

TECHNICAL WORKING PAPER SERIES

THE METHOD OF ENDOGENOUS GRIDPOINTS
FOR SOLVING DYNAMIC STOCHASTIC OPTIMIZATION

Christopher D. Carroll

Technical Working Paper 309
<http://www.nber.org/papers/T0309>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2005

I am grateful to Wouter den Haan, Michael Haliassos, Ken Judd, Albert Marcet, Dimitri Mavridis, Michael Reiter, Victor Rios-Rull, John Rust, Eric Young, and participants in the 2002 meetings of the Society for Computational Economics for discussions of the literature and existing solution techniques. Excellent research assistance has been provided by Zhou Lu and Marc Chan. All errors are my own. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2005 by Christopher D. Carroll. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization
Christopher D. Carroll
NBER Technical Working Paper No. 309
June 2005
JEL No. C6, D9, E2

ABSTRACT

This paper introduces a method for solving numerical dynamic stochastic optimization problems that avoids rootfinding operations. The idea is applicable to many microeconomic and macroeconomic problems, including life cycle, buffer-stock, and stochastic growth problems. Software is provided.

Christopher D. Carroll
Department of Economics
Mergenthaler 440
Johns Hopkins University
Baltimore, MD 21218
and NBER
carroll@jhu.edu

1 The Problem

Consider a consumer whose goal is to maximize discounted utility from consumption

$$\max \sum_{s=t}^T \beta^{t-s} u(C_s) \quad (1)$$

for a CRRA utility function $u(C) = C^{1-\rho}/(1-\rho)$.¹

The consumer's problem will be specialized below to two cases: A standard microeconomic problem with uninsurable idiosyncratic shocks to labor income, and a standard representative agent problem with shocks to aggregate productivity (the 'micro' and the 'macro' models).²

The consumer's initial condition is defined by two state variables: M_t is 'market resources' (macro interpretation: capital plus current output) or 'cash-on-hand' (micro interpretation: net worth plus current income), while P_t is permanent labor productivity in both interpretations.

The transition process for M_t is broken up, for convenience of analysis, into three steps. Assets at the end of the period are market resources minus consumption, equal to

$$A_t = M_t - C_t, \quad (2)$$

and capital at the beginning of the next period is what remains after a depreciation factor $\bar{\gamma}$ is applied,

$$K_{t+1} = A_t \bar{\gamma}, \quad (3)$$

where $\bar{\gamma} = (1 - \delta)$ in the usual macro notation and $\bar{\gamma} = 1$ in the micro interpretation.

The final step can be thought of as the transition from the beginning of period $t + 1$, when capital K_{t+1} but has not yet been used to produce output, and the middle of that period, when output has been produced and incorporated into resources:

$$M_{t+1} = \overbrace{\ell \Theta_{t+1} P_{t+1}}^{\equiv L_{t+1}} \mathcal{W}_{t+1} + K_{t+1} \mathcal{R}_{t+1} \quad (4)$$

where \mathcal{W}_{t+1} is the wage rate; Θ_{t+1} is an iid transitory shock (e.g., unemployment) that satisfies $E_t[\Theta_{t+n}] = 1 \forall n > 0$ (usually $\Theta_t = 1 \forall t$ in the macro interpretation); and ℓ is a placeholder for labor supply, which for purposes of this paper is fixed at $\ell = 1$, but in general could be allowed to vary. The disarticulation of the flow of income into labor and capital components is useful in thinking separately about the effects of productivity growth (captured in the P component of the 'labor income' term) and capital accumulation.

¹Putting leisure in the utility function is straightforward but would distract from the paper's point.

²Different aspects of the setup of the problem will strike micro and macroeconomists as peculiar; with patience, it should become clear how the problem as specified can be transformed into more familiar forms.

Permanent labor productivity (in either interpretation) evolves according to

$$P_{t+1} = G_{t+1}P_t\Psi_{t+1} \quad (5)$$

for a permanent shock that satisfies $E_t[\Psi_{t+n}] = 1 \forall n > 0$ and G_t is exogenous and perfectly predictable (see below for varying interpretations of G).

