

Dynamic Programming

Peter Ireland*

ECON 772001 - Math for Economists
Boston College, Department of Economics

Fall 2018

We have now studied two ways of solving dynamic optimization problems, one based on the Kuhn-Tucker theorem and the other based on the maximum principle. These two methods both lead us to the same sets of optimality conditions; they differ only in terms of how those optimality conditions are derived.

Here, we will consider a third way of solving dynamic optimization problems: the method of dynamic programming. We will see, once again, that dynamic programming leads us to the same set of optimality conditions that the Kuhn-Tucker theorem does; once again, this new method differs from the others only in terms of how the optimality conditions are derived.

While the maximum principle lends itself equally well to dynamic optimization problems set in both discrete time and continuous time, dynamic programming is easiest to apply in discrete time settings. On the other hand, dynamic programming, unlike the Kuhn-Tucker theorem and the maximum principle, can be used quite easily to solve problems in which optimal decisions must be made under conditions of uncertainty.

Thus, in our discussion of dynamic programming, we will begin by considering dynamic programming under certainty; later, we will move on to consider stochastic dynamic programming.

References:

Dixit, Chapter 11.

Acemoglu, Chapters 6 and 16.

Dynamic programming was invented by Richard Bellman in the late 1950s, around the same time that Pontryagin and his colleagues were working out the details of the maximum principle. A famous early reference is:

Richard Bellman. *Dynamic Programming*, 1957.

A very comprehensive reference with many economic examples is

Nancy L. Stokey and Robert E. Lucas, Jr. with Edward C. Prescott. *Recursive Methods in Economic Dynamics*, 1989.

*These lecture notes are licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0) License. <http://creativecommons.org/licenses/by-nc-sa/4.0/>

1 Dynamic Programming Under Certainty

1.1 A Perfect Foresight Dynamic Optimization Problem in Discrete Time

No uncertainty

Discrete time, infinite horizon: $t = 0, 1, 2, \dots$

y_t = stock, or state, variable

z_t = flow, or control, variable

Objective function:

$$\sum_{t=0}^{\infty} \beta^t F(y_t, z_t; t)$$

$1 > \beta > 0$ discount factor

Constraint describing the evolution of the state variable

$$Q(y_t, z_t; t) \geq y_{t+1} - y_t$$

or

$$y_t + Q(y_t, z_t; t) \geq y_{t+1}$$

for all $t = 0, 1, 2, \dots$

Constraint applying to variables within each period:

$$c \geq G(y_t, z_t; t)$$

for all $t = 0, 1, 2, \dots$

Constraint on initial value of the state variable:

$$y_0 \text{ given}$$

The problem: choose sequences $\{z_t\}_{t=0}^{\infty}$ and $\{y_t\}_{t=1}^{\infty}$ to maximize the objective function subject to all of the constraints.

Notes:

- a) It is important for the application of dynamic programming that the problem is additively time separable: that is, the values of F , Q , and G at time t must depend only on the values of y_t and z_t at time t .
- b) Once again, it must be emphasized that although the constraints describing the evolution of the state variable and that apply to the variables within each period can each be written in the form of a single equation, these constraints must hold for all $t = 0, 1, 2, \dots$. Thus, each equation actually represents an infinite number of constraints.

1.2 The Kuhn-Tucker Formulation

Let's begin our analysis of this problem by applying the Kuhn-Tucker theorem. That is, let's begin by setting up the Lagrangian and taking first order conditions.

Set up the Lagrangian, recognizing that the constraints must hold for all $t = 0, 1, 2, \dots$:

$$L = \sum_{t=0}^{\infty} \beta^t F(y_t, z_t; t) + \sum_{t=0}^{\infty} \tilde{\mu}_{t+1} [y_t + Q(y_t, z_t; t) - y_{t+1}] + \sum_{t=0}^{\infty} \tilde{\lambda}_t [c - G(y_t, z_t; t)]$$

It will be convenient to define

$$\begin{aligned} \mu_{t+1} &= \beta^{-(t+1)} \tilde{\mu}_{t+1} \Rightarrow \tilde{\mu}_{t+1} = \beta^{t+1} \mu_{t+1} \\ \lambda_t &= \beta^{-t} \tilde{\lambda}_t \Rightarrow \tilde{\lambda}_t = \beta^t \lambda_t \end{aligned}$$

and to rewrite the Lagrangian in terms of μ_{t+1} and λ_t instead of $\tilde{\mu}_{t+1}$ and $\tilde{\lambda}_t$:

$$L = \sum_{t=0}^{\infty} \beta^t F(y_t, z_t; t) + \sum_{t=0}^{\infty} \beta^{t+1} \mu_{t+1} [y_t + Q(y_t, z_t; t) - y_{t+1}] + \sum_{t=0}^{\infty} \beta^t \lambda_t [c - G(y_t, z_t; t)]$$

FOC for z_t , $t = 0, 1, 2, \dots$:

$$\beta^t F_2(y_t, z_t; t) + \beta^{t+1} \mu_{t+1} Q_2(y_t, z_t; t) - \beta^t \lambda_t G_2(y_t, z_t; t) = 0$$

FOC for y_t , $t = 1, 2, 3, \dots$:

$$\beta^t F_1(y_t, z_t; t) + \beta^{t+1} \mu_{t+1} [1 + Q_1(y_t, z_t; t)] - \beta^t \lambda_t G_1(y_t, z_t; t) - \beta^t \mu_t = 0$$

Now, let's suppose that somehow we could solve for μ_t as a function of the state variable y_t :

$$\begin{aligned} \mu_t &= W(y_t; t) \\ \mu_{t+1} &= W(y_{t+1}; t+1) = W[y_t + Q(y_t, z_t; t); t+1] \end{aligned}$$

Then we could rewrite the FOC as:

$$F_2(y_t, z_t; t) + \beta W[y_t + Q(y_t, z_t; t); t+1] Q_2(y_t, z_t; t) - \lambda_t G_2(y_t, z_t; t) = 0 \quad (1)$$

$$W(y_t; t) = F_1(y_t, z_t; t) + \beta W[y_t + Q(y_t, z_t; t); t+1] [1 + Q_1(y_t, z_t; t)] - \lambda_t G_1(y_t, z_t; t) \quad (2)$$

And together with the binding constraint

$$y_{t+1} = y_t + Q(y_t, z_t; t) \quad (3)$$

and the complementary slackness condition

$$\lambda_t [c - G(y_t, z_t; t)] = 0 \quad (4)$$

we can think of (1) and (2) as forming a system of four equations in three unknown variables y_t , z_t , and λ_t and one unknown function $W(\cdot; t)$. This system of equations determines the problem's solution.

Note that since (3) is in the form of a difference equation, finding the problem's solution involves solving a difference equation.

1.3 An Alternative Formulation

Now let's consider the same problem in a slightly different way.

For any given value of the initial state variable y_0 , define the value function

$$v(y_0; 0) = \max_{\{z_t\}_{t=0}^{\infty}, \{y_t\}_{t=1}^{\infty}} \sum_{t=0}^{\infty} \beta^t F(y_t, z_t; t)$$

subject to

$$y_0 \text{ given}$$

$$y_t + Q(y_t, z_t; t) \geq y_{t+1} \text{ for all } t = 0, 1, 2, \dots$$

$$c \geq G(y_t, z_t; t) \text{ for all } t = 0, 1, 2, \dots$$

More generally, for any period t and any value of y_t , define

$$v(y_t; t) = \max_{\{z_{t+j}\}_{j=0}^{\infty}, \{y_{t+j}\}_{j=1}^{\infty}} \sum_{j=0}^{\infty} \beta^j F(y_{t+j}, z_{t+j}; t+j)$$

subject to

$$y_t \text{ given}$$

$$y_{t+j} + Q(y_{t+j}, z_{t+j}; t+j) \geq y_{t+j+1} \text{ for all } j = 0, 1, 2, \dots$$

$$c \geq G(y_{t+j}, z_{t+j}; t+j) \text{ for all } j = 0, 1, 2, \dots$$

Note that the value function is a maximum value function.

Now consider expanding the definition of the value function by separating out the time t components:

$$v(y_t; t) = \max_{z_t, y_{t+1}} [F(y_t, z_t; t) + \max_{\{z_{t+j}\}_{j=1}^{\infty}, \{y_{t+j}\}_{j=2}^{\infty}} \sum_{j=1}^{\infty} \beta^j F(y_{t+j}, z_{t+j}; t+j)]$$

subject to

$$y_t \text{ given}$$

$$y_t + Q(y_t, z_t; t) \geq y_{t+1}$$

$$y_{t+j} + Q(y_{t+j}, z_{t+j}; t+j) \geq y_{t+j+1} \text{ for all } j = 1, 2, 3, \dots$$

$$c \geq G(y_t, z_t; t)$$

$$c \geq G(y_{t+j}, z_{t+j}; t+j) \text{ for all } j = 1, 2, 3, \dots$$

Next, relabel the time indices:

$$v(y_t; t) = \max_{z_t, y_{t+1}} [F(y_t, z_t; t) + \beta \max_{\{z_{t+1+j}\}_{j=0}^{\infty}, \{y_{t+1+j}\}_{j=1}^{\infty}} \sum_{j=0}^{\infty} \beta^j F(y_{t+1+j}, z_{t+1+j}; t + 1 + j)]$$

subject to

$$\begin{aligned} & y_t \text{ given} \\ & y_t + Q(y_t, z_t; t) \geq y_{t+1} \\ & y_{t+j+1} + Q(y_{t+1+j}, z_{t+1+j}; t + 1 + j) \geq y_{t+1+j+1} \text{ for all } j = 0, 1, 2, \dots \\ & c \geq G(y_t, z_t; t) \\ & c \geq G(y_{t+1+j}, z_{t+1+j}; t + 1 + j) \text{ for all } j = 0, 1, 2, \dots \end{aligned}$$

Now notice that together, the components for $t + 1 + j$, $j = 0, 1, 2, \dots$ define $v(y_{t+1}; t + 1)$, enabling us to simplify the statement considerably:

$$v(y_t; t) = \max_{z_t, y_{t+1}} F(y_t, z_t; t) + \beta v(y_{t+1}; t + 1)$$

subject to

$$\begin{aligned} & y_t \text{ given} \\ & y_t + Q(y_t, z_t; t) \geq y_{t+1} \\ & c \geq G(y_t, z_t; t) \end{aligned}$$

Or, even more simply:

$$v(y_t; t) = \max_{z_t} F(y_t, z_t; t) + \beta v[y_t + Q(y_t, z_t; t); t + 1] \quad (5)$$

subject to

$$\begin{aligned} & y_t \text{ given} \\ & c \geq G(y_t, z_t; t) \end{aligned}$$

Equation (5) is called the Bellman equation for this problem, and lies at the heart of the dynamic programming approach.

