Chapter 1

Perfect Competition

1.1 Introduction

The objective of this chapter is to provide a complete and precise description of perfect competition as the equilibrium of a non-cooperative game with a large but finite numbers of players.

For reasons of tractability, all the analysis takes place within the framework of a static pure exchange economy. Ignoring production is a fairly drastic simplification. The justification is that it allows us to focus on price formation and exchange without unnecessary complications. To include production, one would have to answer a number of awkward questions. What is the objective function of the firm? When and how do households receive income from firms? What is the timing of inputs and outputs when production plans have to be feasible? These are questions that deserve to be answered but are outside the scope of the present study.

In a pure exchange economy, competitive equilibrium is characterized by two properties, efficiency and budget balance. The First Theorem of Welfare Economics tells us that competitive equilibrium allocations are efficient. Under the usual convexity assumptions, the Second Theorem tells us that an efficient allocation has a vector of supporting prices $p \neq 0$ at which each agent's consumption bundle minimizes the cost of achieving that level of utility. If x_i is agent i's consumption bundle and another bundle x'_i is preferred to x_i then $p \cdot x'_i > p \cdot x_i$. However, the value of the agent's consumption bundle may not be equal to the value of his endowment at these prices. To ensure that an efficient allocation is a competitive equilibrium, we need to have budget balance: the value of an agent's consumption bundle, determined using the efficiency prices, must be equal to the value of his endowment: $p \cdot x_i = p \cdot e_i$. Thus we can think of a competitive equilibrium allocation as an efficient allocation that satisfies budget balance.

It is convenient to divide the analysis of competition into two parts, one concerned with efficiency and one with budget balance. The conditions needed for efficiency are somewhat weaker than those required for budget balance and the analysis is somewhat more general. In particular, efficiency is a property of strategic equilibrium in finite economies, whereas budget balance requires large numbers of players.

Many of the arguments used in this chapter are familiar. The novel part of the analysis is the assumption of a finite number of players. The earlier literature assumes a non-atomic continuum of players. The continuum assumption is justified by the claim that a continuum of agents is a good approximation to a large but finite number of players. It is important to test the validity of this assumption by proving a competitive limit theorem, showing that as the number of players becomes unboundedly large, the behavior of the finite game is indeed a good approximation to the continuum game.

The vehicle used for the analysis is a dynamic matching and bargaining game (DMBG). The DMBG is special and has no particular claim to generality. The important advantage of the DMBG is that it is a playable game, in which all the rules and assumptions are well specified. If we can analyze it convincingly we shall have learned a lot.

The main discovery of this chapter will be that in order to prove a competitive limit theorem we need some assumptions that are not imposed on primitives of the model and are not derived directly from properties of equilibrium. The most important of these concerns the idea that in a competitive market a single agent has a negligible effect on the equilibrium. This idea comes immediately into conflict with the well established principle that, in repeated games, any player can have a large effect because other players condition their actions on his. The Folk Theorem of repeated games, which holds for arbitrary numbers of players, is perhaps the best known example of this principle. The games under consideration here are not repeated games, but the same general idea applies. In choosing a strategy, a player has to consider not just its direct impact on his payoffs, but also the reaction of the other players. Because of the possible reactions of other players, it may be possible to sustain many different outcomes as strategic equilibria. Some of these will depart from the perfectly competitive outcome, even in a frictionless market with many players.

In this chapter, I am content to point out the restrictions that must be imposed (somewhat arbitrarily) on equilibrium strategies in order to achieve a competitive limit theorem. In the next chapter I attempt to derive such restrictions endogenously from assumptions about the primitives of the game.

1.2 Pure Exchange Economies

We begin by describing a pure exchange economy. In this economy there is no production. Instead, a finite number of economic agents have exogenously given endowments of a finite number of commodities. Because the agents have different endowments and preferences, there are gains from trade. The agents exchange commodities in order to maximize their preferences. To describe a formal model of an exchange economy, we have to specify the list of commodities, the list of agents, and the agents' preferences and endowments.

Commodities. There is a finite number of commodities, indexed $h = 1, ..., \ell$. Each commodity is assumed to be homogeneous and perfectly divisible. The quantity of any commodity h is represented by a real number x_h and we adopt the usual convention that negative numbers represent "supplies" or "deficits". A **commodity bundle** can then be represented by a vector $x = (x_1, ..., x_\ell)$, where x_h represents the quantity of commodity h. The **commodity space** consisting of all possible commodity bundles is represented by the ℓ-dimensional Euclidean vector space $\mathbf{R}^\ell$.

Agents. The economy consists of a finite number of economic **agents**, indexed $i = 1, ..., m$, who can be thought of as "consumers".

Consumption sets. A **consumption bundle** is a commodity bundle that is feasible for an agent. The set of feasible consumption bundles for agent i is called his **consumption set** and is denoted by $X_i \subset \mathbf{R}^\ell$. For example, the set X_i may consist of non-negative commodity bundles that are compatible with subsistence for agent i. Each consumption set X_i is assumed to be non-empty, closed, convex, and bounded below.

Endowments. Each agent has an initial **endowment** of commodities which he wants to exchange for a more preferred bundle. Agent i's endowment is denoted by the consumption bundle $e_i \in X_i$. The assumption that e_i is a consumption bundle implies that agent i can survive without trading.

Preferences. Agent i's preferences over consumption bundles are represented by the **utility function** $u_i : X_i \to \mathbf{R}$, which assigns the real number $u_i(x_i)$ to each feasible consumption bundle x_i. In a pure exchange economy, agents maximize the value of a utility function by exchanging commodities with other agents. The utility function u_i is assumed to be concave, continuous, and increasing[1]. Concavity implies that agents who maximize expected utility are risk averse, as well as having the usual diminishing marginal rate of substitution.

The pure exchange economy is defined by the array $\mathcal{E} = \{(X_i, e_i, u_i)\}_{i=1}^m$.

1.2.1 Pareto and Pairwise Efficiency

The next step is to define two concepts of efficiency and characterize the corresponding sets of efficient allocations.

The allocation of resources in a pure exchange economy is described by a list of the consumption bundles for all the agents in the economy. Formally, an **allocation** is an m-tuple of consumption bundles $x = (x_1, ..., x_m)$ such that $x_i \in X_i$ for each i. An allocation x is **attainable** if aggregate demand equals aggregate supply, that is, the sum of the commodities allocated to the agents is equal to the sum of the endowments:

$$\sum_{i=1}^{m} x_i = \sum_{i=1}^{m} e_i.$$

Let $\hat{X}$ denote the set of attainable allocations. Since each consumption set is closed and bounded below, it is easy to see that the set of attainable allocations is compact.

An attainable allocation x is (strongly) **Pareto efficient** if there does not exist another attainable allocation x' that makes some agents better off and none worse off, that is, $u_i(x_i') \geq u_i(x_i)$ for any agent i and $u_i(x_i') > u_i(x_i)$ for some agent i. Let $P \subset \hat{X}$ denote the set of Pareto-efficient allocations.

[1] I use the phrase "increasing" to mean increasing. Some writers use the phrase "strictly increasing" for this purpose.

The continuity of the utility functions implies that P is closed and hence compact.

An attainable allocation x is **pairwise-efficient** if it is impossible for any two agents i and j to improve their utilities by trading together, that is, there does not exist an attainable allocation x' and a pair of agents (i, j) such that

(i) $x'_k = x_k, \forall k \neq i, j$
(ii) $u_i(x'_i) \geq u_i(x_i), u_j(x'_j) \geq u_j(x'_j)$, with at least one strict inequality.

In addition to the standard assumptions, some non-standard properties will be needed in the sequel. First, it will be assumed that preferences are smooth:

- for any i, u_i can be extended to an open set G_i containing X_i and $u_i : G_i \to \mathbf{R}$ is C^1.

This assumption is not strictly necessary, but it simplifies the analysis by allowing us to use the Kuhn-Tucker theorem to characterize solutions to maximization problems.

The second non-standard property is that indifference surface through the initial endowment does not intersect the boundary of the consumption set X_i:

- for any i,
$$\{x_i \in G_i | u_i(x_i) \geq u_i(e_i)\} \cap \partial X_i = \varnothing, \tag{1.1}$$
where ∂X_i denotes the boundary of X_i.

A familiar example that satisfies this assumption is provided by the Cobb-Douglas utility function when the consumption set $X_i = \mathbf{R}^\ell_+$. The importance of the property (1.1) is that it ensures that an allocation that is pairwise-efficient is also Pareto-efficient, as will be shown below.

Because trade is voluntary, we can restrict attention to allocations that are individually rational, that is, allocations x such that $u_i(x_i) \geq u_i(e_i)$ for every i. Under property (1.1) individually rational consumption bundles belong to the interior of the consumption set, not to the boundary. Then we do not need to worry about the boundary of the consumption set when we characterize Pareto-efficient allocations. Without the boundary, there will be no "corner solutions" to worry about and the first-order conditions that characterize efficiency will hold as equations.

Now we are ready to state and prove the promised result, namely, that pairwise efficiency and Pareto efficiency are equivalent under the maintained assumptions. The argument is simple. Under condition (1.1), Pareto-efficient allocations are characterized by the condition that all agents have the same marginal rates of subsitution between pairs of commodities. Likewise, pairwise-efficient allocations are characterized by the condition that any two agents have the same marginal rates of substitution between pairs of commodities. These conditions are obviously equivalent so Pareto-efficiency and pairwise efficiency are equivalent.

In general, there may exist allocations that are pairwise-efficient but not Pareto-efficient. In other words, if condition (1.1) is not satisfied or if the allocation is not individually rational, there may exist allocations that can be improved on by making commodity transfers among three or more agents, but not by making commodity transfers between a single pair of agents.

Proposition 1 *Under the maintained assumptions, an (attainable) and individually rational allocation x is Pareto-efficient if and only if it is pairwise-efficient.*

Proof. If x is a pairwise-efficient allocation, then it must solve the following maximization problem for each pair of agents (i, j):

$$\begin{array}{l} \max u_i(x'_i) \\ \text{s.t. } x'_i + x'_j \leq x_i + x_j, u_j(x'_j) \geq u_j(x_j), \end{array}$$

where x'_i and x'_j are restricted to the open sets G_i and G_j. Note that we can dispense with the conditions $x'_i \in X_i$ and $x'_j \in X_j$ since they are guaranteed by the conditions $u_i(x'_i) \geq u_i(x_i)$ and $u_j(x'_j) \geq u_j(x_j)$. This is the only place that property (1.1) is used.

An application of the "necessity" part of the Kuhn-Tucker theorem implies that a solution of this problem satisfies the first-order conditions

$$\nabla u_i(x_i) \propto \nabla u_j(x_j), \forall i, j.$$

But according to the "sufficiency" part of the Kuhn-Tucker theorem, the first-order conditions imply that x solves the problem

$$\begin{array}{l} \max u_i(x'_i) \\ \text{s.t. } \sum_{j=1}^m x'_j \leq \sum_{j=1}^m x_j, u_j(x'_j) \geq u_j(x_j), \forall j \neq i, \end{array}$$

where x_i' is restricted to the open set G_i and x_j' is restricted to the open set G_j for every j. In other words, it is impossible to make one agent better off without making some of the others worse off. Since the utility functions are continuous and increasing, this is equivalent to the Pareto efficiency of x.

Obviously, Pareto efficiency implies pairwise efficiency, so the proof of Proposition 1 is complete. ■

In the sequel we study a game in which exchange occurs between *pairs* of agents. This trading process results in a *pairwise*-efficient final allocation. The equivalence of pairwise efficiency and Pareto efficiency immediately implies that the final allocation is Pareto-efficient.

As was noted above, without property (1.1) there may exist allocations that are pairwise-efficient but not Pareto-efficient. A simple example will make this clear. Suppose there are three agents $i = 1, 2, 3$ and three goods $h = 1, 2, 3$. Each agent i has an endowment of one unit of good $h = i$ and none of the goods $h \neq i$. The consumption set is assumed to be $\mathbf{R}^3_+$ for each agent and the utility function of agent i is defined by putting

$$u_i(x_i) = \begin{cases} x_{ii} + 2x_{ii+1} & i = 1, 2 \\ x_{ii} + 2x_{i1} & i = 3. \end{cases}$$

Then the allocation $x = (x_1, x_2, x_3) = (e_1, e_2, e_3)$ is pairwise efficient but not Pareto-efficient. It is not Pareto-efficient because each agent gets a utility of 1 whereas he would get 2 if he gave up his endowment in exchange for one unit of his preferred good. On the other hand, no pair of agents can increase their utility levels because of the lack of *mutual coincidence of wants*. For example, agent 1 would like to give up one unit of good 1 in exchange for one unit of good 2, but agent 2, who has one unit of good 2, does not value good 1. A moment's thought will show that it is crucial for this example that every agent is on the boundary of his consumption set and that the indifference curves intersect the boundary of the consumption set. (The utility functions in this example are not increasing, but the example could easily be adjusted to satisfy this property without changing the result). For a more detailed discussion of the relationship between pairwise, t-wise and Pareto efficiency, see Goldman and Starr (1982).

1.3 Dynamic Matching and Bargaining Games

The model economy in the preceding section describes the agents' environment, but not the restrictions on their behavior. To do that, we need to define an extensive-form game, a set of rules that tells us exactly what actions players are allowed to take and when they are allowed to take them. The rules of the game also tell us exactly what consequences follow from any sequence of actions adopted by the players.

Game theorists take the business of defining a game very seriously. This is not just a matter of mathematical precision. As the Nash Program described in Chapter 0 makes clear, the goal in defining a non-cooperative game is to make the assumptions and economic analysis of the game as clear and complete as possible. One criterion of whether we have achieved the goals of the Nash Program is to ask whether the game is "playable". Are the instructions provided complete and clear enough so that someone who was only given these instructions could, in principle, play the game? For example, would an experimental economist be able to go into a laboratory and have subjects play this game or would he have to invent additional rules and conventions himself before the experiment could proceed?

Without a "playable" game in this sense, it is not clear that the object of study is well defined. Imagine that someone gives you an unfamiliar board game as a gift, but the instructions are missing. Without knowing what you are supposed to do, how can you play the game? You don't know what the game is. The same problem confronts a game theorist when he encounters a loosely specified economic environment. Without limiting the actions available to the players and specifying the consequences of these actions, it is difficult to say how the game should be played and what the outcome should be.

The Nash Program requires a non-cooperative game. The fact that a game is non-cooperative in this technical sense does not rule out cooperative behavior. It simply says that if negotiation, pre-play communication, binding agreements, and other elements of cooperation occur, then they must be included in the definition of the game. The Nash Program is a commitment to being explicit about assumptions and economic analysis of the game, rather than a commitment to a particular kind of game or particular outcomes.

Having adopted these general principles, there remain many different ways of modeling trade among a group of self-interested agents. One model that has proved particularly useful in analyzing decentralized trade is based on pairwise matching and bargaining. Individual agents search for trading opportunities and meet one other individual at a time. When a pair of agents encounter one another (matching), they determine whether and what to trade (bargaining). Bargaining is represented by the "alternating offers" model of Stahl (1972) and Rubinstein (1972) described in Chapter 0.

Even within this framework, there are many choices to be made. How are matches determined? Who meets whom and how often? When agents are matched, how does the bargaining proceed? Are offers made simultaneously or sequentially? How long are offers on the table before they can be withdrawn or amended? When is a match dissolved? Under what circumstances can an agent change his partner? What do agents know about each other and about the previous play of the game?

