

## Chapter 3.

### Dynamic Programming

This chapter introduces basic ideas and methods of dynamic programming.<sup>1</sup> It sets out the basic elements of a recursive optimization problem, describes the functional equation (the Bellman equation), presents three methods for solving the Bellman equation, and gives the Benveniste-Scheinkman formula for the derivative of the optimal value function. Let's dive in.

#### 3.1. Sequential problems

Let  $\beta \in (0, 1)$  be a discount factor. We want to choose an infinite sequence of "controls"  $\{u_t\}_{t=0}^{\infty}$  to maximize

$$\sum_{t=0}^{\infty} \beta^t r(x_t, u_t), \quad (3.1.1)$$

subject to  $x_{t+1} = g(x_t, u_t)$ , with  $x_0$  given. We assume that  $r(x_t, u_t)$  is a concave function and that the set  $\{(x_{t+1}, x_t) : x_{t+1} \leq g(x_t, u_t), u_t \in R^k\}$  is convex and compact. Dynamic programming seeks a time-invariant *policy function*  $h$  mapping the *state*  $x_t$  into the control  $u_t$ , such that the sequence  $\{u_s\}_{s=0}^{\infty}$  generated by iterating the two functions

$$\begin{aligned} u_t &= h(x_t) \\ x_{t+1} &= g(x_t, u_t), \end{aligned} \quad (3.1.2)$$

starting from initial condition  $x_0$  at  $t = 0$  solves the original problem. A solution in the form of equations (3.1.2) is said to be *recursive*. To find the policy function  $h$  we need to know another function  $V(x)$  that expresses the optimal value of the original problem, starting from an arbitrary initial condition  $x \in X$ . This is called the *value*

---

<sup>1</sup> This chapter is written in the hope of getting the reader to start using the methods quickly. We hope to promote demand for further and more rigorous study of the subject. In particular see Bertsekas (1976), Bertsekas and Shreve (1978), Stokey and Lucas (with Prescott) (1989), Bellman (1957), and Chow (1981). This chapter covers much of the same material as Sargent (1987b, chapter 1).

function. In particular, define

$$V(x_0) = \max_{\{u_s\}_{s=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t r(x_t, u_t), \quad (3.1.3)$$

where again the maximization is subject to  $x_{t+1} = g(x_t, u_t)$ , with  $x_0$  given. Of course, we cannot possibly expect to know  $V(x_0)$  until after we have solved the problem, but let's proceed on faith. If we knew  $V(x_0)$ , then the policy function  $h$  could be computed by solving for each  $x \in X$  the problem

$$\max_u \{r(x, u) + \beta V(\tilde{x})\}, \quad (3.1.4)$$

where the maximization is subject to  $\tilde{x} = g(x, u)$  with  $x$  given, and  $\tilde{x}$  denotes the state next period. Thus, we have exchanged the original problem of finding an infinite *sequence* of controls that maximizes expression (3.1.1) for the problem of finding the optimal value function  $V(x)$  and a function  $h$  that solves the continuum of maximum problems (3.1.4)—one maximum problem for each value of  $x$ . This exchange doesn't look like progress, but we shall see that it often is.

Our task has become jointly to solve for  $V(x), h(x)$ , which are linked by the *Bellman equation*

$$V(x) = \max_u \{r(x, u) + \beta V[g(x, u)]\}. \quad (3.1.5)$$

The maximizer of the right side of equation (3.1.5) is a *policy function*  $h(x)$  that satisfies

$$V(x) = r[x, h(x)] + \beta V\{g[x, h(x)]\}. \quad (3.1.6)$$

Equation (3.1.5) or (3.1.6) is a *functional equation* to be solved for the pair of unknown functions  $V(x), h(x)$ .

Methods for solving the Bellman equation are based on mathematical structures that vary in their details depending on the precise nature of the functions  $r$  and  $g$ .<sup>2</sup>

---

<sup>2</sup> There are alternative sets of conditions that make the maximization (3.1.4) well behaved. One set of conditions is as follows: (1)  $r$  is concave and bounded, and (2) the constraint set generated by  $g$  is convex and compact, that is, the set of  $\{(x_{t+1}, x_t) : x_{t+1} \leq g(x_t, u_t)\}$  for admissible  $u_t$  is convex and compact. See Stokey, Lucas, and Prescott (1989), and Bertsekas (1976) for further details of convergence results. See Benveniste and Scheinkman (1979) and Stokey, Lucas, and Prescott (1989) for the results on differentiability of the value function. In an appendix on functional analysis, chapter A, we describe the mathematics for one standard set of assumptions about  $(r, g)$ . In chapter 5, we describe it for another set of assumptions about  $(r, g)$ .