Note that the maximization on the right-hand side of (5) is a static optimization problem, involving no dynamic elements.

By the Kuhn-Tucker theorem:

$$v(y_t; t) = \max_{z_t} F(y_t, z_t; t) + \beta v[y_t + Q(y_t, z_t; t); t + 1] + \lambda_t [c - G(y_t, z_t; t)]$$

The FOC for z_t is

$$F_2(y_t, z_t; t) + \beta v'[y_t + Q(y_t, z_t; t); t + 1] Q_2(y_t, z_t; t) - \lambda_t G_2(y_t, z_t; t) = 0 \quad (6)$$

And by the envelope theorem:

$$v'(y_t; t) = F_1(y_t, z_t; t) + \beta v'[y_t + Q(y_t, z_t; t); t + 1][1 + Q_1(y_t, z_t; t)] - \lambda_t G_1(y_t, z_t; t) \quad (7)$$

Together with the binding constraint

$$y_{t+1} = y_t + Q(y_t, z_t; t) \quad (3)$$

and complementary slackness condition

$$\lambda_t [c - G(y_t, z_t; t)] = 0, \quad (4)$$

we can think of (6) and (7) as forming a system of four equations in three unknown variables y_t , z_t , and λ_t and one unknown function $v(\cdot, t)$. This system of equations determines the problem's solution.

Note once again that since (3) is in the form of a difference equation, finding the problem's solution involves solving a difference equation.

But more important, notice that (6) and (7) are equivalent to (1) and (2) with

$$v'(y_t; t) = W(y_t; t).$$

Thus, we have two ways of solving this discrete time dynamic optimization problem, both of which lead us to the same set of optimality conditions:

- a) Set up the Lagrangian for the dynamic optimization problem and take first order conditions for z_t , $t = 0, 1, 2, \dots$ and y_t , $t = 1, 2, 3, \dots$
- b) Set up the Bellman equation and take the first order condition for z_t and then derive the envelope condition for y_t .

One question remains: How, in practice, can we solve for the unknown value functions $v(\cdot, t)$?

To see how to answer this question, consider two examples:

Example 1: Optimal Growth - Here, it will be possible to solve for v explicitly.

Example 2: Saving Under Certainty - Here, it will not be possible to solve for v explicitly, yet we can learn enough about the properties of v to obtain some useful economic insights.

2 Example 1: Optimal Growth

Here, we will modify the optimal growth example that we solved earlier using the maximum principle in two ways:

- a) We will switch to discrete time in order to facilitate the use of dynamic programming.

- b) Set the depreciation rate for capital equal to $\delta = 1$ in order to obtain a very special case in which an explicit solution for the value function can be found.

Production function:

$$F(k_t) = k_t^\alpha$$

where $0 < \alpha < 1$

k_t = capital (state variable)

c_t = consumption (control variable)

Evolution of the capital stock:

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

for all $t = 0, 1, 2, \dots$

Initial condition:

$$k_0 \text{ given}$$

Utility or social welfare:

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

The social planner's problem: choose sequences $\{c_t\}_{t=0}^{\infty}$ and $\{k_t\}_{t=1}^{\infty}$ to maximize the utility function subject to all of the constraints.

To solve this problem via dynamic programming, use

k_t = state variable

c_t = control variable

Set up the Bellman equation:

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

Now guess that the value function takes the time-invariant form

$$v(k_t; t) = v(k_t) = E + F \ln(k_t),$$

where E and F are constants to be determined.

Using the guess for v , the Bellman equation becomes

$$E + F \ln(k_t) = \max_{c_t} \ln(c_t) + \beta E + \beta F \ln(k_t^\alpha - c_t) \quad (8)$$

FOC for c_t :

$$\frac{1}{c_t} - \frac{\beta F}{k_t^\alpha - c_t} = 0 \quad (9)$$

Envelope condition for k_t :

$$\frac{F}{k_t} = \frac{\alpha\beta F k_t^{\alpha-1}}{k_t^\alpha - c_t} \quad (10)$$

Together with the binding constraint

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

(8)-(10) form a system of four equations in 4 unknowns: c_t , k_t , E , and F .

Equation (9) implies

$$k_t^\alpha - c_t = \beta F c_t$$

or

$$c_t = \left(\frac{1}{1 + \beta F} \right) k_t^\alpha \quad (11)$$

Substitute (11) into the envelope condition (10):

$$\frac{F}{k_t} = \frac{\alpha\beta F k_t^{\alpha-1}}{k_t^\alpha - c_t} \quad (10)$$

$$F k_t^\alpha - F \left(\frac{1}{1 + \beta F} \right) k_t^\alpha = \alpha\beta F k_t^\alpha$$

$$1 - \left(\frac{1}{1 + \beta F} \right) = \alpha\beta$$

Hence

$$\frac{1}{1 + \beta F} = 1 - \alpha\beta \quad (12)$$

Or, equivalently,

$$1 + \beta F = \frac{1}{1 - \alpha\beta}$$

$$\beta F = \frac{1}{1 - \alpha\beta} - 1 = \frac{\alpha\beta}{1 - \alpha\beta}$$

$$F = \frac{\alpha}{1 - \alpha\beta} \quad (13)$$

Substitute (12) into (11) to obtain

$$c_t = (1 - \alpha\beta) k_t^\alpha \quad (14)$$

which shows that it is optimal to consume the fixed fraction $1 - \alpha\beta$ of output.

Evolution of capital:

$$k_{t+1} = k_t^\alpha - c_t = k_t^\alpha - (1 - \alpha\beta) k_t^\alpha = \alpha\beta k_t^\alpha \quad (15)$$



which is in the form of a difference equation for k_t .

Equations (14) and (15) show how the optimal values of c_t and k_{t+1} depend on the state variable k_t and the parameters α and β . Given a value for k_0 , these two equations can be used to construct the optimal sequences $\{c_t\}_{t=0}^{\infty}$ and $\{k_t\}_{t=1}^{\infty}$. Two examples are illustrated in the figure below.



Numerical Solutions to the Optimal Growth Model with Complete Depreciation

Generated using equations (14) and (15). Each example sets $\alpha = 0.33$ and $\beta = 0.99$.

Example 1: $k(0) = 0.01$



Example 2: $k(0) = 1$



In both examples, $c(t)$ converges to its steady state value of 0.388 and $k(t)$ converges to its steady-state value of 0.188.

For the sake of completeness, substitute (14) and (15) back into (8) to solve for E :

$$E + F \ln(k_t) = \max_{c_t} \ln(c_t) + \beta E + \beta F \ln(k_t^\alpha - c_t) \quad (8)$$

$$E + F \ln(k_t) = \ln(1 - \alpha\beta) + \alpha \ln(k_t) + \beta E + \beta F \ln(\alpha\beta) + \alpha\beta F \ln(k_t)$$

Since (13) implies that

$$F = \alpha + \alpha\beta F,$$

this last equality reduces to

$$E = \ln(1 - \alpha\beta) + \beta E + \beta F \ln(\alpha\beta)$$

which leads directly to the solution

$$E = \frac{\ln(1 - \alpha\beta) + \frac{\alpha\beta}{1 - \alpha\beta} \ln(\alpha\beta)}{1 - \beta}$$

3 Example 2: Saving Under Certainty

Here, a consumer maximizes utility over an infinite horizon, $t = 0, 1, 2, \dots$, earning income from labor and from investments.

A_t = beginning-of-period assets

A_t can be negative, that is, the consumer is allowed to borrow

y_t = labor income (exogenous)

c_t = consumption

saving = $s_t = A_t + y_t - c_t$

r = constant interest rate

Evolution of assets:

$$A_{t+1} = (1+r)s_t = (1+r)(A_t + y_t - c_t)$$

Note:

$$A_t + y_t - c_t = \left(\frac{1}{1+r}\right) A_{t+1}$$

$$A_t = \left(\frac{1}{1+r}\right) A_{t+1} + c_t - y_t$$

Similarly,

$$A_{t+1} = \left(\frac{1}{1+r}\right) A_{t+2} + c_{t+1} - y_{t+1}$$

Combining these last two equalities yields

$$A_t = \left(\frac{1}{1+r}\right)^2 A_{t+2} + \left(\frac{1}{1+r}\right) (c_{t+1} - y_{t+1}) + (c_t - y_t)$$

Continuing in this manner yields

$$A_t = \left(\frac{1}{1+r}\right)^T A_{t+T} + \sum_{j=0}^{T-1} \left(\frac{1}{1+r}\right)^j (c_{t+j} - y_{t+j}).$$

Now assume that the sequence $\{A_t\}_{t=0}^{\infty}$ must remain bounded (while borrowing is allowed, unlimited borrowing is ruled out), and take the limit as $T \rightarrow \infty$ to obtain

$$A_t = \sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j (c_{t+j} - y_{t+j})$$

or

$$A_t + \sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j y_{t+j} = \sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j c_{t+j}. \quad (16)$$

Equation (16) takes the form of an infinite horizon budget constraint, indicating that over the infinite horizon beginning at any period t , the consumer's sources of funds include assets A_t and the present value of current and future labor income, while the consumer's use of funds is summarized by the present value of current and future consumption.

