X})$ corresponds to the optimal decision rules. In the context of the New Keynesian DSGE model, these decision rules are given by $\hat{y}(\cdot)$, $\hat{\pi}(\cdot)$, $\hat{c}(\cdot)$, and $\hat{R}(\cdot)$. Because $\hat{c}(\cdot)$ and $\hat{R}(\cdot)$ can be directly determined by plugging $\hat{y}(\cdot)$ and $\hat{\pi}(\cdot)$ into the equilibrium conditions, in our model $f_0(\cdot) = [\hat{y}(\cdot), \hat{\pi}(\cdot)]$. The precise definition of $\mathcal{H}[f_0(\cdot), \mathbb{X}]$ for our New Keynesian DSGE model is provided in the Online Appendix.

We approximate $f_0(\mathbb{X})$ by PLC decision rules $g(\mathbb{X}; \vartheta) \in \mathcal{G}$, where ϑ contains the free coefficients that are necessary to characterize the PLC function and $\mathcal{G}$ is the set of all PLC functions. To determine ϑ, we minimize the norm of the vector-valued function $\mathcal{H}[g(\mathbb{X}; \vartheta), \mathbb{X}]$ over a set of M grid points $\mathcal{S}$ obtained using a sparse Smolyak grid:

$$\vartheta = \arg\min_{\vartheta} \frac{1}{M} \sum_{\mathbb{X} \in \mathcal{S}} \| \mathcal{H}[g(\mathbb{X}; \vartheta), \mathbb{X}; \theta] \|^2.$$

In the simplified DSGE model in Section 3 we were able to find two sets of decision rule coefficients ϑ that set $\mathcal{H}[\cdot]$ exactly equal to zero for all $\mathbb{X}$ and generate what we called the targeted-inflation equilibrium and the deflation equilibrium. In the numerical illustrations of Section 7 and the empirical application of Section 8 we will focus on the targeted-inflation equilibrium by choosing starting values for the ϑ optimization that generate the decision rules for a linearized version of the model without ELB constraint.[9]

Solution grid. There are two popular ways to choose the solution grid. In the collocation approach, the grid points typically come from a grid that is constructed using a tensor product of grids for each state variable, which in turn are constructed using the roots of a set of complete polynomials. It is well known that tensor product grids used to approximate the solution of nonlinear models suffer from the curse of dimensionality. Maliar and Maliar (2014, 2015) propose a series of techniques based on stochastic simulations to construct lower dimensional grids that represent the ergodic distribution of the model. However, these simulation-based methods require a time-consuming iterative procedure, and, in general, there does not seem to be a guarantee for the convergence of the grid and the approximate solution.

For our application, where we need to solve the model with different parameters tens of thousands of times, neither the collocation approach that uses tensor grids, nor the iterative approach that uses the ergodic distribution seem feasible. Coleman et al. (2018) propose the use of random and quasi-random grids on a fixed hypercube, because they are easier and faster to construct but lack the dimensionality reduction. Smolyak grids (Krueger and Kubler (2004), Malin et al. (2011), Judd et al. (2014)) offer a balance in this trade-off, combining the advantages of a fixed and predetermined domain and the dimensionality reduction of sparse grid methods.

In constructing the grid $\mathcal{S}$, we follow Judd et al. (2014) and build a sparse Smolyak grid.[10] The Smolyak grid is a sparse grid defined on the interval $[-1, 1]$. To use it in an application, it has to be scaled so that it represents the space of $\mathbb{X}$. The scaling of the grid amounts to picking minimum and maximum values for each state variable. The extrema correspond to -1 and 1 in the original domain of the Smolyak grid, respectively. One of the properties of the Smolyak grid is it places grid points at the edges of the domain – at -1 and 1. Thus, we recommend picking values for the scaling that are not too extreme in order to have some mass on both sides of the grid point.

In the context of the New Keynesian DSGE model, we proceed as follows. For the exogenous state variables in $\mathbb{X}$, except for the discount factor shock d_t, we linearly scale the grid so that it starts from the 10th percentile and goes to the 90th percentile of the distribution of each state variable. In the d_t dimension, we use an asymmetric grid where the grid is centered at the unconditional mean of zero, and the upper bound of the grid is at the 10th percentile. Unlike for the other shocks, we set the lower bound to -4 standard deviations below the mean. This asymmetry ensures that our solution is accurate in the range of values for the d_t shock that will be necessary to match the U.S. data in our empirical application.[11]

[9] While some progress has been made in Ascari and Mavroeidis (2020) studying parameter and innovation domain restrictions that guarantee coherency and completeness, formal results for general DSGE models with endogenous state variables and continuously distributed innovations remain elusive.

[10] One interpretation of our approach is that we are using the sum squared residual over the Smolyak grid as a proxy for integrating the squared residual function over the ergodic distribution. Monte Carlo experiments in Heiss and Winschel (2008) in the context of the calculation of the likelihood function of a mixed logit model, which also involves evaluating an integral without a closed-form expression, show that using a Smolyak grid provides superior performance over simulation techniques.

[11] Using the same symmetric grid for d_t leads to essentially no ZLB in the ergodic distribution of the model, which is a very well-known result for these models. See, for example, Boehl and Strobel (2020).

For the endogenous state variable $\hat{y}_{t-1}$, we use the same scaling as the exogenous state $\hat{z}_t$, since we have verified that they have similar dispersion when simulating the model. Finally, for $\hat{R}_{t-1}$, we use the observed nominal interest rate data. Because we want to analyze the ELB, the grid is scaled so that its minimum value matches the ELB with $\hat{R}_{t-1} = -\ln(r_*\pi_*)$, which happens to be the 10th percentile in the data. The maximum value is matched to the 90th percentile of the nominal interest rate in the data. For this variable in particular, we scale the grid so that the middle of the Smolyak grid coincides with the steady state at $\hat{R}_{t-1} = 0$.

Integration and minimization. Expectations in the residual functions as in (14) are computed using the monomial integration rule M2 as in Judd et al. (2010). For a generic expression $\mathbb{E}_t[v(x_{t+1})]$, our implementation with four random variables requires computing $v(\cdot)$ at 33 nodes and taking a weighted average. In our experience, this method produces results that are very similar to using a Gauss-Hermite integration for each random variable. As an example, with 5 nodes per random variable, the latter approach would make it necessary to evaluate $v(\cdot)$ at 625 nodes and increase running time considerably.

To minimize the objective function, we utilize the Levenberg–Marquardt algorithm, which is particularly effective in solving nonlinear least squares problems like ours, with the Jacobians evaluated analytically. As an initial guess for the algorithm, we use the decision rules from a log-linear approximation. Because the log-linear decision rules are a special case of the PLC decision rules—recall that we defined the model variables in log-deviations from the steady state—we can denote them by $g^{(0)}(\mathbb{X})$. We find the free coefficients ϑ_0 (with $\alpha^i_{11} = \alpha^i_{21}$ and $\alpha^i_{12} = \alpha^i_{22}$ for all $i = 1, ..., k$) that generate the same decision rules and use this to initialize the minimization algorithm.

Interpretation of PLC decision rules. We offer two interpretations of the PLC decision rules. First, they can be viewed as an approximation to the optimal decision rules $f_0(\mathbb{X})$. In fact, our motivation for constructing PLC rules was that the decision rules computed in Aruoba et al. (2018) with Chebyshev polynomials for a New Keynesian DSGE model that is essentially identical to the model in Section 2, appeared to be almost piecewise-linear.

Second, the PLC rules can be viewed as describing the behavior of boundedly-rational agents. In principle, bounded rationality can take many forms. The basic notion is that decision making is constrained by agents' abilities to gather, retain, and process decision-relevant information. Boundedly-rational agents may also be unable to solve a complicated mathematical problem. Under this interpretation, the PLC rules can be viewed as more easily computable decisions that have the additional benefit of being linear, except when the constraint in the model becomes binding.

6. Particle filters for PLC models

The state-space representation associated with the PLC solution comprises the nonlinear transition equation (39) and the linear measurement equation (40). The state-space representation provides a joint density for the states s_t and the observations y_t (omitting the o superscript):

$$p(Y_{1:T}, S_{1:T}|\theta) = \prod_{t=1}^{T} p(y_t|s_t, \theta)p(s_t|s_{t-1}, \theta), \tag{44}$$

where $Y_{t_1:t_2}$ and $S_{t_1:t_2}$ denote the sequences $y_{t_1}, \ldots, y_{t_2}$ and $s_{t_1}, \ldots, s_{t_2}$ and θ is the vector of model parameters. Of particular interest are the sequence of estimates $p(s_t|Y_{1:t})$ of the state vector and the likelihood function, which is defined as

$$p(Y_{1:T}|\theta) = \prod_{t=1}^{T} p(y_t|Y_{1:t-1}, \theta) = \prod_{t=1}^{T} \int\int p(y_t|s_t, \theta)p(s_t|s_{t-1}, \theta)p(s_{t-1}|Y_{1:t-1}, \theta)ds_t ds_{t-1}. \tag{45}$$

These objects can be obtained from a nonlinear filter. We describe below how $p(s_t|Y_{1:t}, \theta)$ and $p(y_t|Y_{1:t-1}, \theta)$ can be efficiently approximated by a particle filter.[12]

6.1. Generic particle filter

A particle filter represents the density $p(s_t|Y_{1:t}, \theta)$ through a swarm of particles $\{s^j_t, W^j_t\}^M_{j=1}$. Following the notation in Herbst and Schorfheide (2015) we now use $h(\cdot)$ to denote a function of s_t for which an expected value is supposed to be computed. $\mathbb{E}[h(s_t)|Y_{1:t}, \theta]$ can be approximated by Monte Carlo averages of the form $\frac{1}{M}\sum_{j=1}^{M} h(s^j_t)W^j_t$. The approximation typically holds in the form of a Law of Large Numbers and a Central Limit Theorem. The particle filter can be implemented using the following algorithm[13]:

[12] Surveys and tutorials can be found, for instance, in Arulampalam et al. (2002), Cappé et al. (2007), Doucet and Johansen (2011), and Creal (2012). Kantas et al. (2014) discuss using particle filters in the context of estimating the parameters of state-space models. Textbook treatments of the statistical theory underlying particle filters can be found in Cappé et al. (2005), Liu (2001), and Del Moral (2013).

