

Econ 204B: Section 6

Ryan Sherrard

University of California, Santa Barbara

15 February 2017

Stochastic Dynamic Programming: Introduction

- ▶ Now we will move more seriously into *stochastic* dynamic programming; i.e. we will add uncertainty into the mix
- ▶ Uncertainty can be motivated in numerous ways: it's more realistic, it can help generate things called “business cycles,” etc.
- ▶ First we will look at a fairly generic way of writing the problems down, then we'll look at (or mention) what is utilized in research

A Generic Problem

Let's take a generic problem and suppose that there is uncertainty over the future. Let $r(\cdot)$ be some objective function, let u_t denote the control vector, and let x_t be the state that evolves according to $x_{t+1} = h(x_t, u_t, \varepsilon_{t+1})$. The sequential problem looks like

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t r(x_t, u_t) \quad \text{s.t.} \quad x_{t+1} = h(x_t, u_t, \varepsilon_{t+1}) \quad \forall t, \quad x_0 \text{ given.}$$

Here, ε_t is some stochastic process ("shock") with support ξ and some distribution function $F(\varepsilon)$ that we usually take to be independent and identically distributed or Markov.

Note that, as a matter of notation, you'll often see \mathbb{E}_0 to denote the expectation operator with respect to all period 0 information.

Before we write down the problem recursively, let's think about what an equilibrium looks like in this environment.

- ▶ A sequence $\{u_t\}_{t=0}^{\infty}$ for every possible sequence of realizations for ε 's
- ▶ This is not so bad insofar as, at any given point in time, the problem has an infinite horizon and looks the same
- ▶ The above can be unwieldy, so we can instead find a *policy function* that tells the agent, at any point in time, what they should do given some observed x_t considering what they expect the ε 's to be in the future

Now let's translate the above into a recursive problem.

$$V(x) = \max_u \left\{ r(x, u) + \beta \mathbb{E} \left[\underbrace{V(h(x, u, \varepsilon'))}_{x'} \mid x \right] \right\}$$

$$\text{where} \quad \mathbb{E} \left[V(h(x, u, \varepsilon')) \mid x \right] \equiv \int_{\xi} V(h(x, u, \varepsilon')) dF(\varepsilon')$$

So, how do we deal with this? The answer turns out to be “pretty directly.” Let's begin by taking the FOCs of the optimization problem.

$$\frac{dV(x)}{du} = 0 : \quad r_2(x, u) + \beta \frac{d}{du} \mathbb{E} \left[V(h(x, u, \varepsilon')) \mid x \right] = 0$$

The only roadblock that we have to overcome is passing the derivative through the expectations operator.

If the limits of integration *do not* depend on the control u , we can directly apply **Leibniz's rule** for differentiation under the integral (i.e., you just do it).

$$r_2(x, u) + \beta \mathbb{E} \left[\frac{dV(h(x, u, \varepsilon'))}{dx'} h_2(x, u, \varepsilon') \middle| x \right] = 0$$

Alas, another roadblock: we do not know what $dV(x')/dx'$ is. Now we'll want to apply the Envelope Theorem. That is, we'll want to find $dV(x)/dx$.

This would proceed as it normally would. Without some functional forms, though, there isn't much more we can do. Hypothetically, however, suppose we actually knew the policy function for this problem, $u = g(x)$.

Recall also that we can think of the value function $V(\cdot)$ as the max of some problem, while the policy function $g(\cdot)$ is the associated argmax.

Now, before applying the Envelope Theorem, we can plug in the policy function $u = g(x)$ because, when we differentiate w.r.t. x , we'll be able to account for any effect this has on optimized behavior by agents.

The Bellman Equation becomes

$$V(x) = r(x, g(x)) + \beta \mathbb{E} \left[V(h(x, g(x), \varepsilon')) \mid x \right].$$

And so we can determine

$$\begin{aligned} \frac{dV(x)}{dx} &= r_1(x, g(x)) + r_2(x, g(x)) \frac{dg(x)}{dx} \\ &+ \beta \mathbb{E} \left[\frac{dV(h(x, g(x), \varepsilon'))}{dx'} \left(h_1(x, g(x), \varepsilon') + h_2(x, g(x), \varepsilon') \frac{dg(x)}{dx} \right) \mid x \right]. \end{aligned}$$

If the problem we are working with can be written in such a way such that the transition does not depend on x , this can be greatly simplified to

$$\frac{dV(x)}{dx} = r_1(x, u) \quad \implies \quad \frac{dV(x')}{dx'} = r_1(x', u').$$

Plugging this back into the FOC gives the stochastic EE.

$$r_2(x, u) + \beta \mathbb{E} [r_1(x', u') h_2(x, u, \varepsilon') | x] = 0$$

Let's try to map our previous example into this framework. Suppose that capital evolves according to $k' = (1 - \delta)k + a + \varepsilon$ (where ε is iid), and that there is full depreciation ($\delta = 1$).

Our Old Example

$$V(k, \varepsilon) = \max_{c, k'} \{ \ln(c) + \beta \mathbb{E} [V(k', \varepsilon')] \} \quad \text{s.t.} \quad c = k^\alpha - k' + \varepsilon$$

$$\implies V(k, \varepsilon) = \max_{k'} \{ \ln(k^\alpha - k' + \varepsilon) + \beta \mathbb{E} [V(k', \varepsilon')] \}$$

The FOC is given by

$$\frac{1}{k^\alpha - k' + \varepsilon} = \beta \mathbb{E} \left[\frac{dV(k', \varepsilon')}{dk'} \right],$$

where we passed the derivative through the integral using Leibniz's rule.