All of these structures contain versions of the following four findings. Under various particular assumptions about  $r$  and  $g$ , it turns out that

1. The functional equation (3.1.5) has a unique strictly concave solution.
2. This solution is approached in the limit as  $j \rightarrow \infty$  by iterations on

$$V_{j+1}(x) = \max_u \{r(x, u) + \beta V_j(\tilde{x})\}, \quad (3.1.7)$$

subject to  $\tilde{x} = g(x, u)$ ,  $x$  given, starting from any bounded and continuous initial  $V_0$ .

3. There is a unique and time invariant optimal policy of the form  $u_t = h(x_t)$ , where  $h$  is chosen to maximize the right side of (3.1.5).<sup>3</sup>
4. Off corners, the limiting value function  $V$  is differentiable with

$$V'(x) = \frac{\partial r}{\partial x}[x, h(x)] + \beta \frac{\partial g}{\partial x}[x, h(x)] V'\{g[x, h(x)]\}. \quad (3.1.8)$$

This is a version of a formula of Benveniste and Scheinkman (1979). We often encounter settings in which the transition law can be formulated so that the state  $x$  does not appear in it, so that  $\frac{\partial g}{\partial x} = 0$ , which makes equation (3.1.8) become

$$V'(x) = \frac{\partial r}{\partial x}[x, h(x)]. \quad (3.1.9)$$

At this point, we describe three broad computational strategies that apply in various contexts.

---

<sup>3</sup> The time invariance of the policy function  $u_t = h(x_t)$  is very convenient econometrically, because we can impose a single decision rule for all periods. This lets us pool data across period to estimate the free parameters of the return and transition functions that underlie the decision rule.

### 3.1.1. Three computational methods

There are three main types of computational methods for solving dynamic programs. All aim to solve the functional equation (3.1.4).

**Value function iteration.** The first method proceeds by constructing a sequence of value functions and associated policy functions. The sequence is created by iterating on the following equation, starting from  $V_0 = 0$ , and continuing until  $V_j$  has converged:<sup>4</sup>

$$V_{j+1}(x) = \max_u \{r(x, u) + \beta V_j(\tilde{x})\}, \quad (3.1.10)$$

subject to  $\tilde{x} = g(x, u)$ ,  $x$  given.<sup>5</sup> This method is called *value function iteration* or *iterating on the Bellman equation*.

**Guess and verify.** A second method involves guessing and verifying a solution  $V$  to equation (3.1.5). This method relies on the uniqueness of the solution to the equation, but because it relies on luck in making a good guess, it is not generally available.

**Howard's improvement algorithm.** A third method, known as *policy function iteration* or *Howard's improvement algorithm*, consists of the following steps:

1. Pick a feasible policy,  $u = h_0(x)$ , and compute the value associated with operating forever with that policy:

$$V_{h_j}(x) = \sum_{t=0}^{\infty} \beta^t r[x_t, h_j(x_t)],$$

where  $x_{t+1} = g[x_t, h_j(x_t)]$ , with  $j = 0$ .

2. Generate a new policy  $u = h_{j+1}(x)$  that solves the two-period problem

$$\max_u \{r(x, u) + \beta V_{h_j}[g(x, u)]\},$$

for each  $x$ .

3. Iterate over  $j$  to convergence on steps 1 and 2.

---

<sup>4</sup> See the appendix on functional analysis for what it means for a sequence of functions to converge.

<sup>5</sup> A proof of the uniform convergence of iterations on equation (3.1.10) is contained in the appendix on functional analysis, chapter A.

In the appendix on functional analysis, chapter A, we describe some conditions under which the improvement algorithm converges to the solution of Bellman's equation. The method often converges faster than does value function iteration (e.g., see exercise 2.1 at the end of this chapter).<sup>6</sup> The policy improvement algorithm is also a building block for the methods for studying government policy to be described in chapter 22.

Each of these methods has its uses. Each is "easier said than done," because it is typically impossible analytically to compute even *one* iteration on equation (3.1.10). This fact thrusts us into the domain of computational methods for approximating solutions: pencil and paper are insufficient. The following chapter describes some computational methods that can be used for problems that cannot be solved by hand. Here we shall describe the first of two special types of problems for which analytical solutions *can* be obtained. It involves Cobb-Douglas constraints and logarithmic preferences. Later in chapter 5, we shall describe a specification with linear constraints and quadratic preferences. For that special case, many analytic results are available. These two classes have been important in economics as sources of examples and as inspirations for approximations.

