Stochastic Dynamic Programming

Jesús Fernández-Villaverde
University of Pennsylvania

Introducing Uncertainty in Dynamic Programming

- Stochastic dynamic programming presents a very flexible framework to handle multitude of problems in economics.

- We generalize the results of deterministic dynamic programming.

- Problem: taking care of measurability.

References

- Read chapter 9 of SLP!!!!!!!!!!!!!!

- Problem of SLP: based on Borel sets. Raises issues of measurability. See page 253 and 254 of SLP.

- Bertsekas and Shreve (Stochastic Optimal Control, 1978) redo much of the theory with universal measurability

- Read chapter 10 of SLP: it is full of economic applications.

Environment

- $(X, \mathcal{X})$: universally measurable space for the endogenous state.
- $(Z, \mathcal{Z})$: universally measurable space for the exogenous state.
- $(S, \mathcal{S})$: $(X, \mathcal{X}) \times (Z, \mathcal{Z})$.
- Q: stationary transition function for $(Z, \mathcal{Z})$.
- $\Gamma : X \times Z \rightarrow X$: correspondence constraint.
- $A = \{(x, y, z) \in X \times X \times Z : y \in \Gamma(x, z)\}$: graph of Γ.
- $F : A \rightarrow \mathbb{R}$: one-period return function.
- β: discount factor.

Plans

- $\pi_t : Z^t \to X$ for $t = 1, 2, \ldots$: sequence of measurable functions.

- $\pi = (\pi_0 \in X, \pi_t)$: plan.

- Interpretation of a plan: contingent decision rules.

- A plan π is feasible from $s_0 \in S$ if:

  1. $\pi_0 \in \Gamma(s_0)$.
  2. $\pi_t \in \Gamma\left(\pi_{t-1}\left(z^{t-1}\right), z_t\right)$ for $z^t \in Z^t$, $t = 1, 2, \ldots$

- $\Pi(s_0)$: set of all feasible plans from $s_0 \in S$.

- If π does not depend on t but only on z^t, we call the plan stationary or Markov.

Some Preliminary Results I

- Assumption 1:

  1. Γ is non-empty valued.

  2. A is $(\mathcal{X} \times \mathcal{X} \times \mathcal{Z})$ $-$measurable.

  3. $\exists$ a measurable selection $h : S \rightarrow X$ s.t. $h(s) \in \Gamma(s)$ for $\forall s \in S$.

- Lemma 1: under previous assumption, $\Pi(s_0)$ is nonempty for $\forall s_0 \in S$.

- Lemma 2: $\mathcal{A} = (\mathcal{X} \times \mathcal{X} \times \mathcal{Z})$ is a $\sigma-$algebra.

- Corollary 1: $F\left(\pi_{t-1}\left(z^{t-1}\right), \pi_t\left(z^t\right), z_t\right)$ is $\mathcal{Z}^t-$measurable.

Some Preliminary Results II

- Given Q on $(Z, \mathcal{Z})$ and $s_0 \in S$,

$$\mu^t(z_0, \cdot) : \mathcal{Z}^t \to [0, 1], \; t = 1, 2, ...$$

- Assumption 2: $F : A \to \mathbb{R}$ is $\mathcal{A}-$measurable and either (a) or (b) holds:

  a. $F \geq 0$ or $F \leq 0$.

  b. For each $(x_0, z_0) = s_0 \in S$ and each plan $\pi \in \Pi(s_0)$,

  $F\left(\pi_{t-1}\left(z^{t-1}\right), \pi_t\left(z^t\right), z_t\right)$ is $\mu^t(z_0, \cdot) -$ integrable, $t = 1, 2, ...$ and the limit:

$$F(x_0, \pi_0, z_0) + \lim_{t \to \infty} \sum_{t=1}^{\infty} \int_{Z^t} \beta^t F\left(\pi_{t-1}\left(z^{t-1}\right), \pi_t\left(z^t\right), z_t\right) \mu^t(z_0, \cdot)$$

  exists (though it may be plus or minus infinity).

Sequential Problem

- Define $u_n(\cdot, s_0) : \Pi(s_0) \rightarrow \mathbb{R}$, $n = 0, 1, ...$ by:

$$\begin{aligned} u_0(\pi, s_0) &= F(x_0, \pi_0, z_0) \\ u_n(\pi, s_0) &= F(x_0, \pi_0, z_0) \\ &+ \sum_{t=1}^{n} \int_{Z^t} \beta^t F\left(\pi_{t-1}\left(z^{t-1}\right), \pi_t\left(z^t\right), z_t\right) \mu^t\left(z_0, dz^t\right) \end{aligned}$$

- Define $u(\pi, s_0) : \Pi(s_0) \rightarrow \mathbb{R}_\infty$ by

$$u(\pi, s_0) = \lim_{n \rightarrow \infty} u_n(\pi, s_0)$$

- Define $v^* : S \rightarrow \mathbb{R}_\infty$ by

$$v^*(s) = \sup_{\pi \in \Pi(s)} u(\pi, s_0)$$

Recursive Problem

- Functional equation:

$$v(s) = v(x,z) = \sup_{y \in \Gamma(x,z)} \left[F(x,y,z) + \beta \int v\left(y, z'\right) Q\left(z, dz'\right) \right]$$

- Associate with the functional equation, we have a policy correspondence:

$$G(x,z) = \left\{ y \in \Gamma(x,z) : v(x,z) = F(x,y,z) + \beta \int v\left(y, z'\right) Q\left(z, dz'\right) \right\}$$

- If G is nonempty and if there is a sequence of measurable selections $g_1, ...$ from G, we have the plan generated by G from s_0:

$$\begin{aligned} \pi_0 &= g_0(s_0) \\ \pi_t\left(z^t\right) &= g_t\left[\pi_{t-1}\left(z^{t-1}\right), z^t\right], \forall z^t \in Z^t, t = 1, 2, ... \end{aligned}$$

Transversality Condition

- In general, dynamic programming problems require two boundary conditions: an initial condition and a final condition.

