Lecture 10 March 2008

In this final lecture we discuss some applications of dynamic choice models with "continuous" choices. There are many useful references to the methodology:

Adda and Cooper, chapters 2, 3, and 6

Daron Acemoglu's Growth Theory text (in progress) chapter 16

Christopher Carroll. "Lecture Notes on Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems" (available at his website)

Angus Deaton"Saving and Liquidity Constraints" Econometrica 59 (Sept 1991).

1. A (very) Little Theory (borrowing from Adda-Cooper and Acemoglu)

We will consider stationary discrete time problems with a state variable $x(t) \in X$, a compact subset of $\mathbb{R}^k$. The source of uncertainty is a 1st order Markov process $\{z(t)\}$, with $z(t) \in Z$, a discrete set. (The agent observes $z(t)$ at period t, but does not know future values.) The agent's flow payoff is

$$U(x(t), x(t+1), z(t)).$$

The agent chooses $x(t+1)$ in period t, subject to the constraint that

$$x(t+1) \in G(x(t), z(t)).$$

The so-called "sequence problem" as of time 0 is sometimes written

$$V^*(x(0), z(0)) = \max E_0 \sum_{t=0}^{\infty} \beta^t U(x(t), x(t+1), z(t)).$$

As written the r.h.s. is under-specified. Implicitly, however, what is meant is that at time $t = 0$, the agent is trying to find an optimal choice $x(1)$ for each value of $x(0)$ and $z(0)$, and a corresponding plan for future consumption choices, so as to maximize the discounted future flow value, subject to the constraints that $x(t+1) \in G(x(t), z(t))$, and that in all future periods, $x(t)$ will be selected according to the plan. In the case where $z(t)$ is a first order process any plan will relate the choice of $x(t+1)$ to current and lagged values of $z(t)$, so the sequence problem can be properly stated in terms of a sequence of such plans. (See Acemoglu chapter 16 for a clear way to do it).

The Bellman equation is

$$V(x, z) = \max_{y \in G(x,z)} U(x, y, z) + \beta E[V(y, z')|z]$$

where expectations in the last term are taken with respect to the Markov process that determines $z(t)$. (Note how much easier it is to write this than to even state the sequence problem). Under the assumptions that:

$U(x, y, z)$ is a continuous, bounded, concave function

$\beta < 1$

$G(x, z)$ is non-empty, compact, continuous, convex

there is a unique real-valued solution function $V(x, z)$, continuous and bounded in x for each z, that solves both the sequence problem and the Bellman equation (i.e., $V^*(x, z) = V(x, z)$). V is concave in x for each z. Moreover, there is a stationary policy function $y = \phi(x, z)$ that gives the optimal choice of $x(t+1)$ for each $x(t), z(t)$. The proof of existence and uniqueness of V uses the fact that the operator

$$T(V)(x,z) = \max_{y \in G(x,z)} U(x,y,z) + \beta E[V(y,z')|z]$$

is a contraction mapping. Thus, it is possible to calculate $V(x,z)$ by starting from an arbitrary "guess", and repeatedly applying the operator T until convergence. This is the key idea used in applications.

2. A simple application to consumption

Consider an agent with an infinite life who has a flow utility $u(c_t)$ in period t, and a discount factor $\beta < 1$. At the beginning of period t the agent has a stock of assets A_t inherited from the previous period, and receives income y_t. Future income is uncertain: for now we assume $y_t \in \{y_1, ...y_J\}$ with $P(y_t = y_j\} = \pi_j$ (so there is no serial correlation in the process). Assuming a real interest rate R, one could write out the problem using A_t and y_t as the state variables. This would fit nicely into the framework above.

$$A_{t+1} = R(A_t + y_t - c_t),$$

we could re-write flow utility as

$$u(A_t + y_t - A_{t+1}/R).$$

However, Deaton (1991) noted that it was easier to re-frame the problem in terms of the control variable c_t and the state variable $x_t = A_t + y_t$, which he called "cash-on-hand". Note that

$$x_{t+1} = R(x_t - c_t) + y_{t+1}.$$

Thus, the Bellman equation for this problem is

$$\begin{aligned} V(x) &= \max_c \quad u(c) + \beta E[V(R(x-c) + y') \\ &= \max_c \quad u(c) + \beta \sum_j V(R(x-c) + y_j)\pi_j. \end{aligned}$$

Note that if we had a solution to the Bellman functional equation, we would be able to find the agent's "consumption function" $c = c^*(x)$. In particular,

$$c^*(x) = \arg\max \quad u(c) + \beta \sum_j V(R(x-c) + y_j)\pi_j$$

for each value of x. Note that the problem as stated so far allows the agent's cash-on-hand stock to grow infinitely big or infinitely small. To prevent problems it is convenient to assume $\beta = R$ when studying infinitely lived consumers.

