Comments Welcome

# The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems

Christopher D. Carroll[†]
ccarroll@jhu.edu

January 12, 2005

**Abstract**

This paper introduces a solution method for 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.

Keywords: Dynamic optimization, precautionary saving, stochastic growth model, endogenous gridpoints

I am grateful to Wouter den Haan, Ken Judd, Albert Marcet, 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. All errors are my own.

† Department of Economics, The Johns Hopkins University, Baltimore MD 21218-2685; `http://www.econ.jhu.edu/people/ccarroll`. Software archive at `http://www.econ.jhu.edu/people/ccarroll/EndogenousArchive.tar`.

# 1 The Problem

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

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

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. These will be referred to respectively as the 'micro' and the 'macro' models.[2]

In either case, the consumer's initial condition is defined by two 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 resources minus consumption,

$$A_t \;=\; M_t - C_t, \tag{2}$$

while capital at the beginning of the next period is what remains of assets after the depreciation factor $\daleth$ is applied,

$$K_{t+1} \;=\; A_t \daleth, \tag{3}$$

where $\daleth = (1-\delta)$ in the usual macro notation and $\daleth = 1$ in the micro interpretation. The final step in the transition is

$$M_{t+1} \;=\; R_{t+1} K_{t+1} + W_{t+1} \underbrace{\ell \xi_{t+1} P_{t+1}}_{\equiv L_{t+1}} \tag{4}$$

where $R_{t+1}$ and $W_{t+1}$ are respectively the gross interest factor (including return of capital) and the wage rate, $\xi_{t+1}$ is an iid transitory shock that satisfies $E_t[\xi_{t+n}] = 1 \;\forall\; n > 0$ (usually $\xi_t = 1 \;\forall\; t$ in the macro interpretation). $\ell$ 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.

Permanent labor productivity (in either interpretation) evolves according to

$$P_{t+1} \;=\; G_{t+1} P_t \Psi_{t+1} \tag{5}$$

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

[2] Different 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.

for a permanent shock that satisfies $E_t[\Psi_{t+n}] = 1 \ \forall \ n > 0$ and $G_t$ is exogenous and perfectly predictable (though possibly time-varying). In the life-cycle micro interpretation, $t$ is age and $G_t$ captures the predictable component of productivity growth over the lifetime; in the buffer-stock micro interpretation, or in the macro interpretation, $G_t$ is taken to be constant at $G$, reflecting underlying trend wage or productivity growth.

$R$ and $W$ 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

$$V_t(M_t, P_t) \ = \ \max\left\{u(C_t) + \beta E_t[V_{t+1}(M_{t+1}, P_{t+1})]\right\} \tag{6}$$

subject to the transition equations for $M$ and $P$.

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

$$a_t \ = \ m_t - c_t \tag{7}$$

and with a bit of algebra (4) becomes

$$m_{t+1} \ = \ R_{t+1} \underbrace{a_t \daleth / G\Psi_{t+1}}_{\equiv k_{t+1}} + W_{t+1} \underbrace{\ell \xi_{t+1}}_{\equiv l_{t+1}}. \tag{8}$$

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

$$v_t(m_t) \ = \ \max_{c_t}\left\{u(c_t) + \beta E_t\left[\Gamma_{t+1}^{1-\rho} v_{t+1}(a_t R_{t+1}/\Gamma_{t+1} + l_{t+1}W_{t+1})\right]\right\}. \tag{9}$$

Assume that there is some last period $T$ in which

$$V_T(M_T, P_T) \ = \ P_T^{1-\rho} v_T(M_T/P_T) \tag{10}$$

for some well-behaved $v_T(m)$ (we will be more specific about the terminal $v_T$ below). In this case it is easy to show that the solution to the 'normalized' problem defined by (7)-(9) yields the solution to the original problem via $V_t(M_t, P_t) = P_t^{1-\rho} v_t(M_t/P_t)$ for any $t < T$.[3]

