Chapter 0

Markets and Games

0.1 Strategic Foundations of Perfect Competition

In these lectures I report on a research program that began around fifteen years ago. It is part of a larger effort, underway for a much longer time, to provide strategic foundations for the theory of perfect competition[1]. The theory of competition has held a central place in economic analysis since the time of Adam Smith (1976). By providing strategic foundations for the theory of competition, economists use the principles of game theory to motivate or justify a macroscopic description of markets in which certain behavioral characteristics, such as price-taking behavior, are taken for granted. Game theory begins with individual agents and models their strategic interaction. A strategic foundation for competitive equilibrium must show how strategic interaction by rational agents leads to competitive, price-taking behavior. In practice, this research program includes the following three steps:

- First, describe a market or a whole economy.

In this step, the economist has to specify the commodities traded, the agents (households and firms) that make up the market or economy, their preferences, their resources and the available technology.

[1] When the context makes the meaning reasonably clear, I adopt the usual practice of writing *competition* or *competitive* when I really mean *perfect competition* and *perfectly competitive*.

- Secondly, define an extensive-form market game describing the behavior of the agents in the market or economy.

In this step, the economist has to specify the players, the information available to each player, the strategies available to them, the outcomes resulting from their choices, and the payoffs received.

- Thirdly, analyze the market game to show that, under certain conditions, the equilibrium outcome corresponds to a perfectly competitive equilibrium of the original market or economy.

There are many ways in which this program can be carried out. I shall be discussing one class of models which I find particularly interesting and fruitful, the class of *dynamic matching and bargaining games* (DMBG). Before getting into the details, however, I want to discuss the motivation for this kind of undertaking.

0.2 Why Strategic Foundations?

The first question that ought to occur to anyone encountering this kind of work for the first time is "Why?". Why should anyone want to take the time to build strategic foundations for the theory of perfect competition? After all, the theory of competitive equilibrium is well defined and self-contained in its own right. The cornerstones of this theory were laid over a hundred years ago by Alfred Marshall[2] and Leon Walras[3]. The modern axiomatic theory developed by Kenneth Arrow, Gerard Debreu, Lionel McKenzie and their successors is one of the most complete and definitive constructions that economics has to offer[4]. Instead of pursuing Steps 2 and 3, Arrow, Debreu, McKenzie, et al., simply assume that households maximize utility subject to a budget constraint, that firms maximize profits subject to their production technology, and that prices adjust to clear markets. Why do we need more than this? There are three reasons.

[2] Marshall, Alfred. *Principles of Economics; an introductory volume*, 8th ed. London: Macmillan (1920).

[3] Walras, Leon. *Elements of Pure Economics*. London: Allen and Unwin (1954).

[4] Arrow, Kenneth and Gerard Debreu.
Lionel McKenzie.

0.2.1 Games and Markets

The first reason is provided by the rise of game theory. Since the foundations of general equilibrium theory were laid in the nineteenth century, game theory has assumed an increasingly central role in economics. Game theory, as defined by von Neumann and Morgenstern, offers a general and powerful framework with which to analyze interactive decision making. To the extent that we accept the claim of game theory to be the correct framework for analyzing decision-making by individuals, we should want to use it as a tool for analyzing the behavior of agents in markets. Unfortunately, models of competitive equilibrium are not games in a strict, formal sense. In particular, games have two attractive features that models of market equilibrium do not have.

- In a strategic game, all the endogenous variables are chosen by players in the game.
- In a strategic game, any profile of strategies choosen by the players determines a unique feasible outcome of the game.

A long recognized and embarassing lacuna in the theory of competitive equilibrium is the failure to explain where the prices come from. Sometimes we are reduced to saying that prices are chosen by an "auctioneer", but essentially they are free parameters that are "determined", along with the other endogenous variables, by the market-clearing condition. In other words, unlike strategic games, models of market equilibrium have endogenous variables that are not chosen by the agents.

Another weakness of the classical model of competitive equilibrium is the assumption that agents believe that they can buy and sell as much as they like at the prevailing prices. It is true that agents buy and sell as much as they want in equilibrium, but if an agent deviates from his equilibrium excess demand, he will find that the assumption is violated. Except in the case where agents are literally negligible, any deviation leads to an infeasible set of excess demands. As a result, the market cannot clear and some agent will be disappointed in his attempt to trade. In the language of game theory, there are some strategy profiles to which no feasible outcome corresponds (unless we abandon the assumption that agents can trade as much as they want).

So, by comparison with the framework of game theory, the model of competitive equilibrium has some loose ends when it comes to the treatment of prices and the feasibility of trades.

These loose ends do not mean that the model of competitive equilibrium is a "bad model". On the contrary, it can be regarded as a reduced form of a more complicated model or process that describes the final outcome without giving all the details. One of the advantages of building a strategic foundation for perfect competition is that we will be forced to describe the process completely and explain how the competitive equilibrium outcome is reached. The complete model will, naturally, be an extensive-form game.

0.2.2 When is Perfect Competition Appropriate?

A second reason for wanting a strategic foundation of competitive equilibrium is to provide a theoretical *rationale* for perfect competition. A formal statement of the theory may be aesthetically pleasing, but it does not tell us why or under what circumstances the theory is appropriate. For example, we can define a perfectly competitive equilibrium for an Edgeworth Box economy with two goods and two agents, even though perfect competition is unlikely to be achieved in a market consisting of only two agents[5]. There is nothing in the formal definition to tell us whether the model applies to economies consisting of two, or a hundred, or a million agents.

One advantage of providing strategic foundations for perfect competition is that we are forced to construct a game in which perfect competition is only one of many possible outcomes. Then we show that under certain conditions competition arises as the unique equilibrium outcome. Deriving perfect competition from a more general framework provides us with a rationale for the competitive equilibrium: it shows under what conditions the model of competitive equilibrium is a good description of rational interaction among economic agents.

This aspect of the research program is illustrated by the distinction between *limit theorems* and *theorems in the limit.* Perfect competition is an idealized state, one which is only more or less imperfectly approximated by

[5] Ostroy (1980) and Makowski (1983) have characterized competitive equilibrium in terms of the *no-surplus condition.* They have argued that finite economies such as the Edgeworth Box economy sometimes satisfy the no-surplus condition and hence are perfectly competitive. However, the finite economies that satisfy the no-surplus condition are very special, and might be considered pathological.

reality and then only under certain special conditions. It is a non-trivial task to decide when the theory is a reasonable approximation. Some markets may be approximately perfectly competitive; others are not. How do we know which is which? Consider a sequence of economies, in which the number of agents is growing without bound. In the limit there is a continuum of individually insignificant agents. At what point does the economy become "competitive". Theorems in the limit characterize the exact conditions under which a perfectly competitive outcome may be expected to occur, whereas limit theorems tell us that, as we approach those conditions, the observed outcome will approach the competitive outcome.

By embedding different accounts of equilibrium behavior in a single theoretical framework, we are able to distinguish and classify the conditions under which different forms of competition arise. So another use for strategic foundations of perfect competition is to understand better the conditions under which perfect competition is an appropriate description of market behavior.

0.2.3 Normative Economics

Still another use for strategic foundations of competition comes from the role of the competitive equilibrium as a normative ideal. The classical theorems of welfare economics tell us that, under certain conditions, every perfectly competitve equilibrium allocation is Pareto-efficient and every Pareto-efficient allocation can be decentralized as a perfectly competitive equilibrium with lump sum transfers. But how is perfect competition brought about? What practical institutions will achieve the desired outcome? In order to use the theory of perfect competition as a practical tool for achieving the efficient allocation of resources, we need more than a definition of a perfectly competitive outcome. We also need a theory of the institutions and conditions under which perfect competition may arise as the result of interactive decision making by rational agents. In other words, we need a theory of the strategic foundations of perfect competition. Because game-theoretic models describe more completely the institutions that underly the market, they give us insight into the reasons why perfect competition may or may not be achieved. This knowledge may suggest policies to increase the degree of competition and, in the limit, achieve perfect competition. Alternatively, it may convince us that perfect competition is not the optimum in certain circumstances.

