4 Dynamic Programming under Certainty

Posed in terms of infinite sequences, the problems we are interested in are of the form

$$\text{(SP)} \qquad \sup_{\{x_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t F(x_t, x_{t+1})$$

$$\text{s.t.} \quad x_{t+1} \in \Gamma(x_t), \quad t = 0, 1, 2, \ldots,$$

$$x_0 \in X \text{ given.}$$

Corresponding to any such problem, we have a functional equation of the form

$$\text{(FE)} \qquad v(x) = \sup_{y \in \Gamma(x)} [F(x, y) + \beta v(y)], \quad \text{all } x \in X.$$

In this chapter we establish the relationship between solutions to these two problems and develop methods for analyzing the latter.

Exercise 4.1 a. Show that the one-sector growth model discussed at the beginning of Chapter 3 can be expressed as in (SP).

b. Show that the many-sector growth model

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

$$\text{s.t.} \quad (k_{t+1} + c_t, k_t) \in Y, \quad t = 0, 1, 2, \ldots,$$

$$\text{given } k_0 \in \mathbf{R}^l_+,$$

where $Y \subseteq \mathbf{R}^{2l}_+$ is a fixed production set, can also be written this way.

As we hinted in the last chapter and will show in this one, some very powerful—and relatively simple—mathematical tools can be used to study the functional equation (FE). To take advantage of these, however, we must show that solutions to (FE) correspond to solutions to the sequence problem (SP). In Section 4.1 we rigorously establish the connections between solutions to these two problems, connections that Richard Bellman called the "Principle of Optimality." Section 4.2 then develops the main results of the chapter: existence, uniqueness, and characterization theorems for solutions to (FE) under the assumption that the return function F is bounded. The case where F displays constant returns to scale is treated in Section 4.3, and the case where F is an arbitrary unbounded return function in Section 4.4. Section 4.5 treats the relationship between the dynamic programming approach to optimization over time and the classical (variational) approach. Section 4.6 contains references for further discussion of some of the mathematical and economic ideas. In Chapter 5 we illustrate how the methods developed in Sections 4.2–4.4 can be applied to a wide variety of economic problems.

4.1 The Principle of Optimality

In this section we study the relationship between solutions to the problems (SP) and (FE). (Note that "sup" has been used instead of "max" in both, so that we can ignore—for the moment—the question of whether the optimum is attained.) The general idea, of course, is that the solution v to (FE), evaluated at x_0, gives the value of the supremum in (SP) when the initial state is x_0 and that a sequence $\{x_{t+1}\}_{t=0}^{\infty}$ attains the supremum in (SP) if and only if it satisfies

$$\text{(1)} \qquad v(x_t) = F(x_t, x_{t+1}) + \beta v(x_{t+1}), \quad t = 0, 1, 2, \ldots.$$

Richard Bellman called these ideas the Principle of Optimality. Intuitive as it is, the Principle requires proof. Spelling out precisely the conditions under which it holds is our task in this section.

The main results are Theorem 4.2, establishing that the supremum function v^* for the sequence problem (SP) satisfies the functional equation (FE), and Theorem 4.3, establishing a partial converse. The "partial" nature of the converse arises from the fact that a boundedness condition must be imposed. Theorems 4.4 and 4.5 then deal with the characterization of optimal policies. Theorem 4.4 shows that if $\{x_{t+1}\}_{t=0}^{\infty}$ is a sequence attaining the supremum in (SP), then it satisfies (1) for $v = v^*$. Conversely, Theorem 4.5 establishes that any sequence $\{x_{t+1}\}_{t=0}^{\infty}$ that satisfies (1) for $v = v^*$, and also satisfies a boundedness condition, attains the supremum in (SP). The four theorems taken together thus establish conditions under which solutions to (SP) and to (FE) coincide exactly, and optimal policies are those that satisfy (1).

To begin we must establish some notation. Let X be the set of possible values for the state variable x. In this section we will not need to impose any restrictions on the set X. It may be a subset of a Euclidean space, a set of functions, a set of probability distributions, or any other set. Let $\Gamma: X \rightarrow X$ be the correspondence describing the feasibility constraints. That is, for each $x \in X$, $\Gamma(x)$ is the set of feasible values for the state variable next period if the current state is x. Let A be the graph of Γ:

$$A = \{(x, y) \in X \times X : y \in \Gamma(x)\}.$$

Let the real-valued function $F: A \rightarrow \mathbf{R}$ be the one-period return function, and let $\beta \geq 0$ be the (stationary) discount factor. Thus the "givens" for the problem are X, Γ, F, and β.

First we must establish conditions under which the problem (SP) is well defined. That is, we must find conditions under which the feasible set is nonempty and the objective function is well defined for every point in the feasible set.

Call any sequence $\{x_t\}_{t=0}^{\infty}$ in X a *plan*. Given $x_0 \in X$, let

$$\Pi(x_0) = \{\{x_t\}_{t=0}^{\infty}: x_{t+1} \in \Gamma(x_t), \quad t = 0, 1, \ldots\}$$

be the set of plans that are *feasible from* x_0. That is, $\Pi(x_0)$ is the set of all sequences $\{x_t\}$ satisfying the constraints in (SP). Let $\underline{x} = (x_0, x_1, \ldots)$ denote a typical element of $\Pi(x_0)$. The following assumption ensures that $\Pi(x_0)$ is nonempty, for all $x_0 \in X$.

