

Existence, uniqueness and computation of solutions to dynamic models with occasionally binding constraints.

Tom D. Holden¹, School of Economics, University of Surrey

Abstract: We present the first necessary and sufficient conditions for the existence of a unique perfect-foresight solution, returning to a given steady-state, in an otherwise linear model with occasionally binding constraints. We derive further conditions on the existence of a solution in such models, and provide a proof of the inescapability of the “curse of dimensionality” for them. We also construct the first solution algorithm for these models that is guaranteed to return a solution in finite time, if one exists. When extended to allow for other non-linearities and future uncertainty, our solution algorithm is shown to produce fast and accurate simulations. In an application, we show that widely used New Keynesian models with endogenous states possess multiple perfect foresight equilibrium paths when there is a zero lower bound on nominal interest rates. However, we show that price level targeting is sufficient to restore determinacy in these situations.

Keywords: *occasionally binding constraints, zero lower bound, existence, uniqueness, price targeting, DSGE, linear complementarity problem*

JEL Classification: *C61, C62, C63, E17, E3, E4, E5*

This version: 09 February 2016

The latest version of this paper may be downloaded from:

<https://github.com/tholden/dynareOBC/raw/master/paper.pdf>

¹ Contact address: School of Economics, FASS, University of Surrey, Guildford, GU2 7XH, England

Contact email: thomas.holden@gmail.com Web-site: <http://www.tholden.org/>

The author would particularly like to thank Michael Paetz for joint work on an earlier algorithm for simulating models with occasionally binding constraints. Additionally, the author would like to thank the participants in numerous seminars at which versions of this paper were presented. Further special thanks are due to the following individuals for helpful discussions: Gary Anderson, Martin M. Andreasen, Paul Beaudry, Saroj Bhattacharai, Charles Brendon, Zsolt Csizmadia, Oliver de Groot, Michael B. Devereux, Bill Dupor, Roger Farmer, Michael Funke, William T. Gavin, Fabio Ghironi, Andy Glover, Luca Guerrieri, Pablo Guerrón-Quintana, Matteo Iacoviello, Tibor Illés, Michel Juillard, Hong Lan, Paul Levine, Jesper Lindé, Albert Marcet, Enrique Martínez-García, Antonio Mele, Alexander Meyer-Gohde, Adrienn Nagy, Matthias Paustian, Jesse Perla, Frank Portier, Søren Ravn, Alexander Richter, Jonathan Swarbrick, Simon Wren-Lewis and Carlos Zarazaga. Financial support provided to the author by the ESRC and the EC is also greatly appreciated. Furthermore, the author gratefully acknowledges the use of the University of Surrey FASS cluster for the numerical computations performed in this paper, and the University of Washington for providing the office space in which the bulk of this paper was written.

1. Introduction

Since the financial crisis of 2007-2008, many central banks around the world have chosen to keep their nominal interest rate close to 0%. While in a few countries, rates on some assets have gone slightly negative, central banks are unable to push their target rate to the level a Taylor rule might suggest since agents always have the option of holding cash. In practice then, central banks face a zero lower bound (ZLB) on their policy rate, which limits their ability to provide stimulus in severe recessions. Furthermore, during the crisis, both households, firms and banks have hit their borrowing constraints, which has limited their ability to smooth out its effects. However, traditional approaches to understanding the behaviour of medium-scale DSGE models, such as perturbation, cannot capture occasionally binding constraints (OBCs), meaning that the profession still lacks all of the necessary tools for understanding the behaviour of models with OBCs.

In this paper, we attack the problem of understanding the behaviour of models with occasionally binding constraints from two directions. Much as the seminal paper of Blanchard and Kahn (1980) both provided necessary and sufficient conditions for the existence of a unique solution to a linear model, and provided a practical approach for their simulation, we shall do the same for models with OBCs. Firstly, we provide theoretical results on the existence and uniqueness of solutions to such models. Secondly, we provide a computational algorithm for their robust, accurate and scalable simulation, along with a toolkit (“DynareOBC”) implementing the algorithm.²

Our theoretical results are for perfect-foresight solutions to models that are otherwise linear, apart from their occasionally binding constraints. As was observed by Benhabib, Schmitt-Grohé, and Uribe (2001a; 2001b), in the presence of OBCs, there are often multiple steady-states. For example, a model with a zero lower bound on nominal interest rates and Taylor rule monetary policy when away from the bound will have an additional “bad” deflationary steady-state in which nominal interest rates are zero. The presence of such multiple steady-states means that there can be sunspot equilibria which jump between the neighbourhoods of the two steady-states. Furthermore, if agents put a positive probability on being in the neighbourhood of the “bad” steady-state in future, then since this “bad” steady-state is indeterminate in the ZLB case, by a backwards induction argument, there is indeterminacy now. The consequences of indeterminacy of these kinds have been explored by Schmitt-Grohé and Uribe (2012), Mertens and Ravn (2014) and Aruoba, Cuba-Borda, and Schorfheide (2014), amongst others. In all cases,

² DynareOBC is available from: <https://github.com/tholden/dynareOBC>

the key to generating indeterminacy is that agents' beliefs about the point to which the economy would converge in the absence of future uncertainty are switching from one steady-state to the other.

One might hope, however, that the presence of announced inflation targets in many countries might anchor long-run expectations at the “good” steady-state. Additionally, one might legitimately wonder about the plausibility of the required coordination in beliefs necessary to sustain these sunspot equilibria. It is an interesting question then whether there are still multiple equilibria even when all agents believe that in the long-run, the economy will return to the “good” steady-state. It is on such equilibria that we focus on in this paper, providing necessary and sufficient conditions for the existence of a unique perfect-foresight path, and also examining whether such beliefs are actually consistent with rationality. A restricted class of such equilibria were also examined by Brendon, Paustian, and Yates (2015) who examined multiplicity in specific models when agents believe that with probability one, in one period's time, they will escape the bound and return to the neighbourhood of the “good” steady-state.

We show that many standard New Keynesian models featuring endogenous state variables (e.g. price dispersion), such as those of Fernández-Villaverde et al. (2012) or Smets and Wouters (2003; 2007) do not possess such a unique perfect-foresight path, meaning that even in this best case scenario in which agents' long-run expectations are pinned down, there is still multiplicity of equilibria. Indeed, we show that in these models, there are some initial states from which the economy has one return path that never hits the ZLB, and another that does hit it, so the fact that the ZLB is not violated in a model in which it is not imposed does not mean that it would not be hit were it to be imposed. However, we show that under a price-targeting regime, there is a unique equilibrium path even when we impose the ZLB.

We also provide both necessary and sufficient conditions for the existence of any perfect-foresight solutions which return to the original (“good”) steady-state. When no such equilibria exists, agents must switch their beliefs to the other (“bad”) steady-state, where they will remain in the absence of any way for agents to coordinate back on the “good” steady-state. We show that for standard New Keynesian models with endogenous state variables, there is a positive probability of ending up in a state of the world in which there is no perfect foresight path returning to the “good” steady-state,³

³ This has some similarities to the results of Richter and Throckmorton (2014) and Appendix B of Gavin et al. (2015), who show numerically that a particular solution algorithm does not converge in certain areas of the state/parameter/guess space for a simple NK model. However, our results are theoretical, so whereas Richter and Throckmorton and Gavin et al.'s results may possibly be driven by the particular properties of their solution procedure, ours imply true non-existence, at least for perfect foresight, otherwise linear models. For example, for the model with Rotemberg (1982) type pricing, and no steady-state distortions, that these authors work with, our results imply global existence and uniqueness for the linearized model when the standard Taylor principle is satisfied.

implying that in the stochastic model, agents must always put positive probability on tending to the “bad” steady-state. This in turn implies global indeterminacy in such models, by a backwards induction argument. Once again though, price level targeting is sufficient to restore determinacy.

Additionally, we present theoretical results on the computational complexity of finding a solution to models with occasionally binding constraints. We are able to show that were there an algorithm which could solve general non-stochastic models with OBCs in an amount of time that is polynomial in the number of states in the model, then such an algorithm could also be used for solving in polynomial time any problem the solution of which could be verified in polynomial time. This would imply, for example, efficient methods of breaking all standard forms of cryptography used to secure internet banking. Of course, neither we nor the computer science profession believe that such a polynomial time algorithm can exist, hence this provides a proof by contradiction of the “curse of dimensionality” for this class of models. It also raises doubts about the realism of assuming rational expectations in these situations. However, we show that under price level targeting, the solution may be computed in polynomial time, meaning it is much easier for agents to form expectations under price level targeting than it is under inflation targeting. This gives an additional argument for price level targeting.

We go on to present an algorithm for solving general models with occasionally binding constraints, as efficiently as is possible given the aforementioned theoretical results. In the otherwise linear, perfect foresight case, we are able to represent the problem as the solution to a mixed integer linear programming problem, a problem for which incredibly efficient solvers already exist. The key idea of the algorithm is that an OBC provides a source of endogenous news about the future. When a shock hits, driving the economy to the bound in some future periods, that tells us that in those future periods, the (lower) bounded variable will be higher than it would be otherwise.⁴

Thinking in terms of endogenous news shocks also provides intuition for the presence of multiple equilibria in these models. As an example, consider a New Keynesian model with significant real and nominal frictions. If these frictions are large enough, then learning about a future positive shock to nominal interest rates induces a sufficiently severe downturn that the Taylor rule calls for much lower rates, even in the period in which the shock actually arrives. While positive shocks having negative effects may sound somewhat bizarre, in fact this is a relatively common phenomenon in New Keynesian models. Then, there will be some magnitude of news shock to nominal

⁴ The idea of imposing the zero lower bound by adding news shocks is also present in Holden (2010), Hebden et al. (2011), Holden and Paetz (2012) and Bodenstein et al. (2013). News shocks were introduced to the literature by Beaudry and Portier (2006).

interest rates today at which the news is of precisely the correct magnitude to bring the negative interest rates implied by the Taylor rule up to zero, in that period. A news shock of this magnitude thus becomes a self-fulfilling prophecy, as illustrated in Figure 1. In models with weaker rigidities, multiple equilibria are still possible if there is some combination of future periods such that with appropriate news shocks in each, a similar self-fulfilling prophecy occurs.




Figure 1: Self-fulfilling news shocks

Our algorithm for finding the required news shocks to impose the zero lower bound is guaranteed to return a solution in finite time when one exists, and when there is no solution, the algorithm returns a certificate of this in finite time instead. This contrasts with approaches based on a fixed point iteration for which non-existence is not normally detectable in finite time, since one cannot rule out that the algorithm would converge if only it were left for another hour/day/year. Furthermore, where there are multiple solutions, our algorithm always returns one minimising an intuitive criterion, with a free parameter that enables the user to select the desired “type” of equilibrium. This algorithm may be applied to stochastic models using the idea of the extended path algorithm of Fair and Taylor (1983), and made consistent with rationality following the stochastic extended path algorithm of Adjemian and Juillard (2013), as discussed below.

For models that are non-linear even apart from the constraint, we exploit the convenient properties of pruned perturbation approximations (Kim et al. 2008), which enables the base algorithm to be applied to higher order approximations with only minimal modifications. Although this does not capture the risk of hitting the bound, it does at least help capture the fact that at the bound, the economy is a substantial distance from steady-state, and so the slopes of variables’ responses will have changed.

To ensure consistency with rationality, such as precautionary motives to avoid the bound, we use a modified version of the stochastic extended path algorithm of Adjemian and Juillard (2013) that is designed to exploit both the special properties of our inner

solution algorithm, and the special properties of pruned perturbation. Whereas in the original Adjemian and Juillard (2013) approach, integrating over S periods of future uncertainty required a number of solutions to the perfect foresight problem that was exponential in both S and the number of shocks, we are able to integrate over the same number of periods of future uncertainty with only polynomial in S solutions of the perfect foresight problem. In practice, this means that we can integrate over enough periods of future uncertainty to capture even one hundred periods of future uncertainty, and even in medium scale models.

Our paper is structured as follows. In the following section, section 2 we present our key theoretical results on otherwise linear perfect foresight models. We then discuss the application of these results to New Keynesian models in section 3. In section 4 we present our solution algorithm for otherwise linear perfect foresight models, and discuss the computational complexity of the problem. We then extend this to non-linear, non-perfect foresight models in section 5, and we assess the algorithm's numerical accuracy and speed, and discuss its relationship to other algorithms in the literature. Section 6 concludes. All files needed for the replication of this paper's numerical results are included in the Examples directory of the DynareOBC toolkit.⁵

2. Theoretical results on occasionally binding constraints in otherwise linear models under perfect foresight

In this section, we present our main theoretical results on existence and uniqueness of perfect foresight solutions to models which are linear apart from an occasionally binding constraint. We start by defining the problem to be solved, and examining its relationship both to the problem without OBCs, and to a related problem with news shocks to the bounded variable. Using the news shock representation, we demonstrate that solving the model with OBCs is equivalent to solving a linear complementarity problem. We then discuss some theoretical background on these problems, before presenting the main existence and uniqueness results.

2.1. Problem set-ups

Suppose that for $t \in \mathbb{N}^+$, (i.e. $t \in \mathbb{N}$, $t > 0$), the first order conditions of some model may be represented as:

$$(\hat{A} + \hat{B} + \hat{C})\hat{\mu} = \hat{A}\hat{x}_{t-1} + B\hat{x}_t + \hat{C}\mathbb{E}_t\hat{x}_{t+1} + \hat{D}\varepsilon_t,$$

where $\hat{\mu} \in \mathbb{R}^{\hat{n}}$ and $\hat{x}_t \in \mathbb{R}^{\hat{n}}$, $\varepsilon_t \in \mathbb{R}^m$, $\mathbb{E}_{t-1}\varepsilon_t = 0$ for all $t \in \mathbb{N}^+$, and suppose that \hat{x}_0 is given as an initial condition. Throughout this paper, we will refer to first order

⁵ These files may be viewed online at <https://github.com/tholden/dynareOBC/tree/master/Examples>.

conditions such as these as “the model”, conflating them with the optimisation problem(s) which gave rise to them.

Furthermore, suppose that $\varepsilon_t = 0$ for $t > 1$, as in an impulse response or perfect foresight simulation exercise. Additionally, we assume the existence of a terminal condition of the form $\hat{x}_t \rightarrow \hat{\mu}$ as $t \rightarrow \infty$, coming, for example, from the source model’s transversality constraints.

For $t \in \mathbb{N}^+$, define $x_t := \begin{bmatrix} \hat{x}_t \\ \varepsilon_{t+1} \end{bmatrix}$, $\mu := \begin{bmatrix} \hat{\mu} \\ 0 \end{bmatrix}$, $A := \begin{bmatrix} \hat{A} & \hat{D} \\ 0 & 0 \end{bmatrix}$, $B := \begin{bmatrix} \hat{B} & 0 \\ 0 & I \end{bmatrix}$, $C := \begin{bmatrix} \hat{C} & 0 \\ 0 & 0 \end{bmatrix}$, then, for $t \in \mathbb{N}^+$:

$$(A + B + C)\mu = Ax_{t-1} + Bx_t + Cx_{t+1}, \quad (1)$$

and we have the extended initial condition $x_0 = \begin{bmatrix} \hat{x}_0 \\ \varepsilon_1 \end{bmatrix}$, and the extended terminal condition $x_t \rightarrow \mu$ as $t \rightarrow \infty$. Expectations have disappeared since there is no uncertainty after period 0. Thus, the problem of solving the original model has the same form as that given in:

Problem 1 Suppose that $x_0 \in \mathbb{R}^n$ is given. Find $x_t \in \mathbb{R}^n$ for $t \in \mathbb{N}^+$ such that $x_t \rightarrow \mu$ as $t \rightarrow \infty$, and such that for all $t \in \mathbb{N}^+$, equation (1) holds.

We make the following assumption in all of the following:

Assumption 1 For any given $x_0 \in \mathbb{R}^n$, Problem 1 has a unique solution, which takes the form $x_t = (I - F)\mu + Fx_{t-1}$, for $t \in \mathbb{N}^+$, where $F = -(B + CF)^{-1}A$, and where all of the eigenvalues of F are weakly inside the unit circle.

Sims’s (2002) generalisation of the standard Blanchard-Kahn (1980) conditions is necessary and sufficient for this. Further, to avoid dealing specially with the knife-edge case of exact unit eigenvalues (even if they are constrained to the part of the model that is solved forward), in the following we rule it out with the subsequent assumption, which is, in any case, a necessary condition for perturbation to produce a consistent approximation to a source non-linear model, and which is also necessary for the linear model to have a unique steady-state:

Assumption 2 $\det(A + B + C) \neq 0$.

The combination of Assumption 1 and Assumption 2 imply that all of the eigenvalues of F are strictly inside the unit circle.

We are interested in models featuring occasionally binding constraints. We will concentrate on models featuring a single zero lower bound type constraint in their first equation, which we treat as defining the first element of x_t . Generalising from this special case is straightforward, and is discussed in online appendix E. First, let us write

$x_{1,t}, I_{1,\cdot}, A_{1,\cdot}, B_{1,\cdot}, C_{1,\cdot}$ for the first row of x_t, I, A, B, C (respectively) and $x_{-1,t}, I_{-1,\cdot}, A_{-1,\cdot}, B_{-1,\cdot}, C_{-1,\cdot}$ for the remainders. Likewise, we write $I_{\cdot,1}$ for the first column of I , and so on. Then we are interested in the solution to:

Problem 2 Suppose that $x_0 \in \mathbb{R}^n$ is given. Find $T \in \mathbb{N}$ and $x_t \in \mathbb{R}^n$ for $t \in \mathbb{N}^+$ such that $x_t \rightarrow \mu$ as $t \rightarrow \infty$, and such that for all $t \in \mathbb{N}^+$:

$$x_{1,t} = \max\{0, I_{1,\cdot}\mu + A_{1,\cdot}(x_{t-1} - \mu) + (B_{1,\cdot} + I_{1,\cdot})(x_t - \mu) + C_{1,\cdot}(x_{t+1} - \mu)\},$$

$$(A_{-1,\cdot} + B_{-1,\cdot} + C_{-1,\cdot})\mu = A_{-1,\cdot}x_{t-1} + B_{-1,\cdot}x_t + C_{-1,\cdot}x_{t+1},$$

and such that $x_{1,t} > 0$ for $t > T$.

Note that in this problem we are implicitly ruling out any solutions which get permanently stuck at an alternative steady-state, by assuming that the terminal condition remains as before. Indeed, we are assuming that the bound is only relevant for some finite number of periods T . We continue to assume that there is no uncertainty after period 0, so, in this non-linear model, the path of the endogenous variables will not necessarily match up with the path of their expectation in a richer model in which there was uncertainty after period 0.

In many models, the occasionally binding constraint comes from the KKT conditions of an optimisation problem. We will give in section 2.6 a general procedure for converting such conditions into a problem in the form of that Problem 2, and we provide a simple example of doing this in the second half of section 5.3. The intuition is that one can use the model's equations to find the value the (lower) constrained variable would take were there no constraint and were the Lagrange multiplier on the constraint equal to zero today. This gives a “shadow” value of the constrained variable, and the actual value it takes will be the maximum of the bound and this shadow value.

We will analyse Problem 2 with the help of solutions to the auxiliary problem:

Problem 3 Suppose that $T \in \mathbb{N}$, $x_0 \in \mathbb{R}^n$ and $y_0 \in \mathbb{R}^T$ is given. Find $x_t \in \mathbb{R}^n, y_t \in \mathbb{R}^T$ for $t \in \mathbb{N}^+$ such that $x_t \rightarrow \mu, y_t \rightarrow 0$, as $t \rightarrow \infty$, and such that for all $t \in \mathbb{N}^+$:

$$(A + B + C)\mu = Ax_{t-1} + Bx_t + Cx_{t+1} + I_{\cdot,1}y_{1,t-1},$$

$$y_{T,t} = 0, \quad \forall i \in \{1, \dots, T-1\}, \quad y_{i,t} = y_{i+1,t-1}.$$

This may be thought of as a version of Problem 1 with news shocks up to horizon T added to the first equation. The value of $y_{t,0}$ gives the news shock that hits in period t , i.e. $y_{1,t-1} = y_{t,0}$ for $t \leq T$, and $y_{1,t-1} = 0$ for $t > T$.

2.2. Relationships between the problems

Since $y_{1,t-1} = 0$ for $t > T$, and using Assumption 1, $(x_{T+1} - \mu) = F(x_T - \mu)$, so with $t = T$, defining $s_{T+1} := 0$, $(x_{t+1} - \mu) = s_{t+1} + F(x_t - \mu)$.

Proceeding now by backwards induction on t , note that $0 = A(x_{t-1} - \mu) + B(x_t - \mu) + CF(x_t - \mu) + Cs_{t+1} + I_{\cdot,1}y_{t,0}$, so:

$$\begin{aligned}(x_t - \mu) &= -(B + CF)^{-1}[A(x_{t-1} - \mu) + Cs_{t+1} + I_{\cdot,1}y_{t,0}] \\ &= F(x_{t-1} - \mu) - (B + CF)^{-1}(Cs_{t+1} + I_{\cdot,1}y_{t,0}),\end{aligned}$$

i.e., if we define: $s_t := -(B + CF)^{-1}(Cs_{t+1} + I_{\cdot,1}y_{t,0})$, then $(x_t - \mu) = s_t + F(x_{t-1} - \mu)$. By induction then, this holds for all $t \in \{1, \dots, T\}$.⁶ Hence, we have proved the following lemma:

Lemma 1 There is a unique solution to Problem 3 that is linear in x_0 and y_0 .

For future reference, let $x_t^{(3,k)}$ be the solution to Problem 3 when $x_0 = \mu$, $y_0 = I_{\cdot,k}$ (i.e. a vector which is all zeros apart from a 1 in position k). Then, by linearity, for arbitrary y_0 the solution to Problem 3 when $x_0 = \mu$ is given by:

$$x_t - \mu = \sum_{k=1}^T y_{k,0} (x_t^{(3,k)} - \mu).$$

Let $M \in \mathbb{R}^{T \times T}$ satisfy:

$$M_{t,k} = x_{1,t}^{(3,k)} - \mu_1, \quad \forall t, k \in \{1, \dots, T\}, \quad (2)$$

i.e. M horizontally stacks the (column-vector) relative impulse responses to the news shocks. Then this result implies that for arbitrary y_0 , the path of the first variable in the solution to Problem 3 when $x_0 = \mu$ is given by: $(x_{1,1:T})' = \mu_1 + My_0$, where $x_{1,1:T}$ is the row vector of the first T values of the first component of x_t . Furthermore, for both arbitrary x_0 and y_0 , the path of the first variable in the solution to Problem 3 is given by: $(x_{1,1:T})' = q + My_0$, where $q := (x_{1,1:T}^{(1)})'$ and $x_t^{(1)}$ is the unique solution to Problem 1, for the given x_0 .⁷ This ease in solving Problem 3 given y_0 will be crucial to the efficiency of our eventual solution algorithm for Problem 2.

Now let $x_t^{(2)}$ be a solution to Problem 2 given an arbitrary x_0 . Since $x_t^{(2)} \rightarrow \mu$ as $t \rightarrow \infty$, there exists $T' \in \mathbb{N}$ such that for all $t > T'$, $x_{1,t}^{(2)} > 0$. We assume without loss of generality that $T' \leq T$. We seek to relate the solution to Problem 2 with the solution to Problem 3 for an appropriate choice of y_0 . First, for all $t \in \mathbb{N}^+$, let:

$$e_t := \begin{cases} -[I_{1,\cdot}\mu + A_{1,\cdot}(x_{t-1}^{(2)} - \mu) + (B_{1,\cdot} + I_{1,\cdot})(x_t^{(2)} - \mu) + C_{1,\cdot}(x_{t+1}^{(2)} - \mu)] & \text{if } x_{1,t}^{(2)} = 0 \\ 0 & \text{if } x_{1,t}^{(2)} > 0 \end{cases}, \quad (3)$$

i.e. e_t is the shock that would need to hit the first equation for the positivity constraint on $x_{1,t}^{(2)}$ to be enforced. Note for future reference that by the definition of Problem 2, $e_t \geq 0$ and $x_{1,t}^{(2)}e_t = 0$, for all $t \in \mathbb{N}^+$. From this definition, we also have that for all $t \in \mathbb{N}^+$, $0 = A(x_{t-1}^{(2)} - \mu) + B(x_t^{(2)} - \mu) + C(x_{t+1}^{(2)} - \mu) + I_{\cdot,1}e_t$. Furthermore, if $t > T$, then $t > T'$, and hence $e_t = 0$. Hence, by Assumption 1, $(x_{T+1}^{(2)} - \mu) = F(x_T^{(2)} - \mu)$.

⁶ This representation of the solution to Problem 3 was inspired by that of Anderson (2015).

⁷ This representation was also exploited by Holden (2010) and Holden and Paetz (2012).

Thus, much as before, with $t = T$, defining $\tilde{s}_{T+1} := 0$, $(x_{T+1}^{(2)} - \mu) = \tilde{s}_{T+1} + F(x_T^{(2)} - \mu)$. Consequently, $0 = A(x_{T-1}^{(2)} - \mu) + B(x_T^{(2)} - \mu) + CF(x_T^{(2)} - \mu) + C\tilde{s}_{T+1} + I_{\cdot,1}e_T$, so $(x_T^{(2)} - \mu) = F(x_{T-1}^{(2)} - \mu) - (B + CF)^{-1}(C\tilde{s}_{T+1} + I_{\cdot,1}e_T)$, i.e., if we define: $\tilde{s}_T := -(B + CF)^{-1}(C\tilde{s}_{T+1} + I_{\cdot,1}e_T)$, then $(x_T^{(2)} - \mu) = \tilde{s}_T + F(x_{T-1}^{(2)} - \mu)$. As before, by induction this must hold for all $t \in \{1, \dots, T\}$. By comparing the definitions of s_t and \tilde{s}_t , and the laws of motion of x_t under both problems, we then immediately have that if Problem 3 is started with $x_0 = x_0^{(2)}$ and $y_0 = e'_{1:T}$, then $x_t^{(2)}$ solves Problem 3. Conversely, if $x_t^{(2)}$ solves Problem 3 for some y_0 , then from the laws of motion of x_t under both problems it must be the case that $\tilde{s}_t = s_t$ for all $t \in \mathbb{N}$, and hence from the definitions of s_t and \tilde{s}_t , we have that $y_0 = e'_{1:T}$. This has established the following result:

Lemma 2 For any solution, $x_t^{(2)}$ to Problem 2:

- 1) With $e_{1:T}$ as defined in equation (3), $e_{1:T} \geq 0$, $x_{1,1:T}^{(2)} \geq 0$ and $x_{1,1:T}^{(2)} \circ e_{1:T} = 0$, where \circ denotes the Hadamard (entry-wise) product.
 - 2) $x_t^{(2)}$ is also the unique solution to Problem 3 with $x_0 = x_0^{(2)}$ and $y_0 = e'_{1:T}$.
 - 3) If $x_t^{(2)}$ solves Problem 3 with $x_0 = x_0^{(2)}$ and with some y_0 , then $y_0 = e'_{1:T}$.
-

However, to use the easy solution to Problem 3 to assist us in solving Problem 2 requires a slightly stronger result. Suppose that $y_0 \in \mathbb{R}^T$ is such that $y_0 \geq 0$, $x_{1,1:T}^{(3)} \circ y'_0 = 0$ and $x_{1,t}^{(3)} \geq 0$ for all $t \in \mathbb{N}$, where $x_t^{(3)}$ is the unique solution to Problem 3 when started at x_0, y_0 . We would like to prove that in this case $x_t^{(3)}$ must also be a solution to Problem 2. I.e., we must prove that for all $t \in \mathbb{N}^+$:

$$x_{1,t}^{(3)} = \max\{0, I_{1,\cdot}\mu + A_{1,\cdot}(x_{t-1}^{(3)} - \mu) + (B_{1,\cdot} + I_{1,\cdot})(x_t^{(3)} - \mu) + C_{1,\cdot}(x_{t+1}^{(3)} - \mu)\}, \quad (4)$$

$$(A_{-1,\cdot} + B_{-1,\cdot} + C_{-1,\cdot})\mu = A_{-1,\cdot}x_{t-1}^{(3)} + B_{-1,\cdot}x_t^{(3)} + C_{-1,\cdot}x_{t+1}^{(3)}.$$

By the definition of Problem 3, the latter equation must hold with equality, so there is nothing to prove there. Hence we just need to prove that equation (4) holds for all $t \in \mathbb{N}^+$. So let $t \in \mathbb{N}^+$. Now, if $x_{1,t}^{(3)} > 0$, then $y_{t,0} = 0$, by the complementary slackness type condition ($x_{1,1:T}^{(3)} \circ y'_0 = 0$). Thus, from the definition of Problem 3:

$$\begin{aligned} x_{1,t}^{(3)} &= I_{1,\cdot}\mu + A_{1,\cdot}(x_{t-1}^{(3)} - \mu) + (B_{1,\cdot} + I_{1,\cdot})(x_t^{(3)} - \mu) + C_{1,\cdot}(x_{t+1}^{(3)} - \mu) \\ &= \max\{0, I_{1,\cdot}\mu + A_{1,\cdot}(x_{t-1}^{(3)} - \mu) + (B_{1,\cdot} + I_{1,\cdot})(x_t^{(3)} - \mu) + C_{1,\cdot}(x_{t+1}^{(3)} - \mu)\}, \end{aligned}$$

as required. The only remaining case is that $x_{1,t}^{(3)} = 0$ (since $x_{1,t}^{(3)} \geq 0$ for all $t \in \mathbb{N}$, by assumption), which implies that:

$$\begin{aligned} x_{1,t}^{(3)} = 0 &= A_{1,\cdot}(x_{t-1} - \mu) + B_{1,\cdot}(x_t - \mu) + C_{1,\cdot}(x_{t+1} - \mu) + y_{t,0} \\ &= I_{1,\cdot}\mu + A_{1,\cdot}(x_{t-1} - \mu) + (B_{1,\cdot} + I_{1,\cdot})(x_t - \mu) + C_{1,\cdot}(x_{t+1} - \mu) + y_{t,0}, \end{aligned}$$

by the definition of Problem 3. Thus:

$$I_{1,\cdot}\mu + A_{1,\cdot}(x_{t-1} - \mu) + (B_{1,\cdot} + I_{1,\cdot})(x_t - \mu) + C_{1,\cdot}(x_{t+1} - \mu) = -y_{t,0} \leq 0,$$

where the inequality is an immediate consequence of another of our assumptions. Consequently, equation (4) holds in this case too. Together with Lemma 1, Lemma 2, and our representation of the solution of Problem 3, this completes the proof of the following proposition:

Proposition 1 The following hold:

- 1) Let $x_t^{(3)}$ be the unique solution to Problem 3 when initialized with some x_0, y_0 . Then $x_t^{(3)}$ is a solution to Problem 2 when initialized with x_0 if and only if $y_0 \geq 0$, $y_0 \circ (q + My_0) = 0$, $q + My_0 \geq 0$ and $x_{1,t}^{(3)} \geq 0$ for all $t \in \mathbb{N}$ with $t > T$.
 - 2) Let $x_t^{(2)}$ be any solution to Problem 2 when initialized with x_0 . Then there exists a $y_0 \in \mathbb{R}^T$ such that $y_0 \geq 0$, $y_0 \circ (q + My_0) = 0$, $q + My_0 \geq 0$, such that $x_t^{(2)}$ is the unique solution to Problem 3 when initialized with x_0, y_0 .
-

2.3. The linear complementarity representation

Proposition 1 establishes that providing we initially choose T sufficiently high, to find a solution to Problem 2, it is sufficient to solve the following problem instead:

Problem 4 Suppose $q \in \mathbb{R}^T$ and $M \in \mathbb{R}^{T \times T}$ are given. Find $y \in \mathbb{R}^T$ such that $y \geq 0$, $y \circ (q + My) = 0$ and $q + My \geq 0$. We call this the **linear complementarity problem (LCP)** (q, M) . (Cottle 2009)

These problems have been extensively studied, and so we can import results on the properties of LCPs to derive results on the properties of solutions to models with OBCs.

