CHAPTER 8

Strategic Bargaining in a Market with One-Time Entry

8.1 Introduction

In this chapter we study two strategic models of decentralized trade in a market in which all potential traders are present initially (cf. Model B of Chapter 6). In the first model there is a single indivisible good that is traded for a divisible good ("money"); a trader leaves the market once he has completed a transaction. In the second model there are many divisible goods; agents can make a number of trades before departing from the market. (This second model is close to the standard economic model of competitive markets.)

We focus on the conditions under which the outcome of decentralized trade is competitive; we point to the elements of the models that are critical for a competitive outcome to emerge. In the course of the analysis, several issues arise concerning the nature of the information possessed by the agents. In Chapter 10 we return to the first model and study in detail the role of the informational assumptions in leading to a competitive outcome.

8.2 A Market in Which There Is a Single Indivisible Good

The first model is possibly the simplest model that combines pairwise meetings with strategic bargaining.

Goods A single indivisible good is traded for some quantity of a divisible good ("money").

Time Time is discrete and is indexed by the nonnegative integers.

Economic Agents In period 0, S identical sellers enter the market with one unit of the indivisible good each, and $B > S$ identical buyers enter with one unit of money each. No more agents enter at any later date. Each individual's preferences on lotteries over the price p at which a transaction is concluded satisfy the assumptions of von Neumann and Morgenstern. Each seller's preferences are represented by the utility function p, and each buyer's preferences are represented by the utility function $1 - p$ (i.e. the reservation values of the seller and buyer are zero and one respectively, and no agent is impatient). If an agent never trades then his utility is zero.

Matching In each period any remaining sellers and buyers are matched pairwise. The matching technology is such that each seller meets exactly one buyer and no buyer meets more than one seller in any period. Since there are fewer sellers than buyers, $B - S$ buyers are thus left unmatched in each period. The matching process is random: in each period all possible matches are equally probable, and the matching is independent across periods.

Although this matching technology is very special, the result below can be extended to other technologies in which the probabilities of any particular match are independent of history.

Bargaining After a buyer and a seller have been matched they engage in a short bargaining process. First, one of the matched agents is selected randomly (with probability 1/2) to propose a price between 0 and 1. Then the other agent responds by accepting the proposed price or rejecting it. Rejection dissolves the match, in which case the agents proceed to the next matching stage. If the proposal is accepted, the parties implement it and depart from the market.

Information We assume that the agents have information only about the index of the period and the names of the sellers and buyers in the market. (Thus they know more than just the *numbers* of sellers and buyers in the market.) When matched, an agent recognizes the name of his opponent. However, agents do not remember the past events in their lives. This may be because their memories are poor or because they believe that their personal experiences are irrelevant. Nor do agents receive any information about the events in matches in which they did not take part.

These assumptions specify an extensive game. Note that since the agents forget their own past actions, the game is one of "imperfect recall". We comment briefly on the consequences of this at the end of the next section.

8.3 Market Equilibrium

Given our assumption about the structure of information, a *strategy* for an agent in the game specifies an offer and a response function, possibly depending on the index of the period, the sets of sellers and buyers still in the market, and the name of the agent's opponent. To describe a strategy precisely, note that there are two circumstances in which agent i has to move. The first is when the agent is matched and has been selected to make an offer. Such a situation is characterized by a triple (t, A, j), where t is a period, A is a set of agents that includes i (the set of agents in the market in period t), and j is a member of A of the opposite type to i (i's partner). The second is when the agent has to respond to an offer. Such a situation is characterized by a four-tuple (t, A, j, p), where t is a period, A is a set of agents that includes i, j is a member of A of the opposite type to i, and p is a price in $[0, 1]$ (an offer by j). Thus a strategy for agent i is a pair of functions, the first of which associates a price in the interval $[0, 1]$ with every triple (t, A, j), and the second of which associates a member of the set $\{Y, N\}$ ("accept", "reject") with every four-tuple (t, A, j, p).

The spirit of the solution concept we employ is close to that of sequential equilibrium. An agent's strategy is required to be optimal not only at the beginning of the game but also at every other point at which the agent has to make a decision. A strategy induces a plan of action starting at any point in the game. We now explain how each agent calculates the expected utility of each such plan of action.

First, suppose that agent i is matched and has been selected to make an offer. In such a situation i's information consists of (t, A, j), as described above. The behavior of every other agent in A depends only on t, A, and the agent with whom that agent is matched (if any). Thus the fact that i does not know the events that have occurred in the past is irrelevant, because neither does any other agent, so that no other agent's actions are conditioned on these events. In this case, agent i's information is sufficient, given the strategies of the other agents, to calculate the moves of his future partners, and thus find the expected utility of any plan of action starting at t.

Second, suppose that agent i has to respond to an offer. In this case i's information consists of a four-tuple (t, A, j, p), as described above. If he accepts the offer then his utility is determined by p. If he rejects the offer, then his expected utility is determined by the events in other matches (which determine the probabilities with which he will be matched with any remaining agents) and the other agents' strategies. If p is the offer that is made when all agents follow their equilibrium strategies, then the agent uses these strategies to form a belief about the events in other matches. If p is different from the offer made in the equilibrium—if the play of the game has moved "off the equilibrium path"—then the notion of sequential equilibrium allows the agent some freedom in forming his belief about the events in other matches. We assume that the agent believes that the behavior of all agents in any simultaneous matches, and in the future, is still given by the equilibrium strategies. Even though he has observed an action that indicates that some agent has deviated from the equilibrium, he assumes that there will be no further deviations. Given that the agent expects the other agents to act in the future as they would in equilibrium, he can calculate his expected utility from each possible plan of action starting at that point.

Definition 8.1 A *market equilibrium* is a strategy profile (a strategy for each of the $S + B$ agents), such that each agent's strategy is optimal at every point at which the agent has to make a choice, on the assumption that all the actions of the other agents that he does not observe conform with their equilibrium strategies.

Proposition 8.2 *There exists a market equilibrium, and in every such equilibrium every seller's good is sold at the price of one.*

This result has two interesting features. First, although we do not assume that all transactions take place at the same price, we obtain this as a result. Second, the equilibrium price is the competitive price.