ASSUMPTION 4.1 *$\Gamma(x)$ is nonempty, for all $x \in X$.*

The only additional restriction on X, Γ, F, and β we will need in this section is a requirement that all feasible plans can be evaluated using the objective function F and the discount rate β.

ASSUMPTION 4.2 *For all $x_0 \in X$ and $\underline{x} \in \Pi(x_0)$, $\lim_{n\to\infty}\sum_{t=0}^{n}\beta^t F(x_t, x_{t+1})$ exists (although it may be plus or minus infinity).*

There are a variety of ways of ensuring that Assumption 4.2 holds. Clearly it is satisfied if the function F is bounded and $0 < \beta < 1$. Alternatively, for any $(x, y) \in A$, let

$$F^+(x, y) = \max\{0, F(x, y)\} \quad \text{and} \quad F^-(x, y) = \max\{0, -F(x, y)\}.$$

Then Assumption 4.2 holds if for each $x_0 \in X$ and $\underline{x} \in \Pi(x_0)$, either

$$\lim_{n\to\infty} \sum_{t=0}^{n} \beta^t F^+(x_t, x_{t+1}) < +\infty, \quad \text{or}$$

$$\lim_{n\to\infty} \sum_{t=0}^{n} \beta^t F^-(x_t, x_{t+1}) < +\infty, \quad \text{or both.}$$

Thus a sufficient condition for Assumptions 4.1–4.2 is that F be bounded above or below and $0 < \beta < 1$. Another sufficient condition is that for each $x_0 \in X$ and $\underline{x} \in \Pi(x_0)$, there exist $\theta \in (0, \beta^{-1})$ and $0 < c < \infty$ such that

$$F(x_t, x_{t+1}) \leq c\theta^t, \quad \text{all } t.$$

The following exercise provides a way of verifying that the latter holds.

Exercise 4.2 a. Show that Assumption 4.2 is satisfied if $X = \mathbf{R}^l_+$; $0 < \beta < 1$; there exists $0 < \theta < 1/\beta$ such that $y \in \Gamma(x)$ implies $\|y\| \leq \theta\|x\|$; $F(0, 0) = 0$; F is increasing in its first l arguments and decreasing in its last l arguments; F is concave in its first l arguments; and $0 \in \Gamma(x)$, all x.

b. Show that Assumption 4.2 is satisfied if $X = \mathbf{R}^l_+$; $0 < \beta < 1$; there exists $0 < \theta < 1/\beta$ such that $y \in \Gamma(x)$ implies $F(y, 0) \leq \theta F(x, 0)$; F is increasing in its first l arguments and decreasing in its last l arguments; and $0 \in \Gamma(x)$, all x.

For each $n = 0, 1, \ldots,$ define $u_n: \Pi(x_0) \rightarrow \mathbf{R}$ by

$$u_n(\underline{x}) = \sum_{t=0}^{n} \beta^t F(x_t, x_{t+1}).$$

Then $u_n(\underline{x})$ is the partial sum of the (discounted) returns in periods 0 through n from the feasible plan $\underline{x}$. Under Assumption 4.2 we can also define $u: \Pi(x_0) \to \overline{\mathbf{R}}$ by

$$u(\underline{x}) = \lim_{n\to\infty} u_n(\underline{x}),$$

where $\overline{\mathbf{R}} = \mathbf{R} \cup \{+\infty, -\infty\}$ is the set of extended real numbers. Thus $u(\underline{x})$ is the (infinite) sum of discounted returns from the feasible sequence $\underline{x}$.

If Assumptions 4.1 and 4.2 both hold, then the set of feasible plans $\Pi(x_0)$ is nonempty for each $x_0 \in X$, and the objective function in (SP) is well defined for every plan $\underline{x} \in \Pi(x_0)$. We can then define the *supremum function* $v^*: X \to \overline{\mathbf{R}}$ by

$$v^*(x_0) = \sup_{\underline{x}\in\Pi(x_0)} u(\underline{x}).$$

Thus $v^*(x_0)$ is the supremum in (SP). Note that it follows by definition that v^* is the unique function satisfying the following three conditions:

a. if $|v^*(x_0)| < \infty$, then

$$v^*(x_0) \geq u(\underline{x}), \quad \text{all } \underline{x} \in \Pi(x_0); \tag{2}$$

and for any $\varepsilon > 0$,

$$v^*(x_0) \leq u(\underline{x}) + \varepsilon, \quad \text{some } \underline{x} \in \Pi(x_0); \tag{3}$$

b. if $v^*(x_0) = +\infty$, then there exists a sequence $\{\underline{x}^k\}$ in $\Pi(x_0)$ such that $\lim_{k\to\infty} u(\underline{x}^k) = +\infty$; and

c. if $v^*(x_0) = -\infty$, then $u(\underline{x}) = -\infty$, for all $\underline{x} \in \Pi(x_0)$.

Our interest is in the connections between the supremum function v^* and solutions v to the functional equation (FE). In interpreting the next results, it is important to remember that v^* is always uniquely defined (provided Assumptions 4.1–4.2 hold), whereas (FE) may—for all we know so far—have zero, one, or many solutions.

