Chapter 1

Mathematical and Computational Tools

1.1 Dynamic Optimization

We will start with a very simple dynamic optimization problem, that it is nothing else than an application a finite dimension optimization problem. Then, we will discuss extensions to infinite dimension problems.

Consider a consumer that has to decide how to consume a cake of a size W_1 over T periods. Let c_t denote the consumption level of period t, and let $u(c_t)$ the flow of consumption or utility associated to c_t. Preferences are stationary and not indexed by time. We make the following set of assumptions on $u(\cdot)$

1. Real valued function
2. Differentiable
3. Strictly increasing
4. Strictly concave
5. Inada conditions $\lim_{c\to 0} u'(c) = \infty$.

Consumer preferences are represented by

$$\sum_{t=1}^{T} \beta^{t-1} u(c_t)$$

where $\beta \in (0,1)$ is the discount rate. We assume that the cake does not depreciate over time. The law of motion for the cake is given by

$$W_{t+1} = W_t - c_t, \qquad t = 1, ..., T$$

The objective is to determine the optimal consumption plan $\{c_t^*\}_{t=1.}^T$

1.1.1 Sequence problem approach

A direct approach would solve the constraint optimization problem directly. Formally

$$\max_{\{c_t\}_{t=1}^T, \{W_t\}_{t=2}^{T+1}} \sum_{t=1}^T \beta^{t-1} u(c_t)$$

$$s.to. \qquad W_{t+1} = W_t - c_t, \qquad t = 1, ..., T$$

together with some non-negativity constraints, $c_t \geq 0$, and $W_{t+1} \geq 0$, where W_1 is given. There are two different ways to simplify this problem and reduce the number of choice variables.

1. **Sequential formulation:** Substitute the laws of motion into the objective function. Formally,

$$\max_{\{W_t\}_{t=2}^{T+1}} \sum_{t=1}^T \beta^{t-1} u(W_t - W_{t+1})$$

   and we still have that $W_{t+1} \geq 0$. The non-negativity constraint on consumption implies that $c_t = W_t - W_{t+1} \geq 0$, which amounts to say that the cake just gets smaller over time $W_t \geq W_{t+1}$.

2. **Compact formulation:** Combining the laws of motion for all periods we have

$$\max_{\{c_t\}_{t=1}^T} \sum_{t=1}^T \beta^{t-1} u(c_t)$$

$$\sum_{t=1}^T c_t + W_{t+1} = W_1$$
$$W_{t+1} \geq 0$$

   This is a well-behaved problem with a strictly concave objective function over a compact set. Weierstrass theorem guarantees a solution on the problem. One way of characterizing the solution is using the first-order conditions. Let λ and μ be the Lagrange multiplier of the resource constraint and the non-negativity constraint respectively. Formally

$$\beta^{t-1}u'(c_t) = \lambda$$
$$\lambda = \mu$$

Given that consumption will always be positive $c_t > 0$, the implied multiplier is always positive too, $\lambda > 0$. Consequently $\mu > 0$, and the non-negativity constraint binds in the last period $\mu > 0$ and $W_{t+1} > 0$.

Euler Equations: The results of both formulations yield to a necessary condition for optimality that links consumption across two periods.

$$u'(c_t) = \beta u'(c_{t+1})$$

This condition is called Euler equation. Explain the economic intuition behind the Euler equation....

The Euler equation can also be used to cover deviations that last more than one period. In particular, we can combine it over extended periods of time

$$u'(c_t) = \beta^{\tau} u'(c_{t+\tau})$$

As long as the problem is finite, the fact that the Euler equation holds across adjacent periods implies that any finite deviations from a candidate solution that satisfies the Euler equation will not increase utility.

