

Econ 8107 Recitation 2

Stochastic Dynamic Programming

Zach Swaziek

January 24, 2022

SLP Chapter 9 - Introduction

We will now show how to incorporate stochastic shocks into dynamic programming models.

Example:

- One good, corn, which is consumed or used as seed each period.
- Total output is $zf(k)$, where k is seed from previous period and z is quantity of rainfall.
- Rainfall quantity is exogenous and stochastic

The functional equation for this problem is:

$$v(k, z) = \sup_{y \in [0, zf(k)]} \left\{ U(zf(k) - y) + \beta \int_Z v(y, z') Q(z, dz') \right\} \quad (1)$$

For this FE, we have:

- State variables: k, z
- Control variable: y

The control variable y , how much seed is saved, is chosen after the rainfall shock occurs (e.g. in the fall after the harvest). **The timing of shocks/decision making in any problem with uncertainty is very important.**

The problem described above is the type we will study for the recitation. In these problems, the shock occurs first and then the value of the control variable is chosen. SLP also describes a more general problem, where the agent chooses a value for the control variable but the shock affects the value of next period's endogenous state variable. Thus the the agent does not directly choose next period's endogenous state variable; it is a random variable. You will read about the treatment for these types of problems for the problem set.

SLP Section 9.1 - Principle of Optimality

First we define the objects used in this chapter.

- X : set of possible values of the endogenous state variable
- Z : set of possible values for exogenous shock
- $S = X \times Z$: set of possible states for the system
- $(X, \mathcal{X}), (Z, \mathcal{Z})$: measurable spaces
- $(S, \mathcal{S}) = (X \times Z, \mathcal{X} \times \mathcal{Z})$: product space
- Q : stationary¹ transition function on $(Z, \mathcal{Z})$
- $\Gamma : X \times Z \rightarrow X$: choice correspondence for y ; describes constraints on choice of value of next period's endogenous state variable
- $\Gamma(x, z)$: set of feasible values for next period's endogenous state variable if current state is (x, z)
- $A = \{(x, y, z) \in X \times X \times Z; y \in \Gamma(x, z)\}$: graph of Γ
- A_z and A_{yz} : sections of A
- $\mathcal{A} = \{C \in \mathcal{X} \times \mathcal{X} \times \mathcal{Z} : C \subseteq A\}$: σ -algebra² containing subsets of A
- $F : A \rightarrow \mathbb{R}$: one period return function
- $F(x, y, z)$: current period return if the current state is (x, z) and $y \in \Gamma(x, z)$ is chosen as next period's endogenous state variable.
- $\beta \geq 0$: one period discount factor
- $Z^t = \{(z_1, z_2, \dots, z_t) : z_1 \in Z, z_2 \in Z, \dots, z_t \in Z\}$: set of possible partial histories at period t
- $(Z^t, \mathcal{Z}^t)$: product space³ for $t = 1, 2, \dots$
- $z^t = (z_1, z_2, \dots, z_t) \in Z^t$: partial history of shocks
- $\mu^t(z_0, \cdot) : \mathcal{Z}^t \rightarrow [0, 1]$: probability measures given transition function Q on $(Z, \mathcal{Z})$ and $s_0 = (x_0, z_0)$

This completes the list of objects. The **givens** for any stochastic dynamic programming problem are: $(X, \mathcal{X}), (Z, \mathcal{Z}), Q, \Gamma, F, \beta$

¹I.e. not indexed by time.

²Assuming A9.1 satisfied.

³See page 195 for a section on product spaces.

Now consider the stochastic problem in **sequence form**.

- $s_0 = (x_0, z_0)$, the initial endogenous state variable and initial shock, is known
- The agent wants the choice of x_{t+1} to depend on the **information** known at period t
- **Information** available at period $t \geq 1$ is sequence of shocks $(z_1, z_2, \dots, z_t)$.
- Agent makes contingency plans for all periods and all possible histories.
- Agent chooses sequence of contingency plans to maximize discounted sum of returns
- Expectation is with respect to the realization of shocks.