The consumer's problem: choose the sequences $\{s_t\}_{t=0}^{\infty}$ and $\{A_t\}_{t=1}^{\infty}$ to maximize the utility function

$$\sum_{t=0}^{\infty} \beta^t u(c_t) = \sum_{t=0}^{\infty} \beta^t u(A_t + y_t - s_t)$$

subject to the constraints

$$A_0 \text{ given}$$

and

$$(1 + r)s_t \geq A_{t+1}$$

for all $t = 0, 1, 2, \dots$

To solve the problem via dynamic programming, note first that

A_t = state variable

s_t = control variable

Set up the Bellman equation

$$v(A_t; t) = \max_{s_t} u(A_t + y_t - s_t) + \beta v(A_{t+1}; t + 1) \text{ subject to } (1 + r)s_t \geq A_{t+1}$$

$$v(A_t; t) = \max_{s_t} u(A_t + y_t - s_t) + \beta v[(1 + r)s_t; t + 1]$$

FOC for s_t :

$$-u'(A_t + y_t - s_t) + \beta(1 + r)v'[(1 + r)s_t; t + 1] = 0$$

Envelope condition for A_t :

$$v'(A_t; t) = u'(A_t + y_t - s_t)$$

Use the constraints to rewrite these optimality conditions as

$$u'(c_t) = \beta(1 + r)v'(A_{t+1}; t + 1) \tag{17}$$

and

$$v'(A_t; t) = u'(c_t) \tag{18}$$

Since (18) must hold for all $t = 0, 1, 2, \dots$, it implies

$$v'(A_{t+1}; t + 1) = u'(c_{t+1})$$

Substitute this result into (17) to obtain:

$$u'(c_t) = \beta(1 + r)u'(c_{t+1}) \tag{19}$$

Now make 2 extra assumptions:

- a) $\beta(1+r) = 1$ or $1+r = 1/\beta$, the interest rate equals the discount rate
- b) u is strictly concave

Under these 2 additional assumptions, (19) implies

$$u'(c_t) = u'(c_{t+1})$$

or

$$c_t = c_{t+1}$$

And since this last equation must hold for all $t = 0, 1, 2, \dots$, it implies

$$c_t = c_{t+j} \text{ for all } j = 0, 1, 2, \dots$$

Now, return to (16):

$$A_t + \sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j y_{t+j} = \sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j c_{t+j}. \quad (16)$$

$$A_t + \sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j y_{t+j} = c_t \sum_{j=0}^{\infty} \beta^j \quad (20)$$

FACT: Since $|\beta| < 1$,

$$\sum_{j=0}^{\infty} \beta^j = \frac{1}{1-\beta}$$

To see why this is true, multiply both sides by $1 - \beta$:

$$\begin{aligned} 1 &= \frac{1-\beta}{1-\beta} \\ &= (1-\beta) \sum_{j=0}^{\infty} \beta^j \\ &= (1+\beta+\beta^2+\dots) - \beta(1+\beta+\beta^2+\dots) \\ &= (1+\beta+\beta^2+\dots) - (\beta+\beta^2+\beta^3+\dots) \\ &= 1 \end{aligned}$$

Use this fact to rewrite (20):

$$A_t + \sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j y_{t+j} = \left(\frac{1}{1-\beta}\right) c_t$$

or

$$c_t = (1-\beta) \left[A_t + \sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j y_{t+j} \right] \quad (21)$$

Equation (21) indicates that it is optimal to consume a fixed fraction $1-\beta$ of wealth at each date t , where wealth consists of value of current asset holdings and the present discounted value of future labor income. Thus, (21) describes a version of the permanent income hypothesis.

4 Stochastic Dynamic Programming

4.1 A Dynamic Stochastic Optimization Problem

Discrete time, infinite horizon: $t = 0, 1, 2, \dots$

y_t = state variable

z_t = control variable

ε_{t+1} = random shock, which is observed at the beginning of $t + 1$

Thus, when z_t is chosen:

ε_t is known . . .

. . . but ε_{t+1} is still viewed as random.

The shock ε_{t+1} may be serially correlated, but will be assumed to have the Markov property (i.e., to be generated by a Markov process): the distribution of ε_{t+1} depends on ε_t , but not on $\varepsilon_{t-1}, \varepsilon_{t-2}, \varepsilon_{t-3}, \dots$

For example, ε_{t+1} may follow a first-order autoregressive process:

$$\varepsilon_{t+1} = \rho\varepsilon_t + \eta_{t+1}.$$

Now, the full state of the economy at the beginning of each period is described jointly by the pair of values for y_t and ε_t , since the value for ε_t is relevant for forecasting, that is, forming expectations of, future values of ε_{t+j} , $j = 1, 2, 3, \dots$

Objective function:

$$E_0 \sum_{t=0}^{\infty} \beta^t F(y_t, z_t, \varepsilon_t)$$

$1 > \beta > 0$ discount factor

E_0 = expected value as of $t = 0$

Constraint describing the evolution of the state variable

$$y_t + Q(y_t, z_t, \varepsilon_{t+1}) \geq y_{t+1}$$

for all $t = 0, 1, 2, \dots$ and for all possible realizations of ε_{t+1}

Note that the constraint implies that the randomness in ε_{t+1} induces randomness into y_{t+1} as well: in particular, the value of y_{t+1} does not become known until ε_{t+1} is observed at the beginning of $t + 1$ for all $t = 0, 1, 2, \dots$

Note, too, that the sequential, period-by-period, revelation of values for ε_t , $t = 0, 1, 2, \dots$, also generates sequential growth in the information set available to the agent solving the problem:

At the beginning of period $t = 0$, the agent knows

$$I_0 = \{y_0, \varepsilon_0\}$$

At the beginning of period $t = 1$, the agent's information set expands to

$$I_1 = \{y_1, \varepsilon_1, y_0, \varepsilon_0\}$$

And, more generally, at the beginning of period $t = 0, 1, 2, \dots$, the agent's information set is given by

$$I_t = \{y_t, \varepsilon_t, y_{t-1}, \varepsilon_{t-1}, \dots, y_0, \varepsilon_0\}$$

so that conditional expectations of future variables are defined implicitly with respect to this growing information set: for any variable X_{t+j} whose value becomes known at time $t + j$, $j = 0, 1, 2, \dots$:

$$E_t X_{t+j} = E(X_{t+j} | I_t) = E(X_{t+j} | y_t, \varepsilon_t, y_{t-1}, \varepsilon_{t-1}, \dots, y_0, \varepsilon_0)$$

The role of the additive time separability of the objective function, the similar "additive time separability" that is built into the constraints, and the Markov property of the shocks is to make the most recent values of y_t and ε_t sufficient statistics for I_t , so that within the confines of this problem,

$$E_t(X_{t+j} | y_t, \varepsilon_t, y_{t-1}, \varepsilon_{t-1}, \dots, y_0, \varepsilon_0) = E_t(X_{t+j} | y_t, \varepsilon_t).$$

Note, finally, that the randomness in y_{t+1} induced by the randomness in ε_{t+1} also introduces randomness into the choice of z_{t+1} from the perspective of time t :

$$\begin{aligned} & \text{Given } (y_t, \varepsilon_t), \text{ choose } z_t \\ \Rightarrow & \text{ Given } (y_t, z_t) \text{ the realization of } \varepsilon_{t+1} \text{ determines } (y_{t+1}, \varepsilon_{t+1}) \\ \Rightarrow & \text{ Given } (y_{t+1}, \varepsilon_{t+1}), \text{ choose } z_{t+1} \end{aligned}$$

This makes the number of choice variables, as well as the number of constraints, quite large.

The problem: choose contingency plans for z_t , $t = 0, 1, 2, \dots$, and y_t , $t = 1, 2, 3, \dots$, to maximize the objective function subject to all of the constraints.

Notes:

a) In order to incorporate uncertainty, we have really only made two adjustments to the problem:

First, we have added the shock ε_t to the objective function for period t and the shock ε_{t+1} to the constraint linking periods t and $t + 1$.

And second, we have assumed that the planner cares about the expected value of the objective function.

- b) For simplicity, the functions F and Q are now assumed to be time-invariant, although now they depend on the shock as well as on the state and control variable.
- c) For simplicity, we have also dropped the second set of constraints, $c \geq G(y_t, z_t)$. Adding them back is straightforward, but complicates the algebra.
- d) In the presence of uncertainty, the constraint

$$y_t + Q(y_t, z_t, \varepsilon_{t+1}) \geq y_{t+1}$$

must hold, not only for all $t = 0, 1, 2, \dots$, but for all possible realizations of ε_{t+1} as well. Thus, this single equation can actually represent a very large number of constraints.

- e) The Kuhn-Tucker theorem can still be used to solve problems that feature uncertainty. But because problems with uncertainty can have a very large number of choice variables and constraints, the Kuhn-Tucker theorem can become very cumbersome to apply in practice, since one may have to introduce a very large number of Lagrange multipliers. We will return to the Kuhn-Tucker theorem to see how it works under uncertainty in section 6 of the notes, below. But for now, let's see how dynamic programming can be an easier and more convenient way to solve dynamic stochastic optimization problems.

