

OccBin: A Toolkit for Solving Dynamic Models With Occasionally Binding Constraints Easily*

Luca Guerrieri[†]
Federal Reserve Board

Matteo Iacoviello[‡]
Federal Reserve Board

January 29, 2014

Abstract

We describe how to adapt a first-order perturbation approach and apply it in a piecewise fashion to handle occasionally binding constraints in dynamic models. Our examples include a real business cycle model with a constraint on the level of investment, a New Keynesian model subject to the zero lower bound on nominal interest rates, and a model of optimal consumption choice in the presence of liquidity constraints. In each case, we compare the piecewise linear perturbation solution with a high-quality numerical solution that can be taken to be virtually exact. The piecewise linear perturbation method can adequately capture key properties of the models we consider. A key advantage of this method is its applicability to models with a large number of state variables.

KEYWORDS: occasionally binding constraints, DSGE models, regime shifts, first-order perturbation.

JEL CLASSIFICATION: C61, C63

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. The library of routines that accompanies this paper as well as additional documentation are available at the following website: <http://www2.bc.edu/matteo-iacoviello/research.htm>.

[†]Luca Guerrieri, Office of Financial Stability, Federal Reserve Board, 20th and C St. NW, Washington, DC 20551. Email: luca.guerrieri@frb.gov.

[‡]Matteo Iacoviello, Division of International Finance, Federal Reserve Board, 20th and C St. NW, Washington, DC 20551. Email: matteo.iacoviello@frb.gov.

1. Introduction

Inequality constraints that bind occasionally arise in a wide array of economic applications. We describe how to adapt a first-order perturbation approach and apply it in a piecewise fashion to handle occasionally binding constraints. To showcase the wide applicability of our approach, we solve an array of popular dynamic stochastic models. The first model is an RBC model with limitations on the mobility of factors of production. The second model is a canonical New Keynesian model subject to the zero lower bound on nominal interest rates. The third is a model of optimal consumption choice in the presence of liquidity constraints. As is typical for dynamic models, the models we consider do not have a closed-form analytical solution. In each case, we compare the piecewise linear perturbation solution with a high-quality numerical solution that can be taken to be virtually exact.

Our contribution is twofold. First, we outline an algorithm to obtain a piecewise linear solution. While the individual elements of the algorithm are not original, our recombination simplifies the application of this type of solution to a general class of models.¹ We offer a library of numerical routines, *OccBin*, that implements the algorithm and is compatible with *Dynare*, a convenient and popular modeling environment ([Adjemian et al. 2011](#)). Second, we present a systematic assessment of the quality of the piecewise linear perturbation method relative to a virtually exact solution, which has not been attempted by others.²

Because standard perturbation methods only provide a local approximation, they cannot capture occasionally binding constraints without adaptation. Our analysis builds on an insight that has been used extensively in the literature on the effects of attaining the zero-lower bound on nominal interest rates.³ That insight is that occasionally binding constraints can be handled as different regimes of the same model. Under one regime, the occasionally binding constraint is slack. Under the other regime, the same constraint is binding. The piecewise linear solution method involves linking the first-order approximation of the model around the same point under each regime. Importantly, the solution that

¹ Our approach, including the title of this paper, is inspired by the work by [Uhlig \(1995\)](#) who developed an early toolkit to analyze nonlinear dynamic discrete-time stochastic models without occasionally binding constraints.

² The virtually exact solution is obtained either by dynamic programming on a very fine lattice for the state variables of the model or by spectral methods, following [Christiano and Fisher \(2000\)](#).

³ Recent examples of the use of this technique include [Jung, Teranishi, and Watanabe \(2005\)](#), [Eggertsson and Woodford \(2003\)](#), [Christiano, Eichenbaum, and Rebelo \(2011\)](#).

the algorithm produces is not just linear – with two different sets of coefficients depending on whether the occasionally binding constraint is binding or not – but rather, it can be highly nonlinear. The dynamics in one of the two regimes may crucially depend on how long one expects to be in that regime. In turn, how long one expects to be in that regime depends on the state vector. This interaction produces the high nonlinearity.

Our assessment focuses on several aspects of the solution. Following [Christiano and Fisher \(2000\)](#), we compare moments of key variables by reporting mean, standard deviation, and skewness. Following [Taylor and Uhlig \(1990\)](#), we compare plots of stochastic simulations. In addition, we assess the accuracy of the piecewise linear approximation by computing two bounded rationality metrics. The first metric is the Euler equation residual, following [Judd \(1992\)](#). The Euler equation residual quantifies the error in the intertemporal allocation problem using units of consumption. The second metric relies on the broader evaluation of expected utility. Intuitively, the closest approximation to the solution of the model will lead to the highest utility level. The difference in utility implied by two solution methods can also be expressed as a compensating variation in consumption that a utility-maximizing agent would have to be offered in order to continue using the less accurate method. On the basis of these comparisons and assessments, we find that the piecewise linear perturbation method can capture adequately key properties of the models we consider.

We also highlight some limitations of the piecewise linear solution. Namely, just like any linear solution, it discards all information regarding the realization of future shocks. Accordingly, our piecewise linear approach is not able to capture precautionary behavior linked to the possibility that a constraint may become binding in the future, as a result of shocks yet unrealized. However, the piecewise method also inherits some of the key advantages of a first-order perturbation approach. It is computationally fast and applicable to models with a large number of state variables even when the curse of dimensionality renders other higher-quality methods inapplicable.⁴ Moreover, our library of numerical routines accepts a model written in a natural way with no meaningful syntax restrictions.

⁴ The library of routines that accompanies this paper contains additional examples of models that can be solved with a piecewise linear algorithm. One of the examples is the celebrated [Smets and Wouters \(2007\)](#) model, extended to incorporate the zero lower bound on the policy interest rate. As an illustration of the speed of the piecewise linear algorithm, our toolkit solves that model in a fraction of a second.

Accordingly, application of our algorithm to different models requires only minimal programming.

Section 2 outlines the piecewise linear solution algorithm. Section 3 relates our approach to the literature. Section 4 considers a real business cycle model with a constraint on investment. Section 5 considers a New Keynesian model subject to the zero lower bound on nominal interest rates. Section 6 considers a permanent income model with a borrowing constraint. Section 7 concludes.

2. The Solution Algorithm

For clarity of exposition, we confine our attention to a model with only one occasionally binding constraint. Extensions to multiple occasionally binding constraints are implemented in the library of routines.

A model with an occasionally binding constraint is equivalent to one with two regimes. Under one regime, the occasionally binding constraint is slack. Under the other regime, the constraint binds. We linearize the model under each regime around the non-stochastic steady state, although a different point could be chosen. We dub the regime that applies at the point of linearization the “reference” regime, or (M1) below. We dub the other regime “alternative”, or (M2). It is immaterial whether the occasionally binding constraint is slack at the reference regime or at the alternative regime.

There are two important requirements for the application of our algorithm.

1. The conditions for existence of a rational expectations solution in Blanchard and Kahn (1980) hold at the reference regime.
2. If shocks move the model away from the reference regime to the alternative regime, the model will return to the reference regime in finite time under the assumption that agents expect that no future shocks will occur.⁵

Definition of a piecewise Linear Solution. Without loss of generality, when the occasionally binding constraint $g(E_t X_{t+1}, X_t, X_{t-1}) \leq 0$ is slack, the linearized system of necessary conditions for

⁵ This restriction might appear draconian, but it is routinely imposed when solving DSGE models with standard first-order perturbation methods. In fact, the linear approximation to the solution could be equivalently characterized as implementing either the rational expectations restrictions or perfect foresight.

an equilibrium under the reference regime can be expressed as:

$$\mathcal{A}E_tX_{t+1} + \mathcal{B}X_t + \mathcal{C}X_{t-1} + \mathcal{E}\epsilon_t = 0, \quad (\text{M1})$$

where X is a vector of size n that collects all the endogenous variables; E_t is the expectation operator, conditional on information available at time t ; $\mathcal{A}$, $\mathcal{B}$, and $\mathcal{C}$ are $n \times n$ matrices of structural parameters for the model's linearized equations that are conformable with X ; ϵ is a vector of zero mean, i.i.d. exogenous innovations of size m and $\mathcal{E}$ is an $n \times m$ matrix of structural parameters.

When the constraint binds, then $h(E_tX_{t+1}, X_t, X_{t-1}) > 0$, the analogous system of necessary conditions for an equilibrium under the alternative regime, linearized again around the non-stochastic steady state, can be expressed as:

$$\mathcal{A}^*E_tX_{t+1} + \mathcal{B}^*X_t + \mathcal{C}^*X_{t-1} + \mathcal{D}^* + \mathcal{E}^*\epsilon_t = 0. \quad (\text{M2})$$

The matrices $\mathcal{A}^*$, $\mathcal{B}^*$, $\mathcal{C}^*$ are again $n \times n$ matrices of structural parameters. In addition, under (M2) there is a column vector of parameters $\mathcal{D}^*$ whose size is n . The presence of $\mathcal{D}^*$ arises from the fact that the linearization is carried out around a point (the steady state by our choice) in which regime (M1) applies. Finally $\mathcal{E}^*$ is another $n \times m$ matrix of structural parameters. Notice that the conditions implied by the functions g and h above are assumed to be mutually exclusive and collectively exhaustive. We are now in a position to define a solution for our model.

Definition 1. : *A solution for a model with an occasionally binding constraint is a function $f : X_{t-1} \times \epsilon_t \rightarrow X_t$ such that the conditions under system (M1) or the system (M2) hold, depending on the evaluation of the occasionally binding constraint, governed by g and h .*

An alternative way of characterizing the function f relies on matrix expressions which closely mirror the familiar decision rules of a linearized dynamic model. Accordingly, given initial conditions X_0 and the realization of a shock ϵ_1 , the function f can be expressed as a set of matrices $\mathcal{P}_t$, a set of matrices

$\mathcal{R}_t$, and a matrix $\mathcal{Q}_1$, such that:

$$X_t = \mathcal{P}_t X_{t-1} + \mathcal{Q}_1 \epsilon_1 + \mathcal{R}_t, \quad \text{for } t = 1 \quad (1)$$

$$X_t = \mathcal{P}_t X_{t-1} + \mathcal{R}_t \quad \forall t \in \{2, \infty\}. \quad (2)$$

As equations 1 and 2 show, the solution from our piecewise algorithm need not be linear, even if the original system described by (M1) and (M2) is. At each point in time the matrices $\mathcal{P}_t$, $\mathcal{Q}_t$, $\mathcal{R}_t$ are time varying, even if they are functions of X_{t-1} and ϵ_1 only.