The sequence form of the problem described above can be written as:

$$\begin{aligned} \max_{\pi_0, \{\{\pi_t(z^t)\}_{z^t \in Z^t}\}_{t=1}^\infty} & U[z_0 f(x_0) - \pi_0] + \sum_{t=1}^\infty \int_{Z^t} \beta^{t-1} U[z^t f(\pi_{t-1}(z^{t-1})) - \pi_t(z^t)] \mu^t(z_0, dz^t) \\ \text{s.t. } & \pi_0 \in \Gamma(s_0) \\ & \pi_t(z^t) \in \Gamma[\pi_{t-1}(z^{t-1}), z_t] \quad \forall z^t \in Z^t, \quad t = 1, 2, \dots \\ & s_0 = (k_0, z_0) \in S \quad \text{given} \end{aligned}$$

In this problem, $\pi_t(z^t)$ is the value of x_{t+1} that will be chosen at period t if the partial history z^t occurs.

Throughout the remainder of these notes, version (a) or (b) of Definitions, Assumptions, and Theorems will refer to the analogous version in Chapter 4 (the deterministic dynamic programming setting).

Definition 1 A **plan** is a value $\pi_0 \in X$ and a sequence of measurable functions $\pi_t : Z^t \rightarrow X$, $t = 1, 2, \dots$

Definition 1a A **plan** is any sequence $\{x_t\}_{t=0}^\infty$ in X .

In the stochastic setting, the agent chooses a sequence of **functions** denoted by $\pi = \{\pi_t\}_{t=0}^\infty$. Since potentially many histories can happen, the agent has to plan for each of them. Compare this to the deterministic setting, where the agent chooses a sequence of **values**, $\{x_t\}_{t=0}^\infty$. There is no randomness in a deterministic dynamic programming problem, so there is no need for contingency planning.

Definition 2 A plan π is **feasible from** s_0 if:

1. $\pi_0 \in \Gamma(s_0)$,
2. $\pi_t(z^t) \in \Gamma[\pi_{t-1}(z^{t-1}), z_t] \quad \forall z^t \in Z^t, \quad t = 1, 2, \dots$

Definition 2a A plan $\{x_t\}_{t=0}^\infty$ is **feasible from** x_0 if $x_{t+1} \in \Gamma(x_t)$, $t = 0, 1, \dots$

Let $\Pi(s_0)$ denote the set of plans that are feasible from $s_0 = (x_0, z_0)$. In the deterministic case the set of all possible plans given x_0 , $\Pi(x_0)$, was non-empty if Γ was non-empty. In the stochastic case, we also need to address measurability.

Assumption 1 (A9.1) Γ is nonempty-valued and the graph of Γ is $(\mathcal{X} \times \mathcal{X} \times \mathcal{Z})$ -measurable. In addition, Γ has a measurable selection; that is, there exists a measurable function $h : S \rightarrow X$ such that $h(s) \in \Gamma(s)$, all $s \in S$.

Assumption 1a (A4.1) $\Gamma(x)$ is nonempty for all $x \in X$.

Under A9.1, $\Pi(s_0)$ is nonempty for all $s_0 \in S$ (Lemma 1).

Assumption 2 (A9.2) $F : A \rightarrow \mathbb{R}$ is $\mathcal{A}$ -measurable and either (1) or (2) holds:

1. $F \geq 0$ or $F \leq 0$
2. For each $(x_0, z_0) = s_0 \in S$ and each plan $\pi \in \Pi(s_0)$,
 $F[\pi_{t-1}(z^{t-1}), \pi_t(z^t), z_t]$ is $\mu^t(z_0, \cdot)$ -integrable, $t = 1, 2, \dots$,
and the limit
 $F[x_0, \pi_0, z_0] + \lim_{n \rightarrow \infty} \sum_{t=1}^n \int_{Z^t} \beta^t F[\pi_{t-1}(z^{t-1}), \pi_t(z^t), z_t] \mu^t(z_0, dz^t)$
exists (although it may be plus or minus infinity).

Assumption 2a (A4.2) For all $x_0 \in X$ and $\{x_t\}_{t=0}^\infty \in \Gamma(x_0)$,
 $\lim_{n \rightarrow \infty} \sum_{t=0}^\infty \beta^t F(x_t, x_{t+1})$ exists (although it may be plus or minus infinity).

