PhD Macroeconomics:
Intro to Dynamic Programming

David Murakami

University of Milan and University of Pavia

5 April 2024

Introduction

- Here we will look at recursive methods, starting with a deterministic setting.
- With recursive methods, one looks for a policy function, a mapping from the initial conditions, given by the past or the present, to a set of decisions about what to do with the variables we can choose during this periods.
- Because these are normally infinite horizon problems, how we will want to behave in the future matters in determining what we want to do today.
- Since what one will want to do in the future matters and the whole future time path can be determined, the recursive methods we describe are also known as dynamic programming.

States and controls

- It is helpful to separate the set of variable that we are using into state variables and control variables.
- In some period t, the state variables are those whose values are already determined, either by our actions in the past or by some other process (such as nature).
- Normally, the capital stock that we inherit from the past must be considered a state variable. One might also think that the technology level in each period is determined by nature and therefore, in any period, the agents living in that economy must take it as a given.
- The control variables in period t are those variables whose values individuals explicitly choose in that period with the goal of maximising some objective function.

Simple Infinite Horizon Model

Robinson Crusoe model I

- Consider the simple version of the Robinson Crusoe (RC) model. In that model, Robinson Crusoe wants to maximise:

$$\max \sum_{i=0}^{\infty} \beta^i u(c_{t+i}),$$

subject to the constraints:

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

and

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

Robinson Crusoe model II

- We do some familiar rearranging:

$$\begin{aligned} i_t &= k_{t+1} + (1-\delta)k_t, \\ \implies c_t &= f(k_t) - i_t \\ &= f(k_t) + (1-\delta)k_t - k_{t+1}, \end{aligned}$$

  which allows us to write Robinson Crusoe's problem as:

$$\max \sum_{i=0}^{\infty} \beta^i u(f(k_{t+i}) + (1-\delta)k_{t+i} - k_{t+i+1}).$$

- In this case, k_{t+1} is the control variable in period t.

Robinson Crusoe model III

- Whatever our choice of a control variable, there must be enough budget constraints or market conditions so that the values of the rest of the relevant variables in period t are determined.
- What may be surprising is that the choice of control variables can matter in how easily we can solve our models. Some choices will simply be more convenient than others.

The value function I

- Assume that it is possible to calculate the value of the discounted value of utility that an agent receives when that agent is maximising the infinite horizon objective function subject to the budget constraints.
- For Robinson Crusoe, this value is clearly a function of the initial per worker capital stock, k_t. As shown above, we can write out a version of this problem where the RC economy is using the capital stock to be carried over to the next period, k_{t+1}, as the control variable.
- For that example, the value utility is equal to:

$$V(k_t) = \max_{\{k_s\}_{s=t+1}^{\infty}} \sum_{i=0}^{\infty} \beta^i u(f(k_{t+i}) + (1-\delta)k_{t+i} - k_{t+1+i}), \tag{1}$$

where we denote the value of the discounted utility by $V(k_t)$, to stress that it is a function of the value of the initial capital stock, k_t.

The value function II

- For any value of k_t, limited to the appropriate domain, the value of the value function, $V(k_t)$, is the discounted value of utility when the maximisation problem has been solved and when k_t was the initial capital stock.
- Since $V(k_t)$ is a function, its value can be found for any permitted value of k_t.
- The economy is recursive as mentioned above. In period $t+1$, the value of k_{t+1} is given (it's a state variable) and the problem to be solved is simply the maximisation of utility beginning in period $t+1$:

$$V(k_{t+1}) = \max_{\{k_s\}_{s=t+2}^{\infty}} \sum_{i=0}^{\infty} \beta^i u(f(k_{t+i+1}) + (1-\delta)k_{t+i+1} - k_{t+i+2}), \tag{2}$$

and its value, $V(k_{t+1})$, is a function of the stock of capital per worker at $t+1$.

The value function III

- By separating the period t problem from that of future periods, we can rewrite the value function of (1) as:

$$V(k_t) = \max_{k_{t+1}} \left[u\left(f(k_t) + (1-\delta)k_t - k_{t+1}\right) + \max_{\{k_s\}_{s=t+2}^{\infty}} \sum_{i=1}^{\infty} \beta^i u(f(k_{t+i}) + (1-\delta)k_{t+i} - k_{t+i+1} \right]$$

- Adjusting the timing of the second discounted term gives:

$$V(k_t) = \max_{k_{t+1}} \left[u\left(f(k_t) + (1-\delta)k_t - k_{t+1}\right) + \beta \max_{\{k_s\}_{s=t+2}^{\infty}} \sum_{i=0}^{\infty} \beta^i u(f(k_{t+i+1}) + (1-\delta)k_{t+i+1} - k_{t} \right.$$

- The second discounted term is nothing but the value function $V(k_{t+1})$ that we wrote in (2).