These details are important because the outcome of the game may depend critically on the details of the modeling. Some people find the sensitivity of the analysis to the details of the modeling disappointing. Since we have limited information about what the "right" specification of the model is, what confidence can we have in the predictions of our theory? Certainly, it would be nice to develop a theory that gives strong, unambiguous predictions under very general conditions; but failing that the best we can do is to understand the reasons for the sensitivity of the results. At the very least, knowing the sensitivity of our results to specific assumptions can teach us to be cautious. Moreover, it provides a starting point for thinking about the kind of data that one might need to make more precise predictions. Finally, although our models are only examples, seeing how the argument is constructed in concrete settings gives us a better understanding of the logic of the argument. This is a useful lesson even if it does not allow us to predict what the world is like from first principles.

With these caveats we are ready to begin the business of describing the game.

Time. The pure exchange model is timeless and static. The dynamic matching and bargaining game takes time to play. We assume that time is divided into a countable number of periods or **dates**, indexed $t = 1, 2, ...$ The process of matching, bargaining, and exchange occurs at these dates. Consumption occurs after the game is finished (each agent consumes his terminal consumption bundle).

Matching. At each date an ordered pair of agents (i_t, j_t) is matched, with

$i_t \neq j_t$. It is assumed that for any pair (i, j) with $i \neq j$ there is an infinite number of dates at which the pair (i, j) is matched.

Bargaining. Suppose that a pair of agents (i, j) has been chosen (matched) at date t. Agent i is called the **proposer** and j is called the **responder**. The proposer chooses a vector z of feasible net trades. The responder can accept the proposal by saying "Yes", or reject it by saying "No". If the proposal is accepted, the two agents exchange the proposed vector of trades. The proposer gets z and the responder gets $-z$. If the proposal is rejected, no trade takes place and all agents begin the next period with the allocation they had at the beginning of the present period.

It cannot be stressed too often that the game described here is special. Other games would give different results. (I will try to give specific warnings where appropriate). The point of the exercise is not to claim that this is the unique "true" model, but rather to give an example of the kinds of assumptions and arguments that are needed to ensure a particular outcome, namely, perfect competition in a pure exchange economy.

One particular limitation should be noted immediately. Only one pair of agents is matched at any date and only these agents have an active role at a given date. The other agents remain passive until the next round. The assumption that there is no simultaneous trade is restrictive, but perhaps not as restrictive as it appears. Suppose that we had modeled "time" as a continuous variable, with matches distributed according to a Poisson process. Then the probability that two pairs were formed at the same time would be zero. Here we work with discrete time and retain the assumption that simultaneous matching has a zero probability. Ruling out simultaneous moves simplifies the game. One can only hope that it does not oversimplify.

All pairs of agents meet infinitely often, for each assignment of roles in the bargaining game. This connectedness assumption is important to ensure that the agents function as a single integrated economy.

Weaker assumptions would also suffice. For example, it would be enough to assume that for every pair of agents (i, j) there is some finite sequence of meetings that connects each pair of agents and that this sequence of meetings occurs infinitely often. To ensure this, we would have to assume that for any i and j there exists a sequence $\{i_k\}_{k=0}^{K}$ with $i_0 = i$ and $i_K = j$ and the ordered pair (i_k, i_{k+1}) meets infinitely often for each $k < K$. The stronger assumption is maintained here for simplicity, but the results would still hold with the weaker connectedness assumption.

1.3.1 The Play of the Game

The play of the game is described by a path. A path is a complete history of the play of the game, a description of everything that happens at every date. Formally, **path** is a sequence$\{a_t\}_{t=1}^{\infty}$ consisting of ordered triples $a_t = (x_t, z_t, r_t)$, where x_t is the allocation that has been reached at the beginning of date t, z_t is the proposal offered by i_t, and r_t is the response made by j_t.

A feasible path must satisfy several conditions. First, x_t must be an attainable allocation at each date t. Secondly, x_t must be consistent with the actions chosen by the agents. Recall that only the proposer and responder are allowed to exchange commodities at date t and the actual vector of commodities exchanged is equal to the accepted proposal (if the proposal is rejected there is no trade). Then a feasible path $a = \{a_t\}$ must satisfy

$$\begin{aligned} x_{it+1} &= x_{it} + z_t \text{ if } i = i_t \text{ and } r_t = \text{“}yes\text{”} \\ x_{it+1} &= x_{it} - z_t \text{ if } i = j_t \text{ and } r_t = \text{“}yes\text{”} \\ x_{it+1} &= x_{it} \text{ otherwise,} \end{aligned}$$

for every date t. Let A denote the set of feasible paths.

At each date the agents observe the proposal and response. The information available at the beginning of date t consists of the path segment $(a_1, ..., a_{t-1})$. Call this path segment the **history** of the game up until date t and denote it by h_t. At each date t, the proposer moves first and then the responder moves. The proposer knows the history h_t; the responder knows the history h_t and the proposal z_t. So the information set at which the proposer moves has the form (h, x) and the information set at which the responder moves has the form (h, x, z).[2] Let H_i denote the information sets at which agent i controls play.

A strategy for agent i is a decision rule that specifies a feasible action for the player at every information set where he is required to move. If agent i is the proposer at an information set (h, x) then he has to choose a net trade z that is feasible for himself and the responder, that is, $x_i + z \in X_i$ and $x_j - z \in X_j$. If agent i is the responder at an information set (h, x, z) he

[2] The allocation x is a function of h and to this extent there is some redundancy in this notation for information sets, but there is some advantage in making x explicit since I refer to it often.

has to choose "yes" or "no". Formally, a **strategy** for agent i is a function f_i defined on the domain H_i with values in $\mathbf{R}^\ell$ or $\{yes, no\}$ as appropriate. A feasible strategy for i must must satisfy

$$\begin{aligned} x_{jt} + f_i(h_t, x_t) &\in X_j \text{ for } j = i_t \\ x_{jt} - f_i(h_t, x_t) &\in X_j \text{ for } j = j_t \end{aligned}$$

at any information set (h_t, x_t) such that $i = i_t$. The set of feasible strategies for i is denoted by F_i and the set of strategy profiles is denoted by $F = \times_{i=1}^m F_i$.

For any strategy profile $f \in F$ we can define a unique path $a^f = \{a_t^f\}$ and a unique outcome $\xi(a^f) = \{\xi_t(a^f)\}$. If h_t^f is the history at date t, this determines the attainable allocation x_t uniquely by putting

$$x_{it} = \xi_t(a^f) = \begin{cases} x_{it-1} + z_{t-1} & \text{if } i = i_{t-1} \text{ and } r_{t-1} = \text{“}yes\text{”} \\ x_{it-1} - z_{t-1} & \text{if } i = j_{t-1} \text{ and } r_{t-1} = \text{“}yes\text{”} \\ x_{it-1} & \text{otherwise.} \end{cases}$$

Then the proposal z_t is given by

$$z_t = f_{i_t}(h_t^f, x_t)$$

and the response r_t is given by

$$r_t = f_{j_t}(h_t^f, x_t, z_t).$$

Thus, for any history h_t^f the strategy profile f allows us to define the history $h_{t+1}^f = (h_t^f, a_t^f)$ and by induction we can define the entire path a^f.

1.3.2 Payoffs

The construction of a path for each strategy profile f is crucial in defining the payoffs of the game. Roughly speaking, an agent's payoff is the utility he gets from his terminal consumption, but since the game goes on forever his terminal consumption may not be well defined. For this reason, we calculate the utility he would get if he consumed his commodity bundle at date t and take as his payoff the 'lim inf' of this sequence of utilities. For any outcome $\{x_t\}_{t=1}^\infty$ we define the payoff to agent i to be the

$$\liminf_{t\to\infty} u_i(x_{it}) = \sup_{T\geq 1}\{\inf\{u_i(x_{it}) | t \geq T\}\}.$$

Note that there is no discounting here and that the use of the lim inf is conservative in the sense that it takes the lowest possible estimate for the limiting value of the agent's utility. In practice, we will be able to show that the sequence of utilities $\{u_i(x_{it})\}$ has a limit and the sequence of random allocations has a limit too.

In the same way, the construction of an outcome $\xi(a^f)$ for each strategy profile allows us to define the individual payoff functions $v_i : F \to \mathbf{R}$ by putting the payoff for agent i equal to

$$v_i(f) = \liminf_{t \to \infty} u_i(\xi_{it}(a^f))$$

for any strategy profile $f \in F$.

So, finally, we have arrived at a normal form game Γ defined by

- a set of players $\{1, ..., m\}$,
- a set of strategy profiles $F = \times_{i=1}^m F_i$, and
- a payoff function $v = (v_1, ..., v_m)$.

1.4 Equilibrium

Having defined a playable game, the next step in the analysis of the game is to define an equilibrium. The central concept in non-cooperative game theory is the Nash equilibrium, in which each player chooses a strategy that is a best response to the strategy profile chosen by his opponents. In this context, a Nash equilibrium is a strategy profile $f^* \in F$ with the property that for each i the strategy f_i^* maximizes the payoff of agent i given the strategies $f_{-i}^* \equiv (f_1^*, ..., f_{i-1}^*, f_{i+1}^*, ..., f_m^*)$:

$$v_i(f^*) \geq v_i(f_{-i}^*, f_i), \forall f_i \in F_i.$$

In dynamic games, the Nash equilibrium concept is often found to be too weak to give interesting results. In particular, it does not restrict the behavior of the agents off the equilibrium path. As an example, consider the following strategy profile:

- suppose that every agent i rejects every offer whenever he is in the position of the responder;

- every agent i proposes the no-trade vector $z = 0$ whenever he is in the position of the proposer.

The profile of strategies defined in this way is a Nash equilibrium. Since the proposer anticipates that any offer will be rejected, it is optimal for him to make only the no-trade offer. Since the only offers received in equilibrium are no-trade offers, it is optimal for the responder to reject them. Thus, each of the strategies is a best response to what the other players are actually doing.

Note, however, that if a proposer were to deviate from the equilibrium strategy and offer a Pareto-improving trade, it would not be optimal for the responder to reject it. This kind of behavior seems unreasonable, but it is compatible with the definition of Nash equilibrium. The problem is that the definition of Nash equilibrium only requires agent i's strategy to be a best response to what the other agents actually do in equilibrium. This does not restrict his behavior at information sets that are not supposed to arise in equilibrium. An equilibrium strategy for agent i may commit him, at some information sets, to take an action that would not be optimal if the information set were reached. This commitment to take a suboptimal action if the information set were reached may be crucial for the equilibrium.

For this reason, we often make use of the stronger requirement of **subgame perfect equilibrium** (SPE) proposed by Selten (1965). SPE is stronger than Nash equilibrium because it requires each player to choose an optimal action in every situation he could conceivably find himself in, and not just those that occur along the equilibrium path. In particular, it rules out empty or non-credible threats, that is, threats which an agent would not be willing to carry out if he found himself in the situation in which he had threatened to take the action. Formally, for any information set $h \in H_i$ let $\Gamma(h)$ denote the subgame that begins at the information set h. For any history h' of $\Gamma(h)$ and any strategy f_i in F_i let $\langle f_i | h \rangle$ denote the strategy defined by putting

$$\langle f_i | h \rangle (h') = f_i(h, h').$$

That is, the strategy $\langle f_i | h \rangle$ tells a player to behave in the new game when he observes the history h' the same as he would have behaved in the original game if he had observed the combined history (h, h'), that is, h followed by h'. Then a SPE of Γ is a strategy profile f^* such that, for any information set $h \in H_i$, $\langle f_i^* | h \rangle$ is a best response to $\langle f_{-i}^* | h \rangle$. Each agent i is required to choose a strategy that satisfies the Nash equilibrium conditions at every possible information set and not just those that are reached along the equilibrium path. In other words, $\langle f^*|h\rangle$ is a Nash equilibrium of $\Gamma(h)$.

Although the concept of SPE restricts the behavior of the agents in ways that may be important, it still leaves open the door for many kinds of behavior as equilibrium phenomena. The reason is well known from the theory of repeated games. As the Folk Theorem for repeated games shows, when players are extremely patient every "individually rational" payoff can be supported as a SPE (Fudenberg and Maskin (1986)). The present game is not a repeated game. A repeated game consists of the repeated play of identical games at successive dates. A DMBG changes over time because the current allocation changes as agents trade. Technically, a DMBG is a *stochastic game* in which the state is the allocation of commodities and the date (recall that the matching process is a function of the date). Since the state changes from period to period, a different game is being played at each date. Nonetheless, the principle of indeterminacy that applies in repeated games can also affect the analysis of this game. An agent's action at date t affects his payoff in two ways. It affects his payoff directly by changing his current commodity bundle and it affects his payoff indirectly by changing the future behavior of his opponents. If a particular action will lead his opponents to adopt a punishment strategy, he can be deterred from choosing that action even if the direct effect on his payoff might be beneficial. In the same way, his opponents' decision to punish him will be motivated by the punishments they anticipate if they fail to punish. In an infinite-horizon game, this never ending sequence of threats supported by counter-threats supported by counter-counter-threats ... and so on, may undermine the restrictiveness of SPE altogether. In other words, many threats may become credible if they are supported by further threats ad infinitum. The result is a large set of outcomes consistent with SPE behavior.

To avoid this indeterminacy, theorists are sometimes led to restrict their attention to an even smaller set of strategies, namely those which are **memoryless** or **history-independent**. If all agents choose memoryless strategies in equilibrium, the scope for supporting an outcome as an equilibrium is considerably reduced, since an agent knows (when he chooses his action) that it will be forgotten when agents have to make their choices in the future. In the present game, we cannot escape from the effects of memory altogether, since the current allocation is a product of the past and must be known by the agents choosing actions in each period. However, restricting memory as much as possible will be seen to have a crucial effect on the analysis of the game. Formally, a **Markov strategy** f is one in which

- for any information sets $(h, x), (h', x) \in H_i$, where i is the proposer, the strategy chooses the same proposal $z = f(h, x) = f(h', x)$;
- for any information sets $(h, x, z), (h', x, z) \in H_i$, where i is the responder, the strategy chooses the same response $r = f(h, x, z) = f(h', x, z)$.

In other words, a proposer's action at the information set (h, x) depends only on the attainable allocation x and the date t, and not the history h. We have to include t as part of the information because the identity of the proposer i and the responder j are functions of t. Similarly, at the information set (h, x, z) the responder's action depends only on the allocation x, the date t and the proposal z, and not the history h.

The concept of memoryless strategies leads to a further refinement of Nash equilibrium. A **Markov perfect equilibrium** (MPE) is a SPE in which each agent chooses a Markov strategy. Note that we are restricting attention to SPE in which agents choose Markov strategies, which is not the same as restricting an agent's choice to Markov strategies. In a MPE, the agent is allowed to choose any feasible strategy, but it turns out that a Markov strategy is optimal in the set of all strategies.

1.5 The Edgeworth Property

In the nineteen-sixties and nineteen-seventies economists studied a type of allocation process known as the Edgeworth Process (Uzawa (1960)). An allocation process is a dynamic process in which the state variable is an attainable allocation that evolves according to a deterministic or stochastic law of motion. Corresponding to the allocation process is a utility process defined by the vector of utilities associated with the current allocation at each point in time. The defining property of the **Edgeworth process** is that the utility process is non-decreasing over time. (See Negishi (1962) for a more detailed description).

The Edgeworth Process is like a gradient process, except that it is characterized by a non-decreasing vector of utilities, rather than by a scalar objective function. Under certain regularity conditions it can be shown that an Edgeworth Process converges to a Pareto-efficient allocation (or a set of Pareto-efficient allocations). One difficulty that arises is that there may exist inefficient allocations at which it is impossible to find utility increasing trades. For example, there may exist Pareto-inefficient allocations that are pairwise-efficient (Madden (1976)). If trade is assumed to take place between pairs of agents and if one of these allocations is reached, the Edgeworth Process may become stuck there and never reach a Pareto-efficient allocation.

