Chapter 4

Stochastic Dynamic Programming

The aim of this chapter is to extend the framework we introduced in Chapter 3 to include uncertainty. To evaluate decisions, we use the well known expected utility theory.[1] With uncertainty we will face Bellman equations of the following form

$$V(x,z) = \sup_{x' \in \Gamma(x,z)} F(x,x',z) + \beta \mathbf{E}\left[V(x',z') \mid z\right], \tag{4.1}$$

where z is a stochastic component, assumed to follow a (stationary) first order *Markov Process.* A first order Markov process is a sequence of random variables $\{z_t\}_{t=0}^{\infty}$ with the property that the conditional expectations depend only on the last realization of the process, that is if C is a set of possible values for z, then

$$\Pr\left\{z_{t+1} \in C \mid z_t, z_{t-1}, ..., z_0\right\} = \Pr\left\{z_{t+1} \in C \mid z_t\right\}.$$

To make the above statements formally meaningful we need to review some concepts of Probability Theory.

4.1 The Axiomatic Approach to Probability: Basic Concepts of Measure Theory

I am sure you are all familiar with the expression $\Pr\left\{z_{t+1} \in C \mid z_t\right\}$ for conditional probabilities, and with the conditional expectation operator $\mathbf{E}\left[\cdot \mid z\right]$ in (4.1). Probability theory is a special case of the more general and very powerful *Measure Theory,* first formulated in 1901 by Henri Léon Lebesgue.[2]

[1] For a review on the theories of decisions under uncertainty see Machina (1987).

[2] This outstanding piece of work appears in Lebesgue's dissertation, *Intégrale, Longueur, Aire*, presented to the University of Nancy in 1902.

We first introduce a set Z which will be our *sample space.* Any subset E of Z, will be denoted as an *event.* In this way, all results of set theory - unions, intersections, complements, ... - can be directly applied to events as subsets of Z. To each event we also assign a "measure" $\mu(E) = \Pr\{E\}$ called *probability of the event.* These values are assigned according to the function μ which has by assumption the following properties (or axioms):

1. $0 \leq \mu(E) \leq 1$;
2. $\mu(Z) = 1$;
3. For any finite or infinite sequence of disjoint sets (or mutually exclusive events) $E_1, E_2, ...$; such that $E_i \cap E_j = \emptyset$ for any i, j, we have

$$\mu\left(\cup_{i=1}^{N} E_i\right) = \sum_{i=1}^{N} \mu\left(E_i\right) \text{ where } N \text{ possibly equals } \infty.$$

All properties 1-3 are very intuitive for probabilities. Moreover, we would intuitively like to consider E as any subset of Z. Well, if Z is a finite or countable set then E can literally be *any* subset of Z. Unfortunately, when Z is a uncountably infinite set - such as the interval $[0, 1]$ for example - it might be impossible to find a function μ defined on all possible subsets of Z and at the same time satisfying all the three axioms we presented above. Typically, what fails is the last axiom of additivity when $N = \infty$. Lebesgue managed to keep property 3 above by defining the measure function μ only on the so-called *measurable sets* (or events). This is not an important limitation, as virtually all events of any practical interest tuned out to be measurable. Actually, in applications one typically considers only some class of possible events. A subset of the class of all measurable sets.

The reference class of sets $\mathcal{Z}$ represents the *set of possible events,* and will constitute a σ-algebra.[3] Notice that $\mathcal{Z}$ is a set of sets, hence an event E is an element of $\mathcal{Z}$, i.e. in contrast to $E \subset Z$ we will write $E \in \mathcal{Z}$. The pair $(Z, \mathcal{Z})$ constitutes a measurable space while the turple $(\mu, Z, \mathcal{Z})$ is denotes as a measured (or probability) space.

[3]A family $\mathcal{Z}$ of subsets of Z is called a σ algebra if: *(i)* both the empty set $\emptyset$ and Z belong to $\mathcal{Z}$; *(ii)* If $E \in \mathcal{Z}$ then also its complement (with respect to Z) $E^c = Z \backslash E \in \mathcal{Z}$; and *(iii)* for any sequence of sets such that $E_n \in \mathcal{Z}$ for all $n = 1, 2, ...$ we have that the set $(\cup_{n=1}^{\infty} E_n) \in \mathcal{Z}$. It is easy to show that whenever $\mathcal{Z}$ is a σ-algebra then $(\cap_{n=1}^{\infty} E_n) \in \mathcal{Z}$ as well. When Z is a set of real numbers, we can consider our set of possible events as the *Borel* σ-algebra. Which is σ-algebra 'generated' by the set of all open sets.

