Dave Backus / NYU

Linear-Quadratic Approximations to Dynamic Programs

September 30, 2004

These notes review linear-quadratic methods of approximating dynamic programming problems, with the stochastic growth model as a recurring example.

1 A dynamic program

We'll be concerned with the stationary dynamic program,

$$\max_{\{u_t, x_{t+1}\}} \quad E_0 \sum_{t=0}^{\infty} \beta^t f(x_t, u_t), \tag{1}$$

subject to

$$x_{t+1} \;=\; g(x_t, u_t, \varepsilon_{t+1}) \tag{2}$$

with x_0 given and $\{\varepsilon_t\}$ iid with zero mean and unit variance. We refer to x as the state, u as the control, f as the return function, and g as the law of motion. Both are smooth, concave, etc. This structure doesn't cover all cases of interest, but it works for lots of them, including the stochastic growth model. It's the standard model in Sargent's red book (*Dynamic Macroeconomic Theory*).

A recursive approach to the solution follows from the Bellman equation

$$J(x) \;=\; \max_u \{f(x,u) + \beta E J[g(x,u,\varepsilon')]\} \tag{3}$$

The first-order condition is

$$f_u(x,u) + \beta E\left\{g_u(x,u,\varepsilon')^\top J_x[g(x,u,\varepsilon')]\right\} \;=\; 0 \tag{4}$$

and the envelope condition is

$$J_x(x) \;=\; f_x(x,u) + \beta E\left\{g_x(x,u,\varepsilon')^\top J_x[g(x,u,\varepsilon')]\right\}. \tag{5}$$

The final terms in the last two equations reflect the chain rule, but they look a little different than usual because x and u are vectors. Here f_z is a (column) vector with elements $\partial f/\partial z_i$ ($z = x$ or u), g_z is a matrix with elements $\partial g_i/\partial z_j$, and J_x is a vector with elements $\partial J/\partial x_i$ (and $J_x^\top$ is its transpose). The expectation is over ε'.

Typically we solve the first-order condition for $u = h(x)$, referring to h as the decision rule. In what follows, we look for a linear function h that approximates the solution to the original problem. Given h, the dynamic properties of the model follow from $x' = g[x, g(x), \varepsilon']$.

2 Linear-quadratic problems

Linear-quadratic or LQ problems are those with quadratic returns and linear laws of motion. As a result, the value functions are quadratic functions of the state and the first-order conditions and decision rules are linear. LQ problems have the property that the decision rules for stochastic problems are the same as those for the associated deterministic problem ("certainty equivalence").

The canonical discounted linear-quadratic problem looks like this: The return function is

$$\begin{aligned} f(x,u) &= z^\top \begin{bmatrix} Q & S \\ S^\top & R \end{bmatrix} z \\ &= x^\top Q x + u^\top R u + 2x^\top S u, \end{aligned}$$

where $z = (x, u)$ and Q and R are symmetric. The law of motion is

$$x' = Ax + Bu + C\varepsilon'.$$

Since it has no impact on the answer, we'll solve the problem with $\varepsilon = 0$. We guess that the value function is quadratic: $J(x)$ can be written $x^\top P x + p$ for some matrix P and scalar p. The Bellman equation is then

$$\begin{aligned} x^\top P x &= \max_u \left\{ x^\top Q x + u^\top R u + 2x^\top S u + \beta (Ax + Bu)^\top P (Ax + Bu) \right\} \qquad (6) \\ &= \max_u \left\{ x^\top (Q + \beta A^\top P A) x + u^\top (R + \beta B^\top P B) u + 2x^\top (S + \beta A^\top P B) u \right\}. \end{aligned}$$

To find the first-order condition, note the following rules for differentiating matrix expressions (for vectors a, x, and a matrix A):

$$\begin{aligned} \frac{\partial a^\top x}{\partial x} &= \frac{\partial x^\top a}{\partial x} = a \\ \frac{\partial x^\top A x}{\partial x} &= (A + A^\top) x \;\; [= \; 2Ax \text{ if } A \text{ is symmetric}]. \end{aligned}$$