The processes generated by the DMBG have a similar property, which I will call the Edgeworth Property. The critical difference between my use of the term here and its use in the earlier literature is that the agents in an Edgeworth Process are assumed to be myopic. They care about the utility of the commodity bundle they currently hold, even though the process is intended to be a *tâtonnement* process in which consumption does not occur until trade stops. In a DMBG, by contrast, agents are far-sighted. They care about the utility of the commodity bundle they receive asymptotically as time goes to infinity. What we can show is that the equilibrium payoff is non-decreasing.

Suppose that f^* is a MPE and that $\{x_t\}$ is the equilibrium outcome. The equilibrium payoff to agent i is given by

$$v_i(f^*) = \lim \inf_{t\to\infty} u_i(x_{it}).$$

At the beginning of date t, before a proposal has been made, the equilibrium payoff is a function of the initial allocation x_t and the date t. Trade is voluntary in the sense that an agent can always guarantee no trade by offering a proposal $z = 0$ whenever he is the proposer and rejecting all offers when he is the responder. At the worst, he can guarantee himself $u_i(x_{it})$, so

$$v_i(f^*) \geq u_i(x_{it})$$

for every agent i and date i. From this it follows immediately that

$$\lim \inf_{t\to\infty} u_i(x_{it}) = v_i(f^*) \geq \lim \sup_{t\to\infty} u_i(x_{it}),$$

which implies that

$$v_i(f^*) = \lim_{t\to\infty} u_i(x_{it}).$$

Proposition 2 *Let f^* be a MPE of Γ and let $\{x(t)\}_{t=1}^{\infty}$ be the equilibrium outcome. Then*

$$v_i(f^*) = \lim_{t\to\infty} u_i(x_{it}), \forall i.$$

Establishing convergence is important. In such a complex game, it is hard to know where to begin to analyze the equilibrium strategies. The fact that payoffs are eventually constant gives us a place to start. As we see in the next section, it implies that gains from trade eventually vanish. This allows us to say something about the limiting allocation and then we can work backwards to deduce properties of the entire equilibrium.

1.6 Efficiency

The essential idea developed in this section is that if agents meet repeatedly and bargain over the gains from pairwise trade, the resulting equilibrium allocation must be Pareto-efficient. Stated in this way, the result may sound almost tautological. One of the things to be learned from this analysis is that the result is not trivial and contains some important subtleties.

The intuitive part of the argument is quite easy. Imagine that in the limit as $t \to \infty$ the allocation converges to x. The equilibrium payoff for i is $v_i(f^*) = u_i(x_i)$. As we have already seen, $v_i(f^*) \geq u_i(e_i)$ so x_i is individually rational. Then if x is not Pareto-efficient, it is not pairwise efficient, according to Proposition 1. This means that two agents, i and j say, can increase their utility through trade with each other. Sooner or later i and j will meet and one of them can propose to the other a Pareto-improving trade. But this means that they can get more than the equilibrium payoffs, contradicting the conditions for a Nash equilibrium.

There are a couple of technical difficulties in the proof of this result. First, the fact that payoffs and utilities converge along a particular sample path does not imply convergence of the corresponding allocations. There may be non-negligible trade in the limit even though utilities are not changing. Secondly, the heuristic argument above applies in the limit, whereas we need to show that a contradiction occurs for a large but finite date t when convergence has not yet occurred.

More important than these technical issues is a strategic issue. Without the assumption of MPE, it does not follow that a responder will accept a Pareto-improving offer. Suppose that i and j meet late in the game, when their commodity bundles are x_i and x_j, respectively, and their equilibrium payoffs are $\bar{v}_i$ and $\bar{v}_j$, respectively. Suppose further that i proposes a trade z that is feasible and will give both i and j utilities that are greater than their current equilibrium payoffs, that is,

$$u_i(x_i + z) > \bar{v}_i, u_j(x_j - z) > \bar{v}_j.$$

Should j accept? Not if j thinks that by rejecting i's offer of $-z$ he will get an even higher payoff in the future. But if i anticipates a rejection from j then i has no incentive to make the proposal in the first place. So it may be an equilibrium for the agents not to trade and to receive the payoffs $\bar{v}_i$ and $\bar{v}_j$ in equilibrium. The crucial assumption in this argument is that an offer by i will be remembered in the remainder of the game and used as a signal to the players to change their equilibrium play so that j is rewarded for rejecting i.

In a MPE, a rejected offer is not "remembered" in the continuation game. Strategies only depend on the current state, so the past is "remembered" only to the extent that it has an impact on the current allocation. A rejected offer has no effect on the allocation and so the players in the continuation game cannot distinguish it from a no-trade offer. Since j cannot benefit from rejecting the offer, he must accept the Pareto-improving offer in the example above.

From a purely strategic perspective, the assumption of MPE may seem arbitrary. From a competitive market perspective, it seems more reasonable. In a large market, it is not likely that agents will observe all offers and there is no incentive for third parties to condition their behavior on rejected offers. On the other hand, it could be argued that there are social norms that operate in markets, for example, the notion of a fair price, and that violators will be punished by society at large. This whole question needs to be investigated at length in the next chapter. For the moment it is enough to appreciate the importance of the Markov assumption in establishing this result.

For the proof of Proposition 3, we need to strengthen the assumptions on preferences. Specifically, we assume that:

- Each utility function u_i is strictly concave.

This assumption ensures that allocations converge if utilities converge. If utilities converge but trade does not converge to zero, there must in the limit exist a non-zero trade between two agents that leaves utility unchanged. But if two agents can exchange a non-zero commodity bundle that leaves utility unchanged, they can increase utility by trading half as much. The existence of such a trade will be shown to contradict the conditions for a MPE.

The next proposition characterizes the limiting set of allocations on the equilibrium path. It shows first that the allocation converges and secondly that the limiting allocation is Pareto-efficient.

Proposition 3 *The MPE outcome* $\{x_t\}$ *converges to* $x_\infty = \lim_{t\to\infty} x_t \in P$.

Proof. The first step in the proof is to show that $\{x_t\}$ is a Cauchy sequence. The proof is by contradiction. Suppose, contrary to what we want to prove, that for some subsequence (using the same notation) $||x_t - x_{t+1}|| \geq \varepsilon$ for some $\varepsilon > 0$ and all t. Choose some further subsequence such that $i_t = i$ and $j_t = j$ for all t. It is always possible to choose this subsequence because every ordered pair of agents (i, j) is matched infinitely often. Then choose a further subsequence such that $x_t \to y$ and $x_{it+1} - x_{it} \to z$ along this subsequence. It is possible to choose such a subsequence because the set of attainable allocations is compact. Since

$$\lim_{t\to\infty} u_k(x_{kt}) = \lim_{t\to\infty} u_k(x_{kt+1}) = v_k(f^*),$$

for $k = i, j$, it must be the case that

$$\begin{aligned} u_i(y_i) &= u_i(y_i + z) = v_i(f^*) \\ u_j(y_j) &= u_j(y_j - z) = v_j(f^*). \end{aligned}$$

By strict concavity,

$$\begin{aligned} u_i(y_i + z/2) &> v_i(f^*) \\ u_j(y_j - z/2) &> v_j(f^*). \end{aligned}$$

By continuity,

$$\begin{aligned} u_i(x_{it} + z/2) &> v_i(f^*) \\ u_j(x_{jt} - z/2) &> v_j(f^*), \end{aligned}$$

for all t sufficiently large. This contradicts the equilibrium conditions since i can offer j to trade $z/2$ and make them both better off.

This last step is not entirely obvious, since we are considering a deviation from the equilibrium path. We have shown that i can offer j a trade which can make them both better off than they would be along the equilibrium path. However, j may reject this offer if, by doing so, he can achieve an outcome that is better still. Can j do better still? In a Markov equilibrium, the answer is no. Since strategies are memoryless, if j rejects i's offer the allocation at the beginning of the next period is exactly the same as it was at the beginning of the present period. Since f^* is a MPE, the equilibrium payoff depends only on the next state $(x_t, t+1)$ and can be denoted by $v_j(x_t, t+1)$. On the one hand, since trade is voluntary, $v_j(x_t, t+1)$ must satisfy

$$v_j(x_t, t+1) \leq v_j(f^*);$$

otherwise j could do better than his equilibrium payoff by refusing any offer at date t and this would contradict the equilibrium conditions. On the other hand, if he accepts the offer of z, his payoff must be at least

$$v_j(x', t+1) \geq u_j(x_{jt} - z/2) > v_j(f^*),$$

where the new allocation x' is defined by

$$x'_k = \begin{cases} x_{it} + z/2 & \text{if } k = i \\ x_{jt} - z/2 & \text{if } k = j \\ x_{kt} & \text{otherwise.} \end{cases}$$

So in a MPE agent j must accept the offer of z. Then i can achieve a payoff of at least $u_i(x_{it} + z/2) > v_i(f^*)$ and must deviate.

Note that in order to show that j must accept i's offer it is only necessary to look ahead one period. This is because we are assuming that the strategies f^* constitute an equilibrium. So if j rejects the offer, the resulting payoff is determined by the equilibrium strategies beginning with the current allocation at the next date. Since this payoff is less than he was offered, j's equilibrium strategy must be to accept i's offer.

So the hypothesis that $\{x_t\}$ is not Cauchy leads to a contradiction of the equilibrium conditions. This contradiction establishes that $\{x_t\}$ is a Cauchy sequence, so $\{x_t\}$ converges to a limit allocation x_∞.

To show that $x_\infty \in P$ we use a similar argument. Suppose that $x_\infty \notin P$. Then x_∞ is not pairwise-efficient, according to Proposition 1, and for some pair (i, j), there exists a feasible trade z such that

$$\begin{aligned} u_i(x_{i\infty} + z) &> u_i(x_{i\infty}) = v_i(f^*) \\ u_j(x_{j\infty} - z) &> u_j(x_{j\infty}) = v_j(f^*) \end{aligned}$$

and by continuity

$$\begin{aligned} u_i(x_{it} + z) &> v_i(f^*) \\ u_j(x_{jt} - z) &> v_j(f^*), \end{aligned}$$

for all t sufficiently large. Since the pair (i, j) is matched infinitely often on ω, we can use the previous argument show that i must deviate by offering z to j, contradicting the equilibrium conditions. This shows that $x_\infty \in P$. ■

1.7 Competitive Sequences of Economies

To get any further with our story, we have to allow the number of agents to grow unboundedly large. There are several ways to do this. One simple way is to assume there is a fixed sequence of agents $i = 1, 2, ...$, each characterized by a consumption set X_i, an endowment $e_i \in X_i$ and a utility function $u_i : X_i \to \mathbf{R}$. For each positive integer m, define a pure exchange economy $\mathcal{E}^m$ consisting of the first m agents. Using this economy and a matching probability π^m, one can define a matching and bargaining game Γ^m consisting of the players $I^m = \{1, ..., m\}$, the strategy sets $F^m = \times_{i=1}^m F_i$ and the payoff functions $\{v_i^m\}_{i=1}^m$ defined in the usual way. The set of Markov perfect equilibria corresponding to the game Γ^m is denoted by $MPE(\Gamma^m)$.

The objective of this section is to study the behavior of a sequence of equilibria $\{f^m\}_{m=1}^\infty$, where $f^m \in MPE(\Gamma^m)$ for each m.

For each equilibrium f^m, let $x^m = \{x_t^m\}_{t=1}^\infty$ denote the sequence of attainable allocations observed along the equilibrium path, and let

$$y^m = \lim_{t \to \infty} x_t^m.$$

In other words, y^m is the limit allocation. Since $y^m \in P$ the Second Welfare Theorem and the usual concavity assumptions imply that there exists a supporting price vector p^m such that

$$u_i(y_i) \geq u_i(y_i^m) \;\&\; y_i \neq y_i^m \Longrightarrow p^m \cdot y_i > p^m \cdot y_i^m,$$

for every $i = 1, ..., m$. To see this, note that efficiency implies the existence of a common vector p^m proportional to each gradient vector $\nabla u_i(y_i^m)$; strict concavity and the gradient inequality imply the rest.

We can normalize the prices so that $\|p^m\| = 1$ for every m. Since the sequence $\{p^m\}$ is bounded, it possesses a subsequence converging to a limiting price vector p^0. Denote the convergent subsequence by $m \in \mathcal{M}$. This subsequence is the focus of attention in what follows.

Because of the complexity of the sequence of games and corresponding equilibria, two additional assumptions are required in order to characterize equilibria in the limit. One concerns the continuity of the equilibrium play of the game; the other concerns the curvature of the individual utility functions.

1.7.1 Continuity

As the number of agents gets very large, a single individual becomes negligible in terms of his endowment and his potential contribution to the general welfare. However, this does not ensure that his effect on the limit allocation will be negligible, even when the number of agents gets very large. In a dynamic game, the possibility of a large number of agents conditioning their responses on the actions of a single player endogenously generates the possibility that a single player has a non-negligible effect in equilibrium. To avoid this possibility, a continuity assumption will imposed. The essential idea is simple. The continuity assumption requires that a single player have a negligible impact on the limiting allocation.

Before describing the continuity assumption it is worth making several general points about the nature and role of the assumption. First, it is an assumption about endogenous variables. As such it directly restricts the kinds of equilibria that will be considered, rather than the primitives of the model. Such assumptions are often regarded with suspicion because it is not at all clear what is being ruled out. Secondly, it is an assumption about strategies, insofar as it restricts the response of players to deviations from equilibrium play. Thirdly, it is a complex assumption. The situation to which the continuity assumption applies involves two limits, one as time goes to infinity and the other as the number of players grows without bound. Uniform continuity is needed with respect to both limits.

For any strategy profile $f \in F^m$ there is a unique sequence of random allocations $\xi(a^f) = \{\xi_t(a^f)\}_{t=1}^{\infty}$, where $\xi_t(a^f)$ is the (random) allocation at date t. Obviously, the allocation at date t is defined on Ω and is $\mathcal{F}_t$-measurable. Suppose that agent j deviates from the equilibrium by choosing an arbitrary strategy $\hat{f}_j \in F_j$. The resulting path is denoted by $a(f_{-j}, \hat{f}_j)$ and the outcome is $\xi(a(f_{-j}, \hat{f}_j))$. The impact of the deviation on the outcome of the game can be measured by the distance of the current allocation at any date from the equilibrium limit allocation. If y is the equilibrium limit allocation and the current allocation at date t is $\xi_t(a(f_{-j}, \hat{f}_j))$ then the distance between the current and the limit allocations is

$$\frac{1}{m}\sum_{i=1}^{m}\|\xi_{it}(a(f_{-j}, \hat{f}_j)) - y_i\|.$$

Note that we take the average distance. What we need to ensure is that a deviation by a single agent j does not cause a large change in the outcome as measured by this distance.

- Let $\{f^m\}$ be a fixed but arbitrary sequence of MPE. The **continuity assumption** requires that for any $\epsilon > 0$ there exist M and T such that for any agent j and any strategy $f_j \in F_j^m$,

  $$\frac{1}{m}\sum_{i=1}^{m}\|\xi_{it}(a(f_{-j}^m, f_j)) - y_i^m\| < \epsilon, \forall m > M, \forall t > T,$$

  where it is understood that $m \in \mathcal{M}$.

In other words, any deviation by a single player cannot change the average allocations by a large amount for sufficiently large values of t and m. This does not imply that the deviating player cannot change the consumption bundles of some players, himself included, by a non-negligible amount. What it does imply is that he cannot change the consumption bundles of most players by a non-negligible amount. Also, note that we have bounded this distance *uniformly* in t and m. Uniformity is necessary; otherwise the order of limits would matter.

