SZG Macro 2011
Lecture 3:
Dynamic Programming

Background

- Our previous discussion of optimal consumption over time and of optimal capital accumulation suggest studying the general decision problems on the next pages, where

  c = control (consumption)

  k = state

  x = exogenous variable

- In each case, we might solve the dynamic optimization problem using a “discrete time optimal control” approach as in prior lectures.

Finite horizon problem

$$\max_{\{c_t\},\{k_t\}} \sum_{t=0}^{T} \beta^t u(k_t, x_t, c_t) \qquad u: \ \text{momentary objective}$$

$$s.t. \qquad k_{t+1} - k_t = g(k_t, x_t, c_t) \qquad g: \ \text{state law of motion}$$

$$k_0$$

$$k_{T+1} \geq b(x_{T+1})$$

Infinite horizon problem

$$\max_{\{c_t\},\{k_t\}} \sum_{t=0}^{\infty} \beta^t u(k_t, x_t, c_t)$$

$$s.t. \qquad k_{t+1} - k_t = g(k_t, x_t, c_t) \qquad t = 0,1,\ldots\infty$$

$$k_0$$

We might attack the finite horizon problem by forming

$$L = \sum_{t=0}^{T} \beta^t u(k_t, x_t, c_t)$$

$$+ \sum_{t=0}^{T} \beta^t \lambda_t [k_t + g(k_t, x_t, c_t) - k_{t+1}]$$

$$+ \Theta_{T+1} [k_{t+1} - b(x_{t+1})]$$

We might attack the infinite horizon problem by forming

$$L = \sum_{t=0}^{\infty} \beta^t u(k_t, x_t, c_t) + \sum_{t=0}^{\infty} \beta^t \lambda_t [k_t + g(k_t, x_t, c_t) - k_{t+1}]$$

Either case

- Outcome is sequence of optimal control $\{c_t\}$, optimal state $\{k_t\}$, and optimal shadow prices $\{\lambda_t\}$ that satisfy FOCs and TC give the path $\{x_t\}$.
- Desirable because we have both real outcomes and a means of deriving "supporting" prices.

We will study an alternative, dynamic programming, next

- Why?
  - Alternative tool for toolkit
  - DP is better for certain problems with uncertainty
  - DP logic applied in other, equilibrium contexts
  - LS: “the imperialism of recursive methods”

Outline

1. Certainty optimization problem used to illustrate:
   a. Restrictions on exogenous variables $\{x_t\}$
   b. Value function
   c. Policy function
   d. The Bellman equation and an associated Lagrangian
   e. The envelope theorem
   f. The Euler equation

Outline Cont’d

2. Adding uncertainty

3. Applications
   optimal consumption over time
   optimal consumption under uncertainty.

1. A certainty dynamic problem and the DP approach

- Maximize $$\sum_{t=0}^{\infty} \beta^t u(k_t, x_t, c_t)$$

- Subject to $$k_{t+1} - k_t = g(k_t, x_t, c_t) \quad \text{controlled state}$$

and $$x_t = x(\varsigma_t) \qquad \text{exogenous variable}$$
$$\varsigma_t = m(\varsigma_{t-1}) \qquad \text{exogenous state}$$

What's different from background setup?

- Immediate jump to infinite horizon problem, not essential but matches presentation in LS chapter (note differences in notation, though).
- The exogenous (x) variable(s) are now functions of a vector of *exogenous state variables*, which evolve according to a difference equation (perhaps nonlinear, perhaps in a vector).
- The latter is a key part of the vision of Richard Bellman, the inventor of DP: his experience in other areas (such as difference equations) led him to think in terms of describing dynamics in terms of state variables.