0.3 Cooperative Market Games

Attempts to provide a game-theoretic foundation for competition are almost as old as the theory of competition itself.

In *Mathematical Psychics*, Francis Ysidro Edgeworth (1881) addresses a broad theme, the applicability of mathematics in the social sciences. In particular, he addresses the question of whether the behavior of individuals is "determinate", by which he means "predictable using mathematical models". He argues that social processes, which might be indeterminate when small numbers of economic agents are involved, became determinate and hence subject to mathematical analysis when the number of agents is large. As an example, Edgeworth studies trade between two agents in a setting we have since come to call an Edgeworth Box economy. He argues that the outcome of such trade is indeterminate, because only efficiency and individual rationality could be counted on. But as the number of traders increases, the outcome of trade becomes increasingly determinate. The possibilities for recontracting among a large number of agents restrict the possible outcomes and, in the limit, allow for only competitive outcomes.

In *Mathematical Psychics*, Edgeworth follows the three steps of the program listed above. He describes an economy, represents the behavior of the agents by a coalition-forming game and shows that under certain conditions the solution of the game is corresponds to a perfectly competitive equilibrium outcome. In modern terminology, he describes a model of an exchange economy consisting of a finite number of economic agents. Each agent has an initial endowment of the commodities available in the economy, a consumption set that specifies the possible commodity bundles the agent can consume, and preferences over his consumption set. The allocation of resources is undertaken by coalitions of agents. Formally, a coalition is any non-empty set of agents. An allocation is attainable for a coalition if the commodity bundle assigned to each agent belongs to his consumption sets and the bundles add up to the coalition's total endowment. A coalition can improve on an attainable allocation if there exists an allocation that is attainable for the coalition and makes every member of the coalition better off. Edgeworth introduces the concept of the *contract curve* to describe the outcomes of the recontracting process. The contract curve is what we now call the *core* of the market game. It consists of the set of all attainable allocations that cannot be improved upon by any coalition. Edgeworth shows that as the number of agents grows unboundedly large, the core of the market game shrinks until it contains only the perfectly competitive equilibrium allocations.

This profound result, which stands virtually alone in nineteenth-century contributions to economic theory for its depth and beauty, was re-discovered by Shubik (1959) who showed the equivalence of the core and the contract curve. For the next fifteen years the theory was extended and refined until it reached a more or less definitive state with the publication of Hildenbrand (1974).

Now we have theorems that show that the Shapley value (Aumann and Shapley (1974)), the bargaining set (Mas Colell (1989)), the set of fair trades (Schmeidler and Vind (1972)), and other solution concepts also shrink to the set of Walras allocations as the number of players becomes very large. The fact that a variety of different solution concepts lead to the same result is a strong argument for the robustness of the competitive equilibrium.

All of these attempts to provide strategic foundations for perfect competition make use of *cooperative* game theory[6]. The cooperative approach to game theory has its limitations. Some of these are peculiar to particular cooperative solution concepts such as the core. Others are common to cooperative games in general.

At the general level, one of the attractions of cooperative game theory is that it provides a criterion for strategic stability that leads directly to a solution of the game, without the tedious business of specifying an extensive-form game. There is no need to specify the strategy set of each player, the order of moves, the information sets, or the players' assumptions about their opponents' behavior. All one needs is a convenient definition of what counts as a plausible outcome of the game.

In particular, there is no need to specify a well defined maximization problem for each individual player. However, this could also be considered a weakness. As John Hicks remarked in his "A Suggestion for Simplifying Monetary Theory" (reprinted in Hicks (1967)), anything seems to go when no one is required to maximize anything. A well defined maximization problem for each agent is one of the characteristic building blocks of modern economics. Cooperative game theory lacks the discipline that comes from having to specify a maximization problem for each agent. As a result, it

[6] For present purposes we can think of a cooperative game as being defined by a set of players, a specification of the coalitions or non-empty subsets of players that can form, and a specification of the actions or payoffs that can be achieved by each coalition of players. A solution of the game is a set of profiles of actions or payoffs that satisfies some plausible criterion of stability.

begs a number of questions.

This can be seen when we look at particular applications of cooperative game theory to market games. Take the core, for example. The formal definition of the core provides a criterion for stability, not a description of the process of coalition formation. An allocation belongs to the core if no coalition can improve on it. The definition suggests that if an allocation does not belong to the core it could never be an "equilibrium" because the improving coalition, whose existence is guaranteed by the definition, would do something to prevent it. Why the improving coalition should do that is not clear. Indeed, it is hard to think about this question without knowing how an allocation comes to be chosen in the first place; but suppose, for the sake of argument, that a non-core allocation has somehow occurred. It is easy to find allocations that (i) do not belong to the core, (ii) make some agents better off than they would be in any core allocation and (iii) can only be improved on by coalitions which include the agents who will be worse off at any core allocation. In this case, the improving coalition would be very myopic to veto this allocation if it really believed in the core as a solution concept, because it will certainly be worse off when a core allocation is finally reached. It appears that the core concept requires agents to behave myopically, rushing to join improving coalitions so that they can cut their own throats.

More sophisticated cooperative solution concepts try to eliminate such myopia; but the cooperative framework itself is the obstacle to a consistent theory, because it does not provide each agent with a well defined maximization problem. Without an extensive-form game, many of these problems can never be resolved satisfactorily.

For this reason, Nash (1951) proposed that cooperative games should be reduced to non-cooperative games. He interpreted a cooperative game as one in which there is unlimited pre-play communication and binding agreements can be entered into before the game is played. A non-cooperative game is one without pre-play communication or agreements. [This classification is clearly not exhaustive and there are other features that make it less than satisfactory, but the terminology has stuck so I will continue to use it.] The communication and commitment that precede the play of a cooperative game are not explicitly modeled, but they should be regarded as part of the game and analyzed using the same principles as the formal game itself. To do this, we model explicitly the pre-play communication and commitments, and then analyze the behavior of players in this game in the same way as if they were playing a non-cooperative game. This procedure, of reducing the cooperative game to a non-cooperative game by making explicit the informal parts of the cooperative theory, is known as the *Nash program*.

The Nash program is something that should appeal to economists, because it adopts Hicks's principle that every agent should be given a well defined maximization problem to solve. The use of non-cooperative game theory to provide a strategic foundation for competition is a natural extension of the earlier use of cooperative game theory. Ultimately, a satisfactory foundation for competitive markets requires a non-cooperative (extensive-form) game. This is what I shall be doing here, using non-cooperative game theory to provide a more complete description of what goes on in markets and then deriving the familiar competitive outcomes as an implication of particular conditions and assumptions.

0.4 Non-Cooperative Market Games

The first non-cooperative approach to competition antedates the core by about fifty years. It began with the analysis of duopoly by Antoine Augustin Cournot (1838, 1960). Cournot analyzes the problem of the owners of two mineral springs who want to maximize the profits from markets their spring water. A duopoly does not satisfy the conditions for perfect competition. The non-cooperative equilibrium of Cournot's model predicts that each supplier will restrict his output of spring water in order to raise the price and increase profits. However, as more and more suppliers are added one can, under certain conditions, show that the non-cooperative equilibrium converges to a competitive equilibrium in the limit as the number of suppliers becomes infinite. More importantly, the Cournot model is a well defined game for any number of players, taking as given the market demand curve faced by the suppliers.

