Lecture Notes on Dynamic Programming

Economics 200E, Professor Bergin, Spring 1998

Adapted from lecture notes of Kevin Salyer and from Stokey, Lucas and Prescott (1989)

Outline

1. A Typical Problem

Consider the problem of optimal growth (Cass-Koopmans Model). Recall that in the Solow model the saving rate is imposed, and there is no representation of preferences. The optimal growth model adds preferences for households, and derives an optimal saving rate. Utility is maximized for the representative agent, given the technology that they're faced with. The social planner's problem may be described as follows.

Preferences are summarized in the utility function:

$$U\left(c_t\right), t = 0, ..., \infty$$

Utility is assumed to be time-separable, that is, marginal utility of consumption today depends on today's consumption only. When households evaluate utility in the future, they discount it by a constant factor $\beta < 1$, assuming consumption in the future is not valued as much as consumption today. The objective is to maximize the present discounted value of future utility:

$$\sum_{t=0}^{\infty} \beta^t U\left(c_t\right)$$

Consider the technology. Output is produced using capital as the input:

$$y_t = f\left(k_t\right)$$

We could also include labor as a factor. Note that we will use k_t to represent capital available at the beginning of period t; this is capital that was accumulated in the previous period $t-1$.

Capital accumulation takes place through investment and saving. The law of motion for the capital stock is:

$$k_{t+1} = k_t\left(1-\delta\right) + i_t$$

Where i_t is the amount of investment expenditure in a period toward building up the capital stock for the following period. Depreciation is represented by δ, and is the faction of existing capital stock that decays away each period. We will initially assume there is 100% depreciation.

$$\delta = 1$$

Investment expenditure and consumption expenditure must both come from current production, so the resource constraint for the social planner is

$$c_t + i_t = f(k_t)$$

Combining this with the law of motion, we will be using as our budget constraint the following:

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

With complete depreciation, the capital stock available in a period is derived completely from saving in the previous period. So it is convenient to regard to the variable k_{t+1} simply as the level of saving in period t.

The social planner problem may be written:

$$\max_{\{c_t, k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t U[c_t] \tag{1}$$

$$s.t.\ c_t + k_{t+1} = f(k_t)$$

The social planner chooses consumption and saving in each period to maximize utility of the representative household. The solution is a sequence of variables c_t and k_{t+1} for all time periods $t = 0, ..., \infty$. Finding an infinite sequence of variables is a big task. But fortunately, this problem has a recursive structure: the problem facing the social planner in each period is the same as that he faced last period or next period.

We will characterize the solution by a function called a policy rule, which tells what the optimal choice is as a function of the current state of the economy. In this case we will find a rule for choosing c_t and k_{t+1} as a function of k_t, which applies in each and every period.

2. A Deterministic Finite Horizon Problem

2.1 Finding necessary conditions

To develop some intuition for the recursive nature of the problem, it is useful first to consider a version of the problem for a finite horizon. Assume you die in a terminal period T. We will then consider using as a solution for the infinite horizon problem the solution we found for the finite horizon problem, when we take a limiting case as $T \to \infty$.

The problem now may be written, where we substitute the constraint into the objective function, and thereby eliminate the consumption as a variable to choose:

$$\max_{\{k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{T} \beta^t U \left[f \left(k_t \right) - k_{t+1} \right]$$

Look at a section of the sum, which pertains to a generic period:

$$... + \beta^t U \left(f \left(k_t \right) - k_{t+1} \right) + \beta^{t+1} U \left(f \left(k_{t+1} \right) - k_{t+2} \right) + ...$$

This includes all the appearances for k_{t+1}. Take a derivative with respect to k_{t+1}.

$$\beta^t U' \left(f \left(k_t \right) - k_{t+1} \right) (-1) + \beta^{t+1} U' \left(f \left(k_{t+1} \right) - k_{t+2} \right) f' \left(k_{t+1} \right) = 0 \text{ for } t < T$$

or cancelling extra β :

$$U' \left(f \left(k_t \right) - k_{t+1} \right) = \beta U' \left(f \left(k_{t+1} \right) - k_{t+2} \right) f' \left(k_{t+1} \right) \text{ for } t < T$$

This is a necessary first order condition that applies to all periods prior to the terminal period. It describes the nature of the intertemporal decision I am making about whether I should consume or save. It says that I will raise consumption until the point where if I raise consumption one more unit today, the gain in utility no longer exceeds the loss in utility tomorrow, because there is less capital and hence less output and consumption tomorrow.