[13] Exposition and notation are based on Herbst and Schorfheide (2015).

Algorithm 1 *(Generic particle filter).*

1. **Initialization.** Draw the initial particles from the distribution $s_0^j \overset{iid}{\sim} p(s_0|\theta)$ and set $W_0^j = 1$, $j = 1, \ldots, M$.
2. **Recursion.** For $t = 1, \ldots, T$:

   (a) **Forecasting** s_t**.** Draw $\tilde{s}_t^j$ from density $g_t(\tilde{s}_t|s_{t-1}^j, \theta)$ and define the importance weights

   $$\omega_t^j = p(\tilde{s}_t^j|s_{t-1}^j, \theta) \big/ g_t(\tilde{s}_t^j|s_{t-1}^j, \theta). \tag{46}$$

   (b) **Forecasting** y_t**.** Define the incremental weights

   $$\tilde{w}_t^j = p(y_t|\tilde{s}_t^j, \theta)\omega_t^j. \tag{47}$$

   The predictive density $p(y_t|Y_{1:t-1}, \theta)$ can be approximated by

   $$\hat{p}(y_t|Y_{1:t-1}, \theta) = \frac{1}{M}\sum_{j=1}^{M} \tilde{w}_t^j W_{t-1}^j. \tag{48}$$

   (c) Define the normalized weights

   $$\tilde{W}_t^j = \tilde{w}_t^j W_{t-1}^j \Big/ \frac{1}{M}\sum_{j=1}^{M} \tilde{w}_t^j W_{t-1}^j. \tag{49}$$

   (d) **Selection.** Resample the particles, for instance, via multinomial resampling. Let $\{s_t^j\}_{j=1}^M$ denote M iid draws from a multinomial distribution characterized by support points and weights $\{\tilde{s}_t^j, \tilde{W}_t^j\}$ and set $W_t^j = 1$ for $j =, 1 \ldots, M$. An approximation of $\mathbb{E}[h(s_t)|Y_{1:t}, \theta]$ is given by $\bar{h}_{t,M} = \frac{1}{M}\sum_{j=1}^{M} h(s_t^j) W_t^j$.

3. **Likelihood Approximation.** The approximation of the log-likelihood function is given by

$$\ln \hat{p}(Y_{1:T}|\theta) = \sum_{t=1}^{T} \ln\left(\frac{1}{M}\sum_{j=1}^{M} \tilde{w}_t^j W_{t-1}^j\right). \tag{50}$$

The most important choice in the configuration of the algorithm is the proposal density $g_t(\tilde{s}_t|s_{t-1}^j, \theta)$. Different choices of the proposal density lead to different versions of the particle filter.

6.2. Bootstrap particle filter

The BSPF was originally proposed by Gordon and Salmond (1993). It uses the state-transition equation as the proposal density, that is, $g_t(\tilde{s}_t|s_{t-1}^j, \theta) = p(\tilde{s}_t|s_{t-1}^j, \theta)$. This choice is attractive because it is straightforward to implement the forecasting step by forward simulation of the transition equation and the importance weights simplify to $\omega_t^j = 1$. A well-known disadvantage is that the proposal distribution is blind and hence ignores information about s_t contained in the current observation y_t. This can lead to a large variance of the incremental weights $\tilde{\omega}_t^j$. This problem is exacerbated if the measurement error variance is small and $p(y_t|\tilde{s}_t^j)$ has thin tails, or if the model is inappropriately parameterized or misspecified and therefore has difficulties predicting y_t one step ahead. Because the BSPF has been used in the DSGE model literature (see Fernández-Villaverde and Rubio-Ramírez (2007), An and Schorfheide (2007) and Herbst and Schorfheide (2015)), we will include it as a benchmark.

6.3. Conditionally optimal particle filter

The proposal density for the COPF utilizes information in y_t with the goal of minimizing the variance of the incremental weights $\tilde{\omega}_t^j$. It is given by

$$g_t^*(\tilde{s}_t|s_{t-1}^j, \theta) = p(\tilde{s}_t|y_t, s_{t-1}^j, \theta) \propto p(y_t|\tilde{s}_t, \theta)p(\tilde{s}_t|s_{t-1}^j, \theta). \tag{51}$$

Combining the formula for $g_t^*(\tilde{s}_t|s_{t-1}^j, \theta)$ with the expressions for the importance weights ω_t^j in (46) and the incremental weights $\tilde{\omega}_t^j$ in (47), we obtain

$$\tilde{\omega}_t^j = \frac{p(y_t|\tilde{s}_t^j,\theta)p(\tilde{s}_t^j|s_{t-1}^j,\theta)}{p(\tilde{s}_t^j|y_t,s_{t-1}^j,\theta)} = p(y_t|s_{t-1}^j). \tag{52}$$

The second equality follows from Bayes Theorem. It can be shown that conditional on $\{s_{t-1}^j\}$ the proposal density $g_t^*(\tilde{s}_t|s_{t-1}^j,\theta)$ minimizes the variance of the incremental weight $\tilde{w}_t^j$ in (47).

While direct sampling from the conditionally optimal proposal density is elusive for most nonlinear state-space models, we can derive a convenient formula for the piecewise-linear state-transition equation (39). Note that, conditional on s_{t-1}, the current state s_t is determined by η_t. It turns out, that it is more convenient to derive a conditionally optimal proposal density for η_t, denoted by $g_t^*(\eta_t|s_{t-1}^j,\theta)$.

In order to state the result, we have to define the following objects:

$$\begin{aligned}
\nu_t^j(\cdot) &= y_t - A_0 - A_s\big(\Phi_0(\cdot) - \Phi_1(\cdot)s_{t-1}^j\big) \\
\bar{\eta}_t^j(\cdot) &= \Big(\varsigma I + \Phi_\eta'(\cdot)A_s'\Sigma_u^{-1}A_s\Phi_\eta(\cdot)\Big)^{-1}\Phi_\eta'(\cdot)A_s'\Sigma_u^{-1}\nu_t^j(\cdot) \\
\bar{\Omega}(\cdot) &= \varsigma\Big(\varsigma I + \Phi_\eta'(\cdot)A_s'\Sigma_u^{-1}A_s\Phi_\eta(\cdot)\Big)^{-1}.
\end{aligned} \tag{53}$$

I is the identity matrix. Recall that ς scales the measurement error variance. We use the argument $(\cdot)$ to indicate that the expressions are obtained either based on $\big(\Phi_0(n),\Phi_1(n),\Phi_\eta(n)\big)$ or $\big(\Phi_0(b),\Phi_1(b),\Phi_\eta(b)\big)$. Here $\nu_t^j(\cdot)$ is the error made in forecasting y_t based on s_{t-1}^j. $\bar{\eta}_t^j(\cdot)$, and $\bar{\Omega}(\cdot)$ are the posterior mean vector and covariance matrix of $\eta_t|(y_t,s_{t-1}^j)$ absent any truncation—that is, for $\zeta(s_{t-1}^j)$ being $+\infty$ or $-\infty$. Moreover, let

$$\begin{aligned}
D_t^j(n) &= (2\pi)^{-n_y/2}|\Sigma_u|^{-1/2}|\varsigma I + \Phi_\eta(n)'A_s'\Sigma_u^{-1}A_s\Phi_\eta(n)|^{1/2} \\
&\quad\times\exp\left\{-\frac{1}{2}\nu_t^j(n)'(\varsigma\Sigma_u + A_s\Phi_\eta(n)\Phi_\eta'(n)A_s')^{-1}\nu_t^j(n)\right\} \\
&\quad\times\Phi_N\big((\zeta(s_{t-1}) - \bar{\eta}_{1,t}^j(n)/\sqrt{\bar{\Omega}_{11}(n)}\big), \\
D_t^j(b) &= (2\pi)^{-n_y/2}|\Sigma_u|^{-1/2}|\varsigma I + \Phi_\eta(b)'A_s'\Sigma_u^{-1}A_s\Phi_\eta(b)|^{1/2} \\
&\quad\times\exp\left\{-\frac{1}{2}\nu_t^j(b)'(\varsigma\Sigma_u + A_s\Phi_\eta(b)\Phi_\eta'(b)A_s')^{-1}\nu_t^j(b)\right\} \\
&\quad\left(1 - \Phi_N\left((\zeta(s_{t-1}) - \bar{\eta}_{1,t}^j(b))/\sqrt{\bar{\Omega}_{11}(b)}\right)\right).
\end{aligned} \tag{54}$$

It can be shown that $p(y_t|s_{t-1}^j) = D_t^j(n) + D_t^j(b)$.