I am sure it is well known to you that the expectation operator $\mathbf{E}[\cdot]$ in (4.1) is nothing more than an integral, or a summation when z takes finitely or countably many values. For example, assume p_i is the probability that $z = z_i$. The expectation of the function f can be computed as follows

$$\mathbf{E}[f(z)] = \sum_{i=1}^{N} p_i f(z_i).$$

One of the advantages of the Lebesgue theory of integration is that, for example, it includes both summations and the usual concept of (Riemann) integration in an unified framework. We will be able to compute expectations[4]

$$\mathbf{E}[f(z)] = \int_Z f(z) d\mu(z)$$

no matter how Z is and no matter what is the distribution μ of the events. For example, we can deal with situations where Z is the interval $[0,1]$ and the event $z = 0$ has a positive probability $\mu(0) = p_0$. Since the set of all measurable events $\mathcal{Z}$ does not include all possible subsets of Z, we must restrict the set of functions f for which we can take expectations (integrals) as well.

Definition 32 *A real valued function f is measurable with respect to $\mathcal{Z}$ if for every real number x the set*

$$E_f^x = \{z \in Z : f(z) \geq x\}$$

belongs to the set of events $\mathcal{Z}$.

Sometimes we do a sort of inverse operation. We have in mind a class of real valued functions $\mathcal{F}$, each one defined over the set of events Z. We define a σ-algebra $\mathcal{Z}_{\mathcal{F}}$ so that to have every $f \in \mathcal{F}$ measurable, and any such function $f \in \mathcal{F}$ is called as *random variable*.

Definition 33 *The Lebesgue integral of a measurable positive function $f \geq 0$ is defined as follows*

$$\int_Z f(z) d\mu(z) = \sup_{0 \leq \phi \leq f} \int_Z \phi(z) d\mu(z) = \inf_{\phi \geq f \geq 0} \int_Z \phi(z) d\mu(z).$$

[4] When at z the measure μ has a density, the notation $d\mu(z)$ corresponds to the more familiar $f_\mu(z)\, dz$. When μ does not admits density, $d\mu(z)$ it is just the notation we use for its analogous concept.

In the definition, ϕ is any *simple (positive) function* (in its standard representation), that is, ϕ is a finite weighted sum of indicator functions[5]

$$\begin{aligned}\phi(z) &= \sum_{i=1}^{n} a_i I_{E_i}(z);\ a_i \geq 0;\ \text{and its integral is} \\ \int_Z \phi(z) d\mu(z) &= \sum_{i=1}^{n} a_i \mu(E_i),\end{aligned}$$

where for each i, j $E_i \cap E_j = \emptyset$; and $\cup_{i=1}^{n} E_i = Z$.

The Lebesgue integral of f is hence (uniquely) defined as the supremum of integrals of nonnegative dominated simple functions ϕ : such that for all $z,\ 0 \leq \phi(z) \leq f(z)$; which in turn coincides with the infimum over all the dominating simple functions: $\phi \geq f$. We do not have space here to discuss the implications of this definition[6] however, one should recall from basic analysis that the Riemann integral, that is the "usual" integral we saw in our undergraduate studies, can be defined in a similar way; where instead of simple functions one uses step functions. One can show that each function f which is Riemann integrable it is also Lebesgue integrable, and that there are simple examples where the converse is false.[7]

4.2 Markov Chains and Markov Processes

Markov Chains We now analyze in some detail conditional expectations for the simple case where Z is finite. So, assume that the stochastic component z can take finitely many values, that is $z \in Z \equiv \{z_1,\ z_2,\ ...,\ z_N\}$, with corresponding *conditional* probabilities

$$\pi_{ij} = \Pr\{z' = z_j \mid z = z_i\},\ i, j = 1, 2, ..., N.$$

[5]The indicator function of a set E is defined as

$$I_E(z) = \begin{cases} 1 \text{ if } z \in E \\ 0 \text{ otherwise.} \end{cases}$$

[6]See for example SLP, Ch. 7.

[7]One typical counter-example is the function $f : [0, 1] \to [0, 1]$ defined as follows

$$f(z) = \begin{cases} 1 \text{ if } z \text{ is rational} \\ 0 \text{ otherwise.} \end{cases}$$

This function is Lebesgue integrable with $\int f(x)dx = 0$, but it is not Riemann integrable.

