5 Dynamic Programming

5.1 The Principle of Optimality

In the last section we showed that under certain conditions, the functional equation (FE)

$$v(x) = \sup_{y \in \Gamma(x)} \{F(x,y) + \beta v(y)\}$$

has a unique solution which is approached from any initial guess v_0 at geometric speed. What we were really interested in, however, was a problem of sequential form (SP)

$$\begin{aligned} w(x_0) &= \sup_{\{x_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t F(x_t, x_{t+1}) \\ s.t.\ x_{t+1} &\in \Gamma(x_t) \\ x_0 &\in X \text{ given} \end{aligned}$$

Note that I replaced max with sup since we have not made any assumptions so far that would guarantee that the maximum in either the functional equation or the sequential problem exists. In this section we want to find out under what conditions the functions v and w are equal and under what conditions optimal sequential policies $\{x_{t+1}\}_{t=0}^{\infty}$ are equivalent to optimal policies $y = g(x)$ from the recursive problem, i.e. under what conditions the principle of optimality holds. It turns out that these conditions are very mild.

In this section I will try to state the main results and make clear what they mean; I will not prove the results. The interested reader is invited to consult Stokey and Lucas or Bertsekas. Unfortunately, to make our results precise additional notation is needed. Let X be the set of possible values that the state x can take. X may be a subset of a Euclidean space, a set of functions or something else; we need not be more specific at this point. The correspondence $\Gamma : X \Rightarrow X$ describes the feasible set of next period's states y, given that today's state is x. The graph of Γ, A is defined as

$$A = \{(x,y) \in X \times X : y \in \Gamma(x)\}$$

The period return function $F : A \to \mathbf{R}$ maps the set of all feasible combinations of today's and tomorrow's state into the reals. So the fundamentals of our analysis are (X, F, β, Γ). For the neoclassical growth model F and β describe preferences and X, Γ describe the technology.

We call any sequence of states $\{x_t\}_{t=0}^{\infty}$ a plan. For a given initial condition x_0, the set of feasible plans $\Pi(x_0)$ from x_0 is defined as

$$\Pi(x_0) = \{\{x_t\}_{t=1}^{\infty} : x_{t+1} \in \Gamma(x_t)\}$$

Hence $\Pi(x_0)$ is the set of sequences that, for a given initial condition, satisfy all the feasibility constraints of the economy. We will denote by $\bar{x}$ a generic element of $\Pi(x_0)$. The two assumptions that we need for the principle of optimality are basically that for any initial condition x_0 the social planner (or whoever solves the problem) has at least one feasible plan and that the total return (the total utility, say) from all feasible plans can be evaluated. That's it. More precisely we have

Assumption 1: $\Gamma(x)$ is nonempty for all $x \in X$

Assumption 2: For all initial conditions x_0 and all feasible plans $\bar{x} \in \Pi(x_0)$

$$\lim_{n\to\infty} \sum_{t=0}^{n} \beta^t F(x_t, x_{t+1})$$

exists (although it may be $+\infty$ or $-\infty$).

Assumption 1 does not require much discussion: we don't want to deal with an optimization problem in which the decision maker (at least for some initial conditions) can't do anything. Assumption 2 is more subtle. There are various ways to verify that assumption 2 is satisfied, i.e. various sets of sufficient conditions for assumption 2 to hold. Assumption 2 holds if

1. F is bounded and $\beta \in (0,1)$. Note that boundedness of F is not enough. Suppose $\beta = 1$ and $F(x_t, x_{t+1}) = \begin{cases} 1 \text{ if } t \text{ even} \\ -1 \text{ if } t \text{ odd} \end{cases}$ Obviously F is bounded, but since $\sum_{t=0}^{n} \beta^t F(x_t, x_{t+1}) = \begin{cases} 1 \text{ if } n \text{ even} \\ 0 \text{ if } n \text{ odd} \end{cases}$, the limit in assumption 2 does not exist. If $\beta \in (0,1)$ then $\sum_{t=0}^{n} \beta^t F(x_t, x_{t+1}) = \begin{cases} 1 - \beta^{\frac{n}{2}} + \beta^n \text{ if } n \text{ even} \\ 1 - \beta^{\frac{n}{2}} \text{ if } n \text{ odd} \end{cases}$ and therefore $\lim_{n\to\infty} \sum_{t=0}^{n} \beta^t F(x_t, x_{t+1})$ exists and equals 1. In general the joint assumption that F is bounded and $\beta \in (0,1)$ implies that the sequence $y_n = \sum_{t=0}^{n} \beta^t F(x_t, x_{t+1})$ is Cauchy and hence converges. In this case $\lim y_n = y$ is obviously finite.

2. Define $F^+(x,y) = \max\{0, F(x,y)\}$ and $F^-(x,y) = \max\{0, -F(x,y)\}$. Then assumption 2 is satisfied if for all $x_0 \in X$, all $\bar{x} \in \Pi(x_0)$, either
$$\begin{aligned} \lim_{n\to\infty} \sum_{t=0}^{n} \beta^t F^+(x_t, x_{t+1}) &< +\infty \text{ or} \\ \lim_{n\to\infty} \sum_{t=0}^{n} \beta^t F^-(x_t, x_{t+1}) &< +\infty \end{aligned}$$
or both. For example, if $\beta \in (0,1)$ and F is bounded above, then the first condition is satisfied, if $\beta \in (0,1)$ and F is bounded below then the second condition is satisfied.

3. Assumption 2 is satisfied if for every $x_0 \in X$ and every $\bar{x} \in \Pi(x_0)$ there are numbers (possibly dependent on $x_0, \bar{x}$) $\theta \in (0, \frac{1}{\beta})$ and $0 < c < +\infty$ such that for all t
$$F(x_t, x_{t+1}) \leq c\theta^t$$

Hence F need not be bounded in any direction for assumption 2 to be satisfied. As long as the returns from the sequences do not grow too fast (at rate higher than $\frac{1}{\beta}$) we are still fine .