Recursive policies

- Suppose controls are functions of states,

$$c_t = \pi(k_t, \varsigma_t) \quad \text{policy function} \Rightarrow$$

$$k_{t+1} = k_t + g(k_t, x_t, c_t)$$

$$= k_t + g(k_t, x(\varsigma_t), \pi(k_t, \varsigma_t))$$

- Then, the state vector evolves according to a recursion

$$s_{t+1} = \begin{bmatrix} k_{t+1} \\ \varsigma_{t+1} \end{bmatrix} = \begin{bmatrix} k_t + g(k_t, x(\varsigma_t), \pi(k_t, \varsigma_t)) \\ m(\varsigma_t) \end{bmatrix} \equiv M(s_t)$$

that can be used to generate future states from given initial conditions

Evaluating the objective

- Under any recursive policy, we can see that all of the terms which enter in the objective are a function of the initial state (s_0) so that the objective is also a function of the initial state

$$\sum_{t=0}^{\infty} \beta^t u(k_t, x_t, c_t)$$

$$= \sum_{t=0}^{\infty} \beta^t u(k_t, x(\varsigma_t), \pi(k_t, x(\varsigma_t))) = W(k_0, \varsigma_0) = W(s_0)$$

Recursion for objective under arbitrary recursive policy

$$W(s_0;\pi) = w(s_0;\pi) + \beta W(s_1;\pi)$$

$$w(s;\pi) = u(k, x(\varsigma), \pi(k, x(\varsigma)))$$

$$s_{t+1} = M(s_t)$$

Notice the switch

- Given that there is a policy function (π), the objective is now a function of the state vector.
- We have made the change – we are now thinking in terms of functions rather than sequences.
- But we haven't optimized yet! We could be calculating the objective with a very bad policy.

Bellman's core idea

- Subdivide complicated intertemporal problems into many "two period" problems, in which the trade-off is between the present "now" and "later".
- Specifically, the idea was to find the optimal control and state "now", taking as given that latter behavior would itself be optimal.

The Principle of Optimality

- “An optimal policy has the property that, whatever the state and optimal first decision may be, the remaining decisions constitute an optimal policy with respect to the state originating from the first decisions”—Bellman (1957, pg. 83)

Following the principle,

- The natural maximization problem is

$$\max_{c_t, k_{t+1}} \{u(c_t, k_t, x(\varsigma_t)) + \beta V(k_{t+1}, \varsigma_{t+1})\}$$

$$s.t.\;\; k_{t+1} = k_t + g(k_t, x_t, c_t)$$

$$\varsigma_{t+1} = m(\varsigma_t)$$

- Where the right hand side is the current momentary objective (u) plus the consequences (V) for the discounted objective of behaving *optimally* in the future.

Noting that time does not enter in an essential way

- We sometimes write this as (with ‘ meaning next period)

$$\max_{c,k'}\{u(c,k,x(\varsigma))+\beta V(k',\varsigma')\}$$

$$s.t.\ \ k' = k_t + g(k,x(\varsigma),c)$$

$$\varsigma' = m(\varsigma)$$

- So then the *Bellman equation* is written as

$$V(k,\varsigma) = \max_{c,k'}\{u(c,k,x(\varsigma))+\beta V(k',\varsigma')\}$$

$$s.t.\ \ k' = k + g(k,x(\varsigma),c)$$

$$\varsigma' = m(\varsigma)$$

After the maximization

- We know the optimal policy (which we will call π as above, but with the proviso that it is optimal) and can calculate the associated value, so that there is now a Bellman equation of the form

$$V(k,\varsigma) = \{u(\pi(k,\varsigma),k,x(\varsigma)) + \beta V(k + g(k,x(\varsigma),\pi(k,\varsigma)),\varsigma')\}$$

- A functional equation is defined, colloquially, as an equation whose unknowns are functions. In our context, the unknowns are the policy and value functions.

How to do the optimization?

- You are free to choose, depending on the application
- Sometimes we take the Euler route, substituting in the constraint and maximizing directly over k'
- Other times we want to use a Lagrange approach, putting a multiplier on the constraint governing k'

The associated Lagrangian

- Takes the form

$$L = \{u(c,k,x(\varsigma)) + \beta V(k',\varsigma')\}$$
$$+\lambda[k + g(k,x(\varsigma),c) - k']$$

- The optimal policy, state evolution and related multiplier are obtained by maximizing with respect to c,k' and minimizing with respect to λ. Hence these are all functions of the state variables.