Since π_{ij} describes the probability of the system to move to state z_j if the previous state was z_i, they are also called *transition probabilities* and the stochastic process form a *Markov chain*. To be probabilities, the π_{ij} must satisfy

$$\pi_{ij} \geq 0, \text{ and } \sum_{j=1}^{N} \pi_{ij} = 1 \text{ for } i = 1, 2, ..., N,$$

that is, they must belong to a $(N-1)-$dimensional simplex Δ^N. It is typically convenient to arrange the transition probabilities in a square array as follows

$$\Pi = \begin{bmatrix} \pi_{11} & \pi_{12} & \cdots & \pi_{1N} \\ \pi_{21} & \pi_{22} & \cdots & \cdots \\ \cdots & \cdots & \pi_{ij} & \cdots \\ \pi_{N1} & \cdots & \cdots & \pi_{NN} \end{bmatrix}$$

Such an array is called *transition matrix* or *Markov matrix*, or *stochastic matrix.* If the probability distribution over the state in period t is $p^t = (p_1^t, p_2^t, ...p_N^t)$, the distribution over the state in period $t+1$ is $p^t\Pi = \left(p_1^{t+1}, p_2^{t+1}, ...p_N^{t+1}\right)$, where

$$p_j^{t+1} = \sum_{i=1}^{N} p_i^t \pi_{ij}, \ j = 1, 2, ..., N.$$

For example, suppose we want to know what is the distribution of the next period states if in the current period the is z_i. Well, this means that the initial distribution is a degenerate one, namely $p^t = e_i = (0, \ ..., \ 1, \ ..., \ 0)$. As a consequence, the probability distribution over the next period state is the $i-$th row of $\Pi : e_i\Pi = (\pi_{i1}, \ \pi_{i2}, ...\pi_{iN})$. Similarly, if p^t is the period t distribution, then by the properties of the matrix multiplication, $p^t\Pi^n = p(\Pi \cdot \Pi \cdot ...\Pi)$ is the $t+n$ period distribution p^{t+n} over the states. It is easy to see that if Π is a Markov matrix then so is Π^n. A set of natural question then arises. Is there a stationary distribution, that is a probability distribution p^* with the property $p^* = p^*\Pi$? Under what conditions can we be sure that if we start from any initial distribution p^0, the system converges to a unique limiting probability $p^* = \lim_{n\to\infty} \{p^0\Pi^n\}$?

The answer to the first question turns out to always be affirmative for Markov chains.

Theorem 18 *Given a stochastic matrix* Π, *there always exists at least one stationary distribution* p^* *such that* $p^* = p^*\Pi$, *with* $p_i^* \geq 0$ *and* $\sum_{i=1}^{N} p_i^* = 1$.

Proof. Notice that a solution to the system of equations $p^* = p^*\Pi$ corresponds to solving $p^*(I - \Pi) = 0$, where I is the N dimensional identity matrix. Transposing both sides of the above equation gives

$$(I - \Pi')p^{*\prime} = 0.$$

So p^* is a nonnegative eigenvector associated with a unit eigenvalue of Π', normalized to satisfy $\sum_i p_i^* = 1$. So we can use linear algebra to show this result. Thanks to the Leontief's Input-Output analysis, during the 50s and 60s economics (re)discovered many important theorems about matrices with nonnegative elements. Any matrix with nonnegative elements has a Frobenius root $\lambda \geq 0$ with associated a nonnegative eigenvector. This existence result is the most difficult part of the proof and is due to Frobenius (1912) (See also Takayama, 1996, Theorem 4.B.2, pp. 375). Fisher (1965) and Takayama (1960) showed that when the elements of each column of the a matrix with nonnegative elements sum to one then its Frobenius root equals one, i.e. $\lambda = 1$ (Takayama, 1996, Theorem 4.C.11, pp. 388). The proof of this last statement is simple: let $p^* \geq 0$ the eigenvector associated with λ. By definition $\lambda p^* = \Pi' p^*$, that is, $\lambda p_i^* = \sum_j \pi'_{ij} p_j^*$, $i = 1, 2, ..., N$. Summing up over i, we obtain

$$\lambda \sum_{i=1}^{N} p_i^* = \sum_{i=1}^{N} \sum_{j=1}^{N} \pi'_{ij} p_j^* = \sum_{j=1}^{N} p_j^* \left(\sum_{i=1}^{N} \pi'_{ij} \right)$$

since Π' is the transpose of Π, $\sum_{i=1}^{N} \pi'_{ij} = \sum_{j=1}^{N} \pi_{ij} = 1$. Hence $\lambda = \frac{\sum_{j=1}^{N} p_j^*}{\sum_{i=1}^{N} p_i^*} = 1$. **Q.E.D.**

Consider now the second question. Can we say that p^* is unique? Unfortunately, in order to guarantee that the sequence of matrices converges to a unique matrix P^* with identical rows p^*, (so that for any p we have $pP^* = p^*$), we need some further assumptions, as the next exercises shows.

Exercise 41 *Assume that a Markov chain (with $Z = \{z_1, z_2\}$) is summarized by the following transition matrix*

$$\Pi = \begin{bmatrix} \pi_{11} & 1 - \pi_{11} \\ 1 - \pi_{22} & \pi_{22} \end{bmatrix}.$$

A stationary distribution is hence a vector $(q^*, 1 - q^*)$ with $1 \geq q^* \geq 0$ such that $(q^*, 1 - q^*) \cdot \Pi = (q^*, 1 - q^*)$. We know from above that at least one such q^* must exists.
(a) *Using simple algebra show that q^* solves $(2 - \pi_{22} - \pi_{11}) q^* = (1 - \pi_{22})$, and discuss conditions where q^* might take multiple values.*
(b) *Now set $\pi_{11} = \pi_{22} = \pi$ and state conditions for q^* to be unique.*

Here is a set of sufficient conditions for uniqueness.