The Solution Algorithm. Given the conditions for an equilibrium in $M1$ and $M2$, and given the occasionally binding constraint expressed in g and h , the following algorithm characterizes the piecewise linear solution f , defined above. The output of the algorithm is a time varying decision rule whose general form is given in Equations (1) and (2). Accordingly, the algorithm shows how to compute the matrices $\mathcal{P}_t$, $\mathcal{Q}_t$, and $\mathcal{R}_t$ given initial conditions X_0 and given the realization of a shock ϵ_1 .

The algorithm employs a guess-and-verify approach. First, we guess the periods in which each regime applies. Second, we proceed to verify and, if necessary, update the initial guess as follows:

1. Let T be the date when the current guess implies that the model will return to regime (M1).

Then for any $t \geq T$, using standard perturbation methods, one can characterize the linear approximation to the decision rule for X_t , given X_{t-1} , as:

$$X_t = \mathcal{P} X_{t-1} + \mathcal{Q} \epsilon_t, \quad (\text{M1DR})$$

where $\mathcal{P}$ and $\mathcal{Q}$ are $n \times n$ and $n \times m$ matrices of reduced-form parameters, respectively. Then, using the notation of Equation 2, for any $t \geq T$, $\mathcal{P}_t = \mathcal{P}$, $\mathcal{R}_t = 0$.

2. Using $X_T = \mathcal{P} X_{T-1}$ and Equation (M2), coupled with the assumption that agents expect no shocks beyond the first period, the solution in period $T - 1$ will satisfy the following matrix equation:

$$\mathcal{A}^* \mathcal{P} X_{T-1} + \mathcal{B}^* X_{T-1} + \mathcal{C}^* X_{T-2} + \mathcal{D}^* = 0. \quad (3)$$

Solve the equation above for X_{T-1} to obtain the decision rule for X_{T-1} , given X_{T-2} :

$$X_{T-1} = -(\mathcal{A}^*\mathcal{P} + \mathcal{B}^*)^{-1}(\mathcal{C}^*X_{T-2} + \mathcal{D}^*). \quad (4)$$

Accordingly, $\mathcal{P}_{T-1} = -(\mathcal{A}^*\mathcal{P} + \mathcal{B}^*)^{-1}\mathcal{C}^*$ and $\mathcal{R}_{T-1} = -(\mathcal{A}^*\mathcal{P} + \mathcal{B}^*)^{-1}\mathcal{D}^*$

3. Using $X_{T-1} = \mathcal{P}_{T-1}X_{T-2} + \mathcal{R}_{T-1}$ and either (M1) or (M2), as implied by the current guess of regimes, solve for X_{T-2} given X_{T-3} .
4. Iterate back in this fashion until X_0 is reached, applying either (M1) or (M2) at each iteration, as implied by the current guess of regimes.
5. Depending on whether regime (M1) or (M2) is guessed to apply in period 1, $\mathcal{Q}_1 = -(\mathcal{A}\mathcal{P}_2 + \mathcal{B})^{-1}\mathcal{E}$, or $\mathcal{Q}_1 = -(\mathcal{A}^*\mathcal{P}_2 + \mathcal{B}^*)^{-1}\mathcal{E}^*$. Trivially, in the special case in which regime (M1) is guessed to apply in all periods, one can see that $\mathcal{Q}_1 = \mathcal{Q}$, consistent with equation (M1DR).
6. Using the guess for the solution obtained in steps 1 to 5, compute paths for X to verify the current guess of regimes. If the guess is verified, stop. Otherwise, update the guess for when regimes (M1) and (M2) apply and return to step 1.

Given X_0 and ϵ_1 , an expedient initial guess of regimes can be obtained by applying the standard first-order perturbation solution to (M1). In general, the guess will have to be updated, because a switch in regimes is associated with a change in the paths of the endogenous variables. A choice for the updating scheme in step 6 that we have found resilient in practice is to use the path for X from the previous iteration to infer a new guess of regimes. As an alternative, one may choose to dampen the iterations by shrinking (or expanding) the number of periods when a certain regime applies only gradually, in a fashion analogous to the Gauss-Jacobi algorithm.

Computation of the solution requires a series of inversions for the matrix $\mathcal{J}_t \equiv (\mathcal{A}^*\mathcal{P}_t + \mathcal{B}^*)$, for $t = 2$ to $t = T$. Contingent on a guess for a sequence of regimes, non-invertibility of the matrix $\mathcal{J}_t$ implies the existence of multiple paths that lead back from the point X_T to the point X_0 . In that case, the application of a pseudo-inverse, as suggested by [Chen, Cúrdia, and Ferrero \(2012\)](#), arbitrarily

selects one of these paths.

Notice that for multiple solutions to exist, non-invertibility of $\mathcal{J}_t$ is neither sufficient nor necessary. It is not sufficient because one needs to also verify that the given guess of regimes is consistent with the occasionally binding constraint of interest after calculation of a full path for X . It is not necessary because distinct sequences of regimes may support multiple solutions.

Implementation of the Algorithm. The paper is complemented by a library of numerical routines, *OccBin*, that implements the piecewise linear solution algorithm using the Matlab programming language. Our routines are designed as an add-on to Dynare, a widely used set of programs for the solution and estimation of DSGE models. Dynare lets users specify a DSGE model using a readable syntax that imposes only trivial requirements in the way a model is specified. The programs we devised take as inputs two Dynare model files. One file specifies the model ($M1$) at the reference regime. The other file specifies the model at the alternative regime ($M2$). We use the analytical derivatives computed by Dynare to construct $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{E}, \mathcal{A}^*, \mathcal{B}^*, \mathcal{C}^*, \mathcal{D}^*$, and $\mathcal{E}^*$ in equations ($M1$) and ($M2$). We also use the routines in Dynare to construct the matrices $\mathcal{P}$ and $\mathcal{Q}$ in equation ($M1DR$).

3. Related Approaches

A recent review of solution algorithms that mitigate the curse of dimensionality, is provided by [den Haan, Judd, and Juillard \(2011\)](#) and references therein. [Judd, Maliar, and Maliar \(2012\)](#) extend that review by considering methods appropriate for the solution of models with occasionally binding constraints. They do not dwell on piece-wise linear solutions. We do not attempt to replicate a comprehensive review of the literature but focus on the connection between the piecewise linear algorithm presented above and alternative algorithms that ameliorate the curse of dimensionality.

The idea of concatenating decision rules from multiple regimes starting from the last period in which the reference regime is expected to apply in perpetuity can be traced back to [Jung, Teranishi, and Watanabe \(2005\)](#). Their focus is on an economy subject to the zero lower bound on nominal interest rate. Our formulation of the piecewise linear solution algorithm applies to any linear model under a general form for the specification of the occasionally binding constraints.

One extension of the basic piecewise linear solution in [Jung, Teranishi, and Watanabe \(2005\)](#) is due to [Eggertsson and Woodford \(2003\)](#). They consider a model with a shock to the natural rate of interest subject to a Markov process with only two states. In one state, the natural rate is so low that the zero lower bound binds. Under this stark stochastic structure, they compute a rational expectation solution, instead of a perfect-foresight solution. However, in their model the expected duration of the alternative regime, i.e. the policy rate at the lower bound, is always fixed, as determined by the choice of Markov process. By contrast, in our setup the duration of the alternative regime is dependent on the realization of shocks. In turn, the expectation of how long a regime is expected to last affects the value of the endogenous variables contemporaneously.

Building on the work of [Laséen and Svensson \(2009\)](#), [Holden and Paetz \(2012\)](#) provide a solution method that allows for occasionally binding constraints based on introducing anticipated shocks. With first-order perturbation, their method would produce paths for the endogenous variables identical to the ones of our piece-wise approach.⁶ The choice of anticipated shocks that mimic occasionally binding constraints is specific to each model and is not amenable to a general specification, such as the one achieved for our algorithm.

Upon linearization of the model, an extended path algorithm, as the one proposed by [Fair and Taylor \(1983\)](#) and further developed by [Adjemian and Juillard \(2011\)](#), would also yield the same path for the endogenous variables as our piecewise linear algorithm. One advantage of the extended path algorithm is that it can also handle nonlinear perfect-foresight models, avoiding linearization altogether. However, in practice, convergence of the algorithm may be difficult without a high-quality initial guess. An advantage of our piecewise linear method is that it greatly simplifies the search process. Instead of searching for the paths of all the endogenous variables, the piecewise linear algorithm only needs to search for a sequence of regimes.

The extended path algorithm relies on derivative-based methods to search for a solution. This search is complicated by the fact that occasionally-binding constraints introduce a discontinuity in the derivatives of the conditions for an equilibrium. Substitution of the kink implied by the occasionally

⁶ An explanation for the equivalence of the two approaches and a discussion of their relative merits is provided in the appendices of [Bodenstein, Guerrieri, and Erceg \(2009\)](#) and of [Bodenstein, Guerrieri, and Gust \(2013\)](#).

binding constraint with a smooth polynomial approximation may yield a reformulation of the model more easily amenable to derivative-based solution methods. Our attempts at pursuing this strategy revealed undesirable side effects. As we increased the order of the polynomial to get a better approximation to the kink implied by an occasionally binding constraint, the polynomial generated wild oscillations when moving away from the area immediately surrounding the kink.

An alternative way of masking the discontinuity implied by occasionally binding constraints is offered by [McGrattan \(1996\)](#), [Preston and Roca \(2007\)](#), and [Kim, Kollmann, and Kim \(2010\)](#). The insight is to penalize agents' utility when a particular constraint is hit. While this method has the advantage of converting a model with occasionally binding constraints into a model that is solvable by perturbation methods, it suffers from undesirable drawbacks. First, the solution will change with the size and the shape of the penalty (the barrier parameter). Moreover, any high-order perturbation method will generate a smooth solution that in some instances will violate the inequality constraint.