### 3.1.2. Cobb-Douglas transition, logarithmic preferences

Brock and Mirman (1972) used the following optimal growth example.<sup>7</sup> A planner chooses sequences  $\{c_t, k_{t+1}\}_{t=0}^{\infty}$  to maximize

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

subject to a given value for  $k_0$  and a transition law

$$k_{t+1} + c_t = Ak_t^\alpha, \quad (3.1.11)$$

where  $A > 0, \alpha \in (0, 1), \beta \in (0, 1)$ .

---

<sup>6</sup> The quickness of the policy improvement algorithm is linked to its being an implementation of Newton's method, which converges quadratically while iteration on the Bellman equation converges at a linear rate. See chapter 4 and the appendix on functional analysis, chapter A.

<sup>7</sup> See also Levhari and Srinivasan (1969).

This problem can be solved “by hand,” using any of our three methods. We begin with iteration on the Bellman equation. Start with  $v_0(k) = 0$ , and solve the one-period problem: choose  $c$  to maximize  $\ln(c)$  subject to  $c + \tilde{k} = Ak^\alpha$ . The solution is evidently to set  $c = Ak^\alpha, \tilde{k} = 0$ , which produces an optimized value  $v_1(k) = \ln A + \alpha \ln k$ . At the second step, we find  $c = \frac{1}{1+\beta\alpha} Ak^\alpha, \tilde{k} = \frac{\beta\alpha}{1+\beta\alpha} Ak^\alpha, v_2(k) = \ln \frac{A}{1+\alpha\beta} + \beta \ln A + \alpha\beta \ln \frac{\alpha\beta A}{1+\alpha\beta} + \alpha(1 + \alpha\beta) \ln k$ . Continuing, and using the algebra of geometric series, gives the limiting policy functions  $c = (1 - \beta\alpha)Ak^\alpha, \tilde{k} = \beta\alpha Ak^\alpha$ , and the value function  $v(k) = (1 - \beta)^{-1} \{ \ln[A(1 - \beta\alpha)] + \frac{\beta\alpha}{1 - \beta\alpha} \ln(A\beta\alpha) \} + \frac{\alpha}{1 - \beta\alpha} \ln k$ .

Here is how the guess-and-verify method applies to this problem. Since we already know the answer, we’ll guess a function of the correct form, but leave its coefficients undetermined.<sup>8</sup> Thus, we make the guess

$$v(k) = E + F \ln k, \quad (3.1.12)$$

where  $E$  and  $F$  are undetermined constants. The left and right sides of equation (3.1.12) must agree for all values of  $k$ . For this guess, the first-order necessary condition for the maximum problem on the right side of equation (3.1.10) implies the following formula for the optimal policy  $\tilde{k} = h(k)$ , where  $\tilde{k}$  is next period’s value and  $k$  is this period’s value of the capital stock:

$$\tilde{k} = \frac{\beta F}{1 + \beta F} Ak^\alpha. \quad (3.1.13)$$

Substitute equation (3.1.13) into the Bellman equation and equate the result to the right side of equation (3.1.12). Solving the resulting equation for  $E$  and  $F$  gives  $F = \alpha/(1 - \alpha\beta)$  and  $E = (1 - \beta)^{-1} [\ln A(1 - \alpha\beta) + \frac{\beta\alpha}{1 - \alpha\beta} \ln A\beta\alpha]$ . It follows that

$$\tilde{k} = \beta\alpha Ak^\alpha. \quad (3.1.14)$$

Note that the term  $F = \alpha/(1 - \alpha\beta)$  can be interpreted as a geometric sum  $\alpha[1 + \alpha\beta + (\alpha\beta)^2 + \dots]$ .

Equation (3.1.14) shows that the optimal policy is to have capital move according to the difference equation  $k_{t+1} = A\beta\alpha k_t^\alpha$ , or  $\ln k_{t+1} = \ln A\beta\alpha + \alpha \ln k_t$ . That  $\alpha$  is less than 1 implies that  $k_t$  converges as  $t$  approaches infinity for any positive initial value  $k_0$ . The stationary point is given by the solution of  $k_\infty = A\beta\alpha k_\infty^\alpha$ , or  $k_\infty^{\alpha-1} = (A\beta\alpha)^{-1}$ .