Theorem 19 *Assume that $\pi_{ij} > 0$ for all $i, j = 1, 2, ...N$. There exists a limiting distribution p^* such that*

$$p_j^* = \lim_{n\to\infty} \pi_{ij}^{(n)},$$

where $\pi_{ij}^{(n)}$ is the (i, j) element of the matrix Π^n. And p_j^* are the unique nonnegative solutions of the following system of equations

$$\begin{aligned} p_j^* &= \sum_{k=1}^N p_k^* \pi_{kj}; \text{ or } p^* = p^*\Pi; \text{ and} \\ \sum_{j=1}^N p_j^* &= 1. \end{aligned}$$

Proof. See below. **Q.E.D.**

The application of the transition matrix on a probability distribution p can be seen as a mapping of the $(N-1)-$dimensional simplex into itself. In fact, under some conditions, the operator

$$\begin{aligned} T_\Pi &: \Delta^N \to \Delta^N \\ T_\Pi p &= p\Pi \end{aligned} \tag{4.2}$$

defines a contraction on the metric space $\left(\Delta^N, |\cdot|_N\right)$ where

$$|x|_N \equiv \sum_{i=1}^N |x_i|.$$

Exercise 42 *(i) Show that $\left(\Delta^N, |\cdot|_N\right)$ is a complete metric space. (ii) Moreover, show that if $\pi_{ij} > 0$, $i, j = 1, 2, ...$; the mapping T in (4.2) is a contraction of modulus $\beta = 1-\varepsilon$, where $\varepsilon = \sum_{j=1}^N \varepsilon_j$ and $\varepsilon_j = \min_i \pi_{ij} > 0$.*

When some $\pi_{ij}^{(n)} = 0$, we might loose uniqueness. However, following the same line of proof one can show that the stationary distribution is unique as long as $\varepsilon = \sum_{j=1}^N \varepsilon_j > 0$. Could you explain intuitively why this is the case?

Moreover, from the contraction mapping theorem, it is easy to see that the above proposition remains valid if the assumption $\pi_{ij} > 0$ is replaced with: there exists a $n \geq 1$ such that $\pi_{ij}^{(n)} > 0$ for all i, j. (see Corollary 2 of the contraction mapping Theorem (Th. 3.2) in SLP).

Notice that the sequence $\{\Pi^n\}_{n=0}^\infty$ might not always converge. For example, consider $\Pi = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. It is easy to verify that the sequence jumps from $\Pi^{2n} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ and

$\Pi^{2n+1} = \Pi$. However, the fact that in a Markov chain the state space is finite implies that the long-run averages

$$\left\{ \frac{1}{T} \sum_{t=0}^{T-1} \Pi^t \right\}_{T=1}^{\infty}$$

do always converge to a stochastic matrix P^*, and the sequence $p^t = p^0 \Pi^t$ converges to

$$\lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} p^t = p^0 P^*.$$

In the example we saw above one can easily verify that $\frac{1}{T} \sum_{t=0}^{T-1} \Pi^t \to P^* = \begin{bmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{bmatrix}$, and the unique stationary distribution is $p^* = (1/2, 1/2)$.

In other cases, the rows of the limit matrix P^* are not necessarily always identical to each other. For example, consider now the transition matrix $\Pi = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. It is obvious that in this case $P^* = \Pi$, which has two different rows. It is also clear that both rows constitute a stationary distribution. This is true in general: any row of the limit matrix P^* is an invariant distribution for the transition matrix Π.

What is perhaps less obvious is that any convex combination of the rows of P^* constitute a stationary distribution, and that all invariant distributions for Π can be derived by making convex combinations of the rows of P^*.

Exercise 43 *(i) Consider first the above example with* $P^* = \Pi = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. *Show that any vector* $p^*_\lambda = (\lambda, 1 - \lambda)$ *obtained as a convex combination of the rows of* P^* *constitutes a stationary distribution for* Π. *Provide an intuition for the result.* (ii) *Now consider the general case, and let* p^* *and* p^{**} *two stationary distributions for a Markov chains defined by a generic stochastic matrix* Π. *Show that any convex combination* p_λ *of* p^* *and* p^{**} *constitute a stationary distribution for* Π.