Note that (1) holding implies (2). If A9.2 is satisfied, we can define functions that describe the expected discounted returns up to period t for any feasible plan $\pi \in \Pi(s_0)$ given any initial state $s_0 \in S$. The functions $u_n(\cdot, s_0) : \Pi(s_0) \rightarrow \mathbb{R}$, $n = 0, 1, 2, \dots$ are defined by

$$\begin{aligned} u_0(\pi, s_0) &= F[x_0, \pi_0, z_0] \\ u_n(\pi, s_0) &= F[x_0, \pi_0, z_0] + \sum_{t=1}^n \int_{Z^t} \beta^t F[\pi_{t-1}(z^{t-1}), \pi_t(z^t), z_t] \mu^t(z_0, dz^t) \end{aligned} \quad (2)$$

Similarly, define the infinite sum of expected discounted returns from plan π given

initial state $s_0 \in S$, $u(\pi, s_0) : \Pi(s_0) \rightarrow \bar{\mathbb{R}}$, by

$$u(\pi, s_0) = \lim_{n \rightarrow \infty} u_n(\pi, s_0) \quad (3)$$

If A9.1 and A9.2 are satisfied, then $u(\pi, s_0)$ is well defined on the nonempty set $\Pi(s_0)$ for all $s_0 \in S$. Thus we can define the supremum function $v^* : S \rightarrow \mathbb{R}$

$$v^*(s_0) = \sup_{\pi \in \Pi(s_0)} u(\pi, s_0) \quad (4)$$

Now, consider the FE that corresponds to the sequence problem in (5):

$$v(s) = v(x, z) = \sup_{y \in \Gamma(x, z)} \left[F(x, y, z) + \beta \int v(y, z') Q(z, dz') \right] \quad (5)$$

If there exists a v that satisfies the FE in (6), then we can define the policy correspondence G by

$$G(x, z) = \{y \in \Gamma(x, z) : v(x, z) = F(x, y, z) + \beta \int v(y, z') Q(z, dz')\} \quad (6)$$

We say that a plan π is **generated by G from s_0** by taking a sequence of measurable selections from G , $g_0, g_1, \dots$ and define π by

$$\begin{aligned} \pi_0 &= g_0(s_0) \\ \pi_t(z^t) &= g_t[\pi_{t-1}(z^{t-1}), z_t], \quad \text{all } z^t \in Z^t, \quad t = 1, 2, \dots \end{aligned}$$

Notice that if there was no uncertainty (e.g. $z = 1$ for all periods and histories), (6) and (7) would look just like their deterministic counterparts from Chapter 4.

Now we are prepared for the two main theorems related to the Principle of Optimality. The first, Theorem 9.2, is the analogue of Theorems 4.3 and 4.5. The second, Theorem 9.4, is an analogue of Theorem 4.4. There is no analogue to Theorem 4.2.⁴

⁴Theorem 4.2: Let X , Γ , F , and β satisfy Assumptions 4.1-4.2. Then the function v^* satisfies (FE).

Theorem 1 (T9.2) *Let $(X, \mathcal{X})$, $(Z, \mathcal{Z})$, Q , Γ , F , and β be given. Let Assumptions 9.1 and 9.2 hold, and let v^* be defined by (5). let v be a measurable function satisfying the functional equation (6), and such that:*

$$\lim_{t \rightarrow \infty} \int_{Z^t} \beta^t v[\pi_{t-1}(z^{t-1}), z_t] \mu^t(z_0, dz^t) = 0$$

all $\pi \in \Pi(s_0)$, all $(x_0, z_0) = s_0 \in S$

Let G be the correspondence defined by (7), and suppose that G is nonempty and permits a measurable selection.

Then $v = v^*$, and any plan π^* generated by G attains the supremum in (5).

Theorem 1a (T4.3) *Let X , Γ , F , and β satisfy Assumptions 4.1-4.2. If v is a solution to (FE) and satisfies*

$$\lim_{n \rightarrow \infty} \beta^n v(x_n) = 0, \quad \text{all } \{x_t\}_{t=0}^\infty \in \Pi(x_0), \quad \text{all } x_0 \in X.$$

then $v = v^*$.

Theorem 1b (T4.5) *Let X , Γ , F , and β satisfy Assumptions 4.1-4.2. Let $\{x_t^*\}_{t=0}^\infty \in \Pi(x_0)$ be a feasible plan from x_0 satisfying*

$$v^*(x_t^*) = F(x_t^*, x_{t+1}^*) + \beta v^*(x_{t+1}^*), \quad t = 0, 1, 2, \dots$$

and with

$$\limsup_{t \rightarrow \infty} \beta^t v^*(x_t^*) \leq 0.$$

Then $\{x_t^*\}_{t=0}^\infty$ attains the supremum in (SP) for initial state x_0 .