4.2 The Dynamic Programming Formulation

Once again, for any values of y_0 and ε_0 , define

$$v(y_0, \varepsilon_0) = \max_{\{z_t\}_{t=0}^{\infty}, \{y_t\}_{t=1}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t F(y_t, z_t, \varepsilon_t)$$

subject to

y_0 and ε_0 given

$$y_t + Q(y_t, z_t, \varepsilon_{t+1}) \geq y_{t+1} \text{ for all } t = 0, 1, 2, \dots \text{ and all } \varepsilon_{t+1}$$

More generally, for any period t and any values of y_t and ε_t , define

$$v(y_t, \varepsilon_t) = \max_{\{z_{t+j}\}_{j=0}^{\infty}, \{y_{t+j}\}_{j=1}^{\infty}} E_t \sum_{j=0}^{\infty} \beta^j F(y_{t+j}, z_{t+j}, \varepsilon_{t+j})$$

subject to

y_t and ε_t given

$$y_{t+j} + Q(y_{t+j}, z_{t+j}, \varepsilon_{t+j+1}) \geq y_{t+j+1} \text{ for all } j = 0, 1, 2, \dots \text{ and all } \varepsilon_{t+j+1}$$

Note once again that the value function is a maximum value function.

Now separate out the time t components:

$$v(y_t, \varepsilon_t) = \max_{z_t, y_{t+1}} [F(y_t, z_t, \varepsilon_t) + \max_{\{z_{t+j}\}_{j=1}^{\infty}, \{y_{t+j}\}_{j=2}^{\infty}} E_t \sum_{j=1}^{\infty} \beta^j F(y_{t+j}, z_{t+j}, \varepsilon_{t+j})]$$

subject to

y_t and ε_t given

$$y_t + Q(y_t, z_t, \varepsilon_{t+1}) \geq y_{t+1} \text{ for all } \varepsilon_{t+1}$$

$$y_{t+j} + Q(y_{t+j}, z_{t+j}, \varepsilon_{t+j+1}) \geq y_{t+j+1} \text{ for all } j = 1, 2, 3, \dots \text{ and all } \varepsilon_{t+j+1}$$

Relabel the time indices:

$$v(y_t, \varepsilon_t) = \max_{z_t, y_{t+1}} [F(y_t, z_t, \varepsilon_t) + \beta \max_{\{z_{t+1+j}\}_{j=0}^{\infty}, \{y_{t+1+j}\}_{j=1}^{\infty}} E_t \sum_{j=0}^{\infty} \beta^j F(y_{t+1+j}, z_{t+1+j}, \varepsilon_{t+1+j})]$$

subject to

y_t and ε_t given

$$y_t + Q(y_t, z_t, \varepsilon_{t+1}) \geq y_{t+1} \text{ for all } \varepsilon_{t+1}$$

$$y_{t+j+1} + Q(y_{t+1+j}, z_{t+1+j}, \varepsilon_{t+1+j+1}) \geq y_{t+1+j+1} \text{ for all } j = 0, 1, 2, \dots \text{ and all } \varepsilon_{t+1+j+1}$$

FACT (Law of Iterated Expectations): For any random variable X_{t+j} , realized at time $t+j$, $j = 0, 1, 2, \dots$:

$$E_t E_{t+1} X_{t+j} = E_t X_{t+j}.$$

To see why this fact holds true, consider the following example:

Suppose ε_{t+1} follows the first-order autoregression:

$$\varepsilon_{t+1} = \rho \varepsilon_t + \eta_{t+1}, \text{ with } E_t \eta_{t+1} = 0$$

Hence

$$\varepsilon_{t+2} = \rho \varepsilon_{t+1} + \eta_{t+2}, \text{ with } E_{t+1} \eta_{t+2} = 0$$

or

$$\varepsilon_{t+2} = \rho^2 \varepsilon_t + \rho \eta_{t+1} + \eta_{t+2}.$$

It follows that

$$E_{t+1} \varepsilon_{t+2} = E_{t+1} (\rho^2 \varepsilon_t + \rho \eta_{t+1} + \eta_{t+2}) = \rho^2 \varepsilon_t + \rho \eta_{t+1}$$

and therefore

$$E_t E_{t+1} \varepsilon_{t+2} = E_t (\rho^2 \varepsilon_t + \rho \eta_{t+1}) = \rho^2 \varepsilon_t.$$

It also follows that

$$E_t \varepsilon_{t+2} = E_t (\rho^2 \varepsilon_t + \rho \eta_{t+1} + \eta_{t+2}) = \rho^2 \varepsilon_t.$$

So that in this case as in general

$$E_t E_{t+1} \varepsilon_{t+2} = E_t \varepsilon_{t+2}$$

Using this fact:

$$v(y_t, \varepsilon_t) = \max_{z_t, y_{t+1}} [F(y_t, z_t, \varepsilon_t) + \beta \max_{\{z_{t+1+j}\}_{j=0}^{\infty}, \{y_{t+1+j}\}_{j=1}^{\infty}} E_t E_{t+1} \sum_{j=0}^{\infty} \beta^j F(y_{t+1+j}, z_{t+1+j}, \varepsilon_{t+1+j})]$$

subject to

y_t and ε_t given

$$y_t + Q(y_t, z_t, \varepsilon_{t+1}) \geq y_{t+1} \text{ for all } \varepsilon_{t+1}$$

$$y_{t+j+1} + Q(y_{t+1+j}, z_{t+1+j}, \varepsilon_{t+1+j+1}) \geq y_{t+1+j+1} \text{ for all } j = 0, 1, 2, \dots \text{ and all } \varepsilon_{t+1+j+1}$$

Now use the definition of $v(y_{t+1}, \varepsilon_{t+1})$ to simplify:

$$v(y_t, \varepsilon_t) = \max_{z_t, y_{t+1}} F(y_t, z_t, \varepsilon_t) + \beta E_t v(y_{t+1}, \varepsilon_{t+1})$$

subject to

y_t and ε_t given

$$y_t + Q(y_t, z_t, \varepsilon_{t+1}) \geq y_{t+1} \text{ for all } \varepsilon_{t+1}$$

Or, even more simply:

$$v(y_t, \varepsilon_t) = \max_{z_t} F(y_t, z_t, \varepsilon_t) + \beta E_t v[y_t + Q(y_t, z_t, \varepsilon_{t+1}), \varepsilon_{t+1}] \quad (22)$$

Equation (22) is the Bellman equation for this stochastic problem.

Thus, in order to incorporate uncertainty into the dynamic programming framework, we only need to make two modifications to the Bellman equation:

- a) Include the shock ε_t as an additional argument of the value function.
- b) Add the expectation term E_t in front of the value function for $t+1$ on the right-hand side.

Note that the maximization on the right-hand side of (22) is a static optimization problem, involving no dynamic elements.

Note also that by substituting the constraints into the value function, we are left with an unconstrained problem. Unlike the Kuhn-Tucker approach, which requires many constraints and many multipliers, dynamic programming in this case has no constraints and no multipliers.

The FOC for z_t is

$$F_2(y_t, z_t, \varepsilon_t) + \beta E_t\{v_1[y_t + Q(y_t, z_t, \varepsilon_{t+1}), \varepsilon_{t+1}]Q_2(y_t, z_t, \varepsilon_{t+1})\} = 0 \quad (23)$$

The envelope condition for y_t is:

$$v_1(y_t, \varepsilon_t) = F_1(y_t, z_t, \varepsilon_t) + \beta E_t\{v_1[y_t + Q(y_t, z_t, \varepsilon_{t+1}), \varepsilon_{t+1}][1 + Q_1(y_t, z_t, \varepsilon_{t+1})]\} \quad (24)$$

Equations (23)-(24) coincide exactly with the first-order conditions for z_t and y_t that we would have derived through a direct application of the Kuhn-Tucker theorem to the original, dynamic stochastic optimization problem.

Together with the binding constraint

$$y_{t+1} = y_t + Q(y_t, z_t, \varepsilon_{t+1}) \quad (25)$$

we can think of (23) and (24) as forming a system of three equations in two unknown variables y_t and z_t and one unknown function v . This system of equations determines the problem's solution, given the behavior of the exogenous shocks ε_t .

Note that (25) is in the form of a difference equation; once again, solving a dynamic optimization problem involves solving a difference equation.

5 Example 3: Saving with Multiple Random Returns

This example extends example 2 by:

- a) Introducing $n \geq 1$ assets
- b) Allowing returns on each asset to be random

As in example 2, we will not be able to solve explicitly for the value function, but we will be able to learn enough about its properties to derive some useful economic results.

Since we are extending the example in two ways, assume for simplicity that the consumer receives no labor income, and therefore must finance all of his or her consumption by investing.

A_t = beginning-of-period financial wealth

c_t = consumption

s_{it} = savings allocated to asset $i = 1, 2, \dots, n$

Hence,

$$A_t = c_t + \sum_{i=1}^n s_{it}$$

R_{it+1} = random gross return on asset i , not known until $t + 1$

Hence, when s_{it} is chosen:

R_{it} is known ...

... but R_{it+1} is still viewed as random.

Hence

$$A_{t+1} = \sum_{i=1}^n R_{it+1} s_{it}$$

does not become known until the beginning of $t+1$, even though the s_{it} must be chosen during t .

Utility:

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_t) = E_0 \sum_{t=0}^{\infty} \beta^t u(A_t - \sum_{i=1}^n s_{it})$$

The problem can now be stated as: choose contingency plans for s_{it} for all $i = 1, 2, \dots, n$ and $t = 0, 1, 2, \dots$ and A_t for all $t = 1, 2, 3, \dots$ to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t u(A_t - \sum_{i=1}^n s_{it})$$

subject to

$$A_0 \text{ given}$$

and

$$\sum_{i=1}^n R_{it+1} s_{it} \geq A_{t+1}$$

for all $t = 0, 1, 2, \dots$ and all possible realizations of R_{it+1} for each $i = 1, 2, \dots, n$.

As in the general case, the returns can be serially correlated, but must have the Markov property.

