BANK OF CANADA
BANQUE DU CANADA

Staff Discussion Paper/Document d’analyse du personnel — 2021-5

Last updated: March 23, 2021

Occasionally Binding Constraints in Large Models: A Review of Solution Methods

by Jonathan Swarbrick

Canadian Economic Analysis Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9

jswarbrick@bankofcanada.ca

ISSN 1914-0568

Acknowledgements

This paper is based largely on a course delivered with Tom Holden for the Centre for International Macroeconomic Studies at the University of Surrey. Accompanying code is available at https://github.com/jonswarbrick/OBC_example_code. The views expressed in this paper are those of the author. No responsibility for them should be attributed to the Bank of Canada.

Abstract

This practical review assesses several approaches to solving medium- and large-scale dynamic stochastic general equilibrium (DSGE) models featuring occasionally binding constraints. In such models, global solution methods are not possible because of the curse of dimensionality. This causes the modeller to look elsewhere for methods that can handle the significant non-linearities and non-differentiable functions that inequality constraints represent. The paper discusses methods—including Newton-type solvers under perfect foresight, the piecewise linear algorithm (OccBin), regime-switching models (RISE) and the news shocks approach (DynareOBC)—and compares the results from a simple borrowing constraints model obtained using projection methods, providing example MATLAB code. The study focuses on the news shocks method, which I find produces higher accuracy than other methods and allows the modeller to study multiple equilibria and determinacy issues.

Topics: Economic models; Business fluctuations and cycles
JEL code: C6

1 Introduction

A dynamic stochastic general equilibrium (DSGE) model is a system of stochastic difference equations composed of first-order conditions and feasibility constraints. These are typically non-linear systems without a closed-form solution, and while researchers can use global solution methods to simulate and analyze small models, they generally use local approximation methods. Using perturbation to approximate the model delivers a number of benefits, such as allowing quick and easy computation and estimation and, perhaps most appealingly, providing reasonably accurate analysis around the steady state with clear policy recommendations. However, a significant drawback, and one that I focus on in this review paper, is an inability to deal with non-differentiable functions, including occasionally binding constraints (OBCs), such as those that include `min` and `max` operators. This type of function can be found in nearly all macroeconomic models, whether it be a positivity constraint on investment, borrowing constraints on households or firms, or the zero lower bound (ZLB) on nominal interest rates. However, common practice is to either ignore the presence of the constraints or assume they are always binding. In this paper, I briefly discuss the various OBCs commonly encountered in macroeconomic research. I then look at the different ways of solving models with OBCs, focusing on using the DynareOBC tool kit described in Holden (2019, 2016). As this paper shows, this method allows the modeller the freedom to strike the right balance between speed and accuracy and, in contrast to other out-of-the-box methods, allows the easy analysis of multiple equilibria. This study is a practical guide to using these methods, mostly within the Dynare environment, and it assumes some familiarity with Dynare (for more information on Dynare, see Adjemian et al. (2011)).

The two most common methods Dynare uses to solve models are perturbation methods when the command `stoch_simul()` is used and Newton-type methods when `simul()` is used. Specifically, when `stoch_simul()` is used, Dynare solves a Taylor approximation of the decision and transition functions up to third order. Higher-order approximations will give a better fit to the non-linearities of the model but also play a role in how uncertainty affects model dynamics. At first order, certainty equivalence holds and there is no risk premium. At second order, risk is constant. And at third order, risk is linear in the state; this is the lowest order at which risk will be time varying. If `simul()` is used, Dynare uses Newton methods to solve the non-linear model over a finite time period. In this case, agents in the model are assumed to know the full future path of the economy, and so risk plays no part. Dynare can also solve a stochastic, non-linear model using the extended path algorithm proposed by Fair and Taylor (1983). This is solved when the command `extended_path()` is used. The order the user chooses determines the accuracy of evaluating future uncertainty. For orders greater than zero, Dynare uses a Gaussian quadrature, which scales exponentially in the number of shocks and the order. Because of this, attempting to achieve accuracy using this method is likely to be prohibitively slow.

Models with OBCs can be simulated using global solution methods, such as through dynamic programming, or using projection methods, which are able to capture all model non-linearities. The main drawback is that models solved using global techniques do not scale well to larger models; as a result, many models solved using this method make restrictive assumptions to limit the state space, such as using exogenous endowment instead of an endogenous production sector. Global methods also cannot guarantee convergence for non-optimal problems, such as the Taylor rule with a ZLB, and require a high level of technical skill to set up. There are a number of methods for solving larger models with OBCs, including fast ones (such as Guerrieri and Iacoviello (2015)) and accurate ones (such as Maliar and Maliar (2015)). The choice to focus on the method proposed in Holden (2019, 2016) is guided by the balance between accuracy and speed in simulation. The method also provides necessary and sufficient conditions for a unique solution and the ability to select between alternative solutions when there are multiple. I compare simulation results obtained using the following methods: projection methods; the inbuilt Dynare perfect foresight and stochastic extended path methods; the piecewise linear method using the OccBin tool kit; a regime-switching approach using the RISE tool kit; and DynareOBC, which relies on news shocks. Details of the methods are kept as an overview in this review paper; see Guerrieri and Iacoviello (2015) for further information on OccBin and Binning and Maih (2017) for information on RISE. For DynareOBC, see a theoretical paper on the existence and uniqueness conditions in Holden (2019), and a paper detailing the computational algorithm in Holden (2016). I begin with a discussion of one OBC that is often the focus of attention: the ZLB on nominal interest rates.

1.1 The zero lower bound on nominal interest rates

Researchers have long recognized that the ZLB could be a significant constraint on a central bank's attempt to combat deflation. However, it is only since the global financial crisis (GFC) that the issue has risen to prominence in the macroeconomic literature, emphasized by an enormous amount of recent research into the ZLB and related topics.

Many authors have studied the dynamics of the New Keynesian (NK) model under a ZLB constraint (e.g., Braun and Körber 2011; Fernández-Villaverde et al. 2015). Of these, Coibion, Gorodnichenko and Wieland (2012) question whether the inflation target should be increased, Bilbiie, Monacelli and Perotti (2019) analyze government spending at the ZLB and Wieland (2019) examines the effect of supply shocks. Some papers measure the impact of the ZLB on expectations of longer-term rates (Wright 2012; Swanson and Williams 2014; Bauer and Rudebusch 2016); others study the link between financial stability and the ZLB (Fischer 2016); some look at uncertainty (Nakata 2017; Plante, Richter and Throckmorton 2017); and others study credibility and commitment in a ZLB environment (Adam and Billi 2007; Bodenstein et al. 2012). Many papers have also studied the efficacy of unconventional policies when constrained by the ZLB (see Hamilton and Wu 2012; Gambacorta, Hoemann and Peersman 2014; Gilchrist, López-Salido and Zakrajšek 2015; Wu and Xia 2016).

The Japan experience since 1995 led some researchers to work in the area, notably Eggertsson and Woodford (2003), who describe previous interest in the subject as a "theoretical curiosity." The experience of the United States, the United Kingdom and other advanced economies following the GFC led to a rush of research into what had become a problem for countries trying to combat high debt without standard monetary policy to counter the contractionary effects. Even as interest rates began to rise, there was increased concern that hitting the ZLB was becoming a regular occurrence, as "normal" rates seemed lower than in the past. On this issue, contributions have argued that the fiscal spending multiplier is particularly large when the ZLB is binding (e.g., Woodford 2011; Eggertsson 2010; Erceg and Lindé 2014), whereas others have argued that there is no multiplier effect (see, e.g., Braun and Körber 2011; Braun, Boneva and Waki 2016; Mertens and Ravn 2014). The intuition for a higher multiplier is as follows: suppose an economy is experiencing an episode with nominal interest rates at their ZLB and the government begins a fiscal expansion. If the expansion raises inflation expectations, then, because the nominal interest rate is fixed, these are directly passed through to lower expected real rates. This stimulates activity further than if nominal rates were also to rise in response. An alternative channel discussed in Eggertsson (2010) works through the labour market: increased inflation expectations reduce real incomes, prompting increased labour supply; this causes the national income to rise. In contrast to the arguments for a high multiplier, Braun, Boneva and Waki (2016; 2013) study the ZLB in an NK model and find fiscal multipliers close to one. They argue that the result depends on whether a linear or non-linear solution is used and which price-setting friction is used. Eggertsson and Singh (2019) respond by showing that a positive multiplier is present in both linear and non-linear settings and disappears with certain types of price-setting friction only when applied in an unrealistic way. The question is yet to be resolved.

Another debated issue relates to the possible presence of multiple equilibria and indeterminacy. While the uniqueness of equilibria can be guaranteed in the linear NK model, this is not obviously true when interest rates are unable to turn negative. Consider the standard 3-equation NK model with a ZLB:

$$\hat{\pi}_t = \beta \mathbb{E}_t \left[\hat{\pi}_{t+1}\right] + \kappa \hat{y}_t + \epsilon_t^s \tag{1}$$

$$\hat{y}_t = \mathbb{E}_t \left[\hat{y}_{t+1}\right] - \frac{1}{\sigma}\left(\hat{r}_t - \mathbb{E}_t \left[\hat{\pi}_{t+1}\right]\right) + \epsilon_t^d \tag{2}$$

$$\hat{r}_t = \begin{cases} \phi_\pi \hat{\pi}_t + \phi_y \hat{y}_t + \epsilon_t^m & \text{if } r_t \equiv \hat{r}_t + \bar{r} \geq 0 \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

This implies two long-run equilibria dependant on expectations: a normal/inflationary equilibrium with $r = \bar{r}$, $\pi = \pi^*$ and a deflationary one with $r = 0$, $\pi = -\bar{r}$. This is one cause of the so-called liquidity trap (see Benhabib et al. 2014) and opens up the possibility of non-fundamental[1] shifts in expectations that cause the economy to jump between equilibria (sunspots). Because being in the deflationary equilibrium implies that households and firms believe the economy is converging to the *bad* steady state, many papers rule this out explicitly on both theoretical and empirical grounds. Significant evidence shows that central banks are able to signal their commitment to the *good* equilibrium in the long run, and potentially employ alternative policy measures to ensure this is the case. Furthermore, while the inflationary equilibrium is locally stable, the deflationary one is not. However, even when ruling out the long-run deflationary equilibrium, Holden

[1] That is, unrelated to "fundamental" shock, such as to technology and preferences.

(2019) shows that under many specifications, the model can still be indeterminate when a ZLB is present. For example, suppose that in a given quarter, the expectation of the ZLB binding in the next quarter is deflationary; this expectation could lead to a self-fulfilling ZLB episode as the interest rate is cut in response.[2] So, in a similar way to the sunspots, there is the possibility of non-fundamental shifts in expectations. This is the case for many well-known NK models, such as that of Smets and Wouters (2007). Holden shows that uniqueness can be restored by replacing inflation targeting with a price-level targeting regime.

1.2 Other models of occasionally binding constraints

After the ZLB, perhaps the most prominent class of OBC in the literature is borrowing constraints experienced by households, banks and firms. The role of household borrowing constraints has been studied by Guerrieri and Iacoviello (2017), who show how the presence of occasionally binding collateral constraints can drive macroeconomic asymmetries. In a borrower-saver DSGE model, they show that during a housing boom, borrowing constraints become slack as the value of housing collateral increases. When house prices fall, borrowing constraints tighten, leading to much sharper downturns. A growing literature is studying other occasionally binding financial constraints, such as adverse selection in investment (Swarbrick 2019), bank borrowing constraints (Holden, Levine and Swarbrick 2020), equity constraints (Brunnermeier and Sannikov 2014; He and Krishnamurthy 2013) and cash-in-advance constraints (Dixon and Pourpourides 2016).

Another OBC found in the literature is irreversible investment (see Bernanke 1983; Caballero and Pindyck 1996; Bloom et al. 2007). Although its irreversibility is often overlooked in the macroeconomic literature, investment is a sunk cost and cannot be recovered if economic conditions change significantly. This implies an opportunity cost of investment not present in the standard model that introduces an uncertainty effect. Consider a firm choosing to invest and observing the distribution of possible states-of-the-world. When times are good, the firm will want to invest more; but when times are bad, it may want to disinvest. If we define higher uncertainty as an increase in the dispersion of these states of the world, it follows that higher uncertainty implies a higher probability

[2] To see why a future ZLB episode could be deflationary, note that the implication of the ZLB binding is that the interest rate is higher than it would be without the bound. Therefore, the expectation of the ZLB binding is equivalent to the expectation of future monetary shocks.

of disinvestment. This increased probability will cause firms to insure against this by waiting to invest.[3] That being said, an important consideration is that aggregate investment never gets close to zero; so, for irreversible investment to have a role in aggregate business cycles, one must consider some form of firm heterogeneity where the idiosyncratic, firm-level shocks are much larger than the aggregate macroeconomic shocks. The same is partly true with respect to household borrowing constraints, giving rise to the heterogeneous agent New Keynesian (HANK) literature (see, e.g., Kaplan, Moll and Violante 2018). However, rather than introducing full heterogeneity, researchers commonly model two types of households (equivalently, TANK), such as those with patient savers and impatient borrowers, which is the approach that Guerrieri and Iacoviello (2017) take. Debortoli and Gali (2017) show that using two agents can be a good approximation to a fully heterogeneous agent model.

