

ACCURACY, SPEED AND ROBUSTNESS OF POLICY FUNCTION ITERATION*

Alexander W. Richter

Nathaniel A. Throckmorton

Todd B. Walker

August 28, 2013

ABSTRACT

Policy function iteration methods for solving and analyzing dynamic stochastic general equilibrium models are powerful from a theoretical and computational perspective. Despite obvious theoretical appeal, significant startup costs and a reliance on grid-based methods have limited the use of policy function iteration as a solution algorithm. We reduce these costs by providing a user-friendly suite of MATLAB functions that introduce multi-core processing and Fortran via MATLAB's executable function. Within the class of policy function iteration methods, we advocate using time iteration with linear interpolation. We examine a canonical real business cycle model and a new Keynesian model that features regime switching in policy parameters, Epstein-Zin preferences, and monetary policy that occasionally hits the zero-lower bound on the nominal interest rate to highlight the attractiveness of our methodology. We compare our advocated approach to other familiar iteration and approximation methods, highlighting the tradeoffs between accuracy, speed and robustness.

Keywords: Policy function iteration, Zero lower bound, Epstein-Zin preferences, Markov switching, Chebyshev polynomials, Real business cycle model, New Keynesian model

JEL Classifications: C63; C68; E52; E62

*Richter, Department of Economics, Auburn University, 0332 Haley Center, Auburn, AL (arichter@auburn.edu); Throckmorton, Department of Economics, Indiana University, 100 S. Woodlawn, Wylie Hall 105, Bloomington, IN (nathrock@indiana.edu); Walker, Department of Economics, Indiana University, 100 S. Woodlawn, Wylie Hall 105, Bloomington, IN (walkertb@indiana.edu). We thank Troy Davig for many helpful discussions and Bulent Guler and a referee for helpful comments. Walker acknowledges support from the National Science Foundation under grant SES 096221.

1 INTRODUCTION

The Great Recession, the prospect of exponentially rising government debt, interest rates at the zero lower bound, and potential sudden changes to monetary and fiscal policy make clear that nonlinearities are a crucial element to contemporary macroeconomic analysis. Successfully modeling these scenarios requires large and persistent deviations from the non-stochastic equilibrium. Linear approximations around a deterministic steady state poorly capture these equilibrium properties. A nonlinear analysis is needed.

Policy function iteration methods for solving and analyzing dynamic stochastic general equilibrium models are powerful from a theoretical and computational perspective. Despite obvious theoretical appeal, significant startup costs and a reliance on grid-based methods have limited the use of policy function iteration as solution algorithm. We reduce these costs by providing a user-friendly suite of MATLAB functions. Within the class of policy function iteration methods, we advocate using time iteration with linear interpolation, since it provides a flexible and accurate way to solve dynamic stochastic general equilibrium models with substantial nonlinearity.

We demonstrate the usefulness of our approach by examining several topical examples. We begin with a simple real business cycle model and a standard New Keynesian model. These canonical models are useful starting points because their solutions and properties are well known. They also provide useful benchmarks for speed and accuracy. [Section 3](#) provides step-by-step instructions for how to solve these models using our advocated approach in MATLAB. We introduce multi-core processing using the Parallel Computing Toolbox (PCT) and integrate Fortran through MATLAB executable functions (MEX). Using a 6-core processor (3.47GHz each) and MEX, our suite reduces computational time by a factor of 12 in the RBC model and by a factor of 24 in the NK model relative to non-parallelized code that does not use MEX. Additional stochastic components further increase the speeds gains associated with MEX and parallelization. [Section 4](#) analyzes the tradeoffs between accuracy and speed using alternative iteration and approximation methods.

We demonstrate the flexibility of our advocated approach using the canonical New Keynesian model. [Section 5](#) adds Epstein-Zin preferences to show that our suite can be used to study asset pricing facts. [Section 6](#) introduces regime switching in monetary and fiscal policy parameters with an emphasis on understanding the expectational effects generated by regime switches. [Section 7](#) adds a zero lower bound on the nominal interest rate set by the monetary authority, which introduces a kink in the policy functions. We provide MATLAB code with extensive documentation for each example. Richter and Throckmorton (2012) provide additional supporting code.

The primary benefit of our advocated approach over perturbation methods [Gaspar and Judd (1997); Judd and Guu (1993, 1997); Schmitt-Grohe and Uribe (2004)] is its ability to easily account for sudden policy changes and other inherent nonlinearities (i.e., zero interest rate bound, default, irreversible investment, *etc.*). The flexibility and simplicity of the algorithm has led a recent segment of the literature to use this approach to estimate monetary and fiscal policy regimes, quantify expectational effects of policy changes, and study key counterfactual policies.¹ The suite of programs described in this paper can be easily adapted to handle models with: [i] endogenous regime change, [ii] temporarily nonstationary processes, [iii] binding collateral constraints, [iv] stochastic volatility, [v] news shocks, and [vi] heterogeneous agents. Although this method is grid-based, our suite of programs can solve models with state spaces that contain more than one million

¹See Basu and Bundick (2012); Bi (2012); Bi et al. (2013); Chung et al. (2007); Davig and Leeper (2006, 2008); Davig et al. (2010, 2011); Gavin et al. (2013); Kumhof and Ranciere (2010); Mertens and Ravn (2013); Richter (2012).

nodes and multiple stochastic processes in roughly one hour on a standard 6-core computer.

We remind readers that policy function iteration methods are a numerical byproduct of using monotone operators to prove existence and uniqueness of equilibria [Coleman (1991)]. This provides an additional benefit to using policy function iteration methods, as powerful approximation results and proofs of existence and uniqueness can be employed in conjunction with the numerical algorithm. We briefly review the theory of monotone operators in an online appendix.

2 PRACTICAL GUIDE

All of the routines required to implement the algorithm are written in MATLAB or compiled as MEX (Fortran 90) functions. This code and the code used to solve the models below are publicly available at <http://www.auburn.edu/~awr0007/>. The algorithm is broken down into a sequence of easily implementable steps that are executed with the following set of functions:

- `script.m` - Main script that establishes all user-specified inputs in the structure, `O`, and executes all functions. It contains a parallel for-loop that distributes the optimization routine at each point (node) in the discretized state space across all locally available processors, which considerably improves computation time. Optimization is performed with Chris Sims' `csolve.m` by finding updated policy functions that satisfy the equilibrium conditions of the model until a user-specified convergence criterion is met.
- `parameters.m` - Outputs a structure, `P`, containing the baseline calibration of the model.
- `steadystate.m` - Outputs a structure, `S`, with input `P`, containing the deterministic steady state and implied parameters of the model.
- `grids.m` - Outputs a structure, `G`, with inputs `O` and `P`, containing the discretized state space. The structure `O` contains the number of grid points and bounds of each state variable.
- `guess.m` - Outputs an initial conjecture for each policy function with inputs `O`, `P`, `S`, and `G`. The structure `O` specifies whether the initial conjectures for the policy functions are drawn from the log-linear solution to the model or the most recent iteration, which allows the user to resume the algorithm if it is interrupted prior to convergence.
- `variables.m` - Outputs a structure, `V`, containing an index of variables, forecast errors, and shocks and the corresponding variable descriptions.
- `linmodel.m` - Outputs the linear transition matrix, `T`, the impact matrix of the stochastic realizations, `M`, and a 2-element vector of flags, `eu`, indicating existence and uniqueness of the linear solution. The linear model is solved using Chris Sims' `gensys.m` algorithm and requires `P`, `S`, and `V` as inputs. `gensys.m` is user written MATLAB code that is designed to solve stochastic linear rational expectations models. The key to using this code is to map the model into the form

$$T_0 X_{t+1} = T_1 X_t + \Psi \varepsilon_{t+1} + \Pi \eta_{t+1},$$

where X is a vector of variables (exogenous and endogenous), ε is a vector of shocks, and η is a vector of forecast errors whose elements satisfy

$$\eta_{t+1}^x = x_{t+1} - E_t x_{t+1}$$

for some $x \in X$. Assuming the shocks are normally distributed and mean zero, `gensys` outputs the coefficients, T and M , of the following equation:

$$X_{t+1} = TX_t + M\varepsilon_{t+1}.$$

- `eqm.m` - Outputs a matrix, R , containing the residuals to a subsystem of expectational equations that are constrained by the remaining equations in the equilibrium system. This function requires as inputs the structures P , S and G , a vector of initial conjectures, the state at each node, and the current policy functions (required for interpolation/extrapolation).
- `allinterp.m` - Performs linear interpolation and extrapolation. MATLAB and Fortran routines (`Fallinterp.f90`) are provided. The naming convention uses a sequence of three integers representing the number of state variables, policy functions, and continuous stochastic variables. These functions can evaluate multiple points, but, in the interest of speed, are not generalized and may require straightforward modifications for alternative combinations of states, policies, and stochastic variables.²
- `pf.mat` - Data array containing the current policy functions and all structures.

2.1 NUMERICAL ALGORITHM This section outlines how to implement policy function iteration with time iteration and linear interpolation. There are a number of model-specific initializations: model parameters, steady state values, and grid parameters (e.g., the number of points in each dimension and the distance between them) specified by the programmer. Through trial and error, we found that the initializations are best performed in the following order.

First, set the baseline calibration of the model in `parameters.m`. In addition, specify a convergence criterion for the policy functions, which affects the accuracy and speed of the solution.

Second, calculate (using a nonlinear solver, if necessary) the deterministic steady state and any implied parameters in `steadystate.m`. Steady state values are required for solving the linear model and are helpful for setting up the discretized state space.

Third, specify the number of points and the bounds for each grid, which are contained in structure O and specified in `script.m`. The built-in MATLAB function `linspace` establishes a grid for each state variable in `grids.m`. The built-in MATLAB function `ndgrid` creates an array for each state variable where every position represents a unique permutation of the discretized state variables. The total number of elements in each array is the product of the number of points assigned to each state variable and represents the number of nodes in the discretized state space.

Fourth, establish an optimal set of variables as numerical policies. The set of possible policy functions is not unique, but as a rule of thumb, establish policy functions over the minimum set of variables required to solve for all time t variables given the state of the economy. If possible, reduce the number of policy variables to the number of equations with expectations operators.

Finally, obtain initial conjectures (guesses) for each policy function. Some models are sensitive to the guess, but, in general, a linear solution provides a sufficient approximation.³ We use

²Richter and Throckmorton (2012) provide the Fortran and MATLAB source code, as well as the compiled 32- and 64-bit MEX functions, for each of the examples discussed in the paper.

³This is not true for all models. For example, if the model contains discrete variables, such as state dependent parameters, then the conditional linear solution may poorly approximate the global solution. In this case, one can obtain a linear solution for each realization of the parameter(s). A linear combination of those solutions typically provides a good initial conjecture for the state-dependent nonlinear model.

Sims' (2002) `gensys.m` algorithm.⁴ Enter the log- or level-linear system into `linmodel.m`. `guess.m` calls `linmodel.m` to solve the linear model and generate initial conjectures.

After obtaining initial conjectures for each of the policy functions, set up the equilibrium system in `eqm.m`. Using the original policy function guesses or the solution from the previous iteration, solve for all time t variables using all of the equilibrium conditions, except those that contain expectations and require numerical integration. Next, calculate updated (time $t + 1$) values for each of the policy variables using linear interpolation/extrapolation.⁵ We provide two computational routines for this step—one written in MATLAB (`allinterp*.m`) and one written in Fortran 90 (`Fallinterp*.mex*`) but called as a MATLAB executable (MEX) function. The Fortran code is much faster, but for those who are unfamiliar with MEX, the MATLAB function is sufficient for relatively small models—those with no more than four state variables and only one stochastic component. Further details about these routines are given in an online appendix.

After computing the updated values of each policy variable, solve for the remaining time $t + 1$ variables needed to calculate time t expectations by applying numerical integration using the trapezoid rule or Gauss-Hermite quadrature. Additional details about these integration methods are provided in an online appendix. Finally, using Chris Sims' root finder, `csolve.m`, solve for the zeros of the equations with embedded expectations, subject to each of the remaining equilibrium conditions. The output of `csolve.m` on each node are policy values that satisfy the equilibrium system of equations to a specified tolerance level. This set of values characterizes the updated policy functions for the next iteration. If the distance between the guess and the updated policy values is less than the convergence criterion on all nodes, then the policies have converged to their equilibrium values. Otherwise, use the updated policy functions as the new guess until convergence.

3 EXAMPLES

We first consider two conventional models—the real business cycle (RBC) model and the new Keynesian (NK) model with textbook treatments provided by McCandless (2008) and Walsh (2010).

3.1 RBC MODEL The representative household chooses sequences, $\{c_t, k_t, n_t\}_{t=0}^{\infty}$, that maximize expected lifetime utility, $E_0 \sum_{t=0}^{\infty} \beta^t \{c_t^{1-\sigma}/(1-\sigma) - \chi n_t^{1+\eta}/(1+\eta)\}$, where β is the subjective discount factor, $1/\sigma$ is the intertemporal elasticity of substitution, $1/\eta$ is the Frisch elasticity of labor supply, c is consumption, and n is labor hours. These choices are constrained by

$$\begin{aligned} c_t + k_t &= w_t n_t + r_t^k k_{t-1} + (1 - \delta)k_{t-1} \\ k_t &= (1 - \delta)k_{t-1} + i_t, \end{aligned}$$

where w is the real wage, r^k is the real rental price of capital, k is the capital stock, and i is investment. The household's optimality conditions imply

$$1 = \beta E_t \{ (c_t/c_{t+1})^\sigma (r_{t+1}^k + 1 - \delta) \} \quad \text{and} \quad w_t = \chi n_t^\eta c_t^\sigma.$$

Output is produced according to $y_t = z_t k_{t-1}^\alpha n_t^{1-\alpha}$, where $0 < \alpha < 1$. Productivity follows

$$z_t = (1 - \rho)\bar{z} + \rho z_{t-1} + \varepsilon_t,$$

⁴Other methods include Uhlig's (1997) toolkit and Klein's (2000) algorithm.

⁵Alternatively, Chebyshev polynomials or cubic splines could be used for interpolation/extrapolation. We advocate for linear interpolation since it is faster and more stable, but we consider alternative methods in [section 4](#).

where \bar{z} is the average productivity level, and $\varepsilon_z \sim i.i.d. N(0, \sigma_\varepsilon^2)$. A perfectly competitive firm chooses $\{k_t, n_t\}$ to maximize $y_t - w_t n_t - r_t^k k_{t-1}$. The firm's optimality conditions imply

$$r_t^k = \alpha y_t / k_{t-1} \quad \text{and} \quad w_t = (1 - \alpha) y_t / n_t.$$

The aggregate resource constraint is given by $c_t + i_t = y_t$. A competitive equilibrium is given by the household's and firm's optimality conditions, the production function, the law of motion for capital, the aggregate resource constraint, and the productivity process.⁶

3.1.1 SOLVING THE MODEL The following are detailed instructions on how to setup and solve the nonlinear model with time iteration and linear interpolation.

