Chapter 4

Introduction to Dynamic Programming

An approach to solving dynamic optimization problems alternative to optimal control was pioneered by Richard Bellman beginning in the late 1950s. Bellman emphasized the economic applications of dynamic programming right from the start. Unlike optimal control, dynamic programming has been fruitfully applied to problems in both continuous and discrete time. It is generally more powerful than optimal control for dealing with stochastic problems, and it does not always require some of the differentiability and continuity assumptions inherent to optimal control. Dynamic programming can also deal with problems that arise concerning time inconsistency, in ways that are difficult to deal with in optimal control.

In this chapter we lay out the ground work for dynamic programming in both deterministic and stochastic environments. We will see how to characterize a dynamic programming problem and how to solve it. We will also present a series of theorems that are extremely useful for characterizing the properties of solution for the many cases in which an explicit analytical solution cannot be obtained. Subsequent chapters present numerous applications of the methods developed here.

1. Deterministic Finite-Horizon Problems

Consider the following finite-horizon consumption problem:

$$\max_{\{c_t\}_{t=0}^{T}} \sum_{t=0}^{T} \beta^t u(c_t), \tag{1.1}$$

subject to

$$k_{t+1} = f\left(k_t\right) - c_t \,. \tag{1.2}$$

Capital depreciates at the rate of 100 percent per period. Equation (1) is maximized subject to the further constraint that

$$0 \le k_{t+1} \le f\left(k_t\right), \quad 0 \le t \le T\,, \;\; k_0 \text{ given,} \tag{1.3}$$

which states that capital can neither be negative nor exceed output. Substituting (1.2) into (1.1) yields

$$\max_{\{k_{t+1}\}_{t=0}^{T}} \sum_{t=0}^{T} \beta^t u\left(f\left(k_t\right) - k_{t+1}\right), \tag{1.4}$$

so we have changed the problem from maximizing by choice of consumption in each period to one of choosing next period's capital stock.[1] One approach to solving this problem is by brute-force optimization. This is possible because there are a finite number, T, of choices to make. To see this, maximize (1.4) with respect to k_{t+1} to obtain the first-order condition[2]

$$-\beta^t u'\left(f\left(k_t\right) - k_{t+1}\right) + \beta^{t+1} u'\left(f\left(k_{t+1}\right) - k_{t+2}\right) f'\left(k_{t+1}\right) = 0\,,$$

or

$$u'\left(f\left(k_t\right) - k_{t+1}\right) = \beta u'\left(f\left(k_{t+1}\right) - k_{t+2}\right) f'\left(k_{t+1}\right). \tag{1.5}$$

This first-order condition must be satisfied for each t=0,1, . . . , T–1. It is clear that the optimal solution for k_{T+1} is zero, since it only appears in the term $\left[f\left(k_T\right) - k_{T+1}\right]$. Equation (1.5) thus represents T equations in T unknowns. The variables k_0 and k_{T+1} appear in two of these equations, but we already know what they are.

To interpret (1.5), replace $f\left(k_t\right) - k_{t+1}$ with c_t to get

$$u'\left(c_t\right) = \beta u'\left(c_{t+1}\right) f'\left(k_{t+1}\right). \tag{1.6}$$

---

[1] This is not necessary to do, but it often makes the problem easier to deal with algebraically.

[2] We are assuming that $f(k)$ and $u(c)$ have the forms necessary to ensure an interior solution, so we do not need to worry about the Kuhn-Tucker inequality constraints. (What are these assumptions?)

The left-hand side is the marginal utility of consumption in period t. The right-hand side consists of the product of present value of the marginal utility consumption in period t+1 and the marginal productivity of capital. One unit of consumption foregone in period t increases the capital stock in period t+1 by one unit, and this raises output in period t+1 by an amount equal to the marginal product of capital. Converting this to utility measures and discounting back to period t, (1.6) states that the marginal unit in consumption must have equal value across two adjacent periods.

We will soon be interested in extending this model to allow for an infinite planning horizon. The difficulty is that the terminal condition k_{T+1}=0 goes away, leaving us with T equations and T+1 unknowns. However, it turns out that there is an alternative approach to solving this finite-horizon problem that is useful not only for the problem at hand, but also for extending the model to the infinite-horizon case. This is the dynamic programming approach.

Suppose we obtained the solution to the period-1 problem,

$$\max_{\{k_{t+1}\}_{t=1}^{T}} \sum_{t=1}^{T} \beta^{t-1} u\Big(f\left(k_t\right) - k_{t+1}\Big),$$

k_1>0 given. Whatever the solution to this problem is, let $V_1\left(k_1\right)$ denote the value obtained from period 1 onwards. Note that the value depends on the initial capital stock. It then follows that the period-0 problem can be written as

$$V_0\left(k_0\right) = \max_{k_1} \Big\{u\Big(f\left(k_0\right) - k_1\Big) + \beta V_1\left(k_1\right)\Big\}.$$

In fact, for any t we can define an analogous equation

$$V_t\left(k_t\right) = \max_{k_{t+1}} \Big\{u\Big(f\left(k_t\right) - k_{t+1}\Big) + \beta V_{t+1}\left(k_{t+1}\right)\Big\}, \tag{1.7}$$

subject to $0 \le k_{t+1} \le f\left(k_t\right)$, k_t given, for t=T, T–1, . . . , 0. Equation (1.7) is a particular application of **Bellman's Principle of Optimality**:

Theorem 1.1 (Bellman's Principle of Optimality). *An optimal policy has the property that, whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the initial decision.*

Bellman and Dreyfus (1962), among others, provide a proof of the principle, but it is so intuitive that we won't bother to show it here.

The recursive sequence in (1.7) begins in the final time period with $V_{T+1}(k_T) = 0$. Note that solving (1.7) sequentially will yield exactly the same set of equations as (1.5). In period T, we have

$$V_T(k_T) = \max_{k_{T+1}} \left\{ u\left(f(k_T) - k_{T+1}\right)\right\}, \tag{1.8}$$

which implies that k_{T+1}=0. In period T–1 we have

$$V_{T-1}(k_{T-1}) = \max_{k_T} \left\{ u\left(f(k_{T-1}) - k_T\right) + \beta V_T(k_T)\right\},$$

which gives the first-order condition

$$\begin{aligned} 0 &= -u^{\prime}\left(f(k_{T-1}) - k_T\right) + \beta V_T^{\prime}(k_T) \\ &= -u^{\prime}\left(f(k_{T-1}) - k_T\right) + \beta u^{\prime}\left(f(k_T)\right) f^{\prime}(k_T), \end{aligned}$$

where the second line comes from differentiating (1.8). Repeatedly solving (1.7) for each time period yields the system of T equations in (1.5).

EXERCISE 1.1 (Cake eating). *Suppose you have a cake of size x_t, with x_0 given. In each period, t=*1, 2, 3, . . . , *T, you can consume some of the cake and save the remainder. Let c_t be your consumption in period t and let $u(c_t)$ represent the flow of utility from this consumption. Assume that $u()$ is differentiable, strictly increasing and concave, with $\lim_{c\to 0} u^{\prime}(c) = \infty$. Let lifetime utility be represented by $\sum_{t=0}^{\infty} \beta^t u(c)$. Characterize the optimal path of consumption* $\{c_t\}_{t=0}^{T}$, (a) *by the direct method,* (b) *by the method of dynamic programming.*

2. Deterministic Infinite-Horizon Problems

So how does the dynamic programming approach help us in the infinite-horizon case? Consider again the period-1 version of the consumption problem, but now written for an infinite planning horizon:

$$V_1(k_1) = \max_{\{k_{t+1}\}_{t=1}^{\infty}} \sum_{t=1}^{\infty} \beta^{t-1} u\left(f\left(k_t\right) - k_{t+1}\right). \tag{2.1}$$

Assume for the moment that a solution to this problem exists. Let the maximized value of the objective function be $V_1(k_1)$. Then, according to Bellman's principle of optimality, the period-0 problem can be written as

$$V_0\left(k_0\right) = \max_{k_1} \left\{u\left(f\left(k_0\right) - k_1\right) + \beta V_1\left(k_1\right)\right\} \tag{2.2}$$

subject to $0 \leq k_1 \leq f(k_0)$, with k_0 given. Note that we could re-index time in (2.1) by replacing t with s+1 to obtain

$$\beta \sum_{s=0}^{\infty} \beta^s u\left(f(k_s) - k_{s+1}\right) = V_0(k_0). \tag{2.3}$$

It then becomes clear that $V_0(k_0)$ and $V_1(k_1)$ must be exactly the same function because (2.1) and (2.3) differ only by notation. That is, if a solution exists, it must satisfy

$$V\left(k_0\right) = \max_{k_1} \left\{u\left(f\left(k_0\right) - k_1\right) + \beta V\left(k_1\right)\right\}.$$

Because time does not matter directly in this problem, we can drop the subscript notation and let k' denote next period's value of k:

$$V\left(k\right) = \max_{k'} \left\{u\left(f\left(k\right) - k'\right) + \beta V\left(k'\right)\right\}, \tag{2.4}$$

subject to $0 \leq k' \leq f(k)$, k_0 given. Equation (2.4) is usually referred to as the **Bellman equation** of dynamic programming. The first-order condition for this maximization problem is

$$u'\left(f(k) - k'\right) = \beta V'\left(k'\right), \tag{2.5}$$

which is not too helpful as it stands because we do not know the function $V(k')$. However, we can use the **envelope theorem** to make some more progress. Differentiate the value function in (2.4) with respect to k, yielding[3]

$$\begin{aligned} V'(k) &= u'\left(f(k)-k'\right)f'(k)+\left[-u'\left(f(k)-k'\right)+\beta V'(k')\right]\frac{dk'}{dk}. \\ &= u'\left(f(k)-k'\right)f'(k). \end{aligned} \tag{2.6}$$

The term in square brackets is equal to zero from the first-order condition (2.5) (this is the application of the envelope theorem). Update (2.6) by one period,

$$V'(k') = u'\left(f(k')-k''\right)f'(k'),$$

and substitute into (2.5) to obtain

$$u'\left(f(k)-k'\right) = \beta u'\left(f(k')-k''\right)f'(k'). \tag{2.7}$$

In terms of date subscripts, we have

$$u'\left(f(k_t)-k_{t+1}\right) = \beta u'\left(f\left(k_{t+1}\right)-k_{t+2}\right)f'\left(k_{t+1}\right),$$

which is the solution (1.5) we arrived at before for the finite horizon case.