Terminal conditions: Why the Euler equation is only sufficient? There could be solutions that satisfy all the Euler equations but satisfy this property $W_T > c_T$, so there is cake left. Clearly, this is not an optimal solution because we can increase total utility by increasing the consumption, on the last period or even better a bit on every period. Therefore the optimal solution has to satisfy

$$u'(c_t) = \beta u'(c_{t+1})$$

for all t, as well as $W_{T+1} \geq 0$. This last constraint has to purposes. First, it does no allow the consumer to try to set $W_{T+1} = -\infty$ in the last period and obtain unbounded utility. Second, it ensures that in the limit there is no cake left, $\mu W_{T+1} = 0$.

The Euler equation together with the non-negativity constraint $W_{T+1} \geq 0$ and the initial condition $W_1 > 0$ determine a second-order difference equation with two boundary conditions that we need to solve for.

Example (Sequential Formulation): It is useful to have an example that has closed form solution. Assume that preferences are of the form $u(c) = \log c$. In this particular case, the optimization problem is

$$v_T(W_1) = \max \textstyle\sum_{t=1}^{T} \beta^{t-1} \log(c_t)$$
$$s.t. \qquad \textstyle\sum_{t=1}^{T} c_t = W_1,$$

Let λ be the Lagrange multiplier of the resource constraint, then we can write down the Lagrangian as

$$v_T(W_1) = \max_{\{\{c_t\}_{t=1}^{T},\lambda\}} \textstyle\sum_{t=1}^{T} \beta^{t-1} \log(c_t) + \lambda[W_1 - \textstyle\sum_{t=1}^{T} c_t],$$

The Euler equation implies

$$\frac{1}{c_t} = \frac{\beta}{c_{t+1}}$$

or

$$\gamma_c = \frac{c_{t+1}}{c_t} = \beta < 1$$

The growth rate of consumption is decreasing, that implies $c_1 > c_2 > .... > c_T$. In general, we can write consumption in period c_t as a function of the initial period c_1

$$c_t = \beta^{t-1} c_1,$$

Now, we can replace this value in the budget constraint to calculate the optimal consumption at c_1. Formally,

$$\textstyle\sum_{t=1}^{T} \beta^{t-1} c_1 = W_1,$$

or

$$c_1^* = \frac{1}{\sum_{t=1}^{T} \beta^{t-1}} W_1,$$

Now, we can calculate the optimal consumption value for any period t. Formally,

$$c_t^* = \frac{\beta^{t-1}}{\sum_{t=1}^{T} \beta^{t-1}} W_1$$

This expression can be used to calculate the optimal consumption sequence as a function of β and W_1, that is $\{c_t^*\}_{t=1}^T$. The optimal consumption sequence can be used to calculate the value function, that determines the optimal return function given an initial value W_1. That is

$$v_T(W_1) = \sum_{t=1}^T \beta^{t-1} \log(\frac{\beta^{t-1}}{\sum_{t=1}^T \beta^{t-1}} W_1).$$

Rearranging terms,

$$v_T(W_1) = \sum_{t=1}^T \beta^{t-1} \log(\frac{\beta^{t-1}}{\sum_{t=1}^T \beta^{t-1}}) + \sum_{t=1}^T \beta^{t-1} \log(W_1)$$

where the value function is log-linear in W_1. We can rewrite this expression as

$$v_T(W_1) = A_T + B_T \log(W_1)$$

where $A_T = \sum_{t=1}^T \beta^{t-1} \log(\frac{\beta^{t-1}}{\sum_{t=1}^T \beta^{t-1}})$ and $B_T = \sum_{t=1}^T \beta^{t-1}$. The value function depends on the horizon of the problem. Since the utility function is concave, the utility function is increasing in t, that follows directly from the definition of concavity $u((1-\beta)c_1 + \beta c_2) < (1-\beta)u(c_1) + \beta u(c_2)$. We can see some specific cases:

1. $T = 1$: The implied value function is

$$v_1(W_1) = 1 \log(1) + \beta^0 \log(W_1) = 0 + 1 \log(W_1) = \log(W_1)$$

2. $T = 2$: In this case,

$$v_2(W_1) = \ln(\frac{1}{1+\beta}) + \beta \ln(\frac{\beta}{1+\beta}) + (1+\beta) \log(W_1),$$