The value function IV

- Making the substitution, the value function in (1) can be written recursively as:

$$V(k_t) = \max_{k_{t+1}} \left[u\left(f(k_t) + (1-\delta)k_t - k_{t+1}\right) + \beta V(k_{t+1}) \right] \tag{3}$$

- Equations in this form are known as Bellman equations (Bellman, 1957).
- It presents exactly the same problem as in (1)), but written in recursive form.
- The value of the choice variable, k_{t+1}, is being chosen to maximise an objective function of only a single period. The period is reduced from one of infinite dimensions (picking many future k's) to one of only one dimension.
- But there is a problem: both the time t one-period problem, $u(\cdot)$, and the discounted value function evaluated at k_{t+1}, $\beta V(k_{t+1})$, are included.
- The value of $V(k_{t+1})$ is not yet known. If it were known, then the value of the function $V(k_t)$ would also be known – it is the same function – and solving the maximisation problem at time t would be trivial.

The envelope theorem I

- To proceed, we assume that the value function $V(\cdot)$ exists and has a first derivative.
- We can then proceed with the one-period maximisation problem (3) by taking the derivative with respect to k_{t+1} to yield the FOC:

$$\frac{\partial V(k_t)}{\partial k_{t+1}} = -u'\left(f(k_t) + (1-\delta)k_t - k_{t+1}\right) + \beta V'(k_{t+1}) = 0. \tag{4}$$

- Unfortunately, this isn't very helpful. We do not know how or what $V'(\cdot)$ looks like.
- Under certain conditions, and this model has been written so that the conditions hold, one can find the derivative of $V(\cdot)$ simply by taking the partial derivative of the value function as written in equation (3) with respect to k_t.
- Theorems that provide the sufficient conditions for getting a derivative and that tell us how to find it are called envelope theorems.

The envelope theorem II

- This partial derivative is:

$$\frac{\partial V(k_t)}{\partial k_t} = u'\left(f(k_t) + (1-\delta)k_t - k_{t+1}\right)\left(f'(k_t) + (1-\delta)\right).$$

- Now we can substitute this result into (4) for $V'(k_{t+1})$:

$$\frac{\partial V(k_t)}{\partial k_{t+1}} = -u'\left(f(k_t) + (1-\delta)k_t - k_{t+1}\right) + \beta u'(f(k_t)+(1-\delta)k_t - k_{t+1})(f'(k_t)+(1-\delta)) = 0,$$

  which gives us the familiar consumption Euler equation:

$$\frac{u'(c_t)}{u'(c_{t+1})} = \beta(f'(k_{t+1}) + (1-\delta)).$$

The envelope theorem III

- In the steady state, where $c_t = c_{t+1} = \bar{c}$, this Euler equation yields:

$$\frac{1}{\beta} - (1 - \delta) = f'(\bar{k}).$$

- Using recursive methods, we find that for a stationary state, the rental rate on capital is equal to the net interest rate implicit in the discount factor plus the depreciation rate.

A general version I

- Let $\boldsymbol{X}_t$ be a vector of the period t state variables and let $\boldsymbol{Y}_t$ be a vector of the control variables.
- Let $F(\boldsymbol{X}_t, \boldsymbol{Y}_t)$ be the time t value of the objective function that is to be maximised.
- Given initial values of the state variables, $\boldsymbol{X}_t$, the problem to be solved is:

$$V(\boldsymbol{X}_t) = \max_{\{\boldsymbol{Y}_s\}_{s=t+1}^{\infty}} \sum_{i=0}^{\infty} \beta^i F(\boldsymbol{X}_{t+i}, \boldsymbol{Y}_{t+i}),$$

subject to the set of budget constraints given by:

$$\boldsymbol{X}_{s+1} = G(\boldsymbol{X}_s, \boldsymbol{Y}_s),$$

for $s \geq t$.