The remarkable recent work of [Judd, Maliar, and Maliar \(2012\)](#) also provides a solution algorithm that can handle both a sizable number of state variables and occasionally binding constraints. Their innovation is to use a simulation-based approach to construct the approximation grid for projection methods, which ameliorates the curse of dimensionality. However, the computational burden of this method may remain too high for models oriented towards empirical realism. For instance, [Judd, Maliar, and Maliar \(2012\)](#) highlight that a simplified version of the Smets-Wouters model with an added zero lower bound constraint can be solved in 25 minutes (with serial processing in Matlab).

4. A Model with a Constraint on Investment

For its simplicity and widespread use, the RBC model is a staple of the literature that has compared the performance of different solution techniques (see for instance, [Taylor and Uhlig 1990](#)). In our variant of this canonical model, the choice of investment is subject to an occasionally binding constraint. This constraint prevents investment from falling below an exogenously fixed lower bound in every period. This exogenous lower bound could be set to imply that investment cannot be negative. Accordingly, our model nests a model in which capital is irreversible.

Model Overview. A central planner maximizes households' utility

$$\max E_0 \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma} - 1}{1-\gamma},$$

subject to the constraints in Equations (5) to (7) below:

$$C_t + I_t = A_t K_{t-1}^{\alpha}, \quad (5)$$

$$K_t = (1 - \delta) K_{t-1} + I_t, \quad (6)$$

$$I_t \geq \phi I_{SS}. \quad (7)$$

The planner chooses consumption, C_t , investment, I_t , and capital, K_t . Equation (5) is the resource constraint and $A_t K_{t-1}^{\alpha}$ is the economy's output in period t . Technology A_t evolves according to

$$\ln A_t = \rho \ln A_{t-1} + \sigma \epsilon_t, \quad (8)$$

where ρ and σ are parameters and ϵ_t is an exogenous innovation distributed as standard normal. Equation (6) is the capital accumulation equation, with depreciation rate δ . Finally, Equation (7) is an occasionally binding constraint that prevents investment from falling below a fraction ϕ of investment in the non-stochastic steady state, denoted by I_{SS} . When the parameter ϕ equals 0, this last constraint implies that capital is irreversible. In the numerical experiments below, we set ϕ at a value well above zero which ensures that the constraint binds frequently.

Denoting with λ_t the Lagrange multiplier on the investment constraint given by (7), the equations describing the necessary conditions for an equilibrium are (5), (6), (8) together with the consumption Euler equation and the Kuhn-Tucker condition for the investment constraint:

$$C_t^{-\gamma} - \lambda_t = \beta E_t \left(C_{t+1}^{-\gamma} (1 - \delta + \alpha A_{t+1} K_t^{\alpha-1}) - (1 - \delta) \lambda_t \right) \quad (9)$$

$$\lambda_t (I_t - \phi I_{SS}) = 0. \quad (10)$$

These equations form a dynamic system of five equations in the five variables $\{C_t, I_t, K_t, A_t, \lambda_t\}$.

Mapping these conditions into the notation used in Section 2, in this case (M1) and (M2) only differ because of one equation. The complementary slackness condition for the optimization problem implies that $\lambda_t = 0$ when the constraint is slack. Conversely, when the constraint binds, $I_t = \phi I_{SS}$. The conditions in (M1) encompass $\lambda_t = 0$ and the function g captures $I_t \geq \phi I_{SS}$. The conditions in (M2) encompass $I_t = \phi I_{SS}$, and the function h captures $\lambda_t > 0$.

Calibration and Policy Functions. Table 1 summarizes the calibration, which reflects a choice of yearly frequency. Most parameter choices are standard. The risk aversion parameter γ is set to 2: we discuss sensitivity to alternative choices below. We set $\phi = 0.975$, which implies that the constraint binds about 40% of the time. We set $\alpha = 0.33$, $\delta = 0.1$, and $\beta = 0.96$. Finally, we set $\sigma = 0.013$ and $\rho = 0.9$, this parameter choices imply a standard deviation of log output around 4 percent.

In the absence of an analytical closed-form solution for the model, we use alternatively projection methods and dynamic programming to characterize a high-quality, fully-nonlinear solution.⁷ The resulting investment function of the full nonlinear solution is shown in the top panel of Figure 1. Regardless of the initial level of capital, low realizations of technology trigger investment (in deviation from its steady state) to hit its lower bound given by $-(1 - \phi)$. The bottom panel compares the nonlinear solution to the piecewise solution obtained using our method. Given our benchmark calibration, investment is slightly lower under the piecewise solution when the irreversibility constraint does not bind. The higher level of investment in the nonlinear solution comes from the effect of uncertainty on precautionary saving. However, the policy functions are remarkably close.

Assessing Performance: Impulse Responses and Moments. The transitional dynamics are illustrated in Figure 2, which shows the model response of the variables to two shocks. The first shock occurs in period 10, and brings down the level of technology by 4 percent (close to a 3 standard deviation shock). The second shock occurs in period 60 and pushes up the level of technology by 4 percent. The solid lines in the figure denote the piecewise linear solution. The dashed lines denote the dynamic programming solution. The dash-dotted lines denote the first-order perturbation solution.

⁷ A detailed description of the algorithms for the fully-nonlinear solution is given in Appendix A.

The decline in productivity leads to a decline in investment large enough for the investment constraint to bind. The responses obtained from the piecewise linear and the full nonlinear are strikingly close. As investment cannot fall more than 2.5 percent relative to its steady-state value, the drop in consumption is exacerbated relative to a model without an investment constraint. The first-order perturbation solution ignores the constraint altogether, and the responses from the first-order solution exhibit a markedly smaller contraction in consumption.

When productivity rises, the responses from the three solution methods track each other closely. One difference is that the full-nonlinear solution implies a slightly higher accumulation of capital, in line with precautionary motives stemming from the concavity of the utility function. Neither the piecewise linear nor the linear method can capture such precautionary motives. By period 60, the paths of capital from the piecewise linear solution and from the linear solution essentially coincide. When the state variables coincide, by construction the piecewise linear and the linear solutions also coincide, as long as the constraint does not bind. Accordingly, the effects of an increase in productivity are about the same from these two solution methods.

Table 2 compares key moments.⁸ Overall, the moments from the piecewise linear and the nonlinear solution methods are strikingly close. The piecewise linear method captures first, second, and third moments of the distribution of key variables. In particular, it captures the skewness in the distribution of consumption and investment derived from the occasionally binding constraint, which is missed by the first-order perturbation method. Furthermore, the piecewise linear method matches closely the frequency with which the constraint binds. Both the piecewise linear and fully nonlinear solutions imply that the constraint on investment binds, on average, 41 out of every 100 periods.

Assessing Performance: Euler Residuals. Besides a comparison of moments, it is possible to check the accuracy of the piecewise linear solution in economic terms, using the bounded rationality metric in Judd (1992). Moving from the Euler equation for consumption, we define the Euler equation

⁸ Santos and Peralta-Alva (2005) show that simulated moments from numerical approximations to dynamic stochastic models converge to their exact values as the approximation errors of the solutions converge to zero.

error (expressed as a fraction of units of consumption) as:

$$err_t = \frac{-C_t + \left\{ \lambda_t + E_t \beta \left[C_{t+1}^{-\gamma} ((1 - \delta) + \alpha A_{t+1} K_t^{\alpha-1}) - (1 - \delta) \lambda_{t+1} \right] \right\}^{-\frac{1}{\gamma}}}{C_t}. \quad (11)$$

When sizing the errors for different solution methods, we use the decision rules for capital implied by each method, coupled with the full set of nonlinear constraints implied by Equations (5) to (7).

Figure 3 shows Euler equation errors for different levels of technology and different solution methods. The top panel reports Euler residuals for the piecewise linear method. The middle panel relates to the linear method for the same model without the constraint on investment.⁹ All panels report the absolute value of the Euler residuals expressed in logarithmic scale with base 10. Under that scale, the interpretation of a value of, say, -4 is that the Euler error is sized at \$1 per \$10,000 of consumption. The range in the abscissae was chosen to encompass most of the mass of the ergodic distribution for capital under the baseline calibration.

For the levels of technology shown, the errors in the top panel stay uniformly below -3 and dip well below -4 for part of the range of capital. The Euler errors in the middle panel are consistent with results in [Aruoba, Fernandez-Villaverde, and Rubio-Ramirez \(2006\)](#), who also discuss the performance of the log-linear solution algorithm for the standard RBC model. Strikingly, the Euler residuals for the piecewise linear algorithm used for the top panel remain of a similar order of magnitude as for the first-order perturbation method used for the middle panel. In fact, in the case of “Low technology,” the piecewise linear algorithm even implies smaller solution errors. This finding is not too surprising, since for low values of capital and technology, the piecewise linear decision rule coincides with the fully-nonlinear rule. Given that the solution from the piecewise linear algorithm is nearly coincident with the nonlinear solution for low levels of capital and technology, it might be puzzling why the Euler residuals do not plummet even further. To understand why, it is useful to realize that the Euler residuals as defined in Equation (11) also include expectation terms for higher values of technology. These future terms introduce an additional source of Euler error distinct from misses in the current period’s decision.

⁹ An equivalent interpretation of the middle panel of Figure 3 is that it relates to the piecewise linear solution method for an alternative calibration of the parameter ϕ , so low as to make the constraint on investment irrelevant

The bottom panel of Figure 3 shows Euler residuals for a fully nonlinear collocation method. The figure confirms that Euler errors are of an economic negligible magnitude for the fully nonlinear solution, consistent with results presented in [Christiano and Fisher \(2000\)](#). The contours shown stay well below -6 , dropping to around -14 at the collocation nodes.