3. $T = 3$: In this case,

$$v_3(W_1) = \log(\frac{1}{1+\beta+\beta^2}) + \beta \log(\frac{\beta}{1+\beta+\beta^2}) + \beta^2 \log(\frac{\beta^2}{1+\beta+\beta^2}) + (1+\beta+\beta^2) \log(W_1),$$

Example: We could use a different utility function, with constant elasticity of substitution $u(c) = c^{1-\sigma}(1-\sigma)^{-1}$ where $\sigma \geq 0$. The Euler equation implies

$$\left(\frac{1}{c_t}\right)^{-\sigma} = \left(\frac{\beta}{c_{t+1}}\right)^{-\sigma}$$

or

$$\gamma_c = \frac{c_{t+1}}{c_t} = \beta^{\frac{1}{\sigma}} < 1$$

Now the growth rate of consumption can be affected by the parameter σ, but we still have a decreasing sequence of consumption over time. We can redefine the discount rate to incorporate σ, so we have $\widetilde{\beta} = \beta^{\frac{1}{\sigma}}$, and obtain a similar solution. Formally

$$c_1 = \frac{1}{1 + \widetilde{\beta} + \widetilde{\beta}^2 + \widetilde{\beta}^3 + ... + \widetilde{\beta}^{T-1}} W_1$$

or more generally

$$c_t = \frac{(\beta^{\frac{1}{\sigma}})^{t-1}}{\sum_{t=1}^{T-1} (\beta^{\frac{1}{\sigma}})^{t-1}} W_1$$

In particular, when $\sigma = 1$ we obtain the previous solution with logarithmic preferences. When $\sigma \to 0$, then we have that $\gamma_c = \frac{c_{t+1}}{c_t} = \beta^{\frac{1}{\sigma}} = 0$, which can only be true when $c_{t+1} = 0$. With linear preferences we have

$$U = c_1 + \beta c_2 + \beta^2 c_3 + ... + \beta^{T-1} c_T$$

Clearly, it is optimal to consume all the cake in the first-period, because that maximizes the consumer life-time utility.

The examples illustrate that the solutions are consumption functions that depend on the initial size of the cake. Substituting all the optimal consumption sequence in the objective function we obtain a function that depends on W_1 and the number of periods. We call this a value function $v_T(W_1)$ and represents the maximum utility flow you obtain from the T-period optimization problem. Formally,

$$v_T(W_1) = \sum_{t=1}^{T} \beta^{t-1} u(c_t^*(W_1))$$

where $c_t^*(W_1)$ denotes the optimal level of consumption given the initial cake size W_1.

Envelope condition: Clearly, if we increase the size of the cake the lifetime utility should increase. The questions is how much? We can find it by taking the derivative with respect to W_1 and obtain

$$v_T'(W_1) = \lambda = \beta^{t-1} u'(c_t)$$

for all $t = 1, 2, ..., T$. It is increased by a constant amount given by the Lagrange multiplier of the resource constraint, and it does not matter when the additional cake is consumed because consumers are going to act optimally

1.1.2 Dynamic programming approach

Finite Horizon

Suppose we add an additional period $t = 0$ to our optimization problem. One way of solving it would be add the additional period and solve the sequence problem again. An alternative way would use the information that we already have from period 1 onwards. That would convert the $T + 1$ optimization problem into a two period problem. Given W_0 we have to solve

$$v_{T+1}(W_0) = \max_{c_0} u(c_0) + \beta v_T(W_1)$$

$$s.to. \qquad W_1 = W_0 - c_0$$

where W_0 is the initial size of the cake that we take as given. The choice of W_1 determines the size of the cake next period W_1. Once we know these two values we are done because we already know the solution from then on, because it is given by $v_T(W_1)$. It is not important how the cake will be consumed in the future. what it matters is that the agent will be optimizing and choosing $v_T(W_1)$ optimally. This is know as the **Principle of optimality**. If we substitute the constraint into the objective function we obtain

$$v_{T+1}(W_0) = \max_{c_0} u(c_0) + \beta V_T(W_0 - c_0)$$

and the first-order condition assuming that V_T is differentiable and concave is

$$u'(c_0) + \beta v'_T(W_1)(-1) = 0$$

or

$$u'(c_0) = \beta v'_T(W_1)$$

The change of consumption today affects the size of the cake in the future. The optimal choice implies that there cannot be any utility gains from deviating. We know from the envelope condition that

$$v_T'(W_1) = u'(c_1) = \beta^t u'(c_{t+1})$$

for all t. Combining this expressions we obtain the very familiar necessary condition for optimality, the Euler equation.