A general version II

- Using the same recursive argument that we used above, we can write the value function as a Bellman equation:

$$V(\boldsymbol{X}_t) = \max_{\boldsymbol{Y}_t} \left[F(\boldsymbol{X}_s, \boldsymbol{Y}_s) + \beta V(\boldsymbol{X}_{t+1}) \right],$$

subject to the budget constraints:

$$\boldsymbol{X}_{s+1} = G(\boldsymbol{X}_s, \boldsymbol{Y}_s).$$

Or, we can combine the equations to write the problem with a single equation:

$$V(\boldsymbol{X}_t) = \max_{\boldsymbol{Y}_t} \left[F(\boldsymbol{X}_t, \boldsymbol{Y}_t) + \beta V(G(\boldsymbol{X}_t, \boldsymbol{Y}_t)) \right]. \tag{5}$$

The policy function I

- The solution to this problem gives the values of the control variables as a function of the time t state variables:

$$\boldsymbol{Y}_t = H(\boldsymbol{X}_t),$$

  which we call a policy function, since it describes how the control variables behave as a function of the current state variables.
- Equation (5) is really a functional equation, since it must hold for every value of $\boldsymbol{X}_t$ within the permitted domain.
- Since the policy function optimises the choice of the controls for every permitted value of $\boldsymbol{X}_t$, it must fulfil the condition:

$$V(\boldsymbol{X}_t) = F(\boldsymbol{X}_t, H(\boldsymbol{X}_t)) + \beta V(G(\boldsymbol{X}_t, \boldsymbol{Y}_t)), \tag{6}$$

  where maximisation is no longer required because it is implicit in the policy function, $H(\boldsymbol{X}_t)$.

The policy function II

- To find the policy function, $H(\boldsymbol{X}_t)$, we find the FOCs for the problem in equation (5) with respect to the control variables. The FOCs are:

$$\mathbf{0} = F_Y(\boldsymbol{X}_t, \boldsymbol{Y}_t) + \beta V'(G(\boldsymbol{X}_t, \boldsymbol{Y}_t))G_Y(\boldsymbol{X}_t, \boldsymbol{Y}_t), \tag{7}$$

  where $F_Y(\boldsymbol{X}_t, \boldsymbol{Y}_t)$ is the vector of derivatives of the objective function with respect to the control variables, $V'(G(\boldsymbol{X}_t, \boldsymbol{Y}_t))$ is the vector of derivatives of the value function with respect to the time $t+1$ state variables, and $G_Y(\boldsymbol{X}_t, \boldsymbol{Y}_t)$ is the vector of derivatives of the budget constraints with respect to the control variables.
- Take the derivative of the value function (5) with respect to taking time t state variables, $\boldsymbol{X}_t$, one gets the envelope theorem:

$$V'(\boldsymbol{X}_t) = F_X(\boldsymbol{X}_t, \boldsymbol{Y}_t) + \beta V'(G(\boldsymbol{X}_t, \boldsymbol{Y}_t))G_X(\boldsymbol{X}_t, \boldsymbol{Y}_t).$$

The policy function III

- If, as can frequently be done, the controls have been chosen so that $G_X(\boldsymbol{X}_t, \boldsymbol{Y}_t) = \mathbf{0}$, then it is possible to simplify this expression to:

$$V'(\boldsymbol{X}_t) = F_X(\boldsymbol{X}_t, \boldsymbol{Y}_t).$$

- The FOCs of (7) can be written as:

$$\begin{aligned} \mathbf{0} &= F_Y(\boldsymbol{X}_t, \boldsymbol{Y}_t) + \beta F_X(\boldsymbol{X}_{t+1}, \boldsymbol{Y}_{t+1}) G_Y(\boldsymbol{X}_t, \boldsymbol{Y}_t) \\ &= F_Y(\boldsymbol{X}_t, \boldsymbol{Y}_t) + \beta F_X(G(\boldsymbol{X}_t, \boldsymbol{Y}_t), \boldsymbol{Y}_{t+1}) G_Y(\boldsymbol{X}_t, \boldsymbol{Y}_t). \end{aligned}$$

- If the function $F_X(G(\boldsymbol{X}_t, \boldsymbol{Y}_t), \boldsymbol{Y}_{t+1})$ is independent of $\boldsymbol{Y}_{t+1}$, then this equation can be solved for the implicit function, $\boldsymbol{Y}_t = H(\boldsymbol{X}_t)$, which is the required policy function.