I would conclude that assumption 2 is rather weak (I can't think of any interesting economic example where assumption1 is violated, but let me know if you come up with one). A final piece of notation and we are ready to state some theorems.

Define the sequence of functions $u_n : \Pi(x_0) \to \mathbf{R}$ by

$$u_n(\bar{x}) = \sum_{t=0}^{n} \beta^t F(x_t, x_{t+1})$$

For each feasible plan u_n gives the total discounted return (utility) up until period n. If assumption 2 is satisfied, then the function $u : \Pi(x_0) \to \bar{\mathbf{R}}$

$$u(\bar{x}) = \lim_{n\to\infty} \sum_{t=0}^{n} \beta^t F(x_t, x_{t+1})$$

is also well-defined, since under assumption 2 the limit exists. The range of u is $\bar{\mathbf{R}}$, the extended real line, i.e. $\bar{\mathbf{R}} = \mathbf{R} \cup \{-\infty, +\infty\}$ since we allowed the limit to be plus or minus infinity. From the definition of u it follows that under assumption 2

$$w(x_0) = \sup_{\bar{x}\in\Pi(x_0)} u(\bar{x})$$

Note that by construction, whenever w exists, it is unique (since the supremum of a set is always unique). Also note that the way I have defined w above only makes sense under assumption 1. and 2., otherwise w is not well-defined.

We have the following theorem, stating the principle of optimality.

Theorem 40 *Suppose (X, Γ, F, β) satisfy assumptions 1. and 2. Then*

1. *the function w satisfies the functional equation (FE)*
2. *if for all $x_0 \in X$ and all $\bar{x} \in \Pi(x_0)$ a solution v to the functional equation (FE) satisfies*

$$\lim_{n\to\infty} \beta^n v(x_n) = 0 \tag{32}$$

then $v = w$

I will skip the proof, but try to provide some intuition. The first result states that the supremum function from the sequential problem (which is well-defined under assumption 1. and 2.) solves the functional equation. This result, although nice, is not particularly useful for us. We are interested in solving the sequential problem and in the last section we made progress in solving the functional equation (not the other way around).

But result 2. is really key. It states a condition under which a solution to the functional equation (which we know how to compute) is a solution to the sequential problem (the solution of which we desire). Note that the functional equation (FE) may (or may not) have several solution. We haven't made enough assumptions to use the CMT to argue uniqueness. However, only one of these potential several solutions can satisfy (32) since if it does, the theorem tells us that it has to equal the supremum function w (which is necessarily unique). The condition (32) is somewhat hard to interpret (and SLP don't even try), but think about the following. We saw in the first lecture that for infinite-dimensional optimization problems like the one in (SP) a transversality condition was often necessary and (even more often) sufficient (jointly with the Euler equation). The transversality condition rules out as suboptimal plans that postpone too much utility into the distant future. There is no equivalent condition for the recursive formulation (as this formulation is basically a two period formulation, today vs. everything from tomorrow onwards). Condition (32) basically requires that the continuation utility from date n onwards, discounted to period 0, should vanish in the time limit. In other words, this puts an upper limit on the growth rate of continuation utility, which seems to substitute for the TVC. It is not clear to me how to make this intuition more rigorous, though.

A simple, but quite famous example, shows that the condition (32) has some bite. Consider the following consumption problem of an infinitely lived household. The household has initial wealth $x_0 \in X = \mathbf{R}$. He can borrow or lend at a gross interest rate $1 + r = \frac{1}{\beta} > 1$. So the price of a bond that pays off one unit of consumption is $q = \beta$. There are no borrowing constraints, so the sequential budget constraint is

$$c_t + \beta x_{t+1} \leq x_t$$

and the nonnegativity constraint on consumption, $c_t \geq 0$. The household values discounted consumption, so that her maximization problem is

$$\begin{aligned} w(x_0) &= \sup_{\{(c_t, x_{t+1})\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t c_t \\ s.t.\ 0 &\leq c_t \leq x_t - \beta x_{t+1} \\ & \quad x_0 \text{ given} \end{aligned}$$

Since there are no borrowing constraint, the consumer can assure herself infinite utility by just borrowing an infinite amount in period 0 and then rolling over the debt by even borrowing more in the future. Such a strategy is called a Ponzi-scheme -see the hand-out. Hence the supremum function equals $w(x_0) = +\infty$ for all $x_0 \in X$. Now consider the recursive formulation (we denote by x current period wealth x_t, by y next period's wealth and substitute out for consumption $c_t = x_t - \beta x_{t+1}$ (which is OK given monotonicity of preferences)

$$v(x) = \sup_{y \leq \frac{x}{\beta}} \{x - \beta y + \beta v(y)\}$$

Obviously the function $w(x) = +\infty$ satisfies this functional equation (just plug in w on the right side, since for all x it is optimal to let y tend to $-\infty$ and hence $v(x) = +\infty$. This should be the case from the first part of the previous theorem. But the function $\check{v}(x) = x$ satisfies the functional equation, too. Using it on the right hand side gives, for an arbitrary $x \in X$

$$\sup_{y \leq \frac{x}{\beta}} \{x - \beta y + \beta y\} = \sup_{y \leq \frac{x}{\beta}} x = x = \check{v}(x)$$

Note, however that the second part of the preceding theorem does not apply for $\check{v}$ since the sequence $\{x_n\}$ defined by $x_n = \frac{x_0}{\beta^n}$ is a feasible plan from $x_0 > 0$ and

$$\lim_{n\to\infty} \beta^n v(x_n) = \lim_{n\to\infty} \beta^n x_n = x_0 > 0$$

Note however that the second part of the theorem gives only a sufficient condition for a solution v to the functional equation being equal to the supremum function from (SP), but not a necessary condition. Also w itself does not satisfy the condition, but is evidently equal to the supremum function. So whenever we can use the CMT (or something equivalent) we have to be aware of the fact that there may be several solutions to the functional equation, but at most one the several is the function that we look for.

