Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER Journal of Economic Dynamics & Control 31 (2007) 2698–2712
JOURNAL OF Economic Dynamics & Control
www.elsevier.com/locate/jedc

A generalization of the endogenous grid method

Francisco Barillas[a], Jesús Fernández-Villaverde[b,*]

[a]*New York University, USA*
[b]*Department of Economics, 213 Social Sciences, Duke University, Durham, NC 27708, USA, NBER, and CEPR*

Received 24 February 2006; accepted 8 August 2006
Available online 1 November 2006

Abstract

This paper extends Carroll's [2005. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. National Bureau of Economic Research Technical Working Paper 309]. endogenous grid method to perform value function iteration in models with more than one control variable. We propose to mix the endogenous grid method with standard value function iteration to achieve higher efficiency. We illustrate the method using the stochastic neoclassical growth model with and without labor–leisure choice. We report important gains in efficiency that make value function iteration an attractive computational method in terms of both computing time and accuracy.

JEL classification: C63; C68; E37

Keywords: DSGE models; Computational methods; Value function iteration

1. Introduction

Value function iteration (VFI hereafter) is, perhaps, the most popular approach to solving dynamic stochastic optimization models in discrete time. There are several reasons for its popularity. First, we know many properties of the algorithm, which is

*Corresponding author.
E-mail address: jesusfv@duke.edu (J. Fernández-Villaverde).

0165-1889/$ - see front matter
doi:10.1016/j.jedc.2006.08.005

most convenient for assessing its convergence and performance (see, among many others, Santos and Vigo, 1998; Santos, 1999, and references herein). Second, VFI is well adapted to handle problems with discontinuities or nondifferentiabilities such as borrowing limits or discrete choices, situations where other methods easily run into difficulties. Finally, by controlling the number of grid points, we can achieve any desired level of numerical accuracy.

However, VFI requires large amounts of computational time. Also, it suffers from a strong case of the curse of dimensionality, which makes it unattractive or even infeasible in many applications of interest to economists. These drawbacks are especially acute when we want to nest a VFI within an estimation algorithm, a procedure that may require the repeated solution of the problem for many different parameter values.

The literature, aware of this problem, has proposed many routes to enhance the speed of VFI. Examples include parameterizing the value function (Den Haan and Marcet, 1990), using deterministic and stochastic multigrid schemes (Chow and Tsitsiklis, 1991; Rust, 1997), skipping the maximization step in most iterations (Judd, 1998), or relying on aggressive interpolation algorithms (Krueger and Kubler, 2003).

One recent and promising advance in recursive methods is the contribution of Carroll (2005), who has proposed changing the timing convention of the state variables that define the value function. The idea is as follows. Imagine, for instance, that we want to solve the neoclassical growth model. In the standard approach, we use as the state variable capital at the beginning of the period. However, we can rewrite the problem using the total amount of resources available in the economy (the output of the period plus the undepreciated capital) as the state variable. The solution to both versions of the problem, the one expressed in terms of capital at the beginning of the period and the one expressed in terms of total resources, is equivalent. However, when we iterate, there is a key difference between the two versions of the problem: in the second one, we do not need to find a numerical root of a nonlinear equation, saving a tremendous amount of computational time. The details of why it is so will be clear momentarily.

In this paper, we build upon Carroll's contribution, which he calls the endogenous gridpoints method (EGM hereafter). First, we apply EGM to infinite horizon problems, in particular to the canonical example of the stochastic neoclassical growth model. This is only a trivial extension, and we do not make much of it, but since all applications in Carroll's paper have a finite horizon, our application has a pedagogical interest. Second, we handle the case of endogenous labor choice. Negotiating with this case is a bit more involved than it appears at first sight, and it will lead us to propose a combination of EGM with conventional VFI that we find intriguing. We call this combination of EGM and standard VFI, the *generalized EGM*. Third, we present data on computational time and numerical accuracy, as well as a comparison with the performance of VFI that helps to evaluate the advantages of EGM.

We work with two versions of the neoclassical growth model, one without leisure choice and one with endogenous labor supply. Our main finding is that EGM is astonishingly faster than standard VFI. In the first version of the model, without labor–leisure choice, computational time plummets from 2105 s to 179 s for a

roughly equal level of accuracy (in fact, the accuracy of EGM is slightly better). In the second version of the model, with labor–leisure choice, the reduction in time is from 6.6 h to 7 min.

EGM has, then, the promise to make feasible the computation of complex models with great accuracy. More important, the huge gains in efficiency could enable a researcher to estimate models using a likelihood-based approach that previously could only be estimated by a method of moments. Finally, if we want to conduct welfare analysis, we can do so much more reliably given the great precision of the method in approximating the value function for a fixed amount of computation time.

We organize the rest of the paper as follows. We present in Section 2 the neoclassical growth model with and without labor–leisure choice. This economy will serve as our laboratory to illustrate the performance of EGM. We describe in Section 3 the new algorithm used to approximate the policy functions of the two versions of the neoclassical growth model introduced in Section 2. We report our numerical results in Section 4. We offer some concluding remarks in Section 5.

2. The stochastic neoclassical growth model

Our first task is to set up a model to illustrate the working of EGM. We find natural it to use the stochastic neoclassical growth model, the basic workhorse of modern macroeconomics, to do so. The use of a model without constraints or kinks illustrates the power of EGM even in those situations where some of its advantages (like its ability to deal with corner solutions) are not exploited.