$$u'(c_0) = \beta^t u'(c_t)$$

Since the Euler equation holds for the other periods underlie the creation of the value function, one should suspect that solution to the $T+1$ problem using dynamic programming is identical to the sequence problem. The FOC of both problems are identical, thus the strict concavity of $u(c)$ ensures that the solutions will be identical as well.

Problems: The problem was simple because we had some information about $v_T(W_1)$. I general we will not have information about $v(\cdot)$. There are several ways we could construct the value function:

1. Start with a single period problem, and build the value function recursively by $v_1(W_1)$, then $v_2(W_1)$, and so on $v_T(W_1)$ for any T.
2. Guess a value function, and verify your guess. We will discuss this method later on.

Example: Next, we show how to construct the value function and the optimal decision rules recursively, when $u(c) = \ln c$. We begin with the last period and solve it backwards.

- $T = 1$,

$$v_1(W_1) = \max_{c_1} \ln(c_1)$$

$$s.to. \qquad W_2 = W_1 - c_1$$

The optimal choice implies, $c_1^* = W_1$ and $W_2 = 0$. The implied value function is

$$v_1(W_1) = \ln(W_1)$$

- $T = 2$

$$v_2(W_1) = \max_{W_2}\{\ln(W_1 - W_2) + \beta v_1(W_2)\}$$

given that we know the value function on the last period

$$v_2(W_1) = \max_{W_2}\{\ln(W_1 - W_2) + \beta \ln(W_2)\}$$

the FOC imply

$$\frac{1}{W_1 - W_2} = \frac{\beta}{W_2},$$

or

$$W_2^* = \frac{\beta}{1+\beta} W_1,$$

This is the optimal savings function in the first-period $W_2 = g(W_1)$, where in this case is a linear function. Note that in this case if $\beta = 0$, the consumer does not value consumption in the second period and the optimal savings for the next period is zero, that is $W_2 = 0$. We can compute the implied consumption decision rule for period 1 and 2

$$c_1^* = W_1 - W_2 = W_1 - \frac{\beta}{1+\beta} W_1$$

or

$$c_1^* = \frac{1}{1+\beta} W_1$$

and we can use the intertemporal constraint to compute second period consumption, that is $c_1 + c_2 = W_1$. But we already know the in the last period we have

$$c_2^* = W_2 = \frac{\beta}{1+\beta} W_1,$$

both consumption levels satisfy the resource constraint. Now we can compute the value function for the two-period problem.

$$v_2(W_1) = \ln(\frac{1}{1+\beta} W_1) + \beta \ln(\frac{\beta}{1+\beta} W_1)$$

rearranging terms we have

$$v_2(W_1) = A_2 + B_2 \ln(W_1)$$

where $A_2 = \ln(\frac{1}{1+\beta}) + \beta \ln(\frac{\beta}{1+\beta})$ and $B_2 = (1 + \beta)$. It is important that the value function does not include the $max(\cdot)$ operator because we are using the optimal decisions to construct it.