This approach to competitive equilibrium was generalized by Shubik (1973) (see also Shapley and Shubik (1977)) with his Cournotian market games. Here is one version of a *Cournot-Shubik Market Game* (CSMG). Start with an exchange economy with a finite number of agents. Each agent has an endowment which he takes to the exchange floor and puts in the center. The agents are given different amounts of a token money which they can use to bid for quantities of the commodities in the center. A strategy for an agent is a vector of non-negative bids, one for each commodity, summing to his endowment of money. Once all the agents have chosen their bids, the commodities are divided in proportion to the bids, that is, an agent receives a fraction of the total endowment of each commodity equal to the ratio of his bid for that commodity to the total of the bid from all agents for that commodity. There are no explicit prices and each agent perceives that his final bundle of commodities depends not only on his own bids but also on the bids of all the other agents. Furthermore, each agent in a finite economy has market power, that is, the ability to change the implicit rate at which goods exchange for money by changing his bid. However, as the number of agents becomes unboundedly large, the Nash equilibrium outcomes of the CSMG converge, under certain regularity conditions, to the Walras allocations.

There are many other variants of these models, with commodity money, credit, bankruptcy, "wash trading", and other complications. The essential point is that the CSMG is a well defined extensive-form game that extends the partial equilibrium story of Cournot to encompass an economy in general equilibrium.

The CSMG provides a strategic foundation for competition and, as such, it is very much in the spirit of the research reported in these lectures. However, there are some critical features of the CSMG that separate them from the DMBG. The first is that although prices do not enter explicitly into the story, the definition of the CSMG imposes the restriction that all trades in a given commodity take place at the same price. More precisely, each agent gets a fraction of the total endowment of a given commodity and that fraction equals the ratio of his bid to the total bid. Each agent is trading money for the commodity at the same rate, that is, at the same price. In that sense, prices are built into the story in the same way as they were in the original Cournot duopoly model.

A second, more subtle feature is that the institutional structure is highly centralized, again by definition. Every agent is bidding against every other bidder and all the bids are aggregated (added up) in order to determine the final outcome. This is analogous to assuming that everyone trades on a centralized market and that the market mechanism treats every agent symmetrically. Furthermore, we can think of the ratio of the endowment to the total bids as a price. It is only the relaxation of the price-taking assumption that distinguishes this model from the Walrasian model of competitive equilibrium.

So the definition of the CSMG already incorporates some of the features of the Walrasian theory of competitive equilibrium. This is not a criticism of the CSMG as a strategic foundation for competitive equilibrium; however, there may be contexts in which we want to start with more primitive concepts, for example, to allow for decentralized trade or to allow different agents to trade at different prices. In that case, a framework like the DMBG will be needed.

0.5 Dynamic Matching and Bargaining Models

One successful modeling strategy is the class of *dynamic matching and bargaining games* (DMBG). It brings together elements of two branches of economic theory, search theory and bargaining theory. Economic agents are assumed to search at random for trading opportunities and when they meet the terms of trade are determined by bargaining. At one time markets actually worked like this (think of barter in primitive societies or the early stock exchanges). Although most markets have evolved since then, there is still enough realism in this framework to make it relevant. In addition, the flexibility of this framework makes it a good laboratory for "thought experiments", investigating questions about the behavior of actual markets and the design of better ones. In this section I review some of the achievements of the DMBG literature.

0.5.1 The alternating-offers bargaining model

The early nineteen-eighties was a period of tremendous activity in the theory of bargaining. The event that sparked this creative surge was the publication of Rubinstein's (1982) paper on the alternating offers bargaining model. Although the main ideas of the model had been proposed by Stahl (1972) ten years earlier and the analysis of Stahl's model with a finite horizon was already in some textbooks (Moulin (1986)), it was the successful analysis of the *infinite-horizon* bargaining problem by Rubinstein that grabbed the attention of theorists everywhere.

You probably know the elements of the Stahl-Rubinstein model (more general versions can be found in Binmore, Rubinstein and Wolinsky (1986)). Two individuals want to divide a "cake". Without loss of generality we can assume that the cake is worth one dollar and that any division $(x, 1-x)$ giving a share $x \geq 0$ to player 1 and $1-x \geq 0$ to player 2 is feasible. The players bargain as follows. First player 1 proposes a division $(x, 1-x)$ and player 2 accepts or rejects this proposal. If he accepts, the game ends and the two individuals split the cake as agreed. If he rejects, then player 2 proposes a division $(y, 1-y)$ and player 1 accepts or rejects. The bargaining process continues in this way until an agreement is reached. If agreement is never reached, they both receive nothing.

The individuals are asssumed to be impatient and for every period that agreement is delayed, their payoffs are discounted. The discount factor of player i is denoted by $0 < \delta_i < 1$. If the players agree on the division $(x, 1-x)$ at date t then player 1's payoff is $\delta_1^{t-1}x$ and player 2's payoff is $\delta_2^{t-1}(1-x)$. This discounting or shrinkage represents the costs of bargaining: other things being equal, the longer it takes to reach agreement, the less "cake" there is for both players.

As you will remember, Edgeworth (1981) thought of the bargaining problem as indeterminate. Nash (1953) himself proposed a simultaneous demand game in which any Pareto-efficient division was an equilibrium outcome. The remarkable result proved by Stahl and Rubinstein is that their bargaining problems are determinate. More precisely, the alternating-offers bargaining game has a unique, *subgame-perfect equilibrium*[7] in which the players reach agreement immediately and the division is

$$\left(\frac{1-\delta_2}{1-\delta_1\delta_2}, \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2}\right).$$

The outcome is asymmetric both because the players have different discount rates and because player 1 has a first mover advantage. If we assume that the players have the same discount rate $\delta_1 = \delta_2 = \delta$ then the equilibrium division is $(1/(1+\delta), \delta/(1+\delta))$. Player 1 gets the larger share because he is the first mover.

Now suppose that the bargaining process is "speeded up" by making the time interval between successive rounds smaller and smaller. When the period length becomes vanishingly small, the first-mover advantage disappears. For example, let τ be the period length and then the discount factor δ is given by $e^{-\rho\tau}$, where ρ is the instantaneous rate of time preference. Holding ρ constant, $\delta \to 1$ as $\tau \to 0$. It can easily be seen that the equilibrium division of the cake converges to the symmetric Nash bargaining solution $(1/2, 1/2)$ as δ converges to 1.

[7] In a Nash equilibrium each player chooses a best response at each information set that occurs in the equilibrium play of the game. In a subgame perfect equilibrium each player chooses a best response at every information set, whether or not it occurs in the equilibrium play of the game.

The uniqueness result does not survive some simple relaxations of the assumptions presented here. For example, if there are three players then the number of subgame perfect equilibria may be infinite (Shaked (1994)). If the set of possible distributions is discrete rather than continuous then, again, there may be multiple subgame perfect equilibria.

There is a lot more that can be said about this interesting framework (see Osborne and Rubinstein (1990) for a masterful treatment), but these are the essentials that are needed for what follows.

0.5.2 Marshallian markets

The application of non-cooperative bargaining theory to competitive equilibrium came in a seminal paper by Rubinstein and Wolinsky (1985). Rubinstein and Wolinsky use the alternating-offers bargaining model as a strategic foundation for "competitive" price determination in a Marshallian market.

The model contains a continuum of traders, M buyers and N sellers of an indivisible good. Each buyer wants at most one unit of the good, which he values at one dollar. The sellers each have one unit of the good, which they value at zero dollars. If a buyer and a seller trade, they realize a surplus of one dollar. These gains from trade are the "cake" over which the buyers and sellers will bargain.

In each period, the buyers and sellers are matched at random to form pairs of agents, each consisting of a buyer and a seller. One agent is chosen at random to propose a price at which to trade. The other can accept or reject this price. If an agreement is reached (a price is proposed and accepted), the pair of agents trade and leave the market. Each consumes his share of the surplus after leaving the market. If an agreement is not reached, they remain in the market until the next period. An unmatched agent remains passive until the next period.

A pair of agents can remain together and bargain for several periods, but if one of them is matched with another agent it is assumed that he leaves his former partner and begins bargaining with his new partner[8]. This is the

[8] It is interesting to note that it is optimal for agents to switch whenever they meet a new agent. Since all buyers (resp. sellers) are identical and the equilibrium is stationary, there is no advantage for a seller (resp. buyer) in remaining with his current partner if he has the option of switching. There are other equilibria, however. If agents choose to remain with their current partners regardless of whether they are matched with new partners, this is also optimal. Then the analysis of the model effectively reduces to the

only way that partnerships can be broken up.