Assessing Performance: Welfare. Intuitively, a superior approximation to the solution of the model should yield a higher level of utility regardless of the initial conditions. To express the differences in utility implied by the piecewise linear solution and by the fully nonlinear solution, we focus on the constant proportional increase in consumption, the subsidy rate, that would have to be promised in order to make the representative agent indifferent between using the inferior piecewise linear decision rule instead of the full nonlinear decision rule. Denoting with $C_{NL,t}$ and by $C_{PL,t}$ the consumption levels implied by the nonlinear decision rule and by the piecewise linear rule respectively, we size the accuracy of the piecewise linear decision rule by the subsidy rate τ , where τ is such that:

$$E_t \sum_{t=0}^{\infty} \beta^t \frac{C_{NL,t}^{1-\gamma} - 1}{1-\gamma} = E_t \sum_{t=0}^{\infty} \beta^t \frac{(C_{PL,t}(1+\tau))^{1-\gamma} - 1}{1-\gamma}. \quad (12)$$

We compute the expected utility of the representative agent implied by the decision rules from the full nonlinear solution and from the piecewise linear solution. Both decision rules can be expressed in terms of capital and the level of technology. Then, from the capital accumulation equation, one can back out the level of investment. After enforcing the occasionally-binding investment constraint in Equation (7), one can compute consumption using the resource constraint. We obtain the value function from the decision rule using the Howard improvement algorithm as described in [Ljungqvist and Sargent \(2004\)](#). We find that the value of the subsidy for the baseline parameterization is \$1 per about \$14,500,000 of consumption. Such a small subsidy implies that the piecewise linear approximation works remarkably well. This statement can be put in context by contrasting our approximation with a clearly suboptimal rule that always sets the capital stock to its previous value, so that $K_t = K_{t-1}$. In that case, consumption moves in lockstep with movements in technology, and the welfare cost of not optimizing is orders of magnitude larger, \$1 per about \$3,000 of consumption.

In robustness experiments (expanded on in Appendix C), we also consider sensitivity with respect to the choice of the value for two key parameters, the discount factor β and the risk aversion coefficient γ . The welfare cost of using the piecewise method increases as the discount factor rises, from \$1 per about \$14,500,000 of consumption in the baseline to \$1 per about \$3,000,000 when $\beta=0.98$. We conjecture that in the plain vanilla RBC model the nonlinearities become more pronounced as the risk free rate becomes lower, thus penalizing linearization in general over a fully nonlinear solution algorithm. Moreover, the welfare cost of using our piecewise solution method is a non-monotonic function of risk aversion. In Appendix C we discuss further the intuition for this result, highlighting the subtle, model-specific differences between our solution method and the fully nonlinear one.

5. A New Keynesian Model with the Zero Lower Bound

We consider a textbook version of the New Keynesian model, such as the one described in [Galí \(2008\)](#). For ease of comparison, the notation and calibration hew closely to the version presented in [Fernández-Villaverde et al. \(2012\)](#), who also consider the consequences of attaining the zero lower bound on nominal interest rates using fully nonlinear solution techniques.

In the model, a representative household provides labor (the only input in production) to intermediate firms and consumes. A continuum of intermediate firms that produce differentiated products subject to monopolistic competition adjust their prices according to Calvo-type contracts. The intermediate products are repackaged by competitive final firms. A government sector consumes part of the final good and sets monetary policy according to a Taylor rule subject to the zero lower bound.

Model Overview. A representative household chooses consumption and labor streams C_t and L_t to maximize:

$$\max_{C_t, L_t, B_t} E_0 \sum_{t=0}^{\infty} \left(\prod_{i=0}^t \beta_i \right) \left(\log C_t - \psi \frac{L_t^{1+\vartheta}}{1+\vartheta} \right),$$

where the discount factor β_t follows the process

$$\ln \beta_t = (1 - \rho) \log \beta + \rho \log \beta_{t-1} + \sigma \epsilon_t. \quad (13)$$

The term ϵ_t is an exogenous innovation distributed as standard normal, and σ_b is the standard deviation of the innovation. The budget constraint is given by:

$$C_t + B_t/P_t = w_t L_t + R_{t-1} B_{t-1}/P_t + T_t + F_t. \quad (14)$$

For simplicity, we do not describe the full set of Arrow-Debreu securities available to households in addition to the non-state contingent government bond B_t , which pays the nominal gross interest rate R_t . The price level is P_t . The terms T_t and F_t represent lump-sum taxes and an aliquot share of the profits/losses of intermediate firms.

Competitive final firms repackage intermediate goods Y_{it} to produce a final good Y_t according to $Y_t = \left(\int_0^1 Y_{it}^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}$. Profit maximization yields the demand schedule $Y_{it} = \left(\frac{P_{it}}{P_t} \right)^{-\epsilon} Y_t$ for each intermediate variety, where P_{it} is the price of variety i . Taking the demand from final firms as given, intermediate firms choose their price to maximize profits, subject to Calvo-type restrictions. Each period, a fraction $1 - \theta$ of firms is selected to re-optimize its price (while all other firms keep the old price). The firms selected solve:

$$\max_{P_{it}} E_t \sum_{\tau=0}^{\infty} \theta^{\tau} \left(\prod_{i=0}^{\tau} \right) \frac{\lambda_{t+\tau}}{\lambda_t} \left(\frac{P_{it}}{P_{t+\tau}} - mc_{t+\tau} \right) \left(\frac{P_{it}}{P_t} \right)^{-\epsilon} Y_t, \quad (15)$$

where λ_t is the Lagrangian multiplier on the household's budget constraint for period t and mc_t is the real marginal cost of production. Given that the production technology is $Y_{it} = L_{it}$, the term mc_t equals the wage rate w_t .

The government budget is balanced every period ($B_t = 0 \forall t$), and spending is financed by lump sum taxes T_t . Government spending is a constant share of aggregate output, given by $G_t = s_g Y_t$. Monetary policy is implemented as follows:

$$Z_t = R \left(\frac{\Pi_t}{\Pi} \right)^{\phi_{\pi}} \left(\frac{Y_t}{Y} \right)^{\phi_y} \quad (16)$$

$$R_t = \max(Z_t, 1) \quad (17)$$

where Z_t is the notional policy rate and R_t is the actual policy rate, both expressed in gross terms.

Equation (16) is a Taylor-type rule for setting the interest rate Z_t . Equation (17) is the occasionally binding constraint stating that the actual policy rate cannot fall below 1. Above, Π is the steady-state target level of inflation, R is the steady-state nominal gross return of bonds (equal to Π divided by β), and Y is steady-state output.

For reasons of space, we only emphasize key conditions for an equilibrium. In particular, the conditions that involve intertemporal terms are of special interest because the fully nonlinear collocation solution is obtained by parameterizing the expectations of future variables. For completeness, Appendix B lists all the necessary conditions for an equilibrium and describes the collocation method used to obtain the fully nonlinear solution.

Following Yun (2005), aggregate supply Y_t can be shown to be equal to:

$$Y_t = L_t/v_t \quad (18)$$

where v_t is a measure of price dispersion for intermediate producers equal to $v_t = \int_0^1 \left(\frac{P_{it}}{P_t} \right)^{-\epsilon} di$. In turn, the evolution of dispersion is a state variable given by:

$$v_t = \theta \Pi_t^\epsilon v_{t-1} + (1 - \theta)(\Pi_t^*)^\epsilon. \quad (19)$$

The term Π_t is defined as $\frac{P_t}{P_{t-1}}$, and the term Π_t^* is defined as $\frac{P_t^*}{P_t}$, where P_t^* is the price selected by firms that can re-optimize in period t .

The remaining intertemporal conditions involve expectations: the Euler equation for consumption

$$\frac{1}{C_t} = E_t \left(\beta_t \frac{1}{C_{t+1}} \frac{R_t}{\Pi_{t+1}} \right); \quad (20)$$

and two additional equations that are obtained from the profit maximization problem for intermediate firms and that involve two auxiliary variables x_{1t} and x_{2t}

$$x_{1t} = mc_t Y_t / C_t + \theta E_t \beta_t \Pi_{t+1}^\epsilon x_{1t+1}; \quad (21)$$

$$x_{2t} = \Pi_t^* Y_t / C_t + \theta E_t \beta_t \Pi_{t+1}^{\epsilon-1} \Pi_t^* x_{2t+1} / \Pi_{t+1}^*. \quad (22)$$

To map the conditions for an equilibrium into the notation used in Section 2 one only needs to recognize that, again, (M1) and (M2) only differ because of one equation so as to capture the max operator in Equation (17). The set of conditions under (M1) encompasses $R_t = Z_t$, implying that the actual and notional interest rates coincide. The set of conditions under (M2) encompasses $R_t = 1$, implying a gap between the actual and notional interest rates. The function $h(\cdot)$ is defined as $-R_t + 1$; hence, when the system under (M1) implies that $R_t < 1$, the set of conditions under (M2) applies, ensuring the actual gross interest rate cannot fall below 1. Furthermore, the function $g(\cdot)$ is defined as $-Z_t + 1$; hence, when $Z_t \geq 1$ the set of conditions under (M1) applies, implying that actual and notional rates coincide only when the gross notional rate is at or above 1.

Calibration. The calibration is summarized in Table 3. We follow closely the standard choices in Fernández-Villaverde et al. (2012), but we adapt them to reflect that we simplified the stochastic structure of the model. They consider an array of shocks, while we focus only on shocks to the discount factor β_t , which are often used to bring the model to the zero lower bound in stylized New Keynesian settings, see for instance Christiano, Eichenbaum, and Rebelo (2011). We increase σ , the standard deviation for this shock, from 0.0025 to 0.005 and its persistence ρ at 0.8. With this choice, the model attains the zero lower bound with an empirically realistic frequency in the 5 to 10 percent range. Another departure from the calibration in Fernández-Villaverde et al. (2012) pertains θ , the probability that an intermediate firm will have to keep its price unchanged. To curb the volatility of inflation with a nod to empirical realism, we push this parameter from 0.75 to 0.9. This change, in conjunction with an interest rate rule that responds aggressively to inflation, with a coefficient ϕ_π set at 2.5, prevents large disinflations from occurring at the zero lower bound, in line with recent U.S. experience.

Assessing Performance: Impulse Responses, Moments, and Euler Residuals. Figure 4 shows the response to two unexpected shocks to the discount factor whose size is symmetric around the steady state. The first one occurs in period 6 and brings β up to 1.019 (a positive innovation to the shock process equal to 4 standard deviations). After the first shock has died back down, 25 quarters

later, the second shock hits and brings β_t down to 0.969. The figure compares the response of the model as implied by three solution methods, the piecewise linear method implemented with OccBin, a fully nonlinear collocation solution, and a first order perturbation that disregards the zero lower bound. The responses implied by the piecewise linear and the nonlinear solution lie close to each other, especially for the shock that reduces the discount factor. The differences are more pronounced for the shock that pushes β up. In response to this first shock, the nonlinear solution implies a spell at the zero lower bound lasting 4 periods; the path of the policy rate implied by the piecewise linear solution lifts off one quarter early. The contraction in output and inflation implied by the nonlinear solution when the zero lower bound is attained are both a tad larger relative to the paths for the piecewise linear solution.