The basic approach to solving for V has 2 steps:

1) discretize the state space x into n values $x_1, x_2, ...x_n$

2) starting from an initial guess for the value of V at each point in the discretized state space, V_i^1 iterate the following "inner loop":

a) interpolate between grid points to get $V^k(x)$ for the entire range of x

b) update the guess of the value function at each gridpoint using the contraction mapping iteration:

$$V_i^{k+1} = \max_c \quad u(c) + \beta \sum_j V^k(R(x_i - c) + y_j)\pi_j.$$

There are 2 ways to perform the max at step 2. One is to discretize the possible values for c into some (relatively fine) grid, and do a simple search. An alternative is to conduct a numerical optimization. Note that $u(c)$ is a known function, and at each iteration

$$\Omega^k(c) = \sum_j V^k(R(x_i - c) + y_j)\pi_j$$

is also a known function, so the numerical optimization has to find a c to maximize

$$u(c) + \beta\Omega^k(c)$$

Assuming u is concave, and V^k is concave in x, this is a 1-dimensional concave programming problem (i.e., very easy).

Some Details.

a) How do we discretize x? With $\beta = R$ and a stationary i.i.d. income process consumers will not allow cash on hand to get too large or too small. So the grid of values for x does not have to be much wider than the range between y_1 and y_J. A standard approach is to adopt a relatively course grid, get the problem working, then narrow the grid and see how much the answer changes.

b) How do we interpolate V between the grid points? A linear interpolation is not very good because the f.o.c. for optimum c at grid point i requires

$$u'(c) = \beta\Sigma_j[V'(R(x_i - c) + y_j)\pi_j$$

Linear interpolation implies that the "derivatives" of the interpolated V function are step functions, so it will not be possible to find a unique c that is optimal for x_i. Adda and Cooper)pp. 54-55) recommend cubic splines that are fit so the interpolated V is continuous and has continuous first and second derivatives at the spline points (the $x_i's$). The coefficients of the cubic for each interval $[x_i, x_{i+1}]$ can be solved easily. (Some assumption is needed for the first and last interval.)

3. A consumption problem with serially correlated income shocks.

a) Prologue

Suppose that income in period t is not i.i.d., but follows a correlated process. The standard approach (introduced by G. Tauchen, 1986 Economics Letters) is to discretize y_t (as above) and assume a 1st order Markov process that "approximates" a serially correlated continuous process. For example, suppose we want to approximate an AR-1 income process:

$$y_t = a + \rho y_{t-1} + \epsilon_t$$

where $\epsilon_t \sim N(0, \sigma^2)$. Note that for this process $E[y_t] = \mu_y = a/(1-\rho)$, and $var[y_t] = \sigma_y^2 = \sigma^2/(1-\rho^2)$. To approximate this with a discrete 1st order markov model with N points of support, first find $N-1$ cut points k_j $(j = 1, ..N-1)$ such that

$$\Phi[\frac{k_{j+1} - \mu_y}{\sigma_y}] - \Phi[\frac{k_j - \mu_y}{\sigma_y}] = \frac{1}{N}$$

with $k_0 = -\infty$, and $k_N = \infty$. (This defines the boundaries so that the probability a $N(\mu_y, \sigma_y^2)$ falls in each bin is 1/N). Next, find the mean value of a $N(\mu_y, \sigma_y^2)$ within each bin. These values will be the points of support for the discrete process. If $\rho = 0$ we can stop. Otherwise, the last step is to define "transition probabilities" π_{ij} such that

$$\pi_{ij} = P(k_i < y_t < k_{i+1} | k_j < y_{t-1} < k_{j+1})$$

assuming that

$$\begin{pmatrix} y_{t-1} \\ y_t \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_y \\ \mu_y \end{pmatrix}, \sigma_y^2 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right)$$

This can be computed using the usual formulas (e.g. in Johnson and Kotz).

b) The Consumption Model

With serially correlated incomes, the state space has to include y, the discretized current value of income. The Bellman equation becomes

$$\begin{aligned} V(x, y_j) &= \max_c \quad u(c) + \beta E[V(R(x-c) + y', y') \mid y_j] \\ &= \max_c \quad u(c) + \beta \sum_\ell V(R(x-c) + y_\ell, y_\ell) \pi_{\ell j}, \end{aligned}$$

where $\pi_{\ell j}(y) = P(y_{t+1} = y_\ell | y_t = y_j)$. The agent's consumption function is $c = c^*(x, y)$. As in the uncorrelated income case, we discretize x into a set of points $\{x_i\}$. Now the discretized value function includes a second dimension, over the (discrete) values for y_j. At the k^{th} iteration of the contraction mapping the value function at $x = x_i, y = y_j$ is $V_{i,j}^k$. Updating proceeds as before:

a) interpolate between grid points x_i to get $V_{\cdot,j}^k(x)$ at all x

b) update the guess of the value function

$$V_{i,j}^{k+1} = \max_c \quad u(c) + \beta \sum_\ell V_{\cdot,j}^k(R(x_i - c) + y_\ell) \pi_{\ell j}.$$

