Lecture 10

Chapter 4: Dynamic Programming

6/22, 2023

Basic Idea

Finite Horizon

- Consider a two-period version of the Neoclassical growth model:

$$\max_{c_0, c_1, k_1 \geq 0, k_2 \geq 0} \ln c_0 + \beta \ln c_1$$

- Subject to the initial condition $k_0 > 0$ and

$$k_{t+1} = Ak_t^\alpha - c_t$$

- This problem is equivalent to:

$$\max_{k_1 \geq 0, k_2 \geq 0} \ln[Ak_0^\alpha - k_1] + \beta \ln[Ak_1^\alpha - k_2]$$

- We shall call it the Sequence Problem (SP).

Finite Horizon

- The first-order condition with respect to k_1 is

$$\frac{-1}{Ak_0^{\alpha} - k_1} + \beta \frac{\alpha A k_1^{\alpha-1}}{Ak_1^{\alpha} - k_2} = 0$$

- Evidently, $k_2 = 0$ is optimal because increasing k_2 will only decrease the level of utility. Thus,

$$\frac{Ak_1^{\alpha} - 0}{Ak_0^{\alpha} - k_1} = \frac{c_1}{c_0} = \beta \alpha A k_1^{\alpha-1}$$

- This is the Euler equation.

Finite Horizon

- Solve the Euler equation for k_1 to obtain

$$k_1 = \frac{\alpha\beta}{1+\alpha\beta} A k_0^{\alpha}$$

- Thus,

$$c_0 = A k_0^{\alpha} - k_1 = \frac{1}{1+\alpha\beta} A k_0^{\alpha}$$
$$c_1 = A k_1^{\alpha} - 0 = A k_1^{\alpha}$$

- The solution is given by the sequences of numbers.

Finite Horizon

- Let us now solve the same problem differently.
  - By the method of **backward induction**.
- Suppose that we are <u>in the terminal period</u>.
- The problem:

$$\max_{c_1, k_2 \geq 0} \ln c_1$$

- Subject to

$$k_2 = Ak_1^{\alpha} - c_1$$

- In this period, k_1 is given because the level has been determined by <u>your own actions in the past</u>.

Finite Horizon

- The problem reduces to:
$$\max_{k_2 \geq 0} \ln[Ak_1^\alpha - k_2]$$
- Evidently, the optimal choice is $k_2 = 0$.
- The maximized utility (i.e., indirect utility) is:
$$\ln[Ak_1^\alpha] \equiv v_1(k_1)$$
- Indirect utility is a function of the state variable k_1.
  - Given the predetermined level of capital, the household makes the best choice.
- This function is called the **value function**.

Finite Horizon

- Now consider the problem in the initial period.
- The household takes into account that c_1 and k_2 will be chosen optimally in the next period, and the result is summarized by $v_1(k_1)$.
- Thus, the problem becomes:

$$\max_{c_0, k_1 \geq 0} \ln c_0 + \beta v_1(k_1)$$

- Subject to

$$k_1 = Ak_0^{\alpha} - c_0$$

- In this period, k_1 can be chosen by the household.

Finite Horizon

- The problem is equivalent to:

$$\max_{k_1 \geq 0} \ln[Ak_0^\alpha - k_1] + \beta \ln[Ak_1^\alpha]$$

- The first-order condition with respect to k_1 is

$$\frac{-1}{Ak_0^\alpha - k_1} + \beta \frac{\alpha A k_1^{\alpha-1}}{Ak_1^\alpha} = 0$$

- Solving it for k_1 to obtain

$$k_1 = \frac{\alpha\beta}{1 + \alpha\beta} Ak_0^\alpha$$

Finite Horizon

- The optimal consumption levels are:

$$c_0 = \frac{1}{1+\alpha\beta} A k_0^{\alpha}$$
$$c_1 = A k_1^{\alpha}$$

- The solutions from the two different methods are the same.
- In principle, for <u>any</u> large T, a T-period problem can be solved backward from the terminal period.
- For infinite-horizon problems, there is <u>no terminal period</u>. Backward induction method does not work.