We present two versions of the model: one without labor–leisure choice, which we call MODEL 1, and one with endogenous labor supply, which we call MODEL 2. Just to fix some notation, we go through an accelerated exposition of the models (see, for extended details, Cooley and Prescott, 1995). In MODEL 1, there is a representative agent with utility function from consumption, c_t:

$$U = \mathrm{E}_0 \sum_{t=1}^{\infty} \beta^{t-1} \frac{c_t^{1-\tau}}{1-\tau}.$$

In MODEL 2, the representative household has the following utility function:

$$U = \mathrm{E}_0 \sum_{t=1}^{\infty} \beta^{t-1} \frac{\left(c_t^{\theta}(1-l_t)^{1-\theta}\right)^{1-\tau}}{1-\tau},$$

where $\beta \in (0,1)$ is the discount factor, τ is the coefficient of relative risk aversion, θ controls labor supply, and E_0 is the conditional expectation operator. Note that this function nests a log utility as $\tau \to 1$.

There is one good in the economy, produced according to $y_t = \mathrm{e}^{z_t} k_t^{\alpha}$ for MODEL 1 and $y = \mathrm{e}^{z_t} k_t^{\alpha} l_t^{1-\alpha}$ for MODEL 2, where k_t is the aggregate capital stock, l_t is aggregate labor, and z_t is a stochastic process representing random technological progress. The technology follows the process $z_t = \rho z_{t-1} + \varepsilon_t$ with $|\rho| \leqslant 1$ and $\varepsilon_t \sim \mathcal{N}(0, \sigma^2)$. Capital evolves according to the law of motion $k_{t+1} = (1-\delta)k_t + i_t$,

where δ is the depreciation rate and i_t investment. The economy must satisfy the aggregate resource constraint $y_t = c_t + i_t$. The model is similar to the infinitely lived agent version of Carroll's (2005) macro model except that we do not have a growth rate of technology and that we allow $|\rho| \leqslant 1$.

Both welfare theorems hold in this economy. Consequently, we can solve directly for the social planner's problem where we maximize the utility of the household subject to the production function, the evolution of technology, the law of motion for capital, the resource constraint, and some initial k_0 and z_0.

The solution to MODEL 1 is fully characterized by the equilibrium conditions:

$$c_t^{-\tau} = \beta \mathrm{E}_t\{c_{t+1}^{-\tau}(\alpha \mathrm{e}^{z_{t+1}} k_{t+1}^{\alpha-1} + 1 - \delta)\}, \tag{1}$$

$$c_t + k_{t+1} = \mathrm{e}^{z_t} k_t^{\alpha} + (1-\delta)k_t, \tag{2}$$

$$z_t = \rho z_{t-1} + \varepsilon_t \tag{3}$$

while MODEL 2 is fully characterized by the only marginally more involved equilibrium conditions:

$$\frac{(c_t^{\theta}(1-l_t)^{1-\theta})^{1-\tau}}{c_t} = \beta \mathrm{E}_t \left\{ \frac{(c_{t+1}^{\theta}(1-l_{t+1})^{1-\theta})^{1-\tau}}{c_{t+1}} (1 + \alpha \mathrm{e}^{z_{t+1}} k_{t+1}^{\alpha-1} l_{t+1}^{1-\alpha} - \delta) \right\}, \tag{4}$$

$$\frac{c_t}{1-l_t} = \frac{\theta}{1-\theta}(1-\alpha)\mathrm{e}^{z_t} k_t^{\alpha} l_t^{-\alpha}, \tag{5}$$

$$c_t + k_{t+1} = \mathrm{e}^{z_t} k_t^{\alpha} l_t^{1-\alpha} + (1-\delta)k_t, \tag{6}$$

$$z_t = \rho z_{t-1} + \varepsilon_t. \tag{7}$$

Eqs. (1) and (4) are the standard Euler equations forr MODEL 1 and MODEL 2, respectively, which relate current and future marginal utilities from consumption. Eq. (5), without equivalent in MODEL 1, is the static first order condition between labor and consumption. The last two equations of each set of equilibrium conditions are the resource constraint of the economy and the law of motion of technology.

Solving for the equilibrium of this economy amounts to finding three policy functions for consumption $c(\cdot,\cdot)$, labor $l(\cdot,\cdot)$, and next period's capital $k'(\cdot,\cdot)$ that deliver the optimal choice of the variables as functions of the two state variables, capital and technology.

Also, with Eqs. (1)–(7), we can compute the steady state of both versions of the model. For the model without leisure, the steady state is given by $k_{\mathrm{ss}} = (1/\alpha(1/\beta - 1 + \delta))^{1/(\alpha-1)}$, $c_{\mathrm{ss}} = k_{\mathrm{ss}}^{\alpha} - \delta k_{\mathrm{ss}}$, and $y_{\mathrm{ss}} = k_{\mathrm{ss}}^{\alpha}$. For the model with labor, the steady state is given by $k_{\mathrm{ss}} = \Psi/\Omega + \varphi\Psi$, $l_{\mathrm{ss}} = \varphi k_{\mathrm{ss}}$, $c_{\mathrm{ss}} = \Omega k_{\mathrm{ss}}$, and $y_{\mathrm{ss}} = k_{\mathrm{ss}}^{\alpha} l_{\mathrm{ss}}^{1-\alpha}$, where $\varphi = (1/\alpha(1/\beta - 1 + \delta))^{1/(1-\alpha)}$, $\Omega = \varphi^{1/\alpha} - \delta$, and $\Psi = \theta/(1-\theta)(1-\alpha)\varphi^{-\alpha}$.