Now define and then differentiate 'Gothic v' as

$$\begin{aligned} \mathfrak{v}_t(a_t) \ &= \ \beta E_t[\Gamma_{t+1}^{1-\rho} v_{t+1}(a_t R_{t+1}/\Gamma_{t+1} + l_{t+1}W_{t+1})] & (11) \\ \mathfrak{v}_t'(a_t) \ &= \ \beta E_t\left[R_{t+1}\Gamma_{t+1}^{-\rho} v_{t+1}'(a_t R_{t+1}/\Gamma_{t+1} + l_{t+1}W_{t+1})\right] & (12) \end{aligned}$$

so that (9) can be rewritten as

$$v_t(m_t) \ = \ \max_{\{c_t\}}\left\{u(c_t) + \mathfrak{v}_t(m_t - c_t)\right\}, \tag{13}$$

[3] See Carroll [2004] for a proof.

with first order condition

$$u'(c_t) \;=\; \mathfrak{v}'_t(m_t - c_t), \tag{14}$$

and for future use note that the usual envelope theorem logic implies that

$$u'(c_t) \;=\; v'_t(m_t). \tag{15}$$

# 2 Recursion

Generically, problems like this can be solved by specifying a final-period consumption rule $\mathbf{c}_T(m)$ (boldface indicates that this is a function) and a procedure for recursion (obtaining $\mathbf{c}_t(m)$ from $\mathbf{c}_{t+1}(m)$ for any $t+1$). Here we specify the recursion; below we specify choices for $\mathbf{c}_T(m)$.

As a preliminary, note that since $k_{t+1} = a_t/\Gamma_{t+1}$ the $t+1$ version of (15) implies that (12) can be written as

$$\mathfrak{v}'_t(a_t) \;=\; \beta E_t \left[ R_{t+1} u' \left( \Gamma_{t+1} \mathbf{c}_{t+1}(k_{t+1} R_{t+1} + l_{t+1} W_{t+1}) \right) \right]. \tag{16}$$

with the natural interpretation of $l$ and $k$ as normalized labor and capital.

## 2.1 A Standard Solution Method

The absence of a closed-form solution means that optimal behavior must be calculated at a finite grid of possible values of $m_t$. Call some ordered set of such values $\vec{\mu}$, with individual elements $\{\mu_1, \mu_2, ..., \mu_n\}$.

The usual solution procedure is to specify a $\vec{\mu}$ and, for each $\mu_i \in \vec{\mu}$, to use a numerical rootfinding routine and (16) to find the $\chi_i$ that satisfies (14),

$$u'(\chi_i) \;=\; \mathfrak{v}'_t(\mu_i - \chi_i). \tag{17}$$

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

One of the most computationally burdensome steps in this approach is the numerical solution of (17) for each specified $\mu_i$. 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 $i$ a numerical rootfinding operation still must evaluate a substantial number of candidate values for $\chi_i$ before finding a value that satisfies (17) 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, the Euler equation, and the budget constraint to construct the unique values of beginning-of-period $m_t$ and $c_t$ associated with those $a_t$ values.

Specifically, define a set of values of $a_t$ collected in $\vec{\alpha} = \{\alpha_1, \alpha_2, \ldots, \alpha_n\}$. For each $\alpha_j$ calculate the associated

$$\chi_j \;=\; u'^{-1}\left(\mathfrak{v}'_t(\alpha_j)\right). \tag{18}$$

Note that the budget constraint implies that $c_t + a_t = m_t$ so we can obtain

$$\mu_j \;=\; \alpha_j + \chi_j. \tag{19}$$

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

The key distinction between this approach and the standard one is that the gridpoints for the beginning of the period are not predetermined; instead they are endogenously generated from a predetermined grid of values of end-of-period assets (hence the method's name).

Note that one reason the method is efficient is that expectations are never computed for any $\{\mu, \chi\}$ pair not used in the final interpolating function; standard methods 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 $F(K, P) = K^{\gamma} P^{1-\gamma}$, after normalizing by permanent labor supply (and assuming a constant value of $G$), under the usual assumptions of perfect competition etc. if there is no aggregate transitory shock ($l_{t+1} = 1$) we have

$$R_{t+1} \;=\; 1 + \gamma k_{t+1}^{\gamma-1} \tag{20}$$

$$W_{t+1} \;=\; (1-\gamma) k_{t+1}^{\gamma} \tag{21}$$

and market resources are the sum of capital and production,

$$m_{t+1} \;=\; k_{t+1} R_{t+1} + W_{t+1} \tag{22}$$

$$\;=\; k_{t+1} + k_{t+1}^{\gamma}. \tag{23}$$

We specify the terminal consumption function as

$$\mathbf{c}_T(m) \;=\; m, \tag{24}$$

which is very far from the converged infinite horizon consumption rule, but easy to verify as satisfying the assumption (10) imposed earlier because $\mathbf{C}_T(M,P)=M$ then $\mathbf{c}_T(m)=m$. 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 ($\xi$) and permanent ($\Psi$) shocks to idiosyncratic labor productivity. We continue to assume that $c_T=m_T$, and the permanent shocks as specified for the macro problem are retained.[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 $\beta$.