(In each case, the derivative is a column vector with elements $\partial / \partial x_i$. You can prove both results by simply writing out the relevant terms.) The first-order condition is

$$\left(R + \beta B^\top P B \right) u + \left(S^\top + \beta B^\top P A \right) x = 0,$$

which gives us a decision rule of the form $u = -Fx$, with

$$F = \left(R + \beta B^\top P B \right)^{-1} \left(\beta B^\top P A + S^\top \right).$$

The Bellman equation then gives us

$$P = Q + \beta A^\top P A - \left(\beta A^\top P B + S \right) \left(R + \beta B^\top P B \right)^{-1} \left(\beta B^\top P A + S^\top \right), \qquad (7)$$

which is referred to as the matrix Riccati equation. If we find a P satisfying this equation, we've effectively found a solution to the problem.

There are lots of ways to find P. One is brute force: start with an initial guess for P; plug the guess into the right-hand side of (7) to get a new value of P on the left; repeat until successive estimates of P are sufficiently similar. This works, but if β is close to one, it's pretty slow. I prefer a doubling algorithm that's a lot faster; see McGrattan (1990) or Anderson, Hansen, McGrattan, and Sargent (1996).

[Add: summary of the doubling algorithm.]

3 Linear-quadratic approximation

Since LQ problems are easily solved, we often use them as approximations of non-LQ problems. Typically this is based on a quadratic approximation of the return function and a linear approximation of the laws of motion, both approximations done in the neighborhood of the steady state. In this section, we take the more general problem of Section 1 and describe a method of approximating it by an LQ problem.

Steady state

We define the steady state as the stationary point of the controlled system with the noise turned off (ie, with $\varepsilon = 0$). We find the steady state by solving the law of motion, the first-order condition, and the envelope condition for $[\bar{x}, \bar{u}, J_x(\bar{x})]$. The law of motion (2) defines a stationary point as

$$\bar{x} \;=\; g(\bar{x}, \bar{u}, 0).$$

The first-order condition (again, with $\varepsilon = 0$) implies

$$f_u(\bar{x}, \bar{u}) + \beta g_u(\bar{x}, \bar{u}, 0)^\top J_x(\bar{x}) \;=\; 0.$$

And the envelope condition implies

$$J_x(\bar{x}) \;=\; f_x(\bar{x}, \bar{u}) + \beta g_x(\bar{x}, \bar{u}, 0)^\top J_x(\bar{x}).$$

That gives us the equations need to find the unknowns. An example follows shortly.

LQ approximation

We start with a quadratic approximation of the return function:

$$\begin{aligned} f(z) \;&\cong\; f(\bar{z}) + f_z(z)^\top (z - \bar{z}) + \frac{1}{2}(z - \bar{z})^\top f_{zz}(z)(z - \bar{z}) \\ &=\; f(\bar{z}) + f_x(z)(x - \bar{x}) + f_x(z)(x - \bar{x}) + \frac{1}{2}(x - \bar{x})^\top f_{xx}(x, u)(x - \bar{x}) \\ &\quad + \frac{1}{2}(x - \bar{x})^\top f_{uu}(x, u)(x - \bar{x}) + \frac{1}{2}(u - \bar{u})^\top f_{ux}(x, u)(x - \bar{x}) \end{aligned}$$

This doesn't quite fit the LQ setup, since we have the steady state return f and linear terms f_z. We incorpoprate them by augmenting the state with the additional variable "1". With this change, the approximate problem fits into the LQ setup:

$$\begin{bmatrix} Q & S \\ S^\top & R \end{bmatrix} = \begin{bmatrix} f & \frac{1}{2}f_x^\top & \frac{1}{2}f_u^\top \\ \frac{1}{2}f_x & \frac{1}{2}f_{xx} & \frac{1}{2}f_{ux} \\ \frac{1}{2}f_u & \frac{1}{2}f_{xu} & \frac{1}{2}f_{uu} \end{bmatrix}.$$

This step can be automated by computing the derivatives numerically or analytically using a symbolic math program. [More on this some other time.]