---

<sup>8</sup> This is called the method of undetermined coefficients.

### 3.1.3. Euler equations

In many problems, there is no unique way of defining states and controls, and several alternative definitions lead to the same solution of the problem. Sometimes the states and controls can be defined in such a way that  $x_t$  does not appear in the transition equation, so that  $\partial g_t / \partial x_t \equiv 0$ . In this case, the first-order condition for the problem on the right side of the Bellman equation in conjunction with the Benveniste-Scheinkman formula implies

$$\frac{\partial r_t}{\partial u_t}(x_t, u_t) + \frac{\partial g_t}{\partial u_t}(u_t) \cdot \frac{\partial r_{t+1}(x_{t+1}, u_{t+1})}{\partial x_{t+1}} = 0, \quad x_{t+1} = g_t(u_t).$$

The first equation is called an *Euler equation*. Under circumstances in which the second equation can be inverted to yield  $u_t$  as a function of  $x_{t+1}$ , using the second equation to eliminate  $u_t$  from the first equation produces a second-order difference equation in  $x_t$ , since eliminating  $u_{t+1}$  brings in  $x_{t+2}$ .

### 3.1.4. A sample Euler equation

As an example of an Euler equation, consider the Ramsey problem of choosing  $\{c_t, k_{t+1}\}_{t=0}^{\infty}$  to maximize  $\sum_{t=0}^{\infty} \beta^t u(c_t)$  subject to  $c_t + k_{t+1} = f(k_t)$ , where  $k_0$  is given and the one-period utility function satisfies  $u'(c) > 0$ ,  $u''(c) < 0$ ,  $\lim_{c_t \searrow 0} u'(c_t) = \infty$ ; and where  $f'(k) > 0$ ,  $f''(k) < 0$ . Let the state be  $k$  and the control be  $k'$ , where  $k'$  denotes next period's value of  $k$ . Substitute  $c = f(k) - k'$  into the utility function and express the Bellman equation as

$$v(k) = \max_{\tilde{k}} \{u[f(k) - \tilde{k}] + \beta v(\tilde{k})\}. \quad (3.1.15)$$

Application of the Benveniste-Scheinkman formula gives

$$v'(k) = u'[f(k) - \tilde{k}] f'(k). \quad (3.1.16)$$

Notice that the first-order condition for the maximum problem on the right side of equation (3.1.15) is  $-u'[f(k) - \tilde{k}] + \beta v'(\tilde{k}) = 0$ , which, using equation (3.1.16), gives

$$u'[f(k) - \tilde{k}] = \beta u'[f(\tilde{k}) - \hat{k}] f'(k'), \quad (3.1.17)$$

where  $\hat{k}$  denotes the "two-period-ahead" value of  $k$ . Equation (3.1.17) can be expressed as

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

an Euler equation that is exploited extensively in the theories of finance, growth, and real business cycles.

### 3.2. Stochastic control problems

We now consider a modification of problem (3.1.1) to permit uncertainty. Essentially, we add some well-placed shocks to the previous non-stochastic problem. So long as the shocks are either independently and identically distributed or Markov, straightforward modifications of the method for handling the nonstochastic problem will work.

Thus, we modify the transition equation and consider the problem of maximizing

$$E_0 \sum_{t=0}^{\infty} \beta^t r(x_t, u_t), \quad 0 < \beta < 1, \quad (3.2.1)$$

subject to

$$x_{t+1} = g(x_t, u_t, \epsilon_{t+1}), \quad (3.2.2)$$

with  $x_0$  known and given at  $t = 0$ , where  $\epsilon_t$  is a sequence of independently and identically distributed random variables with cumulative probability distribution function  $\text{prob}\{\epsilon_t \leq e\} = F(e)$  for all  $t$ ;  $E_t(y)$  denotes the mathematical expectation of a random variable  $y$ , given information known at  $t$ . At time  $t$ ,  $x_t$  is assumed to be known, but  $x_{t+j}$ ,  $j \geq 1$  is not known at  $t$ . That is,  $\epsilon_{t+1}$  is realized at  $(t + 1)$ , after  $u_t$  has been chosen at  $t$ . In problem (3.2.1)–(3.2.2), uncertainty is injected by assuming that  $x_t$  follows a random difference equation.

Problem (3.2.1)–(3.2.2) continues to have a recursive structure, stemming jointly from the additive separability of the objective function (3.2.1) in pairs  $(x_t, u_t)$  and from the difference equation characterization of the transition law (3.2.2). In particular, controls dated  $t$  affect returns  $r(x_s, u_s)$  for  $s \geq t$  but not earlier. This feature implies that dynamic programming methods remain appropriate.