Defining lower case variables as the upper-case variable scaled by the level of permanent labor productivity, e.g. $a_t = A_t/P_t$, we have

$$a_t = m_t - c_t \quad (6)$$

while with a bit of algebra the state transition becomes

$$m_{t+1} = \underbrace{\ell\Theta_{t+1}}_{\equiv l_{t+1}}\mathcal{W}_{t+1} + \underbrace{(a_t \lrcorner / G_{t+1} \Psi_{t+1})}_{=k_{t+1}}\mathcal{R}_{t+1}. \quad (7)$$

The interest and wage factors are assumed not to depend on anything other than capital and productive labor input; together with the iid assumption about the structure of the shocks, this implies that the problem has a Bellman equation representation (henceforth boldface indicates functions)

$$\mathbf{V}_t(M_t, P_t) = \max_{C_t} \{u(C_t) + \beta E_t[\mathbf{V}_{t+1}(M_{t+1}, P_{t+1})]\} \quad (8)$$

subject to the transition equations.

Defining $\Gamma_{t+1} \equiv G_{t+1}\Psi_{t+1}$, consider the related problem

$$\mathbf{v}_t(m_t) = \max_{c_t} \left\{ u(c_t) + \beta E_t \left[\Gamma_{t+1}^{1-\rho} \mathbf{v}_{t+1} \left(\overbrace{l_{t+1}\mathcal{W}_{t+1} + a_t\mathcal{R}_{t+1} \lrcorner / \Gamma_{t+1}}^{=m_{t+1}} \right) \right] \right\}. \quad (9)$$

Assume that there is some last period T in which

$$\mathbf{V}_T(M_T, P_T) = P_T^{1-\rho} \mathbf{v}_T(M_T/P_T) \quad (10)$$

for some well-behaved \mathbf{v}_T (we will be more specific about the terminal value function below). In this case it is easy to show that the solution to the ‘normalized’ problem defined by (9) yields the solution to the original problem via $V_t = P_t^{1-\rho} v_t$ for any $t < T$.³

Now define an end-of-period value function ‘Gothic v’ as

$$\mathbf{v}_t(a_t) = \beta E_t[\Gamma_{t+1}^{1-\rho} \mathbf{v}_{t+1}(l_{t+1}\mathcal{W}_{t+1} + \lrcorner a_t \mathcal{R}_{t+1} / \Gamma_{t+1})] \quad (11)$$

with derivative

$$\begin{aligned} \mathbf{v}_t^a(a_t) &= \beta E_t [\Gamma_{t+1}^{1-\rho} \mathbf{v}_{t+1}^m(\mathcal{R}_{t+1} a_t \lrcorner / \Gamma_{t+1} + l_{t+1}\mathcal{W}_{t+1}) \mathcal{R}_{t+1} \lrcorner / \Gamma_{t+1}] \\ &= \lrcorner \beta E_t [\Gamma_{t+1}^{-\rho} \mathbf{v}_{t+1}^m(\mathcal{R}_{t+1} a_t \lrcorner / \Gamma_{t+1} + l_{t+1}\mathcal{W}_{t+1}) \mathcal{R}_{t+1}] \end{aligned} \quad (12)$$

³See Carroll (2004) for a proof.

and (11) and (6) imply that (9) can be rewritten using \mathbf{v}_t as

$$\mathbf{v}_t(m_t) = \max_{\{a_t\}} \{u(m_t - a_t) + \mathbf{v}_t(a_t)\}, \quad (13)$$

and the envelope theorem can be applied

$$\mathbf{v}_t^m(m_t) = u'(c_t) \quad (14)$$

while the first order condition yields the Euler equation

$$\begin{aligned} u'(c_t) &= \mathbf{v}_t^a(a_t) \\ &= \neg\beta E_t[u'(\Gamma_{t+1}c_{t+1})\mathcal{R}_{t+1}]. \end{aligned} \quad (15)$$

2 Recursion

Generically, problems like this can be solved by specifying a final-period decision rule \mathbf{c}_T and a procedure for recursion (obtaining \mathbf{c}_t from \mathbf{c}_{t+1}). Here we specify the recursion; below we specify choices for the terminal decision rule.

