CHAPTER

19

General Equilibrium Under Uncertainty

19.A Introduction

In this chapter, we apply the general equilibrium framework developed in Chapters 15 to 18 to economic situations involving the exchange and allocation of resources under conditions of uncertainty. In a sense, this chapter offers the equilibrium counterpart of the decision theory presented in Chapter 6 (and which we recommend you review at this point).

We begin, in Section 19.B, by formalizing uncertainty by means of *states of the world* and then introducing the key idea of a *contingent commodity*: a commodity whose delivery is conditional on the realized state of the world. In Section 19.C we use these tools to define the concept of an *Arrow–Debreu equilibrium*. This is simply a Walrasian equilibrium in which contingent commodities are traded. It follows from the general theory of Chapter 16 that an Arrow–Debreu equilibrium results in a Pareto optimal allocation of risk.

In Section 19.D, we provide an important reinterpretation of the concept of Arrow–Debreu equilibrium. We show that, under the assumptions of *self-fulfilling*, or *rational*, *expectations*, Arrow–Debreu equilibria can be implemented by combining trade in a certain restricted set of contingent commodities with *spot trade* that occurs *after* the resolution of uncertainty. This results in a significant reduction in the number of ex ante (i.e., before uncertainty) markets that must operate.

In Section 19.E, we generalize our analysis. Instead of trading contingent commodities prior to the resolution of uncertainty, agents now trade *assets*; and instead of an Arrow–Debreu equilibrium we have the notion of a *Radner equilibrium*. We also discuss here the important notion of *arbitrage* among assets. The material of this section lies at the foundations of a very rich body of finance theory [good introductions are Duffie (1992) and Huang and Litzenberger (1988)].

In Section 19.F, we briefly illustrate some of the welfare difficulties raised by the possibility of *incomplete markets*, that is, by the possibility of there being too few asset markets to guarantee a fully Pareto optimal allocation of risk.

Section 19.G is devoted to the issue of the objectives of the firm under conditions of uncertainty. In particular, it gives sufficient conditions for shareholders to agree unanimously on the objective of *market value maximization*.

Section 19.H takes a close look at the informational requirements of the theory developed in this chapter. We see that the theory applies well to situations of *symmetric* information across consumers (reviewed in Section 19.H); but its applicability is more problematic in situations of *asymmetric* information. This provides a further argument for the techniques developed in Chapters 13 and 14 for the study of asymmetric information problems.

For additional material and references on the topic of this chapter, see the textbooks of Huang and Litzenberger (1988) and Duffie (1992) already mentioned, or, at a more advanced level, Radner (1982) and Magill and Shafer (1991).

19.B A Market Economy with Contingent Commodities: Description

As in our previous chapters, we contemplate an environment with L physical commodities, I consumers, and J firms. The new element is that technologies, endowments, and preferences are now *uncertain*.

Throughout this chapter, we represent uncertainty by assuming that technologies, endowments, and preferences depend on the *state of the world*. The concept of state of the world was already introduced in Section 6.E. A state of the world is to be understood as a complete description of a possible outcome of uncertainty, the description being sufficiently fine for any two distinct states of the world to be mutually exclusive. We assume that an exhaustive set S of states of the world is given to us. For simplicity we take S to be a finite set with (abusing notation slightly) S elements. A typical element is denoted $s = 1, \ldots, S$.

We state in Definition 19.B.1 the key concepts of a (*state-*)*contingent commodity* and a (*state-*)*contingent commodity vector*. Using these concepts we shall then be able to express the dependence of technologies, endowments, and preferences on the realized states of the world.

Definition 19.B.1: For every physical commodity $\ell = 1, \ldots, L$ and state $s = 1, \ldots, S$, a unit of (*state-*)*contingent commodity* ℓs is a title to receive a unit of the physical good ℓ if and only if s occurs. Accordingly, a (*state-*)*contingent commodity vector* is specified by

$$x = (x_{11}, \ldots, x_{L1}, \ldots, x_{1S}, \ldots, x_{LS}) \in \mathbb{R}^{LS},$$

and is understood as an entitlement to receive the commodity vector $(x_{1s}, \ldots, x_{Ls})$ if state s occurs.[1]

We can also view a contingent commodity vector as a collection of L *random variables*, the ℓth random variable being $(x_{\ell 1}, \ldots, x_{\ell S})$.