Proof of Proposition 8.2. We first exhibit a market equilibrium in which all units of the good are sold at the price of one. In every event all agents offer the price one, every seller accepts only the price one, and every buyer accepts any price. The outcome is that all goods are transferred, at the price of one, to the buyers who are matched with sellers in the first period. No agent can increase his utility by adopting a different strategy. Suppose, for example, that a seller is confronted with the offer of a price less than one (an event inconsistent with equilibrium). If she rejects this offer, then she will certainly be matched in the next period. Under our assumption that she believes, despite the previous inconsistency with equilibrium, that all agents will behave in the future according to their equilibrium strategies, she believes that she will sell her unit at the price one in the next period. Thus it is optimal for her to reject the offer.

We now prove that there is no other market equilibrium outcome. We use induction on the number of sellers in the market. First consider the case of a market with a single seller ($S = 1$). In this case the set of agents in the market remains the same as long as the market continues to operate. Thus if no transaction has taken place prior to period t, then at the beginning of period t, before a match is established, the expected utilities of the agents depend only on t. For any given strategy profile let $V_i^b(t)$ and $V^s(t)$ be these expected utilities of buyer i and the seller, respectively.

Let m be the infimum of $V^s(t)$ over all market equilibria and all t. Fix a market equilibrium. Since there is just one unit of the good available in the economy, we have $\sum_{i=1}^{B} V_i^b(t) \leq 1 - m$ for all t. Thus for each t there is a buyer for whom $V_i^b(t+1) \leq (1-m)/B$. Suppose the seller adopts the strategy of proposing the price $1 - \epsilon - (1-m)/B$, and rejecting all lower prices, for some $\epsilon > 0$. Eventually she will meet, say in period t, a buyer for whom $V_i^b(t+1) \leq (1-m)/B$. The optimality of this buyer's strategy demands that he accept this offer, so that the seller obtains a utility of $1 - \epsilon - (1-m)/B$. Thus $V^s(t) \geq 1 - \epsilon - (1-m)/B$. Therefore $m \geq 1 - \epsilon - (1-m)/B$, and hence $m \geq 1 - \epsilon B/(B-1)$ for any $\epsilon > 0$, which means that $m = 1$.

Now assume the proposition is valid if the number of sellers in the markets is strictly less than $\overline{S}$. Fix a set of sellers of size $\overline{S}$. For any given strategy profile let $V_j^s(t)$ and $V_i^b(t)$ be the expected utilities of seller j and buyer i, respectively, at the beginning of period t (before any match is established) if all the $\overline{S}$ sellers in the set and all B buyers remain in the market. We shall show that for all market equilibria in a market containing the $\overline{S}$ sellers and B buyers we have $V_j^s(0) = 1$ for every seller j. Let m be the infimum of $V_j^s(t)$ over all market equilibria, all t, and all j. Fix a market equilibrium. For all t we have $\sum_{i=1}^{B} V_i^b(t) \leq (1-m)\overline{S}$. Therefore, in any period t there exists some buyer i such that $V_i^b(t+1) \leq (1-m)\overline{S}/B$. Consider a seller who adopts the strategy of demanding the price $1-\epsilon-(1-m)\overline{S}/B$ and not agreeing to less as long as the market contains the $\overline{S}$ sellers and B buyers. Either she will be matched in some period t with a buyer for whom $V_i^b(t+1) \leq (1-m)\overline{S}/B$ who will then agree to that price, or some other seller will transact beforehand. In the first case the seller's utility will be $1 - \epsilon - (1-m)\overline{S}/B$, while in the second case it will be 1 by the inductive hypothesis. Since a seller can always adopt this strategy, we have

$V_j^s(t) \geq 1 - \epsilon - (1-m)\overline{S}/B$. Therefore $m \geq 1 - \epsilon - (1-m)\overline{S}/B$, and hence $m \geq 1 - \epsilon B/(B - \overline{S})$ for any $\epsilon > 0$, which means that $m = 1$. □

There are three points to notice about the result. First, it does not state that there is a unique market equilibrium—only that the price at which each unit of the good is sold in every market equilibrium is the same. There are in fact other market equilibria—for example, ones in which all sellers reject all the offers made by a particular buyer. Second, the proof remains unchanged if we assume that agents do not recognize the name of their opponents. The informational assumptions we have made allow us to conclude that, at the beginning of each period, the expected utilities of being in the market depend only on the index of the period. Assuming that agents cannot recognize their opponents does not affect this conclusion. Third, the proof reveals the role played by the surplus of buyers in determining the competitive outcome. The probability that a seller is matched in any period is one, while this probability is less than one for a buyer. Although there is no impatience in the model, the situation is somewhat similar to that of a sequential bargaining game in which the seller's discount factor is 1 and the buyer's discount factor is $S/B < 1$.

As we mentioned above, the model is a game with imperfect recall. Each agent forgets information that he possessed in the past (like the names of agents with whom he was matched and the offers that were made). The only information that an agent recalls is the time and the set of agents remaining in the market. The issue of how to interpret the assumption of imperfect recall is subtle; we do not discuss it in detail (see Rubinstein (1991) for more discussion). We simply remark that the assumption we make here has implications beyond the fact that the behavior of an agent can depend only on time and the set of agents remaining in the market. The components of an agent's strategy that specify his actions after arbitrary histories can be interpreted as reflecting his beliefs about what other agents expect him to do in such cases. Thus our assumption means also that no event in the past leads an agent to change his beliefs about what other agents expect him to do.

8.4 A Market in Which There Are Many Divisible Goods

The main differences between the model we study here and that of the previous two sections are that the market here contains many divisible goods, rather than a single indivisible good, and that agents may make many transactions before departing from the market. We begin with an outline of the model.