We will say that v^* satisfies the functional equation if three conditions hold:

a. If $|v^*(x_0)| < \infty$, then

$$v^*(x_0) \geq F(x_0, y) + \beta v^*(y), \quad \text{all } y \in \Gamma(x_0), \tag{4}$$

and for any $\varepsilon > 0$,

$$v^*(x_0) \leq F(x_0, y) + \beta v^*(y) + \varepsilon, \quad \text{some } y \in \Gamma(x_0); \tag{5}$$

b. if $v^*(x_0) = +\infty$, then there exists a sequence $\{y^k\}$ in $\Gamma(x_0)$ such that

$$\lim_{k\to\infty} [F(x_0, y^k) + \beta v^*(y^k)] = +\infty; \tag{6}$$

c. if $v^*(x_0) = -\infty$, then

$$F(x_0, y) + \beta v^*(y) = -\infty, \quad \text{all } y \in \Gamma(x_0). \tag{7}$$

Before we prove that the supremum function v^* satisfies the functional equation, it is useful to establish a preliminary result.

LEMMA 4.1 *Let X, Γ, F, and β satisfy Assumption 4.2. Then for any $x_0 \in X$ and any $(x_0, x_1, \ldots) = \underline{x} \in \Pi(x_0)$,*

$$u(\underline{x}) = F(x_0, x_1) + \beta u(\underline{x}'),$$

where $\underline{x}' = (x_1, x_2, \ldots)$.

Proof. Under Assumption 4.2, for any $x_0 \in X$ and any $\underline{x} \in \Pi(x_0)$,

$$\begin{aligned} u(\underline{x}) &= \lim_{n\to\infty} \sum_{t=0}^{n} \beta^t F(x_t, x_{t+1}) \\ &= F(x_0, x_1) + \beta \lim_{n\to\infty} \sum_{t=0}^{n} \beta^t F(x_{t+1}, x_{t+2}) \\ &= F(x_0, x_1) + \beta u(\underline{x}'). \quad \blacksquare \end{aligned}$$

THEOREM 4.2 *Let X, Γ, F, and β satisfy Assumptions 4.1–4.2. Then the function v^* satisfies (FE).*

Proof. If $\beta = 0$, the result is trivial. Suppose that $\beta > 0$, and choose $x_0 \in X$.

Suppose $v^*(x_0)$ is finite. Then (2) and (3) hold, and it is sufficient to show that this implies (4) and (5) hold. To establish (4), let $x_1 \in \Gamma(x_0)$ and $\varepsilon > 0$ be given. Then by (3) there exists $\underline{x}' = (x_1, x_2, \ldots) \in \Pi(x_1)$ such that $u(\underline{x}') \geq v^*(x_1) - \varepsilon$. Note, too, that $\underline{x} = (x_0, x_1, x_2, \ldots) \in \Pi(x_0)$. Hence it follows from (2) and Lemma 4.1 that

$$v^*(x_0) \geq u(\underset{\sim}{x}) = F(x_0, x_1) + \beta u(\underset{\sim}{x}') \geq F(x_0, x_1) + \beta v^*(x_1) - \beta\varepsilon.$$

Since $\varepsilon > 0$ was arbitrary, (4) follows.

To establish (5), choose $x_0 \in X$ and $\varepsilon > 0$. From (3) and Lemma 4.1, it follows that one can choose $\underset{\sim}{x} = (x_0, x_1, \ldots) \in \Pi(x_0)$, so that

$$v^*(x_0) \leq u(\underset{\sim}{x}) + \varepsilon = F(x_0, x_1) + \beta u(\underset{\sim}{x}') + \varepsilon,$$

where $\underset{\sim}{x}' = (x_1, x_2, \ldots)$. It then follows from (2) that

$$v^*(x_0) \leq F(x_0, x_1) + \beta v^*(x_1) + \varepsilon.$$

Since $x_1 \in \Gamma(x_0)$, this establishes (5).

If $v^*(x_0) = +\infty$, then there exists a sequence $\{\underset{\sim}{x}^k\}$ in $\Pi(x_0)$ such that $\lim_{k\to\infty} u(\underset{\sim}{x}^k) = +\infty$. Since $x_1^k \in \Gamma(x_0)$, all k, and

$$u(\underset{\sim}{x}^k) = F(x_0, x_1^k) + \beta u(\underset{\sim}{x}'^k) \leq F(x_0, x_1^k) + \beta v^*(x_1^k), \quad \text{all } k,$$

it follows that (6) holds for the sequence $\{y^k = x_1^k\}$ in $\Gamma(x_0)$.

If $v^*(x_0) = -\infty$, then

$$u(\underset{\sim}{x}) = F(x_0, x_1) + \beta u(\underset{\sim}{x}') = -\infty, \quad \text{all } (x_0, x_1, x_2, \ldots) = \underset{\sim}{x} \in \Pi(x_0),$$

where $\underset{\sim}{x}' = (x_1, x_2, \ldots)$. Since F is real-valued (it does *not* take on the values $-\infty$ or $+\infty$), it follows that

$$u(\underset{\sim}{x}') = -\infty, \quad \text{all } x_1 \in \Gamma(x_0), \text{ all } \underset{\sim}{x}' \in \Pi(x_1).$$

Hence $v^*(x_1) = -\infty$, all $x_1 \in \Gamma(x_0)$. Since F is real-valued and $\beta > 0$, (7) follows immediately. ■

The next theorem provides a partial converse to Theorem 4.2. It shows that v^* is the only solution to the functional equation that satisfies a certain boundedness condition.