3. The endogenous grid method

3.1. No labor–leisure choice (MODEL 1)

In the case of no labor–leisure choice, we follow closely the algorithm proposed by Carroll (2005) with the rather trivial modifications required to deal with the infinite

horizon case. The problem has a Bellman equation representation:

$$\mathbb{V}(k_t, z_t) = \max_{k_{t+1}} \left\{ \frac{(e^{z_t} k_t^{\alpha} + (1-\delta)k_t - k_{t+1})^{1-\tau}}{1-\tau} + \beta \mathrm{E}_t \mathbb{V}(k_{t+1}, z_{t+1}) \right\}$$
$$\text{s.t.} \quad z_{t+1} = \rho z_t + \varepsilon_{t+1},$$

where $\mathbb{V}(\cdot,\cdot)$ is the value function of the problem.

This problem has two state variables: k_t, which is endogenous, and the exogenous stochastic process z_t. Carroll (2005) uses a state variable called 'market resources' or 'cash-on-hand,' instead of k_t. For the stochastic neoclassical growth model, the market resources are given by

$$Y_t = c_t + k_{t+1} = y_t + (1-\delta)k_t = e^{z_t} k_t^{\alpha} + (1-\delta)k_t$$

i.e., the production function plus the undepreciated capital. We use a capital Y_t to denote the total market resources and a lower y_t for the production function.

As a result, we can write the problem recursively with the following Bellman equation:

$$V(Y_t, z_t) = \max_{k_{t+1}} \left\{ \frac{(Y_t - k_{t+1})^{1-\tau}}{1-\tau} + \beta \mathrm{E}_t V(Y_{t+1}, z_{t+1}) \right\}$$
$$\text{s.t.} \quad z_{t+1} = \rho z_t + \varepsilon_{t+1},$$

where $V(\cdot,\cdot)$ is the value function of the problem. Note the difference in notation between 'bold-face' $\mathbb{V}(\cdot,\cdot)$, the value function defined on capital at the beginning of the period, and 'regular' $V(\cdot,\cdot)$, the value function defined on market resources. Since Y_{t+1} is only a function of k_{t+1} and z_{t+1}, we can write the expectation term in the Bellman equation as follows:

$$\tilde{V}(k_{t+1,} z_t) = \beta \mathrm{E}_t V(Y_{t+1}, z_{t+1})$$

to get

$$V(Y_t, z_t) = \max_{k_{t+1}} \left\{ \frac{(Y_t - k_{t+1})^{1-\tau}}{1-\tau} + \tilde{V}(k_{t+1,} z_t) \right\}.$$

The first order condition of consumption in this alternative formulation of the Bellman equation is

$$(c_t^*)^{-\tau} = \tilde{V}_{k_{t+1}}(k_{t+1}^*, z_t),$$

where $c_t^* = Y_t - k_{t+1}^*$. So, if we know $\tilde{V}(k_{t+1,} z_t)$, the beauty of the problem expressed in terms of market resources is that optimal consumption can be readily computed as follows:

$$c_t^* = (\tilde{V}_{k_{t+1}}(k_{t+1}, z_t))^{-1/\tau} \tag{8}$$

for each point in a grid for k_{t+1} and z_t. Then, given c_t^* and k_{t+1}, we can find $Y_t^* = c_t^* + k_{t+1}$ and obtain

$$V(Y_t^*, z_t) = \left\{ \frac{(c_t^*)^{1-\tau}}{1-\tau} + \tilde{V}(k_{t+1,} z_t) \right\},$$

where we can drop the max operator, since we have already computed the optimal level of consumption. Since $Y_t^* = e^{z_t}(k_t^*)^\alpha + (1-\delta)k_t^*$, an alternative interpretation of the algorithm is that, during the iterations, the grid on k_{t+1} is fixed, but the values of k_t change endogenously in each iteration. Hence, the name of EGM.

How does EGM compare with the standard approach? In the standard VFI, the first order condition is given by

$$(c_t^*)^{-\tau} = \beta E_t \mathbb{V}_k(k_{t+1}^*, z_{t+1}).$$

Since by the resource constraint $c_t = e^{z_t}k_t^\alpha + (1-\delta)k_t - k_{t+1}$, we have to solve

$$(e^{z_t}k_t^\alpha + (1-\delta)k_t - k_{t+1}^*)^{-\tau} = \beta E_t \mathbb{V}_k(k_{t+1}^*, z_{t+1}) \tag{9}$$

a nonlinear equation on k_{t+1}^* for each point in a grid for k_t.

The key difference is, thus, that the endogenous grid method defines a fixed grid over the values of k_{t+1} instead of over the values of k_t. This implies that we already know what values the policy function for next period's capital take and, thus, we can skip the root-finding in (9). Solving for this root is an expensive procedure in terms of time and prone to numerical instabilities.

The EGM is also closely related to time iteration (Coleman, 1990; Judd, 1988). In time iteration, we use tomorrow's policy function to find today's policy function and we iterate until convergence. In EGM, we use the value function at the end of the period to compute today's policy function and we iterate until convergence. However, the usual way in which time iteration is implemented (Coleman, 1990, p. 28, second column), we define a grid on k_t and we solve a nonlinear equation for c_t. Hence, we do not have the steep increase in speed provided by EGM. Nevertheless, it is straightforward to modify time iteration to have a grid on k_{t+1} and avoid the nonlinear equation problem (see Krueger and Ludwig, 2006; Rendhal, 2006 who developed this idea).