- There are five structural parameters. Using a quarterly calibration, these parameters are set to $\beta = 0.99$ (4% real interest rate), $\delta = 0.025$, $\alpha = 0.33$, $\sigma = 1$, and $\eta = 1$. The productivity process is persistent with $\rho = 0.95$ and $\bar{z} = 1$. Productivity shocks are normally distributed with mean zero and standard deviation $\sigma_\varepsilon = 0.0025$. The convergence criterion is 10^{-10} .
- Given the steady state labor share, $\bar{n} = 0.33$, the deterministic steady state has a straightforward analytical solution.
- There are two continuous state variables, k_{t-1} and z_t , and one continuous stochastic variable, ε_{t+1} , that are discretized. The bounds of the state space and the number of points for each state variable are contained in the structure `O` and inputs of `grids.m`. The capital and productivity grids are evenly spaced with bounds that are 5 percent from their steady state values. These bounds minimize extrapolation and ensure that the model simulates on the discretized state space. With only two state variables, we can afford dense grids. We specify 41 points for each state variable (k_{t-1} and z_t), which implies 1,681 nodes. No specific number of grid points is required for any of the continuous state variables (minimum 3 points), but using more grid points on state variables that contribute to the curvature of the policy functions improves accuracy. However, there is delicate balance between increasing the number of grid points and keeping the problem numerically tractable.
- We enter the log-linear equilibrium system (8 equations) into `linmodel.m` to produce initial conjectures to the nonlinear model.
- We choose n_t as a policy since it allows us to conveniently calculate the time t variables, but this choice is not unique. Given $\{k_{t-1}, z_t, n_t\}$, it is easy to calculate y_t from the production function, i_t from the aggregate resource constraint, k_t from the law of motion for capital, and c_t by combining the firm's and household's optimality conditions for labor. To obtain n_{t+1} interpolate/extrapolate for all realizations of ε_{t+1} . Given $\{k_t, z_{t+1}, n_{t+1}\}$, calculate c_{t+1} and r_{t+1}^k (from the firm's optimality condition), which enter expectations. We use Gauss Hermite quadrature to integrate across ε_{t+1} . `csolve.m` searches for labor policy values that satisfy the equilibrium system in `eqm.m` by finding the zero of the consumption Euler equation.
- After entering the equilibrium system in `eqm.m`, execute `script.m` to run the algorithm. To take advantage of the speed gains associated with the parallel toolbox, execute the command `matlabpool`, which pools all locally available processors.

Processors	All MATLAB (No MEX)		Interpolation using MEX	
	Structures	No Structures	Structures	No Structures
1	129 (1.0)	100 (1.3)	90.5 (1.4)	62.9 (2.1)
2	65.6 (2.0)	52.4 (2.5)	47.1 (2.7)	34.2 (3.8)
6	27.5 (4.7)	23 (5.6)	20.7 (6.2)	16.6 (7.8)

Table 1: RBC model solution times—MEX, parallelization, structures comparison (in seconds). Based on a state space of 1,681 nodes (41 points on k_{t-1} , 41 points on z_t). We specify 10 realizations of ε_{t+1} . The routines were computed with an Intel Xeon X5690 6-core processor (3.47GHz) operating 64-bit Windows 7. The values in parentheses are the speed gains relative to the slowest setup.

3.1.2 SOLUTION TIMES The major drawback with time iteration is its reliance on a nonlinear solver that executes on each node, which is computationally expensive. The introduction of multi-core processing and Fortran via MEX reduces this computational burden (see table 1). Multi-core processing in MATLAB is easily facilitated with the PCT, which creates a “worker pool” consisting of locally available processors. Without the PCT, MATLAB only uses one processor even though additional processors are available in most computers. In this case, computational resources are not maximized and each node in the discretized state space is evaluated sequentially. The PCT allows MATLAB to evaluate multiple nodes simultaneously, which produces near linear speed gains.

The interpolation/extrapolation step imposes the most significant slowdown. Our algorithm achieves further speed gains by using a Fortran 90 MEX function that is called in MATLAB. For relatively small models (two or fewer state variables), full implementation using only MATLAB routines is relatively inexpensive. For larger models, such as the NK model, the benefits of MEX clearly outweigh the costs. The two pitfalls of programming in Fortran are the need for a MEX gateway function and the noninteractive nature of the editor. The MEX gateway function initializes the function’s inputs and outputs and calls other subroutines that perform the actual task. There is a fixed cost in understanding how to initialize various functions with MATLAB’s application programming interface, allocate memory, and create matrices and arrays that properly interact with MATLAB. An additional cost is the inconvenience of debugging. MATLAB’s editor is proactive in addressing problem areas in scripts; Fortran is obviously not since it is a compiled language. Even if a Fortran function successfully compiles via MEX there still may be problems in the code that require further attention. Our Fortran functions are written using the free-format, which is relatively easy to learn since it is very similar to MATLAB syntax.⁷ Once the Fortran function is compiled using MEX (which requires Intel Visual Fortran), a new MEX function is created, which is called in MATLAB in the exact same way as MATLAB’s built-in functions.

MATLAB structures are a convenient way to group data—such as parameters, steady state values, grids, and model-specific options—and reduce the number of inputs when calling functions. The drawback is that MATLAB requires more time to access values contained in structures. Thus, we often forfeit convenience and simplicity in favor of faster code by removing the values

⁶We also provide code to solve the canonical RBC model with three conventional frictions—capital adjustment costs, variable utilization rates, and external habit persistence.

⁷By default, MEX only interprets fixed-format (f77) Fortran code. Since our fortran code is written in the free-format (f90), MATLAB batch files need to be modified. Navigate to the mexopts folder in the MATLAB directory (e.g., ... \MATLAB\R2010a\bin\win64\mexopts), and open the batch file corresponding to the installed version of Intel Visual Fortran (IVF) in a text editor. Delete the ‘/fixed’ flag and save the batch file. Then use `mex -setup` to select IVF as the compiler in MATLAB. MATLAB must be restarted anytime the batch file is changed.

contained in the structures before the `parfor` loop.

Table 1 reports (in seconds) how the solution times differ across the three modifications discussed above—multi-core processing, interpolation/extrapolation via MEX, and MATLAB structures. Using two processors instead of only one (which is the default) reduces the solution time by nearly 50 percent regardless of the choice to use MEX or structures. Using MEX also cuts the solution time in half, which is independent of the number of processors or the use of structures. In both cases, adopting structures simplifies the code, but it imposes a fixed cost of roughly 30 seconds, which decreases with the number of processors.

The values in parentheses are the speed gains relative to the slowest setup. Notice that it does not increase one-for-one with the number of processors, which is indicative of communication overhead. Additional processors still provide evident speed gains, but this benefit is limited if the code is not optimized for speed. For example, code that employs six processors without MEX while retaining structures yields nearly the same speed gain as code that employs just two processors, but uses MEX and removes the structures. Overall, using six processors with optimized code (MEX and no structures) reduces the solution time by a factor of nearly 8.

3.2 NEW KEYNESIAN MODEL The representative household chooses $\{c_t, n_t, M_t, k_t\}_{t=0}^{\infty}$ to maximize expected lifetime utility,

$$E_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{c_t^{1-\sigma}}{1-\sigma} - \chi \frac{n_t^{1+\eta}}{1+\eta} + \nu \frac{(M_t/P_t)^{1-\kappa}}{1-\kappa} \right\}, \quad \chi, \nu > 0$$

where P_t is the aggregate price index, M_t is nominal money balances, and $1/\kappa$ is the interest (semi) elasticity of money demand. The household's choices are constrained by

$$\begin{aligned} c_t + m_t + i_t + b_t &= (1 - \tau_t)(w_t n_t + r_t^k k_{t-1}) + (m_{t-1} + r_{t-1} b_{t-1})/\pi_t + d_t, \\ k_t &= (1 - \delta)k_{t-1} + i_t, \end{aligned}$$

where lower case letters denote real quantities ($x_t = X_t/P_t$), $\pi_t = P_t/P_{t-1}$ is the gross inflation rate, b_t is the stock of real government bonds, τ_t is a proportional tax rate levied against capital and labor earnings, and d_t is the share of real firm profits. Optimality implies

$$\begin{aligned} \chi n_t^\eta c_t^\sigma &= (1 - \tau_t)w_t, \\ \nu m_t^{-\kappa} &= (1 - 1/r_t)c_t^{-\sigma}, \\ 1 &= \beta r_t E_t \{ (c_t/c_{t+1})^\sigma / \pi_{t+1} \}, \\ 1 &= \beta E_t \{ (c_t/c_{t+1})^\sigma ((1 - \tau_{t+1})r_{t+1}^k + 1 - \delta) \}. \end{aligned}$$

The production sector consists of monopolistically competitive intermediate goods producing firms who produce a continuum of differentiated inputs and a representative final goods producing firm. Each firm $i \in [0, 1]$ in the intermediate goods sector produces a differentiated good, $y_t(i)$, with identical technologies given by $y_t(i) = k_{t-1}(i)^\alpha n_t(i)^{(1-\alpha)}$, where $k(i)$ and $n(i)$ are the levels of capital and employment used by firm i . Each intermediate firm chooses capital and labor to minimize its operating costs, $r_t^k k_{t-1}(i) + w_t n_t(i)$, subject to its production function.

Using a Dixit and Stiglitz (1977) aggregator, the representative final goods producer purchases $y_t(i)$ units from each intermediate firm to produce the final good, $y_t \equiv [\int_0^1 y_t(i)^{(\theta-1)/\theta} di]^\theta / (\theta-1)$,

where $\theta > 1$ measures the elasticity of substitution between the intermediate goods. Maximizing profits for a given level of output yields the demand function for intermediate inputs given by $y_t(i) = (p_t(i)/P_t)^{-\theta} y_t$, where $P_t = [\int_0^1 p_t(i)^{1-\theta} di]^{1/(1-\theta)}$ is the price of the final good. Following Rotemberg (1982), each firm faces a cost to adjusting its price, which emphasizes the potentially negative effect that price changes can have on customer-firm relationships. Using the functional form in Ireland (1997), real profits of firm i are

$$d_t(i) = \left[\left(\frac{p_t(i)}{P_t} \right)^{1-\theta} - \Psi_t \left(\frac{p_t(i)}{P_t} \right)^{-\theta} - \frac{\varphi}{2} \left(\frac{p_t(i)}{\bar{\pi} p_{t-1}(i)} - 1 \right)^2 \right] y_t,$$

where $\varphi \geq 0$ determines the magnitude of the adjustment cost, Ψ is real marginal costs, and $\bar{\pi}$ is the steady state gross inflation rate. Each intermediate goods producing firm chooses their price level, $p_t(i)$, to maximize the expected discounted present value of real profits $E_t \sum_{k=t}^{\infty} q_{t,k} d_k(i)$, where $q_{t,t} \equiv 1$, $q_{t,t+1} = \beta(c_t/c_{t+1})^\sigma$, and $q_{t,k} \equiv \prod_{j=t+1}^k q_{j-1,j}$ is the stochastic discount factor between periods t and $k > t$. In a symmetric equilibrium, all intermediate goods producing firms make the same decisions and the optimality condition becomes

$$\varphi \left(\frac{\pi_t}{\bar{\pi}} - 1 \right) \frac{\pi_t}{\bar{\pi}} = (1 - \theta) + \theta \Psi_t + \varphi E_t \left[q_{t,t+1} \left(\frac{\pi_{t+1}}{\bar{\pi}} - 1 \right) \frac{\pi_{t+1}}{\bar{\pi}} \frac{y_{t+1}}{y_t} \right].$$

In the absence of costly price adjustments (i.e., $\varphi = 0$), the real marginal cost of producing a unit of output equals $(\theta - 1)/\theta$, which is the inverse of the firm's markup of price over marginal cost.

The fiscal authority finances a constant level of discretionary spending, \bar{g} , through proportional taxes on capital and labor, seigniorage revenues, and by issuing one-period nominal government debt. The government's flow budget constraint is given by

$$m_t + b_t + \tau_t(w_t n_t + r_t^k k_{t-1}) = \bar{g} + (m_{t-1} + r_{t-1} b_{t-1})/\pi_t.$$

Following Leeper (1991), the monetary and fiscal authorities set policy according to

$$r_t = \bar{r}(\pi_t/\pi^*)^\phi \quad \text{and} \quad \tau_t = \bar{\tau}(b_{t-1}/b^*)^\gamma \exp(\varepsilon_{\tau,t}),$$

where π^* and b^* are the target levels of debt and inflation, ϕ and γ are parameters controlling the policy responses to inflation and debt, and $\varepsilon_{\tau,t} \sim i.i.d.N(0, \sigma_\tau^2)$.

The aggregate resource constraint is given by $c_t + i_t + \bar{g} = [1 - \varphi(\pi_t/\bar{\pi} - 1)^2/2]y_t$. Equilibrium is characterized by the household's and firm's optimality conditions, the government's budget constraint, the monetary and fiscal policy rules, and the aggregate resource constraint.

3.2.1 SOLVING THE MODEL The following instructions provide a complete description on how to solve the nonlinear model with time iteration and linear interpolation.

- There are seven structural parameters. Given an annual calibration, these parameters are set to $\beta = 0.9615$ (4% real interest rate), $\sigma = 1$, $\eta = 1$, $\kappa = 1$, $\theta = 7.66$ (15% price markup), $\delta = 0.1$, and $\varphi = 10$. The monetary policy parameter, $\phi = 1.5$, and the fiscal policy parameter, $\gamma = 0.15$, which guarantees a unique and bounded solution. Fiscal policy shocks are normally distributed with mean zero and standard deviation $\sigma_\tau = 0.001$.

Processors	No MEX	Interpolation using MEX	Equilibrium system in MEX
1	144.8 (1.0)	59.5 (2.4)	22.6 (6.4)
2	75.6 (1.9)	31.8 (4.6)	13.0 (11.1)
6	28.0 (5.2)	13.2 (11.0)	6.1 (23.7)

Table 2: NK model solution times—MEX and parallelization comparison (in seconds). Based on a state space of 2,401 nodes (7 points on each continuous state variable). We specify 10 realizations of ε_τ . The routines were computed with an Intel Xeon X5690 6-core processor (3.47GHz) operating 64-bit Windows 7. The values in parentheses are the speed gains relative to the slowest setup.

- The preference parameters χ and ν are pinned down by the steady state labor share, $\bar{n} = 0.33$, and the velocity of money, $v = 3.8$. Given the steady state tax rate, $\bar{\tau} = 0.21$, and the share of government spending to output, $\bar{g}/\bar{y} = 0.17$, the remaining steady state values have a closed-form solution.
- The set of state variables for the linearized model, m_{t-1} , r_{t-1} , b_{t-1} , and k_{t-1} , can be reduced to a_t , b_{t-1} , and k_{t-1} in the nonlinear model, where $a_t \equiv m_{t-1} + r_{t-1}b_{t-1}$ is real government liabilities. However, it is necessary to include the fiscal policy shock, $\varepsilon_{\tau,t}$, in the state to solve for the time t tax rate. Thus, the minimum state in the nonlinear model consists of $\{a_t, b_{t-1}, k_{t-1}, \varepsilon_{\tau,t}\}$.
- The NK model allows less flexibility for choosing policy functions. The presence of Rotemberg adjustment costs imply that inflation is a required policy to analytically solve for all time t variables. We specify labor as a policy function to pin down output. Additionally, we choose capital as a policy to minimize the number of computations. Given these policies and the state, all time t variables fall out naturally from the equilibrium conditions.
- We specify seven points on each continuous state variable, which implies 2,401 nodes. The number of grid points is subjective, but the more dense the grid, the greater the accuracy of the solution. In models this size, we recommend first solving the model with relatively sparse grids and gradually increasing the density of the grids until there is no noticeable change in the policy functions. One way to do this is to use the solution with a sparse grid to interpolate the policy functions on a denser grid.

3.2.2 SOLUTION TIMES Relative to the RBC model presented in [section 3.1](#), the NK model contains two features that increase computational time. First, there are two additional state variables. With seven grid points on each continuous state variable, this increases the number of nodes in the discretized state space by about 40 percent. Second, there are two additional policy functions, which double the amount of interpolation/extrapolation.