2 Model solutions methods

If a particular OBC is thought to play an important role in macroeconomic dynamics, then linear solution and estimation are likely to perform poorly. Modelling the ZLB in DSGE models, for instance, has become necessary for analysts to produce accurate forecasts and recommend alternative measures, such as fiscal stimulus or unconventional monetary policy. I identify several approaches in the theoretical literature used to solve models with OBCs and, at least briefly, look at all of them in this paper. The first, as has already been discussed, is to use global approximation methods to capture all model non-linearities, including the precautionary behaviour, although use of these methods is restricted to small-scale models given the rapidly increasing computation cost.

A popular alternative that retains the model non-linearities is to use a deterministic solution that assumes perfect foresight and in which the model must return to equilibrium in finite time. In one sense, this method is more accurate than approaches based on local approximation techniques under uncertainty, as it captures the full non-linearities of the model—but, of course, at the cost of losing agents' precautionary behaviour inherent to a stochastic environment. One approach that can incorporate the role of risk is to use a functional approximation, or a penalty function approach, that approximates the

[3] Gilchrist, Sim and Zakrajšek (2014) find evidence for a positive relationship between uncertainty and credit spreads, consistent with the option value.

constraint in a perturbation solution (see, e.g., Brzoza-Brzezina et al. 2015). Another approach, inspired by Eggertsson and Woodford (2003), was developed further by Jung et al. (2005) and implemented into Dynare by Guerrieri and Iacoviello (2017) in the OccBin tool box. This relies on a two-state (or piecewise) linear approximation that is driven by a Markov chain with an absorbing state. Related to this is a regime-switching approach pioneered by Binning and Maih (2017) and implemented in Junior Maih's tool kit, RISE. The final approach, and the one that I focus on, is to use anticipated news shocks to impose the constraints. I will discuss in detail the computational method proposed in Holden (2016) and review some of the theoretical results and considerations outlined in Holden (2019).

These different approaches can be used to introduce many different types of OBCs into macroeconomic models, such as those summarized above. The benefits of these approaches vary with respect to the trade-off between speed and accuracy, and depend on the following: the number of constraints and how they interact; how often they are expected to bind; and whether they are endogenous in the sense of being internalized in the decisions of economic agents (e.g., borrowing constraints) or not (e.g., the ZLB). In this discussion, I will explore the alternative methods presented above and consider some of these factors. I will examine some methods in greater detail, with a particular focus on the implementation in Dynare.

To provide some comparison between methods, I will analyze the results from the following simple model with a fixed real interest rate in which households face a borrowing constraint when selling bonds but can buy bonds without limit. As the interest rate is fixed and there is an unlimited supply of bonds, the model can be seen as a small open economy. The representative household solves

$$\max_{C_{t+s},H_{t+s},b_{t+s}} \mathbf{E}_t \sum_{s=0}^{\infty} \beta^s \left[\log C_{t+s} + \chi \log\left(1 - H_{t+s}\right) - \frac{\delta}{2} b_{t+s}^2\right] \tag{4}$$

$$\text{s.t.} \quad C_t + b_t = \exp\left(z_t\right) H_t + R b_{t-1} \tag{5}$$

$$b_t \geq \underline{b}, \tag{6}$$

where R is an exogenous rest-of-world interest rate; H_t is hours worked; $\underline{b} \leq 0$ is the borrowing limit; and $\frac{\delta}{2} b_t^2$ is a cost of holding or selling bonds, introduced to ensure stationarity. δ can be set very small. z_t is the exogenous productivity level following an

AR(1) process:

$$z_t = \rho z_{t-1} + \epsilon_t, \quad \epsilon_t \sim N\left(0, \sigma^2\right). \tag{7}$$

With a little rearranging of the resulting first-order conditions, this leads to saving, consumption and leisure decisions characterized by

$$\frac{1}{C_t} = \beta \mathbb{E}_t \left[\frac{1}{C_{t+1}}\right] R + \mu_t - \delta b_t \tag{8}$$

$$C_t = \frac{\exp(z_t) + R b_{t-1} - b_t}{1+\chi} \tag{9}$$

$$H_t = 1 - \chi \frac{C_t}{\exp(z_t)} \tag{10}$$

$$\mu_t \left(b_t - \underline{b}\right) = 0 \tag{11}$$

$$\mu_t \geq 0 \tag{12}$$

$$b_t \geq \underline{b}. \tag{13}$$

The last three equations are the Kuhn-Tucker conditions, and μ_t is the Kuhn-Tucker multiplier, equivalent to a Lagrange multiplier. This has the same interpretation as giving the value of relaxing the borrowing constraint. As the constraint is well-behaved, the last three conditions can be summarized with

$$\min\left\{\mu_t, b_t - \underline{b}\right\} = 0. \tag{14}$$

2.1 Choice of benchmark

Of the methods mentioned above, the global approximation can be considered the closest to the "true" rational expectations solution to the model. For this reason, in this analysis, I treat the solution found using global methods as the benchmark and loosely assess the other approaches conditional on the distance from the projection solution.[4]

3 Projection methods

If the dimension of the model is small—such as with the borrowing constraints model defined by equations (7–10, 14), which has only one state variable and one shock—then it

[4] A more typical approach to assessing accuracy is to compute the approximation error using Euler equation residuals (see Santos 2000). The Euler error is commonly found to be of the same order of magnitude as the approximation error, but OBCs introduce an additional source of approximation from which this choice may not be as robust.

is possible to compute a global approximation of the policy function. Various methods are examined in a large literature on the topic; but, as a benchmark example, a small model with an occasionally binding constraint can easily be solved with projection methods (see Judd 1992).

The objective is to find a function that characterizes the household consumption, saving and labour supply decisions. This unknown function can be written as

$$b_t = g\left(b_{t-1}, z_t\right) \tag{15}$$

to give a backward-looking law of motion for the economy. The true "policy" function $g\left(\cdot\right)$ can be approximated using some "basis" function, a polynomial for instance, by specifying a grid over b_{t-1} and z_t and making sure the basis function satisfies the model at each grid point. Whether this is satisfied exactly or whether residuals at each grid point are minimized depends on the choice of basis function.

In the borrowing constraints model here, the "functional" equation to approximate is the Euler equation (8). This is the only equation with an expectation operator, and focusing on this is sufficient to solve the household saving problem, with the remaining equations determining consumption and hours in closed form. Using the as-yet-unknown policy function, $g\left(\cdot\right)$, the Euler equation can be given by

$$0 = \frac{1+\chi}{\exp\left(z\right) + Rb - g\left(b, z\right)} - \mathbb{E}_t\left[\frac{1+\chi}{\exp\left(z'\right) + Rg\left(b, z\right) - g\left(b', z'\right)}\right] - \mu + \delta g\left(b, z\right), \tag{16}$$

where we are using the notation $z = z_t$, $z' = z_{t+1}$, $b = b_{t-1}$ and $b' = b_{t+1}$. Notice, then, that if the approximation of the policy function is denoted as $\tilde{g}\left(\cdot\right)$, then the error implied by the approximation can be computed using

$$\begin{aligned} e = {} & \frac{1+\chi}{\exp\left(z\right) + Rb - \tilde{g}\left(b, z\right)} \\ & - \sum_{z'} p(z')\left[\frac{1+\chi}{\exp\left(z'\right) + R\tilde{g}\left(b, z\right) - \tilde{g}\left(\tilde{g}(b, z), z'\right)}\right] - \mu + \delta\tilde{g}\left(b, z\right), \end{aligned} \tag{17}$$

where we have nested the policy function in place of b'' and replaced the expectation operator with a numerical integration over possible shocks z' where $p(z')$ is the probability of observing a particular z'. Over a defined set of values for b and z, we wish to solve $\tilde{g}\left(b, z\right)$ that minimizes e. In practice, and in what we do here, this might be to choose the set of parameters that minimizes the sum of squared errors at each grid point, interpolating between each grid point with cubic splines.[5] To do so, we rearrange equation (17) to

$$\tilde{g}_{n+1}(b,z) = \exp(z) + Rb - \frac{1}{\sum_{z'} p(z')\left[\frac{1}{\exp(z') + R\tilde{g}_n(b,z) - \tilde{g}_n\left(\tilde{g}_n(b,z), z'\right)}\right] + \mu - \frac{\delta}{1+\chi}\tilde{g}_n(b,z)} \tag{18}$$

with an implied zero error. The n indicates the iteration number; if we begin with $\tilde{g}_1(b,z) = b$, the choice of bonds is set equal to the existing stock, and we can evaluate (18) to solve $\tilde{g}_2(b,z)$. On the face of it, we still have an unknown variable μ. But recall that this is the value of relaxing the borrowing constraint. If the borrowing constraint is slack, then μ must equal zero. If it binds, then $\mu > 0$, but $\tilde{g}_n(b,z) = b' = \underline{b}$. We can therefore rewrite (18) as

$$\tilde{g}_{n+1}(b,z) = \min\left\{\underline{b}, \left(\exp(z) + Rb - \frac{1}{\sum_{z'} p(z')\left[\frac{1}{\exp(z') + R\tilde{g}_n(b,z) - \tilde{g}_n\left(\tilde{g}_n(b,z), z'\right)}\right] - \frac{\delta}{1+\chi}\tilde{g}_n(b,z)}\right)\right\}. \tag{19}$$

Given the new policy function, we check the error from equation (17). If it is within chosen tolerance, then we have found our policy function; otherwise, use $\tilde{g}_{n+1}(b,z)$ on the right-hand side in the next iteration.[6]

To solve this in MATLAB, given the model parameters and the choice of projection method, we first make decisions on the number of grid points, error tolerance and the method of quadrature. This example has linearly spaced grid points, and so we define this using `linspace()`.[7] With N_b grid points in the state and N_z in the shock, we end up with a $N_b \times N_z$ grid for both b and z, with each element representing a possible state of the world. Note that for a higher dimension problem—for example, if we also

[5] As discussed in Miranda and Fackler (2002), one could choose among different basis functions (e.g., Chebychev polynomials) and several projection methods (e.g., the Galerkin method and collocation). For example, Chebychev collocation works particularly well for smooth policy functions. In our case, we want to allow for kinks stemming from the constraint.

[6] Of course, for some models and starting guesses, this may not converge to a solution. In this case, I refer the reader to Judd (1992) for alternative solutions.

[7] One may wish to concentrate grid points in the state where the economy spends the majority of time, or at a point at which there is a kink. The `funnode()` function in Miranda and Fackler's (2002) CompEcon tool kit will generate grids concentrating nodes toward the end points.

had habits in consumption so that C_{t-1} was a state variable—we could conveniently use `[B_mesh,C_mesh,z_mesh] = ndgrid(B,C,z);`, given the arrays for b, C and z.[8]

```
Nb = 100;
Nz = 100;
min_b = b_limit;
max_b = 0.8;
min_z = -0.1;
max_z = 0.1;
err_tol = 1e-8;
% Grids for the state
b = linspace( min_b , max_b , Nb )';
z = linspace( min_z , max_z , Nz );
b_mesh = repmat( b , 1 , Nz );
z_mesh = repmat( z , Nb , 1 );
```

For the integration, we use Gauss-Hermite quadrature, with the nodes and weights given by the following:

```
% Nodes and weights for numerical integration
q_pts = 50;
[q_n,q_w]= hernodes(q_pts);
```

Using this, for each of the 50 nodes, z' will equal $\rho z + \sigma q_n \sqrt{2}$, and the probabilities will equal $p\left(z'\right) = q_w / \sqrt{\pi}$. We define a function, `ex_mu()`, that takes the grids and the current guess as inputs and returns the expectation term in (19), using the following:

```
for i_q = 1:q_pts
z_p = rho * z_mesh + eps_p( i_q );
b_p_p = interp2( z_mesh , b_mesh , b_p , z_p , b_p , 'spline' );
ex_up_p = ex_up_p + ( q_w(i_q) * up( ( exp( z_p ) + r * b_p - b_p_p ) / (
    1 + chi ) ) )./ sqrt(pi);
end
```

[8] Alternatively, if using the CompEcon tool kit, `fundefn()` can be used to set up the function space, `funnode()` to create the grid points, and `gridmake()` to create the grids. See Miranda and Fackler (2002) for more details.

The function `interp2()` is a standard MATLAB function. The first two arguments (b and z) are inputs to the policy function that returns the third argument (b'). The third and fourth arguments (b' and z') are then inputs to the policy function that returns the output from `interp2()` (b''). The function has exact values for b' at the grid points of b and z and then interpolates between them with cubic splines, that is, third-order polynomials. Given the initial guess for the policy function, `b_p = b_mesh;`, we compute an initial value for the expectation term `mu_p = ex_mu( b , z , b_p , params );` then loop through the following:

```
while ssr>err_tol
b_p = max( exp( z_mesh ) + r * b_mesh -  ( 1 + chi ) ./ ( mu_p - delta *
    b_p ) , b_limit );
mu_p_new = ex_mu( b , z , b_p , params );
error = mu_p_new- mu_p;
ssr = sum( sum( error.^2 ) );
mu_p = mu_p_new;
end
```

On each iteration, we first calculate a new b' from the Euler equation, compute the expectation term given this new value and compare with that from the previous iteration. If the sum of squared residuals falls below the defined tolerance, then we finish; otherwise we repeat.