With the help of the concept of contingent commodity vectors, we can now describe how the characteristics of economic agents depend on the state of the world. To begin, we let the endowments of consumer $i = 1, \ldots, I$ be a contingent commodity vector:

$$\omega_i = (\omega_{11i}, \ldots, \omega_{L1i}, \ldots, \omega_{1Si}, \ldots, \omega_{LSi}) \in \mathbb{R}^{LS}.$$

1. As usual, a negative entry is understood as an obligation to deliver.

The meaning of this is that if state s occurs then consumer i has endowment vector $(\omega_{1si}, \ldots, \omega_{Lsi}) \in \mathbb{R}^L$.

The preferences of consumer i may also depend on the state of the world (e.g., the consumer's enjoyment of wine may well depend on the state of his health). We represent this dependence formally by defining the consumer's preferences over contingent commodity vectors. That is, we let the preferences of consumer i be specified by a rational preference relation $\succsim_i$ defined on a consumption set $X_i \subset \mathbb{R}^{LS}$.

Example 19.B.1: As in Section 6.E, the consumer evaluates contingent commodity vectors by first assigning to state s a probability π_{si} (which could have an objective or a subjective character), then evaluating the physical commodity vectors at state s according to a Bernoulli state-dependent utility function $u_{si}(x_{1si}, \ldots, x_{Lsi})$, and finally computing the expected utility.[2] That is, the preferences of consumer i over two contingent commodity vectors $x_i, x_i' \in X_i \subset \mathbb{R}^{LS}$ satisfy

$$x_i \succsim_i x_i' \quad \text{if and only if} \quad \sum_s \pi_{si} u_{si}(x_{1si}, \ldots, x_{Lsi}) \geq \sum_s \pi_{si} u_{si}(x'_{1si}, \ldots, x'_{Lsi}).$$

■

It should be emphasized that the preferences $\succsim_i$ are in the nature of ex ante preferences: the random variables describing possible consumptions are evaluated before the resolution of uncertainty.

Similarly, the technological possibilities of firm j are represented by a production set $Y_j \subset \mathbb{R}^{LS}$. The interpretation is that a *(state-)contingent production plan* $y_j \in \mathbb{R}^{LS}$ is a member of Y_j if for every s the input–output vector $(y_{1sj}, \ldots, y_{Lsj})$ of physical commodities is feasible for firm j when state s occurs.

Example 19.B.2: Suppose there are two states, s_1 and s_2, representing good and bad weather. There are two physical commodities: seeds ($\ell = 1$) and crops ($\ell = 2$). In this case, the elements of Y_j are four-dimensional vectors. Assume that seeds must be planted before the resolution of the uncertainty about the weather and that a unit of seeds produces a unit of crops if and only if the weather is good. Then

$$y_j = (y_{11j}, y_{21j}, y_{12j}, y_{22j}) = (-1, 1, -1, 0)$$

is a feasible plan. Note that since the weather is unknown when the seeds are planted, the plan $(-1, 1, 0, 0)$ is not feasible: the seeds, if planted, are planted in both states. Thus, in this manner we can imbed into the structure of Y_j constraints on production related to the timing of the resolution of uncertainty.[3] ■

To complete the description of an economy in a manner parallel to Chapters 16 and 17 it only remains to specify ownership shares for every consumer i and firm j. In principle, these shares could also be state-contingent. It will be simpler, however, to let $\theta_{ji} \geq 0$ be the share of firm j owned by consumer i whatever the state. Of course $\sum_j \theta_{ji} = 1$ for every i.

2. The discussion in Section 6.E was for $L = 1$. It extends straightforwardly to the current case of $L \geq 1$.

3. A similar point could be made on the consumption side. If, for a particular commodity ℓ, any vector $x_i \in X_i$ is such that all entries $x_{\ell si}$, $s = 1, \ldots, S$, are equal, then we can interpret this as asserting that the consumption of ℓ takes place before the resolution of uncertainty.

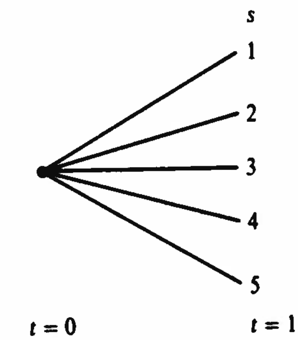

Figure 19.B.1 **Two periods.** Perfect information at $t = 1$.

Information and the Resolution of Uncertainty

In the setting just described, time plays no explicit formal role. In reality, however, states of the world unfold over time. Figure 19.B.1 captures the simplest example. In the figure, we have a period 0 in which there is no information whatsoever on the true state of the world and a period 1 in which this information has been completely revealed.