EXERCISE 2.1 (Cake eating forever). *Extend the cake-eating example* (1.1) *to an infinite planning horizon.*

(a) *Derive the Bellman equation and use it to characterize the optimal policy.*

(b) *Assume that utility is given by* $u(c_t)=\ln(c_t)$. *Use the method of undetermined coefficients to show that the value function takes the linear form* $V(x)=A+B\ln(x)$.

(c) *Show that the optimal policy is to eat a constant fraction* $1-\beta$ *of the cake in each period.*

(d) *What is the optimal policy when* $u(c)=c$?

[3] This requires, of course, that $V(k)$ be differentiable. It turns out that if u is differentiable then $V(k)$ is also differentiable under quite general conditions. The result was established by Benveniste and Scheinkman (1979), so (2.6) is sometimes referred to in the literature as the Benveniste-Scheinkman condition.

In other infinite-horizon dynamic programs, it may not be reasonable to assume that time does not matter, and so the time subscripts on the problem are fundamentally important. For example, in the consumption problem we have just seen, the resource constraint could take the form

$$k_{t+1} = A_t f(k_t) - c_t ,$$

where A_t is a technology parameter that fluctuates with time. In this case, the value function depends on time, which we denote in the following way:

$$V_t(k) = \max_{k'} \left\{ u\left(A_t f(k) - k'\right) + \beta V_{t+1}(k') \right\}.$$

Despite the apparent added complexity, the general approach to finding the optimal policy remains the same. The first-order condition is

$$u'\left(A_t f(k) - k'\right) = \beta V'_{t+1}(k').$$

The envelope theorem tells us that

$$V'_t(k) = u'\left(A_t f(k) - k'\right) A_t f'(k).$$

Updating by one period,

$$V'_{t+1}(k') = u'\left(A_{t+1} f(k') - k''\right) A_{t+1} f'(k'),$$

and substituting into the first-order condition yields

$$u'\left(A_t f(k) - k'\right) = \beta u'\left(A_{t+1} f(k') - k''\right) A_{t+1} f'(k'). \tag{2.8}$$

With the exception of A_t and A_{t+1}, this is exactly the same as the result obtained in (2.7).

The next example is a somewhat more complicated application of these ideas.

EXAMPLE 2.2 (Eekhout and Jovanovic [2002]). Firms produce output, y, using human capital, k, according to the production function $y = A_t(k)k$. The term $A_t(k)$ is a productivity parameter that changes over time, and captures knowledge spillovers to the firm from its competitors. The greater a firm's own level of human capital, the less it has to learn from others. Hence, it is assumed that $A'_t(k) < 0$. Firms face a cost of adjustment to k that is proportional to output. Given k units this period, the firm can have k' units next period at a cost of $yc(k'/k)$. It is assumed that $c' > 0$ and $c'' > 0$.

(a) *Derive the Bellman equation for this problem.*

$$V_t(k) = \max_{k'} \left\{ \left(1 - c\left(\frac{k'}{k}\right)\right) A_t(k)k + \beta V_{t+1}(k') \right\}.$$

Because $A_t(k)$ may vary from period to period, the time subscripts on the value function are important.

(b) *Derive a difference equation in k that characterizes the optimal policy.*

The first-order condition is

$$c'\left(\frac{k'}{k}\right) A_t(k) = \beta V'_{t+1}(k')$$

The envelope theorem says

$$V'_t(k) = A_t(k)\left[\frac{k'}{k} c'\left(\frac{k'}{k}\right) + \left(1 - c\left(\frac{k'}{k}\right)\right)\left(\frac{A'_t(k)k}{A_t(k)} + 1\right)\right].$$

Updating by one period and substituting back into the first-order condition gives

$$c'\left(\frac{k'}{k}\right) A_t(k) = \beta A_{t+1}(k')\left[\frac{k''}{k'} c'\left(\frac{k''}{k'}\right) + \left(1 - c\left(\frac{k'}{k'}\right)\right)\left(\frac{A'_{t+1}(k')k'}{A_t(k')} + 1\right)\right], \qquad (2.9)$$

which implicitly defines a second-order difference equation in k.

(c) *Let*

$$\varepsilon = \left|\frac{kA'_t(k)}{A_t}\right|$$

denote the absolute value of the elasticity of A_t with respect to k. Let $x_k = k_{t+1}/k_t$ denote the growth factor for human capital, and let $x_A = A_{t+1}/A_t$ denote the growth factor for productivity. Assume that ε, x_k and x_A are constant for all t. Derive a stationary solution relating the elasticity ε to x_k.

Substitute $\varepsilon = \left| A^{/}_{t+1}(k')k'/A_{t+1}(k') \right|$ into (2.9), divide throughout by $A_t(k)$; replace $A_{t+1}(k')/A_t(k)$ with x_A, and let $k''/k' = k'/k = x_k$:

$$\varepsilon = 1 - \frac{\left(\dfrac{1}{\beta x_A} - x_k\right) c^{/}\left(x_k\right)}{1 - c\left(x_k\right)}. \tag{2.10}$$

(d) *Why is it reasonable to assume $\beta x_A x_k < 1$?*

Output is $y = A(k)k$. If A and k grow too fast, the present value of output will become infinite. Constraining x_A and x_k ensures that the present value of output diminishes to zero for periods far enough in the future.

(e) *Show that, if ε is large enough, $dx_k/d\varepsilon < 0$. Interpret this finding.*

Direct differentiation of (2.10) gives

$$\frac{d\varepsilon}{dx_k} = \frac{c^{/}}{1-c} - \left(\frac{1}{\beta x_A} - x_k\right)\left(\frac{(1-c)c^{//} + \left(c^{/}\right)^2}{(1-c)^2}\right).$$

Rearranging and making use of (2.10) allows us to write

$$\frac{dx_x}{d\varepsilon} = \frac{\dfrac{1}{\beta x_A} - x_k}{(1-\varepsilon) - (1-\varepsilon)^2\left[\left(\dfrac{1}{\beta x_A} - x_k\right)^2 \dfrac{c^{//}}{c^{/}} + \left(\dfrac{1}{\beta x_A} - x_k\right)\right]}.$$

As $\beta x_A x_k < 1$ the term in square brackets is unambiguously positive, as is the numerator. Thus, if $\varepsilon > 1$, $dx_k/d\varepsilon < 0$. There are two effects of human capital growth. First, for given A, output is increased. Second, A is reduced as it becomes more difficult to absorb knowledge from other firms. If the latter effect is large enough (i.e. if ε is large enough), the firm reduces its investment in human capital, preferring to free ride on the knowledge developed by other firms. Eekhout and Jovanovic use this insight to develop an equilibrium model of inequality. •

Although we can usually make good progress in characterizing optimal policies defined implicitly by equations such as (2.7) and (2.8), in most cases it will not be possible to obtain an explicit solution for the optimal policy. This is unfortunate because we would usually like an explicit solution in order to solve for the value function. When explicit solutions are not available we must take a more indirect route to ask some of our basic questions, including

- Can we prove formally existence and uniqueness of the value function?
- Can we prove there is a unique optimal policy for the choice or state variable?
- What other properties of the value function can be derived?

We will address these questions in the remainder of the section. It should be noted that, in the interests of tractability, we will be stating theorems that may be more restrictive than necessary. The standard treatment of the following material at its most general level is to be found in Stokey and Lucas (1989), a rather difficult and time-consuming book.

A Contraction Mapping Theorem for Bounded Functional Equations

Recall from Chapter 3 the following contraction mapping theorem for fixed point expressions:

THEOREM [ch. 3] 3.7 (Contraction mapping theorem). *Let $f(x)$ denote a continuous function which maps a value x from a closed, bounded interval into a closed, bounded interval. If $f(x)$ is a contraction mapping, then there exists exactly one fixed point $x^* = f(x^*)$.*

To explore uniqueness and existence of a solution to the Bellman equation, we will replace x_t and x_{t-1} in our difference equations, with *functions*, $f(x)$ and $g(x)$. That is, we write

$$g(x) = Tf(x'), \tag{2.11}$$

where T denotes some operation on the first function f that yields the second function g. Equation (2.11) is called a **functional equation**, because from it we want to solve for the *function* $g(x)$ for all values of x rather than for any particular *value* of x. For example, in the consumption problem we have

$$V(k) = \max_{k'} u\left(f(k) - k'\right) + \beta V\left(k'\right). \tag{2.12}$$

In (2.12), the operator T is the act of multiplying V by β, adding $u\left(f(k) - k'\right)$ to it, and then maximizing the resulting function by choice of k'. We would like find a unique function $V(k)$ that satisfies the recursive representation in (2.12) for all k. This may be very difficult so, before we proceed, we would like to know if one exists. Fortunately, the contraction mapping theorem also applies to such functional equations.