There is a continuum of agents in the market, trading m divisible goods. Time is discrete and is indexed by the nonnegative integers. All agents enter the market simultaneously in period 0; each brings with him a bundle of goods, which may be stored costlessly. In period 0 and all subsequent periods there is a positive probability that any given agent is matched with a trading partner. Once a match is formed, one of the parties is selected at random to propose a trade (an exchange of goods). The other agent may accept or reject this proposal. If he rejects it then he may, if he wishes, leave the market. Agents who remain in the market are matched anew with positive probability each period and may execute a sequence of transactions. All matches cease after one period: even if an agent who is matched in period t is not matched with a new partner in period $t+1$, he must abandon his old partner. An agent obtains utility from the bundle he holds when he leaves the market. Note that agents may not leave the market immediately after accepting an offer; they may leave *only* after rejecting an offer. Although this assumption lacks intuitive appeal, it formalizes the idea that an agent who is about to depart from the market always has a "last chance" to receive an offer.

We now spell out the details of the model.

Goods There are m divisible goods; a bundle of goods is a member of $\mathcal{R}^m_+$.

Time Time is discrete and is indexed by the nonnegative integers.

Economic Agents There is a continuum of agents in the market. Each agent is characterized by the initial bundle with which he enters the market and his von Neumann–Morgenstern utility function over the union of the set $\mathcal{R}^m_+$ of feasible consumption bundles and the event D of staying in the market forever. Each agent chooses the period in which to consume, and is indifferent about the timing of his consumption (i.e. is not impatient). The agents initially present in the market are of a finite number K of *types*. All members of any given type k have the same utility function $u_k: \mathcal{R}^m_+ \cup \{D\} \to \mathcal{R} \cup \{-\infty\}$ and the same initial bundle $\omega_k \in \mathcal{R}^m_+$. For each type k there is initially the measure n_k of agents in the market (with $\sum_{k=1}^{K} n_k = 1$). Each utility function u_k is restricted as follows. There is a continuous function $\phi_k: \mathcal{R}^m_+ \to \mathcal{R}$ that is increasing and strictly concave on the interior of $\mathcal{R}^m_+$ and satisfies $\phi_k(x) = 0$ if x is on the boundary of $\mathcal{R}^m_+$. Let $\overline{\phi} > 0$ be a number, and let $X_k = \{x \in \mathcal{R}^m_+: \phi_k(x) \geq \overline{\phi}\}$. Then u_k is given by $u_k(x) = \phi_k(x)$ if $x \in X_k$ and $u_k(x) = -\infty$ for all other x (including $x = D$). (The number $\overline{\phi}$ can be interpreted as the minimal utility necessary to survive. The assumption that $u_k(D) = -\infty$ means that agents must leave the market eventually.) Further, we assume that

$\omega_k \in X_k$. An interpretation of the concavity of the utility functions is that each agent is risk-averse. We make two further assumptions on the utility functions.

1. For each k there is a unique tangent to each indifference curve of u_k at every point in X_k.
2. Fix some type k and some nonzero vector $p \in \mathcal{R}^m_+$. Consider the set $S(k,p)$ of bundles c for which the tangent to the indifference curve of u_k through c is $\{x: px = pc\}$ (i.e. $S(k,p)$ is k's "income-expansion" path at the price vector p). Then for every vector $z \in \mathcal{R}^m$ for which $pz > 0$ there exists a positive integer L such that $u_k(c + z/L) > u_k(c)$ for every c in $S(k,p)$.

The first assumption is weaker than differentiability of u_k on X_k (since it relates only to the indifference curves of u_k). Note that it guarantees that for each vector $z \in \mathcal{R}^m$ and each bundle c in $S(k,p)$ we can find an integer L such that $u_k(c + z/L) > u_k(c)$. The second assumption imposes the stronger condition that for each vector $z \in \mathcal{R}^m$ we can find a single L such that $u_k(c+z/L) > u_k(c)$ for *all* c in $S(k,p)$. This second assumption is illustrated in Figure 8.1. (It is related to Gale's (1986c) assumption that the indifference curves of the utility function have uniformly bounded curvature.)

Matching In every period each agent is matched with a partner with probability $0 < \alpha < 1$ (independent of all past events). Matches are made randomly; the probability that any given agent is matched in any given period with an agent in a given set is proportional to the measure of that set in the market in that period. Notice that since the probability of an agent being matched is less than one, in every period there are agents who have never been matched. Thus even though agents leave the market as time passes, at any finite time a positive measure of every type remains.

Bargaining Once a match is established, each party learns the type (i.e. utility function and initial bundle) and current bundle of his opponent. The members of the match then conduct a short bargaining session. First, one of them is selected to propose a vector z of goods, to be transferred to him from his opponent. (That is, an agent who holds the bundle x and proposes the trade z will hold the bundle $x+z$ if his proposal is accepted.) This vector will typically contain positive and negative elements; it must have the property that it is feasible, in the sense that the bundles held by both parties after the exchange are nonnegative. The probability of each party being selected to make a

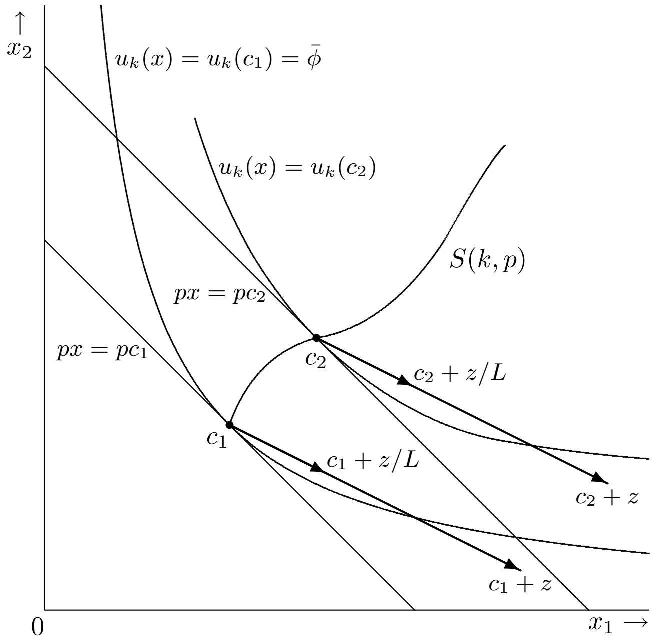

Figure 8.1 An illustration of Assumption 2 on the utility functions.

proposal is 1/2, independent of all past events. After a proposal is made, the other party either accepts or rejects the offer.