Now we want to establish a similar equivalence between the sequential problem and the recursive problem with respect to the optimal policies/plans. The first observation. Solving the functional equation gives us optimal policies $y = g(x)$ (note that g need not be a function, but could be a correspondence). Such an optimal policy induces a feasible plan $\{\hat{x}_{t+1}\}_{t=0}^{\infty}$ in the following fashion: $x_0 = \hat{x}_0$ is an initial condition, $\hat{x}_1 \in g(\hat{x}_0)$ and recursively $\hat{x}_{t+1} = g(\hat{x}_t)$. The basic question is how a plan constructed from a solution to the functional equation relates to a plan that solves the sequential problem. We have the following theorem.

Theorem 41 *Suppose (X, Γ, F, β) satisfy assumptions 1. and 2.*

1. *Let $\bar{x} \in \Pi(x_0)$ be a feasible plan that attains the supremum in the sequential problem. Then for all $t \geq 0$*

$$w(\bar{x}_t) = F(\bar{x}_t, \bar{x}_{t+1}) + \beta w(\bar{x}_{t+1})$$

2. *Let $\hat{x} \in \Pi(x_0)$ be a feasible plan satisfying, for all $t \geq 0$*

$$w(\hat{x}_t) = F(\hat{x}_t, \hat{x}_{t+1}) + \beta w(\hat{x}_{t+1})$$

and additionally[26]

$$\lim_{t\to\infty} \sup \beta^t w(\hat{x}_t) \leq 0 \tag{33}$$

Then $\hat{x}$ attains the supremum in (SP) for the initial condition x_0.

---

[26] The limit superior of a bounded sequence $\{x_n\}$ is the infimum of the set V of real numbers v such that only a finite number of elements of the sequence strictly exceed v. Hence it is the largest cluster point of the sequence $\{x_n\}$.

What does this result say? The first part says that any optimal plan in the sequence problem, together with the supremum function w as value function satisfies the functional equation for all t. Loosely it says that any optimal plan from the sequential problem is an optimal policy for the recursive problem (once the value function is the right one).

Again the second part is more important. It says that, for the "right" fixed point of the functional equation w the corresponding policy g generates a plan $\hat{x}$ that solves the sequential problem if it satisfies the additional limit condition. Again we can give this condition a loose interpretation as standing in for a transversality condition. Note that for any plan $\{\hat{x}_t\}$ generated from a policy g associated with a value function v that satisfies (32) condition (33) is automatically satisfied. From (32) we have

$$\lim_{t\to\infty} \beta^t v(x_t) = 0$$

for any feasible $\{x_t\} \in \Pi(x_0)$, all x_0. Also from Theorem 32 $v = w$. So for any plan $\{\hat{x}_t\}$ generated from a policy g associated with $v = w$ we have

$$w(\hat{x}_t) = F(\hat{x}_t, \hat{x}_{t+1}) + \beta w(\hat{x}_{t+1})$$

and since $\lim_{t\to\infty} \beta^t v(\hat{x}_t)$ exists and equals to 0 (since v satisfies (32)), we have

$$\limsup_{t\to\infty} \beta^t v(\hat{x}_t) = 0$$

and hence (33) is satisfied. But Theorem 33.2 is obviously not redundant as there may be situations in which Theorem 32.2 does not apply but 33.2 does. Let us look at the following example, a simple modification of the saving problem from before. Now however we impose a borrowing constraint of zero.

$$\begin{aligned} w(x_0) &= \max_{\{x_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t (x_t - \beta x_{t+1}) \\ s.t.\ 0 &\leq x_{t+1} \leq \frac{x_t}{\beta} \\ & \quad x_0 \text{ given} \end{aligned}$$

Writing out the objective function yields

$$\begin{aligned} w_0(x_0) &= (x_0 - \beta x_1) + (x_1 - \beta x_2) + \ldots \\ &= x_0 \end{aligned}$$

Now consider the associated functional equation

$$v(x) = \max_{0 \leq x' \leq \frac{x}{\beta}} \{x - \beta x' + v(x')\}$$

Obviously one solution of this functional equation is $v(x) = x$ and by Theorem 32.1 is rightly follows that w satisfies the functional equation. However, for

v condition (32) fails, as the feasible plan defined by $x_t = \frac{x_0}{\beta^t}$ shows. Hence Theorem 32.2 does not apply and we can't conclude that $v = w$ (although we have verified it directly, there may be other examples for which this is not so straightforward). Still we can apply Theorem 33.2 to conclude that certain plans are optimal plans. Let $\{\hat{x}_t\}$ be defined by $\hat{x}_0 = x_0, \hat{x}_t = 0$ all $t > 0$. Then

$$\limsup_{t\to\infty} \beta^t w(\hat{x}_t) = 0$$

and we can conclude by Theorem 33.2 that this plan is optimal for the sequential problem. There are tons of other plans for which we can apply the same logic to shop that they are optimal, too (which shows that we obviously can't make any claim about uniqueness). To show that condition (33) has some bite consider the plan defined by $\hat{x}_t = \frac{x_0}{\beta^t}$. Obviously this is a feasible plan satisfying

$$w(\hat{x}_t) = F(\hat{x}_t, \hat{x}_{t+1}) + \beta w(\hat{x}_{t+1})$$

but since for all $x_0 > 0$

$$\limsup_{t\to\infty} \beta^t w(\hat{x}_t) = x_0 > 0$$

Theorem 33.2 does not apply and we can't conclude that $\{\hat{x}_t\}$ is optimal (as in fact this plan is not optimal).