Although we will leave out the technical details associated with the theorem, we do need to introduce a new distance function, known as the **supremum norm** and denoted by $\|x-y\|$. Let $f(x)$ and $g(x)$ denote two functions of $x \in [a,b]$, then the supremum norm, $\|f-g\|$ denotes the maximum absolute difference between the two functions observed as x takes on different values in the closed interval $[a,b]$. An operator on a function is a contraction mapping whenever applying the operator to two such functions brings them closer together for any admissible values of x. Using the supremum norm as our measure of distance, if T is a contraction mapping then $\|Tf(x) - Tg(x)\| < \|f(x) - g(x)\|$.[4] This will require that the functions are continuous. Additionally, for the supremum norm to exist, the functions $f(x)$ and $g(x)$ must have well-defined maxima and minima, and the contraction mapping theorem applies to sets of functions that have them. That is, the contraction mapping theorem applies to sets of continuous functions mapping closed bounded intervals into closed bounded intervals. For such sets, the supremum norm always exists.[5]

---

[4] The idea here is that as the supremum norm goes to zero the two functions get closer and closer together and eventually become the same function.

[5] A set of functions with these properties is called a "complete metric space." We don't need to get into where this name comes from, but if you come across the term now you know what it means.

The simplest way to ensure that a maximum exists is to bound the one-period payoffs in some way. For example, in (2.12) above, there could be some maximum value, $\overline{u} < \infty$, of the utility function regardless of how much capital the consumer has. Given discounting, it then follows that $V < \overline{u}/(1-\beta) < \infty$ so V is bounded from above. Even so, we still need to ensure that $V(k)$ can actually attain its maximum, and this requires that k must be able to attain the value that maximizes $V(k)$. Imagine that this is at the boundary of the interval from which k is drawn. Then we need to ensure that k can actually attain the value at this boundary. That is, we need to ensure that the interval for k includes its own boundary. Put another way, k must be drawn from a closed, bounded interval. For example, there might be some maximum feasible amount of k, say $\overline{k}$, such that $u\left(f\left(\overline{k}\right)-k'\right) \leq \overline{u} < \infty$ for any $k' \in \left[0,\overline{k}\right]$. If k is bounded, then assuming that u is continuous everywhere ensures that u is bounded. One also needs to worry about the lower bound. Imagine, for example, that $u(c)$=ln(c). Then $u(0) \to -\infty$ and V will not be bounded below. However, this is often not a practical concern because we have a maximization problem. In the consumption problem, for any k>0, we will always want to ensure that c>0 for all t. Hence, as u=ln(c) is continuous and c can be bounded above zero, $u(c)$ is bounded below and so is V.

Assuming such bounds leads to the following existence and uniqueness theorem:

THEOREM 2.1 (Contraction mapping theorem for bounded returns). *Let $C[a,b]$ be the set of all continuous functions mapping values from a bounded closed interval into a bounded closed interval. Let $f(x)$ and $g(x)$ be two arbitrary functions from this set. Now consider an operator, T, on that function, such that $g(x)=Tf(x)$. If T is a contraction mapping then there exists exactly one function satisfying $f(x) = Tf(x)$.*

PROOF. We will not prove existence, which is tricky, but uniqueness is easy. Suppose there were two functions $f^*(x)$ and $g^*(x)$, satisfying $f^*(x)=T\,f^*(x)$ and $g^*(x)=T\,g^*(x)$ for all x. Then, as T is a contraction we have $\|Tf^*(x)-Tg^*(x)\| < \|f^*(x)-g^*(x)\|$. But as $f^*(x)=T\ f^*(x)$ and $g^*(x)=T\ g^*(x)$, this implies $\|f^*(x)-g^*(x)\| = \|Tf^*(x)-Tg^*(x)\|$ $< \|f^*(x)-g^*(x)\|$, a contradiction. Thus there cannot be two fixed points. •

The contraction mapping theorem is, of course, a special type of fixed point theorem. In fact, another name for it is the **Banach fixed point theorem**. The main difficulty is that deciding whether the operator T is a contraction mapping directly is likely to be a very hard problem. And that is why it is so nice to be helped out by a handy little theorem from Mr. Blackwell:

THEOREM 2.2 (Blackwell's contraction mapping theorem). (a) *The operator T is a contraction if it has the following two properties*:

- (Monotonicity). *If $f(x) \leq g(x)$ for all x, then $Tf(x) \leq Tg(x)$ for all x.*
- (Discounting). *Let a be a positive constant. There exists some $\delta \in (0,1)$ such that $T(f+a)(x) \leq Tf(x) + \delta a$.*

(b) *If T is a contraction mapping, and the other assumptions of Theorem 2.1 are satisfied, the unique fixed point of the functional equation $g(x) = Tf(x)$, which we will denote by $f^*(x) = Tf^*(x)$, can be found by repeatedly applying the operator T to the function f. That is, $f^*(x) = \lim_{n\to\infty} T^n f(x)$.*

PROOF. If $f(x) \leq g(x)$ for all admissible values of x, then $f(x) \leq g(x) + \|f(x) - g(x)\|$, because $\|f(x) - g(x)\|$ is a distance function and therefore is non-negative. If monotonicity holds, we have $Tf(x) \leq T\left(g(x) + \|f(x) - g(x)\|\right)$, and if discounting holds this inequality further implies $Tf(x) \leq Tg(x) + \beta\|f(x) - g(x)\|$ for some $\beta < 1$. Subtracting $Tg(x)$ from both sides of the inequality, we have $Tf(x) - Tg(x) \leq \beta\|f(x) - g(x)\|$. This inequality holds for any admissible x, including the value that makes for the largest difference between $Tf(x)$ and $Tg(x)$. Thus, applying the supremum norm, $\|Tf(x) - Tg(x)\| \leq \beta\|f(x) - g(x)\|$, which is the definition of a contraction mapping. •

It is easiest to show what this theorem means by example. Example 2.2 is particularly simple. Example 2.3 continues the consumption problem. Both examples deal with the existence question in part (a) of the theorem.

EXAMPLE 2.2. Let C be the set of all continuous and bounded functions on the closed interval [0,1] and equipped with the supremum norm. Let the functional equation be given by $g(x)$=0.5$f(x)$, where f and g are any two function drawn from the set C. We will first use Blackwell's theorem to show that the operator T in $Tf(x)$=0.5$f(x)$ is a contraction mapping (it is obvious that this is the case, because $\|Tf(x) - Tg(x)\| = 0.5\|f(x) - g(x)\| < \|f(x) - g(x)\|$, but we will go ahead and use the theorem anyway.

- Step one is to verify that the candidate functions f and g satisfy the requirements of Theorem 2.1. They do by assumption in this example, but we will normally have to verify that this is the case (for example, if f and g are both utility functions, we will have to check that they are bounded).

- Step 2 is to check Blackwell's monotonicity condition. Consider a pair of functions such that $f(x) \geq g(x)$ for all $x \in [0,1]$. Then, it must be the case that $0.5f(x) \geq 0.5g(x)$ in the same domain, and thus that $Tf(x) \geq Tg(x)$ for all $x \in [0,1]$. Thus, monotonicity holds.

- Step 3 is to show that discounting holds. $T(f+a)(x) = 0.5(f(x)+a)$ $= 0.5f(x) + 0.5a < Tf(x) + \delta a$ for any $\delta \in (0.5,1)$. Thus, discounting holds.

- We therefore conclude that T is a contraction mapping and that there exists a fixed point function satisfying $f(x)$=$Tf(x)$ for all x. Part (b) tells us how to find this function, as $f^*(x) = \lim_{n\to\infty} T^n f(x) = \lim_{n\to\infty} 0.5^n f(x) = 0$. Thus, the only bounded function $f^*(x)$ that satisfies $f(x)$=$Tf(x)$ for all $x \in [0,1]$ is the zero function, f^*=0 for all x. •

EXAMPLE 2.3. Continuing our consumption example,

$$V(k) = \max_{k'} u\left(f(k) - k'\right) + \beta V\left(k'\right).$$

First, we assume that f and u are such that u is bounded below $\overline{u}$. Then, as V cannot exceed the discounted present value of receiving $\overline{u}$ forever, it follows that $V \leq \overline{u}/(1-\beta) < \infty$, so V is bounded and therefore it has a maximum. Next, we show monotonicity, which states that if there exist two functions $V(k)$ and $Q(k)$ such that $V(k) \geq Q(k)$ for all k, then $TV \geq TQ$ for all k. This is straightforward to establish because of the maximization involved in the problem. Let k'_Q denote the optimal choice of k' when it is the function $Q(k)$ that is being maximized. Then, because we are in fact maximizing $V(k)$, it must be the case that

$$
\begin{aligned}
T(V)(k') &= \max_{k'} u\left(f(k)-k'\right)+\beta V(k') \\
&\geq u\left(f(k)-k'_Q\right)+\beta V\left(k'_Q\right), \qquad \text{because } k'_Q \text{ is not the maximum of } V \\
&\geq u\left(f(k)-k'_Q\right)+\beta Q\left(k'_Q\right), \qquad \text{because } V\left(k'_Q\right) \geq Q\left(k'_Q\right) \\
&\equiv T(Q)(k').
\end{aligned}
$$

Thus, monotonicity holds. Finally, we need to show that discounting holds. This is again easy, in this case because we have discounting in our problem. Let a be some positive constant. Then

$$
\begin{aligned}
T(V+a)(k') &= \max_{k'} u\left(f(k)-k'\right)+\beta\left(V(k')+a\right) \\
&= TV(k')+\beta a.
\end{aligned}
$$

Hence, we have shown that there exists a unique solution to the functional equation. •

It should be apparent from this example that the monotonicity and discounting conditions of Blackwell's theorem can virtually be confirmed by casual inspection of the model. In essence, if you have a dynamic maximization problem with discounting of future returns, then Blackwell's theorem will apply to any problem in which the undiscounted returns are bounded and the state variable can take on any value in a closed bounded interval.