[Table 2](#) reports speed test results across the number of processors and the amount of computations performed with MEX/Fortran. Once again, introducing multi-core processing produces near-linear speed gains regardless of the involvement of MEX. We report three different levels of MEX involvement: 1. none of the computations are performed using MEX, 2. only the interpolation/extrapolation step is written in Fortran and called using MEX, and 3. the entire eqm function is written in Fortran and called by `csolve.m` as a MEX function. Programming the entire equilibrium system in Fortran significantly reduces computational time. With six cores, using MEX for the interpolation/extrapolation step cuts the computational time in half, but when the entire

Acronym	FM	FC	TL	TC
Iteration Method	Fixed-Point	Fixed-Point	Time	Time
Approx. Method	Monomial Basis	Chebyshev Basis	Linear Interp.	Chebyshev Interp.

Table 3: Alternative iteration and approximation methods.

equilibrium system is written entirely in Fortran the algorithm is additional 2.5 times faster. The biggest drawback with programming the entire equilibrium system in Fortran and compiling it with MEX is that it is tedious to debug and requires a model-specific Fortran function. The interpolation/extrapolation MEX functions, on the other hand, are general and can be used to solve different models. For users who are less familiar with Fortran, the additional start-ups costs may outweigh the significant speed advantage of using additional amounts of MEX.

In total, multi-core processing and MEX reduces computational time from 144 to only 6 seconds—a speed factor of nearly 24 in the NK model. When we compile the equilibrium system to the RBC model in MEX, computational time falls to 10 second, a speed factor of 12 over the slowest setup. The reason for the significant improvement in relative speed gains over the RBC model is due to the `for`-loop in the interpolation/extrapolation step over two additional policies and the `for`-loop required by the non-linear solver, `csolve.m`.⁸ Thus, these speed factors would be even larger with additional state variables or continuous stochastic variables.

Coding the entire time iteration/linear interpolation algorithm in Fortran and compiling it as a MEX function achieves further speed gains. With the exact same setup the model solves in about 0.5 seconds, which is fast enough to estimate nonlinear models with a particle filter. However, this approach imposes additional startup costs. First, it requires the user to learn how to call IMSL libraries, since few functions are intrinsic to Fortran. Second, the user would have to parallelize the code using either the Fortran MPI or OpenMP instead of the PCT. Fortunately, OpenMP is straightforward to implement, since it does not require explicit memory management like the MPI, but it cannot parallelize code across networked computers. Third, it requires the user to either program `csolve` in Fortran or use a built-in solver in the IMSL package such as `dnegnf`. Given these costs and the challenges of debugging the code, we advise programming the entire algorithm in Fortran only in cases where maximum speed gains are required. Regardless of the amount of computations we execute with MEX functions, we always use MATLAB as our interface, since it has a convenient GUI and it is easy to store data in structures, run simulations, and plot results.

4 COMPARISON WITH ALTERNATIVE SOLUTION TECHNIQUES

This section considers alternative solution techniques. We first solve the model using a log-linear approximation around its deterministic steady state. This method is very attractive because first-order Taylor approximations are straightforward to obtain, the solution to the equilibrium system is quick to compute, existence and uniqueness conditions are well established, and large models (in terms of the number of shocks and states) do not increase the computational burden. However, these benefits can come at a steep cost—inaccuracy outside of a close neighborhood around the deterministic steady state and an inability to solve models with inherent nonlinearities.

⁸Solution times are not directly comparable across models because the speed of convergence may differ. For example, the NK model is larger in terms of the number of policy functions and the size of the state space, but the solution time is similar to the RBC model since the RBC model takes four times as many iterations to converge.

Within the class of policy function iteration methods, we compare time iteration (our advocated approach) with fixed-point iteration [Judd (1998)]. The downside of time iteration is that it is costly to call a nonlinear solver on every node. Thus, the solution times can be slow relative to fixed-point iteration. Fixed-point iteration evaluates the policy functions at current and future values of the state to back out the policy function updates from the Euler equations. In sharp contrast, time iteration only evaluates the policy functions at future values of the state and requires a nonlinear solver to solve for the policy functions at current values of the state. Moreover, fixed-point iteration does not guarantee that the equilibrium system of equations is satisfied on each iteration to a specified tolerance level, making it potentially less stable than time iteration.

We also consider alternative approximation methods. Linear interpolation locally approximates the policy functions at each node in the state space. As an alternative to this local approximation method, we adopt projection methods, which build global approximations of the policy functions, but still rely on grid-based iterative techniques [Judd (1992, 1998)]. Projection methods postulate that the policies can be written as a function of a user-specified basis. We consider monomial and Chebyshev bases, whose coefficients are updated by minimizing the sum of squared residuals.

The accuracy of global solution methods depends on the choice of a basis function. Given a chosen basis, each iteration implies new least squares estimates of the coefficients that globally approximate the policy function. As alternatives to our advocated approach (TL), we apply fixed-point iteration with a monomial (FM) and Chebyshev polynomial (FC) basis and time iteration with Chebyshev interpolation (TC). Henceforth, we distinguish between methods using an acronym where the first letter is the iteration technique (time or fixed) and the second letter is the approximation method (linear interpolation, Chebyshev polynomial basis/interpolation, monomial basis). See table 3 for details. Solving for the policy functions requires calculating the value of these functions off the nodes (i.e., interpolation between nodes and extrapolation outside the grid), which is inaccurate when there is curvature in the policy functions and the grids on the state variables are sparse. Thus, orthogonal bases, such as Chebyshev polynomials, often yield more accurate solutions, but at a greater computational cost, which is increasing in the order of the polynomial used to approximate the policy functions.

4.1 LEAST SQUARES PROJECTION The following algorithm implements the least squares projection method, given a specified basis. The basis is evaluated at a particular state and forms the dependent variables used to obtain the least squares estimates. Let X denote the basis evaluated using the original discretized state space and X_t the basis evaluated at the time t state.

We obtain an approximation of the policy functions, $\hat{\Lambda}_t$ (an M -by- p matrix where M is the number of nodes and p the number of policy functions), using the least-squares estimates, $\hat{\eta}_t$.⁹ We use $\hat{\Lambda}_t$ to calculate the updated state variables, which forms X_t . $\hat{\Lambda}_{t+1}$ is a function of X_t and is used to calculate the $t + 1$ variables necessary to evaluate the terms within the expectation operators. The expectations are evaluated using Gauss-Hermite quadrature.

Once the expectations are evaluated, we calculate Λ_t as implied by the equilibrium system of equations. Given Λ_t , new estimates of the coefficients are obtained using the least squares estimator, $\hat{\eta}_t$. The algorithm iterates on these estimates. Since this method minimizes the sum of the squared residuals, the approximation error of the policy functions is minimized for a given basis. The least-squares estimates of the coefficients are updated using a convex combination of

⁹We obtain the initial least-squares estimates from the log-linear solution.

the old and new estimates, $\hat{\eta}_{t+1} = \lambda \hat{\eta}_{t'} + (1 - \lambda) \hat{\eta}_t$ for $\lambda \in [0, 1]$, which helps maintain stability, especially at the beginning of the algorithm. This iterative process continues until $|\hat{\eta}_{t'} - \hat{\eta}_t|$ reaches a pre-specified tolerance criterion.

4.1.1 MONOMIAL BASIS Following Heer and Maussner (2005), we choose the set of monomials corresponding to an n^{th} order Taylor approximation of the policy function. The set of basis functions is given by

$$\mathcal{P}_n^s = \left\{ (x_1^{k_1} \cdots x_s^{k_s}) \left| \sum_{i=1}^s k_i = j, k_i \geq 0, j = 0, 1, \dots, p \right. \right\},$$

where s is the number of state variables. The monomials are evaluated at each of the M nodes in the discretized state-space and stored column-wise in X . The approximated policy functions are given by $\hat{\Lambda}_t = X \hat{\eta}_t$. New estimates of the coefficients are obtained every iteration using the least squares estimator, $\hat{\eta}_{t'} = (X'X)^{-1} X' \Lambda_t$.

4.1.2 CHEBYSHEV POLYNOMIAL BASIS As an alternative to the monomial basis, we approximate the policy functions using a Chebyshev polynomial basis. The first two Chebyshev polynomials are $T_0(x) = 1$ and $T_1(x) = x$. The remaining polynomials are given by the recursive formulation, $T_j(x) = 2xT_{j-1}(x) - T_{j-2}(x)$, for $j \geq 2$. We use the MATLAB function `ChebyshevCoeff.m`, written by David Terr, to obtain the coefficients for each polynomial. The MATLAB function `chebpoly.m` outputs these coefficients and the corresponding powers of x , which allows us to construct the polynomials with element-by-element matrix operations.

`chebpoly.m` also outputs the corresponding zeros of the Chebyshev polynomials. Discretizing the state-space so the nodes coincide with the zeros of the polynomials minimizes the error of the approximating function. The zeros of the Chebyshev polynomials are given by

$$\bar{x}_{k_i} \equiv \cos \left(\frac{2k_i - 1}{2m_i} \pi \right),$$

where $k_i = 1, \dots, m_i$ and $i = 1, \dots, s$. Choose m_i as the number of points in the i^{th} dimension of the state space. Since $\bar{x}_{k_i} \in [-1, 1]$, transform the grid to the interval $[a_i, b_i]$ by applying

$$\bar{z}_{k_i} \equiv \frac{\bar{x}_{k_i}(b_i - a_i)}{2} + a_i.$$

The next step is to estimate the least squares coefficients of the approximating function.

Define $\bar{\Lambda}_{k_1, \dots, k_s}$ as the value of the approximating function, $\hat{\Lambda}(\bar{z}_{k_1}, \dots, \bar{z}_{k_s})$. We seek η_{j_1, \dots, j_s} that minimizes

$$\sum_{k_1=1}^{m_1} \cdots \sum_{k_s=1}^{m_s} \left[\bar{\Lambda}_{k_1, \dots, k_s} - \sum_{j_1=0}^{n_1} \cdots \sum_{j_s=0}^{n_s} \eta_{j_1, \dots, j_s} T_{j_1}(\bar{x}_{k_1}) \cdots T_{j_s}(\bar{x}_{k_s}) \right]^2$$

for $n_i \leq m_i$. This yields the least-squares estimator

$$\hat{\eta}_{j_1, \dots, j_s} = \frac{1 + \mathbf{1}(j_1 > 0)}{m_1} \cdots \frac{1 + \mathbf{1}(j_s > 0)}{m_s} \sum_{k_1=1}^{m_1} \cdots \sum_{k_s=1}^{m_s} \bar{\Lambda}_{k_1, \dots, k_s} T_{j_1}(\bar{x}_{k_1}) \cdots T_{j_s}(\bar{x}_{k_s})$$

and is the output of the MEX function `Fchebweights*`. The approximating function is evaluated in the MEX function `Fallcheb*` according to

$$\hat{\Lambda}(z_1, \dots, z_s) = \sum_{j_1=0}^{n_1} \cdots \sum_{j_s=0}^{n_s} \hat{\eta}_{j_1, \dots, j_s} T_{j_1}(x_1) \cdots T_{j_s}(x_s),$$

where $x_i = 2(z_i - a_i)/(b_i - a_i) - 1$.

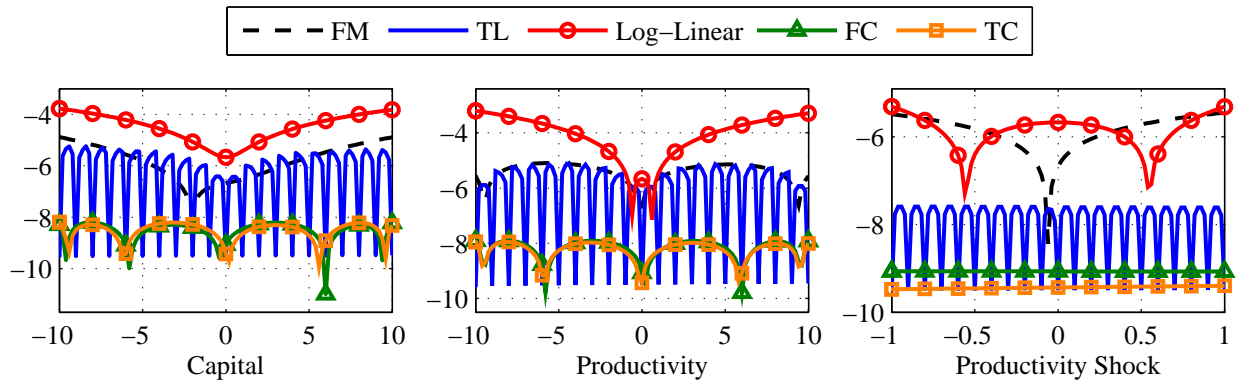


Figure 1: RBC model consumption Euler equation errors in base 10 logarithms. FM: Fixed point iteration with a monomial basis, TL: Time iteration with linear interpolation, FC: Fixed point iteration with Chebyshev polynomial basis, TC: Time iteration with a Chebyshev interpolation. FC and TC are based on a 4th order Chebyshev polynomial.

4.2 EULER EQUATION ERRORS Figure 1 compares errors for the consumption Euler equation across the log-linear, TL, TC, FM and FC solution methods. For ease of presentation, we show absolute errors in base 10 logarithms. This means that if the Euler equation error is -4 , the household makes an error of one consumption good for every 10,000 units of consumption goods. In general, our findings are in line with Aruoba et al. (2006).

The log-linear method (circle markers) is the least accurate. FM (dashed line) relies on an arbitrarily-specified basis and is the least computationally intensive global approximation method. To capture nonlinearities in the policy functions, we write the approximating functions as $\Lambda = X\eta$, where $\Lambda = c_t$ is the sole policy function, $X = [1 \ k \ z \ kz \ k^2 \ z^2]$ is a collection of basis functions, and η is a 6×1 matrix of coefficients. We considered several alternative collections of basis functions with higher-order terms, but found that they had little effect on the magnitude of the errors. We find that FM is as much as two orders of magnitude more accurate than the log-linear method.

Like the log-linear method, TL (solid line) solves for a *local* approximation of the policy functions. Thus, the errors can meet any user-specified tolerance criterion on each node, but lose accuracy when policy function values are interpolated between nodes or extrapolated outside the state space. The key difference from the log-linear method is that TL outperforms FM off the grid for the productivity shock, improving the accuracy of the consumption Euler equation by as much as two orders of magnitude. FC and TC are even more accurate than TL. Regardless of the iteration technique, using Chebyshev polynomials to approximate the policy functions consistently satisfies the consumption Euler equation given any user-specified tolerance level—both on and off the grid. However, this result is sensitive to the order of the polynomial. The errors with TC and FC are based on a 4th-order Chebyshev polynomial. We found that lower order polynomials significantly reduce accuracy, but higher order polynomials only marginally improve accuracy.

The NK model contains five state variables—lagged real debt, real money balances, the nominal interest rate, capital, and the fiscal policy shock—and three policy functions—labor, inflation, and capital. In figure 2, we report errors for the consumption Euler equation, the firm pricing equation, and the bond Euler equation. For ease of presentation, we restrict our attention to errors as a function of capital and the tax shock.¹⁰ Although this model is more complicated than the RBC model, the ordering of the Euler equation errors remains unchanged. Off the grid, TL consistently performs two orders of magnitude better than log-linear methods. TC and FC further increases accuracy, reducing Euler equation errors by an additional two orders of magnitude.