For an optimum (off corners)

- We must have

$$\frac{\partial L}{\partial c} = \frac{\partial u(c,k,x(\varsigma))}{\partial c} + \lambda \frac{\partial g(k,x(\varsigma),c)}{\partial c} = 0$$

$$\frac{\partial L}{\partial k'} = -\lambda + \beta \frac{\partial V(k',\varsigma')}{\partial k'} = 0$$

$$\frac{\partial L}{\partial \lambda} = [k + g(k,x(\varsigma),c) - k'] = 0$$

- And, at the values which solve these equations, V=L

The envelope theorem (Benveniste-Scheinkman)

- Question: what is the effect of an infinitessimal change in k on V?
- Answer: It is given by

$$\frac{\partial V}{\partial k} = \frac{\partial u(c,k,x(\varsigma))}{\partial k} + \lambda[\frac{\partial g(k,x(\varsigma),c)}{\partial k} + 1]$$

when we evaluate at the optimal policy and the associated multiplier. As in LS, this may also be written a form which does not involve the multiplier, $\frac{\partial V}{\partial k} = \frac{\partial u(c,k,x(\varsigma))}{\partial k} + \beta\frac{\partial V(k',\varsigma')}{\partial k'}[\frac{\partial g(k,x(\varsigma),c)}{\partial k} + 1]$

Outline of proof

- Nontrivial to show differentiability of V
- But if we have this (as we will frequently assume) then

$$\frac{\partial V}{\partial k} = \frac{\partial L}{\partial k} = \{\frac{\partial u(c,k,x(\varsigma))}{\partial c}\frac{\partial c}{\partial k} + \frac{\partial u(c,k,x(\varsigma))}{\partial k}\}$$

$$+\beta\frac{\partial V(k',\varsigma')}{\partial k'}\frac{\partial k'}{\partial k}$$

$$+\frac{\partial \lambda}{\partial k}[k + g(k,x(\varsigma),c) - k']$$

$$+\lambda[1 + \frac{\partial g(k,x(\varsigma),c)}{\partial k}] + \lambda[\frac{\partial g(k,x(\varsigma),c)}{\partial k}\frac{\partial c}{\partial k} - \frac{\partial k'}{\partial k}]$$

- While this looks ugly, all terms involving behavior are multiplied by coefficients that are set to zero by the FOCs.

Details on ET

$$\frac{\partial V}{\partial k} = \frac{\partial u(c,k,x(\varsigma))}{\partial k} + \lambda[1 + \frac{\partial g(k,x(\varsigma),c)}{\partial k}]$$

$$+\{\frac{\partial u(c,k,x(\varsigma))}{\partial c} + \frac{\partial g(c,k,x(\varsigma))}{\partial c}\}\frac{\partial c}{\partial k}$$

$$+\{\beta \frac{\partial V(k',\varsigma')}{\partial k'} - \lambda\}\frac{\partial k'}{\partial k}$$

$$+\{\frac{\partial \lambda}{\partial k}[k + g(k,x(\varsigma),c) - k']\}$$

Iterating on the Bellman Equation

- Under specific conditions on the functions u and g, the Bellman equation has a unique, strictly concave (in k) solution.
- Under these conditions, it can be calculated by considering the limit

$$V_{j+1}(k,\varsigma) = \max_{c,k'}\{u(k,x(\varsigma),c) + \beta V_j(k',\varsigma')\}$$

$$s.t. \quad k' = k + g(k,x(\varsigma),c)$$

- These iterations are interpretable as calculating the value functions for a class of finite horizon problems, with successively longer horizons.

3. A Stochastic dynamic problem and the DP approach

- Maximize $$E\{\sum_{t=0}^{\infty}\beta^t u(k_t, x_t, c_t)\} \mid (k_0, \varsigma_0)$$

- Subject to $$k_{t+1} - k_t = g(k_t, x_t, c_t)$$

and Markovian exogenous state variables

$$x_t = x(\varsigma_t)$$

$$\Upsilon(\underline{\varsigma}, B) = prob(\varsigma_{t+1} \in B \mid \varsigma_t = \underline{\varsigma})$$

Markov examples

- Markov chains (LS, Chapter 1)
- Linear state space systems
- Nonlinear difference equations with iid shocks,

$$\varsigma_{t+1} = m(\varsigma_t, e_{t+1})$$

- We won't be more explicit until necessary.
- Key point: states are enough to compute expectations.