All of the results in the mathematical literature rest on properties of the matrix M , thus we must first establish if the structure of our particular M implies it has any special properties. Unfortunately, it seems that M has no general properties. We show this by constructing, in online appendix A, a model for each matrix in $\mathbb{R}^{T \times T}$, such that the M matrix the model produces is precisely the matrix we started with. Our results from that appendix are summarised in the following proposition:

Proposition 2 For any matrix $\mathcal{M} \in \mathbb{R}^{T \times T}$, there exists a model in the form of Problem 2 with a number of state variables given by a quadratic in T , such that $M = \mathcal{M}$ for that model, where M is defined as in equation (2), and such that for all $\varphi \in \mathbb{R}^T$, there exists an initial state for which $q = \varphi$, where q is the path of the bounded variable when constraints are ignored.

We now introduce some definitions of matrix properties that are necessary for the statement of our key existence and uniqueness results. The ultimate properties of the solutions to the OBC model are determined by which of these matrix properties M possesses. In each case, we give the definitions in a constructive form which makes

clear both how the property might be verified computationally, and the links between definitions. These are not necessarily in the form which is standard in the original literature, however. For both the original definitions, and the proofs of equivalence between the ones below and the originals, see Cottle, Pang, and Stone (2009a) and Xu (1993) (for the characterisation of sufficient models).

Definition 1 (Principal sub-matrix, Principal minor) For a matrix $M \in \mathbb{R}^{T \times T}$, the **principal sub-matrices** of M are the matrices:

$$\left\{ [M_{i,j}]_{i,j=k_1, \dots, k_S} \mid S, k_1, \dots, k_S \in \{1, \dots, T\}, k_1 < k_2 < \dots < k_S \right\},$$

i.e. the **principal sub-matrices** of M are formed by deleting the same rows and columns. The **principal minors** of M are the collection of values:

$$\left\{ \det \left([M_{i,j}]_{i,j=k_1, \dots, k_S} \right) \mid S, k_1, \dots, k_S \in \{1, \dots, T\}, k_1 < k_2 < \dots < k_S \right\},$$

i.e. the **principal minors** of M are the determinants of the principal sub-matrices of M .

Definition 2 (P_0 -matrix) A matrix $M \in \mathbb{R}^{T \times T}$ is called a **P-matrix (P₀-matrix)** if the principal minors of M are all strictly (weakly) positive. *Note: for symmetric M , M is a P_0 -matrix if and only if all of its eigenvalues are strictly (weakly) positive.*

Definition 3 (General positive (semi-)definite) A matrix $M \in \mathbb{R}^{T \times T}$ is called **general positive (semi-)definite** if $M + M'$ is a P-matrix (P₀-matrix). *If M is symmetric, then, M is general positive (semi-)definite if and only if it is positive (semi-)definite.*

Definition 4 (Sufficient matrices) Let $M \in \mathbb{R}^{T \times T}$. M is called **column sufficient** if M is a P₀-matrix, and for each principal sub-matrix $W := [M_{i,j}]_{i,j=k_1, \dots, k_S}$ of M , with zero determinant, and for each proper principal sub-matrix $[W_{i,j}]_{i,j=l_1, \dots, l_R}$ of W ($R < S$), with zero determinant, the columns of $[W_{i,j}]_{\substack{i=1, \dots, S \\ j=l_1, \dots, l_R}}$ do not form a basis for the column space of W .⁸ M is called **row sufficient** if M' is column sufficient. M is called **sufficient** if it is column sufficient and row sufficient.

Definition 5 (S_0 -matrix) A matrix $M \in \mathbb{R}^{T \times T}$ is called an **S-matrix (S₀-matrix)** if there exists $y \in \mathbb{R}^T$ such that $y > 0$ and $My \gg 0$ ($My \geq 0$).⁹

Definition 6 ((Strictly) Semi-monotone) A matrix $M \in \mathbb{R}^{T \times T}$ is called **(strictly) semi-monotone** if each of its principal sub-matrices is an **S₀-matrix (S-matrix)**.

⁸ This may be checked via the singular value decomposition.

⁹ These condition may be rewritten as $\sup\{\zeta \in \mathbb{R} \mid \exists y \geq 0 \text{ s.t. } \forall t \in \{1, \dots, T\}, (My)_t \geq \zeta \wedge y_t \leq 1\} > 0$, and $\sup\{\sum_{t=1}^T y_t \mid y \geq 0, My \geq 0 \wedge \forall t \in \{1, \dots, T\}, y_t \leq 1\} > 0$, respectively. As linear-programming problems, these may be verified in time polynomial in T using the methods described in e.g. Roos, Terlaky, and Vial (2006). Alternatively, by Ville's theorem of the alternative (Cottle, Pang, and Stone 2009b), M is not an S₀-matrix if and only if $-M'$ is an S-matrix.

Definition 7 ((Strictly) Copositive) A matrix $M \in \mathbb{R}^{T \times T}$ is called **(strictly) copositive** if $M + M'$ is (strictly) semi-monotone.¹⁰

Cottle, Pang, and Stone (2009a) note the following relationships between these classes (amongst others):

Lemma 3 The following hold:

- 1) All general positive semi-definite matrices are copositive and sufficient.
 - 2) P_0 includes skew-symmetric matrices, general positive semi-definite matrices, sufficient matrices and P-matrices.
 - 3) All P_0 -matrices, and all copositive matrices are semi-monotone, and all P-matrices, and all strictly copositive matrices are strictly semi-monotone (and hence also S-matrices).
-

Additionally, from considering the 1×1 principal sub-matrices of M , we have the following restrictions on the diagonal of M :

Lemma 4 All general positive semi-definite, semi-monotone, sufficient, P_0 and copositive matrices have non-negative diagonals, and all general positive definite, strictly semi-monotone, P and strictly copositive matrices have positive diagonals.

For many macroeconomic models, this simple condition is sufficient to rule out membership of these matrix classes, as medium-scale DSGE models¹¹ with a ZLB frequently have negative elements on the diagonal of their M matrix, when T is large enough. Thus, following the intuition of Figure 1, such models will satisfy the conditions to have multiple equilibria, though they will not be the only such models.

Unfortunately, for all of these matrix classes except the classes of general positive (semi-)definite matrices, and $S_{(0)}$ -matrices, no algorithm which runs in an amount of time that is polynomial in T is known, thus verifying class membership may not be feasible with large T . However, disproving class membership only requires finding one principal sub-matrix which fails to have the required property, and for this, starting with the 1×1 principal sub-matrices (e.g. the diagonal), then considering the 2×2 ones (etc.) is often a good strategy.¹²

A common intuition is that in models without state variables, M must be both a P matrix, and an S matrix. In fact, this is not true. Indeed, there are even purely static

¹⁰ Väliäho (1986) contains an alternative characterisation which avoids solving any linear programming problems.

¹¹ This applies, for example, to the Smets and Wouters (2003) model, as we will show in section 3.5.

¹² The facts that all of the eigenvalues of a $T \times T$ P-matrix have complex arguments in the interval $(-\pi + \frac{\pi}{T}, \pi - \frac{\pi}{T})$, and all of the eigenvalues of a $T \times T$ P_0 -matrix have complex arguments in the interval $[-\pi + \frac{\pi}{T}, \pi - \frac{\pi}{T}]$ (Fang 1989) may also assist in ruling out these matrix classes.

models for which M is not in either of these classes. For example, in online appendix B, we construct a purely static model for which $M_{1:\infty,1:\infty} = -I_{\infty \times \infty}$, which is neither a P-matrix, nor an S-matrix, for any T .

2.4. Existence results

We start by considering necessary or sufficient conditions for the existence of a solution to a model with occasionally binding constraints. Ideally, we would like the solution to exist for any possible path the bounded variable might have taken in the future were there no OBC, i.e. for any possible q . To see this, note that under a perfect foresight exercise we are ignoring the fact that shocks might hit the economy in future. More properly, we ought to integrate over future uncertainty, as in the stochastic extended path approach of Adjemian and Juillard (2013). A crude way to do this would just be to draw lots of samples of future shocks for periods $1, \dots, S$, and average over these draws. However, in a linear model with shocks with unbounded support, providing at least one shock has an impact on a given variable, the distribution of future paths of that variable has positive support over the entirety of \mathbb{R}^S . Thus, ideally we would like M to be such that for any q , the linear complementarity problem (q, M) has a solution.

Definition 8 (Feasible LCP) Suppose $q \in \mathbb{R}^T$ and $M \in \mathbb{R}^{T \times T}$ are given. The LCP corresponding to M and q is called **feasible** if there exists $y \in \mathbb{R}^T$ such that $y \geq 0$ and $q + My \geq 0$.

By construction, if an LCP (q, M) has a solution, then it is feasible, i.e. being feasible is a necessary condition for existence. Checking feasibility is straightforward for any particular (q, M) , since to find a feasible solution we just need to solve a standard linear programming problem, which is possible in an amount of time that is polynomial in T .

Note that if the LCP (q, M) is not feasible, then for any $\hat{q} \leq q$, if $y \geq 0$, then $\hat{q} + My \leq q + My < 0$ since (q, M) is not feasible, so the LCP (\hat{q}, M) is also not feasible. Consequently, if there are any q for which the LCP is non-feasible, then there is a positive measure of such q . Thus, in a model with uncertainty, if there are some q for which the model has no solution satisfying the terminal condition, even with arbitrarily large T , then the model will have no solution satisfying the terminal condition with positive probability. This in turn means that it is not consistent with rationality for agents to believe that our terminal condition is satisfied with certainty, so they would have to place some positive probability on getting stuck in an alternative steady-state.

The following proposition gives an easily verified necessary condition for the global existence of a solution to the model with occasionally binding constraints, given some fixed horizon T :

Proposition 3 The LCP (q, M) is feasible for all $q \in \mathbb{R}^T$ if and only if M is an S-matrix. (Cottle, Pang, and Stone 2009a)¹³

Of course, it may be the case that the M matrix is only an S-matrix when T is very large, so we must be careful in using this condition to imply non-existence of a solution. Furthermore, it may be the case that although there exists some $y \in \mathbb{R}^T$ with $y \geq 0$ such that $M_{1:T,1:T}y \gg 0$, where we are indexing the M matrix by its size for clarity, for any such y , $\inf_{t \in \mathbb{N}^+} M_{t,1:T}y < 0$, so for some $q \in \mathbb{R}^{\mathbb{N}^+}$, the infinite LCP $(q, M_{1:\infty,1:\infty})$ is not feasible under the additional restriction that $y_t = 0$ for $t > T$. Strictly, it is this infinite LCP which we ought to be solving, subject to the additional constraint that y has only finitely many non-zero elements, which is implied by our terminal condition.

By Proposition 3, this infinite problem is feasible if and only if:

$$\zeta := \sup_{\substack{y \in [0,1]^{\mathbb{N}^+} \\ \exists T \in \mathbb{N} \text{ s.t. } \forall t > T, y_t = 0}} \inf_{t \in \mathbb{N}^+} M_{t,1:\infty}y > 0.$$

Consequently, if $\zeta > 0$ then for every $q \in \mathbb{R}^{\mathbb{N}^+}$, for sufficiently large T , the finite problem $(q_{1:T}, M_{1:T,1:T})$ will be feasible, which is a sufficient condition for solvability. In order to evaluate this limit, we first need to derive constructive bounds on the M matrix for large T . We do this in the online appendix C, where we prove that the rows and columns of M are converging to 0 (with constructive bounds), and that the k^{th} diagonal of the M matrix is converging to the value $d_{1,k}$, to be defined (again with constructive bounds), where diagonals are indexed such that the principal diagonal is index 0, and indices increase as one moves up and to the right in the M matrix. To explain the origins of $d_{1,k}$ we note the following lemma proved in online appendix C:

Lemma 5 The (time-reversed) difference equation $A\hat{d}_{k+1} + B\hat{d}_k + C\hat{d}_{k-1} = 0$ for all $k \in \mathbb{N}^+$ has a unique solution satisfying the terminal condition $\hat{d}_k \rightarrow 0$ as $k \rightarrow \infty$, given by $\hat{d}_k = H\hat{d}_{k-1}$, for all $k \in \mathbb{N}^+$, for some H with eigenvalues in the unit circle.

Then, we define $d_0 := -(AH + B + CF)^{-1}I_{\cdot,1}$, $d_k = Hd_{k-1}$, for all $k \in \mathbb{N}^+$, and $d_{-t} = Fd_{-(t-1)}$, for all $t \in \mathbb{N}^+$, so d_k follows the time reversed difference equation for positive indices, and the original difference equation for negative indices. This is opposite to what one might perhaps expect since time is increasing as one descends the rows of M , but diagonal indices are decreasing as one descends in M .

Using the resulting bounds on M , we can construct upper and lower bounds on ζ , which are described in the following propositions, also proven in online appendix C:

¹³ Most of the results on LCPs in both this and the following section are restatements of (assorted) results contained in Cottle, Pang, and Stone (2009a) and Väliäho (1986) (for the characterisation of “copositive-plus” matrices), and the reader is referred to those works for proofs and further references.

Proposition 4 There exists $\underline{\zeta}_T, \bar{\zeta}_T \geq 0$, defined in the online appendix C, computable in time polynomial in T , such that $\underline{\zeta}_T \leq \zeta \leq \bar{\zeta}_T$, and $|\underline{\zeta}_T - \bar{\zeta}_T| \rightarrow 0$ as $T \rightarrow \infty$.

These conditions give simple tests for feasibility or non-feasibility with sufficiently large T .

We now turn to sufficient conditions for the existence of a solution for some finite T .

Proposition 5 The LCP (q, M) is solvable if it is feasible and, either:

1. M is row-sufficient, or,
2. M is copositive and for all non-singular principal sub-matrices W of M , all non-negative columns of W^{-1} possess a non-zero diagonal element.

(Cottle, Pang, and Stone 2009a; Väliäho 1986)

If either condition 1 or condition 2 of Proposition 5 is satisfied, then to check existence for any particular q , we only need to solve a linear programming problem to see if a solution exists for a particular q . As this may be substantially faster than solving the LCP, this may be helpful in practice.

Proposition 6 The LCP (q, M) is solvable for all $q \in \mathbb{R}^T$, if at least one of the following conditions holds:

1. M is an S-matrix, and either condition 1 or condition 2 of Proposition 5 are satisfied.
2. M is copositive with non-zero principal minors.
3. M is a P-matrix, a strictly copositive matrix or a strictly semi-monotone matrix.

(Cottle, Pang, and Stone 2009a)

If condition 1, 2 or 3 of Proposition 6 is satisfied, then we know that the LCP will always have a solution. Therefore, for any path of the bounded variable in the absence of the bound, we will also be able to solve the model when the bound is imposed. Monetary policy makers should always choose a policy rule that produces a model that satisfies one of these three conditions, if they can, since otherwise there is a positive probability that only solutions converging to the “bad” steady-state will exist in some state of the world.

Ideally, we might have liked conditions for the existence of a solution that are both necessary and sufficient, but unfortunately at present no such conditions exist in full generality. However, in the special case of M matrices with nonnegative entries, we have the following result:

Proposition 7 If M is a matrix with nonnegative entries, then the LCP (q, M) is solvable for all $q \in \mathbb{R}^T$, if and only if M has a strictly positive diagonal. (Cottle, Pang, and Stone 2009a)

2.5. Uniqueness results

While no fully general necessary and sufficient conditions have been derived for existence, such conditions are available for uniqueness, in particular:

Proposition 8 The LCP (q, M) has a unique solution for all $q \in \mathbb{R}^T$, if and only if M is a P-matrix. If M is not a P-matrix, then the LCP (q, M) has multiple solutions for some q . (Samelson, Thrall, and Wesler 1958; Cottle, Pang, and Stone 2009a)

This proposition is the equivalent for models with OBCs of the key proposition of Blanchard and Kahn (1980). By testing whether our matrix M is a P-matrix we can immediately determine if the model possesses a unique solution in any state of the world, and for any sequence of future shocks, for a fixed T . In our experience, this condition is satisfied in efficient models, such as models of irreversible investment, as one would expect, but is not generally satisfied in medium-scale New-Keynesian models with a ZLB on nominal interest rates. Given that if M is a P-matrix, so too are all its principal sub-matrices, if we see that M is not a P-matrix for some T , then we know that with larger T it would also not be a P-matrix. Thus, if for some T , M is not a P-matrix, then we know that the model does not have a unique solution, even for arbitrarily large T . Alternatively, we can prove that with large T some M is not a P-matrix by using the analytic formula for the limit of its diagonal given in the previous section, i.e. $d_{0,1} = -I_{1,1}(AH + B + CF)^{-1}I_{1,1}$. If this value is negative, then we know that with sufficiently large T , M will not be a P-matrix.

Since some classes of models almost never possess a unique solution when at the zero lower bound, we might reasonably require a lesser condition, namely that at least when the solution to the model without a bound is a solution to the model with the bound, then it ought to be the unique solution. This is equivalent to requiring that when q is non-negative, the LCP (q, M) has a unique solution. Conditions for this are given in the following propositions:

Proposition 9 The LCP (q, M) has a unique solution for all $q \in \mathbb{R}^T$ with $q \gg 0$ ($q \geq 0$) if and only if M is (strictly) semi-monotone. (Cottle, Pang, and Stone 2009a)

Hence, by verifying that M is (strictly) semi-monotone, we can reassure ourselves that merely introducing the bound will not change the solution away from the bound. When this condition is violated, even when the economy is a long way from the bound, there may be solutions which jump to the bound. Again, since principal sub-matrices of (strictly) semi-monotone are (strictly) semi-monotone, a failure of (strict) semi-monotonicity for some T implies a failure for all larger T . Furthermore, if $d_{0,1} < 0$ then again for sufficiently large T , M cannot be semi-monotone.

Where there are multiple solutions, we might like to be able to select one via some objective function. This is particularly tractable when either the number of solutions is finite, or the solution set is convex. Conditions for this are given in online appendix D.

2.6. Results from dynamic programming

Alternative existence and uniqueness results for the infinite T problem can be established via dynamic programming methods, under the assumption that Problem 2 comes from the first order conditions solution of a social planner problem. These have the advantage that their conditions are potentially much easier to evaluate, though they also have somewhat limited applicability. We focus here on uniqueness results, since these are generally of greater interest.

Suppose that the social planner in some economy solves the following problem:

Problem 5 Suppose $\mu \in \mathbb{R}^n$, $\Psi^{(0)} \in \mathbb{R}^{c \times 1}$ and $\Psi^{(1)} \in \mathbb{R}^{c \times 2n}$ are given, where $c \in \mathbb{N}$.

Define $\tilde{\Gamma}: \mathbb{R}^n \rightarrow \mathbb{P}(\mathbb{R}^n)$ (where \mathbb{P} denotes the power-set operator) by:

$$\tilde{\Gamma}(x) = \left\{ z \in \mathbb{R}^n \mid 0 \leq \Psi^{(0)} + \Psi^{(1)} \begin{bmatrix} x - \mu \\ z - \mu \end{bmatrix} \right\}, \quad (5)$$

for all $x \in \mathbb{R}^n$. (Note: $\tilde{\Gamma}(x)$ will give the set of feasible values for next period's state if the current state is x . Equality constraints may be included by including an identical lower bound and upper bound.) Define:

$$\tilde{X} := \{x \in \mathbb{R}^n \mid \tilde{\Gamma}(x) \neq \emptyset\}, \quad (6)$$

and suppose without loss of generality that for all $x \in \mathbb{R}^n$, $\tilde{\Gamma}(x) \cap \tilde{X} = \tilde{\Gamma}(x)$. (Note: this means that the linear inequalities bounding \tilde{X} are already included in those in the definition of $\tilde{\Gamma}(x)$. It is without loss of generality as the planner will never choose an $\tilde{x} \in \tilde{\Gamma}(x)$ such that $\tilde{\Gamma}(\tilde{x}) = \emptyset$.) Further define $\tilde{\mathcal{F}}: \tilde{X} \times \tilde{X} \rightarrow \mathbb{R}$ by:

$$\tilde{\mathcal{F}}(x, z) = u^{(0)} + u^{(1)} \begin{bmatrix} x - \mu \\ z - \mu \end{bmatrix} + \frac{1}{2} \begin{bmatrix} x - \mu \\ z - \mu \end{bmatrix}' \tilde{u}^{(2)} \begin{bmatrix} x - \mu \\ z - \mu \end{bmatrix}, \quad (7)$$

for all $x, z \in \tilde{X}$, where $u^{(0)} \in \mathbb{R}$, $u^{(1)} \in \mathbb{R}^{1 \times 2n}$ and $\tilde{u}^{(2)} = \tilde{u}^{(2)'} \in \mathbb{R}^{2n \times 2n}$ are given.

Finally, suppose $x_0 \in \tilde{X}$ is given and $\beta \in (0, 1)$, and choose x_1, x_2, \dots to maximise:

$$\liminf_{T \rightarrow \infty} \sum_{t=1}^T \beta^{t-1} \tilde{\mathcal{F}}(x_{t-1}, x_t) \quad (8)$$

subject to the constraints that for all $t \in \mathbb{N}^+$, $x_t \in \tilde{\Gamma}(x_{t-1})$.

To ensure the problem is well behaved, we make the following assumption:

Assumption 3 $\tilde{u}^{(2)}$ is negative-definite.

In online appendix F, we establish the following result:

Proposition 10 If either \tilde{X} is compact, or, $\tilde{\Gamma}(x)$ is compact valued and $x \in \tilde{\Gamma}(x)$ for all $x \in \tilde{X}$, then for all $x_0 \in \tilde{X}$, there is a unique path $(x_t)_{t=0}^{\infty}$ which solves Problem 5.

We wish to use this result to establish the uniqueness of the solution to the first order conditions. The Lagrangian for our problem is given by:

$$\sum_{t=1}^{\infty} \beta^{t-1} \left[\tilde{\mathcal{F}}(x_{t-1}, x_t) + \lambda'_{\Psi,t} \left[\Psi^{(0)} + \Psi^{(1)} \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix} \right] \right], \quad (9)$$

for some KKT-multipliers $\lambda_t \in \mathbb{R}^c$ for all $t \in \mathbb{N}^+$. Taking the first order conditions leads to the following necessary KKT conditions, for all $t \in \mathbb{N}^+$:

$$0 = u_{:,2}^{(1)} + \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix}' \tilde{u}_{:,2}^{(2)} + \lambda_t' \Psi_{:,2}^{(1)} + \beta \left[u_{:,1}^{(1)} + \begin{bmatrix} x_t - \mu \\ x_{t+1} - \mu \end{bmatrix}' \tilde{u}_{:,1}^{(2)} + \lambda_{t+1}' \Psi_{:,1}^{(1)} \right], \quad (10)$$

$$0 \leq \Psi^{(0)} + \Psi^{(1)} \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix}, \quad 0 \leq \lambda_t, \quad 0 = \lambda_t \circ \left[\Psi^{(0)} + \Psi^{(1)} \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix} \right], \quad (11)$$

where subscripts 1 and 2 refer to blocks of rows or columns of length n . Additionally, for μ to be the steady-state of x_t and $\bar{\lambda}$ to be the steady-state of λ_t , we require:

$$0 = u_{:,2}^{(1)} + \bar{\lambda}' \Psi_{:,2}^{(1)} + \beta \left[u_{:,1}^{(1)} + \bar{\lambda}' \Psi_{:,1}^{(1)} \right], \quad (12)$$

$$0 \leq \Psi^{(0)}, \quad 0 \leq \bar{\lambda}, \quad 0 = \bar{\lambda} \circ \Psi^{(0)}. \quad (13)$$

In online appendix G we prove the following result:

Proposition 11 Suppose that for all $t \in \mathbb{N}$, $(x_t)_{t=1}^{\infty}$ and $(\lambda_t)_{t=1}^{\infty}$ satisfy the KKT conditions given in equations (10) and (11), and that as $t \rightarrow \infty$, $x_t \rightarrow \mu$ and $\lambda_t \rightarrow \bar{\lambda}$, where μ and λ satisfy the steady-state KKT conditions given in equations (12) and (13). Then $(x_t)_{t=1}^{\infty}$ solves Problem 5. If, further, either condition of Proposition 10 is satisfied, then $(x_t)_{t=1}^{\infty}$ is the unique solution to Problem 5, and there can be no other solutions to the KKT conditions given in equations (10) and (11) satisfying $x_t \rightarrow \mu$ and $\lambda_t \rightarrow \bar{\lambda}$ as $t \rightarrow \infty$.

Now, it is possible to convert the KKT conditions given in equations (10) and (11) into a problem in the form of the multiple-bound generalisation of Problem 2 quite generally.

To see this, first note that we may rewrite equation (10) as:

$$0 = u_{:,2}^{(1)'} + \tilde{u}_{2,1}^{(2)}(x_{t-1} - \mu) + \tilde{u}_{2,2}^{(2)}(x_t - \mu) + \Psi_{:,2}^{(1)'} \lambda_t + \beta \left[u_{:,1}^{(1)'} + \tilde{u}_{1,1}^{(2)}(x_t - \mu) + \tilde{u}_{1,2}^{(2)}(x_{t+1} - \mu) + \Psi_{:,1}^{(1)'} \lambda_{t+1} \right].$$

Now, $\tilde{u}_{2,2}^{(2)} + \beta u_{1,1}^{(2)}$ is negative definite, hence it is valid to define $\mathcal{V} := \Psi_{:,2}^{(1)} [\tilde{u}_{2,2}^{(2)} + \beta \tilde{u}_{1,1}^{(2)}]^{-1}$, and equation (9) implies that:

$$\begin{aligned} & \Psi^{(0)} + \Psi^{(1)} \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix} \\ &= \Psi^{(0)} + (\Psi_{:,1}^{(1)} - \mathcal{V} \tilde{u}_{2,1}^{(2)})(x_{t-1} - \mu) - \mathcal{V} \left[u_{:,2}^{(1)'} + \beta \left[u_{:,1}^{(1)'} + \tilde{u}_{1,2}^{(2)}(x_{t+1} - \mu) + \Psi_{:,1}^{(1)'} \lambda_{t+1} \right] \right] \\ & \quad - \Psi_{:,2}^{(1)} [\tilde{u}_{2,2}^{(2)} + \beta \tilde{u}_{1,1}^{(2)}]^{-1} \Psi_{:,2}^{(1)'} \lambda_t. \end{aligned} \quad (14)$$

Moreover, equation (11) implies that if the k^{th} element of $\Psi^{(0)} + \Psi^{(1)} \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix}$ is strictly positive, then the k^{th} element of λ_t is zero, so:

$$\Psi^{(0)} + \Psi^{(1)} \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix} = \max\{0, z_t\}, \quad (15)$$

where:

$$\begin{aligned}
z_t := & \Psi^{(0)} + (\Psi_{:,1}^{(1)} - \mathcal{V} \tilde{u}_{2,1}^{(2)})(x_{t-1} - \mu) \\
& - \mathcal{V} \left[u_{:,2}^{(1)'} + \beta \left[u_{:,1}^{(1)'} + \tilde{u}_{1,2}^{(2)}(x_{t+1} - \mu) + \Psi_{:,1}^{(1)'} \lambda_{t+1} \right] \right] \\
& - \left[\Psi_{:,2}^{(1)} [\tilde{u}_{2,2}^{(2)} + \beta \tilde{u}_{1,1}^{(2)}]^{-1} \Psi_{:,2}^{(1)'} + \mathcal{W} \right] \lambda_t,
\end{aligned}$$

and $\mathcal{W} \in \mathbb{R}^{c \times c}$ is an arbitrary, strictly positive diagonal matrix. A natural choice is:

$$\mathcal{W} := -\text{diag} \text{diag} \left[\Psi_{:,2}^{(1)} [\tilde{u}_{2,2}^{(2)} + \beta \tilde{u}_{1,1}^{(2)}]^{-1} \Psi_{:,2}^{(1)'} \right],$$

providing this is strictly positive (it is weakly positive at least as $\tilde{u}_{2,2}^{(2)} + \beta \tilde{u}_{1,1}^{(2)}$ is negative definite), where the diag operator maps matrices to a vector containing their diagonal, and maps vectors to a matrix with the given vector on the diagonal, and zeros elsewhere.

We claim that we may replace equation (11) with equation (15) without changing the model. We have already shown that equation (11) implies equation (15), so we just have to prove the converse. We continue to suppose equation (9) holds, and thus, so too does equation (14). Then, from subtracting equation (14) from equation (15), we have that $\mathcal{W} \lambda_t = \max\{-z_t, 0\}$. Hence, as \mathcal{W} is a strictly positive diagonal matrix, and the right hand side is weakly positive, $\lambda_t \geq 0$. Furthermore, the k th element of λ_t is non-negative if and only if the k th element of z_t is non-positive (as \mathcal{W} is a strictly positive diagonal matrix), which in turn holds if and only if the k th element of $\Psi^{(0)} + \Psi^{(1)} \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix}$ is equal to zero, by equation (15). Thus equation (11) is satisfied.

Combined with our previous results, this gives the following proposition:

Proposition 12 Suppose we are given a problem in the form of Problem 5. Then, the KKT conditions of that problem may be placed into the form of the multiple-bound generalisation of Problem 2. Let (q_{x_0}, M) be the infinite LCP corresponding to this representation, given initial state $x_0 \in \tilde{X}$. Then, if y is a solution to the LCP, $q_{x_0} + My$ gives the stacked paths of the bounded variables in a solution to Problem 5. If, further, either condition of Proposition 10 is satisfied, then this LCP has a unique solution for all $x_0 \in \tilde{X}$, which gives the unique solution to Problem 5, and, for sufficiently large T^* , the finite LCP $(q_{x_0}^{(T^*)}, M^{(T^*)})$ has a unique solution $y^{(T^*)}$ for all $x_0 \in \tilde{X}$, where $q_{x_0}^{(T^*)} + M^{(T^*)}y^{(T^*)}$ gives the first T^* periods of the stacked paths of the bounded variables in a solution to Problem 5.

This proposition provides some evidence that the LCP will have a unique solution when it is generated from a dynamic programming problem with a unique solution. In online appendix H, we derive similar results for models with more general constraints and objective functions. The proof of this proposition also showed how one can convert KKT conditions into equations of the form handled by our methods.

3. Applications to New Keynesian models

Brendon, Paustian, and Yates (2015) (henceforth: BPY) consider multiple equilibria in a simple New Keynesian (NK) model with an output growth rate term in the Taylor rule. They show that with sufficiently large reaction to the growth rate, there can be multiple equilibria today, even when the policy rule used to form tomorrow's expectations is held fixed. This is equivalent to the existence of multiple equilibria even when $T = 1$. In the first subsection here, we give an alternative analytic proof of this using our results, and discuss the generalisation to higher T .

BPY go on to show numerical results from the model with persistence in the shadow nominal interest rate (i.e. the rate which would obtain were it not for the ZLB). In the second subsection here we numerically examine this case with larger T , illustrating how multiple equilibria tend to become easier to support when T is large.

With an appropriately constructed limit, price level targeting is the result of increasing the persistence in the monetary rule to unity. In new analytic results in the third subsection here, we show that under price level targeting, with $T = 1$, M is a P-matrix regardless of the coefficients on prices and output. We also show that this continues to hold for large T , and that the model is always feasible for sufficiently large T .

However, we do not want to give the impression that multiplicity and non-existence are only caused by the central bank responding to the growth rate, or that they are only a problem in carefully constructed theoretical examples. In subsection 3.4, we show that a standard NK model with positive steady-state inflation and a ZLB possesses multiple equilibria in some states, and no solutions in others, even with an entirely standard Taylor rule. We also show that here too price level targeting is sufficient to restore determinacy. Finally, in the last sub-section we show that these conclusions also carry through to the posterior-modes of the Smets and Wouters (2003; 2007) models.