The characterization of the conditionally optimal proposal density is summarized in Proposition 1. A proof of the proposition is provided in the Online Appendix.

Proposition 1. *Suppose the state-transition equation is given by (39), $\eta_t \sim N(0,I)$, $\eta_{1,t}$ is a scalar, and the measurement equation is given by (40). Draws from the conditionally optimal proposal densities $g_t^*(\tilde{s}_t|s_{t-1}^j,\theta)$, $j=1,\ldots,M$, defined in (51) can be generated as follows:*

1. *Let*

$$\xi_t^j = \begin{cases} \text{'n'} & \text{with prob. } \lambda_t^j \\ \text{'b'} & \text{with prob. } 1-\lambda_t^j \end{cases}, \quad \text{where} \quad \lambda_t^j = \frac{D_t^j(n)}{D_t^j(n)+D_t^j(b)}.$$

2. *If $\xi_t^j =$ 'n' then generate η_t from the distribution*

$$\eta_{1,t}^j \sim N\big(\bar{\eta}_{1,t}^j(n),\bar{\Omega}_{11}(n)\big)\mathbb{I}\{\eta_{1,t}^j \le \zeta(s_{t-1}^j)\}, \quad \eta_{2,t}^j|\eta_{1,t}^j \sim N(\bar{\eta}_{2|1}^j(n,\eta_{1,t}^j),\bar{\Omega}_{2|1}(n)) \tag{55}$$

and let

$$\tilde{s}_t^j = \Phi_0(n) + \Phi_1(n)s_{t-1}^j + \Phi_\eta(n)\eta_t^j.$$

If $\xi_t^j =$ 'b' then generate η_t^j from the distribution

$$\eta_{1,t}^j \sim N\big(\bar{\eta}_1^j(b),\bar{\Omega}_{11}(b)\big)\mathbb{I}\{\eta_{1,t}^j > \zeta(s_{t-1}^j)\}, \quad \eta_{2,t}^j|\eta_{1,t}^j \sim N\big(\bar{\eta}_{2|1}^j(b,\eta_{1,t}^j),\bar{\Omega}_{2|1}(b)\big) \tag{56}$$

and

$$\tilde{s}_t^j = \Phi_0(b) + \Phi_1(b)s_{t-1}^j + \Phi_\eta(b)\eta_t^j.$$

3. *The incremental particle weight is* $\tilde{\omega}_t^j = D(n) + D(b)$. □

Vanishing measurement errors. It is instructive to examine what happens as $\varsigma \longrightarrow 0$. For the conditional density of $y_t|s_{t-1}$ to be nonsingular in the limit, it has to be the case that the number of rotated structural innovations is at least n_y. Formally, the covariance matrices $A_s\Phi_\eta(\cdot)\Phi_\eta'(\cdot)A_s'$ have to be non-singular. Suppose that $n_y = n_\eta$ and $(A_s\Phi_\eta(\cdot))$ are invertible $n_y \times n_y$ matrices. This means, ignoring the truncation, under the invertibility assumption, we can solve for the innovations η_t as a function of (y_t, s_{t-1}^j):

$$\eta_{t*}^j(\cdot) = (A_s\Phi_\eta(\cdot))^{-1}(y_t - A_s(\Phi_0(\cdot) + \Phi_1(\cdot)s_{t-1})).$$

Now consider what happens if we let the measurement error variance converge to zero. First, the expressions in (53) remain well defined in the limit as $\varsigma \longrightarrow 0$:

$$\bar{\eta}_t^j \longrightarrow \eta_{t*}^j(\cdot), \quad \bar{\Omega}(\cdot) \longrightarrow 0.$$

The posterior variance converges to zero and the posterior mean converges to the innovation $\eta_{t*}^j(\cdot)$ that generates the observed y_t conditional on s_{t-1}^j. For the limit behavior on $D_t^j(n)$, the crucial term is

$$\lim_{\varsigma \longrightarrow 0} \Phi_N\left((\zeta(s_{t-1}^j) - \bar{\eta}_{1,t}^j(n))/\sqrt{\bar{\Omega}_{11}(n)}\right) = \begin{cases} 1 & \text{if } \zeta(s_{t-1}^j) - \bar{\eta}_{1,t}^j(n) \geq 0 \\ 0 & \text{otherwise} \end{cases}.$$

This term measures whether it is possible to explain y_t using the (n) coefficients, accounting for the fact that the n regime is only active if $\bar{\eta}_{1,t}^j(n) \leq \zeta(s_{t-1}^j)$. A similar analysis can be conducted for the term $D_t^j(b)$. Thus, for each particle j, there are four possible cases (ignoring equalities):

$$\begin{aligned}
&\text{Case 1}: \quad \bar{\eta}_{1,t}^j(n) < \zeta(s_{t-1}^j), \quad \bar{\eta}_{1,t}^j(b) < \zeta(s_{t-1}^j) \\
&\text{Case 2}: \quad \bar{\eta}_{1,t}^j(n) > \zeta(s_{t-1}^j), \quad \bar{\eta}_{1,t}^j(b) > \zeta(s_{t-1}^j) \\
&\text{Case 3}: \quad \bar{\eta}_{1,t}^j(n) < \zeta(s_{t-1}^j), \quad \bar{\eta}_{1,t}^j(b) > \zeta(s_{t-1}^j) \\
&\text{Case 4}: \quad \bar{\eta}_{1,t}^j(n) > \zeta(s_{t-1}^j), \quad \bar{\eta}_{1,t}^j(b) < \zeta(s_{t-1}^j).
\end{aligned}$$

In Case 1 $D_t^j(b) = 0$ and $\lambda_t^j = 1$. Here, only the (n) decision rules can rationalize the data conditional on s_{t-1}^j. Case 2 is the opposite: $D_t^j(n) = 0$, $\lambda_t^j = 0$, and only the (b) decision rules can rationalize the data. Under Case 3, both $D_t^j(n)$ and $D_t^j(b)$ are strictly positive, $0 < \lambda_t^j < 1$, and both decision rules could explain the data. Finally, in Case 4 y_t is inconsistent with s_{t-1}^j, and none of the decision rules can explain the data. If each $j = 1, \ldots, M$ falls into Case 4, then the particle-filter based likelihood approximation for this particular parameterization of the DSGE model will be zero. Note that, if the measurement error variance is strictly greater than zero, (potentially very large) measurement errors could also rationalize the data under Case 4.[14]

The previous calculations highlight that, unlike for the BSPF, the weights of the COPF do not degenerate if one decreases the measurement error variance. In this case, if $A_s\Phi_\eta(\cdot)$ is a square matrix, the COPF specializes to the inversion filter that solves for the innovations as a function of y_t and s_{t-1}^j. Because our model is piecewise-linear, this inversion may have one, two, or no solution(s).

Perfectly observed regimes. Our ELB application has the special feature that the observation y_t identifies the regime. Let $y_t = [y_{1,t}', y_{2,t}]'$, where $y_{2,t}$ corresponds to the nominal interest rate. Suppose the ELB is binding in the b regime and non-binding in the n regime. Then, the ex post regime probability λ^j is independent of s_{t-1}^j and given by the indicator function.

$$\lambda^j = \mathbb{I}\{y_{2,t} > c\}. \tag{57}$$

While the distribution of y_t is continuous in the n regime, for the binding regime the continuous part of the distribution concentrates in the lower-dimensional subspace defined by $y_{2,t} = c$. Thus, the formulas for $D_t^j(b)$ in (54) and the moments of the distribution of η_t^j in Proposition 1 in the b regime have to be adjusted to account for the reduced dimensionality of the continuous part of the y_t distribution. Further details are provided in the Online Appendix.

[14] The four cases distinguished here are closely connected to the coherency and completeness conditions discussed in Section 3. Previously, we asked whether conditional on a realization of η_t the vector s_t, using the notation of the canonical form (39), is uniquely determined. If the condition is not satisfied, then $p(s_t|s_{t-1}, \theta)$ is not well defined. In addition, for the filtering it matters whether in the absence of measurement errors, there exist one or more η_t's that can rationalize the data.

Table 1
DGP and prior.