- $T = 3$, now we solve

$$v_3(W_1) = \max_{W_2}\{\ln(W_1 - W_2) + \beta(A_2 + B_2 \ln(W_2))\}$$

the FOC imply

$$\frac{1}{(W_1 - W_2)} = \frac{\beta B_2}{W_2}$$

or

$$W_2 = \frac{\beta B_2}{1 + \beta B_2} W_1$$

but we know that $B_2 = 1 + \beta$, so we have,

$$W_2^* = \frac{\beta + \beta^2}{1 + \beta + \beta^2} W_1,$$

again, we have calculated the optimal decision rule given W_1. Now, we can compute the consumption levels for all three periods by using the sequential resource constraint

$$\begin{aligned} W_2 &= W_1 - c_1 \\ W_3 &= W_2 - c_2 \\ W_4 &= W_3 - c_3 \end{aligned}$$

Then, we obtain

$$\begin{aligned} c_1^* &= \frac{1}{1 + \beta + \beta^2} W_1 \\ c_2^* &= \frac{\beta}{1 + \beta + \beta^2} W_1 \\ c_3^* &= \frac{\beta^2}{1 + \beta + \beta^2} W_1 \end{aligned}$$

Again, with the optimal decision rules we can compute the new value function

$$v_3(W_1) = A_3 + B_3 \ln(W_1)$$

where $A_3 = \ln(\frac{1}{1+\beta+\beta^2}) + \beta \ln(\frac{\beta}{1+\beta+\beta^2}) + \beta^2 \ln(\frac{\beta^2}{1+\beta+\beta^2})$ and $B_2 = (1 + \beta + \beta^2)$.

However, there is an important aspect of the problem that we might have missed. The dynamic programming approach gives us recursive decision rules, that only depend on last period size of the cake. In particular, another way to look at decision rules is

$$\begin{array}{lll} W_2 = \frac{\beta+\beta^2}{1+\beta+\beta^2} W_1 & \Rightarrow & c_1 = \frac{1}{1+\beta+\beta^2} W_1 \\ W_3 = \frac{\beta}{1+\beta} W_2 & \Rightarrow & c_2 = \frac{1}{1+\beta} W_2 \\ W_4 = 0 \cdot W_3 & \Rightarrow & c_3 = W_3 \end{array}$$

So when $\beta = 1$, we have

$$\begin{array}{lll} W_2 = \frac{2}{3} W_1 & \Rightarrow & c_1 = \frac{1}{3} W_1 \\ W_3 = \frac{1}{2} W_2 & \Rightarrow & c_2 = \frac{1}{2} W_2 \\ W_4 = 0 & \Rightarrow & c_3 = W_3 \end{array}$$

If we want to have every thing in terms of first-period cake size, we only need to move across the optimal decision rules. In particular

$$\begin{array}{l} c_1 = \frac{1}{3} W_1 \\ c_2 = \frac{1}{2}(\frac{2}{3} W_1) = \frac{1}{3} W_1 \\ c_3 = (\frac{1}{2}(\frac{2}{3} W_1)) = \frac{1}{3} W_1 \end{array}$$

As expected, with no discount it is optimal split the cake evenly over time. With a concave utility function this is optimal.

Infinite Horizon

Next, we assume that $T = \infty$. It is always useful consider the infinite sequence problem before studying the recursive formulation. In particular, our cake-eating problem becomes

$$v(W_0) = \max_{\{c_t\}_{t=0}^{\infty}, \{W_t\}_{t=1}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_t)$$

$$s.to. \qquad W_{t+1} = W_t - c_t, \qquad t = 1, ...,$$

It is direct to construct the Bellman equation from the sequential objective function. Developing the sum just one period ahead we obtain.

$$v(W_0) = u(c_0) + \sum_{t=1}^{\infty} \beta^t u(c_t) = u(c_0) + \beta \underbrace{\sum_{t=0}^{\infty} \beta^t u(c_{t+1})}_{v(W_1)}$$

so we have

$$v(W_0) = u(c_0) + \beta v(W_1)$$

In general, we can specify the problem for any point in the W space

$$v(W) = \max_{c \in [0,W]} \{u(c) + \beta v(W - c)\}$$

We use the variables with primes to denote future variables. The consumer starts the period with a cake of size W, and it has to choose present consumption c and that directly determines W'. The starting value for next period is given by the transition equation $W' = W - c$, and the future value for the size of the cake left is $V(W - c)$. The relevant variables for a dynamic programming problem are given by:

- **State variable:** The size of the cake given at the start of the period, W. The state completely summarizes all the information from the past needed for the forward looking optimization problem.

