November 25th 2009 - Lecture Nineteen

Econ205A Macro - 20091125

Outline – Stochastic Models Continued. Dynamic Programming with Stochastic Models. Markov Processes as State of Nature Variable. Discrete and Continuous Time. Setting up a Stochastic Bellman Equation.

Stochastic Models Continued – Introduction to today's lecture

Stochastic Models allow us to talk about many different things that happen in the economy.

- Stochastic elements in people's preferences – bad weather results in different consumption than good
- Stochastic elements in firm's production – positive and negative shocks to technology. Weather, discoveries.
- Stochastic elements in government policy – taxes, government spending shocks, times of war.

One thing to keep in mind is that the old framework of solving for competitive equilibrium was like solving a stochastic problem, but for just one realized state. We now have potentially many different states, meaning potentially many different problems to solve individually for. Complexity added for each new state.

Dynamic Programming with Stochastic Models

The next step is to solve these models with dynamic programming – in a recursive fashion. The dynamic program is a nicer way to solve the problem – if you can keep track of the current state that you're at (k, h ...), as well as the current state of nature (s) and all other state variables, you're solving in effect for the entire group – adding little to now complication that we didn't have in a non-stochastic model.

Issue One – are there restrictions that need to be placed on the types of stochastic processes that go into this framework.

Last class we talked about how – if the stochastic processes that we have have a Markov Property – then that would be okay. Our stocks will be called z instead of s

$$z_t \text{ a first order Markov process with a stationary transition functions} \Rightarrow$$
$$z_{t+1} \text{ depends only on } z_t$$
$$-\text{note: this isn't overly restrictive :}$$
$$\text{any finite order Markov process can be written as a 1st order Markov Process}$$
$$\text{by expanding the state space.}$$

How do we do that?

$$\text{if the distribution of } z_{t+1} \text{ depends on } (z_t, z_{t-1}, \dots, z_{t-n}) \equiv \widehat{Z}_t$$

Z_t is the vector of various period's $z_t's$. That means that $\hat{z}_t$ is just a first order markov process.

Ideas continued on next page:…

In particular,

$$\hat{Z}_{t+1} = (z_{t+1},, z_t, \ldots, z_{t-n+1})$$

This means that $\widehat{Z_t}$ will contain all the information you need to get Z_{t+1}, and $\widehat{Z_{(t+1)}}$ – because it contains all of those other elements.

Example of an easy way to put a stochastic process into a recursive formulation.
We will consider two different cases that will work:

1) $z = \{z_1, z_2, \ldots, z_n\}$ *is a countable (finite) set.*

The state space z is equal to a countable (finite) set. An example of this might be two state-contingent utility functions, one for sunny days and one for rainy days.

In this case the state space and the shocks are going to be discrete. You can also define a transition matrix,

$$\begin{aligned}&\textit{We can define a transition matrix } [q_{ij}] = Q\\&\textit{where } q_{ij} \textit{ is the probability of moving from state i to state j}\end{aligned}$$

A matrix composed of these elements q_{ij}'s = Q
This is like a Markov transition matrix

Restrictions Needed

$$0 \le q_{ij} \le 1$$

$$\sum_j q_{ij} = 1 \quad \forall i$$

The sum of the q's over $j = 1$ –in fact the summation implies the inequality property above

This has been fore the discrete state version,

Now we'll do the continuous time version

2) $x \subset \mathbb{R}^k$ *is a rectangle that is* $\forall z \in Z$,

$$\textit{there exists a density function } q(z, z') \textit{ such that}$$

$$q(z, z') \ge 0 \ \textit{ and } \int q(z, z')dz' = 1$$

For a z, there exists a density function q that, given your state today z, q is going to give you the density or probability that you move from z to z'. Such that $q(z, z')$ is greater than or equal to zero & the integral of $q(z, z')$ over all of the states that you're transitioning to is equal to one.

$$\textit{also,} \quad \forall z', \quad q(\cdot, z') \textit{ is continuous in its first arguement}$$

In either of these cases we are able to use the tools we learned in dynamic programming.

$$Q = \begin{bmatrix} & prob\ of\ go\ to\ state\ 1 & prob\ go\ to\ state\ 2 \\ possible\ paths\ 1 & 1/3 & 2/3 \\ possible\ paths\ 2 & 1/2 & 1/2 \end{bmatrix}$$

What does the distribution of transitions probabilities looks like n states in the future?

Above is the one period transition probability. It tells you if you're in state one today what do your transition probabilities look like in the next day.

If you want a two period transition matrix you'd multiply by this matrix twice.