2.1 A Standard Solution Method

The absence of a closed-form solution means that optimal decision functions (e.g. the consumption function) must be constructed by calculating their values a finite grid of possible values of the state variables. Call some ordered set of such values $\mu_i \in \vec{\mu} \equiv \{\mu_1, \mu_2, \dots, \mu_I\}$.

With \mathbf{c}_{t+1} in hand, the usual solution procedure is to specify a $\vec{\mu}$ and, for each element, to use a numerical rootfinding routine to find the χ_i that satisfies (15),

$$u'(\chi_i) = \mathbf{v}_t^a(\mu_i - \chi_i). \quad (16)$$

The points $\{\mu_i, \chi_i\}$ are then used to construct an interpolating approximation to \mathbf{c}_t . (Choice of interpolation method is separable from the point of this paper; see Judd (1998) for a discussion of choices). Given the interpolated \mathbf{c}_t function the solution for earlier periods is found by recursion.

One of the most computationally burdensome steps in this approach is the numerical solution of (16) for each specified state gridpoint. Even if efficient methods are used for constructing the expectations (cf. the parameterized expectations method of den Haan and Marcet (1990)) and shrewd choices are made for the points to include in $\vec{\mu}$, for each gridpoint a numerical rootfinding operation still must evaluate a substantial number of candidate values for the control variable before finding values that satisfy (16) to an acceptable degree of precision.

2.2 Endogenous Gridpoints Solution Method

This paper's key contribution is to introduce an alternative approach that does not require numerical rootfinding. The trick is to begin with *end-of-period* assets a_t and to use the end-of-period marginal value function \mathbf{v}_t^a , the first order condition, and the budget constraint to construct the unique values of middle-of-period m_t generated by those a_t values.

Specifically, define an exogenous, time-invariant ordered set of values of a_t collected in $\alpha_i \in \vec{\alpha} \equiv \{\alpha_1, \alpha_2, \dots, \alpha_I\}$. For each end-of-period state α_i the marginal value $\mathbf{v}_t^a(\alpha_i)$ is easy to calculate; inverting the consumption first order condition, the α 's generate

$$\chi_i = u'^{-1}(\mathbf{v}_t^a(\alpha_i)). \quad (17)$$

Note that the budget constraint implies that

$$\mu_i = \alpha_i + \chi_i. \quad (18)$$

We now have a collection of $\{\mu_i, \chi_i\}$ pairs in hand and can interpolate as before to generate an approximation to \mathbf{c}_t . This completes the recursion.

The key distinction between this approach and the standard one is that the gridpoints for the policy functions are not predetermined; instead they are endogenously generated from a predetermined grid of values of end-of-period assets (hence the method's name). One reason the method is efficient is that expectations are never computed for any gridpoint not used in the final interpolating function; the standard method may compute expectations for many unused gridpoints.

3 Macro Specialization

We first specialize to a macroeconomic stochastic growth model. Assuming a Cobb-Douglas aggregate production function $F(K, P) = K^\varepsilon P^{1-\varepsilon}$, after normalizing by productivity P (and assuming a constant value G for the the labor productivity growth factor), under the usual assumptions of perfect competition etc. if there is no aggregate transitory shock ($l_{t+1} = 1$) we have

$$\mathcal{R}_{t+1} = 1 + \varepsilon k_{t+1}^{\varepsilon-1} \quad (19)$$

$$\mathcal{W}_{t+1} = (1 - \varepsilon)k_{t+1}^\varepsilon \quad (20)$$

and market resources are the sum of capital and production,

$$m_{t+1} = k_{t+1}\mathcal{R}_{t+1} + \mathcal{W}_{t+1} \quad (21)$$

$$= k_{t+1} + k_{t+1}^\varepsilon. \quad (22)$$

We specify the terminal consumption function as

$$\mathbf{c}_T(m) = m, \quad (23)$$

which is very far from the converged infinite horizon consumption rule, but easy to verify as satisfying the assumption (10) imposed earlier. More efficient choices are available, but for our purposes simplicity trumps efficiency.