3.1. The simple Brendon, Paustian, and Yates (2015) (BPY) model

The equations of the simple Brendon, Paustian, and Yates (2015) model are as follows:

$$x_{i,t} = \max\{0, 1 - \beta + \alpha_{\Delta y}(x_{y,t} - x_{y,t-1}) + \alpha_{\pi}x_{\pi,t}\},$$

$$x_{y,t} = \mathbb{E}_t x_{y,t+1} - \frac{1}{\sigma}(x_{i,t} + \beta - 1 - \mathbb{E}_t x_{\pi,t+1}), \quad x_{\pi,t} = \beta \mathbb{E}_t x_{\pi,t+1} + \gamma x_{y,t},$$

where $x_{i,t}$ is the nominal interest rate, $x_{y,t}$ is the deviation of output from steady-state, $x_{\pi,t}$ is the deviation of inflation from steady-state, and $\beta \in (0,1)$, $\gamma, \sigma, \alpha_{\Delta y} \in (0, \infty)$, $\alpha_{\pi} \in (1, \infty)$ are parameters. In online appendix I, we prove the following:

Proposition 13 The BPY model is in the form of Problem 2, and satisfies Assumptions 1 and 2. With $T = 1$, $M < 0$ ($M = 0$) if and only if $\alpha_{\Delta y} > \sigma \alpha_{\pi}$ ($\alpha_{\Delta y} = \sigma \alpha_{\pi}$).

For a 1×1 matrix, checking the conditions from section 2.3 is trivial. In particular, we have that if $\alpha_{\Delta y} < \sigma \alpha_{\pi}$, M is a general positive definite, strictly semi-monotone, strictly co-positive, sufficient, P, S matrix; if $\alpha_{\Delta y} \leq \sigma \alpha_{\pi}$, M is a general positive semi-definite, semi-monotone, co-positive, sufficient, P_0 , S_0 matrix. Hence, when $T = 1$, if $\alpha_{\Delta y} < \sigma \alpha_{\pi}$, the model has a unique solution for all q ; if $\alpha_{\Delta y} \leq \sigma \alpha_{\pi}$, the model has a unique solution whenever $q > 0$, and at least one solution when $q = 0$. When $\alpha_{\Delta y} > \sigma \alpha_{\pi}$, M is negative, and so for any positive q , there exists $y > 0$ such that $q + My = 0$, so the model has multiple solutions. I.e. there are solutions that jump to the bound, even when the nominal interest rate would always be positive were there no bound at all.

We illustrate this by adding a shock to the Euler equation, and showing impulse responses for alternative solutions. In particular, we replace the Euler equation with:

$$x_{y,t} = \mathbb{E}_t x_{y,t+1} - \frac{1}{\sigma} (x_{i,t} + \beta - 1 - \mathbb{E}_t x_{\pi,t+1} - (0.01)\varepsilon_t),$$

and take the parameterisation $\sigma = 1$, $\beta = 0.99$, $\gamma = \frac{(1-0.85)(1-\beta(0.85))}{0.85} (2 + \sigma)$, $\rho = 0.5$, following BPY, and we additionally set $\alpha_{\pi} = 1.5$ and $\alpha_{\Delta y} = 1.6$, to ensure we are in the region with multiple solutions. In Figure 2, we show two alternative solutions to the impulse response to a magnitude 1 shock to ε_t . The solid line in the left plot gives the solution which minimises $\|y\|_{\infty}$. This solution never hits the bound, and is moderately expansionary. The solid line in the right plot gives the solution which minimises $\|q + My\|_{\infty}$. (The dotted line in the right plot repeats the left plot, for comparison.) This solution stays at the bound for two periods, and is strongly contractionary, with a magnitude around 100 times larger than the other solution.

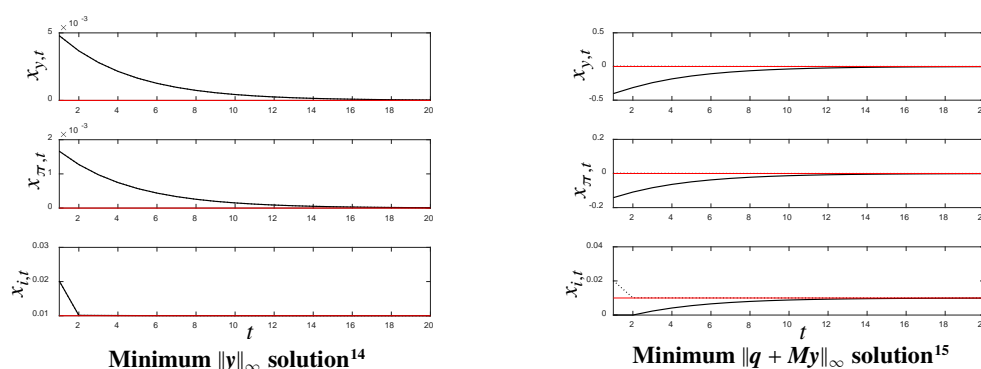


Figure 2: Alternative solutions following a magnitude 1 impulse to ε_t

When $T > 1$, the previous results imply that if $\alpha_{\Delta y} > \sigma \alpha_{\pi}$, then M is neither P_0 , general positive semi-definite, semi-monotone, co-positive, nor sufficient, since the top-left 1×1 principal sub-matrix of M is the same as when $T = 1$. Thus, if anything, when $T > 1$, the parameter region in which there are multiple solutions (when away

¹⁴ Calculated by setting $\omega = 1000$, in the terminology of section 4.2.

¹⁵ Calculated by setting $\omega = 0.01$, in the terminology of section 4.2.

from the bound or at it) is larger. However, numerical experiments suggest that this parameter region in fact remains the same as T increases, which is unsurprising given the weak persistence of this model. Thus, if we want more interesting results with higher T , we need to consider a model with a stronger persistence mechanism.

3.2. The BPY model with shadow interest rate persistence

We introduce persistence in the shadow interest rate by replacing the previous Taylor rule with $x_{i,t} = \max\{0, x_{d,t}\}$, where $x_{d,t}$, the shadow nominal interest rate is given by:

$$x_{d,t} = (1 - \rho)(1 - \beta + \alpha_{\Delta y}(x_{y,t} - x_{y,t-1}) + \alpha_{\pi}x_{\pi,t}) + \rho x_{d,t-1}.$$

It is easy to verify that this may be put in the form of Problem 2, and that with $\beta \in (0,1)$, $\gamma, \sigma, \alpha_{\Delta y} \in (0, \infty)$, $\alpha_{\pi} \in (1, \infty)$, $\rho \in (-1,1)$, Assumption 2 is satisfied. For our numerical exercise, we again set $\sigma = 1$, $\beta = 0.99$, $\gamma = \frac{(1-0.85)(1-\beta(0.85))}{0.85}(2 + \sigma)$, $\rho = 0.5$, following BPY.

In Figure 3, we plot the regions in $(\alpha_{\Delta y}, \alpha_{\pi})$ space in which M is a P-matrix (P₀-matrix) when $T = 2$ or $T = 4$. For this model, these correspond to the regions in which M is strictly semi-monotone (semi-monotone). As may be seen, in the smaller T case, the P-matrix region is much larger. This relationship appears to continue to hold for both larger and smaller T , with the equivalent $T = 1$ plot being almost entirely shaded, and the large T plot apparently tending to the equivalent plot from the model without monetary policy persistence. Intuitively, the persistence in the shadow nominal interest rate dampens the immediate response of nominal interest rates to inflation and output growth, making it harder to induce a zero lower bound episode over short-horizons.




Figure 3: Regions in which M is a P-matrix (shaded grey) or a P₀-matrix (shaded grey, plus the black line), when $T = 2$ (left) or $T = 4$ (right).

Further evidence that the long-horizon behaviour is the same as in the model without persistence is provided by the fact that with $\alpha_{\pi} = 1.5$ and $\alpha_{\Delta y} = 1.05$,¹⁶ then M is a P-matrix, and from Proposition 4 we have that $\zeta > 6.131 \times 10^{-8}$, so M is an S-matrix for

¹⁶ Results for larger $\alpha_{\Delta y}$ were impossible due to numerical errors.

all sufficiently large T . Furthermore, with $\alpha_\pi = 1.5$ and $\alpha_{\Delta y} = 1.51$, then with $T = 200$, M is not an S-matrix,¹⁷ and from Proposition 4, $\zeta \leq 0 + \text{numerical error}$, providing strong numerical evidence that for all sufficiently large T , the LCP (q, M) is not feasible for some q , and hence that the model does not always possess a solution.

3.3. The BPY model with price targeting

An alternative way to introduce persistence to the shadow interest rate is to set:

$$\begin{aligned} x_{d,t} &= (1 - \rho) \left(1 - \beta + \frac{\alpha_{\Delta y}}{1 - \rho} (x_{y,t} - x_{y,t-1}) + \frac{\alpha_\pi}{1 - \rho} x_{\pi,t} \right) + \rho x_{d,t-1} \\ &= (1 - \rho)(1 - \beta) + (\alpha_{\Delta y}(x_{y,t} - x_{y,t-1}) + \alpha_\pi x_{\pi,t}) + \rho x_{d,t-1}, \end{aligned}$$

which is as before apart from a missing $(1 - \rho)$ multiplying the second bracketed term. In the limit as $\rho \rightarrow 1$, this tends to:

$$x_{d,t} = 1 - \beta + \alpha_{\Delta y} x_{y,t} + \alpha_\pi x_{p,t}$$

where $x_{p,t}$ is the price level, so $x_{\pi,t} = x_{p,t} - x_{p,t-1}$. This is a level targeting rule, with nominal GDP targeting as a special case with $\alpha_{\Delta y} = \alpha_\pi$. Note that the omission of the $(1 - \rho)$ coefficient on $\alpha_{\Delta y}$ and α_π is akin to having a “true” response to output growth of $\frac{\alpha_{\Delta y}}{1 - \rho}$ and a “true” response to inflation of $\frac{\alpha_\pi}{1 - \rho}$, so in the limit as $\rho \rightarrow 1$, we effectively have an infinitely strong response to these quantities. It turns out that this is sufficient to produce determinacy for all $\alpha_{\Delta y}, \alpha_\pi \in (0, \infty)$.

In particular, given the model:

$$\begin{aligned} x_{i,t} &= \max\{0, 1 - \beta + \alpha_{\Delta y} x_{y,t} + \alpha_\pi x_{p,t}\}, \\ x_{y,t} &= \mathbb{E}_t x_{y,t+1} - \frac{1}{\sigma} (x_{i,t} + \beta - 1 - \mathbb{E}_t x_{p,t+1} + x_{p,t}), \\ x_{p,t} - x_{p,t-1} &= \beta \mathbb{E}_t x_{p,t+1} - \beta x_{p,t} + \gamma x_{y,t}, \end{aligned}$$

we prove in online appendix J that the following proposition holds:

Proposition 14 The BPY model with price targeting is in the form of Problem 2, and satisfies Assumptions 1 and 2. With $T = 1$, $M > 0$ for all $\alpha_\pi \in (0, \infty)$, $\alpha_{\Delta y} \in [0, \infty)$.

Furthermore, with $\sigma = 1$, $\beta = 0.99$, $\gamma = \frac{(1-0.85)(1-\beta(0.85))}{0.85} (2 + \sigma)$, as before, and $\alpha_{\Delta y} = 1$, $\alpha_\pi = 1$, if we check our lower bound on ζ with $T = 20$, we find that $\zeta > 0.042$. Hence, this model is always feasible for any sufficiently large T . Given that $d_0 > 0$ for this model, and that for $T = 20$, M is a P-matrix, this is strongly suggestive of the existence of a unique solution for any q and for arbitrarily large T .

3.4. The linearized Fernández-Villaverde et al. (2012) model

The discussion of BPY might lead one to believe that multiplicity and non-existence is solely a consequence of overly aggressive monetary responses to output growth, and overly weak monetary responses to inflation. However, it turns out that in basic New

¹⁷ This was verified a second way by checking that $-M'$ was an S_0 -matrix, as discussed in footnote 9.

Keynesian models with positive inflation in steady-state, and hence price dispersion, even without any monetary response to output growth, and even with extremely aggressive monetary responses to inflation, there are still multiple equilibria in some states of the world, and no solutions in others. Price level targeting is again sufficient to fix these problems though.

We show these results in the Fernández-Villaverde et al. (2012) model, which is a basic non-linear New Keynesian model without capital or price indexation of non-resetting firms, but featuring (non-valued) government spending and steady-state inflation (and hence price-dispersion). We refer the reader to the original paper for the model's equations. After substitutions, the model has four non-linear equations which are functions of gross inflation, labour supply, price dispersion and an auxiliary variable introduced from the firms' price-setting first order condition. Of these variables, only price dispersion enters with a lag. We linearize¹⁸ the model around its steady-state, and then reintroduce the “max” operator which linearization removed from the Taylor rule.¹⁹ All parameters are set to the values given in Fernández-Villaverde et al. (2012). There is no term featuring output growth in the Taylor rule, so any multiplicity or non-existence in this model cannot be a consequence of the mechanism highlighted by BPY.

For this model, numerical calculations reveal that with $T \leq 14$, M is a P-matrix. However, with $T = 15$, M is neither a P nor an S matrix, and thus there are certainly some states of the world in which the model has multiple solutions, and others in which it has no solution at all.²⁰ This also implies that M is not a P-matrix for all larger T . Furthermore, with $T = 1000$, our upper bound on ζ from Proposition 4 implies that $\zeta \leq 0 + \text{numerical error}$, providing evidence that M is not an S-matrix for large T either.²¹

However, if we replace inflation in the monetary rule with the price level relative to its linear trend, which evolves according to:

$$x_{p,t} = x_{p,t-1} + x_{\pi,t} - x_{\pi}, \quad (16)$$

then with $T = 200$, we have that M is an S-matrix, and the lower bound from Proposition 4 implies that $\zeta > 0.003$, and hence that for all sufficiently large T , M is an S-matrix, so there is always a feasible solution.

¹⁸ Prior to linearization, we first transform the model's variables so that the transformed variables may take values on the entire real line. I.e. we work with the logarithms of labour supply, price dispersion and the auxiliary variable. For inflation, we note that inflation is always less than $\theta^{\frac{1}{1-\varepsilon}}$ (in the notation of Fernández-Villaverde et al. (2012)). Thus we work with a logit transformation of inflation over $\theta^{\frac{1}{1-\varepsilon}}$. This is generally more accurate than working with the logarithm of inflation.

¹⁹ This procedure is discussed in more detail in section 5.1.

²⁰ That New Keynesian models might have no solution at all in some states of the world has also been discussed by Basu and Bundick (2015), though their mechanism only applies in the stochastic model.

²¹ Since these results depend on the presence of the endogenous state, price dispersion, they are not directly related to the results of Davig and Leeper (2007). Further differences include the endogeneity of ZLB episodes here, and the fact that we are not making any restrictions on the solution space, which they do, as observed by Farmer, Waggoner, and Zha (2010).

3.5. The Smets and Wouters (2003) and Smets and Wouters (2007) models

Smets and Wouters (2003) and Smets and Wouters (2007) are the canonical medium-scale linear DSGE models, featuring assorted shocks, habits, price and wage indexation, capital (with adjustment costs), (costly) variable utilisation and quite general monetary policy reaction functions. The former model is estimated on Euro area data, while the latter is estimated on US data. The latter model also contains trend growth (permitting its estimation on non-detrended data), and a slightly more general aggregator across industries. However, overall, they are quite similar models, and any differences in their behaviour chiefly stems from differences in the estimated parameters. Since both models are incredibly well known in the literature, we omit their equations here, referring the reader to the original papers for further details.

To assess the likelihood of multiple equilibria at or away from the zero lower bound, we augment each model with a ZLB on nominal interest rates, and evaluate the properties of each model's M matrix with large T , at the estimated posterior-modes from the original papers. Note that we do not introduce an auxiliary for shadow nominal interest rates, so the monetary rules take the form of $x_{r,t} = \max\{0, (1 - \rho_r)(\dots) + \rho_r x_{r,t-1} + \dots\}$, in both cases.




Figure 4: The diagonals of the M matrices for the Smets and Wouters (2003; 2007) models

As shown in Lemma 4, if the diagonal of the M matrix ever goes negative, then the M matrix cannot be general positive semi-definite, semi-monotone, sufficient, P_0 or copositive, and hence the model will sometimes have multiple solutions even when away from the zero lower bound (i.e. for some strictly positive q). In Figure 4, we plot the diagonal of the M matrix for each model in turn,²² i.e. the impact on nominal interest rates in period t of news in period 1 that a positive, magnitude one shock will hit nominal interest rates in period t . Immediately, we see that while in the US model, these impacts remain positive at all horizons, in the Euro area model, these impacts turn

²² The MOD files for the Smets and Wouters (2003) model were derived from the Macro Model Database (Wieland et al. 2012). The MOD files for the model were derived from files provided by Johannes Pfeifer here: <http://goo.gl/CP53x5>

negative after just a few periods, and remain so at least up to period 40. Therefore, in the ZLB augmented Smets and Wouters (2003) model, there is not always a unique equilibrium. Furthermore, there are sequences of predicted future shocks (with positive density) for which the model without the ZLB would always feature positive interest rates, but for which the model with the ZLB could hit zero.

It remains for us to assess whether M is a $P_{(0)}$ -matrix or (strictly) semi-monotone for the Smets and Wouters (2007) model. Numerical calculations reveal that for $T < 9$, M is a P-matrix, and hence is strictly semi-monotone. However, with $T \geq 9$, M contains a 6×6 principal sub-matrix (with indices 1,2,4,6,7,9) with negative determinant, which is neither an S nor an S_0 -matrix. Thus, for $T \geq 9$, M is not a $P_{(0)}$ -matrix or (strictly) semi-monotone, and hence this model also has multiple equilibria, even when away from the bound. Given that the US has been at the ZLB for over eight years, that T ought to be greater than eight quarters seems uncontroversial. Hence, in both the Euro area and the US, we ought to take seriously the possibility that the existence of the ZLB produces non-uniqueness. Furthermore, it turns out that for neither model is M an S-matrix even with $T = 1000$, and thus for both models there are some $q \in \mathbb{R}^{1000}$ for which no solution exists. This is strongly suggestive of non-existence for some q even for arbitrarily large T . While placing a larger coefficient on inflation in the Taylor rule can make the Euro area picture more like the US one, with a strictly positive diagonal to the M matrix, even with incredibly large coefficients, M remained a non-P-matrix.

Alternatively, suppose we replace the monetary rule in both models by:

$$x_{r,t} = \max\{0, (1 - \rho_r)(x_{y,t} + x_{p,t}) + \rho_r x_{r,t-1}\}$$

where ρ_r is as in the respective original model, where the price level $x_{p,t}$ again evolves according to equation (16), and where $x_{y,t}$ is output relative to its linear trend. Then, for both models, for all T tested, M was a P-matrix, and for the Euro area model we have that $\zeta > 3 \times 10^{-7}$ and for the US model we have that $\zeta > 0.002$. As one would expect, this result is also robust to departures from equal, unit, coefficients. Thus, price level targeting again appears to be sufficient for determinacy in the presence of the ZLB.

4. Computation of solutions in the otherwise linear case

4.1. On the difficulty of the problem

We start with a note of caution. If no properties of the matrix M are known a priori, then Problem 4 is provably a computationally difficult problem; more formally, it may be shown to be “strongly-NP complete” (Chung 1989), and this remains true even if M is restricted to be a P_0 matrix (Kojima et al. 1991). This means that even if the inputs q and M have descriptions which are of a polynomial length in T , then if we could solve

Problem 4 in an amount of time that was polynomial in T , we could also solve in polynomial time any problem for which the solution could be verified in polynomial time. In the language of computer science, this would mean that “P=NP”, something almost all computer scientists believe to be false. The strength of computer scientists’ conviction that this is not true is best exemplified by the fact that were P=NP, all commonly used forms of cryptography (such as those used to secure internet banking) could potentially be defeated.

Since there is a bijection between solutions for Problem 4 and solutions for Problem 2, this means that while forming expectations in linear models without occasionally binding constraints is computationally easy (polynomial algorithms exist for it), in models with OBCs, forming expectations may be incredibly difficult. It also means that we should be sceptical of claims of computational efficiency from other algorithms for solving models with OBCs. A proof that such algorithms actually ran in time polynomial in the number of state variables of the model, on all models, would again function as a proof that “P=NP”, since we showed in Proposition 2 that there is a model corresponding to any $T \times T$ M matrix, featuring polynomial in T state variables.

A natural response to this is that in macroeconomics, we are only concerned with approximate solutions, whereas the previous computational complexity results were for exact solutions. In fact, allowing approximation error will not change these results. Note that a sufficiently accurate approximation to the solution would tell us when the constraint binds in the exact solution. However, the difficulty of the exact LCP comes from the fact that there are 2^T possible combinations of periods in which the constraint might bind, so no solution procedure can “quickly” tell us the periods constraint binds. Hence, the approximate problem cannot be easier than the original problem. More formally, we establish the following proposition in online appendix K.

Proposition 15 For any problem in the form of Problem 2, let $\mathcal{D}^* \subseteq \mathbb{R}^n$, and $p^*: \mathcal{D}^* \rightarrow \mathcal{D}^*$ be an exact policy function for Problem 2, by which we mean that:

1. For all $x \in \mathcal{D}^*$:

$$x_1 = \max\{0, I_{1,\cdot} \mu + A_{1,\cdot} (x - \mu) + (B_{1,\cdot} + I_{1,\cdot}) (p^*(x) - \mu) + C_{1,\cdot} (p^*(p^*(x)) - \mu)\},$$

$$(A_{-1,\cdot} + B_{-1,\cdot} + C_{-1,\cdot}) \mu = A_{-1,\cdot} x + B_{-1,\cdot} p^*(x) + C_{-1,\cdot} p^*(p^*(x)).$$

2. For all $x_0 \in \mathcal{D}^*$, if $x_t = p^*(x_{t-1})$ for all $t \in \mathbb{N}^+$, then $x_t \rightarrow \mu$ as $t \rightarrow \infty$.

3. For all $x_0 \notin \mathcal{D}^*$, there is no $(x_t)_{t=1}^\infty \subseteq \mathbb{R}^n$ solving the given instance of Problem 2.

Suppose that for all $\kappa, \epsilon > 0$, and for any problem in the form of Problem 2, we can calculate in time polynomial in n a set $\mathcal{D}_{\kappa,\epsilon} \subseteq \mathbb{R}^n$ with $\mathcal{D}^* \subseteq \mathcal{D}_{\kappa,\epsilon}$ and an approximate policy function $p_{\kappa,\epsilon}: \mathcal{D}_{\kappa,\epsilon} \rightarrow \mathbb{R}^n$, where membership of $\mathcal{D}_{\kappa,\epsilon}$ may be tested in time polynomial in n , where $p_{\kappa,\epsilon}$ may be evaluated in time polynomial in n ,

and where for all $x \in \mathcal{D}^*$ with $\|x - \mu\|_\infty < \kappa$, $\|p^*(x) - p_{\kappa,\epsilon}(x)\|_\infty < \epsilon$, then $P=NP$, i.e. all problems verifiable in polynomial time may be solved in polynomial time.

Providing one believes (along with almost all of the computer science profession) that $P \neq NP$, this provides a reductio ad absurdum of our assumption that there was a general procedure capable of providing the policy function p_ϵ in polynomial time. Thus, for example, global methods will never escape the curse of dimensionality in general models with OBCs, even using methods explicitly designed to do this such as that of Judd, Maliar, and Maliar (2012). In fact, even proving the finiteness of algorithms for solving these problems is non-trivial (see e.g. Csizmadia and Illés 2006), and, for example, there is no reason to believe that the iterations in Guerrieri and Iacoviello (2015) will converge in finite time on all models.

Admittedly, for some special classes further discussed in online appendix L, it has been shown that the problem is solvable in polynomial time in T . However, as discussed further in that appendix, it turns out that checking whether M is in one of the relevant special cases is itself not possible in polynomial time, so this is of little use. The results there imply that were there a general algorithm running in time polynomial in the number of state variables for testing if a particular model, always had a unique solution or always had a unique solution when away from the bound, then that algorithm would also serve as a proof that $P=NP$, since Proposition 2 implies that a model could be constructed producing any given $T \times T$ M matrix, with polynomial in T state variables.

4.2. The mixed integer linear programming representation

Given that there is no reason to believe that there is a polynomial time algorithm to solve the LCPs we encounter, it is important that we choose an algorithm, which, although it may not complete in polynomial time in the worst case, is nonetheless as computationally efficient as possible, particularly on average. One way to do this is to reduce the problem of finding an LCP to the solution of a problem for which highly efficient algorithms and computational libraries are available. One such problem is mixed integer linear programming (MILP), for which algorithms are included in most major optimisation suites (e.g. CPLEX, Gurobi, XPress MP, MOSEK, etc.). Conveniently, we can reduce the LCP problem to the MILP one in a way that gives not only a solution when one exists, but also a definite answer on whether or not there is a solution. This is an improvement over more naïve approaches, such as those of Holden (2010), Holden and Paetz (2012) or Guerrieri and Iacoviello (2015), for which a failure of convergence may just mean that the optimiser got stuck at some local minimum.

To motivate the MILP representation, suppose that y solves the LCP (q, M) . Then $y \geq 0$, $0 \leq q + My$ and if $y_j > 0$ then $(q + My)_j = 0$. Now let $\tilde{\omega} > 0$ be an arbitrary

constant, let $\alpha := \min\{\|y\|_\infty^{-1}, \tilde{\omega}\|q + My\|_\infty^{-1}\} > 0$, (where $\|\cdot\|_\infty$ is the usual sup norm), let $\hat{y} := \alpha y$, let $1_{T \times 1}$ be a $T \times 1$ vector of ones, and let $z \in \{0,1\}^T$ be such that for all $j \in \{1, \dots, T\}$, $z_j = 1$ if and only if $y_j > 0$ (i.e. z is an indicator for being away from the bound). Then $0 \leq \hat{y} \leq 1_{T \times 1}$ and $0 \leq \alpha q + M\hat{y} \leq \alpha\|q + My\|_\infty \leq \tilde{\omega}1_{T \times 1}$. Now, if $z_j = 0$ for some $j \in \{1, \dots, T\}$, then $y_j = 0$. Hence, in fact, $0 \leq \hat{y} \leq z$. Likewise, if $z_j = 1$, then $y_j > 0$, so since y solves the LCP, $0 = \alpha(q + My)_j = (\alpha q + M\hat{y})_j$. Hence, similarly, $0 \leq \alpha q + M\hat{y} \leq \tilde{\omega}(1_{T \times 1} - z)$. Moreover, for any $\tilde{\alpha} > \alpha$, we claim that there is no $z \in \{0,1\}^T$ such that $0 \leq \tilde{\alpha}y \leq z$ and $0 \leq \tilde{\alpha}q + M(\tilde{\alpha}y) \leq \tilde{\omega}(1_{T \times 1} - z)$. To see this, suppose for a contradiction that there were. Then $\tilde{\alpha}\|y\|_\infty \leq 1$, so $\alpha < \tilde{\alpha} \leq \|y\|_\infty^{-1}$. Hence, $\alpha = \tilde{\omega}\|q + My\|_\infty^{-1}$. But, by assumption $\tilde{\alpha}\|q + My\|_\infty \leq \tilde{\omega}$, hence $\tilde{\omega}\|q + My\|_\infty^{-1} = \alpha < \tilde{\alpha} \leq \tilde{\omega}\|q + My\|_\infty^{-1}$, which gives the required contradiction. Therefore, α, \hat{y}, z are feasible for the following MILP problem (though they may not necessarily be the solution):

Problem 6 Suppose $\tilde{\omega} > 0$, $q \in \mathbb{R}^T$ and $M \in \mathbb{R}^{T \times T}$ are given. Find $\alpha \in \mathbb{R}$, $\hat{y} \in \mathbb{R}^T$, $z \in \{0,1\}^T$ to maximise α subject to the following constraints: $\alpha \geq 0$, $0 \leq \hat{y} \leq z$, $0 \leq \alpha q + M\hat{y} \leq \tilde{\omega}(1_{T \times 1} - z)$. We call this the **mixed integer linear programming (MILP) representation of the LCP** (q, M) .

A version of this representation with $\tilde{\omega} = 1$ was first given by Pardalos and Rosen (1988), and its properties in that special case were proven by Rosen (1990).

We now establish that solutions of the MILP representation are solutions of the LCP. Suppose that α, \hat{y}, z solve Problem 6. If $\alpha = 0$, then there is no $\alpha > 0$ such that $0 \leq \hat{y} \leq z$, $0 \leq \alpha q + M\hat{y} \leq \tilde{\omega}(1_{T \times 1} - z)$. Now, we showed above that if the LCP (q, M) had a solution, then there would be an $\alpha > 0$, \hat{y} and z which were feasible for Problem 6, hence, this $\alpha > 0$ provides a lower bound on the solution to Problem 6. Thus, if $\alpha = 0$, the LCP cannot have a solution. Alternatively, suppose $\alpha > 0$. Then if for some $j \in \{1, \dots, T\}$, $z_j = 1$, then $0 = (\alpha q + M\hat{y})_j$, and if for some $j \in \{1, \dots, T\}$, $z_j = 0$, then $\hat{y}_j = 0$. Thus, $\hat{y} \circ (\alpha q + M\hat{y}) = 0$. Finally, define $y := \frac{\hat{y}}{\alpha} \geq 0$, hence $\hat{y} = \alpha y$, $0 \leq q + My$ and $y \circ (q + My) = 0$, i.e. y solves the LCP (q, M) . This establishes the following result:

Proposition 16 Suppose $\tilde{\omega} > 0$, $q \in \mathbb{R}^T$ and $M \in \mathbb{R}^{T \times T}$ are given. If y solves Problem 4, then the solution to Problem 6 has $\alpha \geq \min\{\|y\|_\infty^{-1}, \tilde{\omega}\|q + My\|_\infty^{-1}\}$. If the solution to Problem 4 is unique, then this last inequality holds with equality, and $\hat{y} = \alpha y$, $z = \begin{cases} 1 & \text{if } y_j > 0 \\ 0 & \text{if } y_j = 0 \end{cases}$ in the solution to Problem 6. Conversely, if α, \hat{y}, z solve Problem 6, then if $\alpha = 0$, Problem 4 has no solution, and if $\alpha > 0$, then $y := \frac{\hat{y}}{\alpha}$ solves Problem 4.

This result establishes that we can use the MILP representation both to find out if the LCP problem has a solution, and to find a solution when one exists. Furthermore, by

varying $\tilde{\omega}$ we can determine which solution is returned, when there are multiple. In the limit as $\tilde{\omega} \rightarrow 0$, the MILP solver will return the solution which minimises $\|q + My\|_\infty$, and in the limit as $\tilde{\omega} \rightarrow \infty$, the MILP solver will return the solution to the LCP which minimises $\|y\|_\infty$. The former objective would ensure that the returned solution does not generate large fluctuations in the path of the constrained variable, and the latter would reduce the fluctuations in other variables. Intermediate values of $\tilde{\omega}$ result in a solution being returned that features balanced concern for these two extremes. In practice, we suggest choosing $\tilde{\omega} = \omega \|q\|_\infty$ where $\omega > 0$ is another constant, to ensure that the solution returned scales appropriately with q . We suggest $\omega = 1000$, which places heavy weight on minimising $\|y\|_\infty$, without inducing numerical instabilities. It would also be possible to make ω stochastic in order to capture sunspot solutions to the model.

We can further constrain the solution returned in the presence of multiplicity by solving Problem 6 first with $T = 0$ (i.e. testing if $q \geq 0$), then with $T = 1$, and so on. Doing this ensures that the time to finally escape the bound is minimised.

5. Algorithms for general non-linear models, without perfect foresight

Up to now, we have solely been concerned with the perfect foresight solution of models which were linear apart from the occasionally binding constraint. In this section, we will apply these insights to the solution of general non-linear models, allowing for future uncertainty, i.e. we attempt to solve the following general problem:

Problem 7 Suppose that $x_0 \in \mathbb{R}^n$ is given and that $f: \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^c \times \mathbb{R}^m \rightarrow \mathbb{R}^n$, $g, h: \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^c \times \mathbb{R}^m \rightarrow \mathbb{R}^c$ are given continuously $d \in \mathbb{N}^+$ times differentiable functions. Find $x_t \in \mathbb{R}^n$ and $r_t \in \mathbb{R}^c$ for $t \in \mathbb{N}^+$ such that for all $t \in \mathbb{N}^+$:

$$0 = \mathbb{E}_t f(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t),$$

$$r_t = \mathbb{E}_t \max\{h(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t), g(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t)\}$$

where $\varepsilon_t \sim \text{NIID}(0, \Sigma)$, where the max operator acts elementwise on vectors, and where the information set is such that for all $t \in \mathbb{N}^+$, $\mathbb{E}_{t-1} \varepsilon_t = 0$ and $\mathbb{E}_t \varepsilon_t = \varepsilon_t$.