But this condition does not apply to the terminal period. In the last period of your life, you will not save for the future, but instead consume all that you produce:

$$k_{T+1} = 0$$

This is regarded as a boundary condition for the problem.

2.2 A special case

This problem can be solved analytically for a specific case with particular functional forms. Consider a log utility:

$$U(c_t) = \ln c_t$$

And a Cobb-Douglas production function:

$$f(k_t) = k_t^{\alpha}$$

The first order condition above then becomes:

$$\frac{1}{k_t^{\alpha} - k_{t+1}} = \beta \left(\frac{1}{k_{t+1}^{\alpha} - k_{t+2}} \right) \alpha k_{t+1}^{\alpha-1}$$

This is a second-order difference equation which is difficult to solve. We need to make it into a fist-order difference equation by using a change in variable.

$$\text{Define } z_t = \text{savings rate at time t } = \frac{k_{t+1}}{k_t^{\alpha}}$$

This turns out to be useful here, because the utility function here implies a constant saving rate, regardless of the level of income.

The first order condition then can be written:

$$\frac{1}{k_t^{\alpha}} \left(\frac{1}{1 - z_t} \right) = \frac{\beta}{k_{t+1}^{\alpha}} \left(\frac{1}{1 - z_{t+1}} \right) \alpha k_{t+1}^{\alpha-1}$$

which may be simplified:

$$\frac{z_t}{1 - z_t} = \alpha\beta \left(\frac{1}{1 - z_{t+1}} \right)$$

or

$$z_{t+1} = 1 + \alpha\beta - \frac{\alpha\beta}{z_t}$$

This is a relation between tomorrow's savings rate and today's savings rate. It is not a solution in itself, because it is a relation between optimal choices in different periods, not a rule telling use the optimal choice as a function of the current state of the economy.

Graph this relationship as a convex curve on a set of axes marked z_t and z_{t+1}. Graph also a 45-degree line showing where $z_t = z_{t+1}$. There are two points where the lines intersect. If the economy finds its way to one of these points, it will stay there. One such point is where the saving rate is 1.0. This is implausible, since it says all output is saved and none is consumed. The other point where the saving rate is $\alpha\beta$, a fraction of income less than one. Verify that $z_t = \alpha\beta$ satisfies the condition above.

Recall that the problem is constrained by the boundary condition, that saving is zero in the terminal period. In the graph, this means we start at the point $z_T = 0$ and work our way backward in time, to find out what the solution is for the current period. The graph suggests we will converge to the point where $z_T = \alpha\beta$. Provided that the terminal period when you die is far enough in the future, this saving rate will be solution to the optimal saving problem for the current period.

2.3 Solving recursively

We can perform this recursive operation explicitly. Start at the boundary point: $z_T = 0$. Now solve for saving in the previous period, z_{T-1}using the first order condition above.

$$z_T = 0 = 1 + \alpha\beta - \frac{\alpha\beta}{z_{T-1}}$$

so

$$z_{T-1} = \frac{\alpha\beta}{1 + \alpha\beta}$$

Now plug this back into the first order condition for the previous period:

$$z_{T-1} = \frac{\alpha\beta}{1+\alpha\beta} = 1 + \alpha\beta - \frac{\alpha\beta}{z_{T-2}}$$

which produces:

$$z_{T-2} = \frac{\alpha\beta\left(1+\alpha\beta\right)}{1+\alpha\beta\left(1+\alpha\beta\right)}$$

If we keep moving backwards:

$$\lim_{T\to\infty} z_t = \lim_{T\to\infty} \frac{\sum_{s=1}^{T-t}\left(\alpha\beta\right)^s}{1+\sum_{s=1}^{T-t}\left(\alpha\beta\right)^s} = \frac{\alpha\beta/\left(1-\alpha\beta\right)}{1/\left(1-\alpha\beta\right)} = \alpha\beta$$

Since this is the solution for the current period for when the terminal period is infinitely far away, it would be a natural conjecture that this is also the solution to an infinite horizon problem, in which there is no terminal period. We will investigate this conjecture below.