An arbitrary specification of the process for permanent productivity shocks is a three point distribution defined by $\vec{\Psi} = \{0.9, 1.0, 1.1\}$ with probabilities $\Pr(\vec{\Psi}) = \{0.25, 0.50, 0.25\}$.⁴

The top panel of figure 1 plots the converged consumption function that emerges from this solution method for the benchmark set of parameter values specified in Table 1, along with the consumption function for the standard perfect foresight version of the model ($\vec{\Psi} = \Pr(\vec{\Psi}) = \{1\}$).

4 Micro Specialization

In the microeconomic literature, the usual approach is to take aggregate interest and wage rates as exogenous, and to focus on transitory (Θ) and permanent (Ψ) shocks to idiosyncratic labor productivity. We again start the recursion with $\mathbf{c}_T(m) = m$, and the permanent shocks are retained exactly as specified for the macro problem.⁵

4.1 Life Cycle Models

Life cycle models specify a stereotypical pattern of lifetime income growth defined by G_t where t is age rather than time and T is the maximum possible lifespan;⁶ mortality uncertainty can be accommodated by age-varying values of β .

4.2 Buffer Stock Models

If R, W, G and β are constant, $\bar{\gamma} = 1$, and the impatience condition

$$R\beta E[(G\Psi)^{-\rho}] < 1 \quad (24)$$

holds, Deaton (1991) and Carroll (2004) show that the problem defines a contraction mapping so that the consumption functions defined by the problem converge from any well-behaved initial starting function $\mathbf{c}_T(m)$; the converged function is defined as

$$\mathbf{c}(m) = \lim_{n \rightarrow \infty} \mathbf{c}_{T-n}(m). \quad (25)$$

We solve for the converged consumption function for two versions.

⁴With careful choice of points and weights, small-dimensional discrete representations like this do a good job of approximating commonly-used continuous distributions like a lognormal, cf. Judd (1998).

⁵Of course, appropriate calibrations for macro and micro permanent shocks are very different, but appropriate calibration is not the point of this paper.

⁶This is the context in which the assumption that $\mathbf{c}_T(m) = m$ actually makes economic sense, as distinct from merely providing a convenient starting point for recursion.

4.2.1 Version With Unemployment

Assume that in future periods there is a small probability \wp that income will be zero (corresponding to a substantial spell of unemployment):

$$\Theta_{t+1} = \begin{cases} 0 & \text{with probability } \wp > 0 \\ \Xi_{t+1}/(1 - \wp) & \text{with probability } (1 - \wp) \end{cases} \quad (26)$$

where $\Xi = \{0.9, 1.0, 1.1\}$ and $\Pr(\Xi|\Theta > 0) = \{0.25, 0.50, 0.25\}$ (the same structure of non-unemployment transitory shocks as for the permanent shocks).

Carroll (2004) shows that in this model,

$$\lim_{m_t \rightarrow 0} \mathbf{c}_t(m_t) = 0. \quad (27)$$

This implies that the minimum value in $\vec{\alpha}$ should be $\alpha_1 = 0$, which will generate $\{\mu_1, \chi_1\} = \{0., 0.\}$ as the first point in the set of interpolating points. The resulting converged $\mathbf{c}(m)$ is shown as the thin solid locus in the bottom panel of figure 1; see the software for details of how the remaining values in $\vec{\alpha}$ were chosen.

4.2.2 Version With Liquidity Constraints

Microeconomic models often include a liquidity constraint in addition to the usual transition equations, and capturing the constraint often induces much additional code.