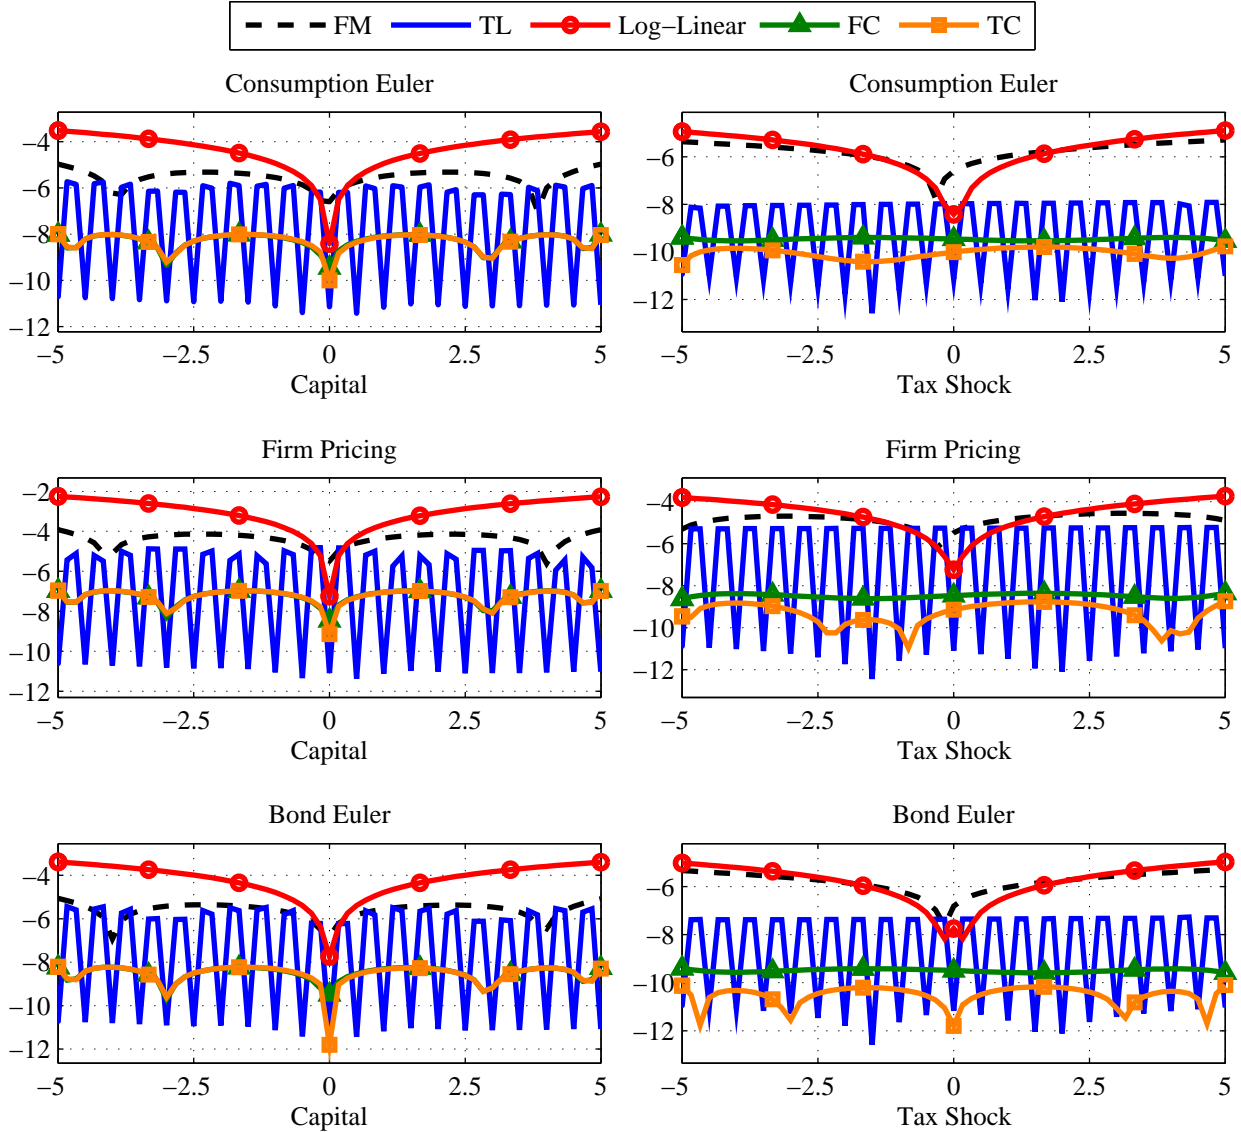


Figure 2: NK model Euler equation errors in base 10 logarithms. FM: Fixed point iteration with a monomial basis, TL: Time iteration with linear interpolation, FC: Fixed point iteration with a Chebyshev polynomial basis, TC: Time iteration with Chebyshev interpolation. FC and TC are based on a 4th order Chebyshev polynomial.

¹⁰Alternative state variables do not impact the ordering of the Euler equation errors. We decided not to average across the linear state to obtain a better comparison between the linear and nonlinear solution techniques.

Method	RBC Model		NK Model	
	No MEX	Basis in MEX	No MEX	Basis in MEX
FM	0.5	N/A	0.9	N/A
FC	27.2	6.5	302.0	31.2
TL	27.3	20.6	37.8	22.5
TC	45.9	23.9	508.2	60.8

Table 4: RBC and NK model solution times across the alternative solution methods (in seconds). The RBC solution is based on a state space of 1,681 nodes (41 points on k_{t-1} , 41 points on z_t) and 10 realizations of $\varepsilon_{z,t+1}$. The NK solution is based on a state space of 2,401 nodes (7 points on a_{t-1} , b_{t-1} , k_{t-1} , and $\varepsilon_{\tau,t}$) and 10 realizations of $\varepsilon_{\tau,t+1}$. FC and TC are based on a 4th order Chebyshev polynomial. The routines were computed with an Intel Xeon X5690 6-core processor (3.47GHz) operating 64-bit Windows 7. All 6 locally available processors were used in every speed test. FM: Fixed point iteration with a monomial basis, TL: Time iteration with linear interpolation, FC: Fixed point iteration with a Chebyshev polynomial basis, TC: Time iteration with Chebyshev interpolation.

4.3 SPEED COMPARISONS Table 4 reports solution times for the RBC and NK models across the alternative solution methods. Linear methods are fast and easy to apply, and there are numerous toolboxes for obtaining solutions. However, these methods are less accurate and unable to handle intrinsic nonlinearities or capture potentially important expectational effects (for example, in models with Markov-switching policy parameters or binding constraints.)

Figures 1 and 2 indicate that TC and FC offer a clear increase in accuracy (as much as two orders of magnitude), but table 4 shows that greater accuracy comes at the expense of a greater computational burden in terms of running time and implementation. Chebyshev interpolation is more intensive than linear interpolation, which MEX alleviates. The MEX implementation of FC is relatively quick and is four orders of magnitude more accurate than FM, but the speeds of both FC and TC relative to TL exponentially decrease for larger models and higher order Chebyshev polynomials. For the RBC and NK models, fixed-point iteration is faster than time iteration. However, fixed-point iteration solution times depend on the policy function update weight, λ . These simple models permit $\lambda = 1$, but complex models may require a lower λ to maintain stability of the algorithm. Time iteration does not suffer from this instability since the nonlinear solver minimizes error each iteration. Our findings suggest that TL provides the best balance between speed and accuracy, offering a 41 percent (63 percent) speed decrease from TC in the RBC (NK) model with a modest reduction in accuracy. TL is also more robust than its alternatives, since it is stable in every macroeconomic model and calibration we have tested.

4.4 MAXIMUM RESIDUAL AND EULER EQUATION ERRORS INTEGRAL Following Aruoba et al. (2006) and Caldara et al. (2012), in tables 5 we provide the integral of the Euler equation errors and the maximum residual for each solution method. These statistics provide complementary measures of accuracy and are defined over a square that is ± 10 percent on either side of steady state capital and productivity in the RBC model and ± 5 percent on either side of steady state capital and the tax shock in the NK model. We simulate each model for 100,000 periods from its stochastic steady state and create a distribution by counting the number of realizations in evenly spaced intervals and dividing by the simulation length.

The maximum Euler equation errors correspond to the largest errors plotted in figures 1 (RBC model) and 2 (NK model). Once again, we report base 10 logarithms of the errors. There are two distinct groups. Listed from least to most accurate, the log-linear method, FM, and TL have the

Method	Capital	Productivity
FM	-8.4 (-4.9)	-8.7 (-5.1)
TL	-8.0 (-5.2)	-8.8 (-5.1)
Log-linear	-7.2 (-3.8)	-8.4 (-3.2)
FC	-10.3 (-8.2)	-11.3 (-7.9)
TC	-10.5 (-8.2)	-11.3 (-7.9)

(a) RBC model

Method	Capital			Tax Shock		
	Capital FOC	Firm Pricing	Bond FOC	Capital FOC	Firm Pricing	Bond FOC
FM	-9.5 (-5.0)	-8.4 (-3.9)	-9.8 (-5.1)	-9.4 (-5.3)	-8.3 (-4.5)	-9.6 (-5.3)
TL	-14.1 (-5.7)	-14.0 (-4.8)	-14.1 (-5.4)	-11.1 (-7.9)	-8.4 (-5.2)	-10.5 (-7.3)
Log-linear	-11.4 (-3.5)	-10.2 (-2.2)	-10.7 (-3.4)	-11.0 (-4.9)	-9.6 (-3.7)	-10.6 (-5.0)
FC	-12.4 (-8.0)	-11.4 (-7.0)	-12.4 (-8.2)	-12.2 (-9.4)	-11.2 (-8.4)	-12.3 (-9.4)
TC	-13.0 (-8.0)	-12.1 (-7.0)	-14.8 (-8.2)	-12.8 (-9.8)	-11.9 (-8.7)	-14.1 (-10.1)

(b) NK model

Table 5: Euler equation error integrals (maximums). Values in base 10 logarithms. FM: Fixed point iteration with a monomial basis, TL: Time iteration with linear interpolation, FC: Fixed point iteration with a Chebyshev polynomial basis, TC: Time iteration with Chebyshev interpolation. FC and TC are based on a 4th order Chebyshev polynomial.

largest errors. FC and TC perform equally well and provide a clear improvement in accuracy.

The integrals take into consideration the frequency at which the errors occur and provide measures of the welfare loss associated with a particular solution method. Compared to the maximum residual, the relative accuracy of the various methods does not change. These integrals highlight a smaller difference in accuracy between TL and FM along the simulated path than suggested by the maximum error over the entire interval. However, FC and TC still show a clear comparative advantage in accuracy with at least a 4th order Chebyshev polynomial.

5 NEW KEYNESIAN MODEL WITH RECURSIVE PREFERENCES

This section considers a cashless version of the NK model laid out in section 3.2, but where households have preferences that distinguish between risk aversion and intertemporal substitution. This demonstrates the flexibility of the time iteration/linear interpolation algorithm and shows that our suite can be used to study asset pricing facts. Following Giovannini and Weil (1989) and Epstein and Zin (1989, 1991), we adopt a recursive structure for intertemporal utility given by

$$U(u_t, E_t U_{t+1}^{1-\eta}) = \left\{ (1-\beta)u_t^{(1-\eta)/\chi} + \beta(E_t U_{t+1}^{1-\eta})^{1/\chi} \right\}^{\chi/(1-\eta)}, \quad (1)$$

where η determines relative risk aversion, σ is the elasticity of intertemporal substitution, and $\chi = (1-\eta)/[1-\sigma^{-1}]$. Time- t utility is given by

$$u_t \equiv u(c_t, n_t) = c_t^\nu (1-n_t)^{1-\nu}, \quad \nu \in (0, 1). \quad (2)$$

The household's choices are constrained by

$$c_t + i_t + b_t = (1-\tau_t)(w_t n_t + r_t^k k_{t-1}) + r_{t-1} b_{t-1} / \pi_t + d_t, \quad (3)$$

$$k_t = (1-\delta)k_{t-1} + i_t. \quad (4)$$

Given prices, the representative household chooses a sequence of quantities, $\{c_t, n_t, B_t, k_t\}_{t=0}^{\infty}$, to maximize (1) subject to (2)-(4). Optimality yields the following first order conditions

$$\begin{aligned} (1 - \tau_t)w_t &= \frac{1 - \nu}{\nu} \frac{c_t}{1 - n_t}, \\ 1 &= r_t E_t \{q_{t,t+1} / \pi_{t+1}\}, \\ 1 &= E_t \{q_{t,t+1} [(1 - \tau_{t+1})r_{t+1}^k + (1 - \delta)]\}, \end{aligned}$$

where the stochastic discount factor, q , is given by

$$q_{t,t+1} = \beta \left(\frac{u_{t+1}}{u_t} \right)^{\frac{1-\eta}{\chi}} \frac{c_t}{c_{t+1}} \left(\frac{V_{t+1}^{1-\eta}}{E_t V_{t+1}^{1-\eta}} \right)^{1-\frac{1}{\chi}}$$

and, at optimum, the value function, V , can be written

$$V_t = \left\{ (1 - \beta)u_t^{(1-\eta)/\chi} + \beta[E_t V_{t+1}^{1-\eta}]^{1/\chi} \right\}^{\chi/(1-\eta)}.$$

When $\eta = 1/\sigma$, the value function, V_t , disappears from the first order conditions. Moreover, the constant of relative risk aversion, η , only alters the dynamics of higher order approximations, since, to a first-order, the term $(V_{t+1}^{1-\eta}/E_t V_{t+1}^{1-\eta})^{1-1/\chi}$ cancels out in expectation. The firm's problem and the policy specification are identical to [section 3.2](#).

The model contains four state variables—lagged real debt, the nominal interest rate, capital, and the fiscal policy shock—and four policy functions—labor, inflation, capital, and the time- t value of the optimal value function. [Figure 3](#) reports errors for the consumption Euler equation, the bond Euler equation, the firm pricing equation, and the household's optimal value function. Once again, we restrict our attention to errors as a function of capital and the tax shock and report them in base 10 logarithms. Our findings are in line with [Caldara et al. \(2012\)](#).

Outside of a close neighborhood of steady state, the log-linear method continues to perform at least two orders of magnitude worse than its nearest nonlinear competitor. However, the accuracy of the FM solution technique increases relative to the canonical NK model, performing equally as well as the TL solution method. TC and FC remain the most accurate, improving accuracy by an additional two orders of magnitude over FM and TL.