Parameter	DGP	Prior distribution				
		Density	P(1)	P(2)	HPD Low	HPD High
τ	2.00	$\mathcal{G}$	2.00	0.20	1.67	2.32
κ	0.13	$\mathcal{B}$	0.10	0.05	0.02	0.17
ψ_1	2.60			fixed at 2.60		
ψ_2	0.98			fixed at 0.98		
ρ_R	0.80	$\mathcal{B}$	0.80	0.10	0.65	0.96
ρ_g	0.91	$\mathcal{B}$	0.80	0.10	0.65	0.96
ρ_d	0.97	$\mathcal{B}$	0.80	0.10	0.65	0.96
ρ_z	0.37	$\mathcal{B}$	0.40	0.20	0.08	0.73
σ_R	0.0019	$\sqrt{\mathcal{IG}}$	0.005	4.00	0.003	0.010
σ_g	0.0025	$\sqrt{\mathcal{IG}}$	0.005	4.00	0.003	0.010
σ_d	0.020	$\sqrt{\mathcal{IG}}$	0.01	4.00	0.005	0.020
σ_z	0.0058	$\sqrt{\mathcal{IG}}$	0.01	4.00	0.005	0.020
η	0.72			fixed at 0.72		
ν	0.10			fixed at 0.10		
χ_H	1.00			fixed at 1.00		
g^*	1.27	$\mathcal{G}$	1.20	0.20	0.88	1.63
rAnet	0.22	$\mathcal{G}$	1.00	0.40	0.39	1.64
gamQnet	0.33	$\mathcal{N}$	0.50	0.25	0.09	0.91
piAnet	0.75	$\mathcal{N}$	2.50	1.00	0.87	4.14

Notes: $\mathcal{G}$ is Gamma distribution; $\mathcal{B}$ is Beta distribution; $\mathcal{IG}$ is Inverse Gamma distribution; and $\mathcal{N}$ is Normal distribution. P(1) and P(2) are mean and standard deviations for Beta, Gamma, and Normal distributions. The $\mathcal{IG}$ distribution is parameterized as scaled inverse χ^2 distribution with density $p(\sigma^2|s^2,\tilde{\nu}) \propto (\sigma^2)^{-\tilde{\nu}/2-1}\exp[-\tilde{\nu}s^2/(2\sigma^2)]$, where P(1) is $\sqrt{s^2}$ and P(2) is $\tilde{\nu}$. The density of σ is obtained by the change of variables $\sigma = \sqrt{\sigma^2}$. HPD(Low, High) refers to the boundaries of 90% highest prior density intervals. We use the following parameter transformations: $\beta = \exp\{-\text{rANet}/400\}$, $\gamma = \exp\{\text{gamQnet}/100\}$, and $\pi_* = \exp\{\text{piAnet}/400\}$.

7. Numerical illustrations

We now illustrate the proposed filtering method based on data simulated from the DSGE model of Section 2. First, we compare the distribution of the stochastic likelihood approximation conditional on a particular parameter θ between the proposed COPF and the basic BSPF. It has been shown in the literature that likelihood approximations of particle filters are unbiased; see Herbst and Schorfheide (2015). In view of the unbiasedness, a low-variance approximation is preferable to a high-variance approximation. Jensen's inequality implies that *log* likelihood approximations are downward biased. The magnitude of the bias is connected to the variability of the likelihood approximation: the larger the variance, the larger also the downward bias.

Second, we embed the particle filter approximation of the likelihood function into a Metropolis-Hastings algorithm. It has also been shown in the literature that if the exact likelihood function is replaced by a noisy but unbiased estimate, MCMC algorithms still converge to the exact posterior distribution. However, there is no free lunch: the larger the variability of the likelihood approximation, the slower the convergence. In the numerical illustrations below, we will compare the speed of convergence as measured by the degree of serial correlation in the parameter draws, across different versions of the particle filter.

Throughout this section, we assume that all variables, including interest rates, are measured with error and that therefore regime (ELB binding versus non-binding) is not perfectly observed. Thus, conditional on the states s_t, y_t has a continuous distribution; see (40) and (41). All computations reported below are executed on a single core of PC with an Intel Xeon CPU E5-2687Wv3 at 3.10 GHz running Windows 10 (64bit) and JuliaPro 1.5.3.

7.1. Accuracy of likelihood approximation

We simulate a sample of $T = 140$ observations from the DSGE model, loosely parameterized based on the empirical estimates reported in Section 8. The parameter values are summarized in the second column of Table 1. The sample size matches the one used in the empirical application in Section 8. We select a subsample from a long simulation in which the ELB binds approximately 10% of the periods.

We begin by comparing the accuracy of the particle-filter-based likelihood approximation using the same parameter values that were used to generate the data. The likelihood evaluation is conditional on a vector of initial states which are

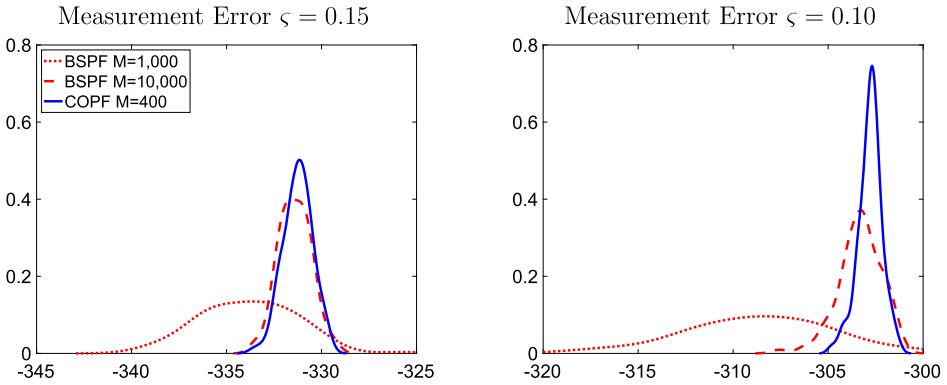

Notes: Density plots are based on $N_{run} = 100$ runs of the BSPF ($M = 1,000$ is dotted and $M = 10,000$ is dashed) and COPF ($M = 400$, solid), respectively.

Fig. 1. Density of log-likelihood approximations.

the true initial states associated with the simulated observations. We consider two values for the scaling parameter for the measurement errors, ς: 0.15 and 0.10. The smaller ς is, the larger is the penalty for deviations between model-predicted and actual observations. Thus, small measurement errors can be provided as a stress test for the filter in case the model (or its parameterization) is at odds with the data.

The accuracy of the filter depends on the number of particles M. For the COPF, we set $M = 400$, and, for the BSPF, we consider $M = 1,000$ and $M = 10,000$. Starting from the benchmark of the BSPF with $M = 1,000$, the choice of $M = 400$ for the COPF approximately equalizes the run times of the two filters, which is about 0.9 seconds for $\varsigma = 0.15$.[15] We run the COPF and the BSPF $N_{run} = 100$ times, respectively, and construct kernel estimates of the likelihood from the repeated runs, which are depicted in Fig. 1. The more concentrated the densities, the more accurate the likelihood approximation. Holding the run time fixed, the COPF is substantially more accurate than the BSPF. Raising the number of particles from 1,000 to 10,000 increases the run time of the BSPF roughly tenfold. For $\varsigma = 0.15$ this leads to roughly the same accuracy as the COPF. In other words, for a pre-specified level of accuracy the COPF is ten times faster than the BSPF.

Because the BSPF ignores the information in y_t when generating the proposal draws $\tilde{s}_t^j$, the variance of the particle weights increases as ς falls. This, in turn, translates into an increase in the variance (and by virtue of the concave transformation also the bias) of the log-likelihood approximation, which is clearly visible by comparing the dotted densities across panels for the two different values of ς. The precision of the COPF, on the other hand, increases as ς falls. When $\varsigma \longrightarrow 0$, it is possible to uniquely determine the innovations η_t^j conditional on s_{t-1}^j during non-ELB periods, because the number of observables equals the number of shocks.

Fig. 2 depicts standard deviations of log-likelihood approximations as a function of the mean log-likelihood value across $N_{run} = 100$ runs of the filter for $\varsigma = 0.15$ and $\varsigma = 0.10$. Each dot (or asterisk) in the two scatter plots corresponds to a different parameter value θ^i.[16] Two important findings emerge. First, as we have seen already from Fig. 1, the COPF likelihood approximation is less dispersed than the BSPF approximation. The accuracy gain from the conditionally optimal proposal density increases as the measurement error variance decreases, because the BSPF performance deteriorates. Second, while the accuracy of the COPF is independent of the log-likelihood value associated with the posterior draw θ^i, the accuracy of the BSPF approximation deteriorates the further the θ^i draw is in the tails of the posterior distribution.

7.2. Particle MCMC

We now embed the particle filter likelihood approximations in a standard single-block RWMH algorithm. The particle RWMH algorithm operates on an enlarged probability space that includes all the random variables that are generated when running the particle filter; see Herbst and Schorfheide (2015) for a textbook exposition and Andrieu et al. (2010) for a formal analysis. The use of an enlarged probability space leads to an increase in the persistence of the Markov chain generated by the posterior sampler. The noisier the likelihood approximation, the larger the persistence of the resulting Markov chain. A high degree of serial correlation is undesirable, because it leads to very noisy Monte Carlo approximations of posterior moments. We will show that the use of an accurate particle filter such as the COPF can alleviate this problem.

[15] The absolute run times depend on the programming language – Julia in our case – and the efficiency of the code. We are mainly interested in the relative run times. The code for the COPF and BSPF is almost identical, except that the COPF requires a few additional steps to evaluate the expressions in Proposition 1, which implies that for the same M, the COPF is slower than the BSPF.

[16] The parameter values are obtained from a preliminary run of the particle Markov chain Monte Carlo (PMCMC) algorithm in Section 7.2.