3. A Deterministic Infinite Horizon Problem

3.1 Recursive Formulation

Lets consider again the infinite horizon problem:

$$\max_{\{c_t, k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t U\left[c_t\right] \qquad \text{(problem 1)}$$

$$s.t.\; c_t + k_{t+1} = f\left(k_t\right)$$

In general, we can't just find the solution to the infinite horizon problem by taking the limit of the finite-horizon solution as $T \to \infty$. This is because we cannot assume we can interchange the limit and the max operators:

$$\max \lim_{T\to\infty} \sum_{t=0}^{T} U\left(c_t\right) \neq \lim_{T\to\infty} \max \sum_{t=0}^{T} U\left(c_t\right)$$

So we will take a different approach, taking advantage of the problem's recursive nature, called dynamic programming. Although we stated the problem as choosing an infinite sequences for consumption and saving, the problem that faces the household in period $t = 0$ can be viewed simply as a matter of choosing today's consumption and tomorrows beginning of period capital. The rest can wait until tomorrow.

Recall that the solution we found for the finite-horizon problem suggested that the desired level of saving is a function of the current capital stock. That is, the rule specifies the choice as a function of the current state of the economy:

$$k_{t+1} = g\left(k_t\right)$$

Our goal in general will be to solve for such a function g, called a policy function.

Define a function $v(k_0)$, called the value function. It is the maximized value of the objective function, the discounted sum of all future utilities, given an initial level of capital in period $t = 0$

for k_0.

$$v(k_0) = \max_{\{c_t,k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t U[c_t]$$

Then $v(k_1)$ is the value of utility that can be obtained with a beginning level of capital in period $t = 1$ of k_1, and $\beta v(k_1)$ would be this discounted back to period $t = 0$.

So rewrite problem 1 above as:

$$\begin{aligned} v(k_0) =& \max_{c_0,k_1} \left[U(c_0) + \max_{\{c_t,k_{t+1}\}_{t=1}^{\infty}} \sum_{t=1}^{\infty} \beta^t U[c_t] \right] \\ =& \max_{c_0,k_1} \left[U(c_0) + \beta v(k_1) \right] \qquad \text{(problem 2)} \\ & s.t.\ c_0 + k_1 = f(k_0) \end{aligned}$$

If we knew what the true value function was, we could plug it into problem 2 above, and do the optimization over it, and solve for the policy function for the original problem 1. But we do not know the true value function.

For convenience, rewrite with constraint substituted into objective function:

$$v(k_0) = \max \{ U(f(k_0) - k_1) + \beta v(k_1) \}$$

This is called Bellman's equation. We can regard this as an equation where the argument is the function v, a "functional equation". It involves two types of variables. First, state variables are a complete description of the current position of the system. In this case the capital stock going into the current period, k_0 is the state variable. Second, control variables are the variables that must be chosen in the current period; here this is the amount of saving k_1. If consumption c_0 had not been substituted out in the equation above, it too would be a control variable.

The first order condition for the equation above is:

$$U'[f(k_0) - k_1] = \beta v'[k_1]$$

This equates the marginal utility of consuming current output to the marginal utility of allocating it to capital and enjoying augmented consumption next period.

3.2 Envelope theorem

We would like to get rid of the term v' in the necessary condition. Assume a solution for the problem exists, and it is just a function of the state variable:

$$k_1 = g(k_0)$$

So

$$v(k_0) = \max\{U(f(k_0) - k_1) + \beta v(k_1)\}$$

becomes (where we can drop the max operator):

$$v(k_0) = U(f(k_0) - g(k_0)) + \beta v(g(k_0))$$

Now totally differentiate (where everything is a function of k_0):

$$v'(k_0) = U'(f(k_0) - g(k_0)) f'(k_0) - U'(f(k_0) - g(k_0)) g'(k_0) + \beta v'(g(k_0)) g'(k_0)$$

$$v'(k_0) = U'(f(k_0) - g(k_0)) f'(k_0) - [U'(f(k_0) - g(k_0)) - \beta v'(g(k_0))] g'(k_0)$$

The FOC says the second term equals zero, so

$$v'(k_0) = U'(f(k_0) - k_1) f'(k_0)$$

Update one period

$$v'(k_1) = U'(f(k_1) - k_2) f'(k_1)$$

Or in a more compact form easier to remember, the envelope condition here is:

$$v'_{t+1} = U'_{t+1} f'(k_t)$$

We can use this to get rid of term in first order condition. The FOC then becomes:

$$U'[f(k_0) - k_1] = \beta U'(f(k_1) - k_2) f'(k_1)$$

This tells us that the marginal value of current capital, in terms of total discounted utility, is given by the marginal utility of using the capital in current production and allocating its return to current consumption.