Finally, the theorem also gives us a way to solve the dynamic programming problem, which may be useful in certain settings. Define an arbitrary function $Q(k)$ and apply the contraction repeatedly to obtain $V(k) = \lim_{n\to\infty} T^n Q(k)$. This function so obtained will satisfy the fixed point functional equation $V(k) = TV(k)$ *and* be the unique solution to the dynamic programming problem. But how useful is this solution technique? In practice it often is not very useful, because no one really has enough time to do an infinite amount of algebra! However, for problems with specific functional forms, it can work if you can make a guess of the general form the solution will take. Then, after a few iterations you may see a pattern arising, allowing you to jump the remaining (infinite number of) steps .

A Theorem for Unbounded Returns

In Example 2.2 we just assumed that $u(c)$ was bounded. But what if it is not? In principle, capital can grow without bound and so can utility, and then it is not obvious that the value function will be bounded (which is, after all, what we really care about). The problem is that boundedness is an essential component of Theorem 2.2. Stokey and Lucas (1989) discuss this case in some detail (see their Theorem [4.14]). We provide here a more restrictive theorem that will, for many applications, suffice.

THEOREM 2.3 (A theorem for unbounded returns). *Consider the general dynamic program*

$$TV(x') = \sup_{x'} \left\{ h(x, x') + \beta V(x') \right\}.[6]$$

Assume that the term $\sum_{t=0}^{\infty} \beta^t h\left(x_t, x_{t+1}\right)$ *exists and is finite for any feasible path* $\left\{x_t\right\}_{t=0}^{\infty}$ *given* x_0. *Then, if* T *is a contraction mapping, there is a unique solution to the dynamic optimization problem.*

---

[6] The term ‘sup’ refers to supremum. Until we know that the value function is bounded, we cannot be sure that the maximum value exists. If V is unbounded, we can get arbitrarily close to infinity, but we cannot actually attain it. The term supremum applies to such cases. Once we know that the value function is bounded, we can replace ‘sup’ with ‘max’.

Theorem 2.3 restricts the admissible one-period payoffs to sequences that cannot grow too rapidly relative to the discount factor. As $\sum_{t=0}^{\infty}\beta^t h(x_t, x_{t+1}) < \infty$ by assumption, then $V(x_0) = \max\sum_{t=0}^{\infty}\beta^t h(x_t, x_{t+1})$ is also bounded. Thus, $V(x_0)$ has a maximum and the remainder of the theorem can be applied. The only difficulty with Theorem 2.3 is this: you need to solve the dynamic programming problem to find the path of the state variable, yet you don't know if the solution technique works until you have shown that the path of the one-period payoff function is finite in present value terms. The way out of this impasse can be shown by example:

EXAMPLE 2.4. We return to the consumption example again, this time without assuming that returns are bounded. Assume that $f(k) = k^{\alpha}$ for some α<1 and $u(c)$=ln(c). Then, as $u(c_t) = \ln(c_t) = \ln(k_t^{\alpha} - k_{t+1})$, and $k_{t+1} = k_t^{\alpha} - c_t$ we can make two observations. First, the most rapidly that the capital stock can grow is to choose zero consumption at each point in time. This implies an upper bound to the capital stock given by

$$\ln k_{t+1} \le \alpha \ln k_t \le \alpha^{t+1} \ln k_0 . \tag{2.13}$$

Second, the largest one-period pay-off is found by consuming all the output, so that

$$u(c_t) = \ln(c_t) = \ln(k_t^{\alpha} - k_{t+1}) \le \ln(k_t^{\alpha}) = \alpha \ln k_t . \tag{2.14}$$

So, if we combine the most rapid growth in capital (2.13) with the largest payoff in each period (2.14), we have

$$u(c_t) \le \alpha \ln k_t \le \alpha^{t+1} \ln k_0 .$$

It then follows that

$$\sum_{t=0}^{\infty}\beta^t u(c_t) \le \sum_{t=0}^{\infty}\alpha(\alpha\beta)^t \ln k_0 = \frac{\alpha \ln k_0}{1-\alpha\beta} < \infty \tag{2.15}$$

for any finite k_0. Note that the actual sequence of payoffs must be less than this in present value. We have combined in (2.15) a sequence of utilities from consuming everything with a sequence of capital stocks from consuming nothing, and one cannot have both simultaneously. So any feasible sequence of payoffs must be bounded in present value, and this implies that the value function must also be bounded. The function $V(k)$ therefore has a maximum, and the reminder of the theorem can be applied as before. •

EXERCISE 2.2. *An agent can produce two goods, x and y according to the production functions* $x_t = l_t^x$ *and* $y_t = l_t^y$. *The agent is endowed with one unit of labor time in each period, so* $l_t^x + l_t^y = 1$. *Good x cannot be stored, but good y is indefinitely storable. Production of good x is immediate, so that quantities of x produced in period t are consumed in period t. Production of good y takes time, so that production of y in period t cannot be consumed until period* $t+1$ *at the earliest. Utility in each period is given by* $u_t = \sqrt{c_t^x c_t^y}$, *and the agent maximizes the discounted lifetime utility over the infinite horizon, with a discount factor* $\beta<1$.

(a) *Show that the value function is bounded.* [*Hint*: *construct an infeasible sequence of returns that must exceed any feasible sequence*].

(b) *Now assume that stored y depreciates at the rate* δ *per time period. Show, this time using a more direct approach, that the value function is bounded for this case.*

A corollary to Theorem 2.3 in Stokey and Lucas (1989) also gives us a way to solve for the value function regardless of whether we can show that T is a contraction mapping:

THEOREM 2.4 (Solving models with unbounded returns). *Consider the general dynamic program*

$$TV(x') = \sup_{x'} \left\{ h(x, x') + \beta V(x') \right\},$$

and, for any given x_0, *let* $\hat{V}(x_0) = \max \sum_{t=0}^{\infty} \beta^t h\left(x_t, x_{t+1}\right) < \infty$. *Then, if* $T\hat{V}(x) \leq \hat{V}(x)$ *for all admissible x, and* $V(x) = \lim_{n\to\infty} T^n \hat{V}(x)$ *yields a well-defined finite-valued function, then* $V(x)$ *is the unique solution to the dynamic programming problem.*

Theorem 2.4 says that we can find the solution to the dynamic optimization problem by first defining a function that we know gives a value for any x that is greater than the solution. Then, repeatedly applying the operator T to this function, we see if we converge onto a well-defined finite-valued function. This will work as long as $T\hat{V}(x) \leq \hat{V}(x)$.

This theorem offers a solution technique under different assumptions than we saw for from Theorem 2.2. The earlier theorem stated that if T were a contraction mapping, you could start with *any* value function and iterate to find the unique solution. But doing so would only work if T is a contraction mapping. Theorem 2.4 states that all you need is that T returns a function that is smaller at each value of x. But this will only work if you know you are starting with a function that is larger than $V(x)$ for each value of x.

EXAMPLE 2.5. We will apply Theorem 2.4 to the consumption example with $f(k) = k^{\alpha}$ $u(c)=\ln(c)$. From Example 2.4, define $\hat{V}(k) = \alpha \ln k / (1-\alpha\beta)$, and recall that the operator is defined by

$$T\hat{V}(k') = \max_{0 \leq k' \leq k} \ln\left(k^{\alpha} - k'\right) + \beta\hat{V}(k').$$

$$= \max_{0 \leq k' \leq k} \ln\left(k^{\alpha} - k'\right) + \frac{\alpha\beta \ln k'}{1-\alpha\beta}. \tag{2.16}$$

The maximum of this expression if found upon setting $k' = \alpha\beta k^{\alpha}$, so that on substituting back into (2.16) we get

$$T\hat{V}(k') = \ln(1-\alpha\beta) + \frac{\alpha\beta}{1-\alpha\beta}\ln(\alpha\beta) + \frac{\alpha}{1-\alpha\beta}\ln k.$$

As $\alpha\beta \in (0,1)$, $\ln(1-\alpha\beta) < 0$ and $\ln(\alpha\beta) < 0$, clearly $T\hat{V}(k) < \hat{V}(k)$ as Theorem 2.4 requires. Now, apply the operator a second time: update the equation, replacing k with k', multiply by β, add $\ln\left(k^{\alpha} - k'\right)$, and again take the maximum with respect to k':

$$T^2\hat{V}(k') = \max_{k'} \left\{\ln\left(k^{\alpha} - k'\right) + \beta\left[\ln(1-\alpha\beta) + \frac{\alpha\beta}{1-\alpha\beta}\ln(\alpha\beta) + \frac{\alpha}{1-\alpha\beta}\ln k'\right]\right\},$$

which again gives the optimality condition $k' = \alpha\beta k^{\alpha}$. Substituting back yields

$$T^2\hat{V}(k') = (1+\beta)\left[\ln(1-\alpha\beta) + \frac{\alpha\beta}{1-\alpha\beta}\ln(\alpha\beta)\right] + \frac{\alpha}{1-\alpha\beta}\ln k.$$

