Lectures Notes on Deterministic Dynamic Programming*

Craig Burnside†

October 2006

1 The Neoclassical Growth Model

1.1 An Infinite Horizon Social Planning Problem

Consider a model in which there is a large fixed number, H, of identical households. The total population is L_t, so each household has L_t/H members. Each household has the following utility function

$$U = \sum_{t=0}^{\infty} \beta^t u(c_t) \frac{L_t}{H}, \tag{1}$$

where $0 < \beta < 1$, c_t is the consumption of each member of the household at time t, and $u(c)$, the instantaneous utility function, has the following properties: $u'(c) > 0$, $u''(c) < 0$ and $\lim_{c \to 0} u'(c) = \infty$.[1] From now on I will assume that $L_t = H$ for all t.

The economy's total production of output, Y_t, is given by

$$Y_t = F(K_t, L_t) \tag{2}$$

where K_t is the capital stock. Assuming that F is constant returns to scale (CRTS), $F(K_t, L_t) = L_t f(k_t)$ where $k_t \equiv K_t/L_t$ and $f(k_t) \equiv F(k_t, 1)$. We assume that $f(0) = 0$, $f'(k) > 0$, $f''(k) < 0$, $\lim_{k \to 0} f'(k) = \infty$ and $\lim_{k \to \infty} f'(k) = 0$.

Output can either been consumed or invested in the form of new capital. The aggregate resource constraint is

$$C_t + K_{t+1} - (1 - \delta)K_t = Y_t, \text{ for } t \geq 0, \tag{3}$$

where $C_t = L_t c_t$ is total consumption and $0 < \delta < 1$ is the rate of depreciation.

---

†Department of Economics, Duke University.

[1] Sometimes I will refer to the parameter $\rho = \beta^{-1} - 1$ as the rate of time preference. I assume that the household shares its consumption allocation equally among its members.

Imagine that there is a social planner who maximizes the utility of a representative household, (1), subject to the technology, (2) and the aggregate resource constraint, (3). Because the population is constant and each household has one member, this problem is equivalent to the social planner choosing $\{c_t, k_{t+1}\}_{t=0}^{\infty}$ to maximize

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

subject to $k_{t+1} \geq 0$, $c_t \geq 0$, for $t \geq 0$, $k_0 > 0$ given and

$$c_t + k_{t+1} = f(k_t) + (1-\delta)k_t, \text{ for } t \geq 0.$$

I will define $g(k) \equiv f(k) + (1-\delta)k$ and rewrite the social planning problem as

$$\max_{\{k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u[g(k_t) - k_{t+1}] \tag{4}$$

subject to $0 \leq k_{t+1} \leq g(k_t)$, for $t \geq 0$, $k_0 > 0$ given.[2] Assuming an interior solution, the first order condition for k_{t+1} is

$$-\beta^t u'[g(k_t) - k_{t+1}] + \beta^{t+1} u'[g(k_{t+1}) - k_{t+2}] g'(k_{t+1}) = 0.$$

Rearranged, this gives us the familiar Euler equation

$$u'[g(k_t) - k_{t+1}] = \beta u'[g(k_{t+1}) - k_{t+2}] g'(k_{t+1}), \quad t = 0, 1, 2 \ldots \tag{5}$$

Sometimes it's helpful to substitute back in the fact that $c_t = g(k_t) - k_{t+1}$ to write (5) as

$$u'(c_t) = \beta u'(c_{t+1}) g'(k_{t+1}), \quad t = 0, 1, 2 \ldots$$

A special aspect of (4) is that the planner's problem is infinite dimensional. That is, he chooses the optimal value of an infinite sequence, $\{k_{t+1}\}_{t=0}^{\infty}$. Dynamic programming turns out to be an ideal tool for dealing with the theoretical issues this raises. But as we will see, dynamic programming can also be useful in solving finite dimensional problems, because of its recursive structure.

1.2 A Finite Horizon Analog

Consider the analogous finite horizon problem

$$\max_{\{k_{t+1}\}_{t=0}^{T}} \sum_{t=0}^{T} \beta^t u[g(k_t) - k_{t+1}] \tag{6}$$

[2] Unfortunately, my notation is similar but different to that used by Stokey and Lucas who use the notation $f(k)$ to represent $F(k,1) + (1-\delta)k$.

subject to $0 \leq k_{t+1} \leq g(k_t)$, for $0 \leq t \leq T$, $k_0 > 0$ given. It's easy to characterize the solution of this problem using the Kuhn-Tucker conditions because (i) it's a standard optimization problem with a finite number of choice variables, (ii) the objective function for the problem is concave in the vector of choice variables $(k_1, k_2, \ldots, k_{T+1})$ and (iii) the constraints are quasi-convex in that vector.

To see why it's not immediately obvious how to extend the finite horizon problem to the infinite horizon consider the optimality conditions that emerge from the finite horizon problem. Since k_{T+1} only appears in the $[g(k_T) - k_{T+1}]$ term, it is clear that the optimal solution for k_{T+1} is $k_{T+1} = 0$. The remaining choice variables $(k_1, k_2, \ldots, k_T)$ are determined according to the familiar Euler equations:

$$u'[g(k_t) - k_{t+1}] = \beta u'[g(k_{t+1}) - k_{t+2}]g'(k_{t+1}), \quad t = 0, 1, \ldots, T-1. \tag{7}$$

(7) represents T equations in the T unknowns, k_1, k_2, $\ldots$, k_T. The variables k_0 and k_{T+1} also appear in these equations but k_0 is given and we have already shown that $k_{T+1} = 0$.

In the infinite horizon problem we have the same Euler equations, but an infinite number of them. We lose the end condition $k_{T+1} = 0$, and it's not obvious what it's replaced by, if anything. Dynamic programming is an approach to optimization that deals with these issues. I will illustrate the approach using the finite horizon problem. Then I will show how it is used for infinite horizon problems.

1.3 Solving the Finite Horizon Problem Recursively

Dynamic programming involves taking an entirely different approach to solving the planner's problem. Rather than getting the full set of Kuhn-Tucker conditions and trying to solve T equations in T unknowns, we break the optimization problem up into a recursive sequence of optimization problems.