Before presenting pseudo-code for the algorithm, we mop up a few remaining issues. First, we define a grid on capital tomorrow, $G_{k_{t+1}} \equiv \{k_1, k_2, \ldots, k_M\}$. The usual recommendations regarding how to place the grid points (for example, by covering more thoroughly those parts of the domain of the value function where the curvature is higher) apply here. Second, to solve the expectation with respect to z_{t+1}, a simple choice is to follow Tauchen's (1986) method to discretize the stochastic process z, $G_z \equiv \{z_1, z_2, \ldots, z_N\}$, with Π^N being the resulting transition matrix with generic element $\pi_{i,j}^N \equiv \Pr(z' = z_j | z' = z_i)$. Third, it is convenient to compute a grid of values of market resources in the next period implied by $G_{k_{t+1}}$ as $G_{Y_{t+1}} = e^{z_{t+1}}k^\alpha + (1-\delta)k$ for all $k \in G_k$ and store those in the memory before running the iterations. Fourth, the algorithm requires an initial $\tilde{V}^0(k_{t+1}, z_t)$. This initial guess on the value function has to be increasing in capital so that (8) has a solution. Since the standard VFI is well known (see a description in Aruoba et al., 2006), we do not present it here and proceed directly to outline the new algorithm.

The algorithm to iterate on the value function for a given grid of capital tomorrow is

1. Set $n = 0$ and guess values for $\tilde{V}^0(k_{t+1}, z_t)$.
2. For every value of $\tilde{V}^n(k_{t+1}, z_t)$, compute the derivative at the points on $G_{k_{t+1}}$ to obtain $\tilde{V}_k^n(k_{t+1}, z_t)$ and compute the optimal level of consumption using (8). To

compute the derivatives, we used the average of the slopes of the linearly interpolated value function. Given this optimal level of consumption, compute the value of the endogenously determined market resources as $Yend_t(kend_t, z_t) = c_t^*(kend_t, z_t) + G_{k_{t+1}}$.

3. Update the value function as follows:

$$V^n(Yend_t, z_t) = \frac{c_t^*(kend_t, z_t)^{1-\tau}}{1-\tau} + \tilde{V}^n(k_{t+1,z_t}).$$

Our new guess should be defined on tomorrow's market resources grid $G_{Y_{t+1}}$. Therefore, we interpolate $V^{n+1}(Yend_t, z_t)$ on $G_{Y_{t+1}}$ with the value of $Yend_t$ computed in step 2 to obtain $V^{n+1}(Y_{t+1}, z_{t+1})$.

4. Compute $\tilde{V}^{n+1}(k_{t+1,z_t}) = \beta \mathrm{E}_z V^{n+1}(Y_{t+1}, z_{t+1})$.
5. If $\sup_{i,j} |\tilde{V}^{n+1}(k_i, z_j) - \tilde{V}^n(k_i, z_j)| \geqslant 1.0\mathrm{e}^{-6}$, then $n \rightsquigarrow n+1$ and go to 2.
6. After the value function has converged, use a nonlinear equation solver to retrieve $kend_t(z_i)$ for all values of $z \in G_z$ from $Yend_t(kend_{t,z_i}) = \mathrm{e}^{z_i} kend_t^{\alpha} + (1-\delta)kend_t$. This step results in different grids for capital, one for each point on the productivity grid.

The pseudo-code illustrates how EGM is more computationally efficient than the standard algorithm mainly for two reasons. First, by predetermining the optimal value of our endogenous state variable k_{t+1}, the model circumvents the need to use a nonlinear equation solver to maximize the Bellman operator, the most time-consuming step in the standard algorithm. Second, we compute expectations only once for every iteration in step 4. In the standard algorithm, we have to compute the expectation multiple times during the maximization of the value function.

3.2. Labor–leisure choice (MODEL 2)

This section shows how we extend Carroll's (2005) idea to solve the stochastic neoclassical growth model with labor–leisure choice. For reasons that will become clearer in what follows, the algorithm presented in Section 3.1 is no longer applicable in this case. However, we show that by keeping the policy function for labor constant at its steady-state value during the iteration, it is possible to still use the algorithm presented in the previous section. In this case, once EGM has converged, we need one iteration of the standard VFI algorithm to recover the policy function for labor.

Our numerical results show that when we apply this combination of EGM and one iteration of the standard VFI, the solution is not as precise as the standard algorithm. The reason is that we are holding labor constant at its steady-state value during the iterations. This is, nonetheless, a good approximation, since for a standard calibration of the growth model, labor choice varies little across different levels of capital and productivity.

Because there could be other models in which a control variable like labor changes substantially or because we are especially concerned about minimizing numerical error, we modify EGM to attain the same level of accuracy as the standard VFI, but in a substantially reduced amount of time. The idea consists of combining the EGM

with the standard VFI by running multiple iterations of the endogenous grid, with periodic iterations of the standard VFI to update the policy function for labor nested within an EGM. As we mentioned before, we call this nesting of EGM and standard VFI, the *generalized EGM*.

Before formally presenting the generalized EGM, we explain the problem of the basic EGM when labor is endogenous. The Bellman equation representation of the stochastic neoclassical growth model with labor–leisure choice is

$$\mathbb{V}(k_t, z_t) = \max_{\{k_{t+1}, l_t\}} \left\{ \frac{(c_t^\theta (1 - l_t)^{1-\theta})^{1-\tau}}{1-\tau} + \beta \mathrm{E}_t \mathbb{V}(k_{t+1}, z_{t+1}) \right\}$$
$$\text{s.t.} \quad c + k_{t+1} = \mathrm{e}^{z_t} k_t^\alpha l_t^{1-\alpha} + (1-\delta)k_t.$$

The model also has the following Bellman equation representation using market resources:

$$V(Y_t, z_t) = \max_{\{k_{t+1}, l_t\}} \left\{ \frac{(c_t^\theta (1 - l_t)^{1-\theta})^{1-\tau}}{1-\tau} + \beta \mathrm{E}_t V(Y_{t+1}, z_{t+1}) \right\},$$
$$Y_t = c_t + k_{t+1} = \mathrm{e}^{z_t} k_t^\alpha l_t^{1-\alpha} + (1-\delta)k_t, \tag{10}$$

where we follow the same convention with respect to $\mathbb{V}(\cdot,\cdot)$ and $V(\cdot,\cdot)$ as before.