Markov Processes The more general concept corresponding to a Markov chain, where Z can take countably or uncountably many values, is denoted as a Markov Process. Similarly to the case where Z is finite, a Markov process is defined by a *transition function* (or kernel) $Q : Z \times \mathcal{Z} \to [0, 1]$ such that: *(i)* for each $z \in Z$ $Q(z, \cdot)$ is a probability measure; and *(ii)* for each $C \in \mathcal{Z}$ $Q(\cdot, C)$ is a measurable function.

Given Q, one can compute conditional probabilities

$$\Pr \{z_{t+1} \in C \mid z_t = c\} = Q(c, C)$$

and conditional expectations in the usual way

$$\mathbf{E}\left[f \mid z\right] = \int_Z f(z')dQ(z,z').$$

Notice that Q can be used to map probability measure into probability measures since for any μ on $(Z, \mathcal{Z})$ we get a new $\mu\prime$ by assigning to each $C \in \mathcal{Z}$ the measure

$$\left(T_Q\mu\right)(C) = \mu'(C) = \int_Z Q(z,C)d\mu(z),$$

and T is denoted as *Markov operator*.

We now define a very useful property for Q.

Definition 34 *Q has the **Feller property** if for any bounded and continuous function f the function*

$$g(z) = (\mathbf{P}_Q f)(z) = \mathbf{E}\left[f \mid z\right] = \int f(z')dQ(z,z') \textit{ for any } z$$

is still bounded and continuous.

The above definition first of all shown another view of Q. It also defines an operator (sometimes called transition operator) that in general maps bounded and measurable functions into bounded measurable functions. When Q has the feller property the operator $\mathbf{P}_Q$ preserves continuity.

Technical Digression (optional). It turns out that the Feller property characterizes continuous Markov transitions. The rigorous idea is simple. Let $\mathbf{M}$ be the set of all probability measures on Borel sets $\mathcal{Z}$ over a metrizable space Z, and for each z, let $Q(z,\cdot)$ a member of $\mathbf{M}$. The usual topology defined in the space of Borel measures is the topology of convergence in distribution (or weak topology).[8] It is now useful to make pointwise considerations. For each z the probability measure $Q(z,\cdot)$ can be seen as a linear mapping from the set of bounded and measurable functions into the real numbers according to $x = \langle f, Q(z,\cdot)\rangle = \int f(z')dQ(z,z')$.

It turns out that a transition function $Q : Z \to \mathbf{M}$ is continuous if and only if it has the Feller property. The fact that a continuous Q has the Feller property is immediate: By definition of the topology defined on $\mathbf{M}$ (weak topology), via the map $F_f(\mu) = \langle f, \mu\rangle$ each

[8] In this topology, a sequence $\{\mu_n\}$ in M converges to μ if and only if $\int f d\mu_n \to \int f d\mu$ for all continuous and bounded functions f.

continuous and bounded function $f : Z \to \mathbb{R}$ defines a continuous real valued function $F_f : \mathbf{M} \to \mathbb{R}$.[9] Now note that when μ is $Q(z, \cdot)$ we have $F_f(Q(z, \cdot)) = (\mathbf{P}_Q f)(z)$. Now, continuity of Q means that as $z_n \to z$ we have $Q(z_n, \cdot) \to Q(z, \cdot)$ in $\mathbf{M}$. Equivalently, if we let $\mu_n(\cdot) = Q(z_n, \cdot)$ and use the usual topology on $\mathbf{M}$, continuity of Q means that $F_f(Q(z_n, \cdot)) \to F_f(Q(z, \cdot))$ (interpreted now as sequence of real numbers). We have hence established that $(\mathbf{P}_Q f) = g$ is a continuous function in Z, i.e. that Q has the Feller property. In order to show rigorously that the Feller property implies continuity - although it is intuitive - one needs some more work.[10]

We can now study the issue of existence and uniqueness of a stationary distribution. A stationary distribution for Q is a measure μ^* on $(Z, \mathcal{Z})$ such that for any $C \in \mathcal{Z}$

$$\mu^*(C) = \int_Z Q(z, C) d\mu^*(z),$$

that is μ^*, is a fixed point of the Markov operator T_Q. There are many results establishing existence and uniqueness of a stationary distribution. Here is a result which is among the easiest to understand, and that uses the Feller property of Q.

Theorem 20 *If Z is a compact set and Q has the Feller property then there exists a stationary distribution $\mu^* : \mu^* = T_Q \mu^*$, where $\mu = \lambda$ if and only if $\int f d\mu = \int f d\lambda$ for each continuous and bounded function f.*

Proof. See SLP, Theorem 12.10, page 376-77. The basic idea of the proof can also be get as an application of one of the infinite dimensional extensions of the Brower fixed point theorem (usually called Brower-Shauder-Tyconoff fixed point). We saw above that whenever Q has the Feller property, the associated Markov operator T_Q is a continuous map from the compact convex (locally convex Hausdorff) space of distributions Λ into itself. [See Aliprantis and Border (1994), Corollary 14.51, page 485] **Q.E.D.**

Similarly to the finite state case, this invariant measure can be obtained by looking at the sequence $\left\{ \frac{1}{T} \sum_{t=1}^{T-1} T_Q^t \lambda_0 \right\}_{T=1}^{\infty}$ of T-period averages.

