1 Two- and Three-Period Optimization Problem

In this section, we are going to introduce dynamic optimization (i.e., optimization *over time*) by solving a simple two-period optimization problem. Consider, for example, the problem of solving:

$$\max_{c_1,a_2,c_2,a_3} log(c_1) + \beta log(c_2), \tag{1}$$

$$\text{s.t. } c_1 = \bar{a} - a_2 \text{ (period 1 budget constraint)}$$

$$\text{and } c_2 = Ra_2 - a_3 \text{ (period 2 budget constraint)}$$

where c_t is consumption in period t, a_{t+1} is savings chosen at time t to be brought into period $t+1$, and β is a discount factor. This discount factor generally accounts for individuals' preference towards future consumption, or their measure of patience. The remaining parameters are initial wealth, $\bar{a}$, and the gross rate of return on savings, R.

We can do a few things that simplify our problem. First, note that $a_3^* = 0$, since optimal savings in the last period of life is zero. Second, note that we can use substitution to rewrite the problem as:

$$\max_{a_2} \; log(\bar{a} - a_2) + \beta log(Ra_2). \tag{2}$$

This reduces the problem to a single-dimensional optimization problem. Notice that marginal utility is infinite near zero and zero at infinity. This problem satisfies the so-called *Inada Conditions* that keep us from having to worry about corner solutions. Instead, we can just focus on the first-order condition, which gives:

$$\frac{-1}{\bar{a} - a_2} + \frac{\beta R}{Ra_2} = 0. \tag{3}$$

The solving this problem for a_2 gives optimal solution:

$$a_2^* = \frac{\beta \bar{a}}{1+\beta}. \tag{4}$$

Then, we can plug a_2^* into the budget constraints to get:

$$c_1^* = \bar{a} - \frac{\beta \bar{a}}{1+\beta} \tag{5}$$

$$c_2^* = R\frac{\beta \bar{a}}{1+\beta}. \tag{6}$$

Finally, we can plug c_1^* and c_2^* into the utility function to get the *indirect utility function*:

$$V(\bar{a}, R, \beta) = log(\bar{a} - \frac{\beta \bar{a}}{1+\beta}) + \beta log(R\frac{\beta \bar{a}}{1+\beta}). \tag{7}$$

Whereas the optimal solutions $c_1^*(\bar{a}, R, \beta)$ and $c_2^*(\bar{a}, R, \beta)$ tell us how behavior responds to a change in the parameters, the indirect utility function $V(\bar{a}, R, \beta)$ tells us how optimal utility responds to a change in the parameters.

Notice that a_2 gave us our starting wealth in period 2, and $\bar{a}$ gave us our starting wealth in period 1. As it turns out, we can solve the three-period optimization problem starting in period 0 simply by converting the parameter $\bar{a}$ into the variable a_1 and using the indirect utility function in the place of period 1 and 2 utility as follows:

$$\max_{a_1} \quad log(c_0) + \beta V(a_1) \tag{8}$$

$$\text{s.t.} \;\; c_0 = a_0 - a_1 \;\; \text{(period 0 budget constraint)},$$

where a_0 is given, and we dropped β and R from $V(\cdot)$ for simpler notation. In general, we can write the period 0 indirect utility function for an arbitrary utility function $u(\cdot)$ as the solution to the optimization problem:

$$V_0(a_0) = \max_{a_1} \quad u(c_0) + \beta V_1(a_1) \tag{9}$$

$$\text{s.t.} \;\; c_0 = a_0 - a_1 \;\; .$$

2 T-Period Problem

Notice in the period 0 problem (9) that we needed V_1 before we could solve for V_0. We will always need the value function from period $t+1$ before we can solve the period t problem. The process of solving the problem in reverse sequence is known as *backwards induction.* We can generalize (9) to represent the problem in any period t as follows:

$$V_t(a_t) = \max_{a_{t+1}} \quad u(c_t) + \beta V_{t+1}(a_{t+1}) \tag{10}$$

$$\text{s.t.} \;\; c_t = a_t - a_{t+1} \;\; .$$

In the language of dynamic optimization, (10) is know as *Bellman's equation.*

We can start to think of dynamic optimization problems more generally by thinking about the common features of Bellman's equations. First, consider the wealth brought into period t, denoted as a_t in equation (10). At the time that the agent optimizes, there is nothing that he or she can do about a_t. In other words, a_t is a *sunk cost* that reflects previous decisions or realizations which the agent can not control at period t. We call a_t a *state variable.* The best that the agent can do is take a_t as given and choose a_{t+1} to maximize remaining lifetime utility. Choosing a_{t+1} not only determines the amount of utility received by the agent in period t, but it will also determine the agent's state when it comes time to optimize in period $t+1$. We call a_{t+1} a *choice variable* or an *action variable.* We will also have generic names for the functions in our Bellman's equations. The indirect utility function in period t, $V_t(\cdot)$, is called the *value function*, whereas the future indirect utility function in period t, $V_{t+1}(\cdot)$, is called the *continuation value.* The utility function $u(\cdot)$ is generally called the *contemporaneous payoff function.* Finally, once we have solve for the optimal choice, a^*_{t+1}, we generally call it the *policy function*, and we write it as a function of any state variables: $a^*_{t+1}(a_t)$.

To summarize, we generally want to optimize a value function, which depends on the state variables. The value function is the sum of a contemporaneous payoff function and discounted continuation value optimized over the choice space.