The algorithm presented in Section 3.1 is not applicable anymore to the value function in problem (10). The reason is that the time-invariant optimal policy function for labor $l(k_t, z_t)$ affects the market resources next period, since $Y_{t+1} = \mathrm{e}^{z_{t+1}} k_{t+1}^\alpha l_{t+1}^{1-\alpha} + (1-\delta)k_{t+1}$ and, therefore, we cannot fix a grid on Y_{t+1} any longer. However, if we knew the time-invariant policy function for labor, we could still calculate optimal consumption using the first order condition for this problem:

$$\theta \frac{(c_t^\theta (1 - l_t)^{1-\theta})^{1-\tau}}{c_t} = \beta \mathrm{E}_t V_Y(Y_{t+1}, z_{t+1}) = \tilde{V}_k(k_{t+1,} z_t),$$

where, as in the previous subsection, $\tilde{V}(k_{t+1,} z_t) = \beta \mathrm{E}_t V(Y_{t+1}, z_{t+1})$. So if we knew $\tilde{V}(k_{t+1,} z_t)$, optimal consumption could be readily computed as follows:

$$c_t^* = \left(\frac{\tilde{V}_{k_{t+1}}(k_{t+1,} z_t) \times (1 - l_t)^{(\theta-1)(1-\tau)}}{\theta} \right)^{-1/(\theta(1-\tau)-1)}. \tag{11}$$

Our generalized EGM is based precisely on this insight. In the first step, it keeps the labor choice constant at its steady-state value l_{ss}. Then, once convergence has been achieved, we apply standard VFI once in a second step to obtain the policy function for labor. We can think then of EGM in the first step as a procedure to generate an extraordinarily accurate guess of a value function for the standard VFI. To achieve further accuracy, we use the policy function from step 2 to plug in (11)

and obtain c_t^* for each iteration of a new EGM. We can iterate steps 2 and 3 until we converge to any desired level of numerical error.

As we did before, we define a grid on capital tomorrow, $G_{k_{t+1}} \equiv \{k_1, k_2, \ldots, k_M\}$, and use Tauchen's (1986) method to discretize the stochastic process z, $G_z \equiv \{z_1, z_2, \ldots, z_N\}$, with Π^N being the resulting transition matrix with generic element $\pi_{i,j}^N \equiv \Pr(z' = z_j | z' = z_i)$. Also, we compute a grid of values of market resources in the next period implied by $G_{k_{t+1}}$ as $G_{Y_{t+1}} = \mathrm{e}^{z_{t+1}} k^{\alpha} l_{\mathrm{ss}}^{1-\alpha} + (1-\delta)k$ for all $k \in G_k$. This G_k is a grid for capital today that we will use for the standard VFI step. We can choose the same grid as $G_{k_{t+1}}$, being careful, however, with the interpretation of the variable in this second grid.

The algorithm to iterate on the value function for a given grid is given by

1. Implement the algorithm described in Section 3.1 with $l_t = l_{\mathrm{ss}}$ until convergence is achieved.
2. Use the converged value function interpolated on the grid $G_{k_{t+1}}$ or from step 3 below, $\mathbb{V}^0(k_t, z_t)$, as an input for the standard VFI algorithm and perform one iteration to recover the policy function for labor, capital, and consumption. Also compute the degree of precision of this iteration as $\sup_{i,j} |\mathbb{V}^1(k_i, z_j) - \mathbb{V}^0(k_i, z_j)| = \varepsilon$. If $\varepsilon < 1.0\mathrm{e}^{-6}$, exit the algorithm.
3. Apply EGM using the policy function for labor from step 2 as follows:

   (a) From a previous iteration of step 2, we obtained policy functions that are defined over the points on the grid $G_{k_{t+1}}$. To apply a similar algorithm as in Section 3.1, we need to find the values of capital that give us as a choice points on the grid $G_{k_{t+1}}$. Therefore, we use an interpolation method and solve for all $kend_{i,j}$ such that $k'(kend_{i,j}, z_j) = G_{k_{t+1}}$.

   (b) Take the policy function for labor from a previous iteration of step 2, $l(G_{k_{t+1}}, z_j)$, and use the endogenous grid points $kend_{i,j}$ to obtain $l(kend_{ij}, z_j)$. This is necessary here because the policy function for labor in step 2 is defined on the grid $G_{k_{t+1}}$.

   (c) Set $n = 0$ and compute values for $\tilde{V}^0(k_{t+1,} z_t)$ as follows:

$$\tilde{V}^0(k_{t+1,} z_t) = \beta \mathrm{E}_t V^0(k_{t+1,} z_{t+1}).$$

   (d) For every value of $\tilde{V}^n(k_{t+1,} z_t)$, compute the derivative at the points on $G_{k_{t+1}}$ to obtain $\tilde{V}_k^n(k_{t+1,} z_t)$ and compute the optimal level of consumption:

$$c_t^*(kend_t, z_t) = \left(\frac{\tilde{V}_{k_{t+1}}(k_{t+1,} z_t) \times (1 - l(kend_{i,j}, z_j))^{(\theta-1)(1-\tau)}}{\theta} \right)^{-1/(\theta(1-\tau)-1)}. \tag{12}$$

   Then, using the budget constraint, solve for the $kend_{i,j}$ implied by this level of optimal consumption and labor choice by solving for $kend_t$ in the following nonlinear equation $c_t^*(kend_t, z_t) + G_{k_{t+1}} = \mathrm{e}^{z_i} kend_t^{\alpha} l(kend_{ij}, z_j)^{1-\alpha} + (1-\delta)kend_t$.