6 NEW KEYNESIAN MODEL WITH REGIME SWITCHING

This section adds monetary and fiscal policy switching to the canonical NK model, but where the fiscal authority only has access to lump-sum taxes to isolate the expectational effects of switching regimes. Unlike the previous examples, the discontinuity in the monetary and fiscal policy rule implies that linear interpolation is more accurate than Chebyshev approximation, which cannot accurately approximate the policy functions with a low order polynomial. The policy rules are

$$\begin{aligned} r_t &= \bar{r}(\pi_t/\pi^*)^{\phi(s_t)} \exp(\varepsilon_{r,t}) \\ \tau_t &= \bar{\tau}(b_{t-1}/b^*)^{\gamma(s_t)} \exp(\varepsilon_{\tau,t}), \end{aligned}$$

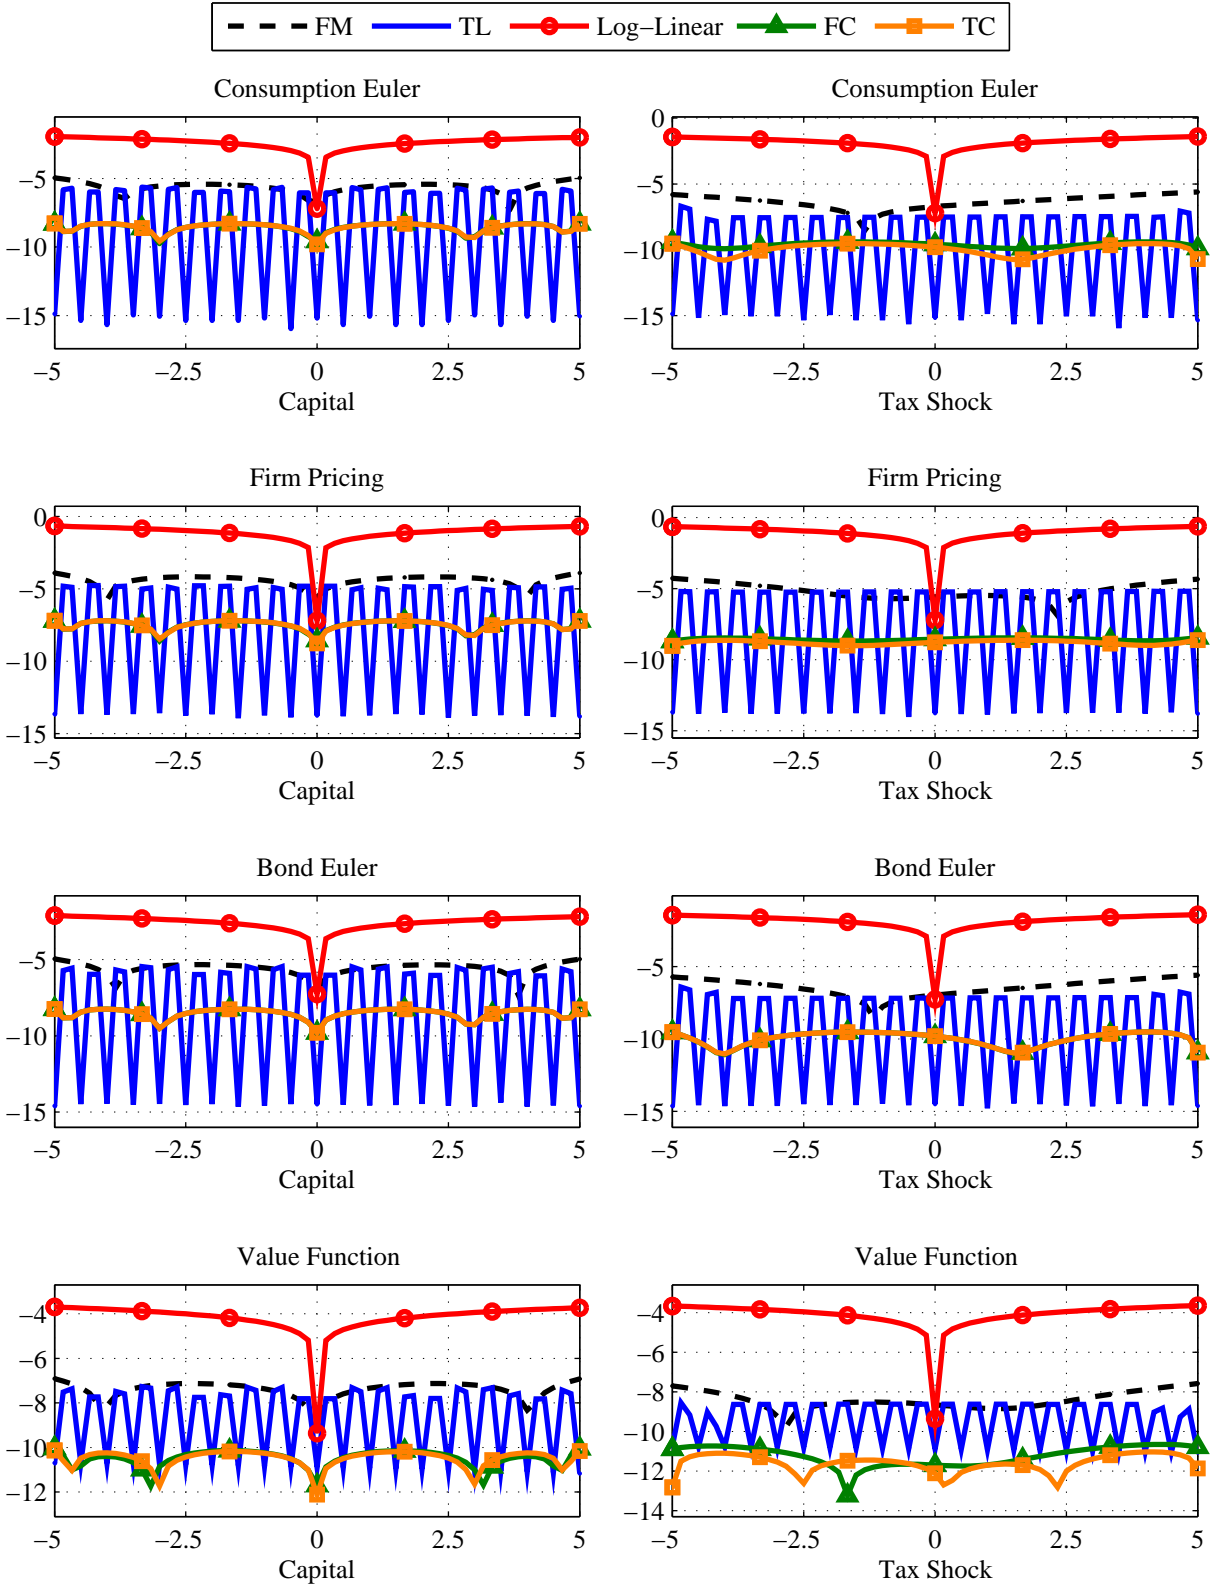


Figure 3: NK model with recursive preferences Euler equation errors in base 10 logarithms. FM: Fixed point iteration with a monomial basis, TL: Time iteration with linear interpolation, FC: Fixed point iteration with Chebyshev interpolation, TC: Time iteration with a Chebyshev polynomial basis.

where $s_t \in \{1, 2\}$ and the regime-dependent reaction coefficients are given by

$$\phi(s_t) = \begin{cases} \phi & \text{for } s_t = 1, \\ 0 & \text{for } s_t = 2, \end{cases} \quad \gamma(s_t) = \begin{cases} \gamma & \text{for } s_t = 1, \\ 0 & \text{for } s_t = 2. \end{cases}$$

The policy mix evolves according to a first-order two-state Markov chain given by

$$\begin{bmatrix} \Pr[s_t = 1 | s_{t-1} = 1] & \Pr[s_t = 2 | s_{t-1} = 1] \\ \Pr[s_t = 2 | s_{t-1} = 1] & \Pr[s_t = 2 | s_{t-1} = 2] \end{bmatrix} = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix}.$$

Following Leeper (1991), we label state 1 as active monetary and passive fiscal policy (AM/PF) and state 2 as passive monetary and active fiscal policy (PM/AF).

6.1 SOLUTION TECHNIQUE AND EULER EQUATION ERRORS The two policy regimes imply very different linear solutions. Thus, using the linear solution in only one of the states does not provide a good initial conjecture for both policy states. Instead, we use a linear combination of the linear solution in each state with weights equal to the transition probabilities to form a guess for the full nonlinear model. Although this is not required in all regime switching models, we found this is the best approach because it approximates the expectational effect of switching between regimes.

We interpolate/extrapolate for every permutation of the discretized continuous stochastic variables, $\varepsilon_{\tau,t}$ and $\varepsilon_{r,t}$. We facilitate this with a nested loop in `allterp*` that outputs a matrix of all possible realizations of the policy variables at time $t + 1$, where the rows correspond to the values of $\varepsilon_{\tau,t}$ and the columns correspond to the values of $\varepsilon_{r,t}$.

To evaluate expectations, we perform numerical integration using a two-step process that first applies the Trapezoid rule to each continuous stochastic variable and subsequently applies Markov-Chain integration to the discrete stochastic variable. To integrate across the continuous stochastic variables, first replicate the vector of $\varepsilon_{\tau,t}$ weights across the m realizations of $\varepsilon_{r,t}$ to create an $n \times m$ weighting matrix. Then weight each expectation and apply the Trapezoid rule across the $\varepsilon_{\tau,t}$ -dimension (the rows) to collapse each expectation to a vector of realizations. Finally, weight each of these outcomes using the $\varepsilon_{r,t}$ weights and once again apply the Trapezoid rule. This step produces expectations conditional on the realizations of the discrete stochastic variable, s_t . To integrate across each of these outcomes, simply weight each conditional expectation by its likelihood and sum across all realizations (see the online appendix for details). With three or more continuous stochastic variables, the product and trapezoid rules become exponentially costly to evaluate. In this case, we recommend using monomial rules [Judd (1998); Stroud (1971)].

Figure 4 compares the accuracy of the TM, TL, and TC solutions to the New Keynesian model with switching in the policy parameters.¹¹ We specify $p_{11} = p_{22} = 0.8$. We plot the residuals for the AM/PF ($s_t = 1$) and PM/AF ($s_t = 2$) regimes, which are roughly the same. With a 4th-order polynomial, TC is less accurate than TL. This is in sharp contrast with our findings in the non-switching NK model, where Chebyshev interpolation is consistently more accurate than linear interpolation. This is because Chebyshev interpolation is a global method and cannot capture the discontinuity as well as linear interpolation, which is a local method.

¹¹We only provide the consumption Euler equation errors, since the firm pricing and bond Euler equation errors imply the same qualitative results. The other errors are available from the authors upon request.

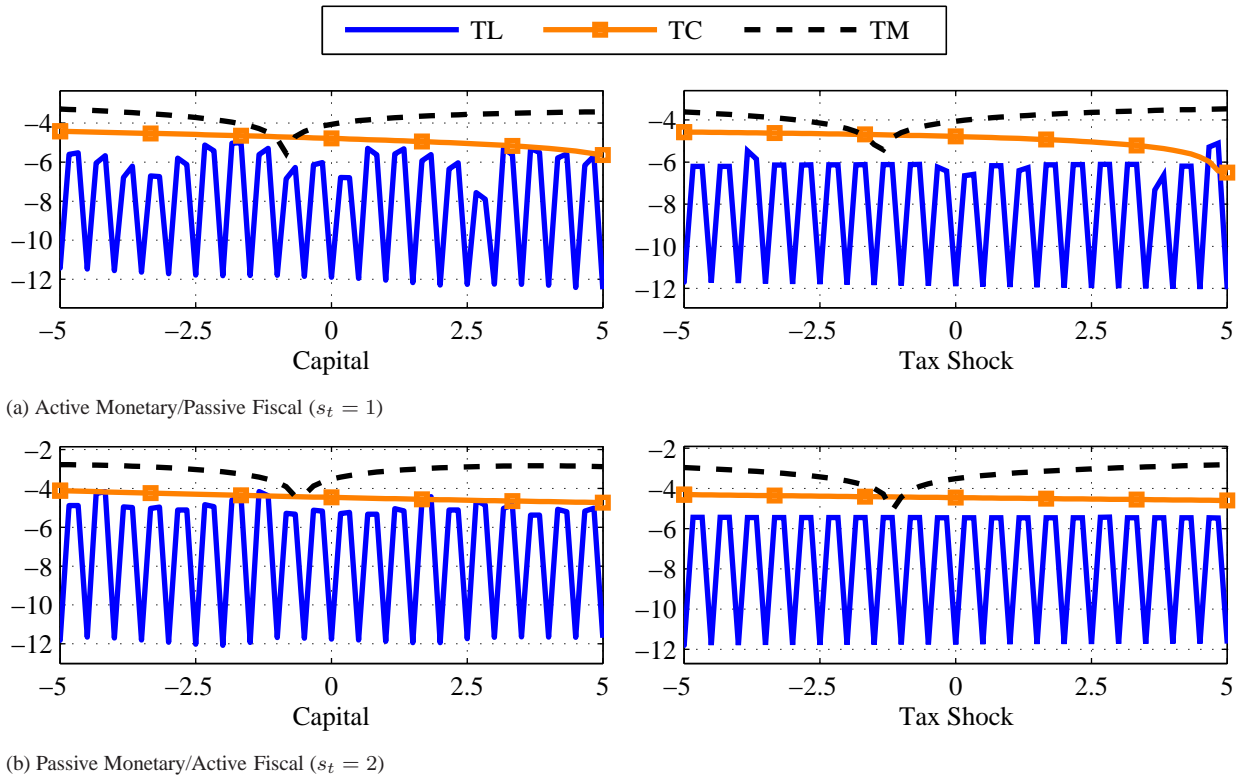


Figure 4: NK model with regime switching in the policy parameters consumption Euler equation errors in base 10 logarithms. TL: Time iteration with linear interpolation, TC: Time iteration with Chebyshev interpolation, TM: Time iteration with a monomial basis. TC is based on a 4th order Chebyshev polynomial and $p_{11} = p_{22} = 0.8$.

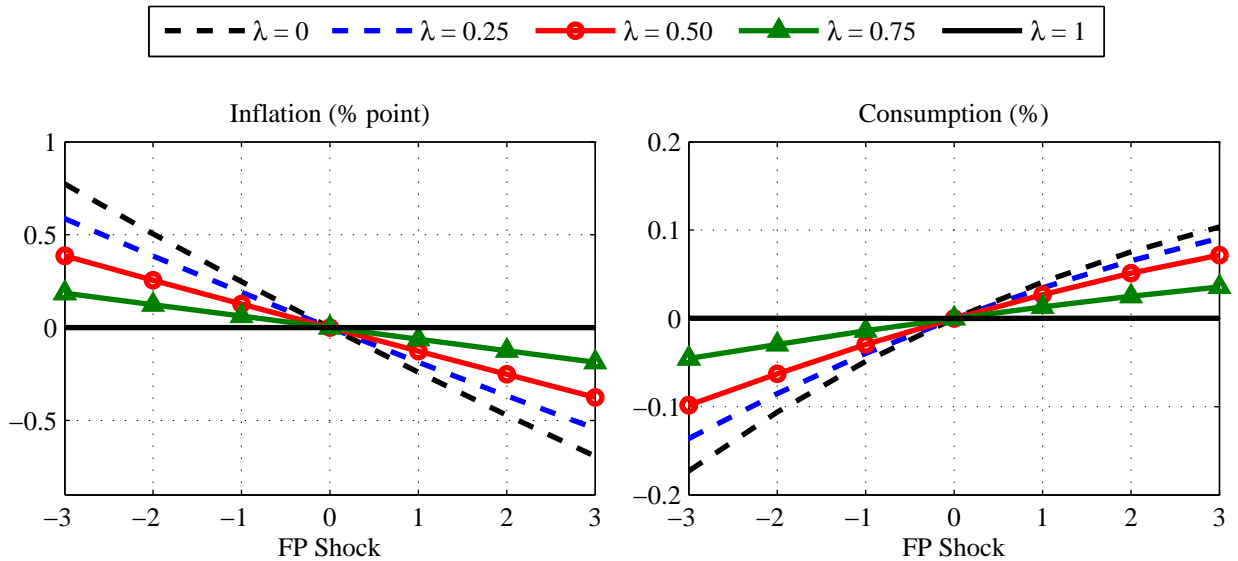


Figure 5: Policy functions for the NK model with regime switching in the policy parameters, based on a time iteration/linear interpolation technique (TL). λ measures the average duration of time spent in the active monetary/passive tax regime (AM/PF, $s_t = 1$) in the ergodic distribution.

6.2 EXPECTATIONAL EFFECTS The effect of monetary and fiscal policy shocks is dependent on the probabilities in the transition matrix. The average duration of time spent in the active monetary/passive tax regime (AM/PF, $s_t = 1$) in the ergodic distribution is given by

$$\xi = \frac{1 - p_{22}}{2 - p_{11} - p_{22}}.$$

Following Chung et al. (2007), [figure 5](#) shows how the policy functions in the AM/PF regime change as ξ declines, which demonstrates the nonlinearities that the TL solution method captures. In the passive monetary/active fiscal regime (PM/AF, $s_t = 2$), a tax shock elicits no response from the fiscal authority and the price level adjusts to stabilize debt. When $\xi = 1$, there is no chance of moving to the PM/AF regime. Thus, Ricardian equivalence holds and tax shock have no effect on the policy functions. As the expected duration of time spent in the PM/AF regime rises, negative tax shocks increase inflation, which, with costly price adjustments, lowers output and consumption.