We construct our algorithm in two steps. First, we present an algorithm which accounts for the effects of uncertainty that would be present even without the bound, but which treats hitting the bound as a probability zero event. Then, we extend the aforementioned algorithm to capture the risk of hitting the bound in the future. The algorithms are implemented in the author's open source "DynareOBC" toolkit,²³ which extends Dynare (Adjemian et al. 2011) with the ability to deal with OBCs.

²³ Available from <http://github.org/tholden/dynareOBC>.

5.1. Dealing with non-linearity other than the bounds

Setup and assumptions Given a non-linear, but $d \in \mathbb{N}$ times continuously differentiable model, a common practice in macroeconomics is to take a perturbation approximation to the model around its deterministic steady-state. Given that high order perturbation approximations are often unstable, the use of a “pruned” approximation (Kim et al. 2008) is usually advisable. We will proceed along similar lines, taking a perturbation approximation to the model ignoring the bound, and then imposing the bound on the approximated model. The advantage for our purposes of the pruned approximation is that the result is linear in an augmented state space, which will assist us transferring results from the linear case to the non-linear one.

We start by making a further assumption that is necessary for us to be able to construct a perturbation approximation to the model without the bound.

Assumption 4 In the setup of Problem 7, there exists $\mu_x \in \mathbb{R}^n$ and $\mu_r \in \mathbb{R}^c$ such that:

$$0 = f(\mu_x, \mu_x, \mu_x, \mu_r, 0),$$

$$\mu_r = \max\{h(\mu_x, \mu_x, \mu_x, \mu_r, 0), g(\mu_x, \mu_x, \mu_x, \mu_r, 0)\},$$

and such that for all $a \in \{1, \dots, c\}$, $(h(\mu_x, \mu_x, \mu_x, \mu_r, 0))_a \neq (g(\mu_x, \mu_x, \mu_x, \mu_r, 0))_a$.

This is necessary because if any of the constraints just bind in steady-state, then the equation defining the corresponding element of μ_r is not differentiable at $(\mu_x, \mu_x, \mu_x, \mu_r, 0)$, preventing us from taking a perturbation approximation.

Henceforth, we suppose without loss of generality that $h(\mu_x, \mu_x, \mu_x, \mu_r, 0) \ll g(\mu_x, \mu_x, \mu_x, \mu_r, 0)$. We claim that we may further assume without loss of generality that $h(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t) \equiv 0$. First, note we can rewrite the equation defining r_t as:
 $r_t = h(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t) + \mathbb{E}_t \max\{0, g(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t) - h(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t)\}$.
Then if we define $\hat{x}_t := \begin{bmatrix} x_t \\ r_t \end{bmatrix}$ for all $t \in \mathbb{N}$, and:

$$\hat{g}(\hat{x}_{t-1}, \hat{x}_t, \hat{x}_{t+1}, \hat{r}_t, \varepsilon_t) := g(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t) - h(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t),$$

$$\hat{f}(\hat{x}_{t-1}, \hat{x}_t, \hat{x}_{t+1}, \hat{r}_t, \varepsilon_t) := f(x_{t-1}, x_t, x_{t+1}, g_1(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t) + \hat{r}_t, \varepsilon_t),$$

then for all $t \in \mathbb{N}^+$:

$$0 = \mathbb{E}_t \hat{f}(\hat{x}_{t-1}, \hat{x}_t, \hat{x}_{t+1}, \hat{r}_t, \varepsilon_t), \quad \hat{r}_t = \mathbb{E}_t \max\{0, \hat{g}(\hat{x}_{t-1}, \hat{x}_t, \hat{x}_{t+1}, \hat{r}_t, \varepsilon_t)\},$$

which is again in the form of Problem 7. Thus, without loss of generality, we can indeed assume that $h(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t) \equiv 0$, meaning that we work with the system:

$$0 = \mathbb{E}_t f(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t), \quad r_t = \mathbb{E}_t \max\{0, g(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t)\},$$

where $g(\mu_x, \mu_x, \mu_x, \mu_r, 0) \gg 0$.

First order approximations Now, suppose that we believe that a first order approximation gives adequate accuracy away from the bound. This system is locally $d \geq 1$ times differentiable in a neighbourhood of $(\mu_x, \mu_x, \mu_x, \mu_r, 0)$, so we can certainly

take a first order approximation around this point. Doing this gives the following approximation for the equation for r_t :

$$r_t = \mu_r + g_1(x_{t-1} - \mu_x) + g_2(x_t - \mu_x) + g_3\mathbb{E}_t(x_{t+1} - \mu_x) + g_4(r_t - \mu_r) + g_5\varepsilon_t,$$

where g_1, \dots, g_5 are the matrices of partial derivatives of g with respect to its first to fifth arguments, respectively. This approximation obviously completely ignores the bound. Thus, we propose to increase its accuracy by imposing the bound on the linearized equations, i.e. by instead working with the equation:

$$r_t = \max\{0, \mu_r + g_1(x_{t-1} - \mu_x) + g_2(x_t - \mu_x) + g_3\mathbb{E}_t(x_{t+1} - \mu_x) + g_4(r_t - \mu_r) + g_5\varepsilon_t\}.$$

This gives a system of equations in nearly the same form as that for which we developed a solution algorithm in section 4, the only difference being the presence of expectations operators and uncertainty. In our base approach, we deal with these following the extended path algorithm of Fair and Taylor (1983). I.e., if we are currently in period t of a simulation, we assume that the agents in the model believe that for all $s > t$, $\varepsilon_s = 0$. Thus, in each period of a simulation run, we merely have to solve a perfect foresight problem of the form of Problem 2, using the methods of section 4. We then advance one period, draw new shocks, and repeat the process.

We can also use a slightly modified form of the representation of Problem 3 to track the endogenous “news” that is coming from the bound, following Holden and Paetz (2012). In particular, we are effectively replacing the bounded equations with equations of the form $r_{a,t} = \mathbb{E}_t(g(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t))_a + I_{1,\cdot}y_t^{(a)}$, for all $a \in \{1, \dots, c\}$, where, for all $a \in \{1, \dots, c\}$:

$$y_{T,t}^{(a)} = \eta_{T,t}^{(a)}, \quad \forall i \in \{1, \dots, T-1\}, \quad y_{i,t}^{(a)} = y_{i+1,t-1}^{(a)} + \eta_{i,t}^{(a)},$$

implying that $y_{1,t}^{(a)} = \sum_{i=1}^T \eta_{i,t-i}^{(a)}$, where $\eta_{i,t}^{(a)}$ contains the news about the likelihood of the a^{th} bound binding in period $t+i$. Whereas the y s found by the LCP solver will always be positive, the implied $\eta_{i,t}^{(a)}$ need not necessarily be positive, as shocks may hit today which result in the economy moving away from the bound.

Higher order approximations Applying the methods of section 4 to models solved with a higher order approximation is slightly more difficult than in the linear case.

Recall that in linear models, we started by introducing news shocks to the bounded equation(s) and stacking the impulse responses to these news shocks into the M matrix. In order to use these impulse responses to tell us about the path of the bounded variable, we exploited the fact that the impulse response to a linear combination of shocks is the same linear combination of the individual impulse responses. It was this linearity that gave the $q + My$ representation of the path of the bounded variable.

Now, consider what would happen in a pruned or non-pruned second order approximation following a similar linear combination of shocks. Under such an

approximation, it is no longer true in general that the impulse response to a linear combination of shocks is a linear combination of the impulse responses, since the second order approximation captures interactions between the shocks. However, if the partial derivative of f and g with respect to each of the shocks being combined is zero, then the shocks only have second or higher order effects, hence, any interaction between them would be a fourth order effect or higher, and so would not be captured contemporaneously by the second order approximation. The period after the shocks hit, though, linearity would again be broken if a non-pruned second order approximation had been taken, since the slope of the response of the states to their lags vary with the states' levels. This is not true under a second order pruned perturbation approximation though, since under such an approximation, the solution takes the form:

$$x_t = \mu_x + x^{(0)} + x_t^{(1)} + x_t^{(2)}, \quad x_t^{(1)} = \alpha x_{t-1}^{(1)} + \beta_0 \varepsilon_t, \\ x_t^{(2)} = \alpha x_{t-1}^{(2)} + \frac{1}{2} \beta_{22} (x_{t-1}^{(1)} \otimes x_{t-1}^{(1)}) + \beta_{20} (x_{t-1}^{(1)} \otimes \varepsilon_t) + \frac{1}{2} \beta_{00} (\varepsilon_t \otimes \varepsilon_t),$$

where $x^{(0)} \in \mathbb{R}^n$ is a constant, $x_t^{(1)}$ is the first order component of the approximation, $x_t^{(2)}$ is the second order component of the approximation, and where β_0 has zero columns corresponding to each shock with respect to which the partial derivatives of f and g is zero (Kim et al. 2008). Thus, $x_t^{(1)}$ does not respond to any shocks for which the partial derivatives of f and g are zero, and hence $x_t^{(2)}$ and x_t are linear in such shocks.

In light of this discussion, in order to preserve the $q + My$ representation, we just need to define M as stacking the impulse responses of the bounded equation(s) to news shocks which hit the bounded equation(s) raised to the power of two, rather than in levels. This generalises to higher order pruned perturbation approximations as one would expect. Hence, in a d^{th} order pruned perturbation approximation, we replace the bounded equations with equations of the form $r_{a,t} = \mathbb{E}_t(g(x_{t-1}, x_t, x_{t+1}, r_t, \varepsilon_t))_a + I_{1,y_t^{(a)}}$, for all $a \in \{1, \dots, c\}$, where, now, for all $a \in \{1, \dots, c\}$:

$$y_{T,t}^{(a)} = \kappa (\eta_{T,t}^{(a)})^d, \quad \forall i \in \{1, \dots, T-1\}, \quad y_{i,t}^{(a)} = y_{i+1,t-1}^{(a)} + \kappa (\eta_{i,t}^{(a)})^d,$$

where κ is a very small constant in order to ensure that the presence of the additional shocks does not have an unwanted risk effect elsewhere in the model. We obviously scale the impulse responses which make up M by $\frac{1}{\kappa}$ to correct for this. In practice, we do not actually need to augment the model we approximate (in, e.g. Dynare) with all of these additional equations. This is because in the limit as $\kappa \rightarrow 0$, the generated impulse responses tend to the impulse responses used to construct the M matrix at first order.²⁴

Hence, in order to impose the bound in non-linear models solved by perturbation, we can proceed much as we did at first order. At each time step, we first evaluate the

²⁴ As the instant response must be the same, and the subsequent response is given by $x_t - x^{(0)} = \alpha(x_{t-1} - x^{(0)})$ in both cases.

expected path of the bounded variable(s), and stack the results in q . Thanks to the augmented state space representation of pruned perturbation solutions, this is possible without any Monte-Carlo simulation, as we show in online appendix M. Then, we use the M matrix derived from the first order approximation, and calculated as in section 2.2 in the LCP (q, M) . Finally, we use the solution to this LCP to calculate the required offsets to each variable this period, again based on the first order approximation to the model. Since we are not actually augmenting the model's state space, even for high degree approximations to the model, imposing the bound will not slow down simulation much more than it does at order one.

5.2. Integrating over future uncertainty

The downside to the approach discussed in the previous section is that the news shocks that hit the bounded equation(s) will not be conditionally mean zero, i.e. $\mathbb{E}_{t-1} \eta_{i,t}^{(a)} \neq 0$. For example, in a model with a zero lower bound on nominal interest rates, stochastic discount factors, and persistence in these discount factors, the higher is the state of the discount factor, the higher the chances of hitting the bound today, thus $\mathbb{E}_{t-1} \eta_{i,t}^{(a)}$ will be decreasing in this state. This failure of rationality with respect to expectations at the bound stems from the fact that we are still treating the bound in a pseudo-perfect foresight manner. In each period, agents act as if they believed that no future news shocks would ever hit the bounded equation. Due to the strict convexity of the $x \mapsto \max\{0, x\}$ mapping and Jensen's inequality, this manifests itself as a systematic downward bias in expectations of r_t .

To rectify this bias, we need to integrate over future uncertainty to calculate the expectation of the cumulated news shocks (the elements of y). We do this following the stochastic extended path approach of Adjemian and Juillard (2013). However, in our context this will be much easier than in the general fully non-linear context of Adjemian and Juillard (2013). In particular, in the basic algorithm of that paper, to integrate over S periods of future uncertainty, in a model with m shocks, they have to solve the perfect foresight model p^{mS} times, for some constant $p \geq 2$. While they are able to reduce this somewhat through the removal of low weighted quadrature nodes to produce a sparse tree of shocks, the resulting distribution of paths will significantly under-estimate the true variance of the model, and they still have to solve the perfect foresight model many more times when the number of shocks is high.

In our context, we will be able to do much better. In particular, we will be able to attain comparable accuracy with the evaluation of only a polynomial in S number of solutions of the perfect foresight problem, regardless of the number of shocks in the model. Furthermore, since solving an LCP is much easier than solving a general fully

non-linear perfect foresight problem, each of these solutions is orders of magnitude faster for us. The key to our invariance to the number of shocks in the model is the fact that in the absence of any bounds, we are able to write down a closed form expression for the conditional covariance of the bounded variables, thanks to the properties of pruned perturbation solutions. To be slightly more specific, suppose that $w_{t,s}$ is the value the bounded variables would take at s if the constraints did not apply from period t onwards. Then, we are able to calculate $\text{cov}_t(w_{t,t+i}, w_{t,t+j})$, for $t, i, j \in \mathbb{N}$, without any numerical integration. The derivation of this is contained in online appendix N.

For clarity of presentation, let us assume that there is a single bounded variable. As ever, the generalisation to multiple bounded variables will be straightforward. Now, ideally we would like to integrate over infinitely many periods of future uncertainty, but clearly this is not practical in reality. Instead, Adjemian and Juillard (2013) advocate integrating over $S \in \mathbb{N}^+$ periods of future uncertainty, and then ignoring uncertainty from period $S + 1$ onwards. By introducing a “discontinuity” in time of this sort, we would risk getting spurious movement in the expected path of variables around S periods into the future. Indeed, this occurred in some early numerical experiments that took this approach. Instead then, we apply a smooth windowing function to the variance of shocks. In particular, if the “true” shock covariance matrix is Σ , then when considering uncertainty at horizon k , we instead use:

$$\widehat{\Sigma}_k := \frac{1}{2} \left(1 + \cos \left(\pi \frac{\min\{k-1, S\}}{S} \right) \right) \Sigma.$$

Even with a time varying covariance matrix, it is still straightforward to calculate:

$$\Omega_t := \text{var}_t[[w_{t,t+1} \ \cdots \ w_{t,t+S}]'],$$

following the calculations in online appendix N. Then for the purposes of integration, we make the approximation that:

$$[w_{t,t+1} \ \cdots \ w_{t,t+S}]' \sim \mathcal{N}(\mathbb{E}_t[w_{t,t+1} \ \cdots \ w_{t,t+S}]', \Omega_t).$$

Of course the cosine window is ad hoc, but so too is the step-function window used by Adjemian and Juillard (2013). The legitimacy of both come from the fact that as $S \rightarrow \infty$, the error in this approximation would go to zero were $[w_{t,t+1} \ \cdots \ w_{t,t+S}]'$ really normally distributed. One further argument in favour of our cosine window is that it is widely used in signal processing due to its low distortion in the frequency domain (see e.g. Harris 1978). In this literature, it is termed the Hann or Hanning window. We also note that for most DSGE models, the additional error coming from the normal approximation will be minimal, since it is exact at first order, and higher order approximations are usually dominated by their first order terms.

Given this normal approximation, integration is then relatively straightforward. We first take the Schur decomposition of Ω_t , giving $\Omega_t = UDU'$, where U is an orthogonal

matrix and $D \geq 0$ is diagonal, with the elements sorted in decreasing order. To reduce the integration dimension without overly affecting accuracy, we set any elements of D which are less than some small multiple (e.g. 1%) of the maximum element of D to zero, as these components are unlikely to have a big impact. Indeed, when S is very large, it may be advisable to set all but the \hat{S} largest elements of D to zero, which means the cost of integration will scale in \hat{S} not S . After these steps we have that $D = \begin{bmatrix} D_{11} & 0 \\ 0 & 0 \end{bmatrix}$, where $\dim D_{11} = \hat{S} \times \hat{S}$ for some $\hat{S} \leq S$. Conformably partitioning U as $U = [U_{\cdot 1} \quad U'_{\cdot 1}]$, and defining $\Lambda := U_{\cdot 1} \sqrt{D_{11}}$, we then have that $\Lambda \Lambda' \approx \Omega_t$. Then if $\zeta \sim N(0, I_{\hat{S}})$, then $(\mathbb{E}_t[w_{t,t+1} \quad \dots \quad w_{t,t+S}'] + \Lambda \zeta)$ has approximately the same distribution as $[w_{t,t+1} \quad \dots \quad w_{t,t+S}]'$. We have thus transformed the problem of integrating over the distribution of $[w_{t,t+1} \quad \dots \quad w_{t,t+S}]'$ to that of integrating over the \hat{S} independent standard normals making up ζ . We do this using either quasi-Monte Carlo methods, the sparse nested Gaussian cubature rules of Genz and Keister (1996), or the equal weight degree 3 monomial cubature rule with $2\hat{S} + 1$ nodes²⁵. Of course, ideally we would like to break the domain of integration into pieces on which the integrand was differentiable, but this is not computationally practical for even moderately large \hat{S} .²⁶ We discuss the merits of these rules in online appendix O.

Whichever approach to integration is taken, we end with an approximation to the expected value of the “y” vector of cumulated news shocks needed to impose the bound. This y will imply a set of news shocks that hit today, just as it does when we ignore future uncertainty. We can thus proceed with the simulation exactly as we do in the case without integrating over future uncertainty.

5.3. Performance of our algorithm

The accuracy of the numerical algorithm presented in this paper is almost an immediate consequence of results from the prior literature. In particular, Guerrieri and Iacoviello (2015) showed the surprising accuracy of a perfect foresight solution to an otherwise linear approximation to a stochastic non-linear model with occasionally binding constraints. When a first order approximation is taken to the underlying model, and there is a unique solution, our method will produce exactly the same answers as that of Guerrieri and Iacoviello (2015).²⁷ Relative to their method, our method improves

²⁵ While there is a degree 3 monomial cubature rule with only $2\hat{S}$ nodes, including the 0 node generally increases accuracy at no cost, since we are evaluating the point anyway. As well as the 0 node, we evaluate at $\pm \frac{1}{2} \sqrt{2 + 4\hat{S}}$ with respect to each coordinate, which is easily shown to give a degree 3 rule. The use of monomial rules has been promoted by e.g. Judd and Skrainka (2011).

²⁶ For example, with $\hat{S} = 1$, there are at least as many discontinuities as there are non-zero elements in Λ . While we could get the full set of discontinuities at arbitrary dimension using a parametric linear complementarity problem solver such as that of Jones and Morrari (2006) which works providing that M is general positive semi-definite, this is computationally intractable for \hat{S} or T bigger than (about) ten, and integrating over all of these regions separately is computationally intractable even for much smaller \hat{S} .

²⁷ This point was noted by Guerrieri and Iacoviello (2015) with respect to the solution algorithm by Holden and Paetz (2012).

along four dimensions. Firstly, our method is guaranteed to produce a result in finite time. Secondly, it gives guarantees about which solution is selected when there are multiple. Thirdly, it also applies to higher order pruned perturbation solutions to the underlying model. Given the evidence that higher order pruned perturbations solutions are considerably more accurate than first order approximations (see e.g. Lan and Meyer-Gohde 2013b), this is likely to produce substantial accuracy gains, particularly as OBCs are usually located far from the steady-state. Finally, our method takes future uncertainty into account, meaning that it captures precautionary effects, unlike the Guerrieri and Iacoviello (2015) method. That integrating over future uncertainty usually increases accuracy in non-linear models has been established by Adjemian and Juillard (2013), so here too substantial accuracy gains are almost certain. Nonetheless, in the rest of this section, we give some further brief indications of the accuracy and speed of our approach, by applying it to two indicative non-linear models with OBCs.

A model with a closed form solution We first apply our method to a model for which we can calculate an exact closed form solution, giving very reliable accuracy measures. The model has the property that integrating over a single period of uncertainty is enough for accuracy, making it an ideal test of the performance of different cubature rules.

Suppose the representative household in an economy chooses consumption C_t and zero net supply bond holdings B_t to maximise $\mathbb{E}_t \sum_{k=0}^{\infty} \beta^k \frac{C_{t+k}^{1-\gamma} - 1}{1-\gamma}$ subject to the restriction that $A_t + R_{t-1}B_{t-1} = C_t + B_t$ for all $t \in \mathbb{Z}$, where A_t 's evolution is given by: $\log A_t = \log A_{t-1} + g_t$, where $g_t = \max\{0, (1-\rho)\bar{g} + \rho g_{t-1} + \sigma \varepsilon_t\}$ and $\varepsilon_t \sim \mathcal{N}(0,1)$. This specification may be thought of as capturing the fact that technologies cannot be un-invented. Market clearing implies $A_t = C_t$ and $B_t = 0$ for all $t \in \mathbb{Z}$, and from this, a closed form expression for R_t may be derived.²⁸ Using this, we define simulation errors as the gap between the true value of $\log R_t$ and the simulated value.²⁹ In Table 1, we report errors along simulated paths of length 1000, after discarding an initial 100 periods of burn-in, where for simulation, we used the following parameters: $\beta := 0.99$, $\gamma := 5$, $\bar{g} := 0.05$, $\rho := 0.95$, $\sigma := 0.07$. All cubature runs involve integrating over a single period of future uncertainty. We also report errors from the model with the bound removed, for comparison.

As may be seen from Table 1, our algorithm is generally very fast, as imposing the bound only doubles the running time without cubature. While this is a very simple model, the advantage of our algorithm is that running times are almost independent of

²⁸ In particular, $R_t = \left[\frac{\beta}{2} \left[1 - \operatorname{erf} \left(\frac{\sqrt{\beta}}{2\sigma} \mu_t \right) + \left(1 + \operatorname{erf} \left(\frac{\sqrt{\beta}}{2\sigma} (\mu_t - \gamma \sigma^2) \right) \right) \exp \left(\frac{\sigma^2 \gamma^2}{2} - \gamma \mu_t \right) \right] \right]^{-1}$, where $\mu_t = (1-\rho)\bar{g} + \rho g_t$.

²⁹ We also recorded errors in g_t , but these were essentially zero for all simulation runs.

the complexity of the model, so similar running times can be expected with even much larger models. These results also show that our algorithm is quite accurate, providing cubature is used. Without cubature, accuracy is below the accuracy of the first order approximation to the model without a bound. Since without cubature, at first order, our method will give identical answers to that of Guerrieri and Iacoviello (2015), this suggests that neither their algorithm nor ours without cubature can deliver comparable accuracy to that delivered by linearization in models without bounds. However, with cubature and a second order approximation, we can deliver errors that are lower in the model with the bound to those in the first order approximation to the model without the bound. As might be expected, quasi-Monte Carlo with many points is the most accurate integration method, but it does come at a significant time cost.

<i>Bound in</i>					<i>Root Mean</i>	<i>Mean Abs</i>	
<i>Model</i>	<i>Order</i>	<i>Cubature</i>	<i>Seconds</i> ³⁰	<i>Mean Abs</i>	<i>Squared</i>	<i>Max Abs</i>	<i>Error at</i>
				<i>Error</i>	<i>Error</i>	<i>Error</i>	<i>Bound</i> ³¹
No	1	N/A	66	6.13E-04	6.13E-04	6.13E-04	
No	2	N/A	62	1.52E-17	2.36E-17	1.67E-16	
No	3	N/A	53	1.99E-17	2.72E-17	1.11E-16	
Yes	1	No	141	3.67E-03	6.05E-03	1.31E-02	1.31E-02
Yes	2	No	139	3.76E-03	6.39E-03	1.37E-02	1.37E-02
Yes	3	No	140	3.76E-03	6.39E-03	1.37E-02	1.37E-02
Yes	1	Monomial, Degree 3	274	7.32E-04	8.45E-04	1.88E-03	7.40E-04
Yes	2	Monomial, Degree 3	1537	4.18E-04	6.73E-04	1.97E-03	1.28E-04
Yes	3	Monomial, Degree 3	1397	4.18E-04	6.73E-04	1.97E-03	1.28E-04
Yes	2	Sparse, Degree 3	1794	9.65E-04	1.67E-03	3.85E-03	3.85E-03
Yes	2	Sparse, Degree 5	1840	9.65E-04	1.67E-03	3.85E-03	3.85E-03
Yes	2	Sparse, Degree 7	2009	5.25E-04	9.30E-04	2.17E-03	2.17E-03
Yes	2	QMC, 15 Points	1965	9.12E-04	1.27E-03	2.17E-03	2.17E-03
Yes	2	QMC, 31 Points	2214	5.98E-04	8.17E-04	1.39E-03	1.39E-03
Yes	2	QMC, 63 Points	3184	4.04E-04	5.49E-04	9.55E-04	9.55E-04
Yes	2	QMC, 1023 Points	5197	1.57E-04	2.30E-04	4.45E-04	4.45E-04

Table 1: Accuracy in the model of bounded productivity growth

A model for which log-linearization gives an exact solution in the absence of bounds We next apply our procedure to a model for which log-linearization gives the exact solution when bounds are removed, but which features quite different behaviour when bounds are included, due to precautionary effects. This allows us to isolate the inaccuracy coming from the presence of an occasionally binding constraint.

³⁰ All timings are “wall” time, and include time spent starting the parallel pool, time spent compiling code (although written in MATLAB, DynareOBC generates and compiles C code for key routines), and time spent calculating accuracy. Code was run on one of the following (very similar) twenty core machines: 2x E5-2670 v2 2.5GHz, 64GB RAM; 2x E5-2660 v3 2.6GHz, 128GB RAM. Use of machines with network attached storage means that there may be some additional variance in these timings.

³¹ Errors conditional on the bounded variable being less than 0.0001. The numbers for this column would be identical had we used root mean squared errors or maximum absolute errors, conditional on being at the bound.

The model is a real business cycle model with a 100% depreciation rate (following Brock and Mirman (1972)), endogenous labour supply, and a lower bound on capital. The bound on capital may be thought of as an extreme case of a capital adjustment cost, with output falling to zero if the new level of capital is below the bound. To be more specific, the social planner chooses consumption, C_t , labour effort, L_t , and capital, K_t , to maximise $\mathbb{E}_t \sum_{k=0}^{\infty} \beta^k \left[\log C_{t+k} - \frac{L_{t+k}^{1+\nu}}{1+\nu} \right]$, subject to the capital constraint $K_t \geq \theta K_{t-1}$, and to the budget constraint $C_t + K_t = Y_t = A_t K_{t-1}^\alpha L_t^{1-\alpha}$, where productivity, A_t , evolves according to $A_t = A_{t-1}^\rho \exp \varepsilon_t$, where $\varepsilon_t \sim \mathcal{N}(0, \sigma^2)$. The first order conditions of the model imply that $(1 - \alpha) \frac{Y_t}{C_t} = L_t^{1+\nu}$ and that $\frac{1}{C_t} - \lambda_t = \alpha \beta \mathbb{E}_t \frac{Y_{t+1}}{C_{t+1} K_t} - \beta \theta \mathbb{E}_t \lambda_{t+1}$, where λ_t is the KKT-multiplier on the borrowing constraint. To convert the KKT conditions into a form amenable to simulation by our algorithm, note that the bound implies that $\frac{1}{C_t} = \frac{1}{Y_t - K_t} \geq \frac{1}{Y_t - \theta K_{t-1}}$, and the positivity of λ_t implies that $\frac{1}{C_t} \geq \alpha \beta \mathbb{E}_t \frac{Y_{t+1}}{C_{t+1} K_t} - \beta \theta \mathbb{E}_t \lambda_{t+1}$. Furthermore, by the KKT conditions, at least one of these constraints always binds. Hence, $\frac{1}{C_t} = \max \left\{ \frac{1}{Y_t - \theta K_{t-1}}, \alpha \beta \mathbb{E}_t \frac{Y_{t+1}}{C_{t+1} K_t} - \beta \theta \mathbb{E}_t \lambda_{t+1} \right\}$. In the following, we set $\alpha = 0.3$, $\beta = 0.99$, $\nu = 2$, $\theta = 0.99$, $\rho = 0.95$ and $\sigma = 0.01$.

In order to have a comparator for accuracy tests, we first solve the model globally to a high degree of accuracy, using value function iteration on a fine grid. Full details of the global solution procedure are given in appendix P, where we also plot the value and policy functions for the problem. Note, that as $K_{t-1} \rightarrow \infty$, in order to avoid violating the constraint, the solution must feature $L_t \rightarrow \infty$. Thus, whereas in the original model without bound, the value function is monotonic increasing in capital, in the model with bound, the value function is decreasing in capital for large enough capital levels. This significantly changes behaviour when away from the bound, as the planner will increase consumption now to avoid having too much capital in future.

To assess the extent to which integrating over future uncertainty captures these effects, we compare the value of consumption implied by the global solution procedure to that implied by our solution procedure with different values for S (the number of periods of future uncertainty considered). In particular, we construct a sparse grid by drawing 1023 Quasi-Monte Carlo (Sobol 1967) points from the stationary joint distribution of $[\log K_{t-1}, \log A_t]$ in the model without bounds,³² then for each grid point, our error measure is the difference between the value of $\log C_t$ implied by our algorithm at order 2,³³ and that implied by the global procedure. The mean absolute values of these errors over the grid are shown in Figure 5 for the fastest and slowest integration rules used for

³² This is very similar to the stationary distribution of the model with the bound, but is much easier to sample from.

³³ While the model without bounds may be exactly simulated via log-linearization, the bound introduces additional non-log-linearities which means there are substantial gains from higher order approximations.

the previous model, with $S = 0, \dots, 20$. It may be seen that integrating over future uncertainty produces reasonable accuracy gains in this model, with accuracy initially increasing as more periods of future uncertainty are considered. Accuracy eventually plateaus out as other sources of inaccuracy come to dominate, such as the limited order of perturbation approximation, and the imperfections of the integration rules. Surprisingly, for this model, the faster, monomial rule actually produces more accurate results at these horizons. This is perhaps due to the fact that the monomial rule places more weight on the tails of the distribution, better capturing the area in which the model is at the bound. These results suggest that this fast monomial rule may often have adequate performance in practice, removing the need to use slower integration rules, which may be prohibitively costly in large models, such as those used for policy.

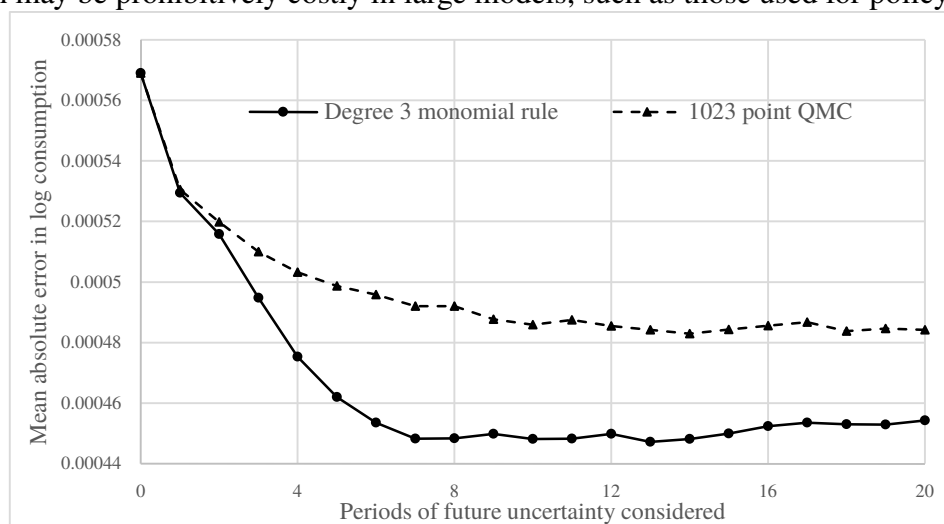


Figure 5: Effect on accuracy of increasing the number of periods of uncertainty considered

6. Conclusion

This paper provides the first theoretical results on existence and uniqueness for otherwise linear models with occasionally binding constraints. It also presents the first solution algorithm for such models that is guaranteed to return an answer in finite time. As such, it may be thought of as doing for models with OBCs what Blanchard and Kahn (1980) did for linear models.