Infinite Horizon

- Consider the neoclassical growth model:

$$\max_{\{c_t\}_{t=0}^{\infty},\{k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_t)$$

- Subject to given k_0 and

$$k_{t+1} = Ak_t^{\alpha} - c_t$$

  - The depreciation rate is $\delta = 1$ for simplicity.
- This problem reduces to:

$$\max_{\{k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(Ak_t^{\alpha} - k_{t+1})$$

Infinite Horizon

- Suppose that we somehow find the optimal sequence $\{k_{t+1}\}_{t=0}^{\infty}$.
- We can then define the **value function** v as:

$$v(k_0) = \max_{\{k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(Ak_t^{\alpha} - k_{t+1})$$

- This is the maximized life-time utility as <u>a function of the initial capital stock</u>.

Infinite Horizon

- For now, let us forget about maximization.
- Take <u>any feasible sequence</u> $\{k_{t+1}\}_{t=0}^{\infty}$.
- Define $\bar{v}$ nearly identical to v (except for max) by:

$$\bar{v}(k_0) = \sum_{t=0}^{\infty} \beta^t u(Ak_t^{\alpha} - k_{t+1})$$

$$= u(Ak_0^{\alpha} - k_1) + \beta \sum_{t=1}^{\infty} \beta^{t-1} u(Ak_t^{\alpha} - k_{t+1})$$

$$= u(Ak_0^{\alpha} - k_1) + \beta \bar{v}(k_1)$$

Infinite Horizon

- We obtain:

$$\bar{v}(k_t) = u(Ak_t^\alpha - k_{t+1}) + \beta \bar{v}(k_{t+1})$$

- Very nice expression!
- But, can we introduce maximization here?
- In other words, is it OK to write the following?

$$v(k_t) = \max_{k_{t+1}} u(Ak_t^\alpha - k_{t+1}) + \beta v(k_{t+1})$$

- (With some cautions) it is OK.
- We call it the **Bellman equation**.

Infinite Horizon

- More generally, the Bellman equation is
$$v(k_t) = \max_{k_{t+1}} F(k_t, k_{t+1}) + \beta v(k_{t+1})$$
- $F(k_t, k_{t+1})$ is (generally defined) payoff function.
- Value function v is unknown.
  - The solution to the Bellman equation is the shape of v.
  - Bellman equation is a functional equation.
- Our objective is to find the value **function** $v(k_t)$ and the (time-invariant) policy **function**:
$$c_t = g(k_t)$$

Infinite Horizon

- We are <u>not</u> looking for the optimal sequence.
- We are looking for the <u>best response function</u>.
- Thus, time subscript is **irrelevant**.
- The Bellman equation is generally written as:
$$v(x) = \max_{y} F(x, y) + \beta v(y)$$
- v is used to define v itself (**recursive**)
- v is the unknown (**functional equation**).
- How can we solve it?

Mathematical Background

Solution as the Limit of Convergent Sequence of Numbers

- Finding a solution to $2x + 1 = 0$ is easy.
- We say that an equation has a **closed-form solution** if the solution is obtained by a finite number of operations (addition, multiplication, log transformation, etc.).
- Let us find a solution to the following:
$$-x^3 - x^5 + 3 = 0$$
- This equation has <u>no</u> closed-form solution.

Solution as the Limit of Convergent Sequence of Numbers

- Let me use Mathematica (or Wolfram Alpha web) to numerically find the solution.
- By the figure, we are 100% sure that there is a solution.
- How can a computer find the solution?