TL is particularly attractive for solving models with discrete random variables. As [figure 5](#) illustrates, this method captures the curvature in the policy functions, even with linear interpolation. Spectral methods increase accuracy, but rely on smoothness in the policy functions and are far less stable and less accurate when applied to models with regime-switching. We were unable to achieve convergence using TC when $\xi < 0.5$ or when the polynomial order exceeds 4.

7 ZERO LOWER BOUND ON THE NOMINAL INTEREST RATE

Most of the existing literature that studies the zero lower bound (ZLB) uses log-linearized New Keynesian models.¹² However, using log-linearized models creates the potential for large approximation errors [Braun et al. (2012); Fernández-Villaverde et al. (2012)]. This section shows that time iteration with linear interpolation is a stable and accurate way to solve models with an occasionally binding constraint. Chebyshev interpolation is far less stable and less accurate, since it cannot accurately capture the kink in the policy functions that the ZLB constraint imposes.

We augment the simple Taylor rule in the New Keynesian model to include an automatic stabilizer component, which increases the frequency of ZLB events. Given the inequality constraint on the nominal interest rate, the monetary policy rule becomes

$$r_t = \max\{1, \bar{r}(\pi_t/\pi^*)^{\phi_\pi}(y_t/\bar{y})^{\phi_y}\}. \quad (5)$$

To simplify the analysis, we assume the economy is at its cashless limit and the government balances its budget each period by levying lump-sum taxes (government debt is in zero net supply). The economy is subject to demand and supply shocks. The demand shock propagates through the discount factor and productivity is a proxy for the supply shock. They evolve according to

$$\begin{aligned} \beta_t &= \beta(\beta_{t-1}/\beta)^{\rho_\beta} \exp(\varepsilon_{\beta,t}), \\ z_t &= \bar{z}(z_{t-1}/\bar{z})^{\rho_z} \exp(\varepsilon_{z,t}), \end{aligned}$$

where $\varepsilon_x \sim N(0, \sigma_x^2)$, $x \in \{\beta, z\}$.

Richter and Throckmorton (2013) show that the calibration of the model (including the parameters of the stochastic processes) impacts determinacy, since the boundary of the determinacy

¹²Recent papers that solve the nonlinear model include Aruoba and Schorfheide (2013); Basu and Bundick (2012); Fernández-Villaverde et al. (2012); Gavin et al. (2013); Gust et al. (2012); Judd et al. (2011); Mertens and Ravn (2013).

region imposes a clear tradeoff between the expected frequency and average duration of episodes at the ZLB. We calibrate the model to ensure a determinate solution. Unless specified, the parameters are identical to [section 3.2.1](#). Given a quarterly calibration, $\beta = 0.99$ (1% real interest rate), $\delta = 0.025$, and $\varphi = 77.67$ (75 percent of firms cannot adjust prices each period). The productivity and discount factor processes are semi-persistent with autoregressive coefficients, $\rho_z = \rho_\beta = 0.75$. The shocks are normally distributed with mean zero and standard deviations $\sigma_z = 0.005$ and $\sigma_\beta = 0.0025$. The monetary authority response coefficients, ϕ_π and ϕ_y , are set to 1.5 and 0.1. We normalize steady state output to 1 and solve for the implied steady state productivity level.

Although the ZLB imposes a kink in the policy functions, the linear solution to the NK model *without* a ZLB constraint typically provides a good initial conjecture for the constrained nonlinear model. The only difference from the solution procedure described in [section 3.2.1](#) is that we set the gross nominal interest rate equal to 1 on any node where the Taylor rule implies $r < 1$. The nonlinear solver searches for policy function values on each node that satisfy the ZLB constraint.

Both the linear and nonlinear models contain a unique set of state variables—capital, productivity, and the discount factor. We discretize each state variable with 41 points, which implies 68,921 nodes. Once again, the number of grid points is subjective, but denser grids are required to accurately locate the kinks in the policy functions. There is still flexibility when choosing the policy functions. We specify policies over inflation (required), labor, and capital.

7.1 ZLB RESULTS We solve the constrained nonlinear model using TL and TC. Variants of both methods have been recently used in the literature, although TC is far more common. TL is just as flexible as the model without a ZLB. It converges for a broad range of model parameterizations and grid specifications. Although TC is stable and accurate when the ZLB does not bind, it is very unstable when the ZLB binds on even a small percentage of the state space, since it is impossible for the algorithm to accurately locate the kinks in the policy functions with a continuous approximation. Even with a high order Chebyshev polynomial, TC did not converge for most parameterizations and grid specifications. In cases where TC did converge, the solutions implied by the TL and TC algorithms are very different. With a 5th order Chebyshev polynomial, the ZLB binds on 6.5 percent of the nodes in the state space, whereas the ZLB binds on only 2.1 percent of the nodes with TL. Wider bounds on any of the state variables will increase these percentages.

[Figure 6](#) shows a productivity-discount rate cross-section of the state space. We fix capital at 6 percent above steady state to concentrate on a region of the state space where the ZLB occasionally binds and there is a kink in the policy functions. Both solution methods imply that the ZLB binds in the upper right-hand corner of the state space, since higher productivity (positive supply shock) and more patient households (negative demand shock) decrease inflation and drive the nominal interest rate to its ZLB.¹³ This figure makes clear that there are stark differences between the two solution methods. The TL solution implies that the ZLB binds on only 2 percent of the nodes in this cross section (star markers), while the TC solution implies that nearly 30 percent of the nodes bind (square markers). These difference have important implications for the conditional and unconditional probabilities of hitting and staying at the ZLB.

[Figure 7](#) reports errors for the consumption Euler equation, firm pricing equation, and bond Euler equation as a function of the capital, productivity, and the discount factor states in base 10 logarithms. We fix capital, productivity, and the discount factor at 6, 7, and 1.5 percent above their

¹³For a more detailed discussion on how productivity and discount factor shocks impact the model solutions and dynamics in models with and without capital see Gavin et al. (2013).

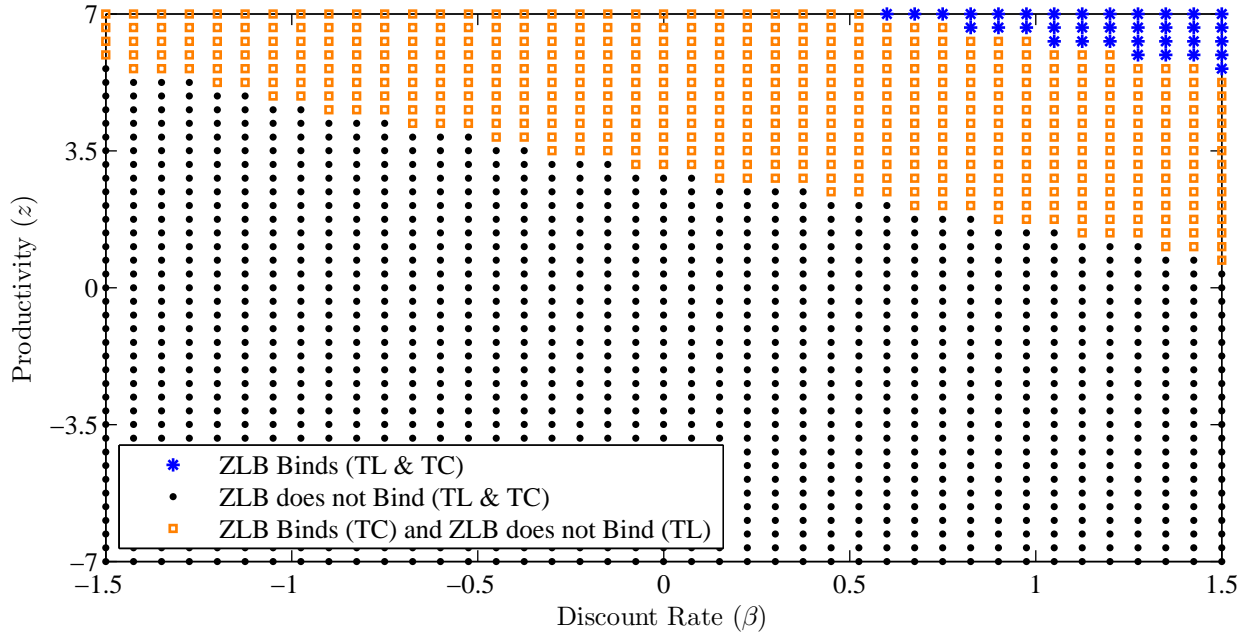


Figure 6: Productivity-discount factor cross-section of the state space where capital is 6 percent above steady state.

steady state values. The darker (entire) shaded region indicates where the ZLB binds when the model is solved with TL (TC). Both methods are less accurate when the ZLB binds. However, TL outperforms TC with a 5th order Chebyshev polynomial regardless of whether the ZLB binds, which is consistent with the results in [section 6](#). Even in alternative cross-sections of the state space where the ZLB never binds, TL is always more accurate than TC. This is due to the kinks in the policy functions, which make it very difficult for a global approximation method, such as TC, to find a polynomial that well-approximates regions of the state space where the ZLB binds and does not bind. In the canonical models described above, TC and TL presented a clear tradeoff between accuracy and speed. In models with an inherent nonlinearity, such as those with occasionally binding constraints or Markov switching, TL is more accurate, faster, and more stable.

Accurately locating the kink of the policy functions has important implications for economic dynamics. [Figure 8](#) shows the consumption and labor policy functions based on the TL (solid line) and TC (dashed line) solution methods. Once again, we fix capital at 6 percent above steady state. The inequality constraint induces a kink in the policy functions that does not exist in other NK models, even when they contain real frictions. As the economy moves to higher productivity and discount factor states where the ZLB does not bind, output rises and inflation falls as firms revise prices downward. In states where the ZLB binds, the real interest rate rises sharply, which reduces demand. This causes output and inflation to fall. Although both solution method produces similar qualitative dynamics in states away from and at the ZLB, the quantitative dynamics of the model are drastically different. This figure makes clear that the TC solution method has trouble locating the kink in the policy functions, as the policy functions oscillate across the state space.

A byproduct of inaccurately capturing the kinks in the policy functions is that extrapolation will be very inaccurate. [Figure 9](#) demonstrates that a global approximation with Chebyshev polynomials will lead inaccurate expectations. The vertical dashed lines are the bounds of the discretized state space. Values of the approximating functions that lie outside these bounds are extrapolated,

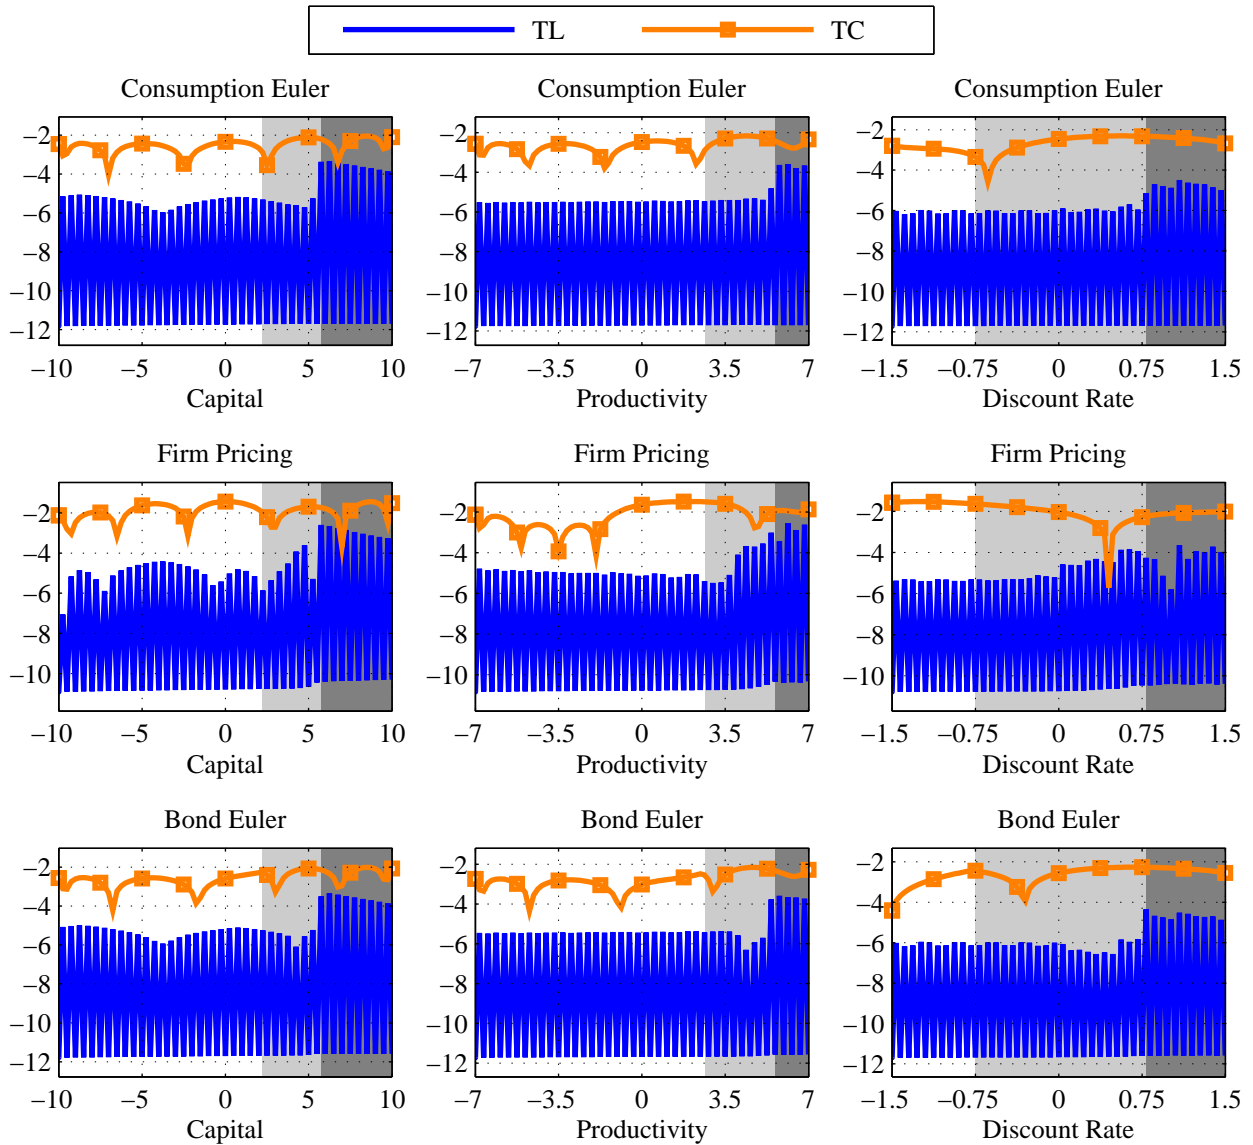


Figure 7: NK ZLB model Euler equation errors in base 10 logarithms. Unless specified, capital, productivity, and the discount factor are fixed at 6, 7, and 1.5 percent above steady state. TL: Time iteration with linear interpolation, TC: Time iteration with Chebyshev interpolation. TC is based on a 5th order Chebyshev polynomial. The dark (entire) shaded region indicates where the ZLB binds when the model is solved with TL (TC).