To solve this problem via dynamic programming, let

A_t = state variable

$s_{it}, i = 1, 2, \dots, n$ = control variables

$R_t = [R_{1t}, R_{2t}, \dots, R_{nt}]$ = vector of random returns

The Bellman equation is

$$v(A_t, R_t) = \max_{\{s_{it}\}_{i=1}^n} u(A_t - \sum_{i=1}^n s_{it}) + \beta E_t v(\sum_{i=1}^n R_{it+1} s_{it}, R_{t+1})$$

FOC:

$$-u'(A_t - \sum_{i=1}^n s_{it}) + \beta E_t R_{it+1} v_1(\sum_{i=1}^n R_{it+1} s_{it}, R_{t+1}) = 0$$

for all $i = 1, 2, \dots, n$

Envelope condition:

$$v_1(A_t, R_t) = u'(A_t - \sum_{i=1}^n s_{it})$$

Use the constraints to rewrite the FOC and envelope conditions more simply as

$$u'(c_t) = \beta E_t R_{it+1} v_1(A_{t+1}, R_{t+1})$$

for all $i = 1, 2, \dots, n$ and

$$v_1(A_t, R_t) = u'(c_t)$$

Since the envelope condition must hold for all $t = 0, 1, 2, \dots$, it implies

$$v_1(A_{t+1}, R_{t+1}) = u'(c_{t+1})$$

Hence, the FOC imply that

$$u'(c_t) = \beta E_t R_{it+1} u'(c_{t+1}) \tag{26}$$

must hold for all $i = 1, 2, \dots, n$

Equation (26) generalizes (19) to the case where there is more than one asset and where the asset returns are random. It must hold for all assets $i = 1, 2, \dots, n$, even though each asset may pay a different return ex-post.

In example 2, we combined (19) with some additional assumptions to derive a version of the permanent income hypothesis. Similarly, we can use (26) to derive a version of the famous capital asset pricing model.

For simplicity, let

$$m_{t+1} = \frac{\beta u'(c_{t+1})}{u'(c_t)}$$

denote the consumer's intertemporal marginal rate of substitution.

Then (26) can be written more simply as

$$1 = E_t R_{it+1} m_{t+1} \tag{27}$$

Keeping in mind that (27) must hold for all assets, suppose that there is a risk-free asset, with return R_{t+1}^f that is known during period t . Then R_{t+1}^f must satisfy

$$1 = R_{t+1}^f E_t m_{t+1}$$

or

$$E_t m_{t+1} = \frac{1}{R_{t+1}^f} \tag{28}$$

FACT: For any two random variables x and y ,

$$\text{cov}(x, y) = E[(x - \mu_x)(y - \mu_y)], \text{ where } \mu_x = E(x) \text{ and } \mu_y = E(y).$$

Hence,

$$\begin{aligned} \text{cov}(x, y) &= E[xy - \mu_x y - x \mu_y + \mu_x \mu_y] \\ &= E(xy) - \mu_x \mu_y - \mu_x \mu_y + \mu_x \mu_y \\ &= E(xy) - \mu_x \mu_y \\ &= E(xy) - E(x)E(y) \end{aligned}$$

Or, by rearranging,

$$E(xy) = E(x)E(y) + \text{cov}(x, y)$$

Using this fact, (27) can be rewritten as

$$1 = E_t R_{it+1} m_{t+1} = E_t R_{it+1} E_t m_{t+1} + \text{cov}_t(R_{it+1}, m_{t+1})$$

or, using (28),

$$\begin{aligned} R_{t+1}^f &= E_t R_{it+1} + R_{t+1}^f \text{cov}_t(R_{it+1}, m_{t+1}) \\ E_t R_{it+1} - R_{t+1}^f &= -R_{t+1}^f \text{cov}_t(R_{it+1}, m_{t+1}) \end{aligned} \quad (29)$$

Equation (29) indicates that the expected return on asset i exceeds the risk-free rate only if R_{it+1} is negatively correlated with m_{t+1} .

Does this make sense?

Consider that an asset that acts like insurance pays a high return R_{it+1} during bad economic times, when consumption c_{t+1} is low. Therefore, for this asset:

$$\begin{aligned} \text{cov}_t(R_{it+1}, c_{t+1}) < 0 &\Rightarrow \text{cov}_t[R_{it+1}, u'(c_{t+1})] > 0 \\ &\Rightarrow \text{cov}_t(R_{it+1}, m_{t+1}) > 0 \\ &\Rightarrow E_t R_{it+1} < R_{t+1}^f. \end{aligned}$$

This implication seems reasonable: assets that work like insurance often have expected returns below the risk-free return.

Consider that common stocks tend to pay a high return R_{it+1} during good economic times, when consumption c_{t+1} is high. Therefore, for stocks:

$$\begin{aligned} \text{cov}_t(R_{it+1}, c_{t+1}) > 0 &\Rightarrow \text{cov}_t[R_{it+1}, u'(c_{t+1})] < 0 \\ &\Rightarrow \text{cov}_t(R_{it+1}, m_{t+1}) < 0 \\ &\Rightarrow E_t R_{it+1} > R_{t+1}^f. \end{aligned}$$

This implication also seems to hold true: historically, stocks have had expected returns above the risk-free return.

Recalling once more that (29) must hold for all assets, consider in particular the asset whose return happens to coincide exactly with the representative consumer's intertemporal marginal rate of substitution:

$$R_{t+1}^m = m_{t+1}.$$

For this asset, equation (29) implies

$$\begin{aligned} E_t R_{t+1}^m - R_{t+1}^f &= -R_{t+1}^f \text{cov}_t(R_{t+1}^m, m_{t+1}) \\ E_t m_{t+1} - R_{t+1}^f &= -R_{t+1}^f \text{cov}_t(m_{t+1}, m_{t+1}) = -R_{t+1}^f \text{var}_t(m_{t+1}) \end{aligned}$$

or

$$-R_{t+1}^f = \frac{E_t m_{t+1} - R_{t+1}^f}{\text{var}_t(m_{t+1})} \quad (30)$$

Substitute (30) into the right-hand side of (29) to obtain

$$E_t R_{it+1} - R_{t+1}^f = \frac{\text{cov}_t(R_{it+1}, m_{t+1})}{\text{var}_t(m_{t+1})} (E_t m_{t+1} - R_{t+1}^f)$$

or

$$E_t R_{it+1} - R_{t+1}^f = b_{it} (E_t m_{t+1} - R_{t+1}^f), \quad (31)$$

where

$$b_{it} = \frac{\text{cov}_t(R_{it+1}, m_{t+1})}{\text{var}_t(m_{t+1})}$$

is like the slope coefficient from a regression of R_{it+1} on m_{t+1} .

Equation (31) is a statement of the consumption-based capital asset pricing model, or consumption CAPM. This model links the expected return on each asset to the risk-free rate and the representative consumer's intertemporal marginal rate of substitution.

6 Dynamic Programming In Continuous Time

Although dynamic programming is most useful in the discrete-time case, particularly as a way of handling stochastic problems, it can also be applied in continuous time. Working through the derivations in continuous time shows why the maximum principle provides what is typically the more convenient approach. But doing so is still useful, partly because it yields some additional insights into the maximized Hamiltonian and its derivatives and what they measure.

6.1 A Perfect Foresight Dynamic Optimization Problem in Continuous Time

Let's return to the case of perfect foresight, and start with a discrete-time formulation.

Infinite horizon $t = 0, 1, 2, \dots$

y_t = stock, or state, variable

z_t = flow, or control, variable

Objective function:

$$\sum_{t=0}^{\infty} \beta^t F(y_t, z_t; t)$$

$1 > \beta > 0$ discount factor

Constraint describing the evolution of the state variable:

$$Q(y_t, z_t; t) \geq y_{t+1} - y_t$$

for all $t = 0, 1, 2, \dots$

For simplicity, drop the constraint $c \geq G(y_t, z_t; t)$ imposed on the choice of z_t given y_t for each $t = 0, 1, 2, \dots$

Constraint on the initial value of the state variable:

$$y_0 \text{ given}$$

The discrete-time problem is then to choose sequences $\{z_t\}_{t=0}^{\infty}$ and $\{y_t\}_{t=1}^{\infty}$ to maximize the objective function

$$\sum_{t=0}^{\infty} \beta^t F(y_t, z_t; t)$$

subject to the constraints y_0 given and

$$Q(y_t, z_t; t) \geq y_{t+1} - y_t$$

for all $t = 0, 1, 2, \dots$

The continuous-time analog to this problem has

$y(t)$ = stock, or state, variable

$z(t)$ = flow, or control, variable

Objective function:

$$\int_0^{\infty} e^{-\rho t} F(y(t), z(t); t) dt$$

$\rho > 0$ discount rate

Constraint describing the evolution of the state variable:

$$Q(y(t), z(t); t) \geq \dot{y}(t)$$

for all $t \in [0, \infty)$

Initial condition

$$y(0) \text{ given}$$

The continuous-time problem is to choose continuously differentiable functions $z(t)$ and $y(t)$ for $t \in [0, \infty)$ to maximize the objective function

$$\int_0^{\infty} e^{-\rho t} F(y(t), z(t); t) dt$$

subject to the constraints $y(0)$ given and

$$Q(y(t), z(t); t) \geq \dot{y}(t)$$

for all $t \in [0, \infty)$

6.2 The Dynamic Programming Formulation

For the discrete-time problem, the Bellman equation is

$$v(y_t; t) = \max_{z_t} F(y_t, z_t; t) + \beta v(y_{t+1}; t + 1),$$

where

$$y_{t+1} = y_t + Q(y_t, z_t; t)$$

To translate the Bellman equation into continuous time, consider that in discrete-time, the interval between time periods is $\Delta t = 1$. Hence

$$v(y_t; t) = \max_{z_t} F(y_t, z_t; t)\Delta t + e^{-\rho\Delta t}v(y_{t+\Delta t}; t + \Delta t),$$

where

$$y_{t+\Delta t} = y_t + Q(y_t, z_t; t)\Delta t$$

Consider a first-order Taylor approximation of the second term on the right-hand side of the discrete-time Bellman equation, viewed as a function of $y_{t+\Delta t}$ and Δt , around y_t and $\Delta t = 0$:

$$\begin{aligned} e^{-\rho\Delta t}v(y_{t+\Delta t}; t + \Delta t) &\approx v(y_t; t) + v_y(y_t; t)(y_{t+\Delta t} - y_t) + v_t(y_t; t)\Delta t - \rho v(y_t; t)\Delta t \\ &= v(y_t; t) + v_y(y_t; t)Q(y_t, z_t; t)\Delta t + v_t(y_t; t)\Delta t - \rho v(y_t; t)\Delta t \end{aligned}$$

where v_y and v_t denote the partial derivatives of the value function v with respect to its first and second arguments.