We have already seen (Example 19.B.2) how, by conveniently defining consumption and production sets, we can accomodate within our setup the temporal structure of Figure 19.B.1: a commodity that has as part of its physical description its availability at $t = 0$ should never appear in differing amounts across states.

The same methodology can be used to incorporate into the formalism a much more general temporal structure. Suppose we have $T + 1$ dates $t = 0, 1, \ldots, T$ and, as before, S states, but assume that the states emerge gradually through a *tree*, as in Figure 19.B.2. These trees are

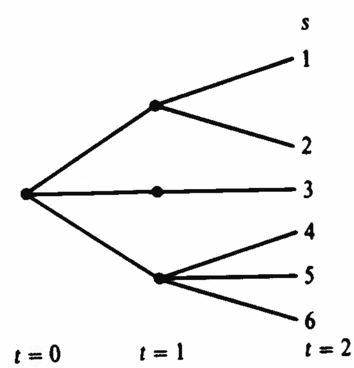

Figure 19.B.2 **An information** tree: **gradual** release of **information**.

similar to those described in Chapter 7. Here final nodes stand for the possible states realized by time $t = T$, that is, for complete histories of the uncertain environment. When the path through the tree coincides for two states, s and s', up to time t, this means that in all periods up to and including period t, s and s' cannot be distinguished.

Subsets of S are called *events*. A collection of events $\mathcal{S}$ is an *information structure* if it is a partition, that is, if for every state s there is $E \in \mathcal{S}$ with $s \in E$ and for any two $E, E' \in \mathcal{S}$, $E \neq E'$, we have $E \cap E' = \emptyset$. The interpretation is that if s and s' belong to the same event in $\mathcal{S}$ then s and s' cannot be distinguished in the information structure $\mathcal{S}$.

To capture formally a situation with sequential revelation of information we look at a *family* of information structures: $(\mathcal{S}_0, \ldots, \mathcal{S}_t, \ldots, \mathcal{S}_T)$. The process of information revelation makes the $\mathcal{S}_t$ increasingly fine: once one has information sufficient to distinguish between two states, the information is not forgotten.

Example 19.B.3: Consider the tree in Figure 19.B.2. We have

$$\mathcal{S}_0 = (\{1, 2, 3, 4, 5, 6\}),$$
$$\mathcal{S}_1 = (\{1, 2\}, \{3\}, \{4, 5, 6\}),$$
$$\mathcal{S}_2 = (\{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \{6\}). \blacksquare$$

The partitions could in principle be different across individuals. However, except in the last section of this chapter (Section 19.H), we shall assume that the information structure is the same for all consumers.

A pair (t, E) where t is a date and $E \in \mathcal{S}_t$ is called a *date-event*. Date-events are associated with the nodes of the tree. Each date-event except the first has a *unique predecessor*, and each date-event not at the end of the tree has one or more *successors*.

With this temporal modeling it is now necessary to be explicit about the time at which a physical commodity is available. Suppose there is a number H of basic physical commodities (bread, leisure, etc.). We will use the double index ht to indicate the time at which a commodity h is produced, appears as endowment, or is available for consumption. Then x_{hts} stands for an amount of the physical commodity h available at time t along the path of state s.

Fortunately, this multiperiod model can be formally reduced to the timeless structure introduced above. To see this, we define a new set of $L = H(T + 1)$ physical commodities, each of them being one of these double-indexed (i.e., ht) commodities. We then say that a vector $z \in \mathbb{R}^{LS}$ is *measurable* with respect to the family of information partitions $(\mathcal{S}_0, \ldots, \mathcal{S}_T)$ if, for every hts and hts', we have that $z_{hts} = z_{hts'}$ whenever s,s' belong to the same element of the partition $\mathcal{S}_t$. That is, whenever s and s' cannot be distinguished at time t, the amounts assigned to the two states cannot be different. Finally, we impose on endowments $\omega_i \in \mathbb{R}^{LS}$, consumption sets $X_i \subset \mathbb{R}^{LS}$ and production sets $Y_j \subset \mathbb{R}^{LS}$ the restriction that all their elements be measurable with respect to the family of information partitions. With this, we have reduced the multiperiod structure to our original formulation.