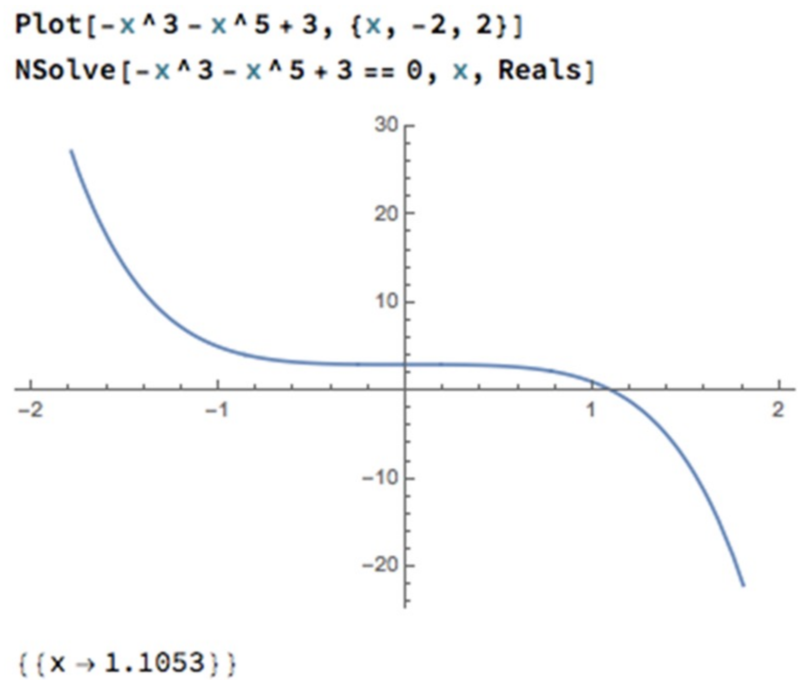

Solution as the Limit of Convergent Sequence of Numbers

- Equation is generally written as:
$$f(x) = 0$$
- In other words, finding a solution is equivalent to finding a **zero point** of function f.
- Another representation is
$$F(x) = x$$
- Here, finding a solution is equivalent to finding a **fixed point** of map F.
- If we define $f(x) = F(x) - x$, then we can switch these two representations.

Solution as the Limit of Convergent Sequence of Numbers

- To understand the basic idea of finding a solution as the limit of convergent sequence, consider **Newton's method**, a famous algorithm for finding a numerical solution using computer.
- Let x^* be the solution to $f(x) = 0$.
- Suppose that we have an initial guess about the solution, x_n.
- Apply linear approximation on $f(x)$, evaluated at x_n, to obtain:

$$f(x) \approx f(x_n) + f'(x_n)(x - x_n) \equiv g(x)$$

- We can easily solve the linear equation $g(x) = 0$.

Solution as the Limit of Convergent Sequence of Numbers

- The solution is

$$x = x_n - \frac{f(x_n)}{f'(x_n)}$$

- This is our new "guess" from the initial guess x_n.
- Consider the sequence generated by the following difference equation:

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}$$

- As $n \to \infty$, $x_n \to x^*$. The solution is found as the limit of a convergent **sequence** of numbers.

Convergent Sequence of Functions (?)

- Key idea: When a closed-form solution is not available, we look for a convergent **sequence**.
- Suppose we have an initial guess about the value **function** v_n, not a number.
- We construct a new guess v_{n+1} by
$$v_{n+1}(x) = \max_y F(x, y) + \beta v_n(y)$$
- Our hope is, as $n \to \infty$, $v_n \to v$.
- This requires us to work with a sequence of functions on a space filled with functions.

Metric Space (Distance Space)

- The core concept is how we measure the <u>distance</u> between two functions.
  - Otherwise, we cannot talk about convergence.
- Definition: A **metric space (or distance space)** is a set S, together with a metric (distance function) $\rho: S \times S \rightarrow \boldsymbol{R}$, such that for all $x, y, z \in S$:

1. $\rho(x, y) \geq 0$, with equality if and only if $x = y$;
2. $\rho(x, y) = \rho(y, x)$ (Symmetry); and
3. $\rho(x, z) \leq \rho(x, y) + \rho(y, z)$ (Triangle inequality).

Examples of Metric Space

- $\boldsymbol{R^1}$: The set of all real numbers with distance
$$\rho(x, y) = |x - y|$$
- This one is very easy.
- When we measure the distance between two real numbers, we use $|x - y|$ as our measure.
- It is also straightforward to prove that this measure satisfies all the three properties of a metric space.