In the finite horizon problem, (6), we are asked to solve the planner's problem at date 0. Suppose, instead, we solve the planner's problem at date T. This is easy to do, since the planner's objective at date T will just be

$$\max_{k_{T+1}} u[g(k_T) - k_{T+1}]$$

subject to $0 \leq k_{T+1} \leq g(k_T)$, $k_T > 0$ given. The solution is $k_{T+1} = h_{T+1}(k_T) = 0$ for all possible k_T. The objective function at the optimum, whose value depends on k_T, is

$$v_T(k_T) = u[g(k_T) - h_{T+1}(k_T)] = u[g(k_T)]. \tag{8}$$

Now consider the planner's problem at date $T-1$, which is

$$\max_{k_T, k_{T+1}} u[g(k_{T-1}) - k_T] + \beta u[g(k_T) - k_{T+1}] \tag{9}$$

subject to $0 \leq k_T \leq g(k_{T-1})$, $0 \leq k_{T+1} \leq g(k_T)$, $k_{T-1} > 0$ given. The idea is to break this problem of choosing two variables into two problems in which only one variable is chosen. In particular, we will instead solve

$$\max_{k_T} \left\{ u[g(k_{T-1}) - k_T] + \beta \max_{k_{T+1}} u[g(k_T) - k_{T+1}] \right\}. \tag{10}$$

where the inner optimization is done subject to subject to $0 \leq k_{T+1} \leq g(k_T)$, $k_T > 0$ given, and the outer optimization is done subject to $0 \leq k_T \leq g(k_{T-1})$, $k_{T-1} > 0$ given. The two optimization problems, (9) and (10), are clearly equivalent. Since $v_T(k_T) = \max_{k_{T+1}} u[g(k_T) - k_{T+1}]$, we can rewrite (10) as

$$v_{T-1}(k_{T-1}) = \max_{k_T} u[g(k_{T-1}) - k_T] + \beta v_T(k_T), \tag{11}$$

subject to $0 \leq k_T \leq g(k_{T-1})$, $k_{T-1} > 0$ given.

Now consider the planner's problem at date $T - 2$, which is

$$\max_{k_{T-1}, k_T, k_{T+1}} u[g(k_{T-2}) - k_{T-1}] + \beta u[g(k_{T-1}) - k_T] + \beta^2 u[g(k_T) - k_{T+1}],$$

subject to $0 \leq k_t \leq g(k_{t-1})$, $t = T-1, T, T+1$, and $k_{T-2} > 0$ given. It is immediately clear that this too can be written recursively as

$$v_{T-2}(k_{T-2}) = \max_{k_{T-1}} u[g(k_{T-2}) - k_{T-1}] + \beta v_{T-1}(k_{T-1}). \tag{12}$$

In fact, we can write the time s problem recursively in terms of the time $s+1$ problem as

$$v_s(k_s) = \max_{k_{s+1}} u[g(k_s) - k_{s+1}] + \beta v_{s+1}(k_{s+1}). \tag{13}$$

Proceeding in this fashion we would stop when we came to $s = 0$ because we would then have solved the time 0 problem.

1.4 Solving the Infinite Horizon Problem Recursively

It would be useful if the optimization problem (4) could be characterized recursively using an equation such as (13). Recall that the notation $v_s(k_s)$ in the finite horizon problem was just

$$v_s(k_s) = \max_{\{k_{t+1}\}_{t=s}^{T}} \sum_{t=s}^{T} \beta^{t-s} u[g(k_t) - k_{t+1}] \tag{14}$$

subject to $0 \leq k_{t+1} \leq g(k_t)$, for $s \leq t \leq T$, $k_s > 0$ given. In the infinite horizon problem this suggests that we use the notation

$$v_s^{\infty}(k_s) = \max_{\{k_{t+1}\}_{t=s}^{\infty}} \sum_{t=s}^{\infty} \beta^{t-s} u[g(k_t) - k_{t+1}]. \tag{15}$$

Suppose that the equivalent of (13) holds for the infinite horizon problem. Then

$$v_s^\infty(k_s) = \max_{k_{s+1}} u[g(k_s) - k_{s+1}] + \beta v_{s+1}^\infty(k_{s+1}). \tag{16}$$

Notice, however, that unlike in the finite horizon problem the *functions* v_s^∞ and v_{s+1}^∞ must be the same.[3]

Dropping subscripts and the ∞ notation, we have

$$v(k) = \max_{k'} u[g(k) - k'] + \beta v(k') \tag{17}$$

subject to $0 \le k' \le g(k)$, with k given. (17) is the *Bellman equation*. There are two subtleties we will deal with later:

(i) we have not shown that a v satisfying (17) exists,

(ii) we have not shown that such a v actually gives us the correct value of the planner's objective at the optimum.

1.5 Optimality Conditions in the Recursive Approach

To get the optimality conditions that coincide with (17) we will defer some details until later. In particular we will not prove here that the value function is differentiable, nor that the solution k' always lies in the interior of the set $[0, g(k)]$. For the moment we will simply assume that these statements are true.

When there is an interior solution, and v is differentiable the first-order condition for the maximization problem in (17) is

$$u'[g(k) - k'] = \beta v'(k'). \tag{18}$$

In and of itself this doesn't look too useful because we don't yet know the shape of the value function. However, when there is an interior solution for k' and v is differentiable, we also have the envelope condition:

$$v'(k) = u'[g(k) - k']g'(k). \tag{19}$$

[3] To see why notice that we would write

$$v_{s+1}^\infty(k_{s+1}) = \max_{\{k_{t+1}\}_{t=s+1}^\infty} \sum_{t=s+1}^\infty \beta^{t-s-1} u[g(k_t) - k_{t+1}]$$

Now make a change of variables, defining $j = s + 1$. Notice that we end up with

$$v_j^\infty(k_j) = \max_{\{k_{t+1}\}_{t=j}^\infty} \sum_{t=j}^\infty \beta^{t-j} u[g(k_t) - k_{t+1}].$$

But this is the same as (15) with different letters.

A Digression on the Envelope Theorem. You should recall the envelope theorem from standard 1st-year micro problems. The basic principle at work is as follows. Let $k' = h(k)$ be the optimal policy function. Substituting this into (17) we see that

$$v(k) = u[g(k) - h(k)] + \beta v[h(k)].$$

Hence

$$\begin{aligned} v'(k) &= u'[g(k) - h(k)][g'(k) - h'(k)] + \beta v'[h(k)]h'(k) \\ &= u'[g(k) - h(k)]g'(k) + \{\beta v'[h(k)] - u'[g(k) - h(k)]\}\, h'(k) \end{aligned}$$

Notice however, that we can use (18) to reduce this to

$$v'(k) = u'[g(k) - h(k)]g'(k),$$

which is the same as (19).