Repeating this process n times[7], we get

$$T^n \hat{V}(k') = \left(\sum_{i=0}^{n-1} \beta^i\right)\left[\ln(1-\alpha\beta) + \frac{\alpha\beta}{1-\alpha\beta}\ln(\alpha\beta)\right] + \frac{\alpha}{1-\alpha\beta}\ln k$$
$$= \frac{1-\beta^{n+1}}{1-\beta}\left[\ln(1-\alpha\beta) + \frac{\alpha\beta}{1-\alpha\beta}\ln(\alpha\beta)\right] + \frac{\alpha}{1-\alpha\beta}\ln k \ ,$$

which converges as $n \to \infty$ to

$$\lim_{n\to\infty} T^n \hat{V}(k') \equiv V(k) = \frac{1}{1-\beta}\left[\ln(1-\alpha\beta) + \frac{\alpha\beta}{1-\alpha\beta}\ln(\alpha\beta)\right] + \frac{\alpha}{1-\alpha\beta}\ln k \ . \qquad (2.17)$$

By Theorem 2.4, this represents the solution to our fixed point problem. Of course, now that we know what $V(k)$ is, we can easily solve for the optimal policy upon noting that

$$V(k) = \max_{k'} \ln\left(k^\alpha - k'\right) + \beta\left\{\frac{1}{1-\beta}\left[\ln(1-\alpha\beta) + \frac{\alpha\beta}{1-\alpha\beta}\ln(\alpha\beta)\right] + \frac{\alpha}{1-\alpha\beta}\ln k'\right\},$$

and the first-order condition yields $k' = \alpha\beta k^\alpha$. •

No one said that explicitly solving dynamic programming problems would be easy! There is, however, an alternative way to solve the problem in Example 2.4, using a technique with which you are already familiar: the method of undetermined coefficients.

EXAMPLE 2.6. We will solve the consumption problem from Example 2.5 using the method of undetermined coefficients. We guess a solution of the form

$$V(k) = A + B\ln k$$

for constant A and B to be determined (from equation [2.17] we know this guess is correct). Given this guess, the next task is to derive the optimal policy. The Bellman equation must take the form

$$A + B\ln k = \max_{k'} \left\{\ln\left(k^\alpha - k'\right) + \beta A + \beta B\ln k'\right\}, \qquad (2.18)$$

[7] This is very tedious and its easy to make algebraic mistakes. However, after two or three rounds you will spot a pattern that allows you to write T^n.

so the first-order condition yields

$$k' = \left(\frac{\beta B}{1+\beta B}\right) k^{\alpha} . \tag{2.19}$$

Next, substitute (2.19) into (2.18) to obtain

$$A + B\ln k = \ln\left(1 - \frac{\beta B}{1+\beta B}\right) + \alpha \ln k + \beta A + \beta B \ln\left(\frac{\beta B}{1+\beta B}\right) + \alpha\beta B \ln k .$$

This expression must hold for *any* k. Hence, matching coefficients on $\ln k$, we get

$$B = \frac{\alpha}{1-\alpha\beta} .$$

Matching coefficients on the constants, we get

$$\begin{aligned} A &= \frac{1}{1-\beta}\left[\beta B \ln\left(\frac{\beta B}{1+\beta B}\right) + \ln\left(\frac{1}{1+\beta B}\right)\right] \\ &= \frac{1}{1-\beta}\left[\ln(1-\alpha\beta) + \frac{\alpha\beta}{(1-\alpha\beta)} \ln(\alpha\beta)\right] . \end{aligned}$$

Hence,

$$V(k) = A + B\ln k = \frac{1}{1-\beta}\left[\ln(1-\alpha\beta) + \frac{\alpha\beta}{1-\alpha\beta}\ln(\alpha\beta)\right] + \frac{\alpha}{1-\alpha\beta}\ln k ,$$

which is the same as in (2.17). •

So now we have two ways that we may be able to use to find an explicit solution to a dynamic programming with specific functional forms. The first is to iterate an infinite number of times using the operator $T\hat{V}(k)$. The second is to guess the functional form and use the method of undetermined coefficients. In most cases, neither of these approaches is easy. The infinite iteration approach is algebraically tedious and requires a bit of luck: you need to spot a pattern developing before things gets too messy. The benefit of this approach is that you don't have to guess the form of the solution in advance, although you either have to find a function $\hat{V}(k)$ satisfying $\hat{V}(k) > V(k)$ for all k, or show that T is a contraction mapping. The second method is algebraically easier, but it requires luck (and experience) in guessing the functional form. In Example 2.5, we correctly guessed that $V = A + B \ln k$ only because we had already seen the solution. Without that rather large help, how many guesses would you have tried before hitting on the right functional form?

EXERCISE 2.3 *Consider the two-good production problem of Exercise 2.2, without depreciation of the stored good. Guess that the solution takes the form*

$$V(y) = A\sqrt{B + y}$$

for unknown constants A and B. Find the parameters A and B as a function of the discount factor. Show that the production of x is inversely related to the discount factor.

Conditions for Uniqueness of the Policy Function

The contraction mapping theorem gives conditions for existence and uniqueness of the value function. However, it need not generate a unique policy function. In this section, we provide a theorem that gives the conditions under which the policy function is also unique.

The general dynamic programming problem,

$$V(x) = \max_{x'} \left\{ f(x, x') + \beta V(x') \right\},$$

has the first-order condition

$$f_{x'}(x, x') + \beta V'(x') = 0 . \tag{2.20}$$

For this to constitute a uniquely-defined maximum, we would naturally turn to the second order condition,

$$f_{x'x'}(x, x') + \beta V''(x') < 0 . \tag{2.21}$$

So, one thing we need for uniqueness is clearly that the function f be strictly concave. So the only difficulty is checking the conditions under which $V''(x') < 0$. Conditions under which V is concave are easy to come by. It is also very generally true that we can differentiate V *once* (and hence that our whole solution technique is valid). However, V may not be twice differentiable, so a statement such as (2.21) may not have much meaning.

However, even in this case we can provide conditions for concavity of V, and (2.20) will continue to define a unique maximum.

It turns out that if f is strictly concave then V will also be a strictly concave function with one additional assumption, that the set X from which x and x' are drawn is strictly convex. Before we state the theorem formally and prove it, it might be useful to offer the following reminder of the meaning of concavity of a function and convexity of a set. A function f is strictly concave if, for any valid inputs into the function, $\{x_0, x'_0\}$ and $\{x_1, x'_1\}$, and any third set of inputs $\{x_\theta, x'_\theta\}$ satisfying $x_\theta = \theta x_0 + (1-\theta)x_1$ and $x'_\theta = \theta x'_0 + (1-\theta)x'_1$ for any $\theta \in (0,1)$, then

$$f\left(x_\theta, x'_\theta\right) > \theta f\left(x_0, x'_0\right) + (1-\theta) f\left(x_1, x'_1\right).$$

(Plot this for a concave function with a single argument). A set X is convex if, for any x_0 and x_1 belonging to the set, then x_θ also belongs to the set. Intuitively, the boundary of a convex set is a concave function, and a convex set has no holes. By far the most important and common convex set we will deal with in economic modeling is a bounded interval of real numbers: if two numbers belong in an interval then so does a weighted average of them.

THEOREM 2.5 (Concavity of the value function and uniqueness of the policy function). *Given the general dynamic programming problem* $V(x) = \max_{x'} \left\{ f\left(x, x'\right) + \beta V\left(x'\right) \right\}$, *if f is a strictly concave function, the set X of admissible values for x is convex, and the optimal sequence for* $\{x_t\}_{t=0}^{\infty}$ *involves an interior solution in every period, then* (i) *$V(x)$ is a strictly concave function, and* (ii) *the optimal policy is unique.*

PROOF. Let $x_\theta = \theta x_0 + (1-\theta)x_1$, and assume that x_0 and x_1 are admissible values for the state variable. Then, as the set of admissible values is convex, x_θ is also admissible and we can write

$$\begin{aligned} TV\left(x_\theta\right) &= \left\{ f\left(x_\theta, x'_\theta\right) + \beta V^{/}\left(x'_\theta\right) \right\} \\ &> \theta\left(f\left(x_0, x'_0\right) + \beta V\left(x'_0\right)\right) + (1-\theta)\left(f\left(x_1, x'_1\right) + \beta V\left(x'_1\right)\right) \quad \text{(strict concavity of } f\text{)} \end{aligned}$$

$$= \theta TV(x_0, x'_0) + (1-\theta)TV(x_1, x'_1),$$

so the operator is also strictly concave. This proves part (i). To prove part (ii), note that the sum of two strictly concave functions is also strictly concave. Hence, $f(x,x') + \beta V(x')$ is strictly concave. Moreover (and this will be familiar from standard optimization problems), if a strictly concave function has a maximum, the maximum is unique. Hence, the maximum identified by (2.18) is unique if f is strictly concave. •

Further Properties of the Value Function

Two more useful properties can be established when we have a unique solution to the dynamic programming problem. We state these without proof.

THEOREM 2.6 (Further properties of the value function). (i) *If the one-period payoff function $f(x,x')$ is monotonically increasing [decreasing] in the current value of the state variable, x, then V(x) is also monotonically increasing [decreasing] in x.* (ii) *If there exists a parameter, α, such that $f(x,x';\alpha)$ is monotonically increasing [decreasing] in α, then $V(x;\alpha)$ is also monotonically increasing [decreasing] in α..*

PROOF. We will provide a proof of part (i), which is easy. Let x'_i denote the optimal value of next period's state when today's value is x_i, and consider two values for today's state, $x_1 < x_2$. Then,

$$\begin{aligned} V(x_1) &= f(x_1, x'_1) + \beta V(x'_1) \\ &< f(x_2, x'_1) + \beta V(x'_1). \\ &\le f(x_2, x'_2) + \beta V(x'_2) \\ &= V(x_2). \end{aligned}$$

The first inequality is because $x_1 < x_2$ and f is strictly increasing in x. The second is because the value function obtained on responding optimally to a current value x_2 must exceed any value function obtained by responding suboptimally. •

This section has developed a lot of concepts. Becoming comfortable with their use requires practice and will take time. It will therefore be useful to see the concepts in action. To that end, we close this section with an example that makes use of much of the material developed here.