An assumption of this kind, which restricts the values of endogenous variables, is unfortunately obscure. We do not know what this assumption is ruling out or ruling in in terms of the primitives of the model. We do not know whether the assumption is likely to be satisfied by most MPE or not. For the moment, we have to accept it on the basis of intuitive plausibility and the fact that something like this is needed to make the analysis tractable and to get the result we want. In the following chapter we investigate the basis for this kind of assumption in more detail.

A similar assumption has been used by Green (1984) to show that the Nash equilibria of an n-person game converge to Nash equilibria of a continuum game as the value of n increases without bound.

An assumption of this kind appears to be needed in order to show that the Nash equilibria of a Cournot oligopoly game converge to competitive equilibria as the number of firms increases without bound. Roberts (1980) provides a counter-example to the competitive limit result. If the equilibrium price correspondence, which maps quantities chosen by firms into equilibrium prices, is multi-valued then a small change in quantity by one firm can lead to a discontinuous change in price. This discontinuous effect of a single firm's action can prevent the achievement of perfect competition in the limit, even though each firm is "small" in terms of its impact on total output.

1.7.2 Curvature

The second assumption requires uniformity in the curvature of the agents' preferences. As we let the number of players grow without bound, the degree of substitutability along an agent's indifference curves may be falling as the index of the agent gets very large. The assumption introduced in this section puts a bound on how little substitutability there can be. This assumption is not strictly necessary, for reasons indicated below in Section 1.7.4, but it makes life a lot easier.

The first part of the assumption guarantees that the indifference surface has positive curvature, that is, the indifference surface is not kinked. This is guaranteed by the assumption that, for any consumption bundle x_i, there is a ball with positive radius contained in the set of bundles preferred to x_i and tangent to the indifference surface at x_i. The larger this ball is, the flatter the indifference surface must be. For present purposes, it is also necessary to ensure some degree of uniformity in the curvature of the indifference surfaces. This can be done by assuming that for every agent i and every bundle x_i, the ball can be chosen to have a fixed minimum radius $r > 0$. The rest of this section deals with details that can be skipped by the non-technical reader.

For any agent i and consumption bundle $x_i \in X_i$, let

$$H_i(x_i) \equiv \{x_i' \in X_i | u_i(x_i') \geq u_i(x_i)\}$$

denote the set of consumption bundles that are at least as good as x_i. Define the normal function $g_i : X_i \to \mathbf{R}^\ell$ by putting

$$g_i(x_i) = -\nabla u_i(x_i)/\|\nabla u_i(x_i)\|.$$

Then we note that for any $\alpha > 0$, the point $x_i - \alpha g_i(x_i)$ belongs to $H_i(x_i)$ and lies at a distance α from the point x_i. Let

$$B(x, r) = \{y \in \mathbf{R}^\ell | \ \|x - y\| \leq r\}$$

denote the ball with center x and radius r.

- Our **curvature assumption** is that for some $\alpha > 0$, for any i and for any $x_i \in X_i$,
$$B(x_i - \alpha g_i(x_i), \alpha) \subseteq H_i(x_i),$$
that is, every point in the ball with radius α and center $x_i - \alpha g_i(x_i)$ is weakly preferred to x_i.

As before, we could restrict this assumption to hold only for individually rational bundles x_i such that $u_i(x_i) \geq u_i(e_i)$.

Note that the strength of this assumption comes from the fact that α is chosen independently of i and the consumption bundle x_i. This means that there is a uniform bound on the curvature of the indifference curves of every agent at every point in his consumption set. As usual, we only need this assumption to apply to the individually rational bundles for each agent.

1.7.3 Linearity

The preceding sections have introduced two special assumptions that are needed to continue the development of a competitive limit theorem. The next step is to show that budget balance is satisfied. More precisely, it will be proved that any net trade lying in the competitive budget set $\{z \in \mathbf{R}^\ell | p^0 \cdot z \leq 0\}$ is achievable in an asymptotically pure MPE for m sufficiently large. In other words, we shall have demonstrated that each agent can achieve the same utility that he would obtain in a competitive equilibrium with price vector p^0.

The strategy of the proof, which is an adaptation of the proof used in Gale (1986), consists of a number of steps. First, we note that since the limiting allocation is efficient and has a supporting price vector, for any net trade z such that $p^0 \cdot z < 0$ we can find a number n such that $-z/n$ is a preferred trade for any agent j holding the limiting bundle y_j^m. This follows from the curvature assumption. Then we claim that the following strategy must allow agent i to achieve the net trade z. Simply offer to trade z/n whenever the opportunity arises and refuse all other trades. Repeat this pattern until the trade z/n has been executed n times. Any agent j must be willing to trade $-z/n$ for t sufficiently large because as t gets large the agent's current bundle converges to y_j^m so $-z/n$ is a preferred trade. Since agent i has an infinite number of opportunities to make these trades, with probability one, he will surely achieve the net trade z.

A number of assumptions are needed to make this argument go through, however. We have already mentioned the curvature assumption. Another is the continuity assumption. It is crucial that the agent's deviation does not change the limit allocation, by very much, for most players. If it did, the supporting prices p^0 and the set of preferred trades z would be altered, so we could not be sure that anyone would accept $-z/n$. Finally, the Markov assumption is important, as it was in the proof of efficiency, because it ensures that if an agent j refuses to trade $-z/n$ with agent i he is not going to be rewarded for that refusal in the continuation game.

Lemma 4 *Let $\{f^m\}$ be a competitive sequence of asymptotically pure MPE satisfying the maintained assumptions. For any i and z such that $e_i + z \in X_i$ and $p^0 \cdot z < 0$, $v_i^m(f^m) \geq u_i(e_i + z)$ for all $m \in \mathcal{M}$ sufficiently large.*

Proof. Fix i and choose z such that $p^0 \cdot z < 0$. Choose M so that $p^m \cdot z < 0$ for all $m > M$ and then choose $0 < \lambda \leq 1$ and $\epsilon > 0$ so that for all $j \leq m$ and $m > M$, the ball with radius ϵ and center $y_j^m - \lambda z$ is contained in the ball with radius α and center $y_j^m + \alpha p^m$ and hence contained in $H_i(y_j^m)$. The possibility of this construction is guaranteed by the curvature assumption.

Now consider the following strategy f_i for some fixed but arbitrary m. Choose a number $n > 1/\lambda$ and have agent i offer to trade z/n at every opportunity until he has made the aggregate trade z. In other words, put

$$f_i(h) = \begin{cases} z/n & \text{if } i \text{ is the proposer and he has traded } kz/n \text{ for } k < n; \\ 0 & \text{if } i \text{ is the proposer and he has traded something other} \\ & \text{than } kz/n \text{ for } k < n; \\ \text{“No”} & \text{if } i \text{ is the responder.} \end{cases}$$

The exact pattern of trades depends on the order of play, but agent i will meet every agent $j \neq i$ an infinite number of times. The continuity assumption implies that there exist numbers M and T such that for all $m > M$ and $t > T$,

$$\frac{1}{m} \sum_{j=1}^{m} \|\xi_{jt}(f_{-i}^m, f_i) - y_j^m)\| < \epsilon/2.$$

This implies that for at least half the agents,

$$\|\xi_{jt}(f_{-i}^m, f_i) - y_j^m)\| < \epsilon. \tag{1.2}$$

Let J denote the set of agents satisfying the inequality (1.2). Then for all j belonging to this set and for $m > M$ and $t > T$,

$$u_j(\xi_{jt}(f_{-i}^m, f_i) - z/n) > u_j(\xi_{jt}(f^m)).$$

This follows from the fact that $\xi_{jt}(f_{-i}^m, f_i) - z/n$ belongs to the ball with center $y_j^m - z/n$ and radius ϵ and $u_j(\xi_{jt}(f^m))$ converges to $u_j(y_j^m)$ as $t \to \infty$.

Then the Markov property implies that an agent $j \in J$ must accept an offer of z/n for $t > T$ sufficiently large if he has not already done so.

Thus, agent i can in the limit carry out the trade z for almost every realization ω and so $\xi_{it}(f_{-i}^m, f_i) \to e_i + z$ as $t \to \infty$. ■

Lemma 4 shows that, in the limit, the outcome of a MPE of the game maximizes each agent's utility subject to a linear budget constraint. This brings us to the conclusion of our argument that, in the limit as the number of agents gets very large, the dynamic matching and bargaining game achieves a competitive equilibrium allocation. The allocations are always attainable and Lemma 4 shows that the terminal consumption bundle of each agent gives him a utility equal to the competitive equilibrium indirect utility $v_i(p^0, p^0 \cdot e_i)$ at prices p^0 with the wealth $p^0 \cdot e_i$.

1.7.4 A Counter-Example with Limited Substitution

It is now clear why uniform curvature was needed for the preceding argument. We want to show that an agent can trade any vector z such that $p \cdot z < 0$ by breaking z up into n pieces z/n and trading each of these pieces with a different agent. If the agents do not have uniform curvature, there may not be n of them for whom $-z/n$ is a preferred trade. Of course, there could be other ways of achieving the net trade z. One could find a sequence $\{\rho_i\}$ of positive numbers summing to one and make a sequence of trades $\{\rho_i z\}$. If the numbers $\rho_i \to 0$ were chosen appropriately it might be possible to make the desired trade without assuming uniform curvature. Clearly, the agents who have limited substitutability would be offered the smallest trades. However, even if this approach were adopted, some bound on substitutability would be needed to ensure that the trade z could be achieved. To see how lack of uniform curvature can prevent convergence to a competitive equilibrium allocation, consider what would happen if most agents' preferences had no curvature at all.

Suppose that there are two commodities $\ell = 2$ and two agents $m = 2$. Agent 1 has Cobb-Douglas preferences

$$u_1(x_1) = \log_e x_{11} + \log_e x_{12}$$

and agent 2 has Leontief preferences

$$u_2(x_2) = \min\{x_{21}, x_{22}\}.$$

Suppose that the initial endowment is given by $e_1 = (3, 2)$ and $e_2 = (2, 3)$. There is a unique competitive equilibrium in which each agent consumes $(2.5, 2.5)$ but this is not the only limiting equilibrium allocation of the DMBG. For example, the allocation in which agent 1 gets $(2.5 - \varepsilon, 2.5 - \varepsilon)$ and agent 2 gets $(2.5 + \varepsilon, 2.5 + \varepsilon)$, for sufficiently small $\varepsilon > 0$, can be supported as an equilibrium of the DMBG, as is shown in Section 1.11.

It is no surprise that competition fails with two agents, but the same kind of allocation can be achieved as a MPE of a DMBG with any number of agents. Suppose that, in addition to agents 1 and 2 there are a large number of additional agents with Leontief preferences and endowments $(1, 1)$. The same bundles for agents 1 and 2 are still attainable. Furthermore, the presence of the additional agents does not affect the equilibrium play of the game. There is no possibility of trade between agent 1 and agents $i > 2$, so they might as well not be there.

By contrast, if the additional agents had Cobb-Douglas preferences identical to agent 1's and endowments equal to $(1, 1)$ then the argument of Lemma 4 would apply. By making small offers to trade a small vector $z/n = (-0.5, 0.5)/n$ to n different agents, agent 1 could achieve his competitive trade z.

1.7.5 The Competitive Limit Theorem

The preceding analysis can be summed up in a **competitive limit theorem**. Take an infinite sequence of agents $i = 1, 2, ...$ and define the exchange economy $\mathcal{E}^m$ to comprise the first m agents $i = 1, 2, ..., m$, for every finite number m. Then define the game Γ^m by specifying an order of play $\pi^m = \{(i_t^m, j_t^m)\}$. The sequence $\{\Gamma^m\}$ is called a **competitive sequence of games** if the following assumptions are satisfied for every i:

- $X_i \subset \mathbf{R}^\ell$ is non-empty, and $X_i \equiv \{x_i \in G_i | u_i(x_i) \geq c_i\}$ for some constant c_i (hence X_i is closed);
- $e_i \in X_i$;
- $u_i : X_i \to \mathbf{R}$ is strictly concave, increasing, and C^1 on an open superset of X_i;

and the sequence satisfies

- the curvature assumption introduced above in section 1.7.2;
- the sets $\{X_i\}$ are uniformly bounded below and the mean endowments $m^{-1} \sum_{i=1}^m e_i$ are uniformly bounded above.

The last assumption is the only new one. It guarantees that the resources available to agents do not grow on average as the size of the economy expands without limit.

Now let $\{f^m\}$ be a sequence of equilibrium strategy profiles such that $f^m \in MPE(\Gamma^m)$ for each m. A subsequence $m \in \mathcal{M}$ is called a **competitive sequence of equilibria** if the following properties are satisfied:

- for every m, f^m is a MPE and $x_t^m \to y^m$ for almost every ω;
- for every $m \in \mathcal{M}$ there exists a price vector p^m such that $\|p^m\| = 1$ and $u_i(x_i) > u_i(y_i^m)$ implies $p^m \cdot x_i > p^m \cdot y_i^m$, for $i = 1, ..., m$;
- $\lim_{m \in \mathcal{M}} y_i^m = y_i^0$ and $\lim_{m \in \mathcal{M}} p^m = p^0$;

and

- the subsequence satisfies the continuity assumption introduced above in section 1.7.1.

Call (y^0, p^0) a **competitive limit equilibrium** if, for every i,

$$y_i^0 \in \arg\max\{u_i(x_i)|\ x_i \in X_i, p^0 \cdot x_i \leq p^0 \cdot e_i\}.$$

In other words, at the limiting allocation y^0 each agent is maximizing his utility subject to a budget constraint defined by the limiting price system p^0. This property is guaranteed by Lemma 4.

Although the limit allocation (with respect to time) $y^m = (y_1^m, ..., y_m^m)$ is attainable in the sense that $\sum_{i=1}^m (y_i^m - e_i) = 0$, there is no guarantee that the limit allocation (with respect to numbers of agents) $y^0 = (y_1^0, y_2^0, ...)$ will be **attainable** in the sense that

$$m^{-1} \sum_{i=1}^{m} (y_i^0 - e_i) \to 0. \tag{1.3}$$

The reason y^0 may not be attainable is that the convergence of $\{y_i^m\}_{m=1}^\infty$ to y_i^0 may not be uniform with respect to i. Of course, if we *assume* uniform convergence then we can be sure that the limiting allocation y^0 will be attainable. For any $\varepsilon > 0$ and all m sufficiently large, uniform convergence implies that $||y_i^0 - y_i^m|| \leq \varepsilon$ for $i = 1, ..., m$. From this it easily follows that

$$\begin{aligned} m^{-1} || \sum_{i=1}^{m} (y_i^0 - e_i) || &= m^{-1} || \sum_{i=1}^{m} (y_i^0 - y_i^m + y_i^m - e_i) || \\ &= m^{-1} || \sum_{i=1}^{m} y_i^0 - y_i^m || \\ &\leq m^{-1} \sum_{i=1}^{m} || y_i^0 - y_i^m || \leq \varepsilon \end{aligned}$$

for all m sufficiently large. Thus,

$$\lim_{m \to \infty} m^{-1} || \sum_{i=1}^{m} (y_i^0 - e_i) || \leq \varepsilon.$$

Since the choice of ε was arbitrary, the limit is in fact zero.

Call (y^0, p^0) a **strong competitive limit equilibrium** if it is a competitive limit equilibrium and y^0 is attainable in the sense of condition (1.3).