The policy function IV

- One can substitute this policy function into equation (6) and solve for the implicit value function $V(\cdot)$.
- If $F_X(G(\boldsymbol{X}_t, \boldsymbol{Y}_t), \boldsymbol{Y}_{t+1})$ is not independent of $\boldsymbol{Y}_{t+1}$, then one can solve for the stationary state as we did before, using the condition that $\boldsymbol{Y}_t = \boldsymbol{Y}_{t+1} = \bar{\boldsymbol{Y}}$.
- If it is not the case that $G_X(\boldsymbol{X}_t, \boldsymbol{Y}_t) = \boldsymbol{0}$, then an alternative solution method is to find an approximation to the value function numerically.
- Consider some initial guess for the value function, $V_0(\boldsymbol{X}_t)$. it doesn't matter very much what this initial guess is, and a convenient one is to assume that it has a constant value of zero.
- One can then calculate an updated value function, $V_1(\boldsymbol{X}_t)$, using the formula:

$$V_1(\boldsymbol{X}_t) = \max_{\boldsymbol{Y}_t} \left[F(\boldsymbol{X}_t, \boldsymbol{Y}_t) + \beta V_0(G(\boldsymbol{X}_t, \boldsymbol{Y}_t)) \right],$$

  and doing the maximisation numerically over a sufficiently dense set of values from the domain of $\boldsymbol{X}_t$.

The policy function V

- This maximisation defines, approximately, the function $V_1(\boldsymbol{X}_t)$. Using this new function, one can update again and get a new approximate value function $V_2(\boldsymbol{X}_t)$ using:

$$V_2(\boldsymbol{X}_t) = \max_{\boldsymbol{Y}_t} \left[F(\boldsymbol{X}_t, \boldsymbol{Y}_t) + \beta V_1(G(\boldsymbol{X}_t, \boldsymbol{Y}_t)) \right].$$

- Repeated application of this process results in a sequence of approximate value functions $\{V_i(\boldsymbol{X}_t)\}_{i=0}^{\infty}$. This sequence converges to the value function, $V(\boldsymbol{X}_t)$.

Returning to the Robinson Crusoe model I

- It is useful to write out the example economy showing how each component matches with the general version.
- For the Robinson Crusoe example we had k_t as the state variable so $\boldsymbol{X}_t = k_t$, and the capital stock at time $t+1$ was the control variable so $\boldsymbol{Y}_t = k_{t+1}$.
- The objective function is:

$$F(\boldsymbol{X}_t, \boldsymbol{Y}_t) = u(f(k_t) + (1-\delta)k_t - k_{t+1}),$$

  and the budget constraint is written so that the time $t+1$ state variable is:

$$k_{t+1} = \boldsymbol{X}_{t+1} = G(\boldsymbol{X}_t, \boldsymbol{Y}_t) = \boldsymbol{Y}_t = k_{t+1}.$$

Returning to the Robinson Crusoe model II

- The FOC for Robinson Crusoe is:

$$\begin{aligned} 0 &= F_Y(\boldsymbol{X}_t, \boldsymbol{Y}_t) + \beta V'(F(\boldsymbol{X}_t, \boldsymbol{Y}_t))G_Y(\boldsymbol{X}_t, \boldsymbol{Y}_t) \\ &= -u'(f(k_t) + (1-\delta)k_t - k_{t+1}) + \beta V'(G(\boldsymbol{X}_t, \boldsymbol{Y}_t)) \times 1. \end{aligned} \tag{8}$$

- Recall that the envelope theorem gives:

$$V'(\boldsymbol{X}_t) = F_X(\boldsymbol{X}_t, \boldsymbol{Y}_t) + \beta V'(G(\boldsymbol{X}_t, \boldsymbol{Y}_t))G_X(\boldsymbol{X}_t, \boldsymbol{Y}_t).$$

Returning to the Robinson Crusoe model III

- For our example, the derivative of the budget constraint with respect to the time t state variable is simply:

$$G_X(\boldsymbol{X}_t, \boldsymbol{Y}_t) = \frac{\partial \boldsymbol{X}_{t+1}}{\partial \boldsymbol{X}_t} = \boldsymbol{0},$$

  so that the envelope theorem condition can be simplified to:

$$V'(\boldsymbol{X}_t) = F_X(\boldsymbol{X}_t, \boldsymbol{Y}_t) = u'(f(k_t) + (1-\delta)k_t - k_{t+1})(f'(k_t) + (1-\delta)),$$

  and the derivative of the value function is defined in terms of functions that we known.
- We substitute this into (8) and get:

$$0 = -u'(f(k_t) + (1-\delta) - k_{t+1}) + \beta \left[u'(f(k_{t+1}) + (1-\delta)k_{t+1} - k_{t+2})(f'(k_{t+1}) + (1-\delta))\right].$$