Two period transitions:

$$Q^2 = \begin{bmatrix} \frac{4}{9} & \frac{4}{9} \\ \frac{5}{12} & \frac{7}{12} \end{bmatrix}$$

n period transitions

$$Q^n = \cdots$$

If you take the limit as n approaches infinity of the Q^n transition matrix – you'll get the long run stable distribution of how much time people are spending in states.

Also, let's say you have shocks that are like IID (independent and identically-distributed) shocks. Then you'd have a matrix where every row is identical. The transition matrix doesn't depend on the state that you're in.

$$Under\ either\ case\ the\ Feller\ property\ holds$$

$$(Mf)(y,z) = \int f(y,z')\, q(z,z')\, dz'$$

$$If\ f\ is\ continuous, then\ (Mf)\ is\ continuous$$

This is a technical detail. If f is your value function, this is going to be your value function with your choice y tomorrow, your state tomorrow over all possible realizations of the stochastic process. This is like the second term in the Bellman equation (the first term is still the current period objective function). This is saying that the second term – the discounted value function with the random realization of shock – is going to be able to preserve the continuity property that we use in the non-random version.

$$Moreover, if\ f\ is\ (strictly)\ increasing\ in\ y\ then\ (Mf)\ is\ also\ strictly\ increasing\ -\ and\ if\ f\ is\ concave\ (Mf)\ is\ also\ concave$$

Example – Let-Us Consider the follow Stochastic Sequence Problem

$$the\ state\ is{:}\ \ s = (x, z)$$

x is the deterministic of the state and z is going to be the random part. If we had a deterministic neo classical growth model x would just be capital and z would be something that is random like state-contingent productivity or preference. z reflects shocks.

$$the\ objective\ function\ is{:}\ F(x, y, z)$$

The objective function depends on $x, y, \&\ z$. y is going to be the control, or (more accurately in this formulation) the state next period determined by the control.

We'll write down the sequence problem first & then write the Bellman equation

$$V^*(x_0, z_0) = \max_{\pi} E_0 \left\{ \sum_{t=0}^{\infty} \beta^t\ F\left(\pi_t(z^{t-1}), \pi_{t+1}(z^t), z_t\right) |z_0 \right\} + F(x_0, \pi_0, z_0)$$

$$where\ \ z^t = (z_0, z_1, \dots, z_t)$$

Writing down the maximized value, given x_0 (an initial x) and initial z. Maximizing over $\pi,$ where pi is an optimal choice of you state next period. Given randomness we need expectations at time zero of the sum from $time = 0\ to\ \infty$ - maximizing over the discounted sequence of these $\pi's.$ plus the current period's value function.

This ($\sum_{t=0}^{\infty} \beta^t\ F\left(\pi_t(z^{t-1}), \pi_{t+1}(z^t), z_t\right) |z_0$) is your entire discounted sum of the objective functions.

Normally with a basic Bellman we are maximizing over choices of capital in each period. Now, one's choices of capital in each period depend on the realizations of the state. That is, the amount of capital that one has will depend in general on the entire history of past states. –IF you wrote down a discrete time version of the neo classical growth model all that you'd write down is that we're maximizing over k_t. Then you'd get $k_t\ \&\ k_{t+1}$ etc etc. When we are doing this, the value of capital that you have in period t depends on the entire history of past shocks. (this is because, say you've had period after period of bad shocks resulting in a very low level of capital coming into period t). The amount of capital you depend where you are in the tree – which depends on all of these past realizations.

What we are solving for will be only one optimal plan as you go through the tree. Thus the notation we need to do that is that we need to have an optimal choice – so that in period zero we have an optimal choice of our capital going forward into period one. In period one, (and that depends on x_0) we have this optimal choice that we already made ($that's\ \ \pi_0(x_0)$) –so we already made the optimal choice there. We have a new realization z_1. And then we are going to make a new optimal choice π_t that is going to depend on the last couple of shocks because your π_{t-1} decision was made given previous $z's.$

This problem is recursive and we can also write it having π_t as a function of $z_t\ \&\ \ \pi_{t-1}$ (pi t minus one: the value of capital that we came into the period with)

E_0 is the value at time zero of what you expect to get in all future periods. It includes all the future realizations of z-shocks. And the information you have then is just z_0.