(e) Update the value function. Since we already have the optimal level of consumption, we can drop the *max* operator from the Bellman equation and simply update the value function as follows:

$$V^n(Kend_t, z_t) = \frac{(c_t^*(kend_t, z_t)^\theta (1 - l(kend_{ij}, z_j))^{1-\theta})^{1-\tau}}{1-\tau} + \tilde{V}^n(k_{t+1}, z_t).$$

However, our new guess should be defined on tomorrow's capital grid $G_{k_{t+1}}$, and therefore, it is necessary to interpolate $V^{n+1}(kend_t, z_t)$ on $G_{k_{t+1}}$ using the value of $kend_t$ computed in step (d) to obtain $V^{n+1}(k_{t+1}, z_{t+1})$.

(f) Take the policy function for labor from a previous iteration of step 2, $l(G_{k_{t+1}}, z_j)$ and use the endogenous grid $kend_{i,j}$ computed in step (d) to obtain $l(kend_{ij}, z_j)$ via interpolation.

(g) Compute $\tilde{V}^{n+1}(k_{t+1}, z_t) = \beta \mathrm{E}_z V^{n+1}(k_{t+1}, z_{t+1})$.

(h) If $\sup_{i,j} |V^{n+1}(k_i, z_j) - V^n(k_i, z_j)| \geqslant 0.1\varepsilon$, then $n \rightsquigarrow n+1$ and go to step (d). If convergence has been achieved, go to step 2.

The value of ε in step (h) refers to the maximum difference in the value functions from step 2, $\sup_{i,j} |\mathbb{V}^1(k_i, z_j) - \mathbb{V}^0(k_i, z_j)| = \varepsilon$. The modified algorithm should run past that threshold in order for this extra step to have the desired effect on accuracy. Here we suggest running it until the change in the value functions is smaller than 0.1ε. For our particular calibration, we found that doing just one iteration of the standard VFI between different runs of EGM is enough to obtain results as precise as those obtained by using the standard algorithm. Nonetheless, in other applications, it may prove convenient to run several iterations of the standard VFI before switching back to EGM.

The basic idea of the generalized EGM, combining steps of EGM and of standard VFI, is reminiscent of the accelerator principle described by Judd (1998). When we compute a standard VFI, we do need to update the policy function (i.e., solving the maximization operator) in each iteration, since little accuracy is gained and much time lost. Instead, we can update the policy function, for example, only one of each 10 iterations. The role of EGM is similar to the role of the nine iterations in standard VFI where the max operator is not applied. However, since EGM solves the exact problem in each iteration (just in a different grid), the accuracy and speed of the procedure is superior.

The generalized EGM with a labor–leisure choice is more computationally efficient than the standard algorithm for the same two reasons outlined in the case without labor–leisure choice. In this instance, however, we cannot circumvent solving a nonlinear equation because we are required to update the policy function for labor with the endogenous grid for capital and the knowledge of the endogenous market resources, Y_t , cannot achieve that goal. As we will document below, despite this additional task, our generalized EGM dramatically outperforms the standard algorithm: it reduces computational time from hours to a few minutes without sacrificing accuracy. Also, the generalized EGM can be readily applied to models, like MODEL 1, where it is not strictly necessary. In those situations, the generalized EGM may achieve greater accuracy than any of the two solution approaches, EGM and standard VFI, separately.

4. Calibration and numerical results

4.1. Calibration

We select two set of calibration values. First, to ensure comparability with previous exercises, we reproduced Carroll's (2005) macro model specialization (which does not have a labor/leisure choice) using his parameter values as stated in Table 1.

Also, we followed Carroll (2005) by choosing a unit root process for technology with growth factor 1.01, and three permanent shock realizations $(0.9, 1, 1.1)$ with probabilities $(0.25, 0.5, 0.25)$.

For our second calibration, targeted to match some U.S. economy observations, we picked parameter values for the model as follows. The discount factor $\beta = 0.9896$ matches an annual interest rate of 4%. The risk aversion $\tau = 2$ is a common choice in the literature. $\theta = 0.357$ matches labor supply to 31% of available time in the steady state. We set $\alpha = 0.4$ to match labor share of national income (after the adjustments to National Income and Product Accounts suggested by Cooley and Prescott, 1995). The depreciation rate $\delta = 0.0196$ fixes the investment/output ratio. Values of $\rho = 0.95$ and $\sigma = 0.007$ match the stochastic properties of the Solow residual of the U.S. economy. The chosen values are summarized in Table 2.

4.2. Numerical results

As explained before, our first exercise reproduces Carroll's (2005) macro specialization model. We compute the model using standard VFI and the EGM with a 100 points grid. We find that the EGM delivers a nearly identical solution in less than 25% of time. The finding is relevant because it confirms Carroll's results and it shows the power of EGM.