When the state space is not finite, we may define several different concepts of converge for distributions. The most known ones are weak convergence (commonly denoted convergence in distribution) and strong convergence (or convergence in total variation norm, also denoted as setwise convergence). We are not dealing with these issues in these class notes. The concept of weak convergence is in most cases all that we care about in the

[9] Let $x_{\mu_n} = F_f(\mu_n)$. By definition of weak topology, if $\mu_n \to \mu$ then $F_f(\mu_n) \to F_f(\mu)$.

[10] The interest reader can have a look at Aliprantis and Border (1994), Theorem 15.14, page 531-2.

context of describing the dynamics of an economic system. Theorem 20 deals with weak convergence. The most known results of uniqueness use some monotonicity conditions on the Markov operator, together with some mixing conditions. For a quite general treatment of monotonic Markov operators, with direct applications to economics and dynamic programming, see Hopenhayn and Prescott (1992).

If we require strong convergence, one can guarantee uniqueness under conditions similar to those of Theorem 19, using the contraction mapping theorem. See Chapter 11 in SLP, especially Theorem 11.12.

4.3 Bellman Principle in the Stochastic Framework

The Finite Z case. When the shocks belong to a finite set all the results we saw for the deterministic case are true for the stochastic environment as well. The Bellman Principle of optimality remains true since both Lemma 1 and 2 remain true. Expectation are simply a weighted sums of the continuation values. In this case Theorem 12 remains true under the same conditions as in the deterministic case. From the proof of Theorem 13 and 14 it is easy to see that also the verification and sufficiency theorems can easily be extended to the stochastic case with finite shocks. We just need to require boundedness to be true for all z. Even the Theorems 15 and 16 are easily extended to the stochastic case following the same lines of proof we proposed in Chapter 3.1. In order to show you that there is practically no difference between the deterministic and the stochastic case when Z is finite, let me be a bit boring and consider for example the stochastic extension of Theorem 15. Assume w.l.o.g. that z may take N values, i.e. $Z = (z_1,\ z_2,\ ...,\ z_N)$. We can always consider our fixed point

$$V(x, z_i) = \sup_{x' \in \Gamma(x, z_i)} F(x, x', z_i) + \beta \sum_{j=1}^{N} \pi_{ij} V(x', z_j),\ \forall i$$

in the space $\mathcal{C}_N(X)$ of vectors of real valued functions:

$$\mathbf{V}(x) = (V(x, z_1), ..., V(x, z_N)) = (V_1(x), ..., V_N(x))$$

which are continuous and bounded in X with the metric d_∞^N, where[11]

$$d_\infty^N(\mathbf{V},\mathbf{W}) = \sum_{i=1}^{N} d_\infty\left(V_i, W_i\right) = \sum_{i=1}^{N} \sup_x \left|V(x,z_i) - W(x,z_i)\right|.$$

One can easily show that such metric space of functions is complete, and that the same conditions for a contraction in the deterministic case can be used here to show that the operator

$$\begin{aligned} T &: \mathcal{C}_N(X) \to \mathcal{C}_N(X) \\ T\mathbf{V}(x) &= \left\{ \begin{array}{c} \sup_{x'\in\Gamma(x,z_1)} F(x,x',z_1) + \beta \sum_{j=1}^{N} \pi_{1j} V(x',z_j) \\ \sup_{x'\in\Gamma(x,z_2)} F(x,x',z_2) + \beta \sum_{j=1}^{N} \pi_{2j} V(x',z_j) \\ \dots \\ \sup_{x'\in\Gamma(x,z_N)} F(x,x',z_N) + \beta \sum_{j=1}^{N} \pi_{Nj} V(x',z_j) \end{array} \right. \end{aligned}$$

is a contraction with modulus β. It is easy to see that both boundedness and - by the Theorem of the Maximum - continuity is preserved under T. Similarly, given that (conditional) expectations are nothing more than convex combinations, concavity is preserved under T, and the same conditions used for the deterministic case can be assumed here to guarantee the stochastic analogous to Theorem 16.

The General case When Z is continuous, we need to use measure theory. We need to assume some additional technical restrictions to guarantee that the integrals involved in the expectations and the limits inside those integrals are well defined.

Unfortunately, these technical complications prevent the possibility of having a result on the lines of Theorem 12. The reason is that we one cannot be sure that the true value function is measurable. As a consequence, the typical result in this case are in form of the verification or sufficiency theorems. Before stating formally the result we need to introduce some notation.