Of course, if we combine (18) and (19) we get the familiar Euler equation

$$u'[g(k) - k'] = \beta u'[g(k') - k'']g'(k'). \tag{20}$$

Rewriting this with date subscripts you get back (5).

A Digression on the Ljungqvist-Sargent Approach Ljungqvist and Sargent and the text I call "New Sargent" both use slightly different language—I'll call it the language of optimal control—in describing the Bellman equation and the optimality conditions. They write the problem in terms of both the *state variable*, k, and the *control variable*, c. In a sense, all this means is that they do not substitute out the resource constraint, which happens to describe the law of motion of k. In particular for the growth model they would write

$$v(k) = \max_c u\,(c) + \beta v(k')$$

subject to $k' = g(k) - c$.[4] Using the resource constraint they would then substitute the future state out of the Bellman equation to get

$$v(k) = \max_c u\,(c) + \beta v\,[g(k) - c]\,. \tag{21}$$

The first order condition with respect to c is:

$$u'(c) = \beta v'\,[g(k) - c]\,. \tag{22}$$

[4] Ljungqvist and Sargent refer to the last equation as the *transition law*.

Notice that since $c = g(k) - k'$, (18) and (22) are equivalent.

In place of the envelope condition, Sargent and Ljungqvist refer to the Benveniste and Scheinkman condition. These are effectively the same thing. The Benveniste and Scheinkman condition is obtained by imagining that you have solved for the optimal value of the control, c, as a function of the state, k. Here I will denote this function as $c = \eta(k)$. Substituting this into (21) you get

$$v(k) = u\left[\eta(k)\right] + \beta v\left[g(k) - \eta(k)\right].$$

Assuming v and η are differentiable you get

$$v'(k) = u'\left[\eta(k)\right]\eta'(k) + \beta v'\left[g(k) - \eta(k)\right]\left[g'(k) - \eta'(k)\right].$$

Notice that (22) implies that the terms attached to $\eta'(k)$ drop out (this will be true in more general problems), so that we are left with the Benveniste and Scheinkman condition:

$$v'(k) = \beta v'\left[g(k) - c\right]g'(k). \tag{23}$$

Notice that this condition can easily be derived using the "quick and dirty" method where you differentiate with respect to the state variable on the right-hand side of (21) without subtituting in the policy function, η, to obtain v'.

Although the envelope condition, (19), and the Benveniste and Scheinkman condition, (23), are not identical, notice that if we combine (22) and (23) we get

$$v'(k) = u'(c)g'(k),$$

which is equivalent to (19). So the two approaches yield equivalent optimality conditions.

1.6 Important Issues to Deal with Later

In considering (17), and in deriving conditions that characterize its solution, we have ignored several important questions:

1. Can we be sure that a v that satisfies (17) exists?
2. Can we be sure that there is a single-valued, and differentiable policy function, $h(k)$, that describes the optimal value of k'?
3. Does the solution to (17) correspond to the solution to (4)?
4. Since the optimality conditions, (20), are the same as the ones we obtained using the Lagrangean method, how do we deal with the issue of the missing end condition?

To answer these questions the next section considers a general dynamic optimization problem.

2 A General Dynamic Optimization Problem

Stokey and Lucas write down a general dynamic optimization problem

$$\max_{\{x_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t F(x_t, x_{t+1}) \tag{24}$$

subject to $x_{t+1} \in \Gamma(x_t)$, x_0 given. In our optimal growth model, where the equivalent of x_{t+1} was k_{t+1}, we had $0 \leq k_{t+1} \leq g(k_t)$, so that $\Gamma(k_t)$ would have been given by the set $[0, g(k_t)]$. In general, $\Gamma(x_t)$ is the set of feasible values of x_{t+1}, given x_t.

By analogy to what we previously saw with the growth model, we would expect the following dynamic programming problem to be equivalent to (24):

$$v(x) = \max_{y \in \Gamma(x)} F(x, y) + \beta v(y) \tag{25}$$

with x given.

The Bellman equation, (25), is a *functional equation*. That is, it is an equation defined over functions not over vectors. A standard equation solving problem is to find the value of an unknown vector x. An example is $x = f(x)$. On one side of the equation we have the unknown *vector* x. On the other side of the equation we have a *function*, f, applied to that vector. A functional equation will often be written in a similar (but not always identical way). An example of a functional equation is $f = T(f)$, where f is a *function* (it's common to drop its dependence on any arguments, such as x when writing the functional equation) and T is an *operator* that is applied to f.[5]

To write the Bellman equation as a functional equation we formally define a the operator T:

$$(Tw)(x) \equiv \max_{y \in \Gamma(x)} F(x, y) + \beta w(y) \tag{26}$$

with $x \in X$ given, and where, for the moment, w is some arbitrary function of the variable x.[6] T is an operator that multiplies a function, w, by β, then adds another function, $F(x, y)$, to it, and then maximizes the resulting function by choice of y subject to $y \in \Gamma(x)$ and x given.

The Bellman equation is a functional equation because it can be written as $v = T(v)$ or even more compactly as

$$v = Tv. \tag{27}$$

[5] A function maps vectors in some space, to vectors in some other space (often the same space). In our example with $x = f(x)$ we might imagine that $f : \mathbb{R}^n \to \mathbb{R}^n$. Similarly an operator maps a function in some space to a function in some other space (often the same space). When the two spaces are the same we can write $T : \mathcal{F} \to \mathcal{F}$, where $\mathcal{F}$ is some space of functions.

[6] X simply represents the set of all possible values x can take on. In our growth model example we could think of it as $[0, \tilde{k}]$, where $\tilde{k}$ is the maximal value of k.

2.1 Two Useful Theorems

There are two crucially useful theorems in Chapter 3 of Stokey and Lucas: *Blackwell's Theorem* and the *Contraction Mapping Theorem*.[7]

Blackwell's Theorem gives sufficient conditions for an operator T to be a *contraction mapping*. Knowing that T is a contraction mapping is very helpful because it lets you invoke the Contraction Mapping Theorem. It establishes that if an operator T is a contraction mapping then (i) it has a unique fixed point, i.e. there exists a unique function v such that $Tv = v$, and (ii) $T^n v_0 \to v$, as $n \to \infty$.[8] The contraction mapping theorem is incredibly powerful. It not only tells you that there is a unique solution, it tells you how to find it!