The differences highlighted in Figure 4 are also reflected in the comparison of moments in Table 4. Overall, the key moments obtained with the full nonlinear solution method line up well with those from the piecewise linear solution, both at the ZLB and away from it. One notable difference is that, under the collocation solution, the ZLB hits more frequently on average (7 against 4 percent) and the volatility of output is slightly larger.

There are two main forces shaping the differences between the solution produced by OccBin and the nonlinear solution: an uncertainty effect, and a price dispersion effect. These effects can reinforce or offset each other in ways that depend on the calibration, on the size of the shock considered, and on the linearization point.

The uncertainty effect implies that negative shocks at the ZLB produce larger contractions than away from it, since monetary policy is unable to offset them. In turn, the expectation of negative shocks when already at the ZLB further reduces prices and output, since agents expect that monetary policy will be unable to accommodate these shocks. As a consequence, when uncertainty is explicitly taken into account, the ZLB hits more frequently, policy is more accommodative, and output is more skewed to the left. The uncertainty effects imply larger responses at the ZLB than captured by the piecewise linear solution method, which ignores uncertainty. This effect has been highlighted by Nakata (2013), among others.

Even when uncertainty is ruled out, the piecewise solution may overstate or understate the response of price dispersion due to nonlinearities. These nonlinearities can be important especially for large

shocks that take output close to the ZLB as highlighted by [Braun and Waki \(2010\)](#). In our application, the size and direction of the misses are a function of the size of the shock, since price dispersion is a U-shaped function of inflation.¹⁰ As implied by equation 18, aggregate supply is negatively related to dispersion, as high dispersion implies that a few firms, stuck with lower prices, are inefficiently capturing a disproportionate fraction of aggregate demand. In the example of Figure 4, price dispersion drops more in the nonlinear solution, temporarily reducing the inefficiency, and cushioning the drop in output relative to the piecewise linear solution. This effect partly offsets the miss related to the uncertainty effect.¹¹

Given that the uncertainty and price dispersion effects may offset or reinforce each other, it is especially important to assess the performance of the piecewise linear method for different calibrations, as well as across a wide range of values for dispersion and for different values of the shock process β_t . Figure 5 shows Euler equation errors for the baseline calibration, expressed as a share of consumption as for the previous model considered.¹² We find it remarkable that, at worst, the Euler errors stay close to \$1 per \$1000 of consumption for extreme values of dispersion even when the zero lower bound binds. This magnitude is similar to the misses for the linear method in a model that disregards the zero lower bound – as shown in the middle panel of the figure – the workhorse model and solution method for a vast slice of the monetary policy literature. The bottom panel confirms that the Euler errors are of an economic negligible magnitude for the fully nonlinear solution.

Additional robustness checks can be found in Table B in the Appendix, where we compare key moments for different calibrations of the model that focus on varying the monetary policy rule and steady-state inflation. We find that the piecewise linear model continues to perform adequately: if anything, across experiments, it tends to always underestimate the frequency of ZLB episodes and the volatility of output.

¹⁰ For the linear and piecewise linear solutions considered, we linearize around a non-zero inflation rate. Perturbation solutions around a zero inflation point would imply that price dispersion is constant. In that special case, dispersion entirely drops out of the linearized set of conditions for an equilibrium.

¹¹ In a general equilibrium setting, the response of dispersion is also influenced by precautionary motives. We confirmed that the piecewise linear solution understates the decrease in dispersion in response to the increase in the discount factor shown in Figure 4 by checking a nonlinear solution under perfect foresight.

¹² For this model, there are three intertemporal errors associated with equations (20), (21), and (22), respectively. The errors for (21) and (22), not shown, have magnitude and patterns similar that of the Euler equation errors.

We conclude our discussion of the New Keynesian model with a word of caution: while the range of parameter values for which we can solve and find a unique solution for the piecewise linear model is extensive, our numerical routines for the fully nonlinear solution encountered convergence problems for very persistent shocks, for low values of the price rigidity, and for monetary policy rules with a small inflation response. We conjecture – as many others have already done – that the New Keynesian model might be afflicted by several pathologies near the zero lower bound that can make the identification of global solutions especially challenging. These issues have been analyzed and discussed by, among others, [Benhabib, Schmitt-Grohe, and Uribe \(2001\)](#), [Braun and Waki \(2010\)](#), [Aruoba and Schorfheide \(2013\)](#). These pathologies seem to afflict the New Keynesian model in particular and are not a general feature of all models with occasionally binding constraints.

6. A Model with a Borrowing Constraint

Occasionally binding borrowing constraints arise in a wide variety of models where households can “self-insure” by holding and managing an asset, up to some borrowing limit, that can be used to buffer consumption against adverse shocks. In these models, one can distinguish situations when a household is not constrained in the current period, and the traditional Euler equation for consumption holds; and situations when the household is credit constrained, current consumption is too low relative to next period, and the Euler equation for consumption does not hold. This behavioral asymmetry introduced by borrowing constraints, made popular by [Zeldes \(1989\)](#) and [Deaton \(1992\)](#), can be studied using our solution method.

Model Overview. A consumer maximizes

$$\max E_0 \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma} - 1}{1-\gamma}$$

subject to the budget constraint, and to an (occasionally binding) constraint stating that borrowing B_t cannot exceed a fraction m of income Y_t :

$$C_t + RB_{t-1} = Y_t + B_t, \quad (23)$$

$$B_t \leq mY_t. \quad (24)$$

Above, R denotes the gross interest rate. The discount factor β is assumed to satisfy the restriction that $\beta R < 1$, so that in the deterministic steady state the borrowing constraint is binding. Given initial conditions, the impatient household prefers a consumption path that is falling over time, and attains this path by borrowing today. If income is constant, the household will eventually be borrowing constrained and will roll its debt over forever, and consumption will settle at a level given by income less the steady state debt service.

The log of income follows an AR(1) stochastic process of the form

$$\ln Y_t = \rho \ln Y_{t-1} + \sigma \epsilon_t \quad (25)$$

where ϵ_t is an exogenous innovation distributed as standard normal, and σ its standard deviation.

Denoting with λ_t the Lagrange multiplier on the borrowing constraint given by equation (24), the set of equations describing the system of necessary conditions for an equilibrium is given by a system of four equations in the four unknowns $\{C_t, B_t, \lambda_t, Y_t\}$ which includes equation (23), equation (25), together with the consumption Euler equation and the Kuhn-Tucker conditions, given respectively by

$$C_t^{-\gamma} = \beta RE_t \left(C_{t+1}^{-\gamma} \right) + \lambda_t \quad (26)$$

$$\lambda_t (B_t - mY_t) = 0. \quad (27)$$

The transitional dynamics of this model will depend in an important way on the gap between the discount rate and the interest rate, which can be measured as $g = 1/\beta - R$. In our setup, when the gap is small, the economy can be characterized as switching between two regimes. In the first regime,

more likely to apply when income and assets are relatively low, the borrowing constraint binds. Then, borrowing moves in lockstep with income, and consumption is more volatile than income. In the second regime, more likely to apply when income and assets are relatively high, the borrowing constraint is slack, and current consumption can be high relative to future consumption even if borrowing is below the maximum amount allowed. We focus our attention on this case, since it presents an asymmetry that can be studied using our solution method.¹³

In the reference regime, the borrowing constraint binds, and the multiplier is greater than zero. In the alternative regime, the borrowing constraint is slack, and the multiplier is zero. Mapping these conditions into the notation used in Section 2, (M1) and (M2) differ because of one equation. The optimization problem implies that $B_t = mY_t$ when the borrowing constraint binds. Conversely, when the constraint is slack, the complementary slackness condition implies that $B_t \leq mY_t$ and $\lambda_t = 0$. The conditions in (M1) encompass $B_t = mY_t$, and the function g captures $\lambda_t > 0$. The conditions in (M2) encompass $\lambda_t = 0$, and the function h captures $B_t \leq mY_t$.

Calibration and Policy Functions. Table 5 summarizes the baseline calibration, which reflects a yearly frequency. We set $\gamma = 1$, so that utility is logarithmic in consumption. We set the maximum borrowing at one year of income, so that $m = 1$. For the income process, we set $\rho = 0.90$ and $\sigma = 0.0131$, so that the standard deviation of $\ln Y$ is 3 percent. Finally, we set $R = 1.05$ and $\beta = 0.945$. Under this calibration, the borrowing constraint, which binds in the reference regime, is slack about 30 percent of the time using the full nonlinear solution.

We use dynamic programming to characterize a high-quality fully-nonlinear solution. We display the policy functions in terms of the optimal consumption chosen by the agent as a function of income and debt, the two state variables of the problem. The top panel of Figure 6 shows that for lower-than-average realizations of income (and high initial debt) the agent hits the borrowing constraint, the consumption function is relatively steep, and consumption is very sensitive to changes in income.

¹³ Depending on the calibration, other types of solutions may arise. If the discount rate is high and the gap g is large, the borrowing constraint may bind always. Moreover, if the variance of the income process is sufficiently high and the gap approaches zero, consumption may not converge to any finite limit. Even when consumption converges, the stochastic steady state may be drastically different from the deterministic one because the household can accumulate enough assets so that the borrowing constraint is never a concern. Our calibration rules these possibilities out.

For higher-than-average income, consumption is sufficiently high today relative to the future that it pays off to save for the future: the borrowing constraint becomes temporarily slack, and consumption becomes less sensitive to changes in income. The bottom panel compares the policy function obtained via dynamic programming with that obtained using our piecewise approach. The two policy functions are very similar. The only slight difference is that – for given level of initial debt – the anticipation of future shocks leads the agent to save more (consume less) at all income levels.