Definition 35 *A plan π is an initial value $\pi_0 \in X$ and a sequence of (h^t-measurable) functions*[12]

$$\pi_t : H^t \to X$$

[11] Another possibility is to use $d_\infty^{\max}$

$$d_\infty^{\max}(\mathbf{V},\mathbf{W}) = \max_i \left\{d_\infty\left(V_i, W_i\right)\right\} = \max_i \left\{ \sup_x \left|V(x,z_i) - W(x,z_i)\right| \right\}.$$

[12] A function is said to be h^t-measurable when it is measurable with respect to the $\sigma-$algebra generated by the set of all possible h^t histories H^t.

for all $t \geq 1$, where H^t is the set of all length-t histories of shocks: $h^t = (z_0, z_1, ...z_t)$, $z_t \in Z$.

That is, $\pi_t(h^t)$ is the value of the endogenous state x_{t+1} that is chosen in period t, when the (partial) history up to this moment is h^t. So, in a stochastic framework agents are taking contingent plans. They are deciding what to do for any possible history, even though some of these histories are never going to happen. Moreover, for any partial history $h^t \in H^t$ one can define a probability measure $\mu^t : \mu^t(C) = \Pr\{h^t \in C \subseteq H^t\}$. In this environment, feasibility is defined similarly to the deterministic case. We say that the plan π is feasible, and write $\pi \in \Pi(x_0, z_0)$ if $\pi_0 \in \Gamma(x_0, z_0)$ and for each $t \geq 1$ and h^t we have $\pi_t(h^t) \in \Gamma(\pi_{t-1}(h^{t-1}), z_t)$. We will always assume that F, Γ, β and μ are such that $\Pi(x_0, z_0)$ is nonempty for any $(x_0, z_0) \in X \times Z$, and that the objective function

$$\begin{aligned} \mathbf{U}(\pi) &= \lim_{T\to\infty} F(x_0, \pi_0, z_0) + \sum_{t=1}^{T} \beta^t \int_{H^t} F\left(\pi_{t-1}(h^{t-1}), \pi_t(h^t), z_t\right) d\mu^t(h^t) \\ &= \lim_{T\to\infty} F(x_0, \pi_0, z_0) + \sum_{t=1}^{T} \beta^t \mathbf{E}_0 \left[F\left(\pi_{t-1}(h^{t-1}), \pi_t(h^t), z_t\right)\right] \end{aligned}$$

is well defined for any $\pi \in \Pi(x_0, z_0)$ and (x_0, z_0). Similarly to the compact notation for the deterministic case, the true value function V^* is defined as follows

$$V^*(x_0, z_0) = \sup_{\pi\in\Pi(x_0,z_0)} \mathbf{U}(\pi). \tag{4.3}$$

Let me first state a verification theorem for the stochastic case.

Theorem 21 *Assume that $V(x, z)$ is a measurable function which satisfies the Bellman equation (4.1). Moreover, assume that*

$$\lim_{t\to\infty} \beta^{t+1}\mathbf{E}_0 \left[V(\pi_t(h^t), z_{t+1})\right] = 0$$

for every possible contingent plan $\pi \in \Pi(x_0, z_0)$ for all $(x_0, z_0) \in X \times Z$; and that the policy correspondence

$$G(x, z) = \left\{x' \in \Gamma(x, z) : V(x, z) = F(x, x', z) + \beta \int_Z V(x', z')dQ(z, z')\right\} \tag{4.4}$$

is non empty and permits a measurable selection. Then $V = V^*$ and all plans generated by G are optimal.

Proof. The idea of the proof follows very closely the lines of Theorems 13 and 14. A plan that solves the Bellman equation and that does not have any left-over value at infinity, is optimal. Of course, we must impose few additional technical conditions imposed by measure theory.[13] For details the reader can see Chapter 9 of SLP. **Q.E.D.**

In order to be able to recover Theorem 12 we need to make an assumption on the *endogenous* V^* :

Theorem 22 *Let* F *be bounded and measurable. If the value function* $V^*(x_0, z_0)$ *defined in (4.3) is measurable and assume that the correspondence analogous to (4.4) admits a measurable selection. Then* $V^*(x_0, z_0)$ *satisfies the functional equation (4.1) for all* (x_0, z_0) , *and any optimal plan* π^* *(which solves (4.3)) also solves*

$$V^*(\pi^*_{t-1}(h^{t-1}), z_t) = F(\pi^*_{t-1}(h^{t-1}), \pi^*_t(h^t), z_t) + \beta \int V^*(\pi^*_t(h^t), z_{t+1}) dQ\,(z_t, z_{t+1}),$$

$\mu^t(\cdot)$ *almost surely for all* t *and* h^t *emanating from* z_0.