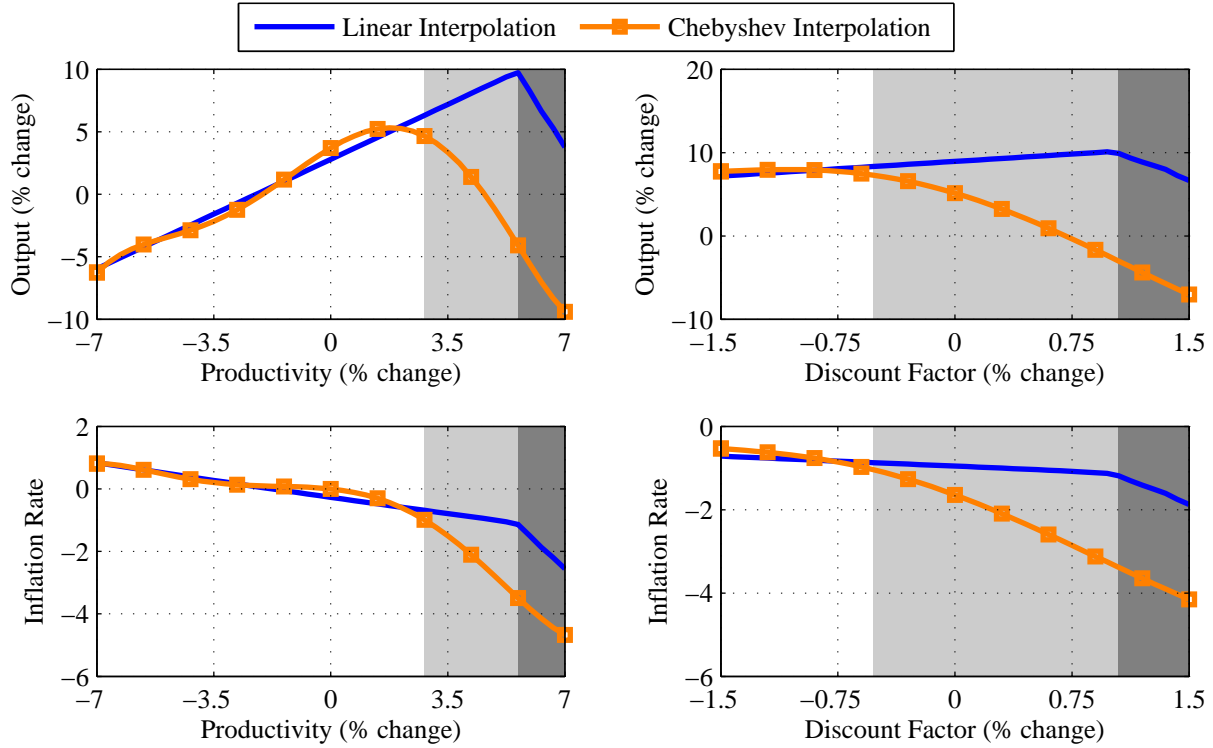


Figure 8: Policy functions based on the TL (solid line) and TC (dashed line) solution methods. Unless specified capital, productivity, and the discount factor are fixed at 6, 7, and 1.5 percent above steady state. The dark (entire) shaded region indicates where the ZLB binds when the model is solved with TL (TC).

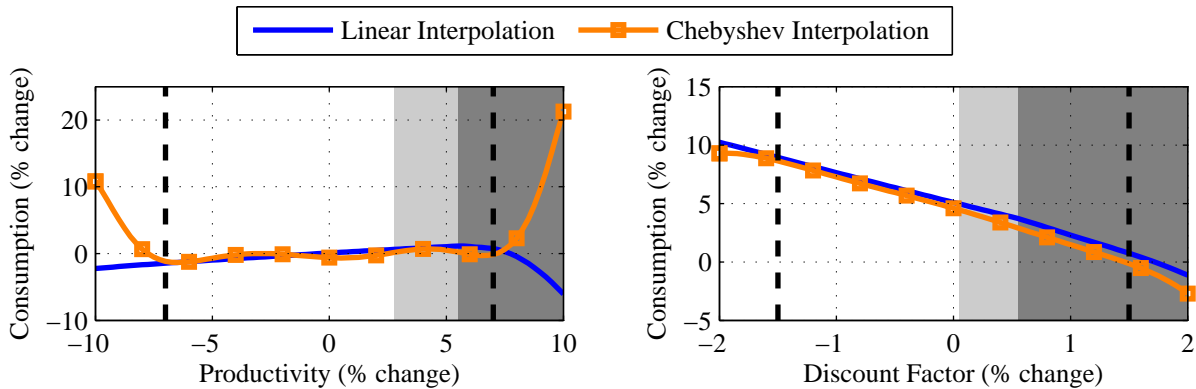


Figure 9: Policy functions based on the TL (solid line) and TC (dashed line) solution methods. Unless specified capital, productivity, and the discount factor are fixed at 6, 7, and 1.5 percent above steady state. The vertical dashed lines indicate the lower and upper bound of the discretized state space.

and will enter the expectation operator. If future variables are computed using extrapolated values, then significant error enters the expectation operators in the equilibrium system. This additional error is another cause of inaccuracy of Chebyshev approximation on the discretized state space.

8 CONCLUSION

Policy function iteration methods have long been known as a reliable way to solve dynamic models nonlinearly. They are particularly useful for studying the economic consequences of a wide variety of potential policy outcomes. Despite the considerable benefits of this algorithm, it suffers from several drawbacks. The most prohibitive feature of this algorithm is its reliance on grid-based techniques, which exponentially increases the size of the problem with the number of state variables and the number of continuous stochastic components. Perturbation solution methods are far less computationally expensive and are more appropriate for models that do not contain recurring regime change or equilibrium paths that deviate far from the deterministic steady state.

Another concern is whether these solution methods consistently satisfy transversality conditions, since it only iterates on the policy functions and has no formal mechanism for imposing these restrictions. As a safeguard, however, it is easy to simulate a model for thousands of periods and check that its average asset levels (e.g., capital, bonds, *etc.*) are convergent. Moreover, simulated paths in models that explicitly violate the transversality condition will typically diverge even if the algorithm converges. Although these exercises do not provide proof, they do provide reasonable confidence that transversality conditions are met.

Lastly, in models that contain significant curvature in the policy functions, it *can* be difficult to obtain initial conjectures. The best starting point to generate a guess for the nonlinear model is always the solution to the linearized model. In almost all cases, the linear solution provides a good enough guess ensure that the policy functions converge in the nonlinear model. In cases where it is not, we recommend two potential solutions: 1. use the linear solution as an initial conjecture for a variant of the nonlinear model, which may provide a better guess for the nonlinear model of interest, or 2. solve the model with a parameterization that induces less curvature in the policy function and slowly iterate on the parameter of interest.

Any numerical algorithm imposes direct costs onto the programmer. We reduce the costs associated with policy function iteration methods by providing a user-friendly suite of MATLAB functions that introduce multi-core processing and Fortran via MATLAB's executable function. Within the class of policy function iteration methods, we advocate using time iteration with linear interpolation. We apply this method to conventional RBC and NK models and carefully document how to chose policy functions, discretize the state space, interpolate/extrapolate future values, and perform numerical integration. The use of multi-core processing alongside optimized code that takes advantage of Fortran's comparative advantage at evaluating loops decreases solutions times by a factor of 8 in the RBC model and a factor of 24 in the NK model. Moreover, comparing time iteration with linear interpolation to alternative solution techniques demonstrates it is accurate, able to capture important nonlinearities, and robust to models with discontinuities.

REFERENCES

- ARUOBA, S. AND F. SCHORFHEIDE (2013): “Macroeconomic Dynamics Near the ZLB: A Tale of Two Equilibria,” Manuscript, University of Pennsylvania.
- 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, 2477–2508.
- BASU, S. AND B. BUNDICK (2012): “Uncertainty Shocks in a Model of Effective Demand,” NBER Working Paper 18420.
- BAXTER, M. (1991): “Approximating Suboptimal Dynamic Equilibria: An Euler Equation Approach,” *Journal of Monetary Economics*, 28, 173–200.
- BAXTER, M., M. CRUCINI, AND K. ROUWENHORST (1990): “Solving the Stochastic Growth Model by a Discrete-State-Space, Euler-Equation Approach,” *Journal of Business & Economic Statistics*, 8, 19–21.
- BI, H. (2012): “Sovereign Default Risk Premia, Fiscal Limits and Fiscal Policy,” *European Economic Review*, 56, 389–410.
- BI, H., E. M. LEEPER, AND C. LEITH (2013): “Uncertain Fiscal Consolidations,” *Economic Journal*, 123, 31–63.
- BRAUN, R. A., L. M. KÖRBER, AND Y. WAKI (2012): “Some Unpleasant Properties of Log-Linearized Solutions When the Nominal Rate is Zero,” FRB Atlanta Working Paper 2012-5a.
- CALDARA, D., J. FERNANDEZ-VILLAVERDE, J. RUBIO-RAMIREZ, AND W. YAO (2012): “Computing DSGE Models with Recursive Preferences and Stochastic Volatility,” *Review of Economic Dynamics*, 15, 188–206.
- CHUNG, H., T. DAVIG, AND E. M. LEEPER (2007): “Monetary and Fiscal Policy Switching,” *Journal of Money, Credit and Banking*, 39, 809–842.
- COLEMAN, II, W. J. (1991): “Equilibrium in a Production Economy with an Income Tax,” *Econometrica*, 59, 1091–1104.
- DATTA, M., L. MIRMAN, O. MORAND, AND K. REFFETT (2005): “Markovian Equilibrium in Infinite Horizon Economies with Incomplete Markets and Public Policy,” *Journal of Mathematical Economics*, 41, 505–544.
- DATTA, M., L. MIRMAN, AND K. REFFETT (2002): “Existence and Uniqueness of Equilibrium in Distorted Dynamic Economies with Capital and Labor,” *Journal of Economic Theory*, 103, 377–410.
- DAVIG, T. (2004): “Regime-switching Debt and Taxation,” *Journal of Monetary Economics*, 51, 837–859.

- DAVIG, T. AND E. M. LEEPER (2006): “Fluctuating Macro Policies and the Fiscal Theory,” in *NBER Macroeconomics Annual 2006, Volume 21*, ed. by D. Acemoglu, K. Rogoff, and M. Woodford, MIT Press, Cambridge, MA, 247–298.
- (2008): “Endogenous Monetary Policy Regime Change,” in *NBER International Seminar on Macroeconomics 2006*, ed. by L. Reichlin and K. D. West, MIT Press, Cambridge, MA, 345–391.
- DAVIG, T., E. M. LEEPER, AND T. B. WALKER (2010): “‘Unfunded Liabilities’ and Uncertain Fiscal Financing,” *Journal of Monetary Economics, Carnegie-Rochester Conference Series on Public Policy*, 57, 600–619.
- (2011): “Inflation and the Fiscal Limit,” *European Economic Review, Special Issue on Monetary and Fiscal Interactions in Times of Economic Stress*, 55, 31–47.
- DIXIT, A. K. AND J. E. STIGLITZ (1977): “Monopolistic Competition and Optimum Product Diversity,” *American Economic Review*, 67, 297–308.
- EPSTEIN, L. G. AND S. E. ZIN (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57, 937–69.
- (1991): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis,” *Journal of Political Economy*, 99, 263–86.
- FERNÁNDEZ-VILLAVERDE, J., G. GORDON, P. A. GUERRÓN-QUINTANA, AND J. RUBIO-RAMÍREZ (2012): “Nonlinear Adventures at the Zero Lower Bound,” NBER Working Paper 18058.
- GASPAR, J. AND K. JUDD (1997): “Solving Large-Scale Rational-Expectations Models,” *Macroeconomic Dynamics*, 1, 45–75.
- GAVIN, W. T., B. D. KEEN, A. W. RICHTER, AND N. A. THROCKMORTON (2013): “Global Dynamics at the Zero Lower Bound,” FRB of St. Louis Working Paper 2013-007B.
- GIOVANNINI, A. AND P. WEIL (1989): “Risk Aversion and Intertemporal Substitution in the Capital Asset Pricing Model,” NBER Working Paper 2824.
- GREENWOOD, J. AND G. W. HUFFMAN (1995): “On the Existence of Nonoptimal Equilibria in Dynamic Stochastic Economies,” *Journal of Economic Theory*, 65, 611–623.
- GUST, C., D. LÓPEZ-SALIDO, AND M. E. SMITH (2012): “The Empirical Implications of the Interest-Rate Lower Bound,” CEPR Discussion Paper 9214.
- HEER, B. AND A. MAUSSNER (2005): *Dynamic General Equilibrium Modelling: Computational Methods and Applications*, Springer.
- IRELAND, P. N. (1997): “A Small, Structural, Quarterly Model for Monetary Policy Evaluation,” *Carnegie-Rochester Conference Series on Public Policy*, 47, 83–108.

- JUDD, K. L. (1992): “Projection methods for solving aggregate growth models,” *Journal of Economic Theory*, 58, 410 – 452.
- (1998): *Numerical Methods in Economics*, Cambridge, MA: The MIT Press.
- JUDD, K. L. AND S.-M. GUU (1993): “Perturbation Solution Methods for Economic Growth Models,” in *Economic and Financial Modeling with Mathematica*, ed. by H. R. Varian, Springer, New York, 80–103.
- (1997): “Asymptotic methods for aggregate growth models,” *Journal of Economic Dynamics and Control*, 21, 1025–1042.
- JUDD, K. L., L. MALIAR, AND S. MALIAR (2011): “A Cluster-Grid Algorithm: Solving Problems with High Dimensionality,” Manuscript, Stanford University.
- KLEIN, P. (2000): “Using the Generalized Schur Form to Solve a Multivariate Linear Rational Expectations Model,” *Journal of Economic Dynamics and Control*, 24, 1405–1423.
- KUMHOF, M. AND R. RANCIERE (2010): “Inequality, Leverage and Crises,” IMF Working Paper 10-268.
- LEEPER, E. M. (1991): “Equilibria Under ‘Active’ and ‘Passive’ Monetary and Fiscal Policies,” *Journal of Monetary Economics*, 27, 129–147.
- MCCANDLESS, G. (2008): *The ABCs of RBCs: An Introduction to Dynamic Macroeconomic Models*, Cambridge, MA: Harvard University Press.
- MERTENS, K. AND M. O. RAVN (2013): “Fiscal Policy in an Expectations Driven Liquidity Trap,” Manuscript, Cornell university.
- MIRMAN, L., O. MORAND, AND K. REFFETT (2008): “A Qualitative Approach to Markovian Equilibrium in Infinite Horizon Economies with Capital,” *Journal of Economic Theory*, 139, 75–98.
- RICHTER, A. (2012): “The Fiscal Limit and Non-Ricardian Consumers,” Manuscript, Auburn University.
- RICHTER, A. AND N. THROCKMORTON (2012): “MATLAB/MEX/Fortran Toolbox,” Available online. http://auburn.edu/~awr0007/AWR_files/MatlabMEXtoolbox.zip.
- RICHTER, A. W. AND N. A. THROCKMORTON (2013): “The Zero Lower Bound: Frequency, Duration, and Determinacy,” Manuscript, Auburn University.
- ROTEMBERG, J. J. (1982): “Sticky Prices in the United States,” *Journal of Political Economy*, 90, 1187–1211.
- SCHMITT-GROHE, S. AND M. URIBE (2004): “Solving dynamic general equilibrium models using a second-order approximation to the policy function,” *Journal of Economic Dynamics and Control*, 28, 755–775.

SIMS, C. A. (2002): “Solving Linear Rational Expectations Models,” *Computational Economics*, 20, 1–20.

STROUD, A. (1971): *Approximate Calculation of Multiple Integrals*, Englewood Cliffs, NJ: Prentice Hall.

UHLIG, H. (1997): “A Toolkit for Analyzing Nonlinear Dynamic Stochastic Models Easily,” Manuscript, Center for Economic Research.

WALSH, C. E. (2010): *Monetary Theory and Policy*, 3ed., Cambridge, MA: The MIT Press.