So basically we have a prescription what to do once we solved our functional equation: pick the right fixed point (if there are more than one, check the limit condition to find the right one, if possible) and then construct a plan from the policy corresponding to this fixed point. Check the limit condition to make sure that the plan so constructed is indeed optimal for the sequential problem. Done.

Note, however, that so far we don't know anything about the number (unless the CMT applies) and the shape of fixed point to the functional equation. This is not quite surprising given that we have put almost no structure onto our economy. By making further assumptions one obtains sharper characterizations of the fixed point(s) of the functional equation and thus, in the light of the preceding theorems, about the solution of the sequential problem.

5.2 Dynamic Programming with Bounded Returns

Again we look at a functional equation of the form

$$v(x) = \max_{y\in\Gamma(x)} \{F(x,y) + \beta v(y)\}$$

We will now assume that $F : X \times X$ is bounded and $\beta \in (0,1)$. We will make the following two assumptions throughout this section

Assumption 3: X is a convex subset of $\mathbf{R}^L$ and the correspondence $\Gamma : X \rightrightarrows X$ is nonempty, compact-valued and continuous.

Assumption 4: The function $F : A \to \mathbf{R}$ is continuous and bounded, and $\beta \in (0,1)$

We immediately get that assumptions 1. and 2. are satisfied and hence the theorems of the previous section apply. Define the policy correspondence connected to any solution to the functional equation as

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

and the operator T on $C(X)$

$$(Tv)(x) = \max_{y \in \Gamma(x)} \{F(x, y) + \beta v(y)\}$$

Here $C(X)$ is the space of bounded continuous functions on X and we use the sup-metric as metric. Then we have the following

Theorem 42 *Under Assumptions 3. and 4. the operator T maps $C(X)$ into itself. T has a unique fixed point v and for all $v_0 \in C(X)$*

$$d(T^n v_0, v) \leq \beta^n d(v_0, v)$$

The policy correspondence G belonging to v is compact-valued and upper-hemicontinuous

Now we add further assumptions on the structure of the return function F, with the result that we can characterize the unique fixed point of T better.

Assumption 5: For fixed y, $F(., y)$ is strictly increasing in each of its L components.

Assumption 6: Γ is monotone in the sense that $x \leq x'$ implies $\Gamma(x) \subseteq \Gamma(x')$.

The result we get out of these assumptions is strict monotonicity of the value function.

Theorem 43 *Under Assumptions 3. to 6. the unique fixed point v of T is strictly increasing.*

We have a similar result in spirit if we make assumptions about the curvature of the return function and the convexity of the constraint set.

Assumption 7: F is strictly concave, i.e. for all $(x, y), (x', y') \in A$ and $\theta \in (0, 1)$

$$F[\theta(x, y) + (1 - \theta)(x', y')] \geq \theta F(x, y) + (1 - \theta)F(x', y')$$

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

Assumption 8: Γ is convex in the sense that for $\theta \in [0, 1]$ and $x, x' \in X$, the fact $y \in \Gamma(x), y' \in \Gamma(x')$

$$\theta y + (1 - \theta)y' \in \Gamma(\theta x + (1 - \theta)x')$$

Again we find that the properties assumed about F extend to the value function.

Theorem 44 *Under Assumptions 3.-4. and 7.-8. the unique fixed point of v is strictly concave and the optimal policy is a single-valued continuous function, call it g.*

Finally we state a result about the differentiability of the value function, the famous envelope theorem (some people call it the Benveniste-Scheinkman theorem).[27]

Assumption 9: F is continuously differentiable on the interior of A.

Theorem 45 *Under assumptions 3.-4. and 7.-9. if $x_0 \in int(X)$ and $g(x_0) \in int(\Gamma(x_0))$, then the unique fixed point of T, v is continuously differentiable at x_0 with*

$$\frac{\partial v(x_0)}{\partial x^i} = \frac{\partial F(x_0, g(x_0))}{\partial x^i}$$

This theorem gives us an easy way to derive Euler equations from the recursive formulation of the neoclassical growth model. Remember the functional

---

[27] You may have seen the envelope theorem stated a bit differently by Prof. Sargent. He sets up the recursive problem as

$$\begin{aligned} V(x) &= \max_u F(x,u) + \beta V(y) \\ s.t. \quad y &= g(x,u) \end{aligned}$$

Substituiting for y we get

$$V(x) = \max_u F(x,u) + \beta V(g(x,u))$$

The difference between his formulation and mine is that inhis formulation a current period control variable u is chosen, which, jointly with today's state x determines next period's state y. In my formulation we substituted out the control u and chose next period's state y. This yields a different statement of the enveope theorem. Let us briefly derive Prof. Sargent's statement.

The first order condition (always assuming interiority) is

$$\frac{\partial F(x,u)}{\partial u} + \beta V'(g(x,u)) \frac{\partial g(x,u)}{\partial u} = 0$$

Let the solution to the FOC be denoted by $u = h(x)$, i.e. h satisfies for every x

$$\frac{\partial F(x,h(x))}{\partial u} + \beta V'(g(x,h(x)) \frac{\partial g(x,h(x))}{\partial u} = 0$$

Now we differentiate the value function to obtain

$$\begin{aligned} V'(x) &= \frac{\partial F(x,h(x))}{\partial x} + \frac{\partial F(x,h(x))}{\partial u} h'(x) \\ &\quad + \beta V'(g(x,h(x))) \left[\frac{\partial g(x,h(x))}{\partial x} + \frac{\partial g(x,h(x))}{\partial u} h'(x) \right] \\ &= \frac{\partial F(x,h(x))}{\partial x} + \beta V'(g(x,h(x))) \frac{\partial g(x,h(x))}{\partial x} \\ &\quad + h'(x) \left[\frac{\partial F(x,h(x))}{\partial u} + \beta V'(g(x,h(x))) \frac{\partial g(x,h(x))}{\partial u} \right] \end{aligned}$$