Bellman Equation

- Uncertainty case is *minor* modification of certainty case

$$V(k,\varsigma) = \max_{c,k'}\{u(c,k,x(\varsigma)) + \beta EV(k',\varsigma') \mid (k,\varsigma)\}$$

$$s.t. \;\; k' = k + g(k,x(\varsigma),c)$$

Proceeding as above

- Lagrangian $$L = \{u(c,k,x(\varsigma)) + \beta EV(k',\varsigma') \,|\, (k,\varsigma)\}$$ $$+\lambda[k + g(k,x(\varsigma),c) - k']$$
- FOCs $$\frac{\partial L}{\partial c} = \frac{\partial u(c,k,x(\varsigma))}{\partial c} + \lambda \frac{\partial g(k,x(\varsigma),c)}{\partial c} = 0$$ $$\frac{\partial L}{\partial k'} = -\lambda + \beta \frac{\partial EV(k',\varsigma')}{\partial k'} = 0$$ $$\frac{\partial L}{\partial \lambda} = [k + g(k,x(\varsigma),c) - k'] = 0$$
- ET is unchanged

Implications for optimal policies and state evolution

- Functions of states

$$c_t = \pi(k_t, \varsigma_t)$$
$$k_{t+1} - k_t = g(k_t, x(\varsigma_t), \pi(k_t, \varsigma_t))$$
$$\lambda_t = \lambda(k_t, \varsigma_t)$$

- State evolution is now a larger Markov process. For example,

$$s_{t+1} = \begin{bmatrix} k_{t+1} \\ \varsigma_{t+1} \end{bmatrix} = \begin{bmatrix} k_t + g(k_t, x(\varsigma_t), \pi(k_t, \varsigma_t)) \\ m(\varsigma_t, e_{t+1}) \end{bmatrix} = M(s_t, e_{t+1})$$

Value Function

- Since c,k,x depend on states, the value function also is V(s).
- It is the maximized RHS of the Bellman equation.

What we’ve covered in this lecture

- Introduction to DP under certainty
- Bellman Equation
- Associated Lagrangian
- FOCs and the ET
- DP with exogenous variables that are functions of a Markov process (exogenous state vector)
- What follows:
  - Optimal consumption over time via dynamic programming: calculation of policy and value functions in a simple case
  - Setting up optimal consumption problem with uncertain income

3A. Optimal consumption over time

- Simple case (no k,x in u)

$$\sum_{t=0}^{\infty} \beta^t u(c_t)$$

- Accumulation of assets

$$a_{t+1} = R[a_t + y_t - c_t]$$
$$y_{t+1} = y + \rho(y_t - y)$$

- And βR=1 (level consumption)

Bellman Equation

$$V(a,y) = \max_{c,a'}\{u(c) + \beta V(a',y')\}$$

$$s.t.\ a' = R[a + y - c]$$

$$y' - \underline{y} = \rho(y - \underline{y})$$

Taking an Euler Route

$$V(a,y) = \max_{c,a'}\{u(a+y-\frac{1}{R}a') + \beta V(a',y')\}$$

$$s.t.\ \ y'-\underline{y} = \rho(y-\underline{y})$$

$$EE: 0 = -u_c(a+y-\frac{1}{R}a')\frac{1}{R} + \beta\frac{\partial V(a',y')}{\partial a'}$$

$$ET: \frac{\partial V(a,y)}{\partial a} = u_c(a+y-\frac{1}{R}a')$$

Learning about consumption

- Update ET and insert in EE to get

$$u_c(a+y-\frac{1}{R}a')=u_c(a'+y'-\frac{1}{R}a'')\Rightarrow c=c'$$

- Suppose there is a linear policy function

$$c=\kappa+\theta_y(y-\underline{y})+\theta_a a$$

$$\begin{aligned} c' &= \kappa+\theta_y(y'-\underline{y})+\theta_a a' \\ &= \kappa+\theta_y\rho(y-\underline{y})+\theta_a R[a+y-c] \\ &= \kappa+\theta_y\rho(y-\underline{y})+\theta_a R[a+y-\kappa-\theta_y(y-\underline{y})-\theta_a a] \end{aligned}$$