A ONLINE APPENDIX (NOT FOR PUBLICATION)

A.1 BRIEF REVIEW OF THE THEORY This section briefly reviews the theory behind monotone operators as applied to DSGE models, using the results of Greenwood and Huffman (1995) (GH, hereafter). We do not convey any new theoretical results but simply demonstrate how the monotone map can be used to prove existence and uniqueness of an equilibrium. We follow GH closely because they proved existence of equilibrium in a very general setup. Moreover, as advocated by Datta et al. (2002), Datta et al. (2005), Mirman et al. (2008), the theoretical properties of the monotone map can be extended to more complex setups. Proofs of existence using monotone operators are constructive in the sense that the numerical algorithm is a byproduct of the proof, which only adds to the appeal of policy function iteration algorithm methods.

A.1.1 ECONOMIC ENVIRONMENT The economic environment is standard. The model consists of a continuum of measure one identical agents with preferences

$$E_0 \left[\sum_{t=0}^{\infty} \beta^t U(c_t) \right],$$

where the momentary utility function is assumed to be strictly increasing, strictly concave and twice differentiable, with $U'(0) = \infty$. The production function is given by

$$y_t = F(k_t, K_t, \eta_t),$$

where output, y_t , is produced with the individual agent's capital stock, k_t , the aggregate capital stock, K_t , and is subject to a random productivity shock, η_t . The productivity shock is assumed to be drawn from a Markov distribution function, $G(\eta_{t+1}|\eta_t)$, with bounded support. The production function is assumed to satisfy the Inada conditions, $\lim_{K \rightarrow 0} F_1(K, K, \eta) = \infty$, be strictly increasing and strictly concave in its first argument, and twice differentiable in its first two arguments. Moreover, GH also impose the following somewhat nonstandard assumptions:

1. $\exists \bar{K} \ni F(\bar{K}, \bar{K}, \eta) \leq \bar{K}$
2. $\forall K \in (0, \bar{K}], F_1(K, K, \eta) + F_2(K, K, \eta) \geq 0$ and $F_{11}(K, K, \eta) + F_{21}(K, K, \eta) < 0$

Assumption 1 places an upper bound on the level of output. Assumption 2 requires that the sum of the marginal products of the individual and aggregate capital stock be positive (along the equilibrium path $k = K$). As noted by GH, these assumptions are innocuous and hold for a wide range of economies.

The agent's dynamic programming problem is given by

$$V(k, K, \eta) = \max_{k'} \left\{ U(F(k, K, \eta) - k') + \beta \int V(k', K', \eta') dG(\eta'|\eta) \right\}, \quad (6)$$

where aggregate capital, K , has the following law of motion $K' = Q(K, \eta)$. Let the optimal policy function associated with (6) be given by $k' = q(k, K, \eta)$. By standard arguments, one can derive the corresponding Euler equation

$$U'(F(k, K, \eta) - k') = \beta \int U'(F(k', K', \eta') - k'') F_1(k', K', \eta') dG(\eta'|\eta)$$

A stationary equilibrium is a pair of functions, $k' = q(k, K, \eta)$ and $K' = Q(k, \eta)$, that satisfy optimality (i.e., solves (6)) and consistency, $q(K, K, \eta) = Q(K, \eta)$. We are now able to state the main proposition of GH.

Proposition 1 (GH, pg. 615). *There exists a nontrivial stationary equilibrium for the economy described above.*

The method of proof in GH follows that of Coleman (1991) and is our primary interest because it uses Euler equation iteration and properties of monotone operators. For these reasons we repeat the proof here. Let the sequence of aggregate laws of motion, $\{H^j(K, \eta)\}_{j=0}^{\infty}$, evolve according to $H^0(K, \eta) \equiv 0$, and let $H^{j+1}(K, \eta)$ for $j \geq 0$ be defined as the solution for x in the Euler equation

$$U'(F(K, K, \eta) - x) = \beta \int U'(F(x, x, \eta') - H^j(x, \eta')) F_1(x, x, \eta') dG(\eta'|\eta). \quad (7)$$

(7) defines a sequential operator mapping H^j into H^{j+1} . GH show that the left-hand side of (7) is strictly increasing in x , while the right-hand side is strictly decreasing in x .¹⁴ This monotonic mapping along with assumptions 1 and 2 imply the existence of a solution to (7).

The intuition behind the result is straightforward: the sequence $\{H^j(K, \eta)\}_{j=0}^{\infty}$ produces a monotonically increasing sequence for the aggregate capital stock, which is bounded above by \bar{K} . GH prove that the pointwise limit of this sequence of functions is the aggregate policy function $\lim_{j \rightarrow \infty} H^j(K, \eta) = Q(K, \eta)$ and that the aggregate law of motion is nondegenerate (i.e., a degenerate law of motion is one that satisfies $Q(K, \eta) = F(K, K, \eta)$ for all K and η).

This mapping serves as the basis for numerical algorithms discussed in this paper, among many others [Coleman (1991); Baxter (1991); Baxter et al. (1990); Davig (2004); Davig et al. (2010); Bi (2012)]. While the purpose of this paper is to provide resources to reduce the cost of implementing the computational algorithm down, the theoretical aspect of the monotone map is very appealing.

A.2 LINEAR INTERPOLATION/EXTRAPOLATION To get an idea of how linear interpolation and extrapolation works, consider the following example with two state variables, x_1 and x_2 . The nearest perimeter around the point, (x'_1, x'_2) , is formed by the four points $(x_{1,i}, x_{2,i})$, $(x_{1,i}, x_{2,i+1})$, $(x_{1,i+1}, x_{2,i})$, and $(x_{1,i+1}, x_{2,i+1})$, where i signifies the position on the grid. We want the policy function value, $f(x'_1, x'_2)$, but we only have policy function values for the four nearest points on the grid (off the grid, we extrapolate using the nearest four points that form a square on the edge of the state space). First, holding x_2 fixed, interpolate/extrapolate in the x_1 direction to obtain

$$\begin{aligned} f(x'_1, x_{2,i}) &= f(x_{1,i}, x_{2,i}) + (x'_1 - x_{1,i}) \frac{f(x_{1,i+1}, x_{2,i}) - f(x_{1,i}, x_{2,i})}{x_{1,i+1} - x_{1,i}} \\ &= \frac{x_{1,i+1} - x'_1}{x_{1,i+1} - x_{1,i}} f(x_{1,i}, x_{2,i}) + \frac{x'_1 - x_{1,i}}{x_{1,i+1} - x_{1,i}} f(x_{1,i+1}, x_{2,i}) \end{aligned} \quad (8)$$

$$f(x'_1, x_{2,i+1}) = \underbrace{\frac{x_{1,i+1} - x'_1}{x_{1,i+1} - x_{1,i}}}_{\omega_{1,i}} f(x_{1,i}, x_{2,i+1}) + \underbrace{\frac{x'_1 - x_{1,i}}{x_{1,i+1} - x_{1,i}}}_{\omega_{1,i+1}} f(x_{1,i+1}, x_{2,i+1}). \quad (9)$$

¹⁴In order to prove the right-hand side is strictly decreasing, the additional assumption, $0 \leq \partial H^j(K, \eta)/\partial K \leq [F_1(K, K, \eta) + F_2(K, K, \eta)]$, needs to be imposed.

Then interpolate/extrapolate in the x_2 direction to obtain

$$\begin{aligned}
 f(x'_1, x'_2) &= f(x'_1, x_{2,i}) + (x'_2 - x_{2,i}) \frac{f(x'_1, x_{2,i+1}) - f(x'_1, x_{2,i})}{x_{2,i+1} - x_{2,i}} \\
 &= \underbrace{\frac{x_{2,i+1} - x'_2}{x_{2,i+1} - x_{2,i}}}_{\omega_{2,i}} f(x'_1, x_{2,i}) + \underbrace{\frac{x'_2 - x_{2,i}}{x_{2,i+1} - x_{2,i}}}_{\omega_{2,i+1}} f(x'_1, x_{2,i+1}). \tag{10}
 \end{aligned}$$

Combining (8), (9) and (10) yields

$$\begin{aligned}
 f(x'_1, x'_2) &= \omega_{2,i} (\omega_{1,i} f(x_{1,i}, x_{2,i}) + \omega_{1,i+1} f(x_{1,i+1}, x_{2,i})) \\
 &\quad + \omega_{2,i+1} (\omega_{1,i} f(x_{1,i}, x_{2,i+1}) + \omega_{1,i+1} f(x_{1,i+1}, x_{2,i+1})) \\
 &= \omega_{1,i} \omega_{2,i} f(x_{1,i}, x_{2,i}) + \omega_{1,i+1} \omega_{2,i} f(x_{1,i+1}, x_{2,i}) \\
 &\quad + \omega_{1,i} \omega_{2,i+1} f(x_{1,i}, x_{2,i+1}) + \omega_{1,i+1} \omega_{2,i+1} f(x_{1,i+1}, x_{2,i+1}) \\
 &= \sum_{j_1=0}^1 \sum_{j_2=0}^1 \omega_{1,i+j_1} \omega_{2,i+j_2} f(x_{1,i+j_1}, x_{2,i+j_2}),
 \end{aligned}$$

which can be easily extended to any number of state variables. We assume that the points for any one dimension in the state space are uniformly spaced, which simplifies evaluation of the policy functions. If unevenly spaced nodes are desired, then `Fallterp*` must be modified to correctly locate the nearest nodes.

A.3 INTEGRATION A model with both continuous and discrete stochastic variables requires two types of numerical integration. For continuous stochastic variables we apply either the Trapezoid rule or Gauss-Hermite quadrature, and for discrete random variables we use the corresponding transition matrix to weight each outcome by its likelihood.

A.3.1 TRAPEZOID RULE Suppose there are m realizations of the stochastic component, ε , in the process for some continuous variable z . Since these realizations show up in agents' expectations, we perform numerical integration to average across each of these m realizations. The trapezoid rule is one method of numerical integration. Assuming uniformly spaced realizations of ε , the formula for the trapezoid rule is given by

$$\begin{aligned}
 E_t[\Phi(\cdot, z_{t+1})] &\approx \frac{\Pr(\varepsilon_1)\Phi(\cdot, z_{t+1}(\varepsilon_1)) + \Pr(\varepsilon_2)\Phi(\cdot, z_{t+1}(\varepsilon_2))}{2} \Delta\varepsilon \\
 &\quad + \frac{\Pr(\varepsilon_2)\Phi(\cdot, z_{t+1}(\varepsilon_2)) + \Pr(\varepsilon_3)\Phi(\cdot, z_{t+1}(\varepsilon_3))}{2} \Delta\varepsilon \\
 &\quad + \frac{\Pr(\varepsilon_{m-1})\Phi(\cdot, z_{t+1}(\varepsilon_{m-1})) + \Pr(\varepsilon_m)\Phi(\cdot, z_{t+1}(\varepsilon_m))}{2} \Delta\varepsilon \\
 &= \frac{\Delta\varepsilon}{2} \left[2 \sum_{i=1}^m \Pr(\varepsilon_i) \Phi(\cdot, z_{t+1}(\varepsilon_i)) - \Pr(\varepsilon_1)\Phi(\cdot, z_{t+1}(\varepsilon_1)) - \Pr(\varepsilon_m)\Phi(\cdot, z_{t+1}(\varepsilon_m)) \right],
 \end{aligned}$$

where $\Delta\varepsilon$ is the distance between stochastic realizations, $\Pr(\varepsilon_i)$ is the probability of realization i , and Φ is the value of the contents of the expectation operator, given the state of the economy. To obtain the weights (the probabilities) in the trapezoid rule, truncate the distribution of the stochastic variable. For normal random variables, we recommend truncating the distribution at no less than four standard deviations, since omitting more of the distribution often leads to inaccurate results.

A.3.2 GAUSS-HERMITE QUADRATURE Another commonly employed method of numerical integration is Gauss-Hermite quadrature. Suppose a shock, u , to a continuous variable, z , is normally distributed with mean μ and variance σ^2 . Then expectations can be written as

$$E_t[\Phi_{t+1}(\cdot, z_{t+1}(u))] = (2\pi\sigma^2)^{-1/2} \int_{-\infty}^{\infty} \Phi_{t+1}(\cdot, z_{t+1}(u)) e^{-(u-\mu)^2/(2\sigma^2)} du.$$

Applying the change of variables, $\varepsilon = (u - \mu)/(\sqrt{2}\sigma)$, the Gauss-Hermite quadrature rule is

$$\begin{aligned} E_t[\Phi_{t+1}(\cdot, z_{t+1}(u))] &= \pi^{-1/2} \int_{-\infty}^{\infty} \Phi_{t+1}(\cdot, z_{t+1}(\sqrt{2}\sigma\varepsilon + \mu)) e^{-\varepsilon^2} d\varepsilon \\ &\approx \pi^{-1/2} \sum_{i=1}^n \omega_i \Phi_{t+1}(\cdot, z_{t+1}(\sqrt{2}\sigma\varepsilon_i + \mu)), \end{aligned}$$

where ε_i are the realizations of the standard normal shock, Φ is the value of the contents of the expectation operator, and ω_i are Gauss-Hermite weights given by $\omega_i = 2^{n+1} n! \sqrt{\pi} [H_{n+1}(\varepsilon_i)]^{-2}$. H_{n+1} is the physicists' Hermite polynomial of order $n + 1$.¹⁵

We usually adopt the Trapezoid rule over Gauss-Hermite quadrature because it is more stable. Moreover, with dense and wide enough grids (at least 10 points and 4 standard deviations) for the continuous shocks, the optimal policy functions under these two methods of numerical integration are virtually identical, even though the Trapezoid rule relies on a truncated distribution.

A.3.3 MARKOV CHAIN INTEGRATION Suppose a discrete stochastic variable, z , evolves according to an m -state first-order Markov chain.¹⁶ Once again, these realizations show up in agents' expectations, and we must integrate across these m realizations conditional on the previous state. Suppose the transition matrix is given by

$$P = \begin{bmatrix} \Pr[s_t = 1|s_{t-1} = 1] & \Pr[s_t = 2|s_{t-1} = 1] & \cdots & \Pr[s_t = m|s_{t-1} = 1] \\ \Pr[s_t = 1|s_{t-1} = 2] & \Pr[s_t = 2|s_{t-1} = 2] & \cdots & \Pr[s_t = m|s_{t-1} = 2] \\ \vdots & \vdots & \ddots & \vdots \\ \Pr[s_t = 1|s_{t-1} = m] & \Pr[s_t = 2|s_{t-1} = m] & \cdots & \Pr[s_t = m|s_{t-1} = m] \end{bmatrix} = \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1m} \\ p_{21} & p_{22} & \cdots & p_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ p_{m1} & p_{m2} & \cdots & p_{mm} \end{bmatrix},$$

where $0 \leq p_{ij} \leq 1$ and $\sum_{j=1}^m p_{ij} = 1$ for all $i \in \{1, 2, \dots, m\}$. Then the conditional expectation can be written as

$$E_t[\Phi_{t+1}(\cdot, z_{t+1})|s_t = i] = [p_{i1} \quad p_{i2} \quad \cdots \quad p_{im}] \begin{bmatrix} \Phi_{t+1}(\cdot, z_{1,t+1}, s_t = i) \\ \Phi_{t+1}(\cdot, z_{2,t+1}, s_t = i) \\ \vdots \\ \Phi_{t+1}(\cdot, z_{m,t+1}, s_t = i) \end{bmatrix}.$$

If a model contains both continuous and discrete stochastic variables, first integrate across the continuous random variables to obtain a set of values, conditional on the realizations of the discrete stochastic variable. Then weight each of these values by their corresponding likelihood. This process yields an expected value across all stochastic components in the model.

¹⁵We provide a function, `ghquad.m`, to compute the Gauss-Hermite weights. To calculate the coefficients of the Hermite polynomial, `ghquad.m` requires `HermitePoly.m`, which is written by David Terr and readily available on the MATLAB file exchange.

¹⁶Recall that higher order Markov chains can always be described by a first-order transition matrix.