The next step is to use a linear approximation of the law of motion:

$$g(x,u,\varepsilon) \;\cong\; g(\bar{x},\bar{u},0) + g_x(x-\bar{x}) + g_u(u-\bar{u}) + g_\varepsilon \varepsilon.$$

Using the same trick as before (adding "1" to the state vector x), we use this to fill in the matrices (A, B, C):

$$A \;=\; \begin{bmatrix} 1 & 0 \\ g & g_x \end{bmatrix}, \quad B \;=\; \begin{bmatrix} 0 \\ g_u \end{bmatrix}, \quad C \;=\; \begin{bmatrix} 0 \\ g_\varepsilon \end{bmatrix},$$

where the top rows pertain to the constant 1.

You might guess that this approach is equivalent to using a quadratic approximation of the Bellman equation, but it's not. A quadratic approximation of $J[g(x,u)]$ includes second derivatives of g, which we did not use above. In some examples (see exercise 3) this makes a big difference.

4 Example: the stochastic growth model

Consider a variant of the stochastic growth model. Utility is $u(c)$, output is $y = f(k,z) = c + x$, and the laws of motion for capital and the technology shock are

$$k_{t+1} \;=\; (1-\delta)k_t + x_t \tag{8}$$

$$z_{t+1} \;=\; (1-\varphi)\bar{z} + \varphi z_t + \sigma\varepsilon_{t+1}, \tag{9}$$

with $\{\varepsilon_t\} \sim$ NID(0,1) and (z_0, k_0) given. We generally use the output equation to eliminate c from the problem and use x as the control. Thus the problem might be represented recursively by

$$J(k,z) \;=\; \max_x \; \{u[f(k,z)-x] + \beta E J[(1-\delta)k + x, (1-\varphi)\bar{z} + \varphi z + \sigma\varepsilon']\}\,. \tag{10}$$

The expectation is over ε'.

Steady state

To find the steady state we have to solve the programming problem, at least in part. I do this by dynamic programming on the nonstochastic problem with constant z. The Bellman equation is

$$J(k) \;=\; \max_x \;\{u[f(k,z)-x] + \beta J\left[(1-\delta)k + x\right]\}. \tag{11}$$

The first-order condition is

$$-u_c(c) + \beta J_k(k') = 0 \tag{12}$$

and the envelope condition is

$$J_k(k) = u_c f_k + \beta J_k(k')(1-\delta). \tag{13}$$

We now solve these relations for the steady state defined by $k = k'$. The envelope condition (5) gives us

$$1 = \beta\left[(1-\delta) + f_k(k,z)\right],$$

which defines the steady state value of k. The law of motion for capital, equation (8), then gives us the steady state value of c:

$$\delta k = f(k,z) - c.$$

Thus we find the steady state values of the state variable k and the control variable c, given values of the various parameters. Output y is implicitly defined by $y = f(k,z)$.

In practice we often do this in reverse, choosing parameters that accord with observed steady state values ("calibration"). From quarterly postwar data for US we find that the mean share of consumption of private output (ie, excluding government) is about 0.75, with the complementary share of 0.25 for investment, and the mean capital-output ratio is about 10 (2.5 for annual output). The law of motion for capital then implies a depreciation rate of

$$\delta \;=\; \frac{1-\bar{c}/\bar{y}}{\bar{k}/\bar{y}} \;=\; \frac{0.25}{10} = 0.025,$$

which is a reasonable approximation to what we see in the data. (Alternatively, we could choose δ from depreciation data and then compute the implied capital-output ratio.) The envelope condition, together with a normalization like $\bar{y} = 1$, then tells us what β must be.

Example. Suppose (as above) that the steady state ratios are $k/y = 10$ and $c/y = 0.75$, and a subset of parameters is preset: $\alpha = 0.36$ (capital's share), and $\gamma = 2$ (risk aversion). I'll ignore φ and σ, which don't appear in the deterministic problem. We then find that the capital-output ratio and consumption share imply $\delta = 0.025$ and $\beta = 0.9891$. With the normalization $y = 1$, this implies steady state $z = 0.4365$, which we hit by adjusting μ (an otherwise worthless parameter).

