An Application to Growth Theory

First let's review the concepts of *solution function* and *value function* for a maximization problem. Suppose we have the problem

$$\max_{x \in X} F(x, \alpha) \;\text{ subject to }\; G(x, \beta) \leqq 0, \tag{P}$$

where $\alpha \in A$ and $\beta \in B$ are parameters, and $F : X \times A \to \mathbb{R}$ and $G : X \times B \to \mathbb{R}$. If for each $\alpha \in A$ and $\beta \in B$ the problem P has a unique solution, then we have the two functions

$s(\cdot) : A \times B \to X$, the **solution function**, where $s(\alpha, \beta)$ is the solution of the problem P for the parameter values α and β; and

$v(\cdot) : A \times B \to \mathbb{R}$, the **value function**, defined by $v(\alpha, \beta) := F(s(\alpha, \beta), \alpha)$, *i.e.*, $v(\alpha, \beta)$ is the value of F at the solution $s(\alpha, \beta)$.

Now suppose we have an economy in which capital is used to produce output according to the production function $y = f(x)$, and assume that the output will be allocated between consumption and investment — *i.e.*, $c + i \leqq y$. Suppose, furthermore, that the amount of capital available tomorrow, x_1, will be equal to the amount invested today — *i.e.*, $x_1 = i_0$. (We're assuming, for simplicity, that today's capital depreciates fully by the time tomorrow arrives. We could instead assume that capital depreciates only at the rate δ — *i.e.*, that $x_1 = (1 - \delta)x_0 + i_0$. But assuming that $\delta = 1$ makes the analysis more transparent without changing anything conceptually.) Finally, assume that we start off with a given amount of capital, x_0, today, and that we care only about the two periods $t = 0$ and $t = 1$ (*i.e.*, "today" and "tomorrow"): we evaluate consumption streams (c_0, c_1) according to a utility function $U(c_0, c_1)$. Thus, we have the following constrained maximization problem:

$$\max_{c_0, c_1, x_1} U(c_0, c_1) \;\text{ subject to }\; c_0 + x_1 \leqq f(x_0) \;\text{ and }\; c_1 \leqq f(x_1). \tag{1}$$

It's clear that a solution must satisfy both inequalities exactly (if U is increasing), so let's write the problem's constraints as equations:

$$\max_{c_0, c_1, x_1} U(c_0, c_1) \;\text{ subject to }\; c_0 + x_1 = f(x_0) \;\text{ and }\; c_1 = f(x_1). \tag{2}$$

Note that both c_0 and c_1 are determined by our choice of just the one variable x_1 (or equivalently, i_0) — which is not surprising, since the problem has three decision variables and two constraints.

Now suppose there are three periods instead of two: $t = 0, 1, 2$. We again start with a given capital stock, x_0, and now we choose the variables c_0, c_1, c_2, x_1, x_2 subject to the three constraints

$$c_0 + x_1 = f(x_0)\ ,\ c_1 + x_2 = f(x_1)\ ,\ \text{and}\ \ c_2 = f(x_2).$$

As before, the choice of the capital stocks x_1 and x_2 (or the investment levels i_0 and i_1) determine the values of c_0, c_1, c_2.

In fact, for any finite number of periods — say $t = 0, 1, \ldots, T$ — and any starting capital stock $x_0 \in \mathbb{R}_+$, we have a straightforward constrained maximization problem, to choose $x_1, x_2, \ldots, x_T$ to maximize $U(c_0, \ldots, c_T)$, where the variables c_t and x_t satisfy the equations

$$c_t + x_{t+1} = f(x_t)\ \text{ for }\ t = 0, 1, \ldots, T-1, \ \text{ and }\ c_T = f(x_T).$$

Notice, in particular, that the last period is different from all the preceding periods: in the last period we don't have to give up any consumption in order to leave some capital for the future periods — *i.e.*, to provide for consumption in future periods — because there are no future periods.