The problem is to maximize expression (3.2.1) subject to equation (3.2.2) by choice of a “policy” or “contingency plan”  $u_t = h(x_t)$ . The Bellman equation (3.1.5) becomes

$$V(x) = \max_u \{r(x, u) + \beta E[V[g(x, u, \epsilon)]|x]\}, \quad (3.2.3)$$

where  $E\{V[g(x, u, \epsilon)]|x\} = \int V[g(x, u, \epsilon)]dF(\epsilon)$  and where  $V(x)$  is the optimal value of the problem starting from  $x$  at  $t = 0$ . The solution  $V(x)$  of equation (3.2.3) can



be computed by iterating on

$$V_{j+1}(x) = \max_u \{r(x, u) + \beta E [V_j [g(x, u, \epsilon)] | x]\}, \quad (3.2.4)$$

starting from any bounded continuous initial  $V_0$ . Under various particular regularity conditions, there obtain versions of the same four properties listed earlier.<sup>9</sup>

The first-order necessary condition for the problem on the right side of equation (3.2.3) is

$$\frac{\partial r(x, u)}{\partial u} + \beta E \left[ \frac{\partial g}{\partial u}(x, u, \epsilon) V' [g(x, u, \epsilon)] | x \right] = 0,$$

which we obtained simply by differentiating the right side of equation (3.2.3), passing the differentiation operation under the  $E$  (an integration) operator. Off corners, the value function satisfies

$$V'(x) = \frac{\partial r}{\partial x} [x, h(x)] + \beta E \left\{ \frac{\partial g}{\partial x} [x, h(x), \epsilon] V'(g[x, h(x), \epsilon]) | x \right\}.$$

In the special case in which  $\partial g / \partial x \equiv 0$ , the formula for  $V'(x)$  becomes

$$V'(x) = \frac{\partial r}{\partial x} [x, h(x)].$$

Substituting this formula into the first-order necessary condition for the problem gives the stochastic Euler equation

$$\frac{\partial r}{\partial u}(x, u) + \beta E \left[ \frac{\partial g}{\partial u}(x, u, \epsilon) \frac{\partial r}{\partial x}(\tilde{x}, \tilde{u}) | x \right] = 0,$$

where tildes over  $x$  and  $u$  denote next-period values.

---

<sup>9</sup> See Stokey and Lucas (with Prescott) (1989), or the framework presented in the appendix on functional analysis, chapter A.

### 3.3. Concluding remarks

This chapter has put forward basic tools and findings: the Bellman equation and several approaches to solving it; the Euler equation; and the Beneveniste-Scheinkman formula. To appreciate and believe in the power of these tools requires more words and more practice than we have yet supplied. In the next several chapters, we put the basic tools to work in different contexts with particular specification of return and transition equations designed to render the Bellman equation susceptible to further analysis and computation.

## Exercise

### Exercise 3.1 Howard's policy iteration algorithm

Consider the Brock-Mirman problem: to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t \ln c_t,$$

subject to  $c_t + k_{t+1} \leq Ak_t^\alpha \theta_t$ ,  $k_0$  given,  $A > 0$ ,  $1 > \alpha > 0$ , where  $\{\theta_t\}$  is an i.i.d. sequence with  $\ln \theta_t$  distributed according to a normal distribution with mean zero and variance  $\sigma^2$ .

Consider the following algorithm. Guess at a policy of the form  $k_{t+1} = h_0(Ak_t^\alpha \theta_t)$  for any constant  $h_0 \in (0, 1)$ . Then form

$$J_0(k_0, \theta_0) = E_0 \sum_{t=0}^{\infty} \beta^t \ln (Ak_t^\alpha \theta_t - h_0 Ak_t^\alpha \theta_t).$$

Next choose a new policy  $h_1$  by maximizing

$$\ln (Ak^\alpha \theta - k') + \beta E J_0(k', \theta'),$$

where  $k' = h_1 Ak^\alpha \theta$ . Then form

$$J_1(k_0, \theta_0) = E_0 \sum_{t=0}^{\infty} \beta^t \ln (Ak_t^\alpha \theta_t - h_1 Ak_t^\alpha \theta_t).$$

Continue iterating on this scheme until successive  $h_j$  have converged.

Show that, for the present example, this algorithm converges to the optimal policy function in one step.