We can compute impulse response functions (IRFs) to a large negative productivity shock. To do so, we simulate the model a large number of times over a horizon T for which we plan to drop the first t_d periods. We repeat each simulation twice for the same sequence of shocks; but for one of them, we add a technology shock of the desired size in period t_d+1. We take the difference, then average over all replication to give an average, or expected IRF. In the code, first, we draw a random sequence of shocks ϵ_t to give a time-series for z. This is repeated r number of times:

```
% IRF simulation
irf_horizon = irf_periods+irf_drop;
% Draw shocks
eps_mat = randn( replic, irf_horizon );
```

For each $i \in \{1, r\}$, we loop over the simulated the model, repeating the simulation twice on each iteration but with the added shock:

```
for ii=1:replic
eps_1 = eps_mat( ii , :);
eps_2 = eps_1;
% Add shock of interest
eps_2( irf_drop+1 ) = eps_2( irf_drop+1 )-IRF_shockscale;
% Simulation
for t=2:irf_horizon
z_1(t) = rho * z_1(t-1) + sigma * eps_1(t);
b_1(t) =  interp2( z_mesh , b_mesh , b_p , z_1(t) , b_1(t-1) , 'spline' );
z_2(t) = rho * z_2(t-1) + sigma * eps_2(t);
b_2(t) =  interp2( z_mesh , b_mesh , b_p , z_2(t) , b_2(t-1) , 'spline' );
end
c_1(2:end) = ( exp( z_1(2:end) ) + r .* b_1(1:end-1) - b_1(2:end) ) ./ (1 + chi);
h_1 = 1 - chi .* c_1 ./ exp( z_1 );
c_2(2:end) = ( exp( z_2(2:end) ) + r .* b_2(1:end-1) - b_2(2:end) ) ./ (1 + chi);
h_2 = 1 - chi .* c_2 ./ exp( z_2 );

z_mat( ii , : ) = z_2 - z_1;
b_mat( ii , : ) = b_2 - b_1;
c_mat( ii , : ) = c_2 - c_1;
h_mat( ii , : ) = h_2 - h_1;
end
```

The size of shock we are interested in is defined by `IRF_shockscale`; in this case, we choose 2.

Following a negative productivity shock, households wish to substitute in leisure for consumption as the returns to work fall but choose to borrow to ensure a smooth consumption path. Results are shown in Chart 1. As expected, all the variables drop following the shock. When households face a borrowing constraint, they are unable to borrow from

Chart 1: Impulse response functions following technology shock using projection method

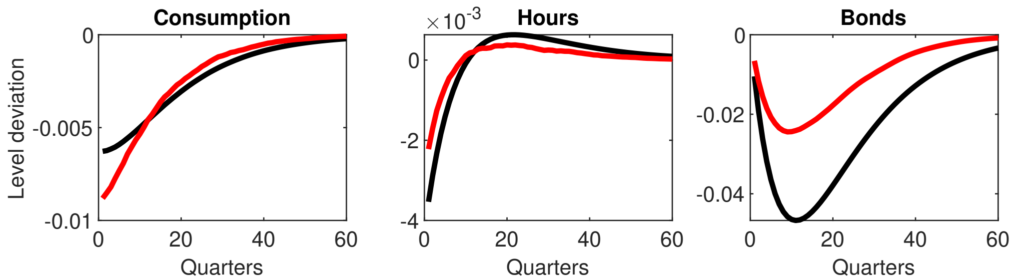

Note: The red line is the model with borrowing constraint and the black line is a model with unconstrained borrowing. Plots show level deviation from the ergodic mean.

Table 1: Simulated moments

		Mean	Standard deviation	Skewness
No constraint	C	0.667	0.012	0.08
With OBC		0.667	0.013	-0.34
No constraint	H	0.667	0.003	-0.03
With OBC		0.666	0.002	-0.07
No constraint	b	0.0008	0.10	-0.005
With OBC		0.059	0.06	1.14

Note: The table presents simulated moments of the model variables with and without the borrowing constraint using projection method. The v ariables are expressed in levels.

their future consumption as much as they would like. So instead, they substitute leisure for higher consumption, working more to make up some of the difference. Note that bonds do not appear bound at $\underline{b} = 0.01$ in this simulation. This is because the chart shows deviation around the ergodic mean, which is found to be approximately 0.06. The fact that bonds do not fall further in expectation is a result of the positive skewness stemming from the presence of a lower bound.

We can simulate time series using either a single draw of shocks over a long horizon or many repeated simulations, and compute moments using the simulated data. We can then compare moments to assess the impact of the borrowing constraint. Table 1 displays the first three moments from the simulated time series. The idea here is not to attempt to match the empirical business cycle, but to emphasize the theoretical impact of the borrowing constraint. When binding, the constraint generates a relative fall in consumption and higher hours. The constraint does not significantly alter the ergodic mean of consumption and hours, but it does affect borrowing, which is 10 times higher in the model with the constraint. Together with the large skewness, this highlights the asymmetry in bonds, as the household can increase saving as much as possible in booms but is limited in borrowing during downturns. I now turn to solution methods available using Dynare. These are particularly relevant to solving larger models, for which global methods are not feasible.

4 Solution methods in Dynare

When the expression `stoch_simul()` is used in the `.mod` file, Dynare computes a Taylor approximation of the decision and transition functions of the model. This is computed around the deterministic steady state and at first order will be of the form

$$b_t = b + A(b_{t-1} - b) + Bz_t, \tag{20}$$

where b is the steady-state value of b_t. The approach generalizes to higher orders and is available in Dynare up to order three. For further reading on the technical details for the derivation of the policy function (20) and the higher-order equivalents, see Collard and Juillard (2001).

Suppose we wanted to include the occasionally binding borrowing constraint (14); then, although it is possible to write the equation `0 = min(mu,b-b_limit)`, if `stoch_simul()` is used, Dynare will perform a local approximation around the deterministic steady state. The equation will then be an approximation of either $b_t = b$ or $\mu_t = 0$, depending on whether the constraint binds in steady state. The inequality constraint will be lost. While the perturbation techniques employed by Dynare to solve stochastic models lose this type of constraint, the same is not true for perfect foresight or extended path solutions.

4.1 Dynare's perfect foresight solver

When `simul()` is used in the `.mod` file, Dynare uses a Newton-type algorithm that preserves these non-linearities. Details of the various algorithms available can be found at Adjemian et al. (2011). In such a model, the economy is in some defined equilibrium up until period 1—the deterministic steady state, for example—at which point the agents learn about any shocks in the current or future periods. The simulation is the path of the economy conditional on a set of shocks over a finite horizon that are known in the first period. The final equilibrium may or may not be the same as the initial one before period 1, and the path to final equilibrium is imposed in finite time rather than asymptotically. We can write the model as

$$\begin{bmatrix} \frac{1}{C_t} - \beta \frac{1}{C_{t+1}} R - \mu_t + \delta b_t \\ \frac{\exp(z_t) + RB_{t-1} - B_t}{1+\chi} - C_t \\ 1 - \chi \frac{C_t}{\exp(z_t)} - H_t \\ \min\left\{\mu_t, b_t - \underline{b}\right\} \\ \rho z_{t-1} + \epsilon_t - z_t \end{bmatrix} = 0 \tag{21}$$

or

$$f(x_{t-1}, x_t, x_{t+1}, \epsilon_t) = 0, \tag{22}$$

where $x_t \equiv \begin{bmatrix} b_t & C_t & H_t & \mu_t & z_t \end{bmatrix}$ is the endogenous state and ϵ_t the exogenous shock. A perfect foresight solution—given the initial conditions, x_0, and sequence of shocks, ϵ_t—is a path $x_0, x_1, x_2, \cdots$ converging to steady state, or another defined end condition, and satisfying

$$\begin{aligned} f(x_0, x_1, x_2, \epsilon_1) &= 0 \\ f(x_1, x_2, x_3, \epsilon_2) &= 0 \\ f(x_2, x_3, x_4, \epsilon_3) &= 0 \\ &\vdots \end{aligned}$$

When `simul()` is called, Dynare stacks the model over T periods,

$$F(X) = \left\{ \begin{matrix} f(x_0, x_1, x_2, \epsilon_1) \\ f(x_1, x_2, x_3, \epsilon_2) \\ \vdots \\ f(x_{T-1}, x_T, x_T, \epsilon_3) \end{matrix} \right\} = 0 \tag{23}$$

T is chosen sufficiently high such that the model is able to converge to the terminal condition x_T.[9] The default algorithm then works as follows. Every element of x is set

[9] It is common to choose $x_T = x_0 = \bar{x}$ so that the model begins and ends at the deterministic steady state. There should be a sufficient burn-off period with no shocks at the end of the simulation to allow the economy time to converge.

to the steady-state value $\bar{x}$ before beginning the iterative Newton method to solve all x simultaneously. For the nth iteration, every element is set to

$$x(n+1) = x(n) - \frac{f(x(n))}{f'(x(n))}. \tag{24}$$

The iterations repeat until convergence.

Dynare offers a number of alternative algorithms, mostly variations of the Newton method described. To enter the OBC in the model file, we can write, for example,

```
0 = min(mu,b-b_limit);
```

invoking the following command:

```
simul( periods=400 , stack_solve_algo = 0 );
```

Options of interest that are available in this setting include:[10]

- `periods = INTEGER`: This is the number of periods of the simulation.
- `stack_solve_algo = 0|1|2|3|4|5|6`: This allows the user to decide which of the available seven Newton type algorithms to use.
  An alternative is to choose `stack_solve_algo = 7` to use an alternative algorithm defined using option `solve_algo = 0|1|2|3|4|5|6|7|8|9|10|11`.

Choosing option `stack_solve_algo = 0` requests that Dynare use the Newton method with sparse matrices. An alternative approach, which is perhaps better suited to models with OBCs, is to treat the problem as a Levenberg-Marquardt mixed complementarity problem (LMMCP), invoked with option `lmmcp`.[11] When a model has a condition such as a borrowing constraint or ZLB, there is a discontinuity in the derivative (i.e., a singularity in the Jacobian), which can lead to problems solving (24). We can use the algorithm proposed in Kanzow and Petra (2004) to evaluate a generalized Jacobian, removing this singularity.

Because using a deterministic set-up allows the analysis of the full implications of non-linearities, it is a useful starting point—although there are some drawbacks. The model is solved by stacking the system over T periods, where T is large enough to return the

[10] Refer to the Dynare manual for full description of all options.
[11] This is equivalent to using `stack_solve_algo = 7, solve_algo = 10`.

system to equilibrium. The Newton approach solves all the equations over every period simultaneously. This approach leads to a very large Jacobian; for a simulation over T periods, a model with n endogenous variables will require a Jacobian matrix of order n x T, leading to high computational demands. However, the Jacobian can be handled as one large sparse matrix, cutting down on these demands. The particular solution returned can also raise questions. It is unclear whether there are multiple solutions to the model and why the one returned is chosen. And it is impossible to even know if there is a solution or not if the method fails to converge. Another drawback, and the reason why the perfect foresight technique can only be a starting point, is that the role of uncertainty in the behaviour of agents is lost. This could introduce bias into the precautionary actions that arise from the risk of hitting the bound and, if the constraint is likely to bind in the near future, the bias could be large.

As before, we can compute the response of the simple borrowing constraint model to a large negative productivity shock. To simulate an IRF, we specify the shock as follows:

```
shocks;
var epsz; periods 1;  values -2;
end;
```

Because Dynare does not produce IRF plots automatically when calling `simul()`, we can include the following code at the end of the .mod file:

```
figure_per_shock=floor((M_.endo_nbr-1)/9)+1;
for i=1:M_.endo_nbr
VarName = strtrim( M_.endo_names(i,:) );
j=floor((i-1)/9)+1;
p=i-(j-1)*9;
figure( figure_per_shock*(1-1)+j );
subplot(3,3,p)
plot( oo_.endo_simul(i,1:20) );
title( VarName );
if p==1
legend('epsz')
end
end
```

Results are shown in Chart 2. As expected, all the variables drop following the shock. When households face a borrowing constraint, they are unable to borrow from their future consumption as much as they would like. So instead, they substitute leisure for consumption, working more to make up some of the difference.

We can compare the perfect foresight IRFs with the global solution under uncertainty. Because of the lack of precautionary saving under perfect foresight, there is much less room for b to fall following the shock. This exaggerates the impact on consumption relative to a solution method that captures precautionary behaviour.