- Transversality condition plays the role of the second condition.

- To ensure the equivalence of the sequential and recursive problem, we also need then a transversality condition:

$$\lim_{t\to\infty} \beta^t \int v\left(\pi_{t-1}\left(z^{t-1}\right), z^t\right) \mu^t\left(z_0, dz^t\right) = 0,\ \forall \pi \in \Pi\left(s_0\right),\ s_0 \in S$$

Equivalence of Sequential and Recursive Problem

- Under our previous assumptions:

  1. $v = v^*$
  2. Any plan π^* generated by G obtains the supremum in $v^*(s) = \sup_{\pi \in \Pi(s)} u(\pi, s_0)$

- Under our previous assumptions and an additional boundness condition, a plan is optimal only if it is generated a.e. by G.

- Our results are equivalent to theorems 4.2-4.5 in SLP for the deterministic case.

Bounded Returns

- As in the deterministic case, we want to show further results.

- Assumptions:

  1. F is bounded and continuous.
  2. $\beta < 1$.
  3. X is a compact set in $\mathbb{R}^l$ and $\mathcal{X}$ is a universally measurable $\sigma-$algebra.
  4. Z is a compact set in $\mathbb{R}^k$ and $\mathcal{Z}$ is a universally measurable $\sigma-$algebra.
  5. Q has the Feller property.

- Intuition: integration will preserve properties of the return function.

Results I

Under these assumptions, we can prove that:

1. The Bellman operator:

$$(Tf)(x,z) = \sup_{y \in \Gamma(x,z)} \left[F(x,y,z) + \beta \int v\left(y,z'\right) Q\left(z,dz'\right) \right]$$

   has a unique fixed point.

2. Contractivity: $\|T^n v_0 - v\| \leq \beta^n \|v_0 - v\|$, $n = 1, 2, ...$

3. The policy correspondence

$$G(x,z) = \left\{ y \in \Gamma(x,z) : v(x,z) = F(x,y,z) + \beta \int v\left(y,z'\right) Q\left(z,dz'\right) \right\}$$

   is non-empty, compact-valued, and u.h.c.

4. The value function will inherit increasing properties from F and Q.

Concavity

- Assumption concavity 1: For each $z \in Z$, $F(\cdot, \cdot, z) : A_z \to \mathbb{R}$ satisfies:

$$\begin{aligned} F\left(\theta(x,y) + (1-\theta)\left(x',y'\right), z\right) &\geq \theta F(x,y,z) + (1-\theta) F\left(x',y',z\right) \\ \forall\theta &\in (0,1), \forall(x,y), \left(x',y'\right) \in A_z \end{aligned}$$

and the inequality is strict if $x \neq x'$.

- Assumption concavity 2: For $\forall z \in Z$ and $\forall x, x' \in X$, $y \in \Gamma(x,z)$ and $y' \in \Gamma(x',z)$

$$\theta y + (1-\theta) y' \in \Gamma\left(\theta x + (1-\theta) x', z\right), \ \forall\theta \in (0,1)$$

Results II

1. Under previous assumptions, $v(\cdot, z) : X \rightarrow \mathbb{R}$ is strictly concave and $G(\cdot, z) : X \rightarrow X$ is a continuous, single-valued function.

2. Let $v_n = T v_{n-1}$ and

$$g_n(x, z) = \arg \max_{y \in \Gamma(x,z)} \left\{ F(x, y, z) + \beta \int v\left(y, z'\right) Q\left(z, dz'\right) \right\} \text{ for } n = 1, 2, \ldots$$

   Then, $g_n \rightarrow g$ uniformly.

3. If $x_0 \in int(X)$ and $g(x_0, z_0) \in int(\Gamma(x_0, z_0))$, $v(\cdot, z_0)$ is continuously differentiable in x at x_0 with derivatives given by:

$$v_i(x_0, z_0) = F_i[x_0, g(x_0, z_0), z_0], \; i = 1, ..., l$$

Unbounded Returns

- What if returns, like in most applications of interest in economics, are unbounded?

- This was already an issue in the deterministic set-up.

- We can get most of the substance of previous results if F is constant returns to scale.

- In the case of CRRA utility functions, we would need to do some ad-hoc work.

Policy Functions and Transition Functions I

- Let us imagine that the decision maker follows $g(x, z)$ given an initial condition s_0.

- The policy function generates a sequence $\{s_t\}$.

- What do we know about $\{s_t\}$?

- Read chapters 11-14 of SLP.

Policy Functions and Transition Functions II

- Let $(X, \mathcal{X})$, $(Z, \mathcal{Z})$, and $(S, \mathcal{S})$: $(X, \mathcal{X}) \times (Z, \mathcal{Z})$ be universally measurable spaces; let Q be a transition function on $(Z, \mathcal{Z})$; and let $g : S \to X$ be a measurable function. Then:

$$P\left[(x, z), A \times B\right] = \begin{cases} Q(z, B) \text{ if } g(x, z) \in A \\ 0 \text{ otherwise} \end{cases}$$

  for $\forall x \in X, z \in Z, A \in \mathcal{X}$, and $B \in \mathcal{Z}$, defines a transition function on $(S, \mathcal{S})$.

- If g is continuous, then P has the Feller property.

- Characterizing long run behavior of the model.