Buyers and sellers discount the future using the same discount factor δ. If agreement is reached on a price p after t periods of bargaining, the seller receives $\delta^{t-1}p$ and the buyer receives $\delta^{t-1}(1-p)$.

Whenever an agent leaves the market, he is immediately replaced with an identical agent who has yet to trade, so the number of buyers and sellers is constant over time.

The model described above is not a game in the sense of von Neumann and Morgenstern (1980). In particular, there is no initial node and there is a continuum of players. Rubinstein and Wolinsky study what they call a *quasi-stationary subgame perfect equilibrium* which is like a subgame perfect equilibrium in which some possible nodes are not considered. They show that there exists a unique equilibrium in which agents reach agreement as soon as they are matched. The equilibrium division depends on the matching probabilities as well as the discount factor and the identity of the proposer. If a seller meets a new buyer each period with probability $0 < \alpha < 1$ and a buyer meets a new seller each period with probability $0 < \beta < 1$, then the equilibrium price is

$$x^* = \frac{2(1-\delta) + \delta\alpha - \delta(1-\delta)(1-\alpha)(1-\beta)}{2(1-\delta) + \delta\alpha + \delta\beta}$$

if the seller proposes and

$$y^* = \frac{\delta\alpha + \delta(1-\delta)(1-\alpha)(1-\beta)}{2(1-\delta) + \delta\alpha + \delta\beta}$$

if the buyer proposes.

The length of a time period in this model represents both the time between successive matches and the time between successive offers by any pair of bargaining agents. Reducing the period length is formally equivalent to reducing the rate at which agents discount the future. As the period length converges to zero,it can be shown that all trade takes place at the same price, that is,

$$\lim_{\delta\to 1} x^* = \lim_{\delta\to 1} y^* = \frac{\alpha}{\alpha+\beta}.$$

The limiting equilibrium price depends only on the matching probabilities since we have assumed buyers and sellers share a common discount factor.

examination of series of two-person bargaining problems. The equilibrium payoffs are the same as for the Rubinstein (1982) game.

The crucial point is that, for fixed values of α and β, the equilibrium price is bounded away from one and zero, so that both sides of the market get a positive share of the gains from trade.

In the Rubinstein-Wolinsky model, agents have to search for trading partners and search takes time. Because of the time taken to search for trading partners, this is not a "frictionless" market and for that reason alone we might not expect it to be perfectly competitive. Allowing the period length to converge to zero, so that the discount factor converges to one, is a way of reducing the transaction costs or "frictions" so that the market becomes frictionless in the limit. However, Rubinstein and Wolinsky argue that even in the limit, when the market becomes frictionless, their model does not produce the competitive outcome.

Rubinstein and Wolinsky suggest that we take as the perfectly competitive benchmark an auction market with M buyers and N sellers. The demand correspondence $D(p)$ for such a market is defined by

$$D(p) = \begin{cases} M & \text{if } p < 1 \\ [0, M] & \text{if } p = 1 \\ 0 & \text{if } p > 1 \end{cases}$$

and the supply correspondence $S(p)$ is defined by

$$S(p) = \begin{cases} [0, N] & \text{if } p = 0 \\ N & \text{if } p > 0. \end{cases}$$

The unique market-clearing price is $p^* = 0$ if $M < N$ and $p^* = 1$ if $M > N$. In other words, if the number of buyers exceeds the number of sellers, the sellers get all the surplus; if the number of sellers exceeds the number of buyers, the buyers get all the surplus. Only if the number of buyers and sellers are equal can there be a market-clearing price between zero and one.

The interpretation of this result is complicated by the fact that there is, over the infinite history of the market, an infinite number of buyers and sellers. It is not clear how to define a market-clearing price in a market with an infinite number of buyers and sellers. Since every agent eventually gets to trade at the common price, it could be argued that the equilibrium of the Rubinstein-Wolinsky model is "competitive". Although the limiting price $\alpha/(\alpha+\beta)$ does not clear an auction market with M buyers and N sellers, but it does clear the market with an infinite number of buyers and sellers flowing through it.

Gale (1987) develops this theme in a paper that analyzes a Marshallian market with many types of buyers and sellers, each of whom wants to trade one unit of the indivisible good. In Gale's model there is a non-atomic continuum of buyers and sellers, with M_i buyers of type $i = 1, ..., I$ and N_j sellers of type $j = 1, ..., J$. Buyers of type i value one unit of the good at u_i and sellers of type j value the good at v_j. All agents discount the future using the common factor δ. If a buyer of type i agrees to buy one unit at a price p at date t, his payoff is $\delta^{t-1}(u_i - p)$. Similarly, if a seller of type j agrees to sell one unit of the good at a price p at time t, his payoff is $\delta^{t-1}(p - v_j)$.

Agents are randomly matched in pairs consisting of one buyer and one seller in each period. Matches last exactly one period and the buyer and the seller each have probability 1/2 of being chosen to make a proposal. The proposer names a price and the responder accepts or rejects. If an agreement is reached, a trade is executed at the agreed price and both agents leave the market. Otherwise they wait until the next period, when they will be matched again with new partners, and the bargaining process continues.

The assumption that matches only last one period does not allow for alternating-offers bargaining of the kind that occurs in Rubinstein (1982) and Rubinstein-Wolinsky (1985). This probably does not affect the outcome in a stationary equilibrium, since the buyer (seller) is replaced by an identical agent in the next period. However, the fact that bargaining is brought to an end exogenously does matter, as we saw earlier (see footnote 8). Another assumption that is crucial is that agents bargain under complete information. The agents' types are assumed to be common knowledge, so there is no uncertainty about the value of the good to the buyer and seller when they bargain.

In Gale's model, trades may occur at many different prices, depending of the types of buyers and sellers involved as well as the identity of the proposer. However, it can be shown that there is a unique stationary equilibrium of the model and as the discount factor converges to unity the stationary equilibrium prices converge to a common limit, that is, all trades take place at the same price. Furthermore, every agent who wants to trade at this price is able to do so in equilibrium. Since there is no discounting in the limit, each agent gets the payoff he would receive in a competitive equilibrium at this price.

Gale's model is not, strictly speaking, a generalization of the Rubinstein-Wolinsky model because matches are dissolved after one period, but in other respects it is similar. The limiting outcome in this case looks more like a competitive equilibrium because there are many types of agents trading at a single price. However, there exists the same problem of how to interpret the market-clearing condition. Suppose we take as a benchmark an auction market with M_i buyers of type i and N_j sellers of type j. If p is the limiting price, assuming that no one is indifferent to trading, the demand in the auction market is

$$\sum_{\{i:u_i>p\}} M_i$$

and the supply is

$$\sum_{\{j:v_j<p\}} N_j$$

and there is no reason why these two should be equal. So the limiting price p does not correspond to the market-clearing price of an auction market with M_i buyers of type i and N_j sellers of type j. Every agent does get to trade at the limit price p and, since there is no discounting, there is no cost of delay (every agent gets the same utility as if he were able to trade immediately at the price p). It is only by comparison with the static auction market that one is led to question whether the outcome is perfectly competitive.

A peculiarity of the Rubinstein-Wolinsky model is that buyers and sellers are replaced by identical agents whenever they leave the market. This ensures that the stock of agents in the market remains constant over time. Presumably, it is this constancy of the population of agents in the market at any date that suggests the comparison with a static auction market having the same population of agents.