Returning to the Robinson Crusoe model IV

- This second-order difference equation can be solved for the steady state, where $k_t = k_{t+1} = k_{t+2} = \bar{k}$, to give:

$$f'(\bar{k}) = \frac{1}{\beta} - (1 - \delta). \tag{9}$$

Robinson Crusoe model with a twist I

- The RC model can be written with different choices for the control variables.
- The state variable in this version is still time t capital, so $\boldsymbol{X}_t = k_t$, but one can choose time t consumption to be the time t control variable, $\boldsymbol{Y}_t = c_t$.
- So, our objective function is now:

$$F(\boldsymbol{X}_t, \boldsymbol{Y}_t) = u(c_t),$$

and the budget constraint is:

$$k_{t+1} = \boldsymbol{X}_{t+1} = G(\boldsymbol{X}_t, \boldsymbol{Y}_t) = f(k_t) + (1 - \delta)k_t - c_t.$$

Robinson Crusoe model with a twist II

- Writing out the model, we have the Bellman equation:

$$V(k_t) = \max_{c_t} \left[u(c_t) + \beta V(f(k_t) + (1-\delta)k_t - c_t) \right],$$

  where we have replaced the time $t+1$ state variable, $\boldsymbol{X}_{t+1} = k_{t+1}$, by the budget constraint in the Bellman equation.
- It should be clear that the problem given above is the exact same economic problem that we solved previously.
- This version is somewhat less convenient than the earlier RC model when we try to write out the condition from the envelope theorem.
- When we take the derivative of the budget constraint with respect to the time t state variable, we get:

$$\frac{\partial G(\boldsymbol{X}_t, \boldsymbol{Y}_t)}{\partial \boldsymbol{X}_t} = f'(k_t) + (1-\delta),$$

  and this is generally not equal to zero.

Robinson Crusoe model with a twist III

- If we then write out the envelope theorem condition, we get:

$$\begin{aligned} V'(\boldsymbol{X}_t) &= F_X(\boldsymbol{X}_t, \boldsymbol{Y}_t) + \beta V'(G(\boldsymbol{X}_t, \boldsymbol{Y}_t))G_X(\boldsymbol{X}_t, \boldsymbol{Y}_t) \\ &= \beta V'(f(k_t) + (1-\delta)k_t - c_t)(f'(k_t) + (1-\delta)), \end{aligned}$$

  and we have the derivative of the value function in terms of the derivative of the value function and some other terms, which is no improvement.
- One of the important tricks of working with the Bellman equation is to write out the objective function and the budget constraints so that one gets a convenient expression of the envelope theorem, that is, so that $G_X(\boldsymbol{X}_t, \boldsymbol{Y}_t) = \mathbf{0}$.
- Doing this usually means putting as much of the model as possible into the objective function and requires keeping the time t state variable out of the budget constraint.

Approximating the value function

- Now let's try to actually code this up.
- Suppose we have:
$$f(k_t) = k_t^{\alpha}, \quad \alpha \in (0, 1),$$
and
$$u(c_t) = \ln c_t.$$
- With: $\delta = 0.1$, $\alpha = 0.36$, and $\beta = 0.98$.

Recursive Stochastic Models

A simple stochastic growth model I

- Now let's add realisations of possible states of nature – denoted as elements from Ω.
- Say that Ω is comprised of finite elements A: $\{A_1, A_2\}$. Each can occur with some fixed probability.
- Let's return to the case of Robinson Crusoe, but now the production function is:

$$y_t = A^t f(k_t),$$

where we apply the usual assumptions to $f(k_t)$ and A^t can take on two values:

$$A^t = \begin{cases} A_1 & \text{w.p. } p_1, \\ A_2 & \text{w.p. } p_2, \end{cases} \tag{10}$$

and where $A_1 > A_2$.

A simple stochastic growth model II

- Capital grows with the following law of motion:

$$k_{t+1} = A^t f(k_t) + (1-\delta)k_t - c_t.$$

- At time 0, Robinson Crusoe wants to maximise an expected discounted utility function of the form:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t).$$

- Future consumption paths are represented by a kind of tree. Given some initial capital k_0, in period 0 there are two possible technology levels that could occur and two different amounts of production, represented by the ordered pair $\{A_1 f(k_0), A_2 f(k_0)\}$, with probabilities $\{p_1, p_2\}$.