To understand Blackwell's theorem you need to know a little about *metric spaces*, *norms*, and contraction mappings.

Definition of a Metric Space (Stokey and Lucas p. 44) A *metric space* is a set $\mathcal{S}$ and a function $\rho : \mathcal{S} \times \mathcal{S} \to \mathbb{R}$, such that

1. $\rho(f, g) \geq 0$ for all $f,\ g \in \mathcal{S}$ (positivity)
2. $\rho(f, g) = 0$ iff $f = g$ (strict positivity)
3. $\rho(f, g) = \rho(g, f)$ for all $f,\ g \in \mathcal{S}$ (symmetry)
4. $\rho(f, h) \leq \rho(f, g) + \rho(g, h)$ for all $f,\ g,\ h \in \mathcal{S}$ (triangle inequality).$\square$

You can see that the function ρ, which is called a *metric*, is a concept of distance which shares some of the features of Euclidean distance. It's often the case, in practice, that the chosen metric is one for which $\rho(f, g) = \rho(f - g, \theta)$, where θ is a zero element of $\mathcal{S}$. By a zero element I mean that $\theta \in \mathcal{S}$ is such that $f + \theta = f$, $0f = \theta$, if 0 is the scalar zero.

Definition of a Normed Vector Space (Stokey and Lucas p. 45, 46) Skipping a little of the formality in Stokey and Lucas, I define a *normed vector space*, to be a set $\mathcal{S}$, and a norm (a special form of metric), $\|\cdot\| : \mathcal{S} \to \mathbb{R}$, such that

1. $\|f\| \geq 0$ for all $f \in \mathcal{S}$,
2. $\|f\| = 0$ iff $f = \theta$,
3. $\|\alpha f\| = |\alpha|\,\|f\|$, for all $f \in \mathcal{S}$, $\alpha \in \mathbb{R}$,
4. $\|f + g\| \leq \|f\| + \|g\|$, for all $f, g \in \mathcal{S}$.$\square$

[7] Blackwell's Theorem is described in section 3.3 (p. 54) of Stokey and Lucas. It also appears in section A.1 (p. 1010) of Ljungqvist and Sargent. The Contraction Mapping Theorem is described in section 3.2 (p.50) of Stokey and Lucas, and is mentioned indirectly in section A.2 (p. 1012) of Ljungqvist and Sargent.

[8] We will be more precise about the convergence concept being used here, later.

One useful metric space we will use is

$$\mathcal{S} = C[a,b]$$

$$\rho(f,g) = d_\infty(f,g) \equiv \sup_{x \in X = [a,b]} |f(x) - g(x)|.$$

Here, $[a,b] \subset \mathbb{R}$ is a closed interval on the real line, and $\mathcal{S}$ is the space of continuous functions with domain $X = [a,b]$. The metric d_∞ is often called the *sup norm*. Another metric that works with the same space of functions is

$$d_p(x,y) = \left[\int_a^b |f(x) - g(x)|^p dx \right]^{1/p},$$

where p is an integer (often 2). Another useful metric space would be the set of continuously differentiable functions with domain $[a,b]$ which is usually denoted $C^1[a,b]$. We will see how these metric spaces will be useful in a while.

Definition of a Contraction Mapping Let $(\mathcal{S}, \rho)$ be a metric space and let $T : \mathcal{S} \to \mathcal{S}$. T is a *contraction mapping* with modulus β if there exists a real number $0 \leq \beta < 1$ such that $\rho(Tf, Tg) \leq \beta \rho(f,g)$ for all f, $g \in \mathcal{S}$.$\Box$

Basically this means that applying a contraction mapping to two elements of $\mathcal{S}$, brings them closer together.

In the following theorem we will assume that there is an operator T which maps from a set of functions to the same set. You should keep in mind, for the moment, that we do not yet know whether the operator T in Bellman's equation has this property. So one thing we will try to show later is that there is a set of functions, $\mathcal{S}$, for which $w \in \mathcal{S}$ implies $Tw \in \mathcal{S}$.

In the following theorem you will also see the notation $\geq$ or $\leq$ in comparisons of function. What does this notation mean? Let f and g be any two functions with some common domain X. If $f(x) \geq g(x)$ for all $x \in X$, then we write $f \geq g$. In other words f must lie everywhere (nonstrictly) above g. Similarly for $\leq$.

Blackwell's Theorem Let $T : \mathcal{F} \to \mathcal{F}$ be an operator on a metric space $(\mathcal{F}, d_\infty)$, where $\mathcal{F}$ is a set of bounded functions and d_∞ is the sup norm. Assume that T has the following properties:

1) Monotonicity: For any f, $g \in \mathcal{F}$, $f \geq g \Rightarrow Tf \geq Tg$.

2) Discounting: For any constant real number $c > 0$, and every $f \in \mathcal{F}$, $T(f+c) \leq Tf + \beta c$, for some $0 \leq \beta < 1$.

Then T is a contraction mapping with modulus β.$\square$

Proof: Consider any f, $g \in \mathcal{F}$. We need only consider $f \neq g$ since if $f = g$, $d_\infty(Tf, Tg) = \beta d_\infty(f, g) = 0$. Notice that

$$\begin{aligned} f &= g + f - g \\ &\leq g + |f - g| \\ &\leq g + d_\infty(f, g) \end{aligned}$$

Property (1) implies that $Tf \leq T[g + d_\infty(f, g)]$. Property (2) implies that $T[g + d_\infty(f, g)] \leq Tg + \beta d_\infty(f, g)$. Therefore, $Tf \leq Tg + \beta d_\infty(f, g)$, or $Tf - Tg \leq \beta d_\infty(f, g)$.

Similarly you can reverse the roles of f and g to show that $Tg - Tf \leq \beta d_\infty(f, g)$. So

$$d_\infty(Tf, Tg) = \sup_{x \in X} |(Tf)(x) - (Tg)(x)| \leq \beta d_\infty(f, g).\blacksquare$$

Blackwell's theorem is useful because it is often easier to show that an operator is a contraction mapping by demonstrating that it has the monotonicity and discounting properties than it is to use the definition of a contraction mapping. Our next step is to look at the Contraction Mapping Theorem. Before doing so, however, we need a little more mathematical background. In particular, we need the definitions of Cauchy and convergent sequences as well as the definition of a complete metric space, and a Banach space.