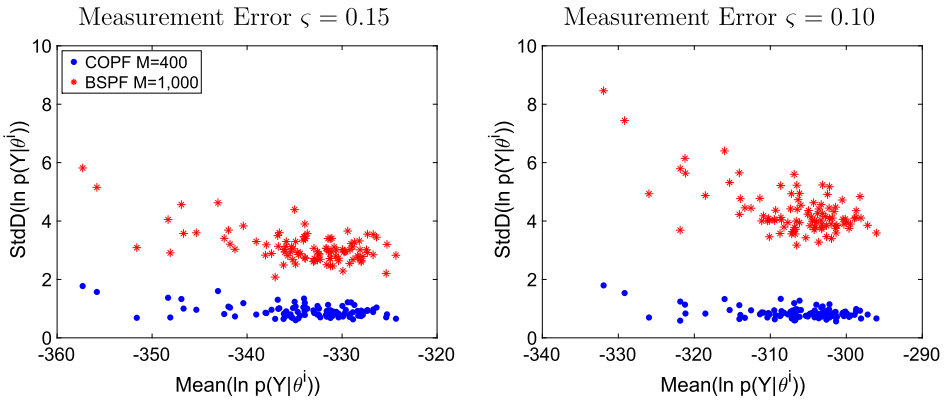

Notes: Standard deviations of log-likelihood approximations are based on $N_{run} = 100$ runs of the two filters. Each dot (or asterisk) corresponds to a particular θ^i. We are using $M = 1,000$ particles for the BSPF (asterisks) and $M = 400$ particles for the COPF (dots).

Fig. 2. Comparison of log-likelihood approximations.

Under some regularity conditions, the sequence of posterior draws generated from a RWMH algorithm satisfies a Central Limit Theorem (CLT) for dependent processes. The numerical accuracy of the Monte Carlo approximation of posterior means depends on the long-run covariance matrix of the sequence of parameter draws θ^i, $i = 1, \ldots, N$. The larger the autocorrelation of these draws, the less precise the Monte Carlo approximation.

The RWMH algorithm requires a covariance matrix for the proposal distribution (we use a multivariate normal distribution) that is constructed as follows. We start from a log-linearized version of the DSGE model that ignores the ELB constraint and sets the measurement error variance to zero ($\varsigma = 0$). In this case the likelihood function can be evaluated with the Kalman filter (KF). We conduct two preliminary MCMC runs using the linearized model. The first run is based on a diagonal covariance matrix with scaled prior variances on the diagonal and the second run is based on the posterior covariance matrix of the first run. Finally, we compute the posterior covariance matrix from the second run and denote it by $\hat{\bar{V}}_\theta$. The particle MCMC runs using the nonlinear DSGE model are based on $c\hat{\bar{V}}_\theta$ where $c = 0.1$.

For the estimation of the nonlinear model we set the scale of the measurement error variance to $\varsigma = 0.15$ and keep the number of particles at $M = 400$ for the COPF and $M = 1,000$ for the BSPF. Thus, according to the left panel of Fig. 2, the standard deviation of the log-likelihood approximation of the COPF is around 1 whereas the standard deviation for the BSPF ranges from 3 to 5. We generate 22,000 draws from the posterior distribution and discard the first 2,000 draws. The overall run time using the two filters is approximately 5 hours and 45 minutes for both filters. Note that the codes for the COPF-RWMH and the BSPF-RWMH algorithms are identical, except for the more complicated way of generating draws from $g_t(\tilde{s}_t|s_{t-1}^j, \theta)$ and evaluating $\tilde{\omega}_t^j$ in the COPF likelihood approximation.

Using a scaling of $c = 0.1$, the acceptance rate for the proposed draws is 32% for the COPF-RWMH algorithm. Herbst and Schorfheide (2015) documented that for the estimation of small-scale linearized DSGE models, an acceptance rate between 15% and 30% is associated with the most accurate Monte Carlo approximations of posterior means. Replacing the COPF by the BSPF while keeping the scale factor $c = 0.1$, the acceptance rate drops to 4.3%. The lower acceptance rate is caused by the noisier likelihood approximation.[17]

The first row of Fig. 3 compares posterior densities constructed from the output of the COPF- and BSPF-RWHM algorithms for two representative parameters: τ and ρ_d. The posterior densities for τ look similar and both peak near the "true" parameter values depicted by the solid vertical line. For the parameter ρ_d, the posterior density obtained from the BSPF is slightly to the left of the COPF density and its peak is further away from the "true" value. Apparently the BSPF chain has not fully converged yet. It moves slowly due to the low acceptance rate and needs more draws to fully explore the high-density area of the posterior.

The second row of Fig. 3 shows the autocorrelation functions for the τ^i and ρ_d^i sequences. Here, stark differences emerge. While under the BSPF-based sampler the autocorrelation at lag 100 is still around 0.8 for the ρ_d^i sequence, it is only 0.28 for draws from the COPF-RWMH. Suppose the autocorrelation function of the draws is given by ρ^j, where j is the temporal displacement among the draws. Then the variance of the mean of the draws is given by $\mathbb{V}_\pi[h(\theta)](1+\rho)/(1-\rho)$, where $\mathbb{V}_\pi[h(\theta)]$ is the posterior variance of $h(\theta)$ and $(1+\rho)/(1-\rho)$ can be viewed as an inefficiency factor that arises due to the serial correlation of the draws. A reduction from $\rho = 0.8^{1/100} = 0.998$ to $\rho = 0.28^{1/100} = 0.987$ lowers the inefficiency factor from approximately 1,000 to 150. This implies only 15% of the draws are required to achieve the same accuracy of

[17] The acceptance rate could be increased by decreasing c, but it does not cure the persistence problem because the chain moves very slowly due to the smaller size of the accepted steps.

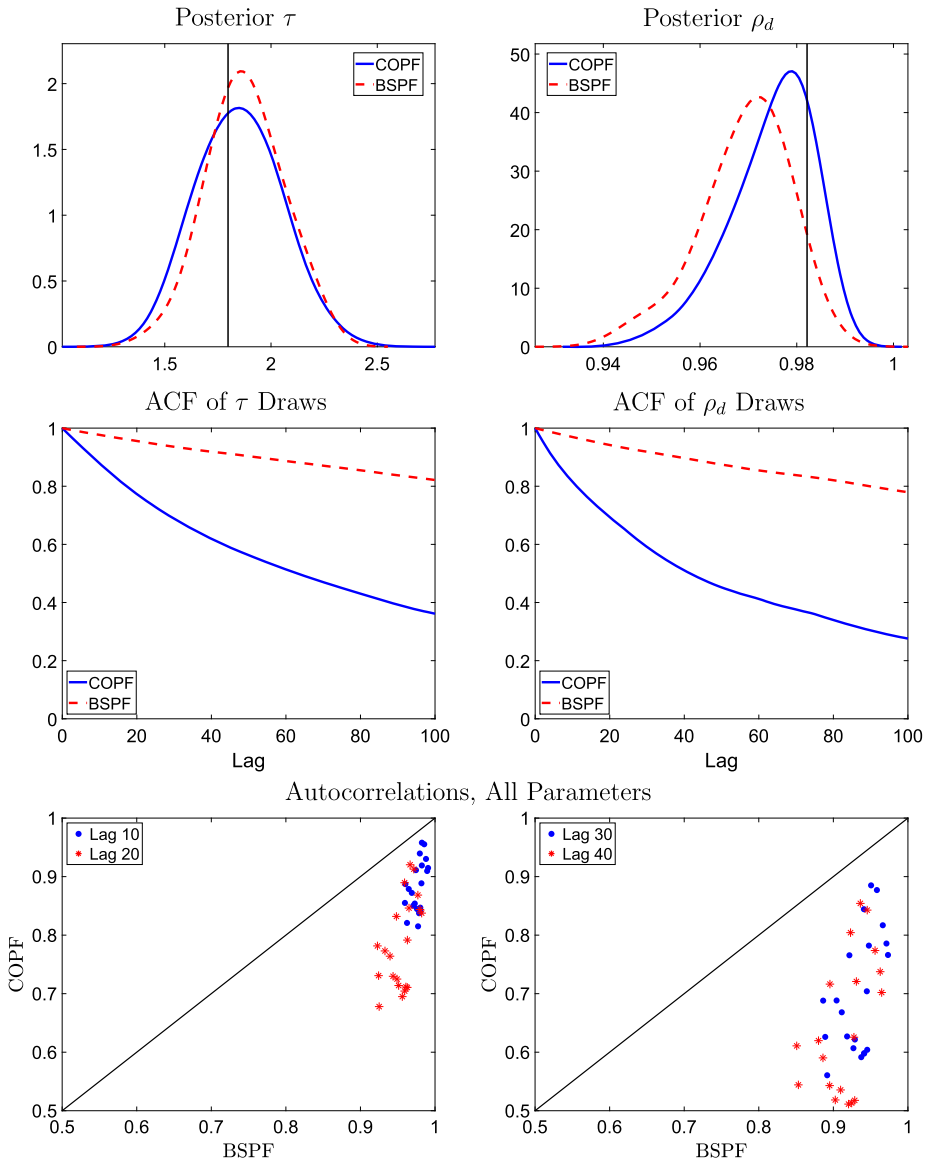

Notes: ME scale $\varsigma = 0.15$, proposal covariance scale $c = 0.1$, $N = 22,000$ draws (drop first 10%). COPF: number of particles $M = 400$. BSPF: number of particles $M = 1,000$. Top row: kernel density estimates of posterior distributions based on MCMC output. Vertical lines indicate true parameter values. Center row: autocorrelation functions of posterior draws based on COPF and BSPF. Bottom row: scatter plots of autocorrelations BSPF vs. COPF for various lags. Solid line is 45 degree line.