Here, a plan is

$\pi_0 \in \Gamma(x_0, z_0)$

$Then \quad y_{t+1} = \pi_t(z_t) \in \Gamma(\pi_{t-1}(z^{t-1}, z_t)$

You choice at time zero has to be an element of the feasible set – which depends on your state at time zero (x_0, z_0)

Keep in mind that we are choosing the sequence of $\pi' \ s \quad (\pi_1, \pi_2, \ldots)$. The optimal choice of those is the choice of $x_1, x_2, x_3, \ldots\ldots$ – so the choice of capital and other state variables.

the plan is a sequence of functions:

$\Pi = \{\pi_t(z^t)\}_{t=0}^{\infty}$

The ideal is that here we have a stochastic sequence problem. And this probably highlights that solving a sequence problem like this is a pain because you are solving for a particular contingent plan over time at each step. You have to keep track of what's happened in the past to know how capital has evolved to the current point where we're at. And depending on some random realization of states-of-nature we have to think about what we want to do with capital for the future.

Solving a problem like this in terms of contingent plans is much much more complicated than solving the deterministic problems we trained on.

The logic of this problem:

If you what the shock is today and you know what the value of x is today, that's the information needed to one will do tomorrow.

Since F, Γ, β, q *are time invariant ,*
then we can write the Bellman Equation

$$V(x,z) = \max_{y \in \Gamma(x,z)} \left[F(x,y,z) + \beta \int V(y,z')\, q(z,z')\, dz' \right]$$

· F - your objective function
· Γ - your feasibility constraint that depends on the state that you're at. But the constraint itself doesn't depend on time.
· β - your discounting. It doesn't change with time.
· q - the transition probability on the shocks don't change with time.

If all of these are time invariant then we can write the Bellman Equation.

$V(x,z) = \rightarrow$ It says that the value of being at a state (a state is the pair x, z) equals taking the maximum over the choice tomorrow (variable y – choosing the best y that's feasible) of

· $F(x,y,z)$
· $\beta \int V(y,z')\, q(z,z')\, dz'$ $\rightarrow$ we need to integrate over our value tomorrow of our choice tomorrow ($V(y,z')$). And the probability of getting z' tomorrow is $q(z,z')$.

Alternatively, in discrete space with q_{ij}, information set at time t

$$\beta \sum_j q_{ij}\ V(y, z_j) \qquad \xrightarrow[\text{more generally}]{\text{or}} \quad (\beta\, E_t(V_{t+1}(y, z_{t+1})))$$

(alternative to $\beta \int V(y,z')\, g(z,z')\, dz'$) - we are summing over any state that we can get tomorrow – over the probability of that state occurring ($\sum_j q_{ij}$) times your value if that state occurs ($V(y, z_j)$)

This whole term is the expected value tomorrow, which equals 1:20:00

This whole term is the expected value tomorrow – how do you calculate that? It's the probability of a certain state tomorrow (q_{ij}) multiplied by the value of realizing that state ($V(y, z_j)$),

To Summerize

$If\ F, \Gamma, \beta\ \ are\ as\ before\ \ (for\ deterministic\ dynamic\ programming)$
$and\ q\ satisfies\ required\ properties\ \ (case\ 1\ or\ case\ 2\ from\ earlier\ in\ class\ today)$
$then\ if\ V^*\ is\ \ the\ solution\ to\ the\ Bellman\ Equation$
$then\ it\ is\ unique,$
$V^* is\ bounded\ and\ continuous$
$and\ the\ optimal\ policy\ function\ g\ is$
$compact\ valued\ and$
$upper - hemming\ continuous$

$if\ \ F(x, y, z)\ is\ increasing\ in\ z$
$and\ z < \hat{z} \Rightarrow \Gamma(x, z)\ \subset \Gamma(x, \hat{z})$
$\forall\ x\ \ ,\qquad then\ V\ is\ increasing$

This says that if z were increasing that is doesn't make our choice set worse off

The plan for next lecture:

- We'll talk briefly about an alternative set-up
  - One restriction that we've placed so far is that we are choosing y next period. And the random shock gets realized. One thing that might be the case is that one's state next period could also depend on the shock next period. So if you choose y it enters your objective function, but it doesn't go directly into the state next period. That if you state next period in this alternate set-up depends on both your choice of y and what shock is realized next period.

  Gorry will show to us that that is actually not a problem. We can just define a law of motion that depends on the shock. The law of motion will just get plugged into the value function.

  We're going to start off next class by writing down a set up like that.
- We'll go through and summarize the results on how dynamic programming works.
- And what extra properties do we need on q for us to have V^*'s bounded and continuous. We want properties so that we have monotonicity, convexity, etc.
- Gorry wants to go over an example of writing down the Bellman Equation for the growth model. And doing a recursive competitive equilibrium. – That means both writing down the dynamic program for the social planner's problem stochastic growth model - and then going recursive competitive equilibrium.
- That will tie up the section on stochastic dynamic programming.
- Finally – an additional set of examples on this type of stuff – labor search. Set up some simple search models.
- ALL THIS NEXT WEEK SUCKERS