THEOREM 4.3 *Let X, Γ, F, and β satisfy Assumptions 4.1–4.2. If v is a solution to (FE) and satisfies*

(8) $$\lim_{n\to\infty} \beta^n v(x_n) = 0, \quad \text{all } (x_0, x_1, \ldots) \in \Pi(x_0), \text{ all } x_0 \in X,$$

then $v = v^*$.

Proof. If (8) holds, then v cannot take on the values $+\infty$ or $-\infty$. Hence v satisfies (4) and (5), and it is sufficient to show that this implies v satisfies (2) and (3).

If v satisfies (FE), then (4) implies that for all $x_0 \in X$ and $\underset{\sim}{x} \in \Pi(x_0)$

$$\begin{aligned} v(x_0) &\geq F(x_0, x_1) + \beta v(x_1) \\ &\geq F(x_0, x_1) + \beta F(x_1, x_2) + \beta^2 v(x_2) \\ &\;\;\vdots \\ &\geq u_n(\underset{\sim}{x}) + \beta^{n+1} v(x_{n+1}). \end{aligned}$$

Taking the limit on the right as $n \to \infty$ and using (8), we find that v satisfies (2).

To see that v satisfies (3), let $x_0 \in X$ and $\varepsilon > 0$ be given. Choose a sequence $\{\delta_n\}_{n=1}^{\infty}$ in $\mathbf{R}_+$ such that $\sum_{n=1}^{\infty} \beta^{n-1}\delta_n \leq \varepsilon/2$. From (5) it follows that one can choose $x_1 \in \Gamma(x_0)$ so that

$$v(x_0) \leq F(x_0, x_1) + \beta v(x_1) + \delta_1,$$

and choose $x_2 \in \Gamma(x_1)$ so that

$$v(x_1) \leq F(x_1, x_2) + \beta v(x_2) + \delta_2.$$

From these two inequalities it follows that

$$v(x_0) \leq F(x_0, x_1) + \beta F(x_1, x_2) + \beta^2 v(x_2) + \delta_1 + \beta\delta_2.$$

Continuing in this way, one defines $\underset{\sim}{x} = (x_0, x_1, x_2, \ldots) \in \Pi(x_0)$ such that

$$\begin{aligned} v(x_0) &\leq u_n(\underset{\sim}{x}) + \beta^{n+1} v(x_{n+1}) + (\delta_1 + \beta\delta_2 + \ldots + \beta^n \delta_{n+1}) \\ &\leq u_n(\underset{\sim}{x}) + \beta^{n+1} v(x_{n+1}) + \varepsilon/2, \quad n = 1, 2, \ldots. \end{aligned}$$

Since (8) implies that for n sufficiently large the second term is also less than $\varepsilon/2$, it follows that

$$v(x_0) \leq u_n(\underline{x}) + \varepsilon, \quad \text{all } n \text{ sufficiently large.}$$

Since $\varepsilon > 0$ was arbitrary, it then follows that (3) holds. ∎

It is an immediate consequence of Theorem 4.3 that the functional equation (FE) has at most one solution satisfying (8).

In summary, we have established two main results about solutions to (FE). Theorem 4.2 shows that v^* satisfies (FE). The functional equation may have other solutions as well, but Theorem 4.3 shows that these "extraneous" solutions always violate (8). Hence a solution to (FE) that satisfies (8) is v^*. The following example is a case where (FE) has an "extraneous" solution in addition to v^*.

Consider a consumer whose objective function is simply discounted consumption. The consumer has initial wealth $x_0 \in X = \mathbf{R}$, and he can borrow or lend at the interest rate $\beta^{-1} - 1$, where $\beta \in (0, 1)$. There are no constraints on borrowing, so his problem is simply

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

$$\text{s.t.} \quad 0 \leq c_t \leq x_t - \beta x_{t+1}, \quad t = 0, 1, \ldots,$$

$$x_0 \text{ given.}$$

Since consumption is unbounded, the supremum function is obviously $v^*(x) = +\infty$, all x. Now consider the recursive formulation of this problem. The return function is $F(x, y) = x - \beta y$, and the correspondence describing the feasible set is $\Gamma(x) = (-\infty, \beta^{-1}x]$; so the functional equation is

$$v(x) = \sup_{y \leq \beta^{-1}x} [x - \beta y + \beta v(y)].$$

The function $v^*(x) = +\infty$ satisfies this equation, as Theorem 4.2 implies, but the function $v(x) = x$ does, too. But since the sequence $x_t = \beta^{-t}x_0$, $t = 0, 1, \ldots$, is in $\Pi(x_0)$, (8) does not hold and Theorem 4.3 does not apply.

The next exercise gives two variations on Theorem 4.3 that are sometimes useful when (8) does not hold.