Substitute this expression into the Bellman equation

$$v(y_t; t) \approx \max_{z_t} F(y_t, z_t; t)\Delta t + v(y_t; t) + v_y(y_t; t)Q(y_t, z_t; t)\Delta t + v_t(y_t; t)\Delta t - \rho v(y_t; t)\Delta t$$

Subtract $v(y_t, t)$ from both sides

$$0 \approx \max_{z_t} F(y_t, z_t; t)\Delta t + v_y(y_t; t)Q(y_t, z_t; t)\Delta t + v_t(y_t; t)\Delta t - \rho v(y_t; t)\Delta t$$

Divide through by Δt :

$$0 \approx \max_{z_t} F(y_t, z_t; t) + v_y(y_t; t)Q(y_t, z_t; t) + v_t(y_t; t) - \rho v(y_t; t)$$

And move the terms that don't depend on the control variable z_t to the left-hand side:

$$\rho v(y_t; t) - v_t(y_t; t) \approx \max_{z_t} F(y_t, z_t; t) + v_y(y_t; t)Q(y_t, z_t; t)$$

Finally, note that as $\Delta t \rightarrow 0$, the Taylor approximation for $y_{t+\Delta t}$ becomes exact. Use this fact to replace the \approx with an equal sign, and convert the rest of the notation to continuous time to obtain

$$\rho v(y(t); t) - v_t(y(t); t) = \max_{z(t)} F(y(t), z(t); t) + v_y(y(t); t)Q(y(t), z(t); t) \quad (32)$$

Equation (32) is called the Hamilton-Jacobi-Bellman (HJB) equation for the continuous-time problem. But as we have just seen, it is really just the continuous-time analog to the discrete-time Bellman equation.

The HJB equation takes the form of a partial differential equation, linking the unknown value function $v(y(t), t)$ to its partial derivatives $v_y(y(t), t)$ and $v_t(y(t), t)$. Since partial differential equations are often quite difficult to work with, the dynamic programming approach is not used as frequently as the maximum principle to solve continuous-time dynamic optimization in most areas of economics. In financial economics, however, the stochastic version of (32) has served quite usefully as a starting point for numerous analyses that draw on analytic and computational methods for characterizing the solutions to stochastic differential equations in which randomness in asset prices is driven by continuous-time random walks, or Brownian motions. The Nobel prize-winning economist Robert Merton was a pioneer in initiating this line of research.

Recall that the dynamic programming approach takes a “current value” view of the dynamic optimization problem. For this same problem, the current-value formulation of the maximized Hamiltonian is

$$\tilde{H}(y(t), \theta(t); t) = \max_{z(t)} F(y(t), z(t); t) + \theta(t)Q(y(t), z(t); t) \quad (33)$$

Comparing (32) and (33) suggests that across the two formulations

$$\theta(t) = v_y(y(t); t), \quad (34)$$

consistent with our earlier interpretation of $\theta(t)$ as measuring the current value, at time t , if having an additional unit of the stock variable $y(t)$.

Comparing (32) and (33) also suggests that

$$\tilde{H}(y(t), \theta(t); t) = \rho v(y(t); t) - v_t(y(t); t). \quad (35)$$

To see that this equality also holds, consider the definition of the value function $v(y(t), t)$ as the maximized value of the objective function from time t forward, given the predetermined value $y(t)$ of the stock or state variable:

$$v(y(t), t) = \max_{z(s), s \in [t, \infty)} \int_t^\infty e^{-\rho(s-t)} F(y(s), z(s); s) ds$$

subject to $y(t)$ given and

$$Q(y(s), z(s), s) \geq \dot{y}(s)$$

for all $s \in [t, \infty)$.

Differentiate $v(y(t), t)$ by t to obtain

$$\begin{aligned} & v_y(y(t), t)\dot{y}(t) + v_t(y(t); t) \\ &= \max_{z(s), s \in [t, \infty)} \left[-F(y(t), z(t); t) + \rho \int_t^\infty e^{-\rho(s-t)} F(y(s), z(s); s) ds \right] \end{aligned}$$

or in light of (34) and the definition of $v(y(t), t)$,

$$\theta(t)Q(y(t), z(t); t) + v_t(y(t); t) = \max_{z(s), s \in [t, \infty)} -F(y(t), z(t); t) + \rho v(y(t); t).$$

Note: in these expressions, “ $\max_{z(s), s \in [t, \infty)}$ ” simply means that the functions that follow are being evaluated at the values of $z(s)$, $s \in [t, \infty)$ that solve the problem, not that these functions are themselves being maximized by choice of $z(s)$.

Rearranging terms and simplifying yields

$$\rho v(y(t); t) - v_t(y(t); t) = \max_{z(t)} F(y(t), z(t); t) + \theta(t)Q(y(t), z(t); t)$$

which, in light of (33), coincides with (35).

Returning to the Hamilton-Jacobi-Bellman equation

$$\rho v(y(t); t) - v_t(y(t); t) = \max_{z(t)} F(y(t), z(t); t) + v_y(y(t); t)Q(y(t), z(t); t), \quad (32)$$

the first-order condition for the optimal choice of $z(t)$ on the right-hand side is

$$F_z(y(t), z(t); t) + v_y(y(t); t)Q_z(y(t), z(t); t) = 0, \quad (36)$$

which in light of (34) coincides with the first-order condition that one could derive, instead, with the help of the maximum principle.

Now differentiate both sides of (32) with respect to $y(t)$, invoking the envelope theorem to ignore the dependence of the optimal $z(t)$ on $y(t)$:

$$\begin{aligned} & \rho v_y(y(t); t) - v_{ty}(y(t); t) \\ &= F_y(y(t), z(t); t) + v_{yy}(y(t); t)Q(y(t), z(t); t) + v_y(y(t); t)Q_y(y(t), z(t); t). \end{aligned} \quad (37)$$

Using (34) once again,

$$\theta(t) = v_y(y(t); t), \quad (34)$$

implies

$$\dot{\theta}(t) = v_{yy}(y(t); t)\dot{y}(t) + v_{yt}(y(t); t) = v_{yy}(y(t); t)Q(y(t), z(t); t) + v_{yt}(y(t); t) \quad (38)$$

Equation (34) and (38) reveal that (37) is equivalent to

$$\rho\theta(t) = F_y(y(t), z(t); t) + \dot{\theta}(t) + \theta(t)Q_y(y(t), z(t); t)$$

or

$$\dot{\theta}(t) = \rho\theta(t) - \tilde{H}_y(y(t), \theta(t); t),$$

which again coincides with the optimality conditions that one could derive with the help of the maximum principle.

Yet again we see that the Kuhn-Tucker theorem, the maximum principle, and dynamic programming are just different approaches to deriving the same optimality conditions that characterize the solution to any given constrained optimization problem.

7 Example 4: Optimal Growth

Although the maximum principle is easier to apply in this case, let's consider solving the optimal growth model in continuous time using dynamic programming instead.

In this model, a benevolent social planner or a representative consumer chooses continuously differentiable functions $c(t)$ and $k(t)$ for $t \in [0, \infty)$ to maximize the utility function

$$\int_0^{\infty} e^{-\rho t} \ln(c(t)) dt$$

subject to the constraints $k(0)$ given and

$$k(t)^\alpha - \delta k(t) - c(t) \geq \dot{k}(t)$$

for all $t \in [0, \infty)$.

For this problem, the value function depends on $k(t)$ and not separately on t . Hence, the Hamilton-Jacobi-Bellman equation specializes from

$$\rho v(y(t); t) - v_t(y(t); t) = \max_{z(t)} F(y(t), z(t); t) + v_y(y(t); t)Q(y(t), z(t); t) \quad (32)$$

to

$$\rho v(k(t)) = \max_{c(t)} \ln(c(t)) + v'(k(t))[k(t)^\alpha - \delta k(t) - c(t)]$$

The first-order condition for $c(t)$ is

$$\frac{1}{c(t)} = v'(k(t))$$

and the envelope condition for $k(t)$ is

$$\rho v'(k(t)) = v''(k(t))[k(t)^\alpha - \delta k(t) - c(t)] + v'(k(t))[\alpha k(t)^{\alpha-1} - \delta]$$

As always, the problem with these optimality conditions is that they make reference to the unknown value function $v(k(t))$ – in this case, the first and second derivatives of $v(k(t))$. But the first-order condition gives us a “solution” for $v'(k(t))$. And if we differentiate the first-order condition with respect to t to obtain

$$0 = v''(k(t))c(t)\dot{k}(t) + v'(k(t))\dot{c}(t)$$

or, using the binding constraint to eliminate $\dot{k}(t)$,

$$0 = v''(k(t))[k(t)^\alpha - \delta k(t) - c(t)] + v'(k(t)) \left[\frac{\dot{c}(t)}{c(t)} \right].$$

This last equation can be substituted into the envelope condition to yield

$$\rho v'(k(t)) = -v'(k(t)) \left[\frac{\dot{c}(t)}{c(t)} \right] + v'(k(t))[\alpha k(t)^{\alpha-1} - \delta]$$

Divide through by $v'(k(t))$ and rearrange to obtain

$$\dot{c}(t) = c(t)[\alpha k(t)^{\alpha-1} - \delta - \rho]$$

which together with the constraint

$$\dot{k}(t) = k(t)^\alpha - \delta k(t) - c(t)$$

yields the same system of two differential equations in $c(t)$ and $k(t)$ that we analyzed previously with the help of the phase diagram.