3.3 Apply to our special case:

Lets use again $f(k_t) = k_t^\alpha$, and $u(c_t) = \ln c_t$. This time, let's solve by a Lagrangian instead of substituting the constraint into the objective function.

The problem is stated:

$$v(k_t) = \max_{c_t, k_{t+1}} \{\ln(c_t) + \beta v(k_{t+1})\}$$

$$s.t.\ c_t + k_{t+1} = k_t^\alpha$$

Rewrite this in Bellman form:

$$v(k_t) = \max_{c_t, k_{t+1}} \{\ln(c_t) + \beta v(k_{t+1})\} + \lambda_t (k_t^\alpha - c_t - k_{t+1})$$

Differentiate to derive the first order conditions:

$$\frac{1}{c_t} - \lambda_t = 0$$

$$\beta v'(k_{t+1}) - \lambda_t = 0$$

or combining them:

$$\frac{1}{c_t} = \beta v'(k_{t+1})$$

Let's derive the envelope condition in this case. In general you can use a short cut, but let's do it the longer way here. Write the solution to all variables as functions of the state variable: $c_t = c(k_t)$, $k_{t+1} = k(k_t)$. Then we may write the Bellman equation with all arguments as a function of k_t:

$$v(k_t) = \ln(c(k_t)) + \beta v(k(k_t)) + \lambda(k_t)(k_t^\alpha - c(k_t) - k(k_t))$$

Now differentiate with respect to k_t:

$$v'(k_t) = \frac{c'{}_t}{c_t} + \beta v'(k_{t+1}) k'_{t+1} + \lambda'_t [k_t^\alpha - c(k_t) + k_{t+1}] + \lambda_t \left[\alpha k_t^{\alpha-1} - c'{}_t - k'_{t+1}\right]$$

Eliminate the term that equals zero because of the constraint, and regroup:

$$v'(k_t) = c'{}_t \left(\frac{1}{c_t} - \lambda_t\right) + k'_{t+1} (\beta v'(k_{t+1}) - \lambda_t) + \lambda_t \alpha k_t^{\alpha-1}$$

The first two terms here are zero also, because of the first order conditions. Substitute for λ_t and update one period:

$$v'(k_{t+1}) = \frac{1}{c_{t+1}} \alpha k_{t+1}^{\alpha-1}$$

This is the envelope condition here.

A simpler way to get the envelope condition would be just to take the derivative of the original problem

$$v(k_t) = \max_{c_t, k_{t+1}} \{U(c_t) + \beta v(k_{t+1})\} + \lambda_t (f(k_t) - c_t - k_{t+1})$$

with respect to k_t. This gives you the same result, since all the other terms drop out in the end. Shifted up one period, this also would give us:

$$v'_{t+1} = \lambda_{t+1} f'(k_{t+1})$$

And when combined with the FOC: $\lambda_t = U'_t$, we get the same envelope condition.

$$v'_{t+1} = U'_{t+1} f'(k_t)$$

Now substitute the envelope condition into FOC:

$$\frac{1}{c_t} = \alpha\beta \frac{1}{c_{t+1}} k_{t+1}^{\alpha-1}$$

So the two necessary conditions for this problem are the equation above and the budget constraint:

$$c_t + k_{t+1} = k_t^\alpha$$

3.4 Solution by iterative substitution in our special case

In this special case we can solve explicitly for a solution. Rewrite the FOC and budget constraint:

$$\frac{k_{t+1}}{c_t} = \alpha\beta \frac{k_{t+1}^\alpha}{c_{t+1}}$$
$$\frac{k_t^\alpha}{c_t} - 1 = \frac{k_{t+1}}{c_t}$$

Substitute the FOC into the constraint to get a consolidated condition:

$$\frac{k_t^\alpha}{c_t} = 1 + \alpha\beta \frac{k_{t+1}^\alpha}{c_{t+1}}$$

Update one period the consolidated condition above and substitute it into itself:

$$\frac{k_t^\alpha}{c_t} = 1 + \alpha\beta \left(1 + \alpha\beta \frac{k_{t+2}^\alpha}{c_{t+2}}\right)$$

Do this recursively. Note that it is a geometric progression:

$$\frac{k_t^\alpha}{c_t} = 1 + \alpha\beta + (\alpha\beta)^2 + (\alpha\beta)^3 + ... = \frac{1}{1-\alpha\beta}$$

So the policy function is:

$$c_t = (1-\alpha\beta)\, k_t^\alpha$$

Note that this is the answer we guessed earlier, based on the finite horizon problem.