Arrow–Debreu Equilibrium

We have seen in Section 19.B how an economy where uncertainty matters can be described by means of a set of states of the world S, a consumption set $X_i \subset \mathbb{R}^{LS}$, an endowment vector $\omega_i \in \mathbb{R}^{LS}$, and a preference relation $\succsim_i$ on X_i for every consumer i, together with a production set $Y_j \subset \mathbb{R}^{LS}$ and profit shares $(\theta_{j1}, \ldots, \theta_{jI})$ for every firm j.

We now go a step further and make a strong assumption. Namely, we postulate the existence of a market for every contingent commodity ℓs. These markets open *before* the resolution of uncertainty, at date 0 we could say. The price of the commodity is denoted $p_{\ell s}$. What is being purchased (or sold) in the market for the contingent commodity ℓs is commitments to receive (or to deliver) amounts of the physical good ℓ if, and when, state of the world s occurs. Observe that although deliveries are contingent, the payments are not. Notice also that for this market to be well defined it is indispensable that all economic agents be able to recognize the occurrence of s. That is, information should be *symmetric* across economic agents. This informational issue will be discussed further in Section 19.H.

Formally, the market economy just described is nothing but a particular case of the economies we have studied in previous chapters. We can, therefore, apply to our market economy the concept of Walrasian equilibrium and, with it, all the theory developed so far. When dealing with contingent commodities it is customary to call the Walrasian equilibrium an *Arrow–Debreu equilibrium.*[4]

Definition 19.C.1: An allocation

$$(x_1^*, \ldots, x_I^*, y_1^*, \ldots, y_J^*) \in X_1 \times \cdots \times X_I \times Y_1 \times \cdots \times Y_J \subset \mathbb{R}^{LS(I+J)}$$

and a system of prices for the contingent commodities $p = (p_{11}, \ldots, p_{LS}) \in \mathbb{R}^{LS}$ constitute an *Arrow–Debreu equilibrium* if:

(i) For every j, y_j^* satisfies $p \cdot y_j^* \geq p \cdot y_j$ for all $y_j \in Y_j$.

(ii) For every i, x_i^* is maximal for $\succsim_i$ in the budget set

$$\{x_i \in X_i \colon p \cdot x_i \leq p \cdot \omega_i + \textstyle\sum_j \theta_{ij} p \cdot y_j^*\}.$$

(iii) $\sum_i x_i^* = \sum_j y_j^* + \sum_i \omega_i$.

The welfare and positive theorems of Chapters 16 and 17 apply without modification to the Arrow–Debreu equilibrium. Recall from Chapter 6, especially Sections 6.C and 6.E, that, in the present context, the convexity assumption takes on an interpretation in terms of risk aversion. For example, in the expected utility setting of Example 19.B.1, the preference relation $\succsim_i$ is convex if the Bernoulli utilities $u_{si}(x_{si})$ are concave (see Exercise 19.C.1).

The Pareto optimality implication of Arrow–Debreu equilibrium says, effectively, that the possibility of trading in contingent commodities leads, at equilibrium, to an efficient allocation of risk.

It is important to realize that at any production plan the profit of a firm, $p \cdot y_j$, is a nonrandom amount of dollars. Productions and deliveries of goods do, of course, depend on the state of the world, but the firm is active in all the contingent markets and manages, so to speak, to insure completely. This has important implications for the justification of profit maximization as the objective of the firm. We will discuss this point further in Section 19.G.

Example 19.C.1: Consider an exchange economy with $I = 2$, $L = 1$, and $S = 2$. This lends itself to an Edgeworth box representation because there are precisely two contingent commodities. In Figures 19.C.1(a) and 19.C.1(b) we have $\omega_1 = (1, 0)$, $\omega_2 = (0, 1)$, and utility functions of the form $\pi_{1i} u_i(x_{1i}) + \pi_{2i} u_i(x_{2i})$, where (π_{1i}, π_{2i}) are the subjective probabilities of consumer i for the two states. Since $\omega_1 + \omega_2 = (1, 1)$ there is no aggregate uncertainty, and the state of the world determines only which consumer receives the endowment of the consumption good. Recall from Section 6.E (especially the discussion preceding Example 6.E.1) that for this model [in which the $u_i(\cdot)$ do not depend on s], the marginal rate of substitution of consumer i at any point where the consumption is the same in the two states equals the probability ratio π_{1i}/π_{2i}.