Exercise 4.3 Let X, Γ, F, and β satisfy Assumptions 4.1–4.2. Let v be a solution to (FE) with

$$\limsup_{n\to\infty} \beta^n v(x_n) \leq 0, \quad \text{all } x_0 \in X, \text{ all } (x_0, x_1, \ldots) \in \Pi(x_0),$$

a. Show that $v \leq v^*$.

b. Suppose in addition that for each $x_0 \in X$ and $\underline{x} \in \Pi(x_0)$, there exists $\underline{x}' = (x_0, x_1', x_2', \ldots) \in \Pi(x_0)$ such that $\lim_{n\to\infty} \beta^n v(x_n') = 0$ and $u(\underline{x}') \geq u(\underline{x})$. Show that $v = v^*$.

Our next task is to characterize feasible plans that attain the optimum, if any do. Call a feasible plan $\underline{x} \in \Pi(x_0)$ an *optimal plan from* x_0 if it attains the supremum in (SP), that is, if $u(\underline{x}) = v^*(x_0)$. The next two theorems deal with the relationship between optimal plans and those that satisfy the policy equation (1) for $v = v^*$. The next theorem shows that optimal plans satisfy (1).

THEOREM 4.4 *Let X, Γ, F, and β satisfy Assumption 4.1–4.2. Let $\underline{x}^* \in \Pi(x_0)$ be a feasible plan that attains the supremum in (SP) for initial state x_0. Then*

$$v^*(x_t^*) = F(x_t^*, x_{t+1}^*) + \beta v^*(x_{t+1}^*), \quad t = 0, 1, 2, \ldots. \tag{9}$$

Proof. Since $\underline{x}^*$ attains the supremum,

$$\begin{aligned} v^*(x_0^*) &= u(\underline{x}^*) = F(x_0, x_1^*) + \beta u(\underline{x}^{*\prime}) \\ &\geq u(\underline{x}) = F(x_0, x_1) + \beta u(\underline{x}'), \quad \text{all } \underline{x} \in \Pi(x_0). \end{aligned} \tag{10}$$

In particular, the inequality holds for all plans with $x_1 = x_1^*$. Since $(x_1^*, x_2, x_3, \ldots) \in \Pi(x_1^*)$ implies that $(x_0, x_1^*, x_2, x_3, \ldots) \in \Pi(x_0)$, it follows that

$$u(\underline{x}^{*\prime}) \geq u(\underline{x}'), \quad \text{all } \underline{x} \in \Pi(x_1^*).$$

Hence $u(\underline{x}^{*\prime}) = v(x_1^*)$. Substituting this into (10) gives (9) for $t = 0$. Continuing by induction establishes (9) for all t. ∎

The next theorem provides a partial converse to Theorem 4.4. It shows that any sequence satisfying (9) and a boundedness condition is an optimal plan.

THEOREM 4.5 *Let X, Γ, F, and β satisfy Assumptions 4.1–4.2. Let* $\underline{x}^* \in \Pi(x_0)$ *be a feasible plan from* x_0 *satisfying* (9), *and with*

(11) $$\limsup_{t\to\infty} \beta^t v^*(x_t^*) \le 0.$$

Then $\underline{x}^*$ *attains the supremum in (SP) for initial state* x_0.

Proof. Suppose that $\underline{x}^* \in \Pi(x_0)$ satisfies (9) and (11). Then it follows by an induction on (9) that

$$v^*(x_0) = u_n(\underline{x}^*) + \beta^{n+1} v^*(x_{n+1}^*), \quad n = 1, 2, \ldots.$$

Then using (11), we find that $v^*(x_0) \le u(\underline{x}^*)$. Since $\underline{x}^* \in \Pi(x_0)$, the reverse inequality holds, establishing the result. ■

The consumption example used after Theorem 4.3 can be modified to illustrate why (11) is needed. Let preferences be as specified before, so that $c_t = x_t - \beta x_{t+1} = F(x_t, x_{t+1})$, all t. However, let us prohibit indebtedness by requiring $x_t \ge 0$, all t. Then in sequence form the problem is

$$\max_{\{x_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t (x_t - \beta x_{t+1})$$

$$\text{s.t. } 0 \le x_{t+1} \le \beta^{-1} x_t, \quad t = 0, 1, \ldots,$$

$$x_0 \text{ given.}$$

If we cancel all of the offsetting terms in the objective function, it follows immediately that the supremum function is $v^*(x_0) = x_0$, all $x_0 \ge 0$. It is also clear that v^* satisfies the functional equation

$$v^*(x) = \max_{y \in [0, \beta^{-1}x]} [(x - \beta y) + \beta v^*(y)], \quad \text{all } x,$$

as Theorem 4.2 implies.

Now consider plans that attain the optimum. Given any $x_0 \ge 0$, the set of feasible plans $\Pi(x_0)$ consists of the sequences

$$(x_0, 0, 0, 0, \ldots), (x_0, \beta^{-1}x_0, 0, 0, \ldots),$$

$$(x_0, \beta^{-1}x_0, \beta^{-2}x_0, 0, \ldots), \text{ etc.},$$

and all convex combinations thereof. Hence every feasible plan satisfies (9). It is straightforward to verify that, as Theorem 4.5 implies, any plan that satisfies (11) as well yields utility $v^*(x_0) = x_0$. (Essentially, it does not matter when consumption occurs as long as it occurs in finite time.) On the other hand, the feasible plan $x_t = \beta^{-t}x_0$, $t = 0, 1, \ldots$, (in each period invest everything and consume nothing) yields discounted utility of zero, for all $x \ge 0$. For $x > 0$, however, it violates (11), so Theorem 4.5 does not apply.