This version of Gale's model adopts the Rubinstein-Wolinsky endogenous replacement assumption. Another way to model entry and exit is to assume a constant flow of potential entrants to the market, in which case a stationary state can only be maintained if the equilibrium price adjusts so that the numbers of buyers and sellers entering the market at each date are equal. If p is the equilibrium price, a buyer of type i will not enter the market if $u_i < p$ and a seller of type j will not enter the market if $v_j > p$. The absence of these types is important because it affects the matching probabilities and hence the equilibrium prices. Furthermore, stationarity requires that equal numbers of buyers and sellers enter the market in each period. Thus, the equilibrium price must equate the flows of buyers and sellers into the market. Assuming that no type is indifferent, stationarity requires that

$$\sum_{\{i:u_i>p\}} M_i = \sum_{\{j:v_j<p\}} N_j,$$

where M_i (resp. N_j) is the flow of potential buyers of type i (resp. potential sellers of type j) into the market at each date. This is exactly the Marshallian notion of equilibrium for an auction market with M_i buyers of type i and N_j sellers of type j.

The long side of the market adjusts so that the equilibrium price satisfies this condition. For example, if more sellers than buyers are entering the market at each date, the sellers will be on the long side. The number of sellers waiting to trade will increase until the price is reduced to the point where the number of sellers entering at each date is just equal to the number of buyers entering. Similarly, if the number of buyers entering the market is greater than the number of sellers, the price will have to rise. Here price is fulfilling its Marshallian role of adjusting to equate the flow of demand and supply, but the price itself is being determined by the relative bargaining power of buyers and sellers, which in turn is influenced by the relative numbers of buyers and sellers in the market.

0.5.3 Walrasian models of exchange

The problems of interpretation that arose in the previous section are caused by the stationarity assumption, which implies that an infinite measure of agents flows through the market over time. An economy with an infinite measure of agents does not have a well defined competitive equilibrium to serve as a benchmark. We can avoid this difficulty by working with a finite measure of agents, but this entails analyzing a non-stationary equilibrium.

An example of this approach is found in Gale (1986a,b,c) which take the basic idea of the dynamic matching and bargaining framework from Rubinstein-Wolinsky and apply it to a Walrasian model of exchange.[9]

Begin with an exchange economy consisting of a finite number of types of agents $i = 1, ..., I$ and a continuum of each type. All agents have the same consumption set $\mathbf{R}_+^\ell$, where ℓ is the number of commodities, and agents of type i have the same initial endowment e_i and the same utility function $u_i(x)$.

Trade takes place at a sequence of dates. At each date, agents have a probability $0 < \alpha < 1$ of being matched with another agent. Agents are distinguished by their type and by the bundle of commodities they are currently holding, so a typical agent is represented by an ordered pair (i, x)(agents of

[9] An extension of the simple Rubinstein-Wolinsky framework with a single type of buyer and seller to deal with non-stationary equilibria is found in Binmore and Herrero (1988a,b).

the same type typically have different commodity bundles). The state of the economy at any date t can then be described by a measure μ_t on the set of ordered pairs (i, x) representing types and current holdings . The probability of being matched with agents belonging to a measurable set A is proportional to the measure $\mu_t(A)$.

Suppose that a pair of agents with current commodity bundles x and y are matched at some date. The two agents' types and commodity bundles are common knowledge when the bargaining begins. Each of them has an equal chance of being chosen as the proposer. Once the proposer has been chosen, he offers a net trade z to the other agent. This net trade must be feasible for both agents. If x is the proposer's commodity bundle and y is the responder's, then feasibility requires that $x + z \geq 0$ and $y - z \geq 0$. The responder can accept or reject the offer or decide to leave the game. If the offer is accepted, then the proposed trades are executed and the agents proceed to the next date with their new bundles $x + z$ and $y - z$ respectively. Otherwise there is no trade.

Matches only last for one period. Agents remain in the market until they choose to exit and they only have the option to exit after they have been matched. Because there is always a positive probability of being unmatched in any period ($\alpha < 1$), there will always be a positive measure of agents who have not yet been matched and so remain in the game.

There is no consumption during the game. After exiting with a commodity bundle x an agent of type i consumes the bundle and receives the payoff $u_i(x)$. Note that there is no discounting. An agent who does not exit is assumed to receive an infinitely negative payoff, so every agent will choose to exit with probability one. Once an agent exits, he cannot re-enter and the game is effectively over for him. Since agents are not replaced, the number of active agents in the DMBG decreases over time. The model is therefore non-stationary, in contrast to the Rubinstein-Wolinsky model.

A number of regularity assumptions, which will not be discussed here, are used in the course of the analysis of this model. Taking these for granted, we can summarize the conclusion of the analysis as follows. A subgame-perfect equilibrium of the DMBG always results in a Walras allocation, that is, each player receives the commodity bundle that he would receive in a Walrasian equilibrium. Conversely, for any Walras allocation there exists a subgame-perfect equilibrium of the DMBG in which almost every agent left the market with the corresponding commodity bundle. In this sense, the dynamic matching and bargaining game implements the Walras allocations.

0.5.4 Recent developments

McLennan and Sonnenschein (1991) provide an elegant extension of Gale's model version of this model in which many of the regularity conditions required by Gale were relaxed. The McLennan-Sonnenschein version applies only to stationary environments, in which the population and its characteristics remain constant over time, while traders enter and exit the market. McLennan and Sonnenschein use an approach based on the idea of *fair net trades*. They show that a non-cooperative equilibrium of their model satisfies the fair net trade axioms proposed by Schmeidler and Vind (1972). This implies that the resulting allocation is competitive in the sense that, for some price vector p, every agent behaves as if he were maximizing his utility subject to the usual budget constraint defined by the price vector p. The interpretation of the McLennan-Sonnenschein model presents the same difficulty as the original Rubinstein-Wolinsky model. Since an infinite measure of agents pass through the stationary model, the market-clearing condition is not well defined. In fact, as Dagan, Serrano and Volij (to appear) have recently argued, there is something unsatisfactory about the feasibility requirements in their model. They require cumulative trades to be feasible, rather than requiring feasibility period by period. In this case, any allocation is feasible because one can operate Ponzi schemes.

Dagan et al. have further expanded the class of environments in which one can demonstrate the equivalence of DNMB equilibria and Walrasian equilibria, but they do so by dropping one of the key ingredients of the earlier theory. They allow finite coalitions of any size to form. One member is chosen to make a proposal and the rest accept or reject. Although Dagan et al. define a non-cooperative game, it is one that has a strong resemblence to the cooperative notion of the core of an exchange economy insofar as arbitrary finite coalitions are allowed to form and trade among themselves. Finite coalitions can be quite large and large numbers have a "convexifying" effect that explains why many of the restrictive assumptions needed in the analysis of pairwise matching and bargaining are not needed in the analysis of this game.

The recent developments in the theory of DMBG have eliminated many of the technical limitations of the early theory, but there remain a number of important open questions. These questions, to which I turn next, provide the motivation for the lectures on which the following chapters are based.

0.6 Open Questions

0.6.1 The continuum assumption

All of the papers referred to above assume that the economy or market consists of a non-atomic continuum of agents[10]. This *continuum assumption*, as I shall call it, is of course one of the standard assumptions underlying the theory of perfect competition. A competitive market is one in which no firm or consumer has significant market power. Now, this is always an assumption that can at best hold in an approximate sense. In studying economies with a continuum of agents, we hope that the continuum economy is a good approximation to an economy with a large but finite number of agents. To establish the validity of the theory we need to prove limit theorems rather than theorems in the limit. To do this it is necessary to assume the existence of a finite number of agents, let the number increase without bound, and see whether the equilibria of the finite games converge to the appropriate allocation in the limit. (The theory referred to earlier contains a number of limit theorems, of course, but they are limits with respect to the period length or discount factor, not the number of agents).

The first topic I deal with is the construction of a finite counterpart of the theory based on the continuum assumption. This is done in Chapter 1. Dropping the continuum assumption requires us to rebuild the theory in a substantially different way. Some parts of the analysis are straightforward. Others are difficult. The payoff in either case is that one begins to clarify the nature of the arguments and the assumptions that are required to establish competitive limit theorems.