Definition of a Cauchy Sequence Let $(\mathcal{S}, \rho)$ be a metric space. A sequence $\{f_n\}_{n=0}^\infty$, with $f_n \in \mathcal{S}$ for all n, is a Cauchy sequence if for each $\epsilon > 0$, there exists an $N(\epsilon)$ such that $\rho(f_n, f_m) < \epsilon$ for all n, $m \geq N(\epsilon)$.$\square$

Definition of a Convergent Sequence Let $(\mathcal{S}, \rho)$ be a metric space. A sequence $\{f_n\}_{n=0}^\infty$, with $f_n \in \mathcal{S}$ for all n, converges to $f \in \mathcal{S}$ if for each $\epsilon > 0$, there exists an $N(\epsilon)$ such that $\rho(f_n, f) < \epsilon$ for all $n \geq N(\epsilon)$. I.e. $\lim_{n \to \infty} \rho(f_n, f) = 0$.$\square$

Definition of a Complete Metric Space A metric space $(\mathcal{S}, \rho)$ is complete if every Cauchy sequence in $\mathcal{S}$ is a convergent sequence in $\mathcal{S}$.$\square$

Definition of a Banach Space A *complete normed vector space* is called a *Banach space*.$\square$

Contraction Mapping Theorem Let $(\mathcal{S}, \rho)$ be a complete metric space and let $T : \mathcal{S} \to \mathcal{S}$ be a contraction mapping with modulus β. Then:

1) there is a unique point $f \in \mathcal{S}$ such that $Tf = f$,

2) for any $f_0 \in \mathcal{S}$, the sequence $\{f_n\}_{n=0}^{\infty}$ defined by $f_n = Tf_{n-1}$, $n = 1, 2, \ldots$, satisfies $\rho(f_n, f) \leq \beta^n (f_0, f)$ for all n.$\square$

Proof: *Preliminaries.* Choose any $f_0 \in \mathcal{S}$, and define $f_n = Tf_{n-1}$, $n = 1, 2, \ldots$. It is clear, by induction, that $f_n \in \mathcal{S}$. From the fact that T is a contraction

$$\rho(f_{n+1}, f_n) \leq \beta\rho(f_n, f_{n-1}) \leq \beta^2\rho(f_{n-1}, f_{n-2}) \leq \cdots \leq \beta^n\rho(f_1, f_0). \tag{28}$$

We have not yet established convergence of the sequence f_n.[9]

Now take any n and any $m > n$. It follows from the triangle inequality that

$$\begin{aligned} \rho(f_m, f_n) &\leq \rho(f_m, f_{m-1}) + \rho(f_{m-1}, f_n) \\ \cdots &\leq \rho(f_m, f_{m-1}) + \rho(f_{m-1}, f_{m-2}) + \cdots + \rho(f_{n+1}, f_n). \end{aligned}$$

Using (28) this means

$$\begin{aligned} \rho(f_m, f_n) &\leq \left(\beta^{m-1} + \beta^{m-2} + \cdots + \beta^n\right)\rho(f_1, f_0) \\ &= \beta^n\left(1 + \beta + \cdots + \beta^{m-n-1}\right)\rho(f_1, f_0) \\ &= \beta^n \frac{1-\beta^{m-n}}{1-\beta}\rho(f_1, f_0) \\ &< \frac{\beta^n}{1-\beta}\rho(f_1, f_0). \end{aligned}$$

From this last result it is clear that f_n is a Cauchy sequence. Since $(\mathcal{S}, \rho)$ is a complete metric space this means $\exists f \in S$ such that $f_n \to f$, i.e. $\lim_{n\to\infty} \rho(f_n, f) = 0$.

Proof of Part (1): For all n and all $f_0 \in \mathcal{S}$ we have

$$0 \leq \rho(Tf, f) \leq \rho(Tf, T^n f_0) + \rho(T^n f_0, f)$$

[9] If, at this stage, we took the limit as $n \to \infty$, we would not have a proof that the sequence f_n converges. Rather, we would have a proof that consecutive terms in the sequence become closer and closer to one another. The classic example where the latter does not imply the former is

$$x_n = 1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{n}$$

which has no limit, whereas

$$x_{n+1} - x_n = \frac{1}{n+1}$$

does become arbitrarily small.

from the triangle inequality. But then, from the definition of a contraction, we have

$$\rho(Tf, T^n f_0) \leq \rho(f, T^{n-1} f_0).$$

Hence

$$0 \leq \rho(Tf, f) \leq \rho(f, T^{n-1} f_0) + \rho(T^n f_0, f) = \rho(f, f_{n-1}) + \rho(f_n, f). \tag{29}$$

Since $\lim_{n\to\infty} \rho(f_n, f) = 0$, and (29) holds for all n, including n arbitrarily large, it must be the case that $\rho(Tf, f) = 0$. Hence $Tf = f$.

Next we need to show that f is the *unique* fixed point. Suppose, to the contrary, that there is some other function, $\hat{f} \in \mathcal{S}$, $\hat{f} \neq f$, such that $T\hat{f} = \hat{f}$. Let $a = \rho(f, \hat{f})$. Clearly $a > 0$ since $\hat{f} \neq f$. Notice that since $f = Tf$ and $\hat{f} = T\hat{f}$, we have

$$a = \rho(f, \hat{f}) = \rho(Tf, T\hat{f}) \leq \beta\rho(f, \hat{f}) = \beta a,$$

where the inequality follows from the definition of a contraction. Since $\beta < 1$, this immediately implies a contradiction.

Proof of Part (2): $\rho(f_n, f) = \rho(Tf_{n-1}, Tf) \leq \beta\rho(f_{n-1}, f) \leq \beta^2\rho(f_{n-2}, f) \leq \cdots \leq \beta^n\rho(f_0, f)$.■

2.2 The Theorem of the Maximum

Blackwell's theorem is obviously powerful, but in order to use it we will need to show that the assumptions made by the theorem hold for the T that appears in Bellman's equation. A first step towards this goal is to study the Theorem of the Maximum.

Define

$$v(x) = \max_{y \in \Gamma(x)} f(x, y) \text{ given } x \in X. \tag{30}$$

Also let

$$H(x) = \{y \in \Gamma(x) | f(x, y) = v(x)\}. \tag{31}$$

Notice that v is a type of value function because it gives the maximized value of f for any x. $H(x)$ is the set of all optimal values of y given x, because it's the set of y's for which f is as high as v. Here, a capital letter is used simply to warn you of the possibility that H is not single-valued. If there were not a unique maximizer of f, H would be a correspondence.