3.5 Other solution methods: Solution by conjecture:

In general, the functional forms will not allow us to get an analytical solution this way. There are several other standard solution methods:

1) solution by conjecture

2) solution by iteration

3) others we will talk about later

Suppose we suspect because of the form of the utility function that the amount that the household saves should be a constant fraction of their income, but we don't know what this fraction, θ, is:

$$k_{t+1} = \theta k_t^{\alpha}$$

or equivalently

$$c_t = (1 - \theta) k_t^{\alpha}$$

Divide the two equations above:

$$\frac{k_{t+1}}{c_t} = \frac{\theta}{1 - \theta}$$

and substitute this into the FOC:

$$\frac{k_{t+1}}{c_t} = \alpha\beta \frac{k_{t+1}^{\alpha}}{c_{t+1}}$$
$$\alpha\beta \frac{k_{t+1}^{\alpha}}{c_{t+1}} = \frac{\theta}{1 - \theta}$$

Substitute the consumption function for c_{t+1}:

$$\alpha\beta \frac{k_{t+1}^{\alpha}}{(1 - \theta) k_{t+1}^{\alpha}} = \frac{\theta}{1 - \theta}$$

so

$$\alpha\beta = \theta$$

We again reach the same solution as before:

$$c_t = (1 - \alpha\beta) k_t^{\alpha}$$

3.6 Solution by value-function iteration:

Another solution method is based on iteration of the value function. The value function actually will be different in each period, just as we earlier found the function $g(k)$ was different depending on how close we were to the terminal period. But it can be shown (but we do not show this here) that as we iterate through time, the value function converges, just as $g(k)$ converged in our earlier example as we iterated back further away from the terminal period. This suggests that if we iterate on an initial guess for the value function, even a guess we know is incorrect, the iterations eventually will converge to the true function.

Suppose we start with a guess for some period $T+1$:

$$v_0\left(k_{T+1}\right)=0$$

This is similar to assuming a terminal period in which I die. Our guess for the value function implies that the discounted value of all future utility is zero, which implies that I consume all wealth and save nothing in this period: $c_T=k_T^{\alpha}$, and $k_{T+1}=0$.

So in previous period:

$$v_1\left(k_T\right)=\ln k_T^{\alpha}$$

We put a subscript on the value function, because this is changing over time, as it converges to some function.

In the period before that, the problem is:

$$v_2\left(k_{T-1}\right)=\max\left\{\ln\left(c_{T-1}\right)+\beta v_1\left(k_T\right)\right\}$$

$$s.t.\ c_{T-1}+k_T=k_{T-1}^{\alpha}$$

or

$$v_2\left(k_{T-1}\right)=\max\left\{\ln\left(c_{T-1}\right)+\beta\ln k_T^{\alpha}\right\}$$
$$s.t.\ c_{T-1}+k_T=k_{T-1}^{\alpha}$$

Do the optimization:

$$L=\ln\left(c_{T-1}\right)+\beta\ln k_T^{\alpha}+\lambda\left(k_{T-1}^{\alpha}-c_{T-1}-k_T\right)$$

FOCs are:

$$\lambda=\frac{1}{c_{T-1}}$$

and

$$\lambda = \alpha\beta\frac{1}{k_T}$$

or

$$\frac{1}{c_{T-1}} = \frac{\alpha\beta}{k_T}$$

Substitute into budget constraint and get:

$$k_T = \frac{\alpha\beta}{1+\alpha\beta}k_{T-1}^{\alpha}$$

and

$$c_{T-1} = \frac{1}{1+\alpha\beta}k_{T-1}^{\alpha}$$

Now plug these solutions into the value function v_2:

$$\begin{aligned} v_2\left(k_{T-1}\right) &= \ln\left(\frac{1}{1+\alpha\beta}k_{T-1}^{\alpha}\right) + \beta\ln\left(\frac{\alpha\beta}{1+\alpha\beta}k_{T-1}^{\alpha}\right)^{\alpha} \\ &= \alpha\beta\ln\left(\alpha\beta\right) - \left(1+\alpha\beta\right)\ln\left(1+\alpha\beta\right) + \left(1+\alpha\beta\right)\ln k_{T-1}^{\alpha} \end{aligned}$$