Exit In the event an agent rejects an offer, he chooses whether or not to stay in the market. An agent who makes an offer, accepts an offer, or who is unmatched, must stay in the market until the next period: he may not exit. An agent who exits obtains the utility of the bundle he holds at that time.

8.5 Market Equilibrium

A *strategy* for an agent is a plan that prescribes his bargaining behavior for each period, each bundle he currently holds, and each type and current bundle of his opponent. An agent's bargaining behavior is specified by the offer to be made in case he is chosen to be the proposer and, for each possible offer, one of the actions "accept", "reject and stay", or "reject and exit".

An assumption that leads to this definition of a strategy is that each agent observes the index of the period, his current bundle, and the current bundle and type of his opponent, but no past events. Events in the life of the agent (like the type of agents he met in the past, the offers that were made, and the sequence of trades) cannot affect his behavior except insofar as they influence his current bundle. Gale (1986a, Proposition 1) derives the restriction from more primitive assumptions. The idea is the following. Given that there is a continuum of agents, the probability of an agent meeting any particular individual is zero, so that an agent can learn from his personal history about only a finite number of other agents—a set of measure zero. Further, the matching technology forces partners to separate at the end of each period. Thus even if an agent records the entire list of past events, there is no advantage in conditioning his strategy on this information.

We restrict attention to the case in which all agents of a given type use the same strategy. As trade occurs, the bundle held by each agent changes. Different agents of the same type, even though they use the same strategy, may execute different trades. Thus the number of different bundles held by agents may increase. However, the number of different bundles held by agents is finite at all times. Thus in any period the market is given by a finite list $(k_i, c_i, \nu_i)_{i=1,\ldots,I}$, where ν_i is the measure of agents who are still in the market, currently hold the bundle c_i, and are of type k_i. We call such a list a *state of the market*. We say that an agent of type k who holds the bundle c is *characterized by* (k, c).

With each K-tuple σ of strategies is associated a state of the market $\rho(\sigma, t)$ in each period t. Although each agent faces uncertainty, the presence of a continuum of agents allows us to define ρ in a deterministic fashion. For example, since in each period the probability that any given agent is matched is α, we take the fraction of agents with any given characteristic who are matched to be precisely α.

Formally, $\rho(\sigma, t+1)$ is generated from $\rho(\sigma, t) = (k_i, c_i, \nu_i)_{i=1,\ldots,I}$ by the following transition rules. The set of agents characterized by (k_j, c_j) who are matched with agents characterized by (k_h, c_h) and are selected to make an offer has measure $\alpha\nu_j\nu_h/2\sum_{i=1}^{I}\nu_i$. If σ instructs these agents to offer a trade z that, according to σ, is accepted, then the measure $\alpha\nu_j\nu_h/2\sum_{i=1}^{I}\nu_i$ of agents is transferred from (k_j, c_j) to $(k_j, c_j + z)$, and the measure $\alpha\nu_j\nu_h/2\sum_{i=1}^{I}\nu_i$ of agents is transferred from (k_h, c_h) to $(k_h, c_h - z)$. If σ instructs the responders to reject z and exit, then the measure of agents characterized by (k_h, c_h) is reduced by $\alpha\nu_j\nu_h/2\sum_{i=1}^{I}\nu_i$. Otherwise the measures of agents remain the same.

As an illustration of the determination of $\rho(\sigma, t)$, consider a market in which there are two types, each comprising half of the population. Both types have the same utility function. There are two goods; each agent of type 1 initially owns the bundle $(2,0)$, while each agent of type 2 owns the bundle $(0,2)$. Suppose that the agents use the following pair of strategies. An agent of type 1 offers and accepts only the trade $(-1,1)$ whenever he holds the bundle $(2,0)$; in all other cases he offers $(0,0)$ and rejects all offers. An agent of type 2 offers and accepts only the trade $(1,-1)$ whenever he holds the bundle $(0,2)$; in all other cases he offers $(0,0)$ and rejects all offers. An agent leaves the market if and only if he holds the bundle $(1,1)$, is matched with a partner, and is chosen to respond to an offer.

In any period, the bundle held by each agent is $(2,0)$, $(0,2)$, or $(1,1)$. Suppose that in period t the measures of agents holding these three bundles are p, q, and r. Let $s = p + q + r$. The measures of agents holding these bundles in period $t+1$ can be found as follows. The measure αr of those holding $(1,1)$ will be matched in period $t+1$; the measure $\alpha r/2$ will be chosen to respond, and hence will leave the market. The remainder of those holding $(1,1)$ (the measure $r(1-\alpha)/2$) will stay in the market through period $t+1$, making the null offer $(0,0)$ if matched. Of the agents holding $(2,0)$, the measure $\alpha pq/s$ will be matched with agents holding $(0,2)$, and will trade and join the set of agents holding $(1,1)$. The remainder will retain $(2,0)$. Thus the total measure of agents holding $(2,0)$ in period $t+1$ is $p(1-\alpha q/s)$. Similarly, the total measure of agents holding $(0,2)$ in period $t+1$ is $q(1-\alpha p/s)$. The total measure of agents holding $(1,1)$ in period $t+1$ is $2\alpha pq/s + r(1-\alpha/2)$.

We emphasize that although we take the evolution of the state of the market to be deterministic, each agent still faces a nondegenerate stochastic process. Given a strategy profile σ, for all pairs (k,c) the state of the market $\rho(\sigma,t)$ induces a well-defined probability that any agent will be matched in period t with an agent characterized by (k,c).

The notion of equilibrium we use is the following.

Definition 8.3 A *market equilibrium* is a K-tuple σ^* of strategies, one for each type, each of which satisfies the following condition for any trade z, bundles c and c', type k, and period t. The behavior prescribed by each agent's strategy from period t on is optimal, given that in period t the agent holds c and has either to make an offer or to respond to the offer z made by his opponent, who is of type k and holds the bundle c', given the strategies of the other types, and given that the agent believes that the state of the market is $\rho(\sigma^*,t)$.