Theorem of the Maximum Let $X \subseteq \mathbb{R}^l$ and $Y \subseteq \mathbb{R}^m$, let $f : X \times Y \to \mathbb{R}$ be a continuous function, and let $\Gamma : X \to Y$ be a compact-valued and continuous correspondence. Then the function v, defined in (30), is continuous, and the correspondence H, defined in (31), is nonempty, compact-valued, and upper-hemi continuous.□

Here compact-valued simply means that for each x, $\Gamma(x)$ is a compact set. A continuous correspondence is one that is both lower hemi-continuous and upper-hemi continuous. Without going into quite the detail Stokey and Lucas do, this means that any point $y \in \Gamma(x)$ can be reached as the limit of a sequence $y_n \in \Gamma(x_n)$ with $x_n \to x$, and that every sequence $y_n \in \Gamma(x_n)$, with $x_n \to x$, has a limit point in $\Gamma(x)$. To take an example, if $x \in \mathbb{R}$ and $y \in \mathbb{R}$, and $\Gamma(x)$ is defined in terms of inequalities, continuity of the boundaries defined by those inequalities is sufficient for lower- and upper-hemi continuity.

Basically what this theorem says is that if the function being maximized is continuous, so is the value function. Plus we get some nice properties for the policy function. If you make some more assumptions you get even more results.

Continuity of the Policy Function This makes two additional assumptions relative to the above theorem. Let Γ also be convex-valued and let f be strictly concave as a function of y for each $x \in X$. Then

$$h(x) = \{y \in \Gamma(x) | f(x, y) = v(x)\}$$

is single-valued, continuous function. An equivalent way of defining h is to state that

$$h(x) = \arg \max_{y \in \Gamma(x)} f(x, y) \text{ given } x \in X. \square$$

Limiting Policy Function is the Policy Function of the Limiting Problem Let $\{f_n\}$ be a sequence of functions, and f a function, which all have the properties assumed in the previous two theorems. Let $f_n \to f$ uniformly (in sup norm). Then define

$$h_n(x) = \arg \max_{y \in \Gamma(x)} f_n(x, y) \text{ given } x \in X,$$

$$h(x) = \arg \max_{y \in \Gamma(x)} f(x, y) \text{ given } x \in X.$$

If follows that $h_n \to h$, pointwise, or uniformly if X is a compact set.$\square$

2.3 Applying the Theorems to T

The main results are established in section 4.2 of Stokey and Lucas. Theorem 4.6 establishes that the theorem of the maximum applies to the problem on the right-hand side of T. This result allows the theorem to also establish that the conditions of Blackwell's Theorem are satisfied and that, therefore, T is a contraction mapping. Using the Contraction Mapping Theorem this delivers a very powerful result: there is a unique solution to (25).

The theorems subsequent to 4.6 establish other properties of the value and policy functions under additional assumptions. These assumptions are easily verified for the neoclassical growth model.

Define $A = \{(x, y) \in X \times X | y \in \Gamma(x)\}$. Notice the requirement that X be defined in such a way that it encompasses at least all the possible values of y, given all the possible values of $x \in X$. In our optimal growth model, we know that k and k' both must lie within the set $[0, \tilde{k}]$, where $f(\tilde{k}) = \delta\tilde{k}$ so A is $\{(k, k') \in [0, \tilde{k}] \times [0, \tilde{k}] | 0 \leq k' \leq g(k)\}$.

In what follows I will use the same numbering of assumptions and theorems as Stokey and Lucas.

Assumption 4.3. X is a convex subset of $\mathbb{R}^l$ and the correspondence $\Gamma : X \rightarrow X$ is nonempty, compact-valued and continuous.□

Assumption 4.4. The function $F : A \rightarrow \mathbb{R}$ is bounded and continuous, and $0 < \beta < 1$.□

It's easy to show that assumptions 4.3 and 4.4 hold for our growth model. You should do so as an exercise.

Theorem 4.6 (Existence and Uniqueness of the Value Function). Let assumptions 4.3 and 4.4 hold. Let $C(X)$ be the space of bounded continuous functions on X, with the sup norm. Then T [defined in (26)] is such that $T : C(X) \rightarrow C(X)$, T has a unique fixed point $v \in C(X)$, and for all $v_0 \in C(X)$,

$$||T^n v_0 - v|| \leq \beta^n ||v_0 - v||, \; n = 0, 1, \ldots.$$

Moreover, the optimal policy correspondence

$$H(x) = \{y \in \Gamma(x) | v(x) = F(x, y) + \beta v(y)\}$$

is compact-valued and upper hemi-continuous.□

Sketch of the proof: Assumptions 4.3 and 4.4 imply that a maximum exists if you solve the problem on the right-hand side of (26). If the function w on the right-hand side of (26) is continuous and bounded over X, i.e. it is an element of $C(X)$, it then follows from the theorem of the maximum that Tw is also continuous and bounded. Hence T maps from $C(X)$ to $C(X)$. This is crucial because it means T we can try to apply Blackwell's theorem. Conditions (a) and (b) of Blackwell's theorem are easy to establishfor T. This means that

T is a contraction, and that the remaining results stated in Theorem 4.6 follow from the Contraction Mapping Theorem.■

The next 4 theorems establish further properties of the value and policy functions. Some of these properties require us to make additional assumptions.

Assumption 4.5. For each y, $F_x(x, y) > 0$ (recall that x is a vector in $\mathbb{R}^l$ so this is a statement about the sign of l derivatives).□
Assumption 4.6. Γ is monotone in the sense that $x \leq x'$ implies $\Gamma(x) \subseteq \Gamma(x')$.□

These are both assumptions that our growth model example satisfies.