Using the first order condtions yields the envelope theorem for Prof. Sargent's setup of the problem.

$$V'(x) = \frac{\partial F(x,h(x))}{\partial x} + \beta V'(g(x,h(x))) \frac{\partial g(x,h(x))}{\partial x}$$

equation

$$v(k) = \max_{0 \leq k' \leq f(k)} U(f(k) - k') + \beta v(k')$$

Taking first order conditions with respect to k' (and ignoring corner solutions) we get

$$U'(f(k) - k') = \beta v'(k')$$

Denote by $k' = g(k)$ the optimal policy. The problem is that we don't know v'. But now we can use Benveniste-Scheinkman to obtain

$$v'(k) = U'(f(k) - g(k))f'(k)$$

Using this in the first order condition we obtain

$$\begin{aligned} U'(f(k) - g(k)) &= \beta v'(k) = \beta U'(f(k') - g(k'))f'(k') \\ &= \beta f'(g(k))U'(f(g(k)) - g(g(k)) \end{aligned}$$

Denoting $k = k_t$, $g(k) = k_{t+1}$ and $g(g(k)) = k_{t+2}$ we obtain our usual Euler equation

$$U'(f(k_t) - k_{t+1}) = \beta f(k_{t+1})U'(f(k_{t+1}) - k_{t+2})$$

6 Models with Uncertainty

In this section we will introduce a basic model with uncertainty, in order to establish some notation and extend our discussion of efficient economy to this important case. Then, as a fist application, we will look at the stochastic neoclassical growth model, which forms the basis for a particular theory of business cycles, the so called "Real Business Cycle" (RBC) theory. In this section we will be a bit loose with our treatment of uncertainty, in that we will not explicitly discuss probability spaces that form the formal basis of our representation of uncertainty.

6.1 Basic Representation of Uncertainty

The basic novelty of models with uncertainty is the formal representation of this uncertainty and the ensuing description of the information structure that agents have. We start with the notion of an event $s_t \in S$. The set $S = \{\eta^1, \eta^2, \ldots \eta^N\}$ of possible *events* that can happen in period t is assumed to be finite and the same for all periods t. For example S may consist of all weather conditions than can happen in the economy, with $s_t = 1$ indicating sunshine in period t, $s_t = 2$ indicating cloudy skies, $s_t = 3$ indicating rain and so forth.[28] As another example, consider the economy from Section 2, but now with random endowments. In each period one of the two agents has endowment 0 and the other has endowment 2, but who has what is random, with $s_t = 1$ indicating that agent 1 has high endowment and $s_t = 2$ indicating that agent 2 has high endowment at period t. The set of possible events is given by $S = \{1, 2\}$

An event history $s^t = (s_0, s_1, \ldots s_t)$ is a vector of length $t+1$ summarizing the realizations of all events up to period t. Formally, with $S^t = S \times S \times \ldots \times S$ denoting the $t+1$-fold product of S, event history $s^t \in S^t$ lies in the set of all possible event histories of length t.

By $\pi(s^t)$ let denote the probability of a particular event history. We assume that $\pi(s^t) > 0$ for all $s^t \in S^t$, for all t. For our example economy, if $s^2 = (1, 1, 2)$ then $\pi(s^2)$ is the probability that agent 1 has high endowment in period $t = 0$ *and* $t = 1$ *and* agent 2 has high endowment in period 2. Figure 5 summarizes the concepts introduced so far, for the case in which $S = \{1, 2\}$ is the set of possible events that can happen in every period. Note that the sets S^t of possible events of length t become fairly big very rapidly, which poses computational problems when dealing with models with uncertainty.

All commodities of our economy, instead of being indexed by time t as before, now also have to be indexed by event histories s^t. In particular, an allocation for the example economy of Section 2, but now with uncertainty, is given by

$$(c^1, c^2) = \{c_t^1(s^t), c_t^2(s^t)\}_{t=0, s^t \in S^t}^{\infty}$$

---

[28] Technically speaking s_t is a random variable with respect to some underlying probability space ($, \mathcal{A}, P$), where is some set of basis events with generic element ω, $\mathcal{A}$ is a sigma algebra on and P is a probability measure.

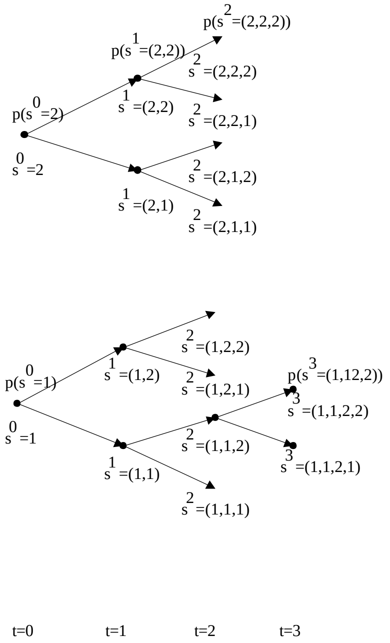

Figure 5:

with the interpretation that $c_t^i(s^t)$ is consumption of agent i in period t if event history s^t has occurred. Note that consumption in period t of agents are allowed to (and in general will) vary with the history of events that have occurred in the past .