## 4.2 Buffer Stock Models

If $R, W, G$ and $\beta$ are constant, $\daleth=1$, and the impatience condition

$$R\beta E[(G\Psi)^{-\rho}] < 1 \tag{25}$$

holds, Deaton [1991] and Carroll [2004] show that the consumption functions defined by the problem converge from any well-behaved initial starting function $\mathbf{c}_T(m)$; the

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

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

[6]This is the context in which the assumption that $c_T=m_T$ actually makes economic sense, as distinct from merely providing a starting point for recursion.

converged function is defined as

$$\mathbf{c}(m) \;=\; \lim_{n\to\infty} \mathbf{c}_{T-n}(m). \tag{26}$$

We solve for the converged consumption function for two versions.

#### 4.2.1 Version With Unemployment

Following Carroll [2004], assume that in future periods there is a small probability $p$ that income will be zero (corresponding to a substantial spell of unemployment):

$$\xi_{t+1} = \begin{cases} 0 & \text{with probability } p > 0 \\ \Theta_{t+1}/q & \text{with probability } q \equiv (1-p) \end{cases} \tag{27}$$

where $\vec{\Theta} = \{0.9, 1.0, 1.1\}$ and $\Pr(\vec{\Theta}) = \{0.25, 0.50, 0.25\}(1-p)$ (the same structure of non-unemployment transitory shocks as for the permanent shocks).

Carroll [2004] shows that in this model,

$$\lim_{m_t \to 0} \mathbf{c}_t(m_t) \;=\; 0. \tag{28}$$

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 dealing with the liquidity constraint using the standard method often requires 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

$$\chi_1 = \mu_1 = u'^{-1}(\mathfrak{v}'_t(0)), \tag{29}$$

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

$$\mathbf{c}_t(m) \;=\; \min(m, \hat{\mathbf{c}}_t(m)). \tag{30}$$

Table 1: Parameter Values

| | | |
|---|---|---|
| Parameters Common to All Models | | |
| $\rho$ | 2 | Relative Risk Aversion |
| $\beta$ | 0.96 | Annual Discount Factor |
| $\vec{\Psi}$ | $\{0.90, 1.00, 1.10\}$ | Permanent Shock Realizations |
| $\Pr(\vec{\Psi})$ | $\{0.25, 0.50, 0.25\}$ | Permanent Shock Probabilities |
| Macro Model Parameters | | |
| $\daleth$ | 0.975 | Depreciation Factor |
| $\gamma$ | 0.36 | Capital Share in Production |
| $G$ | 1.01 | Exogenous Aggregate Productivity Growth Factor |
| Micro Model Parameters | | |
| $G$ | 1.03 | Trend Individual Wage Growth Factor |
| $R$ | 1.04 | Real Interest Rate |
| $W$ | 1.00 | Wage Rate |
| $\vec{\Theta}$ | $\{0.90, 1.00, 1.10\}$ | Transitory Shock Realizations |
| $\Pr(\vec{\Theta})$ | $\{0.25, 0.50, 0.25\}$ | Transitory Shock Probabilities |
| Parameter Unique to Unemployment Model | | |
| $p$ | 0.005 | Probability of Unemployment Spell |

If the consumption function is defined as a piecewise linear spline interpolation among the $\{\mu_j, \chi_j\}$ 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). Its effect is generally to reduce by one the number of variables with respect to which a rootfinding operation needs to be performed. The case described here is particularly simple, as there was only one variable with respect to which rootfinding needed to be performed, so all rootfinding was eliminated.