EXAMPLE 2.7 (*Convex investment costs*). In this example I describe a general investment problem, and then see what I can say about its properties. The example highlights the use of the theorems in making precise statements about quite general problems. In this case, also, checking that the value function is bounded is a little difficult.

The cost of investment, $c(i)$, is strictly increasing, strictly convex and differentiable with $c(0)=0$. The firm produces output according to the production function $f(k)$, with $k \geq 0$ and $f(0)=0$, and where f is differentiable, strictly increasing and strictly concave. The production function further satisfies $\lim_{k \to 0} f^{/}(k) = +\infty$, so we can restrict attention to interior solutions, and $\lim_{k \to \infty} f^{/}(k) = 0$. Capital must be purchased one period ahead of its use, and it depreciates at the constant rate $\delta \in (0,1)$. The price of output is p, the discount factor is $\beta \in (0,1)$, the interest rate is r, and used capital can always be sold at the price q.

The firm's problem is

$$\max_{\{k_{t+1}\}_{t=0}^{\infty}} \quad \sum_{t=0}^{\infty} \beta^t \left\{ pf\left(k_t\right) - c\left(k_{t+1} - (1-\delta)k_t\right) \right\},$$

and the associated Bellman equation is

$$V(k) = \max_{k^{/}} \left\{ pf\left(k\right) - c\left(k' - (1-\delta)k\right) + \beta V\left(k'\right) \right\}. \tag{2.22}$$

To show existence and uniqueness of a solution, I first need to show that (2.22) maps bounded continuous functions into bounded continuous functions. This is a little tricky in this case, and I have to think about the nature of the optimal solution before I actually solve the model. Note that capital bought in period t can be resold in period $t+1$ for a price $(1-\delta)q$. Thus capital will only be accumulated as long as $V^{/}(k) > (1-\delta)q$. I need to show that this inequality cannot hold for any quantities of capital, but I am going to have to do it in a roundabout away. I claim that the following inequality holds:

$$V^{/}(k') < \frac{pf^{/}(k)}{1-\beta}. \tag{2.23}$$

If this claim is true, I can show that $V(k)$ is bounded. I will then characterize the solution to the model assuming it is true, and use this characterization to verify the claim later. Given (2.23), continued investment in capital requires that

$$f^{/}(k) > \frac{(1-\beta)(1-\delta)q}{p}. \tag{2.24}$$

However, $\lim_{k\to\infty} f^{/}(k) = 0$, so there must exist some $\overline{k} < \infty$ such that (2.24) is no longer satisfied. Thus, k is bounded between zero (by assumption) and $\overline{k} < \infty$. As the functions f and c are continuous, boundedness of capital implies that the one-period return is bounded, while discounting then implies that the value function is bounded. Thus, (2.22) maps bounded continuous functions into bounded continuous functions. Then, by Theorem 2.1 (contraction mapping theorem for bounded returns), if the operator defined by (2.22) is a contraction mapping, the function $V(k)$ is uniquely determined.

I can therefore use Theorem 2.2 (Blackwell's Theorem) to verify existence and uniqueness of a fixed point. Monotonicity and discounting are both satisfied for this model, the former because the operator involves maximization, and the latter because we are discounting future returns by the factor $\beta<1$ (these claims can be verified by exactly following the steps in Example 2.2). Consequently, there is a unique value function satisfying (2.22).

The one-period return, $pf(k) - c\left(k' - (1-\delta)k\right)$ is a strictly concave function of $k^{/}$ (because c is convex, $-c$ is concave). Hence, by Theorem 2.5, $V(k')$ is strictly concave, and the policy function obtained from the first order condition defines a unique investment strategy. The first-order condition is given by

$$c^{/}\left(k' - (1-\delta)k\right) = \beta V^{/}(k'). \tag{2.25}$$

Before applying the envelope theorem, I will use the first-order condition to show uniqueness of the policy function directly. As c is differentiable and strictly convex, the left hand side of (2.25) is continuous and strictly increasing in k'. As $V(k')$ is at least once differentiable and strictly concave, the right hand side of (2.25) is continuous and strictly decreasing in k'. Thus, there exists a unique k' satisfying (2.25). Moreover, the left hand side of (2.23) is decreasing in k for any k'. I have now shown that optimal policy, k', is increasing in k. (You might like to draw the graph to verify these arguments).

The envelope theorem gives

$$V'(k) = pf'(k) + (1-\delta)c'\left(k' - (1-\delta)k\right). \tag{2.26}$$

Updating one period,

$$V'(k') = pf'(k') + (1-\delta)c'\left(k'' - (1-\delta)k'\right),$$

and substituting into the first-order condition yields

$$c'\left(k' - (1-\delta)k\right) = \beta pf'(k') + (1-\delta)\beta c'\left(k'' - (1-\delta)k'\right),$$

a second-order difference equation that fully characterizes the unique time path of the optimal investment strategy.

Finally, I need to use these results to verify claim (2.23). Substitute (2.25) into (2.26) to eliminate c':

$$V'(k) = pf'(k) + (1-\delta)\beta V'(k').$$

I do not need to worry about the case where $k > k'$ because if it were ever optimal to reduce the capital stock the desired quantity could be sold immediately at a price q. Hence, restricting attention to the case where, $k \le k'$, concavity of the value function implies that $V'(k') \le V'(k)$. That is,

$$V'(k') \le pf'(k) + (1-\delta)\beta V'(k'),$$

so that

$$V'(k') \le \frac{pf'(k)}{1-(1-\delta)\beta} < \frac{pf'(k)}{1-\beta},$$

as claimed in (2.23).

Although there is relatively little structure to the model, we have been able to establish some important properties. To do so, we made use of Theorems 2.1, 2.2, and 2.5. Having established that the one-period return function was bounded, Theorems 2.3 and 2.4 for unbounded returns were not relevant. We found that the one-period return is increasing in the capital stock. By Theorem 2.6, then, the value of a firm is also increasing in the size of its capital stock. Moreover, we have shown that the value of next period's capital stock is increasing in the size of the stock this period. Thus, there is persistence in firm size – if a firm were to receive a positive shock to its capital stock today, that shock would persist for some time. We have also shown that there is an upper limit to the amount of capital that a firm will accumulate, and hence that there is an upper bound to firm size and value. This finding tells us that, as long as demand is sufficiently large, no one firm would get to dominate any industry exhibiting diminishing returns and convex adjustment costs. •

3. Dynamic Programming and Optimal Control [8]

Although dynamic programming most often is carried out in discrete-time settings, it can also be used in continuous time. In this section we show the equivalence of dynamic programming and optimal control solutions to continuous-time, deterministic, dynamic optimization problems coincide.

Consider the following familiar continuous-time investment problem for a firm:

$$\max_{x(t)} \int_0^T u\big(k(t), x(t), t\big)\, dt\,,$$

subject to

$$\dot{k}(t) = f\big(k(t), x(t), t\big)\,, \qquad k(0) = k_0\,. \tag{3.1}$$

Define $V\big(t_0, k(t_0)\big)$ as the best value for the firm that can be attained at time t_0 given that the capital stock at time t_0 is $k(t_0)$. This function is defined for all $t_0 \in [0, T]$ and any feasible $k(t_0)$. That is,

$$V\big(t_0, k(t_0)\big) = \max_x \int_{t_0}^T u\big(k(t), x(t), t\big)\, dt\,, \tag{3.2}$$

[8] This section can be omitted without loss of continuity.

subject to (3.1). Note that $V\left(T, k(t_0)\right) = 0$ by definition.

Break up the integral in (3.2) into two parts, one covering the short interval $\left[t_0, t_0 + \Delta t\right]$, and the other covering the interval $\left(t_0 + \Delta t, T\right]$:

$$V\left(t_0, k(t_0)\right) = \max_x \left(\int_{t_0}^{t_0+\Delta t} u\left(k(t), x(t), t\right) dt + V\left(t_0, k(t_0)\right) + \int_{t_0+\Delta t}^{T} u\left(k(t), x(t), t\right) dt \right),$$

By Bellman's Principle of Optimality, the investment path $x(t),\ t \in \left(t_0 + \Delta t, T\right]$, must be optimal for the problem beginning at time $t_0 + \Delta t$. That is,

$$V\left(t_0, k(t_0)\right) = \max_{x(t), t_0 \le t \le \Delta t} \left\{ \int_{t_0}^{t_0+\Delta t} u\left(k(t), x(t), t\right) dt + \max_{x(t), t_0 + \Delta t < t \le T} \left(\int_{t_0+\Delta t}^{T} u\left(k(t), x(t), t\right) dt \right) \right\},$$

subject to (3.1). Put another way,

$$V\left(t_0, k(t_0)\right) = \max_{x(t), t_0 \le t \le \Delta t} \left\{ \int_{t_0}^{t_0+\Delta t} u\left(k(t), x(t), t\right) dt + V\left(t_0 + \Delta t, k(t_0) + \Delta k\right) \right\}, \qquad (3.3)$$

which states that the value of the optimal policy is equal to the return to choosing an optimal policy over the interval $\left[t_0, t_0 + \Delta t\right]$ plus the return from continuing optimally thereafter.