This notion of equilibrium is not directly equivalent to any game-theoretic notion. However, as in the previous model, it is closely related to the notion of sequential equilibrium. Each agent's strategy has to be optimal after *every* event, including events that are inconsistent with the equilibrium. (These events are: (1) being matched in period t with an agent of type k holding a bundle c when no agent of type k holds c in period t if all agents follows σ^*; (2) being confronted with an offer that the opponent does not make if he adheres to σ^*; (3) having an offer rejected when σ^* calls for the opponent to accept; (4) making a move that is different from that dictated by σ^*.) In order to test the optimality of his strategy, an agent must form a belief about the state of the market, which determines the probabilities with which he meets the various types of agents. If no unexpected event has occurred up to period t, then the equilibrium state of the market in period t, namely $\rho(\sigma^*, t)$, provides this belief. However, once an event that is inconsistent with equilibrium has occurred, an agent must make a conjecture about the current state of the market. The definition of equilibrium requires that each agent believe that, after any sequence of events, the state of the market is the same as it is in equilibrium. This excludes the possibility that an agent interprets out-of-equilibrium behavior by other agents as a signal that the behavior of a *positive* measure of agents was different than in equilibrium, so that the state of the market has changed. This assumption is close to that of the previous model.

8.6 Characterization of Market Equilibrium

An *allocation* is a K-tuple of bundles $(x_1, \ldots, x_K)$ for which $\sum_{k=1}^{K} n_k x_k = \sum_{k=1}^{K} n_k \omega_k$. An allocation $(x_1, \ldots, x_K)$ is *competitive* if there exists a price vector $p \in \mathcal{R}^m_{++}$ such that for all k the bundle x_k maximizes u_k over the budget set $\{x \in X_k \colon px \leq p\omega_k\}$.

The result below establishes a close relationship between competitive allocations and the allocations induced by market equilibria. Before stating this result we need to introduce some terminology. Suppose that the market equilibrium calls for agents characterized by (k, c) who are matched in period t with agents characterized by (k', c') to reject some offer z and leave the market. Then we say that all agents characterized by (k, c) are *ready to leave the market* in period t.

Proposition 8.4 *For every market equilibrium there is a competitive allocation $(x_1, \ldots, x_K)$ such that each agent of type k $(= 1, \ldots, K)$ leaves the market with the bundle x_k with probability one.*

Proof. Consider a market equilibrium; all of our statements are relative to this equilibrium. All agents of type k who hold the bundle c at the beginning of period t (before their match has been determined) face the same probability distribution of future trading opportunities. Thus in the equilibrium all such agents have the same expected utility; we denote this utility by $V_k(c,t)$.

Step 1. $V_k(c,t) \geq u_k(c)$ for all values of k, c, and t.

Proof. Suppose that an agent of type k who holds the bundle c in period t makes the null offer whenever he is matched and is chosen to propose a trade, and rejects every offer and leaves the market when he is matched and chosen to respond. Since he is matched and chosen to respond to an offer in finite time with probability one, this strategy guarantees him a payoff of $u_k(c)$. (Recall that all agents are indifferent about the timing of consumption.) Thus $V_k(c,t) \geq u_k(c)$.

Step 2. $V_k(c,t) \geq V_k(c,t+1)$ for all values of k, c, and t.

Proof. The assertion follows from the fact that by proposing the null trade and rejecting any offer and staying in the market, any agent in the market in period t is sure of staying in the market until period $t+1$ with his current bundle.

Step 3. For an agent of type k who holds the bundle c and is ready to leave the market in period t we have $V_k(c,t+1) = u_k(c)$.

Proof. By Step 1 we have $V_k(c,t+1) \geq u_k(c)$. If $V_k(c,t+1) > u_k(c)$ and the circumstances that make the agent leave the market are realized (in which case he would leave with the bundle c), then he is better off by deviating and staying in the market until period $t+1$.

Step 4. Suppose that an agent of type k holds the bundle c and is ready to leave the market in period t. Then it is optimal for him to accept any offer z (of a transfer *from* him *to* the proposer) for which $u_k(c-z) > u_k(c)$.

Proof. If he accepts the offer, then his expected utility $V_k(c-z,t+1)$ in the continuation is at least $u_k(c-z)$ (by Step 1), and this exceeds his expected utility in the continuation if he rejects the offer, which is $V_k(c,t+1) = u_k(c)$ (see Step 3).

Step 5. For any period t and any given agent, the probability that in some future period the agent will be chosen to make an offer to an agent who is ready to leave the market is one.

Proof. Let Q_s be the measure of the set of agents in the market in period s, and let E_s be the measure of the set of agents who are ready to leave the market in period s. The probability that an agent in the market is matched with an agent who is ready to leave is $\alpha E_s/Q_s$, in which case the agent will be chosen with probability $1/2$ to make an offer. Thus the probability of not being able to make an offer to an agent who is ready to leave is $1 - \alpha E_s/2Q_s$. The measure of agents who actually leave is at most $\alpha E_s/2$. (Recall that an agent who is ready to leave does so only under some circumstances, not necessarily whenever he has to respond to an offer.) Hence $Q_s - \alpha E_s/2 \leq Q_{s+1}$, so that $1 - \alpha E_s/2Q_s \leq Q_{s+1}/Q_s$. Thus the probability of not being able to make an offer from period t through period s to an agent who is ready to leave the market, where $s > t$, is at most Q_{s+1}/Q_t. Since the utility of staying in the market forever is $-\infty$, $Q_s \to 0$ as $s \to \infty$, and thus the probability that an agent will be able to make an offer in some future period to an agent who is ready to leave the market is one.

Step 6. There is a vector $p \in \mathcal{R}^m_{++}$, unique up to multiplication by a nonnegative scaler, such that, for all k, if each member of a set of agents of positive measure of type k leaves the market in some period with the bundle c, then the tangent to the indifference curve $\{x \in X_k: u_k(x) = u_k(c)\}$ at c is $\{x: px = pc\}$ (i.e. $pz > 0$ for all z such that $u_k(c+z) > u_k(c)$).