Fig. 3. Posterior draws: density and autocorrelation.

posterior mean approximations with the COPF-based sampler as with the BSPF-based sampler, or, holding the numbers of draws fixed across samplers, the COPF-based sampler delivers Monte Carlo approximations that are six times as accurate (in terms of sampling variance).

The last row of Fig. 3 compares autocorrelations for lags 10, 20, 30, and 40 for all estimated parameters. The solid lines are 45-degree lines. The two panels show that the COPF is able to reduce the autocorrelation for all estimated parameters, and hence it drastically improves the performance of the MCMC algorithm.

To summarize, in our likelihood approximation exercise we found that for a pre-specified level of accuracy, the COPF is at least ten times faster than the BSPF. When used inside a RWMH algorithm we showed that the COPF is able to reduce the persistence of the Markov chain considerably which leads to substantially more accurate Monte Carlo approximations of posterior moments. The COPF enables us to accurately estimate a DSGE model with an occasionally-binding constraint without supercomputing capabilities in a relatively short amount of time.

8. Empirical application

We now estimate the small-scale New Keynesian DSGE model based on quarterly U.S. data using the previously developed model solution and filtering techniques and conduct a fiscal policy experiment. The estimation results are summarized in Section 8.1, and the fiscal policy analysis appears in Section 8.2.

8.1. Estimation

The DSGE model is estimated based on data on quarterly GDP growth (quarterly %), the log consumption-GDP ratio (scaled by 100), GDP deflator inflation (annualized %), and nominal interest rates (annualized %) with data from 1984:Q1 to 2018:Q4. The data for the estimation was extracted from the FRB St. Louis FRED database (vintage 2019-10-30). Output growth is defined as real gross domestic product (GDPC1) growth converted into per capita terms. Our measure of population is Civilian Noninstitutional Population (CNP16OV). We compute population growth rates as log differences and apply an eight-quarter backward-looking moving average filter to the growth rates to smooth out abrupt changes in the population growth series. In constructing a measure of the consumption-GDP ratio, we define consumption as the difference between output and government spending. Thus, our measure of consumption includes investment and net exports. Government spending is constructed as real government consumption expenditures and gross investment (GCEC1). We remove a linear trend from the log consumption-GDP ratio to correct for different time trends in the price deflators of the GDP components. Inflation is defined as the log difference in the GDP deflator (GDPDEF), and the interest rate is the average effective federal funds rate (FEDFUNDS) within each quarter. During the period 2009:Q1 to 2015:Q4, when the effective federal funds rate was between 0 and 25 basis points, we set the interest rate exactly equal to zero and regard the ELB as binding.

The prior distribution used for the estimation is identical to the one in Table 1. We absorb the initial values of the latent state variables into the parameter vector and specify prior distributions over the initial states; see Table A-1 in the Online Appendix.[18] We fix a number of parameters prior to estimation. Because our sample does not include observations on labor market variables, we fix the Frisch labor supply elasticity. Based on Ríos-Rull et al. (2012), who provide a detailed discussion of parameter values that are appropriate for DSGE models of U.S. data, we set $\eta = 0.72$. The parameter ν, which captures the elasticity of substitution between intermediate goods, is not separately identifiable from the slope of the Phillips curve κ which in turn determines the adjustment cost parameter ϕ. We set $\nu = 0.1$, which generates a markup of 10%. We fix the preference parameter at $\chi_H = 1$. It determines steady-state hours worked and is neither relevant for the model dynamics nor identifiable based on our observables. We also fix the monetary policy coefficients ψ_1 and ψ_2 at 2.60 and 0.98, respectively, which are values estimated in Aruoba and Schorfheide (2016).

We start out by estimating the log-linearized version of the DSGE model that ignores the ELB constraint, setting the measurement error variances of the state-space model to zero. Draws from the posterior are generated by a single-block RWMH algorithm. In an initial run, we use a diagonal matrix with the prior variances to configure the covariance matrix of the proposal distribution. In the main run, we use the estimated posterior covariance matrix $\hat{\hat{V}}$ from the initial run to construct a proposal covariance matrix Σ with the scaling factor $c = 0.3$. We generate $N = 110,000$ draws, discarding the first 10,000.

For the estimation of the nonlinear version of the model we assume that interest rates are observed without error. Thus, conditional on R_t^o, it is known whether the ELB is binding or not; see (57). We maintain the assumption that the remaining variables are observed subject to a measurement error and set the scale factor for their measurement error variances to $\varsigma = 0.001$. We generate $N = 55,000$ draws from the posterior distribution of θ using the particle RWMH algorithm, discarding the first 5,000 draws. We use a scaling of $c = 0.1$ for the covariance matrix of the proposal distribution. The likelihood function is approximated using the COPF with $M = 250$ particles. The resulting acceptance rate of the particle RWMH algorithm is 17%, and the run time is approximately 18 hours on a single core, which comes to about 1.2 seconds per draw.[19]

The parameter estimates are summarized in Table 2. The table reports posterior means and lower and upper endpoints of highest posterior density (HPD).[20] The parameter estimates are similar to the ones reported elsewhere in the literature for variants of the small-scale New Keynesian DSGE model. The estimated slope of the Phillips curve is $\hat{\kappa} = 0.18$. The government spending shock is close to a unit-root process ($\hat{\rho}_g = 0.97$), and the estimated autoregressive parameter of the discount factor shock is $\hat{\rho}_d = 0.93$. Thus, innovations to these processes will have a long-lasting effect. One of the outputs of the estimation is the set of filtered values for the exogenous variables. We use these explicitly in our policy experiment and discuss them in the next section.

[18] This approach has the advantage that uncertainty about the initial state does not add to the variability of the particle-filter-based likelihood approximation conditional on a parameter θ^i.

[19] If we reduce the number of particles from $M = 250$ to $M = 150$ the run time falls to approximately 12 hours and 30 minutes, yielding very similar approximation of the posterior distribution.

[20] We compared these estimates with those obtained from a linearized model using the Kalman Filter and data that exclude the ELB episode. The most noteworthy difference is in σ_d, which needs to be larger when the ELB episode is used in order to deliver large (negative) and persistent shocks that take the economy to the ELB.

Table 2
Posterior distribution (PLC / COPF).

Parameter	Mean	HPD Low	HPD High
τ	1.98	1.74	2.24
κ	0.18	0.13	0.22
ρ_R	0.79	0.77	0.82
ρ_g	0.97	0.95	0.99
ρ_d	0.93	0.92	0.95
ρ_z	0.32	0.17	0.48
$100\sigma_R$	0.18	0.15	0.20
$100\sigma_g$	0.26	0.23	0.28
$100\sigma_d$	3.16	2.48	3.81
$100\sigma_z$	0.61	0.53	0.69
g^*	1.27	1.26	1.29
rAnet	0.74	0.35	1.15
gamQnet	0.38	0.26	0.49
piAnet	2.56	2.22	2.89

Notes: The estimation period is 1984:Q1 to 2018:Q4. The following parameters are fixed during the estimation: $\psi_1 = 2.6$, $\psi_2 = 0.98$, $\eta = 0.72$, $\nu = 0.10$, and $\chi_H = 1.00$. We use the following parameter transformations: $\beta = \exp\{-\text{rANet}/400\}$, $\gamma = \exp\{\text{gamQnet}/100\}$, and $\pi_* = \exp\{\text{piAnet}/400\}$. HPD(Low, High) refers to the boundaries of 90% highest posterior density intervals. COPF configuration: number of particles $M = 250$, ME scale $\varsigma = 0.001$, proposal covariance scale $c = 0.1$, $N = 55,000$ draws (drop first 10%), acceptance rate is 17%, run time is approximately 18 hours.

8.2. Fiscal policy analysis

The recent literature has emphasized that the effects of expansionary fiscal policies on output may be larger if the economy is at or near the ELB; see, for instance, Eggertsson (2011) and Christiano and Eichenbaum (2012). In the absence of the ELB, a typical interest rate feedback rule implies that the central bank raises nominal interest rates in response to rising inflation and output caused by an increase in government spending. This monetary contraction raises the real interest rate, reduces private consumption, and overall dampens the stimulating effect of the fiscal expansion. If, however, the economy remains at the ELB despite the expansionary fiscal policy, then the increase in inflation that results from the fiscal expansion reduces the real rate. In turn, current-period demand is stimulated, amplifying the positive effect on output. We will use our model to provide a quantitative assessment of this effect.

Because our model solution is nonlinear, the effect of a fiscal intervention depends on the initial condition and the size of the intervention. We use the Great Recession and the subsequent period in the U.S. as our laboratory and consider a fiscal intervention that is calibrated to a portion of the ARRA of February 2009 as we explain below. Our analysis is conducted from an ex post perspective, where we extract the historical shocks that make our model match the realized U.S. data, which include both a fiscal and monetary intervention, and ask what would have happened if one or both of the policy interventions were not implemented.