Figure 1: Micro and Macro Consumption Functions

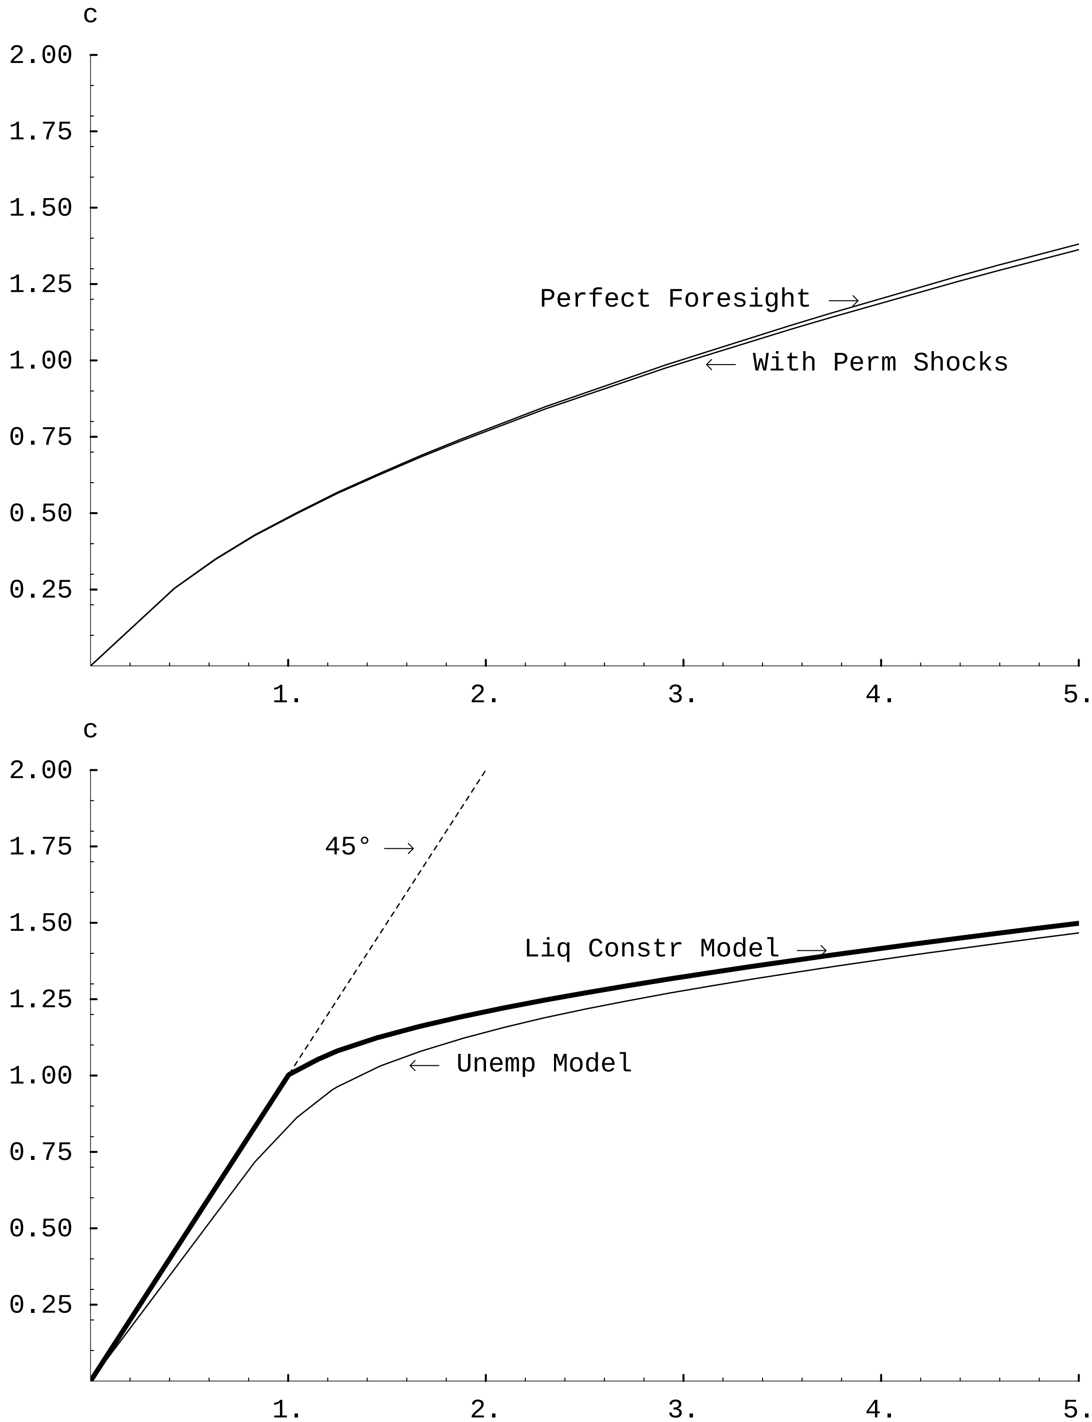

# References

Christopher D. Carroll. Lecture notes on solving microeconomic dynamic stochastic optimization problems. *Manuscript, Johns Hopkins University*, 2000. Available at `http://www.econ.jhu.edu/people/ccarroll/solvingmicrodsops.pdf`.

Christopher D. Carroll. Theoretical foundations of buffer stock saving. *NBER Working Paper Number 10867*, 2004. `http://www.econ.jhu.edu/people/ccarroll/BufferStockProofsNew.pdf`.

Angus S. Deaton. Saving and liquidity constraints. *Econometrica*, 59:1221–1248, 1991.

Wouter J den Haan and Albert Marcet. Solving the stochastic growth model by parameterizing expectations. *Journal of Business and Economic Statistics*, 8(1):31–34, January 1990. available at `http://ideas.repec.org/a/bes/jnlbes/v8y1990i1p31-34.html`.

Kenneth L. Judd. *Numerical Methods in Economics*. The MIT Press, 1998.

# 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.

## Common.nb

```
{β, ρ, n, ℓ} = {0.96, 2, 20, 1};

uP[c_] := c^ - ρ;
nP[z_] := z^ - (1/ρ);

υP[at_] := β If[InadaAta0 && at == 0, ∞, Sum[
 Ψtp1 = ΨVec[[ΨLoop]];
 Γtp1 = Γ Ψtp1;
 ξtp1 = ξVec[[ξLoop]];
 ktp1 = at/Γtp1;
 ltp1 = ξtp1 ℓ;
 ΨVecProb[[ΨLoop]] ξVecProb[[ξLoop]]*