Proof. Suppose that each member of a set of positive measure of agents of type k_1 leaves the market in period t_1 with the bundle c_1, and each member of a set of positive measure of agents of type k_2 leaves the market in period t_2 with the bundle c_2. Assume, contrary to the claim, that the tangent to the indifference curve $\{x \in X_{k_1}: u_{k_1}(x) = u_{k_1}(c_1)\}$ at c_1 is different from the tangent to the indifference curve $\{x \in X_{k_2}: u_{k_2}(x) = u_{k_2}(c_2)\}$ at c_2. Then (by the assumption that each indifference curve has a unique tangent at every point) there is a trade z between an agent of type k_1 holding the bundle c_1 and an agent of type k_2 holding the bundle c_2 that makes both agents better off. More precisely, $c_1 + z \in X_{k_1}$, $c_2 - z \in X_{k_2}$, $u_{k_1}(c_1 + z) > u_{k_1}(c_1)$, and $u_{k_2}(c_2 - z) > u_{k_2}(c_2)$.

First assume that $t_1 < t_2$. Consider an agent of type k_1 who holds the bundle c_1 in period t_1. By our hypothesis he is ready to leave the market. We will show that the following is a profitable deviation. Instead of leaving the market, he stays until period t_2 (by proposing the null trade and rejecting all offers as necessary). In period t_2 there is a positive probability that he is matched with an agent of type k_2 who holds c_2 (and thus is ready to leave the market). In this event he proposes the mutually beneficial trade z. In every other event he departs from the market at the first opportunity. By Step 4 the agent of type k_2 accepts the offer, so that the agent of type k_1 either achieves the bundle $c_1 + z$ in period t_2 (with positive probability) or holds the bundle c_1 in that period. Thus the agent of type k_1 achieves an expected utility in excess of $u_{k_1}(c_1)$, so that the deviation is profitable.

If $t_1 = t_2 = t$ then an agent of type k_1 who faces the circumstances in which he plans to leave the market can deviate from the equilibrium and postpone his departure by one period. In period $t+1$ there is a positive probability that he will be matched with an agent of type k_2 who was ready to leave the market in period t but did not do so. Suppose the agent of type k_1 offers the trade z to such an agent of type k_2. If the latter accepts this offer then by Step 1 that agent's expected utility in the continuation is at least $u_{k_2}(c_2 - z)$, while if he rejects the offer then he either leaves the market with the bundle c_2 or enters period $t+2$ with that bundle. But $V_{k_2}(c_2, t+2) \leq V_{k_2}(c_2, t+1) = u_{k_2}(c_2)$ by Steps 2 and 3. So the fact that $u_{k_2}(c_2 - z) > u_{k_2}(c_2)$, and the requirement that each agent's strategy prescribe optimal actions after every history, demand that the agent of type k_2 accept the offer. Hence, as in the previous case, the agent of type k_1 has a profitable deviation from his purported equilibrium strategy.

We conclude that the tangents to the indifference curves of agents who leave the market, at the bundles with which they depart, coincide.

Step 7. Let p be the vector defined in Step 6. Then for all k, c, and t we have $V_k(c,t) \geq \max_{x \in X_k}\{u_k(x){:}\, px \leq pc\}$.

Proof. Assume to the contrary that $V_\kappa(c,t) < \max_{x \in X_\kappa}\{u_\kappa(x){:}\, px \leq pc\}$ for some κ, c, and t. Then there is a vector z such that $V_\kappa(c,t) < u_\kappa(c+z)$ and $pz < 0$. (See Figure 8.2.) We shall argue that an agent of type κ who holds the bundle c has a deviation that yields him the utility $u_\kappa(c+z)$. By Assumption 2 (p. 158) for each $k = 1, \ldots, K$ there exists a positive integer L_k such that $u_k(c_k - z/L_k) > u_k(c_k)$ whenever c_k is a bundle with which agents of type k leave the market (using Step 6). By Step 4, an agent of type k who is ready to leave the market thus accepts an offer of the trade $-z/L_k$. Hence there exists a positive integer L such that all agents (of whatever type) who are ready to leave the market would accept the trade $-z/L$ before doing so. Now, by Step 5 the probability that in some future period a given agent of type κ will be able to make an offer to an agent who is ready to leave the market is one. Thus the probability that he will be able to make L such offers is also one. Hence an agent of type κ who holds the bundle c in period t can profitably deviate from his original strategy and with certainty carry out L trades of z/L before he leaves the market, thereby attaining the utility $u_\kappa(c+z)$, which exceeds $V_\kappa(c,t)$.

Step 8. For an agent of type k who leaves the market with the bundle c we have $u_k(c) = \max_{x \in X_k}\{u_k(x){:}\, px \leq p\omega_k\}$.

Proof. By Step 7 we have $V_k(\omega_k, 0) \geq \max_{x \in X_k}\{u_k(x){:}\, px \leq p\omega_k\}$. Let $\mathbf{y}_k$ be the random bundle with which an agent of type k leaves the market. We show that for all k the random variable $\mathbf{y}_k$ is degenerate and $V_k(\omega_k, 0) = \max_{x \in X_k}\{u_k(x){:}\, px \leq p\omega_k\}$. Assume this is not so for

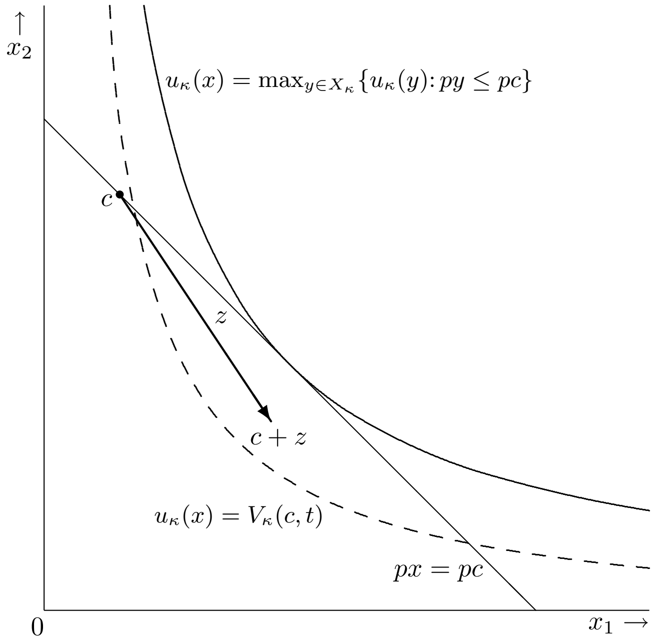

Figure 8.2 A vector z for which $V_\kappa(c, t) < u_\kappa(c + z)$ and $pz < 0$.