Theorem 4.7. (The Value Function is Increasing) Let assumptions 4.3–4.6 hold. Then, v, the unique solution to (25), is strictly increasing in x.□

Assumption 4.7. F is strictly concave.□
Assumption 4.8. If $y \in \Gamma(x)$ and $y' \in \Gamma(x')$, then $\lambda y + (1 - \lambda)y' \in \Gamma[\lambda x + (1 - \lambda)x']$ for all x, $x' \in X$ and any $0 \leq \lambda \leq 1$.□

Theorem 4.8. (The Value Function is Concave and the Policy Function is Continuous and Single Valued) Under assumptions 4.3, 4.4, 4.7 and 4.8, v, the solution to (25) is strictly concave and

$$h(x) = \{y \in \Gamma(x) | v(x) = F(x, y) + \beta v(y)\} = \arg \max_{y \in \Gamma(x)} F(x, y) + \beta v(y) \tag{32}$$

is a continuous and single-valued function.□

Theorem 4.9 (Convergence of the Policy Functions when you use Value Function Iteration) Let assumptions 4.3, 4.4, 4.7 and 4.8 hold. Let v satisfy (25) and h satisfy (32). Let $C'(X)$ be the space of bounded, continuous, concave functions, $f : X \to \mathbb{R}$, and let $v_0 \in C'(X)$. Then define v_n and h_n according to $v_{n+1} = Tv_n$, and

$$h_n(x) = \arg \max_{y \in \Gamma(x)} [F(x, y) + \beta v_n(y)], \quad \forall x \in X.$$

Then $h_n \to h$, pointwise. If X is compact, the convergence is uniform.□

Assumption 4.9. F is continuously differentiable on $\text{int}(A)$.□

Theorem 4.11 (The Envelope Theorem). Let assumptions 4.3, 4.4, and 4.7–4.9 hold, and let v and h be the solutions to (25) and (32). If $x_0 \in \text{int}(X)$ and $h(x_0) \in \text{int}(X)$, then v is continuously differentiable at x_0 with derivatives given by $v_x(x_0) = F_x[x_0, h(x_0)]$.□

2.4 Are the Planner's Problem and Bellman's Equation Equivalent?

Section 4.1 of Stokey and Lucas establishes that, indeed, (25) and (24) are equivalent. Theorems 4.2 and 4.3 work with modified versions of the Bellman equation and the planner's problem with "sup" replacing "max". For each x_0 the modified planner's problem has a unique supremum (by the definition of supremum), which is denoted $v^*(x_0)$. Theorem 4.2 establishes that v^* also satisfies the modified Bellman equation. Theorem 4.3 establishes that while there may be more than one v that satisfies the modified Bellman equation, the only one that satisfies a particular boundedness condition on the objective function is v^*. To summarize these two theorems, they say that v^* is the value associated with both the planner's problem and Bellman's equation.

Theorems 4.4 and 4.5 make similar statements, but about the policy functions rather than the value functions. Theorem 4.4 states that any plan (i.e. a chosen path for $\{x_t\}_{t=1}^{\infty}$) that achieves the sup of the modified planner's problem (in other words is the argmax) can be generated from the policy correspondence of Bellman's equation. Theorem 4.5 states that any plan that can be generated from the policy correspondence of Bellman's equation, maximizes the planner's objective, as long as it satisfies a particular boundedness condition.

Assumption 4.1. The correspondence $\Gamma(x)$ is nonempty for all $x \in X$.$\square$

Define the set of feasible plans given x_0: $\Pi(x_0) = \{\{x_t\}_{t=0}^{\infty} | x_{t+1} \in \Gamma(x_t)\ \forall t\}$. Let any element of this set be denoted $\mathbf{x} \in \Pi(x_0)$.

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

Define $u(\mathbf{x}) \equiv \lim_{n\to\infty} \sum_{t=0}^{n} \beta^t F(x_t, x_{t+1})$.

Lemma 4.1. Let X, Γ, F and β be such that assumption 4.2 is satisfied. Then for any $x_0 \in X$ and any plan $\mathbf{x} \in \Pi(x_0)$,

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

where $\mathbf{x}' = \{x_t\}_{t=1}^{\infty}$.$\square$

It's sort of a trivial lemma (it basically says what I said before when I asserted that v_s^{∞} had ot be the same function as v_{s+1}^{∞}). Now define

$$v^*(x_0) \equiv \sup_{\mathbf{x}\in\Pi(x_0)} u(\mathbf{x}) = \sup_{\mathbf{x}\in\Pi(x_0)} \lim_{n\to\infty} \sum_{t=0}^{n} \beta^t F(x_t, x_{t+1}) \tag{33}$$

for each $x_0 \in X$. Notice that (33) is a modified version of the social planner's problem with "sup" replacing "max", and v^* is its unique solution (this is a property of "sup"). The next theorem, similarly, works with a modified version of Bellman's equation with "sup" replacing "max".

Theorem 4.2. Let X, Γ, F and β be such that assumptions 4.1 and 4.2 are satisfied. Then $v^*(x)$ satisfies the Bellman equation with "max" replaced by "sup", i.e. v^* is one solution to:

$$v(x) = \sup_{y \in \Gamma(x)} F(x, y) + \beta v(y), \quad \forall x \in X. \square \tag{34}$$

At this point we know that v^* defined in (33) is a solution to the modified version of the Bellman equation, (34), but we don't know if it's the unique solution. The next theorem shows that v^* is the only solution that leads to a particular kind of boundedness in utility.

Theorem 4.3. Under assumptions 4.1 and 4.2, if v satisfies the modified Bellman equation and if $\lim_{n\to\infty} \beta^n v(x_n) = 0$ for all $\mathbf{x} \in \Pi(x_0)$ and x_0, then $v = v^*$.$\square$

So if we found a solution to the modified Bellman equation, we'd know it was the solution to modified planner's problem as long as we checked that the boundedness condition held.

The problem with "sups" is that there may not be feasible plans that reach them. Theorem 4.4 states that *if* there is an optimal plan $\mathbf{x}$ that reaches the sup in (33) (in other words, the max exists), it is also consistent with the optimal policy correspondence of (34). Theorem 4.5 states an approximate converse.

Theorem 4.4. Make assumptions 4.1 and 4.2. Suppose there exists a plan $\mathbf{x}^* \in \Pi(x_0)$ that obtains the sup in (33) [in other words a solution to (24) exists]. Then $v^*(x_t^*) = F(x_t^*, x_{t+1}^*) + \beta v^*(x_{t+1}^*)$ $\forall t$, with $x_0^* = x_0$. This means any plan $\mathbf{x}^*$ that is optimal for the social planner, can be generated from the optimal policy correspondence H^* associated with Bellman's equation where

$$H^*(x) = \{y \in \Gamma(x) | v^*(x) = F(x, y) + \beta v^*(y)\}. \square$$

Theorem 4.5. Let $\mathbf{x}^* \in \Pi(x_0)$ be any feasible plan such that $v^*(x_t^*) = F(x_t^*, x_{t+1}^*) + \beta v(x_{t+1}^*)$ $\forall t$, with $x_0^* = x_0$, and such that $\lim_{t\to\infty} \sup \beta^t v^*(x_t^*) \leq 0$ [in other words let $\mathbf{x}^*$ be in the optimal policy correspondence associated with Bellman's equation]. Then $\mathbf{x}^*$ obtains the sup in (33) [in other words it is an optimal plan for the social planner].$\square$

You should note that although Theorems 4.2–4.5 seem to leave open the possibility of multiple solutions to Bellman's equation and that the "max" of either problem might not be achievable, these loose ends are tied up in the later sections of the text, which we have already covered. Remember that Theorem 4.6 proves that a unique solution to the "max" version of Bellman's equation exists, but under more stringent conditions than those imposed in Section 4.1 of Stokey and Lucas.