4. Borrowing constraints.

Let's return to the case of i.i.d income and assume now that the agent can never have $c_t > x_t$. Deaton (1991) discusses this case at some length (see also A-C, pp. 156-159). Formally, this means we impose the condition $0 \leq c \leq x$ in the definition of the Bellman equation:

$$\begin{aligned} V(x) &= \max_{0 \leq c \leq x} \quad u(c) + \beta E[V(R(x-c) + y') \\ &= \max_{0 \leq c \leq x} \quad u(c) + \beta \sum_j V(R(x-c) + y_j) \pi_j. \end{aligned}$$

The solution for V can be approached as above using contraction-mapping iterations on V. The wrinkle is that the search for the optimal c at each discretized value x_i has to impose the restriction $c \leq x_i$. With this restriction in place, for low values of x, $c = x$. The grid for x has to be relatively fine in the lower tail in order to capture the critical value at which this occurs.

Deaton actually solves the problem a different way, looking for a contraction mapping that defines the marginal utility of income with different levels of cash-on-hand. Define

$\lambda(c) = u'(c)$. From the envelope theorem, if $c_t = c^*(x_t)$ is the optimal consumption choice given cash-on-hand x_t, then

$$V'(x_t) = u'(c^*(x_t)) = \lambda(c^*(x_t)).$$

Moreover, if the condition $c_t \leq x_t$ is not binding then c_t satisfies the intertemporal first order condition (Euler condition):

$$u'(c_t) = \beta R\ E[V'(x_{t+1})] = \beta R\ E_t[u'(c_{t+1})]$$

But if the constraint is binding,$c^*(x_t) = x_t$ and

$$u'(c_t) = u'(x_t) > \beta R\ E_t[u'(c_{t+1})]$$

Therefore

$$\lambda(c^*(x_t)) = \max[\ \ \lambda(x_t),\ \beta R\ E_t[\lambda(c^*(x_{t+1})]\ \ ].$$

Deaton looks for a stationary "marginal utility of money" function $p(x) = \lambda(c^*(x))$ that has this property, i.e.

$$\begin{aligned} p(x) &= \max[\ \lambda(x)\ ,\ \beta R \sum_j V'(R(x - c^*(x)) + y_j)\pi_j\ \ ] \\ &= \max[\ \lambda(x)\ ,\ \beta R \sum_j p(R(x - \lambda^{-1}(p(x)) + y_j)\pi_j\ ] \end{aligned}$$

Deaton argues that this functional equation is a contraction mapping, and iterates on successive values of $p(x)$.

5. Other Income Processes

Simple i.i.d or AR-1 income processes are not very good descriptions of individual income generating functions. One standard extension (introduced by Zeldes, 1989) is the following. Let Y_t represent income in period t. Assume

$$\begin{aligned} Y_t &= P_t U_t, \\ P_t &= G_t P_{t-1} N_t, \end{aligned}$$

where P_t represents "permanent income", which has a non-stochastic growth rate $\log G_t$, and an innovation N_t (which is usually assumed to be i.i.d. normal), and U_t represents a transitory shock. In the literature there are two processes that have been explored for U_t. Carroll (1999) considers the case where $\log U_t$ is normal and $E\ U_t = 1$. Carroll (1992) considered the case where $U_t = 0$ with probability p, and $U_t \sim \log Normal$, with probability $1 - p$.

As before, cash on hand evolves as

$$X_{t+1} = R(X_t - C_t) + Y_{t+1},$$

where C_t is consumption in period t. In this model there are 2 state variables, P_t and X_t. However, note that if we divide everything by P_t things work out pretty nicely. In particular

$$\begin{aligned} \frac{X_t}{P_t} &= \frac{R(X_{t-1} - C_{t-1})}{P_t} + \frac{Y_t}{P_t} \\ &= R\left(\frac{X_{t-1}}{P_{t-1}} - \frac{C_{t-1}}{P_{t-1}}\right)\frac{P_{t-1}}{P_t} + \frac{Y_t}{P_t} \\ &= \frac{R}{G_t N_t}\left(\frac{X_{t-1}}{P_{t-1}} - \frac{C_{t-1}}{P_{t-1}}\right) + \frac{Y_t}{P_t} \\ x_t &= r_t(x_{t-1} - c_{t-1}) + y_t \end{aligned}$$

where small letters denote values divided by permanent income at time t. Note that cash on hand per unit of permanent income depends on lagged cash on hand per unit of permanent income, the realization of current income per unit of permanent income, and the "interest rate" r_t. This transformation gets everything in terms of ratios, and eliminates the inherent non-stationarity of a lifecycle growth model, for example.