Then write for previous period:

$$v_3\left(k_{T-2}\right) = \max\left\{\ln\left(c_{T-2}\right) + \beta v_2\left(k_{T-1}\right)\right\}$$

$$s.t.\ c_{T-2} + k_{T-1} = k_{T-2}^{\alpha}$$

and so on...

It can be shown that this sequence of functions converges to:

$$v\left(k_t\right) = \max\left\{\ln c_t + \beta\left(\left(1-\beta\right)^{-1}\left[\ln\left(1-\alpha\beta\right) + \frac{\alpha\beta}{1-\alpha\beta}\ln\alpha\beta\right] + \frac{\alpha}{1-\alpha\beta}\ln k_{t+1}\right)\right\}$$

Once we know the true value function, we can solve for the policy function.

The FOC still holds that says:

$$U'\left(c_t\right) = \beta v'\left(k_{t+1}\right)$$

Now we can say that:

$$v'(k_{t+1}) = \frac{\alpha}{1-\alpha\beta}\frac{1}{k_{t+1}}$$

So

$$\frac{1}{c_t} = \frac{\alpha\beta}{1-\alpha\beta}\frac{1}{k_{t+1}}$$

or

$$\frac{c_t}{k_{t+1}} = \frac{1-\alpha\beta}{\alpha\beta}$$

Combine with the resource constraint:

$$\frac{c_t}{k_{t+1}} = \frac{k_t^\alpha}{k_{t+1}} - 1$$

to get:

$$\frac{1-\alpha\beta}{\alpha\beta} = \frac{k_t^\alpha}{k_{t+1}} - 1$$

and so:

$$k_{t+1} = \alpha\beta k_t^\alpha$$

which is same solution as before.

Although this solution method is very cumbersome and computationally demanding, because it relies on explicit iteration of the value function, it has the advantage that it will always work. It is common to set up a computer routine to perform the iteration for you.

4. A stochastic Problem

4.1 Introducing uncertainty

One benefit of analyzing problems using dynamic programming, is that the method extends very easily to a stochastic setting, in which there is uncertainty. We will introduce uncertainty as affecting the production technology only, by including a random term in the production function:

$$y_t = z_t f\left(k_t\right) \text{ where } z_t \text{ is i.i.d.}$$

Here $\{z_t\}$ is a sequence of independently and identically distributed random variables. This technology shock varies over time, but it's deviations in different periods are uncorrelated. These shocks may be interpreted as technological innovations, crop failures, etc.

We assume that households maximize the expected utility over time. Assume utility takes the same additively separable form as in the deterministic case, but now future consumption is uncertain.

The social planner problem now is:

$$\max_{\{c_t, k_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t U\left[c_t\right]$$

$$s.t.\ c_t + k_{t+1} = z_t f\left(k_t\right)$$

Where the expectations operator E_0 indicates the expected value with respect to the probability distribution of the random variables, c_t, k_{t+1}, z_t over all t, based on information available in period $t = 0$.

Assume the timing of information and action is as follows: at the beginning of the period t the current value of z_t the exogenous shock is realized. So the value of technology and hence total output is known when consumption takes place and when the end-of-period capital k_{t+1} is accumulated. The state variables are now k_t and z_t. The control variables are c_t and $k_{t+1.}$

As before we can think about a social planner in the initial period choosing a whole sequence of consumption -saving pairs $\{c_t, k_{t+1}\}_{t=0}^{\infty}$. But now in the stochastic case, this is not a sequence of numbers but rather a sequence of contingency plans, one for each period. The choice of consumption and saving in each period is contingent on the realization of the technology and capital stock in that period. This information is available during each period when consumption and saving are executed for that period, but it is not available in the initial period when the initial decision is being made for all future periods. A solution is a rule that shows how the control variables are determined as a function of these state variables, both the capital stock and the exogenous technology shock.

$$c_t = c(k_t, z_t)$$

$$k_{t+1} = k(k_t, z_t)$$

Note that all the variables involved here are random variables now, because they are functions of the exogenous shock, which is a random variable.