We provided necessary and sufficient conditions for the existence of a unique equilibrium, as well as such conditions for uniqueness when away from the bound. In our application to New Keynesian models, we showed that these conditions were violated in entirely standard models, rather than being an artefact of strange policy rules, as one might have inferred from the results of Brendon, Paustian, and Yates (2015). In the presence of multiplicity, there is the potential for additional endogenous volatility from sunspots, so the welfare benefits of avoiding multiplicity may be substantial.

Additionally, as we saw in Figure 2, the additional equilibria may feature huge drops in output, giving further welfare reasons for their avoidance. The possibility of self-fulfilling jumps and returns from the ZLB also gives an alternative rationale for the neo-Fisherian view that argues that raising interest rates may raise inflation at the ZLB.³⁴

Luckily, our results suggest that a determinate equilibrium may be produced in standard New Keynesian models if the central bank switches to targeting the price level, rather than the inflation rate. This provides an additional argument for price level targeting in the presence of a zero lower bound to those made by Basu and Bundick (2015) and Coibion, Gorodnichenko, and Wieland (2012). Indeed, it is possible that Coibion, Gorodnichenko, and Wieland's results on the welfare benefits of price level targeting were actually driven by having inadvertently selected one of the worse equilibria under inflation targeting, since they use a solution algorithm for the otherwise linear model which gives no guarantees on the returned equilibrium.

In addition, we provided conditions for existence of any solution converging to the “good” steady-state at all, and showed that under inflation targeting, standard New Keynesian models again failed to satisfy these conditions in some states of the world. Whereas the literature started by Benhabib, Schmitt-Grohé, and Uribe (2001a; 2001b) showed that the existence of a “bad” steady-state may imply additional volatility if agents long-run beliefs are not pinned down by the inflation target, here we showed that in some states of the world, under inflation targeting there is no way for the economy to converge to the “good” steady-state. This in turn implies that agents cannot place prior certainty on converging to the “good” steady-state, thus rationalising the beliefs required to get the kind of global multiplicity at the zero lower bound that these and other authors have focussed on. Once again though, we showed that price level targeting is sufficient to restore existence and determinacy.

We also showed that price level targeting ensures that computing expectations is a computationally “easy” problem (i.e. polynomial time), whereas under inflation targeting, it is as difficult as the computational problems used to secure private communications, increasing the chance that agents would make mistakes in forming expectations. Given that there are often substantial welfare costs to failures in rationality, this provides yet further arguments in favour of price level targeting.

This paper has also presented the first algorithm for perfect foresight solutions of otherwise linear models with occasionally binding constraints that always completes in finite time, and which is able to detect when the model has no solution. The paper finished by exploiting this underlying solver to produce a robust, accurate and scalable

³⁴ Theoretical and empirical evidence for this view is presented in Cochrane (2015).

simulation algorithm for general nonlinear models with occasionally binding constraints that was consistent with rational expectations. Code implementing all of the algorithms discussed here is contained in the author’s “DynareOBC” toolkit which augments the abilities of Dynare (Adjemian et al. 2011) with the ability to solve models with OBCs.³⁵

7. References

- Adjemian, Stéphane, Houtan Bastani, Frédéric Karamé, Michel Juillard, Junior Maih, Ferhat Mihoubi, George Perendia, Johannes Pfeifer, Marco Ratto, and Sébastien Villemot. 2011. *Dynare: Reference Manual Version 4*. CEPREMAP.
- Adjemian, Stéphane, and Michel Juillard. 2013. ‘Stochastic Extended Path Approach’.
- Aliprantis, C.D., and K. Border. 2013. *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Studies in Economic Theory. Springer Berlin Heidelberg.
- Anderson, Gary. 2015. ‘A Solution Strategy for Occasionally Binding Constraints’. Unpublished.
- Aruoba, S. Borağan, Pablo Cuba-Borda, and Frank Schorfheide. 2014. *Macroeconomic Dynamics Near the ZLB: A Tale of Two Countries*. Penn Institute for Economic Research, Department of Economics, University of Pennsylvania.
- Basu, Susanto, and Brent Bundick. 2015. ‘Endogenous Volatility at the Zero Lower Bound: Implications for Stabilization Policy’. *National Bureau of Economic Research Working Paper Series* No. 21838.
- Beaudry, Paul, and Franck Portier. 2006. ‘Stock Prices, News, and Economic Fluctuations’. *American Economic Review* 96 (4): 1293–1307.
- Benhabib, Jess, Stephanie Schmitt-Grohé, and Martín Uribe. 2001a. ‘Monetary Policy and Multiple Equilibria’. *American Economic Review* 91 (1). *American Economic Review*: 167–186.
- . 2001b. ‘The Perils of Taylor Rules’. *Journal of Economic Theory* 96 (1-2). *Journal of Economic Theory*: 40–69.
- Benigno, Pierpaolo, and Michael Woodford. 2012. ‘Linear-Quadratic Approximation of Optimal Policy Problems’. *Journal of Economic Theory* 147 (1). *Journal of Economic Theory*: 1–42.
- Blanchard, Olivier Jean, and Charles M. Kahn. 1980. ‘The Solution of Linear Difference Models under Rational Expectations’. *Econometrica* 48 (5): 1305–1311.
- Bodenstein, Martin, Luca Guerrieri, and Christopher J. Gust. 2013. ‘Oil Shocks and the Zero Bound on Nominal Interest Rates’. *Journal of International Money and Finance* 32 (February): 941–967.
- Brendon, Charles, Matthias Paustian, and Tony Yates. 2015. ‘Self-Fulfilling Recessions at the Zero Lower Bound’. *Unpublished*.
- Brock, William A, and Leonard J Mirman. 1972. ‘Optimal Economic Growth and Uncertainty: The Discounted Case’. *Journal of Economic Theory* 4 (3) (June): 479–513.
- Chung, S.J. 1989. ‘NP-Completeness of the Linear Complementarity Problem’. *Journal of Optimization Theory and Applications* 60 (3) (March 1): 393–399.
- Cochrane, John H. 2015. ‘Do Higher Interest Rates Raise or Lower Inflation?’ Unpublished.
- Coibion, Olivier, Yuriy Gorodnichenko, and Johannes 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) (October 1): 1371–1406.
- Cools, Ronald. 2003. ‘An Encyclopaedia of Cubature Formulas’. *Oberwolfach Special Issue* 19 (3) (June): 445–453.
- Cottle, Richard W. 2009. ‘Linear Complementarity Problem’. In *Encyclopedia of Optimization*, edited by Christodoulos A. Floudas and Panos M. Pardalos, 1873–1878. Springer US.
- Cottle, Richard W., Jong-Shi Pang, and Richard E. Stone. 2009a. ‘3. Existence and Multiplicity’. In *The Linear Complementarity Problem*, 137–223. Classics in Applied Mathematics. Society for Industrial and Applied Mathematics.

³⁵ Further notes on DynareOBC’s internals are contained in online appendix Q. The toolkit may be freely downloaded from <http://github.org/tholden/dynareOBC>.

- . 2009b. ‘2. Background’. In *The Linear Complementarity Problem*, 43–136. Classics in Applied Mathematics. Society for Industrial and Applied Mathematics.
- Coxson, Gregory E. 1994. ‘The P-Matrix Problem Is Co-NP-Complete’. *Mathematical Programming* 64 (1-3) (March 1): 173–178.
- Csizmadia, Zsolt, and Tibor Illés. 2006. ‘New Criss-Cross Type Algorithms for Linear Complementarity Problems with Sufficient Matrices’. *Optimization Methods and Software* 21 (2) (April 1): 247–266.
- Davig, Troy, and Eric M. Leeper. 2007. ‘Generalizing the Taylor Principle’. *American Economic Review* 97 (3): 607–635.
- Fair, Ray C, and John B Taylor. 1983. ‘Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models’. *Econometrica* 51 (4) (July): 1169–85.
- Fang, Li. 1989. ‘On the Spectra of P- and P0-Matrices’. *Linear Algebra and Its Applications* 119 (July): 1–25.
- Farmer, Roger E. A., Daniel F. Waggoner, and Tao Zha. 2010. ‘Generalizing the Taylor Principle: Comment’. *American Economic Review* 100 (1): 608–17.
- Fernández-Villaverde, Jesús, Grey Gordon, Pablo Guerrón-Quintana, and Juan F. Rubio-Ramírez. 2012. ‘Nonlinear Adventures at the Zero Lower Bound’. NBER Working Papers.
- Fukuda, Komei. 2015. ‘Toward a Strongly Polynomial Algorithm for the Linear Complementarity Problems with Sufficient Matrices’. *Unpublished* (September 22).
- Gavin, William T., Benjamin D. Keen, Alexander W. Richter, and Nathaniel A. Throckmorton. 2015. ‘The Zero Lower Bound, the Dual Mandate, and Unconventional Dynamics’. *Journal of Economic Dynamics and Control* 55 (June): 14–38.
- Genz, Alan, and Bradley D. Keister. 1996. ‘Fully Symmetric Interpolatory Rules for Multiple Integrals over Infinite Regions with Gaussian Weight’. *Journal of Computational and Applied Mathematics* 71 (2): 299–309.
- Guerrieri, Luca, and Matteo Iacoviello. 2015. ‘OccBin: A Toolkit for Solving Dynamic Models with Occasionally Binding Constraints Easily’. *Journal of Monetary Economics* 70 (March): 22–38.
- Harris, F.J. 1978. ‘On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform’. *Proceedings of the IEEE* 66 (1) (January): 51–83.
- Hebden, James, Jesper Lindé, and Lars Svensson. 2011. ‘Optimal Monetary Policy in the Hybrid New Keynesian Model under the Zero Lower Bound’. Unpublished.
- Herceg, Martin, Michal Kvasnica, Colin N. Jones, and Manfred Morari. 2013. ‘Multi-Parametric Toolbox 3.0’. In *Control Conference (ECC), 2013 European*, 502–510. IEEE.
- Holden, Tom. 2010. *Products, Patents and Productivity Persistence: A DSGE Model of Endogenous Growth*. Economics Series Working Paper. University of Oxford, Department of Economics.
- Holden, Tom, and Michael Paetz. 2012. ‘Efficient Simulation of DSGE Models with Inequality Constraints’. School of Economics Discussion Papers.
- Illés, Tibor, Marianna Nagy, and Tamás Terlaky. 2010. ‘A Polynomial Path-Following Interior Point Algorithm for General Linear Complementarity Problems’. *Journal of Global Optimization* 47 (3) (July 1): 329–342.
- Jin, Hehui, and Kenneth L. Judd. 2002. ‘Perturbation Methods for General Dynamic Stochastic Models’.
- Jones, C.N., and M. Morrari. 2006. ‘Multiparametric Linear Complementarity Problems’. In *Decision and Control, 2006 45th IEEE Conference on*, 5687–5692.
- Judd, Kenneth L., Lilia Maliar, and Serguei Maliar. 2012. *Merging Simulation and Projection Approaches to Solve High-Dimensional Problems*. National Bureau of Economic Research, Inc.
- Judd, Kenneth L., and Ben Skrainka. 2011. *High Performance Quadrature Rules: How Numerical Integration Affects a Popular Model of Product Differentiation*. Centre for Microdata Methods and Practice, Institute for Fiscal Studies.
- Kamihigashi, Takashi. 2014. ‘Elementary Results on Solutions to the Bellman Equation of Dynamic Programming: Existence, Uniqueness, and Convergence’. *Economic Theory* 56 (2) (June 1): 251–273.
- Kamihigashi, Takashi, and Santanu Roy. 2003. *A Nonsmooth, Nonconvex Model of Optimal Growth*. Discussion Paper Series. Research Institute for Economics & Business Administration, Kobe University.
- Karp, Richard M. 1972. ‘Reducibility among Combinatorial Problems’. In *Complexity of Computer Computations*, edited by Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger, 85–103. The IBM Research Symposia Series. Springer US.
- Kim, Sunghyun Henry, Jinill Kim, Ernst Schaumburg, and Christopher A. Sims. 2008. ‘Calculating and Using Second Order Accurate Solutions of Discrete Time Dynamic Equilibrium Models’. *Journal of Economic Dynamics and Control* 32 (11): 3397–3414.

- Kojima, M., N. Megiddo, T. Noma, and A. Yoshise. 1991. *A Unified Approach to Interior Point Algorithms for Linear Complementarity Problems*. Springer.
- Kojima, M., Shinji Mizuno, and Akiko Yoshise. 1989. 'A Polynomial-Time Algorithm for a Class of Linear Complementarity Problems'. *Mathematical Programming* 44 (1-3) (May 1): 1–26.
- Lan, Hong, and Alexander Meyer-Gohde. 2013a. 'Solving DSGE Models with a Nonlinear Moving Average'. *Journal of Economic Dynamics and Control* 37 (12). *Journal of Economic Dynamics and Control*: 2643–2667.
- . 2013b. *Pruning in Perturbation DSGE Models - Guidance from Nonlinear Moving Average Approximations*. SFB 649 Discussion Papers. Sonderforschungsbereich 649, Humboldt University, Berlin, Germany.
- Levine, Paul, Joseph Pearlman, and Richard Pierse. 2008. 'Linear-Quadratic Approximation, External Habit and Targeting Rules'. *Journal of Economic Dynamics and Control* 32 (10). *Journal of Economic Dynamics and Control*: 3315–3349.
- Löfberg, J. 2004. 'YALMIP: A Toolbox for Modeling and Optimization in MATLAB'. In *Proceedings of the CACSD Conference*. Taipei, Taiwan.
- Magnus, Jan R., and H. Neudecker. 1979. 'The Commutation Matrix: Some Properties and Applications'. *Ann. Statist.* (2) (March): 381–394.
- Mertens, Karel RSM, and Morten O. Ravn. 2014. 'Fiscal Policy in an Expectations-Driven Liquidity Trap'. *The Review of Economic Studies*: rdu016.
- Pardalos, P. M., and J. B. Rosen. 1988. 'Global Optimization Approach to the Linear Complementarity Problem'. *SIAM J. Sci. Stat. Comput.* 9 (2): 341–354.
- Richter, Alexander W., and Nathaniel A. Throckmorton. 2014. 'The Zero Lower Bound: Frequency, Duration, and Numerical Convergence'. *The B.E. Journal of Macroeconomics* 15 (1): 157.
- Roos, C., T. Terlaky, and J. P. Vial. 2006. *Interior Point Methods for Linear Optimization*. Springer.
- Rosen, J. B. 1990. 'Minimum Norm Solution to the Linear Complementarity Problem'. *Functional Analysis, Optimization and Economics*: 208–216.
- Rotemberg, J. 1982. 'Monopolistic Price Adjustment and Aggregate Output'. *Review of Economic Studies* 49: 517–531.
- Samelson, Hans, Robert M. Thrall, and Oscar Wesler. 1958. 'A Partition Theorem for Euclidean N-Space'. *Proceedings of the American Mathematical Society* 9 (5): 805–807.
- Schmitt-Grohé, Stephanie, and Martín Uribe. 2012. *The Making Of A Great Contraction With A Liquidity Trap and A Jobless Recovery*. National Bureau of Economic Research, Inc.
- Sims, Christopher A. 2002. 'Solving Linear Rational Expectations Models'. *Computational Economics* 20 (1-2) (October 1): 1–20.
- Smets, Frank, and Rafael Wouters. 2003. 'An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area'. *Journal of the European Economic Association* 1 (5): 1123–1175.
- . 2007. 'Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach'. *American Economic Review* 97 (3) (June): 586–606.
- Sobol, Ilya M. 1967. 'On the Distribution of Points in a Cube and the Approximate Evaluation of Integrals'. *USSR Computational Mathematics and Mathematical Physics* (7): 86–112.
- Stokey, N., R. J. Lucas, and E. Prescott. 1989. *Recursive Methods in Economic Dynamics*. Harvard University Press.
- Trefethen, L.N., and M. Embree. 2005. *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*. Princeton University Press.
- Tsatsomeros, Michael J., and Lei Li. 2000. 'A Recursive Test for P-Matrices'. *BIT Numerical Mathematics* 40 (2) (June 1): 410–414.
- Tseng, Paul. 2000. 'Co-NP-Completeness of Some Matrix Classification Problems'. *Mathematical Programming* 88 (1) (June 1): 183–192.
- Väliäho, Hannu. 1986. 'Criteria for Copositive Matrices'. *Linear Algebra and Its Applications* 81 (September 1): 19–34.
- Wieland, Volker, Tobias Cwik, Gernot J. Müller, Sebastian Schmidt, and Maik Wolters. 2012. 'A New Comparative Approach to Macroeconomic Modeling and Policy Analysis'. *The Great Recession: Motivation for Re-Thinking Paradigms in Macroeconomic Modeling* 83 (3) (August): 523–541.
- Wright, Thomas G, and Lloyd N Trefethen. 2001. 'Large-Scale Computation of Pseudospectra Using ARPACK and Eigs'. *SIAM Journal on Scientific Computing* 23 (2): 591–605.
- Xu, Song. 1993. 'Notes on Sufficient Matrices'. *Linear Algebra and Its Applications* 191 (September 15): 1–13.

Online Appendices to: “Existence, uniqueness and computation of solutions to dynamic models with occasionally binding constraints.”

Tom D. Holden, School of Economics, University of Surrey

A. Construction of a model with arbitrary M matrix

Let $\mathcal{M} \in \mathbb{R}^{T \times T}$. Consider a model with the following equations:

$$\begin{aligned} a_t &= \max\{0, \ell_t\}, \\ a_t &= 1 + \sum_{j=1}^T \sum_{k=1}^T \mathcal{M}_{j,k} (c_{j-1,k-1,t} - c_{j,k,t}) + d_{0,t}, \\ c_{0,0,t} &= a_t - \ell_t, \\ c_{0,k,t} &= \mathbb{E}_t c_{0,k-1,t+1}, \quad \forall k \in \{1, \dots, T\}, \\ c_{j,k,t} &= c_{j-1,k,t-1}, \quad \forall j \in \{1, \dots, T\}, k \in \{0, \dots, T\}, \\ d_{k,t} &= d_{k+1,t-1}, \quad \forall k \in \{0, \dots, T-1\}, \\ d_{T,t} &= 0 \end{aligned}$$

with steady-state $a_{\cdot} = \ell_{\cdot} = 1$, $c_{j,k,\cdot} = 0$, $d_{k,\cdot} = 0$ for all $j, k \in \{0, \dots, T\}$. Defining:

$$x_t := [a_t \quad \ell_t \quad (\text{vec } c_{\cdot,\cdot,t})' \quad d'_{\cdot,t}]'$$

and dropping expectations, this model is then in the form of Problem 2.

Now consider the model's Problem 3 type equivalent, in which for $t \in \mathbb{N}^+$:

$$a_t = \begin{cases} \ell_t + y_{t,0} & \text{if } t \leq T \\ \ell_t & \text{if } t > T \end{cases}$$

where $y_{\cdot,\cdot}$ is defined as in Problem 3. Thus, if $c_{j,k,0} = 0$ and $d_{k,0} = 0$ for all $j, k \in \{0, \dots, T\}$, then for all $t \in \mathbb{N}^+$, $j, k \in \{0, \dots, T\}$:

$$c_{0,k,t} = \begin{cases} y_{t+k,0} & \text{if } t+k \leq T \\ 0 & \text{if } t+k > T \end{cases}$$

$$c_{j,k,t} = \begin{cases} c_{0,k,t-j} & \text{if } t-j > 0 \\ 0 & \text{if } t-j \leq 0 \end{cases} = \begin{cases} y_{t+k-j,0} & \text{if } t-j > 0, t+k-j \leq T \\ 0 & \text{otherwise} \end{cases}$$

Hence, for all $t \in \mathbb{N}^+$, $j, k \in \{1, \dots, T\}$:

$$c_{j-1,k-1,t} - c_{j,k,t} = \begin{cases} y_{t+k-j,0} & \text{if } t-j = 0, t+k-j \leq T \\ 0 & \text{otherwise} \end{cases} = \begin{cases} y_{k,0} & \text{if } t = j \\ 0 & \text{otherwise} \end{cases}$$

Therefore, for all $t \in \{1, \dots, T\}$:

$$a_t - 1 = \sum_{k=1}^T \mathcal{M}_{t,k} y_{k,0}.$$

Consequently, if $y_{k,0} = I_{\cdot,l}$ for some $l \in \{1, \dots, T\}$, then $a_t - 1 = \mathcal{M}_{t,l}$ (i.e. the relative impulse response to a news-shock at horizon l) is the l^{th} column of \mathcal{M} .

Finally, note that in the model's Problem 1 equivalent, if $c_{j,k,0} = 0$ for all $j, k \in \{0, \dots, T\}$, then for all $t \in \mathbb{N}^+$, $a_t = b_t = d_{0,t} = d_{t,0}$. Hence, if $d_{\cdot,0} = q$ for some $q \in \mathbb{R}^T$, then $q = q$ for this model.

B. Construction of a static model with no dynamic solution in some states

Consider the model:

$$a_t = \max\{0, b_t\}, \quad a_t = 1 - c_t, \quad c_t = a_t - b_t.$$

The model has steady-state $a = b = 1, c = 0$. Furthermore, in the model's Problem 3 type equivalent, in which for $t \in \mathbb{N}^+$:

$$a_t = \begin{cases} b_t + y_{t,0} & \text{if } t \leq T \\ b_t & \text{if } t > T \end{cases}$$

where $y_{\cdot,0}$ is defined as in Problem 3, we have that:

$$c_t = \begin{cases} y_{t,0} & \text{if } t \leq T \\ 0 & \text{if } t > T \end{cases}$$

so:

$$b_t = \begin{cases} 1 - 2y_{t,0} & \text{if } t \leq T \\ 1 & \text{if } t > T \end{cases}$$

implying:

$$a_t = \begin{cases} 1 - y_{t,0} & \text{if } t \leq T \\ 1 & \text{if } t > T \end{cases}$$

thus, $M = -I$ for this model.

C. Proof of sufficient conditions for feasibility with $T = \infty$

First, define $G := -C(B + CF)^{-1}$, and note that if L is the lag (right-shift) operator, the model from Problem 1 can be written as:

$$L^{-1}(ALL + BL + C)(x - \mu) = 0.$$

Furthermore, by the definitions of F and G :

$$(L - G)(B + CF)(I - FL) = ALL + BL + C,$$

so the stability of the model from Problem 1 is determined by the solutions for $z \in \mathbb{C}$ of the polynomial:

$$0 = \det(Az^2 + Bz + C) = \det(Iz - G) \det(B + CF) \det(I - Fz).$$

Now by Assumption 1, all of the roots of $\det(I - Fz)$ are strictly outside of the unit circle, and all of the other roots of $\det(Az^2 + Bz + C)$ are weakly inside the unit circle (else there would be indeterminacy), thus, all of the roots of $\det(Iz - G)$ are weakly inside the unit circle. Therefore, if we write $\rho_{\mathcal{M}}$ for the spectral radius of some matrix \mathcal{M} , then, by this discussion and Assumption 2, $\rho_G < 1$.

Next, let $s_t^*, x_t^* \in \mathbb{R}^{n \times \mathbb{N}^+}$ be such that for any $y \in \mathbb{R}^{\mathbb{N}^+}$, the k^{th} columns of $s_t^* y$ and $x_t^* y$ give the value of s_t and x_t following a magnitude 1 news shock at horizon k , i.e. when $x_0 = \mu$ and y_0 is the k^{th} row of $I_{\mathbb{N}^+ \times \mathbb{N}^+}$. Then:

$$\begin{aligned} s_t^* &= -(B + CF)^{-1} [I_{\cdot,1} I_{t,1:\infty} + G I_{\cdot,1} I_{t+1,1:\infty} + G^2 I_{\cdot,1} I_{t+2,1:\infty} + \dots] \\ &= -(B + CF)^{-1} \sum_{k=0}^{\infty} (GL)^k I_{\cdot,1} I_{t,1:\infty} = \\ &= -(B + CF)^{-1} (I - GL)^{-1} I_{\cdot,1} I_{t,1:\infty}, \end{aligned}$$

where the infinite sums are well defined as $\rho_G < 1$, and where $I_{t,1:\infty} \in \mathbb{R}^{1 \times \mathbb{N}^+}$ is a row vector with zeros everywhere except position t where there is a 1. Thus:

$$s_t^* = [0_{n \times (t-1)} \quad s_1^*] = L^{t-1} s_1^*.$$

Furthermore,

$$(x_t^* - \mu^*) = \sum_{j=1}^t F^{t-j} s_j^* = \sum_{j=1}^t F^{t-j} L^{j-1} s_1^*,$$

i.e.:

$$(x_t^* - \mu^*)_{\cdot,k} = \sum_{j=1}^t F^{t-j} s_{1^*,k+1-j}^* = - \sum_{j=1}^{\min\{t,k\}} F^{t-j} (B + CF)^{-1} G^{k-j} I_{\cdot,1},$$

and so the k^{th} offset diagonal of M ($k \in \mathbb{Z}$) is given by the first row of the k^{th} column of:

$$L^{-t} (x_t^* - \mu^*) = L^{-1} \sum_{j=1}^t (FL^{-1})^{t-j} s_1^* = L^{-1} \sum_{j=0}^{t-1} (FL^{-1})^j s_1^*,$$

where we abuse notation slightly by allowing L^{-1} to give a result with indices in \mathbb{Z} rather than \mathbb{N}^+ , with padding by zeros. Consequently, for all $k \in \mathbb{N}^+$, $M_{t,k} = O(t^n \rho_F^t)$, as $t \rightarrow \infty$, for all $t \in \mathbb{N}^+$, $M_{t,k} = O(t^n \rho_G^k)$, as $k \rightarrow \infty$, and for all $k \in \mathbb{Z}$, $M_{t,t+k} - \lim_{\tau \rightarrow \infty} M_{\tau,\tau+k} = O(t^{n-1} (\rho_F \rho_G)^t)$, as $t \rightarrow \infty$. Hence,

$$\sup_{y \in [0,1]^{\mathbb{N}^+}} \inf_{t \in \mathbb{N}^+} M_{t,1:\infty} y$$

exists and is well defined, and so:

$$\zeta = \sup_{\substack{y \in [0,1]^{\mathbb{N}^+} \\ \exists T \in \mathbb{N} \text{ s.t. } \forall t > T, y_t = 0}} \inf_{t \in \mathbb{N}^+} M_{t,1:\infty} y = \sup_{y \in [0,1]^{\mathbb{N}^+}} \inf_{t \in \mathbb{N}^+} M_{t,1:\infty} y,$$

since every point in $[0,1]^{\mathbb{N}^+}$ is a limit (under the supremum norm) of a sequence of points in the set:

$$\{y \in [0,1]^{\mathbb{N}^+} \mid \exists T \in \mathbb{N} \text{ s.t. } \forall t > T, y_t = 0\}.$$

Thus, we just need to provide conditions under which $\sup_{y \in [0,1]^{\mathbb{N}^+}} \inf_{t \in \mathbb{N}^+} M_{t,1:\infty} y > 0$.

To produce such conditions, we need constructive bounds on M , even if they have slightly worse convergence rates. For any matrix, $\mathcal{M} \in \mathbb{R}^{n \times n}$ with $\rho_{\mathcal{M}} < 1$, and any $\phi \in (\rho_{\mathcal{M}}, 1)$, let:

$$\mathcal{E}_{\mathcal{M},\phi} := \sup_{k \in \mathbb{N}} \|(\mathcal{M} \phi^{-1})^k\|_2.$$

Furthermore, for any matrix, $\mathcal{M} \in \mathbb{R}^{n \times n}$ with $\rho_{\mathcal{M}} < 1$, and any $\epsilon > 0$, let:

$$\rho_{\mathcal{M},\epsilon} := \max\{|z| \mid z \in \mathbb{C}, \sigma_{\min}(\mathcal{M} - zI) = \epsilon\},$$

where $\sigma_{\min}(\mathcal{M} - zI)$ is the minimum singular value of $\mathcal{M} - zI$, and let $\epsilon^*(\mathcal{M}) \in (0, \infty]$ solve:

$$\rho_{\mathcal{M},\epsilon} = 1.$$

(This has a solution in $(0, \infty]$ by continuity as $\rho_{\mathcal{M}} < 1$.) Then, by Theorem 16.2 of Trefethen and Embree (2005), for any $K \in \mathbb{N}$ and $k > K$:

$$\|(\mathcal{M}\phi^{-1})^k\|_2 \leq \|(\mathcal{M}\phi^{-1})^K\|_2 \|(\mathcal{M}\phi^{-1})^{k-K}\|_2 \leq \frac{\|(\mathcal{M}\phi^{-1})^K\|_2}{\epsilon^*(\mathcal{M}\phi^{-1})}.$$

Now, $\|(\mathcal{M}\phi^{-1})^K\|_2 \rightarrow 0$ as $K \rightarrow \infty$, hence, there exists some $K \in \mathbb{N}$ such that:

$$\sup_{k=0,\dots,K} \|(\mathcal{M}\phi^{-1})^k\|_2 \geq \frac{\|(\mathcal{M}\phi^{-1})^K\|_2}{\epsilon^*(\mathcal{M}\phi^{-1})} \geq \sup_{k>K} \|(\mathcal{M}\phi^{-1})^k\|_2,$$

meaning $\mathcal{C}_{\mathcal{M},\phi} = \sup_{k=0,\dots,K} \|(\mathcal{M}\phi^{-1})^k\|_2$. The quantity $\rho_{\mathcal{M},\epsilon}$ (and hence $\epsilon^*(\mathcal{M})$) may

be efficiently computed using the methods described by Wright and Trefethen (2001), and implemented in their EigTool toolkit³⁶. Thus, $\mathcal{C}_{\mathcal{M},\phi}$ may be calculated in finitely many operations by iterating over $K \in \mathbb{N}$ until a K is found which satisfies:

$$\sup_{k=0,\dots,K} \|(\mathcal{M}\phi^{-1})^k\|_2 \geq \frac{\|(\mathcal{M}\phi^{-1})^K\|_2}{\epsilon^*(\mathcal{M}\phi^{-1})}.$$

From the definition of $\mathcal{C}_{\mathcal{M},\phi}$, we have that for any $k \in \mathbb{N}$ and any $\phi \in (\rho_{\mathcal{M}}, 1)$:

$$\|\mathcal{M}^k\|_2 \leq \mathcal{C}_{\mathcal{M},\phi} \phi^k.$$

Now, fix $\phi_F \in (\rho_F, 1)$ and $\phi_G \in (\rho_G, 1)$,³⁷ and define:

$$\mathcal{D}_{\phi_F,\phi_G} := \mathcal{C}_{F,\phi_F} \mathcal{C}_{G,\phi_G} \|(B + CF)^{-1}\|_2,$$

then, for all $t, k \in \mathbb{N}^+$:

$$\begin{aligned} |M_{t,k}| &= |(x_t^* - \mu^*)_{1,k}| \leq \|(x_t^* - \mu^*)_{\cdot,k}\|_2 \leq \sum_{j=1}^{\min\{t,k\}} \|F^{t-j}\|_2 \|(B + CF)^{-1}\|_2 \|G^{k-j}\|_2 \\ &\leq \mathcal{D}_{\phi_F,\phi_G} \sum_{j=1}^{\min\{t,k\}} \phi_F^{t-j} \phi_G^{k-j} = \mathcal{D}_{\phi_F,\phi_G} \phi_F^t \phi_G^k \frac{(\phi_F \phi_G)^{-\min\{t,k\}} - 1}{1 - \phi_F \phi_G}. \end{aligned}$$

Additionally, for all $t \in \mathbb{N}^+$, $k \in \mathbb{Z}$:

$$\begin{aligned} |M_{t,t+k} - \lim_{\tau \rightarrow \infty} M_{\tau,\tau+k}| &= \left| (L^{-t}(x_t^* - \mu^*))_{1,k} - \left(\lim_{\tau \rightarrow \infty} L^{-\tau}(x_\tau^* - \mu^*) \right)_{1,k} \right| \\ &\leq \left\| \left(L^{-1} \sum_{j=0}^{t-1} (FL^{-1})^j s_1^* - L^{-1} \sum_{j=0}^{\infty} (FL^{-1})^j s_1^* \right)_{\cdot,k} \right\|_2 \\ &= \left\| \left(\sum_{j=\max\{t,-k\}}^{\infty} F^j s_{1,\cdot,j+k+1}^* \right)_{\cdot,0} \right\|_2 \end{aligned}$$

³⁶ This toolkit is available from <https://github.com/eigtool/eigtool>, and is included in dynareOBC.