Chart 2: Perfect foresight impulse response functions following technology shock

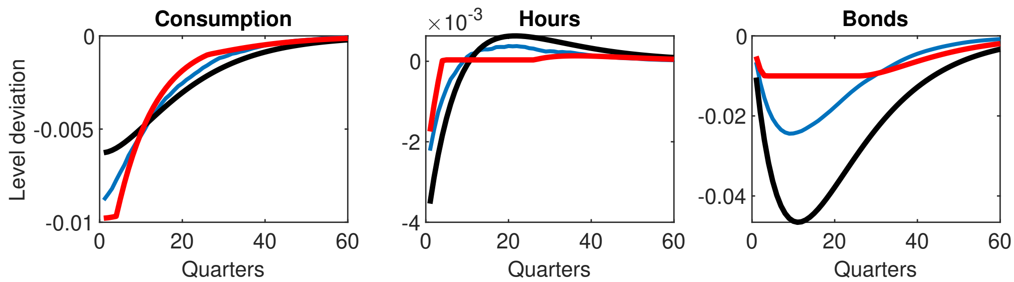

Note: The red line is the model with borrowing constraint, the black line is a model with unconstrained borrowing, and the blue line is the borrowing constraints model solved using projection methods. Plots show level deviation from the ergodic mean.

4.2 Stochastic extended path solution

When `extended_path()` is used in the `.mod` file, Dynare simulates the model using the stochastic version of the Fair and Taylor (1983) extended path algorithm as described in Adjemian and Juillard (2013). In this approach, as under perfect foresight, the dimension of approximation is on the distribution of shocks. Under perfect foresight, the approximation is that future shocks are known with certainty. Another approximation could be to assign zero probability weight to future non-zero shocks. This is the case at "order" zero, creating a stochastic simulation *as if* decisions are made conditional only on the current shocks being non-zero. Because agents do not expect future shocks, uncertainty plays no role in the model. By increasing the order, a higher number of periods of uncertainty are considered and the approximation error falls. For example, at order one, the algorithm allows for the expectation of non-zero shocks in next period but assumes agents believe no shocks will arrive in two periods' time; at order two, two future periods of non-zero shocks are expected, and so on. The approximation is solved by integrating over the shocks up to the specified horizon. Beyond this time horizon, the assumption is that agents believe all shocks will forevermore be zero and so consequently the model will ignore the long-run risk of hitting the bound.

To achieve accuracy with the stochastic extended path, we must integrate over all shocks of the model to a horizon of the same magnitude as the decay time of the model. As Dynare uses a Gaussian quadrature, which scales exponentially in the number of shocks and the order, this becomes infeasible except for very small models. In practice, when the command `extended_path(order=k,periods=T)` is used, Dynare first draws a random path for the exogenous variables over T periods before computing the corresponding path for endogenous variables, taking the steady state as starting point. The main options are:

- `periods = INTEGER`: This is the number of periods for which the simulation is to be computed. This must be set by the user as there is no default value.
- `order = INTEGER`: This is the order of the solution, that is, the number of future periods for which expectations are calculated.

Order set to zero would mean that every period, agents believe that only the current exogenous variables are non-zero and so are surprised by the shocks. Order equal to four would solve the model with agents believing that there can be shocks occurring in the next four periods but with no more future shocks five periods ahead or more.

The solution method builds on the underlying perfect foresight algorithm of the last section, the difference being that $\mathbb{E}_t\left[\{x_{t+s}\}_{s=1}^{k}\right]$ will not necessarily equal $\{x_{t+s}\}_{s=1}^{k}$. To compute the expectations, we follow the iterative procedure:

1. Start with the expectation horizon $j = 1$.
2. Beginning with expectations equal to the steady-state values $\mathbb{E}_t\left[x_{t+s}\right] = \bar{x}$ for $s = 0, ..., j$, obtain a new set of expectations by solving the non-linear model dynamically. Every period t, evaluate the expectations of the future state up to period $t + j$ using a Gaussian quadrature.

3. Setting the new expectations as the starting values, repeat step 2 until convergence.
4. Increase j by 1 and repeat steps 2–3.
5. Repeat up to $j = k$.

This captures the non-linearities of the system and differs from the standard stochastic simulation, which computes approximations to the transition and policy equations. Notably, there is no policy function to speak of, just simulated time series of the model variables from which moments can be computed.

To compute time series for the borrowing constraints model, we set the order $k = 0$ and periods $T = 10000$. Chart 3 shows the paths of consumption, hours, bond holdings and the multiplier, μ, over 100 periods. Note that borrowing is constrained during several periods of the simulation. When the borrowing constraint binds, households work more and cut consumption relative to the model without a borrowing limit, as they are unable to substitute as much consumption across time.

As before, we can compare moments using the simulated data. These are shown in Table 2. The table shows similar patterns to those observed above except the average saving is not as high as with uncertainty and the skewness of bonds is higher. This highlights the precautionary aspect of savings in the previous table, as opposed to just stemming

Chart 3: Extended path simulations (order = 0)

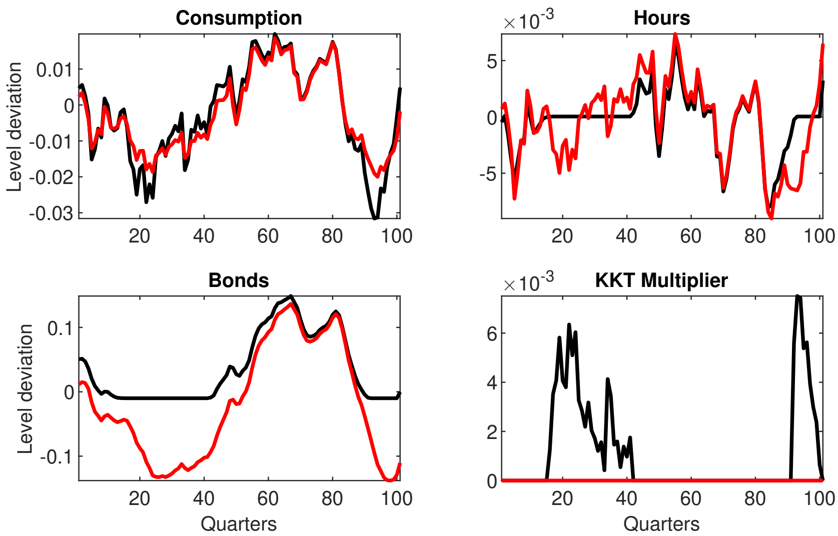

Table 2: Simulated moments

		Mean	Standard deviation	Skewness
No constraint	C	0.667	0.012	0.074
With OBC		0.667	0.013	-0.44
No constraint	H	0.667	0.003	-0.14
With OBC		0.666	0.002	-0.20
No constraint	B	0.004	0.10	-0.001
With OBC		0.046	0.06	1.45

Note: The table presents simulated moments of the model variables with and without the borrowing constraint using extended path approach. The variables are expressed in levels.

from the skewness introduced by the constraint in these results. At order zero, there is neither future uncertainty nor precautionary behaviour. Up to the period the constraint binds, the paths of the variables are identical because the agents ignore any risk of the constraint binding in future. To deal with this, we would need to increase the order and allow the agents to evaluate future uncertainty. It is possible to do this up to orders two to three, but it is prohibitively slow beyond this because the dimension of integration increases exponentially in the order.

5 Piecewise linear and regime switching

Another approach used to solve models with OBCs is to treat the economy as switching between regimes. For example, in the borrowing constraints model, one regime would be an unconstrained regime when $\mu_t = 0$, and a second would be a binding regime when $b_t = \underline{b}$. One can then approximate policy functions in each regime using a choice of method, typically perturbation. The resulting policy function within each regime will be a function of the policy function within all other regimes. This introduces some difficulties in both choosing the point around which the solution is perturbed and solving the probabilities of switching between different regimes. The transition probabilities will be endogenous and a function of the state. In the borrowing constraints model, for instance, the probability of switching to the borrowing constrained regime will clearly increase as b approaches $\underline{b}$.

One solution to these issues taken by Guerrieri and Iacoviello (2015) and implemented

Chart 4: Impulse response functions following technology shock using piecewise linear method

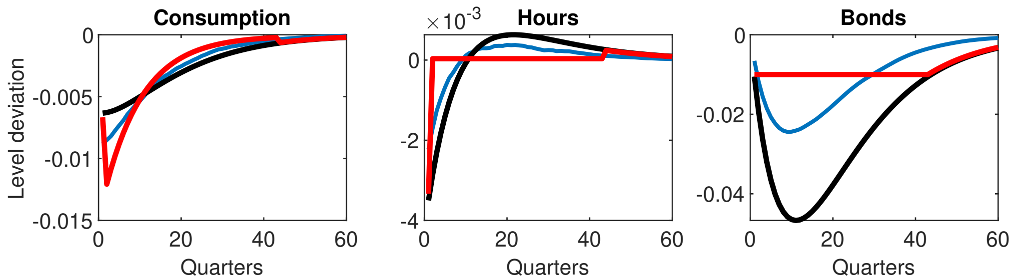

Note: The red line is the model with borrowing constraint, the black line is a model with unconstrained borrowing, and the blue line is the borrowing constraints model solved using projection methods. The plots show level deviation from the ergodic mean for projection and from the deterministic steady state for piecewise linear.

in the OccBin tool kit[12] is to take a first-order approximation to each regime around the deterministic steady state. One regime is considered the "reference" regime, which would likely be an absorbing equilibrium and would need to be saddle-path stable. In the case of the borrowing constraints model, the non-binding constraint would represent the reference regime and the constrained borrowing would represent the second regime. We then abstract from the transition probabilities and their effect on precautionary behaviour, as the policy function implied by the linearized solution to the reference regime ignores the possibility of the constraint binding. When the model is simulated and the constraint binds, the regime switches. The path of the economy is solved subject to the condition that agents do not expect any future shocks; in other words, they act as if under perfect foresight, much like the order-zero extended path simulation described above. IRFs of technology shocks are shown in Chart 4. The approach faces similar drawbacks to that of the perfect foresight simulation, as precautionary saving does not occur. As a result, there is an extended period spent at the borrowing constraint and an exaggerated effect on consumption.

Maih (2015) proposes a more general solution that is implemented in MATLAB in the RISE tool kit.[13] Rather than treat the presence of the constraint as a surprise, the modeller defines the transition probabilities either as constants (which is the natural option in a Markov-switching parameter model) or as functions of the state for models

[12] Available at https://github.com/lucaguerrieri/occbin.

[13] Available at https://github.com/jmaih/RISE_toolbox.

Chart 5: Impulse response functions following technology shock using regime-switching approach

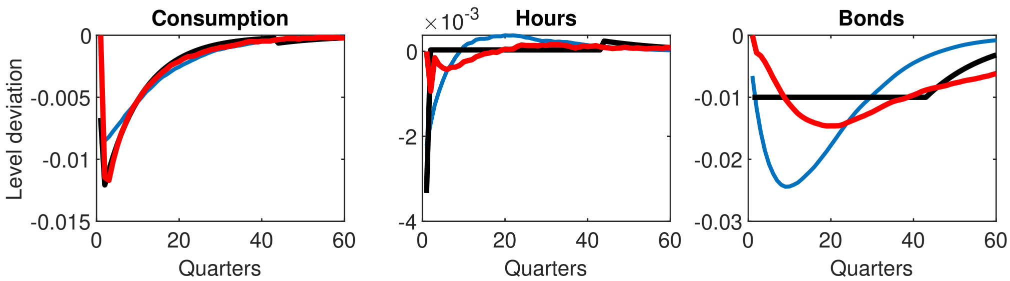

Note: The red lines are the regime switching simulations. The blue lines are impulse response functions solved using projection methods, and the black lines are those solved with the piecewise linear algorithm. The plots show level deviation from the ergodic mean.

of OBCs in which regime switching is endogenous. The modeller chooses the functional form and parametrization of these probabilities, and it is feasible to write an iterative process to simulate-update-simulate-update and so on, until the probabilities are close to the true solution under rational expectations. The policy functions in each regime are then perturbed around regime-specific steady states as if the economy will stay in that regime forever. In addition to allowing for endogenous switching in expectation, the policy functions can be approximated up to the fifth order. While still having drawbacks, this more general approach improves on the piecewise linear solution. It relies on the user defining sensible transition probabilities, and there may be some approximation error as a result of the fixed point being far from the true solution. But it is an interesting and useful approach, particularly from the perspective of different policy regimes. Suppose that nominal interest rates are at zero, or the economy is in a constrained "crisis" regime, and the policy-maker uses alternative rules and policy programs. In such cases, the regime-switching method could be a particularly useful tool, with differing policy rules determining macroeconomic outcomes. Impulse responses to the same shock as shown for previous methods are displayed in Chart 5. In contrast to the piecewise linear approach through which we solved perfect foresight IRFs, generalized or average IRFs are computed via Monte Carlo simulation using RISE. The results seem to improve on piecewise linear approach, and the remaining differences between the projection method and regime switching are likely to be largely related to the parametrization of the transition probability functions and so illustrate the difficulty of achieving good accuracy.