Now we are ready to specify to remaining elements of the economy. With respect to endowments, these also take the general form

$$(e^1, e^2) = \{e_t^1(s^t), e_t^2(s^t)\}_{t=0, s^t \in S^t}^{\infty}$$

and for the particular example

$$\begin{aligned} e_t^1(s^t) &= \begin{cases} 2 & \text{if } s_t = 1 \\ 0 & \text{if } s_t = 2 \end{cases} \\ e_t^2(s^t) &= \begin{cases} 0 & \text{if } s_t = 1 \\ 2 & \text{if } s_t = 2 \end{cases} \end{aligned}$$

i.e. for the particular example endowments in period t only depend on the realization of the event s_t, not on the entire history. Nothing, however, would prevent us from specifying more general endowment patterns.

Now we specify preferences. We assume that households maximize *expected* lifetime utility where expectations E_0 is the expectation operator at period 0, prior to any realization of uncertainty (in particular the uncertainty with respect to s_0). Given our notation just established, assuming that preferences admit a von-Neumann Morgenstern utility function representation[29] we represent households' preferences by

$$u(c^i) = \sum_{t=0}^{\infty} \sum_{s^t \in S^t} \beta^t \pi(s^t) U(c_t^i(s^t))$$

This completes our description of the simple example economy.

6.2 Definitions of Equilibrium

Again there are two possible market structures that we can work with. The Arrow-Debreu market structure turns out to be easier than the sequential markets market structure, so we will start with it. Again there is an equivalence theorem between these two economies, once we allow the asset market structure for the sequential markets market structure to be rich enough.

6.2.1 Arrow-Debreu Market Structure

As usual with Arrow-Debreu, trade takes place at period 0, *before* any uncertainty has been realized (in particular, before s_0 has been realized). As with allocations, Arrow-Debreu prices have to be indexed by event histories in addition to time, so let $p_t(s^t)$ denote the price of one unit of consumption, quoted at period 0, delivered at period t if (and only if) event history s^t has been realized. Given this notation, the definition of an AD-equilibrium is identical to the case without uncertainty:

Definition 46 *A (competitive) Arrow-Debreu equilibrium are prices* $\{\hat{p}_t(s^t)\}_{t=0,s^t \in S^t}^{\infty}$ *and allocations* $(\{\hat{c}_t^i(s^t)\}_{t=0,s^t \in S^t}^{\infty})_{i=1,2}$ *such that*

1. *Given* $\{\hat{p}_t(s^t)\}_{t=0,s^t \in S^t}^{\infty}$*, for* $i = 1, 2$*,* $\{\hat{c}_t^i(s^t)\}_{t=0,s^t \in S^t}^{\infty}$ *solves*

$$\max_{\{c_t^i(s^t)\}_{t=0,s^t \in S^t}^{\infty}} \sum_{t=0}^{\infty} \sum_{s^t \in S^t} \beta^t \pi(s^t) U(c_t^i(s^t)) \tag{34}$$

$$\text{s.t.}$$

$$\sum_{t=0}^{\infty} \sum_{s^t \in S^t} p_t(s^t) c_t^i(s^t) \leq \sum_{t=0}^{\infty} \sum_{s^t \in S^t} p_t(s^t) e_t^i(s^t) \tag{35}$$

$$c_t^i(s^t) \geq 0 \text{ for all } t \tag{36}$$

---

[29] Felix Lubler will discuss in great length what is required for this.

2.

$$\hat{c}_t^1(s^t) + \hat{c}_t^2(s^t) = e_t^1(s^t) + e_t^2(s^t) \text{ for all } t, \text{ all } s^t \in S^t \tag{37}$$

Note that there is again only one budget constraint, and that market clearing has to hold date, by date, event history by event history. Also note that, when computing equilibria, one can normalize the price of only one commodity to 1, and consumption at the same date, but for different event histories are different commodities. That means that if we normalize $p_0(s_0 = 1) = 1$ we can't also normalize $p_0(s_0 = 2) = 1$. Finally, there are no probabilities in the budget constraint. Equilibrium prices will reflect the probabilities of different event histories, but there is no room for these probabilities in the Arrow-Debreu budget constraint directly.

The definition of Pareto efficiency is identical to that of the certainty case; the first welfare theorem goes through without any changes (in particular, the proof is identical, apart from changes in notation). We state both for completeness

Definition 47 *An allocation* $\{(c_t^1(s^t), c_t^2(s^t))\}_{t=0,s^t \in S^t}^{\infty}$ *is feasible if*

1.

$$c_t^i(s^t) \geq 0 \text{ for all } t, \text{ all } s^t \in S^t, \text{ for } i = 1, 2$$

2.

$$c_t^1(s^t) + c_t^2(s^t) = e_t^1(s^t) + e_t^2(s^t) \text{ for all } t, \text{ all } s^t \in S^t$$

Definition 48 *An allocation* $\{(c_t^1(s^t), c_t^2(s^t))\}_{t=0,s^t \in S^t}^{\infty}$ *is Pareto efficient if it is feasible and if there is no other feasible allocation* $\{(\tilde{c}_t^1(s^t), \tilde{c}_t^2(s^t))\}_{t=0,s^t \in S^t}^{\infty}$ *such that*

$$\begin{aligned} u(\tilde{c}^i) &\geq u(c^i) \text{ for both } i = 1, 2 \\ u(\tilde{c}^i) &> u(c^i) \text{ for at least one } i = 1, 2 \end{aligned}$$

Proposition 49 *Let* $(\{\hat{c}_t^i(s^t)\}_{t=0,s^t \in S^t}^{\infty})_{i=1,2}$ *be a competitive equilibrium allocation. Then* $(\{\hat{c}_t^i(s^t)\}_{t=0,s^t \in S^t}^{\infty})_{i=1,2}$ *is Pareto efficient.*

6.2.2 Sequential Markets Market Structure

Now let trade take place sequentially in each period (more precisely, in each period, event-history pair). With certainty, we allowed trade in consumption and in one-period IOU's. For the equivalence between Arrow-Debreu and sequential markets with uncertainty, this is not enough. We introduce one period contingent IOU's, financial contracts bought in period t, that pay out one unit of the consumption good in $t+1$ only for a particular realization of $s_{t+1} = s^j$

tomorrow.[30] So let $q_t(s^t, s_{t+1} = \eta^j)$ denote the price at period t of a contract that pays out one unit of consumption in period $t+1$ if (and only if) tomorrow's event is $s_{t+1} = \eta^j$. These contracts are often called Arrow securities, contingent claims or one-period insurance contracts. Let $a^i_{t+1}(s^t, s_{t+1})$ denote the quantities of these Arrow securities bought (or sold) at period t by agent i.