Dealing with a liquidity constraint using the method of endogenous gridpoints is simple. The key observation is that when the constraint is on the cusp of binding, the marginal value of consumption is equal to the marginal value of saving exactly zero (assuming the constraint is of the form that requires a to be nonnegative; generalization to more elaborate kinds of constraints is straightforward). If the first value in the ordered set $\vec{\alpha}$ is $\alpha_1 = 0$, then the method will produce

$$\chi_1 = \mu_1 = u'^{-1}(\mathbf{v}_t^a(0)), \quad (28)$$

and if we define $\hat{\mathbf{c}}_t(m)$ as the function produced by interpolation among the points generated by $\vec{\alpha}$, the consumption function imposing the constraint will be

$$\mathbf{c}_t(m) = \min(m, \hat{\mathbf{c}}_t(m)). \quad (29)$$

If the consumption function is defined as a piecewise linear spline interpolation among the $\{\mu, \chi\}$ points, the constraint can be handled simply by adding the point $\{\mu_0, \chi_0\} = \{0, 0\}$ to the set of points that constitute the interpolation data.

The converged solution is shown as the bold locus in the bottom panel of figure 1.

5 Conclusion



The method of endogenous gridpoints can be extended to problems with multiple state variables and multiple controls, e.g. a micro consumer with a portfolio choice problem,

or a labor supply decision; or a macro consumer with a utility function that exhibits habit formation (see Carroll (2000) for examples). The method is useful both because it is simpler than the standard method and because it reduces computational demands by eliminating a rootfinding operation.

Table 1: Parameter Values

Parameters Common to All Models		
ρ	2	Relative Risk Aversion
β	0.96	Annual Discount Factor
ℓ	1	Labor supply (fixed)
$\vec{\Psi}$	{0.90, 1.00, 1.10}	Permanent Shock Realizations
$\Pr(\vec{\Psi})$	{0.25, 0.50, 0.25}	Permanent Shock Probabilities
Macro Model Parameters		
τ	0.90	Depreciation Factor
G	1.01	Exogenous Aggregate Productivity Growth Factor
ε	0.36	Capital Share in Production
Micro Model Parameters		
τ	1	Depreciation Factor
G	1.03	Trend Individual Wage Growth Factor
R	1.04	Real Interest Rate
W	1.00	Wage Rate
$\vec{\Xi}$	{0.90, 1.00, 1.10}	Transitory Shock Realizations for Employed
$\Pr(\vec{\Xi} \Theta > 0)$	{0.25, 0.50, 0.25}	Transitory Shock Probabilities for Employed
Parameter Unique to Unemployment Model		
\wp	0.005	Probability of Unemployment Spell

Figure 1: Macro and Micro Consumption Functions



References

- Carroll, Christopher D., 2000, Lecture notes on solving microeconomic dynamic stochastic optimization problems, Manuscript, Johns Hopkins University. Available at <http://econ.jhu.edu/people/ccarroll/solvingmicrodsops.pdf>.
- Carroll, Christopher D., 2004, Theoretical foundations of buffer stock saving, NBER Working Paper Number 10867. <http://econ.jhu.edu/people/ccarroll/BufferStockProofsNew.pdf>.
- Deaton, Angus S., 1991, Saving and liquidity constraints, *Econometrica* 59, 1221–1248.
- denHaan, Wouter J, and Albert Marcet, 1990, Solving the stochastic growth model by parameterizing expectations, *Journal of Business and Economic Statistics* 8(1), 31–34. available at <http://ideas.repec.org/a/bes/jnlbes/v8y1990i1p31-34.html>.
- Judd, Kenneth L., 1998, *Numerical Methods in Economics* (The MIT Press, Cambridge, Massachusetts).