6 The news shock method

An alternative approach, and one that I focus on in this review, is to treat the constraint as an endogenous source of news, ensuring that where disturbances would cause bounds to be violated, anticipated news shocks return the bounded variable to the constraint. The method is described fully in Holden (2016), with a companion paper, Holden (2019), discussing the necessary and sufficient conditions for the existence and uniqueness of solutions at the bound.[14] The method is implemented in the tool kit "DynareOBC."

Just as one can view the stochastic extended path algorithm as building upon an underlying perfect foresight algorithm, at the core of Holden's (2016) description is the perfect foresight solver with an extended path approach used to integrate over future uncertainty. Below, I discuss the perfect foresight solver and the algorithm used to impose the bound before describing the extended path method used. I finish with a discussion of the conditions that guarantee the existence and uniqueness of a solution and look at some examples of multiple equilibria in the NK model.

6.1 Layer 1: The perfect foresight solver

At the heart of the Holden (2016) algorithm is a fast perfect foresight solver. The speed is necessary, as the solver will be invoked many times in the algorithm. Recalling equation (19), implementing the borrowing constraint under perfect foresight involves ensuring that $b_t = \underline{b}$ during periods when borrowing would otherwise violate the constraint, and having borrowing determined by the Euler equation otherwise. The method outlined below endogenously determines when the constraint will bind by introducing anticipated news shocks to b_t when it would otherwise violate the constraint, ensuring that $b_t = \underline{b}$.

I will describe the method in a simple linear model with a single constraint from which it generalizes to higher-order pruned perturbation and multiple bounds.[15] Consider the basic problem in computing the IRF under a perfect foresight simulation, that is, with

[14] An earlier version of the idea was first described in the appendix to Holden (2010) and further developed in Holden and Paetz (2012) prior to these papers.

[15] Pruning is a way of dealing with explosive roots that can occur in models approximated to orders greater than one (see e.g., Andreasen et al. 2018). The idea is to substitute lower-order approximations into higher-order roots, i.e., to use a first-order approximation in second-order terms.

no other future shocks. Let us write the model as

$$B\left(x_t-\mu\right)=A\left(x_{t-1}-\mu\right)+C\mathbb{E}_t\left(x_{t+1}-\mu\right)+D\varepsilon_t, \tag{25}$$

where $\mathbb{E}_{t-1}\varepsilon_t=0$ and $\varepsilon_t=0$ for $t>1$. x_t is a vector of model variables, ε_t a vector of shocks, and μ a vector of constants where the ith element of μ is the steady-state value of the ith element of x_t. x_0 is given as an initial condition and we assume a terminal condition $x\to\mu$ as $t\to\infty$ holds. For $t>1$, this becomes

$$\begin{aligned}B\left(x_1-\mu\right)&=A\left(x_0-\mu\right)+C\left(x_2-\mu\right)+D\varepsilon_1\\ B\left(x_t-\mu\right)&=A\left(x_{t-1}-\mu\right)+C\left(x_{t+1}-\mu\right), \quad \text{for } t>1,\end{aligned} \tag{26}$$

where the expectation operator and future shocks disappear because agents know $\varepsilon_t=0$ for $t>1$. As in the earlier perfect foresight section, solving impulse responses without a bound means finding a path of $x_t\in\mathbb{R}^n$ that satisfies (26).

But suppose there is a ZLB on variable $x_{1,t}$, where the subscript indicates it is the first element of vector x_t. The first row of (26) becomes

$$x_{1,t}=\max\left\{0,\mu_1+A_1\left(x_{t-1}-\mu\right)+B_1\left(x_t-\mu\right)+C_1\left(x_{t+1}-\mu\right)+D_1\varepsilon_t\right\}, \tag{27}$$

with the matrix subscripts indicating row one. If we consider a shock that would cause the bound to be violated, we want to find a path for x_t that satisfies equations (25) and (27), where $x\to\mu$ as $t\to\infty$.[16] As the model returns to steady state asymptotically, there is some horizon T by which the constraint will no longer be violated.

News shocks will be used to impose the bound, and so equation (27) is rewritten as

$$x_{1,t}=\mu_1+A_1\left(x_{t-1}-\mu\right)+B_1\left(x_t-\mu\right)+C_1\left(x_{t+1}-\mu\right)+D_1\varepsilon_t+y_{1,t}, \tag{28}$$

where $y_{i,t}$ is a news shock known at period i that hits at period $i+t-1$. $y_{1,t}$ can be positive during periods $t\le T$; but for $t>T$, the news shocks are known to be $y_{1,t}=0$. The problem of the algorithm becomes finding a path for $x_t\in\mathbb{R}^n$ and $y_1=\begin{bmatrix} y_{1,1} & y_{1,2} & \cdots & y_{1,T}\end{bmatrix}'\in\mathbb{R}^T$, given initial conditions x_0, that satisfies

$$B\left(x_1-\mu\right)=A\left(x_0-\mu\right)+C\left(x_2-\mu\right)+D\varepsilon_1+I_1'y_{1,t}, \tag{29}$$

[16]This rules out switching to another steady state. The introduction of OBCs can introduce multiple steady-state equilibria. In the ZLB example with inflation targeting, for instance, Benhabib et al. (2001) highlight a "bad" deflation steady state in the New Keynesian ZLB to which the economy might converge. Ruling out multiple steady states in DynareOBC in this case implies the long-run credibility of the policy-maker.

which modifies the first row of (25) to equation (28).

6.1.1 The news shock problem as a linear complementarity problem

As the model is linear, DynareOBC exploits the fact that the impulse response to two shocks is equal to the sum of the impulse response to each individual shock. Consider first the path of $x_{1,t}$ given an (arbitrary) vector of news shocks $y_0 \in \mathbb{R}^T$. Let $m_k \in \mathbb{R}^T$ be a column vector with the impulse response of $x_{1,t}$ to a news shock of size 1 at period k with $x_0 = \mu$, and let

$$M \equiv \begin{bmatrix} m_1 & m_2 & \cdots & m_T \end{bmatrix} \tag{30}$$

horizontally stack these relative IRFs. It follows that the path of $x_{1,t}$ given $x_0 = \mu$ and that an arbitrary vector of new shocks y_0 is given by My_0.

Let $q \in \mathbb{R}^T$ be the path of $x_{1,t}$ up to period T that satisfies equation (26)—that is, the model without the constraint—given any x_0. The path of $x_{1,t}$ for any x_0 and $y_0 \in \mathbb{R}^T$ that satisfies equation (26) can be written as

$$q + My_0. \tag{31}$$

Let us get specific about how the news shocks can be utilized to impose the bound. For a given x_0, we must find a vector y_0 and path for x_t that satisfy equation (26) and the ZLB on $x_{1,t}$. The path for $x_{1,t}$ will be given by equation (31). The news shocks are used only to impose the bound, so when $x_{1,t} > 0$, $y_{1,t} = 0$, this implies firstly that the solution must satisfy

$$y_0 \circ (q + My_0) = 0, \tag{32}$$

and secondly that the news shocks can only push the variable *up* to the bound, that is, $y_0 \geq 0$. Finally, the solution must impose the bound: so $q + My_0 \geq 0$.

The news shock problem is then characterized as a linear complementarity problem LCP(q, M) (see Cottle (2009)): for a given q and M, the LCP(q, M) finds $y \in \mathbb{R}^T$ to satisfy

$$\begin{aligned} q + My &\geq 0 \\ y &\geq 0 \\ y \circ (q + My) &= 0. \end{aligned} \tag{33}$$

The linear complementarity problem is then solved by representing it as a mixed integer linear programming problem for which there are many commercial solvers available. A convenient property of this representation is that it allows us to numerically test whether there is a unique solution and whether there is a solution at all. This contrasts with standard perfect foresight solvers. For those, the failure of the solver tells nothing about whether there is a solution, since the solver may have just become stuck in a local minimum. Additionally, standard perfect foresight solvers may return "non-fundamental" solutions even when fundamental solutions exist, and there is no way to either force the solver to always return a fundamental solution or easily tell whether a fundamental solution has even been returned.

The necessary and sufficient conditions for the existence and stability of a solution depend on the matrix M. These are outlined in detail in Holden (2019).

6.1.2 Some generalizations

The problem extends to multiple bounds. For n bounds, the vectors $q \in \mathbb{R}^{nT}$ and $y \in \mathbb{R}^{nT}$ stack the impulse responses, ignoring the bound for each bounded variable and the impulse response to the news shocks, respectively. Matrix $M \in \mathbb{R}^{nT \times nT}$ is a block matrix where block $M_{i,j}$ is the response of variable x_i to the news shock on variable x_j.

The problem as defined above relates to a ZLB constraint that does not bind in steady state. When either the bound is not at zero, the OBC is an upper bound, or the OBC is binding in steady state, we can amend the equation to generalize into the required ZLB form: that is,

$$z_{1,t} = \max\left(z_{2,t}, z_{3,t}\right) \quad \text{becomes} \quad z_{1,t} - z_{2,t} = \max\left(0, z_{3,t} - z_{2,t}\right) \tag{34}$$

$$z_{1,t} = \min\left(z_{2,t}, z_{3,t}\right) \quad \text{becomes} \quad -z_{1,t} = \max\left(-z_{2,t}, -z_{3,t}\right). \tag{35}$$

6.1.3 Simulation and higher-order approximations

In the first-order (linear) case, the model was first approximated around the steady state, yielding a model approximation of the form in equation (25); while at higher orders, DynareOBC takes a pruned perturbation approximation using the Lan and Meyer-Gohde (2013) algorithm. The algorithm as described readily extends to higher-order approximations when the shadow shocks are of the form y^d, since pruned perturbation approximations of order d are linear in shocks to the power of d.

6.2 Layer 2: The stochastic extended path type solver

Because $y_t \geq 0$, the news shocks cannot be mean zero: so $\mathbb{E}_t\left[y_{t+s,k}\right] > 0 \forall s, k$. Furthermore, the expected value of y will depend on the current state. If this is ignored, bias will be introduced in the expectation of the bounded variable. This could be particularly detrimental if the constraint is likely to bind frequently. To extend to stochastic simulation, the perfect foresight solver described in the previous section is embedded in an extended path type solver. Recall that the order of integration increased exponentially in the number of shocks and the number of periods of uncertainty using the Adjemian and Juillard (2013) stochastic extended path algorithm. Now that we are working with a pruned perturbation approximation to the model, the process of integrating over future uncertainty is greatly simplified.

In particular, due to some nice properties of pruned perturbation solutions, we can derive a closed-form formula for the covariance of the expected future path of the bounded variables in the absence of the bound. Using this, we can take a Gaussian approximation to the future distribution of the bounded variables in the absence of the bound and then integrate over this distribution using Gaussian cubature techniques. In this way, Dynare-OBC only has to solve the perfect foresight problem a number of times that is polynomial in the periods of uncertainty, independent of the number of shocks. This means that we can readily integrate over many periods of uncertainty with minimal computational cost. Additionally, providing we take a higher-order perturbation approximation, we capture the effects of even long-run risk not stemming from the bound for free, again in contrast to the Adjemian and Juillard (2013) algorithm.

A further difference is that rather than assuming that the news shock variance goes to zero beyond the horizon S, we use a cosine windowing function to scale the shock variance. Particularly, if the covariance matrix is Σ, the covariance matrix used when considering uncertainty at horizon k is given by

$$\hat{\Sigma}_k = \frac{1}{2}\left[1 + \cos\left(\pi \frac{\min\{k-1, S\}}{S}\right)\right] \Sigma. \tag{36}$$

The derivation of the covariance of the expected future path of the bounded variables ignoring the constraint is unchanged, using the modified covariance matrix $\hat{\Sigma}_k$ instead.[17] Let Ω denote this covariance matrix, and let $w_{t,t+i}$ be the value the bounded

[17] See appendix E in Holden (2016) for details.

variable would take at time $t+i$ if the bound no longer applied after time t. Then $\left[\begin{array}{ccc} w_{t,t+1} & \cdots & w_{t,t+S} \end{array}\right]'$ is assumed to be normally distributed, which it will be at first order and is a close approximation at higher orders. Using a Schur decomposition of Ω will give $\Omega = UDU'$, where $D \geq 0$ will be a diagonal matrix. By setting any very low values of D to zero, we reduce the cost of integration, only scaling in $\hat{S} < S$, where $\hat{S}$ is the number of remaining non-zero elements of D. Using this step and the normal assumption, we arrive at an approximation of the distribution of $\left[\begin{array}{ccc} w_{t,t+1} & \cdots & w_{t,t+S} \end{array}\right]'$ of $\mathbb{E}_t \left[\begin{array}{ccc} w_{t,t+1} & \cdots & w_{t,t+S} \end{array}\right]' + \Lambda\zeta$, where $\Lambda \equiv U_1\sqrt{D_{11}}$, with $D_{11} \in \mathbb{R}^{\hat{S}}$ denotes the matrix block in D that includes the $\hat{S}$ non-zero elements and U_1 the corresponding elements of U, and where $\zeta \sim \mathcal{N}\left(0, I_{\hat{S}}\right)$. This simplifies the integration problem to one of integrating over $\hat{S}$ standard normals.

