

14.451 Lecture Notes 9

Guido Lorenzoni

Fall 2009

1 Stochastic dynamic programming

Payoff function depends on x_t, x_{t+1} and an exogenous shock z_t :

$$F(x_t, x_{t+1}, z_t).$$

The constraint correspondence also depends on z_t :

$$x_{t+1} \in \Gamma(x_t, z_t).$$

The shock z_t follows a Markov chain: a finite set Z and transition probabilities with $P(z'|z) \geq 0$ for all $z, z' \in Z$ and

$$\sum_{z' \in Z} P(z'|z) = 1.$$

Now we will use the pair (x_t, z_t) as state variable for a recursive approach and write the functional equation

$$V(x, z) = \max_{y \in \Gamma(x, z)} F(x, y, z) + \beta \sum_{z' \in Z} V(y, z') P(z'|z). \quad (\text{FE})$$

Defining the sequence problem is cumbersome here: a plan is a map from sequences $z^t = (z_0, \dots, z_t)$ to actions x_t , we can write them as $x(z^t)$ and then write

$$E \left[\sum_{t=0}^{\infty} \beta^t F(x(z^t), x(z^{t+1}), z_t) \mid z_0 \right]$$

where the conditional expectation is formed using the probability distribution on z^t generated by the transition probabilities above (e.g. the probability of $z^t = (z_0, \dots, z_t)$ is equal to $P(z_t|z_{t-1}) P(z_{t-1}|z_{t-2}) \dots P(z_1|z_0)$).

Let us assume:

- X is a compact subset of R^l ;
- F is bounded and continuous;
- $\Gamma(x, z)$ is non-empty, compact-valued and continuous.

With a finite set Z the treatment of the principle of optimality is very similar to the non-stochastic case. Moreover, under our assumptions, we can use the contraction mapping theorem as in the non-stochastic case to show that a unique solution to (FE) $V(x, z)$ exists, is bounded and continuous. In this case the principle of optimality can be easily applied to show that $V(x, z)$ corresponds to the value function for the sequence problem.

Also, we can use the contraction mapping approach to characterize the value function $V(x, z)$ (making more assumptions we can establish monotonicity, concavity, differentiability).

Things are technically more challenging if we want to allow for continuous random variables, then we have a Markov process for z , with transition function: $Q(z, A)$ is the probability that the shock tomorrow is $z' \in A$ if the shock today is z . That is, $Q(z, \cdot)$ is a probability measure on the space Z (this requires that we define a σ -algebra \mathcal{Z} of measurable sets on Z , so we can properly define $Q(z, A)$ for all sets $A \in \mathcal{Z}$). SLP has a full treatment of the general case.

1.1 Stochastic Dynamics

The crucial difference here is the analysis of the dynamics: the solution to (FE) gives us a policy which now is a function of x_t and z_t (suppose from now on that the policy is unique)

$$x_{t+1} = g(x_t, z_t).$$

How do we analyze this policy?

- start at x_0, z_0 , and find $x_1 = g(x_0, z_0)$,
- then x_2 is a random variable, that takes the values $g(x_1, z_1)$ with probabilities $P(z_1|z_0)$,
- continue... $x_3, x_4, x_5 \dots$ all random variables,
- find probability distribution $\Pr((x_t, z_t) | (x_0, z_0))$ for all t .

Question is $\Pr((x_t, z_t) | (x_0, z_0))$ converging somewhere? Does it depend on the initial condition (x_0, z_0) ? What does it mean for a probability distribution to converge?

First, notice that the optimal policy, being recursive, gives us a new Markov process, not just on the space Z but on the space $S \equiv X \times Z$, with transition function:

$$Q((x, z), A) = \begin{cases} \sum_{\{z':(x',z') \in A\}} P(z'|z) & \text{if } x' = g(x, z) \\ 0 & \text{otherwise} \end{cases}.$$

So all our questions can be addressed studying the properties of this Markov process.

1.1.1 Finite state space (Markov Chains)

Let us start from the case where also X is finite, so the Markov process on $S \equiv X \times Z$ boils down to a Markov chain with transition probabilities

$$\pi((x', z') | (x, z)) = \begin{cases} P(z'|z) & \text{if } x' = g(x, z) \\ 0 & \text{otherwise} \end{cases}.$$

List the states $(x, z) \in S$ and index them by $i = 1, 2, \dots, m$ and then we have a $m \times m$ transition matrix Π and we can compute Π^n for any n .

Call $\pi_{i,j}^{(n)}$ the i, j element of Π^n : this is the probability of going from i to j in n periods.

Various behaviors:

- there can be *transient* states: a state j such that if $p_i = 0$ then $[p\Pi]_i = 0$ (once you are not in i you can never go back);
- there are self-replicating sets of states: a set of states A such that if $\sum_{i \in A} p_i = 1$ then $\sum [p\Pi]_i = 1$ (S is one);
- there are *ergodic* sets: a self-replicating set of states E such no proper subset of E is also self-replicating (at least one must exist);
- there can be cyclically moving subsets.

Some properties about convergence of Markov chains can be established in general. Here we focus on a convergence result which is nice because it let's us use once more the contraction mapping theorem.