We will call any nonempty correspondence $G: X \to X$, with $G(x) \subseteq \Gamma(x)$, all $x \in X$, a *policy correspondence*, since the set $G(x)$ is a feasible set of actions if the state is x. If G is single-valued, we will call it a *policy function* and denote it by a lowercase g. If a sequence $\underline{x} = (x_0, x_1, \ldots)$ satisfies $x_{t+1} \in G(x_t)$, $t = 0, 1, 2, \ldots$, we will say that $\underline{x}$ is *generated from* x_0 *by* G. Finally, we will define the *optimal policy correspondence* G^* by

$$G^*(x) = \{y \in \Gamma(x): v^*(x) = F(x, y) + \beta v^*(y)\}.$$

Then Theorem 4.4 shows that every optimal plan $\{x_t^*\}$ is generated from G^*, and Theorem 4.5 shows that any plan $\{x_t^*\}$ generated from G^*—if, in addition, it satisfies (11)—is an optimal plan.

4.2 Bounded Returns

In this section we study functional equations of the form

(1) $$v(x) = \max_{y \in \Gamma(x)} [F(x, y) + \beta v(y)],$$

under the assumption that the function F is bounded and the discount factor β is strictly less than one.

As above, let X be the set of possible values for the state variable; let $\Gamma: X \to X$ be the correspondence describing the feasibility constraints; let

$A = \{(x, y) \in X \times X: y \in \Gamma(x)\}$ be the graph of Γ; let $F: A \to \mathbf{R}$ be the return function; and let $\beta \geq 0$ be the discount factor. Throughout this section, we will impose the following two assumptions on X, Γ, F, and β.

ASSUMPTION 4.3 *X is a convex subset of $\mathbf{R}^l$, and the correspondence $\Gamma: X \to X$ is nonempty, compact-valued, and continuous.*

ASSUMPTION 4.4 *The function $F: A \to \mathbf{R}$ is bounded and continuous, and $0 < \beta < 1$.*

It is clear that under Assumptions 4.3–4.4, Assumptions 4.1–4.2 hold, so the sequence problem corresponding to (1) is well defined. Moreover, Theorems 4.2–4.5 imply that under these assumptions solutions to (1) coincide exactly—in terms of both values and optimal plans—to solutions of the sequence problem.

The requirement that X be a subset of a finite-dimensional Euclidean space could be relaxed in much of what follows, but at the expense of a substantial additional investment in terminology and notation. (Recall that the definitions of u.h.c. and l.h.c. provided in Chapter 3 applied only to correspondences from one Euclidean space to another.) The reader who is interested in applications where X is not a Euclidean space should note, however, that most of the arguments in this section apply much more broadly.

If B is a bound for $|F(x, y)|$, then the supremum function v^* satisfies $|v^*(x)| \leq B/(1 - \beta)$, all $x \in X$. In this case it is natural to seek solutions to (1) in the space $C(X)$ of bounded continuous functions $f: X \to \mathbf{R}$, with the sup norm: $\|f\| = \sup_{x \in X}|f(x)|$. Clearly, any solution to (1) in $C(X)$ satisfies the hypothesis of Theorem 4.3 and hence is the supremum function. Moreover, given a solution $v \in C(X)$ to (1), we can define the policy correspondence $G: X \to X$ by

$$(2) \qquad G(x) = \{y \in \Gamma(x): v(x) = F(x, y) + \beta v(y)\},$$

and Theorems 4.4 and 4.5 imply that for any $x_0 \in X$, a sequence $\{x_t^*\}$ attains the supremum in the sequence problem if and only if it is generated by G.

The rest of the section proceeds as follows. Define the operator T on $C(X)$ by

$$(3) \qquad (Tf)(x) = \max_{y \in \Gamma(x)}[F(x, y) + \beta f(y)],$$

so (1) becomes $v = Tv$. First, if we use only the boundedness and continuity restrictions in Assumptions 4.3 and 4.4, Theorem 4.6 establishes that $T: C(X) \to C(X)$, that T has a unique fixed point in $C(X)$, and that the policy correspondence G defined in (2) is nonempty and u.h.c. Theorem 4.7 establishes that under additional monotonicity restrictions on F and Γ, v is strictly increasing. Theorem 4.8 establishes that under additional concavity restrictions on F and convexity restrictions on Γ, v is strictly concave and G is a continuous (single-valued) function. Theorem 4.9 shows that if $\{v_n\}$ is a sequence of approximations defined by $v_n = T^n v_0$, with v_0 appropriately chosen, then the sequence of associated policy functions $\{g_n\}$ converges uniformly to the optimal policy function g given by (2). Finally, Theorem 4.11 establishes that if F is continuously differentiable, then v is, too.