The period t, event history s^t budget constraint of agent i is given by

$$c^i_t(s^t) + \sum_{s_{t+1}\in S} q_t(s^t, s_{t+1}) a^i_{t+1}(s^t, s_{t+1}) \leq e^i_t(s^t) + a^i_t(s^t)$$

Note that agents purchase Arrow securities $\{a^i_{t+1}(s^t, s_{t+1})\}_{s_{t+1}\in S}$ for all contingencies $s_{t+1} \in S$ that can happen tomorrow, but that, once s_{t+1} is realized, only the $a^i_{t+1}(s^{t+1})$ corresponding to the particular realization of s_{t+1} becomes his asset position with which he starts the current period. We assume that $a^i_0(s_0) = 0$ for all $s_0 \in S$.

We then have the following

Definition 50 *A SM equilibrium is allocations* $\{\left(\hat{c}^i_t(s^t), \left\{\hat{a}^i_{t+1}(s^t, s_{t+1}\right\}_{s_{t+1}\in S}\right)_{i=1,2}\}^{\infty}_{t=0,s^t\in S^t}$, *and prices for Arrow securities* $\{\hat{q}_t(s^t, s_{t+1})\}^{\infty}_{t=0,s^t\in S^t,s_{t+1}\in S}$ *such that*

1. *For* $i = 1, 2$, *given* $\{\hat{q}_t(s^t, s_{t+1})\}^{\infty}_{t=0,s^t\in S^t,s_{t+1}\in S}$, *for all* i, $\{\hat{c}^i_t(s^t), \left\{\hat{a}^i_{t+1}(s^t, s_{t+1}\right\}_{s_{t+1}\in S}\}^{\infty}_{t=0,s^t\in S^t}$ *solves*

$$\begin{aligned}
& \max_{\{c^i_t(s^t), \left\{a^i_{t+1}(s^t, s_{t+1}\right\}_{s_{t+1}\in S}\}^{\infty}_{t=0,s^t\in S^t}} \quad u(c^i) \\
& s.t \\
c^i_t(s^t) + \sum_{s_{t+1}\in S} \hat{q}_t(s^t, s_{t+1}) a^i_{t+1}(s^t, s_{t+1}) &\leq e^i_t(s^t) + a^i_t(s^t) \\
c^i_t(s^t) &\geq 0 \textit{ for all } t, s^t \in S^t \\
a^i_{t+1}(s^t, s_{t+1}) &\geq -\bar{A}^i \textit{ for all } t, s^t \in S^t
\end{aligned}$$

2. *For all* $t \geq 0$

$$\begin{aligned}
\sum_{i=1}^{2} \hat{c}^i_t(s^t) &= \sum_{i=1}^{2} e^i_t(s^t) \textit{ for all } t, s^t \in S^t \\
\sum_{i=1}^{2} \hat{a}^i_{t+1}(s^t, s_{t+1}) &= 0 \textit{ for all } t, s^t \in S^t \textit{ and all } s_{t+1} \in S
\end{aligned}$$

---

[30] A full set of one-period Arrow securities is sufficient to make markets "sequentially complete", in the sense that any (nonnegative) consumption allocation is attainable with an appropriate sequence of Arrow security holdings $\{a_{t+1}(s^t, s_{t+1})\}$ satisfying all sequential markets budget constraints.

Note that we have a market clearing condition in the asset market for each Arrow security being traded for period $t+1$. Define

$$q_t(s^t) = \sum_{s_{t+1} \in S} q_t(s^t, s_{t+1})$$

The price $q_t(s^t)$ can be interpreted as the price, in period t, event history s^t, for buying one unit of consumption delivered for sure in period $t+1$ (we buy one unit of consumption for each contingency tomorrow). The risk free interest rate (the counterpart to the interest rate for economies without uncertainty) between periods t and $t+1$ is then given by

$$\frac{1}{1 + r_{t+1}(s^t)} = q_t(s^t)$$

6.2.3 Equivalence between Market Structures

[To Be Completed]

6.3 Markov Processes

So far we haven't specified the exact nature of uncertainty. In particular, in no sense have we assumed that the random variables s_t and s_τ, $\tau > t$ are independent or dependent in a simple way. Our theory is completely general along this dimension; to make it implementable (analytically or numerically), however, one has to assume a particular structure of the uncertainty.

In particular, it simplifies matters a lot if one assumes that the s_t's follow a discrete time (time is discrete), discrete state (the number of values s_t can take is finite) time homogeneous Markov chain. Let by

$$\pi(\eta^j | \eta^i) = \text{prob}(s_{t+1} = \eta^j | s_t = \eta^i)$$

denote the conditional probability that the state in $t+1$ equals $\eta^j \in S$ if the state in period t equals $s_t = \eta^i \in S$. Time homogeneity means that π is not indexed by time. Given that $s_{t+1} \in S$ and $s_t \in S$ and S is a finite set, $\pi(.|.)$ is an $N \times N$-matrix of the form

$$\pi = \begin{pmatrix} \pi_{11} & \pi_{12} & \cdots & \vdots & \cdots & \pi_{1N} \\ \pi_{21} & & & \vdots & & \vdots \\ \vdots & & & \vdots & & \vdots \\ \pi_{i1} & \cdots & \cdots & \pi_{ij} & \cdots & \pi_{iN} \\ \vdots & & & \vdots & & \vdots \\ \pi_{N1} & \cdots & \cdots & \vdots & \cdots & \pi_{NN} \end{pmatrix}$$

with generic element $\pi_{ij} = \pi(\eta^j|\eta^i)$ =prob$(s_{t+1} = \eta^j | s_t = \eta^i)$. Hence the i-th row gives the probabilities of going from state i today to all the possible states tomorrow, and the j-th column gives the probability of landing in state j tomorrow conditional of being in an arbitrary state i today. Since $\pi_{ij} \geq 0$ and $\sum_j \pi_{ij} = 1$ for all i (for all states today, one has to go somewhere for tomorrow), the matrix π is a so-called stochastic matrix.