The tool kit uses three alternative methods to integrate over future uncertainty: a "fast" degree 3 monomial rule; the Gaussian cubature rule proposed in Genz and Keister (1996); and the Quasi-Monte Carlo method. The first is with $2\hat{S}+1$ nodes and equal positive weights that provide a robust approximation, giving reasonable accuracy despite fast simulation times. The second is a $2K+1$ monomial rule with $O\left(\hat{S}^K\right)$ nodes, where $K \leq 25$ can be chosen by the modeller. The rule will feature negative weights if $K > 0$ and $\hat{S} > 1$, which might prevent the upward bias that positive weights can introduce by evaluating the integral far from the steady state. The higher degree is also likely to give greater accuracy, and because the rule is nested, we can use an adaptive integration degree that will improve computation times. The final method uses a Sobol sequence (Sobol 1967) to generate $\left(2^{1+l}-1\right)$ points, where $l \in \mathbb{N}$ is chosen by the modeller. This approach should be the most accurate; however, for well-behaved functions it will require a far greater number of evaluations of the integrand than the Genz and Keister (1996) approach for similar accuracy.

Which of the latter two methods delivers the most desirable mix of speed and accuracy depends on the model. As a rule of thumb, if the bound binds very frequently, Genz and Keister's (1996) cubature is likely to dominate, whereas when the constraint binds with moderate probability, quasi-Monte Carlo may be better. Whichever approach is chosen, the result is an expected value of the vector y needed to impose the bound. These news shocks will impact the expected value of x_{t+1} by moving up the expectation of the bounded variable $\mathbb{E}_t\left[x_{1,t+1}\right]$.

6.3 Impulse response functions in DynareOBC

When cubature is *not* used, at first order the IRFs are simply perfect foresight simulations. At higher orders, Dynare computes average IRFs using Monte Carlo simulation, while DynareOBC uses the closed-form covariance matrices due to the pruning algorithm so computes average IRFs without repeated samples. In the presence of OBCs, DynareOBC will compute IRFs, adding the news shocks to impose the constraint. This is as if agents thought the expected sequence of news shocks were known with certainty, even if future shocks were to cause the constraint to be violated and different news shocks to materialize.

With cubature, DynareOBC integrates out the effects of future shocks on whether the model is at the bound in future, but it does not integrate out the effect of current uncertainty on the bound. This means that while the base path takes into account current uncertainty, current uncertainty is ignored when calculating the expected news shocks. DynareOBC offers an improvement with a Monte Carlo simulation that provides a true average given the behavioural approximation. Combining with cubature to calculate agent expectations will provide IRFs of very high accuracy, but at some cost in speed.

7 Using DynareOBC to solve models with OBCs

The algorithms are designed to be easily implemented into Dynare. The OBC tool kit, DynareOBC, needs to be downloaded and added to the MATLAB path. The tool kit is available from:

https://github.com/tholden/dynareOBC

The recommended download process is as follows:

1. Download and install GitHub Desktop, from https://desktop.github.com/.
2. Click on the "Clone or download" button on the DynareOBC website.
3. Click "Open in desktop."

This method allows the user to view change history and contribute fixes. The alternative method is to download a DynareOBC release in a zip file either from https://github.com/tholden/dynareOBC/releases or by clicking "Download ZIP" instead of

"Open in desktop" in the steps. Once downloaded, the root DynareOBC folder, containing `dynareOBC.m`, needs to be added to the user's MATLAB path.

The user can read information on the tool kit by typing `dynareOBC` into the MATLAB command window. The first step is to test the Dynare installation by typing the following:

```
dynareOBC testsolvers
```

This will attempt to solve a number of computational problems and check whether the required solvers are working correctly.

7.1 The basics of DynareOBC usage

The user simply has to write the bounded variable in the .mod file—as we have above. For example, in our borrowing constraints model, we enter the constraint as we did when using the perfect foresight solver:

```
0 = min(mu,b-b_limit);
```

and use the standard `stoch_simul()` command. The model is then run from the MATLAB command window with the command `dynareOBC [FILENAME].mod [OPTIONS]`. Rather than running Dynare with the .mod file as input, MATLAB runs DynareOBC, which solves the model using Dynare but searches for non-differentiable functions and performs the necessary steps in the algorithm to include the inequality constraint. Note that not all `stoch_simul` options are supported by DynareOBC, and currently no warning will be generated if unsupported options are used. Details are found by typing `dynareOBC` into the command window and hitting enter.

To solve the borrowing constraints model, the user can try typing

```
dynareOBC soe_borrowing_constraint shockscale=-2
```

from the command window when in the **Model 1/news-shocks/** directory of the accompanying code.[18] This will try to simulate the model without using cubature. If the user chooses `stoch_simul(irf=40)` in the .mod file, IRFs will be computed, and if, for

[18] Available at `https://github.com/jonswarbrick/OBC_example_code`.

Chart 6: Impulse response functions following technology shock using news shock perfect foresight solver

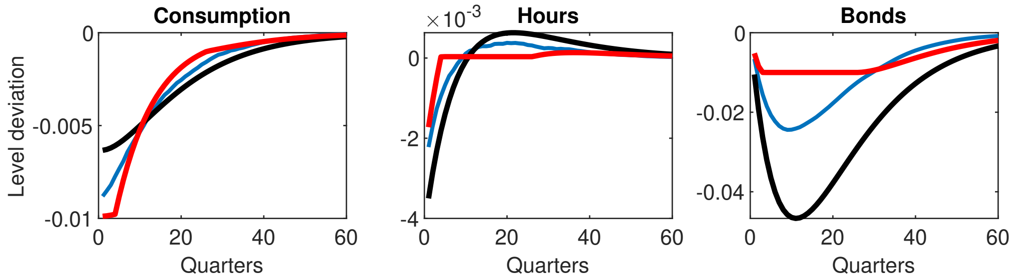

Note: The red line is the model with borrowing constraint, the black line is a model with unconstrained borrowing, and the blue line is the borrowing constraints model solved using projection methods. Plots show level deviation from the ergodic mean.

example, they also include `periods=1000`, then the model will use the extended path type method to simulate the model over 1,000 periods. The option `shockscale=−2` scales the technology shock by a factor of -2 but does not affect the solution or simulation. Note that when the user lists options after the model file names that have '=', the user must make sure to not include spaces around the equal sign. The options are not case sensitive.

After running the code above, the user will see the following error: `Impossible problem encountered. Try increasing TimeToEscapeBounds, or reducing the magnitude of shocks`. This error occurs when there is definitely no solution to the bounds problem. This could mean there is no solution in general, which contrasts with other methods discussed, where the failure to converge to a solution might simply reflect a failure of the algorithm. However, in this case, it is because there is no solution in which b is again unconstrained within horizon T. This can be fixed by entering the following:

```
» dynareOBC soe_borrowing_constraint shockscale=−2
  timetoescapebounds=40
```

Plots are shown in Chart 6. The results are virtually identical to those using Dynare's built-in perfect foresight solver.

Users will notice that when calling DynareOBC, Dynare is called three times and prints output similar to the screen. Diagnostics about the existence and uniqueness of the solution will appear below this. (For further technical detail on these conditions, see

Holden (2019).) In addition to the `oo_.` and `M_.` structures, DynareOBC stores results and data in a structure, `dynareOBC_.`. For instance, the IRFs without the bounds are stored in `dynareOBC_.IRFsWithoutBounds`, and the M matrix is stored in `dynareOBC_.MMatrix`.

As discussed above, several integration techniques are available to evaluate future uncertainty. To use the Genz and Keister (1996) method with a maximum degree 7 polynomial, the user can enter the following command:

```
dynareOBC soe_borrowing_constraint shockscale=-2
timetoescapebounds=40 GaussianCubatureDegree=7
```

Changing the integer to something less than or equal to 51 will set the maximum degree of integration. The periods of uncertainty given by S in the description above is 16 by default and is set using `PeriodsOfUncertainty=16`. Another option that might be useful is `CubatureSmoothing=0.1`; since the Genz and Keister (1996) rules sometimes oscillate from too high to too low as the degree increases, this option can be used to smooth the results of adjacent degrees by setting to a float in the unit interval. To speed up simulation, DynareOBC will not integrate to the full degree if adjacent rules are sufficiently close; the choice of `CubatureTolerance` defines sufficiently close. Further options to improve computational times of cubature include `CubaturePruningCutOff=0.01`, where the number is a float on the unit interval. Eigenvalues of the covariance matrix of the distribution from which we integrate that are below this number times the maximum eigenvalue are pruned to zero, in order to increase integration speed. The maximum dimension over which to integrate is set using `MaxCubatureDimension=128`, where the number is an integer. If the algorithm needs to integrate over a larger space, it will keep only the largest eigenvalues of the covariance matrix, "pruning" all remaining to zero. This option sets the maximum number of eigenvalues to keep. To use the degree 3 monomial rule, the user can enter the following:

```
dynareOBC soe_borrowing_constraint shockscale=-2
timetoescapebounds=40 FastCubature
```

This tends to give much faster computation speeds and is often of competitive accuracy. The quasi-Monte Carlo can be invoked by using:

```
» dynareOBC soe_borrowing_constraint shockscale=-2
  timetoescapebounds=40 QuasiMonteCarloLevel=8
```

where `QuasiMonteCarloLevel` is the parameter l from the description where a Sobol sequence with $2^{1+l} - 1$ points is used.

Finally, as mentioned at the end of the previous section, although using cubature captures the effect of future uncertainty of the constraint, it misses current uncertainty. To improve on this, Monte Carlo simulation can be performed to compute generalized IRFs (GIRFs) and capture the average effect of news shocks on the current period, which should have important level effects on the IRFs.[19] To compute GIRFs, the user invokes the option `SlowIRFs`. The standard Dynare option `replic=50` sets the number of draws, as it does with normal Dynare. Ideally, to capture the precautionary effects of the bounds, the user should invoke Monte Carlo simulation with cubature. In practice, I find that Monte Carlo on its own seems to deliver very good accuracy with appealing computation times. All DynareOBC options are printed to screen when `dynareOBC` is entered in the command window with no model file or other argument listed.

Chart 7 shows IRFs from the borrowing constraints model, comparing across the Dynare-OBC integration options and the projection method solution. All solutions with cubature improve on the shape but miss the magnitude of the response. This result is perhaps due to the level around which we solve IRFs. Chart 8 shows IRFs solved using only quasi-Monte Carlo integration, using both this and Monte Carlo, and using only Monte Carlo together with those solved with the projection method. The result is impressive. When Monte Carlo simulation is used, DynareOBC produces IRFs that almost exactly match those solved using projection methods.

As well as computing IRFs, the extended path method can be used to simulate the model over a large number of periods to evaluate simulated moments and cross-correlations. Perhaps this is something the user can attempt: to compare simulated moments with and without integration and across different integration methods. Without integrating over future uncertainty, the simulation will be similar to that offered by the extended

[19]Note that at first order without the bound, Monte Carlo is not required, as the effects of shocks are symmetric and independent of shock size and the current state. This is no longer true at higher orders of approximation, and so Monte Carlo simulation is required to compute generalized IRFs with or without the bound.

Chart 7: Comparison of impulse response to a technology shock in the borrowing constraints model using different solution methods

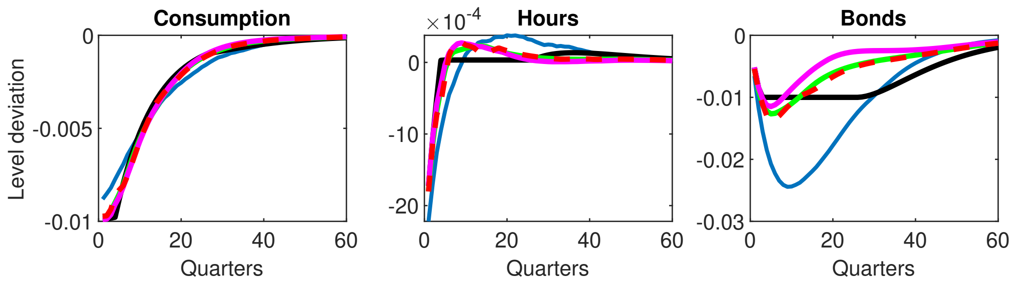

Note: The blue line is projection method and the remaining news shocks method; the black line is perfect foresight; the red line is a fast degree 3 monomial; the magenta line is the Genz and Keister (1996) cubature rule; and the green line is quasi–Monte Carlo.