- **Control variables:** This is the variable that is being chosen. In this particular case the consumption level, c that lies in a compact set, $c \in [0, W]$. The size of the cake is irrelevant, in the sense that all the important information for future choices in summarized by the state variable.

- **Transition equation:** This equation determines the relation between present control and state variable today on tomorrow state variables. In particular, $W' = W - c$.

Alternatively, we can state the problem so we choose tomorrow's state variable W'.

$$v(W) = \max_{W' \in [0,W]} \{u(W - W') + \beta v(W')\}$$

Either specification should yield the same results, but in some cases it is easier to work with the last case. The present value function is a **functional equation (or Bellman equation)** that is the object that we want to solve. With an infinite number of periods, we cannot solve the problem backwards from the last period. To solve it we will use the fact that the value function appear in both sides, such that the functional equation always has to satisfy

$$v(W) = u(W - g(W)) + \beta v(g(W))$$

where $W' = g(w)$ is the policy function of the optimal decision rule associated to the problem. We will use iterative methods to find the correct value function. In particular, we could start with a guess of the value function $v_0(W)$, solve the optimization problem, and check if the implied value function is the same. Formally,

$$v_{i+1}(W) = u(W - g_i(W)) + \beta v_i(g_i(W))$$

if this is not the case, update the value function with the new one and iterate until it converges according to some norm $\sup |v_{i+1}(W) - v_i(W)| < \varepsilon$.

Now, we go back to our original problem and explore some properties of the optimal conditions.

$$v(W) = \max_{W' \in [0,W]} \{u(W - W') + \beta v(W')\}$$

the first-order conditions are given by

$$u'(c) = \beta v'(W')$$

This condition will hold if the value function is differentiable (might not be always the case). However, we can calculate the value using the envelope condition

$$v'(W) = u'(W - W') = u'(c)$$

since this condition holds for any period and for all W, it must also hold for W', so we have

$$v'(W') = u'(c')$$

Combining both terms we obtain the Euler equation.

$$u'(c) = \beta u'(c')$$

The optimal decision rules have to satisfy the necessary condition for all W. The solution of our dynamic programming problem are two functions that depend on the state of the economy W.

$$c = g_1(W),$$
$$W' = g_2(W) \equiv W - g_1(W),$$

The optimal policy functions have to satisfy the Euler equation. This is one procedure that we can use to check the accuracy of the solution, because the Euler equations are not generally used to compute the value function. Therefore, we only need to check that the implied errors are small.

$$u'(g_1(W)) = \beta u'(g_1(W')) = \beta u'(g_2(W - g_1(W))),$$

or

$$|u'(g_1(W)) - \beta u'(g_2(W - g_1(W)))| < \varepsilon$$

The equations need to hold for any point in the state space W.

Example: Next, we consider the infinite horizon version of the previous problem. The utility function is $u(c) = \ln c$. We don't know the true value function, but we conjecture based on the finite horizon solution that it has this form

$$v(W) = A + B\ln(W)$$

for all W. So we only need to determine the coefficients A and B that makes both sides of the value function be satisfied. The functional equation is given by

$$A + B\ln(W) = \max_{W' \in [0,W]} \{\ln(W - W') + \beta(A + B\ln(W'))\}$$

the first-order conditions imply

$$\frac{1}{W - W'} = \frac{\beta B}{W'}$$

or

$$W' = \frac{\beta B}{1 + \beta B} W$$

and

$$c = W - \frac{\beta B}{1+\beta B} W = W \left[1 - \frac{\beta B}{1+\beta B} \right]$$

finally

$$c = \frac{1}{1+\beta B} W$$

this is the policy function that depends on the state variable, and an unknown coefficient B. We need to find the coefficients that ensures both sides are equal. Substituting the optimal choice in the value function

$$A + B \ln(W) = \ln \frac{W}{1+\beta B} + \beta \left(A + B \ln \left(\frac{\beta B W}{1+\beta B} \right) \right)$$