Proof. The idea of the proof is similar to that of Theorem 12. For the several details however, the reader is demanded to Theorem 9.4 in SLP. **Q.E.D.**

Let finally state the corresponding of Theorems 15 and 16 for the stochastic environment allowing for continuous shocks.

Theorem 23 *Assume* F *is continuous and bounded;* Γ *compact valued and continuous;* Q *possesses the Feller property,* $\beta \in [0, 1)$ *and* X *is a closed and convex subset of* $\mathbb{R}^l$. *Then the Bellman operator* T

$$(TW)(x, z) = \max_{x' \in \Gamma(x,z)} F(x, x', z) + \beta \int_Z W(x', z') dQ(z, z')$$

has a unique fixed point V *in the space of continuous and bounded functions.*

Proof. Once we have noted that the Feller property of Q guarantees that if W is bounded and continuous function then $\int_Z W(x', z')dQ(z, z')$ is also bounded and continuous for all (x', z), we can apply basically line by line the proof of Theorem 15. **Q.E.D.**

Theorem 24 *Assume* F *is concave continuous and bounded;* Γ *is continuous and with convex graph;* Q *possesses the Feller property,* $\beta \in [0, 1)$ *and* X *is a closed and convex subset of* $\mathbb{R}^l$. *Then the Bellman operator has a unique fixed point* V *in the space of concave, continuous and bounded functions.*

[13] For example, the policy correspondence G permits a measurable selection if there exists a function $h : X \times Z \to X$, such that $h(x, z) \in G(x, z)$ for all $(x, z) \in X \times Z$.

Proof. Again the proof is similar to the deterministic case. Once we have noted that the linearity of the integral preserves concavity (since $\int_Z dQ(z, z') = 1$) we can basically apply line by line the proof of Theorem 16. **Q.E.D.**

It is important to notice that whenever the conditions of Theorem 23 are met, the boundedness of V and an application of the Maximum Theorem imply the conditions of Theorem 21 are also satisfied, hence $V = V^*$ which is a continuous function (hence measurable). In this case the Bellman equation fully characterizes the optimization problem also with uncountably many possible levels of the shock.

4.4 The Stochastic Model of Optimal Growth

Consider the stochastic version of the optimal growth model

$$\begin{aligned} V(k_0, z_0) &= \sup_{\{k_{t+1}\}_{t=0}^{\infty}} \mathbf{E}_0 \left[\sum_{t=0}^{\infty} \beta^t u\left(f(z_t, k_t) - k_{t+1}\right) \right] \\ s.t.\ 0 &\leq k_{t+1} \leq f(z_t, k_t) \quad \text{for all } t \\ k_0 &\in X,\ z_0 \in Z, \end{aligned}$$

where the expectation is over the sequence of shocks $\{z_t\}_{t=0}^{\infty}$. Assume that $\{z_t\}_{t=0}^{\infty}$ is an i.i.d. sequence of shocks, each drawn according to the probability measure μ on $(Z, \mathcal{Z})$.

Exercise 44 *Let $u(c) = \ln c$ and $f(z, k) = zk^{\alpha}$, $0 < \alpha < 1$ (so $\delta = 1$). I tell you that the optimal policy function takes the form $k_{t+1} = \alpha\beta z_t k_t^{\alpha}$ for any t and z_t. (i) Use this fact to calculate an expression for the optimal policy $\pi_t^*(h^t)$ [recall that $h^t = (z_0, ...z_t)$] and the value function $V^*(k_0, z_0)$ for any initial values (k_0, z_0), and verify that V^* solves the following Bellman equation*

$$V(k, z) = \max_{0 \leq k' \leq zk^{\alpha}} \ln(zk^{\alpha} - k') + \beta \mathbf{E}\left[V(k', z')\right].$$

(ii) Now show that a solution to the above functional equation is

$$V(k, z) = A(z) + \frac{\alpha}{1 - \beta\alpha} \ln k,$$

and discuss the relationship between V^* and V.

This model can be extended in many directions. This model with persistent shocks and non inelastic labor supply has been used in the Real Business Cycles literature to study the effects of technological shocks on aggregate variables like consumption and employment. This line of research started in the 80s, and for many macroeconomists is still the building block for any study about the aggregate real economy. RBC will be the next topic of these notes. Moreover, since most interesting economic problem do not have closed forms, you must first learn how to use numerical methods to approximate V and perform simulations.

Bibliography

[1] Bertsekas, D. P. (1976) *Dynamic Programming and Stochastic Control*, Academic Press, NY.

[2] Bertsekas, D. P., and S. E. Shreve (1978) *Stochastic Control: The Discrete Time Case.* New York: Academic Press.

[3] Rudin, W. *Principles of Mathematical Analysis*, McGraw Hill.

[4] Stokey, N., R. Lucas and E. Prescott (1991), *Recursive Methods for Economic Dynamics*, Harvard University Press.