Chart 8: Comparison of impulse response to a technology shock in the borrowing constraints model using different solution methods

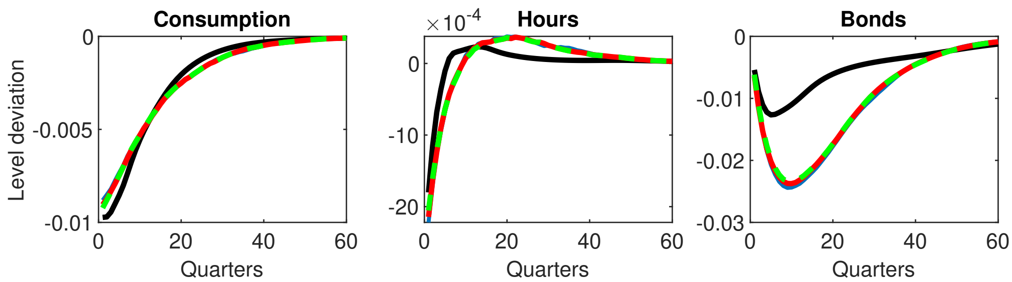

Note: The blue line is projection method and the remaining news shocks method; the black line is quasi-Monte Carlo (QMC); the red line is QMC and Monte Carlo; and the green line is just Monte Carlo.

path at order zero. As discussed above and in Holden (2016), the DynareOBC simulation without integration notably improves the accuracy and speed over the existing extended path method in Dynare.

7.2 Determinacy and multiple equilibria in the NK model with a ZLB

Another consideration is that a unique solution to the LCP in equation (33)—giving a path for the news shocks—can only be found if the "M" matrix exhibits a number of properties (see Holden 2019). Indeed, these conditions are often not satisfied. In this case, there could be multiple equilibria consistent with the true RE solution to the model. To gain some intuition, suppose that the expectation in one quarter of the ZLB binding in the next is deflationary; then this expectation could lead to a self-fulfilling ZLB episode as the interest rate is cut in response. So, as with sunspot shocks, there is the possibility of non-fundamental shifts in expectations. When there are multiple solutions, DynareOBC always chooses the one that minimizes a linear combination of the L_1-norm of the shadow-shocks and the L_1-norm of the return path. As mentioned earlier, the LCP is expressed as a mixed integer linear programming (MILP) problem. The specific form the MILP takes is as follows: given $\tilde{\omega} > 0$, $q \in \mathbb{R}^T$ and $M \in \mathbb{R}^{T \times T}$, we find $\alpha \in \mathbb{R}$, $\hat{y} \in \mathbb{R}^T$ and $z \in \{0,1\}^T$ that satisfy

$$0 \leq \alpha \tag{37}$$

$$0 \leq \hat{y} \leq z \tag{38}$$

$$0 \leq \alpha q + M\hat{y} \leq \tilde{\omega}\left(1_{T\times 1} - z\right). \tag{39}$$

The solution we are interested in, if it exists, will be $y = \frac{\hat{y}}{\alpha}$. Given this form of problem, $\tilde{\omega}$ controls the linear combination that is minimized to choose a solution. As $\tilde{\omega} \to 0$, the solution given will be one that minimizes $\|q + My\|_{\infty}$, while as $\tilde{\omega} \to \infty$, the solution given will be one that minimizes $\|y\|_{\infty}$. DynareOBC sets $\tilde{\omega} = \omega\, \|q\|_{\infty}$, where the modeller chooses ω with the option `omega=10000`, for example. As a rule of thumb, choosing a high value such as 10,000 will find a solution minimizing time at the bound, whereas a low value, 0.0001 for instance, will select a solution with a longer time at the bound. To gain some intuition, let us suppose we have a model with a shock causing a ZLB to bind. If we believe the bound will be escaped after two periods, the linear complementarity problem is to find a solution,

$$y = \begin{bmatrix} y_1 \\ y_2 \end{bmatrix}, \tag{40}$$

that satisfies

$$y_1\,,\,y_2 \geq 0 \tag{41}$$

$$(q_1 + M_{11}y_1 + M_{12}y_2)\,,\,(q_2 + M_{21}y_1 + M_{22}y_2) \geq 0 \tag{42}$$

$$y_1\,(q_1 + M_{11}y_1 + M_{12}y_2) = y_2\,(q_2 + M_{21}y_1 + M_{22}y_2) = 0. \tag{43}$$

These impose that the news shocks must be positive (41), that the bound is not violated (42) and that the constraint must be binding when there are non-zero news shocks (43). These conditions imply two quadratics, so up to four solutions:

1. $y_1 = y_2 = 0$. This solution exists if $q_1 \geq 0$ and $q_2 \geq 0$.
2. $y_1 = -\frac{q_1}{M_{11}}, y_2 = 0$. This solution exists if $\frac{q_1}{M_{11}} \leq 0$ and $M_{11}q_2 \geq M_{21}q_1$.
3. $y_1 = 0, y_2 = -\frac{q_2}{M_{22}}$. This solution exists if $\frac{q_2}{M_{22}} \leq 0$ and $M_{22}q_1 \geq M_{12}q_2$.
4. $y_1 = \frac{M_{12}q_2 - M_{22}q_1}{M_{11}M_{22} - M_{12}M_{21}}, y_2 = \frac{M_{21}q_1 - M_{11}q_2}{M_{11}M_{22} - M_{12}M_{21}}$. This solution exists if $M_{11}M_{22} - M_{12}M_{21} \leq 0$.

That is, there are multiple solutions for at least some combination of $q_1, q_2 \geq 0$ if

1. $M_{11}, M_{21} \leq 0$, or
2. $M_{22}, M_{12} \leq 0$, or
3. $M_{11}M_{22} - M_{12}M_{21} \leq 0$.

In the NK model, the first of these occurs if a positive monetary policy shock, or news of one, can cause interest rates to fall. Perhaps news of a shock occurring in the subsequent period might be expected to cause a cut today, but it is also possible that a positive shock today can cause a fall in rates depending on the coefficients in the Taylor rule. Suppose the rule reacts strongly to deviations of output and inflation; then the response of the economy to a monetary shock might cause an even larger policy reaction in the opposite direction. Absent the ZLB, there is a well-discussed literature on the determinacy in the NK model with inflation targeting; the so-called Taylor Principle implies that, providing the nominal interest rate is adjusted more than one-for-one with inflation, the model

Chart 9: Comparison of impulse response to a Euler equation shock in the Brendon et al. (2020) NK model

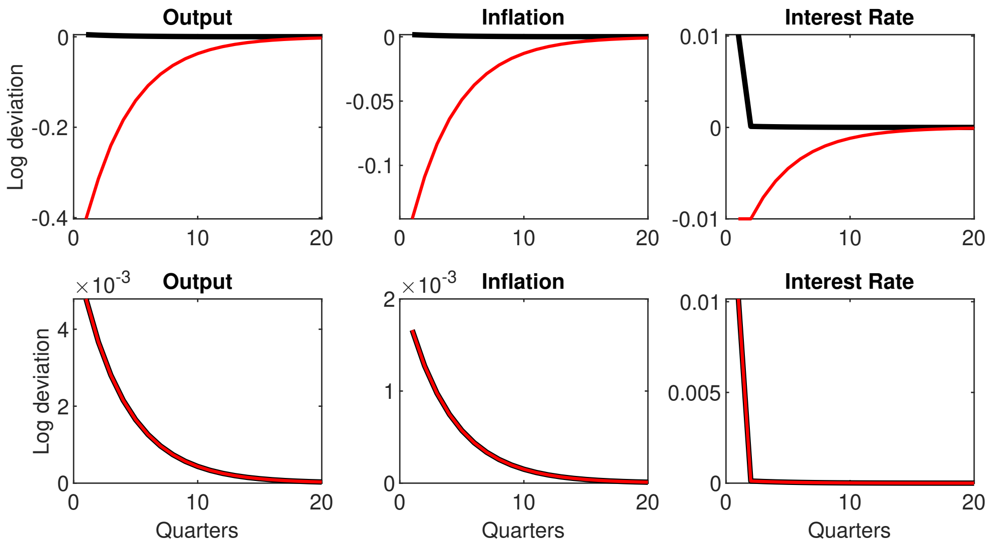

Note: The top panel is with omega = 10,000, bottom with omega = 0.0001. The black line is without a ZLB, and the red line is the ZLB model.

has a unique solution (see Taylor 1999; Woodford 2001). Even when the "bad" long-run equilibrium is ruled out, Holden (2019) shows that in the presence of the ZLB there is no unique perfect foresight path regardless of whether the Taylor Principle is satisfied. Holden finds that replacing inflation targeting with price-level targeting restores determinacy. As an example, Chart 9 shows impulse responses to a shock in the Euler equation using the NK model from Brendon et al. (2020). Without the ZLB, the shock causes a rise in economic activity and so an interest rate rise in response. With the ZLB present, this is still a possible equilibrium path; but just the presence of the constraint opens up the possibility of a second path in which economic activity falls sharply and interest rates drop to the ZLB.

There are further options to help select between equilibria. These include `FullHorizon`, `ReverseSearch` and `SkipFirstSolutions=INTEGER`. By default, DynareOBC searches for a solution that minimizes time at the bound; it does so by trying to find a solution by iterating through the exact length of time at the bound, T^*, first trying $T^* = 1$, then $T^* = 2$, $T^* = 3$ and so on. Under this default behaviour, solutions are typically found with a minimal time spent at the bound, regardless of the choice of `omega`. When

Chart 10: Comparison of impulse response to a discount factor shock in the Smets and Wouters (2003) NK model

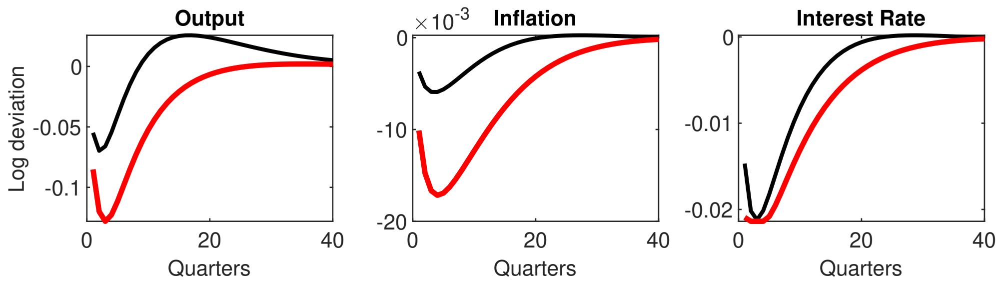

Note: The black line is first solution found; the red line is the solution found when `SkipFirstSolutions=1` is used.

`ReverseSearch` is invoked, DynareOBC instead begins from the maximum T^* defined in `TimeToEscapeBounds`, i.e., $T^* = T$, then $T^* = T - 1$ and so on. With this option, we are more likely to find solutions where the constraint binds for a longer period. The full horizon option searches for a solution for any $T^* \in (0, \bar{T})$. The `SkipFirstSolutions=INTEGER` ignores the first `INTEGER` solutions it finds. If fewer solutions are found than the number requested, then the last found will be returned. This last option is demonstrated first in Chart 10, which shows the two solutions found for IRFs to a shock to the discount factor using the Smets and Wouters (2003) model. In the second, the economy spends a little longer at the ZLB but exhibits much larger drops in aggregate demand. Chart 11 shows the response to a large combination of shocks using the Smets and Wouters (2007) version. In the first, the economy spends just a couple of periods at the ZLB, whereas in the second solution found, conditional on exactly the same fundamental shocks, the economy spends over three years at the bound and there is a very deep fall in economic activity and inflation.

It is not clear that there is a general rule on placing probabilities on the occurrence of particular equilibrium paths. I would argue that ignoring the presence of alternative equilibrium paths is the least desirable option, but particular paths could also represent unlikely outcomes. For example, consider an NK model that goes to the ZLB for a couple of periods following a shock. Now suppose that an alternative model-consistent equilibrium path leaves the economy at the bound for 100 periods. One might rule this out on the basis that it would represent significant coordination failure. Some forward guidance

Chart 11: Comparison of impulse response to a combination of shocks in the Smets and Wouters (2007) NK model

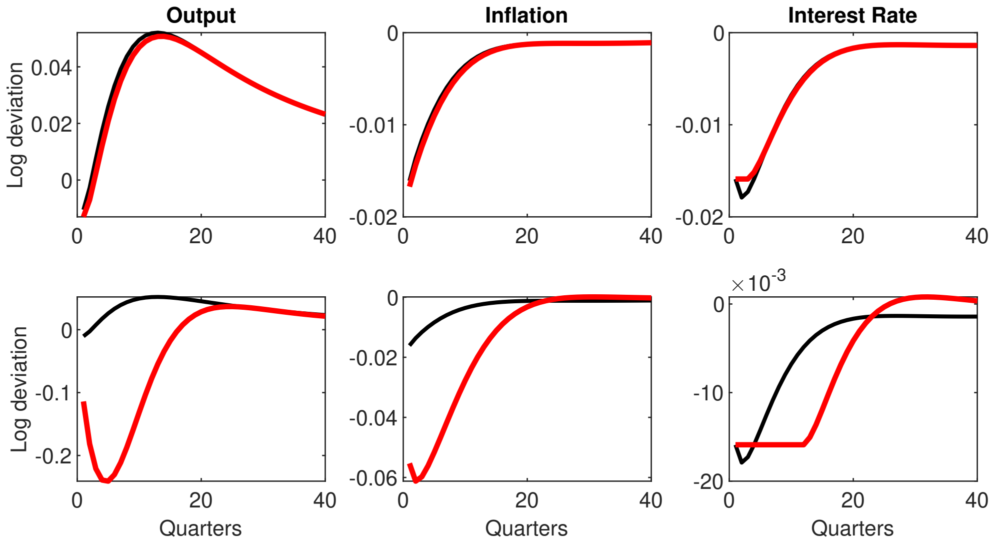

Note: The black line is the first solution found; the red line is the solution found when `SkipFirstSolutions=1` is used.

policy or another alternative monetary policy action could shift inflation expectations and move the economy from the ZLB. In any case, checking for multiple equilibria is a model-by-model issue that warrants consideration.