A simple stochastic growth model III

- Depending on which state occurs in period 0, RC will choose some time 1 capital stocks of $\{k_1^1, k_1^2\}$.
- In period 1, production will be one of these four possibilities: $\{A_1 f(k_1^1), A_2 f(k_1^1), A_1 f(k_1^2), A_2 f(k_1^2)\}$, with probabilities $\{p_1 p_1, p_1 p_2, p_2 p_1, p_2 p_2\}$.
- Suppose that one can write the value of the maximum expected discounted utility given an initial capital stock of k_0, when the time 0 realisation of technology is A_1, as:

$$V(k_0, A_1) = \max_{\{c_t\}_{t=0}^{\infty}} \mathbb{E} \sum_{t=0}^{\infty} \beta^t u(c_t),$$

subject to the budget constraint for $t = 0$:

$$k_1 = A_1 f(k_0) + (1 - \delta) k_0 - c_0,$$

A simple stochastic growth model IV

and those for $t \geq 1$:

$$k_{t+1} = A^t f(k_t) + (1-\delta)k_t - c_t,$$

and the independent realisations of A^t given by (10). One could write a similar setup by replacing A_1 with A_2.

- Expected utility is a function of two state variables: capital inherited and the realised technology level. As shown previously, this expression can be written recursively as:

$$V(k_0, A^0) = \max_{c_0} \left[u(c_0) + \beta \mathbb{E}_0 V(k_1, A^1) \right],$$

subject to:

$$k_1 = A^0 f(k_0) + (1-\delta)k_0 - c_0.$$

A simple stochastic growth model V

- There is a subtle change in how the value function is written. It is now written as a function of the time 0 realisation of the technology shock. As this function is written, k_0 and A^0 are the state variables and c_0 is the control variable.
- The second part of the value function is written with the expectations term because given a choice for c_0 (and through the budget constraint of k_1), it will have a value of $V(k_1, A_1)$ with probability p_1, and a value of $V(k_1, A_2)$ with probably p_2.
- For any particular choice of $\hat{k}_1$ of the time 1 capital stock, the expectations expression is equal to:

$$\mathbb{E}_0 V(\hat{k}_1, A^1) = p_1 V(\hat{k}_1, A_1) + p_2 V(\hat{k}_1, A_2).$$

A simple stochastic growth model VI

- For any initial time period t, the problem can be written as:

$$V(k_t, A^t) = \max_{c_t} \left[u(c_t) + \beta \mathbb{E}_t V(k_{t+1}, A^{t+1}) \right],$$

subject to:

$$k_{t+1} = A^t f(k_t) + (1-\delta)k_t - c_t.$$

- Or, by making k_{t+1} as the control/choice variable, we can write the problem as:

$$V(k_t, A^t) = \max_{k_{t+1}} \left[u(A^t f(k_t) + (1-\delta)k_t - k_{t+1}) + \beta \mathbb{E}_t V(k_{t+1}, A^{t+1}) \right], \tag{11}$$

and the budget constraint (using the definition in the previous section) is:

$$k_{t+1} = G(\boldsymbol{X}_t, \boldsymbol{Y}_t) = k_{t+1}.$$

A simple stochastic growth model VII

- The solution to a stochastic recursive problem finds a function that gives the values of the control variables that maximises the value function over the domain of the state variables.
- Since the state variables include both the results of previous choices of control variables and the results of nature's choices of the value for the stochastic state variables, we call the solution function a plan and write it as:

$$k_{t+1} = H(k_t, A^t).$$

- The plan gives the optimising choice of the control variables in every period as a function of the regular state variables and of the states of nature. A plan fulfils the condition that:

$$V(k_t, A^t) = u(A^t f(k_t) + (1-\delta)k_t - H(k_t, A^t)) + \beta \mathbb{E}_t V(H(k_t, A^t), A^{t+1}).$$

A general version I

- Using the notation previously, we can write the value function as:

$$V(\boldsymbol{X}_t, \boldsymbol{Z}_t) = \max_{\{\boldsymbol{Y}_s\}_{s=t}^{\infty}} \mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} F(\boldsymbol{X}_s, \boldsymbol{Y}_s, \boldsymbol{Z}_s),$$

  subject to the budget constraints given by:

$$\boldsymbol{X}_{s+1} = G(\boldsymbol{X}_s, \boldsymbol{Y}_s, \boldsymbol{Z}_s),$$

  for $s \geq t$, where $\boldsymbol{X}_t$ is the set of regular state variables, $\boldsymbol{Z}_t$ is the set of state variables determined by nature (stochastic state variables), and $\boldsymbol{Y}_t$ are the control variables.
- As before, $F(\cdot)$ is the objective function and $G(\cdot)$ are the budget constraints.