But now let's assume there will always be future periods — so there *is* no last period. We now have an infinite horizon maximization problem, to choose a sequence $\mathbf{x} = (x_1, x_2, \ldots) \in \mathbb{R}^\infty_+$ to maximize a utility function $U(c_0, c_1, \ldots)$ subject to the constraints

$$c_t + x_{t+1} = f(x_t)\ \text{ for }\ t = 0, 1, 2, \ldots, \tag{3}$$

where x_0 is the initial capital stock, which is given.

Let's henceforth assume that the utility function is a sum of period-by-period utilities, discounted at the rate β:

$$U(c_0, c_1, \ldots) = \sum_{t=0}^{\infty} \beta^t u(c_t). \tag{4}$$

Notice that we can write the problem entirely in terms of the capital-stock decision variables x_t:

$$\begin{aligned} \max_{\mathbf{x}\in\mathbb{R}^\infty_+}\ \widetilde{U}(x_1, x_2, \ldots) &= \sum_{t=0}^{\infty} \beta^t u(f(x_t) - x_{t+1}) \\ \text{subject to} \quad 0 &\leqq x_{t+1} \leqq f(x_t),\ t = 0, 1, 2, \ldots. \end{aligned} \tag{5}$$

Let's assume, for now, that whatever is the level of the initial capital stock, $x_0 \in \mathbb{R}_+$, the problem (5) has a unique solution $\widehat{\mathbf{x}} = (\widehat{x}_1, \widehat{x}_2, \ldots) \in \mathbb{R}^\infty_+$, which of course depends on x_0. Then we have both a solution function $\widehat{\mathbf{x}}(\cdot) : \mathbb{R}_+ \to \mathbb{R}^\infty_+$ and a value function $v(\cdot) : \mathbb{R}_+ \to \mathbb{R}$:

$$\widehat{\mathbf{x}}(x_0) = (\widehat{x}_1, \widehat{x}_2, \ldots) \in \mathbb{R}^\infty_+ \quad \text{and} \quad v(x_0) = \widetilde{U}(\widehat{\mathbf{x}}(x_0)). \tag{6}$$

Notice that while we want to choose a plan $\mathbf{x} = (x_1, x_2, \ldots)$ for the entire future, the only part of the plan that we can actually implement today is the pair (c_0, i_0), or equivalently, the capital stock x_1 that will be carried forward to period $t = 1$. If we knew the value function $v(\cdot)$ in (6), we could write today's problem as a simple maximization problem with two variables and one constraint:

$$\max_{c_0, x_1} [u(c_0) + \beta v(x_1)] \text{ subject to } c_0 + x_1 = f(x_0), \tag{7}$$

because the problem we will face tomorrow is the same as today's problem, but with today's capital stock x_0 replaced by tomorrow's capital stock, x_1. And the value of the solution of the problem (7) must be $v(x_0)$, so we have the equation

$$v(x_0) = \max_{c_0, x_1} [u(c_0) + \beta v(x_1)] \text{ subject to } c_0 + x_1 = f(x_0), \tag{8}$$

and this equation must hold for every $x_0 \in \mathbb{R}_+$. We can rewrite (8) as follows, using the constraint equation to replace the variable c_0:

$$\forall x_0 \in \mathbb{R}_+ : \quad v(x_0) = \max_{x_1} [u(f(x_0) - x_1) + \beta v(x_1)] \text{ subject to } 0 \leqq x_1 \leqq f(x_0). \tag{9}$$

And finally, notice that we don't really need the time subscripts: the equation in (9) must hold at every period, for any capital stock $x_0 \in \mathbb{R}_+$ with which we enter the period; and the variable in the maximization problem does not have to have a subscript, it just has to be distinguished from the variable x_0. So we can rewrite (9) as follows:

$$\forall x \in \mathbb{R}_+ : \quad v(x) = \max_{z} [u(f(x) - z) + \beta v(z)] \text{ subject to } 0 \leqq z \leqq f(x). \tag{10}$$