Requiring c=c', we have equations that restrict undetermined coefficients

$$\kappa + \theta_y (y - \underline{y}) + \theta_a a$$
$$= \kappa + \theta_y \rho (y - \underline{y}) + \theta_a R[a + (y - \underline{y}) + \underline{y} - \kappa - \theta_y (y - \underline{y}) - \theta_a a]$$

$$\kappa = \kappa + \theta_a R(\underline{y} - \kappa) \Rightarrow \kappa = \underline{y}$$
$$\theta_y = \theta_y \rho + \theta_a R(1 - \theta_y) \Rightarrow \theta_y = \theta_a R / [1 - \rho + \theta_a R]$$
$$\theta_a = \theta_a R[1 - \theta_a] \Rightarrow \theta_a = (\frac{R-1}{R})$$

$$\theta_y = \theta_a \frac{1}{(1 - \frac{\rho}{R})}$$

Economic Rules

- Consume the normal level of income ($\underline{y}$)
- Consume the interest from asset stock, leaving the asset stock unchanged period to period (except as noted next)
- Consume based on the "present value" of deviations from normal income, treating this as if it were another source of wealth; allow variations in asset position on this basis.

Could have gotten these rules more directly

$$\sum_{j=0}^{\infty}\left(\frac{1}{R}\right)^{j} c = a + \sum_{j=0}^{\infty}\left(\frac{1}{R}\right)^{j}\left[\underline{y} + \rho^{j}(y - \underline{y})\right]$$

$$\frac{1}{1-\frac{1}{R}} c = a + \frac{1}{1-\frac{1}{R}}\underline{y} + \frac{1}{1-\frac{\rho}{R}}(y - \underline{y})$$

$$c = \underline{y} + \frac{R-1}{R}\left[a + \frac{1}{1-\frac{\rho}{R}}(y - \underline{y})\right]$$

Questions & Answers

- If we could have gotten them more easily, then why do we need DP?
  - Because there are many problems that we cannot solve so easily and DP is a procedure for solving them.
- What is the value function?

$$V(a,y)=\frac{1}{1-\beta}u\left(a+\frac{1}{1-\frac{1}{R}}\underline{y}+\frac{1}{1-\frac{\rho}{R}}(y-\underline{y})\right)$$

  - Easy to determine in this case because c is constant over time; V inherits properties of u
  - Check: take this v, insert in Bellman equation as v’, show optimal form c has specified form, show v has this form.

3B. Optimal consumption with fluctuating income: setting up a DP

- Simple case (no k,x in u)

$$E\{\sum_{t=0}^{\infty}\beta^t u(c_t)\} \mid s_0$$

- Accumulation of assets (don't necessarily restrict R)

$$a_{t+1} = R[a_t + y_t - c_t]$$

- Income process

$$y(\varsigma_t)$$

$$\varsigma_t : Markov$$

One version of the Bellman equation

$$V(a,\varsigma) = \max_{c,a'}\{(u(c) + \beta EV(a',\varsigma'))\}$$

$$s.t.\ [a + y(\varsigma) - c - \frac{1}{R}a'] = 0$$

FOCs and ET

- Make sure you can work these out following the recipe above,

$$c: \quad u_c(c) - \lambda = 0$$

$$a': \; -\frac{1}{R}\lambda + \beta E\{\frac{\partial EV(a',\varsigma')}{\partial a'}\} = 0$$

$$\lambda: \; [a + y(\varsigma) - c - \frac{1}{R}a'] = 0$$

$$ET: \frac{\partial EV(a,\varsigma)}{\partial a} = \lambda$$

Implications for policies

- Optimal consumption depends on (a) wealth; and (b) the variables that are useful for forecasting future income.

$$c(a, \varsigma)$$

- But solving for this function is no longer easy. Rationalizes SL's discussion of numerical methods, a topic that we will consider further later.

Implication for value function

- Value function is objective evaluated at optimal consumption policy, which is a function of a Markov process, so that

$$V(a_0,\varsigma_0) = E\{\sum_{t=0}^{\infty}\beta^t u(\pi(a_t,\varsigma_t))\} \mid (a_0,\varsigma_0)$$

- Value function satisfies the Bellman functional equation.

$$V(a,\varsigma) = \max_{c,a'}\{(u(c) + EV(a',\varsigma')\}$$

$$s.t.\ [a + y(\varsigma) - c - \frac{1}{R}a' = 0]$$

$$= (u(\pi(a,\varsigma)) + EV(R[a + y(\varsigma) - \pi(a,\varsigma)],\varsigma') \mid (a,\varsigma)$$