 R[ktp1] uP[Γtp1 Last[cInterpFunc][ktp1 R[ktp1] + ltp1 W[ktp1]] ]
, {ΨLoop, Length[ΨVec]}
, {ξLoop, Length[ξVec]}
]];

cInterpFunc = {Interpolation[{{0., 0.}, {1000., 1000.}}, InterpolationOrder- > 1]};

SolveAnotherPeriod := Block[{},
 AppendTo[cInterpFunc,
  Interpolation[
   Union[
    Chop[
     Prepend[
      Table[
       α = αVec[[αLoop]];
       χ = nP[υP[α]];
       μ = α + χ;
       {μ, χ}
      , {αLoop, Length[αVec]}]
     , {0., 0.}] (* Prepending {0, 0} handles potential liquidity constraint *)
    ]     (* Chop cuts off numerically insignificant digits *)
   ]      (* Union removes duplicate entries *)
  , InterpolationOrder- > 1]
 ];
];
```

# Micro.nb

```
{G, p} = {1.03, 0.005};
Γ = G;
InadaAta0 = True;
<< Common.nb;

αVec = Table[Exp[αLoop] - 1 //N, {αLoop, 0, Log[10], Log[10]/(n - 1)}];
ΘVec  = ΨVec   = {0.9, 1., 1.1};
ΘVecProb = ΨVecProb = {0.25, 0.5, 0.25};
ξVec   = Prepend[ΘVec  / (1 - p), 0.];
ξVecProb = Prepend[ΘVecProb (1 - p), p ];

R[k_] := 1.04;
W[k_] := 1.;
```

# Macro.nb

```
{ℸ, γ, G} = {0.95, 0.36, 1.01};
Γ = G/ℸ;
InadaAta0 = True;

<< Common.nb;

kBar = ((((Γ^ρ)/β) - 1)/γ)^(1/(γ - 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};
ξVec   = {1.};
ξVecProb = {1.};

R[k_] := 1 + γ k^(γ - 1);
W[k_] := (1 - γ) k^γ;
```