Theorem 5 *Let $\{\Gamma^m\}$ be a competitive sequence of games, let $\{f^m\}_{m \in \mathcal{M}}$ be a corresponding competitive sequence of equilibria and let y^0 and p^0 be the limiting allocation and price system. Under the maintained assumptions (y^0, p^0) is a competitive limit equilibrium and if the convergence of $\{y_i^m\}$ to y^0 is uniform then (y^0, p^0) is a strong competitive equilibrium.*

This brings us to the end of our first attempt to characterize the MPE of the dynamic matching and bargaining game. As we have seen, some strong assumptions are needed to carry through the program of providing strategic foundations for the theory of competitive equilibrium. These will be investigated in more detail in the next chapter. In the remainder of this chapter, we explore some other aspects of the theory, beginning with existence.

1.8 Existence

So far, we have focused on the characterization of MPE without worrying to much about the restrictiveness of the Markov property. Do MPE even exist? A complete characterization of existence is not possible in general, but in a special case we can say quite a lot. It turns out that MPE do exist for a rich class of economic environments. Furthermore, they have a nice structure that makes them natural objects to study and allows us to pursue a constructive approach to existence.

The strategy for proving that a MPE equilibrium exists is quite simple. In the theory of repeated games and DMBGs with a continuum of agents, e.g., Gale (1986), one constructs an equilibrium by guessing the equilibrium outcome and then proving that it can be supported as a MPE by appropriate strategies. Here too the proof is constructive. I assume that the outcome of the MPE converges to a competitive equilibrium allocation for the economy. The first step is to find trading rules that will allow agents to trade commodities to get from their initial endowments to their competitive equilibrium consumption bundles while keeping the value of their current bundles constant. These trading rules can be used as the basis of the equilibrium strategies: whatever allocation has been reached, the agents take as their target a competitive equilibrium relative to that allocation and use the trading rules to trade towards their target. The difficult step is to show that these strategies constitute a MPE, that they prevent anyone from doing better than his equilibrium payoff. In fact this can only be shown in special cases, so it is necessary to make some special assumptions. This should not be surprising, given the fact that in a finite economy each of the agents has some 'market power'. In general, we should expect the competitive equilibrium to be the outcome of a MPE when the number of agents is unboundedly large.

Let $\mathcal{E} = \{(X_i, u_i, e_i)\}_{i=1}^m$ be a fixed but arbitrary economy and for any attainable allocation $x \in \hat{X}$ let $\mathcal{E}(x)$ denote the exchange economy formed by taking x as the initial endowment, that is, $\mathcal{E}(x) = \{(X_i, u_i, x_i)\}_{i=1}^m$. Let $W(\mathcal{E}(x))$ denote the set of Walras allocations of $\mathcal{E}(x)$, that is, the set of allocations x^* such that for some price system p^* the ordered pair (p^*, x^*) is a competitive equilibrium of $\mathcal{E}(x)$. Under certain conditions, it can be shown that there exists a MPE of Γ that implements the Walras allocations of $\mathcal{E}$. This ensures that the discussion of the preceding sections has not been vacuous. More interestingly, it shows that the competitive equilibrium can be implemented by a MPE even for finite games. This conclusion would not be too surprising if we considered SPE rather than MPE, because SPE allow for the use of trigger strategies. For this reason, the set of SPE may be much larger than the set of MPE. In any case, the proof of existence of MPE is not without difficulty.

Under the maintained assumptions, every efficient allocation $x \in P$ has associated with it an essentially unique vector of supporting prices. Let $\Delta = \{p \in \mathbf{R}_+^\ell | \sum_{h=1}^\ell p_h = 1\}$ denote the $(\ell - 1)$-dimensional simplex. For any Pareto-efficient allocation $x \in P$, let $\pi(x)$ denote the unique price vector in Δ such that

$$p \propto \frac{\partial u_i(x_i)}{\partial x_i}, \forall i = 1, ..., m.$$

By Proposition 1 we know that the gradients of the utility functions are proportional. So the price vector exists and is positive under the maintained assumptions and the normalization ensures that it is unique.

For any Pareto-efficient allocation x, let $H(x)$ denote the hyperplane through x defined by

$$H(x) = \{x' \in \mathbf{R}^{\ell m} | \pi(x) \cdot x_i' = \pi(x) \cdot x_i, i = 1, ..., m\}$$

and let $B(x)$ denote the set of attainable allocations in $H(x)$, that is, $B(x) = H(x) \cap X$. The importance of the set $B(x)$ is that starting from an initial endowment in $B(x)$, the allocation x is a Walras allocation.

Proposition 6 *Suppose that x is a Pareto-efficient allocation and $x' \in B(x)$. Then x is a Walras allocation of $\mathcal{E}(x')$, that is, $x \in W(\mathcal{E}(x'))$.*

Proof. By construction and the gradient inequality, for every agent i, $u_i(x_i'') > u_i(x_i)$ implies that $\pi(x) \cdot x_i'' > \pi(x) \cdot x_i = \pi(x) \cdot x_i'$. Thus,

$$x_i \in \arg\max\{u_i(x_i'') | x_i'' \in X_i, \pi(x) \cdot x_i'' \leq \pi(x) \cdot x_i'\}, \forall i = 1, ..., m,$$

and $(\pi(x), x)$ is a competitive equilibrium for $\mathcal{E}(x')$. ■

The second fact that we need is that the sets $\{B(x)|x \in P\}$ form a partition of the set of attainable allocations, that is, each attainable allocation belongs to one and only one set in $\{B(x)|x \in P\}$. To ensure this we need to assume that the competitive equilibrium allocation of $\mathcal{E}(x)$ is unique for every attainable x.

Uniqueness Assumption: For any attainable allocation x, the set of Walras allocations $W(\mathcal{E}(x))$ is a singleton.

There are well known conditions such as gross substitutability that can guarantee uniqueness. There is no need to pursue the details here.

Now note that if

$$x'' \in B(x) \cap B(x')$$

for two Pareto-efficient allocations $x \neq x'$, then x and x' are both Walras allocations for the exchange economy $\mathcal{E}(x'')$, contradicting our assumption. Thus, uniqueness of equilibrium implies that the sets $\{B(x) : x \in P\}$ are non-intersecting. Secondly, any attainable allocation x' must belong to $B(x)$ for some Pareto-efficient allocation x. This is because under the maintained assumptions, each exchange economy $\mathcal{E}(x')$ has a Walras allocation x and every Walras allocation is Pareto-efficient. Thus, we have the following result.

Proposition 7 *The set of attainable allocations $\hat{X}$ is partitioned by $\{B(x) : x \in P\}$.*

With these two propositions established, we are ready to define an equilibrium strategy for Γ. Define the function $\phi : \hat{X} \to P$ by putting

$$\phi(x) = B^{-1}(x)$$

for any attainable allocation x. In the equilibrium we are going to construct, $\phi(x)$ will be the target allocation once the current allocation x has been achieved. In other words, for any attainable allocation x that is reached during the play of the game, the MPE will implement a Walras allocation $\phi(x) \in W(\mathcal{E}(x))$ in the continuation game. To see exactly how this will be done, it is necessary to specify a **trading rule**. Intuitively, when agents i and j meet, the trading rule maximizes a measure of their trade subject to a number of constraints.

Let x be the current allocation and (i,j) be the matched pair during the current period. Let $z=\zeta(x,i,j)$ denote the proposal made by agent i, where $\zeta(x,i,j)$ is defined by putting

$$\begin{aligned}\zeta(x,i,j) \quad = \quad & \arg\max \sum_{h=2}^{\ell} -\exp\{z_h\} \\ & \text{s.t. } \; |z_h| \leq |\phi_{kh}(x)-x_{kh}|, k=i,j; h=2,...,\ell; \\ & \quad sign\{z_h\} = sign\{\phi_{ih}(x)-x_{ih}\} = -sign\{\phi_{jh}(x)-x_{jh}\}, \\ & \quad \text{for } k=i,j,h=2,...,\ell; \\ & \quad \pi(x)\cdot z = 0; \\ & \quad x_i+z\in X_i, x_j-z\in X_j.\end{aligned}$$

Since $\phi(x)$ is a Walras allocation relative to the initial endowment x, we can think of $\phi(x)-x$ as the vector of agents' excess demands. The first constraint says that the absolute value of the trade in commodities $h=2,...,\ell$ cannot exceed the absolute value of the excess demands of agents i and j for the corresponding commodity. The second constraint says that the trade in commodities $h=2,...,\ell$ must have the same sign as the agents' excess demands. The final constraint says that the trade must respect the agents' budget constraints at the equilibrium prices $\pi(x)$. Note the special role of commodity 1. It acts as a means of payment and agents may be forced to trade in a direction that increases or changes the sign of their excess demands in order to balance the budget constraint.

Now the trading rule ϕ can be used to define an equilibrium strategy profile $f=(f_1,...,f_m)$ as follows. For any attainable allocation x and ordered pair (i,j), put

$$f_i(x,i,j)=\zeta(x,i,j) \tag{1.4}$$

and for any net trade z put

$$f_j(x,i,j,z)=\begin{cases}\text{“yes”} & \text{if } \pi(x)\cdot z\leq 0, x_j-z\in X_j, \\ \text{“no”} & \text{otherwise.}\end{cases} \tag{1.5}$$

The first step in establishing that f is a MPE is to show that f leads to an outcome $\phi(x)$ starting from any initial allocation x. The only obstacle to this outcome is the possiblity that the process will get stuck because of an inadequate amount of the numeraire good. To avoid this possibility, we make the following assumption.

Interiority Assumption: For each i and any attainable allocation x we assume that either (i) $\phi_i(x) = x_i$ or (ii) for some $\varepsilon > 0$, $\phi(x) - (\varepsilon, 0, ..., 0) \in X_i$.

There will exist efficient allocations on the boundary of X where condition (ii) cannot be satisfied, but in that case we assume that no trade is optimal, so that agent i effectively does not participate in the trading process. Note that if the allocation $\phi(x)$ is individually rational relative to the initial endowment e then $\phi(x)$ belongs to the interior of X anyway, but since we have to define a continuation equilibrium for every possible subgame, i.e., for every possible allocation x, we cannot rely on individual rationality relative to e to ensure the trading process does not get stuck.

Under the interiority assumption, if all agents adopt the strategies in f then the allocation will converge almost surely to the Walras allocation $\phi(x)$.

Proposition 8 *For any initial allocation x let $\{z_t\}$ denote the sequence of excess demands generated along the equilibrium path by the strategy profile f. Then $\{z_t\}$ converges 0.*

Proof. By construction, the absolute values of the excess demands for commodities $h = 2, ..., \ell$ are monotonically non-decreasing. Thus, they must converge. Suppose the excess demands z_t converge to some limit z_∞, say.

If $z_\infty = 0$ then there is nothing to prove, so suppose that $z_\infty \neq 0$. Then there are two cases to be considered. Either (i) every agent has achieved his demand for commodity 1 and hence can afford to give up a small amount of commodity 1 in exchange for other commodities or (ii) some agent i has an excess supply of commodity 1. In case (i) market-clearing (attainability) and the budget constraints imply that there exist at least two agents i and j and some commodity $h \geq 2$ such that $z_{ih} < 0$ and $z_{jh} > 0$. By construction, $\zeta(x, i, j)$ is bounded away from 0 for any $x \in B(x^*)$ close to the limit value $x^* - z_\infty$, where $x^* = \phi(x)$ is the Walras allocation. Since (i, j) are matched infinitely often we have a contradiction of the assumption of convergence. In case (ii), $z_{i1}^0 < 0$ for agent i and the budget constraint implies that $z_{ih\infty} > 0$ for some commodity h. Then market clearing implies that there is some agent j such that $z_{jh\infty} < 0$ and by construction, $\zeta(x, i, j)$ is bounded away from 0 for any $x \in B(x^*)$ close to the limit value $x^* - z_\infty$. Since i and j meet infinitely often this contradicts the convergence assumption. This establishes that $z_\infty = 0$. ■

For any attainable allocation x, Proposition 8 shows that agent i's equilibrium payoff in the continuation game beginning at x is given by

$$v_i(x) = u_i(\phi_i(x)), \forall i = 1, ..., m.$$

To show that f is a MPE we need to show that the strategies defined by (1.4) and (1.5) are best responses at every information set. A sufficient condition for this is the following:

Independence Assumption: The competitive equilibrium price vector is independent of the attainable allocation x, that is, $\pi(x) = p$ for all x.

This assumption is quite strong. It would be satisfied, for example, in a representative agent economy, where all agents had identical, homothetic preferences. The independence assumption would also be satisfied if agents had transferable utility. Strictly speaking, transferable utility is not consistent with other assumptions we have made, though the results would probably continue to hold without them.

To show that the proposer's strategy is optimal, suppose that the current allocation is x, the pair (i, j) is matched, and the vector z is traded. That is, agent i receives the bundle $x_i + z$ and agent j receives the bundle $x_j - z$. Let x' denote the new allocation and let $p = \pi(x) = \pi(x')$. If $p \cdot z < 0$, then

$$\begin{aligned} p \cdot \phi_i(x') &= p \cdot x_i' \\ &< p \cdot x_i \\ &= p \cdot \phi_i(x) \end{aligned}$$

which implies that $v_i(x) = u_i(\phi_i(x)) > v_i(x') = u_i(\phi_i(x'))$, so that agent i is worse off. On the other hand, the responder will not accept any vector z such that $p \cdot z > 0$. Thus, the only possibility for i to increase his payoff is to offer a trade z such that $p \cdot z = 0$. But we know from the definition of f that any such z leads to $\phi(x)$ and so does not increase the ultimate payoff. Thus, offering $\zeta(x, i, j)$ is weakly optimal for the proposer.

Now we can show that the responder's strategy is optimal. The responder accepts any vector z such that $p \cdot z \leq 0$. By the preceding argument we can show that any trade z such that $p \cdot z > 0$ makes j worse off, so it is optimal to reject such offers, and any trade z such that $p \cdot z = 0$ leads to the same outcome so it is weakly optimal to accept such offers. On the other hand, if

$p \cdot z < 0$ then

$$\begin{aligned} v_j(x') &= u_j(\phi(x')) \\ &\geq \arg\max\{u_j(\hat{x}_j)|\hat{x}_j \in X_j, p \cdot \hat{x}_j \leq p \cdot x'_j\} \\ &> \arg\max\{u_j(\hat{x}_j)|\hat{x}_j \in X_j, p \cdot \hat{x}_j \leq p \cdot x_j\} \\ &= u_j(\phi_j(x)) = v_j(x), \end{aligned}$$

so it is optimal to accept such offers. This completes the proof of the desired result.

Theorem 9 *If the Uniqueness, Interiority and Independence Assumptions are satisfied, then the strategy profile f defined in (1.4) and (1.5) is a MPE and the equilibrium outcome $\{x_t\}$ converges to the asymptotic allocation $x^* = \phi(e)$.*

The key elements of the argument presented above seem to be the Uniqueness Assumption, which ensures that the strategy f is Markov, and the Independence Assumption, which ensures that value-increasing (value-reducing) trades are payoff-increasing (payoff-reducing). An interesting question is whether and how the Independence Assumption can be weakened.

The Independence Assumption is clearly not a necessary condition. As long as the effect of price changes is not too great a similar argument would hold. For example, for any efficient allocation x let

$$B_i(x) = \{x'_i \in X_i|\pi(x) \cdot x'_i = \pi(x) \cdot x_i\}.$$

Then the following assumption, which looks like a strengthening of the Uniqueness Assumption, will do the trick.

Non-Intersection Property: For any $x, x' \in P$ and any agent i, either $B_i(x) \cap B_i(x') = \varnothing$ or $B_i(x) = B_i(x')$.

The Non-Intersection Property is clearly implied by the Uniqueness and Independence Assumptions. It implies uniqueness, but it is weaker than independence because it allows for the supporting prices $\pi(x)$ to vary with x.