As Δt is assumed to be small, then the following approximations are reasonable (as they will be exact in a moment when we let $\Delta t \to \infty$).

$$\int_{t_0}^{t_0+\Delta t} u\left(k(t), x(t), t\right) dt \approx u\left(k(t_0), x(t_0), t\right) \Delta t\ ,$$

$$\frac{1}{\Delta t} \int_{t_0}^{t_0+\Delta t} x(t) dt \approx x_0\ .$$

That is, as Δt is a small interval, then we can approximate the two integrals by assuming that u and x are constant over the interval.

Now, in discrete time modeling we would let $\Delta t = 1$, and assume that within each period of length 1, the chosen policy must be constant. Doing so yields

$$V\left(t, k(t)\right) = \max_{x(t)} \left\{ u\left(k(t), x(t), t\right) + V\left(t+1, k(t+1)\right) \right\}, \qquad (3.4)$$

where u now measures the payoff during a single period from choosing investment $x(t)$ and beginning with capital stock $k(t)$. This is the key functional equation for discrete time that we have already seen. But what I want to do right now is to think about the continuous-time problem and relate it to optimal control. In fact, we can go from (3.3) to optimal control with the simple assumption that $V\left(t, k(t)\right)$ is twice differentiable. The assumption allows us to take a Taylor expansion of (3.3) around $t{=}t_0$:

$$V\left(t_0, k(t_0)\right) \approx \max_{x(t_0)} \left\{ u\left(k(t_0), x(t_0), t_0\right) \Delta t + V\left(t_0, k(t_0)\right) + V_t\left(t_0, k(t_0)\right) \Delta t + V_k\left(t_0, k(t_0)\right) \Delta k \right\}.$$

Subtract $V\left(t_0, k(t_0)\right)$ from both sides and divide through by Δt:

$$0 = \max_{x(t_0)} \left\{ u\left(k(t_0), x(t_0), t_0\right) + V_t\left(t_0, k(t_0)\right) + V_k\left(t_0, k(t_0)\right) \frac{\Delta k}{\Delta t} \right\}.$$

Finally, we let $\Delta t \to 0$, yielding

$$0 = \max_{x(t)} \left\{ u\left(k(t), x(t), t\right) + V_t\left(t, k(t)\right) + V_k\left(t, k(t)\right) \dot{k}(t) \right\}, \tag{3.5}$$

where we can now, without inaccuracy, drop the zero subscript on current time. Let $\lambda(t)$ denote the costate variable from optimal control. We know that $\lambda(t)$ has the meansing of the marginal value of the state variable, and hence that $\lambda(t){=}V_k(k(t),t)$. Using this fact in (3.5), we have

$$\begin{aligned} -V_t\left(t, k(t)\right) &= \max_{x(t)} \left\{ u\left(k(t), x(t), t\right) + \lambda(t) \dot{k}(t) \right\} \\ &= \max_{x(t)} \left\{ u\left(k(t), x(t), t\right) + \lambda(t) f\left(k(t), x(t), t\right) \right\}. \end{aligned} \tag{3.6}$$

Equation (3.6) is known as the **Hamilton-Jacobi-Bellman** equation, and represents the fundamental *partial* differential equation obeyed by the optimal value function. Note, that the right hand side of (3.6) must be maximized by choice of – in the language of optimal control – the control variable, $x(t)$. But the right hand side is simply the Hamiltonian of optimal control, and its first-order condition is

$$u_x + \lambda f_x = 0\,. \tag{3.7}$$

The optimality condition for $\lambda(t)$ is also readily derived. Equation (3.6) must be true even if $k(t)$ is modified. Thus, we can differentiate (3.6) with respect to $k(t)$ to get

$$-V_{tk} = u_k + \lambda_k f + \lambda f_k \,. \tag{3.8}$$

where the term $-V_{tk} = -V_{kt} = -\dot{\lambda}(t)$. Now, noting that the total derivative of $V_k(t, k(t))$ with respect to time is

$$\dot{\lambda}(t) = \frac{dV_k\left(t, k(t)\right)}{dt} = V_{kt} + V_{kk}\dot{k}(t) = V_{kt} + V_{kk} f = V_{kt} + \lambda_k f \,, \tag{3.9}$$

we can combine (3.8) and (3.9) to get

$$-\dot{\lambda} + \lambda_k f = u_k + \lambda_k f + \lambda f_k \,,$$

or

$$u_k + \lambda f_k = -\dot{\lambda} \,.$$

Hence, if $V(k,t)$ is twice differentiable, then optimal control and dynamic programming give equivalent optimality conditions.

4. Stochastic Dynamic Programming

One of the most attractive features of dynamic programming is the relative ease with which stochastic elements can be incorporated. We therefore now extend the methods of section 2 to incorporate stochastic features into our models. The extension is in principle straightforward: one adds a judiciously-placed random variable such that at time t past realizations are known but future realizations are not. Thus, the current value function depends upon the distribution of future values of the random shock, and the way in which the shock affects future returns. This uncertainly is handled with the expectations operator.

Consider, for example, the cake-eating problem of Exercise 1.1. In this problem, the natural source of uncertainty concerns random variations in the agent's preferences. For example, it may be that utility in each period is given by

$$v\left(c_t\right) = z_t u\left(c_t\right).$$

where z_t is a random variable. The correct formulation for the Bellman equation in this problem depends upon what we assume is known about the random process. One assumption is that the realized value of z_t is known when period-t consumption is chosen, in which case we write

$$V\left(x_t\right)=\max_{c_t}\left\{z_t u\left(c_t\right)+\beta E_t\left[V\left(x_{t+1}\right)\right]\right\},$$

(recall that x is the size of cake remaining). An alternative assumption is that the taste shock for the current period is not known at the time the consumption decision is being made, in which case we write

$$V\left(x_t\right)=\max_{c_t}\left\{E_t\left[z_t u\left(c_t\right)+\beta V\left(x_{t+1}\right)\right]\right\}.$$

and this time we cannot take the expectations operator inside to the second term. In both cases, E_t denotes the expectation of the value function conditional upon information that is known when period-t decisions are made. It is up to the modeler to be clear about what belongs in the period-t information set, because different assumptions may lead to drastically different behavior.

The cake-eating example adds shocks to the returns in each period but, conditional on the consumption choice, the evolution of the state variable is deterministic. A second common way to introduce stochastic elements is to suppose that the payoff function is deterministic once the value of the state variable is given, but the evolution of the state variable is subject to random shocks. Consider a stochastic version of the simple infinite-horizon consumption problem with capital accumulation (equation [2.1]):

$$\max\ E_0\sum_{t=1}^{\infty}\beta^t u\left(z_t f\left(k_t\right)-k_{t+1}\right).$$

Here, output is subject to random productivity shocks, z_t, so that $k_{t+1}=z_t f\left(k_t\right)-c_t$. However, once c_t and k_t are given, the one-period return is fixed. We assume here that z_t is known at the time c_t is chosen, so that next period's capital stock is also known. However, next period's value function remains stochastic because it will depend upon the realization of z_{t+1}.Thus, the Bellman equation is:

$$\begin{aligned}V\left(k\right)&=\max_{c}\left\{u(c)+\beta E\left[V\left(k'\right)\right]\right\}\\&=\max_{c}\left\{u(c)+\beta E\left[V\left(zf(k)-c\right)\right]\right\}.\end{aligned}\tag{4.1}$$

Because k' can be controlled directly once z_t is known, we could make a substitution of k' for c:

$$V(k) = \max_{k'} \left\{ u(zf(k) - k') + \beta E\left[V(k')\right] \right\}.$$

This substitution would not be possible if z were not known when c is chosen.

We have now written two stochastic dynamic programs, but we have made no attempt to solve them. It turns out that the solution principle is no different for stochastic problems than it is for deterministic problems although, as we will see, the expectations operator often makes life rather more difficult. Before thinking about solutions, however, we need to delimit the sorts of problems we are prepared to tackle. and we need to consider what assumptions are necessary to ensure that the dynamic programming approach will yield meaningful and unique solutions.

Markov Decision Problems

We will restrict our attention as always to dynamic optimization problems in which the stream of payoffs enter additively. More important, we will restrict the types of stochastic processes we consider to a special class of stochastic processes known as **Markov Processes**.

DEFINITION (Markov Process). *A random process whose future probabilities are determined by its most recent value. A stochastic process $x(t)$ is* Markov *if for every n and $t_1 < t_2 < \cdots < t_n$, we have*

$$\Pr\left\{x(t_n) \le x_n \mid x(t_{n-1}), \ldots, x(t_1)\right\} = \Pr\left\{x(t_n) \le x_n \mid x(t_{n-1})\right\}.$$

If $x(t)$ takes on only discrete values, then such a process is called a Markov chain. *If $x(t)$ is a continuous random variable and is Markov, then the process is known as a* Markov sequence. *The n^{th} element in the sequence has a conditional distribution satisfying*

$$F_n\left(x \mid x(t_{n-1}), \ldots, x(t_1)\right) = F_n\left(x \mid x(t_{n-1})\right).$$

Markov processes have the property that the current value of the random process is all you need to know to characterize the distribution of the next element in the sequence. History does not matter in the sense that, if the current value is z_t, it does not matter how you got there.

A dynamic optimization problem in which stochastic elements are Markov processes and in which the stream of payoffs enter additively, is known as a **Markov decision process.** The enormously simplifying feature of Markov decision processes is that the value function and the optimal policy can be expressed as a function of the most recently observed random variables and the current value of the state variable alone.