The setup here is very similar to the deterministic case, except that the variables involved are random, and the decision involves expectations over these. We again can write the problem recursively in Bellman form:

$$v(k_t, z_t) = \max_{c_t, k_{t+1}} \{U(c_t) + \beta E_t v(k_{t+1}, z_{t+1})\}$$

$$s.t.\ c_t + k_{t+1} = z_t f(k_t)$$

Note that the expectations operator may be thought of as integrating over the probability distribution of the shock:

$$v(k_t, z_t) = \max_{c_t, k_{t+1}} \left\{U(c_t) + \beta \int v(k_{t+1}, z_{t+1}) h(z_{t+1}) dz_{t+1}\right\}$$

where $h(z_{t+1})$ is the distribution of the shock.

Write this as a Lagrangian:

$$v(k_t, z_t) = \max_{c_t, k_{t+1}} \{U(c_t) + \beta E v(k_{t+1}, z_{t+1})\} + \lambda_t [z_t f(k_t) - c_t - k_{t+1}]$$

Take first order conditions:

$$c_t : \; U'(c_t) - \lambda_t = 0$$
$$k_{t+1} : \; \beta E_t v'_{k_{t+1}}(k_{t+1}, z_{t+1}) - \lambda_t = 0$$

so

$$U'(c_t) = \beta E_t v'_{k_{t+1}}(k_{t+1}, z_{t+1})$$

Envelope condition:

$$\begin{aligned} v'_{k_t}(k_t, z_t) &= \lambda_t z_t f'(k_t) \\ &= U'(c_t) z_t f'(k_t) \end{aligned}$$

So the necessary condition becomes:

$$U'(c_t) = \beta E_t [U'(c_{t+1}) z_{t+1} f'(k_{t+1})]$$

This has the same interpretation as always: you equate the marginal benefit of extra consumption today to the marginal cost in terms of lost production and hence consumption tomorrow.

4.2 Our special case again

Again we can demonstrate a solution analytically for the special case of log utility and Cobb-Douglas technology:

$$\begin{aligned} f(k_t) &= z_t k_t^{\alpha} \\ \text{and } u(c_t) &= \ln c_t \end{aligned}$$

The necessary condition then becomes:

$$\frac{1}{c_t} = \beta E_t \left[\frac{1}{c_{t+1}} z_{t+1} \alpha k_{t+1}^{\alpha-1} \right]$$

In the deterministic version we found a solution was a constant saving rate of $\alpha\beta$, so that:

$$\begin{aligned} k_{t+1} &= \alpha\beta k_t^{\alpha} \\ \text{and } c_t &= (1-\alpha\beta)\, k_t^{\alpha} \end{aligned}$$

Let's conjecture a similar solution for the stochastic problem, of some constant saving rate, θ, where we take into consideration that income varies with the exogenous shock, z_t:

$$\begin{aligned} k_{t+1} &= \theta z_t k_t^{\alpha} \\ \text{and } c_t &= (1-\theta)\, z_t k_t^{\alpha} \end{aligned}$$

Plug these into the necessary condition above, to see if it is satisfied.

$$\begin{aligned} \frac{1}{(1-\theta)\, z_t k_t^{\alpha}} &= \alpha\beta E_t\left[\frac{1}{(1-\theta)\, z_{t+1} k_{t+1}^{\alpha}} z_{t+1} k_{t+1}^{\alpha-1}\right] \\ &= \alpha\beta \frac{1}{(1-\theta)\, k_{t+1}} \\ &= \alpha\beta \frac{1}{(1-\theta)\, \theta z_t k_t^{\alpha}} \end{aligned}$$

Which holds if we choose the constant saving rate to be $\theta = \alpha\beta$. (Note the expectations operator was dropped because the variable $k_{t+1}$is known in period t.)

4.3 Finding distributions

Because these solutions are for random variables, we would like to be able to characterize their distributions. To do this we need to assume a convenient distribution for the underlying stochastic shock. We assume that it follows a log normal distribution, with mean μ and variance σ.

$$\ln z_t \tilde{} N(\mu, \sigma^2)$$

Because z_t is distributed log normal and the saving level, k_{t+1}, is a function of it, saving must also be distributed log normal. We can characterize the distribution in terms of its mean and variance, which will be functions of μ and σ. First let's find the mean of the saving variable.