8 Conclusion

To briefly conclude, I have reviewed several common ways to solve models featuring an occasionally binding constraint. While I have presented projection methods as a benchmark approach, I have also outlined alternative approaches for use with larger models when global solution methods are not available. MATLAB code for all methods is available at `https://github.com/jonswarbrick/OBC_example_code`.

Using a simple borrowing constraints model, I find that the news shock approach available as a Dynare tool kit, DynareOBC, performs very well relative to the alternatives. Furthermore, this approach has additional appeal: it can confirm whether a solution to a particular problem exists or not, and it allows the user to carefully analyze multiple equilibria and choose between equilibrium paths when there are more than one.

9 References

Adam, K. and R. M. Billi. 2007. "Discretionary Monetary Policy and the Zero Lower Bound on Nominal Interest Rates." *Journal of Monetary Economics* 54 (3): 728–752.

Adjemian, S., H. Bastanie, F. Karamé, M. Juillard, J. Maih, F. Mihoubi, G. Perendia, J. Pfeifer, M. Ratto and S. Villemot. 2011. "Dynare: Reference Manual Version 4." Dynare Working Paper No. 1.

Adjemian, S. and M. Juillard. 2013. "Stochastic Extended Path Approach." Centre for Economic and International Studies.

Andreasen, M. M., J. Fernández-Villaverde and J. F. Rubio-Ramírez. 2018. "The Pruned State-Space System for Non-Linear DSGE Models: Theory and Empirical Applications." *The Review of Economic Studies* 85 (1): 1–49.

Bauer, M. D. and G. D. Rudebusch. 2016. "Monetary Policy Expectations at the Zero Lower Bound." *Journal of Money, Credit and Banking* 48 (7): 1439–1465.

Benhabib, J., G. W. Evans, and S. Honkapohja. 2014. "Liquidity Traps and Expectation Dynamics: Fiscal Stimulus or Fiscal Austerity?" *Journal of Economic Dynamics and Control* 45: 220–238.

Benhabib, J., S. Schmitt-Grohé and M. Uribe. 2001. "The Perils of Taylor Rules." *Journal of Economic Theory* 96 (1): 40–69.

Bernanke, B. 1983. "Irreversibility, Uncertainty, and Cyclical Investment." *Quarterly Journal of Economics* 98 (1): 85–106.

Bilbiie, F. O., T. Monacelli and R. Perotti. 2019. "Is Government Spending at the Zero Lower Bound Desirable?" *American Economic Journal: Macroeconomics* 11 (3): 147–173.

Binning, A. and J. Maih. 2017. "Modelling Occasionally Binding Constraints Using Regime-Switching." Norges Bank Working Paper No. 2017/23.

Bloom, N., S. Bond and J. Van Reenen. 2007. "Uncertainty and Investment Dynamics." *Review of Economic Studies* 74 (2): 391–415.

Bodenstein, M., J. Hebden and R. Nunes. 2012. "Imperfect Credibility and the Zero

Lower Bound." *Journal of Monetary Economics* 59 (2): 135–149.

Braun, A., L. Korber and Y. Waki. 2013. "Small and Orthodox Fiscal Multipliers at the Zero Lower Bound." Federal Reserve Bank of Atlanta Working Paper No. 2013-13.

Braun, R. A., L. M. Boneva and Y. Waki. 2016. "Some Unpleasant Properties of Log-Linearized Solutions When the Nominal Rate Is Zero." *Journal of Monetary Economics* 84 (C): 216–232.

Braun, R. A. and L. M. Körber. 2011. "New Keynesian Dynamics in a Low Interest Rate Environment." *Journal of Economic Dynamics and Control* 35(12): 2213–2227.

Brendon, C., M. Paustian and T. Yates. 2020. "Self-Fulfilling Recessions at the Zero Lower Bound." *Journal of Monetary Economics* 115: 213–232.

Brunnermeier, M. and Y. Sannikov. 2014. "A Macroeconomic Model with a Financial Sector." *American Economic Review* 104 (2): 379–421.

Brzoza-Brzezina, M., M. Kolasa and K. Makarski. 2015. "A Penalty Function Approach to Occasionally Binding Credit Constraints." *Economic Modelling* 51: 315–327.

Caballero, R. J. and R. C. Pindyck. 1996. "Uncertainty, Investment, and Industry Evolution." *International Economic Review* 37 (3): 641–662.

Coibion, O., Y. Gorodnichenko and J. Wieland. 2012. "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?" *The Review of Economic Studies* 79 (4): 1371–1406.

Collard, F. and M. Juillard. 2001. "A Higher-Order Taylor Expansion Approach to Simulation of Stochastic Forward-Looking Models with an Application to a Nonlinear Phillips Curve Model." *Computational Economics* 17 (2-3): 125–139.

Cottle, R. W. 2009. "Linear Complementarity Problem." In Encyclopedia of Optimization, edited by C. A. Floudas and P. M. Pardalos, 1873– 1878. Boston, MA: Springer US.

Debortoli, D. and J. Gali. 2017. "Monetary Policy with Heterogeneous Agents: Insights from TANK Models." Universitat Pompeu Fabra Economics Working Papers No. 1686.

Dixon, H. and P. M. Pourpourides. 2016. "On Imperfect Competition with Occasionally Binding Cash-in-Advance Constraints." *Journal of Macroeconomics* 50: 72–85.

Eggertsson, G. B. 2010. "What Fiscal Policy Is Effective at Zero Interest Rates?" In National Bureau of Economic Research Macroeconomics Annual 2010, Volume 25, 59–12. Chicago: University of Chicago Press.

Eggertsson, G. and S. Singh, S., 2019. "Log-Linear Approximation Versus an Exact Solution at the ZLB in the New Keynesian Model." *Journal of Economic Dynamics and Control* 105: 21–43.

Eggertsson, G. B. and M. Woodford. 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity* 34 (1): 139–235.

Erceg, C. and J. Lindé. 2014. "Is There a Fiscal Free Lunch in a Liquidity Trap?" *Journal of the European Economic Association* 12 (1): 73–107.

Fair, R. C. and J. B. Taylor. 1983. "Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models." *Econometrica* 51 (4): 1169–1185.

Fernández-Villaverde, J., G. Gordon, P. Guerrón-Quintana and J. F. Rubio-Ramírez. 2015. "Nonlinear Adventures at the Zero Lower Bound." *Journal of Economic Dynamics and Control* 57: 182–204.

Fischer, S. 2016. "Monetary Policy, Financial Stability, and the Zero Lower Bound." *American Economic Review* 106 (5): 39–42.

Gambacorta, L., B. Hoemann and G. Peersman. 2014. "The Effectiveness of Unconventional Monetary Policy at the Zero Lower Bound: A Cross-Country Analysis." *Journal of Money, Credit and Banking* 46 (4): 615–642.

Genz, A. and B. D. Keister. 1996. "Fully Symmetric Interpolatory Rules for Multiple Integrals over Abstract Infinite Regions with Gaussian Weight." *Journal of Computational and Applied Mathematics* 71: 299–309.

Gilchrist, S., D. López-Salido and E. Zakrajšek. 2015. "Monetary Policy and Real Borrowing Costs at the Zero Lower Bound." *American Economic Journal: Macroeconomics* 7 (1): 77–109.

Gilchrist, S., J. Sim and E. Zakrajšek. 2014. "Uncertainty, Financial Frictions, and

Irreversible Investment." National Bureau of Economic Research Working Paper No. 20038.

Guerrieri, L. and M. Iacoviello. 2015. "OccBin: A Toolkit for Solving Dynamic Models with Occasionally Binding Constraints Easily." *Journal of Monetary Economics* 70: 22–38.

Guerrieri, L. and M. Iacoviello. 2017. "Collateral Constraints and Macroeconomic Asymmetries." *Journal of Monetary Economics* 90: 28–49.

Hamilton, J. D. and J. C. Wu. 2012. "The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment." *Journal of Money, Credit and Banking* 44 (s1): 3–46.

He, Z. and A. Krishnamurthy. 2013. "Intermediary Asset Pricing." *American Economic Review* 103 (2): 732–770.

Holden, T. 2010. "Products , Patents and Productivity Persistence: A DSGE Model of Endogenous Growth." University of Oxford Discussion Paper No. 512.

Holden, T. D. 2016. "Computation of Solutions to Dynamic Models with Occasionally Binding Constraints." ZBW—Leibniz Information Centre for Economics, EconStor Preprints No. 144569.

Holden, T. D. 2019. "Existence and Uniqueness of Solutions to Dynamic Models with Occasionally Binding Constraints." ZBW—Leibniz Information Centre for Economics, Econstor Preprints No. 144570.

Holden, T. D., P. Levine and J. M. Swarbrick. 2020. "Credit Crunches from Occasionally Binding Bank Borrowing Constraints." *Journal of Money, Credit and Banking* 52 (2-3): 549–582.

Holden, T. and M. Paetz. 2012. "Efficient Simulation of DSGE Models with Inequality Constraints", University of Surrey School of Economics Discussion Paper No. 1612.

Judd, K. L. 1992. "Projection Methods for Solving Aggregate Growth Models." *Journal of Economic Theory* 58 (2): 410–452

Jung, T., Y. Teranishi and T. Watanabe. 2005. "Optimal Monetary Policy at the Zero-Interest-Rate Bound." *Journal of Money, Credit, and Banking* 37 (5): 813–835.

Kanzow, C. and S. Petra. 2004. "On a Semismooth Least Squares Formulation of Complementarity Problems with Gap Reduction." *Optimization Methods and Software* 19 (5): 507–525.

Kaplan, G., B. Moll and G. L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review* 108 (3): 697–743.

Lan, H. and A. Meyer-Gohde. 2013. "Pruning in Perturbation DSGE Models—Guidance from Nonlinear Moving Average Approximations." Humboldt University SFB Discussion Paper No. 649.

Maih, J. 2015. "Efficient Perturbation Methods for Solving Regime-Switching DSGE Models." Norges Bank Working Paper No. 2015/01.

Maliar, L. and S. Maliar. 2015. "Merging Simulation and Projection Approaches to Solve High-Dimensional Problems with an Application to a New Keynesian Model." *Quantitative Economics* 6: 1–47.

Mertens, K. R. S. M. and M. O. Ravn. 2014. "Fiscal Policy in an Expectations-Driven Liquidity Trap." *Review of Economic Studies* 81 (4): 1637–1667.

Miranda, M. J. and P. L. Fackler. 2002. Applied Computational Economics and Finance. Cambridge, MA: MIT Press.

Nakata, T. 2017. "Uncertainty at the Zero Lower Bound." *American Economic Journal: Macroeconomics* 9 (3): 186–221.

Plante, M., A. W. Richter and N. A. Throckmorton. 2017. "The Zero Lower Bound and Endogenous Uncertainty." *The Economic Journal* 128 (611): 1730–1757.

Santos, M. S. 2000. "Accuracy of Numerical Solutions Using the Euler Equation Residuals." *Econometrica* 68 (6): 1377–1402.

Smets, F. and R. Wouters. 2003. "An Estimated Dynamic Stochastic General Equilibrium Model." *Journal of the European Economic Association* 1 (5): 1123–1175.

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

Sobol, I. M. 1967. "On the Distribution of Points in a Cube and the Approximate Evaluation of Integrals." *USSR Computational Mathmatics and Mathematical Physics*

7(4): 784–802.

Swanson, E. T. and J. C. Williams. 2014. "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates." *American Economic Review* 104 (10): 3154–3185.

Swarbrick, J. 2019. "Lending Standards, Productivity and Credit Crunches." Bank of Canada Staff Working Paper No. 2019-25.

Taylor, J. B. 1999. "A Historical Analysis of Monetary Policy Rules." In Monetary Policy Rules, NBER Chapters, Volume 1, edited by J. B. Taylor, 319–348. Chicago: University of Chicago Press.

Wieland, J. F., 2019. "Are Negative Supply Shocks Expansionary at the Zero Lower Bound?" *Journal of Political Economy* 127 (3): 973–1007.

Woodford, M. 2001. "The Taylor Rule and Optimal Monetary Policy." *American Economic Review* 91 (2): 232–237.

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

Wright, J. H. 2012. "What Does Monetary Policy Do to Long-term Interest Rates at the Zero Lower Bound?" *The Economic Journal* 122 (564): F447–F466.

Wu, J. C. and F. D. Xia. 2016. "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound." *Journal of Money, Credit and Banking* 48 (2-3): 253–291.