Examples of Metric Space

- $\boldsymbol{R^n}$: n-dimensional Euclidean space is a metric space with distance

$$\rho(x, y) = \sqrt{\Sigma_{k=1}^{n}(x_k - y_k)^2}$$

- Any high-school student knows how to measure the distance between two points on the two-dimensional space.
- This is just an extension to a higher-dimensional space.

Examples of Metric Space

- Measuring the distance between two functions can be tricky because there are too many points (or an infinity of points) to consider.
- **Function space** $C_{[a,b]}$: The set of all continuous functions defined on the closed interval $[a, b]$ with distance

$$\rho(f, g) = \max_{a \le t \le b} |f(t) - g(t)|$$

- The idea is that if the largest gap between the two functions is nearly zero, then it is safe to say that the two functions are sufficiently similar.

Cauchy Sequence

- Definition: A sequence $\{x_n\}_{n=0}^{\infty}$ in S **converges** to $x \in S$, if for each $\varepsilon > 0$, there exits N_ε such that $\rho(x_n, x) < \varepsilon$ for all $n \geq N_\varepsilon$.
- Definition: A sequence $\{x_n\}_{n=0}^{\infty}$ in S is a **Cauchy sequence** if for each $\varepsilon > 0$, there exits N_ε such that $\rho(x_n, x_m) < \varepsilon$ for all $n, m \geq N_\varepsilon$.
- Key idea: Cauchy offers a convergence concept that <u>does not</u> require our knowledge about the limit point.

Complete Metric Space

- Definition: A metric space (S, ρ) is **complete** if every Cauchy sequence in S converges to an element in S.
  - In a complete metric space, we can verify the existence of a limit point by showing that a sequence is Cauchy.
- Fact: The set of real numbers $\boldsymbol{R}$ with metric $\rho(x, y) = |x - y|$ is a complete metric space.

Contraction Mapping

- Definition: Let (S, ρ) be a metric space and $T: S \to S$ be a function mapping S into itself. T is a **contraction mapping** if for some $\beta \in (0,1)$,
$$\rho(Tx, Ty) \leq \beta\rho(x, y)$$
for all $x, y \in S$.
- Mapping and function are the same meaning. We use the term "mapping" because we are talking about a function that transforms a function, and it sounds confusing.
  - T can also be a function that transforms numbers.

Example

- Consider $f(x) = 0.9x + 1$.
- To verify whether f is a contraction, consider two elements (numbers), x and y.
- The question is whether we can find $\beta \in (0,1)$ such that

$$\rho\big(f(x), f(y)\big) \leq \beta\rho(x, y)$$

Example

- The distance between two numbers $f(x)$ and $f(y)$ is
$$\rho\big(f(x), f(y)\big) = |0.9x - 0.9y\ | = 0.9|x - y|$$
- Thus, for any β satisfying $0.9 \leq \beta < 1$, we can show that
$$\rho\big(f(x), f(y)\big) = 0.9|x - y| \leq \beta|x - y| = \beta\rho(x, y)$$
- Thus, f is a contraction mapping.

Contraction Mapping

Every contraction mapping is uniformly continuous.

Proof) Suppose $\rho(x, y) < \delta$ for some $\delta > 0$. Then,

$$\beta\rho(x, y) < \beta\delta$$

Any contraction mapping satisfies

$$\rho(Tx, Ty) \leq \beta\rho(x, y) < \beta\delta$$

If we define $\varepsilon \equiv \beta\delta$, then, for any $\varepsilon > 0$, there is δ $(= \varepsilon/\beta)$ such that

$$\rho(x, y) < \delta \Rightarrow \rho(Tx, Ty) < \varepsilon$$

This is a definition of continuity. End of the proof.