THEOREM 4.6 *Let X, Γ, F, and β satisfy Assumptions 4.3 and 4.4, and let $C(X)$ be the space of bounded continuous functions $f: X \to \mathbf{R}$, with the sup norm. Then the operator T maps $C(X)$ into itself, $T: C(X) \to C(X)$; T has a unique fixed point $v \in C(X)$; and for all $v_0 \in C(X)$,*

$$(4) \qquad \|T^n v_0 - v\| \leq \beta^n \|v_0 - v\|, \quad n = 0, 1, 2, \ldots.$$

Moreover, given v, the optimal policy correspondence $G: X \to X$ defined by (2) is compact-valued and u.h.c.

Proof. Under Assumptions 4.3 and 4.4, for each $f \in C(X)$ and $x \in X$, the problem in (3) is to maximize the continuous function $[F(x, \cdot) + \beta f(\cdot)]$ over the compact set $\Gamma(x)$. Hence the maximum is attained. Since both F and f are bounded, clearly Tf is also bounded; and since F and f are continuous, and Γ is compact-valued and continuous, it follows from the Theorem of the Maximum (Theorem 3.6) that Tf is continuous. Hence $T: C(X) \to C(X)$.

It is then immediate that T satisfies the hypotheses of Blackwell's sufficient conditions for a contraction (Theorem 3.3). Since $C(X)$ is a Banach space (Theorem 3.1), it then follows from the Contraction Mapping Theorem (Theorem 3.2), that T has a unique fixed point $v \in C(X)$, and (4) holds. The stated properties of G then follow from the Theorem of the Maximum, applied to (1). ■

It follows immediately from Theorem 4.3 that under the hypotheses of Theorem 4.6, the unique bounded continuous function v satisfying

(1) is the supremum function for the associated sequence problem. That is, Theorems 4.3 and 4.6 together establish that under Assumptions 4.3–4.4 the supremum function is bounded and continuous. Moreover, it then follows from Theorems 4.5 and 4.6 that there exists at least one optimal plan: any plan generated by the (nonempty) correspondence G is optimal.

To characterize v and G more sharply, we need more information about F and Γ. The next two results show how Corollary 1 to the Contraction Mapping Theorem can be used to obtain more precise characterizations of v and G.

ASSUMPTION 4.5 *For each y, $F(\cdot, y)$ is strictly increasing in each of its first l arguments.*

ASSUMPTION 4.6 *Γ is monotone in the sense that $x \leq x'$ implies $\Gamma(x) \subseteq \Gamma(x')$.*

THEOREM 4.7 *Let X, Γ, F, and β satisfy Assumptions 4.3–4.6, and let v be the unique solution to (1). Then v is strictly increasing.*

Proof. Let $C'(X) \subset C(X)$ be the set of bounded, continuous, nondecreasing functions on X, and let $C''(X) \subset C'(X)$ be the set of strictly increasing functions. Since $C'(X)$ is a closed subset of the complete metric space $C(X)$, by Theorem 4.6 and Corollary 1 to the Contraction Mapping Theorem (Theorem 3.2), it is sufficient to show that $T[C'(X)] \subseteq C''(X)$. Assumptions 4.5 and 4.6 ensure that this is so. ∎

ASSUMPTION 4.7 *F is concave; that is,*

$$F[\theta(x, y) + (1 - \theta)(x', y')] \geq \theta F(x, y) + (1 - \theta)F(x', y'),$$

$$\text{all } (x, y), (x', y') \in A, \quad \text{and all } \theta \in (0, 1).$$

In addition, the inequality is strict if $x \neq x'$.

ASSUMPTION 4.8 *Γ is convex in the sense that for any $0 \leq \theta \leq 1$, and $x, x' \in X$,*

$$y \in \Gamma(x) \quad \text{and} \quad y' \in \Gamma(x') \quad \text{implies}$$

$$\theta y + (1 - \theta)y' \in \Gamma[\theta x + (1 - \theta)x'].$$

Assumption 4.8 implies that for each $x \in X$, the set $\Gamma(x)$ is convex and there are no "increasing returns." Note that since X is convex, Assumption 4.8 is equivalent to assuming that the graph of Γ (the set A) is convex.

THEOREM 4.8 *Let X, Γ, F, and β satisfy Assumptions 4.3–4.4 and 4.7–4.8; let v satisfy (1); and let G satisfy (2). Then v is strictly concave and G is a continuous, single-valued function.*

Proof. Let $C'(X) \subset C(X)$ be the set of bounded, continuous, weakly concave functions on X, and let $C''(X) \subset C'(X)$ be the set of strictly concave functions. Since $C'(X)$ is a closed subset of the complete metric space $C(X)$, by Theorem 4.6 and Corollary 1 to the Contraction Mapping Theorem (Theorem 3.2), it is sufficient to show that $T[C'(X)] \subseteq C''(X)$.

To verify that this is so, let $f \in C'(X)$ and let

$$x_0 \neq x_1, \quad \theta \in (0, 1), \quad \text{and} \quad x_\theta = \theta x_0 + (1 - \theta)x_1.$$

Let $y_i \in \Gamma(x_i)$ attain $(Tf)(x_i)$, for $i = 0, 1$. Then by Assumption 4.8, $y_\theta = \theta y_0 + (1 - \theta)y_1 \in \Gamma(x_\theta)$. It follows that

$$\begin{aligned}(Tf)(x_\theta) &\geq F(x_\theta, y_\theta) + \beta f(y_\theta)\\ &> \theta[F(x_0, y_0) + \beta f(y_0)] + (1 - \theta)[F(x_1, y_1) + \beta f(y_1)]\\ &= \theta(Tf)(x_0) + (1 - \theta)(Tf)(x_1),\end{aligned}$$

where the first line uses (3) and the fact that $y_\theta \in \Gamma(x_\theta)$; the second uses the hypothesis that f is concave and the concavity restriction on F in Assumption 4.7; and the last follows from the way y_0 and y_1 were selected. Since x_0 and x_1 were arbitrary, it follows that Tf is strictly concave, and since f was arbitrary, that $T[C'(X)] \subseteq C''(X)$.