Now for the Envelope Theorem.

$$\frac{dV(k, \varepsilon)}{dk} = \frac{\alpha k^{\alpha-1}}{k^\alpha - k' + \varepsilon} \quad \Rightarrow \quad \frac{dV(k', \varepsilon')}{dk'} = \frac{\alpha k'^{\alpha-1}}{k'^\alpha - k'' + \varepsilon'}$$

Plugging this back into the FOC, we have the EE (which we can rewrite however we want).

$$\frac{1}{k^\alpha - k' + \varepsilon} = \beta \mathbb{E} \left[\frac{\alpha k'^{\alpha-1}}{k'^\alpha - k'' + \varepsilon'} \right]$$

$$\frac{1}{c} = \beta \mathbb{E} \left[\frac{\alpha k'^{\alpha-1}}{c'} \right]$$

Commonly Used Stochastic Processes

Here I provide a simple list of commonly utilized stochastic processes.

- ▶ Anything iid
- ▶ **Markov**
- ▶ AR(1)
- ▶ Brownian Motion
- ▶ Poisson Process (very useful in search)

Markov Process

Let the **state space** S be a finite set with k elements: $\{x_1, \dots, x_k\}$. A random process $\{X_t\}$ possesses the **Markov** property (or is called called a Markov chain) if, for some $y \in S$ we have that

$$\Pr(X_{t+1} = y | X_t, X_{t-1}, \dots) = \Pr(X_{t+1} = y | X_t).$$

That is, the probability of going to the state value of y (next period) only depends on the current state. Note that, for $y, z \in S$, we can succinctly write that

$$p_{zy} \equiv \Pr(X_{t+1} = y | X_t = z).$$

Furthermore, we can also summarize the transition probabilities in a matrix,

$$P \equiv \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1k} \\ p_{21} & p_{22} & \cdots & p_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ p_{k1} & p_{k2} & \cdots & p_{kk} \end{bmatrix},$$

where we must have that

$$\sum_{j=1}^k p_{ij} = 1 \quad \forall i = 1, \dots, k.$$

That is, each row must sum to 1.

The Markov property is very useful, and can be very powerful in simplifying otherwise difficult problems. Consider, for presentation, a simplified 2-state Markov process.

Let the state space be $S = \{0, 1\}$, and the associated transition matrix be given by

$$P = \begin{bmatrix} p_{11} & p_{10} \\ p_{01} & p_{00} \end{bmatrix}.$$

We know from before that, for example, the probability of going to state 1 from state 1 is p_{11} . What is the probability of being in state 1 in *two* periods (given we are in state 1 now)?

$$\begin{aligned} \Pr(S = 1 \text{ in two periods}) &= p_{11} \times \Pr(S = 1 \text{ tomorrow} | S = 1 \text{ today}) \\ &\quad + p_{01} \times \Pr(S = 0 \text{ tomorrow} | S = 1 \text{ today}) \\ &= p_{11}^2 + p_{01}p_{10} \end{aligned}$$

Note, though, that we could have used matrix multiplication instead.

$$P^2 = PP = \begin{bmatrix} p_{11} & p_{10} \\ p_{01} & p_{00} \end{bmatrix} \begin{bmatrix} p_{11} & p_{10} \\ p_{01} & p_{00} \end{bmatrix} = \begin{bmatrix} p_{11}^2 + p_{10}p_{01} & p_{11}p_{10} + p_{10}p_{00} \\ p_{01}p_{11} + p_{00}p_{01} & p_{01}p_{10} + p_{00}^2 \end{bmatrix}$$

Indeed, this pattern will hold for any k -state Markov process any number of periods down the road.

Want to know the probability of being in state 0, starting from state 1, in 50 periods? Select element “ p_{10} ” from the transition matrix given by P^{50} : $p_{10}^{(50)}$.

To this end, we can actually look at the probability distribution across states at certain points in the future (that is, the probability of being in each state at some point in time)...

Define the **probability distribution** across states in period t as $\pi^{(t)}$. Suppose that we know the probability distribution across states in period 0 (the current period). In the two-state example, if we are in the high state, we'll have $\pi^{(0)} = [\pi_1^{(0)}, \pi_0^{(0)}] = [1, 0]$ (with probability 1 we are in the high state).

After 1 period, the probability distribution will be given by $\pi^{(1)} = \pi^{(0)}P$. After two periods, we'll similarly have $\pi^{(2)} = \pi^{(1)}P = \pi^{(0)}P^2$. Generally, this becomes

$$\pi^{(t)} = \pi^{(0)}P^t.$$

The **limiting** or **invariant** distribution of a Markov chain is given by taking the above in the limit (as $t \rightarrow \infty$). Denoting this distribution as π , this is also sometimes written as

$$\pi = \pi P.$$

With our two-state example, this is fairly easy to find analytically (with pen and paper). Notice that the last expression defines a system of equations (which we can solve).

$$\pi_1 = \pi_1 p_{11} + \pi_0 p_{01}$$

$$\pi_0 = \pi_1 p_{10} + \pi_0 p_{00}$$

If we note that $\pi_1 + \pi_0 = 1$ and that the rows of P must sum to 1, we'll determine that

$$[\pi_1, \pi_0] = \left[\frac{1 - p_{00}}{(1 - p_{11}) + (1 - p_{00})}, \frac{p_{00}}{(1 - p_{11}) + (1 - p_{00})} \right].$$