Corollary 10 *If the Non-Intersection Property and the Interiority Assumption are satisfied, then the strategy profile f defined in (1.4) and (1.5) is a MPE and the equilibrium outcome $\{x_t\}$ converges to the asymptotic allocation $x^* = \phi(e)$.*

Proof. Suppose as in the proof of Theorem 9 that two agents i and j have met and made a trade that changes the allocation from x to x'. If $\pi(x) \cdot \phi_i(x') > (\geq)\pi(x) \cdot \phi_i(x)$ and $\pi(x) \cdot \phi_j(x') \geq (>)\pi(x) \cdot \phi_j(x)$ then the Non-Intersection Property implies that $\pi(x) \cdot x'_i > (\geq)\pi(x) \cdot x_i$ and $\pi(x) \cdot x'_j \geq (>)\pi(x) \cdot x_j$. But this is impossible since $x_i + x_j = x'_i + x'_j$. Hence, one of the agents can only be made better off if the other is made worse off and this prevents any payoff-increasing deviation from the strategy f. ■

In general, we should not expect to achieve a competitive equilibrium relative to the initial endowments e when there is a small number of agents. Even with a large but finite number of agents, the market power possessed by individual agents may lead to distortions that prevent the attainment of the competitive equilibrium. But the market power of agents should get smaller as the number of agents gets larger and this may allow the attainment of an approximate competitive equilibrium as a MPE of the DMBG. Alternatively, the competitive equilibrium may be an ε-equilibrium for a sufficiently large number of agents m. For example, if we could show that for any allocation x that occurs on the equilibrium path and any trade that is possible for agents i and j, it is impossible for them to increase their utilities by more than ε by deviating from the equilibrium path, then the competitive equilibrium could be achieved as an ε-MPE. Whatever trades these agents can make, the equilibrium prices will not change by very much if the number of agents is very large, and consequently they cannot change their payoffs very much by distorting the prices. Note however that the notion of ε-MPE used here is applied to each information set separately. In calculating the increase in payoff from a deviation, an agent assumes that in the future play the other agents will follow the equilibrium strategies and hence he cannot make any payoff-increasing deviations in the future. A sequence of successful deviations might increase his payoff by more than ε but this is not taken into consideration because each of them in isolation is unable to increase payoffs by more than ε and hence will not succeed in being made.

1.9 Efficiency with Discounting

Up until now we have assumed that agents do not discount future utilities. In this section, the analysis is extended to include the possibility of discounting.

There are two reasons for wanting to include discounting. The first reason is that time matters. If convergence to an efficient or Walrasian allocation takes a very long time, then even a small amount of discounting will have a large effect on the payoffs. The assumption that agents do not care whether they consume early or late will not be a good approximation. By introducing discounting, we can test the robustness of the efficiency results when the time taken to converge matters. Of course, introducing a cost of time will change the results somewhat, since convergence will always take some time and discounting will reduce the eventual payoffs. For that reason, we concentrate on the case where the discount rate is small (the period length is short) and ask whether we obtain approximately the same results as in the case with no discounting.

The second reason for the interest in discounting is that it has played a crucial role in the bargaining literature, where it has been shown to yield determinate outcomes under certain circumstances. Beginning with the Stahl-Rubinstein theory, it has been shown that discounting forces agents to reach a determinate agreement, whereas the bargaining problem has an infinite number of solutions in the absence of discounting. So it is an interesting question whether discounting makes a difference in the present context.

In the bargaining literature, there are two interpretations of discounting. One is the usual interpretation that individuals prefer early consumption to late consumption and that their time preferences can be represented by geometrically discounting the utility of future consumption. If the (constant) rate of time preference is ρ, the discount factor applied to future utilities is defined to be $\gamma = (1+\rho)^{-1}$. An agreement reached at date t giving the agent an (undiscounted) utility w is worth $\gamma^{t-1}w$ in terms of present utility.

The second interpretation assumes that there is a positive probability that the bargaining process will be stopped before an agreement is reached. If no agreement is reached, the agents both receive zero, so the equilibrium payoff is equal to the sum over all t of the probability that the game does not end before t times the utility of the agreement reached at t. The probability that the game continues until date $t+1$, conditional on the game lasting until date t, is a constant $0 < \gamma < 1$, so the probability that the game reaches date t is γ^{t-1}. Let w_t denote the equilibrium payoff if agreement is reached at date t. Then the expected payoff is $\sum_{t=1}^{\infty} \gamma^{t-1} w_t$.

So, the two approaches produce the same payoff functions and are formally equivalent. (See Binmore, Rubinstein and Wolinsky (1986) for a subtle analysis of the differences between the two interpretations).

In the present context, the second interpretation is easier to use. We can simply assume that when the game stops at date t every individual consumes the bundle that he is currently holding. If the game stops at date t and agent i is holding the bundle x_{it} then he receives the utility $u_i(x_{it})$. Since the probability that the game stops at date t is $\gamma^{t-1}(1-\gamma)$, the expected utility of agent i from the outcome $\{x_t\}$ is

$$(1-\gamma)\sum_{t=1}^{\infty}\gamma^{t-1}u_i(x_{it}).$$

This interpretation has a number of advantages over the time-preference interpretation. First, under the time-preference interpretation it is not clear exactly when consumption takes place. It could take place any time after trading stops, but trading may continue for different lengths of time for different agents and forever for some. In any case, it is not clear how to apply time-preference discounting. Secondly, when the future is discounted, there may be an incentive for some agents to drop out of the trading process and consume their bundles immediately, if the gains from future trade are sufficiently small. Or there may be an incentive to begin consuming before trade has finished. Both of these possibilities, voluntary exit and continuous consumption, introduce complications that are beyond the scope of the present treatment.

1.9.1 The Equilibrium Path

If f^* is a MPE and $\{x_t\}$ is the equilibrium outcome, then the equilibrium payoff of agent i at the beginning of date t, before a proposal has been made, is a function of the attainable allocation x_t and the date t. Let $v_i(x_t,t)$ denote the equilibrium payoff of player i at date t when the initial allocation at that date is x_t. As in the case without discounting, we can show that the utilities $u_i(x_{it})$ converge to the equilibrium payoffs $v_i(x_t,t)$, but the argument is slightly different with discounting.

The no-trade strategy is always available to the agent, so at any date an agent can ensure that he holds his current bundle until the game stops. Since the game stops with probability one, his equilibrium payoff must be at least as great as the utility from his current bundle. Thus, at each date t and for each agent i,

$$v_i(x_t,t) \geq u_i(x_{it}). \tag{1.6}$$

Suppose now that the game has continued until date t. With probability $(1-\gamma)$ the game stops at date t and the agent is forced to consume his current bundle x_{it}. With probability γ the game continues at least one more period and he begins date $t+1$ with a continuation-game payoff of $v_i(x_{t+1}, t+1)$. Then the equilibrium payoffs must satisfy the recursive relation

$$v_i(x_t, t) = (1-\gamma)u_i(x_{it}) + \gamma v_i(x_{t+1}, t+1). \tag{1.7}$$

These two relationships (1.6) and (1.7) imply that the equilibrium payoffs are non-decreasing,

$$v_i(x_t, t+1) \leq v_i(x_{t+1}, t+1), \tag{1.8}$$

for every t.

From the definition of the payoffs, we know that

$$v_i(x_t, t) = (1-\gamma)\sum_{s=t}^{\infty} \gamma^{s-t-1} u_i(x_{is})$$

Consequently,

$$\begin{aligned} & v_i(x_t, t) - \gamma v_i(x_{t+1}, t+1) \\ = \; & (1-\gamma)\sum_{s=t}^{\infty} \gamma^{s-t-1} u_i(x_{is}) - \sum_{s=t+1}^{\infty} \gamma^{s-t} u_i(x_{is}) \\ = \; & (1-\gamma) u_i(x_{it}), \end{aligned} \tag{1.9}$$

Taking limits in (1.9) and dividing by $(1-\gamma)$ implies that

$$\lim_{t\to\infty} v_i(x_t, t) = \lim_{t\to\infty} u_i(x_{it})$$

Hence we have the following result:

Proposition 11 *Let f^* be a MPE of Γ, let $\{x_t\}_{t=1}^{\infty}$ be the equilibrium outcome and let $\{v_i(x_t, t)\}_{t=1}^{\infty}$ be the equilibrium payoffs. Then*

$$\lim_t u_i(x_i(t)) = \lim_t v_i(x_t, t).$$

and $v_i(x_t, t) \leq \lim_{t\to\infty} v_i(x_t, t)$ for every t.

1.9.2 Asymptotic Efficiency

Proposition 11 can be used to prove the asymptotic efficiency of the equilibrium allocations. The proof is the same as in the case without discounting and will not be proved here.

Proposition 12 *There exists an attainable allocation x_∞ such that*

$$\lim_{t\to\infty} x_t = x_\infty \in P$$

almost surely.

1.9.3 Pareto-Efficiency

Although this proposition parallels the efficiency theorem derived for the game without discounting, it is in fact weaker. Asymptotic efficiency implies nothing about the equilibrium payoffs of the game, viewed from the initial date, because the game terminates in finite time with probability one. The natural case to look at is the one in which $\gamma \to 1$, which one hopes will approximate the "frictionless" case $\gamma = 1$. Here there is a problem. If the convergence of the sequence $\{x_t\}$ is sufficiently fast, then a small risk of termination should not matter too much. But what happens if the rate of convergence is slow? More precisely, suppose that for each value of γ, f^γ is a MPE of $\Gamma(\gamma)$ and let $\{x_t^\gamma\}$ denote the corresponding sequence of attainable allocations along the equilibrium path. In general, one would expect $v_i(f^\gamma) < u_i(x_{i\infty}^\gamma)$, where x_∞^γ is the limit of $\{x_\infty^\gamma\}$ as $t \to \infty$. If the rate of convergence of $\{x_t^\gamma\}$ to x_∞^γ becomes slower as $\gamma \to 1$, then there is no reason to expect that the desired result

$$\lim_{\gamma\to 0}[u_i(x_{i\infty}^\gamma) - v_i(f^\gamma)] = 0$$

holds.

Some bound on the rate of convergence to the asymptotic allocation is needed to ensure that the effect of discounting disappears in the limit as $\gamma \to 0$. It is very difficult to obtain such a bound in general because the strategies can be very complex. There are too many possibilities that cannot, in our present state of knowledge, be ruled out.

Instead of dealing with the complexity of the current game, we can simplify it by assuming that the game has a finite number of states. Then the

Markov property and the finiteness of the state space immediately ensure that convergence in utility occurs in finite time. Since utility is non-decreasing and there are only a finite number of states, the utilities cannot keep changing indefinitely.

There are several possible motivations for the assumption of finiteness.

- Indivisibility can be motivated by realism. In practice, most commodities have an irreducible minimum quantity that can be traded. All trades have to be made in multiples of this unit. As a result, the set of attainable allocations is finite.

- The concept of a perfectly divisible commodity is an idealization of a commodity with a small but finite minimum unit. The introduction of a small minimum unit for each commodity can be thought of as a test of robustness of the model with perfectly divisible commodities. If the introduction of an indivisible unit makes a big change to the results, the perfectly divisible commodity is not a good approximation.

- A continuous strategy space allows for very complex strategies, which may be extremely sensitive to small changes in the state of the game. Imposing finiteness is a way of bounding complexity and thus building in an element of bounded rationality.

For any $\eta > 0$, let X_i^η denote the subset of X_i consisting of bundles consisting of integral multiples of η^{-1} of each commodity:

$$X_i^\eta \equiv \{x_i \in X_i | x_{ih} = m_h \eta^{-1}, \forall h, \exists m_h \in \{0, 1, ..., \ell\}\}.$$

The game with this consumption set substituted for X_i is denoted by $\Gamma(\gamma, \eta)$.

Let $f^{\gamma,\eta}$ denote a Markov perfect equilibrium of $\Gamma(\gamma, \eta)$ and $\{x_t^{\gamma,\eta}\}$ the corresponding outcome. Convergence in utilities is proved in exactly the same way as before (divisibility played no role in the proof) and it must occur in a finite number of steps since there is only a finite number of values that $v_i(x_t, t)$ can attain. Let $N^{\gamma,\eta}$ denote the minimum number of periods before every agent i reaches the limiting payoff $v_i^{\gamma,\eta}$ of the equilibrium $f^{\ \gamma,\eta}$. Finally, let P^ε denote the ε-Pareto-optimal set of the economy, that is, the set of attainable allocations x such that there does not exist an attainable allocation x' such that $u_i(x_i') > u_i(x_i) + \varepsilon$.

The previous arguments can be adapted to show that if $x_\infty^{\gamma,\eta}$ is a limit point of the sequence $\{x_t^{\gamma,\eta}\}$ as $t \to \infty$, then for any $\varepsilon > 0$ and all $\eta > \eta(\varepsilon)$, $x_\infty^{\gamma,\eta}$

must belong to P^ε. If not, then for some fixed $\varepsilon > 0$ we can find arbitrarily large values of η such that $x_\infty^{\gamma,\eta} \notin P^\varepsilon$. Let x_∞^γ be a limit point of $x_\infty^{\gamma,\eta}$ as this subsequence of values of $\eta \to \infty$. Clearly, $x_\infty^\gamma \notin P^\varepsilon$ and this implies that there exists an ordered pair (i, j) who meet infinitely often and a trade z such that $u_i(x_{i\infty}^{\gamma,\eta} + z) > u_i(x_{i\infty}^{\gamma,\eta})$ and $u_i(x_{i\infty}^\gamma + z) > u_i(x_{i\infty}^\gamma)$. Then, by continuity, for sufficiently large values of η we can find a trade $z^{\gamma,\eta}$ such that $z^{\gamma,\eta}$ is a feasible trade and $u_i(x_{i\infty}^{\gamma,\eta} + z^{\gamma,\eta}) > u_i(x_{i\infty}^{\gamma,\eta})$ and $u_i(x_{i\infty}^{\gamma,\eta} + z^{\gamma,\eta}) > u_i(x_{i\infty}^{\gamma,\eta})$. By previous arguments we can show that

$$v_k^{\gamma,\eta}(x_t^{\gamma,\eta}, t) \le \lim_{t\to\infty} v_k^{\gamma,\eta}(x_t^{\gamma,\eta}, t) = u_k(x_{k\infty}^{\gamma,\eta})$$

for $k = i, j$ and this allows us to generate a contradiction of the equilibrium conditions for the game $\Gamma(\gamma, \eta)$ for sufficiently large η and t, since agents i and j can make an improving trade and get a payoff better than $v_i(x_t^{\gamma,\eta}, t)$ and $v_j(x_t^{\gamma,\eta}, t)$, respectively. This contradiction proves the following result.

Lemma 13 *Let $f^{\gamma,\eta}$ be a MPE of $\Gamma(\gamma, \eta)$. For any $\varepsilon > 0$ there exists a number $\eta(\varepsilon)$ such that for all $\eta > \eta(\varepsilon)$, if $x_\infty^{\gamma,\eta}$ is the limit of the equilibrium allocations $\{x_t^{\gamma,\eta}\}$ then $x_\infty^{\gamma,\eta} \in P^\varepsilon$.*

Clearly, for some fixed value of $\eta > \eta(\varepsilon)$ the vector of equilibrium payoffs $v(f^{\gamma,\eta})$ will be ε-Pareto-efficient in the limit as $\gamma \to \infty$. In this sense, we can say that $f^{\gamma,\eta}$ is approximately efficient as $\gamma \to 1$ and $\eta \to 0$ in that order.