Hence the unique fixed point v is strictly concave. Since F is also concave (Assumption 4.7) and, for each $x \in X$, $\Gamma(x)$ is convex (Assumption 4.8), it follows that the maximum in (3) is attained at a unique y value. Hence G is a single-valued function. The continuity of G then follows from the fact that it is u.h.c. (Exercise 3.11). ∎

Theorems 4.7 and 4.8 characterize the value function by using the fact that the operator T preserves certain properties. Thus if v_0 has property

P and if P is preserved by T, then we can conclude that each function in the sequence $\{T^n v_0\}$ has property P. Then, if P is preserved under uniform convergence, we can conclude that v also has property P. The same general idea can be used to establish facts about the policy function g, but we need to establish the sense in which the approximate policy functions—the functions g_n that attain $T^n v_0$—converge to g. The next result draws on Theorem 3.8 to address this issue.

THEOREM 4.9 (*Convergence of the policy functions*) *Let X, Γ, F, and β satisfy Assumptions 4.3–4.4 and 4.7–4.8, and let v and g satisfy (1) and (2). Let $C'(X)$ be the set of bounded, continuous, concave functions $f: X \to \mathbf{R}$, and let $v_0 \in C'(X)$. Let $\{(v_n, g_n)\}$ be defined by*

$$v_{n+1} = Tv_n, \quad n = 0, 1, 2, \ldots, \quad \text{and}$$

$$g_n(x) = \underset{y \in \Gamma(x)}{\operatorname{argmax}} [F(x, y) + \beta v_n(y)], \quad n = 0, 1, 2, \ldots.$$

Then $g_n \to g$ pointwise. If X is compact, then the convergence is uniform.

Proof. Let $C''(X) \subset C'(X)$ be the set of strictly concave functions $f: X \to \mathbf{R}$. As shown in Theorem 4.8, $v \in C''(X)$. Moreover, as shown in the proof of that theorem, $T[C'(X)] \subseteq C''(X)$. Since $v_0 \in C'(X)$, it then follows that every function v_n, $n = 1, 2, \ldots$, is strictly concave. Define the functions $\{f_n\}$ and f by

$$f_n(x, y) = F(x, y) + \beta v_n(y), \quad n = 1, 2, \ldots, \quad \text{and}$$

$$f(x, y) = F(x, y) + \beta v(y).$$

Since F satisfies Assumption 4.7, it follows that each function f_n, $n = 1, 2, \ldots$, is strictly concave, as is f. Hence Theorem 3.8 applies and the desired results are proved. ■

The next exercise establishes a related convergence result that is often useful for computational purposes.

Exercise 4.4 Let X, Γ, F, and β satisfy Assumptions 4.3–4.4 and 4.7–4.8, and let $C(X)$ be as given above. Let H be the set of continuous

functions $h: X \to X$ such that $h(x) \in \Gamma(x)$, all $x \in X$. For any $h \in H$, define the operator T_h on $C(X)$ by

$$(T_h f)(x) = F[x, h(x)] + \beta f[h(x)].$$

a. Show that $T_h: C(X) \to C(X)$; that T_h is monotone; and that T_h is a contraction mapping on $C(X)$.

Let $h_0 \in H$ be given, and consider the following algorithm. Given h_n, let w_n be the unique fixed point of T_{h_n}. Given w_n, consider the problem

$$\max_{y \in \Gamma(x)} [F(x, y) + \beta w_n(y)].$$

b. Show that the optimal policy function for this problem is single-valued and continuous. Call this function h_{n+1}, and note that $h_{n+1} \in H$.

c. Show that the sequence of functions $\{w_n\}$ converges to v, the unique solution to (1). [*Hint.* Show that $w_0 \le Tw_0 \le w_1 \le Tw_1 \le \ldots$.]

An algorithm based on Exercise 4.4 involves applying the operators T_{h_n}—operators that require no maximization—repeatedly and applying T only infrequently. Since maximization is usually the expensive step in these computations, the savings can be considerable.

Once the existence of a unique solution $v \in C(X)$ to the functional equation (1) has been established, we would like to treat the maximum problem in that equation as an ordinary programming problem and use the standard methods of calculus to characterize the policy function g. For example, consider the functional equation for the one-sector growth model:

$$v(x) = \max_{0 \le y \le f(x)} \{U[f(x) - y] + \beta v(y)\}.$$

If we knew that v was differentiable (and that the solution to the maximum problem in (1) was always interior), then the policy function g would be given implicitly by the first-order condition

$$U'[f(x) - g(x)] - \beta v'[g(x)] = 0. \tag{5}$$

Moreover, if we knew that v was twice differentiable, the monotonicity of g could be established by differentiating (5) with respect to x and exam-