Encouraged by the findings from Carroll's baseline model, we then compute both versions of the neoclassical growth model using a much finer grid on capital tomorrow of 1000 points centered around the steady state with a coverage of 25%

Table 1
Calibrated parameters

Parameter	β	τ	α	δ
Value	0.9896	2.0	0.36	0.1

Table 2
Calibrated parameters

Parameter	β	τ	θ	α	δ	ρ	σ
Value	0.9896	2.0	0.357	0.4	0.0196	0.95	0.007

and for a range of technology shocks from −0.065 to 0.065 (with zero being the level of technology in the deterministic case). The value of 0.065 corresponds to roughly the 99.5th percentile of the normal distribution given our parameterization. The interval for capital includes virtually 100% of the stationary distributions in our simulations. We use Tauchen's (1986) method to discretize the stochastic process z into 41 discrete states. Therefore, we iterate on a grid of 41,000 points (1000 points for capital and 41 for the productivity level).

To assess the accuracy of our algorithm, we follow Judd (1992) and we compute normalized Euler equation errors. First, note that in our model the intertemporal condition,

$$u_c'(c(k_t,z_t),l(k_t,z_t)) = \beta \mathrm{E}_t\{u_c'(c(k(k_t,z_t),z_{t+1}),l(k(k_t,z_t),z_{t+1}))R(k_t,z_t,z_{t+1})\}, \tag{13}$$

where

$$R(k_t,z_t,z_{t+1}) = (\alpha \mathrm{e}^{z_{t+1}} k(k_t,z_t)^{\alpha-1} l(k(k_t,z_t),z_{t+1})^{1-\alpha} + 1 - \delta)$$

is the gross return rate of capital, should hold exactly for given k_t, and z_t. Since the solution methods used are only approximations, (13) will not hold exactly when evaluated using the computed decision rules. Instead, for solution method i with associated policy rules $c^i(\cdot,\cdot)$, $l^i(\cdot,\cdot)$, and $k^i(\cdot,\cdot)$, and the implied gross return of capital $R^i(k_t,z_t,z_{t+1})$, we can define the Euler equation error function $EE^i(\cdot,\cdot)$ as

$$EE^i(k_t,z_t) \equiv 1 - \frac{\left(\dfrac{\beta \mathrm{E}_t\{u_c'(c^i(k^i(k_t,z_t),z_{t+1}),l^i(k^i(k_t,z_t),z_{t+1}))R^i(k_t,z_t,z_{t+1})\}}{\theta(1-l^i k_t,z_t)^{(1-\theta)(1-\tau)}}\right)^{1/(\theta(1-\tau)-1)}}{c^i(k_t,z_t)}.$$

This function determines the (unit free) error in the Euler equation as a fraction of the consumption given the current states k_t and z_t and solution method i. Judd and Guu (1997) interpret this error as the relative optimization error incurred by the use of the approximated policy rule. For instance, if $EE^i(k_t,z_t) = 0.01$, then the agent is making a \$1 mistake for each \$100 spent. In comparison, $EE^i(k_t,z_t) = 1\mathrm{e}^{-6}$ implies that the agent is making a \$1 mistake for each one million dollars spent.

The Euler equation error is also important because we know that, under certain conditions, the approximation error of the policy function is of the same order of magnitude as the size of the Euler equation error. Correspondingly, the change in welfare is of the square order of the Euler equation error (Santos, 2000).

We can summarize the information from Euler equation error functions in two complementary ways. First, following Judd and Guu (1997), we report the maximum error in a set around the steady state. The second procedure to summarize Euler equation errors is to combine them with the information from the simulations to find the average error. This integral is a welfare measure of the loss induced by the use of the approximating solution method. We report the absolute errors in base 10 logarithms to ease interpretation. A value of −3 means a \$1 mistake for each \$1000, a value of −4 a \$1 mistake for each \$10,000, and so on.

Table 3 presents the results in terms of time and accuracy for both models. Our computations were carried out with a Pentium IV 3.0 GHz, and the programs were written in `Fortran 95`. All the codes are available at www.econ.upenn.edu/~jesusfv/research.html for download and reproduction of the results of the paper.

We highlight two results. First, the accuracy of the endogenous grid is at least as good as, and in the model without labor choice better, than the accuracy of the standard algorithm. For the model without labor choice, the maximum of the Euler error for standard VFI is −2.54, while for EGM is −3.73. The integral of the Euler error decreases from −4.33 to −4.78. In MODEL 2, the maximum Euler error is −3.31 for standard VFI and −3.32 for EGM, while the integral of the Euler error is −4.19 and −4.18, respectively.

Second, and most importantly, in terms of computation time, the endogenous grid algorithm is incredibly faster than the standard VFI. For MODEL 1, EGM it is approximately 11.75 times faster than regular VFI. For MODEL 2, EGM is 56.5 times faster than VFI. The cost of standard VFI can be seen in the 23,907 s (6.6 h) that the procedure takes even for only 41,000 grid points. The reader should not be surprised: VFI is well known to be slow. Also, it is interesting to note how the introduction of a labor choice increases the computational time by 2.36 times in EGM but a much worse 11.36 times in VFI.

We used as an initial guess for VFI the value function implied by the steady-state values. For EGM, we use the same guess except we introduce a small slope to satisfy the conditions of the algorithm explained above. EGM requires 995 iterations in MODEL 1 and 1490 iterations in MODEL 2. VFI requires 753 iterations in MODEL 1 and 695 in MODEL 2. In both cases, the tolerance level was 1.0e(−6).

It could be argued that it is possible to speed up the convergence of the standard VFI. For example, we can use multigrid methods, skipping the maximization step every certain number of iterations, or rely on aggressive vectorization. However, even doing this does not bring the computation time down to a level comparable with that of the endogenous grid procedure. Moreover, we could also apply multigrid methods and vectorization to EGM. To make the comparison fair, the results presented here do not use any of these procedures for either method.