³⁷ In practice, we try a grid of values, as it is problem dependent whether high ϕ_F and low $\mathcal{K}(\mathcal{M}\phi^{-1})$ is preferable to low ϕ_F and high $\mathcal{K}(\mathcal{M}\phi^{-1})$.

$$\begin{aligned}
&= \left\| \sum_{j=\max\{t,-k\}}^{\infty} F^j (B + CF)^{-1} G^{j+k} I_{.,1} \right\|_2 \\
&\leq \sum_{j=\max\{t,-k\}}^{\infty} \|F^j\|_2 \|(B + CF)^{-1}\|_2 \|G^{j+k}\|_2 \\
&\leq \mathcal{D}_{\phi_F, \phi_G} \sum_{j=\max\{t,-k\}}^{\infty} \phi_F^j \phi_G^{j+k} = \mathcal{D}_{\phi_F, \phi_G} \frac{\phi_F^{\max\{t,-k\}} \phi_G^{\max\{0,t+k\}}}{1 - \phi_F \phi_G},
\end{aligned}$$

so, for all $t, k \in \mathbb{N}^+$:

$$|M_{t,k} - \lim_{\tau \rightarrow \infty} M_{\tau, \tau+k-t}| \leq \mathcal{D}_{\phi_F, \phi_G} \frac{\phi_F^t \phi_G^k}{1 - \phi_F \phi_G}.$$

To evaluate $\lim_{\tau \rightarrow \infty} M_{\tau, \tau+k-t}$, note that this limit is the top element from the $(k-t)$ th column of:

$$\begin{aligned}
d &:= \lim_{\tau \rightarrow \infty} L^{-\tau} (x_{\tau}^* - \mu^*) = L^{-1} (I - FL^{-1})^{-1} s_1^* \\
&= -(I - FL^{-1})^{-1} (B + CF)^{-1} (I - GL)^{-1} I_{.,1} I_{0,-\infty:\infty},
\end{aligned}$$

where $I_{0,-\infty:\infty} \in \mathbb{R}^{1 \times \mathbb{Z}}$ is zero everywhere apart from index 0 where it equals 1. Hence, by the definitions of F and G :

$$AL^{-1}d + Bd + CLd = -I_{.,1} I_{0,-\infty:\infty}.$$

In other words, if we write d_k in place of $d_{.,k}$ for convenience, then, for all $k \in \mathbb{Z}$:

$$Ad_{k+1} + Bd_k + Cd_{k-1} = - \begin{cases} I_{.,1} & \text{if } k = 0 \\ 0 & \text{otherwise} \end{cases}$$

I.e. the homogeneous part of the difference equation for d_{-t} is identical to that of $x_t - \mu$. The time reversal here is intuitive since we are indexing diagonals such that indices increase as we move up and to the right in M , but time is increasing as we move down in M .

It turns out that exploiting the possibility of reversing time is the key to easy evaluating d_k . First, note that for $k < 0$, it must be the case that $d_k = Fd_{k+1}$, since the shock has already ‘‘occurred’’ (remember, that we are going forwards in ‘‘time’’ when we reduce k). Now consider the model in which we are going forwards time when we increase k , i.e. the model with:

$$L(AL^{-1}L^{-1} + BL^{-1} + C)d = 0,$$

subject to the terminal condition that $d_k \rightarrow 0$ as $k \rightarrow \infty$, which must hold as we have already proved that the first row of M converges to zero. Now, let $z \in \mathbb{C}$, $z \neq 0$ be a solution to:

$$0 = \det(Az^2 + Bz + C),$$

and define $\tilde{z} = z^{-1}$, so:

$$\begin{aligned}
0 &= \det(A + B\tilde{z} + C\tilde{z}^2) = z^{-2} \det(Az^2 + Bz + C) \\
&= \det(I - G\tilde{z}) \det(B + CF) \det(I\tilde{z} - F).
\end{aligned}$$

By Assumption 1, all of the roots of $\det(I\tilde{z} - F)$ are inside the unit circle, thus they cannot contribute to the dynamics of the time reversed process, else the terminal

condition would be violated. Thus, the time reversed model has a unique solution satisfying the terminal condition with a transition matrix with the same eigenvalues as G . Consequently, this solution can be calculated via standard methods for solving linear DSGE models, and it will be given by $d_k = Hd_{k-1}$, for all $k > 0$, where $H = -(B + AH)^{-1}C$, and $\phi_H = \phi_G < 1$, by Assumption 2.

It just remains to determine the value of d_0 . By the previous results, this must satisfy:

$$-I_{.,1} = Ad_1 + Bd_0 + Cd_{-1} = (AH + B + CF)d_0.$$

Hence:

$$d_0 = -(AH + B + CF)^{-1}I_{.,1}.$$

This gives a readily computed solution for the limits of the diagonals of M . Lastly, note that:

$$|d_{-t,1}| \leq \|d_{-t}\|_2 = \|F^t d_0\|_2 \leq \|F^t\|_2 \|d_0\|_2 \leq \mathcal{C}_{F, \phi_F} \phi_F^t \|d_0\|_2,$$

and:

$$|d_{t,1}| \leq \|d_t\|_2 = \|H^t d_0\|_2 \leq \|H^t\|_2 \|d_0\|_2 \leq \mathcal{C}_{H, \phi_H} \phi_H^t \|d_0\|_2.$$

We will use these results in producing our bounds on ς .

First, fix $T \in \mathbb{N}^+$, and define a new matrix $\underline{M}^{(T)} \in \mathbb{R}^{\mathbb{N}^+ \times \mathbb{N}^+}$ by $\underline{M}_{1:T,1:T}^{(T)} = M_{1:T,1:T}$, and for all $t, k \in \mathbb{N}^+$, with $\min\{t, k\} > T$, $\underline{M}_{t,k}^{(T)} = d_{k-t,1} - \mathcal{D}_{\phi_F, \phi_G} \frac{\phi_F^t \phi_G^k}{1 - \phi_F \phi_G}$, then:

$$\begin{aligned} \varsigma &\geq \max_{\substack{y \in [0,1]^T \\ y_\infty \in [0,1]}} \inf_{t \in \mathbb{N}^+} M_{t,1:\infty} \begin{bmatrix} y \\ y_\infty \mathbf{1}_{\infty \times 1} \end{bmatrix} \geq \max_{\substack{y \in [0,1]^T \\ y_\infty \in [0,1]}} \inf_{t \in \mathbb{N}^+} \underline{M}_{t,1:\infty}^{(T)} \begin{bmatrix} y \\ y_\infty \mathbf{1}_{\infty \times 1} \end{bmatrix} \\ &= \max_{\substack{y \in [0,1]^T \\ y_\infty \in [0,1]}} \min \left\{ \begin{aligned} &\min_{t=1, \dots, T} \left[M_{t,1:T} y + \sum_{k=T+1}^{\infty} \left(d_{k-t,1} - \mathcal{D}_{\phi_F, \phi_G} \frac{\phi_F^t \phi_G^k}{1 - \phi_F \phi_G} \right) y_\infty \right], \\ &\inf_{t \in \mathbb{N}^+, t > T} \left[\sum_{k=1}^T \left(d_{k-t,1} - \mathcal{D}_{\phi_F, \phi_G} \frac{\phi_F^t \phi_G^k}{1 - \phi_F \phi_G} \right) y_k + \sum_{k=T+1}^{\infty} \left(d_{k-t,1} - \mathcal{D}_{\phi_F, \phi_G} \frac{\phi_F^t \phi_G^k}{1 - \phi_F \phi_G} \right) y_\infty \right] \end{aligned} \right\} \\ &\geq \max_{\substack{y \in [0,1]^T \\ y_\infty \in [0,1]}} \min \left\{ \begin{aligned} &\min_{t=1, \dots, T} \left[M_{t,1:T} y + ((I - H)^{-1} d_{T+1-t})_1 y_\infty - \mathcal{D}_{\phi_F, \phi_G} \frac{\phi_F^t \phi_G^{T+1}}{(1 - \phi_F \phi_G)(1 - \phi_G)} y_\infty \right], \\ &\min_{t=T+1, \dots, 2T} \left[\sum_{k=1}^T \left(d_{-(t-k),1} - \mathcal{D}_{\phi_F, \phi_G} \frac{\phi_F^t \phi_G^k}{1 - \phi_F \phi_G} \right) y_k + ((I - F)^{-1} (d_{-1} - d_{-(t-T)}))_1 y_\infty \right. \\ &\quad \left. + ((I - H)^{-1} d_0)_1 y_\infty - \mathcal{D}_{\phi_F, \phi_G} \frac{\phi_F^t \phi_G^{T+1}}{(1 - \phi_F \phi_G)(1 - \phi_G)} y_\infty \right], \\ &\inf_{t \in \mathbb{N}^+, t > 2T} \left[\sum_{k=1}^T d_{-(t-k),1} y_k + ((I - F)^{-1} (d_{-1} - d_{-(t-T)}))_1 y_\infty \right. \\ &\quad \left. + ((I - H)^{-1} d_0)_1 y_\infty - \mathcal{D}_{\phi_F, \phi_G} \frac{\phi_F^{2T+1} \phi_G}{(1 - \phi_F \phi_G)(1 - \phi_G)} \right] \end{aligned} \right\}. \end{aligned}$$

Now, for $t \geq T$:

$$\begin{aligned} |((I - F)^{-1} d_{-(t-T)})_1| &\leq \|(I - F)^{-1} d_{-(t-T)}\|_2 \leq \|(I - F)^{-1}\|_2 \|d_{-(t-T)}\|_2 \\ &\leq \mathcal{C}_{F, \phi_F} \phi_F^{t-T} \|(I - F)^{-1}\|_2 \|d_0\|_2, \end{aligned}$$

so:

$$\begin{aligned}
& \sum_{k=1}^T d_{-(t-k),1} y_k - ((I-F)^{-1} d_{-(t-T)})_1 y_\infty \\
& \geq - \sum_{k=1}^T \mathcal{C}_{F,\phi_F} \phi_F^{t-k} \|d_0\|_2 - \mathcal{C}_{F,\phi_F} \phi_F^{t-T} \|(I-F)^{-1}\|_2 \|d_0\|_2 y_\infty \\
& = - \mathcal{C}_{F,\phi_F} \frac{\phi_F^t (\phi_F^{-T} - 1)}{1 - \phi_F} \|d_0\|_2 - \mathcal{C}_{F,\phi_F} \phi_F^{t-T} \|(I-F)^{-1}\|_2 \|d_0\|_2 y_\infty,
\end{aligned}$$

thus $\varsigma \geq \underline{\varsigma}$, where:

$$\underline{\varsigma}_T := \max_{\substack{y \in [0,1]^T \\ y_\infty \in [0,1]}} \min \left\{ \begin{array}{l} \min_{t=1,\dots,T} \left[M_{t,1:T} y + ((I-H)^{-1} d_{T+1-t})_1 y_\infty - \mathcal{D}_{\phi_F,\phi_G} \frac{\phi_F^t \phi_G^{T+1}}{(1-\phi_F\phi_G)(1-\phi_G)} y_\infty \right], \\ \min_{t=T+1,\dots,2T} \left[\sum_{k=1}^T \left(d_{-(t-k),1} - \mathcal{D}_{\phi_F,\phi_G} \frac{\phi_F^t \phi_G^k}{1-\phi_F\phi_G} \right) y_k + ((I-F)^{-1} (d_{-1} - d_{-(t-T)}))_1 y_\infty \right. \\ \quad \left. + ((I-H)^{-1} d_0)_1 y_\infty - \mathcal{D}_{\phi_F,\phi_G} \frac{\phi_F^t \phi_G^{T+1}}{(1-\phi_F\phi_G)(1-\phi_G)} y_\infty \right], \\ \left[- \mathcal{C}_{F,\phi_F} \frac{\phi_F^{2T+1} (\phi_F^{-T} - 1)}{1 - \phi_F} \|d_0\|_2 - \mathcal{C}_{F,\phi_F} \phi_F^{T+1} \|(I-F)^{-1}\|_2 \|d_0\|_2 y_\infty + ((I-F)^{-1} d_{-1})_1 y_\infty \right. \\ \quad \left. + ((I-H)^{-1} d_0)_1 y_\infty - \mathcal{D}_{\phi_F,\phi_G} \frac{\phi_F^{2T+1} \phi_G}{(1-\phi_F\phi_G)(1-\phi_G)} \right] \end{array} \right\}$$

It is worth noting that as $T \rightarrow \infty$, the final minimand in this expression tends to:

$$((I-F)^{-1} d_{-1})_1 y_\infty + ((I-H)^{-1} d_0)_1 y_\infty,$$

i.e. a positive multiple of the doubly infinite sum of $d_{1,k}$ over all $k \in \mathbb{Z}$. If this expression is negative, then our lower bound on ς will be negative as well, and hence uninformative.

To construct an upper bound on ς , fix $T \in \mathbb{N}^+$, and define a new matrix $\overline{M}^{(T)} \in \mathbb{R}^{\mathbb{N}^+ \times \mathbb{N}^+}$ by $\overline{M}_{1:T,1:T}^{(T)} = M_{1:T,1:T}$, and for all $t, k \in \mathbb{N}^+$, with $\min\{t, k\} > T$, $\overline{M}_{t,k}^{(T)} = |d_{k-t,1}| + \mathcal{D}_{\phi_F,\phi_G} \frac{\phi_F^t \phi_G^k}{1-\phi_F\phi_G}$. Then:

$$\begin{aligned}
\varsigma & = \sup_{y \in [0,1]^{\mathbb{N}^+}} \inf_{t \in \mathbb{N}^+} M_{t,1:\infty} y \leq \sup_{y \in [0,1]^{\mathbb{N}^+}} \inf_{t \in \mathbb{N}^+} \overline{M}_{t,1:\infty} y \leq \sup_{y \in [0,1]^{\mathbb{N}^+}} \min_{t=1,\dots,T} \overline{M}_{t,1:\infty} y \\
& \leq \max_{y \in [0,1]^T} \min_{t=1,\dots,T} \overline{M}_{t,1:\infty} \begin{bmatrix} y \\ 1_{\infty \times 1} \end{bmatrix} \\
& \leq \max_{y \in [0,1]^T} \min_{t=1,\dots,T} \left[M_{t,1:T} y + \sum_{k=T+1}^{\infty} |d_{k-t,1}| + \sum_{k=T+1}^{\infty} \mathcal{D}_{\phi_F,\phi_G} \frac{\phi_F^t \phi_G^k}{1-\phi_F\phi_G} \right] \\
& \leq \max_{y \in [0,1]^T} \min_{t=1,\dots,T} \left[M_{t,1:T} y + \sum_{k=T+1-t}^{\infty} |d_{k,1}| + \mathcal{D}_{\phi_F,\phi_G} \frac{\phi_F^t \phi_G^{T+1}}{1-\phi_F\phi_G} \sum_{k=0}^{\infty} \phi_G^k \right] \\
& \leq \max_{y \in [0,1]^T} \min_{t=1,\dots,T} \left[M_{t,1:T} y + \mathcal{C}_{H,\phi_H} \|d_0\|_2 \phi_H^{T+1-t} \sum_{k=0}^{\infty} \phi_H^k + \mathcal{D}_{\phi_F,\phi_G} \frac{\phi_F^t \phi_G^{T+1}}{(1-\phi_F\phi_G)(1-\phi_G)} \right] \\
& = \overline{\varsigma}_T := \max_{y \in [0,1]^T} \min_{t=1,\dots,T} \left[M_{t,1:T} y + \frac{\mathcal{C}_{H,\phi_H} \|d_0\|_2 \phi_H^{T+1-t}}{1-\phi_H} + \mathcal{D}_{\phi_F,\phi_G} \frac{\phi_F^t \phi_G^{T+1}}{(1-\phi_F\phi_G)(1-\phi_G)} \right].
\end{aligned}$$

D. Other properties of the solution set

First, let us give one further definition:

Definition 9 ((Non-)Degenerate matrix) A matrix $M \in \mathbb{R}^{T \times T}$ is called a **non-degenerate matrix** if the principal minors of M are all non-zero. M is called a **degenerate matrix** if it is not a non-degenerate matrix.

Then, conditions for having a finite or convex set of solutions are given in the following propositions.

Proposition 17 The LCP (q, M) has a finite (possibly zero) number of solutions for all $q \in \mathbb{R}^T$ if and only if M is non-degenerate. (Cottle, Pang, and Stone 2009a)

Proposition 18 The LCP (q, M) has a convex (possibly empty) set of solutions for all $q \in \mathbb{R}^T$ if and only if M is column sufficient. (Cottle, Pang, and Stone 2009a)

E. Generalisations to richer otherwise linear models

It is straightforward to generalise the results for Problem 2 to less restrictive otherwise linear models with occasionally binding constraints.

Firstly, if the constraint is on a variable other than $x_{1,t}$, or in another equation than the first, then it is immediately clear that all of the results must go through as before (just by relabelling and rearranging). Furthermore, if the constraint takes the form of $z_{1,t} = \max\{z_{2,t}, z_{3,t}\}$, where $z_{1,t}$, $z_{2,t}$ and $z_{3,t}$ are linear expressions in the contemporaneous values, lags and leads of x_t , then, assuming without loss of generality that $z_{3,t} > z_{2,t}$ in steady-state, we have that $z_{1,t} - z_{2,t} = \max\{0, z_{3,t} - z_{2,t}\}$. Hence, adding a new auxiliary variable $x_{n+1,t}$, with the associated equation $x_{n+1,t} = z_{1,t} - z_{2,t}$, and replacing the constrained equation with $x_{n+1,t} = \max\{0, z_{3,t} - z_{2,t}\}$, we have a new equation in the form covered by our original results. Moreover, if rather than a max we have a min, we just use the fact that if $z_{1,t} = \min\{z_{2,t}, z_{3,t}\}$, then $-z_{1,t} = \max\{-z_{2,t}, -z_{3,t}\}$, which is in the form covered by the generalisation just established.

We may also readily deal with multiple occasionally binding constraints, following the representation used in Holden and Paetz (2012). Suppose there are c constrained variables in the model. For $a \in \{1, \dots, c\}$, let $q^{(a)}$ be the path of the a^{th} constrained variable in the absence of all constraints. For $a, b \in \{1, \dots, c\}$, let $M^{(a,b)}$ be the matrix created by horizontally stacking the column vector relative impulse responses of the a^{th} constrained variable to magnitude 1 news shocks at horizon $0, \dots, T - 1$ to the equation

defining the b^{th} constrained variables. For example, if $c = 1$ so there is a single constraint, then we would have that $M^{(1,1)} = M$ as defined in equation (2). Finally, let:

$$q := \begin{bmatrix} q^{(1)} \\ \vdots \\ q^{(c)} \end{bmatrix}, \quad M := \begin{bmatrix} M^{(1,1)} & \dots & M^{(1,c)} \\ \vdots & \ddots & \vdots \\ M^{(c,1)} & \dots & M^{(c,c)} \end{bmatrix},$$

and let y be a solution to the LCP (q, M) . Then the vertically stacked paths of the constrained variables in a solution which satisfies these constraints is given by $q + My$, and again any solution satisfying the constraints corresponds to a solution to the LCP. Thus, in the multiple constraint case, all of our previous results go through (almost) immediately, with this redefined q vector and M matrix.

F. Proof of the sufficient conditions for the existence of a unique solution to the dynamic programming problem

Results when \tilde{X} is possibly non-compact, but $\tilde{\Gamma}(x)$ is compact valued and $x \in \tilde{\Gamma}(x)$ for all $x \in \tilde{X}$ We first note that for all $x, z \in \tilde{X}$:

$$\tilde{\mathcal{F}}(x, z) \leq u^{(0)} - \frac{1}{2}u^{(1)}\tilde{u}^{(2)-1}u^{(1)'},$$

thus our objective function is bounded above without additional assumptions. For a lower bound, we assume that for all $x \in \tilde{X}$, $x \in \tilde{\Gamma}(x)$, so holding the state fixed is always feasible. This is true in very many standard applications. Then, the value of setting $x_t = x_0$ for all $t \in \mathbb{N}^+$ provides a lower bound for our objective function.

More precisely, we define $\mathbb{V} := \{v | v: \tilde{X} \rightarrow [-\infty, \infty)\}$ and $\underline{v}, \bar{v} \in \mathbb{V}$ by:

$$\underline{v}(x) = \frac{1}{1-\beta} \tilde{\mathcal{F}}(x_0, x_0),$$

$$\bar{v}(x) = \frac{1}{1-\beta} \left[u^{(0)} - \frac{1}{2}u^{(1)}\tilde{u}^{(2)-1}u^{(1)'} \right],$$

for all $x \in \tilde{X}$.

Finally, define $\mathcal{B}: \mathbb{V} \rightarrow \mathbb{V}$ by:

$$\mathcal{B}(v)(x) = \sup_{z \in \tilde{\Gamma}(x)} \left[\tilde{\mathcal{F}}(x, z) + \beta v(z) \right] \quad (17)$$

for all $v \in \mathbb{V}$ and for all $x \in \tilde{X}$. Then $\mathcal{B}(\underline{v}) \geq \underline{v}$ and $\mathcal{B}(\bar{v}) \leq \bar{v}$. Furthermore, if some sequence $(x_t)_{t=1}^{\infty}$ satisfies the constraint that for all $t \in \mathbb{N}^+$, $x_t \in \tilde{\Gamma}(x_{t-1})$, and the objective in (8) is finite for that sequence, then it must be the case that $\|x_t\|_{\infty} t \beta^{\frac{t}{2}} \rightarrow 0$ as $t \rightarrow \infty$ (by the comparison test), so:

$$\liminf_{t \rightarrow \infty} \beta^t \underline{v}(x_t) = 0.$$

Additionally, for any sequence $(x_t)_{t=1}^{\infty}$:

$$\limsup_{t \rightarrow \infty} \beta^t \bar{v}(x_t) = 0.$$

Thus, our dynamic programming problem satisfies the assumptions of Theorem 2.1 of Kamihigashi (2014), and so \mathcal{B} has a unique fixed point in $[\underline{v}, \bar{v}]$ to which $\mathcal{B}^k(\underline{v})$

converges pointwise, monotonically, as $k \rightarrow \infty$, and which is equal to the function $v^*: \tilde{X} \rightarrow \mathbb{R}$ defined by:

$$v^*(x_0) = \sup\{\sum_{t=1}^{\infty} \beta^{t-1} \tilde{\mathcal{F}}(x_{t-1}, x_t) \mid \forall t \in \mathbb{N}^+, x_t \in \Gamma(x_{t-1})\}, \quad (18)$$

for all $x_0 \in \tilde{X}$.

Furthermore, if we define $\mathbb{W} := \{v \in V \mid v \text{ is continuous on } \tilde{X}, v \text{ is concave on } \tilde{X}\}$, then as $\tilde{u}^{(2)}$ is negative-definite, $\underline{v} \in \mathbb{W}$. Additionally, under the assumption that $\tilde{\Gamma}(x)$ is compact valued, if $v \in \mathbb{W}$, then $\mathcal{B}(v) \in \mathbb{W}$, by the theorem of the maximum,³⁸ and, furthermore, there is a unique policy function which attains the supremum in the definition of $\mathcal{B}(v)$. Moreover, $v^* = \lim_{k \rightarrow \infty} \mathcal{B}^k(\underline{v})$ is concave and lower semi-continuous on \tilde{X} .³⁹ We just need to prove that v^* is upper semi-continuous.⁴⁰ Suppose for a contradiction that it is not, so there exists $x^* \in \tilde{X}$ such that:

$$\limsup_{x \rightarrow x^*} v^*(x) > \lim_{k \rightarrow \infty} v^*(x^*).$$

Then, there exists $\delta > 0$ such that for all $\epsilon > 0$, there exists $x_0^{(\epsilon)} \in \tilde{X}$ with $\|x^* - x_0^{(\epsilon)}\|_{\infty} < \epsilon$ such that:

$$v^*(x_0^{(\epsilon)}) > \delta + v^*(x^*).$$

Now, by the definition of a supremum, for all $\epsilon > 0$, there exists $(x_t^{(\epsilon)})_{t=1}^{\infty}$ such that for all $t \in \mathbb{N}^+$, $x_t^{(\epsilon)} \in \Gamma(x_{t-1}^{(\epsilon)})$ and:

$$v^*(x_0^{(\epsilon)}) < \delta + \sum_{t=1}^{\infty} \beta^{t-1} \tilde{\mathcal{F}}(x_{t-1}^{(\epsilon)}, x_t^{(\epsilon)}).$$

Hence:

$$\sum_{t=1}^{\infty} \beta^{t-1} \tilde{\mathcal{F}}(x_{t-1}^{(\epsilon)}, x_t^{(\epsilon)}) > v^*(x_0^{(\epsilon)}) - \delta > v^*(x^*).$$

Now, let $\mathcal{S}_0 := \{x \in \tilde{X} \mid \|x^* - x\|_{\infty} \leq 1\}$, and for $t \in \mathbb{N}^+$, let $\mathcal{S}_t := \Gamma(\mathcal{S}_{t-1})$. Then, since we are assuming Γ is compact valued, for all $t \in \mathbb{N}$, \mathcal{S}_t is compact by the continuity of Γ . Furthermore, for all $t \in \mathbb{N}$ and $\epsilon \in (0, 1)$, $x_t^{(\epsilon)} \in \mathcal{S}_t$. Hence, $\prod_{t=0}^{\infty} \mathcal{S}_t$ is sequentially compact in the product topology. Thus, there exists a sequence $(\epsilon_k)_{k=1}^{\infty}$ with $\epsilon_k \rightarrow 0$ as $k \rightarrow \infty$ and such that $x_t^{(\epsilon_k)}$ converges for all $t \in \mathbb{N}$. Let $x_t := \lim_{k \rightarrow \infty} x_t^{(\epsilon_k)}$, and note that $x^* = x_0 \in \mathcal{S}_0 \subseteq \tilde{X}$, and that for all $t, k \in \mathbb{N}^+$, $x_t^{(\epsilon_k)} \in \Gamma(x_{t-1}^{(\epsilon_k)})$, so by the continuity of Γ , $x_t \in \Gamma(x_{t-1})$ for all $t \in \mathbb{N}^+$. Thus, by Fatou's Lemma:

$$v^*(x^*) \geq \sum_{t=1}^{\infty} \beta^{t-1} \tilde{\mathcal{F}}(x_{t-1}, x_t) \geq \limsup_{k \rightarrow \infty} \sum_{t=1}^{\infty} \beta^{t-1} \tilde{\mathcal{F}}(x_{t-1}^{(\epsilon_k)}, x_t^{(\epsilon_k)}) > v^*(x^*),$$

which gives the required contradiction. Thus v^* is continuous and concave, and there is a unique policy function which attains the supremum in the definition of $\mathcal{B}(v^*) = v^*$.

³⁸ See e.g. Theorem 3.6 and following of Stokey, Lucas, and Prescott (1989).

³⁹ See e.g. Lemma 2.41 of Aliprantis and Border (2013).

⁴⁰ In the following, we broadly follow the proof of Lemma 3.3 of Kamihigashi and Roy (2003).

Results when \tilde{X} is compact If \tilde{X} is compact, then Γ is compact valued. Furthermore, \tilde{X} is clearly convex, and Γ is continuous. Thus assumption 4.3 of Stokey, Lucas, and Prescott (1989) (henceforth: SLP) is satisfied. Since the continuous image of a compact set is compact, $\tilde{\mathcal{F}}$ is bounded above and below, so assumption 4.4 of SLP is satisfied as well. Furthermore, $\tilde{\mathcal{F}}$ is concave and Γ is convex, so assumptions 4.7 and 4.8 of SLP are satisfied too. Thus, by theorem 4.6 of SLP, with \mathcal{B} defined as in equation (17) and v^* defined as in equation (18), \mathcal{B} has a unique fixed point which is continuous and equal to v^* . Moreover, by theorem 4.8 of SLP, there is a unique policy function which attains the supremum in the definition of $\mathcal{B}(v^*) = v^*$.

G. Proof of the sufficiency of the KKT and limit conditions

Suppose that $(x_t)_{t=1}^\infty, (\lambda_t)_{t=1}^\infty$ satisfy the KKT conditions given in equations (10) and (11), and that $x_t \rightarrow \mu$ and $\lambda_t \rightarrow \bar{\lambda}$ as $t \rightarrow \infty$. Let $(z_t)_{t=0}^\infty$ satisfy $z_0 = x_0$ and $z_t \in \tilde{\Gamma}(z_{t-1})$ for all $t \in \mathbb{N}^+$. Then, by the KKT conditions and the concavity of:

$$(x_{t-1}, x_t) \mapsto \tilde{\mathcal{F}}(x_{t-1}, x_t) + \lambda'_t \left[\Psi^{(0)} + \Psi^{(1)} \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix} \right],$$

we have that for all $T \in \mathbb{N}^+$:⁴¹

$$\begin{aligned} & \sum_{t=1}^T \beta^{t-1} [\tilde{\mathcal{F}}(x_{t-1}, x_t) - \tilde{\mathcal{F}}(z_{t-1}, z_t)] \\ &= \sum_{t=1}^T \beta^{t-1} \left[\tilde{\mathcal{F}}(x_{t-1}, x_t) + \lambda'_t \left[\Psi^{(0)} + \Psi^{(1)} \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix} \right] - \tilde{\mathcal{F}}(z_{t-1}, z_t) \right] \\ &\geq \sum_{t=1}^T \beta^{t-1} \left[\tilde{\mathcal{F}}(x_{t-1}, x_t) + \lambda'_t \left[\Psi^{(0)} + \Psi^{(1)} \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix} \right] - \tilde{\mathcal{F}}(z_{t-1}, z_t) \right. \\ &\quad \left. - \lambda'_t \left[\Psi^{(0)} + \Psi^{(1)} \begin{bmatrix} z_{t-1} - \mu \\ z_t - \mu \end{bmatrix} \right] \right] \\ &\geq \sum_{t=1}^T \beta^{t-1} \left[\left[u_{\cdot,2}^{(1)} + \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix}' \tilde{u}_{\cdot,2}^{(2)} + \lambda'_t \Psi_{\cdot,2}^{(1)} \right] (x_t - z_t) \right. \\ &\quad \left. + \left[u_{\cdot,1}^{(1)} + \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix}' \tilde{u}_{\cdot,1}^{(2)} + \lambda'_t \Psi_{\cdot,1}^{(1)} \right] (x_{t-1} - z_{t-1}) \right] \\ &= \sum_{t=1}^T \beta^{t-1} \left[\left[u_{\cdot,2}^{(1)} + \begin{bmatrix} x_{t-1} - \mu \\ x_t - \mu \end{bmatrix}' \tilde{u}_{\cdot,2}^{(2)} + \lambda'_t \Psi_{\cdot,2}^{(1)} \right. \right. \\ &\quad \left. \left. + \beta \left[u_{\cdot,1}^{(1)} + \begin{bmatrix} x_t - \mu \\ x_{t+1} - \mu \end{bmatrix}' \tilde{u}_{\cdot,1}^{(2)} + \lambda'_{t+1} \Psi_{\cdot,1}^{(1)} \right] \right] (x_t - z_t) \right] \\ &\quad + \beta^T \left[u_{\cdot,1}^{(1)} + \begin{bmatrix} x_T - \mu \\ x_{T+1} - \mu \end{bmatrix}' \tilde{u}_{\cdot,1}^{(2)} + \lambda'_{T+1} \Psi_{\cdot,1}^{(1)} \right] (z_T - x_T) \\ &= \beta^T \left[u_{\cdot,1}^{(1)} + \begin{bmatrix} x_T - \mu \\ x_{T+1} - \mu \end{bmatrix}' \tilde{u}_{\cdot,1}^{(2)} + \lambda'_{T+1} \Psi_{\cdot,1}^{(1)} \right] (z_T - x_T). \end{aligned}$$

⁴¹ Here, we broadly follow the proof of Theorem 4.15 of Stokey, Lucas, and Prescott (1989).