The behavior of Marshallian markets with a finite number of agents has already been investigated by Rubinstein and Wolinsky (1990). The market consists of a finite number M of buyers and a single seller. The seller has one unit of the indivisible good to sell which he values at zero. Each of the buyers places a value of one dollar on the good. All agents discount the future using the same discount factor δ. An important difference between this game and the "game" studied in Rubinstein and Wolinsky (1985) is that, instead of having exogenous random matching, the seller can choose the identity of the buyer to whom he wishes to offer the good in each period. Rubinstein and

[10] For example, the set of agents can be identified with the points in the unit interval $[0, 1]$. The "number" or "mass" of agents in a set is given by its Lebesgue measure, if this is well defined. In particular, a single agent has measure zero.

Wolinsky show that for any price $0 \leq p \leq 1$ there exists a subgame-perfect equilibrium in which the good is sold for this price.

The construction of the subgame-perfect equilibria is delicate, but the essential idea is this. Let p be the price at which the good is traded in equilibrium. Obviously, the seller would like to trade at a higher price and the buyers would like to trade at a lower price. To support trade at the price p it is assumed that if the seller asks for a higher price, all of the buyers refuse to trade at that price and offer a lower price instead. The seller has no choice but to accept the lower price. The buyer who initially rejects the seller's offer gets the good at the lower price.

Similarly, if one of the buyers offers a price below p, the seller refuses and sells it to another buyer at a higher price. Since the buyers believe that one of them will always trade with the seller in equilibrium, they have no choice but to accept the higher price.

In this way, any deviator is punished and the agent who supports the punishment strategy by rejecting the deviator's proposal is rewarded. Note that because there is only one good, only one buyer can have the good in equilibrium and the other buyers get a payoff of zero. Thus, the buyers will be at least as well off supporting the punishment strategy as they are along the equilibrium path.

According to a well known "folk theorem", repeated games have many subgame-perfect equilibria[11]. Although DMBG are not repeated games[12], they have some of the features of repeated games. In particular, if an agent thinks that a deviation from the equilibrium play of the game will trigger a future punishment by his opponents, he may be deterred from deviating. These trigger strategies can then be used to support a large variety of equilibrium behaviors. Some of these equilibria can be eliminated in a model with a continuum of agents. With a continuum of agents, the action of a single agent has no impact on the payoff of the most other agents, so it is natural to assume that agents ignore the actions of any single agent[13]. In effect, agents

[11] For example, if players do not discount the future, then any feasible and strictly individually rational payoff vector can be achieved by some subgame-perfect equilibrium.

[12] Because the allocation of commodities changes over time, a different "game" is being played at each date.

[13] In a DMBG, a single agent's actions will of course be observed by his trading partners, who can hardly be expected to ignore them. However, a single agent will only ever meet a finite number of other agents, and having met him they will never meet him again. So it seems reasonable to assume that agents do not condition their current behavior on the

are anonymous. They are invisible to most other agents, with whom they never come into contact. Anonymity rules out trigger strategies of the kind described above, so an agent can deviate without worrying about the reaction of his future opponents. This fact allows us to eliminate a large number of equilibria. Anonymity is only an assumption, however, and it is one that is not natural in a finite economy, however large. This presents a serious challenge to the development of a theory based on large finite economies.

A finite exchange economy consists of n agents, indexed $i = 1, ..., n$, each characterized by a consumption set $X_i \subset \mathbf{R}^\ell$, an endowment of commodities e_i and a utility function u_i. Trade takes place at a sequence of dates $t = 1, 2, ...$ At each date, a randomly selected ordered pair of agents (i, j) is matched. The first agent i is the proposer and makes an offer in the form of a feasible net trade vector z to the second agent j, which j can accept or reject. In principle, this process can continue indefinitely. There is no discounting, so agents only care about the utility of the commodity bundle they hold after an unlimited number of trading periods have passed. Agent i's payoff is the limiting value of his utility $u_i(x_{it})$ as t goes to infinity, where x_{it} is the bundle of commodities held by agent i at the end of period t. Individual rationality and voluntary trade imply that the equilibrium payoff at each date is greater than or equal to the utility of the currently held commodity bundle $u_i(x_{it})$. Consequently, the limit of $u_i(x_{it})$ exists.

Attention is restricted to *Markov perfect equilibria* (MPE), that is, subgame-perfect equilibria in which the strategies memoryless (independent of the history of the game). At each date strategies depend only on the current state of the game. One can show that, under certain regularity conditions, the limiting allocation as time goes to infinity must be Pareto-efficient. Pareto-efficiency is a strong result in its own right. In a finite economy each agent has some degree of market power. For the usual reasons, one might expect imperfect competition to lead to distortions in the market as each agent tries to use his market power for his own advantage. In a MPE, however, this is not the case: trade must continue as long as there are unexploited gains from trade,. In the limit, all gains from trade are exhausted so the final allocation must be Pareto-efficient. This result has an obvious resemblence to the Coase Conjecture (see Gul, Sonnenschein, and Wilson (1986)).

Establishing efficiency is only the first step, albeit an important one, on the road to a competitive limit theorem. There are many Pareto-efficient

past actions of individual agents.

allocations and, in a finite economy, there is no reason why the limiting allocation should be a Walras allocation[14]. Large numbers of agents are needed to guarantee this. The next step, then, is to analyze the asymptotic properties of MPE as the number of agents increases without bound. Here there is a problem. Since there is no reason to expect the MPE to be unique (Walrasian equilibria are not unique and so the MPE of the market game cannot be unique either), there is always the possibility that a small deviation by a single agent might trigger a switch to a very different equilibrium of the continuation game. As a result, even in the limit, individual agents may have non-negligible market power. In fact, to make any progress in such a general setting, it is necessary to impose a *continuity assumption* of the kind introduced in the work of Green (1980, 1984). The appropriate continuity assumption here requires that a small change in a single agent's strategy leads to uniformly (as $n \to \infty$) small changes in the asymptotic (as $t \to \infty$) allocation. Under this assumption, one can show that asymptotic (as $t \to \infty$) equilibrium outcomes converge to Walrasian outcomes as the number of agents becomes unboundedly large.

These results assume that there is no discounting and consumption takes place after the the game is over so only the asymptotic allocation matters to the individual agents. This is in contrast to much of the literature on DMBG, where discounting has played an important role. Recall that in Rubinstein (1982), discounting of future utilities is essential uniqueness of the SPE. Without discounting, the model collapses to a Nash demand game in which any Pareto-efficient division is a SPE outcome. Discounting has been used to test the robustness of equilibria in many settings. It is also used as a way of representing transaction costs. Search and bargaining are costly because the longer they continue, the more the payoffs are discounted. So an interesting extension of the theory sketched above is to introduce discounting and see whether it significantly changes the results.

Discounting is formally equivalent to introducing a random stopping date for the trading process. Suppose that at each date there is a probability $0 < \gamma < 1$ that the game continues for one more period, given that it has lasted until the present date. As soon as the trading process ends, agents consume their current commodity bundle. The probability that the game

[14] Under standard convexity conditions, if x is a Pareto-efficient allocation there exists a price vector $p \neq 0$ such that, for every agent i, $p \cdot x_i' > p \cdot x_i$ if $u_i(x_i') > u_i(x_i)$. To ensure that x is a Walras allocation, i.e., that (p, x) is a competitive equilibrium, we need budget balance, i.e., $p \cdot x_i = p \cdot e_i$, for every agent i.

stops at date t is given by $\gamma^{t-1}(1-\gamma)$, that is, the probability γ^{t-1} that it continues at the first $t-1$ periods multiplied by the probability $1-\gamma$ that it stops at the t-th period. If the game stops at date t the agent's payoff is $u_i(x_{it})$, where x_{it} is the commodity bundle held at the end of date t. To find the expected utility of agent i we simply multiply the probability of stopping at date t by the utility if the game stops at date t and sum over all dates $t = 1, 2, ...$ to get

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

Here γ plays the role of a discount factor. As the length of a period gets vanishingly short, we assume that the probability of the game continuing for one more period converges to 1, that is, the "discount factor" γ converges to 1 just as in the Rubinstein-Wolinsky model.