Define set Δ^m as the set of probability distributions on S (i.e. the set of vectors $p \in R^m$ such that $p_i \geq 0$ and $\sum_{i=1}^m p_i = 1$).

The transition matrix Π defines a mapping

$$\Pi : \Delta^m \rightarrow \Delta^m$$

To use the contraction mapping theorem we need a metric on Δ^m we will use the total absolute deviation:

$$\|p - q\| \equiv \sum_{j=1}^m |p_j - q_j|. \quad (1)$$

Theorem 1 *Suppose for some state \hat{j} we have $\pi_{i,\hat{j}} > 0$ for all i , then S has a unique ergodic set with no cyclically moving subsets and $\{p_0 \Pi^n\}$ converges to a unique invariant distribution p^* for all $p_0 \in \Delta^m$.*

Proof. Let $\varepsilon_j = \min_i \pi_{i,j}$. Notice that $\pi_{i,j} - \varepsilon_j \geq 0$. Notice also that

$$\begin{aligned} [p\Pi - q\Pi]_j &= \sum_{i=1}^m p_i \pi_{i,j} - \sum_{i=1}^m q_i \pi_{i,j} = \sum_{i=1}^m (p_i - q_i) (\pi_{i,j} - \varepsilon_j) + \sum_{i=1}^m (p_i - q_i) \varepsilon_j = \\ &= \sum_{i=1}^m (p_i - q_i) (\pi_{i,j} - \varepsilon_j) \end{aligned}$$

because $\sum_{i=1}^m (p_i - q_i) = 1 - 1$. Then we have the following chain of inequalities

$$\begin{aligned} \|p\Pi - q\Pi\| &= \sum_{j=1}^m \left| \sum_{i=1}^m (p_i - q_i) (\pi_{i,j} - \varepsilon_j) \right| \\ &\leq \sum_{j=1}^m \sum_{i=1}^m |p_i - q_i| (\pi_{i,j} - \varepsilon_j) \leq \sum_{i=1}^m |p_i - q_i| \sum_{j=1}^m (\pi_{i,j} - \varepsilon_j) \\ &= (1 - \varepsilon_j) \|p - q\|. \end{aligned}$$

Since $\varepsilon_j > 0$, by assumption, the mapping Π is a contraction with modulus $(1 - \varepsilon_j)$ and has a unique fixed point p^* . Moreover for any p we have

$$\|p\Pi^n - p^*\| \leq (1 - \varepsilon_j)^n \|p - p^*\|$$

which implies that $p\Pi^n$ converges to p^* for any initial condition p . ■

It is easy to extend this result to a more general condition: there is a state \hat{j} such that for some $n \geq 1$ we have $\pi_{i,\hat{j}}^{(n)} > 0$ for all i . That is, there is a state that can be reached with positive probability from any initial condition after n periods.

Then use the fact that the matrix

$$\tilde{\Pi} = \Pi^n$$

is also a transition matrix and apply the theorem above. (See Theorems 11.3 and 11.4 in SLP).

1.1.2 Weak convergence

We won't go into the general case where X (or also Z) are subsets of R^l . However, the idea of the theorem above can be extended (see Theorem 11.12 in SLP) and many other type of convergence results are also available when the conditions for that theorem are not met (check out Chapter 11).

It is useful to know that often people use a convergence notion for probability measures on Euclidean spaces that is a bit different from the convergence notion used above. In particular, people focus on *weak convergence*.

Example: suppose you have a sequence of probability measures λ_n on $[0, 1]$ (with the Borel algebra \mathcal{B}). Each measure λ_n is degenerate with $\lambda_n(\{1/n\}) = 1$. Now it seems natural to say that $\{\lambda_n\}$ converges to a distribution that puts all the probability at 0, i.e. to the probability measure λ with $\lambda(\{0\}) = 1$. However, if we use any metric that looks at the distance between $\lambda_n(A)$ and $\lambda(A)$ for different sets (like (1) was doing) we are in trouble. We have

$$\lambda_n(\{0\}) = 0 \text{ for all } n$$

so

$$\lim_n \lambda_n(\{0\}) = 0 \neq 1 = \lambda(\{0\}).$$

The problem is that this approach does not take into account that $1/n$ and 0 are close to each other when n is large.

Therefore a more appealing approach is to compare measures looking at how they integrate continuous functions. We have then the following:

Definition 2 *A sequence of measures $\{\lambda_n\}$ on the space $([0, 1], \mathcal{B})$ converges weakly to the measure λ (on the same space) iff*

$$\lim_{n \rightarrow \infty} \int_0^1 f d\lambda_n = \int_0^1 f d\lambda \text{ for all continuous functions on } [0, 1].$$

With this notion of convergence the sequence in our example does indeed converge to a degenerate distribution at 0. Chapter 12 of SLP provides conditions for weak convergence of Markov processes on Euclidean spaces and for the existence and uniqueness of an invariant probability measure.

MIT OpenCourseWare
<http://ocw.mit.edu>

14.451 Dynamic Optimization Methods with Applications
Fall 2009

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.