8 Applying the Kuhn-Tucker Theorem in the Case of Uncertainty

Although dynamic programming is usually the most convenient approach to solving dynamic, stochastic optimization problems, the Kuhn-Tucker is applicable, too. For the sake of completeness, let's see how this can be done with an investment in some additional notation.

Consider again the discrete-time, stochastic problem: choose contingency plans for z_t , $t = 0, 1, 2, \dots$, and y_t , $t = 1, 2, 3, \dots$, to maximize the objective function

$$E_0 \sum_{t=0}^{\infty} \beta^t F(y_t, z_t, \varepsilon_t)$$

subject to the constraints that y_0 and ε_0 are given and

$$y_t + Q(y_t, z_t, \varepsilon_{t+1}) \geq y_{t+1}$$

for all $t = 0, 1, 2, \dots$ and all possible realizations of ε_{t+1} .

To describe the “contingency plans” in more detail, without getting too deeply into the measure-theoretic foundations, let

$$H_t = \{\varepsilon_0, \varepsilon_1, \dots, \varepsilon_t\}$$

denote a realized history of shocks through time t .

Then let

$p(H_t)$ = the probability of history H_t

$y_t(H_t)$ = the value of the stock variable at time t following history H_t

$z_t(H_t)$ = the value of the flow variable at time t following the history H_t

$\varepsilon_t(H_t)$ = the value of the shock ε_t realized at the end of history H_t

Ω_t = the set of all possible histories H_t through time t

$H_{t+1} = (H_t, \varepsilon_{t+1})$ = a history through time $t + 1$ that follows the history H_t through time t

$\Omega_{t+1}|H_t$ = the set of all possible histories $H_{t+1} = (H_t, \varepsilon_{t+1})$ at time $t + 1$ that follow the history H_t through t

Note that the period t variables $y_t(H_t)$ and $z_t(H_t)$ depend on H_t , implying that they must be chosen before the realization of ε_{t+1} is observed.

Now restate the problem with the help of this new notation: choose values for $z_t(H_t)$ for all $t = 0, 1, 2, \dots$ and all $H_t \in \Omega_t$ and $y_t(H_t)$ for all $t = 1, 2, 3, \dots$ and all $H_t \in \Omega_t$ to maximize

$$\sum_{t=0}^{\infty} \sum_{H_t \in \Omega_t} \beta^t p(H_t) F[y_t(H_t), z_t(H_t), \varepsilon_t(H_t)]$$

subject to the constraints that $H_0 = \varepsilon_0$ and $y_t(H_0)$ are given and

$$y_t(H_t) + Q[y_t(H_t), z_t(H_t), \varepsilon_{t+1}(H_t, \varepsilon_{t+1})] \geq y_{t+1}(H_t, \varepsilon_{t+1})$$

for all $t = 0, 1, 2, \dots$, all $H_t \in \Omega_t$, and all $H_{t+1} = (H_t, \varepsilon_{t+1}) \in \Omega_{t+1}|H_t$.

Although the new notation hints at the technical details involved in formalizing the choices of $z_t(H_t)$ and $y_t(H_t)$ as random variables “adapted” to the expanding information set that accumulates as time passes and shocks are realized, it also makes clear that the dynamic, stochastic problem is complicated mainly by the fact that there are many choice variables and many constraints.

Still, one can form the Lagrangian in the usual way, taking care to introduce a separate multiplier for each of the many constraints:

$$\begin{aligned}
L &= \sum_{t=0}^{\infty} \sum_{H_t \in \Omega_t} \beta^t p(H_t) F[y_t(H_t), z_t(H_t), \varepsilon_t(H_t)] \\
&+ \sum_{t=0}^{\infty} \sum_{H_t \in \Omega_t} \sum_{(H_t, \varepsilon_{t+1}) \in \Omega_{t+1} | H_t} \tilde{\mu}_{t+1}(H_t, \varepsilon_{t+1}) y_t(H_t) \\
&+ \sum_{t=0}^{\infty} \sum_{H_t \in \Omega_t} \sum_{(H_t, \varepsilon_{t+1}) \in \Omega_{t+1} | H_t} \tilde{\mu}_{t+1}(H_t, \varepsilon_{t+1}) Q[y_t(H_t), z_t(H_t), \varepsilon_{t+1}(H_t, \varepsilon_{t+1})] \\
&- \sum_{t=0}^{\infty} \sum_{H_t \in \Omega_t} \sum_{(H_t, \varepsilon_{t+1}) \in \Omega_{t+1} | H_t} \tilde{\mu}_{t+1}(H_t, \varepsilon_{t+1}) y_{t+1}(H_t, \varepsilon_{t+1})
\end{aligned}$$

Next, let

$$\mu_{t+1}(H_t, \varepsilon_{t+1}) = \frac{\tilde{\mu}_{t+1}(H_t, \varepsilon_{t+1})}{\beta^{t+1} p(H_t, \varepsilon_{t+1})}$$

This transformation extends the similar change in variables that we used in section 1.2 in the perfect foresight case. It puts the Lagrange multipliers in “current value” form so as to make the links to objects from the dynamic programming approach, which also takes a current-value view of the problem, more apparent. The notation, $\mu_{t+1}(H_t, \varepsilon_{t+1})$ also makes clear that the multipliers are also random variables in the stochastic case.

Rewrite the Lagrangian as

$$\begin{aligned}
L &= \\
&\sum_{t=0}^{\infty} \sum_{H_t \in \Omega_t} \beta^t p(H_t) F[y_t(H_t), z_t(H_t), \varepsilon_t(H_t)] \\
&+ \sum_{t=0}^{\infty} \sum_{H_t \in \Omega_t} \sum_{(H_t, \varepsilon_{t+1}) \in \Omega_{t+1} | H_t} \beta^{t+1} p(H_t, \varepsilon_{t+1}) \mu_{t+1}(H_t, \varepsilon_{t+1}) y_t(H_t) \\
&+ \sum_{t=0}^{\infty} \sum_{H_t \in \Omega_t} \sum_{(H_t, \varepsilon_{t+1}) \in \Omega_{t+1} | H_t} \beta^{t+1} p(H_t, \varepsilon_{t+1}) \mu_{t+1}(H_t, \varepsilon_{t+1}) Q[y_t(H_t), z_t(H_t), \varepsilon_{t+1}(H_t, \varepsilon_{t+1})] \\
&- \sum_{t=0}^{\infty} \sum_{H_t \in \Omega_t} \sum_{(H_t, \varepsilon_{t+1}) \in \Omega_{t+1} | H_t} \beta^{t+1} p(H_t, \varepsilon_{t+1}) \mu_{t+1}(H_t, \varepsilon_{t+1}) y_{t+1}(H_t, \varepsilon_{t+1})
\end{aligned}$$

For a given value of $t = 0, 1, 2, \dots$ and a given history $H_t \in \Omega_t$, the first-order condition for $z_t(H_t)$ is

$$\begin{aligned} & \beta^t p(H_t) F_2[y_t(H_t), z_t(H_t), \varepsilon_t(H_t)] \\ & + \sum_{(H_t, \varepsilon_{t+1}) \in \Omega_{t+1} | H_t} \beta^{t+1} p(H_t, \varepsilon_{t+1}) \mu_{t+1}(H_t, \varepsilon_{t+1}) Q_2[y_t(H_t), z_t(H_t), \varepsilon_{t+1}(H_t, \varepsilon_{t+1})] = 0 \end{aligned}$$

For a given value of $t = 1, 2, 3, \dots$ and a given history $H_t \in \Omega_t$, the first-order condition for $y_t(H_t)$ is

$$\begin{aligned} & \beta^t p(H_t) F_1[y_t(H_t), z_t(H_t), \varepsilon_t(H_t)] \\ & + \sum_{(H_t, \varepsilon_{t+1}) \in \Omega_{t+1} | H_t} \beta^{t+1} p(H_t, \varepsilon_{t+1}) \mu_{t+1}(H_t, \varepsilon_{t+1}) \\ & + \sum_{(H_t, \varepsilon_{t+1}) \in \Omega_{t+1} | H_t} \beta^{t+1} p(H_t, \varepsilon_{t+1}) \mu_{t+1}(H_t, \varepsilon_{t+1}) Q_1[y_t(H_t), z_t(H_t), \varepsilon_{t+1}(H_t, \varepsilon_{t+1})] \\ & - \beta^t p(H_t) \mu_t(H_t) = 0 \end{aligned}$$

Divide the first-order condition z_t through by $\beta^t p(H_t)$, and recall that

$$p(\varepsilon_{t+1} | H_t) = \frac{p(H_t, \varepsilon_{t+1})}{p(H_t)}$$

defines the conditional probability that the period $t+1$ shock ε_{t+1} will follow the history H_t :

$$\begin{aligned} & F_2[y_t(H_t), z_t(H_t), \varepsilon_t(H_t)] \\ & + \beta \sum_{(H_t, \varepsilon_{t+1}) \in \Omega_{t+1} | H_t} p(\varepsilon_{t+1} | H_t) \mu_{t+1}(H_t, \varepsilon_{t+1}) Q_2[y_t(H_t), z_t(H_t), \varepsilon_{t+1}(H_t, \varepsilon_{t+1})] = 0 \end{aligned}$$

for all $t = 0, 1, 2, \dots$ and all $H_t \in \Omega_t$.