Thus:

$$\begin{aligned} & \sum_{t=1}^{\infty} \beta^{t-1} [\tilde{\mathcal{F}}(x_{t-1}, x_t) - \tilde{\mathcal{F}}(z_{t-1}, z_t)] \\ & \geq \lim_{T \rightarrow \infty} \beta^T \left[u_{\cdot,1}^{(1)} + \begin{bmatrix} x_T - \mu \\ x_{T+1} - \mu \end{bmatrix}' \tilde{u}_{\cdot,1}^{(2)} + \lambda'_{T+1} \Psi_{\cdot,1}^{(1)} \right] (z_T - x_T) \\ & = \lim_{T \rightarrow \infty} \beta^T [u_{\cdot,1}^{(1)} + \bar{\lambda}' \Psi_{\cdot,1}^{(1)}] (z_T - \mu) = \lim_{T \rightarrow \infty} \beta^T [u_{\cdot,1}^{(1)} + \bar{\lambda}' \Psi_{\cdot,1}^{(1)}] z_T. \end{aligned}$$

Now, suppose $\lim_{T \rightarrow \infty} \beta^T z_T \neq 0$, then since $\tilde{u}^{(2)}$ is negative definite:

$$\sum_{t=1}^{\infty} \beta^{t-1} \tilde{\mathcal{F}}(z_{t-1}, z_t) = -\infty,$$

so $(z_t)_{t=0}^{\infty}$ cannot be optimal. Hence, regardless of the value of $\lim_{T \rightarrow \infty} \beta^T z_T$:

$$\sum_{t=1}^{\infty} \beta^{t-1} [\tilde{\mathcal{F}}(x_{t-1}, x_t) - \tilde{\mathcal{F}}(z_{t-1}, z_t)] \geq 0,$$

which implies that $(x_t)_{t=1}^{\infty}$ solves Problem 5.

H. Results from and for general dynamic programming problems

Here we consider non-linear dynamic programming problems with general objective functions. Consider then the following generalisation of Problem 5:

Problem 8 Suppose $\Gamma: \mathbb{R}^n \rightarrow \mathbb{P}(\mathbb{R}^n)$ is a given compact, convex valued continuous function. Define $X := \{x \in \mathbb{R}^n | \Gamma(x) \neq \emptyset\}$, and suppose without loss of generality that for all $x \in \mathbb{R}^n$, $\Gamma(x) \cap X = \Gamma(x)$. Further suppose that $\mathcal{F}: X \times X \rightarrow \mathbb{R}$ is a given twice continuously differentiable, concave function, and that $x_0 \in X$ and $\beta \in (0,1)$ are given. Choose x_1, x_2, \dots to maximise:

$$\liminf_{T \rightarrow \infty} \sum_{t=1}^T \beta^{t-1} \mathcal{F}(x_{t-1}, x_t),$$

subject to the constraints that for all $t \in \mathbb{N}^+$, $x_t \in \Gamma(x_{t-1})$.

For tractability, we make the following additional assumption, which enables us to uniformly approximate Γ by a finite number of inequalities:

Assumption 5 X is compact.

Then, by theorem 4.8 of Stokey, Lucas, and Prescott (1989), there is a unique solution to Problem 8 for any x_0 . We further assume the following to ensure that there is a natural point to approximate around:⁴²

Assumption 6 There exists $\mu \in X$ such that for any given $x_0 \in X$, in the solution to Problem 8 with that x_0 , as $t \rightarrow \infty$, $x_t \rightarrow \mu$.

⁴² If X is convex, then the existence of a fixed point of the policy function is a consequence of the Brouwer fixed point theorem, but there is no reason the fixed point guaranteed by Brouwer's theorem should be even locally attractive.

Having defined μ , we can let $\tilde{\mathcal{F}}$ be a second order Taylor approximation to \mathcal{F} around μ , which will take the form of equation (7). Assumption 3 will be satisfied for this approximation thanks to the concavity of \mathcal{F} . To apply the previous results, we also then need to approximate the constraints.

Suppose first that the graph of Γ is convex, i.e. the set $\{(x, z) | x \in X, z \in \Gamma(x)\}$ is convex. Since it is also compact, by Assumption 5, for any $\epsilon > 0$, there exists $c \in \mathbb{N}$, $\Psi^{(0)} \in \mathbb{R}^{c \times 1}$ and $\Psi^{(1)} \in \mathbb{R}^{c \times 2n}$ such that with $\tilde{\Gamma}$ defined as in equation (5) and \tilde{X} defined as in equation (6):

- 1) $\mu \in \tilde{X} \subseteq X$,
- 2) for all $x \in X$, there exists $\tilde{x} \in \tilde{X}$ such that $\|x - \tilde{x}\|_2 < \epsilon$,
- 3) for all $x \in \tilde{X}$, $\tilde{\Gamma}(x) \subseteq \Gamma(x)$,
- 4) for all $x \in \tilde{X}$, and for all $z \in \Gamma(x)$, there exists $\tilde{z} \in \tilde{\Gamma}(x)$ such that $\|z - \tilde{z}\|_2 < \epsilon$.

(This follows from standard properties of convex sets.) Then, by our previous results, the following proposition is immediate:

Proposition 19 Suppose we are given a problem in the form of Problem 8 (and which satisfies Assumption 5 and Assumption 6). If the graph of Γ is convex, then we can construct a problem in the form of the multiple-bound generalisation of Problem 2 which encodes a local approximation to the original dynamic programming problem around $x_t = \mu$. Furthermore, the LCP corresponding to this approximation will have a unique solution for all $x_0 \in \tilde{X}$. Moreover, the approximation is consistent for quadratic objectives in the sense that as the number of inequalities used to approximate Γ goes to infinity, the approximate value function converges uniformly to the true value function.

Unfortunately, if the graph of Γ is non-convex, then we will not be able to derive similar results. To see the best we could do along similar proof lines, here we merely sketch the construction of an approximation to the graph of Γ in this case. We will need to assume that there exists $z \in \text{int } \Gamma(x)$ for all $x \in X$, which precludes the existence of equality constraints.⁴³ We first approximate the graph of Γ by a polytope (i.e. n dimensional polygon) contained in the graph of Γ such that all points in the graph of Γ are within $\frac{\epsilon}{2}$ of a point in the polytope. Then, providing ϵ is sufficiently small, for each simplicial surface element of the polytope, indexed by $k \in \{1, \dots, c\}$, we can find a quadratic function $q_k: X \times X \rightarrow \mathbb{R}$ with:

$$q_k = \Psi_k^{(0)} + \Psi_{k,\cdot}^{(1)} \begin{bmatrix} x - \mu \\ z - \mu \end{bmatrix} + \begin{bmatrix} x - \mu \\ z - \mu \end{bmatrix}' \Psi_k^{(2)} \begin{bmatrix} x - \mu \\ z - \mu \end{bmatrix}$$

⁴³ This is often not too much of a restriction, since equality constraints may be substituted out.

for all $x, z \in X$ and such that q_k is zero at the corners of the simplicial surface element, such that q_k is weakly negative on its surface, such that $\Psi_k^{(2)}$ is symmetric positive definite, and such that all points in the polytope are within $\frac{\epsilon}{2}$ of a point in the set:

$$\{(x, z) \in X \times X | \forall k \in \{1, \dots, S\}, 0 \leq q_k(x, z)\}.$$

This gives a set of quadratic constraints that approximate Γ . If we then define:

$$\tilde{u}^{(2)} := u^{(2)} + \sum_{k=1}^c \bar{\lambda}'_{\Psi, k} \Psi_k^{(2)},$$

where $u^{(2)}$ is the Hessian of \mathcal{F} , then the Lagrangian in equation (9) is the same as what would be obtained from taking a second order Taylor approximation to the Lagrangian of the problem of maximising our non-linear objective subject to the approximate quadratic constraints, suggesting it may perform acceptably well for x near μ , along similar lines to the results of Levine, Pearlman, and Pierser (2008) and Benigno and Woodford (2012). However, existence of a unique solution to the original problem cannot be used to establish even the existence of a solution of the approximated problem, since only linear approximations to the quadratic constraints would be imposed by our algorithm, giving a greatly reduced choice set (as the quadratic terms are positive definite).

I. Proof of the properties of the BPY model

Defining $x_t = [x_{i,t} \quad x_{y,t} \quad x_{\pi,t}]'$, the BPY model is in the form of Problem 2, with:

$$A := \begin{bmatrix} 0 & -\alpha_{\Delta y} & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad B := \begin{bmatrix} -1 & \alpha_{\Delta y} & \alpha_{\pi} \\ -\frac{1}{\sigma} & -1 & 0 \\ 0 & \gamma & -1 \end{bmatrix}, \quad C := \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & \frac{1}{\sigma} \\ 0 & 0 & \beta \end{bmatrix}.$$

Assumption 2 is satisfied for this model as:

$$\det(A + B + C) = \det \begin{bmatrix} -1 & 0 & \alpha_{\pi} \\ -\frac{1}{\sigma} & 0 & \frac{1}{\sigma} \\ 0 & \gamma & -1 \end{bmatrix} \neq 0$$

as $\alpha_{\pi} \neq 1$ and $\gamma \neq 0$. Let $f := F_{2,2}$, where F is as in Assumption 1. Then:

$$F = \begin{bmatrix} 0 & \alpha_{\Delta y}(f-1) + \alpha_{\pi} \frac{\gamma f}{1-\beta f} & 0 \\ 0 & f & 0 \\ 0 & \frac{\gamma f}{1-\beta f} & 0 \end{bmatrix}.$$

Hence:

$$f = f^2 - \frac{1}{\sigma} \left(\alpha_{\Delta y}(f-1) + \alpha_{\pi} \frac{\gamma f}{1-\beta f} - \frac{\gamma f^2}{1-\beta f} \right),$$

i.e.:

$$\beta \sigma f^3 - ((\alpha_{\Delta y} + \sigma)\beta + \gamma + \sigma)f^2 + ((1 + \beta)\alpha_{\Delta y} + \gamma\alpha_{\pi} + \sigma)f - \alpha_{\Delta y} = 0. \quad (19)$$

When $f \leq 0$, the left hand side is negative, and when $f = 1$, the left hand side equals $(\alpha_\pi - 1)\gamma > 0$ (by assumption on α_π), hence equation (3) has either one or three solutions in $(0,1)$, and no solutions in $(-\infty, 0]$. We wish to prove there is a unique solution in $(-1,1)$. First note that when $\alpha_\pi = 1$, the discriminant of the polynomial is:

$$\left((1 - \beta)(\alpha_{\Delta y} - \sigma) - \gamma\right)^2 \left((\beta\alpha_{\Delta y})^2 + 2\beta(\gamma - \sigma)\alpha_{\Delta y} + (\gamma + \sigma)^2\right).$$

The first multiplicand is positive. The second is minimised when $\sigma = \beta\alpha_{\Delta y} - \gamma$, at the value $4\beta\gamma\alpha_{\Delta y} > 0$, hence this multiplicand is positive too. Consequently, at least for small α_π , there are three real solutions for f , so there may be multiple solutions in $(0,1)$.

Suppose for a contradiction that there were at least three solutions to equation (3) in $(0,1)$ (double counting repeated roots), even for arbitrary large $\beta \in (0,1)$. Let $f_1, f_2, f_3 \in (0,1)$ be the three roots. Then, by Vieta's formulas:

$$\begin{aligned} 3 > f_1 + f_2 + f_3 &= \frac{(\alpha_{\Delta y} + \sigma)\beta + \gamma + \sigma}{\beta\sigma}, \\ 3 > f_1f_2 + f_1f_3 + f_2f_3 &= \frac{(1 + \beta)\alpha_{\Delta y} + \gamma\alpha_\pi + \sigma}{\beta\sigma}, \\ 1 > f_1f_2f_3 &= \frac{\alpha_{\Delta y}}{\beta\sigma}, \end{aligned}$$

so:

$$\begin{aligned} (2\beta - 1)\sigma &> \beta\alpha_{\Delta y} + \gamma > \gamma > 0 \\ \beta &> \frac{1}{2}, \quad (2\beta - 1)\sigma > \gamma, \\ \beta\sigma &> \beta\alpha_{\Delta y} + \gamma + \sigma(1 - \beta), \\ 2\beta\sigma &> (1 + \beta)\alpha_{\Delta y} + \gamma\alpha_\pi + \sigma(1 - \beta), \\ \beta\sigma &> \alpha_{\Delta y}. \end{aligned}$$

Also, the first derivative of equation (3) must be strictly positive at $f = 1$, so:

$$(1 - \beta)(\alpha_{\Delta y} - \sigma) + (\alpha_\pi - 2)\gamma > 0.$$

Combining all of these inequalities gives the bounds:

$$\begin{aligned} 0 < \alpha_{\Delta y} &< 2\sigma - \frac{\gamma + \sigma}{\beta}, \\ 2 + \frac{(1 - \beta)(\sigma - \alpha_{\Delta y})}{\gamma} &< \alpha_\pi < \frac{(3\beta - 1)\sigma - (1 + \beta)\alpha_{\Delta y}}{\gamma}. \end{aligned}$$

Furthermore, if there are multiple solutions to equation (3), then the discriminant of its first derivative must be weakly positive, i.e.:

$$\left((\alpha_{\Delta y} + \sigma)\beta + \gamma + \sigma\right)^2 - 3\beta\sigma \left((1 + \beta)\alpha_{\Delta y} + \gamma\alpha_\pi + \sigma\right) \geq 0.$$

Therefore, we have the following bounds on α_π :

$$2 + \frac{(1 - \beta)(\sigma - \alpha_{\Delta y})}{\gamma} < \alpha_\pi \leq \frac{\left((\alpha_{\Delta y} + \sigma)\beta + \gamma + \sigma\right)^2 - 3\beta\sigma \left((1 + \beta)\alpha_{\Delta y} + \sigma\right)}{3\beta\sigma\gamma}$$

since,

$$\begin{aligned} & \frac{(3\beta - 1)\sigma - (1 + \beta)\alpha_{\Delta y}}{\gamma} - \frac{\left((\alpha_{\Delta y} + \sigma)\beta + \gamma + \sigma\right)^2 - 3\beta\sigma\left((1 + \beta)\alpha_{\Delta y} + \sigma\right)}{3\beta\sigma\gamma} \\ &= \frac{\left((2\sigma - \alpha_{\Delta y})\beta - \gamma - \sigma\right)\left((4\sigma + \alpha_{\Delta y})\beta + \gamma + \sigma\right)}{3\beta\gamma\sigma} > 0 \end{aligned}$$

as $\alpha_{\Delta y} < 2\sigma - \frac{\gamma + \sigma}{\beta}$. Consequently, there exists $\lambda, \mu, \kappa \in [0, 1]$ such that:

$$\begin{aligned} \alpha_{\pi} &= (1 - \lambda) \left[2 + \frac{(1 - \beta)(\sigma - \alpha_{\Delta y})}{\gamma} \right] \\ &+ \lambda \left[\frac{\left((\alpha_{\Delta y} + \sigma)\beta + \gamma + \sigma\right)^2 - 3\beta\sigma\left((1 + \beta)\alpha_{\Delta y} + \sigma\right)}{3\beta\sigma\gamma} \right], \\ \alpha_{\Delta y} &= (1 - \mu)[0] + \mu \left[2\sigma - \frac{\gamma + \sigma}{\beta} \right], \\ \gamma &= (1 - \kappa)[0] + \kappa[(2\beta - 1)\sigma] \end{aligned}$$

These simultaneous equations have unique solutions for α_{π} , $\alpha_{\Delta y}$ and γ in terms of λ , μ and κ . Substituting these solutions into the discriminant of equation (3) gives a polynomial in $\lambda, \mu, \kappa, \beta, \sigma$. As such, an exact global maximum of the discriminant may be found subject to the constraints $\lambda, \mu, \kappa \in [0, 1]$, $\beta \in [\frac{1}{2}, 1]$, $\sigma \in [0, \infty)$, by using an exact compact polynomial optimisation solver, such as that in the Maple computer algebra package. Doing this gives a maximum of 0 when $\beta \in \{\frac{1}{2}, 1\}$, $\kappa = 1$ and $\sigma = 0$. But of course, we actually require that $\beta \in (\frac{1}{2}, 1)$, $\kappa < 1$, $\sigma > 0$. Thus, by continuity, the discriminant is strictly negative over the entire possible domain. This gives the required contradiction to our assumption of three roots to the polynomial, establishing that Assumption 1 holds for this model.

Now, when $T = 1$, M is equal to the top left element of the matrix $-(B + CF)^{-1}$, i.e.:

$$M = \frac{\beta\sigma f^2 - \left((1 + \beta)\sigma + \gamma\right)f + \sigma}{\beta\sigma f^2 - \left((1 + \beta)\sigma + \gamma + \beta\alpha_{\Delta y}\right)f + \sigma + \alpha_{\Delta y} + \gamma\alpha_{\pi}}$$

Now, multiplying the denominator by f gives:

$$\begin{aligned} & \beta\sigma f^3 - \left((1 + \beta)\sigma + \gamma + \beta\alpha_{\Delta y}\right)f^2 + \left(\sigma + \alpha_{\Delta y} + \gamma\alpha_{\pi}\right)f \\ &= \left[\beta\sigma f^3 - \left((\alpha_{\Delta y} + \sigma)\beta + \gamma + \sigma\right)f^2 + \left((1 + \beta)\alpha_{\Delta y} + \gamma\alpha_{\pi} + \sigma\right)f\right. \\ &\quad \left. - \alpha_{\Delta y}\right] - \left[\beta\alpha_{\Delta y}f - \alpha_{\Delta y}\right] = (1 - \beta f)\alpha_{\Delta y} > 0, \end{aligned}$$

by equation (19). Hence, the sign of M is that of $\beta\sigma f^2 - \left((1 + \beta)\sigma + \gamma\right)f + \sigma$. I.e., M is negative if and only if:

$$\begin{aligned} & \frac{\left((1 + \beta)\sigma + \gamma\right) - \sqrt{\left((1 + \beta)\sigma + \gamma\right)^2 - 4\beta\sigma^2}}{2\beta\sigma} < f \\ & < \frac{\left((1 + \beta)\sigma + \gamma\right) + \sqrt{\left((1 + \beta)\sigma + \gamma\right)^2 - 4\beta\sigma^2}}{2\beta\sigma}. \end{aligned}$$

The upper limit is greater than 1, so only the lower is relevant. To translate this bound on f into a bound on $\alpha_{\Delta y}$, we first need to establish that f is monotonic in $\alpha_{\Delta y}$.

Totally differentiating equation (19) gives:

$$\begin{aligned} & [3\beta\sigma f^2 - 2((\alpha_{\Delta y} + \sigma)\beta + \gamma + \sigma)f + ((1 + \beta)\alpha_{\Delta y} + \gamma\alpha_{\pi} + \sigma)] \frac{df}{d\alpha_{\Delta y}} \\ & = (1 - \beta f)(1 - f) > 0. \end{aligned}$$

Thus, the sign of $\frac{df}{d\alpha_{\Delta y}}$ is equal to that of:

$$3\beta\sigma f^2 - 2((\alpha_{\Delta y} + \sigma)\beta + \gamma + \sigma)f + ((1 + \beta)\alpha_{\Delta y} + \gamma\alpha_{\pi} + \sigma).$$

Note, however, that this expression is just the derivative of the left hand side of equation (19) with respect to f .

To establish the sign of $\frac{df}{d\alpha_{\Delta y}}$, we consider two cases. First, suppose that equation (19) has three real solutions. Then, the unique solution to equation (19) in $(0,1)$ is its lowest solution. Hence, this solution must be below the first local maximum of the left hand side of equation (19). Consequently, at the $f \in (0,1)$, which solves equation (19), $3\beta\sigma f^2 - 2((\alpha_{\Delta y} + \sigma)\beta + \gamma + \sigma)f + ((1 + \beta)\alpha_{\Delta y} + \gamma\alpha_{\pi} + \sigma) > 0$. Alternatively, suppose that equation (19) has a unique real solution. Then the left hand side of this equation cannot change sign in between its local maximum and its local minimum (if it has any). Thus, at the $f \in (0,1)$ at which it changes sign, we must have that $3\beta\sigma f^2 - 2((\alpha_{\Delta y} + \sigma)\beta + \gamma + \sigma)f + ((1 + \beta)\alpha_{\Delta y} + \gamma\alpha_{\pi} + \sigma) > 0$. Therefore, in either case $\frac{df}{d\alpha_{\Delta y}} > 0$, meaning that f is monotonic increasing in $\alpha_{\Delta y}$.

Consequently, to find the critical $(f, \alpha_{\Delta y})$ at which M changes sign, it is sufficient to find the lowest solution with respect to both f and $\alpha_{\Delta y}$ of the pair of equations:

$$\begin{aligned} & \beta\sigma f^2 - ((1 + \beta)\sigma + \gamma)f + \sigma = 0, \\ & \beta\sigma f^3 - ((\alpha_{\Delta y} + \sigma)\beta + \gamma + \sigma)f^2 + ((1 + \beta)\alpha_{\Delta y} + \gamma\alpha_{\pi} + \sigma)f - \alpha_{\Delta y} = 0. \end{aligned}$$

The former implies that:

$$\beta\sigma f^3 - ((1 + \beta)\sigma + \gamma)f^2 + \sigma f = 0,$$

so, by the latter:

$$\alpha_{\Delta y}\beta f^2 - ((1 + \beta)\alpha_{\Delta y} + \gamma\alpha_{\pi})f + \alpha_{\Delta y} = 0.$$

If $\alpha_{\Delta y} = \sigma\alpha_{\pi}$, then this equation holds if and only if:

$$\sigma\beta f^2 - ((1 + \beta)\sigma + \gamma)f + \sigma = 0.$$

Therefore, the critical $(f, \alpha_{\Delta y})$ at which M changes sign are given by:

$$\begin{aligned} & \alpha_{\Delta y} = \sigma\alpha_{\pi}, \\ & f = \frac{((1 + \beta)\sigma + \gamma) - \sqrt{((1 + \beta)\sigma + \gamma)^2 - 4\beta\sigma^2}}{2\beta\sigma}. \end{aligned}$$

Thus, M is negative if and only if $\alpha_{\Delta y} > \sigma\alpha_{\pi}$, and M is zero if and only if $\alpha_{\Delta y} = \sigma\alpha_{\pi}$.

J. Proof of the properties of the BPY model with level targeting

Defining $x_t = [x_{i,t} \quad x_{y,t} \quad x_{p,t}]'$, the model of section 3.3 is in the form of Problem 2, with:

$$A := \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad B := \begin{bmatrix} -1 & \alpha_{\Delta y} & \alpha_{\pi} \\ -\frac{1}{\sigma} & -1 & -\frac{1}{\sigma} \\ 0 & \gamma & -1 - \beta \end{bmatrix}, \quad C := \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & \frac{1}{\sigma} \\ 0 & 0 & \beta \end{bmatrix}.$$

Assumption 2 is satisfied for this model as:

$$\det(A + B + C) = \det \begin{bmatrix} -1 & \alpha_{\Delta y} & \alpha_{\pi} \\ -\frac{1}{\sigma} & 0 & 0 \\ 0 & \gamma & -1 \end{bmatrix} \neq 0$$

as $\alpha_{\Delta y} \neq 0$ and $\alpha_{\pi} \neq 0$. Let $f := F_{3,3}$, where F is as in Assumption 1. Then:

$$F = \begin{bmatrix} 0 & 0 & \frac{f(1-f)(\sigma\alpha_{\pi} - \alpha_{\Delta y})}{\alpha_{\Delta y} + (1-f)\sigma} \\ 0 & 0 & \frac{f(1-f - \alpha_{\pi})}{\alpha_{\Delta y} + (1-f)\sigma} \\ 0 & 0 & f \end{bmatrix},$$

and so:

$$\beta\sigma f^3 - \left((1 + 2\beta)\sigma + \beta\alpha_{\Delta y} + \gamma \right) f^2 + \left((2 + \beta)\sigma + (1 + \beta)\alpha_{\Delta y} + (1 + \alpha_{\pi})\gamma \right) f - (\sigma + \alpha_{\Delta y}) = 0.$$

Now define:

$$\hat{\alpha}_{\Delta y} := \sigma + \alpha_{\Delta y}, \quad \hat{\alpha}_{\pi} := 1 + \alpha_{\pi}$$

so:

$$\beta\sigma f^3 - \left((\hat{\alpha}_{\Delta y} + \sigma)\beta + \gamma + \sigma \right) f^2 + \left((1 + \beta)\hat{\alpha}_{\Delta y} + \gamma\hat{\alpha}_{\pi} + \sigma \right) f - \hat{\alpha}_{\Delta y} = 0.$$

This is identical to the equation for f in the previous section, apart from the fact that $\hat{\alpha}_{\Delta y}$ has replaced $\alpha_{\Delta y}$ and $\hat{\alpha}_{\pi}$ has replaced α_{π} . Hence, by the results of the previous section, Assumption 1 holds for this model as well.

For this model, with $T = 1$:

$$M = \frac{(1-f)(1 + (1-f)\beta)\sigma^2 + \left((1 + (1-f)\beta)\alpha_{\Delta y} + ((1-f) + \alpha_{\pi}f)\gamma \right) \sigma + (1-f)\gamma\alpha_{\Delta y}}{\left((1-f)(1 + (1-f)\beta)\sigma + (1 + (1-f)\beta)\alpha_{\Delta y} + ((1-f) + \alpha_{\pi})\gamma \right) (\sigma + \alpha_{\Delta y})} > 0.$$

K. Proof that the existence of a polynomial time approximate solution would imply P=NP

Suppose that we have some solution procedure which accepts problems in the form of Problem 2 (and possibly other problems) together with a radius κ and an accuracy level ϵ , and, in time polynomial in n , returns a set $\mathcal{D}_{\kappa, \epsilon} \subseteq \mathbb{R}^n$ membership of which

may be evaluated in time polynomial in n , and a policy function $p_{\kappa,\epsilon}: \mathcal{D}_{\kappa,\epsilon} \subseteq \mathbb{R}^n \rightarrow \mathbb{R}^n$ that may be evaluated in time polynomial in n and that satisfies:

$$\|p^*(x) - p_{\kappa,\epsilon}(x)\|_\infty < \epsilon,$$

for all $x \in \mathcal{D}^* \subseteq \mathcal{D}_\epsilon$ with $\|x - \mu\|_\infty < \kappa$, and where $p^*: \mathcal{D}^* \rightarrow \mathcal{D}^*$ is an exact policy function, i.e. a function satisfying:

1. For all $x \in \mathcal{D}^*$:

$$x_1 = \max\{0, I_{1,\cdot} \mu + A_{1,\cdot} (x - \mu) + (B_{1,\cdot} + I_{1,\cdot}) (p^*(x) - \mu) + C_{1,\cdot} (p^*(p^*(x)) - \mu)\},$$

$$(A_{-1,\cdot} + B_{-1,\cdot} + C_{-1,\cdot}) \mu = A_{-1,\cdot} x + B_{-1,\cdot} p^*(x) + C_{-1,\cdot} p^*(p^*(x)).$$

2. For all $x_0 \in \mathcal{D}^*$, if $x_t = p^*(x_{t-1})$ for all $t \in \mathbb{N}^+$, then $x_t \rightarrow \mu$ as $t \rightarrow \infty$.

3. For all $x_0 \notin \mathcal{D}^*$, there is no $x_1, x_2, \dots \in \mathbb{R}^n$ which solve this instance of Problem 2.

Now consider the following “knapsack”-type problem from Chung (1989):

Problem 9 Suppose $a_1, \dots, a_{T-2}, b \in \mathbb{N}^+$ are given. Find $z_1, \dots, z_{T-2} \in \{0,1\}$ such that $\sum_{j=1}^{T-2} a_j z_j = b$.

Chung (1989) shows that there exists $q \in \mathbb{Z}^T$, $\mathcal{M} \in \mathbb{Z}^{T \times T}$, such that Problem 9 has a solution if and only if the LCP (q, \mathcal{M}) has a solution, where q and \mathcal{M} may be computed from a_1, \dots, a_{T-2} and b in time polynomial in T . Furthermore, the details of the proof in Chung (1989) reveal that for any $y \in \mathbb{R}^T$ that solves the LCP (q, \mathcal{M}) , $y_{T-1} = y_T = 0$, and setting $z_t = \frac{y_t}{a_t}$ for $t \in \{1, \dots, T-2\}$ solves Problem 9. Since $a_1, \dots, a_{T-2} \in \mathbb{N}^+$ and $z_1, \dots, z_{T-2} \in \{0,1\}$, this implies that $y_t \in \{0, a_t\} \subseteq \mathbb{N}$ for $t \in \{1, \dots, T-2\}$. Moreover, by Proposition 2, given a_1, \dots, a_{T-2} and b we can thus construct a model in the form of Problem 2 in polynomial time in T , featuring polynomial in T state variables, and such that for an appropriately chosen initial state, setting $z_t = \frac{x_{1,t}}{a_t}$ for $t \in \{1, \dots, T-2\}$ solves Problem 9 if it has a solution (where $x_{1,t}$ gives the path of the bounded variable in the constructed model). Additionally, by inspecting the proof of Proposition 2 from section A, we see that since $q \in \mathbb{Z}^T$, $\mathcal{M} \in \mathbb{Z}^{T \times T}$, and $y_t \in \mathbb{N}$ for all $t \in \{1, \dots, T-2\}$, it must be the case that for all $t \in \mathbb{N}$, $x_t \in \mathbb{Z}^n$ (for some $n \in \mathbb{N}$), so the model is always integer valued.

To complete the proof, we set $\epsilon := \frac{1}{2}$, $\kappa := \|x_0 - \mu\|_\infty + 1$ and construct the policy function $p_{\kappa,\epsilon}$ for the constructed model. By assumption, we can do this in time polynomial in T . Using this we can construct an exact solution for z_t as follows. Set $\hat{x}_0 := x_0$. Now suppose we have defined \hat{x}_{t-1} for some $t \in \mathbb{N}^+$. We first test if $\hat{x}_{t-1} \in \mathcal{D}_{\kappa,\epsilon}$. If it is not, then we terminate the procedure with a “no solution” message. Otherwise, we set each element of \hat{x}_t to be equal to the nearest integer to the corresponding element of $p_{\kappa,\epsilon}(\hat{x}_{t-1})$. We then test if $\|\hat{x}_t - \mu\| < \kappa$. If it is, we proceed,

otherwise, we redefine $\kappa := \|\hat{x}_t - \mu\|_\infty + 1$, and restart. After an amount of time bounded by a polynomial in T , we will have either terminated with a “no solution” message, or have successfully defined \hat{x}_t for $t \in \{0, \dots, T-2\}$, where for all $t \in \{0, \dots, T-2\}$, we will have found some κ for which $\|\hat{x}_t - \mu\| < \kappa$. In this case, we then define $\hat{z}_t = \frac{\hat{x}_{1,t}}{a_t}$ for $t \in \{1, \dots, T-2\}$ and test whether it solves Problem 9. If it does, report a “solution found” message, otherwise, report a “no solution” message. We now prove that this procedure works.