Assessing Performance: Impulse Responses, Moments and Welfare. Figure 7 shows the responses to two shocks, starting from a nonstochastic steady state where income is one and the ratio of debt to income is $m = 1$, the maximum limit. The first shock, in period 2, brings up income by 3 percent (a 2 standard deviations shock). The second shock, in period 21, pushes down income by 3 percent. The solid and dashed lines denote the piecewise linear solution and the dynamic programming solution respectively. The dash-dotted lines denote the first-order perturbation solution, which incorrectly assumes that the borrowing constraint always binds. As the figure shows, the piecewise linear algorithm well captures the asymmetric responses of consumption, debt, and debt-to-income ratio following income shocks. A positive income shock makes the borrowing constraint slack; borrowing rises less than income, and consumption rises less than it would were the constraint binding in all states of the world. Conversely, when income drops, the borrowing constraint binds, borrowing falls in proportion with income, and consumption reacts more than under a positive shock.

Table 6 shows that the moments computed from the piecewise linear and the nonlinear solution method are again strikingly close. The OccBin can capture first, second, and third moments of the distribution of consumption. In particular, it captures the skewness in consumption derived from the occasionally binding constraint, which is missed by the first-order perturbation method.¹⁴ Furthermore, the piecewise linear method comes close to replicating the frequency with which the constraint binds. Under the piecewise linear solution, the borrowing constraint binds 84 percent of the time. Under the fully nonlinear solution, the constraint binds slightly less frequently, 73 times out of 100 periods. The difference reflects the precautionary behavior induced by the anticipation of future shocks which

¹⁴ The linearized model exhibits some small amount of skewness in consumption simply because we write the model in linearized form, but the income shocks are log-linear.

implies higher average saving under the full nonlinear solution.

The differences between the piecewise linear and the full nonlinear solution for this model again highlight aspects of the economic problem that the piecewise method cannot capture. For this particular model, higher income uncertainty, reduced attitudes toward borrowing (caused by higher discount factor or higher risk aversion), and a looser borrowing limit can magnify the differences between the piecewise solution and the global solution. However, in all these cases the piecewise linear solution still performs uniformly better than the solution where the borrowing constraint is assumed to be always binding.

Relative to the full nonlinear solution, the utility cost of using the piecewise linear method is positive, but small.¹⁵ For the baseline calibration, the household suffers a utility cost of \$1 every \$150,000 of consumption. The cost would be five times larger using a policy function based on first-order perturbation, assuming that the constraint is always binding. In other experiments, we find that the utility cost becomes slightly larger with a higher maximum debt-to-income ratio; with higher risk aversion; with higher uncertainty; and with lower impatience. In all these cases, precautionary considerations become somewhat more important, and ignoring them magnifies the differences between the piecewise method and global nonlinear solution. However, even in these cases the improvements afforded by the piecewise method are substantial: compared to the global solution, the welfare cost of using the piecewise method is between five and six times smaller than the cost of using the linearized solution.

7. Conclusion

We presented a simple piece-wise linear solution method that allows to solve models with occasionally binding constraints easily. This solution method has three principal advantages: 1) It is applicable to models with a large number of state variables; 2) It is a general-purpose algorithm whose deployment is standardized across models and requires only modest additional programming; 3) The computational burden is small, resulting in a short solution time.

As documented, the deterioration in the solution quality from the adoption of the piecewise linear algorithm may vary model by model. For instance, the piecewise linear algorithm may be ill-suited

¹⁵ See Table C in the Appendix for a summary of the sensitivity analysis.

if precautionary considerations are a crucial element of the model to be solved. We considered three different workhorse models to showcase the wide applicability of the piecewise linear solution. In the RBC model with a constraint on the mobility of capital, the deterioration in the solution quality from disregarding precautionary motives turned out to be negligible. Precautionary motives at the zero lower bound on nominal interest rate caused the piece-wise linear solution to slightly underpredict the frequency at which the lower bound is attained. Even in the case of the permanent income model with a borrowing constraint, the deterioration in performance associated with the use of a piecewise linear approach instead of a fully nonlinear method remained acceptable.

For simple models for which more accurate solution methods are viable, the piecewise linear algorithm can provide an initial guess and a useful “sanity check.” The deployment costs of our algorithm are minimal, since it can be implemented generally once and for all. We demonstrated a general implementation in the accompanying library of routines. In our experience, we have found general-purpose algorithms especially useful at the experimentation phase of research, when the model of interest can undergo radical changes requiring, otherwise, costly dedicated programming.

For larger models, the inclusion of empirically-realistic features into a model can quickly strain the performance and applicability of other solution methods that handle occasionally binding constraints. Under those circumstances, our algorithm provides an alternative to swallowing unpalatable simplifications to the model of interest.

References

- Abel, A. B. and J. C. Eberly (1999). The effects of irreversibility and uncertainty on capital accumulation. *Journal of Monetary Economics* 44, 339–377. [\[42\]](#)
- Adjemian, S., H. Bastani, F. Karam, M. Juillard, J. Maih, F. Mihoubi, G. Perendia, M. Ratto, and S. Villemot (2011). Dynare: Reference manual version 4. Dynare Working Papers 1, CEPREMAP. [\[2\]](#)
- Adjemian, S. and M. Juillard (2011). Accuracy of the extended path simulation method in a new keynesian model with zero lower bound on the nominal interest rate. Unpublished manuscript. [\[9\]](#)
- Aruoba, S. B., J. Fernandez-Villaverde, and J. F. Rubio-Ramirez (2006). Comparing solution methods for dynamic equilibrium economies. *Journal of Economic Dynamics and Control* 30(12), 2477–2508. [\[14\]](#)

- Aruoba, S. B. and F. Schorfheide (2013, July). Macroeconomic dynamics near the zlb: A tale of two equilibria. NBER Working Papers 19248, National Bureau of Economic Research, Inc. [22]
- Benhabib, J., S. Schmitt-Grohe, and M. Uribe (2001). The perils of taylor rules. *Journal of Economic Theory* 96(1-2), 40–69. [22]
- Blanchard, O. and C. Kahn (1980). The solution of linear difference models under rational expectations. *Econometrica* 48(5), 1305. [4]
- Bodenstein, M., L. Guerrieri, and C. Erceg (2009). The effects of foreign shocks when interest rates are at zero. International Finance Discussion Paper Series 983, Federal Reserve Board. [9]
- Bodenstein, M., L. Guerrieri, and C. Gust (2013). Oil shocks and the zero bound on nominal interest rates. *Journal of International Money and Finance* 32, 941–967. [9]
- Braun, A. and Y. Waki (2010). On the size of the fiscal multiplier when the nominal interest rate is zero. Mimeo, Federal Reserve Bank of Atlanta. [21, 22]
- Chen, H., V. Cúrdia, and A. Ferrero (2012). The macroeconomic effects of large-scale asset purchase programmes. *Economic Journal* 122(564), F289–F315. [7]
- Christiano, L., M. Eichenbaum, and S. Rebelo (2011). When is the government spending multiplier large? *Journal of Political Economy* 119(1), 78 – 121. [2, 19]
- Christiano, L. J. and J. D. M. Fisher (2000). Algorithms for solving dynamic models with occasionally binding constraints. *Journal of Economic Dynamics and Control* 24(8), 1179–1232. [2, 3, 15, 41]
- Deaton, A. (1992). *Understanding consumption*. Oxford University Press. [22]
- den Haan, W. J., K. L. Judd, and M. Juillard (2011). Computational suite of models with heterogeneous agents ii: Multi-country real business cycle models. *Journal of Economic Dynamics and Control* 35, 175–177. [8]
- den Haan, W. J. and A. Marcet (1990). Solving the stochastic growth model by parameterizing expectations. *Journal of Business & Economic Statistics* 8(1), 31–34. [41]
- Eggertsson, G. B. and M. Woodford (2003). The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity* 34(1), 139–235. [2, 9]
- Fair, R. C. and J. B. Taylor (1983). Solution and maximum likelihood estimation of dynamic nonlinear rational expectations models. *Econometrica* 51(4), 1169–85. [9]
- Fernández-Villaverde, J., G. Gordon, P. A. Guerrón-Quintana, and J. Rubio-Ramírez (2012, May). Nonlinear adventures at the zero lower bound. NBER Working Papers 18058, National Bureau of Economic Research, Inc. [16, 19, 42]
- Galí, J. (2008). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton, NJ: Princeton University Press. [16]
- Holden, T. and M. Paetz (2012). Efficient simulation of dsge models with inequality constraints. Working paper, University of Surrey. [9]

- Judd, K. L. (1992). Projection methods for solving aggregate growth models. *Journal of Economic Theory* 58, 410–452. [3, 13]
- Judd, K. L. (1998). *Numerical Methods in Economics*. Cambridge, MA: MIT Press. [40]
- Judd, K. L., L. Maliar, and S. Maliar (2012). Merging simulation and projection approaches to solve high-dimensional problems. Technical report. NBER working paper No. 18501. [8, 10]
- Jung, T., Y. Teranishi, and T. Watanabe (2005). Optimal monetary policy at the zero-interest-rate bound. *Journal of Money, Credit and Banking* 37(5), 813–35. [2, 8, 9]
- Kim, S. H., R. Kollmann, and J. Kim (2010). Solving the incomplete market model with aggregate uncertainty using a perturbation method. *Journal of Economic Dynamics and Control* 34(1), 50–58. [10]
- Laséen, S. and L. E. O. Svensson (2009). Anticipated alternative instrument-rate paths in policy simulations. Technical report, National Bureau of Economic Research, Inc. NBER working paper No. 14902. [9]
- Ljungqvist, L. and T. J. Sargent (2004). *Recursive Macroeconomic Theory, 2nd Edition*. Cambridge, MA: MIT Press. [15, 40]
- McGrattan, E. R. (1996). Solving the stochastic growth model with a finite element method. *Journal of Economic Dynamics and Control* 20(1-3), 19–42. [10]
- Nakata, T. (2013). Uncertainty at the zero lower bound. Finance and Economics Discussion Papers 2013-09, Board of Governors of the Federal Reserve System (U.S.). [20]
- Preston, B. and M. Roca (2007). Incomplete markets, heterogeneity and macroeconomic dynamics. Technical report. [10]
- Santos, M. S. and A. Peralta-Alva (2005, November). Accuracy of simulations for stochastic dynamic models. *Econometrica* 73(6), 1939–1976. [13]
- Smets, F. and R. Wouters (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American Economic Review* 97(3), 586–606. [3]
- Tauchen, G. (1986). Finite state markov-chain approximation to univariate and vector autoregressions. *Economics Letters*, 177–181. [40, 41, 42]
- Taylor, J. B. and H. Uhlig (1990). Solving nonlinear stochastic growth models: A comparison of alternative solution methods. *Journal of Business & Economic Statistics* 8(1), 1–17. [3, 10]
- Uhlig, H. (1995). A toolkit for analyzing nonlinear dynamic stochastic models easily. In *Computational Methods for the Study of Dynamic Economies*, pp. 30–61. Oxford University Press. [2]
- Yun, T. (2005). Optimal monetary policy with relative price distortions. *American Economic Review* 95(1), 89–109. [18]
- Zeldes, S. P. (1989). Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy* 97(2), 305–46. [22]