Table 3
Results

Model	Time (s)	Max. Euler error	Integral of the Euler errors
No labor–leisure choice			
EGM	179	−2.54	−4.33
Standard VFI	2105	−3.73	−4.78
Labor–leisure choice			
EGM	423	−3.31	−4.19
Standard VFI	23 907	−3.32	−4.18

Table 4
Results by stages for the labor–leisure choice endogenous grid procedure

Stage	Time (s)	Max. Euler error	Int. of Euler errors	Tolerance
Step 1	237			
Standard iteration 1	34	−3.21	−4.25	$1.7e^{-5}$
Steps (a)–(g)	118			
Standard iteration 2	34	−3.32	−4.18	$1.2e^{-8}$
Steps (a)–(g)	232			
Standard iteration 3	34	−3.32	−4.18	$5.5e^{-10}$

In the case of MODEL 2, it is insightful to show how the level of accuracy changes during the algorithm and to break the total computation time into the different stages of the algorithm presented in Section 3.2. Table 4 presents these results.

These results illustrate how the algorithm works. Step 1 is extremely fast: it takes only 237 s. When we use a standard iteration to back up the policy functions, the level of tolerance is already almost of the desired magnitude ($1.0e^{-6}$). Then, we perform several more iterations using the modified endogenous grid algorithm (Steps (a)–(g)), in which we no longer hold labor constant at the steady state, and when we use one iteration of the standard algorithm to recover the appropriate policy functions, the algorithm has already converged. These are the results we present in Table 4. As an interesting exercise, we carry on some more iterations using EGM for the third time, but when we recover the policy functions and compute the Euler equation errors, there is no further improvement in terms of accuracy. Another point to highlight here is that this procedure is much slower than that in Section 3.1, mainly because just one iteration of the standard algorithm takes 34 s.

5. Conclusions

In this paper we have presented an extension of the innovative EGM proposed by Carroll (2005). Our generalized EGM applies to problems with more than one state variable and multiple controls. We have shown that the proposed procedure is simple to implement and results in spectacular decreases in computational time at no expense in terms of precision. Because of this, EGM positions value function as a very attractive solution method, in particular for highly nonlinear models.

It is important to stress that the stochastic neoclassical growth model should not be the ultimate target of this exercise. As shown by Aruoba et al. (2006), there are many solution methods that are even faster and as accurate as this one. However, the stochastic neoclassical growth model serves as a useful benchmark to evaluate the accuracy and speed of EGM. Applications where the new algorithm will be more useful, and perhaps the only alternative, include the class of problems in which there are constraints that are occasionally binding over the state space, such as in the case of nonborrowing constraints, and those models with discrete choices.

Acknowledgments

We thank Dirk Krueger for calling our attention to the endogenous grid method for the first time, Chris Carroll for extensive comments on a previous draft of the paper, and Kenn Judd for his editorial help. This material is based upon work supported by the National Science Foundation under Grant No. 0338997.

References

Aruoba, S.B., Fernández-Villaverde, J., Rubio-Ramírez, J., 2006. Comparing solution methods for dynamic equilibrium economies. Journal of Economic Dynamics and Control 30, 2447–2508.

Carroll, C.D., 2005. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. National Bureau of Economic Research Technical Working Paper 309.

Chow, C.S., Tsitsiklis, J.N., 1991. An optimal multigrid algorithm for continuous state discrete time stochastic control. IEEE Transactions on Automatic Control 36, 898–914.

Coleman, W.J., 1990. Solving the stochastic growth model by policy-function iteration. Journal of Business and Economic Statistics 8, 27–30.

Cooley, T.F., Prescott, E.C., 1995. Economic growth and business cycles. In: Cooley, T.F. (Ed.), Frontiers of Business Cycle Research. Princeton University Press, Princeton, NJ, pp. 1–38.

Den Haan, W.J., Marcet, A., 1990. Solving the stochastic growth model by parameterizing expectations. Journal of Business and Economic Statistics 8, 31–34.

Judd, K.L., 1992. Projection methods for solving aggregate growth models. Journal of Economic Theory 58, 410–452.

Judd, K.L., 1998. Numerical Methods in Economics. MIT Press, Cambridge, MA.

Judd, K.L., Guu, S.M., 1997. Asymptotic methods for aggregate growth models. Journal of Economic Dynamics and Control 21, 1025–1042.

Krueger, D., Kubler, F., 2003. Computing equilibrium in OLG models with stochastic production. Journal of Economic Dynamics and Control 28, 1411–1436.

Krueger, D., Ludwig, A., 2006. On the consequences of demographic change for rates of return to capital, and the distribution of wealth and welfare. Mimeo, University of Pennsylvania.

Rendhal, P., 2006. Inequality constraints in recursive economies. Mimeo, European University Institute.

Rust, J., 1997. Using randomization to break the curse of dimensionality. Econometrica 65, 487–516.

Santos, M.S., 1999. Numerical solution of dynamic economic models. In: Taylor, J.B., Woodford, M. (Eds.), Handbook of Macroeconomics, vol. 1a. North-Holland, Amsterdam, pp. 311–386.

Santos, M.S., 2000. Accuracy of numerical solutions using the Euler equation residuals. Econometrica 68, 1377–1402.

Santos, M.S., Vigo, J., 1998. Analysis of error for a dynamic programming algorithm. Econometrica 66, 409–426.

Tauchen, G., 1986. Finite state Markov-chain approximations to univariate and vector autoregressions. Economics Letters 20, 177–181.