Suppose the probability distribution over states today is given by the N-dimensional column vector $P_t = (p_t^1, \ldots, p_t^N)^T$ and uncertainty is described by a Markov chain of the from above. Note that $\sum_i p_t^i = 1$. Then the probability of being in state j tomorrow is given by

$$p_{t+1}^j = \sum_i \pi_{ij} p_t^i$$

i.e. by the sum of the conditional probabilities of going to state j from state i, weighted by the probabilities of starting out in state i. More compactly we can write

$$P_{t+1} = \pi^T P_t$$

A stationary distribution Π of the Markov chain π satisfies

$$\Pi = \pi^T \Pi$$

i.e. if you start today with a distribution over states Π then tomorrow you end up with the *same* distribution over states Π. From the theory of stochastic matrices we know that every π has at least one such stationary distribution. It is the eigenvector (normalized to length 1) associated with the eigenvalue $\lambda = 1$ of π^T. Note that every stochastic matrix has (at least) one eigenvector equal to 1. If there is only one such eigenvalue, then there is a unique stationary distribution, if there are multiple eigenvalues of length 1, then there a multiple stationary distributions (in fact a continuum of them).

Note that the Markov assumption restricts the conditional probability distribution of s_{t+1} to depend only on the realization of s_t, but not on realizations of s_{t-1}, s_{t-2} and so forth. This obviously is a severe restriction on the possible randomness that we allow, but it also means that the nature of uncertainty for period $t+1$ is completely described by the realization of s_t, which is crucial when formulating these economies recursively. We have to start the Markov process out at period 0, so let by $\Pi(s_0)$ denote the probability that the state in period 0 is s_0. Given our Markov assumption the probability of a particular event history can be written as

$$\pi(s^{t+1}) = \pi(s_{t+1}|s_t) * \pi(s_t|s_{t-1}) \ldots * \pi(s_1|s_0) * \Pi(s_0)$$

6.4 Stochastic Neoclassical Growth Model

In this section we will briefly consider a stochastic extension to the deterministic neoclassical growth model. You will have fun with this model in the first problem set. The stochastic neoclassical growth model is the workhorse for half of modern business cycle theory; everybody doing real business cycle theory uses it. I therefore think that it is useful to expose you to this model, even though you may decide not to do RBC-theory in your own research.

The economy is populated by a large number of identical households. For convenience we normalize the number of households to 1. In each period three goods are traded, labor services n_t, capital services k_t and the final output good y_t, which can be used for consumption c_t or investment i_t.

1. Technology:

   $$y_t = e^{z_t} F(k_t, n_t)$$

   where z_t is a technology shock. F is assumed to have the usual properties, i.e. has constant returns to scale, positive but declining marginal products and satisfies the INADA conditions. We assume that the technology shock has unconditional mean 0 and follows a N-state Markov chain. Let $Z = \{z_1, z_2, \dots z_N\}$ be the state space of the Markov chain, i.e. the set of values that z_t can take on. Let $\pi = (\pi_{ij})$ denote the Markov transition matrix and Π the stationary distribution of the chain (ignore the fact that in some of our applications Π will not be unique). Let $\pi(z'|z) = prob(z_{t+1} = z'|z_t = z)$. In most of the applications we will take $N = 2$. The evolution of the capital stock is given by

   $$k_{t+1} = (1 - \delta) k_t + i_t$$

   and the composition of output is given by

   $$y_t = c_t + i_t$$

   Note that the set Z takes the role of S in our general formulation of uncertainty, z^t corresponds to s^t and so forth.

2. Preferences:

   $$E_0 \sum_{t=0}^{\infty} \beta^t u(c_t) \text{ with } \beta \in (0, 1)$$

   The period utility function is assumed to have the usual properties.

3. Endowment: each household has an initial endowment of capital, k_0 and one unit of time in each period. Endowments are not stochastic.

4. Information: The variable z_t, the only source of uncertainty in this model, is publicly observable. We assume that in period 0 z_0 has not been realized, but is drawn from the stationary distribution Π. All agents are perfectly informed that the technology shock follows the Markov chain π with initial distribution Π.

A lot of the things that we did for the case without uncertainty go through almost unchanged for the stochastic model. The only key difference is that now commodities have to be indexed not only by time, but also by histories of productivity shocks, since goods delivered at different nodes of the event tree are different commodities, even though they have the same physical characteristics. For a lucid discussion of this point see Chapter 7 of Debreu's (1959) "Theory of Value".

For the recursive formulation of the social planners problem, note that the current state of the economy now not only includes the capital stock k that the planner brings into the current period, but also the current state of the technology z. This is due to the fact that current production depends on the current technology shock, but also due to the fact that the probability distribution of future shocks $\pi(z'|z)$ depends on the current shock, due to the Markov structure of the stochastic shocks. Also note that even if the social planner chooses capital stock k' for tomorrow today, lifetime utility from tomorrow onwards is uncertain, due to the uncertainty of z'.