A general version II

- This problem can be written recursively as a Bellman equation of the form:

$$V(\boldsymbol{X}_t, \boldsymbol{Z}_t) = \max_{\boldsymbol{Y}_t} \left[F(\boldsymbol{X}_s, \boldsymbol{Y}_s, \boldsymbol{Z}_s) + \beta \mathbb{E}_t V(\boldsymbol{X}_{t+1}, \boldsymbol{Z}_{t+1}) \right], \tag{12}$$

subject to:

$$\boldsymbol{X}_{t+1} = G(\boldsymbol{X}_t, \boldsymbol{Y}_t, \boldsymbol{Z}_t).$$

- The solution is a plan of the form:

$$\boldsymbol{Y}_t = H(\boldsymbol{X}_t, \boldsymbol{Z}_t),$$

where

$$V(\boldsymbol{X}_t, \boldsymbol{Z}_t) = F(\boldsymbol{X}_t, H(\boldsymbol{X}_t, \boldsymbol{Z}_t), \boldsymbol{Z}_t) + \beta \mathbb{E}_t V(G(\boldsymbol{X}_t, H(\boldsymbol{X}_t, \boldsymbol{Z}_t), \boldsymbol{Z}_t), \boldsymbol{Z}_{t+1}),$$

holds for all values of the state variables (including the stochastic state variables).

A general version III

- The FOCs for the problem in equation (12), and its budget constraints are:

$$0 = F_Y(\boldsymbol{X}_t, \boldsymbol{Y}_t, \boldsymbol{Z}_t) + \beta \mathbb{E}_t \left[V_X G(\boldsymbol{X}_t, \boldsymbol{Y}_t, \boldsymbol{Z}_t), \boldsymbol{Z}_{t+1}) G_X(\boldsymbol{X}_t, \boldsymbol{Y}_t, \boldsymbol{Z}_t) \right].$$

- When one is able to choose the controls so that $G_X(\boldsymbol{X}_t, \boldsymbol{Y}_t, \boldsymbol{Z}_t) = \boldsymbol{0}$, the above equation is:

$$V_X(\boldsymbol{X}_t, \boldsymbol{Z}_t) = F_X(\boldsymbol{X}_t, \boldsymbol{Y}_t, \boldsymbol{Z}_t),$$

  and the FOCs give the consumption Euler equation (in stochastic form):

$$0 = F_Y(\boldsymbol{X}_t, \boldsymbol{Y}_t, \boldsymbol{Z}_t) + \beta \mathbb{E}_t \left[F_X(G(\boldsymbol{X}_t, \boldsymbol{Y}_t, \boldsymbol{Z}_t), \boldsymbol{Y}_{t+1}, \boldsymbol{Z}_{t+1}) G_Y(\boldsymbol{X}_t, \boldsymbol{Y}_t, \boldsymbol{Z}_t) \right].$$

The value function for the simple economy I

- Let's write equation (11) as a pair of Bellman equations, one for each of the two possible time t realisations of A^t, as:

$$V(k_t, A_1) = \max_{k_{t+1}} \{u(A_1 f(k_t) + (1-\delta)k_t - k_{t+1}) + \beta [p_1 V(k_{t+1}, A_1) + p_2 V(k_{t+1}, A_2)]\},$$

  and

$$V(k_t, A_2) = \max_{k_{t+1}} \{u(A_2 f(k_t) + (1-\delta)k_t - k_{t+1}) + \beta [p_1 V(k_{t+1}, A_1) + p_2 V(k_{t+1}, A_2)]\},$$

  where we've written out the expected values in terms of their respective probabilistic outcomes.
- The iteration process requires choosing starting functions for both $V_0(k_t, A_1)$ and $V_0(k_t, A_2)$.