Necessary Assumptions

Recall from Section 2 that our main concern for the validity of the dynamic programming approach is that the value function is a continuous function mapping a bounded closed interval into a bounded closed interval. We can ensure that the value function has the necessary properties if (i) the one-period return function is continuous and is bounded, either in every period or in present value terms, and (ii) the state variable comes from a bounded closed interval $[a,b]$. The task for stochastic dynamic programming is to ensure the necessary properties of the value function continue to hold when we take expectations over the exogenous random variable. Fortunately, this is usually the case.

Consider first the case in which z can only take on a finite number of discrete values, z_i with associated probabilities $p_i(x,y,z)$. which may depend on the values of the state and control variables, and the exogenous shock. Then, if the value function is continuous and bounded in an interval for any feasible x and z in this interval, its expectation $\sum_i p_i(x,y,z)V(x',z_i)$ must also lie in the same interval. This can be verified directly upon noting that the summation term that gives the expectation is simply a weighted average of values that lie in the interval.

It is a little harder to verify this for property for continuous distributions. The usual tactic is simply to *assume* that the property holds. That is, if $F\left(z' \mid z\right)$ is the distribution of z' conditional on z, and $V(x',z')$ is a continuous and bounded function taking values in a given interval, then it is assumed that $E\left[V(x',z') \mid z\right] = \int V(x',z')dF\left(z' \mid z\right)$ is also a continuous function taking values in the same interval. When the conditional distribution function $F\left(z' \mid z\right)$ has this property, we say that it has the **Feller property**. In practice, what this means is that the distribution of z must be stable and continuous: small changes in z should only lead to small changes in $F(z' \mid z)$ for any z'. This is not a restrictive assumption. In fact it is difficult to come up with an economically-meaningful example that does not have the Feller property

So, let us put the Feller property out of mind except to note that, if it holds, then our usual solution tactics work, and the properties of the value function already obtained continue to hold. In particular:

- The contraction mapping theorem continues to hold (Theorem 2.1);
- We can continue to use Blackwell's Theorem to see whether the Bellman equation is a contraction mapping (Theorem 2.2);
- If the one-period return is strictly concave and increasing in x, then so is the value function (Theorems 2.5 and 2.6).

The only new feature is that we now have a random shock z in the picture. If the one-period return is strictly increasing in z, can we say the same for $V(x,z)$? The answer is yes, if the conditional expectation, $F(z' \mid z)$ also satisfies a certain monotonicity property:

THEOREM 4.1 (Monotonicity in the exogenous random variable). *Consider the value function*

$$V(x,z) = f(x, x'^{*}(x,z)z) + \beta \int V(x'^{*}, z')dF(z' \mid z),$$

where the asterisk on x' denotes that we have already substituted in the optimal value for x'. If $f(\bullet, z)$ is strictly increasing in z for any x, and $F\left(z' \mid z\right)$ is nonincreasing in z for any z', then $V(\bullet, z)$ is strictly increasing in z.

PROOF. Consider the integral term $\int_{-\infty}^{\infty} V(x', z')dF(z' \mid z)$. Integrate by parts, to obtain

$$V(x' z')F(x', z' \mid z)\Big|_{-\infty}^{\infty} - \int_{-\infty}^{\infty} V_{z'}(x', z')F(z' \mid z)dz'$$

$$= V(x', \infty) - \int_{-\infty}^{\infty} V_{z'}(x', z')F(z' \mid z)dz'$$

Differentiating with respect to z yields

$$-\int_{-\infty}^{\infty} V_{z'}(x', z')F_z(z' \mid z)dz'.$$

As $F_z < 0$, the integral term is increasing in z as long as the value function V is increasing in z'. Now, if this last condition is true then, with f increasing in z by assumption, the left hand size is increasing in z. •

The assumption that $F(z' \mid z)$ is decreasing in z for any z' simply says that the probability that z' is less than a certain number is not increased when z increases. Put another way, shocks to z are assumed not to exhibit negative serial correlation. In the integral, this assumption puts heavier weight onto large values of z', which (if V is increasing in z) raises the value of the integral term. This is consistent, in turn, with the left hand side of the value function being increasing in z.

EXAMPLE 4.1 (*Cake-eating with taste shocks*). Consider the Bellman equation for the stochastic cake-eating example,

$$V(x, z) = \max_c \left\{ zu(c) + \beta E\left[V(x, z)\right]\right\}$$

$$= \max_c \left\{ zu(c) + \beta \int_0^{\overline{z}} V(x', z')\, dF(z' \mid z)\right\}.$$

The utility function is assumed to be strictly increasing, continuous and strictly concave. It is assumed that $z \in [0, \overline{z}]$ and c cannot exceed the size of the cake. These assumptions imply that the one-period return cannot exceed $\overline{z}u(x)$. We further assume that $F(z' \mid z)$

has the Feller property and is nonincreasing in z for any z'. Because the evolution of the size of the cake is not stochastic, we can assume that the agent chooses x' directly. That is

$$V(x,z) = \max_{x'} \left\{ zu(x - x') + \beta \int_0^{\bar{z}} V(x', z')\, dF(z' \mid z) \right\}$$

The first-order condition is

$$zu'(x - x') = \beta \int_0^{\bar{z}} V_{x'}(x', z')\, dF(z' \mid z).$$

The envelope theorem gives

$$V_x(x,z) = zu'(x - x').$$

Updating one period and substituting into the first-order condition gives

$$zu'(x - x') = \beta \int_0^{\bar{z}} z' u'(x' - x'') dF(z' \mid z),$$

or

$$zu'(c) = \beta \int_0^{\bar{z}} z' u'(c') dF(z' \mid z),$$

which says that the marginal utility of consumption today must equal the discounted *expected* marginal utility of consumption tomorrow. •

EXAMPLE 4.2 (*Consumption with stochastic productivity*). The Bellman equation for the capital problem introduced at the beginning of this chapter is $V(k) = \max_c \left\{ u(c) + \beta E\left[V(k')\right] \right\}$, where $k' = zf(k) - c$, and z is not known when c is chosen. In this example, we will assume that z is known, so that, given z and c, k' is deterministic. We write the Bellman equation as

$$V(k,z) = \max_{k'} \left[u\left(zf(k) - k'\right) + \beta EV\left(k' \mid z\right) \right]$$

$$= \max_{k'} \left\{ u\left(zf(k) - k'\right) + \beta \int V(k', z') dF(z' \mid z) \right\},$$

The first-order condition is

$$-u'(zf(k) - k') + \beta \int V_1(k', z') dF(z' \mid z) = 0 ,$$

where the subscript denotes the derivate with respect to the first argument. The envelope theorem gives

$$V_1(k, z) = u'(zf(k) - k') zf'(k) .$$

Update by one period,

$$V_1(k', z') = u'(z' f(k') - k'') z' f'(k') ,$$

and substitute into the first-order condition:

$$u'(zf(k) - k') = \beta \int u'(z' f(k') - k'') z' f'(k') dF(z' \mid z) ,$$

which can also be written as

$$u'(c) = \beta f'(k') \int u'(c') z' dF(z' \mid z) .$$

Because $f'(k')$ is known, it can be taken outside of the expectation operator. The marginal utility of consumption today must equal the expected discounted present value of the product of the marginal utility of consumption tomorrow and the marginal rate of transformation between consumption and capital tomorrow.

It is not usually possible to obtain an explicit solution for this model. But consider the special case where $u(c) = \ln(c)$ and $k' = zk^{\alpha} - c$. Then the first-order condition is

$$\frac{1}{c} = \alpha\beta (k')^{\alpha - 1} \int \frac{z'}{c'} dF(z' \mid z)$$

We guess a solution to this equation of the form $c = \phi z k^{\alpha}$, for a value of ϕ to be determined. If the guess is correct, then the first-order condition satisfies

$$\frac{1}{\phi z k^{\alpha}} = \alpha\beta (k')^{\alpha - 1} \int \frac{z'}{\phi z' (k')^{\alpha}} dF(z' \mid z)$$

$$= \frac{\alpha\beta}{\phi k'} \int dF(z' \mid z)$$

$$= \frac{\alpha\beta}{\phi k'},$$

which solves for $k' = \alpha\beta z k^{\alpha}$. Now, using the transition equation $k' = zk^{\alpha} - c$ along with the guess $c = \phi z k^{\alpha}$ and the provisional solution $k' = \alpha\beta z k^{\alpha}$, we obtain $zk^{\alpha} - fzk^{\alpha} = \alpha\beta z k^{\alpha}$, which solves for $\phi = 1 - \alpha\beta$. Hence the optimal policy is $c = (1-\alpha\beta)zk^{\alpha}$. •

5. Approximations, Algebraic and Numerical

Most dynamic programming problems cannot be solved explicitly. Although we can often obtain a number of interesting properties of the solution we would usually like to have a deeper characterization of the model. There are two approaches one could take. One is to approximate the functions under analysis by means of Taylor expansions. The other is to numerically solve the model for particular parameter values. This section provides a brief introduction to these methods.

This section to be written.

Further Reading

References

Bellman, Richard, and Stuart E. Dreyfus (1962): *Applied Dynamic Programming*. Princeton, NJ: Princeton University Press.

Benveniste, Lawrence M., and Jose Scheinkman (1979): "On the differentiability of the value function in dynamic models of economics." *Econometrica*, **47**:727-732.

Eekhout, Jan, and Boyan Jovanovic (2002): "Knowledge spillovers and inequality." *American Economic Review*, **92**(5):1290-1307.

Stokey, Nancy L., and Robert E. Lucas Jr. (1989): *Recursive Methods in Economic Dynamics*, Cambridge, MA: Harvard University Press.