2.5 Transversality Conditions

The main result of section 4.5 in Stokey and Lucas is Theorem 4.15. It establishes *sufficient* conditions for the optimality of a plan $\mathbf{x}$. In particular, it shows that if a plan, $\mathbf{x}$, satisfies the Euler equations and a transversality condition then it is an optimal plan.

Theorem 4.15. Let $X \subset \mathbb{R}^l_+$ and assumptions 4.3–4.5, 4.7 and 4.9 hold. Then the plan $\mathbf{x}^*$ with $x^*_{t+1} \in \text{int}\Gamma(x^*_t)$, $t = 0, 1, \ldots$, is optimal for (33) given x_0 if it satisfies

$$F_y(x^*_t, x^*_{t+1}) + \beta F_x(x^*_{t+1}, x^*_{t+2}) = 0, \qquad t = 0, 1, \ldots$$

and

$$\lim_{t\to\infty} \beta^t F_x(x^*_t, x^*_{t+1}) \cdot x^*_t = 0.\square$$

Of course, the Euler equation can be derived from the first-order and envelope conditions associated with (25). It's easy to see that those are

$$F_y(x, y) + \beta v'(y) = 0$$

$$v'(x) = F_x(x, y)$$

The Euler equation follows immediately from substituting the expression for v' in the envelope condition into the first-order condition.

In our example of the neoclassical growth model, the analog to $F(x, y)$ is $u[g(k_t) - k_{t+1}]$ where k_t plays the role of x and k_{t+1} plays the role of y. Therefore, for the neoclassical growth model we should write the transversality condition as $\lim_{t\to\infty} \beta^t u'[g(k_t) - k_{t+1}]g'(k_t)k_t = 0$, or equivalently as $\lim_{t\to\infty} \beta^t u'(c_t)[f'(k_t) + (1-\delta)]k_t = 0$.

It is very important to be aware that the transversality condition is not something that needs to be imposed so that a solution to the planning problem exists. Remember, we already know that a solution exists from the theory described prior to this section. As we will see later in more detail, however, the Euler equations are not sufficient to uniquely pin down the solution to the planner's problem. Rather, the transversality condition can be a handy tool for finding the optimal plan. Theorem 4.15 says that if you have a plan that satisfies the Euler equations and satisfies the transversality condition, then it is an optimal plan. The converse need not be true.

3 Stochastic Dynamic Programming

We will not explore stochastic dynamic programming in great detail, as the theoretical machinery required to do this formally is more complex than for deterministic dynamic programming. Instead, we will examine the basic approach.

A simple stochastic extension of the neoclassical model used by Stokey and Lucas involves making the production technology subject to random shocks. For example we can replace the problem (4) with

$$\max_{\{k_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t u[z_t f(k_t) + (1-\delta)k_t - k_{t+1}]$$

subject to $0 \leq k_{t+1} \leq z_t f(k_t) + (1-\delta)k_t$, for $t \geq 0$, and $k_0 > 0$ given.

Here z_t is a sequence of exogenous random variables with some probability distribution (Stokey and Lucas assume that it is an iid sequence in Section 2.2). Here E_t represents the expectation taken with respect to all information known at the beginning of time t.

Of course, at time t the planner cannot literally choose $\{k_{t+1}\}_{t=0}^{\infty}$ because the planner does not know the future realizations of z_t. But there are two perfectly coherent ways in which we can think of the social planner making his choices.

1. The planner chooses k_1 at date 0 given the values of k_0, z_0, and all other information available to the planner that might be useful in forecasting future outcomes. Then at date 1 the planner chooses k_2 given the realizations of k_1, z_1, and all other relevant information, and so on. In this way the planner's problem is treated explicitly as being recursive in nature.

2. The planner chooses k_1 at date 0 given k_0, z_0 and all other information, and chooses *contingency plans* for the entire future sequence of $\{k_{t+1}\}_{t=1}^{\infty}$. These contingency plans determine each k_{t+1} as a function of k_t, z_t, and any other variables that would be useful for forecasting the future from date t forward.

By setting the planner's problem up in a dynamic programming framework we will capture the flavor of both of these interpretations. The recursive structure of the optimization problem as described in (1) sounds very much like the recursive aspect of using dynamic programming. On the other hand, the dynamic programming approach delivers contingency plans (optimal policy rules) that express the next period's capital stock as optimally chosen functions of the current capital stock, the current value of the shock, z, and any other information relevant for forecasting the future. In this way, the dynamic programming approach aslo captures the essence of the description in (2).

Define

$$g(k, z) = zf(k) + (1 - \delta)k.$$

The recursive version of the planning problem is

$$V(k, z) = \max_{k'} u[g(k, z) - k'] + \beta E[V(k^{'}, z')|z]$$

subject to $0 \leq k' \leq g(k, z)$, with k, z given. Here I have implicitly assumed that z is a Markov process, i.e. $F(z_{t+1}|z_t, z_{t-1}, \cdots) = F(z_{t+1}|z_t)$ for all $(z_{t+1}, z_t, \ldots)$. This is what allows me to stop talking about "all other relevant information." If the conditional distribution of z_t depended on more of the history of z than just z_{t-1}, I would need to redefine the vector of state variables to include that extra history.

Under the same assumptions as we used previously (which would be much more difficult to establish) the optimality condition is

$$u'[g(k, z) - k'] = \beta E[V_k(k^{'}, z')|z].$$

The envelope condition is

$$V_k(k, z) = u'[g(k, z) - k']g_k(k, z).$$

Hence the Euler equation is

$$u'[g(k, z) - k'] = \beta E \left\{ u'[g(k', z') - k'']g_k(k', z')|z \right\}.$$

The policy function will be of the form $k' = h(k, z)$.