Or, more compactly,

$$F_2(y_t, z_t, \varepsilon_t) + \beta E_t[\mu_{t+1} Q_2(y_t, z_t, \varepsilon_{t+1})] = 0 \quad (39)$$

Divide the first-order condition for y_t through by $\beta^t p(H_t)$:

$$\begin{aligned} & \mu_t(H_t) \\ = & F_1[y_t(H_t), z_t(H_t), \varepsilon_t(H_t)] \\ & + \beta \sum_{(H_t, \varepsilon_{t+1}) \in \Omega_{t+1} | H_t} p(\varepsilon_{t+1} | H_t) \mu_{t+1}(H_t, \varepsilon_{t+1}) \\ & + \beta \sum_{(H_t, \varepsilon_{t+1}) \in \Omega_{t+1} | H_t} p(\varepsilon_{t+1} | H_t) \mu_{t+1}(H_t, \varepsilon_{t+1}) Q_1[y_t(H_t), z_t(H_t), \varepsilon_{t+1}(H_t, \varepsilon_{t+1})] \end{aligned}$$

for all $t = 1, 2, 3, \dots$ and all $H_t \in \Omega_t$.

Or, again more compactly,

$$\mu_t = F_1(y_t, z_t, \varepsilon_t) + \beta E_t\{\mu_{t+1}[1 + Q_1(y_t, z_t, \varepsilon_{t+1})]\} \quad (40)$$

Together with the binding constraint

$$y_{t+1} = y_t + Q(y_t, z_t, \varepsilon_{t+1}), \quad (41)$$

equations (39) and (40) form a system of three equations in the three unknown variables, z_t , y_t , and μ_t . This system of equations determines the problem's solution, given the behavior of the exogenous shocks ε_t .

Moreover, the three-equation system (39)-(41) that we just derived from the Lagrangian coincides with the three-equation system (23)-(25) that we derived previously using the Bellman equation, but with

$$\mu_t = v_1(y_t, \varepsilon_t)$$

for all $t = 0, 1, 2, \dots$, and all possible values of ε_t .

Once again, the Kuhn-Tucker theorem and the dynamic programming approach lead to exactly the same results. One can choose whichever approach one finds most convenient to characterize the solution to the dynamic, stochastic optimization problem

9 Example 5: Saving Under Uncertainty

This last example presents a dynamic, stochastic optimization problem that is simple enough to allow a relatively straightforward application of the Kuhn-Tucker theorem. The optimality conditions derived with the help of the Lagrangian and the Kuhn-Tucker theorem can then be compared with those that can be derived with the help of the Bellman equation and dynamic programming.

9.1 The Problem

Consider the simplest possible dynamic, stochastic optimization problem with:

Two periods, $t = 0$ and $t = 1$

No uncertainty at $t = 0$

Two possible states at $t = 1$:

Good, or high, state H occurs with probability p

Bad, or low, state L occurs with probability $1 - p$

Notation for a consumer's problem:

y_0 = income at $t = 0$

c_0 = consumption at $t = 0$

s = savings at $t = 0$, carried into $t = 1$ (s can be negative, that is, the consumer is allowed to borrow)

r = interest rate on savings

y_1^H = income at $t = 1$ in the high state

y_1^L = income at $t = 1$ in the low state

$y_1^H > y_1^L$ makes H the good state and L the bad state

c_1^H = consumption at $t = 1$ in the high state

c_1^L = consumption at $t = 1$ in the low state

Expected utility:

$$u(c_0) + \beta E[u(c_1)] = u(c_0) + \beta p u(c_1^H) + \beta(1 - p)u(c_1^L)$$

Constraints:

$$\begin{aligned} y_0 &\geq c_0 + s \\ (1 + r)s + y_1^H &\geq c_1^H \\ (1 + r)s + y_1^L &\geq c_1^L \end{aligned}$$

The problem:

$$\max_{c_0, s, c_1^H, c_1^L} u(c_0) + \beta p u(c_1^H) + \beta(1 - p)u(c_1^L)$$

subject to

$$\begin{aligned} y_0 &\geq c_0 + s \\ (1 + r)s + y_1^H &\geq c_1^H \end{aligned}$$

and

$$(1 + r)s + y_1^L \geq c_1^L$$

Notes:

There are two constraints for period $t = 1$: one for each possible realization of y_1 .

What makes the problem interesting is that savings s at $t = 0$ must be chosen before income y_1 at $t = 1$ is known.

From the viewpoint of $t = 0$, uncertainty about y_1 induces uncertainty about c_1 : the consumer must choose a “contingency plan” for c_1 .

In this simple case, it's not really that much of a problem to deal with all of the constraints in forming a Lagrangian.

In this simple case, it's relatively easy to describe the contingency plan using the notation c_1^H and c_1^L to distinguish between consumption at $t = 1$ in each of the two states.

But, as we've already seen, when the number of periods and/or the number of possible states grow, these notational burdens become increasing tedious, which is what motivates our interest in dynamic programming as a way of dealing with stochastic problems.

9.2 The Kuhn-Tucker Formulation

Set up the Lagrangian, using separate multipliers μ_1^H and μ_1^L for each constraint at $t = 1$:

$$L(c_0, s, c_1^H, c_1^L, \mu_0, \mu_1^H, \mu_1^L) = u(c_0) + \beta p u(c_1^H) + \beta(1-p)u(c_1^L) + \mu_0(y_0 - c_0 - s) \\ + \mu_1^H[(1+r)s + y_1^H - c_1^H] + \mu_1^L[(1+r)s + y_1^L - c_1^L]$$

FOC for c_0 :

$$u'(c_0) - \mu_0 = 0$$

FOC for s :

$$-\mu_0 + \mu_1^H(1+r) + \mu_1^L(1+r) = 0$$

FOC for c_1^H :

$$\beta p u'(c_1^H) - \mu_1^H = 0$$

FOC for c_1^L :

$$\beta(1-p)u'(c_1^L) - \mu_1^L = 0$$

Use the FOC's for c_0 , c_1^H , and c_1^L to eliminate reference to the multipliers μ_0 , μ_1^H , and μ_1^L in the FOC for s :

$$u'(c_0) = \beta p u'(c_1^H)(1+r) + \beta(1-p)u'(c_1^L)(1+r) \quad (42)$$

Together with the binding constraints

$$y_0 = c_0 + s$$

$$(1+r)s + y_1^H = c_1^H$$

and

$$(1+r)s + y_1^L = c_1^L$$

(42) gives us a system of 4 equations in the 4 unknowns: c_0 , s , c_1^H , and c_1^L .

Note also that (42) can be written more compactly as

$$u'(c_0) = \beta(1+r)E[u'(c_1)],$$

which is a special case of the more general optimality condition that we derived previously in the “saving with multiple random returns” example, reflecting that in this simple example:

There are only two periods.

The return on the single asset is known.

9.3 The Dynamic Programming Formulation

Consider “restarting” the problem at $t = 1$, in state H , given that s has already been chosen and y_1^H already determined.

The consumer solves the static problem:

$$\max_{c_1^H} u(c_1^H)$$

subject to

$$(1 + r)s + y_1^H \geq c_1^H.$$

The solution is trivial: set

$$c_1^H = (1 + r)s + y_1^H$$

Hence, if we define the maximum value function

$$v(s, y_1^H) = \max_{c_1^H} u(c_1^H) \text{ subject to } (1 + r)s + y_1^H \geq c_1^H$$

then we know right away that

$$v(s, y_1^H) = u[(1 + r)s + y_1^H]$$

and hence

$$v_1(s, y_1^H) = (1 + r)u'[(1 + r)s + y_1^H] = (1 + r)u'(c_1^H) \quad (43)$$

Likewise, if we restart the problem at $t = 1$ in state L , given that s has already been chosen and y_1^L already determined, then

$$v(s, y_1^L) = \max_{c_1^L} u(c_1^L) \text{ subject to } (1 + r)s + y_1^L \geq c_1^L$$

and we know right away that

$$v(s, y_1^L) = u[(1 + r)s + y_1^L]$$

and

$$v_1(s, y_1^L) = (1 + r)u'[(1 + r)s + y_1^L] = (1 + r)u'(c_1^L) \quad (44)$$

Now back up to $t = 0$, and consider the problem

$$\max_{c_0, s} u(c_0) + \beta E v(s, y_1) \text{ subject to } y_0 \geq c_0 + s$$

or, even more simply

$$\max_s u(y_0 - s) + \beta E v(s, y_1) \quad (45)$$

Equation (45) is like the Bellman equation for the consumer’s problem:

The problem described on the right-hand-side is a static problem: the dynamic programming approach breaks the dynamic program down into a sequence of static problems.

Note, too, that the problem is an unconstrained optimization problem.

And note that in (45), the “maximize with respect to c_1^H and c_1^L ” part of the original dynamic problem has been moved inside the expectation term, sidestepping the need to talk explicitly about “contingency plans” for the future.

Take the FOC for the value of s that solves the problem in (45):

$$-u'(y_0 - s) + \beta E v_1(s, y_1) = 0$$

and rewrite it using (43) and (44) as

$$u'(c_0) = \beta E[(1+r)u'(c_1)] = \beta(1+r)pu'(c_1^H) + \beta(1+r)(1-p)u'(c_1^L) \quad (46)$$

Notes:

Together with the binding constraints

$$y_0 = c_0 + s$$

$$(1+r)s + y_1^H = c_1^H$$

and

$$(1+r)s + y_1^L = c_1^L$$

(46) gives us a system of 4 equations in the 4 unknowns: c_0 , s , c_1^H , and c_1^L .

This system of equations is exactly the same one that we derived earlier with the help of the Lagrangian and the Kuhn-Tucker theorem.