In a DMBG, an agent's equilibrium payoff can be expressed as the sum of a current (flow) payoff and the discounted value of the payoff from the continuation game. The conditions for an optimal strategy can be expressed in terms of a recursive equation similar to Bellman's equation in dynamic programming. From this equation it is possible to show that the equilibrium payoffs for each agent form a gradient process. This allows one to establish that the allocation converges to an asymptotic allocation under quite general conditions. Convergence is crucial because, as we already have seen, trade continues in a MPE until all gains from trade have been exhausted. Under certain regularity assumptions the asymptotic equilibrium allocation must be Pareto-efficient.

In a DMBG with discounting, it is important to distinguish between the Pareto-efficiency of the asymptotic allocation and the Pareto-efficiency of the MPE. The presence of discounting (random termination) implies that with probability one the games ends at some finite date t. The asymptotic allocation is reached with probability zero and so becomes irrelevant to the welfare properties of the equilibrium. Even if we let the period length get vanishingly small, so that the continuation probability γ approaches 1, we cannot be sure that the equilibrium payoffs converge to the utility of the asymptotic allocation. The trouble is that the equilibrium path $\{x_t\}_{t=1}^{\infty}$ depends on the value of γ. As γ converges to 1, the rate at which the sequence $\{x_t\}_{t=1}^{\infty}$ converges to the asymptotic allocation may become slower and slower, so that the probability of ending up with an inefficient allocation in the short run remains non-negligible.

To deal with this challenge, we need to show that the sequence of allocations $\{x_t\}_{t=1}^{\infty}$ converges to the asymptotic allocation uniformly quickly for all values of γ. The argument I use is based on bounded rationality (more precisely, bounded complexity of strategies) and shows that if convergence to an efficient allocation occurs it must occur within a bounded number of periods, so that as the period length becomes vanishingly short (the discount rate converges to zero), the time required to reach the efficient limiting allocation converges to zero. In the limit, as the period length becomes vanishingly short, an efficient allocation is reached with probability one.

To sum up the story so far, the theory developed in Chapter 1 achieves three things:

- It shows that it is possible to extend the theory of competitive markets to apply to large but finite markets.
- In addition, by separating the conditions needed for efficiency and for budget balance, we gain further insight into the robustness or generality of the two properties.
- Finally, note that at each step, extra assumptions (the Markov property, continuity, bounded rationality) are needed.

This last point is, in some ways, the most interesting. It suggests that a theory of competition cannot be based on the equilibrium properties of the DMBG alone. In addition, we need assumptions about markets, like anonymity and continuity (small agents have small effects), that cannot be derived from the analysis of Nash equilibrium.

If these additional assumptions really are necessary, they need to be justified. This is the subject matter of Chapter 2.

0.6.2 Anonymity, continuity, and the Markov assumption

As the number of agents grows without bound, it may be natural to expect that the strategy of a single agent should not matter to the play of the game. As we have seen, this is not a principle of game theory: even when the number of players becomes very large, the actions of a single agent, however insignificant, can trigger a change in the behavior of a large number of the other players, thus making a big difference to the play of the game.

Still, it seems that this is not how markets work. Unless there is something special about a distinguished agent, for example, unless he has some private information that is significant for a large number of agents, it just does not seem right that this agent's actions should matter all that much.

Part of the conventional wisdom about markets is that they are anonymous mechanisms. As Hayek (1945) pointed out, one of the wonderful things about markets is the fact that they economize on information. Prices are all that we need to know in order to take the right decisions. It does not matter who trades what; prices are determined by the aggregate of the agents' demands and supplies. So, if a single agent changes his demand by a small amount, that should have only a small effect on prices. Similarly, identities should not matter. Everyone has the same trading opportunities and everyone should be able to execute the same trade at the same price.

While this kind of reasoning sounds plausible, it is hard to demonstrate that markets have these properties in a strategic context. Having accepted the methodological rules of game theory, we are not in a position to abandon them when they make the analysis difficult. If anonymity and continuity are properties of competitive markets, we should be able to find assumptions about the structure of the DMBG that lead to continuous and anonymous equilibria. Otherwise, we are left clinging to the position that there exist economic principles that cannot be deduced from game-theoretic reasoning. This is a bit like saying that there exist economic principles that cannot be deduced from rationality postulates, which is not a position most economists would support. The possibility of demonstrating continuity and anonymity from first principles is an important issue for the strategic foundations of competitive equilibrium.

The Markov assumption essentially says that only the current state of the economy matters and not the history of the game. There are two ways to relax this assumption. For the efficiency results, it turns out that only certain kinds of information have to be excluded from memory, essentially, rejected offers. For the competitive limit result, the state can be defined very broadly and can be expanded to include some elements of history. What is crucial is that the play of the game should respond continuously to the state of the game and that the state should respond continuously to the actions of individual agents. These assumptions seem reasonable but they do not follow strictly from the definition of equilibrium. However, they can be justified in terms of informational limitations. Equilibria in which a non-negligible number of agents remain pivotal as the total number of agents grows without bound are not robust. They require agents to have a great deal of information about what is going on in the economy. Limited communication, bounded rationality (limited computational ability), or limited memory can all be used to justify the Markov assumption.

One way of formalize the property that agents are not pivotal in the limit is by introducing limited recall. As is well known, repeated games have a large number of equilibria supported by various forms of trigger strategies. If a player deviates from the equilibrium path in one of these games, he is punished in the future play of the game by the response of the other players. If an agent's actions are hidden from other players, it becomes impossible to support so many equilibria. Unfortunately, an agent's actions are likely to lead to some observable effect, on payoffs for example, and this may allow for punishment strategies even if actions are not directly observable. In any case, introducing asymmetric information leads to intractable complications of its own.

An alternative strategy is to limit all agents' ability to recall the past and, hence, their ability to condition future behavior on the history of the game. More precisely, we take a repeated game and turn it into a degenerate type of stochastic game by introducing a set S of *states of the game*. There is a transition probability that maps the current state of the game and the current actions of the players into the next period's state. The strategies of the agents depend only on the current state of the game and not on the history of the game per se. So the state of the game serves as a summary of the past play and the current play depends only on this summary, not on the entire history. It is important to note that the state does not affect the payoffs directly. It serves only to condition strategies. In this sense, the state is an extraneous variable, like a "sunspot". The state evolves in response to the play of the game and, in each period, the players' strategies are functions only of the state of the game. This reconfiguration of a repeated game is not by itself restrictive. It merely provides a framework in which it is easy to specify a criterion of robustness. Specifically, I introduce some noise into the evolution of the state to take account of the fact that recall is imperfect and that small changes in the play of the game are likely to be overlooked. Under relatively mild regularity assumptions, it is possible to show that, in the limit, as the number of players grows without bound, the game becomes anonymous. That is, each player comes to have negligible influence on the evolution of the play of the game. This has the additional consequence that in choosing an optimal strategy at each date a player is not influenced by the fear that his choice might trigger some kind of retaliation. This is not quite the same thing as saying that each agent adopts a Markov strategy. In fact, it is possible that the equilibrium play of the game will be non-Markov, but to the extent that this is true, it is because the equilibria of this game may be correlated in the sense of Aumann (1974, 1987) and not because there are trigger strategies in the conventional sense. Although strategies may depend on the past, they depend on the aggregate past and not on the past play of any individual.

The crucial fact is that limited recall in a large game allows agents to ignore the possibility of punishment in the same way that Markov or memoryless strategies allow them to ignore the possibility of punishment.

For the same reason, because strategies depend on the aggregate history and not on individual histories, the effect of an individual agent is small, ensuring continuity in the limit as the number of agents becomes unboundedly large.

In this way, one can argue that the additional properties required to make markets perfectly competitive are, in fact, properties of robust strategic equilibria.

0.6.3 Equilibrium and disequilibrium

A different kind of question relates to the very notion of equilibrium. Why do we assume that economic systems are in equilibrium at all? It is sometimes argued that without an explanation of how the economy gets "into" a state of equilibrium, equilibrium theory is empty. In other words, we need a theory of disequilibrium to justify the use of equilibrium methods of analysis.