Notice that we do need the max operator because we are substituting the optimal choice in. Collecting terms

$$A+B \ln(W) = (1+\beta B) \ln \left(\frac{1}{1+\beta B} \right) + \beta A + \beta B \ln \beta B + (1+\beta B) \ln(W)$$

it must be the case that

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

The B value that satisfies the second expression is

$$B = 1 + \beta B \Longrightarrow B = \frac{1}{1-\beta}$$

then, we can finally solve for

$$A = \frac{1}{1-\beta} \ln (1-\beta) + \beta A + \frac{\beta}{1-\beta} \ln \frac{\beta}{1-\beta}$$

or

$$A = \frac{1}{(1-\beta)} \left[\ln (1-\beta) + \frac{\beta}{1-\beta} \ln \beta \right]$$

With the optimal coefficients we can compute the optimal policy functions.

$$\begin{aligned} c &= (1-\beta)W, \\ W' &= \beta W, \end{aligned}$$

The optimal policy is always to save a constant fraction β of the cake, and eat the remaining fraction.

The solution of the infinite dimension problem can be usually obtained from the solution of the T period problem where $T = \infty$. Using the optimal decision rule we have

$$B_T = (1 + \beta + \beta^2 + ... + \beta^T) = \lim_{T\to\infty} \sum_{t=1}^{T} \beta^t = \frac{1}{1-\beta}$$

Infinite Horizon: Stochastic version of the cake-eating problem

Dynamic programming provides a convenient way of introducing uncertainty into the model. In particular, we could have to different ways of introducing uncertainty into the model

1. **Taste shocks:** The propensity to consumer changes every period with a certain probability. An alternative formulation could include changes in the discount rate β.

2. **Shocks law of motion or transition equation:** The cake depreciates at random rate, or it could even get bigger in some states of nature.

We focus on taste shocks and assume that the new utility function takes the form $U(c,\varepsilon) = \varepsilon u(c)$ where ε is a random variable. The timing of the problem is important, and we need to decide what does the agent know. Does he know the shock before making any choice? Should the choice be contingent on the shock? These are important elements that need to be defined ahead of time. For simplicity we consider a finite value for the taste shock $\varepsilon \in \{\varepsilon_1, \varepsilon_2, \varepsilon_3, ..., \varepsilon_S\}$ where $\varepsilon_1 < \varepsilon_2 < \varepsilon_3 < ... < \varepsilon_S$. We can assume that the shock follows a first-order Markov process, where $\pi_{ij} = \Pr ob(\varepsilon' = \varepsilon_j | \varepsilon = \varepsilon_i)$ and $\sum_S \pi_{iS} = 1$ for all i. Let Π denote the transition matrix that summarizes all the information about the probability of moving across states of nature.

The Bellman equation for this problem is given by

$$v(\varepsilon, W) = \max_{\{c,W'\}} \{\varepsilon u(c) + \beta E_{\varepsilon'|\varepsilon} v(\varepsilon', W')\}$$

$$\begin{aligned} s.to. \qquad W' &= W - c, \\ c, W' &\geq 0 \end{aligned}$$

The FOC of the problem are given by

$$\varepsilon u'(c) = \beta E_{\varepsilon'|\varepsilon} v'(\varepsilon', W')$$

for all W and ε. Using the envelope condition $v'(\varepsilon, W) = \varepsilon u'(c)$, and updating one period we have an stochastic Euler equation

$$\varepsilon u'(c) = \beta E_{\varepsilon'|\varepsilon} \varepsilon' u'(W' - W'')$$

The optimal decision rule of the problem is

$$W' = g(W, \varepsilon)$$

so we can rewrite the Euler equation as

$$\varepsilon u'(W - g(W, \varepsilon)) = \beta E_{\varepsilon'|\varepsilon} \varepsilon' u'(g(W, \varepsilon) - g(g(W, \varepsilon), \varepsilon'))$$

The optimal decision rule has to satisfy this equation for all point in the state space W and ε.