Take logs of the solution found above for saving:

$$\begin{aligned} k_{t+1} &= \alpha\beta z_t k_t^{\alpha} \\ \ln k_{t+1} &= \ln\alpha\beta + \ln z_t + \alpha\ln k_t \end{aligned}$$

Now exploit the recursive nature of the problem:

$$\begin{aligned} \ln k_t &= \ln\alpha\beta + \ln z_{t-1} + \alpha\ln k_{t-1} \\ &= \ln\alpha\beta + \ln z_{t-1} + \alpha\left(\ln\alpha\beta + \ln z_{t-2} + \alpha\ln k_{t-2}\right) \\ & \quad ... \\ &= \left(1+\alpha+\alpha^2+...+\alpha^{t-1}\right)\ln\alpha\beta \\ & \quad + \left(\ln z_{t-1} + \alpha\ln z_{t-2} + ... + \alpha^{t-1}\ln z_0\right) \\ & \quad +\alpha^t \ln k_0 \end{aligned}$$

As we move in time, the initial value of saving disappears from the expression, and the other expressions become a pair of geometric series. Take the limit as $t \longrightarrow 0$, then take the expected value to find the mean of its distribution.

$$\lim_{t\to\infty} E\left[\ln k_t\right] = \left(\frac{1}{1-\alpha}\right)\ln\alpha\beta + \mu\left(\frac{1}{1-\alpha}\right)$$

This is the mean of distribution for saving. It is a function of the mean of the underlying shock.

Now let's find the variance of the distribution of saving.

$$\text{var}\left(\ln k_t\right) = E\left[\left(\left(\ln k_t\right) - E\left(\ln k_t\right)\right)^2\right]$$

$$
\begin{aligned}
&= E\left[\left(\left(\ln\alpha\beta+\ln z_{t-1}+\alpha\ln k_{t-1}\right)-\left(\left(\frac{1}{1-\alpha}\right)\ln\alpha\beta+\mu\left(\frac{1}{1-\alpha}\right)\right)\right)^{2}\right] \\
&= E\left[\left(\left(\ln\alpha\beta+\ln z_{t-1}+\alpha\left(\ln\alpha\beta+\ln z_{t-2}+\alpha\ln k_{t-2}\right)\right)-\left(\left(\frac{1}{1-\alpha}\right)\ln\alpha\beta+\mu\left(\frac{1}{1-\alpha}\right)\right)\right)^{2}\right] \\
&= E\left[\left(\sum_{i=0}^{t-1}\alpha^{t-1-i}\left(\ln\alpha\beta+\ln z_i\right)-\sum_{i=0}^{t-1}\alpha^{t-1-i}\left(\ln\alpha\beta+\mu\right)\right)^{2}\right] \\
&= E\left[\left(\sum_{i=0}^{t-1}\alpha^{t-1-i}\left(\ln z_i-\mu\right)\right)^{2}\right]
\end{aligned}
$$

multiplying this out and taking limit:

$$
\begin{aligned}
& \lim_{t\to\infty} E\left[\left(\left(\ln k_t\right)-E\left(\ln k_t\right)\right)^{2}\right] \\
= \; & \lim_{t\to\infty}\sum_{i=0}^{t-1}\alpha^{2(t-1-i)}E\left(\ln z_i-\mu\right)^{2}+\sum_{i\neq j}2\alpha^{(t-1-i)}\alpha^{(t-1-j)}E\left(\ln z_i-\mu\right)\left(\ln z_j-\mu\right) \\
= \; & \lim_{t\to\infty}\sum_{i=0}^{t-1}\alpha^{2(t-1-i)}\text{var}\left(z_i\right) \\
= \; & \frac{\sigma^2}{1-\alpha^2}
\end{aligned}
$$

The second set of terms in the first line above is zero, since the shocks were assumed to have a zero covariance across periods. Again, the variance of the saving variable is a function of the variance of the underlying technology shock.

The steps in the analysis above will be applied throughout the course to other situations. Dynamic programming will be used to find the equilibrium policy functions, giving the control variables as functions of current state variables. Then by assuming a convenient distribution for the exogenous state variables, we can use the policy function to describe the distribution of the control variables.