In Figure 19.C.1(a) the subjective probabilities are the same for the two consumers (i.e., $\pi_{11} = \pi_{12}$) and therefore the Pareto set coincides with the diagonal of the box (the box is a square and so the diagonal coincides with the 45-degree line, where the marginal rates of substitution for the two consumers are equal: $\pi_{11}/\pi_{21} = \pi_{12}/\pi_{22}$). Hence, at equilibrium, the two consumers insure completely; that is, consumer i's equilibrium consumption does not vary across the two states. In Figure 19.C.1(b)

4. See Chapter 7 of Debreu (1959) for a succinct development of these ideas.

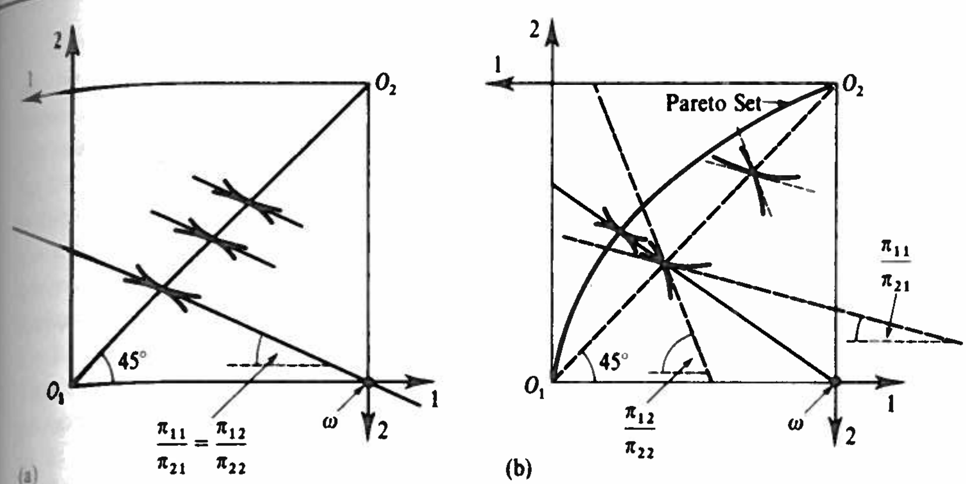

Figure 19.C.1
(a) No aggregate risk: same probability assessments.
(b) No aggregate risk: different probability assessments.

the consumer's subjective probabilities are different. In particular, $\pi_{11} < \pi_{12}$ (i.e., the second consumer gives more probability to state 1). In this case, each consumer's equilibrium consumption is higher in the state he thinks comparatively more likely (relative to the beliefs of the other consumer). ∎

Example 19.C.2: The basic framework is as in Example 19.C.1. The difference is that now there is aggregate risk: $\omega_1 + \omega_2 = (2, 1)$. The utilities are state independent and the probability assessments are the same for the two traders: (π_1, π_2). The corresponding Edgeworth box is represented in Figure 19.C.2. We see that at any point of the Pareto set the common marginal rate of substitution is smaller than the ratio of probabilities (see Exercise 19.C.2). Hence at an equilibrium we must have $p_1/p_2 < \pi_1/\pi_2$, or $p_1/\pi_1 < p_2/\pi_2$. If, say, $\pi_1 = \pi_2 = \frac{1}{2}$, then $p_1 < p_2$: The price of one contingent unit of consumption is larger for the state for which the consumption good is scarcer. This constitutes the simplest version of a powerful theme of finance theory: that contingent instruments (in our case, a unit of contingent consumption) are comparatively more valuable if their returns (in our case, the amount of consumption they give in the different states) are negatively correlated with the "market return" (in our case, the random variable representing the aggregate initial endowment). ∎

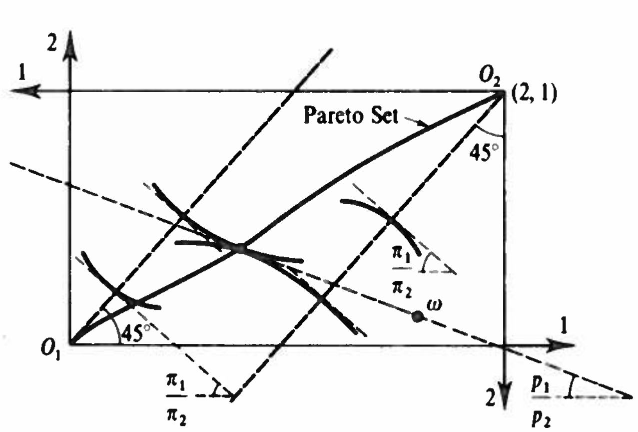

Figure 19.C.2
There is aggregate risk: p_ℓ/π_ℓ negatively correlated with total endowment of commodity ℓ.