Over the years, there have been numerous attempts to explain the disequilibrium behavior of the economy and, in particular, to show that an economy in disequilibrium will eventually approach a state of equilibrium. At one time, the notion of economic equilibrium was considered analogous to the "state of rest" of a physical or mechanical system. This suggests an analogy with statics and dynamics in classical mechanics, where statics are concerned with the determination of states of rest and dynamics are concerned with the processes that lead to states of rest. If equilibrium is a state of rest, then there ought to be a dynamic theory that accounts for how the system is drawn to a state of rest or equilibrium.

The Walrasian theory of the *tâtonnement* (Walras (1954)), Edgeworth's notion of recontracting (Edgeworth (1881)), and the Marshallian dynamics of supply price and demand price (Marshall (1927)) are examples of this kind of dynamic analysis. In the nineteen-fifties and nineteen-sixties, there grew up a large literature on the stability of equilibrium, including both *tâtonnement* and non-*tâtonnement* processes (see Arrow and Hahn (1971) for a survey). In the last thirty years, this sort of analysis seems to have fallen out of favor, mainly because the specification of these processes is largely ad hoc, that is, not based on principles of rational behavior.

Instead, equilibrium dynamics has come to dominate economic theory. See Stokey, Lucas, and Prescott (1989) for a good example of how this kind of theory has matured. There are plenty of dynamic models, but they are models that are always in equilibrium. The dynamics they portray is the evolution or unfolding of an equilibrium, rather than a process leading to equilibrium.

There is something self-defeating about the rejection of an adjustment process because it is not based on rational behavior. Rational behavior presupposes a kind of equilibrium, at least, an equilibrium of the agent's behavior, so it seems a bit much to require a disequilibrium theory to be based on rational behavior. Taking this argument to its logical conclusion, we would have to give up the search for disequilibrium theory altogether. This kind of argument was made quite convincingly by Hahn (1974; reprinted 1984), when he described equilibrium behavior as mutually consistent and optimal behavior by a group of agents. This notion of equilibrium applies to any situation in which rational behavior is simultaneously and consistently exhibited by all economic agents, not just to a situation of rest. It effectively disqualifies disequilibrium behavior by refusing to admit that it is economic behavior in any acceptable sense.

In the nineteen-eighties, game theorists began to feel the need to explain or justify the use of the Nash equilibrium and its various refinements. This urge may have been inspired partly by the same question that led to the analysis of ad hoc adjustment processes in models of market equilibrium: how do the players get to equilibrium? It was also inspired partly by the vast numbers of equilibria found in many well known games, which made the predictive power of game-theoretic analysis rather weak. The hope was that dynamic analysis might be able to discriminate among equilibria in a way that refinements of equilibrium had failed to do (see, for example, Kandori, Mailath and Rob (1993) and Young (1993)).

The attempt to show how equilibrium is arrived at has taken two paths. One in based on ad hoc theories of learning or adaptation, in which a fixed set of players play a game repeatedly and gradually adjust their strategies in response to experience. An old example of this kind of analysis is the well known model of *fictitious play*. In this model, a one-shot game is infinitely repeated. Each player is endowed with beliefs about the mixed strategies of his opponents and chooses a myopic best response to these beliefs at each date. The players' beliefs are assumed to be the same as the relative frequencies with which strategies were adopted in the past. Under certain conditions, this process can be shown to converge to a Nash equilibrium of the stage game. There are numerous variations on this theme. (For a modern example, see Fudenberg and Kreps (1998)).

Another stream of research comes from *evolutionary game theory* (Weibull (1995)). Originally, this branch of game theory was concerned with the evolution of animal species, but it has been adapted by economists to the human condition. Typically, there is a large population of automata, each of whom is programed to play a single strategy. Pairs of automata are selected at random from the population and play a two-person game with their preprogrammed strategies. Then the number of automata using each strategy is changed to reflect the relative success of that strategy. In other words, the number of players using a successful strategy will increase; the number using an unsuccessful strategy will increase less rapidly or will decrease. Different adjustment rules give rise to different dynamics. Again, there are many variations on this theme.

This kind of analysis is subject to the same criticism as the disequilibrium theory based on models of market adjustment processes: it is not solidly based on rational behavior. On the other hand, the cachet of evolutionary theory and widespread interest in learning algorithms and artificial intelligence allows practitioners to claim, with some legitimacy, that they are engaged in a different kind of undertaking, one that has a validity grounded in biology and computer science, whatever economists think of it. Furthermore, the charge of *ad hocery* is not really valid as long as some principles are being used consistently. If one accepts this point of view, then it appears that this kind of analysis provides a (non-economic) foundation for (economic) equilibrium theory, by showing that economic equilibrium is the outcome of an evolutionary or adaptive process.

In the third lecture (Chapter 3), I develop a non-maximizing or disequilibrium model of adaptation in a market context. Instead of maximizing expected utility, the behavior of the agents in the market is characterized by rules of thumb. These behavioral rules then define a stochastic process, whose convergence or non-convergence to a market equilibrium is analyzed. This analysis can be used both to show that market equilibrium does not make unreasonable demands on the rationality of the agents, because even boundedly rational agents can "learn" to play equilibrium strategies. It can also be used to provide insights into the dynamics of markets without assuming that equilibrium has been reached (yet).

0.6.4 Partial Equilibrium, General Equilibrium, and the Coordination Problem

Alfred Marshall, the founder of the Cambridge school of economics, had a very practical approach to the study of economics. In his famous *Principles of Economics* (Marshall (1920)) he defined economics to be "the study of the behavior of men and women in the everyday business of life", a down-to-earth definition that reflected Marshall's hope that economics would be a discipline that would improve the lot of ordinary men and women. It is not surprising then that Marshall is mainly associated with partial equilibrium analysis, that amalgam of handy short cuts that allows economists to isolate particular phenomena and study them on the back of a virtual envelope, ignoring the fact that an economy is a complex system in which "everything affects and is affected by everything else". It may not be pure, but it is very practical.

The founder of the Lausanne school, on the other hand, was a man of a different intellectual temperament. Leon Walras evidently loved the systematic element of economics and his great work *Elements of Pure Economics* (Walras (1954))is devoted to elaborating an elegant theory of general equilibrium in which the interactions of individual agents and individual markets throughout the economy are aggregated to provide a precise account of the equilibrium of the entire economy. It may not be practical, but it is very pure.

An economic theorist today is very glad to have the techniques of both partial equilibrium analysis and general equilibrium analysis in his toolkit. One cannot imagine doing international economics or macroeconomics without a general equilibrium framework, any more than one could imagine industrial organization theory without partial equilibrium analysis. It is perhaps regrettable though that we have come to think of these two very useful forms of economic analysis as separate and distinct, rather than as different aspects of a single theory.

In these lectures, I have been talking about the foundations of the theory of market equilibrium. The theory that has been developed sometimes makes use of a partial equilibrium (Marshallian) framework, and sometimes a general equilibrium (Walrasian) framework. The choice is made for simplicity (or expediency), but it is the same theory that is being developed in either case. This sounds like a good thing, to the extent that it appears to unify the two aspects of the economic analysis of equilibrium. At the same time, there is something a bit worrying about it. Equilibrium in a single market is not the same thing as equilibrium in a vast system of markets. The business of reaching an equilibrium throughout the entire economy is much more complex than the process of reaching equilibrium in a single market, holding the rest of the economy steady. And yet for much of these lectures I am going to pretend that they are more or less the same thing. In other words, that the Walrasian framework is just a multi-dimensional version of the Marshallian framework. Formally, there is nothing to prevent me from doing so. But economic intuition suggests that this cannot be the whole story and, at the end, I shall want to return to this issue in some detail and question whether the ideas explored here are really a satisfactory account of the foundations of general equilibrium. This will lead to a re-consideration of the work of another great Cambridge economist, John Maynard Keynes.