Theorem 14 *Let $f^{\gamma,\eta}$ be a MPE of $\Gamma(\gamma, \eta)$ for every value of (γ, η). For any $\varepsilon > 0$ there exist numbers $\eta(\varepsilon)$ and $\gamma(\varepsilon)$ such that for all $\eta > \eta(\varepsilon)$ and $\gamma > \gamma(\varepsilon)$, $f^{\gamma,\eta}$ is ε-Pareto-optimal, that is, there does not exist an attainable allocation x such that $u_i(x_i) > v_i(f^{\gamma,\eta}) + \varepsilon$ for every i.*

1.10 Random Matching

In many bargaining models, the alternating offers assumption is replaced by the assumption that proposers and responders are chosen at random. This assumption makes the model symmetric and hence allows for the possibility of stationary equilibria. Similarly, in a DMBG, random matching at each date imposes symmetry on the model and allows for stationary equilibria. In the model with a fixed order of play, the structure of the game depends on the date because the identity of the proposers and responders is a function of the date. There is a more compelling reason than this for being interested in random matching, however. Random matching ensures that the outcome of the DMBG is, in principle, random even if the agents are using pure strategies and this randomness raises a number of interesting issues. First, various convergence properties that require only elementary analysis in the deterministic case become much more subtle and require more powerful tools in the stochastic case. Secondly, because agents are risk averse, randomness becomes a source of inefficiency itself and hence can prevent the attainment of the competitive equilibrium.

In place of the deterministic order of play we adopt the following assumption:

Matching. At each date an ordered pair of agents (i, j) is chosen at random. Let π_{ij} be the probability that (i, j) is chosen. We assume that $\pi_{ij} > 0$ for all (i, j) such that $i \neq j$ and $\pi_{ii} = 0$, and $\sum_i \sum_j \pi_{ij} = 1$. The matching probabilities are the same at each date and independent across dates.

The assumption that all ordered pairs (i, j) are formed with positive probability at each date implies that all ordered pairs (i, j) are formed an infinite number of times with probability one. This ensures that the agents are all connected and form a single, integrated economy.

The rest of the model is defined as in Section 1.3 and the definition of MPE is the same as in Section 1.4. Because the outcome of the game is stochastic, it is no longer possible to restrict attention to individually rational allocations. Although an equilibrium must be individually rational ex ante, the outcome need not be individually rational ex post. So the assumption (1.1) has to be strengthened: instead, it is assumed that X_i coincides with one of the upper contour sets of u_i:

- for any i, for some constant c_i,

$$X_i \equiv \{x_i \in G_i | u_i(x_i) \geq c_i\} \tag{1.10}$$

The property (1.10) ensures that pairwise efficiency is equivalent to Pareto efficiency. Weaker conditions would also suffice for this result. For example, if we could ensure that indifference surfaces were tangent to the boundary of the consumption set at each point on the boundary, the same result would hold.

We begin the analysis of equilibrium by giving a more precise description of the equilibrium path. Because the matching process is random, the equilibrium path is a stochastic process even if pure strategies are chosen in equilibrium. Let $(\Omega, \mathcal{F}, P)$ denote the underlying probability space, where Ω is the set of states of nature, $\mathcal{F}$ is the σ-field of measurable subsets of Ω, and P is a probability measure on $(\Omega, \mathcal{F})$. All of the random variables that define the equilibrium path are assumed to be defined on this underlying probability space. For example, we can take Ω to be the set of sequences $\{(i_t, j_t)\}_{t=1}^{\infty}$, $\mathcal{F}$ the σ-field generated by the cylinder sets of Ω and P the probability measure determined by the kernel π. Then we can regard $x(t,\omega)$, $i(t,\omega)$, $j(t,\omega)$, $z(t,\omega)$, and $r(t,\omega)$ as stochastic processes defined on $(\Omega, \mathcal{F}, P)$.

The information available at the beginning of date t is represented by a sub-σ-field of $\mathcal{F}$, which is denoted by $\mathcal{F}_t \subset \mathcal{F}$. The elements of $\mathcal{F}_t$ are the observable events at the beginning of date t. They correspond to the possible histories h_t. The random variables determined at date t are not necessarily measureable with respect to $\mathcal{F}_t$ because they depend on information that only becomes available at date t and is therefore not included in $\mathcal{F}_t$. For example, the identity of the responder j_t is not known at the beginning of date t.

1.10.1 Convergence

Now suppose that f^* is a MPE. The equilibrium payoff to a player at the beginning of date t, before the random choice of proposer and responder has been made, is a function of the attainable allocation $x(t)$ and the date t. Recall that $x(t)$ can be determined by looking at the history h_t and hence is measurable with respect to $\mathcal{F}_t$. The equilibrium payoff of player i is denoted by $v_i(f^*)$ and the payoff conditional on date t and the state of nature ω can be denoted by $v_i(f^*|x(t,\omega),t)$. Since the context normally makes the equilibrium strategy clear, we suppress the reference to f^* and write $v_i(x(t,\omega),t)$ for the payoff at date t in state of nature ω.

Since there is no discounting, the equilibrium payoff at the beginning of date t must be equal to the expected value of the equilibrium payoff at the beginning of date $t+1$, conditional on the information available at the beginning of date t. Thus,

$$E[v_i(x(t,\omega),t)|\mathcal{F}_t] = E[v_i(x(t+1,\omega),t)|\mathcal{F}_t], \forall t. \tag{1.11}$$

We can define a stochastic process $\{V_i(t)\}$ by putting $V_i(t)$ equal to the equilibrium payoff at date t

$$V_i(t) = v_i(x(t), t)$$

and equation (1.11) implies that $\{V_i(t)\}$ is a **martingale** with respect to the filtration $\{\mathcal{F}_t\}_{t=1}^{\infty}$ (Karlin and Taylor (1975), Chapter 6).

The next step is to show that equilibrium payoffs are bounded. Trade is voluntary so the no-trade strategy is available. It follows immediately that in equilibrium,

$$v_i(x(t,\omega), t) \geq u_i(x_i(t,\omega)), \forall t. \tag{1.12}$$

This proves that $\{V_i(t)\}$ is bounded below. Also, since $x(t) \in \hat{X}$ and $\hat{X}$ is compact,

$$v_i(x(t,\omega), t) \leq \sup_{x \in \hat{X}} u_i(x_i) < \infty, \forall \omega,$$

so $\{V_i(t)\}$ is bounded above.

The Martingale Convergence Theorem tells us that a bounded martingale converges to a constant with probability one (Karlin and Taylor (1975) Theorem 5.1, p.278). In this case, there exists a random variable V_i^{∞} such that, for almost every ω, the sequence $\{V_i(\omega, t)\}$ converges to the constant $V_i^{\infty}(\omega)$. This is not the same as saying that there exists a constant c and that $\{V_i(t)\}$ converges almost surely to c. Every sample path becomes constant in the limit, but different sample paths may have different limiting values. The Martingale Convergence Theorem also tells us that the mean is preserved in the limit, so

$$E[V_i(t)|\mathcal{F}_t] = E[V_i^{\infty}|\mathcal{F}_t]. \tag{1.13}$$

The Martingale Convergence Theorem is a powerful tool. It tells us that equilibrium payoffs must converge and it allows us to draw this conclusion with very little information about the nature of the game or the equilibrium strategies.

Once we know that the equilibrium payoffs converge, we can show that the same must be true for the utility associated with the limiting allocations. To see this, recall that agent i's payoff is defined to be the expected value of the lim inf of the utility of his commodity bundle:

$$v_i(f^*) = E[\liminf_{t \to \infty} u_i(\xi_{it}(a^{f^*}))].$$

Thus, the expected payoff conditional on the information available at the beginning of date t must satisfy

$$E[V_i(t)|\mathcal{F}_t] = E[\liminf_{s} u_i(x_i(s))|\mathcal{F}_t]$$

Since the mean is preserved in the limit, according to equation (1.13), we have,

$$E[V_i^\infty|\mathcal{F}_t] = E[\liminf_s u_i(x_i(s))|\mathcal{F}_t]. \tag{1.14}$$

On the other hand, from the individual rationality condition, (1.12) we can see that

$$\begin{aligned} \limsup_t u_i(x_i(t)) &\leq \limsup_t V_i(t) \\ &= V_t^\infty, \end{aligned} \tag{1.15}$$

so, putting equations (1.14) and (1.15) together, we have that for almost every ω

$$\lim_t u_i(x_i(t)) = V_i^\infty$$

In other words, the utility of the commodity bundle held by agent i converges almost surely [a.s.] to a constant. Putting all of these results together, we have the following proposition.

Proposition 15 *Let f^* be a MPE of Γ, let $\{x(t)\}_{t=1}^\infty$ be the equilibrium outcome and let $\{V_i(t)\}_{t=1}^\infty$ be the equilibrium payoffs. Then*

$$\lim_t u_i(x_i(t)) = \lim_t V_i(t) = V_i^\infty \;\; a.s.$$

and $E[V_i(t)|\mathcal{F}_t] = E[V_i^\infty|\mathcal{F}_t]$.

As before, we can use this convergence result to show that the limiting allocation must be pairwise- and hence Pareto-efficient.

1.10.2 Efficiency

The essential idea is the same as in the non-stochastic case, trade continues until all gains from trade are exhausted and the resulting equilibrium allocation must be Pareto-efficient. The additional complication in this case is that the limiting payoffs may be random and depend on the realization of ω, so each sample path has to be considered separately.

For the proof of Proposition 3, we needed to strengthen the assumptions on preferences. Specifically, we assumed that:

- Each utility function u_i is strictly concave.

This assumption is needed here too to show that allocations converge if utilities converge. The next proposition characterizes the limiting set of allocations on the equilibrium path. It shows first that the current allocation converges to a constant and secondly that the limiting allocation is Pareto-efficient. But note that the limiting allocation may be a random variable, that is, it depends on the particular realization of the equilibrium path.

Proposition 16 *There exists a random variable x^∞ on Ω such that*

$$\lim_{t\to\infty} x(t) = x^\infty \in P$$

almost surely.

Proof. Let Ω^* denote the set of states such that for every $\omega \in \Omega$, each pair (i,j) is matched infinitely often and the equilibrium conditions are satisfied on the sequence $\{a(t,\omega)\}_{t=1}^\infty$. Clearly, Ω^* is a set of full measure. The first step in the proof is to show that for any fixed ω in Ω^*, $\{x(t,\omega)\}$ is a Cauchy sequence. The proof is by contradiction. Suppose, contrary to what we want to prove, that for some subsequence (using the same notation) $||x(t,\omega) - x(t+1,\omega)|| \geq \varepsilon$ for some $\varepsilon > 0$ and all t. Choose some further subsequence such that $i(t,\omega) = i$ and $j(t,\omega) = j$ for all t. It is always possible to choose this subsequence because the assumptions on the matching probabilities $[\pi_{ij}]$ guarantee that every ordered pair of agents (i,j) is matched infinitely often for almost every state ω. Then choose a further subsequence such that $x(t,\omega) \to y$ and $x_i(t+1,\omega) - x_i(t,\omega) \to z$ along this subsequence. It is possible to choose such a subsequence because the set of attainable allocations is compact. Since

$$\lim_{t\to\infty} u_k(t,\omega) = \lim_{t\to\infty} u_k(t+1,\omega) = V_k^\infty(\omega),$$

for $k = i, j$, it must be the case that

$$\begin{aligned} u_i(y_i) &= u_i(y_i + z) = V_i^\infty(\omega) \\ u_j(y_j) &= u_j(y_j - z) = V_j^\infty(\omega). \end{aligned}$$

By strict concavity,

$$\begin{aligned} u_i(y_i + z/2) &> V_i^\infty(\omega) \\ u_j(y_j - z/2) &> V_j^\infty(\omega). \end{aligned}$$

By continuity,

$$\begin{aligned} u_i(x_i(t,\omega)+z/2) &> V_i(t,\omega) \\ u_j(x_j(t,\omega)-z/2) &> V_j(t,\omega), \end{aligned}$$

for all t sufficiently large. This contradicts the equilibrium conditions since i can offer j to trade $z/2$ and make them both better off. This last step follows from the Markov property. The argument is the same as the one used in the proof of Proposition 3 for the deterministic case.

So the hypothesis that $\{x(t,\omega)\}$ is not Cauchy leads to a contradiction of the equilibrium conditions. This contradiction establishes that $\{x(t,\omega)\}$ is a Cauchy sequence, so $\{x(t)\}$ converges to a random variable x^∞ almost surely.

To show that $x^\infty(\omega) \in P$ almost surely we use a similar argument. Suppose that for some ω in Ω^*, $x^\infty(\omega) \notin P$. Then $x^\infty(\omega)$ is not pairwise-efficient, according to Proposition 1, and for some pair (i,j), there exists a feasible trade z such that

$$\begin{aligned} u_i(x_i^\infty(\omega)+z) &> V_i^\infty(\omega) \\ u_j(x_j^\infty(\omega)-z) &> V_j^\infty(\omega) \end{aligned}$$

and by continuity

$$\begin{aligned} u_i(x_i(t,\omega)+z) &> V_i(t,\omega) \\ u_j(x_j(t,\omega)-z) &> V_j(t,\omega), \end{aligned}$$

for all t sufficiently large. Since the pair (i,j) is matched infinitely often on ω, we can use the previous argument show that i must deviate by offering z to j, contradicting the equilibrium conditions. This shows that $x^\infty(\omega) \in P$ for all $\omega \in \Omega^*$. ■

Proposition 16 shows that the limit allocation is Pareto-efficient with probability one, but this is not the same thing as showing that the equilibrium is Pareto-efficient. The limit allocation is a random vector x^∞ that belongs to the efficient set P almost surely. Since agents are risk averse, Jensen's Inequality implies that

$$E[u_i(x_i^\infty)] \leq u_i(E[x_i^\infty]),$$

with strict inequality if $u_i(\cdot)$ is strictly concave and x_i^∞ is non-degenerate. Agents would rather have the expected value of the limiting allocation. The set of attainable allocations $\hat{X}$ is convex so $x^\infty \in \hat{X}$ implies that $E[x^\infty] \in \hat{X}$. To ensure Pareto-efficiency ex ante we would have to assume that x^∞ is degenerate, that is, almost surely equal to a constant allocation y.

Call a MPE is **asymptotically pure** if the allocations on the equilibrium path converge to a single allocation almost surely, that is, , for some constant $y \in \hat{X}$

$$x(t, \omega) \to y$$

for almost every ω. This is stronger than restricting attention to pure-strategy equilibria, because even with pure strategies the randomness associated with the matching process might lead to a random allocation in the limit. On the other hand, even if the MPE is asymptotically pure, it does not mean that the equilibrium is non-stochastic. Random matching may still have an effect along the equilibrium path leading to y and there may be randomness off the equilibrium path.

The assumption that a MPE is asymptotically pure is strong. Furthermore, it restricts endogenous variables so it is not entirely clear what it might require in the way of restrictions on the primitives of the model to ensure that the equilibrium is asymptotically pure. An example of a MPE that is not asymptotically pure is given in Section 1.11.

Restricting attention to MPE that are asymptotically pure, one can define a competitive sequence of economies and equilibria and show that in the limit as $m \to \infty$ the asymptotic allocations (as $t \to \infty$) converge to a Walras allocation for the limit economy. The proof is essentially the same as for the non-stochastic case and will not be repeated here.

1.10.3 Existence

The proof of existence given in Section 1.8 can easily be adapted to the case of random matching. The strategies are essentially the same and as long as all pairs (i, j) are matched infinitely often with probability one, the trading rules defined there guarantee convergence to a competitive equilibrium allocation with probability one. The proof that this is an equilibrium is the same as in the non-stochastic case because the asymptotic allocation is independent of the realized matches as long as all pair are formed infinitely often.

1.10.4 Discounting

It is also possible to rework the analysis of the DMBG with discounting for the assumption of random matching.