LQ approximation

The next step is to find an LQ problem that approximates ours. There are a variety of ways to do this, but common practice in real business cycle work is to eliminate consumption from the problem and use investment as the control. The return function becomes

$$u\left[f(k,z)-x\right],$$

and the laws of motion remain

$$\begin{aligned} k' &= (1-\delta)k + x \\ z' &= (1-\varphi)\bar{z} + \varphi z + \sigma\varepsilon'. \end{aligned}$$

The nice thing about this way of expressing the problem is that the laws of motion are already linear, so it fits right into the LQ framework.

We use one additional trick to fit the problem into the LQ mold. You might notice that if we use a second-order Taylor series expansion to approximate the objective function, we have linear as well as quadratic terms. The trick is to handle this by adding a constant "1" to the state vector, so that the cross terms are linear. Let us say, then, that the state vector at date t is $(1, k_t-\bar{k}, z_t-\bar{z})$. The quadratic approximation of the objective function then implies

$$\begin{bmatrix} Q & S \\ S^\top & R \end{bmatrix} = \begin{bmatrix} u & \frac{1}{2}u_c f_k & \frac{1}{2}u_c f_z & -\frac{1}{2}u_c \\ \dots & \frac{1}{2}\left[u_c f_{kk} + u_{cc}(f_k)^2\right] & \frac{1}{2}\left[u_c f_{kz} + u_{cc} f_k f_z\right] & -\frac{1}{2}u_{cc} f_k \\ \dots & \dots & \frac{1}{2}[u_c f_{zz} + u_{cc}(f_z)^2] & -\frac{1}{2}u_{cc} f_z \\ \dots & \dots & \dots & \frac{1}{2}u_{cc} \end{bmatrix}$$

(The matrix is symmetric, so I've skipped the lower triangle.) The understanding is that all functions are evaluated at the steady state. Similarly, the law of motion for the state vector is

$$\begin{bmatrix} 1 \\ k_{t+1}-\bar{k} \\ z_{t+1}-\bar{z} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1-\delta & 0 \\ 0 & 0 & \varphi \end{bmatrix} \begin{bmatrix} 1 \\ k_t - k \\ z_t - z \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} [x_t] + \begin{bmatrix} 0 \\ 0 \\ \sigma \end{bmatrix} [\varepsilon_{t+1}],$$

which defines the matrices A, B, and C.

Example. Using the parameter values of the last section, I find (not guaranteed) that

$$F = [\,0.00000 \quad -0.00110 \quad -1.6746\,],$$

implying the decision rule

$$x_t - \bar{x} \;=\; 0.00110(k_t-\bar{k}) + 1.6746(z_t-\bar{z}).$$

Thus increases in capital and productivity both lead to higher investment, the latter because capital is more productive, the former because it leads to higher saving, hence investment.

Log-LQ approximation

A slight variant is to do an LQ approximation in the logs of the relevant variables. You might try it and compare your answers.

Exercises

Exercise 1. Verify the matrices Q, R, and S in Section 2, and tell me if you find any mistakes.

Exercise 2. Consider a special case of the LQ problem with x and u both scalars and $S = 0$. Show that the riccati equation is quadratic in P and therefore has two solutions. Which solution is the right one, and why?

Exercise 3. Repeat the LQ approximation of growth model using c as the control rather than x. Do you get the same answer?

Exercise 4. Apply the LQ method to a variant of the stochastic growth model in which labor n is added to both utility and production:

$$\begin{aligned} U(c, n) &= c^{1-\gamma}/(1-\gamma) + \eta \log(1-n) \\ F(k, n, z) &= k^{\alpha}(zn)^{1-\alpha}. \end{aligned}$$

What are the matrices Q, R, and W? In what ways do the impulse response functions differ from those reported in Section 4?