$k = \kappa$. By the strict concavity of u_k and Jensen's inequality we have $V_k(\omega_k, 0) = \mathrm{E}[u_k(\mathbf{y}_k)] \le u_k(\mathrm{E}[\mathbf{y}_k])$ (where E is the expectation operator), with strict inequality unless $\mathbf{y}_k$ is degenerate. Let $y_k = \mathrm{E}[\mathbf{y}_k]$. Hence $u_k(y_k) \ge \max_{x \in X_k}\{u_k(x): px \le p\omega_k\}$, with strict inequality for $k = \kappa$. Therefore $py_k \ge p\omega_k$ for all k, and $py_\kappa > p\omega_\kappa$. Thus $p\sum_{k=1}^{K} n_k y_k > p\sum_{k=1}^{K} n_k \omega_k$, contradicting the condition $\sum_{k=1}^{K} n_k y_k = \sum_{k=1}^{K} n_k \omega_k$ for $(y_1, \ldots, y_K)$ to be an allocation. ☐

Note that Assumption 2 (p. 158) is used in Step 7. It is used to show that if $pz < 0$ then there is a trade in the direction $-z$ that makes any agent who is ready to leave the market better off. Thus, by executing a sequence of such trades, an agent who holds the bundle c is assured of eventually obtaining the bundle $c - z$. Suppose the agents' preferences do not satisfy Assumption 2. Then the curvature of the agents' indifference curves at the bundles with which they exit from the market in period t might increase with t, in such a way that the exiting agents are willing to accept only a sequence of successively smaller trades in the direction $-z$, a sequence that never adds up to z itself.

Two arguments are central to the proof. First, the allocation associated with the bundles with which agents exit is efficient (Step 6). The idea is that if there remain feasible trades between the members of two sets of agents that make the members of both sets better off, then by waiting sufficiently long each member of one set is sure of meeting a member of the other set, in which case a mutually beneficial trade can take place. Three assumptions are important here. First, no agent is impatient. Every agent is willing to wait as long as necessary to execute a trade. Second, the matching technology has the property that if in some period there is a positive measure of agents of type k holding the bundle c, then in *every* future period there will be a positive measure of such agents, so that the probability that any other given agent meets such an agent is positive. Third, an agent may not leave the market until he has rejected an offer. This gives every agent a chance to make an offer to an agent who is ready to leave the market. If we assume that an agent can leave the market whenever he wishes then we cannot avoid inefficient equilibria in which all agents leave the market simultaneously, leaving gains from trade unexploited.

The second argument central to the proof is contained in Step 7. Consider a market containing two types of agents and two goods. Suppose that the bundles with which the members of the two types exit from the market leave no opportunities for mutually beneficial trade unexploited. Given the matching technology, in every period there will remain agents of each type who have never been matched and hence who still hold their initial bundles. At the same time, after a number of periods some agents will hold their final bundles, ready to leave the market. If the final bundles are not competitive, then for one of the types—say type 1—the straight line joining the initial bundle and the final bundle intersects the indifference curve through the final bundle. This means that there is some trade z with the property that $u_1(\omega_1 + Lz) > u_1(x_1)$ for some integer L, where x_1 is the final bundle of an agent of type 1, and $u_1(x_1 - z) > u_1(x_1)$. Put differently, a number of executions of z makes an agent of type 1 currently holding the initial bundle better off than he is when he holds the final bundle, and a single execution of $-z$ makes an agent of type 1 who is ready to leave the market better off. Given the matching technology, any agent can (eventually) meet as many agents of type 1 who are ready to leave as he wishes. Thus, given that the matching technology forces some agents to achieve their final bundles before others (rather than all of them achieving the final bundles simultaneously), there emerge unexploited opportunities for trade whenever the final outcome is not competitive, even when it is efficient. Once again we see the role of the three assumptions that the agents are patient, the matching technology leaves a positive measure unmatched in every period, and an agent cannot exit until he has rejected an offer. Another assumption that is significant here is that each agent can make a sequence of transactions before leaving the market. This assumption increases the forces of competition in the market, since it allows an agent to exploit the opportunity of a small gain from trade without prejudicing his chances of participating in further transactions.

8.7 Existence of a Market Equilibrium

Proposition 8.4 leaves open the question of the existence of a market equilibrium. Gale (1986b) studies this issue in detail and establishes a converse of Proposition 8.4: to every competitive equilibrium there is a corresponding market equilibrium. (Thus, in particular, a market equilibrium exists.) We do not provide a detailed argument here. Rather we consider two cases in which a straightforward argument can be made.

First consider a modification of the model in which agents may make "short sales"—that is, agents may hold negative amounts of goods, so that any trade is feasible. This case avoids some difficulties associated with the requirement that trades be feasible and illustrates the main ideas. (It is studied by McLennan and Sonnenschein (1991).) Assume that for every bundle c, type k, and price vector p, the maximizer of $u_k(x)$ over $\{x\colon px \leq pc\}$ is unique, and let $\hat{z}(p,c,k)$ be the difference between this maximizer and c; we refer to $\hat{z}(p,c,k)$ as the *excess demand* at the price vector p of an agent characterized by (k,c). If $\hat{z}(p,c,k) = 0$ then an agent characterized by (k,c) holds the bundle (c) that maximizes his utility at the price vector p. Let p^* be the price vector corresponding to a competitive equilibrium of the market. Consider the strategy profile in which the strategy of an agent characterized by (k,c) is the following. Propose the trade $\hat{z}(p^*,k,c)$. If $\hat{z}(p^*,k,c) \neq 0$ then accept an offer[1] z if $p^*(-z) \geq 0$; otherwise reject z and stay in the market. If $\hat{z}(p^*,k,c) = 0$ then accept an offer z if $p^*(-z) > 0$; otherwise reject z and leave the market. The outcome of this strategy profile is that each agent eventually leaves the market with his competitive bundle (the bundle that maximizes his utility over his budget set at the price p^*). If all other agents adhere to the strategy profile, then any given agent accepts any offer he is faced with; his proposal to trade his excess demand is accepted the first time he is matched and chosen to be the proposer, and he leaves the market in the next period in which he is matched and chosen to be the responder.