Unfortunately, we now have some bad news: we don't have any guarantee that such a function $v(\cdot)$ actually exists. Recall that we obtained (10) by saying "Let's assume for now that the problem (5) has a unique solution." But we don't actually know whether (5) has a solution. On the other hand, if we knew that a function satisfying (10) *does* exist (*and* if it's continuous and $u(\cdot)$ is also continuous), then the Weierstrass Theorem guarantees that for any $x \in \mathbb{R}_+$ the simple maximization problem in (10) will have a solution, say $\widehat{z}$. Note that x is a parameter in the maximization problem, so let's write the solution function for the problem as $\widehat{z} = s(x)$. Now, for any starting capital stock x_0 we can determine the successive capital stocks that solve the maximization problem in (10) by recursively applying the function $s(\cdot)$:

$$x_1 = s(x_0), \; x_2 = s(x_1), \; x_3 = s(x_2), \; \ldots. \tag{11}$$

And then, under some fairly weak conditions, we can show that the sequence $(x_1, x_2, \ldots)$ is actually a solution of our original problem in (5). So our goal now is simply to establish the existence of a function v that satisfies (10).

Let's ignore, for the moment, the function v that appears on the left-hand side of (10), and focus just on the v on the right-hand side. If that function is continuous (and if u and f are continuous, as well), then the maximization problem in (10) will still have a solution for every $x \in \mathbb{R}_+$. Therefore, the right-hand side of (10) defines a real-valued function on $\mathbb{R}_+$, say $\widetilde{v}$, which is is the value function for the maximization problem:

$$\forall x \in \mathbb{R}_+ : \quad \widetilde{v}(x) = \max_z \left[u(f(x) - z) + \beta v(z)\right] \ \text{ subject to } \ 0 \leqq z \leqq f(x). \tag{12}$$

Thus, (12) defines a transformation that transforms any continuous real function v into a new real function $\widetilde{v}$. Later in the course we'll develop a result called the Maximum Theorem, which guarantees that this new function $\widetilde{v}$ is also continuous. We can also impose conditions which guarantee that both v and $\widetilde{v}$ are bounded functions. Thus, we have a transformation $T : C_B(\mathbb{R}_+) \to C_B(\mathbb{R}_+)$, where $C_B(\mathbb{R}_+)$ is the set of all bounded continuous functions on $\mathbb{R}_+$, which is a complete metric space under the *sup*-metric. Now, if we could show that T is a contraction, then we would know that T has a fixed point, a function $v \in C_B(\mathbb{R}_+)$ for which $\widetilde{v} = v$ — *i.e.*, a function v that *does* satisfy (10).

In fact, it's straightforward to verify that T is a contraction, but we won't do that here. We'll just note that the Contraction Mapping Theorem turns out to be exactly the tool we need to solve the problem we started out with — to determine the *optimal policy* for growing the capital stock and its associated consumption stream. We establish that for any $x_0 \in \mathbb{R}_+$ there is an optimal solution $(\widehat{x}_1, \widehat{x}_2, \dots)$ to the problem (5), and that the solution has the recursive structure in (11):

$$\widehat{x}_1 = s(x_0),\ \widehat{x}_2 = s(\widehat{x}_1),\ \widehat{x}_3 = s(\widehat{x}_2),\ \dots. \tag{13}$$

Moreover, the method of successive approximations provides a means of numerically approximating the solution: we can start with any arbitrary function v_0 and recursively determine a sequence of functions v_t:

$$v_1 = \widetilde{v}_0 = T(v_0), \quad v_2 = \widetilde{v}_1 = T(v_1), \quad v_3 = \widetilde{v}_2 = T(v_2),\ \dots, \tag{14}$$

a sequence that has to converge (monotonically-in-distance!) to the unique function v that satisfies (10). We can then determine the solution function s, or at least properties of s, by straightforward application of Kuhn-Tucker/Lagrangean methods to the simple maximization problem in (10).

Notes:

(1) We said that we can write today's problem as the simple maximization problem in (7), using the value function $v(\cdot)$ to summarize the entire future value of arriving at tomorrow's date with a capital stock x_1, and that we can do that because the problem we will face tomorrow is the same as today's problem, but with today's value of the parameter, x_0, replaced by tomorrow's value, x_1. This fact is **The Principle of Optimality**, due to Richard Bellman. While it's extremely intuitive, it does have to be proved, which we don't do here.

(2) The *functional equation* (10), involving the simple maximization problem, is called the **Bellman Equation** for the infinite-dimensional problem (5).