Contraction Mapping Theorem

- Theorem: (a) Every contraction mapping defined on a complete metric space S has a unique fixed point x in S; and (b) for any $x_0 \in S$, $\rho(T^n x_0, x) \leq \beta^n \rho(x_0, x)$ for $n = 0,1,2, \dots$
- Sketch of Proof)
  - Given x_0, construct a sequence: $x_1 = Tx_0, x_2 = Tx_1 = T^2 x_0, \dots, x_n = T^n x_0$.
  - Contraction ⇒ for $n \leq m$, $\rho(x_n, x_m) \leq \beta^n \rho(x_0, x_{m-n})$.
  - Triangle inequality ⇒ $\rho(x_n, x_m) \leq \varepsilon$ for $n, m \geq N_\varepsilon$. Cauchy.
  - <u>Since S is complete, Cauchy implies a limit x.</u>
  - $Tx = T \lim_{n\to\infty} x_n = \lim_{n\to\infty} Tx_n = \lim_{n\to\infty} x_{n+1} = x$.
  - Two fixed points result in a contradiction. Uniqueness proven.

Application: Difference Equation

- Consider a scalar linear difference equation
$$x_{t+1} = ax_t + b \equiv f(x_t)$$
- Let $\rho(x,y) = |x - y|$.
- There is a unique fixed point x if f is a contraction mapping. Thus, if for some $\beta \in (0,1)$,
$$\rho\big(f(x), f(y)\big) \leq \beta\rho(x,y)$$
$$\Leftrightarrow$$
$$|f(x) - f(y)| \leq \beta|x - y|$$
$$\Leftrightarrow$$
$$|ax - ay| \leq \beta|x - y|$$

Application: Difference Equation

- We further rewrite the condition as
$$\frac{|ax - ay|}{|x - y|} = |a| \leq \beta \in (0,1)$$
- Thus, there exists a unique fixed point if $|a| < 1$.
- (b) of the Theorem implies a convergent sequence. This implies that for any initial condition, we have a convergent sequence, and its limit point is the fixed point.
  - To be brief, the fixed point is globally (asymptotically) stable.

Neoclassical Growth Model

- Consider the Bellman equation:
$$v(k_t) = \max_{k_{t+1}} u(Ak_t^\alpha - k_{t+1}) + \beta v(k_{t+1})$$
- Rewrite it as:
$$v(x) = \max_{y} u(Ax^\alpha - y) + \beta v(y)$$
- Using theorems we skip, we can prove that the mapping (including all operations such as maximization) from a function to another one is a contraction mapping. ⇒ unique v exits.

Further Readings

- Simon and Blume, *Mathematics for Economists*, Norton, 1994. Chapter 29.
- A. N. Kolmogorov and S. V. Fomin, *Introductory Real Analysis*, Dover, 1970. Chapter 2.
- Nancy Stokey & Robert Lucas, *Recursive Methods in Economic Dynamics*, Harvard University Press, 1989. Chapters 3 & 4.
- Adda and Cooper, *Dynamic Economics*, MIT Press, 2003. Chapter 2.

Applications

Envelope Condition

- Consider the neoclassical growth model with perfect depreciation $\delta = 1$ (for simplicity):

$$\max_{\{c_t\}_{t=0}^{\infty},\{k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_t)$$

- Subject to given k_0 and

$$k_{t+1} = Ak_t^{\alpha} - c_t$$

- Thus, the sequence problem is:

$$\max_{\{k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(Ak_t^{\alpha} - k_{t+1})$$

Envelope Condition

- From the original (sequence) problem, we obtain the associated Bellman equation:

$$v(k_t) = \max_{k_{t+1}} u(Ak_t^\alpha - k_{t+1}) + \beta v(k_{t+1})$$

- Suppose the utility function is $\ln c$. Then,

$$v(k_t) = \max_{k_{t+1}} \ln(Ak_t^\alpha - k_{t+1}) + \beta v(k_{t+1})$$

- FOC with respect to k_{t+1}:

$$\frac{1}{Ak_t^\alpha - k_{t+1}} = \beta v'(k_{t+1})$$

Envelope Condition

- Remember
$$k_{t+1} = Ak_t^\alpha - c_t$$
- An alternative representation of the Bellman equation is
$$v(k_t) = \max_{c_t} u(c_t) + \beta v(Ak_t^\alpha - c_t)$$
- FOC with respect to c_t is
$$u'(c_t) = \beta v'(Ak_t^\alpha - c_t)$$
- With log utility, this is identical to the FOC on the previous page.