[1] That is, a trade after which the agent holds the bundle $c - z$.

We claim that the strategy profile is a market equilibrium. It is optimal for an agent to accept any trade that results in a bundle that is worth not less than his current bundle, since with probability one he will be matched and chosen to propose in the future, and in this event his proposal to trade his excess demand will be accepted. It is optimal for an agent to reject any trade that results in a bundle that is worth less than his current bundle, since no agent accepts any trade that decreases the value of his bundle. Finally, it is optimal for an agent to propose his excess demand, since this results in the bundle that gives the highest utility among all the trades that are accepted.

We now return to the model in which in each period each agent must hold a nonnegative amount of each good. In this case the trading strategies must be modified to take into account the feasibility constraints. We consider only the case in which there are two goods, the market contains only two types of equal measure, and the initial allocation is not competitive. Then for any competitive price p^* we have $\hat{z}(p^*, 1, \omega_1) = -\hat{z}(p^*, 2, \omega_2) \neq 0$. Consider the strategy profile in which the strategy of an agent characterized by (k, c) is the following.

Proposals Propose the maximal trade in the direction of the agent's optimal bundle that does not increase or change the sign of the responder's excess demand. Precisely, if matched with an agent characterized by (k', c') and if $\hat{z}_1(p^*, k, c)$ has the same sign as $\hat{z}_1(p^*, k', c')$ (where the subscript indicates good 1), then propose $z = 0$. Otherwise, propose the trade $\hat{z}(p^*, k, c)$ if $|\hat{z}(p^*, k, c)| \leq |\hat{z}(p^*, k', c')|$, and the trade $-\hat{z}(p^*, k', c')$ if $|\hat{z}(p^*, k, c)| > |\hat{z}(p^*, k', c')|$, where $|x|$ is the Euclidian norm of x.

Responses If $\hat{z}(p^*, k, c) \neq 0$ then accept an offer z if $p^*(-z) > 0$, or if $p^*(-z) = 0$ and $\hat{z}_i(p^*, k, c - z)$ has the same sign as, and is smaller than $\hat{z}_i(p^*, k, c)$ for $i = 1, 2$. Otherwise reject z and stay in the market. If $\hat{z}(p^*, k, c) = 0$ then accept an offer z if $p^*(-z) > 0$; otherwise reject z and leave the market.

As in the previous case, the outcome of this strategy profile is that each agent eventually leaves the market with the bundle that maximizes his utility over his budget set at the price p^*. If all other agents adhere to the strategy profile, then any given agent realizes his competitive bundle the first time he is matched with an agent of the other type; until then he makes no trade. The argument that the strategy profile is a market equilibrium is very similar to the argument for the model in which the feasibility constraints are ignored. An agent characterized by (k, c) is assured of eventually achieving the bundle that maximizes u_k over $\{x \in X_k : px \leq pc\}$, since he does so after meeting only a finite number of agents of one of the types who have never traded (since any such agent has a nonzero excess demand), and the probability of such an event is one.

8.8 Market Equilibrium and Competitive Equilibrium

Propositions 8.2 and 8.4 show that the noncooperative models of decentralized trade we have defined lead to competitive outcomes. The first proposition, and the arguments of Gale (1986b), show that the converse of the results are also true: every distribution of the goods that is generated by a competitive equilibrium can be attained as the outcome of a market equilibrium.

In both models the technology of trade and the agents' lack of impatience give rein to competitive forces. If, in the first model, a price below 1 prevails, then a seller can push the price up by waiting (patiently) until he has the opportunity to offer a slightly higher price; such a price is accepted by a buyer since otherwise he will be unable, with positive probability, to purchase the good. If, in the second model, the allocation is not competitive, then an agent is able to wait (patiently) until he is matched with an agent to whom he can offer a mutually beneficial trade.

An assumption that is significant in the two models is that agents cannot develop personal relationships. They are anonymous, are forced to separate at the end of each bargaining session, and, once separated, are not matched again. In Chapter 10 we will see that if the agents have personal identities then the competitive outcome does not necessarily emerge.

Notes

The model of Section 8.2 is closely related to the models of Binmore and Herrero (1988a) and Gale (1987, Section 5), although the exact form of Proposition 8.2 appears in Rubinstein and Wolinsky (1990). The model of Section 8.4 and the subsequent analysis is based on Gale (1986c), which is a simplification of the earlier paper Gale (1986a). The existence of a market equilibrium in this model is established in Gale (1986b).

Proposition 8.2 is related to Gale (1987, Theorem 1), though Gale deals with the limit of the equilibrium prices when $\delta \to 1$, rather than with the limit case $\delta = 1$ itself. Gale's model differs from the one here in that there is a finite number of types of agents (distinguished by different reservation prices), and a continuum of agents of each type. Further, each agent can condition his behavior on his entire personal history. However, given the matching technology and the fact that each pair must separate at the end of each period, the only information relevant to each agent is the time and the names of the agents remaining in the market, as we assumed in Proposition 8.2. Thus we view Proposition 8.2 as the analog of Gale's theorem in the case that the market contains a finite number of agents.

Binmore and Herrero (1988a) investigate alternative information structures and define a solution concept that leads to the same conclusion about the relation between the sets of market equilibria and competitive equilibria as the models we have described. The relation between Proposition 8.4 and the theory of General Equilibrium is investigated by McLennan and Sonnenschein (1991), who also prove a variant of the result under the assumption that the behavior dictated by the strategies does not depend on time. Gale (1986e) studies a model in which the agents—workers and firms—are asymmetrically informed. Workers differ in their productivities and in their payoffs outside the market under consideration. These productivities and payoffs are not known by the firms and are positively correlated, so that a decrease in the offered wage reduces the quality of the supply of workers. Gale examines the nature of wage schedule offered in equilibrium.