The value function for the simple economy II

- Given these initial functions, the functions from the first iteration, $V_1(k_t, A_1)$ and $V_1(k_t, A_2)$, are found by simultaneously calculating:

$$V_1(k_t, A_1) = \max_{k_{t+1}} \left\{ u(A_1 f(k_t) + (1-\delta)k_t - k_{t+1}) + \beta \left[p_1 V_0(k_{t+1}, A_1) + p_2 V_0(k_{t+1}, A_2) \right] \right\},$$

and

$$V_1(k_t, A_2) = \max_{k_{t+1}} \left\{ u(A_1 f(k_{t)} + (1-\delta)k_t - k_{t+1}) + \beta \left[p_1 V_0(k_{t+1}, A_1) + p_2 V_0(k_{t+1}, A_2) \right] \right\},$$

over the discrete subset of values of k_t.

The value function for the simple economy III

- To find the results of the next iterations, $V_2(k_t, A_1)$ and $V_2(k_t, A_2)$, we calculate:

$$V_2(k_t, A_1) = \max_{k_{t+1}} \left\{ u(A_1 f(k_t) + (1-\delta)k_t - k_{t+1}) + \beta \left[p_1 V_1(k_{t+1}, A_1) + p_2 V_1(k_{t+1}, A_2) \right] \right\},$$

and

$$V_2(k_t, A_2) = \max_{k_{t+1}} \left\{ u(A_2 f(k_t) + (1-\delta)k_t - k_{t+1}) + \beta \left[p_1 V_1(k_{t+1}, A_1) + p_2 V_1(k_{t+1}, A_2) \right] \right\}.$$

- Repeated iterations result in a sequence of pairs of functions $\{V_j(k_t, A_1), V_j(k_t, A_2)\}_{j=0}^{\infty}$ that converge to the desired pair of value functions, $\{V(k_t, A_1), V(k_t, A_2)\}$.

Let's code this I

- Use our parameters from before: $\delta = 0.1$, $\beta = 0.98$, $\alpha = 0.36$.
- The production function is $f(k_t) = k_t^{\alpha}$, and the utility function is $u(c_t) = \ln c_t$.
- Let $A_1 = 1.75$ with $p_1 = 0.8$ and $A_2 = 0.75$ with $p_2 = 0.2$.
- We choose initial guesses as $V_0(k_t, A_1) = 20$ and $V_0(k_t, A_2) = 20$.
- The first iteration round results in calculations for $V_1(k_t, A^t)$ of:

$$V_1(k_t, A_1 = 1.75) = \max_{k_{t+1}} \ln(1.75k_t^{0.36} + 0.9k_t - k_{t+1}) + 0.98 \times 20,$$

and

$$V_1(k_t, A_2 = 0.75) = \max_{k_{t+1}} \ln(0.75k_t^{0.36} + 0.9k_t - k_{t+1}) + 0.98 \times 20.$$

Let's code this II

- The second round $V_2(k_t, A^t)$ functions are found maximising:

$$V_2(k_t, 1.75) = \max_{k_{t+1}} \left\{ \ln(1.75k_t^{0.36} + 0.9k_t - k_{t+1}) + 0.98\left[0.8V_1(k_{t+1}, 1.75) + 0.2V_1(k_{t+1}, 0.75\right.\right.$$

and

$$V_2(k_t, 0.75) = \max_{k_{t+1}} \left\{ \ln(0.75k_t^{0.36} + 0.9k_t - k_{t+1}) + 0.98\left[0.8V_1(k_{t+1}, 1.75) + 0.2V_2(k_{t+1}, 0.75\right.\right.$$

Comments

- That's really all there is to it!
- The math theorems underlying why the Bellman equation "just works" is complex, but the practical application of it straightforward.
- There is a big problem with dynamic programming or solving models with value function iteration: curse of dimensionality.
- As such, perturbation became a much more favoured solution method to solve dynamic macro models (but it has its own pitfalls).
- Good books to learn more about dynamic programming are: Adda and Cooper (2003), McCandless (2008), and Ljungqvist and Sargent (2018), and check out `QuantEcon`.

References I

Adda, Jérôme, and Russell Cooper. 2003. *Dynamic Economics.* MIT Press.

Bellman, Richard. 1957. *Dynamic Programming.* Princeton University Press.

Ljungqvist, Lars, and Thomas J. Sargent. 2018. *Recursive Macroeconomic Theory.* 4th Edition. MIT Press.

McCandless, George. 2008. *ABCs of RBCs.* Harvard University Press.