Parameter	Value	Parameter	Value
β , Discount Factor	0.96	γ , Relative Risk Aversion	2
δ , Depreciation Rate	0.10	α , Capital Share	0.33
ρ , Persistence of Tech. Shock	0.90	σ , St. Dev. of Tech. Innovation	0.013
ϕ , Threshold for Investment Constraint	0.975		

Table 1: Baseline Calibration of RBC Model with a Constraint on Investment

Log Investment			
Solution Method	Mean	St. dev.	Skewness
Nonlinear	-1.015	0.062	1.18
Piecewise linear	-1.015	0.063	1.33
First-order perturbation	-1.045	0.097	0.00
Log Consumption			
Solution Method	Mean	St. dev.	Skewness
Nonlinear	1.152	0.047	-0.22
Piecewise linear	1.152	0.047	-0.23
First-order perturbation	1.149	0.045	0.03
Correlation between Log Investment and Log Consumption			
Solution Method	Correlation		
Nonlinear	0.81		
Piecewise linear	0.80		
First-order perturbation	0.89		
% of periods at constraint			
Solution Method			
Nonlinear	41		
Piecewise linear	41		
First-order perturbation	0		

Table 2: A Comparison of Key Moments: Model with a Constraint on Investment

Parameter	Value	Parameter	Value
β , Discount Factor	0.994	ε , elasticity of substitution across goods	6
θ , Calvo Parameter	0.90	g , steady-state ratio of G/Y	0.20
ϕ_y , Response to output, mon. pol. rule	0.25	$\bar{\pi}$, steady state inflation rate	1.005
ϕ_π , Response to inflation, mon. pol. rule	2.5	ϕ , labor supply elasticity parameter	1
ρ , Persistence of Discount Rate Shock	0.80	σ , St. Dev. of Discount Rate Shock	0.005

Table 3: Baseline Calibration of Model with Zero Lower Bound on Nominal Interest Rates

	Piecewise Linear	Nonlinear	Linear
Frequency of hitting ZLB (%)	4.2	7.13	
Means			
Interest Rate (percent, AR)	4.43%	4.16%	4.39%
Inflation (percent, AR)	1.99%	1.77%	1.99%
Log Output	0.0125	0.0144	0.0126
Shock innovation (percent)	0.00%	0.00%	0.00%
Standard deviations			
Interest Rate (percent, AR)	2.44%	2.51%	2.52%
Inflation (percent, AR)	0.45%	0.52%	0.45%
Log Output	1.44%	1.54%	1.41%
Log Price Dispersion	0.33%	0.31%	0.32%
Skewness			
Log Output	-0.22	-0.49	-0.04
Interest Rate (percent, AR)	0.16	0.17	-0.02
Moments, conditional on ZLB			
Mean Inflation (percent, AR)	1.03%	0.69%	
Mean Log Output	-0.0206	-0.0189	
Mean, shock innovation (percent)	1.15%	0.99%	
St.dev., Inflation (percent, AR)	0.19%	0.29%	
St.dev. Log Output	0.85%	1.05%	

Table 4: A Comparison of Key Moments: Model with Zero Lower Bound

Parameter	Value	Parameter	Value
β , Discount Factor	0.945	γ , Relative Risk Aversion	1
R , Interest Rate	1.05	m , Borrowing Limit	1
ρ , Persistence of Income Shock	0.90	σ , St. Dev. Income Shock	0.0131

Table 5: Baseline Calibration of Model with Borrowing Constraint

Log Consumption			
Solution Method	Mean	St. dev.	Skewness
Nonlinear	-0.0512	3.40%	-0.24
Piecewise linear	-0.0513	3.46%	-0.22
First-order perturbation	-0.0516	3.60%	-0.04
Log Income			
	Mean	St. dev.	Skewness
Nonlinear	0.0000	3%	0.00
Piecewise linear	0.0000	3%	0.00
First-order perturbation	0.0000	3%	0.00
Correlations			
	$\ln Y, \ln C$	$\ln Y, \ln B$	
Nonlinear	0.96	0.96	
Piecewise linear	0.95	0.98	
First-order perturbation	0.92	1	
% of periods at borrowing constraint			
Nonlinear	73		
Piecewise linear	84		
First-order perturbation	100		

Table 6: A Comparison of Key Moments: Model with Borrowing Constraint.

Figure 1: Policy Function for Investment, RBC Model with Constraint on Investment.




Figure 2: RBC Model with Constraint on Investment: An Unexpected Decrease in Productivity, Followed by an Unexpected Increase in Productivity






Figure 3: RBC Model with Constraint on Investment: Comparison of Euler Equation Residuals Across Solution Methods (residuals expressed as a percent of consumption)




Note: “Median technology” corresponds to $A_t = 1$. “Low technology” corresponds to $A_t = 0.96$. “High technology” corresponds to $A_t = 1.04$. An open circle indicates that the investment constraint is binding.

Figure 4: New Keynesian Model Subject to the Zero Lower Bound: An Unexpected Increase in the Discount Factor, Followed by an Unexpected Decrease in the Discount Factor



Units on the abscissae denote quarters.

Figure 5: New Keynesian Model Subject to the Zero Lower Bound: Comparison of Euler Equation Residuals Across Solution Methods (residuals expressed as a percent of consumption)



Note: “Median Beta” corresponds to $\beta_t = 0.994$. “Low Beta” corresponds to $\beta_t = 0.965$. “High Beta” corresponds to $\beta_t = 1.023$. An open circle indicates that the zero lower bound on the nominal interest rate is binding.

Figure 6: Consumption Function, Model with Borrowing Constraint.





Figure 7: Model with Borrowing Constraint: An Unexpected Increase in Income, Followed by a Decrease



Units on the abscissae denote years.

Appendix

A. Fully-Nonlinear Algorithm for the Solution of the RBC model with a Constraint on Investment

We use two alternative approaches to finding a full nonlinear solution, dynamic programming and projection methods. We verified that the differences in the two approaches are negligible relative to the differences highlighted with the piecewise linear solution.¹⁶

Dynamic Programming Solution. The capital stock K_t is the only state variable in the model. We seek a rule that will map the current state variable K_{t-1} and the realization of the stochastic process A_t into a choice of K_t . We discretize and put boundaries on the support of the decision rule that we seek. We consider a uniformly spaced set of points for K_{t-1} and K_t . We discretize the support of both K_{t-1} and A_t . The lower boundary for K_{t-1} is 5 percent below the non-stochastic steady state for capital. The upper boundary is 40 percent above the non-stochastic steady state for capital. We experimented with different grids with little change in the results. We constrain A to lie within three standard deviations of its process, i.e. $|\ln A_t| \leq 3\sqrt{\frac{\sigma^2}{1-\rho^2}}$. We follow [Tauchen \(1986\)](#) in computing a finite state Markov-chain approximation for $\ln A_t$. The finest discretization we considered involved 75,000 points for capital and 201 points for the stochastic process. Use of shape-preserving splines, allowed us to reduce the number of points in the grid for capital without compromising the quality of the solution.

The dynamic programming algorithm that we use follows closely Chapter 12 of [Judd \(1998\)](#) and Chapter 3 of [Ljungqvist and Sargent \(2004\)](#). The initial choice for the decision rule in the dynamic program is taken to be the linear approximation to the decision rule obtained by standard methods. To accelerate the convergence of the dynamic programming algorithm, we use the Howard improvement algorithm.

Projection Solution. We restate the optimization problem in the model in Lagrangian form as:

$$\begin{aligned} \max_{\{C_t, K_{t+1}, \mu_t, \lambda_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} & \beta^t \frac{C_t^{1-\gamma} - 1}{1-\gamma} \\ & + \beta^t \mu_t (-C_t - K_t + (1-\delta) K_{t-1} + A_t K_{t-1}^\alpha) \\ & + \beta^t \lambda_t (-K_t + (1-\delta) K_{t-1} - \phi I) \end{aligned}$$

From standard manipulation of the first-order conditions for the Lagrangian problem, one obtains:

$$C_t^{-\gamma} - \lambda_t - E_t \left[\beta C_{t+1}^{-\gamma} ((1-\delta) + \alpha A_{t+1} K_t^{\alpha-1}) - \beta (1-\delta) \lambda_{t+1} \right] = 0 \quad (\text{A.1})$$

¹⁶ For instance, the compensating variation for the use of the dynamic programming solution relative to the projection solution is in the order of \$1 in \$100,000,000 for the baseline calibration for the coarsest grid used.

We seek a solution in the form of a function $g(K_{t-1}, A_t)$ that approximates $C_t^{-\gamma} - \lambda_t$ subject to the complementary slackness condition, $\lambda_t(-K_t + (1 - \delta)K_{t-1} - \phi I) = 0$ following the method of parameterized expectations described by [den Haan and Marcet \(1990\)](#) and [Christiano and Fisher \(2000\)](#). We approximate $g(K_{t-1}, A_t)$ with a Chebychev polynomial of order 6 and approximate the process for $\ln A$ with a Markov process, following [Tauchen \(1986\)](#). A Markov process with 10 states usually provides an adequate approximation to the underlying process. In an abundance of caution, we use 51 states. We constrain $\ln A$ to lie within three standard deviations of its process, i.e. $|\ln A_t| \leq 3\sqrt{\frac{\sigma^2}{1-\rho^2}}$.