We need an additional assumption for the analogue to Theorem 4.4.

Assumption 3 *If F takes on both signs, there is a collection of nonnegative, measurable functions $L_t : S \rightarrow R_+$, $t = 0, 1, \dots$, such that for all $\pi \in \Pi(s_0)$ and all $s_0 \in S$*

$$|F(x_0, \pi_0, z_0)| \leq L_0(s_0)$$

$$|F[\pi_{t-1}(z^{t-1}), \pi_t(z^t), z_t]| \leq L_t(s_0), \quad \text{all } z^t \in Z^t, \quad t = 1, 2, \dots$$

and

$$\sum_{t=0}^{\infty} \beta^t L_t(s_0) < \infty$$

Theorem 9.3 is a partial converse to T9.2. This theorem states that a plan π is optimal (it attains the supremum in (5)) only if it is generated (a.e.) by G .

Theorem 2 (T9.3) *Let $(X, \mathcal{X})$, $(Z, \mathcal{Z})$, Q , Γ , F , and β be given. Let Assumptions 9.1-9.3 hold, and define v^* by (5). Assume that v^* is measurable and satisfies the functional equation (6), and define the correspondence G by (7). Assume that G is nonempty and permits a measurable selection. Let $(x_0, z_0) = s_0 \in S$, and let $\pi^* \in \Pi(s_0)$ be a plan that attains the supremum in (5) for initial condition s_0 . Then there exists a plan π^G generated by G from s_0 such that*

$$\begin{aligned}\pi_0^G &= \pi_0^*, \quad \text{and} \\ \pi_t^G(z^t) &= \pi_t^*(z^t), \quad \mu^t(z_0, \cdot) - \text{a.e.}, \quad t = 1, 2, \dots\end{aligned}$$

Theorem 2a (T4.4) *Let X , Γ , F , and β satisfy Assumptions 4.1-4.2. Let $\{x_t^*\}_{t=0}^\infty \in \Pi(x_0)$ be a feasible plan that attains the supremum in (SP) for initial state x_0 . Then*

$$v^*(x_t^*) = F(x_t^*, x_{t+1}^*) + \beta v^*(x_{t+1}^*), \quad t = 0, 1, 2, \dots$$

SLP Section 9.2 - Bounded Returns

Remaining consistent with the previous section, we will focus on functional equations of the following type:

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

In these type of dynamic programming problems the agent directly chooses next period's endogenous state variable, y . These means that any uncertainty is resolved before choosing the value of y . As mentioned in the introduction, there is a more general type of problems, but you can read about those in SLP.⁵

As in Chapter 4.2, we will now add assumptions to the dynamic programming problems that will allow us to use the contraction mapping theorem. Therefore, in this section we will need the following conditions to be satisfied:

1. One period return function F to be bounded and continuous
2. Discount factor $\beta \in (0, 1)$
3. Transition function Q has the Feller property⁶

You will notice that the results and assumptions required in this section are very close to those in Chapter 4.2.

Assumption 4 (A9.4) *X is a convex Borel set in $\mathbb{R}^l$, with its Borel subsets $\mathcal{X}$.*

Assumption 4a (Part of A4.3) *X is a convex subset of $\mathbb{R}^l$.*

Since there is uncertainty, we need to address the state space of the shocks. Clearly there is no analogue to this assumption in Section 4.2.

Assumption 5 (A9.5) *One of the following conditions holds:*

1. *Z is a countable set and $\mathcal{Z}$ is a σ -algebra containing all subsets of Z (in other words, $\mathcal{Z} = 2^Z$, the power set of Z); or*
2. *Z is a compact (Borel) set in $\mathbb{R}^k$, with its Borel subsets $\mathcal{Z}$, and the transition function Q on $(Z, \mathcal{Z})$ has the Feller property.*

⁵In these types of problems, the agent chooses a value of the control variable y today, but next period's state variable has value $\phi(x, y, z')$. This way of writing the problem allows us to introduce uncertainty into the value of next period's state variable.

⁶A transition function Q on $(Z, \mathcal{Z})$ has the **Feller Property** if the associated operator T maps the space of bounded continuous functions on Z onto itself, i.e. $T : C(Z) \rightarrow C(Z)$

A9.4 and A9.5 ensure that the expectation/integral term in the functional equation maps from the space of continuous bounded functions (which we denote by $C(S)$) $f : S \rightarrow \mathbb{R}$ to itself. Note that we once again use the sup norm as our metric: $\|f\| = \sup_{s \in S} |f(s)|$. These two assumptions also ensure that the expectation/integral preserves monotonicity and concavity of the function f .