ARRA of February 2009 consisted of a combination of tax cuts and benefits; entitlement programs; and funding for federal contracts, grants, and loans. We focus on the third component, because it can be interpreted as an increase in g_t. We model the ARRA spending as a one-period positive shock of δ^{ARRA} to the demand shock process, where we calibrated $\delta^{ARRA} = 0.0077$ using data on the disbursement of ARRA funds, as we explain in the Online Appendix. This one-time shock is roughly $2.7\sigma_g$, and, since $\hat{g}_t$ is highly serially correlated, the effect of the shock will slowly decay over time. Since the disbursement of ARRA was spread over the year and was not done instantly, we assume that the ARRA innovation to government spending took place in the third quarter of 2009.

We use the COPF to obtain estimates of the exogenous shock processes for the period 2009:Q3 through 2011:Q2. The subsequent results are based on the mean estimator of the DSGE model parameters. The panels in the first two rows of Fig. 4 show the filtered monetary policy and government spending innovations and the levels of the government spending, technology growth, and discount factor shock processes. Recall that, in the model $\epsilon_{R,t}$ and $\epsilon_{g,t}$ are $N(0,1)$ random variables. The level processes in the second row of the figure are standardized by the unconditional standard deviations of the corresponding AR(1) processes. For the government spending and the monetary policy shocks, we distinguish between intervention and no-intervention paths.

According to the estimated model, the drop in output during the Great Recession is generated by drastic falls in the technology growth and discount factor shock processes. Because of the stylized structure of the DSGE model, these two shocks also absorb the contribution of financial shocks and financial accelerator effects. While the technology growth shock

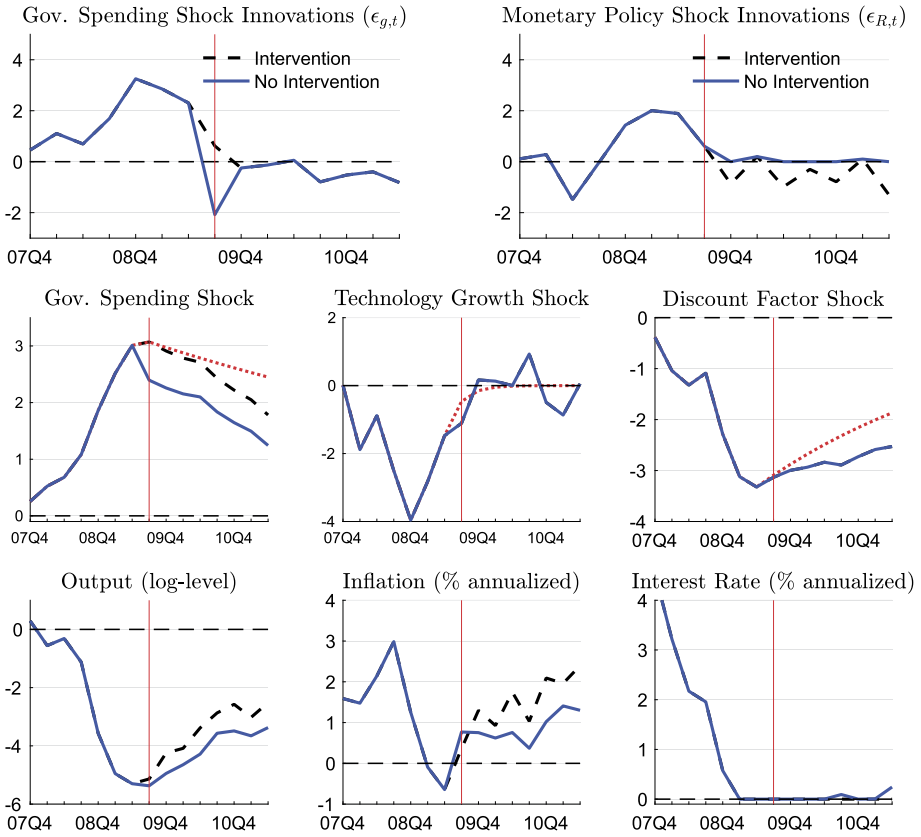

Notes: The vertical line corresponds to 2009:Q3, which is the date of the ARRA intervention. Intervention (solid) versus no-intervention paths (dashed). Along the no-intervention path, we set monetary policy shocks to zero and lower the innovation to the government spending shock in 2009:Q3 by the size of the ARRA intervention. Dotted lines represent paths in the absence of exogenous shock innovations from 2009:Q3 onwards. The level processes in the second row of the figure are standardized by the unconditional standard deviations of the corresponding AR(1) processes. Inflation and the interest rate are expressed in terms of annualized percentage rates.

Fig. 4. Ex post policy analysis.

is not very persistent ($\hat{\rho}_z = 0.32$) and reverts back to zero by the end of 2009, the mean reversion of the discount factor shock is very slow, and it remained below 2 standard deviations until the end of 2010. Meanwhile the government spending process is positive, indicating that fiscal policy started to become expansionary (relative to the historical average) in 2008. The filtered monetary policy innovations $\hat{\epsilon}_{R,t|t}$ turn out to be negative past 2009:Q3, which captures an effort by the Federal Reserve to keep the policy rate lower than what the policy rule implies. This is how our model, which abstracts from explicitly modeling unconventional monetary policies implemented in this period (quantitative easing, forward guidance), handles the existence of these policies in the data.

Because the actual path of the government spending shock already contains the effect of fiscal expansion due to ARRA, we compute the counterfactual path by subtracting the effect of ARRA from the filtered demand shock $\hat{g}_{t|t}$ using

$$\hat{g}^C_{t|t} = \hat{g}_{t|t} - \rho_g^{t-T_*}\delta^{ARRA} \text{ for } t = T_*, T_*+1, ..., T_*+7, \tag{58}$$

where T_* corresponds to 2009:Q3, the period the ARRA intervention is implemented in. The magnitude of the ARRA intervention is reflected in the difference between the intervention (dashed) and no-intervention (solid) government spending innovation depicted in the top left panel of Fig. 4. The ARRA intervention shifts $\hat{g}_t$ persistently downward as shown in the center left panel of the figure. To measure the effect of the combined fiscal and (unconventional) monetary policy, we set the counterfactual monetary policy shocks to zero.

The main finding is depicted in the bottom panels. The ex post effect of the intervention is defined as $X^o - X^C$, where X^o is the observed value of a generic variable and X^C is the counterfactual path along which the policy intervention is removed.[21] In the figure, the ex post effect is given by the gap between the no intervention and the intervention path. In the absence of the ARRA and monetary interventions, output and inflation would have been persistently lower than they

[21] Details on the algorithm to compute the effects of the policy interventions are reported in the Online Appendix.

Table 3
Ex post multipliers.

Intervention	1Q	4Q	8Q
Fiscal only	0.40	0.58	0.63
Fiscal + monetary	0.40	0.93	1.09

Notes: Ex post analysis is conditional on filtered 2009:Q3 - 2011:Q2 shocks. The pure fiscal multipliers are obtained by setting the monetary policy shocks to zero, whereas the combined fiscal and monetary multipliers are based on leaving the monetary policy shocks at their filtered values.

actually were. In particular, annualized inflation would have been 60 basis points lower, and annualized output growth would have been 43 basis points lower, on average, over these eight quarters.

Based on the output and government spending paths, we can also compute cumulative dollar-for-dollar multipliers

$$\mu_H = \frac{\sum_{\tau=1}^{H}(Y^o_{T_*-1+\tau} - Y^C_{T_*-1+\tau})}{\sum_{\tau=1}^{H}(G^0_{T_*-1+\tau} - G^C_{T_*-1+\tau})} \text{ where } H = 1, ..., 8,$$

which are reported in Table 3. The pure fiscal multipliers are obtained by setting the monetary policy shocks to zero, whereas the combined fiscal and monetary multipliers are based on leaving the monetary policy shocks at their (negative) filtered values. After one year, the ex post multiplier associated with the combined fiscal and monetary policy intervention is around 0.9 according to our estimated model. The ex post fiscal-only multiplier is around 0.6, which implies that the accompanying (unconventional) expansionary monetary policy amplified the fiscal multiplier considerably according to our analysis.

The red dotted lines in the second-row panels of Fig. 4 represent the post-2009:Q3 path of the exogenous shock processes in the absence of further innovations, which is the expected path conditional on 2009:Q3 innovation. Subsequent technology innovations had only small effects, keeping its process roughly in line with expectations. However, additional negative discount factor innovations, explain the very low interest rates after 2010:Q2. In the absence of the adverse discount factor shocks, interest rates would have moved away from the ELB, providing additional scope for the Fed to boost the effects of a fiscal expansion by keeping interest rates at zero.

9. Conclusion

Likelihood-based estimation of nonlinear DSGE models is computationally challenging. While it is becoming easier for economists to access powerful computer clusters that enable massive parallel computation, the ability to solve and estimate models on a desktop computer remains useful and desirable. Computations can often be simplified and accelerated considerably by taking shortcuts in regard to model solution or estimation techniques. The goal of this paper has been to develop a new solution method that captures important aspects of the nonlinearity generated by occasionally-binding constraints and, at the same time, allows for efficient filtering and likelihood-based estimation. The piecewise-linearity of the decision rules allows us to solve the model faster and to derive a conditionally optimal proposal distribution for a particle filter. This filter delivers a much more accurate likelihood approximation than a standard bootstrap particle filter and enables us to estimate a nonlinear New Keynesian DSGE model with a ELB constraint in a relatively short amount of time on a single core processor.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.red.2020.12.003.