We solve for the parameters of the Chebychev polynomial function using orthogonal collocation. Given K_t and A_t , we guess $\lambda_t = 0$. Consistent with that guess, $C_t = g(K_{t-1}, A_t)^{-\gamma}$ and we can back out K_t from the resource constraint. We then check whether $K_t - (1 - \delta)K_{t-1} \geq \phi I$. If so, the original guess for λ_t was correct. If not, for the complementary slackness condition to hold, it must be that $K_t = (1 - \delta)K_{t-1} + \phi I$. C_t is then given by the resource constraint and $\lambda_t = g(K_{t-1}, A_t) - C_t^{-\gamma}$.

B. Fully-Nonlinear Algorithm for the Solution of the New Keynesian Model Subject to the Zero Lower Bound

The necessary conditions for an equilibrium are:

$$C_t = \frac{1}{\beta_t E_t \left(\frac{R_t}{C_{t+1} \Pi_{t+1}} \right)} \quad (\text{A.2})$$

$$mc_t = w_t \quad (\text{A.3})$$

$$w_t = \psi L_t^\vartheta C_t \quad (\text{A.4})$$

$$\varepsilon x_{1t} = (\varepsilon - 1) x_{2t} \quad (\text{A.5})$$

$$x_{1t} = \frac{1}{C_t} mc_t Y_t + \theta E_t \beta_t \Pi_{t+1}^\varepsilon x_{1t+1} \quad (\text{A.6})$$

$$x_{2t} = \Pi_t^* \left(\frac{Y_t}{C_t} + \theta \beta_t E_t \frac{\Pi_{t+1}^{\varepsilon-1}}{\Pi_{t+1}^*} x_{2t+1} \right) \quad (\text{A.7})$$

$$Z_t = R \left[\left(\frac{\Pi_t}{\Pi} \right)^{\phi_p} \left(\frac{Y_t}{Y} \right)^{\phi_y} \right] \quad (\text{A.8})$$

$$R_t = \max(Z_t, 1) \quad (\text{A.9})$$

$$G_t = s_g Y_t \quad (\text{A.10})$$

$$1 = \theta \Pi_t^{\varepsilon-1} + (1 - \theta) (\Pi_t^*)^{1-\varepsilon} \quad (\text{A.11})$$

$$v_t = \theta \Pi_t^\varepsilon v_{t-1} + (1 - \theta) (\Pi_t^*)^{-\varepsilon} \quad (\text{A.12})$$

$$Y_t = C_t + G_t \quad (\text{A.13})$$

$$Y_t = \frac{L_t}{v_t} \quad (\text{A.14})$$

$$\ln \beta_t = (1 - \rho) \log \beta + \rho \log \beta_{t-1} + \sigma \epsilon_t. \quad (\text{A.15})$$

Given the balanced budget, $B_t = 0$. The model variables in the system above are C , R , Π , Π^* , mc , w , L , x_1 , x_2 , Z , G , Y , v , and β . We seek a solution in the form of three functions $g_1(v_{t-1}, \beta_t)$, $g_2(v_{t-1}, \beta_t)$

and $g_3(v_{t-1}, \beta_t)$ that approximate, respectively $\frac{1}{C_t \beta_t R_t}$, $\frac{(x_{1t} - \frac{1}{C_t} m c_t Y_t)}{\theta \beta_t}$ and $\frac{(\frac{x_{2t}}{\Pi_t^*} - \frac{Y_t}{C_t})}{\theta \beta_t}$ subject to $R_t = \max(Z_t, 1)$. We approximate g_1 , g_2 , and g_3 with Chebychev polynomials of order 6 and approximate the process for β_t with a Markov process, following [Tauchen \(1986\)](#) and using 51 states. We constrain $\ln \beta_t - \ln \beta$ to lie within 3.5 standard deviations of its process. We also constrain v_t to lie in the interval bounded by 1 and 1.04. We solve for the parameters of the Chebychev polynomial functions using orthogonal collocation. Given these choices, we follow the same guess-and-verify approach described in the appendix of [Fernández-Villaverde et al. \(2012\)](#).

Figure (A) shows the absolute values of the residuals for the three intertemporal equations, equations (A.2), (A.6), and (A.7). The residuals were normalized respectively by C_t , x_{1t} , and x_{2t} . The maximum residual is in the order of 10^{-5} . Figure B shows the residuals for the intertemporal equations for the piecewise linear solution and for a linear solution that disregards the zero lower bound. The residuals for the piecewise linear solution for the model with the ZLB enforced stay close to the residuals for the linear solution for a model that disregards the ZLB.

C. Robustness Analysis

RBC Model. Table A reports the subsidy (as a percent of steady-state consumption) that would compensate the agent for the use of the piecewise linear algorithm over a fully nonlinear method with initial conditions set at the non-stochastic steady state. The welfare cost of using our piecewise solution method is a non-monotonic function of risk aversion; it is “high” when risk aversion is around 1, “low” when risk aversion is at 2, and increasing with risk aversion for values of γ around 3 or higher. This happens because – under the full nonlinear solution – the irreversibility constraint has two opposing effects on the equilibrium average level of capital. The first effect – an *illiquid capital effect* – works to reduce average capital: when capital is irreversible, it is less useful in smoothing consumption when technology is low. The second effect – a *precautionary capital effect* – works to increase average capital through a precautionary saving effect: as capital is irreversible, consumption is more volatile, and more capital is held even if it is less useful in bad states of the world. These two opposing effects – captured by the fully nonlinear solution but not by the piecewise linear method – explain the non-monotonicity. Under the baseline calibration ($\gamma = 2$), the illiquid capital and precautionary capital effect almost offset each other, and the stochastic fixed point for capital¹⁷ happens to be close to its non-stochastic steady state, which does not depend on uncertainty and irreversibility. The precautionary effect dominates for higher values of γ , whereas the illiquidity effect dominates for low risk aversion.¹⁸ These differences in the demand for capital under different parameterizations of the model influence the performance of the piecewise linear solution method. In particular, for high levels of risk aversion, the linear component of the solution cannot capture the increase in demand for capital stemming from precautionary motives, and the performance of the solution algorithm deteriorates, with a the welfare cost of \$1 in about \$100,000 when $\gamma = 5$. Similarly, the linear component of the piecewise linear solution is not able to

¹⁷ The stochastic fixed point for capital is the level attained by capital with all innovations to technology set to 0,

¹⁸ In a partial equilibrium setting, [Abel and Eberly \(1999\)](#) also find that the introduction of capital irreversibility has an ambiguous effect on the long-run level of the capital stock.

capture the drop in capital demand when risk aversion is low, and the performance of the piecewise linear solution also deteriorates.

Table A also considers sensitivity with respect to the choice of the value for the discount factor β . The welfare cost of using the piecewise method increases as the discount factor rises. We conjecture that in the plain vanilla RBC model the nonlinearities become more pronounced as the risk free rate becomes lower, thus penalizing linearization in general over a fully nonlinear solution algorithm.

Model	Solution Method		Solution Method	
	Piecewise linear	\$ 1 every	Constant capital	\$ 1 every
$\gamma = 2, \beta = 0.96$ (Baseline)	0.000007%	\$ 14,556,184	0.04%	\$ 2,842
$\gamma = 1$	0.000036%	\$ 2,793,317	0.02%	\$ 4,653
$\gamma = 3$	0.000113%	\$ 881,804	0.05%	\$ 1,980
$\gamma = 4$	0.000413%	\$ 242,411	0.07%	\$ 1,481
$\gamma = 5$	0.000958%	\$ 104,424	0.09%	\$ 1,157
$\beta = 0.98$	0.000028%	\$ 3,557,192	0.05%	\$ 1,938
$\beta = 0.94$	0.000003%	\$ 29,844,218	0.03%	\$ 3,838
$\gamma = 2, \phi = 0$	0.000014%	\$7,194,352	0.05%	\$ 2,041

Note: The “Piecewise Linear” column indicates the subsidy rate (as a percent of steady-state consumption) that would compensate an agent for the use of the piecewise linear algorithm over a fully nonlinear method with initial conditions set at the non-stochastic steady state. The “Constant Capital” column indicates the subsidy when the agent uses a suboptimal decision rule setting the capital stock to its previous value.

Table A: Utility Cost of the Solution Method: RBC Model with Constraint on Investment.


Model	Solution Method					
	Piecewise linear			Nonlinear		
	% at ZLB	log output	St.dev.	% at ZLB	log output	St.dev.
Baseline	4.2	1.44%	-0.22	7.13	1.54%	-0.49
$\bar{\pi} = 1, \beta = 0.9891$	6.7	1.35%	-0.38	9.35	1.51%	-0.76
$\phi_{\pi} = 5, \phi_y = 0$	6.7	0.86%	-0.66	9.35	0.94%	-1.20
$\phi_{\pi} = 10$	2.91	1.69%	-0.13	4.41	1.72%	-0.14

Table B: Robustness Analysis of Model with ZLB

Model	Solution Method		Solution Method	
	Piecewise linear	\$ 1 every	First-order Perturbation (Always Constrained)	\$ 1 every
Baseline	0.0007%	\$ 149,280	0.0033%	\$ 30,731
High debt limit, $m = 2$	0.0013%	\$ 77,444	0.0071%	\$ 14,082
High risk aversion, $\gamma = 2$	0.0023%	\$ 44,140	0.0131%	\$ 7,657
High uncertainty, $\sigma = 0.0196$	0.0024%	\$ 41,233	0.0116%	\$ 8,650
Low impatience, $\beta = 0.948$	0.001%	\$ 102,989	0.0056%	\$ 17,812


Table C: Utility Cost of the Solution Method: Model with Borrowing Constraint.

Figure A: New Keynesian Model Subject to the Zero Lower Bound: Intertemporal Errors of the Collocation Solution



Note: “Median Beta” corresponds to $\beta_t = 0.994$. “Low Beta” corresponds to $\beta_t = 0.965$. “High Beta” corresponds to $\beta_t = 1.023$. An open circle indicates that the zero lower bound on the nominal interest rate is binding.

Figure B: New Keynesian Model: Intertemporal Errors for the Piecewise Linear Solution (with ZLB enforced) and for the Linear Solution (with ZLB disregarded)



Note: “Median Beta” corresponds to $\beta_t = 0.994$. “Low Beta” corresponds to $\beta_t = 0.965$. “High Beta” corresponds to $\beta_t = 1.023$. An open circle indicates that the zero lower bound on the nominal interest rate is binding.