We first prove that if $x_0 \in \mathcal{D}^* \cap \mathbb{Z}^n$ then $\hat{x}_t = x_t \in \mathcal{D}^* \cap \mathbb{Z}^n$ for all $t \in \mathbb{N}$, by induction on t . The inductive base case is trivial. Suppose for the inductive step that for some $t \in \mathbb{N}^+$, $\hat{x}_{t-1} = x_{t-1} \in \mathcal{D}^* \cap \mathbb{Z}^n$. Then $\hat{x}_{t-1} \in \mathcal{D}_{\kappa, \epsilon}$, so \hat{x}_t is defined and is equal to the (elementwise) nearest integer to $p_{\kappa, \epsilon}(\hat{x}_{t-1})$. Then since as $\hat{x}_{t-1} \in \mathcal{D}^*$ and $\|\hat{x}_{t-1} - \mu\| < \kappa$, $\|p^*(\hat{x}_{t-1}) - p_{\kappa, \epsilon}(\hat{x}_{t-1})\|_\infty < \frac{1}{2}$, and $p^*: \mathcal{D}^* \cap \mathbb{Z}^n \rightarrow \mathcal{D}^* \cap \mathbb{Z}^n$, it must be the case that $\hat{x}_t = p^*(\hat{x}_{t-1}) \in \mathcal{D}^* \cap \mathbb{Z}^n$. This establishes the inductive hypothesis, and hence if we set $\hat{z}_t = \frac{\hat{x}_{1,t}}{a_t}$ for $t \in \{1, \dots, T-2\}$ then test whether it solves Problem 9, we will find that it does, giving a solution, after only polynomial in T calculations.

Now suppose that $x_0 \in \mathbb{Z}^n$, but $x_0 \notin \mathcal{D}^*$. There are two possibilities. Either we will find some $t \in \{1, \dots, T-1\}$ for which $\hat{x}_t \notin \mathcal{D}_\epsilon$, and hence $\hat{x}_t \notin \mathcal{D}^*$, or we will successfully calculate $\hat{x}_{1,t}$ for $t \in \{1, \dots, T-2\}$. By the previous result, in the former case we will have found in polynomial time in T a proof that $x_0 \notin \mathcal{D}^*$, and hence that there is no solution to Problem 9. In the latter case, we can again set $\hat{z}_t = \frac{\hat{x}_{1,t}}{a_t}$ for $t \in \{1, \dots, T-2\}$ and test whether it solves Problem 9, and we will find (in polynomial in T calculations) that it does not, giving an alternative polynomial time in T proof that $x_0 \notin \mathcal{D}^*$, and hence that there is no solution to Problem 9.

We have thus established that the procedure described (calculating $\hat{x}_1, \dots, \hat{x}_{T-2}$ and then testing whether $\hat{z}_t = \frac{\hat{x}_{1,t}}{a_t}$ solves Problem 9), enables us to answer the question of whether Problem 9 has a solution in an amount of time that is polynomial in T . But Problem 9 is NP-complete (Karp 1972), and hence this implies that P=NP.

L. Special cases with polynomial time solutions

Polynomial time algorithms exist for the LCP if M is general positive semi-definite (Kojima, Mizuno, and Yoshise 1989). However, it appears that M is general positive semi-definite in only very few macroeconomic models, so this is of minimal relevance. Furthermore, If either condition 1 or condition 2 of Proposition 5 is known to be satisfied (e.g. M is row sufficient), then we can find out if a solution exists in polynomial time, by solving the feasibility problem. Moreover, a polynomial time

algorithm exists (Illés, Nagy, and Terlaky 2010) which will give a certificate that one of the following is true, for a given q and real number $\bar{\kappa} \geq 0$:

- For any $\kappa \leq \bar{\kappa}$, M is not a member of the matrix class $P^*(\kappa)$, defined in the paper. (Note that for $\kappa_1 < \kappa_2$, $P^*(\kappa_1) \subseteq P^*(\kappa_2)$, and that the class of sufficient matrices is the union of the classes of $P^*(\kappa)$ matrices for all $\kappa \geq 0$.)
- The LCP (q, M) has no solution.
- The LCP (q, M) has the solution y .

Thus for “most” sufficient matrices we can find a solution (or a certificate that there is none), in polynomial time. It has been conjectured that in fact this holds for all sufficient matrices (Fukuda 2015).

Unfortunately, no algorithm is known for finding out if M is sufficient in polynomial time. Indeed, it has also been shown (Coxson 1994; Tseng 2000) that it is “co-NP complete” to test if M is non-degenerate, a P-matrix, a P_0 -matrix, semi-monotone, strictly semi-monotone, column sufficient or row sufficient. This means that were a polynomial time (in T) algorithm available for these things then we would have a proof that $P=NP$.

M. The augmented state-space representation of a pruned perturbation solution

We seek to convert the model into the form:

$$\begin{aligned} z_t &= o + P\tilde{z}_{t-1} + Q\xi_t, \\ x_t &= u + Vz_t, \end{aligned}$$

where $\mathbb{E}_{t-1}\xi_t = 0$, and where throughout, $\tilde{\cdot}$ s over variables denote the subset of state variables. We proceed by taking each order of approximation in turn. We assume that the original model has l state variables. Of the assorted algorithms available for pruning, it appears that Lan and Meyer-Gohde’s (2013a) algorithm is the most accurate (Lan and Meyer-Gohde 2013b), and so both the discussion below, and the implementation in DynareOBC is based on this approach, however, everything we say would also go through with alternative pruning algorithms.

Order 1 At order 1:

$$\begin{aligned} x_t^{(1)} &= \alpha\tilde{x}_{t-1}^{(1)} + \beta_0\varepsilon_t, \\ x_t &= \mu_x + x_t^{(1)}, \end{aligned}$$

so if we define:

$$\begin{aligned} z_t &:= x_t^{(1)}, & \tilde{z}_t &:= \tilde{x}_t^{(1)}, & o &:= 0, & P &:= \alpha, & Q &:= \beta_0, & \xi_t &:= \varepsilon_t, \\ & & u &:= \mu_x, & V &:= I_n, & & & & & & \end{aligned}$$

then we are done.

Order 2 At order 2:

$$\begin{aligned} x_t^{(1)} &= \alpha \tilde{x}_{t-1}^{(1)} + \beta_0 \varepsilon_t, \\ x_t^{(2)} &= \tilde{\alpha} \tilde{x}_{t-1}^{(2)} + \frac{1}{2} \beta_{22} (\tilde{x}_{t-1}^{(1)} \otimes \tilde{x}_{t-1}^{(1)}) + \beta_{20} (\tilde{x}_{t-1}^{(1)} \otimes \varepsilon_t) + \frac{1}{2} \beta_{00} (\varepsilon_t \otimes \varepsilon_t), \\ x_t &= \mu_x + x^{(0)} + x_t^{(1)} + x_t^{(2)}, \end{aligned}$$

for some constant $x^{(0)}$.

Now, note that:

$$\begin{aligned} \tilde{x}_t^{(1)} \otimes \tilde{x}_t^{(1)} &= (\tilde{\alpha} \tilde{x}_{t-1}^{(1)} + \tilde{\beta}_0 \varepsilon_t) \otimes (\tilde{\alpha} \tilde{x}_{t-1}^{(1)} + \tilde{\beta}_0 \varepsilon_t) \\ &= \tilde{\alpha} \tilde{x}_{t-1}^{(1)} \otimes \tilde{\alpha} \tilde{x}_{t-1}^{(1)} + \tilde{\alpha} \tilde{x}_{t-1}^{(1)} \otimes \tilde{\beta}_0 \varepsilon_t + \tilde{\beta}_0 \varepsilon_t \otimes \tilde{\alpha} \tilde{x}_{t-1}^{(1)} + \tilde{\beta}_0 \varepsilon_t \otimes \tilde{\beta}_0 \varepsilon_t \\ &= (\tilde{\alpha} \otimes \tilde{\alpha}) (\tilde{x}_{t-1}^{(1)} \otimes \tilde{x}_{t-1}^{(1)}) + (\tilde{\alpha} \otimes \tilde{\beta}_0) (\tilde{x}_{t-1}^{(1)} \otimes \varepsilon_t) + (\tilde{\beta}_0 \otimes \tilde{\alpha}) (\varepsilon_t \otimes \tilde{x}_{t-1}^{(1)}) \\ &\quad + (\tilde{\beta}_0 \otimes \tilde{\beta}_0) (\varepsilon_t \otimes \varepsilon_t) \\ &= (\tilde{\alpha} \otimes \tilde{\alpha}) (\tilde{x}_{t-1}^{(1)} \otimes \tilde{x}_{t-1}^{(1)}) + ((\tilde{\alpha} \otimes \tilde{\beta}_0) + (\tilde{\beta}_0 \otimes \tilde{\alpha}) K_{m,l}) (\tilde{x}_{t-1}^{(1)} \otimes \varepsilon_t) \\ &\quad + (\tilde{\beta}_0 \otimes \tilde{\beta}_0) (\varepsilon_t \otimes \varepsilon_t) \\ &= (\tilde{\alpha} \otimes \tilde{\alpha}) (\tilde{x}_{t-1}^{(1)} \otimes \tilde{x}_{t-1}^{(1)}) + ((\tilde{\alpha} \otimes \tilde{\beta}_0) + K_{l,l} (\tilde{\alpha} \otimes \tilde{\beta}_0) K_{l,m} K_{m,l}) (\tilde{x}_{t-1}^{(1)} \otimes \varepsilon_t) \\ &\quad + (\tilde{\beta}_0 \otimes \tilde{\beta}_0) (\varepsilon_t \otimes \varepsilon_t) \\ &= (\tilde{\alpha} \otimes \tilde{\alpha}) (\tilde{x}_{t-1}^{(1)} \otimes \tilde{x}_{t-1}^{(1)}) + (I_{l^2} + K_{l,l}) (\tilde{\alpha} \otimes \tilde{\beta}_0) (\tilde{x}_{t-1}^{(1)} \otimes \varepsilon_t) \\ &\quad + (\tilde{\beta}_0 \otimes \tilde{\beta}_0) (\varepsilon_t \otimes \varepsilon_t), \end{aligned}$$

where $K_{l,m} \in \mathbb{R}^{lm \times lm}$ is the commutation matrix for $l \times m$ matrices, i.e. it is the unique matrix such that for all $D \in \mathbb{R}^{l \times m}$, $K_{l,m} \text{vec } D = \text{vec } D'$ (Magnus and Neudecker 1979).

Thus, if we define:

$$\begin{aligned} z_t &:= \begin{bmatrix} x_t^{(1)} \\ x_t^{(2)} \\ \tilde{x}_t^{(1)} \otimes \tilde{x}_t^{(1)} \end{bmatrix}, & \tilde{z}_t &:= \begin{bmatrix} \tilde{x}_t^{(1)} \\ \tilde{x}_t^{(2)} \\ \tilde{x}_t^{(1)} \otimes \tilde{x}_t^{(1)} \end{bmatrix}, \\ P &:= \begin{bmatrix} \alpha & 0 & 0 \\ 0 & \alpha & \frac{1}{2} \beta_{22} \\ 0 & 0 & \tilde{\alpha} \otimes \tilde{\alpha} \end{bmatrix}, & Q &:= \begin{bmatrix} \beta_0 & 0 & 0 \\ 0 & \frac{1}{2} \beta_{00} & \beta_{20} \\ 0 & \tilde{\beta}_0 \otimes \tilde{\beta}_0 & (I_{l^2} + K_{l,l}) (\tilde{\alpha} \otimes \tilde{\beta}_0) \end{bmatrix}, \\ \xi_t &:= \begin{bmatrix} \varepsilon_t \\ \varepsilon_t \otimes \varepsilon_t - \text{vec } \Sigma \\ \tilde{x}_{t-1}^{(1)} \otimes \varepsilon_t \end{bmatrix}, & o &:= \begin{bmatrix} 0 \\ \frac{1}{2} \beta_{00} \text{vec } \Sigma \\ (\tilde{\beta}_0 \otimes \tilde{\beta}_0) \text{vec } \Sigma \end{bmatrix}, \\ u &:= \mu_x + x^{(0)}, & V &:= [I \quad I \quad 0], \end{aligned}$$

then we are done.

Order 3 At order 3:

$$\begin{aligned}
x_t^{(1)} &= \alpha \tilde{x}_{t-1}^{(1)} + \beta_0 \varepsilon_t, \\
x_t^{(2)} &= \alpha \tilde{x}_{t-1}^{(2)} + \frac{1}{2} \beta_{22} (\tilde{x}_{t-1}^{(1)} \otimes \tilde{x}_{t-1}^{(1)}) + \beta_{20} (\tilde{x}_{t-1}^{(1)} \otimes \varepsilon_t) + \frac{1}{2} \beta_{00} (\varepsilon_t \otimes \varepsilon_t), \\
x_t^{(1,\sigma^2)} &= \alpha \tilde{x}_{t-1}^{(1,\sigma^2)} + \frac{1}{2} \beta_{\sigma^2,0} \varepsilon_t + \frac{1}{2} \beta_{\sigma^2,1} \tilde{x}_{t-1}^{(1)}, \\
x_t^{(3)} &= \alpha \tilde{x}_{t-1}^{(3)} + \frac{1}{6} \beta_{333,1} (\tilde{x}_{t-1}^{(1)} \otimes \tilde{x}_{t-1}^{(1)} \otimes \tilde{x}_{t-1}^{(1)}) + \frac{1}{6} \beta_{000} (\varepsilon_t \otimes \varepsilon_t \otimes \varepsilon_t) \\
&\quad + \frac{1}{2} \beta_{330,1} (\tilde{x}_{t-1}^{(1)} \otimes \tilde{x}_{t-1}^{(1)} \otimes \varepsilon_t) + \frac{1}{2} \beta_{300} (\tilde{x}_{t-1}^{(1)} \otimes \varepsilon_t \otimes \varepsilon_t) \\
&\quad + \beta_{22} (\tilde{x}_{t-1}^{(2)} \otimes \tilde{x}_{t-1}^{(1)}) + \beta_{20} (\tilde{x}_{t-1}^{(2)} \otimes \varepsilon_t), \\
x_t &= \mu_x + x^{(0)} + x_t^{(1)} + x_t^{(2)} + x_t^{(1,\sigma^2)} + x_t^{(3)},
\end{aligned}$$

for the same constant $x^{(0)}$ as at order 2, providing the shocks have zero skewness (e.g. they are normally distributed). By similar calculations to those at second order, we then have that if we define:

$$\begin{aligned}
z_t^3 &:= \mathbb{E}_s \begin{bmatrix} x_t^{(1)} \\ x_t^{(2)} \\ \tilde{x}_t^{(1)} \otimes \tilde{x}_t^{(1)} \\ x_t^{(1,\sigma^2)} \\ x_t^{(3)} \\ \tilde{x}_t^{(2)} \otimes \tilde{x}_t^{(1)} \\ [\tilde{x}_t^{(1)} \otimes \tilde{x}_t^{(1)} \otimes \tilde{x}_t^{(1)}] \end{bmatrix}, \\
P &:= \begin{bmatrix} \alpha & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \alpha & \frac{1}{2} \beta_{22} & 0 & 0 & 0 & 0 \\ 0 & 0 & \tilde{\alpha} \otimes \tilde{\alpha} & 0 & 0 & 0 & 0 \\ \frac{1}{2} \beta_{\sigma^2,1} & 0 & 0 & \alpha & 0 & 0 & 0 \\ \frac{1}{2} \beta_{300} (I_l \otimes \text{vec } \Sigma) & 0 & 0 & 0 & \alpha & \beta_{22} & \frac{1}{6} \beta_{333,1} \\ \left(\tilde{\beta}_{20} \otimes \tilde{\beta}_0 + \frac{1}{2} K_{l,l} (\tilde{\alpha} \otimes \tilde{\beta}_{00}) \right) (I_l \otimes \text{vec } \Sigma) & 0 & 0 & 0 & 0 & \tilde{\alpha} \otimes \tilde{\alpha} & \frac{1}{2} \tilde{\beta}_{22} \otimes \tilde{\alpha} \\ \left((I_{l^2} + K_{l,l}) \otimes I_l + K_{l^2,l} \right) (\tilde{\alpha} \otimes \tilde{\beta}_0 \otimes \tilde{\beta}_0) (I_l \otimes \text{vec } \Sigma) & 0 & 0 & 0 & 0 & 0 & \tilde{\alpha} \otimes \tilde{\alpha} \otimes \tilde{\alpha} \end{bmatrix}, \\
Q_{11} &:= \begin{bmatrix} \beta_0 & 0 & 0 \\ 0 & \frac{1}{2} \beta_{00} & \beta_{20} \\ 0 & \tilde{\beta}_0 \otimes \tilde{\beta}_0 & (I_{l^2} + K_{l,l}) (\tilde{\alpha} \otimes \tilde{\beta}_0) \\ \frac{1}{2} \beta_{\sigma^2,0} & 0 & 0 \end{bmatrix}, \\
Q_{22} &:= \begin{bmatrix} \beta_{20} & \frac{1}{2} \beta_{330,1} & \frac{1}{2} \beta_{300} & \frac{1}{6} \beta_{000} \\ \tilde{\alpha} \otimes \tilde{\beta}_0 & \frac{1}{2} \tilde{\beta}_{22} \otimes \tilde{\beta}_0 + (\tilde{\beta}_{20} \otimes \tilde{\alpha}) (I_l \otimes K_{m,l}) & \tilde{\beta}_{20} \otimes \tilde{\beta}_0 + \frac{1}{2} K_{l,l} (\tilde{\alpha} \otimes \tilde{\beta}_{00}) & \frac{1}{2} \tilde{\beta}_{00} \otimes \tilde{\beta}_0 \\ 0 & (I_l \otimes (I_{l^2} + K_{l,l}) + K_{l,l^2}) (\tilde{\alpha} \otimes \tilde{\alpha} \otimes \tilde{\beta}_0) & ((I_{l^2} + K_{l,l}) \otimes I_l + K_{l^2,l}) (\tilde{\alpha} \otimes \tilde{\beta}_0 \otimes \tilde{\beta}_0) & \tilde{\beta}_0 \otimes \tilde{\beta}_0 \otimes \tilde{\beta}_0 \end{bmatrix}, \\
Q &:= \begin{bmatrix} Q_{11} & 0 \\ 0 & Q_{22} \end{bmatrix}, \\
\xi_t &:= \begin{bmatrix} \varepsilon_t \\ \varepsilon_t \otimes \varepsilon_t - \text{vec } \Sigma \\ \tilde{x}_{t-1}^{(1)} \otimes \varepsilon_t \\ \tilde{x}_{t-1}^{(2)} \otimes \varepsilon_t \\ \tilde{x}_{t-1}^{(1)} \otimes \tilde{x}_{t-1}^{(1)} \otimes \varepsilon_t \\ \tilde{x}_{t-1}^{(1)} \otimes (\varepsilon_t \otimes \varepsilon_t - \text{vec } \Sigma) \\ \varepsilon_t \otimes \varepsilon_t \otimes \varepsilon_t \end{bmatrix}, \quad o := \begin{bmatrix} 0 \\ \frac{1}{2} \beta_{00} \text{vec } \Sigma \\ (\tilde{\beta}_0 \otimes \tilde{\beta}_0) \text{vec } \Sigma \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \\
u &:= \mu_x + x^{(0)}, \quad V := [I \ I \ 0 \ I \ I \ 0 \ 0],
\end{aligned}$$

then again we are done.

N. The conditional covariance of future variables under a pruned perturbation solution

First, suppose that:

$$z_t = o + Pz_{t-1} + Q\xi_t$$

where $\mathbb{E}_t \xi_{t+k} = 0$ for $k > 0$. Then:

$$\mathbb{E}_t z_{t+k} = \sum_{j=0}^{k-1} P^j o + P^k z_t,$$

so:

$$z_{t+k} - \mathbb{E}_t z_{t+k} = \sum_{j=1}^k P^{k-j} Q \xi_{t+j}.$$

Consequently:

$$\text{cov}_t(z_{t+a}, z_{t+b}) = \sum_{i=1}^a \sum_{j=1}^b P^{a-i} Q (\mathbb{E}_t \xi_{t+i} \xi'_{t+j}) Q' P'^{b-j}$$

If $\mathbb{E}_t \xi_{t+i} \xi'_{t+j} = 0$ for $i \neq j$, then this simplifies to:

$$\begin{aligned} & \text{cov}_t(z_{t+a}, z_{t+b}) \\ &= P^{a-\min\{a,b\}} \left[\sum_{i=1}^{\min\{a,b\}} P^{\min\{a,b\}-i} Q (\mathbb{E}_t \xi_{t+i} \xi'_{t+i}) Q' P'^{\min\{a,b\}-i} \right] P'^{b-\min\{a,b\}}. \end{aligned}$$

Now, in the previous section of these appendices (M), we showed that at order 1, 2 and 3 the pruned perturbation solutions may be represented in the form:

$$z_t = o + P\tilde{z}_{t-1} + Q\xi_t$$

where $\mathbb{E}_t \xi_{t+k} = 0$ for $k > 0$. It is trivial to add zero columns to P so that we instead have:

$$z_t = o + Pz_{t-1} + Q\xi_t,$$

thus, we just need to evaluate $\mathbb{E}_t(\xi_{t+i} \xi'_{t+j})$ in order to have a closed form expression for $\text{cov}_t(z_{t+a}, z_{t+b})$, then from this and the fact that $x_t = u + Vz_t$, we would have that:

$$\text{cov}_t(x_{t+a}, x_{t+b}) = \text{cov}_t(u + Vz_{t+a}, u + Vz_{t+b}) = V \text{cov}_t(z_{t+a}, z_{t+b}) V'.$$

We now proceed to evaluate $\mathbb{E}_t \xi_{t+i} \xi'_{t+j}$ for orders 1 and 2. We skip the third order case as a second order approximation to the conditional covariance is normally sufficient for reasonable accuracy, and as the third order conditional covariance is very slow to calculate.

Order 1 At order 1, $\xi_t := \varepsilon_t$, thus:

$$\mathbb{E}_t \xi_{t+i} \xi'_{t+j} = \begin{cases} \Sigma & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

Order 2 At order 2:

$$\xi_t := \begin{bmatrix} \varepsilon_t \\ \varepsilon_t \otimes \varepsilon_t - \text{vec } \Sigma \\ \tilde{x}_{t-1}^{(1)} \otimes \varepsilon_t \end{bmatrix},$$

thus:

$$\mathbb{E}_t \zeta_{t+i} \zeta'_{t+j} = 0 \text{ if } i \neq j,$$

and by theorem 4.3 of Magnus and Neudecker (1979):

$$\begin{aligned} \mathbb{E}_s \zeta_t \zeta'_t &= \mathbb{E}_s \begin{bmatrix} 1 \otimes \varepsilon_t \\ \varepsilon_t \otimes \varepsilon_t - \text{vec } \Sigma \\ \tilde{x}_{t-1}^{(1)} \otimes \varepsilon_t \end{bmatrix} \begin{bmatrix} 1 \otimes \varepsilon'_t & \varepsilon'_t \otimes \varepsilon'_t - (\text{vec } \Sigma)' & \tilde{x}_{t-1}^{(1)'} \otimes \varepsilon'_t \end{bmatrix} \\ &= \begin{bmatrix} \Sigma & 0 & \mathbb{E}_s \tilde{x}_{t-1}^{(1)'} \otimes \Sigma \\ 0 & (I_{m^2} + K_{m,m})(\Sigma \otimes \Sigma) & 0 \\ \mathbb{E}_s \tilde{x}_{t-1}^{(1)} \otimes \Sigma & 0 & \mathbb{E}_s (\tilde{x}_{t-1}^{(1)} \tilde{x}_{t-1}^{(1)'} \otimes \Sigma) \end{bmatrix}. \end{aligned}$$

O. Cubature methods

Degree 3 monomial rule The equal weight degree 3 monomial cubature rule with $2\hat{S} + 1$ nodes rule exactly integrates all degree 3 monomials in the components of ν . While a third order approximation to the cumulated news shocks, y , as a function of ζ may do a poor job at capturing this highly non-linear (and even non-differentiable) mapping, in practice the approximation to the integral is often surprisingly accurate. This is in a large part due to the robustness of the integration rule which stems from its equal, positive weights. All known higher degree integration rules that do not use more than polynomial in \hat{S} nodes also feature negative weights on at least some nodes (Cools 2003), which means that their result is not guaranteed to lie within the convex hull of the source evaluations, and, in this case in which we are integrating a positive function (y), it further means the result can have the wrong sign.

Genz and Keister (1996) rules The Genz and Keister (1996) rules allow one to choose the maximum degree of monomial that should be integrated exactly, up to a maximum order of 51. The number of points used is $O(\hat{S}^K)$, where $2K + 1$ is the degree of monomial that is integrated exactly. When $K > 0$ and $\hat{S} > 1$, the rule features negative weights on at least one node, which means it is susceptible to the problems mentioned above. However, it has a few points in its favour. Firstly, by using negative weights, the rule is able to ensure that the maximum over the absolute vectors of integration points is independent of \hat{S} . This contrasts with the aforementioned rule in which the higher is \hat{S} , the further into the tails of the distribution one has to evaluate the integrand. Given the extreme non-linearity of the integrand, evaluating far into the tails can lead the equal weighted integration rule to produce a heavily upwards biased estimate of the integral. Secondly, by using a higher degree rule, we can generally obtain a better approximation to the integrand, despite its non-differentiability. Finally, the Genz and Keister (1996) rules are nested, which means that we can use an adaptive integration degree without wasting evaluations, continuing to increase the degree until approximate convergence.

In practice, the results of these rules often repeatedly flip from biased down to biased up as the degree increases, due to the discontinuities. To lessen this, DynareOBC gives the option of averaging integral estimates of adjacent orders, which still integrates polynomials of the lower of the two orders exactly.

Quasi Monte Carlo The final integration method we consider is quasi Monte Carlo, generating points from a Sobol sequence (Sobol 1967). Given that the functions we are integrating are absolutely continuous (as they are piecewise polynomial, with a finite number of manifolds of non-differentiability), quasi Monte Carlo with $2^{1+l} - 1$ draws will produce an error that decays as $O\left(\frac{l^S}{2^l}\right)$. With the Sobol sequence, the choice of $2^{1+l} - 1$ integration points for some $l \in \mathbb{N}$ also ensures that the points are exactly mean zero, hopefully lessening overall bias. However, on functions that are well approximated by a polynomial, quasi Monte Carlo will generally require far more evaluations of the integrand for a similar accuracy than the Genz and Keister (1996) rules would. Which dominates in practice will depend on the precise integrand, which in turn will depend on the model and its current state. At times where the bound is either highly likely to bind or highly likely not to bind, whatever future shocks hit, it is likely that the Genz and Keister (1996) rules will dominate, however, at times when the bound is only binding with moderate probability, quasi Monte Carlo's "dumb" approach may give it better performance.

P. Global solution procedure for the capital constrained model

The value function is:

$$\begin{aligned}
 V(K, A) &= \max_{\substack{C, L \text{ s.t.} \\ AK^\alpha L^{1-\alpha} - C \geq \theta K}} \left[\log C - \frac{L^{1+\nu}}{1+\nu} + \beta \mathbb{E}V(AK^\alpha L^{1-\alpha} - C, A^\rho \exp \varepsilon) \right] \\
 &= \max_{\substack{C \text{ s.t.} \\ \left((AK^\alpha)^{1+\nu} \left(\frac{1-\alpha}{C} \right)^{1-\alpha} \right)^{\frac{1}{\nu+\alpha}} - C \geq \theta K}} \left[\log C - \frac{1}{1+\nu} \left(\frac{1-\alpha}{C} AK^\alpha \right)^{\frac{1+\nu}{\nu+\alpha}} \right. \\
 &\quad \left. + \beta \mathbb{E}V \left(\left((AK^\alpha)^{1+\nu} \left(\frac{1-\alpha}{C} \right)^{1-\alpha} \right)^{\frac{1}{\nu+\alpha}} - C, A^\rho \exp \varepsilon \right) \right]
 \end{aligned}$$

where the second line comes from substituting in the labour first order condition.

We construct a uniform rectangular grid on $\log A$ and $\log K$ as follows. For both variables, the grid is centred on the model's non-stochastic steady-state. For productivity, the grid extends to ± 4 times the standard deviation of $\log A_t$, which covers more than 99.99% of its stationary distribution. For capital, the grid extends to ± 16 times the standard deviation of $\log K_t$ in the model without bounds, which also covers (a lot) more than 99.99% of its stationary distribution in the model with bounds.

Furthermore, it is wide enough to cover the area in which the value function is highly curved, as it goes from increasing to decreasing. The grid has 256 points along the productivity axis, and 1024 points along the capital axis, making for a total of 262,144 points.

We use linear interpolation/extrapolation for points off the grid. Then, due to piecewise linearity of the integrand in the value function, and the Gaussianity of ε , integration can be performed exactly. For speed, we precompute the associated weights at each productivity level, so exact integration just requires a dot product between the weights, and the value function interpolated to the future capital level.

We initialize the grid to the exact solution in the absence of bounds, e.g.:

$$V = F + G \log K + H \log A,$$

where:

$$F = \frac{\log(1 - \alpha\beta)}{1 - \beta} + \frac{1 - \alpha}{(1 - \alpha\beta)(1 - \beta)(1 + \nu)} \left[\log\left(\frac{1 - \alpha}{1 - \alpha\beta}\right) - 1 \right] + \frac{\alpha\beta \log \alpha\beta}{(1 - \alpha\beta)(1 - \beta)},$$

$$G = \frac{\alpha}{1 - \alpha\beta},$$

$$H = \frac{1}{(1 - \alpha\beta)(1 - \beta\rho)}.$$

To facilitate solving with the bound, in the first iteration, we set $\theta = 0$, and then we increase θ by 0.005 with each iteration until it gets to 0.99, in order to “homotope” from the solution without the bound to the solution with the bound at the correct level. We then continue with conventional fixed point iterations until the maximum absolute change in the value function over the grid ceases to decrease. We report the penultimate value function, i.e. we discard the final one which was a bigger step away. Within these iterations, at each grid node, we first solve for the C at which the constraint binds exactly, then maximise C over the interval from 0 to the found bound. Both procedures will deliver a result accurate to somewhere (roughly) between 10^{-8} and 10^{-16} .

As one indication of accuracy, the final value function iteration step changed the value function by at most 3.91×10^{-8} at all grid points, and the implied policy function changed by at most 2.60×10^{-7} at all grid points. By way of comparison, when the same algorithm was run on the model without a bound, the algorithm made two steps (the minimum possible), the last of which changed the value function by at most 5.95×10^{-12} and which changed the policy function by at most 2.89×10^{-7} .

In Figure 6 and Figure 7 we plot the value and policy functions in terms of $\log K$ for a variety of productivity levels, including both the highest and lowest productivities on the grid. These illustrate the extent of the departure from the model without bounds.




Figure 6: Value function




Figure 7: Policy function

Q. Further details on the DynareOBC toolkit

Code implementing all of the algorithms discussed here is contained in the author’s “DynareOBC” toolkit which is a suite of MATLAB files designed to augment the abilities of Dynare (Adjemian et al. 2011). The toolkit may be freely downloaded from <http://github.org/tholden/dynareOBC>, and this site also contains complete

documentation for its assorted options.⁴⁴ To use it, one merely has to include a “max”, “min” or “abs” in the MOD file describing the DSGE model to be simulated, and then to invoke DynareOBC with the MATLAB command “dynareOBC ModFileName.MOD”.

Internally, DynareOBC uses the “YALMIP” (Löfberg 2004) MATLAB toolkit as an interface to a wide variety of open source and commercial mixed integer linear programming solvers. The distribution of DynareOBC comes with a variety of open source solvers, so DynareOBC is certainly not dependent on any particular commercial packages (other than MATLAB itself). DynareOBC also attempts to obtain a parametric solution to the LCP (q, M) for q which only violate the bound in at most the first few periods, using the MPT toolkit (Herceg et al. 2013), which in turn uses an algorithm due to Jones and Morrari (2006). The resulting parametric solution takes the form of a compiled MEX function, which, when passed a q , returns the y that solves the LCP. This reduces the number of times the LCP needs to be solved in inner loops, increasing performance. Furthermore, DynareOBC includes efficient code for testing whether M is a P-matrix, based on an algorithm of Tsatsomeros and Li (2000), and can also test if M is an S-matrix or (strictly) semi-monotone. Additionally, DynareOBC contains code for facilitating the calculation of Jin and Judd (2002) style accuracy checks. Thus, DynareOBC functions as an easy to use, one stop shop for all queries one might have of a model with OBCs.

⁴⁴ A PDF of the toolkit’s documentation is available from: <https://github.com/tholden/dynareOBC/raw/master/ReadMe.pdf>.