References

Adam, K., Billi, R., 2007. Discretionary monetary policy and the zero lower bound on nominal interest rates. Journal of Monetary Economics 54, 728–752.
An, S., Schorfheide, F., 2007. Bayesian analysis of DSGE models. Econometric Reviews 26, 113–172.
Andrieu, C., Doucet, A., Holenstein, R., 2010. Particle Markov chain Monte Carlo methods. Journal of the Royal Statistical Society, Series B 72, 269–342.
Arulampalam, M.S., Maskell, S., Gordon, N., Clapp, T., 2002. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. IEEE Transactions on Signal Processing 50, 174–188.
Aruoba, S.B., Cuba-Borda, P., Schorfheide, F., 2018. Macroeconomic dynamics near the ZLB: a tale of two countries. The Review of Economic Studies 85, 87–118.
Aruoba, S.B., Schorfheide, F., 2016. Inflation dynamics during and after the zero lower bound. In: Inflation Dynamics and Monetary Policy. 2015 Jackson Hole Symposium Volume Published by the Federal Reserve Bank of Kansas City.
Aruoba, S.B., Schorfheide, F., Villalvazo, S., 2020. SVARs with occasionally-binding constraints. University of Maryland and University of Pennsylvania. Manuscript.
Ascari, G., Mavroeidis, S., 2020. The unbearable lightness of equilibria in a low interest rate environment. Oxford University. Manuscript.

Atkinson, T., Richter, A.W., Throckmorton, N.A., 2020. The zero lower bound and estimation accuracy. Journal of Monetary Economics 115, 249–264.
Benhabib, J., Schmitt-Grohé, S., Uribe, M., 2001. The perils of Taylor rules. Journal of Economic Theory 96, 40–69.
Benigno, G., Foerster, A., Otrok, C., Rebucci, A., 2016. Estimating macroeconomic models of financial crisis: an endogenous regime switching approach. University of Missouri. Manuscript.
Bianchi, F., Melosi, L., 2017. Escaping the great recession. The American Economic Review 107, 1030–1058.
Boehl, G., 2019. Efficient Solution, Filtering and Estimation of Models with OBCs. Working Paper.
Boehl, G., Strobel, F., 2020. US Business Cycle Dynamics at the Zero Lower Bound. Working Paper.
Cappé, O., Godsill, S.J., Moulines, E., 2007. An overview of existing methods and recent advances in sequential Monte Carlo. Proceedings of the IEEE 95, 899–924.
Cappé, O., Moulines, E., Ryden, T., 2005. Inference in Hidden Markov Models. Springer-Verlag.
Chen, H., 2017. The effects of the near-zero interest rate policy in a regime-switching dynamic stochastic general equilibrium model. Journal of Monetary Economics 90, 176–192.
Christiano, L.J., Eichenbaum, M., 2012. Notes on linear approximations, equilibrium multiplicity and E-learnability in the analysis of the zero lower bound. Northwestern University. Manuscript.
Christiano, L.J., Eichenbaum, M., Trabandt, M., 2015. Understanding the great recession. American Economic Journal: Macroeconomics 7, 110–167.
Christiano, L.J., Fisher, J.D., 2000. Algorithms for solving dynamic models with occasionally-binding constraints. Journal of Economic Dynamics & Control 24, 1179–1232.
Coleman, C., Lyon, S., Maliar, L., Maliar, S., 2018. Matlab, Python, Julia: What to Choose in Economics? CEPR Discussion Papers 13210. C.E.P.R. Discussion Papers.
Creal, D., 2012. A survey of sequential Monte Carlo methods for economics and finance. Econometric Reviews 31, 245–296.
Cuba-Borda, P., Guerrieri, L., Iacoviello, M., Zhong, M., 2019. Likelihood evaluation of models with occasionally-binding constraints. Journal of Applied Econometrics, 1–13.
Del Moral, P., 2013. Mean Field Simulation for Monte Carlo Integration. Chapman & Hall/CRC.
Doucet, A., Johansen, A.M., 2011. A tutorial on particle filtering and smoothing: fifteen years later. In: Crisan, D., Rozovsky, B. (Eds.), Handbook of Nonlinear Filtering. Oxford University Press.
Eggertsson, G.B., 2011. What fiscal policy is effective at zero interest rates? In: Acemoglu, D., Woodford, M. (Eds.), NBER Macroeconomics Annual 2010, vol. 25. University of Chicago Press, pp. 59–112.
Eggertsson, G.B., Woodford, M., 2003. The zero bound on interest rates and optimal monetary policy. Brookings Papers on Economic Activity 34, 139–235.
Fair, R.C., Taylor, J.B., 1983. Solution and maximum likelihood estimation of dynamic nonlinear rational expectations models. Econometrica 51, 1169–1185.
Farmer, R.E., Waggoner, D.F., Zha, T., 2011. Minimal state variable solutions to Markov-switching rational expectations models. Journal of Economic Dynamics and Control 35, 2150–2166.
Fernández-Villaverde, J., Gordon, G., Guerrón-Quintana, P., Rubio-Ramírez, J.F., 2015. Nonlinear adventures at the zero lower bound. Journal of Economic Dynamics and Control 57, 182–204.
Fernández-Villaverde, J., Rubio-Ramírez, J.F., 2007. Estimating macroeconomic models: a likelihood approach. The Review of Economic Studies 74, 1059–1087.
Gordon, N., Salmond, D., 1993. A novel approach to nonlinear/non-Gaussian Bayesian state estimation. IEE Proceedings. Part F 140, 107–113.
Gourieroux, C., Laffont, J.J., Monfort, A., 1980. Coherency conditions in simultaneous linear equation models with endogenous switching regimes. Econometrica 48, 675–695.
Guerrieri, L., Iacoviello, M., 2015. OccBin: a toolkit for solving dynamic models with occasionally-binding constraints easily. Journal of Monetary Economics 70, 22–38.
Guerrieri, L., Iacoviello, M., 2017. Collateral constraints and macroeconomic asymmetries. Journal of Monetary Economics 90, 28–49.
Gust, C., Herbst, E., Lopez-Salido, D., Smith, M.E., 2017. The empirical implications of the interest-rate lower bound. The American Economic Review 107, 1971–2006.
Heiss, F., Winschel, V., 2008. Likelihood approximation by numerical integration on sparse grids. Journal of Econometrics 144, 62–80.
Herbst, E., Schorfheide, F., 2015. Bayesian Estimation of DSGE Models. Princeton University Press.
Holden, T.D., 2019. Existence and uniqueness of solutions to dynamic models with occasionally-binding constraints. Deutsche Bundesbank. Manuscript.
Judd, K.L., Maliar, L., Maliar, S., 2010. A Cluster-Grid Projection Method: Solving Problems with High Dimensionality. NBER Working Paper, 15965.
Judd, K.L., Maliar, L., Maliar, S., Valero, R., 2014. Smolyak method for solving dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain. Journal of Economic Dynamics and Control 44, 92–123.
Kantas, N., Doucet, A., Singh, S., Maciejowski, J., Chopin, N., 2014. On particle methods for parameter estimation in state-space models. ArXiv Working Paper. arXiv:1412.8659v1.
Krueger, D., Kubler, F., 2004. Computing equilibrium in OLG models with stochastic production. Journal of Economic Dynamics and Control 28, 1411–1436.
Kulish, M., Morley, J., Robinson, T., 2017. Estimating DSGE models with zero interest rate policy. Journal of Monetary Economics 88, 35–49.
Liu, J.S., 2001. Monte Carlo Strategies in Scientific Computing. Springer-Verlag.
Maliar, L., Maliar, S., 2014. Numerical methods for large-scale dynamic economic models. In: Schmedders, K., Judd, K.L. (Eds.), Handbook of Computational Economics, vol. 3, vol. 3. Elsevier, pp. 325–477.
Maliar, L., Maliar, S., 2015. Merging simulation and projection approaches to solve high-dimensional problems with an application to a New Keynesian model. Quantitative Economics 6, 1–47.
Malin, B.A., Krueger, D., Kubler, F., 2011. Solving the multi-country real business cycle model using a Smolyak-collocation method. Journal of Economic Dynamics and Control 35, 229–239.
Mavroeidis, S., 2020. Identification at the zero lower bound. Oxford University. Manuscript.
Mendes, R., 2011. Uncertainty at the Zero Lower Bound: A Theoretical Analysis. MRPA Working Paper No. 60103.
Mendoza, E.G., Villalvazo, S., 2020. FiPIt: a simple, fast global method for solving models with two endogenous states & occasionally-binding constraints. Review of Economic Dynamics 37, 81–102.
Nakata, T., 2016. Optimal fiscal and monetary policy with occasionally-binding zero bound constraints. Journal of Economic Dynamics and Control 73, 220–240.
Ríos-Rull, J.-V., Schorfheide, F., Fuentes-Albero, C., Kryshko, M., Santaeulalia-Llopis, R., 2012. Methods versus substance: measuring the effects of technology shocks. Journal of Monetary Economics 59, 826–846.
Woodford, M., 2003. Interest and Prices. Princeton University Press.