Appendices: *Mathematica* Code

This appendix contains the core code used to generate the micro and macro model solutions graphed in the figures. `Common.nb` contains the parameters and code that are shared for both micro and macro solutions; `Micro.nb` and `Macro.nb` contain the specific parameterizations and specializations for the respective specific problems. The commands to execute the solutions and graph them are not of general interest and are not included, but are part of a downloadable package available on the author's website. Downloadable MATLAB code is also available on the author's webpage; Michael Haliassos and Dimitri Mavrides have written C++ code that solves a closely related problem, which is available from them.

```

Common.nb
uP[c_] := If[c > 0, (*then*) c-ρ, (*else*) ∞];

nP[z_] := z-1/ρ;

uP[at_] :=  $\Gamma\beta$  Sum[
   $\Psi$ tp1 =  $\Psi$ Vec[[ $\Psi$ Loop]];
   $\Gamma$ tp1 =  $\Gamma$   $\Psi$ tp1;
   $\Theta$ tp1 =  $\Theta$ Vec[[ $\Theta$ Loop]];
  ktp1 =  $\Gamma$  at/ $\Gamma$ tp1;
  ltp1 =  $\Theta$ tp1/;
  mtp1 = If[MacroModel && ktp1 == 0, 0, ktp1 R[ktp1] + ltp1 W[ktp1]];
   $\Psi$ VecProb[[ $\Psi$ Loop]]  $\Theta$ VecProb[[ $\Theta$ Loop]] * R[ktp1] uP[ $\Gamma$ tp1 Last[cInterpFunc][mtp1]]
, { $\Psi$ Loop, Length[ $\Psi$ Vec]}
, { $\Theta$ Loop, Length[ $\Theta$ Vec]}
]; (* End Sum *)
cInterpFunc = {Interpolation[{{0., 0.}, {1000., 1000.}}, InterpolationOrder -> 1]};

SolveAnotherPeriod := Block[{},
AppendTo[cInterpFunc,
Interpolation[
Union[
Chop[
Prepend[
Table[
 $\alpha$  =  $\alpha$ Vec[[ $\alpha$ Loop]];
 $\chi$  = nP[uP[ $\alpha$ ]];
 $\mu$  =  $\alpha$  +  $\chi$ ;
{ $\mu$ ,  $\chi$ }
, { $\alpha$ Loop, Length[ $\alpha$ Vec]}]
, {0., 0.}] (* Prepending {0, 0} handles potential liquidity constraint *)
] (* Chop cuts off numerically insignificant digits *)
] (* Union removes duplicate entries *)
, InterpolationOrder -> 1] (* Piecewise linear interpolation *)
]; (* End of AppendTo *)
]; (* End of SolveAnotherPeriod *)
{ $\beta$ ,  $\rho$ , n, l} = {0.96, 2, 40, 1};

```

Micro.nb

```

{T, G, p} = {1, 1.03, 0.005};
Γ = G;
MacroModel = False;
<< Common.nb;
(* Exponential growth picks a good set of values for α *)

αVec = Table[Exp[αLoop] - 1 // N, {αLoop, 0, Log[10], Log[10]/(n - 1)}];
EVec = ΨVec = {0.9, 1., 1.1};
EVecProb = ΨVecProb = {0.25, 0.5, 0.25};
ΘVec = Prepend[EVec / (1 - p), 0.];
ΘVecProb = Prepend[EVecProb (1 - p), p];
R[k_] := 1.04;
W[k_] := 1.;

```

Macro.nb

```

{T, G, ε} = {0.9, 1.01, 0.36};
Γ = G;
MacroModel = True;
<< Common.nb;
kBar = (((Γ^ρ)/β) - 1)/ε ^ (1/(ε - 1));
αVec = Table[Exp[αLoop] - 1,
  {αLoop, 0, Log[10 kBar], Log[10 kBar]/(n - 1)}];

ΨVec = {0.9, 1., 1.1};
ΨVecProb = {0.25, 0.5, 0.25};
ΘVec = {1.};
ΘVecProb = {1.};
R[k_] := If[k > 0, (*then*) 1 + ε k^(ε - 1), (*else*) ∞];
W[k_] := (1 - ε) k^ε;

```