Envelope Condition

- Note that FOC contains the unknown v'.
- Thanks to the recursive structure, we can calculate v' from the Bellman equation itself:

$$v(k_t) = \max_{k_{t+1}} \ln(Ak_t^\alpha - k_{t+1}) + \beta v(k_{t+1})$$

- The derivative with respect to the current sate is

$$v'(k_t) = \frac{\alpha A k_t^{\alpha-1}}{Ak_t^\alpha - k_{t+1}}$$

- This is called the **Envelope condition**.

Envelope Condition

- From these conditions,

$$\frac{1}{Ak_t^{\alpha} - k_{t+1}} = \beta \frac{\alpha A k_{t+1}^{\alpha-1}}{Ak_{t+1}^{\alpha} - k_{t+2}}$$

- This is the **Euler equation**. To see this,

$$\frac{c_{t+1}}{c_t} = \beta\alpha A k_{t+1}^{\alpha-1} \Leftrightarrow \frac{u'(c_t)}{\beta u'(c_{t+1})} = f'(k_{t+1})$$

- With FOC and the Envelope condition, we can obtain the same set of equations as in the Lagrangian method.

Envelope Condition

- With the budget constraint and the transversality condition, <u>we can study the optimal sequence</u> instead of finding the policy function.
- In other words, (in many, <u>not all</u>, applications) the Bellman equation can be used as a tool to obtain the first-order conditions for optimality just like the Lagrange method.
- For that purpose, the Envelope condition plays a central role.

When to Use DP?

- DP is particularly useful <u>when the state is discrete</u>.
- There are many real-world examples in which the state is discrete:
    - employment status (employed/unemployed)
    - marital status (married/single)
    - success/failure
    - high/low
    - Infected/Not infected
- Some people use DP extensively, others do not use it at all.

Further Readings

- Ljungqvist & Sargent, *Recursive Macroeconomic Theory*, 4th edition, MIT Press, 2018.
  - Any addition is fine.
- It starts with DP.
- Over 1,400 pages long!!
  - It takes forever to read the entire book.

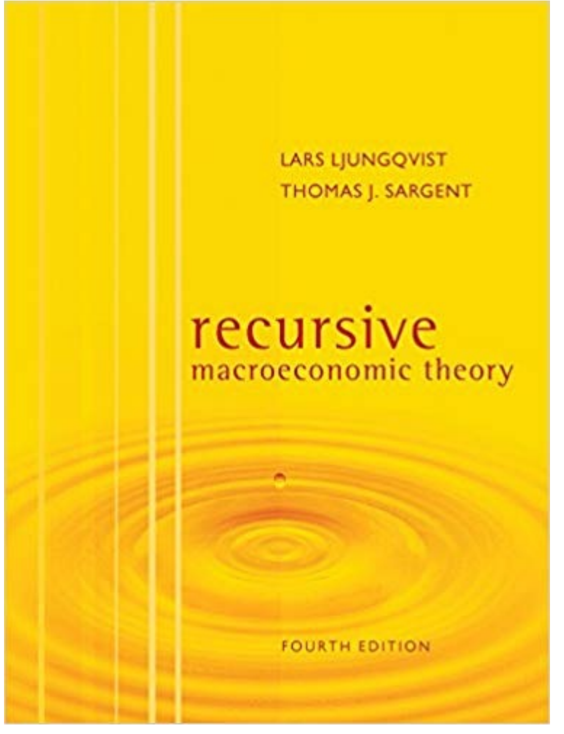

Further Readings

- https://quantecon.org/
- You can find lectures on economic dynamics using Python.
- Good idea to spend some time on this site during the summer.