As before, the equilibrium path can be described as a stochastic process defined on the underlying probability space $(\Omega, \mathcal{F}, P)$. In addition to the random variables $x(t)$, $i(t)$, $j(t)$, $z(t)$, and $r(t)$ we have to define the Markov stopping time $\tau(\omega)$ at which the game stops.

If f^* is a MPE, then the equilibrium payoff to a player at the beginning of date t, before the random choice of proposer and responder has occurred, is a function of the attainable allocation $x(t)$ and the date t. Let $v_i(x(t,\omega),t)$ denote the equilibrium payoff of player i at date t in the state of nature ω. The no-trade strategy is always available to the agent, so at any date an agent can ensure that he holds his current bundle until the game stops. Since the game stops with probability one, his equilibrium payoff must be at least as great as the utility from his current bundle. Thus, at each date t and for each agent i,

$$v_i(x(t),t) \geq u_i(x_i(t)), \tag{1.16}$$

with probability one.

Suppose now that the game has continued until date t. With probability $(1-\gamma)$ the game stops at date t and the agent is forced to consume his current bundle $x_i(t)$. With probability γ the game continues at least one more period and he begins date $t+1$ with a continuation-game payoff of $v_i(x(t+1),t)$. Then the equilibrium payoffs must satisfy the recursive relation

$$v_i(x(t),t) = (1-\gamma)u_i(x_i(t)) + \gamma E[v_i(x(t+1),t+1)|\mathcal{F}_t], \tag{1.17}$$

where this equation is understood to hold with probability one. These two relationships (1.16) and (1.17) imply that the equilibrium payoffs form a **submartingale**, that is,

$$v_i(x(t),t) \leq E[v_i(x(t+1),t+1)|\mathcal{F}_t], \tag{1.18}$$

for every t (Karlin and Taylor (1975), p.248). To see this, suppose to the contrary that for some non-null set $S \in \mathcal{F}_t$,

$$\int_S v_i(x(t,\omega),t)P(d\omega) > \int_S v_i(x(t+1,\omega),t+1)P(d\omega). \tag{1.19}$$

Then, since $\gamma > 0$, we can use the recursive relation (1.17) and (1.19) to conclude that

$$\begin{aligned}\int_S v_i(x(t,\omega),t)P(d\omega) &= \gamma \int_S u_i(x_i(t,\omega))P(d\omega) + (1-\gamma)\int_S v_i(x(t+1,\omega),t+1)P(d\omega) \\ &< \gamma \int_S u_i(x_i(t,\omega))P(d\omega) + (1-\gamma)\int_S u_i(x_i(t,\omega))P(d\omega) \\ &= \int_S u_i(x_i(t,\omega))P(d\omega),\end{aligned}$$

which contradicts (1.16). This contradiction demonstrates (1.18) and proves that the payoffs are a submartingale.

As in the model without discounting, we define a stochastic process $\{V(t)\}$ consisting of the equilibrium payoffs by putting $V_i(t,\omega) = v_i(x(t,\omega),t)$ for every i and t. We have shown that $\{V(t)\}$ is a submartingale with respect to $\{\mathcal{F}_t\}$ and it is straightforward to show that $\{V(t)\}$ is bounded. Then the Martingale Convergence Theorem implies that $\{V_i(t)\}$ converges almost surely to a variable V_i^∞ and $E[V_i(t)|\mathcal{F}_t] \leq E[V_i^\infty|\mathcal{F}_t]$ (Karlin and Taylor (1975), p.278).

From the definition of the payoffs, we know that

$$E[V_i(t)|\mathcal{F}_t] = (1-\gamma)E\left[\sum_{s=t}^{\infty} \gamma^{s-t-1} u_i(x_i(s))|\mathcal{F}_t\right]$$

Consequently,

$$\begin{aligned}&E[V_i(t)|\mathcal{F}_t] - \gamma E[V_i(t+1)|\mathcal{F}_t] \\ =\ & (1-\gamma)E[\sum_{s=t}^{\infty}\gamma^{s-t-1}u_i(x_i(s))|\mathcal{F}_t] - E[\sum_{s=t+1}^{\infty}\gamma^{s-t}u_i(x_i(s))|\mathcal{F}_t] \\ =\ & (1-\gamma)E[u_i(x_i(t))|\mathcal{F}_t] = (1-\gamma)u_i(x_i(t)), \qquad (1.20)\end{aligned}$$

where the expectation operation $E[\cdot|\mathcal{F}_t]$ can be eliminated in the last line because $x_i(t)$ is $\mathcal{F}_t$-measurable. Taking limits in (1.20) and dividing by $(1-\gamma)$ implies that

$$V_i^\infty = \lim_{t\to\infty} u_i(x_i(t)) \text{ a.s.}$$

Hence we have the following result:

Proposition 17 *Let f^* be a MPE of Γ, let $\{x(t)\}_{t=1}^{\infty}$ be the equilibrium allocations and let $\{V_i(t)\}_{t=1}^{\infty}$ be the equilibrium payoffs. Then*

$$\lim_t u_i(x_i(t)) = \lim_t V_i(t) = V_i^{\infty} \;\; a.s$$

and $E[V_i(t)|\mathcal{F}_t] \leq E[V_i^{\infty}|\mathcal{F}_t]$.

1.10.5 Asymptotic Efficiency

Proposition 17 can be used to prove the asymptotic efficiency of the equilibrium allocations. The proof is the same as in the case without discounting and will not be proved here.

Proposition 18 *There exists a random variable x^{∞} on Ω such that*

$$\lim_{t\to\infty} x(t) = x^{\infty} \in P$$

almost surely.

1.10.6 Pareto-Efficiency

As in the model with a non-stochastic order of play (deterministic matching), the asymptotic efficiency result implies nothing about the efficiency of equilibrium, because the game ends in finite time with probability one. So in addition to the need to assume the MPE is asymptotically pure, one also needs to bound the time taken for convergence of the trading process. This can be done as in Section 1.9.3 by using a finite set of sets (allocations) to approximate the original game. The details are similar to those in Section 1.9.3 and will not be repeated here.

1.10.7 Competitive Sequences of Equilibria

Once efficiency has been established, one can analyze competitive sequences of equilibria in the usual way.

1.11 Mixed Equilibria

In the analysis of the model with random matching we focused on asymptotically pure MPE. There can be mixed equilibria as well. In fact, it is hard to rule out mixed equilibria. The essence of the problem arises from the fact that the set of Walras allocations is generically finite but not a singleton (Debreu (1970)). We have seen that, under certain special conditions, any deterministic equilibrium path leads to a Walras allocation. Furthermore, the number of MPE may be as large as the number of different Walras allocations. Mixed equilibria can be generated by randomizing over the MPE. For example, the Walras allocation that is eventually reached can be conditioned on some random element of the allocation process, which acts like a sunspot. If such an equilibrium exists, the final outcome will be a probability distribution over Walras allocations. If consumers are risk averse, a probability distribution over Walras allocations need not be Pareto-efficient, since there are no markets for trading this uncertainty. Ex ante, the game does not implement a Walras allocation even though ex post it will result in one of a finite number of Walras allocations.

To construct a mixed equilibrium with a random outcome, we first consider an exchange economy $\mathcal{E}$ that has multiple competitive equilibria. Then we show that each of these competitive allocations can be attained as the asymptotic outcome of a MPE. We can then construct a mixed equilibrium by using some random move by Nature to select one of the asymptotically pure MPE.

The construction is based on a two-person Edgeworth Box economy $\mathcal{E} = \{(X_i, u_i, e_i)\}_{i=1}^2$. For simplicity, no discounting is assumed. Since $\mathcal{E}$ has multiple competitive equilibria we cannot use the results from Section 1.8. However, when there are only two agents, any competitive equilibrium can be attained asymptotically as the outcome of a MPE. The first step is to give an explicit proof of this proposition.

Suppose that (p^*, x^*) is a competitive equilibrium for $\mathcal{E}$. Let ϕ be a function from $\hat{X}$, the set of attainable allocations, to the set P of Pareto-efficient allocations that satisfies the individual rationality condition:

$$u_i(\phi_i(x)) \geq u_i(x_i), i = 1, 2,$$

and is consistent with x^*:

$$\phi(e) = x^*.$$

We can think of $\phi(x)$ as the target allocation once x has been reached. Note that strict quasi-concavity of u_i implies that $\phi(x)$ is the unique Pareto-efficient allocation that is as good as $\phi(x)$, so

$$\phi(\phi(x)) = \phi(x).$$

Similarly, the payoff to player i conditional on reaching the allocation x is $v_i(x) = u_i(\phi_i(x))$.

Let $f = (f_1, f_2)$ denote an asymptotically pure MPE that implements (p^*, x^*) and define the strategy profile f as follows. If x is an attainable allocation and (i, j) a pair of matched agents such that $i \neq j$, then put

$$\begin{aligned} f_i(x, i, j) &= \phi_i(x) - x_i \\ &= x_j - \phi_j(x) \end{aligned}$$

and

$$f_j(x, i, j, z) = \begin{cases} \text{yes} & \text{if } z = \phi_i(x) - x_i \text{ or } v_j(x - z) > v_j(x) \\ \text{no} & \text{otherwise.} \end{cases}$$

In words, the proposer always makes an offer that, if accepted, will bring the pair to the target allocation $\phi(x)$; the responder accepts that offer or one that is strictly preferred and rejects all others.

The strategy profile f clearly implements the competitive allocation (p^*, x^*) in the sense that the final allocation will be x^* if these strategies are followed. To see that this is an asymptotically pure MPE, we need to show that the agents are choosing best responses at every information set. Suppose that the current allocation at the beginning of some period is x and that the pair (i, j) has been chosen by Nature. The equilibrium proposal is $z = f_i(x, i, j)$. The proposer cannot do any better because the proposal $z = \phi_i(x) - x_i$ is individually rational and the responder will reject any other offer unless it satisfies $v_j(x - z) > v_j(x)$. Since $\phi(x)$ is Pareto-efficient, this last inequality implies that $v_j(x - z) < v_i(x)$, so the proposer cannot do better than to offer z.

Similarly, the responder cannot do better than to follow this strategy along the equilibrium path because the proposal z is individually rational and if he rejects the offer, with probability one he will end up with the same allocation in the future. Off the equilibrium path, it is optimal to accept any proposal z that satisfies $v_j(x - z) > v_j(x)$ and weakly optimal to reject any proposal z that satisfies $v_j(x - z) \leq v_j(x)$, so again f_j is optimal for the responder.

This completes the proof of the following proposition.

Proposition 19 *Let (p^*, x^*) be a competitive equilibrium for the two-person economy $\mathcal{E}$. Then there exists an asymptotically pure MPE f of Γ such that the asymptotic allocation under f is x^*.*

Now suppose that the economy has two competitive equilibria (p^A, x^A) and (p^B, x^B). As shown above, there exist asymptotically pure MPE corresponding to each of the competitive equilibria, call them f^A and f^B respectively. Using these MPE strategy profiles, there is no difficulty in constructing a *subgame* perfect equilibrium (SPE) in which the outcome is a probability distribution over the two Walrasian outcomes (p^A, x^A) and (p^B, x^B). Simply condition the choice of strategy on who is first chosen to be the proposer.

- If $i = 1$ is chosen to be the proposer at date 1 then have both players adopt the strategies f^A corresponding to the MPE that leads to (p^A, x^A).
- If $i = 2$ is chosen to be the proposer at date 1 then have both players adopt the strategies f^B corresponding to the MPE that leads to (p^B, x^B).

It is clear that the strategies defined in this way constitute a SPE and that the outcome produces x^A with probability $1/2$ and x^B with probability $1/2$. Since the agents are risk averse, this random allocation is not Pareto-efficient—they would both prefer to receive $\frac{1}{2}(x^A + x^B)$ for sure—so it cannot be a Walrasian outcome. The equilibrium is not a *MPE*, as I have defined it, because the strategies depend on the move of Nature at the first date.

Theorem 20 *Suppose that the two-person economy $\mathcal{E}$ has two equilibria (p^A, x^A) and (p^B, x^B) with $x^A \neq x^B$. Then there exists a subgame perfect equilibrium such that with probability $1/2$ the asymptotic allocation is x^A and with probability $1/2$ the asymptotic allocation is x^B.*

The arguments we have used in previous sections cannot be used to eliminate the possibility of this equilibrium. Once the game begins with a move by Nature at date 1, it is already too late. The conditioning of the strategies on this random move prevents the achievement of the competitive equilibrium before anyone has a chance to move. The remaining play of the game satisfies all the conditions we could hope for, so equilibrium arguments cannot undo the conditioning.

Note that this conditioning of the play of the game on an initial move by Nature produces a correlated equilibrium in the sense of Aumann (1974), with a probability distribution over pure equilibria. The initial move by Nature plays the role of a "sunspot" or correlation device.

The structure of the equilibrium is very special and it could be argued that we are in a sense implementing the competitive equilibria. From date 1 onwards it appears that the game implements a single, non-stochastic competitive equilibrium relative to the initial endowments. However, there may be other, more complicated equilibria, that do not have this property. For example, there may be SPE in which the limit allocation is not a Walras allocation relative to the original endowments. Or there may be SPE in which, for any finite date t the limit allocation is not a Walras allocation relative to the current allocation at date t. We simply do not know.

If it turns out that randomness is an intractable obstacle to the implementation of competitive equilibria, then the theory would appear to be rather fragile. Random matching is a common ingredient in search models. It is included on the grounds of both realism and the simplicity it offers in many settings. At the very least we should want to include this kind of randomness.

One possible way of preserving the competitive theory in settings with random matching is to appeal again to large numbers, this time to get rid of aggregate uncertainty. Of course, if there is aggregate uncertainty in the world, then perhaps the definition of competitive equilibrium should be expanded to allow for conditioning on this uncertainty. The problem for the theory is small amounts of uncertainty that nonetheless matter a lot. But if there are large numbers of individuals and lots of randomness in the matching process, we might assume that overall it will not matter too much because the distribution of matches will not differ too much across realizations of the process. Of course, showing this may not be easy; but it is not implausible that for well behaved cases we can ignore the effect of this kind of uncertainty on the equilibrium outcome.

1.12 A Summing Up

This completes my attempt to develop a competitive theory based on dynamic matching and bargaining games with a finite number of agents. It has been a bumpy road. Strong assumptions have been required, assumptions imposed not on the primitives of the model, but rather on endogenous variables.

- Attention is restricted to Markov perfect equilibria (although as has been pointed out, it is only a small amount of information that needs to be ignored in order to prove the competitive limit theorem).

- The continuity assumption imposed on competitive sequences of equilibria says that individual agents have a negligible effect (relative to the entire economy).
- Equilibria are assumed to be asymptotically pure.
- A finite-state approximation to the original DMBG is required to ensure that, in the case of discounting, the time taken for convergence is not too long. Otherwise, the effect of discounting does not necessarily disappear as the discount factor γ converges to 1.

Furthermore, the implications of these assumptions are obscure. It is not clear what restrictions on the primitives of the model would be necessary to ensure these properties hold in equilibrium.

Nonetheless, the assumptions themselves are not implausible. The ideas that small agents have little impact on equilibrium, that strategies are Markov, that agents are anonymous, are ideas that economists are quite comfortable with. They seem to accord with our intuitive understanding of large economies and with at least some of the "facts" we know about competitive markets.

One possible conclusion is that competitive theory requires fundamental properties which, for some reason, cannot be derived from game-theoretic analysis. This is an intriguing idea, but one that does not give us much to say in terms of traditional game theory.

Another view is that the model needs to be changed. A richer model may give us more structure and allow us to actually simplify the equilibrium analysis. This approach is explored in the next chapter.