Lemma 1 (L9.5) *Let $(X, \mathcal{X})$, $(Z, \mathcal{Z})$, and Q satisfy Assumptions 9.4 and 9.5. If $f : X \times Z \rightarrow \mathbb{R}$ is bounded and continuous, then Mf defined by*

$$(Mf)(y, z) = \int f(y, z')Q(z, dz'), \quad \text{all } (y, z) \in X \times Z$$

is also bounded and continuous; that is $M : C(S) \rightarrow C(S)$.

If f is increasing in (strictly increasing) in each of its first l arguments, then so is Mf ; and if f is concave (strictly concave) jointly in its first l arguments, then so is Mf .

Because of this lemma, it is straightforward to show that the major results in Chapter 4 carryover to stochastic dynamic programming problems.

Assumption 6 (A9.6) *The correspondence $\Gamma : X \times Z \rightarrow X$ is nonempty, compact-valued, and continuous.*

Assumption 6a (Part of A4.3) *The correspondence $\Gamma : X \rightarrow X$ is nonempty, compact-valued, and continuous.*

Assumption 7 (A9.7) *The function $F : A \rightarrow \mathbb{R}$ is bounded and continuous, and $\beta \in (0, 1)$.*

Assumption 7a *The function $F : A \rightarrow \mathbb{R}$ is bounded and continuous, and $\beta \in (0, 1)$.*

As Exercise 9.6 suggests, you can show that if A9.4-A9.7 are satisfied then so are A9.1-A9.4. Thus we preserve the results from Section 9.1. With these assumptions in hand, we can now state the analogue to Theorem 4.6. Theorem 9.6 tells us that under A9.4-9.7, there exists a unique continuous and bounded solution to the functional equation.

Theorem 3 (T9.6) Let $(X, \mathcal{X})$, $(Z, \mathcal{Z})$, Q , Γ , F , and β satisfy Assumptions 9.4-9.7, and define the operator T on $C(S)$ by

$$(Tf)(x, z) = \sup_{y \in \Gamma(x, z)} \left\{ F(x, y, z) + \beta \int f(y, z') Q(z, dz') \right\}. \quad (7)$$

Then $T : C(S) \rightarrow C(S)$; T has a unique fixed point v in $C(S)$ and for any $v_0 \in C(S)$,

$$\|T^n v_0 - v\| \leq \beta^n \|v_0 - v\|, \quad n = 1, 2, \dots$$

Moreover, the correspondence $G : S \rightarrow X$ defined by

$$G(x, z) = \{y \in \Gamma(x, z) : v(x, z) = F(x, y, z) + \beta \int v(y, z') Q(z, dz')\} \quad (8)$$

is nonempty, compact-valued, and u.h.c.

Theorem 3a (T4.6) Let X , Γ , F , and β satisfy Assumptions 4.3 and 4.5, and let $C(X)$ be the space of bounded continuous functions $f : X \rightarrow \mathbb{R}$, with the sup norm. Then the operator T defined by

$$(Tf)(x) = \max_{y \in \Gamma(x)} [F(x, y) + \beta f(y)]$$

maps $C(X)$ into itself, $T : C(X) \rightarrow C(X)$; T has a unique fixed point $v \in C(X)$; and for all $v_0 \in C(X)$,

$$\|T^n v_0 - v\| \leq \beta^n \|v_0 - v\|, \quad n = 1, 2, \dots$$

Moreover, given v , the optimal policy correspondence $G : X \rightarrow X$ defined by

$$G(x) = \{y \in \Gamma(x) : v(x) = F(x, y) + \beta v(y)\} \quad (9)$$

is compact-valued and u.h.c.

The proof of T9.6 relies on Lemma 9.5, but is otherwise nearly identical to the proof of T4.6. As in section 4.2, we can add assumptions to guarantee that v is strictly increasing and strictly concave in x , as well as continuously differentiable. We can also add assumptions such that v is strictly increasing in z . We call these “inherited properties,” roughly speaking, if the one-period return function F has one of those properties, so will v . Since the assumptions and theorems are nearly identical to those in section 4.2, they are not included in these notes.